

QFX3000-M QFabric System Deployment Guide

Published
2019-12-13

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.

QFX3000-M QFabric System Deployment Guide
Copyright © 2019 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 the Documentation | xvii

Documentation and Release Notes | xvii

Using the Examples in This Manual | xvii

 Merging a Full Example | xviii

 Merging a Snippet | xix

Documentation Conventions | xix

Documentation Feedback | xxii

Requesting Technical Support | xxii

 Self-Help Online Tools and Resources | xxiii

 Creating a Service Request with JTAC | xxiii

1

Overview

Before You Begin | 3

QFabric System Overview | 3

 Legacy Data Center Architecture | 3

 QFX Series QFabric System Architecture | 5

Understanding QFabric System Terminology | 7

Understanding Interfaces on the QFabric System | 15

 Four-Level Interface Naming Convention | 15

 QSFP+ Interfaces | 15

 Link Aggregation | 19

Hardware Architecture Overview | 21

Understanding the QFabric System Hardware Architecture | 21

 QFabric System Hardware Architecture Overview | 21

 QFX3000-G QFabric System Features | 25

 QFX3000-M QFabric System Features | 25

Understanding the Director Group | 26

 Director Group Components | 26

 Director Group Services | 26

Understanding Routing Engines in the QFabric System | 27

- Hardware-Based Routing Engines | 27

- Software-Based External Routing Engines | 27

Understanding Interconnect Devices | 29

- Interconnect Device Introduction | 29

- QFX3008-I Interconnect Devices | 30

- QFX3600-I Interconnect Devices | 31

- QFX5100-24Q Interconnect Devices | 32

Understanding Node Devices | 33

- Node Device Introduction | 34

- QFX3500 Node Devices | 34

- QFX3600 Node Devices | 35

- QFX5100 Node Devices | 36

Understanding Node Groups | 38

- Network Node Groups | 39

- Server Node Groups | 39

Understanding Port Oversubscription on Node Devices | 40

Software Architecture Overview | 45

Understanding the QFabric System Software Architecture | 45

Understanding the Director Software | 46

Understanding Partitions | 47

- QFabric System Default Partition | 48

Understanding the QFabric System Control Plane | 49

- Control Plane Elements | 50

- Control Plane Services | 52

Understanding the QFabric System Data Plane | 53

- Data Plane Components | 53

- QFabric System Fabric | 54

Understanding QFabric Multicast Flow Groups | 55

- How Do I Implement a Flow Group in a Fabric? | 55

- Do Flow Groups Impact QFabric System Performance? | 56

- What Platforms Support Flow Groups? | 56

- How Do Multicast Flow Groups Work? | 56

What are QFabric Multicast Flow Group Limitations? | 57

Software Features | 59

Understanding Software Upgrade on the QFabric System | 59

Operational Software Commands | 60

Operational Reboot Commands | 60

Understanding Nonstop Software Upgrade for QFabric Systems | 61

Understanding Statements and Commands on the QFabric System | 66

Chassis Statements | 66

Chassis Commands | 66

Understanding NTP on the QFabric System | 68

Understanding Network Management Implementation on the QFabric System | 68

Understanding the Implementation of SNMP on the QFabric System | 69

Understanding the Implementation of System Log Messages on the QFabric System | 72

Understanding User and Access Management Features on the QFabric System | 74

Understanding QFabric System Login Classes | 75

Understanding Interfaces on the QFabric System | 76

Four-Level Interface Naming Convention | 76

QSFP+ Interfaces | 76

Link Aggregation | 80

Understanding Layer 3 Features on the QFabric System | 81

Understanding Telnet on the QFabric System | 82

Understanding Security Features on the QFabric System | 83

Understanding Port Mirroring | 84

Port Mirroring Overview | 84

Port Mirroring Instance Types | 85

Port-Mirroring Terminology | 85

Port Mirroring and STP | 87

Port Mirroring Constraints and Limitations | 88

Local and Remote Port Mirroring | 88

Remote Port Mirroring Only | 90

Port Mirroring Constraints on OCX Series Switches | 90

Understanding Fibre Channel Fabrics on the QFabric System | 91

Understanding CoS Fabric Forwarding Class Sets | 93

Default Fabric Forwarding Class Sets | 95

Fabric Forwarding Class Set Configuration and Implementation | 98

Mapping Forwarding Classes to Fabric Forwarding Class Sets | 99

Fabric Forwarding Class Set Implementation | 99

QFabric System CoS | 100

Support for Flow Control and Lossless Transport Across the Fabric | 100

Viewing Fabric Forwarding Class Set Information | 103

Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences | 105

2

Installation

Before You Begin | 109

QFX3000-M QFabric System Installation Overview | 109

Understanding QFX3000-M QFabric System Hardware Configurations | 112

Planning a QFX3000-M QFabric System Deployment | 118

General Site Guidelines | 124

Site Electrical Wiring Guidelines | 124

Environmental Requirements and Specifications for a QFX3100 Director Device | 125

