JUNIPER NETWORKS | Engineering Simplicity

Troubleshooting and Monitoring for QFabric Systems

Published
2021-04-18

JUNOS

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Troubleshooting and Monitoring for QFabric Systems
Copyright © 2021 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Guide | ix

1 **Overview**

General Troubleshooting | 2

Understanding Troubleshooting Resources | **2**

Troubleshooting Overview | **6**

Alarms | 10

Alarm Types and Severity Levels | **10**

Chassis Alarm Messages on a QFX3500 Device | **12**

Interface Alarm Messages | **16**

Alarms | **16**

- System Alarms | **17**
- Configuring Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types | **17**
- System-Wide Alarms and Alarms for Each Interface Type | **18**

2 **Routine Monitoring**

Routine Monitoring Using the CLI | 22

show snmp | **22**

Tracing SNMP Activity on a Device Running Junos OS | **24**

- Configuring the Number and Size of SNMP Log Files | **25**
- Configuring Access to the Log File | **26**
- Configuring a Regular Expression for Lines to Be Logged | **26**
- Configuring the Trace Operations | **27**

Monitoring RMON MIB Tables | **29**

Displaying a Log File from a Single-Chassis System | **30**

Monitoring System Log Messages | **32**

Monitoring Traffic Through the Router or Switch | **34**

Displaying Real-Time Statistics About All Interfaces on the Router or Switch | 34
Displaying Real-Time Statistics About an Interface on the Router or Switch | 36

Pinging Hosts | 38

3 Troubleshooting Features

Configuration and File Management | 42

Returning to a Previously Committed Configuration | 42

Returning to a Configuration Prior to the One Most Recently Committed | 42
Displaying Previous Configurations | 43
Comparing Configuration Changes with a Prior Version | 44

Reverting to the Default Factory Configuration | 46

Reverting to the Rescue Configuration | 47

Freeing Up System Storage Space | 48

Ethernet Switching | 50

Troubleshooting Ethernet Switching | 50

Troubleshooting Layer 2 Protocol Tunneling | 51

Drop Threshold Statistics Might Be Incorrect | 52
Egress Filtering of L2PT Traffic Not Supported | 52

Troubleshooting Private VLANs on QFX Switches | 53

Limitations of Private VLANs | 53
Forwarding with Private VLANs | 54
Egress Firewall Filters with Private VLANs | 55
Egress Port Mirroring with Private VLANs | 57

Troubleshooting Q-in-Q and VLAN Translation Configuration | 58

Firewall Filter Match Condition Not Working with Q-in-Q Tunneling | 58
Egress Port Mirroring with VLAN Translation | 59

Hardware | 60

Troubleshooting QFX3100 Director Device Isolation | 60

High Availability | 63

Troubleshooting VRRP | 63

Interfaces | 65

Troubleshooting an Aggregated Ethernet Interface | **65**

Troubleshooting Network Interfaces | **66**

- Statistics for logical interfaces on Layer 2 interfaces are not accurate | **66**
- The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down | **67**

Troubleshooting Multichassis Link Aggregation | **68**

- MAC Addresses Learned on Multichassis Aggregated Ethernet Interfaces Are Not Removed from the MAC Address Table | **69**
- MC-LAG Peer Does Not Go into Standby Mode | **70**
- Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive | **71**
- Redirect Filters Take Priority over User-Defined Filters | **71**
- Operational Command Output Is Wrong | **72**
- ICCP Connection Might Take Up to 60 Seconds to Become Active | **73**
- MAC Address Age Learned on a Multichassis Aggregated Ethernet Interface Is Reset to Zero | **73**
- MAC Address Is Not Learned Remotely in a Default VLAN | **74**
- Snooping Entries Learned on Multichassis Aggregated Ethernet Interfaces Are Not Removed | **75**
- ICCP Does Not Come Up After You Add or Delete an Authentication Key | **75**
- Local Status Is Standby When It Should Be Active | **76**
- Packets Loop on the Server When ICCP Fails | **76**
- Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change | **77**
- No Commit Checks Are Done for ICL-PL Interfaces | **78**
- Double Failover Scenario | **78**
- Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up | **79**
- Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer | **79**
- Aggregated Ethernet Interfaces Go Down | **80**
- Flooding of Upstream Traffic | **81**
- ARP and MAC Table Entries Become Out of Sync in an MC-LAG Configuration | **81**

Junos OS Basics | 83

System Troubleshooting | **83**

- Saving Core Files Generated by Junos OS Processes | **83**
- Viewing Core Files from Junos OS Processes | **84**

Recovering from a Failed Software Installation | **85**

Recovering the Root Password for Switches | 87

Creating an Emergency Boot Device for QFX Series Switches | 89

Performing a Recovery Installation | 91

Performing a QFabric System Recovery Installation on the Director Group | 94

- (Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive | 95
- Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software | 96

Troubleshooting Network Interfaces | 103

- Statistics for logical interfaces on Layer 2 interfaces are not accurate | 103
- The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down | 104

Troubleshooting an Aggregated Ethernet Interface | 105

Layer 3 Protocols | 107

Troubleshooting Virtual Routing Instances | 107

- Direct Routes Not Leaked Between Routing Instances | 107

Network Management | 109

Understanding Troubleshooting Resources | 109

Troubleshooting Overview | 113

Recovering from a Failed Software Installation | 116

Returning to a Previously Committed Configuration | 119

- Returning to a Configuration Prior to the One Most Recently Committed | 119
- Displaying Previous Configurations | 119
- Comparing Configuration Changes with a Prior Version | 121

Reverting to the Default Factory Configuration | 123

Reverting to the Rescue Configuration | 124

Recovering the Root Password for Switches | 124

Troubleshooting a Deprecated Network Analytics Configuration | 126

Security | 128

Troubleshooting Firewall Filter Configuration | **128**

- Firewall Filter Configuration Returns a No Space Available in TCAM Message | **129**
- Filter Counts Previously Dropped Packet | **131**
- Matching Packets Not Counted | **132**
- Counter Reset When Editing Filter | **133**
- Cannot Include loss-priority and policer Actions in Same Term | **134**
- Cannot Egress Filter Certain Traffic Originating on QFX Switch | **134**
- Firewall Filter Match Condition Not Working with Q-in-Q Tunneling | **135**
- Egress Firewall Filters with Private VLANs | **136**
- Egress Filtering of L2PT Traffic Not Supported | **137**
- Cannot Drop BGP Packets in Certain Circumstances | **137**
- Invalid Statistics for Policer | **138**
- Policers can Limit Egress Filters | **138**

Troubleshooting Policer Configuration | **140**

- Incomplete Count of Packet Drops | **140**
- Counter Reset When Editing Filter | **141**
- Invalid Statistics for Policer | **142**
- Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured | **142**
- Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured | **144**
- Policers Can Limit Egress Filters | **144**

Services | 147

Troubleshooting Port Mirroring | **147**

- Egress Port Mirroring with VLAN Translation | **147**
- Egress Port Mirroring with Private VLANs | **148**

Traffic Management | 150

Troubleshooting Dropped FCoE Traffic | **150**

Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth | **154**

Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth | **155**

Troubleshooting Egress Queue Bandwidth Impacted by Congestion | **157**

Troubleshooting an Unexpected Rewrite Value | **158**

Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic | **160**

Virtual Chassis Fabric | 165

Troubleshooting Virtual Chassis Fabric | **165**

Large-Scale Virtual Chassis Fabric Becomes Unstable When Logging is Enabled | **165**

Virtual Chassis Port Link Does Not Form | **166**

QFX5100 Leaf Device Assumes Routing Engine Role | **167**

About This Guide

Use this guide to troubleshoot the QFabric system.

1

PART

Overview

General Troubleshooting | 2

Alarms | 10

CHAPTER 1

General Troubleshooting

IN THIS CHAPTER

- Understanding Troubleshooting Resources | 2
- Troubleshooting Overview | 6

Understanding Troubleshooting Resources

This topic describes some of the troubleshooting resources available for the QFX Series or OCX Series. These resources include tools such as the Junos OS CLI, Junos Space applications, and the Advanced Insight Scripts (AI-Scripts).

Table 1 on page 2 provides a list of some of the troubleshooting resources.

Table 1: Troubleshooting Resources on the QFX and OCX Series

Troubleshooting Resource	Description	Documentation
Chassis alarms	Chassis alarms indicate a failure on the switch or one of its components. A chassis alarm count is displayed on the LCD panel on the front of the switch.	"Chassis Alarm Messages on a QFX3500 Device" on page 12
Chassis Status LEDs and Fan Tray LEDs	A blinking amber Power, Fan, or Fan Tray LED indicates a hardware component error. A blinking amber Status LED indicates a software error.	Chassis Status LEDs on a QFX3500 Device

Table 1: Troubleshooting Resources on the QFX and OCX Series *(Continued)*

Troubleshooting Resource	Description	Documentation
Interface alarms	A predefined alarm (red or yellow) for an interface type is triggered when an interface of that type goes down.	*Interface Alarm Messages*
System alarms	A predefined alarm is triggered by a missing rescue configuration or problem with the software license.	*Understand Alarms*
System log messages	The system log includes details of system and user events, including errors. Specify the severity and type of system log messages you wish to view or save, and configure the output to be sent to local or remote hosts.	• Overview of Single-Chassis System Logging Configuration • Junos OS System Log Configuration Statements
Junos OS operational mode commands	Operational mode commands can be used to monitor switch performance and current activity on the network. For example, use the **traceroute monitor** command to locate points of failure in a network.	• Monitoring System Process Information • Monitoring System Properties • traceroute monitor
Junos OS automation scripts (event scripts)	Event scripts can be used to automate network troubleshooting and management tasks.	Automation Scripting User Guide

Table 1: Troubleshooting Resources on the QFX and OCX Series *(Continued)*

Troubleshooting Resource	Description	Documentation
Junos OS XML operational tags	XML operational tags are equivalent in function to operational mode commands in the CLI, which you can use to retrieve status information for a device.	*Junos XML API Operational Developer Reference*
NETCONF XML management protocol	The NETCONF XML management protocol defines basic operations that are equivalent to Junos OS CLI configuration mode commands. Client applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as **show**, **set**, and **commit** to perform those operations.	NETCONF XML Management Protocol Developer Guide
SNMP MIBs and traps	MIBs enable the monitoring of network devices from a central location. For example, use the Traceroute MIB to monitor devices remotely.	• SNMP MIBs Support • SNMP Traps Support • Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS

Table 1: Troubleshooting Resources on the QFX and OCX Series *(Continued)*

Troubleshooting Resource	Description	Documentation
AI-Scripts and Advanced Insight Manager (AIM)	AI-Scripts installed on the switch can automatically detect and monitor faults on the switch, and depending on the configuration on the AIM application, send notifications of potential problems and submit problem reports to Juniper Support Systems.	*Advanced Insight Scripts (AI-Scripts) Release Notes*
Junos Space Service Now	This application enables you to display and manage information about problem events. When problems are detected on the switch by Advanced Insight Scripts (AI-Scripts) that are installed on the switch, the data is collected and sent to Service Now for your review and action.	Service Automation
Junos Space Service Insight	This application helps in accelerating operational analysis and managing the exposure to known issues. You can identify devices that are nearing their End Of Life (EOL) and also discover and prevent issues that could occur in your network. The functionality of Service Insight is dependent on the information sent from Service Now.	Service Automation
Juniper Networks Knowledge Base	You can search in this database for Juniper Networks product information, including alerts and troubleshooting tips.	https://kb.juniper.net

Troubleshooting Overview

This topic provides a general guide to troubleshooting some typical problems you may encounter on your QFX Series or OCX Series product.

Table 2 on page 6 provides a list of problem categories, summary of the symptom or problem, and recommended actions with links to the troubleshooting documentation.

Table 2: Troubleshooting on the QFX Series

Problem Category	Symptom or Problem	Recommended Action
Switch hardware components	LCD panel shows a chassis alarm count.	See "Chassis Alarm Messages on a QFX3500 Device" on page 12.
	Fan tray LED is blinking amber.	See Fan Tray LED on a QFX3500 Device.
	Chassis status LED for the power is blinking amber.	See Chassis Status LEDs on a QFX3500 Device.
	Chassis status LED for the fan (on the management board) is blinking amber.	Replace the management board as soon as possible. See Chassis Status LEDs on a QFX3500 Device.
Port configuration	Cannot configure a port as a Gigabit Ethernet port.	Check whether the port is a valid Gigabit Ethernet port (6 through 41). See QFX3500 Device Overview.
	Cannot configure a port as a Fibre Channel port.	Check whether the port is a valid Fibre Channel port (0 through 5 and 42 through 47). See QFX3500 Device Overview.

Table 2: Troubleshooting on the QFX Series *(Continued)*

Problem Category	Symptom or Problem	Recommended Action
	Cannot configure a port as a 10-Gigabit Ethernet port.	If the port is not a 40-Gbps QSFP+ interface, check whether the port is in the range of 0 through 5 or 42 through 47. If one of the ports in that block (0 through 5 or 42 through 47) is configured as a Fibre Channel port, then all ports in that block must also be configured as Fibre Channel ports. If the port is a 40-Gbps QSFP+ interface. make sure the configuration does not exceed the interface limit. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces, but because port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces. See QFX3500 Device Overview.
	Cannot configure a 40-Gbps QSFP+ interface.	The 40-Gbps QSFP+ interfaces can only be used as 10-Gigabit Ethernet interfaces. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces using a breakout cable. However, port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces. See QFX3500 Device Overview.
External devices (USB devices)	Upgrading software from a USB device results in an upgrade failure, and the system enters an invalid state.	Unplug the USB device and reboot the switch.

Table 2: Troubleshooting on the QFX Series *(Continued)*

Problem Category	Symptom or Problem	Recommended Action
Initial device configuration	Cannot configure management Ethernet ports.	Configure the management ports from the console port. You cannot configure the management ports by directly connecting to them. **NOTE**: The management ports are on the front panel of the QFX3500 switch. They are labeled **C0** and **C1** on the front panel. In the CLI they are referred to as **me0** and **me1**. See Configuring a QFX3500 Device as a Standalone Switch.
Software upgrade and configuration	Failed software upgrade.	See *Recovering from a Failed Software Installation*.
	Active partition becomes inactive after upgrade.	
	Problem with the active configuration file.	See the following topics: • Loading a Previous Configuration File • *Reverting to the Default Factory Configuration* • "Reverting to the Rescue Configuration" on page 47 • *Performing a Recovery Installation*
	Root password is lost or forgotten.	Recover the root password. See "Recovering the Root Password for Switches" on page 87.

Table 2: Troubleshooting on the QFX Series *(Continued)*

Problem Category	Symptom or Problem	Recommended Action
Network interfaces	An aggregated Ethernet interface is down.	See "Troubleshooting an Aggregated Ethernet Interface" on page 65.
	Interface on built-in network port is down.	See "Troubleshooting Network Interfaces" on page 66.
	Interface on port in which SFP or SFP + transceiver is installed in an SFP+ uplink module is down.	
Ethernet switching	A MAC address entry in the Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch.	See *Troubleshooting Ethernet Switching*.
Firewall filter	Firewall configuration exceeded available Ternary Content Addressable Memory (TCAM) space.	See Troubleshooting Firewall Filters.

CHAPTER 2

Alarms

IN THIS CHAPTER

- Alarm Types and Severity Levels | 10
- Chassis Alarm Messages on a QFX3500 Device | 12
- Interface Alarm Messages | 16
- Alarms | 16

Alarm Types and Severity Levels

The QFX Series switches support different alarm types and severity levels. Table 3 on page 10 provides a list of alarm terms and definitions that may help you in monitoring the device.

Table 3: Alarm Terms and Definitions

Term	Definition
Alarm	Signal that alerts you to conditions that might prevent normal operation. On the device, alarm indicators might include an LCD panel and LEDs on the device. The LCD panel (if present on the device) displays the chassis alarm message count. Blinking amber or yellow LEDs indicate yellow alarm conditions for chassis components.
Alarm condition	Failure event that triggers an alarm.

Table 3: Alarm Terms and Definitions *(Continued)*

Term	Definition

Alarm severity levels	Seriousness of the alarm. The level of severity can be either major (red) or minor (yellow).
	• Major (red)—Indicates a critical situation on the device that has resulted from one of the following conditions. A red alarm condition requires immediate action.
	  • One or more hardware components have failed.
	  • One or more hardware components have exceeded temperature thresholds.
	  • An alarm condition configured on an interface has triggered a critical warning.
	• Minor (yellow or amber)—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or a degradation in performance. A yellow alarm condition requires monitoring or maintenance. For example, a missing rescue configuration generates a yellow system alarm.

Alarm types	Alarms include the following types: • Chassis alarm—Predefined alarm triggered by a physical condition on the device such as a power supply failure or excessive component temperature. • Interface alarm—Alarm that you configure to alert you when an interface link is down. Applies to **ethernet**, **fibre-channel**, and **management-ethernet** interfaces. You can configure a red (major) or yellow (minor) alarm for the link-down condition, or have the condition ignored. • System alarm—Predefined alarm that might be triggered by a missing rescue configuration, failure to install a license for a licensed software feature, or high disk usage.

Chassis Alarm Messages on a QFX3500 Device

Chassis alarms indicate a failure on the device or one of its components. Chassis alarms are preset and cannot be modified.

The chassis alarm message count is displayed on the LCD panel on the front of the device. To view the chassis alarm message text remotely, use the **show chassis lcd** CLI command.

Chassis alarms on QFX3500 devices have two severity levels:

- Major (red)—Indicates a critical situation on the device that has resulted from one of the conditions described in Table 4 on page 12. A red alarm condition requires immediate action.
- Minor (yellow or amber)—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.

Table 4 on page 12 describes the chassis alarm messages on QFX3500 devices.

Table 4: QFX3500 Chassis Alarm Messages

Component	Alarm Type	CLI Message	Recommended Action
Fans	Major (red)	**Fan/Blower Absent**	The fan is missing. Install a fan.
		Fan Failure	Replace the fan and report the failure to customer support.
		Fan I2C Failure	Check the system log for one of the following messages and report the error message to customer support: • **CM ENV Monitor: Get fan speed failed**. • **CM ENV Monitor: Get fan speed failed *Fan-number* is NOT spinning @ correct speed**, where ***fan-number*** may be 1, 2, or 3.

Table 4: QFX3500 Chassis Alarm Messages *(Continued)*

Component	Alarm Type	CLI Message	Recommended Action
		***fan-number* Not Spinning** Fan	Remove and check the fan for obstructions, and then reinsert the fan. If the problem persists, replace the fan.
Power Supplies	Major (red)	**PEM *pem-number* Airflow not matching Chassis Airflow**	The power supply airflow direction is the opposite of the chassis airflow direction. Replace the power supply with a power supply that supports the same airflow direction as the chassis.
		PEM *pem-number* I2C Failure	Check the system log for one of the following messages and report the error message to customer support: • **I2C Read failed for device *number***, where ***number*** may be from 123 to 125. • **PS *number*: Transitioning from online to offline**, where power supply (PS) ***number*** may be 1 or 2.
		PEM *pem-number* is not supported	Indicates a power supply problem, or the power supply is not supported on the device. Report the problem to customer support.

Table 4: QFX3500 Chassis Alarm Messages *(Continued)*

Component	Alarm Type	CLI Message	Recommended Action
		PEM *pem-number* Not OK	Indicates a problem with the incoming AC or outgoing DC power. Replace the power supply.
	Minor (yellow)	**PEM *pem-number* Absent**	For information only. Indicates the device was powered on with two power supplies installed, but now one is missing. The device can continue to operate with a single power supply. If you wish to remove this alarm message, reboot the device with one power supply.
		PEM *pem-number* is not powered	For information only. Check the power cord connection and reconnect it if necessary.
		PEM *pem-number* Power Supply Type Mismatch	For information only. Indicates that an AC power supply and DC power supply have been installed in the same chassis. If you wish to remove this alarm message, reboot the device with two AC power supplies or two DC power supplies.

Table 4: QFX3500 Chassis Alarm Messages *(Continued)*

Component	Alarm Type	CLI Message	Recommended Action
		PEM *pem-number* Removed	For information only. Indicates the device was powered on with two power supplies installed, but one has been removed. The device can continue to operate with a single power supply. If you wish to remove this alarm message, reboot the device with one power supply.
Temperature Sensors	Major (red)	***sensor-location* Temp Sensor Fail**	Check the system log for the following message and report it to customer support: **Temp sensor *sensor-number* failed**, where *sensor-number* may range from 1 through 10.
		***sensor-location* Temp Sensor Too Hot**	Check environmental conditions and alarms on other devices. Ensure that environmental factors (such as hot air blowing around the equipment) are not affecting the temperature sensor. If the condition persists, the device may shut down.
	Minor (yellow)	***sensor-location* Temp Sensor Too Warm**	For information only. Check environmental conditions and alarms on other devices. Ensure that environmental factors (such as hot air blowing around the equipment) are not affecting the temperature sensor.

RELATED DOCUMENTATION

- Front Panel of a QFX3500 Device
- *Configuring Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types*
- alarm

Interface Alarm Messages

Interface alarms are alarms that you configure to alert you when an interface is down.