Environmental Requirements and Specifications for QFX3600 and QFX3600-I Devices | 126

Environmental Requirements and Specifications for a QFX3500 Device | 128

Environmental Requirements and Specifications for a QFX5100 Device | 129

Environmental Requirements and Specifications for EX Series Switches | 130

Ports and Connectors | 137

Determining Interface Support for the QFX3600 Device | 137

Determining Interface Support for the QFX3500 Device | 138

Determining Interface Support for the QFX5100 Device | 140

Determining Transceiver Support for QFabric Systems | 143

Cable Specifications for Copper-Based Control Plane Connections for the QFabric System | 144

Power | 147

AC Power Specifications for a QFX3100 Director Device | 147

AC Power Cord Specifications for a QFX3100 Director Device | 148

AC Power Specifications for a QFX3600 or QFX3600-I Device | 149

AC Power Specifications for a QFX3500 Device	150
AC Power Specifications for a QFX5100 Device	151
AC Power Cord Specifications for a QFX Series Device	152
AC Power Supply Specifications for EX4300 Switches	154
AC Power Cord Specifications for an EX4300 Switch	155
AC Power Cord Specifications for EX4200 Switches	159
DC Power Specifications for a QFX3600 or QFX3600-I Device	161
DC Power Specifications for a QFX3500 Device	162
DC Power Specifications for a QFX5100 Device	162
DC Power Supply Specifications for EX4300 Switches	164
Installing a QFX3100 Director Device 	165
Installing and Connecting a QFX3100 Director Device	165
Unpacking a QFX3100 Director Device	166
Mounting a QFX3100 Director Device on Two Posts in a Rack or Cabinet	168
Mounting a QFX3100 Director Device on Four Posts in a Rack or Cabinet	170
Connecting AC Power to a QFX3100 Director Device	172
Connecting a QFX Series Device to a Management Console	173
Powering On a QFX3100 Director Device	175
Installing a QFX3600-I Interconnect Device or QFX3600 Node Device 	177
Installing and Connecting a QFX3600 or QFX3600-I Device	177
Unpacking a QFX3600 or QFX3600-I Device	178
Mounting a QFX3600 or QFX3600-I Device on Two Posts in a Rack or Cabinet	180
Mounting a QFX3600 or QFX3600-I Device on Four Posts in a Rack or Cabinet	182
Connecting Earth Ground to QFX3600 or QFX3600-I Devices	187
Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device	188
Connecting DC Power to a QFX3500, QFX3600, or QFX3600-I Device	191
Connecting a QFX Series Device to a Management Console	196

Installing a QFX3500 Node Device | 199

Installing and Connecting a QFX3500 Device | 199

Unpacking a QFX3500 Device | 200

Mounting a QFX3500 Device in a Rack or Cabinet | 201

Before You Begin Rack Installation | 202

Two Mounting Rails Procedure | 203

Four Mounting Rails Procedure | 205

Connecting Earth Ground to a QFX3500 Device | 207

Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device | 209

Connecting DC Power to a QFX3500, QFX3600, or QFX3600-I Device | 211

Connecting a QFX Series Device to a Management Console | 216

Installing a QFX5100 Node Device | 219

Installing and Connecting a QFX5100 Device | 219

Unpacking a QFX5100 Device | 220

Mounting a QFX5100 Device in a Rack or Cabinet | 222

Before You Begin Rack Installation | 222

Four-Post Procedure | 224

Connecting Earth Ground to a QFX5100 Device | 227

Connecting AC Power to a QFX5100 Device | 228

Connecting DC Power to a QFX5100 Device | 231

Connecting a QFX Series Device to a Management Console | 236

Cabling a Copper-Based Control Plane for the QFX3000-M QFabric System | 239

Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy | 239

Connecting QFX3100 Director Devices in a Director Group | 242

Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network | 245

Connecting a QFX3100 Director Device to a Network for Out-of-Band Management | 247

Connecting a QFX5100-24Q Interconnect Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 248

Connecting a QFX3600-I Interconnect Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 252

Connecting a QFX5100 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 256

Connecting a QFX3600 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | **259**

Connecting a QFX3500 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | **263**

Cabling a Fiber-Based Control Plane for the QFX3000-M QFabric System | 267

Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy | **267**

Connecting QFX3100 Director Devices in a Director Group | **270**

Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-M QFabric System Control Plane Network | **273**

Connecting a QFX3100 Director Device to a Network for Out-of-Band Management | **275**

Connecting a QFX5100-24Q Interconnect Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | **276**

Connecting a QFX3600-I Interconnect Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | **281**

Connecting a QFX5100 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | **286**

Connecting a QFX3500 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | **290**