To configure an interface link-down condition to trigger a red or yellow alarm, or to configure the link-down condition to be ignored, use the `alarm` statement at the [**edit chassis**] hierarchy level. You can specify the **ethernet**, **fibre-channel**, or **management-ethernet** interface type.

NOTE: Fibre Channel alarms are valid only on QFX3500 devices.

NOTE: When red alarms or major alarms are issued on QFX5100 or EX4600 switches, the alarm LED glows amber instead of red.

By default, major alarms are configured for interface link-down conditions on the control plane and management network interfaces in a QFabric system. The link-down alarms indicate that connectivity to the control plane network is down. You can configure these alarms to be ignored using the `alarm` statement at the [**edit chassis**] hierarchy level.

NOTE: If you configure a yellow alarm on the QFX3008-I Interconnect device, it is handled as a red alarm.

Alarms

IN THIS SECTION

- System Alarms | 17

Configuring Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types | 17
System-Wide Alarms and Alarms for Each Interface Type | 18

System Alarms

Switches provide predefined system alarms that can be triggered by a missing rescue configuration, failure to install a license for a licensed software feature, or high disk usage. You can display alarm messages by issuing the **show system alarms** operational mode command.

For example: The switch might trigger an alarm when disk usage in the **/var** partition exceeds 75 percent. A usage level between 76 and 90 percent indicates high usage and raises a minor alarm condition, whereas a usage level above 90 percent indicates that the partition is full and raises a major alarm condition.

The following sample output shows the system alarm messages that are displayed when disk usage is exceeded on the switch.

```
user@host> show system alarms

4 alarms currently active
Alarm time               Class  Description
2013-10-08 20:08:20 UTC  Minor  RE 0 /var partition usage is high
2013-10-08 20:08:20 UTC  Major  RE 0 /var partition is full
2013-10-08 20:08:08 UTC  Minor  FPC 1 /var partition usage is high
2013-10-08 20:08:08 UTC  Major  FPC 1 /var partition is full
```

BEST PRACTICE: We recommend that you regularly request a system file storage cleanup to optimize the performance of the switch and prevent generating system alarms.

Configuring Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types

For the different types of PICs, you can configure which conditions trigger alarms and whether they trigger a red or yellow alarm. Red alarm conditions light the **RED ALARM** LED and trigger an audible alarm if one is connected. Yellow alarm conditions light the **YELLOW ALARM** LED and trigger an audible alarm if one is connected.

NOTE: By default, any failure condition on the integrated-services interface (Adaptive Services PIC) triggers a red alarm.

To configure conditions that trigger alarms and that can occur on any interface of the specified type, include the **alarm** statement at the **[edit chassis]** hierarchy level.

```
[edit chassis]
alarm {
    interface-type {
        alarm-name (red | yellow | ignore);
    }
}
```

alarm-name is the name of an alarm.

System-Wide Alarms and Alarms for Each Interface Type

Table 5 on page 18 lists the system-wide alarms and the alarms for each interface type.

Table 5: Configurable PIC Alarm Conditions

Interface/System	Alarm Condition	Configuration Option
SONET/SDH and ATM	Link alarm indication signal	**ais-l**
	Path alarm indication signal	**ais-p**
	Signal degrade (SD)	**ber-sd**
	Signal fail (SF)	**ber-sf**
	Loss of cell delineation (ATM only)	**locd**
	Loss of framing	**lof**

Table 5: Configurable PIC Alarm Conditions *(Continued)*

Interface/System	Alarm Condition	Configuration Option
	Loss of light	**lol**
	Loss of pointer	**lop-p**
	Loss of signal	**los**
	Phase-locked loop out of lock	**pll**
	Synchronous transport signal (STS) payload label (C2) mismatch	**plm-p**
	Line remote failure indication	**rfi-l**
	Path remote failure indication	**rfi-p**
	STS path (C2) unequipped	**uneq-p**
E3/T3	Alarm indicator signal	**ais**
	Excessive numbers of zeros	**exz**
	Failure of the far end	**ferf**
	Idle alarm	**idle**
	Line code violation	**lcv**
	Loss of frame	**lof**

Table 5: Configurable PIC Alarm Conditions *(Continued)*

Interface/System	Alarm Condition	Configuration Option
	Loss of signal	**los**
	Phase-locked loop out of lock	**pll**
	Yellow alarm	**ylw**
Ethernet	Link has gone down	**link-down**
DS1	Alarm indicator signal	**ais**
	Yellow alarm	**ylw**
Integrated services	Hardware or software failure	**failure**
Management Ethernet	Link has gone down	**link-down**

RELATED DOCUMENTATION

- Chassis Conditions That Trigger Alarms
- *Alarm Types and Severity Levels*
- Network Management and Monitoring Guide
- *Freeing Up System Storage Space*
- show system alarms

PART 2

Routine Monitoring

Routine Monitoring Using the CLI | 22

CHAPTER 3

Routine Monitoring Using the CLI

IN THIS CHAPTER

- show snmp | 22
- Tracing SNMP Activity on a Device Running Junos OS | 24
- Monitoring RMON MIB Tables | 29
- Displaying a Log File from a Single-Chassis System | 30
- Monitoring System Log Messages | 32
- Monitoring Traffic Through the Router or Switch | 34
- Pinging Hosts | 38

show snmp

There are several commands that you can access in Junos OS operational mode to monitor SNMP information. Some of the commands are:

- **show snmp health-monitor**, which displays the health monitor log and alarm information.
- **show snmp mib**, which displays information from the MIBs, such as device and system information.
- **show snmp statistics**, which displays SNMP statistics such as the number of packets, silent drops, and invalid output values.
- **show snmp rmon**, which displays the RMON alarm, event, history, and log information

The following example provides sample output from the **show snmp health-monitor** command:

```
user@switch> show snmp health-monitor
Alarm
Index  Variable description                          Value State

 32768 Health Monitor: root file system utilization
       jnxHrStoragePercentUsed.1                        58 active
```

```
 32769 Health Monitor: /config file system utilization
       jnxHrStoragePercentUsed.2                          0 active

 32770 Health Monitor: RE 0 CPU utilization
       jnxOperatingCPU.9.1.0.0                            0 active

 32773 Health Monitor: RE 0 Memory utilization
       jnxOperatingBuffer.9.1.0.0                        35 active

 32775 Health Monitor: jkernel daemon CPU utilization
       Init daemon                                        0 active
       Chassis daemon                                    50 active
       Firewall daemon                                    0 active
       Interface daemon                                   5 active
       SNMP daemon                                       11 active
       MIB2 daemon                                       42 active
       ...
```

The following example provides sample output from the **show snmp mib** command:

```
user@switch> show snmp mib walk system

sysDescr.0    = Juniper Networks, Inc. qfx3500s internet router, kernel
JUNOS 11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC builder@abc.example.net:
/volume/build/junos/11.1/production/20100926.0/obj-xlr/bsd/sys/compile/JUNIPER-
xxxxx
Build date: 2010-09-26 06:00:10 U
sysObjectID.0 = jnxProductQFX3500
sysUpTime.0   = 24444184
sysContact.0  = J Smith
sysName.0     = Lab QFX3500
sysLocation.0 = Lab
sysServices.0 = 4
```

The following example provides sample output from the **show snmp statistics** command:

```
user@switch> show snmp statistics

SNMP statistics:
  Input:
```

```
    Packets: 0, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 0, Total set varbinds: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0, Duplicate request drops: 0
    Output:
    Packets: 0, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 0, Traps: 0
```

RELATED DOCUMENTATION

- health-monitor
- show snmp mib
- show snmp statistics

Tracing SNMP Activity on a Device Running Junos OS

IN THIS SECTION

- Configuring the Number and Size of SNMP Log Files | **25**
- Configuring Access to the Log File | **26**
- Configuring a Regular Expression for Lines to Be Logged | **26**
- Configuring the Trace Operations | **27**

SNMP tracing operations track activity for SNMP agents and record the information in log files. The logged error descriptions provide detailed information to help you solve problems faster.

By default, Junos OS does not trace any SNMP activity. If you include the **traceoptions** statement at the **[edit snmp]** hierarchy level, the default tracing behavior is:

- Important activities are logged in files located in the **/var/log** directory. Each log is named after the SNMP agent that generates it. Currently, the following log files are created in the **/var/log** directory when the **traceoptions** statement is used:
  - chassisd
  - craftd
  - ilmid
  - mib2d
  - rmopd
  - serviced
  - snmpd
- When a trace file named ***filename*** reaches its maximum size, it is renamed ***filename*.0**, then ***filename*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. (For more information about how log files are created, see the System Log Explorer.)
- Log files can be accessed only by the user who configured the tracing operation.

You cannot change the directory (**/var/log**) in which trace files are located. However, you can customize the other trace file settings by including the following statements at the **[edit snmp]** hierarchy level:

```
[edit snmp]
traceoptions {
    file <files number> <match regular-expression> <size size> <world-readable |
no-world-readable>;
    flag flag;
    memory-trace;
    no-remote-trace;
    no-default-memory-trace;
}
```

These statements are described in the following sections:

Configuring the Number and Size of SNMP Log Files

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed ***filename*.0**, then ***filename*.1**, and so on, until there are three trace files. Then the oldest trace file (***filename*.2**) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file files number size size;
```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (***filename***) reaches 2 MB, ***filename*** is renamed ***filename*.0**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0** is renamed ***filename*.1** and ***filename*** is renamed ***filename*.0**. This process repeats until there are 20 trace files. Then the oldest file (***filename*.19**) is overwritten by the newest file (***filename*.0**).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

Configuring Access to the Log File

By default, log files can be accessed only by the user who configured the tracing operation.

To specify that any user can read all log files, include the **file world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file world-readable;
```

To explicitly set the default behavior, include the **file no-world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file no-world-readable;
```

Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged activities.

You can refine the output by including the **match** statement at the **[edit snmp traceoptions file *filename*]** hierarchy level and specifying a regular expression (regex) to be matched:

```
[edit snmp traceoptions]
file filename match regular-expression;
```

Configuring the Trace Operations

By default, only important activities are logged. You can specify which trace operations are to be logged by including the following **flag** statement (with one or more tracing flags) at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
flag {
    all;
    configuration;
    database;
    events;
    general;
    interface-stats;
    nonvolatile-sets;
    pdu;
    policy;
    protocol-timeouts;
    routing-socket;
    server;
    subagent;
    timer;
    varbind-error;
}
```

Table 6 on page 27 describes the meaning of the SNMP tracing flags.

Table 6: SNMP Tracing Flags

Flag	Description	Default Setting
all	Log all operations.	Off
configuration	Log reading of the configuration at the **[edit snmp]** hierarchy level.	Off
database	Log events involving storage and retrieval in the events database.	Off

Table 6: SNMP Tracing Flags *(Continued)*

Flag	Description	Default Setting
events	Log important events.	Off
general	Log general events.	Off
interface-stats	Log physical and logical interface statistics.	Off
nonvolatile-set	Log nonvolatile SNMP set request handling.	Off
pdu	Log SNMP request and response packets.	Off
policy	Log policy processing.	Off
protocol-timeouts	Log SNMP response timeouts.	Off
routing-socket	Log routing socket calls.	Off
server	Log communication with processes that are generating events.	Off
subagent	Log subagent restarts.	Off
timer	Log internal timer events.	Off
varbind-error	Log variable binding errors.	Off

To display the end of the log for an agent, issue the **show log *agent*d | last** operational mode command:

```
[edit]
user@host# run show log agentd | last
```

where ***agent*** is the name of an SNMP agent.

Monitoring RMON MIB Tables

IN THIS SECTION

- Purpose | 29
- Action | 29
- Meaning | 30

Purpose

Monitor remote monitoring (RMON) alarm, event, and log tables.

Action

To display the RMON tables:

```
user@switch> show snmp rmon
Alarm
Index  Variable description                          Value State

     5 monitor
       jnxOperatingCPU.9.1.0.0                           5 falling threshold

Event
Index  Type                        Last Event
    1  log and trap                2010-07-10 11:34:17 PDT
Event Index: 1
    Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 100)
    Time: 2010-07-10 11:34:07 PDT
    Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 5)
    Time: 2010-07-10 11:34:17 PDT
```

Meaning

The display shows that an alarm has been defined to monitor jnxRmon MIB object jnxOperatingCPU, which represents the CPU utilization of the Routing Engine. The alarm is configured to generate an event that sends an SNMP trap and adds an entry to the logTable in the RMON MIB. The log table shows that two occurrences of the event have been generated—one for rising above a threshold of 90 percent, and one for falling below a threshold of 75 percent.

RELATED DOCUMENTATION

- Configuring RMON Alarms and Events
- show snmp rmon
- show snmp rmon history
- clear snmp statistics
- clear snmp history

Displaying a Log File from a Single-Chassis System

To display a log file stored on a single-chassis system, enter Junos OS CLI operational mode and issue the following commands:

```
user@switch> show log log-filename
user@switch> file show log-file-pathname
```

By default, the commands display the file stored on the local Routing Engine.

The following example shows the output from the **show log messages** command:

```
user@switch1> show log messages
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:06  switch1 chassism[952]: CM ENV Monitor: set fan speed is 65
percent for Fan 1
Nov  4 11:30:06  switch1 chassism[952]: CM ENV Monitor: set fan speed is 65
percent for Fan 2
Nov  4 11:30:06  switch1 chassism[952]: CM ENV Monitor: set fan speed is 65
percent for Fan 3
```

```
...
Nov  4 11:52:53  switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Management
process): new instance detected (variable: sysApplElmtRunMemory.5.6.2293)
Nov  4 11:52:53  switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Command-line
interface): new instance detected (variable: sysApplElmtRunMemory.5.8.2292)
...
Nov  4 12:08:30  switch1 rpdf[957]: task_connect: task BGP_100.10.10.1.6+179
addr 10.10.1.6+179: Can't assign requested
address
Nov  4 12:08:30  switch1 rpdf[957]: bgp_connect_start: connect 10.10.1.6
(Internal AS 100): Can't assign requested address
Nov  4 12:10:24  switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov  4 12:10:27  switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
Nov  4 12:10:31  switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages
```

The following example shows the output from the **file show** command. The file in the pathname **/var/log/processes** has been previously configured to include messages from the daemon facility.

```
user@switch1> file show /var/log/processes
Feb 22 08:58:24 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 22 20:35:07 switch1 snmpd[359]: SNMPD_THROTTLE_QUEUE_DRAINED:
trap_throttle_timer_handler: cleared all throttled traps
Feb 23 07:34:56 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 23 07:38:19 switch1 snmpd[359]: SNMPD_TRAP_COLD_START: trap_generate_cold:
SNMP trap: cold start
...
```

RELATED DOCUMENTATION

- Interpreting Messages Generated in Standard Format

Monitoring System Log Messages

IN THIS SECTION

- Purpose | 32
- Action | 32
- Meaning | 33

Purpose

Display system log messages about the QFX Series. By looking through a system log file for any entries pertaining to the interface that you are interested in, you can further investigate a problem with an interface on the switch.

Action

To view system log messages:

```
user@switch1> show log messages
```

Sample Output

command-name

```
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:06  switch1 chassism[952]: CM ENV Monitor: set fan speed is 65
percent for Fan 1
Nov  4 11:30:06  switch1 chassism[952]: CM ENV Monitor: set fan speed is 65
percent for Fan 2
Nov  4 11:30:06  switch1 chassism[952]: CM ENV Monitor: set fan speed is 65
percent for Fan 3
...
Nov  4 11:52:53  switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
```

```
jroute daemon
memory usage (Management process): new instance detected (variable:
sysApplElmtRunMemory.5.6.2293)
Nov  4 11:52:53  switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon
memory usage (Command-line interface): new instance detected (variable:
sysApplElmtRunMemory.5.8.2292)
...
Nov  4 12:10:24  switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov  4 12:10:27  switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
Nov  4 12:10:31  switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages
```

Meaning

The sample output shows the following entries in the **messages** file:

- A new log file was created when the previous file reached the maximum size of

  128 kilobytes (KB).

- The fan speed for Fan 1, 2, and 3 is set at 65 percent.
- Health monitoring activity is detected.
- CLI commands were entered by the user jsmith.

RELATED DOCUMENTATION

> *clear log*
>
> *show log*
>
> syslog

Monitoring Traffic Through the Router or Switch

IN THIS SECTION

- Displaying Real-Time Statistics About All Interfaces on the Router or Switch | 34
- Displaying Real-Time Statistics About an Interface on the Router or Switch | 36

To help with the diagnosis of a problem, display real-time statistics about the traffic passing through physical interfaces on the router or switch.

To display real-time statistics about physical interfaces, perform these tasks:

Displaying Real-Time Statistics About All Interfaces on the Router or Switch

IN THIS SECTION

- Purpose | 34
- Action | 34
- Meaning | 35

Purpose

Display real-time statistics about traffic passing through all interfaces on the router or switch.

Action

To display real-time statistics about traffic passing through all interfaces on the router or switch:

```
user@host> monitor interface traffic
```

Sample Output

command-name

```
user@host> monitor interface traffic
host name                  Seconds: 15                  Time: 12:31:09
Interface    Link  Input packets         (pps)     Output packets         (pps)
 so-1/0/0    Down              0           (0)                 0           (0)
 so-1/1/0    Down              0           (0)                 0           (0)
 so-1/1/1    Down              0           (0)                 0           (0)
 so-1/1/2    Down              0           (0)                 0           (0)
 so-1/1/3    Down              0           (0)                 0           (0)
 t3-1/2/0    Down              0           (0)                 0           (0)
 t3-1/2/1    Down              0           (0)                 0           (0)
 t3-1/2/2    Down              0           (0)                 0           (0)
 t3-1/2/3    Down              0           (0)                 0           (0)
 so-2/0/0      Up         211035            (1)             36778           (0)
 so-2/0/1      Up         192753            (1)             36782           (0)
 so-2/0/2      Up         211020            (1)             36779           (0)
 so-2/0/3      Up         211029            (1)             36776           (0)
 so-2/1/0      Up         189378            (1)             36349           (0)
 so-2/1/1    Down              0           (0)             18747           (0)
 so-2/1/2    Down              0           (0)             16078           (0)
 so-2/1/3      Up              0           (0)             80338           (0)
 at-2/3/0      Up              0           (0)                 0           (0)
 at-2/3/1    Down              0           (0)                 0           (0)
Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

Meaning

The sample output displays traffic data for active interfaces and the amount that each field has changed since the command started or since the counters were cleared by using the **C** key. In this example, the **monitor interface** command has been running for 15 seconds since the command was issued or since the counters last returned to zero.

Displaying Real-Time Statistics About an Interface on the Router or Switch

IN THIS SECTION

- Purpose | 36
- Action | 36
- Meaning | 37

Purpose

Display real-time statistics about traffic passing through an interface on the router or switch.

Action

To display traffic passing through an interface on the router or switch, use the following Junos OS CLI operational mode command:

```
user@host> monitor interface  interface-name
```

Sample Output

command-name

```
user@host> monitor interface so-0/0/1
Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
R1
Interface: so-0/0/1, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: OC3 Traffic statistics:
  Input bytes:                  5856541 (88 bps)
  Output bytes:                 6271468 (96 bps)
  Input packets:                 157629 (0 pps)
  Output packets:                157024 (0 pps)
Encapsulation statistics:
  Input keepalives:               42353
  Output keepalives:              42320
  LCP state: Opened
```

```
Error statistics:
  Input errors:                          0
  Input drops:                           0
  Input framing errors:                  0
  Input runts:                           0
  Input giants:                          0
  Policed discards:                      0
  L3 incompletes:                        0
  L2 channel errors:                     0
  L2 mismatch timeouts:                  0
  Carrier transitions:                   1
  Output errors:                         0
  Output drops:                          0
  Aged packets:                          0
Active alarms : None
Active defects: None
SONET error counts/seconds:
  LOS count                              1
  LOF count                              1
  SEF count                              1
  ES-S                                  77
  SES-S                                 77
SONET statistics:
  BIP-B1                                 0
  BIP-B2                                 0
  REI-L                                  0
  BIP-B3                                 0
  REI-P                                  0
Received SONET overhead:  F1        : 0x00  J0        : 0xZ
```

Meaning

The sample output shows the input and output packets for a particular SONET interface (**so-0/0/1**). The information can include common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors. For more information, see *Checklist for Tracking Error Conditions*.

To control the output of the command while it is running, use the keys shown in Table 7 on page 38.

Table 7: Output Control Keys for the monitor interface Command

Action	Key
Display information about the next interface. The **monitor interface** command scrolls through the physical or logical interfaces in the same order that they are displayed by the **show interfaces terse** command.	**N**
Display information about a different interface. The command prompts you for the name of a specific interface.	**I**
Freeze the display, halting the display of updated statistics.	**F**
Thaw the display, resuming the display of updated statistics.	**T**
Clear (zero) the current delta counters since **monitor interface** was started. It does not clear the accumulative counter.	**C**
Stop the **monitor interface** command.	**Q**

See the CLI Explorer for details on using match conditions with the **monitor traffic** command.

Pinging Hosts

IN THIS SECTION

- Purpose | 39
- Action | 39
- Meaning | 39

Purpose

Use the CLI **ping** command to verify that a host can be reached over the network. This command is useful for diagnosing host and network connectivity problems. The switch sends a series of Internet Control Message Protocol (ICMP) echo (ping) requests to a specified host and receives ICMP echo responses.

Action

To use the **ping** command to send four requests (ping count) to host3:

ping *host* count *number*

Sample Output

command-name

ping host3 count 4

```
user@switch> ping host3 count 4
PING host3.site.net (192.0.2.111): 56 data bytes
64 bytes from 192.0.2.111: icmp_seq=0 ttl=122 time=0.661 ms
64 bytes from 192.0.2.111: icmp_seq=1 ttl=122 time=0.619 ms
64 bytes from 192.0.2.111: icmp_seq=2 ttl=122 time=0.621 ms
64 bytes from 192.0.2.111: icmp_seq=3 ttl=122 time=0.634 ms

--- host3.site.net ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.619/0.634/0.661/0.017 ms
```

Meaning

- The **ping** results show the following information:
  - Size of the ping response packet (in bytes).
  - IP address of the host from which the response was sent.
  - Sequence number of the ping response packet. You can use this value to match the ping response to the corresponding ping request.

- Time-to-live (ttl) hop-count value of the ping response packet.
- Total time between the sending of the ping request packet and the receiving of the ping response packet, in milliseconds. This value is also called round-trip time.
- Number of ping requests (probes) sent to the host.
- Number of ping responses received from the host.
- Packet loss percentage.
- Round-trip time statistics: minimum, average, maximum, and standard deviation of the round-trip time.

PART 3

Troubleshooting Features

Configuration and File Management | 42

Ethernet Switching | 50

Hardware | 60

High Availability | 63

Interfaces | 65

Junos OS Basics | 83

Layer 3 Protocols | 107

Network Management | 109

Security | 128

Services | 147

Traffic Management | 150

Virtual Chassis Fabric | 165

CHAPTER 4

Configuration and File Management

IN THIS CHAPTER

- Returning to a Previously Committed Configuration | 42
- Reverting to the Default Factory Configuration | 46
- Reverting to the Rescue Configuration | 47
- Freeing Up System Storage Space | 48

Returning to a Previously Committed Configuration

IN THIS SECTION

- Returning to a Configuration Prior to the One Most Recently Committed | 42
- Displaying Previous Configurations | 43
- Comparing Configuration Changes with a Prior Version | 44

This topic explains how you can return to a configuration prior to the most recently committed one.

Returning to a Configuration Prior to the One Most Recently Committed

To return to a configuration prior to the most recently committed one, include the configuration number, 0 through 49, in the **rollback** command. The most recently saved configuration is number 0 (which is the default configuration to which the system returns), and the oldest saved configuration is number 49.

```
[edit]
user@host# rollback number
load complete
```

Displaying Previous Configurations

To display previous configurations, including the rollback number, date, time, the name of the user who committed changes, and the method of commit, use the **rollback ?** command.

```
[edit]
user@host# rollback ?
Possible completions:
<[Enter]> Execute this command
<number> Numeric argument
0                    2018-02-27 12:52:10 PST by abc via cli
1                    2018-02-26 14:47:42 PST by def via cli
2                    2018-02-14 21:55:45 PST by ghi via cli
3                    2018-02-10 16:11:30 PST by jkl via cli
4                    2018-02-10 16:02:35 PST by mno via cli
5                    2018-03-16 15:10:41 PST by pqr via cli
6                   2018-03-16 14:54:21 PST by stu via cli
7                    2018-03-16 14:51:38 PST by vwx via cli
8                     2018-03-16 14:43:29 PST by yzz via cli
9                    2018-03-16 14:15:37 PST by abc via cli
  10                   2018-03-16 14:13:57 PST by def via cli
  11                    2018-03-16 12:57:19 PST by root via other
  12                     2018-03-16 10:45:23 PST by root via other
  13                    2018-03-16 10:08:13 PST by root via other
  14                   2018-03-16 01:20:56 PST by root via other
  15                   2018-03-16 00:40:37 PST by ghi via cli
  16                    2018-03-16 00:39:29 PST by jkl via cli
  17                    2018-03-16 00:32:36 PST by mno via cli
  18                     2018-03-16 00:31:17 PST by pqr via cli
  19                    2018-03-15 19:59:00 PST by stu via cli
  20                    2018-03-15 19:53:39 PST by vwx via cli
  21                    2018-03-15 18:07:19 PST by yzz via cli
  22                      2018-03-15 17:59:03 PST by abc via cli
  23                     2018-03-15 15:05:14 PST by def via cli
  24                    2018-03-15 15:04:51 PST by ghi via cli
  25                   2018-03-15 15:03:42 PST by jkl via cli
  26                    2018-03-15 15:01:52 PST by mno via cli
  27                   2018-03-15 14:58:34 PST by pqr via cli
  28                   2018-03-15 13:09:37 PST by root via other
  29                    2018-03-12 11:01:20 PST by stu via cli
  30                     2018-03-12 10:57:35 PST by vwx via cli
  31                    2018-03-11 10:25:07 PST by yzz via cli
```

```
  32                    2018-03-10 23:40:58 PST by abc via cli
  33                    2018-03-10 23:40:38 PST by def via cli
  34                    2018-03-10 23:14:27 PST by ghi via cli
  35                    2018-03-10 23:10:16 PST by jkl via cli
  36                   2018-03-10 23:01:51 PST by mno via cli
  37                   2018-03-10 22:49:57 PST by pqr via cli
  38                    2018-03-10 22:24:07 PST by stu via cli
  39                   2018-03-10 22:20:14 PST by vwx via cli
  40                    2018-03-10 22:16:56 PST by yzz via cli
  41                   2018-03-10 22:16:41 PST by abc via cli
  42                   2018-03-10 20:44:00 PST by def via cli
  43                   2018-03-10 20:43:29 PST by ghi via cli
  44                    2018-03-10 20:39:14 PST by jkl via cli
  45                   2018-03-10 20:31:30 PST by root via other
  46                    2018-03-10 18:57:01 PST by mno via cli
  47                    2018-03-10 18:56:18 PST by pqr via cli
  48                    2018-03-10 18:47:49 PST by stu via cli
  49                   2018-03-10 18:47:34 PST by vw via cli
 | Pipe through a command
[edit]
```

Comparing Configuration Changes with a Prior Version

In configuration mode only, when you have made changes to the configuration and want to compare the candidate configuration with a prior version, you can use the **compare** command to display the configuration. The **compare** command compares the candidate configuration with either the current committed configuration or a configuration file and displays the differences between the two configurations. To compare configurations, specify the **compare** command after the pipe:

```
[edit]
user@host# show | compare (filename| rollback n)
```

- ***filename*** is the full path to a configuration file. The file must be in the proper format: a hierarchy of statements.
- ***n*** is the index into the list of previously committed configurations. The most recently saved configuration is number 0, and the oldest saved configuration is number 49. If you do not specify arguments, the candidate configuration is compared against the active configuration file (**/config/juniper.conf**).

The comparison output uses the following conventions:

- Statements that are only in the candidate configuration are prefixed with a plus sign (+).

- Statements that are only in the comparison file are prefixed with a minus sign (-).
- Statements that are unchanged are prefixed with a single blank space ( ).

The following example shows various changes, then a comparison of the candidate configuration with the active configuration, showing only the changes made at the **[edit protocols bgp]** hierarchy level:

```
[edit]
user@host# edit protocols bgp
[edit protocols bgp]
user@host# show
group my-group {
    type internal;
    hold-time 60;
    advertise-inactive;
    allow 10.1.1.1/8;
}
group fred {
    type external;
    peer-as 33333;
    allow 10.2.2.2/8;
}
group test-peers {
    type external;
    allow 10.3.3.3/8;
}
[edit protocols bgp]
user@host# set group my-group hold-time 90
[edit protocols bgp]
user@host# delete group my-group advertise-inactive
[edit protocols bgp]
user@host# set group fred advertise-inactive
[edit protocols bgp]
user@host# delete group test-peers
[edit protocols bgp]
user@host# show | compare
[edit protocols bgp group my-group]
-hold-time 60;
+hold-time 90;
-advertise-inactive;
[edit protocols bgp group fred]
+advertise-inactive;
```

```
[edit protocols bgp]
-group test-peers {
    -type external;
    -allow 10.3.3.3/8;
}
[edit protocols bgp]
user@host# show
group my-group {
    type internal;
    hold-time 90;
    allow 10.1.1.1/8;
}
group fred {
    type external;
    advertise-inactive;
    peer-as 3333;
    allow 10.2.2.2/8;
}
```

RELATED DOCUMENTATION

Loading a Configuration from a File or the Terminal

Viewing Files and Directories on a Juniper Networks Device

Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
```

```
[edit]
user@switch# commit
```

NOTE: This process clears prior committed configuration parameters, except for those which preserve a Virtual Chassis configuration. This is how you can restore the factory default configuration on a Virtual Chassis (multiple devices configured to work together that look like a single device) without removing anything needed to keep the Virtual Chassis working.

RELATED DOCUMENTATION

Understanding Configuration Files

Reverting to the Rescue Configuration

Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a device and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```

2. Commit your changes.

```
[edit]
user@switch# commit filename
```

RELATED DOCUMENTATION

Reverting to the Default Factory Configuration

Freeing Up System Storage Space

IN THIS SECTION

- Problem | 48
- Solution | 48

Problem

Description

The system file storage space on the device is full. Rebooting the switch does not solve the problem.

The following error message is displayed during a typical operation on the device after the file storage space is full.

```
user@host% cli
user@host> configure
/var: write failed, filesystem is full
```

Solution

Clean up the file storage on the device by deleting system files.

1. Request to delete system files.

```
user@host> request system storage cleanup
```

The list of files to be deleted is displayed.

```
List of files to delete:

         Size Date         Name
    11B Jul 26 20:55 /var/jail/tmp/alarmd.ts
   124B Aug  4 18:05 /var/log/default-log-messages.0.gz
  1301B Jul 26 20:42 /var/log/install.0.gz
```

```
   387B Jun  3 14:37 /var/log/install.1.gz
  4920B Aug  4 18:05 /var/log/messages.0.gz
  20.0K Jul 26 21:00 /var/log/messages.1.gz
  16.3K Jun 25 13:45 /var/log/messages.2.gz
   804B Aug  4 18:05 /var/log/security.0.gz
  16.8K Aug  3 11:15 /var/log/security.1.gz
   487B Aug  4 18:04 /var/log/wtmp.0.gz
   855B Jul 29 22:54 /var/log/wtmp.1.gz
   920B Jun 30 16:32 /var/log/wtmp.2.gz
    94B Jun  3 14:36 /var/log/wtmp.3.gz
 353.2K Jun  3 14:37 /var/sw/pkg/jloader-qfx-11.2I20110303_1117_dc-builder.tgz
 124.0K Jun  3 14:30 /var/tmp/gres-tp/env.dat
     0B Apr 14 16:20 /var/tmp/gres-tp/lock
     0B Apr 14 17:37 /var/tmp/if-rtsdb/env.lck
  12.0K Jul 26 20:55 /var/tmp/if-rtsdb/env.mem
2688.0K Jul 26 20:55 /var/tmp/if-rtsdb/shm_usr1.mem
 132.0K Jul 26 20:55 /var/tmp/if-rtsdb/shm_usr2.mem
2048.0K Jul 26 20:55 /var/tmp/if-rtsdb/trace.mem
   155B Jul 26 20:55 /var/tmp/krt_gencfg_filter.txt
     0B Jul 26 20:55 /var/tmp/rtsdb/if-rtsdb
1400.6K Aug  3 10:13 /var/tmp/sfid.core.0.gz
1398.9K Aug  3 17:01 /var/tmp/sfid.core.1.gz
Delete these files ? [yes,no] (no)
```

2. Enter **yes** to delete the files.

3. Reboot the device.

BEST PRACTICE: We recommend that you regularly request a system file storage cleanup to optimize the performance of the device.

RELATED DOCUMENTATION

request system storage cleanup

CHAPTER 5

Ethernet Switching

IN THIS CHAPTER

- Troubleshooting Ethernet Switching | 50
- Troubleshooting Layer 2 Protocol Tunneling | 51
- Troubleshooting Private VLANs on QFX Switches | 53
- Troubleshooting Q-in-Q and VLAN Translation Configuration | 58

Troubleshooting Ethernet Switching

IN THIS SECTION

- Problem | 50
- Solution | 51

Problem

Description

Sometimes a MAC address entry in the switch's Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch. Typically, the switch does not wait for a MAC address expiration when a MAC move operation occurs. As soon as the switch detects the MAC address on the new interface, it immediately updates the table. Many network devices send a gratuitous ARP packet when switching an IP address from one device to another. The switch updates its ARP cache table after receipt of such gratuitous ARP messages, and then it also updates its Ethernet switching table.

Sometimes silent devices, such as syslog servers or SNMP trap receivers that receive UDP traffic but do not return acknowledgment (ACK) messages to the traffic source, fail to send gratuitous ARP packets

when a device moves. If such a move occurs when the system administrator is not available to explicitly clear the affected interfaces by issuing the **clear ethernet-switching table** command, the entry for the moved device in the Ethernet switching table is not updated.

Solution

Set up the switch to handle unattended MAC address switchovers.

1. Reduce the system-wide ARP aging timer. (By default, the ARP aging timer is set at 20 minutes. The range of the ARP aging timer is from 1 through 240 minutes.)

```
[edit system arp]
user@switch# set aging-timer 3
```

2. Set the MAC aging timer to the same value as the ARP timer. (By default, the MAC aging timer is set to 300 seconds. The range is 60 to 1,000,000 seconds.)

```
[edit protocols l2-learning]
user@switch# set global-mac-table-aging-time 180
```

The ARP entry and the MAC address entry for the moved device expire within the times specified by the aging timer values. After the entries expire, the switch sends a new ARP message to the IP address of the device. The device responds to the ARP message, thereby refreshing the entries in the switch's ARP cache table and Ethernet switching table.

RELATED DOCUMENTATION

arp

global-mac-table-aging-time

Troubleshooting Layer 2 Protocol Tunneling

IN THIS SECTION

Drop Threshold Statistics Might Be Incorrect | 52

- Egress Filtering of L2PT Traffic Not Supported | 52

Drop Threshold Statistics Might Be Incorrect

IN THIS SECTION

- Problem | 52
- Solution | 52

Problem

Description

L2PT processing is done by the CPU, and L2PT traffic to the CPU is rate limited to a maximum of 1000 pps. If traffic is received at a rate faster than this limit, the rate limit causes the traffic to be dropped before it hits the threshold and the dropped packets will not be reported in L2PT statistics. This can also occur if you configure a drop threshold that is less than 1000 pps but traffic is received at a faster rate. For example, if you configure a drop threshold of 900 pps and the VLAN receives traffic at rate of 1100 pps, L2PT statistics will show that 100 packets were dropped. The 100 packets dropped because of the rate limit are not reported. Similarly, if you do not configure a drop threshold and the VLAN receives traffic at rate of 1100 pps, the 100 packets dropped because of the rate limit are not reported.

Solution

This is expected behavior.

Egress Filtering of L2PT Traffic Not Supported

IN THIS SECTION

- Problem | 53
- Solution | 53

Problem

Description

Egress filtering of L2PT traffic is not supported on the QFX3500 switch. That is, if you configure L2PT to tunnel a protocol on an interface, you cannot also use a firewall filter to filter traffic for that protocol on that interface in the output direction. If you commit a configuration for this purpose, the firewall filter is not applied to the L2PT-tunneled traffic.

Solution

This is expected behavior.

RELATED DOCUMENTATION

- Understanding Layer 2 Protocol Tunneling

Troubleshooting Private VLANs on QFX Switches

IN THIS SECTION

- Limitations of Private VLANs | 53
- Forwarding with Private VLANs | 54
- Egress Firewall Filters with Private VLANs | 55
- Egress Port Mirroring with Private VLANs | 57

Use the following information to troubleshoot a private VLAN configuration.

Limitations of Private VLANs

The following constraints apply to private VLAN configurations:

- IGMP snooping is not supported with private VLANs.
- Routed VLAN interfaces are not supported on private VLANs
- Routing between secondary VLANs in the same primary VLAN is not supported.

- If you want to change a primary VLAN to be a secondary VLAN, you must first change it to a normal VLAN and commit the change. For example, you would follow this procedure:

  1. Change the primary VLAN to be a normal VLAN.
  2. Commit the configuration.
  3. Change the normal VLAN to be a secondary VLAN.
  4. Commit the configuration.

  Follow the same sequence of commits if you want to change a secondary VLAN to be a primary VLAN. That is, make the secondary VLAN a normal VLAN and commit that change and then change the normal VLAN to be a primary VLAN.

Forwarding with Private VLANs

IN THIS SECTION

- Problem | 54
- Solution | 55

Problem

Description

- When isolated VLAN or community VLAN tagged traffic is received on a PVLAN trunk port, MAC addresses are learned from the primary VLAN. This means that output from the *show ethernet-switching table* command shows that MAC addresses are learned from the primary VLAN and replicated to secondary VLANs. This behavior has no effect on forwarding decisions.
- If a packet with a secondary VLAN tag is received on a promiscuous port, it is accepted and forwarded.
- If a packet is received on a PVLAN trunk port and meets both of the conditions listed below, it is dropped.
  - The packet has a community VLAN tag.
  - The packet is destined to a unicast MAC address or multicast group MAC address that was learned on an isolated VLAN.

- If a packet is received on a PVLAN trunk port and meets both of the conditions listed below, it is dropped.
  - The packet has an isolated VLAN tag.
  - The packet is destined to a unicast MAC address or multicast group MAC address that was learned on a community VLAN.
- If a packet with a primary VLAN tag is received by a secondary (isolated or community) VLAN port, the secondary port forwards the packet.
- If you configure a community VLAN on one device and configure another community VLAN on a second device and both community VLANs use the same VLAN ID, traffic for one of the VLANs can be forwarded to the other VLAN. For example, assume the following configuration:
  - Community VLAN comm1 on switch 1 has VLAN ID 50 and is a member of primary VLAN pvlan100.
  - Community VLAN comm2 on switch 2 also has VLAN ID 50 and is a member of primary VLAN pvlan200.
  - Primary VLAN pvlan100 exists on both switches.

  If traffic for comm1 is sent from switch 1 to switch 2, it will be sent to the ports participating in comm2. (The traffic will also be forwarded to the ports in comm1, as you would expect.)

Solution

These are expected behaviors.

Egress Firewall Filters with Private VLANs

IN THIS SECTION

Problem | 56
Solution | 56

Problem

Description

If you apply a firewall filter in the output direction to a primary VLAN, the filter also applies to the secondary VLANs that are members of the primary VLAN when the traffic egresses with the primary VLAN tag or isolated VLAN tag, as listed below:

- Traffic forwarded from a secondary VLAN trunk port to a promiscuous port (trunk or access)
- Traffic forwarded from a secondary VLAN trunk port that carries an isolated VLAN to a PVLAN trunk port.
- Traffic forwarded from a promiscuous port (trunk or access) to a secondary VLAN trunk port
- Traffic forwarded from a PVLAN trunk port. to a secondary VLAN trunk port
- Traffic forwarded from a community port to a promiscuous port (trunk or access)

If you apply a firewall filter in the output direction to a primary VLAN, the filter does *not* apply to traffic that egresses with a community VLAN tag, as listed below:

- Traffic forwarded from a community trunk port to a PVLAN trunk port
- Traffic forwarded from a secondary VLAN trunk port that carries a community VLAN to a PVLAN trunk port
- Traffic forwarded from a promiscuous port (trunk or access) to a community trunk port
- Traffic forwarded from a PVLAN trunk port. to a community trunk port

If you apply a firewall filter in the output direction to a community VLAN, the following behaviors apply:

- The filter is applied to traffic forwarded from a promiscuous port (trunk or access) to a community trunk port (because the traffic egresses with the community VLAN tag).
- The filter is applied to traffic forwarded from a community port to a PVLAN trunk port (because the traffic egresses with the community VLAN tag).
- The filter is *not* applied to traffic forwarded from a community port to a promiscuous port (because the traffic egresses with the primary VLAN tag or untagged).

Solution

These are expected behaviors. They occur only if you apply a firewall filter to a private VLAN in the output direction and do not occur if you apply a firewall filter to a private VLAN in the input direction.

Egress Port Mirroring with Private VLANs

IN THIS SECTION

Problem | 57
Solution | 57

Problem

Description

If you create a port-mirroring configuration that mirrors private VLAN (PVLAN) traffic on egress, the mirrored traffic (the traffic that is sent to the analyzer system) has the VLAN tag of the ingress VLAN instead of the egress VLAN. For example, assume the following PVLAN configuration:

- Promiscuous trunk port that carries primary VLANs pvlan100 and pvlan400.
- Isolated access port that carries secondary VLAN isolated200. This VLAN is a member of primary VLAN pvlan100.
- Community port that carries secondary VLAN comm300. This VLAN is also a member of primary VLAN pvlan100.
- Output interface (monitor interface) that connects to the analyzer system. This interface forwards the mirrored traffic to the analyzer.