Connecting a QFX3600 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | **295**

Cabling the Data Plane for the QFX3000-M QFabric System | 301

Connecting a QFX5100 Node Device to a QFX3600-I Interconnect Device | **301**

Connecting a QFX5100 Node Device to a QFX5100-24Q Interconnect Device | **304**

Connecting a QFX3600 Node Device to a QFX3600-I Interconnect Device | **306**

Connecting a QFX3600 Node Device to a QFX5100-24Q Interconnect Device | **309**

Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device | **311**

Connecting a QFX3500 Node Device to an QFX5100-24Q Interconnect Device | **313**

Configuration

Initial Setup | 319

QFabric System Initial and Default Configuration Information | 319

Converting the Device Mode for a QFabric System Component | 322

Example: Configuring EX Series Switches for the QFX3000-M QFabric System Control Plane | 329

Importing a QFX3000-M QFabric System Control Plane EX Series Switch Configuration with a USB Flash Drive | 360

Generating the MAC Address Range for a QFabric System | 362

Performing the QFabric System Initial Setup on a QFX3100 Director Group | 363

Performing an Initial Setup | 364

Restoring a Backup Configuration | 368

QFabric System Configuration | 371

Understanding QFabric System Administration Tasks and Utilities | 371

Gaining Access to the QFabric System Through the Default Partition | 375

Example: Configuring QFabric System Login Classes | 377

Example: Configuring RADIUS Authentication on a QFabric System | 387

Configuring Aliases for the QFabric System | 388

Configuring Node Groups for the QFabric System | 402

Configuring the Port Type on QFX3600 Node Devices | 408

Configuring the QSFP+ Port Type on QFX5100 Devices | 414

Example: Configuring SNMP | 417

Example: Configuring System Log Messages | 421

Configuring Graceful Restart for QFabric Systems | 424

Enabling Graceful Restart | 424

Configuring Graceful Restart Options for BGP | 425

Configuring Graceful Restart Options for OSPF and OSPFv3 | 426

Tracking Graceful Restart Events | 428

Optimizing the Number of Multicast Flows on QFabric Systems | 428

Segregating QFabric Traffic Flows With Flow Groups | 429

Example: Creating a QFabric Flow Group | 431

QFabric System Fabric OAM Configuration | 439

Overview of Internal Fabric Monitoring | 439

Configuring a Fabric Maintenance Association | 441

Configuring Flow Specifications | 442

 Configuring a Unicast Ethernet Flow Specification | 443

 Configuring a Unicast Ethernet IPv4 Flow Specification | 445

 Configuring a Multicast IPv4 Flow Specification | 446

 Configuring a Multicast VLAN Flood Flow Specification | 447

Example: Configuring Internal Fabric OAM Monitoring | 448

Configuration Statements | 457

aliases | 459

archive (QFabric System) | 460

chassis (QFabric System) | 461

device-authentication | 464

director-device (Aliases) | 466

ethernet-frame-size | 467

fabric | 468

fabric (OAM) | 470

fabric-control | 472

fabric-maintenance-associations | 473

fabric-maintenance-end-points | 475

file (QFabric System) | 477

flow-groups | 479

flow-specs | 486

graceful-restart (Enabling Globally) | 488

graceful-restart (Fabric Control) | 490

graceful-restart (Protocols BGP) | 491

graceful-restart (Protocols OSPF) | 493

interconnect-device (Chassis) | 495

interconnect-device (Aliases) | 497

interconnect-device (Flow Groups) | 499

multicast (QFabric Routing Options) | 500

multicast-ipv4 | 501

multicast-vlan-flood | 502

network-domain | 503

no-make-before-break | 504

node-device (Aliases) | 505

node-device (Chassis) | 507

node-device (Flow Groups) | 508

node-device (Resources) | 509

node-group (Chassis) | 510

node-group (Resources) | 512

pic (Port) | 514

protocols (Fabric) | 516

resources | 517

restart-time (Fabric Control) | 518

routing-options (QFabric System) | 519

syslog (QFabric System) | 520

stale-routes-time (Fabric Control) | 521

unicast-ethernet | 522

unicast-ethernet-ipv4 | 524

Administration

Software Upgrade, Downgrade, and Recovery | 529

Nonstop Software Upgrade Checklist for QFabric Systems | 529

- Preparing for an NSSU | 529
- Upgrading Director Groups | 531
- Upgrading QFabric | 532
- Upgrading Network Node Groups | 532

Performing a Nonstop Software Upgrade on the QFabric System | 533

- Backing Up the Current Configuration Files | 535
- Downloading Software Files Using a Browser | 536
- Retrieving Software Files for Download | 537
- Performing a Nonstop Software Upgrade for Director Devices in a Director Group | 537
- Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components | 537
- (Optional) Creating Upgrade Groups for Node Groups | 538