If a packet for pvlan100 enters on the promiscuous trunk port and exits on the isolated access port, the original packet is untagged on egress because it is exiting on an access port. However, the mirror copy retains the tag for pvlan100 when it is sent to the analyzer.

Here is another example: If a packet for comm300 ingresses on the community port and egresses on the promiscuous trunk port, the original packet carries the tag for pvlan100 on egress, as expected. However, the mirrored copy retains the tag for comm300 when it is sent to the analyzer.

Solution

This is expected behavior.

Troubleshooting Q-in-Q and VLAN Translation Configuration

IN THIS SECTION

- Firewall Filter Match Condition Not Working with Q-in-Q Tunneling | 58
- Egress Port Mirroring with VLAN Translation | 59

Firewall Filter Match Condition Not Working with Q-in-Q Tunneling

IN THIS SECTION

- Problem | 58
- Solution | 58

Problem

Description

If you create a firewall filter that includes a match condition of **dot1q-tag** or **dot1q-user-priority** and apply the filter on input to a trunk port that participates in a service VLAN, the match condition does not work if the Q-in-Q EtherType is not 0x8100. (When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network and therefore participate in service VLANs.)

Solution

This is expected behavior. To set the Q-in-Q EtherType to 0x8100, enter the **set dot1q-tunneling ethertype 0x8100** statement at the **[edit ethernet-switching-options]** hierarchy level. You must also configure the other end of the link to use the same Ethertype.

Egress Port Mirroring with VLAN Translation

IN THIS SECTION

- Problem | 59
- Solution | 59

Problem

Description

If you create a port-mirroring configuration that mirrors customer VLAN (CVLAN) traffic on egress and the traffic undergoes VLAN translation before being mirrored, the VLAN translation does not apply to the mirrored packets. That is, the mirrored packets retain the service VLAN (SVLAN) tag that should be replaced by the CVLAN tag on egress. The original packets are unaffected—on these packets VLAN translation works properly, and the SVLAN tag is replaced with the CVLAN tag on egress.

Solution

This is expected behavior.

SEE ALSO

Understanding Q-in-Q Tunneling and VLAN Translation

RELATED DOCUMENTATION

Understanding Q-in-Q Tunneling and VLAN Translation

Example: Setting Up Q-in-Q Tunneling on QFX Series Switches

CHAPTER 6

Hardware

IN THIS CHAPTER

- Troubleshooting QFX3100 Director Device Isolation | 60

Troubleshooting QFX3100 Director Device Isolation

IN THIS SECTION

- Problem | 60
- Resolution | 61

Problem

Description

Both connections between the QFX3100 Director devices are broken so that one of the Director devices in a Director group becomes isolated from the group.

The redundant patch cables interconnecting the Director devices are critical links required for the operation of the Director group. The two inter-Director device links must remain connected when the Director devices are online. After the Director devices are installed and the Director group is active, if a single inter-Director device link loses and regains its connection, the operation of the Director group remains intact. However, the loss of both inter-Director device links causes one Director device to isolate itself from the Director group.

WARNING: Do not reconnect the inter-Director patch cables before properly restarting the isolated Director device. Restarting the active Director device instead of the

isolated Director device can result in both Director devices rebooting, with a subsequent data loss.

Environment

This problem occurs between the two QFX3100 Director devices found in QFabric systems.

Symptoms

Symptoms of this problem include an unscheduled rebooting of one of the Director devices.

Resolution

Determine Which Director Device Is Isolated

Before restoring the inter-Director device links, determine which one of the Director devices is in isolation.

To locate an isolated Director device, use one of the following methods:

- Review logs or management tools for standard SNMP traps issued from the Director group before the Director device became isolated.
  - If eth-2/6 links are down, the Director group cannot communicate. Normally, one of the devices reboots.
  - If both eth-2/6 and eth-7/8/9 links are down, the Director device is isolated from the control plane and is not providing fabric services.
    - Issue **show fabric session-host**.
- Use the CLI to determine the serial numbers of the active Director device.
  - Issue the **show fabric session-host** command.

```
root@qfabric>show fabric session-host
Identifier: 0281042010000013
```

  - Issue the **show fabric administration inventory director-group status | grep “dg0|dg1”** command.

root@qfabrid> **show fabric administration inventory director-group status | grep “dg0|dg1”**

```
dg0 online master  10.94.214.80   0%   13597976k  4  4 days, 22:36 hrs
dg1 online master  10.94.214.81   0%   18677380k  3  4 days, 22:25 hrs
dg0  0281042010000013  online  master
dg1  0281042010000018  online  backup
```

When the Director devices cannot communicate, the **show fabric administration inventory director-group** command only displays the Director device that is online.

Power Off the Isolated Director Device and Restore the Inter-Director Device Links

CAUTION: Be sure you know which Director device is active and which is isolated. If you power off the active Director device, both Director devices reboot and cause potential data loss on the system.

To restore communication within the Director group:

1. Power off the isolated Director device.
2. Restore the inter-Director device links (port **3** to port **3**) by firmly inserting the redundant patch cables.
3. Power on the previously isolated Director device. The Director device reboots.

RELATED DOCUMENTATION

Connecting QFX3100 Director Devices in a Director Group

CHAPTER 7

High Availability

IN THIS CHAPTER

- Troubleshooting VRRP | 63

Troubleshooting VRRP

IN THIS SECTION

- Problem | 63
- Solution | 63

Problem

Description

If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new primary must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new primary).

Solution

Configure a failover delay so that the new primary delays sending gratuitous ARP replies for the period that you set. This allows the new primary to send the ARP replies for all of the VRRP groups simultaneously.

RELATED DOCUMENTATION

- failover-delay

CHAPTER 8

Interfaces

IN THIS CHAPTER

- Troubleshooting an Aggregated Ethernet Interface | 65
- Troubleshooting Network Interfaces | 66
- Troubleshooting Multichassis Link Aggregation | 68

Troubleshooting an Aggregated Ethernet Interface

IN THIS SECTION

- Problem | 65
- Solution | 65

Problem

Description

The **show interfaces terse** command shows that the LAG is down.

Solution

Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet-switching (Layer 2 LAG) or family inet (Layer 3 LAG).

NOTE: Layer 2 LAGs are not supported on OCX Series switches.

- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch.

RELATED DOCUMENTATION

Verifying the Status of a LAG Interface

Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch

Troubleshooting Network Interfaces

IN THIS SECTION

- Statistics for logical interfaces on Layer 2 interfaces are not accurate | 66
- The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down | 67

Statistics for logical interfaces on Layer 2 interfaces are not accurate

IN THIS SECTION

- Problem | 67
- Solution | 67

Problem

Description

On QFX5000 switches, statistics for logical interfaces are not supported on Layer 2 interfaces or on any child member interfaces of Layer 2 aggregated Ethernet (AE) interfaces—that is, output for the **show interfaces *interface-name*** operational-mode command does not provide accurate I/O information for the logical interfaces.

Solution

If you need to see statistics for those logical interfaces, configure firewall filter rules to collect the information.

The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down

IN THIS SECTION

- Problem | 67
- Cause | 68
- Solution | 68

Problem

Description

The switch has an SFP or SFP+ module installed. The interface on the port in which an SFP or SFP+ transceiver is installed is down.

Symptoms

When you check the status with the CLI command **show interfaces *interface-name*** , the disabled port is not listed.

Cause

By default, the SFP or SFP+ module operates in the 10-Gigabit Ethernet mode and supports only SFP or SFP+ transceivers. The operating mode for the module is incorrectly set.

Solution

Only SFP or SFP+ transceivers can be installed in SFP or SFP+ modules. You must configure the operating mode of the SFP or SFP+ module to match the type of transceiver you want to use. For SFP+ transceivers, configure 10-Gigabit Ethernet operating mode.

Troubleshooting Multichassis Link Aggregation

IN THIS SECTION

- MAC Addresses Learned on Multichassis Aggregated Ethernet Interfaces Are Not Removed from the MAC Address Table | **69**
- MC-LAG Peer Does Not Go into Standby Mode | **70**
- Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive | **71**
- Redirect Filters Take Priority over User-Defined Filters | **71**
- Operational Command Output Is Wrong | **72**
- ICCP Connection Might Take Up to 60 Seconds to Become Active | **73**
- MAC Address Age Learned on a Multichassis Aggregated Ethernet Interface Is Reset to Zero | **73**
- MAC Address Is Not Learned Remotely in a Default VLAN | **74**
- Snooping Entries Learned on Multichassis Aggregated Ethernet Interfaces Are Not Removed | **75**
- ICCP Does Not Come Up After You Add or Delete an Authentication Key | **75**
- Local Status Is Standby When It Should Be Active | **76**
- Packets Loop on the Server When ICCP Fails | **76**
- Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change | **77**
- No Commit Checks Are Done for ICL-PL Interfaces | **78**
- Double Failover Scenario | **78**
- Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up | **79**
- Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer | **79**
- Aggregated Ethernet Interfaces Go Down | **80**

Flooding of Upstream Traffic | 81
ARP and MAC Table Entries Become Out of Sync in an MC-LAG Configuration | 81

Use the following information to troubleshoot multichassis link aggregation configuration issues:

MAC Addresses Learned on Multichassis Aggregated Ethernet Interfaces Are Not Removed from the MAC Address Table

IN THIS SECTION

Problem | 69
Solution | 70

Problem

Description

When both of the mulitchassis aggregated Ethernet interfaces on both connected multichassis link aggregation group (MC-LAG) peers are down, the MAC addresses learned on the multichassis aggregated Ethernet interfaces are not removed from the MAC address table.

For example, if you disable the multichassis aggregated Ethernet interface (ae0) on both MC-LAG peers by issuing the **set interfaces ae0 disable** command and commit the configuration, the MAC table still shows the MAC addresses as being learned on the multichassis aggregated Ethernet interfaces of both MC-LAG peers.

```
user@switchA> show ethernet-switching table
Ethernet-switching table: 6 entries, 2 learned, 0 persistent entries
  VLAN               MAC address       Type         Age Interfaces
  v10                *                 Flood          - All-members
  v10                00:00:5E:00:53:00 Learn(L)     3:55 ae0.0 (MCAE)
  v10                00:00:5E:00:53:01 Learn(R)        0 xe-0/0/9.0
  v20                *                 Flood          - All-members
  v30                *                 Flood          - All-members
```

```
  v30               00:00:5E:00:53:03 Static         - Router
```

```
user@switchB> show ethernet-switching table
Ethernet-switching table: 6 entries, 2 learned, 0 persistent entries
  VLAN                MAC address       Type         Age Interfaces
  v10               *                 Flood          - All-members
  v10               00:00:5E:00:53:04 Learn(R)        0 ae0.0 (MCAE)
  v10               00:00:5E:00:53:05 Learn          40 xe-0/0/10.0
  v20               *                 Flood          - All-members
  v30               *                 Flood          - All-members
  v30               00:00:5E:00:53:06 Static         - Router
```

Solution

This is expected behavior.

MC-LAG Peer Does Not Go into Standby Mode

IN THIS SECTION

- Problem | 70
- Solution | 70

Problem

Description

A multichassis link aggregation group (MC-LAG) peer does not go into standby mode if the MC-LAG peer IP address specified in the Inter-Chassis Control Protocol (ICCP) configuration and the IP address specified in the multichassis protection configuration are different.

Solution

To prevent failure to enter standby mode, make sure that the peer IP address in the ICCP configurations and the IP address in multichassis protection configurations are the same.

Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive

IN THIS SECTION

- Problem | 71
- Solution | 71

Problem

Description

When the interchassis control link-protection link (ICL-PL) and multichassis aggregated Ethernet interfaces go down on the primary multichassis link aggregation group (MC-LAG) peer, the secondary MC-LAG peer's multichassis aggregated Ethernet interfaces with status control set to standby become inactive instead of active.

Solution

This is expected behavior.

Redirect Filters Take Priority over User-Defined Filters

IN THIS SECTION

- Problem | 71
- Solution | 72

Problem

Description

Multichassis link aggregation group (MC-LAG) implicit failover redirection filters take precedence over user-configured explicit filters.

Solution

This is expected behavior.

Operational Command Output Is Wrong

IN THIS SECTION

- Problem | 72
- Solution | 72

Problem

Description

After you deactivate Inter-Chassis Control Protocol (ICCP), the **show iccp** operational command output still shows registered client daemons, such as mcsnoopd, lacpd, and eswd.

For example:

```
user@switch> show iccp
Client Application: MCSNOOPD
  Redundancy Group IDs Joined: None

Client Application: lacpd
  Redundancy Group IDs Joined: 1

Client Application: eswd
  Redundancy Group IDs Joined: 1
```

The **show iccp** command output always shows registered modules regardless of whether or not ICCP peers are configured.

Solution

This is expected behavior.

ICCP Connection Might Take Up to 60 Seconds to Become Active

IN THIS SECTION

- Problem | 73
- Solution | 73

Problem

Description

When the Inter-Chassis Control Protocol (ICCP) configuration and the routed VLAN interface (RVI) configuration are committed together, the ICCP connection might take up to 60 seconds to become active.

Solution

This is expected behavior.

MAC Address Age Learned on a Multichassis Aggregated Ethernet Interface Is Reset to Zero

IN THIS SECTION

- Problem | 73
- Solution | 74

Problem

Description

When you activate and then deactivate an interchassis link-protection link (ICL-PL), the MAC address age learned on the multichassis aggregated Ethernet interface is reset to zero. The next-hop interface changes trigger MAC address updates in the hardware, which then triggers aging updates in the Packet Forwarding Engine. The result is that the MAC address age is updated to zero.

For example, the ICL-PL has been deactivated, and the **show ethernet-switching table** command output shows that the MAC addresses have an age of 0.

```
user@switch> show ethernet-switching table
Ethernet-switching table: 3 entries, 2 learned, 0 persistent entries
  VLAN              MAC address       Type          Age Interfaces
  v100              *                 Flood           - All-members
  v100              00:10:00:00:00:01 Learn(L)      0 ae0.0 (MCAE)
  v100              00:10:00:00:00:02 Learn(L)      0 ae0.0 (MCAE)
```

Solution

This is expected behavior.

MAC Address Is Not Learned Remotely in a Default VLAN

IN THIS SECTION

- Problem | 74
- Solution | 74

Problem

Description

On a QFX3500 switch running Junos OS Release 12.3 or earlier, if a multichassis link aggregation group (MC-LAG) peer learns a MAC address in the default VLAN, Inter-Chassis Control Protocol (ICCP) does not synchronize the MAC address with the MAC address of the other MC-LAG peer.

Solution

This is expected behavior.

Snooping Entries Learned on Multichassis Aggregated Ethernet Interfaces Are Not Removed

IN THIS SECTION

- Problem | 75
- Solution | 75

Problem

Description

When multichassis aggregated Ethernet interfaces are configured on a VLAN that is enabled for multicast snooping, the membership entries learned on the multichassis aggregated Ethernet interfaces on the VLAN are not cleared when the multichassis aggregated Ethernet interfaces go down. This is done to speed up convergence time when the interfaces come up, or come up and go down.

Solution

This is expected behavior.

ICCP Does Not Come Up After You Add or Delete an Authentication Key

IN THIS SECTION

- Problem | 75
- Solution | 76

Problem

Description

The Inter-Chassis Control Protocol (ICCP) connection is not established when you add an authentication key and then delete it only at the global ICCP level. However, authentication works correctly at the ICCP peer level.

Solution

Delete the ICCP configuration, and then add the ICCP configuration.

Local Status Is Standby When It Should Be Active

IN THIS SECTION

- Problem | 76
- Solution | 76

Problem

Description

If the multichassis aggregated Ethernet interface is down when the state machine is in a synchronized state, the multichassis link aggregation group (MC-LAG) peer local status is standby. If the multichassis aggregated Ethernet interface goes down after the state machine is in an active state, then the local status remains active, and the local state indicates that the interface is down.

Solution

This is expected behavior.

Packets Loop on the Server When ICCP Fails

IN THIS SECTION

- Problem | 77
- Solution | 77

Problem

Description

When you enable backup liveness detection for a multichassis link aggregation group (MC-LAG), and the backup liveness detection packets are lost because of a temporary failure on the MC-LAG, then both of the peers in the MC-LAG remain active. If this happens, both of the MC-LAG peers send packets to the connected server.

Solution

This is expected behavior.

Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change

IN THIS SECTION

- Problem | 77
- Solution | 77

Problem

Description

After a reboot or after a new Inter-Chassis Control Protocol (ICCP) configuration has been committed, and the ICCP connection does not become active, the Link Aggregation Control Protocol (LACP) messages transmitted over the multichassis aggregated Ethernet interfaces use the default system ID. The configured system ID is used instead of the default system ID only after the MC-LAG peers synchronize with each other.

Solution

This is expected behavior.

No Commit Checks Are Done for ICL-PL Interfaces

IN THIS SECTION

- Problem | 78
- Solution | 78

Problem

Description

There are no commit checks on the interface being configured as an interchassis link-protection link (ICL-PL), so you must provide a valid interface name for the ICL-PL.

Solution

This is expected behavior.

Double Failover Scenario

IN THIS SECTION

- Problem | 78
- Solution | 79

Problem

Description

If the following events happen in this exact order—Inter-Chassis Control Protocol (ICCP) goes down, and the multichassis aggregated Ethernet interface on the multichassis link aggregation group (MC-LAG) peer in active mode goes down—a double failover occurs. In this scenario, the MC-LAG peer in standby mode does not detect what happens on the active MC-LAG peer. The MC-LAG peer in standby mode operates as if the multichassis aggregated Ethernet interface on the MC-LAG in active mode were up and blocks the interchassis link-protection link (ICL-PL) traffic. The ICL-PL traffic is not forwarded.

Solution

This is expected behavior.

Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up

IN THIS SECTION

- Problem | 79
- Solution | 79

Problem

Description

When interchassis link-protection link (ICL-PL) goes down and comes up, multicast traffic is flooded to all of the interfaces in the VLAN. The Packet Forwarding Engine flag Ip4McastFloodMode for the VLAN is changed to MCAST_FLOOD_ALL. This problem only occurs when a multichassis link aggregation group (MC-LAG) is configured for Layer 2.

Solution

This is expected behavior.

Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer

IN THIS SECTION

- Problem | 80
- Solution | 80

Problem

Description

When Inter-chassis Control Protocol (ICCP) is down, the status of a remote MC-LAG peer is unknown. Even if the MC-LAG peer is configured as standby, the traffic is not redirected to this peer because it is assumed that this peer is down.

Solution

This is expected behavior.

Aggregated Ethernet Interfaces Go Down

IN THIS SECTION

- Problem | 80
- Solution | 80

Problem

Description

When a multichassis aggregated Ethernet interface is converted to an aggregated Ethernet interface, it retains some multichassis aggregated Ethernet interface properties. For example, the aggregated Ethernet interface might retain the administrative key of the multichassis aggregated Ethernet interface. When this happens, the aggregated Ethernet interface goes down.

Solution

Restart Link Aggregation Control Protocol (LACP) on the multichassis link aggregation group (MC-LAG) peer hosting the aggregated Ethernet interface to bring up the aggregated Ethernet interface. Restarting LACP removes the multichassis aggregated Ethernet properties of the aggregated Ethernet interface.

Flooding of Upstream Traffic

IN THIS SECTION

- Problem | 81
- Solution | 81

Problem

Description

When MAC synchronization is enabled, the multichassis link aggregation group (MC-LAG) peer can resolve Address Resolution Protocol (ARP) entries for the MC-LAG routed VLAN interface (RVI) with either of the MC-LAG peer MAC addresses. If the downstream traffic is sent with one MAC address (MAC1) but the peer has resolved the MAC address with a different MAC address (MAC2), the MAC2 address might not be learned by any of the access layer switches. Flooding of the upstream traffic for the MAC2 address might then occur.

Solution

Make sure that downstream traffic is sent from the MC-LAG peers periodically to prevent the MAC addresses from aging out.

ARP and MAC Table Entries Become Out of Sync in an MC-LAG Configuration

IN THIS SECTION

- Problem | 82
- Solution | 82

Problem

Description

The ARP and MAC address tables in an MC-LAG configuration normally stay synchronized, but updates might be lost in extreme situations when table updates are happening very frequently, such as when link flapping occurs in an MC-LAG group.

Solution

To avoid ARP and MAC entries becoming out of sync in an MC-LAG configuration, you can configure the arp-l2-validate option on the switch's IRB interface, as follows:

```
user@switch> set interfaces irb arp-l2-validate
```

The **arp-l2-validate** option is available only on QFX Series switches and EX4300 switches starting with Junos OS Release 15.1R4, and EX9200 switches starting with Junos OS Release 13.2R4.

This option turns on validation of ARP and MAC table entries, automatically applying updates if they become out of sync. You might want to enable this option as a workaround when the network is experiencing other issues that also cause loss of ARP and MAC synchronization, but disable it during normal operation because this option might impact performance in scale configurations.