Performing a Nonstop Software Upgrade on a Node Group	539
Verifying Nonstop Software Upgrade for QFabric Systems	540
Verifying a Director Group Nonstop Software Upgrade	540
Verifying a Fabric Nonstop Software Upgrade	558
Verifying a Redundant Server Node Group Nonstop Software Upgrade	560
Verifying a Network Node Group Nonstop Software Upgrade	564
Upgrading Software on a QFabric System	566
Backing Up the Current Configuration Files	566
Downloading Software Files Using a Browser	567
Retrieving Software Files for Download	568
Installing the Software Package on the Entire QFabric System	568
Downgrading Software on a QFabric System	572
(Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive	573
Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software	575
Performing System Backup and Recovery for a QFabric System	581
QFabric System Administration and Maintenance 	583
Overview of Internal Fabric Monitoring	583
QFabric System Operational Mode Commands	585
Filtering Operational Mode Command Output in a QFabric System	587
Operational Mode Commands 	589
ping fabric multicast-flow	591
ping fabric unicast-flow	593
request chassis cb	596
request chassis device-mode	600
request chassis fabric fpc	603
request component login	605
request fabric administration director-group change-master	608
request fabric administration power-off interconnect-device	610
request fabric administration power-off interconnect-device fpc	612
request fabric administration power-off node-device	614
request fabric administration remove	616

request fabric administration system mac-pool add | **618**
request fabric administration system mac-pool delete | **619**
request system halt | **621**
request system reboot | **628**
request system software format-qfabric-backup | **638**
request system software nonstop-upgrade | **639**
request system software system-backup | **647**
set chassis display message | **649**
show chassis device-mode | **653**
show chassis environment cb | **656**
show chassis ethernet-switch interconnect-device cb | **692**
show chassis ethernet-switch interconnect-device fpc | **715**
show chassis fabric connectivity | **749**
show chassis fabric device | **758**
show chassis lcd | **760**
show chassis nonstop-upgrade node-group | **778**
show fabric administration inventory | **780**
show fabric administration inventory director-group status | **786**
show fabric administration inventory infrastructure | **793**
show fabric administration inventory interconnect-devices | **797**
show fabric administration inventory node-devices | **800**
show fabric administration inventory node-groups | **802**
show fabric administration system mac-pool | **805**
show fabric inventory | **807**
show fabric session-host | **811**
show log | **813**
show oam fabric flow-specification | **820**
show oam fabric interfaces | **824**
show system software upgrade status | **826**
traceroute fabric unicast-flow | **828**

Troubleshooting

QFabric System Troubleshooting | 833

Performing System Backup and Recovery for a QFabric System | 833

Performing a QFabric System Recovery Installation on the Director Group | 835

(Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive | 836

Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software | 838

Performing a Recovery Installation | 845

Creating an Emergency Boot Device for QFX Series Switches | 847

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | xvii
- Using the Examples in This Manual | xvii
- Documentation Conventions | xix
- Documentation Feedback | xxii
- Requesting Technical Support | xxii

Use this guide to plan, set up and configure the elements of a QFX3000-M QFabric System, a highly scalable distributed Layer 2 and Layer 3 networking architecture managed as a single entity that provides a high-performance, low-latency, and unified interconnect solution.

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <https://www.juniper.net/documentation/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <https://www.juniper.net/books>.

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xsl;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```


Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {  
    file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]  
user@host# edit system scripts  
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]  
user@host# load merge relative /var/tmp/ex-script-snippet.conf  
load complete
```

For more information about the **load** command, see [CLI Explorer](#).

Documentation Conventions

[Table 1 on page xx](#) defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xx defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions


Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are

covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <https://support.juniper.net/support/requesting-support/>.

1

PART

Overview

Before You Begin | 3

Hardware Architecture Overview | 21

Software Architecture Overview | 45

Software Features | 59

Before You Begin

IN THIS CHAPTER

- QFabric System Overview | 3
- Understanding QFabric System Terminology | 7
- Understanding Interfaces on the QFabric System | 15

QFabric System Overview

IN THIS SECTION


- Legacy Data Center Architecture | 3
- QFX Series QFabric System Architecture | 5

The architecture of legacy data centers contrasts significantly with the revolutionary Juniper Networks data center solution.

Legacy Data Center Architecture

Service providers and companies that support data centers are familiar with legacy multi-tiered architectures, as seen in [Figure 1 on page 4](#).

Figure 1: Legacy Data Center Architecture



The *access layer* connects servers and other devices to a Layer 2 switch and provides an entry point into the data center. Several access switches are in turn connected to intermediate Layer 2 switches at the *aggregation layer* (sometimes referred to as the *distribution layer*) to consolidate traffic. A *core layer* interconnects the aggregation layer switches. Finally, the core switches are connected to Layer 3 routers in the *routing layer* to send the aggregated data center traffic to other data centers or a wide area network (WAN), receive external traffic destined for the data center, and interconnect different Layer 2 broadcast domains within the data center.


The problems that exist with the multi-tiered data center architecture include:

- **Limited scalability**—The demands for electrical power, cooling, cabling, rack space, and port density increase exponentially as the traditional data center expands, which prohibits growth after minimal thresholds are met.
- **Inefficient resource usage**—Up to 50 percent of switch ports in a legacy data center are used to interconnect different tiers rather than support server and storage connections. In addition, traffic that ideally should move horizontally between servers within a data center often must also be sent vertically up through the tiers to reach a router and down through the tiers to reach the required destination server.
- **Increased latency**—By requiring the devices at each tier level to perform multiple iterations of packet and frame processing, the data plane traffic takes significantly longer to reach its destination than if the sending and receiving devices were directly connected. This processing overhead results in potentially poor performance for time-sensitive applications, such as voice, video, or financial transactions.

QFX Series QFabric System Architecture

In contrast to legacy multi-tiered data center architectures, the Juniper Networks QFX Series QFabric System architecture provides a simplified networking environment that solves the most challenging issues faced by data center operators. A fabric is a set of devices that act in concert to behave as a single switch. It is a highly scalable, distributed, Layer 2 and Layer 3 networking architecture that provides a high-performance, low-latency, and unified interconnect solution for next-generation data centers as seen in [Figure 2 on page 5](#).

Figure 2: QFX Series QFabric System Architecture



A QFabric system collapses the traditional multi-tiered data center model into a single tier where all access layer devices (known in the QFabric system model as *Node devices*) are essentially directly connected to all other access layer devices across a very large scale fabric backplane (known in the QFabric system model as the *Interconnect device*). Such an architecture enables the consolidation of data center endpoints

(such as servers, storage devices, memory, appliances, and routers) and provides better scaling and network virtualization capabilities than traditional data centers.

Essentially, a QFabric system can be viewed as a single, nonblocking, low-latency switch that supports thousands of 10-Gigabit Ethernet ports or 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel ports to interconnect servers, storage, and the Internet across a high-speed, high-performance fabric. The entire QFabric system is managed as a single entity through a *Director group*, containing redundant hardware and software components that can be expanded and scaled as the QFabric system grows in size. In addition, the Director group automatically senses when devices are added or removed from the QFabric system and dynamically adjusts the amount of processing resources required to support the system. Such intelligence helps the QFabric system use the minimum amount of power to run the system efficiently, but not waste energy on unused components.

As a result of the QFabric system architecture, data center operators are now realizing the benefits of this next-generation architecture, including:

- **Low latency**—Because of its inherent advantages in this area, the QFabric system provides an excellent foundation for mission-critical applications such as financial transactions and stock trades, as well as time-sensitive applications such as voice and video.
- **Enhanced scalability**—The QFabric system can be managed as a single entity and provides support for thousands of data center devices. As Internet traffic continues to grow exponentially with the increase in high-quality video transmissions and rise in the number of mobile devices used worldwide, the QFabric system can keep pace with the demands for bandwidth, applications, and services offered by the data center.