CHAPTER 9

Junos OS Basics

IN THIS CHAPTER

- System Troubleshooting | 83
- Recovering from a Failed Software Installation | 85
- Recovering the Root Password for Switches | 87
- Creating an Emergency Boot Device for QFX Series Switches | 89
- Performing a Recovery Installation | 91
- Performing a QFabric System Recovery Installation on the Director Group | 94
- Troubleshooting Network Interfaces | 103
- Troubleshooting an Aggregated Ethernet Interface | 105

System Troubleshooting

IN THIS SECTION

- Saving Core Files Generated by Junos OS Processes | 83
- Viewing Core Files from Junos OS Processes | 84

Saving Core Files Generated by Junos OS Processes

By default, when an internal Junos OS process generates a core file, the file and associated context information are saved for debugging purposes in a compressed tar file named **/var/tmp/*process-name*.core.*core-number*.tgz**. The contextual information includes the configuration and system log message files.

- To disable the saving of core files and associated context information:

```
[edit system]
no-saved-core-context;
```

- To save the core files only:

```
[edit system]
saved-core-files number;
```

  Where ***number*** is the number of core files to save and can be a value from 1 through 10.

- To save the core files along with the contextual information:

```
[edit system]
saved-core-context;
```

Viewing Core Files from Junos OS Processes

When an internal Junos OS process generates a core file, you can find the output at **/var/crash/** and **/var/tmp/**. For Junos OS Evolved, you can find the output core files at **/var/core/** for Routing Engine core files and **/var/lib/ftp/in/** for FPC core files. Using these directories provides a quick method of finding core issues across large networks.

Use the CLI command **show system core-dumps** to view core files.

```
root@host> show system core-dumps
-rw-------  1 root  wheel  268369920 Jun 18 17:59 /var/crash/vmcore.0
-rw-rw----  1 root  field    3371008 Jun 18 17:53 /var/tmp/rpd.core.0
-rw-r--r--  1 root  wheel   27775914 Jun 18 17:59 /var/crash/kernel.0
```

SEE ALSO

Saving Core Files from Junos OS Processes

RELATED DOCUMENTATION

Day One: Monitoring and Troubleshooting

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/troubleshooting-qf.html

Recovering from a Failed Software Installation

IN THIS SECTION

- Problem | 85
- Solution | 85

Problem

Description

If the Junos OS appears to have been installed but the CLI does not work, or if the device has no software installed, you can use this recovery installation procedure to install the Junos OS.

Solution

If a Junos OS image already exists on the switch, you can either install the new Junos OS package in a separate partition, in which case both Junos OS images remain on the switch, or you can remove the existing Junos OS image before you start the new installation process.

NOTE: QFX5100, QFX5200, EX4600, QFX10000, and OCX Series switches do not have a separate partition to reinstall a Junos OS image.

A recovery image is created automatically on these switches. If a previously-running switch is powered on and unable to boot using a Junos OS image, you can boot the switch using the recovery Junos OS image by selecting an option in the “Select a recovery image” menu.

We suggest creating a system snapshot on your switch onto the external USB flash drive, and using the snapshot for recovery purposes. The system snapshot feature takes a “snapshot” of the

files currently used to run the device—the complete contents of the **/config** directories, which include the running Juniper Networks Junos OS, the active configuration, and the rescue configuration, as well as the host OS—and copies all of these files into an external USB flash drive. See Creating a Snapshot and Using It to Boot a QFX3500 and QFX3600 Series Switch or Creating a Snapshot and Using It to Boot a QFX Series Switch.

System snapshot is not supported on QFX5200 and QFX10000 switches.

To perform a recovery installation:

1. Power on the switch. The loader script starts.
2. After the message **Loading /boot/defaults/loader.conf** appears, you are prompted with the following message:

   Hit [Enter] to boot immediately, or space bar for command prompt.

   Press the Spacebar to enter the manual loader. The **loader>** prompt appears.

   NOTE: The loader prompt does not appear on QFX5100, QFX5200, EX4600, QFX10000, and OCX Series switches.

   On QFX5100, QFX5200, EX4600, QFX10000, and OCX Series switches only, a recovery image is automatically saved if a previously-running switch is powered on and unable to boot using a Junos OS image.

   The “Select a recovery image” menu appears on the console when one of these switches is booted and unable to load a version of Junos OS. Follow the instructions in the “Select a recovery image” menu to load the recovery version of Junos OS for one of these switches.

   You can ignore the remainder of this procedure if you are using a QFX5100, QFX5200, EX4600, QFX10000, or OCX Series switch.

3. Enter the following command:

   ```
   loader> install [– –format]  [– –external] source
   ```

   where:

   - **format**—Enables you to erase the installation media before installing the installation package. If you do not include this option, the system installs the new Junos OS in a different partition from that of the most recently installed Junos OS.

- **external**—Installs the installation package onto external media (a USB stick, for example).
- ***source***—Represents the name and location of the Junos OS package, either on a server on the network or as a file on an external media, as shown in the following two examples:
  - Network address of the server and the path on the server; for example, **tftp://192.0.2.0/junos/jinstall-qfx-5e-flex-15.1X53-D30.5-domestic-signed.tgz**
  - Junos OS package on a USB device (commonly stored in the root drive as the only file), for example, **file://jinstall-qfx-5e-flex-15.1X53-D30.5-domestic-signed.tgz**).

The installation now proceeds normally and ends with a login prompt.

RELATED DOCUMENTATION

Creating a Snapshot and Using It to Boot a QFX3500 and QFX3600 Series Switch

Creating a Snapshot and Using It to Boot a QFX Series Switch

Recovering the Root Password for Switches

If you forget the root password, you can use the password recovery procedure to reset the root password.

NOTE: The root password cannot be recovered on a QFabric system.

NOTE: You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the device's power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the device into the RJ-45–to–DB-9 serial port adapter supplied with the device.
4. Plug the RJ-45–to–DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the device.

6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
   - Bits per second: 9600
   - Data bits: 8
   - Parity: None
   - Stop bits: 1
   - Flow control: None
9. Power on the device by (if necessary) plugging the power cords into the device's power supply, or turning on the power to the device by switching on the AC power outlet the device is plugged into.

   The terminal emulation screen on your management device displays the device's boot sequence.
10. When the following prompt appears, press the Spacebar to access the device's bootstrap loader command prompt:

```
Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [kernel] in 9 seconds...
```

11. At the following prompt, enter **boot -s** to start up the system in single-user mode.

```
ok boot -s
```

12. At the following prompt, enter **recovery** to start the root password recovery procedure.

```
Enter full pathname of shell or 'recovery' for root password recovery or
RETURN for /bin/sh: recovery
```

13. Enter configuration mode in the CLI.
14. Set the root password. For example:

```
user@switch# set system root-authentication plain-text-password
```

15. At the following prompt, enter the new root password. For example:

```
New password: ABC123
```

```
Retype new password:
```

16. At the second prompt, reenter the new root password.
17. After you have finished configuring the password, commit the configuration.

```
root@host# commit
commit complete
```

18. Exit configuration mode in the CLI.
19. Exit operational mode in the CLI.
20. At the prompt, enter **y** to reboot the device.

```
Reboot the system? [y/n] y
```

RELATED DOCUMENTATION

- Configuring the Root Password

Creating an Emergency Boot Device for QFX Series Switches

Before you begin, you need to download the installation media image for your device and Junos OS release from https://www.juniper.net/customers/support/ .

If Junos OS on the device is damaged in some way that prevents the software from loading properly, you can use an emergency boot device to repartition the primary disk and load a fresh installation of Junos OS. Use the following procedure to create an emergency boot device.

NOTE: You can create the emergency boot device on another Juniper Networks switch or router, or any laptop or desktop PC that supports Linux. The steps you take to create the emergency boot device vary, depending on the device.

To create an emergency boot device:

1. Use FTP to copy the installation media image into the **/var/tmp** directory on the device.
2. Insert a USB storage device into the USB port.
3. From the Junos OS command-line interface (CLI), start the shell:

```
user@device> start shell
%
```

4. Use the gunzip command to unzip the image file.
5. Switch to the root account using the **su** command:

```
% su
Password: password
```

NOTE: The password is the root password for the device. If you logged in to the device as the root user, you do not need to perform this step.

6. Enter the following command on the device:

```
root@device% dd if=/var/tmp/filename of=/dev/da1 bs=1m
```

The device writes the installation media image to the USB storage device:

```
root@device% dd if=install-media-qfx-5e-15.1X53-D30.5-domestic.img of=/dev/
da0 bs=1m
1399+0 records in
1399+0 records out
1466957824 bytes transferred in 394.081902 secs (3722469 bytes/sec)
```

7. Log out of the shell:

```
root@device% exit
% exit
user@device>
```

Performing a Recovery Installation

If Junos OS on your device is damaged in some way that prevents the software from loading correctly, you may need to perform a recovery installation using an emergency boot device (for example, a USB flash drive) to restore the default factory installation. Once you have recovered the software, you need to restore the device configuration. You can either create a new configuration as you did when the device was shipped from the factory, or if you saved the previous configuration, you can simply restore that file to the device.

Starting in Junos OS Release 14.1, you can also use a system snapshot as a bootup option when your Junos OS or configuration is damaged. The system snapshot feature takes a "snapshot" of the files currently used to run the device—the complete contents of the **/config** directories, which include the running Juniper Networks Junos OS, the active configuration, and the rescue configuration, as well as the host OS—and copies all of these files into an external USB flash drive. See Understanding How to Back Up an Installation on Switches.

NOTE: System snapshot is not supported on QFX10002 switches.

If at all possible, you should try to perform the following steps before you perform the recovery installation:

1. Ensure that you have an emergency boot device to use during the installation. See *Creating an Emergency Boot Device for QFX Series Switches* for information on how to create an emergency boot device.

2. Copy the existing configuration in the file **/config/juniper.conf.gz** from the device to a remote system, such as a server, or to an emergency boot device. For extra safety, you can also copy the backup configurations (the files named **/config/juniper.conf.*n***, where *n* is a number from 0 through 9) to a remote system or to an emergency boot device.

WARNING: The recovery installation process completely overwrites the entire contents of the internal flash storage.

3. Copy any other stored files to a remote system as desired.

To reinstall Junos OS:

1. Insert the emergency boot device into the QFX Series device.
2. Reboot the QFX Series device.

NOTE: Do not power off the device if it is already on.

```
[edit system]
user@device> request system reboot
```

If you do not have access to the CLI, power cycle the QFX Series device.

The emergency boot device (external USB install media) is detected. At this time, you can load the Junos OS from the emergency boot device onto the internal flash storage.

3. The software prompts you with the following options:

```
External USB install media detected.
You can load Junos from this media onto an internal drive.
Press 'y' to proceed, 'f' to format and install, or 'n' to abort.
Do you wish to continue ([y]/f/n)? f
```

4. Type **f** to format the internal flash storage and install the Junos OS on the emergency boot device onto the internal flash storage.

If you do no want to format the internal flash storage, type **y**.

The following messages are displayed:

```
Installing packages from external USB drive da1
Packages will be installed to da0, media size: 8G

Processing format options
```

```
Fri September  4 01:18:44 UTC 2012

-- IMPORTANT INFORMATION --
Installer has detected settings to format system boot media.
This operation will erase all data from your system.

Formatting installation disk .. this will take a while, please wait
Disabling platform watchdog - threshold 12 mins

Determining installation slice
Fri September  4 01:27:07 UTC 2012
```

5. The device copies the software from the emergency boot device, occasionally displaying status messages. Copying the software can take up to 12 minutes.

   When the device is finished copying the software, you are presented with the following prompt:

```
*** Fri September  4 01:19:00 UTC 2012***
Installation successful..
Please select one of the following options:
Reboot to installed Junos after removing install media (default) ... 1
Reboot to installed Junos by disabling install media ............... 2
Exit to installer debug shell ..................................... 3
Install Junos to alternate slice .................................. 4
Your choice: 4
NOTE: System installer will now install Junos to alternate slice
Do not power off or remove the external installer media or
interrupt the installation mechanism.
```

6. Select **4** to install Junos OS to the alternate slice of the partition, and then press Enter.

7. Remove the emergency boot device when prompted and then press Enter. The device then reboots from the internal flash storage on which the software was just installed. When the reboot is complete, the device displays the login prompt.

8. Create a new configuration as you did when the device was shipped from the factory, or restore the previously saved configuration file to the device.

Release History Table

Release	Description
14.1	Starting in Junos OS Release 14.1, you can also use a system snapshot as a bootup option when your Junos OS or configuration is damaged.

RELATED DOCUMENTATION

Creating an Emergency Boot Device for QFX Series Switches

Performing a QFabric System Recovery Installation on the Director Group

IN THIS SECTION

- (Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive | **95**
- Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software | **96**

If the software on your QFabric system is damaged in some way that prevents the software from loading correctly, or you need to upgrade the software on your QFabric system, you may need to perform a recovery installation on the Director group.

If possible, perform the following steps before you perform the recovery installation:

1. Ensure that you have an emergency boot device (for example, an external USB flash drive) for each of your Director devices to use during the recovery installation.

   You can either use the external USB flash drive containing the software supplied by Juniper Networks, or you can use an external USB flash drive supplied by Juniper Networks on which you install the QFabric system install media.

2. Because the recovery installation process completely overwrites the entire contents of the Director device, make sure you back up any configuration files and initial setup information on a different external USB flash drive before you begin a recovery installation. You will need to restore this information as part of recovery process.

Use the **request system software configuration-backup** command to back up your configuration files and initial setup information:

```
user@switch> request system software configuration-backup path
```

NOTE: To recover the Director group, you must upgrade both Director devices in parallel. If you are recovering only one Director device in a Director group, and the software version will remain the same between the two Director devices, make sure that the other Director device is powered on and operational. If the software version of the Director device you are recovering will be different, make sure that the other Director device is powered off and is not operational.

(Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive

If you do not have an external USB flash drive preloaded with the software from Juniper Networks to use as an emergency boot device, you can create your own, using a blank external USB flash drive provided by Juniper Networks. Download the install media from the Juniper Networks Support website onto your UNIX workstation, uncompress and untar the software, and then burn the software image onto your Juniper Networks external USB (4-gigabyte) flash drive. Make sure you create two emergency boot devices, one for each Director device, so you can perform a recovery installation in parallel.

1. Using a Web browser, navigate to the https://www.juniper.net/support .
2. Click **Download Software**.
3. In the *Switching*box, click *Junos OS Platforms*.
4. In the *QFX Series* section, click the name of the platform for which you want to download software.
5. Click the *Software* tab and select the release number from the *Release* drop-down list.
6. Select the complete install media you want to download in the *QFabric System Install Media* section.

   A login screen appears.
7. Enter your name and password and press **Enter**.
8. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
9. Log in and save the install media file to your UNIX workstation.
10. Use FTP to access the UNIX workstation where the install media resides.

    ```
    ftp ftp:// hostname/pathname install-media-qfabric-<version>.img.tgz
    ```
11. When prompted, enter your username and password.

12. Make sure you are in binary mode by entering **binary** at the prompt.

```
binary
```

13. Use the **get** command to transfer the installation package from the FTP host to your UNIX workstation.

```
get install-media-qfabric-<version>.img.tgz
```

14. Close the FTP session:

```
bye
```

15. Untar the *install-media-qfabric-<version>* .img.tgz file on your UNIX workstation.

```
tar -xvzf install-media-qfabric-11.3X30.6.img.tgz
```

16. Insert a blank external USB (4-gigabyte) flash drive supplied by Juniper Networks into your UNIX workstation.
17. Erase the bootable partition in the external USB flash drive by issuing the following **dd** command.

```
dd if=/dev/zero of=/dev/sdb count=20
```

18. Burn the software image you just downloaded to your UNIX workstation onto your external USB flash drive by issuing the following **dd** command:

```
dd if=install-media-qfabric-11.3X30.6.img of=/dev/sdb bs=16k
250880+0 records in
250880+0 records out
4110417920 bytes (4.1 GB) copied, 5.10768 seconds, 805 MB/s
```

19. Perform the steps in "Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software" on page 96 to continue with the recovery installation.

Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software

This procedure describes how to perform a recovery installation using an external USB flash drive that contains Junos OS software.

NOTE: Since the recovery installation process completely overwrites the entire contents of the Director device, you will need to restore the required configuration files and initial setup information. The following procedure assumes you previously saved these backup files with the **request system software configuration-backup** command. Ensure that you have these backup files available on an external USB flash drive before you perform the following steps.

1. Insert the external USB flash drive into the Director device.
2. Perform one of the following tasks:
   - If you have access to the default partition, reboot the Director device by issuing the **request system reboot director-group** command.
   - If you do not have access to the default partition, power cycle the Director device.

   The following menu appears on the Director device console when the Director device boots up:

```
Juniper Networks QFabric Director Install/Recovery Media
-  To boot from the local disk, wait 10 seconds or press the Enter key.
-  To reinstall the QFabric software on this Director device, type: install
```

3. Type **install** and then press **Enter** to install the software on the Director device.

   Once the installation process is complete, the Director device reboots, and the following menu appears on the Director device console:

```
Juniper Networks QFabric Director Install/Recovery Media
-  To boot from the local disk, wait 10 seconds or press the Enter key.
-  To reinstall the QFabric software on this Director device, type: install
```

4. Press **Enter** twice.

   The Director device reboots a second time from the local disk that contains the newly installed software.
5. When you see the following prompts, press **Enter**.

```
Starting xinetd: [  OK  ]
Starting atop:  [  OK  ]
```

6. Log in as root on the Director device. Type **root** and press **Enter**.

```
dg0 login: root
```

7. Because the root password has been removed as part of the recovery process, press **Enter** a second time to skip the password entry step.

> **NOTE**: Do not enter a root password at this time.

8. The following menu appears on the Director device console:

```
Before you can access the QFabric system, you must complete the initial
setup of the Director group by using the steps that follow.
If the initial setup procedure does not complete successfully, log out of
the Director device and then log back in to restart
 this setup menu.

Continue?[y/n]
```

9. Enter **n** to bypass the initial setup script and enter the Director device root directory, where you can mount the external USB flash drive containing the configuration files and initial setup information.
10. Issue the **ls /mnt** command to list the *mount* directory.

```
root@dg0 ~]# ls /mnt
```

11. Issue the **mkdir** command to create a directory within the mount directory.

```
root@dg0 ~]# mkdir /mnt/myusb
```

12. Issue the **mount /dev/sdb2 /mnt/myusb/** command to mount the external USB flash drive to the local drive of the Director device.

```
root@dg0 ~]# mount /dev/sdb2 /mnt/myusb/
```

13. Issue the **ls -la /mnt/myusb/** command to verify the contents of your mounted external USB flashdrive.

```
root@dg0 ~]# ls -la /mnt/myusb/
total 1770884
drwxr-xr-x 2 root root       4096 Sep  7 05:16 .
drwxr-xr-x 3 root root       4096 Sep  7 10:15 ..
-rw-r--r-- 1 root root       4249 Sep  7 03:52 mybackup-20110907
```

14. Exit the Director device and log back in as root on the Director device.

    The following menu appears:

```
Before you can access the QFabric system, you must complete the initial
setup of the Director group by using the steps that follow.
If the initial setup procedure does not complete successfully, log out of
the Director device and then log back in to restart
 this setup menu.

Continue?[y/n] y
Initial Configuration

You may enter the configuration manually or restore from a backup.

Specify a backup file? [y/n] : y
Please specify the full path of the configuration backup file. : /mnt/myusb/
mybackup-20110907
```

15. Enter **y** to continue.
16. Enter **y** and specify the path to the backup configuration file located on the external USB flash drive.

```
/mnt/myusb/mybackup-20110907
```

    The following messages appear:

```
Saving temporary configuration...
Configuring peer...
connect error for 1.1.1.2:9001
Configuring local interfaces...
Configuring interface eth0 with [10.49.213.163/24:10.49.213.254]
```

```
Configured interface eth0 with [10.49.213.163/24:10.49.213.254]
Configuring QFabric software with initial pool of 4000 MAC addresses
[00:10:00:00:00:00 - 00:10:00:00:0f:3b]
Configuring QFabric address [10.49.213.50]
Reconfiguring QFabric software static configuration
Applying the new Director Device password
Applying the QFabric component password
First install initial configuration, generating and sharing SSH keys.
First install initial configuration, generating SSH keys.
connect error for 1.1.1.2:9001
Shared SSH keys.
Configuration complete. Director Group services will auto start within 30
seconds.
```

The Director device reboots from the local disk on which the software was just installed. Exit the Director device session and log in to the QFabric default partition CLI.