- **Virtualization-enabled**—The QFabric system was designed to work seamlessly with virtual servers, virtual appliances, and other virtual devices, allowing for even greater scalability, expandability, and rapid deployment of new services than ever before. Migrating to virtual devices also results in significant costs savings, fueled by reduced space requirements, decreased needs for power and cooling, and increased processing capabilities.
- **Simplicity**—Although the QFabric system can scale to hundreds of devices and thousands of ports, you can still manage the QFabric system as a single system.
- **Flexibility**—You can deploy the QFabric system as an entire system or in stages.
- **Convergence**—Because the congestion-free fabric is lossless, all traffic in a QFabric system can be converged onto a single network. As a result, the QFabric system supports Ethernet, Fibre Channel over Ethernet, and native Fibre Channel packets and frames.

Flat, nonblocking, and lossless, the network fabric offered by the QFabric system has the scale and flexibility to meet the needs of small, medium, and large-sized data centers for years to come.

RELATED DOCUMENTATION

Understanding QFabric System Terminology

To understand the QFabric system environment and its components, you should become familiar with the terms defined in [Table 3 on page 7](#).

Table 3: QFabric System Terms

Term	Definition
Clos network fabric	Three-stage switching network in which switch elements in the middle stages are connected to all switch elements in the ingress and egress stages. In the case of QFabric system components, the three stages are represented by an ingress chipset, a midplane chipset, and an egress chipset in an Interconnect device (such as a QFX3008-I Interconnect device). In Clos networks, which are well known for their nonblocking properties, a connection can be made from any idle input port to any idle output port, regardless of the traffic load in the rest of the system.
Director device	Hardware component that processes fundamental QFabric system applications and services, such as startup, maintenance, and inter-QFabric system device communication. A set of Director devices with hard drives can be joined to form a <i>Director group</i> , which provides redundancy and high availability by way of additional memory and processing power. (See also <i>Director group</i> .)
Director group	<p>Set of Director devices that host and load-balance internal processes for the QFabric system. The Director group handles tasks such as QFabric system network topology discovery, Node and Interconnect device configuration, startup, and DNS, DHCP, and NFS services. Operating a Director group is a minimum requirement to manage a QFabric system.</p> <p>The Director group runs the Director software for management applications and runs dual processes in active/standby mode for maximum redundancy and high availability. (See also <i>Director software</i> and <i>Director device</i>.)</p>

Table 3: QFabric System Terms (*continued*)

Term	Definition
Director software	Software that handles QFabric system administration tasks, such as fabric management and configuration. The Junos OS-based Director software runs on the <i>Director group</i> , provides a single, consolidated view of the QFabric system, and enables the main QFabric system administrator to configure, manage, monitor, and troubleshoot QFabric system components from a centralized location. To access the Director software, log in to the default partition. (See also <i>Director device</i> and <i>Director group</i> .)
fabric control Routing Engine	Virtual Junos OS Routing Engine instance used to control the exchange of routes and flow of data between QFabric system hardware components within a partition. The fabric control Routing Engine runs on the Director group.
fabric manager Routing Engine	Virtual Junos OS Routing Engine instance used to control the initialization and maintenance of QFabric system hardware components belonging to the default partition. The fabric manager Routing Engine runs on the Director group.
infrastructure	QFabric system services processed by the virtual Junos Routing Engines operating within the Director group. These services, such as fabric management and fabric control, support QFabric system functionality and high availability.
Interconnect device	QFabric system component that acts as the primary fabric for data plane traffic traversing the QFabric system between Node devices. Examples of Interconnect devices include the QFX3008-I Interconnect device in a QFX3000-G QFabric system, the QFX5100-24Q configured as an Interconnect device, and the QFX3600-I Interconnect device in a QFX3000-M QFabric system. (See also <i>Node device</i> .)
Junos Space	Carrier-class network management system for provisioning, monitoring, and diagnosing Juniper Networks routing, switching, security, and data center platforms.
network Node group	Set of one to eight Node devices that connects to an external network.
network Node group Routing Engine	Virtual Junos OS Routing Engine instance that handles routing processes for a network Node group. The network Node group Routing Engine runs on the Director group.

Table 3: QFabric System Terms (*continued*)

Term	Definition
Node device	<p>Routing and switching device that connects to endpoints (such as servers or storage devices) or external network peers, and is connected to the QFabric system through an Interconnect device. You can deploy Node devices similarly to the way a top-of-rack switch is implemented. Examples of Node devices include the QFX3500 Node device, QFX3600 Node device, and QFX5100 Node device. (See also <i>Interconnect device</i> and <i>network Node group</i>.)</p>
partition	<p>Collection of physical or logical QFabric system hardware components (such as Node devices) that provides fault isolation, separation, and security.</p> <p>In their initial state, all QFabric system components belong to a <i>default partition</i>.</p>
QFabric system	<p>Highly scalable, distributed, Layer 2 and Layer 3 networking architecture that provides a high-performance, low-latency, and unified interconnect solution for next-generation data centers. A QFabric system collapses the traditional multi-tier data center model, enables the consolidation of data center endpoints (such as servers, storage devices, memory, appliances, and routers), and provides better scaling and network virtualization capabilities than traditional data centers.</p> <p>Essentially, a QFabric system can be viewed as a single, nonblocking, low-latency switch that supports thousands of 10-Gigabit Ethernet ports or 2-Gbps, 4-Gbps or 8-Gbps Fibre Channel ports to interconnect servers, storage, and the Internet across a high-speed, high-performance fabric. The QFabric system must have sufficient resources and devices allocated to handle the <i>Director group</i>, <i>Node device</i>, and <i>Interconnect device</i> functions and capabilities.</p>

Table 3: QFabric System Terms (*continued*)

Term	Definition
QFabric system control plane	<p>Internal network connection that carries control traffic between QFabric system components. The QFabric system control plane includes management connections between the following QFabric system hardware and software components:</p> <ul style="list-style-type: none"> • <i>Node devices</i>, such as the QFX3500 Node device. • <i>Interconnect devices</i>, such as the QFX3008-I Interconnect device. • <i>Director group</i> processes, such as management applications, provisioning, and topology discovery. • <i>Control plane Ethernet switches</i> to provide interconnections to all QFabric system devices and processes. For example, you can use EX Series EX4200 switches running in Virtual Chassis mode for this purpose. <p>To maintain high availability, the QFabric system control plane uses a different network than the QFabric system data plane, and uses a fabric provisioning protocol and a fabric management protocol to establish and maintain the QFabric system.</p>
QFabric system data plane	<p>Redundant, high-performance, and scalable data plane that carries QFabric system data traffic. The QFabric system data plane includes the following high-speed data connections:</p> <ul style="list-style-type: none"> • 10-Gigabit Ethernet connections between QFabric system endpoints (such as servers or storage devices) and Node devices. • 40-Gbps quad small form-factor pluggable plus (QSFP+) connections between Node devices and Interconnect devices. • 10-Gigabit Ethernet connections between external networks and a Node device acting as a network Node group. <p>To maintain high availability, the QFabric system data plane is separate from the QFabric system control plane.</p>
QFabric system endpoint	<p>Device connected to a Node device port, such as a server, a storage device, memory, an appliance, a switch, or a router.</p>
QFabric system fabric	<p>Distributed, multistage network that consists of a queuing and scheduling system that is implemented in the Node device, and a distributed cross-connect system that is implemented in Interconnect devices. The QFabric system fabric is part of the QFabric system data plane.</p>

Table 3: QFabric System Terms (*continued*)

Term	Definition
QFX3500 Node device	<p>Node device that connects to either endpoint systems (such as servers and storage devices) or external networks in a QFabric system. It is packaged in an industry-standard 1U, 19-inch rack-mounted enclosure.</p> <p>The QFX3500 Node device provides up to 48 10-Gigabit Ethernet interfaces to connect to the endpoints. Twelve of these 48 interfaces can be configured to support 2-Gbps, 4-Gbps or 8-Gbps Fibre Channel, and 36 of the interfaces can be configured to support Gigabit Ethernet. Also, there are four uplink connections to connect to Interconnect devices in a QFabric system. These uplinks use 40-Gbps quad small form-factor pluggable plus (QSFP+) interfaces. (See also <i>QFX3500 switch</i>.)</p>
QFX3500 switch	<p>Standalone data center switch with 10-Gigabit Ethernet access ports and 40-Gbps quad, small form-factor pluggable plus (QSFP+) uplink interfaces. You can (optionally) configure some of the access ports as 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel ports or Gigabit Ethernet ports.</p> <p>The QFX3500 switch can be converted to a QFabric system Node device as part of a complete QFabric system. The switch is packaged in an industry-standard 1U, 19-inch rack-mounted enclosure. (See also <i>QFX3500 Node device</i>.)</p>
QFX3600 Node device	<p>Node device that connects to either endpoint systems (such as servers and storage devices) or external networks in a QFabric system. It is packaged in an industry-standard 1U, 19-inch rack-mounted enclosure.</p> <p>The QFX3600 Node device provides 16 40-Gbps QSFP+ ports. By default, 4 ports (labeled Q0 through Q3) are configured for 40-Gbps uplink connections between your Node device and your Interconnect device, and 12 ports (labeled Q4 through Q15) use QSFP+ direct-attach copper (DAC) breakout cables or QSFP+ transceivers with fiber breakout cables to support 48 10-Gigabit Ethernet interfaces for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the first eight ports (Q0 through Q7) for uplink connections between your Node device and your Interconnect device, and ports Q2 through Q15 for 10-Gigabit Ethernet connections to either endpoint systems or external networks. (See also <i>QFX3600 switch</i>.)</p>

Table 3: QFabric System Terms (continued)

Term	Definition
QFX3600 switch	<p>Standalone data center switch with 16 40-Gbps quad, small form-factor pluggable plus (QSFP+) interfaces. By default, all the 16 ports operate as 40-Gigabit Ethernet ports. Optionally, you can choose to configure the 40-Gbps ports to operate as four 10-Gigabit Ethernet ports. You can use QSFP+ to four SFP+ breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches.</p> <p>The QFX3600 switch can be converted to a QFabric system Node device as part of a complete QFabric system. The switch is packaged in an industry-standard 1U, 19-inch rack-mounted enclosure. (See also <i>QFX3600 Node device</i>.)</p>