17. Issue the **request system software configuration-restore** command and specify the path to the backup configuration file located on the external USB flash drive to load the previously saved QFabric system configuration.
18. From the default partition, issue the **request system reboot node-group all** command to reboot all of the Node groups in the QFabric system to ensure that all Node devices are running the same version of software as the Director-group.

```
user@switch> request system reboot node-group all
```

19. From the default partition, issue the **request system reboot fabric** command to reboot the Interconnect devices and the other components in the fabric in the QFabric system to ensure that Interconnect devices are running the same version of software as the Director group.

```
user@switch> request system reboot fabric
```

20. Log in to the default partition and issue the **show version component all** command to verify that all components are running the same version of software.

```
user@switch> show version component all
dg1:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3X30.6]
```

```
dg0:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3X30.6]

NW-NG-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FC-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FC-1:
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
```

```
JUNOS Routing Software Suite [11.3X30.6]

DRE-0:
-
Hostname: dre-0
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

FM-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

nodedevice1:
-
Hostname: qfabric
Model: QFX3500
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]

interconnectdevice1:
-
```

```
Hostname: qfabric
Model: QFX3108
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
warning:  from interconnectdevice0: Disconnected
```

RELATED DOCUMENTATION

Performing the QFabric System Initial Setup on a QFX3100 Director Group

Upgrading Software on a QFabric System

request system software configuration-backup

request system software configuration-restore

Troubleshooting Network Interfaces

IN THIS SECTION

- Statistics for logical interfaces on Layer 2 interfaces are not accurate | **103**
- The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down | **104**

Statistics for logical interfaces on Layer 2 interfaces are not accurate

IN THIS SECTION

- Problem | **104**

- Solution | 104

Problem

Description

On QFX5000 switches, statistics for logical interfaces are not supported on Layer 2 interfaces or on any child member interfaces of Layer 2 aggregated Ethernet (AE) interfaces—that is, output for the **show interfaces *interface-name*** operational-mode command does not provide accurate I/O information for the logical interfaces.

Solution

If you need to see statistics for those logical interfaces, configure firewall filter rules to collect the information.

The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down

IN THIS SECTION

- Problem | 104
- Cause | 105
- Solution | 105

Problem

Description

The switch has an SFP or SFP+ module installed. The interface on the port in which an SFP or SFP+ transceiver is installed is down.

Symptoms

When you check the status with the CLI command **show interfaces *interface-name*** , the disabled port is not listed.

Cause

By default, the SFP or SFP+ module operates in the 10-Gigabit Ethernet mode and supports only SFP or SFP+ transceivers. The operating mode for the module is incorrectly set.

Solution

Only SFP or SFP+ transceivers can be installed in SFP or SFP+ modules. You must configure the operating mode of the SFP or SFP+ module to match the type of transceiver you want to use. For SFP+ transceivers, configure 10-Gigabit Ethernet operating mode.

Troubleshooting an Aggregated Ethernet Interface

IN THIS SECTION

- Problem | **105**
- Solution | **105**

Problem

Description

The **show interfaces terse** command shows that the LAG is down.

Solution

Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.

- Verify that a LAG is part of family ethernet-switching (Layer 2 LAG) or family inet (Layer 3 LAG).

  NOTE: Layer 2 LAGs are not supported on OCX Series switches.

- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch.

RELATED DOCUMENTATION

Verifying the Status of a LAG Interface

Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch

CHAPTER 10

Layer 3 Protocols

IN THIS CHAPTER

- Troubleshooting Virtual Routing Instances | 107

Troubleshooting Virtual Routing Instances

IN THIS SECTION

- Direct Routes Not Leaked Between Routing Instances | 107

Direct Routes Not Leaked Between Routing Instances

IN THIS SECTION

- Problem | 107
- Solution | 108

Problem

Description

Direct routes are not exported (leaked) between virtual routing instances For example, consider the following scenario:

- Switch with two virtual routing instances:

- Routing instance 1 connects to downstream device through interface xe-0/0/1.
- Routing instance 2 connects to upstream device through interface xe-0/0/2.

If you enable route leaking between the routing instances (by using the **rib-group** statement, for example), the downstream device cannot connect to the upstream device because the switch connects to the upstream device over a direct route and these routes are not leaked between instances.

> **NOTE**: You can see a route to the upstream device in the routing table of the downstream device, but this route is not functional.

Indirect routes *are* leaked between routing instances, so the downstream device can connect to any upstream devices that are connected to the switch over indirect routes.

Solution

This is expected behavior.

SEE ALSO

- Understanding Virtual Router Routing Instances
- Configuring Virtual Router Routing Instances
- rib-group

RELATED DOCUMENTATION

- Understanding Virtual Router Routing Instances
- Configuring Virtual Router Routing Instances
- rib-group

CHAPTER 11

Network Management

IN THIS CHAPTER

- Understanding Troubleshooting Resources | 109
- Troubleshooting Overview | 113
- Recovering from a Failed Software Installation | 116
- Returning to a Previously Committed Configuration | 119
- Reverting to the Default Factory Configuration | 123
- Reverting to the Rescue Configuration | 124
- Recovering the Root Password for Switches | 124
- Troubleshooting a Deprecated Network Analytics Configuration | 126

Understanding Troubleshooting Resources

This topic describes some of the troubleshooting resources available for the QFX Series or OCX Series. These resources include tools such as the Junos OS CLI, Junos Space applications, and the Advanced Insight Scripts (AI-Scripts).

Table 8 on page 109 provides a list of some of the troubleshooting resources.

Table 8: Troubleshooting Resources on the QFX and OCX Series

Troubleshooting Resource	Description	Documentation
Chassis alarms	Chassis alarms indicate a failure on the switch or one of its components. A chassis alarm count is displayed on the LCD panel on the front of the switch.	"Chassis Alarm Messages on a QFX3500 Device" on page 12

Table 8: Troubleshooting Resources on the QFX and OCX Series *(Continued)*

Troubleshooting Resource	Description	Documentation
Chassis Status LEDs and Fan Tray LEDs	A blinking amber Power, Fan, or Fan Tray LED indicates a hardware component error. A blinking amber Status LED indicates a software error.	Chassis Status LEDs on a QFX3500 Device
Interface alarms	A predefined alarm (red or yellow) for an interface type is triggered when an interface of that type goes down.	*Interface Alarm Messages*
System alarms	A predefined alarm is triggered by a missing rescue configuration or problem with the software license.	*Understand Alarms*
System log messages	The system log includes details of system and user events, including errors. Specify the severity and type of system log messages you wish to view or save, and configure the output to be sent to local or remote hosts.	• Overview of Single-Chassis System Logging Configuration • Junos OS System Log Configuration Statements
Junos OS operational mode commands	Operational mode commands can be used to monitor switch performance and current activity on the network. For example, use the **traceroute monitor** command to locate points of failure in a network.	• Monitoring System Process Information • Monitoring System Properties • traceroute monitor

Table 8: Troubleshooting Resources on the QFX and OCX Series *(Continued)*

Troubleshooting Resource	Description	Documentation
Junos OS automation scripts (event scripts)	Event scripts can be used to automate network troubleshooting and management tasks.	Automation Scripting User Guide
Junos OS XML operational tags	XML operational tags are equivalent in function to operational mode commands in the CLI, which you can use to retrieve status information for a device.	*Junos XML API Operational Developer Reference*
NETCONF XML management protocol	The NETCONF XML management protocol defines basic operations that are equivalent to Junos OS CLI configuration mode commands. Client applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as **show**, **set**, and **commit** to perform those operations.	NETCONF XML Management Protocol Developer Guide
SNMP MIBs and traps	MIBs enable the monitoring of network devices from a central location. For example, use the Traceroute MIB to monitor devices remotely.	• SNMP MIBs Support • SNMP Traps Support • Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS

Table 8: Troubleshooting Resources on the QFX and OCX Series *(Continued)*

Troubleshooting Resource	Description	Documentation
AI-Scripts and Advanced Insight Manager (AIM)	AI-Scripts installed on the switch can automatically detect and monitor faults on the switch, and depending on the configuration on the AIM application, send notifications of potential problems and submit problem reports to Juniper Support Systems.	*Advanced Insight Scripts (AI-Scripts) Release Notes*
Junos Space Service Now	This application enables you to display and manage information about problem events. When problems are detected on the switch by Advanced Insight Scripts (AI-Scripts) that are installed on the switch, the data is collected and sent to Service Now for your review and action.	Service Automation
Junos Space Service Insight	This application helps in accelerating operational analysis and managing the exposure to known issues. You can identify devices that are nearing their End Of Life (EOL) and also discover and prevent issues that could occur in your network. The functionality of Service Insight is dependent on the information sent from Service Now.	Service Automation
Juniper Networks Knowledge Base	You can search in this database for Juniper Networks product information, including alerts and troubleshooting tips.	https://kb.juniper.net

Troubleshooting Overview

This topic provides a general guide to troubleshooting some typical problems you may encounter on your QFX Series or OCX Series product.

Table 9 on page 113 provides a list of problem categories, summary of the symptom or problem, and recommended actions with links to the troubleshooting documentation.

Table 9: Troubleshooting on the QFX Series

Problem Category	Symptom or Problem	Recommended Action
Switch hardware components	LCD panel shows a chassis alarm count.	See "Chassis Alarm Messages on a QFX3500 Device" on page 12.
	Fan tray LED is blinking amber.	See Fan Tray LED on a QFX3500 Device.
	Chassis status LED for the power is blinking amber.	See Chassis Status LEDs on a QFX3500 Device.
	Chassis status LED for the fan (on the management board) is blinking amber.	Replace the management board as soon as possible. See Chassis Status LEDs on a QFX3500 Device.
Port configuration	Cannot configure a port as a Gigabit Ethernet port.	Check whether the port is a valid Gigabit Ethernet port (6 through 41). See QFX3500 Device Overview.
	Cannot configure a port as a Fibre Channel port.	Check whether the port is a valid Fibre Channel port (0 through 5 and 42 through 47). See QFX3500 Device Overview.

Table 9: Troubleshooting on the QFX Series *(Continued)*

Problem Category	Symptom or Problem	Recommended Action
	Cannot configure a port as a 10-Gigabit Ethernet port.	If the port is not a 40-Gbps QSFP+ interface, check whether the port is in the range of 0 through 5 or 42 through 47. If one of the ports in that block (0 through 5 or 42 through 47) is configured as a Fibre Channel port, then all ports in that block must also be configured as Fibre Channel ports. If the port is a 40-Gbps QSFP+ interface. make sure the configuration does not exceed the interface limit. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces, but because port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces. See QFX3500 Device Overview.
	Cannot configure a 40-Gbps QSFP+ interface.	The 40-Gbps QSFP+ interfaces can only be used as 10-Gigabit Ethernet interfaces. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces using a breakout cable. However, port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces. See QFX3500 Device Overview.
External devices (USB devices)	Upgrading software from a USB device results in an upgrade failure, and the system enters an invalid state.	Unplug the USB device and reboot the switch.

Table 9: Troubleshooting on the QFX Series *(Continued)*

Problem Category	Symptom or Problem	Recommended Action
Initial device configuration	Cannot configure management Ethernet ports.	Configure the management ports from the console port. You cannot configure the management ports by directly connecting to them. **NOTE**: The management ports are on the front panel of the QFX3500 switch. They are labeled **C0** and **C1** on the front panel. In the CLI they are referred to as **me0** and **me1**. See Configuring a QFX3500 Device as a Standalone Switch.
Software upgrade and configuration	Failed software upgrade.	See *Recovering from a Failed Software Installation*.
	Active partition becomes inactive after upgrade.	
	Problem with the active configuration file.	See the following topics: • Loading a Previous Configuration File • *Reverting to the Default Factory Configuration* • "Reverting to the Rescue Configuration" on page 47 • *Performing a Recovery Installation*
	Root password is lost or forgotten.	Recover the root password. See "Recovering the Root Password for Switches" on page 87.

Table 9: Troubleshooting on the QFX Series *(Continued)*

Problem Category	Symptom or Problem	Recommended Action
Network interfaces	An aggregated Ethernet interface is down.	See "Troubleshooting an Aggregated Ethernet Interface" on page 65.
	Interface on built-in network port is down.	See "Troubleshooting Network Interfaces" on page 66.
	Interface on port in which SFP or SFP + transceiver is installed in an SFP+ uplink module is down.	
Ethernet switching	A MAC address entry in the Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch.	See *Troubleshooting Ethernet Switching*.
Firewall filter	Firewall configuration exceeded available Ternary Content Addressable Memory (TCAM) space.	See Troubleshooting Firewall Filters.

Recovering from a Failed Software Installation

IN THIS SECTION

- Problem | 117
- Solution | 117

Problem

Description

If the Junos OS appears to have been installed but the CLI does not work, or if the device has no software installed, you can use this recovery installation procedure to install the Junos OS.

Solution

If a Junos OS image already exists on the switch, you can either install the new Junos OS package in a separate partition, in which case both Junos OS images remain on the switch, or you can remove the existing Junos OS image before you start the new installation process.

> **NOTE**: QFX5100, QFX5200, EX4600, QFX10000, and OCX Series switches do not have a separate partition to reinstall a Junos OS image.
>
> A recovery image is created automatically on these switches. If a previously-running switch is powered on and unable to boot using a Junos OS image, you can boot the switch using the recovery Junos OS image by selecting an option in the “Select a recovery image” menu.
>
> We suggest creating a system snapshot on your switch onto the external USB flash drive, and using the snapshot for recovery purposes. The system snapshot feature takes a “snapshot” of the files currently used to run the device—the complete contents of the **/config** directories, which include the running Juniper Networks Junos OS, the active configuration, and the rescue configuration, as well as the host OS—and copies all of these files into an external USB flash drive. See Creating a Snapshot and Using It to Boot a QFX3500 and QFX3600 Series Switch or Creating a Snapshot and Using It to Boot a QFX Series Switch.
>
> System snapshot is not supported on QFX5200 and QFX10000 switches.

To perform a recovery installation:

1. Power on the switch. The loader script starts.
2. After the message **Loading /boot/defaults/loader.conf** appears, you are prompted with the following message:

   Hit [Enter] to boot immediately, or space bar for command prompt.

   Press the Spacebar to enter the manual loader. The **loader>** prompt appears.

NOTE: The loader prompt does not appear on QFX5100, QFX5200, EX4600, QFX10000, and OCX Series switches.

On QFX5100, QFX5200, EX4600, QFX10000, and OCX Series switches only, a recovery image is automatically saved if a previously-running switch is powered on and unable to boot using a Junos OS image.

The “Select a recovery image” menu appears on the console when one of these switches is booted and unable to load a version of Junos OS. Follow the instructions in the “Select a recovery image” menu to load the recovery version of Junos OS for one of these switches.

You can ignore the remainder of this procedure if you are using a QFX5100, QFX5200, EX4600, QFX10000, or OCX Series switch.

3. Enter the following command:

```
loader> install [– –format]  [– –external] source
```

where:

- **format**—Enables you to erase the installation media before installing the installation package. If you do not include this option, the system installs the new Junos OS in a different partition from that of the most recently installed Junos OS.
- **external**—Installs the installation package onto external media (a USB stick, for example).
- ***source***—Represents the name and location of the Junos OS package, either on a server on the network or as a file on an external media, as shown in the following two examples:
  - Network address of the server and the path on the server; for example, **tftp://192.0.2.0/junos/jinstall-qfx-5e-flex-15.1X53-D30.5-domestic-signed.tgz**
  - Junos OS package on a USB device (commonly stored in the root drive as the only file), for example, **file://jinstall-qfx-5e-flex-15.1X53-D30.5-domestic-signed.tgz**).

The installation now proceeds normally and ends with a login prompt.

RELATED DOCUMENTATION

Creating a Snapshot and Using It to Boot a QFX3500 and QFX3600 Series Switch

Creating a Snapshot and Using It to Boot a QFX Series Switch

Returning to a Previously Committed Configuration

IN THIS SECTION

- Returning to a Configuration Prior to the One Most Recently Committed | 119
- Displaying Previous Configurations | 119
- Comparing Configuration Changes with a Prior Version | 121

This topic explains how you can return to a configuration prior to the most recently committed one.

Returning to a Configuration Prior to the One Most Recently Committed

To return to a configuration prior to the most recently committed one, include the configuration number, 0 through 49, in the **rollback** command. The most recently saved configuration is number 0 (which is the default configuration to which the system returns), and the oldest saved configuration is number 49.

```
[edit]
user@host# rollback number
load complete
```

Displaying Previous Configurations

To display previous configurations, including the rollback number, date, time, the name of the user who committed changes, and the method of commit, use the **rollback ?** command.

```
[edit]
user@host# rollback ?
Possible completions:
<[Enter]> Execute this command
<number> Numeric argument
0                    2018-02-27 12:52:10 PST by abc via cli
1                    2018-02-26 14:47:42 PST by def via cli
2                    2018-02-14 21:55:45 PST by ghi via cli
3                    2018-02-10 16:11:30 PST by jkl via cli
4                    2018-02-10 16:02:35 PST by mno via cli
5                    2018-03-16 15:10:41 PST by pqr via cli
6                   2018-03-16 14:54:21 PST by stu via cli
```

```
7                    2018-03-16 14:51:38 PST by vwx via cli
8                     2018-03-16 14:43:29 PST by yzz via cli
9                    2018-03-16 14:15:37 PST by abc via cli
  10                   2018-03-16 14:13:57 PST by def via cli
  11                    2018-03-16 12:57:19 PST by root via other
  12                     2018-03-16 10:45:23 PST by root via other
  13                    2018-03-16 10:08:13 PST by root via other
  14                   2018-03-16 01:20:56 PST by root via other
  15                   2018-03-16 00:40:37 PST by ghi via cli
  16                    2018-03-16 00:39:29 PST by jkl via cli
  17                    2018-03-16 00:32:36 PST by mno via cli
  18                     2018-03-16 00:31:17 PST by pqr via cli
  19                    2018-03-15 19:59:00 PST by stu via cli
  20                    2018-03-15 19:53:39 PST by vwx via cli
  21                    2018-03-15 18:07:19 PST by yzz via cli
  22                      2018-03-15 17:59:03 PST by abc via cli
  23                     2018-03-15 15:05:14 PST by def via cli
  24                    2018-03-15 15:04:51 PST by ghi via cli
  25                   2018-03-15 15:03:42 PST by jkl via cli
  26                    2018-03-15 15:01:52 PST by mno via cli
  27                   2018-03-15 14:58:34 PST by pqr via cli
  28                   2018-03-15 13:09:37 PST by root via other
  29                    2018-03-12 11:01:20 PST by stu via cli
  30                     2018-03-12 10:57:35 PST by vwx via cli
  31                    2018-03-11 10:25:07 PST by yzz via cli
  32                    2018-03-10 23:40:58 PST by abc via cli
  33                    2018-03-10 23:40:38 PST by def via cli
  34                    2018-03-10 23:14:27 PST by ghi via cli
  35                    2018-03-10 23:10:16 PST by jkl via cli
  36                   2018-03-10 23:01:51 PST by mno via cli
  37                   2018-03-10 22:49:57 PST by pqr via cli
  38                    2018-03-10 22:24:07 PST by stu via cli
  39                   2018-03-10 22:20:14 PST by vwx via cli
  40                    2018-03-10 22:16:56 PST by yzz via cli
  41                   2018-03-10 22:16:41 PST by abc via cli
  42                   2018-03-10 20:44:00 PST by def via cli
  43                   2018-03-10 20:43:29 PST by ghi via cli
  44                    2018-03-10 20:39:14 PST by jkl via cli
  45                   2018-03-10 20:31:30 PST by root via other
  46                    2018-03-10 18:57:01 PST by mno via cli
  47                    2018-03-10 18:56:18 PST by pqr via cli
  48                    2018-03-10 18:47:49 PST by stu via cli
  49                   2018-03-10 18:47:34 PST by vw via cli
```

```
  | Pipe through a command
[edit]
```

Comparing Configuration Changes with a Prior Version

In configuration mode only, when you have made changes to the configuration and want to compare the candidate configuration with a prior version, you can use the **compare** command to display the configuration. The **compare** command compares the candidate configuration with either the current committed configuration or a configuration file and displays the differences between the two configurations. To compare configurations, specify the **compare** command after the pipe:

```
[edit]
user@host# show | compare (filename| rollback n)
```

- ***filename*** is the full path to a configuration file. The file must be in the proper format: a hierarchy of statements.
- ***n*** is the index into the list of previously committed configurations. The most recently saved configuration is number 0, and the oldest saved configuration is number 49. If you do not specify arguments, the candidate configuration is compared against the active configuration file (**/config/juniper.conf**).

The comparison output uses the following conventions:

- Statements that are only in the candidate configuration are prefixed with a plus sign (+).
- Statements that are only in the comparison file are prefixed with a minus sign (-).
- Statements that are unchanged are prefixed with a single blank space ( ).

The following example shows various changes, then a comparison of the candidate configuration with the active configuration, showing only the changes made at the **[edit protocols bgp]** hierarchy level:

```
[edit]
user@host# edit protocols bgp
[edit protocols bgp]
user@host# show
group my-group {
    type internal;
    hold-time 60;
    advertise-inactive;
    allow 10.1.1.1/8;
}
```

```
group fred {
    type external;
    peer-as 33333;
    allow 10.2.2.2/8;
}
group test-peers {
    type external;
    allow 10.3.3.3/8;
}
[edit protocols bgp]
user@host# set group my-group hold-time 90
[edit protocols bgp]
user@host# delete group my-group advertise-inactive
[edit protocols bgp]
user@host# set group fred advertise-inactive
[edit protocols bgp]
user@host# delete group test-peers
[edit protocols bgp]
user@host# show | compare
[edit protocols bgp group my-group]
-hold-time 60;
+hold-time 90;
-advertise-inactive;
[edit protocols bgp group fred]
+advertise-inactive;
[edit protocols bgp]
-group test-peers {
    -type external;
    -allow 10.3.3.3/8;
}
[edit protocols bgp]
user@host# show
group my-group {
    type internal;
    hold-time 90;
    allow 10.1.1.1/8;
}
group fred {
    type external;
    advertise-inactive;
    peer-as 3333;
```

```
    allow 10.2.2.2/8;
}
```

RELATED DOCUMENTATION

Loading a Configuration from a File or the Terminal

Viewing Files and Directories on a Juniper Networks Device

Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

NOTE: This process clears prior committed configuration parameters, except for those which preserve a Virtual Chassis configuration. This is how you can restore the factory default configuration on a Virtual Chassis (multiple devices configured to work together that look like a single device) without removing anything needed to keep the Virtual Chassis working.

RELATED DOCUMENTATION

Understanding Configuration Files

Reverting to the Rescue Configuration

Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a device and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```

2. Commit your changes.

```
[edit]
user@switch# commit filename
```

RELATED DOCUMENTATION

Reverting to the Default Factory Configuration

Recovering the Root Password for Switches

If you forget the root password, you can use the password recovery procedure to reset the root password.

NOTE: The root password cannot be recovered on a QFabric system.

NOTE: You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the device's power supplies.

2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the device into the RJ-45–to–DB-9 serial port adapter supplied with the device.
4. Plug the RJ-45–to–DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the device.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
   - Bits per second: 9600
   - Data bits: 8
   - Parity: None
   - Stop bits: 1
   - Flow control: None
9. Power on the device by (if necessary) plugging the power cords into the device's power supply, or turning on the power to the device by switching on the AC power outlet the device is plugged into.

   The terminal emulation screen on your management device displays the device's boot sequence.
10. When the following prompt appears, press the Spacebar to access the device's bootstrap loader command prompt:

```
Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [kernel] in 9 seconds...
```

11. At the following prompt, enter **boot -s** to start up the system in single-user mode.

```
ok boot -s
```

12. At the following prompt, enter **recovery** to start the root password recovery procedure.

```
Enter full pathname of shell or 'recovery' for root password recovery or
RETURN for /bin/sh: recovery
```

13. Enter configuration mode in the CLI.

14. Set the root password. For example:

```
user@switch# set system root-authentication plain-text-password
```

15. At the following prompt, enter the new root password. For example:

```
New password: ABC123
```

```
Retype new password:
```

16. At the second prompt, reenter the new root password.
17. After you have finished configuring the password, commit the configuration.

```
root@host# commit
commit complete
```

18. Exit configuration mode in the CLI.
19. Exit operational mode in the CLI.
20. At the prompt, enter **y** to reboot the device.

```
Reboot the system? [y/n] y
```

RELATED DOCUMENTATION

- Configuring the Root Password

Troubleshooting a Deprecated Network Analytics Configuration

IN THIS SECTION

- Problem | 127
- Cause | 127

Solution | 127

Problem

Description

After a software upgrade to Junos OS Release 13.2X51-D15 from an earlier release, the network analytics configuration is no longer valid and the feature is disabled.

Symptoms

The network analytics configuration used in Junos OS Release 13.2X51-D10 has been deprecated in Release 13.2X51-D15. Issuing the **show services analytics** command results in the following output:

```
root@qfx5100# show services analytics

queue-statistics { ## Warning: 'queue-statistics' is deprecated
    interval 1;
}
```

Cause

Junos OS Release 13.2X51-D15 added enhancements to the network analytics feature, resulting in significant changes in the CLI. The updated **[edit services analytics]** hierarchy level contains some statements that have replaced those that were previously released. As a result, the earlier configuration does not work in the new release.

Solution

Use the new CLI statements to reconfigure the network analytics feature.

RELATED DOCUMENTATION

Network Analytics Overview

analytics

CHAPTER 12

Security

IN THIS CHAPTER

- Troubleshooting Firewall Filter Configuration | **128**
- Troubleshooting Policer Configuration | **140**

Troubleshooting Firewall Filter Configuration

IN THIS SECTION

- Firewall Filter Configuration Returns a No Space Available in TCAM Message | **129**
- Filter Counts Previously Dropped Packet | **131**
- Matching Packets Not Counted | **132**
- Counter Reset When Editing Filter | **133**
- Cannot Include loss-priority and policer Actions in Same Term | **134**
- Cannot Egress Filter Certain Traffic Originating on QFX Switch | **134**
- Firewall Filter Match Condition Not Working with Q-in-Q Tunneling | **135**
- Egress Firewall Filters with Private VLANs | **136**
- Egress Filtering of L2PT Traffic Not Supported | **137**
- Cannot Drop BGP Packets in Certain Circumstances | **137**
- Invalid Statistics for Policer | **138**
- Policers can Limit Egress Filters | **138**

Use the following information to troubleshoot your firewall filter configuration.

Firewall Filter Configuration Returns a No Space Available in TCAM Message

IN THIS SECTION

- Problem | 129
- Solution | 129

Problem

Description

When a firewall filter configuration exceeds the amount of available Ternary Content Addressable Memory (TCAM) space, the system returns the following **syslogd** message:

```
No space available in tcam.
Rules for filter filter-name will not be installed.
```

A switch returns this message during the commit operation if the firewall filter that has been applied to a port, VLAN, or Layer 3 interface exceeds the amount of space available in the TCAM table. The filter is not applied, but the commit operation for the firewall filter configuration is completed in the CLI module.

Solution

When a firewall filter configuration exceeds the amount of available TCAM table space, you must configure a new firewall filter with fewer filter terms so that the space requirements for the filter do not exceed the available space in the TCAM table.

You can perform either of the following procedures to correct the problem:

To delete the filter and its binding and apply the new smaller firewall filter to the same binding:

1. Delete the filter and its binding to ports, VLANs, or Layer 3 interfaces. For example:

```
[edit]
user@switch# delete firewall family ethernet-switching filter ingress-vlan-rogue-block
user@switch# delete vlans employee-vlan description "filter to block rogue devices on employee-vlan"
user@switch# delete vlans employee-vlan filter input ingress-vlan-rogue-block
```

2. Commit the changes:

```
[edit]
user@switch# commit
```

3. Configure a smaller filter with fewer terms that does not exceed the amount of available TCAM space. For example:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-ingress-vlan-rogue-block ...
```

4. Apply (bind) the new firewall filter to a port, VLAN , or Layer 3 interface. For example:

```
[edit]
user@switch# set vlans employee-vlan description "filter to block rogue devices on employee-vlan"
user@switch# set vlans employee-vlan filter input new-ingress-vlan-rogue-block
```

5. Commit the changes:

```
[edit]
user@switch# commit
```

To apply a new firewall filter and overwrite the existing binding but not delete the original filter:

1. Configure a firewall filter with fewer terms than the original filter:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-ingress-vlan-rogue-block...
```

2. Apply the firewall filter to the port, VLAN, or Layer 3 interfaces to overwrite the binding of the original filter—for example:

```
[edit]
user@switch# set vlans employee-vlan description "smaller filter to block rogue devices on employee-vlan"
user@switch# set vlans employee-vlan filter input new-ingress-vlan-rogue-block
```

Because you can apply no more than one firewall filter per VLAN per direction, the binding of the original firewall filter to the VLAN is overwritten with the new firewall filter **new-ingress-vlan-rogue-block**.

3. Commit the changes:

```
[edit]
user@switch# commit
```

NOTE: The original filter is not deleted and is still available in the configuration.

Filter Counts Previously Dropped Packet

IN THIS SECTION

- Problem | 131
- Solution | 132

Problem

Description

If you configure two or more filters in the same direction for a physical interface and one of the filters includes a counter, the counter will be incorrect if the following circumstances apply:

- You configure the filter that is applied to packets first to discard certain packets. For example, imagine that you have a VLAN filter that accepts packets sent to 10.10.1.0/24 addresses and implicitly discards packets sent to any other addresses. You apply the filter to the **admin** VLAN in the output direction, and interface xe-0/0/1 is a member of that VLAN.
- You configure a subsequent filter to accept and count packets that are dropped by the first filter. In this example, you have a port filter that accepts and counts packets sent to 192.168.1.0/24 addresses that is also applied to xe-0/0/1 in the output direction.

The egress VLAN filter is applied first and correctly discards packets sent to 192.168.1.0/24 addresses. The egress port filter is applied next and counts the discarded packets as matched packets. The packets are not forwarded, but the counter displayed by the egress port filter is incorrect.

Remember that the order in which filters are applied depends on the direction in which they are applied, as indicated here:

Ingress filters:

1. Port (Layer 2) filter
2. VLAN filter
3. Router (Layer 3) filter

Egress filters:

1. Router (Layer 3) filter
2. VLAN filter
3. Port (Layer 2) filter

Solution

This is expected behavior.

Matching Packets Not Counted

IN THIS SECTION

- Problem | 132
- Solution | 133

Problem

Description

If you configure two egress filters with counters for a physical interface and a packet matches both of the filters, only one of the counters includes that packet.

For example:

- You configure an egress port filter with a counter for interface xe-0/0/1.
- You configure an egress VLAN filter with a counter for the **admin**VLAN, and interface xe-0/0/1 is a member of that VLAN.

- A packet matches both filters.

In this case, the packet is counted by only one of the counters even though it matched both filters.

Solution

This is expected behavior.

Counter Reset When Editing Filter

IN THIS SECTION

- Problem | 133
- Solution | 133

Problem

Description

If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

Solution

This is expected behavior.

Cannot Include loss-priority and policer Actions in Same Term

IN THIS SECTION

- Problem | 134
- Solution | 134

Problem

Description

You cannot include both of the following actions in the same firewall filter term in a QFX Series switch:

- **loss-priority**
- **policer**

If you do so, you see the following error message when you attempt to commit the configuration: “cannot support policer action if loss-priority is configured.”

Solution

This is expected behavior.

Cannot Egress Filter Certain Traffic Originating on QFX Switch

IN THIS SECTION

- Problem | 135
- Solution | 135

Problem

Description

On a QFX Series switch, you cannot filter certain traffic with a firewall filter applied in the output direction if the traffic originates on the QFX switch. This limitation applies to control traffic for protocols such as ICMP (ping), STP, LACP, and so on.

Solution

This is expected behavior.

Firewall Filter Match Condition Not Working with Q-in-Q Tunneling

IN THIS SECTION

- Problem | 135
- Solution | 135

Problem

Description

If you create a firewall filter that includes a match condition of **dot1q-tag** or **dot1q-user-priority** and apply the filter on input to a trunk port that participates in a service VLAN, the match condition does not work if the Q-in-Q EtherType is not 0x8100. (When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network and therefore participate in service VLANs.)

Solution

This is expected behavior. To set the Q-in-Q EtherType to 0x8100, enter the **set dot1q-tunneling ethertype 0x8100** statement at the **[edit ethernet-switching-options]** hierarchy level. You must also configure the other end of the link to use the same Ethertype.

Egress Firewall Filters with Private VLANs

IN THIS SECTION

- Problem | 136
- Solution | 137

Problem

Description

If you apply a firewall filter in the output direction to a primary VLAN, the filter also applies to the secondary VLANs that are members of the primary VLAN when the traffic egresses with the primary VLAN tag or isolated VLAN tag, as listed below:

- Traffic forwarded from a secondary VLAN trunk port to a promiscuous port (trunk or access)
- Traffic forwarded from a secondary VLAN trunk port that carries an isolated VLAN to a PVLAN trunk port.
- Traffic forwarded from a promiscuous port (trunk or access) to a secondary VLAN trunk port
- Traffic forwarded from a PVLAN trunk port. to a secondary VLAN trunk port
- Traffic forwarded from a community port to a promiscuous port (trunk or access)

If you apply a firewall filter in the output direction to a primary VLAN, the filter does *not* apply to traffic that egresses with a community VLAN tag, as listed below:

- Traffic forwarded from a community trunk port to a PVLAN trunk port
- Traffic forwarded from a secondary VLAN trunk port that carries a community VLAN to a PVLAN trunk port
- Traffic forwarded from a promiscuous port (trunk or access) to a community trunk port
- Traffic forwarded from a PVLAN trunk port. to a community trunk port

If you apply a firewall filter in the output direction to a community VLAN, the following behaviors apply:

- The filter is applied to traffic forwarded from a promiscuous port (trunk or access) to a community trunk port (because the traffic egresses with the community VLAN tag).

- The filter is applied to traffic forwarded from a community port to a PVLAN trunk port (because the traffic egresses with the community VLAN tag).
- The filter is *not* applied to traffic forwarded from a community port to a promiscuous port (because the traffic egresses with the primary VLAN tag or untagged).

Solution

These are expected behaviors. They occur only if you apply a firewall filter to a private VLAN in the output direction and do not occur if you apply a firewall filter to a private VLAN in the input direction.

Egress Filtering of L2PT Traffic Not Supported

IN THIS SECTION

- Problem | 137
- Solution | 137

Problem

Description

Egress filtering of L2PT traffic is not supported on the QFX3500 switch. That is, if you configure L2PT to tunnel a protocol on an interface, you cannot also use a firewall filter to filter traffic for that protocol on that interface in the output direction. If you commit a configuration for this purpose, the firewall filter is not applied to the L2PT-tunneled traffic.

Solution

This is expected behavior.

Cannot Drop BGP Packets in Certain Circumstances

IN THIS SECTION

- Problem | 138
- Solution | 138

Problem

Description

BGP packets with a time-to-live (TTL) value greater than 1 cannot be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface. BGP packets with TTL value of 1 or 0 can be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface.

Solution

This is expected behavior.

Invalid Statistics for Policer

IN THIS SECTION

- Problem | **138**
- Solution | **138**

Problem

Description

If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

Solution

This is expected behavior.

Policers can Limit Egress Filters

IN THIS SECTION

- Problem | **139**
- Solution | **139**

Problem

Description

On some switches, the number of egress policers you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that take up two entries in a 1024-entry TCAM. These are used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you are unable to commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters get used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

Solution

You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

RELATED DOCUMENTATION

Understanding FIP Snooping, FBF, and MVR Filter Scalability

Configuring Firewall Filters

Troubleshooting Policer Configuration

IN THIS SECTION

- Incomplete Count of Packet Drops | **140**
- Counter Reset When Editing Filter | **141**
- Invalid Statistics for Policer | **142**
- Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured | **142**
- Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured | **144**
- Policers Can Limit Egress Filters | **144**

Incomplete Count of Packet Drops

IN THIS SECTION

- Problem | **141**
- Solution | **141**

Problem

Description

Under certain circumstances, Junos OS might display a misleading number of packets dropped by an ingress policer.

If packets are dropped because of ingress admission control, policer statistics might not show the number of packet drops you would expect by calculating the difference between ingress and egress packet counts. This might happen if you apply an ingress policer to multiple interfaces, and the aggregate ingress rate of those interfaces exceeds the line rate of a common egress interface. In this case, packets might be dropped from the ingress buffer. These drops are not included in the count of packets dropped by the policer, which causes policer statistics to underreport the total number of drops.

Solution

This is expected behavior.

Counter Reset When Editing Filter

IN THIS SECTION

- Problem | **141**
- Solution | **142**

Problem

Description

If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

Solution

This is expected behavior.

Invalid Statistics for Policer

IN THIS SECTION

- Problem | 142
- Solution | 142

Problem

Description

If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

Solution

This is expected behavior.

Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

IN THIS SECTION

- Problem | 142
- Solution | 143

Problem

Description

If you configure a policer to rate-limit throughput and apply it on egress to multiple interfaces on a QFX3500 switch or Node, the measured aggregate policed rate might be twice the configured rate,

depending on which interfaces you apply the policer to. The doubling of the policed rate occurs if you apply a policer to multiple interfaces and *both* of the following are true:

- There is at least one policed interface in the range xe-0/0/0 to xe-0/0/23 or the range xe-0/1/1 to xe-0/1/7.
- There is at least one policed interface in the range xe-0/0/24 to xe-0/0/47 or the range xe-0/1/8 to xe-0/1/15.

For example, if you configure a policer to rate-limit traffic at 1 Gbps and apply the policer (by using a firewall filter) to xe-0/0/0 and xe-0/0/24 in the output direction, each interface is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps. The same behavior occurs if you apply the policer to xe-0/1/1 and xe-0/0/24—each interface is rate-limited at 1 Gbps.

If you apply the same policer on egress to multiple interfaces in these groups, each *group* is rate-limited at 1 Gbps. For example, if you apply the policer to xe-0/0/0 through xe-0/0/4 (five interfaces) and xe-0/0/24 through xe-0/0/33 (ten interfaces), each group is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps.

Here is another example: If you apply the policer to xe-0/0/0 through xe-0/0/4 and xe-0/1/1 through xe-0/1/5 (a total of ten interfaces), that group is rate-limited at 1 Gbps in aggregate. If you also apply the policer to xe-0/0/24, that one interface is rate-limited at 1 Gbps while the other ten are still rate-limited at 1 Gbps in aggregate.

Interfaces xe-0/1/1 through xe-0/1/15 are physically located on the QSFP+ uplink ports, according to the following scheme:

- xe-0/1/1 through xe-0/1/3 are on Q0.
- xe-0/1/4 through xe-0/1/7 are on Q1.
- xe-0/1/8 through xe-0/1/11 are on Q2.
- xe-0/1/2 through xe-0/1/15 are on Q3.

The doubling of the policed rate occurs only if the policer is applied in the output direction. If you configure a policer as described above but apply it in the input direction, the total allowed throughput for all interfaces is 1 Gbps.

Solution

This is expected behavior.

Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

IN THIS SECTION

- Problem | 144
- Solution | 144

Problem

Description

You can configure policers to be filter-specific. This means that Junos OS creates only one policer instance no matter how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in ternary content addressable memory (TCAM). If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

Solution

To prevent this unexpected behavior, use the information about TCAM slices presented in Planning the Number of Firewall Filters to Create to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

Policers Can Limit Egress Filters

IN THIS SECTION

- Problem | 145
- Solution | 145

Problem

Description

On some switches, the number of egress policers you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that take up two entries in a 1024-entry TCAM. These are used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you are unable to commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters get used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

Solution

You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

SEE ALSO

- Overview of Policers
- Example: Using Policers to Manage Oversubscription
- Example: Using Two-Color Policers and Prefix Lists

CHAPTER 13

Services

IN THIS CHAPTER

- Troubleshooting Port Mirroring | 147

Troubleshooting Port Mirroring

IN THIS SECTION

- Egress Port Mirroring with VLAN Translation | 147
- Egress Port Mirroring with Private VLANs | 148

Egress Port Mirroring with VLAN Translation

IN THIS SECTION

- Problem | 147
- Solution | 148

Problem

Description

If you create a port-mirroring configuration that mirrors customer VLAN (CVLAN) traffic on egress and the traffic undergoes VLAN translation before being mirrored, the VLAN translation does not apply to the mirrored packets. That is, the mirrored packets retain the service VLAN (SVLAN) tag that should be

replaced by the CVLAN tag on egress. The original packets are unaffected—on these packets VLAN translation works properly, and the SVLAN tag is replaced with the CVLAN tag on egress.

Solution

This is expected behavior.

SEE ALSO

Understanding Q-in-Q Tunneling and VLAN Translation

Egress Port Mirroring with Private VLANs

IN THIS SECTION

Problem | 148
Solution | 149

Problem

Description

If you create a port-mirroring configuration that mirrors private VLAN (PVLAN) traffic on egress, the mirrored traffic (the traffic that is sent to the analyzer system) has the VLAN tag of the ingress VLAN instead of the egress VLAN. For example, assume the following PVLAN configuration:

- Promiscuous trunk port that carries primary VLANs pvlan100 and pvlan400.
- Isolated access port that carries secondary VLAN isolated200. This VLAN is a member of primary VLAN pvlan100.
- Community port that carries secondary VLAN comm300. This VLAN is also a member of primary VLAN pvlan100.
- Output interface (monitor interface) that connects to the analyzer system. This interface forwards the mirrored traffic to the analyzer.

If a packet for pvlan100 enters on the promiscuous trunk port and exits on the isolated access port, the original packet is untagged on egress because it is exiting on an access port. However, the mirror copy retains the tag for pvlan100 when it is sent to the analyzer.

Here is another example: If a packet for comm300 ingresses on the community port and egresses on the promiscuous trunk port, the original packet carries the tag for pvlan100 on egress, as expected. However, the mirrored copy retains the tag for comm300 when it is sent to the analyzer.

Solution

This is expected behavior.

CHAPTER 14

Traffic Management

IN THIS CHAPTER

- Troubleshooting Dropped FCoE Traffic | 150
- Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth | 154
- Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth | 155
- Troubleshooting Egress Queue Bandwidth Impacted by Congestion | 157
- Troubleshooting an Unexpected Rewrite Value | 158
- Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic | 160

Troubleshooting Dropped FCoE Traffic

IN THIS SECTION

- Problem | 150
- Cause | 150
- Solution | 152

Problem

Description

Fibre Channel over Ethernet (FCoE) traffic for which you want guaranteed delivery is dropped.

Cause

There are several possible causes of dropped FCoE traffic (the list numbers of the possible causes correspond to the list numbers of the solutions in the *Solution* section.):

1. Priority-based flow control (PFC) is not enabled on the FCoE priority (IEEE 802.1p code point) in both the input and output stanzas of the congestion notification profile.

2. The FCoE traffic is not classified correctly at the ingress interface. FCoE traffic should either use the default **fcoe** forwarding class and classifier configuration (maps the **fcoe** forwarding class to IEEE 802.1p code point 011) or be mapped to a lossless forwarding class and to the code point enabled for PFC on the input and output interfaces.

3. The congestion notification profile that enables PFC on the FCoE priority is not attached to the interface.

4. The forwarding class set (priority group) used for guaranteed delivery traffic does not include the forwarding class used for FCoE traffic.

   NOTE: This issue can occur only on switches that support enhanced transmission selection (ETS) hierarchical port scheduling. (Direct port scheduling does not use forwarding class sets.)

5. Insufficient bandwidth has been allocated for the FCoE queue or for the forwarding class set to which the FCoE queue belongs.

   NOTE: This issue can occur for forwarding class sets only on switches that support ETS hierarchical port scheduling. (Direct port scheduling does not use forwarding class sets.)

6. If you are using Junos OS Release 12.2, the **fcoe** forwarding class has been explicitly configured instead of using the default **fcoe** forwarding class configuration (forwarding-class-to-queue mapping).

   NOTE: If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless **fcoe** and **no-loss** forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

7. If you are using Junos OS Release 12.3 or later and you are not using the default **fcoe** forwarding class configuration, the forwarding class used for FCoE is not configured with the **no-loss** packet drop attribute. In Junos OS 12.3 or later, explicit forwarding classes configurations must include the **no-loss** packet drop attribute to be treated as lossless forwarding classes.

Solution

The list numbers of the possible solutions correspond to the list numbers of the causes in the *Cause* section.

1. Check the congestion notification profile (CNP) to see if PFC is enabled on the FCoE priority (the correct IEEE 802.1p code point) on both input and output interfaces. Use the **show class-of-service congestion-notification** operational command to show the code points that are enabled for PFC in each CNP.

   If you are using the default configuration, FCoE traffic is mapped to code point 011 (priority 3). In this case, the input stanza of the CNP should show that PFC is enabled on code point 011, and the output stanza should show that priority 011 is mapped to flow control queue 3.

   If you explicitly configured a forwarding class for FCoE traffic, ensure that:

   - You specified the **no-loss** packet drop attribute in the forwarding class configuration
   - The code point mapped to the FCoE forwarding class in the ingress classifier is the code point enabled for PFC in the CNP input stanza
   - The code point and output queue used for FCoE traffic are mapped to each other in the CNP output stanza (if you are not using the default priority and queue, you must explicitly configure each output queue that you want to respond to PFC messages)

   For example, if you explicitly configure a forwarding class for FCoE traffic that is mapped to output queue 5 and to code point 101 (priority 5), the output of the **show class-of-service congestion-notification** looks like:

```
Name: fcoe_p5_cnp, Index: 12183
Type: Input
Cable Length: 100 m
  Priority     PFC          MRU
  000          Disabled
  001          Disabled
  010          Disabled
  011          Disabled
  100          Disabled
  101          Enabled      2500
  110          Disabled
  111          Disabled
Type: Output
  Priority     Flow-Control-Queues
```

```
101
                5
```

2. Use the **show class-of-service classifier type ieee-802.1p** operational command to check if the classifier maps the forwarding class used for FCoE traffic to the correct IEEE 802.1p code point.

3. Ensure that the congestion notification profile and classifier are attached to the correct ingress interface. Use the operational command **show configuration class-of-service interfaces *interface-name***.

4. Check that the forwarding class set includes the forwarding class used for FCoE traffic. Use the operational command **show configuration class-of-service forwarding-class-sets** to show the configured priority groups and their forwarding classes.

5. Verify the amount of bandwidth allocated to the queue mapped to the FCoE forwarding class and to the forwarding class set to which the FCoE traffic queue belongs. Use the **show configuration class-of-service schedulers *scheduler-name*** operational command (specify the scheduler for FCoE traffic as the ***scheduler-name***) to see the minimum guaranteed bandwidth (**transmit-rate**) and maximum bandwidth (**shaping-rate**) for the queue.

   Use the **show configuration class-of-service traffic-control-profiles *traffic-control-profile*** operational command (specify the traffic control profile used for FCoE traffic as the ***traffic-control-profile***) to see the minimum guaranteed bandwidth (**guaranteed-rate**) and maximum bandwidth (**shaping-rate**) for the forwarding class set.

6. Delete the explicit FCoE forwarding-class-to-queue mapping so that the system uses the default FCoE forwarding-class-to-queue mapping. Include the **delete forwarding-classes class fcoe queue-num 3** statement at the **[edit class-of-service]** hierarchy level to remove the explicit configuration. The system then uses the default configuration for the FCoE forwarding class and preserves the lossless treatment of FCoE traffic.

7. Use the **show class-of-service forwarding-class** operational command to display the configured forwarding classes. The *No-Loss* column shows whether lossless transport is enabled or disabled for each forwarding class. If the forwarding class used for FCoE traffic is not enabled for lossless transport, include the **no-loss** packet drop attribute in the forwarding class configuration (**set class-of-service forwarding-classes class *fcoe-forwarding-class-name* queue-num *queue-number* no-loss**).

See *Example: Configuring CoS PFC for FCoE Traffic* for step-by-step instructions on how to configure PFC for FCoE traffic, including classifier, interface, congestion notification profile, PFC, and bandwidth scheduling configuration.

RELATED DOCUMENTATION

show class-of-service congestion-notification

Configuring CoS PFC (Congestion Notification Profiles)

Example: Configuring CoS PFC for FCoE Traffic

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth

IN THIS SECTION

- Problem | 154
- Cause | 154
- Solution | 155

Problem

Description

The maximum bandwidth of a queue when measured at the egress port exceeds the maximum bandwidth rate shaper (**shaping-rate** statement on QFX5200, QFX5100, EX4600, QFX3500, QFX3600, and OCX1100 switches, and on QFabric systems, and **transmit-rate (rate | percentage *percent* exact** statement on QFX10000 switches) configured for the queue.

Cause

When you configure bandwidth for a queue (forwarding class) or a priority group (forwarding class set), the switch accounts for the configured bandwidth as data only. The switch does not rate-shape the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its maximum bandwidth calculations.

The measured egress bandwidth can exceed the configured maximum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.

Solution

When you calculate the bandwidth requirements for queues on which you expect a significant amount of traffic with small packet sizes, consider the shaping rate as the maximum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined maximum data rate (shaping rate) and the preamble and IFG.

If the maximum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to the queue, reduce the shaping rate for that queue.

Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth

IN THIS SECTION

- Problem | 155
- Cause | 156
- Solution | 156

Problem

Description

The guaranteed minimum bandwidth of a queue (forwarding class) or a priority group (forwarding class set) when measured at the egress port exceeds the guaranteed minimum bandwidth configured for the queue (transmit-rate) or for the priority group (guaranteed-rate).

NOTE: On switches that support enhanced transmission selection (ETS) hierarchical scheduling, the switch allocates guaranteed minimum bandwidth first to a priority group using the guaranteed rate setting in the traffic control profile, and then allocates priority group minimum guaranteed bandwidth to forwarding classes in the priority group using the transmit rate setting in the queue scheduler.

On switches that support direct port scheduling, there is no scheduling hierarchy. The switch allocates port bandwidth to forwarding classes directly, using the transmit rate setting in the queue scheduler.

In this topic, if you are using direct port scheduling on your switch, ignore the references to priority groups and forwarding class sets (priority groups and forwarding class sets are only used for ETS hierarchical port scheduling). For direct port scheduling, only the transmit rate queue scheduler setting can cause the issue described in this topic.

Cause

When you configure bandwidth for a queue or a priority group, the switch accounts for the configured bandwidth as data only. The switch does not include the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its minimum bandwidth calculations.

The measured egress bandwidth can exceed the configured minimum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.

NOTE: For ETS, the sum of the queue transmit rates in a priority group should not exceed the guaranteed rate for the priority group. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.) For port scheduling, the sum of the queue transmit rates should not exceed the port bandwidth.

Solution

When you calculate the bandwidth requirements for queues and priority groups on which you expect a significant amount of traffic with small packet sizes, consider the transmit rate and the guaranteed rate as the minimum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined minimum data rate and the preamble and IFG.

If the minimum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to a queue or to a priority group, reduce the transmit rate for that queue and reduce the guaranteed rate of the priority group that contains the queue.

RELATED DOCUMENTATION

- transmit-rate
- Example: Configuring Minimum Guaranteed Output Bandwidth

Troubleshooting Egress Queue Bandwidth Impacted by Congestion

IN THIS SECTION

- Problem | 157
- Cause | 157
- Solution | 157

Problem

Description

Congestion on an egress port causes egress queues to receive less bandwidth than expected. Egress port congestion can impact the amount of bandwidth allocated to queues on the congested port and, in some cases, on ports that are not congested.

Cause

Egress queue congestion can cause the ingress port buffer to fill above a certain threshold and affect the flow to the queues on the egress port. One queue receives its configured bandwidth, but the other queues on the egress port are affected and do not receive their configured share of bandwidth.

Solution

The solution is to configure a drop profile to apply weighted random early detection (WRED) to the queue or queues on the congested ports.

Configure a drop profile on the queue that is receiving its configured bandwidth. This queue is preventing the other queues from receiving their expected bandwidth. The drop profile prevents the queue from affecting the other queues on the port.

To configure a WRED profile using the CLI:

1. Name the drop profile and set the drop start point, drop end point, minimum drop rate, and maximum drop rate for the drop profile:

```
[edit class-of-service]
user@switch# set drop-profile drop-profile-name interpolate fill-level percentage fill-level percentage
drop-probability 0 drop-probability percentage
```

RELATED DOCUMENTATION

- drop-profile
- Example: Configuring WRED Drop Profiles
- Example: Configuring CoS Hierarchical Port Scheduling (ETS)
- Understanding CoS WRED Drop Profiles

Troubleshooting an Unexpected Rewrite Value

IN THIS SECTION

- Problem | 158
- Cause | 159
- Solution | 159

Problem

Description

Traffic from one or more forwarding classes on an egress port is assigned an unexpected rewrite value.

NOTE: For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rules rewrite only the outer VLAN tag.

Cause

If you configure a rewrite rule for a forwarding class on an egress port, but you do not configure a rewrite rule for every forwarding class on that egress port, then the forwarding classes that do not have a configured rewrite rule are assigned random rewrite values.

For example:

1. Configure forwarding classes **fc1**, **fc2**, and **fc3**.
2. Configure rewrite rules for forwarding classes **fc1** and **fc2**, but not for forwarding class **fc3**.
3. Assign forwarding classes **fc1**, **fc2**, and **fc3** to a port.

When traffic for these forwarding classes flows through the port, traffic for forwarding classes **fc1** and **fc2** is rewritten correctly. However, traffic for forwarding class **fc3** is assigned a random rewrite value.

Solution

If any forwarding class on an egress port has a configured rewrite rule, then all forwarding classes on that egress port must have a configured rewrite rule. Configuring a rewrite rule for any forwarding class that is assigned a random rewrite value solves the problem.

TIP: If you want the forwarding class to use the same code point value assigned to it by the ingress classifier, specify that value as the rewrite rule value. For example, if a forwarding class has the IEEE 802.1 ingress classifier code point value **011**, configure a rewrite rule for that forwarding class that uses the IEEE 802.1p code point value **011**.

NOTE: There are no default rewrite rules. You can bind one rewrite rule for DSCP traffic and one rewrite rule for IEEE 802.1p traffic to an interface. A rewrite rule can contain multiple forwarding-class-to-rewrite-value mappings.

1. To assign a rewrite value to a forwarding class, add the new rewrite value to the same rewrite rule as the other forwarding classes on the port:

```
[edit class-of-service rewrite-rules]
user@switch# set (dscp | ieee-802.1) rewrite-name forwarding-class class-name loss-priority priority
code-point (alias | bits)
```

For example, if the other forwarding classes on the port use rewrite values defined in the rewrite rule **custom-rw**, the forwarding class **be2** is being randomly rewritten, and you want to use IEEE 802.1 code point **002** for the **be2** forwarding class:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 custom-rw forwarding-class be2 loss-priority low code-point 002
```

2. Enable the rewrite rule on an interface if it is not already enabled on the desired interface:

```
[edit]
user@switch# set class-of-service interfaces interface-name unit unit rewrite-rules (dscp | ieee-802.1)
rewrite-rule-name
```

For example, to enable the rewrite rule **custom-rw** on interface **xe-0/0/24.0**:

```
[edit]
user@switch# set class-of-service interfaces xe-0/0/24 unit 0 rewrite-rules ieee-802.1 custom-rw
```

RELATED DOCUMENTATION

- interfaces
- rewrite-rules
- Defining CoS Rewrite Rules
- Monitoring CoS Rewrite Rules

Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic

IN THIS SECTION

Problem | 161

Cause | 161
Solution | 161

Problem

Description

In QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets.

Cause

Failure of a queue to transmit packets for 12 consecutive seconds may be due to:

- A strict-high priority queue consuming all of the port bandwidth
- Several queues consuming all of the port bandwidth
- Any queue or port receiving continuous priority-based flow control (PFC) or 802.3x Ethernet PAUSE messages (received PFC and PAUSE messages prevent a queue or a port, respectively, from transmitting packets because of network congestion)
- Other conditions that prevent a queue from obtaining port bandwidth for 12 consecutive seconds

Solution

If the cause is a strict-high priority queue or other queues consuming all of the port bandwidth, you can use rate shaping to configure a maximum rate for the queues that are using all of the port bandwidth and preventing other queues from obtaining bandwidth on the port. You configure a maximum rate by creating a scheduler, using a scheduler map to apply it to a forwarding class (which maps to an output queue), and applying the scheduler map to the port using a forwarding class set and a traffic control profile.

To configure rate shaping using the CLI:

1. Name the existing scheduler or create a scheduler and define the maximum bandwidth as a rate or as a percentage:

```
      [edit class-of-service]
user@switch# set schedulers scheduler-name shaping-rate (rate | percent
percentage)
```

2. Configure a scheduler map to associate the scheduler with the forwarding class (queue) that is consuming all of the port bandwidth:

```
      [edit class-of-service]
user@switch# set scheduler-maps scheduler-map-name forwarding-class forwarding-class-name
scheduler scheduler-name
```

3. Associate the scheduler map with a traffic control profile:

```
[edit class-of-service]
user@switch# set traffic-control-profiles traffic-control-profile-name scheduler-map scheduler-map-
name
```

4. Associate the traffic control profile (and thus the scheduler map that contains the rate shaping queue scheduler) with a forwarding class set and apply them to the interface that is being reset:

```
[edit class-of-service]
user@switch# set interfaces interface-name forwarding-class-set fc-set-name output-traffic-control-
profile traffic-control-profile-name
```

For example, a strict-high priority queue is using all of the bandwidth on interface **shpnode:xe-0/0/10** and preventing other queues from transmitting for 12 consecutive seconds. You decide to set a maximum rate of 7 Gbps on the strict-high priority queue to ensure that at least 3 Gbps of the port bandwidth is available to service other queues. Table 10 on page 163 shows the topology for this example:

Table 10: Components of the Rate Shaping Troubleshooting Example

Component	Settings
Affected interface	**shpnode:xe-0/0/10**
Scheduler (strict-high priority scheduler)	Name: **shp-sched** Shaping rate: **7g** Priority: **strict-high** **NOTE**: This example assumes that the scheduler already exists and has been configured as **strict-high** priority, but that rate shaping to prevent the strict-high priority traffic from using all of the port bandwidth has not been applied.
Scheduler map	Name: **shp-map** Forwarding class to associate with the **shp-sched** scheduler: **strict-high** **NOTE**: This example assumes that a strict-high priority forwarding class has been configured and assigned the name **strict-high**.
Traffic control profile	Name: **shp-tcp** **NOTE**: This example does not describe how to define a complete traffic control profile.
Forwarding class set	Name: **shp-pg**

To configure the scheduler, map it to the strict-high priority forwarding class, and apply it to interface **shpnode:xe-0/0/10** using the CLI:

1. Specify the scheduler for the strict-high priority queue (**shp-sched**) with a maximum bandwidth of 7 Gbps:

```
      [edit class-of-service schedulers]
user@switch# set shp-sched shaping-rate 7g
```

2. Configure a scheduler map (**shp-map**) that associates the scheduler (**shp-sched**) with the forwarding class (**strict-high**):

```
        [edit class-of-service scheduler-maps]
user@switch# set shp-map forwarding-class strict-high scheduler shp-sched
```

3. Associate the scheduler map **shp-map** with a traffic control profile (**shp-tcp**):

```
[edit class-of-service traffic-control-profiles]
user@switch# set shp-tcp scheduler-map shp-map
```

4. Associate the traffic control profile **shp-tcp** with a forwarding class set (**shp-pg**) and the affected interface (**shpnode:xe-0/0/10**):

```
[edit class-of-service]
user@switch# set interfaces shpnode:xe-0/0/10 forwarding-class-set shp-pg output-traffic-control-
profile shp-tcp
```

RELATED DOCUMENTATION

- Understanding CoS Output Queue Schedulers
- Defining CoS Queue Scheduling Priority
- Example: Configuring Queue Schedulers
- Example: Configuring Traffic Control Profiles (Priority Group Scheduling)
- Example: Configuring Forwarding Class Sets
- Example: Configuring CoS Hierarchical Port Scheduling (ETS)

CHAPTER 15

Virtual Chassis Fabric

IN THIS CHAPTER

- Troubleshooting Virtual Chassis Fabric | 165

Troubleshooting Virtual Chassis Fabric

IN THIS SECTION

- Large-Scale Virtual Chassis Fabric Becomes Unstable When Logging is Enabled | 165
- Virtual Chassis Port Link Does Not Form | 166
- QFX5100 Leaf Device Assumes Routing Engine Role | 167

This topic describes troubleshooting some common issues for a Virtual Chassis Fabric (VCF):

Large-Scale Virtual Chassis Fabric Becomes Unstable When Logging is Enabled

IN THIS SECTION

- Problem | 166
- Cause | 166
- Solution | 166

Problem

Description

When detailed system logging or trace operations are enabled in larger-scale VCFs, you observe significant impact on VCF stability, such as:

- Increased VCF convergence time
- Traffic interruption

Cause

System logging and tracing operations place a load on the primary Routing Engine device in a VCF, taking processing cycles away from managing VCF operations. Logging in general, especially higher levels of logging and tracing operations, can have an impact on VCF stability.

Solution

To help ensure good convergence and stable operation in a large-scale VCF, system logging and tracing should always be used with discretion. During normal VCF operation, system logging should be set at or below the **notice** level, and tracing options disabled. When logging or tracing is necessary to troubleshoot a particular issue, use the following guidelines to minimize impact on VCF stability:

- Use the **detail** tracing option or system logging levels at or above the **error** level only for short periods of time during troubleshooting, and disable these settings after gathering enough information to begin analyzing the issue.
- Avoid logging the same level of information to more than one log file, which adds extra processing without the benefit of providing more information. Setting up logging to different files for different levels or facilities is a better option.
- Choose remote logging rather than local logging, and avoid logging to the console.

Virtual Chassis Port Link Does Not Form

IN THIS SECTION

- Problem | **167**
- Cause | **167**
- Solution | **167**

Problem

Description

You connect a 40-Gbps QSFP+ port or a 10-Gbps SFP+ port between a leaf device and a spine device in an autoprovisioned or preprovisioned VCF. You expect the automatic Virtual Chassis port (VCP) conversion feature to convert the link into a VCP link, but the conversion doesn't occur.

The `show virtual-chassis vc-port` output indicates that the status of the interface is **Absent** or one or both of interfaces don't appear in the `show virtual-chassis vc-port` output.

Cause

If one end of a link is configured as a VCP and the other is not configured as a VCP, the VCP link does not form.

The automatic VCP conversion feature, therefore, does not work in the following situations:

- a 40-Gbps QSFP+ or 10-Gbps SFP+ interface on one end of the link is already configured as a VCP.

  If you have previously removed a device from a VCF but haven't used the **request virtual-chassis vc-port delete** command to convert the interface that was connected to the removed device out of VCP mode, the interface is still configured as a VCP.

  If you have removed a device from one Virtual Chassis or VCF and not changed the VCP port setting, the device being added to the VCF might also be configured as a VCP.

- a 40-Gbps QSFP+ port on an EX4300 switch, which is configured as a VCP by default, is interconnecting to a spine device.

Solution

Manually configure the interface that is not configured as a VCP into a VCP using the **request virtual-chassis vc-port set** command.

QFX5100 Leaf Device Assumes Routing Engine Role

IN THIS SECTION

- Problem | **168**
- Solution | **168**

Problem

Description

A QFX5100 device configured as a leaf device assumes the Routing Engine role during VCF setup. The **show virtual-chassis** output confirms the role.

Solution

The device can assume the Routing Engine role for several minutes during setup before it receives the configuration from the primary Routing Engine, but eventually returns to the linecard role with no user intervention.

RELATED DOCUMENTATION

- *Virtual Chassis Fabric Overview*
- *traceoptions*
- Junos OS System Log Configuration Statements
- Junos OS System Logging Facilities and Message Severity Levels