Table 3: QFabric System Terms (*continued*)

Term	Definition
QFX5100 Node device	<p>QFabric system Node device that connects to either endpoint systems (such as servers and storage devices) or external networks. All three supported models are packaged in an industry-standard 1U, 19-inch rack-mounted enclosure. A QFX5100 Node device can be any of these models:</p> <ul style="list-style-type: none"> QFX5100-48S By default, the QFX5100-48S Node device provides 48 10-Gigabit Ethernet interfaces to connect to the endpoints. There are also six 40-Gbps quad small form-factor pluggable plus (QSFP+) interfaces, of which four are uplinks (FTE). QFX5100-48T By default, the QFX5100-48T Node device provides 48 10GBASE-T interfaces to connect to endpoints. There are also six 40-Gbps QSFP+ interfaces, of which four are uplinks (FTE) QFX5100-24Q By default, the QFX5100-24Q Node device provides 24 40-Gigabit Ethernet QSFP+ interfaces to connect to the endpoints. The QFX5100-24Q has two expansion bays. The number of additional interfaces available depends on the expansion module and the System mode configured for the Node device. <p>By default, on the QFX5100-48S Node device and QFX5100-48T Node device, the first 4 ports (labeled fte-0/1/0 through fte-0/1/3) are configured for 40-Gbps uplink connections between your Node device and your Interconnect devices, and 2 ports (labeled xle-0/1/4 and xle-0/1/5) use QSFP+ direct-attach copper (DAC) breakout cables or QSFP+ transceivers with fiber breakout cables to support 8 10-Gigabit Ethernet interfaces for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the middle 2 ports (xle-0/1/2 and xle-0/1/3) for additional connections to either endpoint systems or external networks.</p> <p>(See also <i>QFX3500 Node device</i> and <i>QFX3600 Node device</i>.)</p>
redundant server Node group	<p>Set of two Node devices that connect to servers or storage devices. Link aggregation group (LAG) interfaces can span the Node devices within a redundant server Node group.</p>

Table 3: QFabric System Terms (*continued*)

Term	Definition
rolling upgrade	Method used in the QFabric system to upgrade the software for components in a systematic, low-impact way. A rolling upgrade begins with the Director group, proceeds to the fabric (Interconnect devices), and finishes with the Node groups.
Routing Engine	<p>Juniper Networks-proprietary processing entity that implements QFabric system control plane functions, routing protocols, system management, and user access. Routing Engines can be either physical or virtual entities.</p> <p>The Routing Engine functions in a QFabric system are sometimes handled by Node devices (when connected to endpoints), but mostly implemented by the Director group (to provide support for QFabric system establishment, maintenance, and other tasks).</p>
routing instance	<p>Private collection of routing tables, interfaces, and routing protocol parameters unique to a specific customer. The set of interfaces is contained in the routing tables, and the routing protocol parameters control the information in the routing tables.</p> <p>(See also <i>virtual private network</i>.)</p>
server Node group	Set of one or more Node devices that connect to servers or storage devices.
virtual LAN (VLAN)	Unique Layer 2 broadcast domain for a set of ports selected from the components available in a partition. VLANs allow manual segmentation of larger Layer 2 networks and help to restrict access to network resources. To interconnect VLANs, Layer 3 routing is required.
virtual private network (VPN)	Layer 3 routing domain within a partition. VPNs maintain privacy with a tunneling protocol, encryption, and security procedures. In a QFabric system, a Layer 3 VPN is configured as a <i>routing instance</i> .
flow group	Force redundant multicast streams to flow through different interconnect devices to prevent a single interconnect device from potentially dropping both streams of multicast traffic during a failure.

RELATED DOCUMENTATION

[QFabric System Overview](#) | 3

[Understanding the QFabric System Hardware Architecture | 21](#)

[Understanding the QFabric System Software Architecture | 45](#)

[Understanding Fibre Channel Terminology](#)

[Understanding QFabric Multicast Flow Groups | 55](#)

Understanding Interfaces on the QFabric System

IN THIS SECTION

- [Four-Level Interface Naming Convention | 15](#)
- [QSFP+ Interfaces | 15](#)
- [Link Aggregation | 19](#)

This topic describes:

Four-Level Interface Naming Convention

When you configure an interface on the QFabric system, the interface name needs to follow a four-level naming convention that enables you to identify an interface as part of either a Node device or a Node group. Include the name of the network or server Node group at the beginning of the interface name.

The four-level interface naming convention is:

device-name:type-fpc/pic/port

where *device-name* is the name of the Node device or Node group. The remainder of the naming convention elements are the same as those in the QFX3500 switch interface naming convention.

An example of a four-level interface name is:

node2:xe-0/0/2

QSFP+ Interfaces

The QFX3500 Node device provides four 40-Gbps QSFP+ (quad small form-factor pluggable plus) interfaces (labeled **Q0** through **Q3**) for uplink connections between your Node device and your Interconnect devices.

The QFX3600 Node device provides 16 40-Gbps QSFP+ interfaces. By default, 4 interfaces (labeled **Q0** through **Q3**) are configured for 40-Gbps uplink connections between your Node device and your Interconnect devices, and 12 interfaces (labeled **Q4** through **Q15**) use QSFP+ direct-attach copper (DAC)

breakout cables or QSFP+ transceivers with fiber breakout cables to support 48 10-Gigabit Ethernet interfaces for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the first eight interfaces (Q0 through Q7) for uplink connections between your Node device and your Interconnect devices, and interfaces Q2 through Q15 for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks (see [“Configuring the Port Type on QFX3600 Node Devices” on page 408](#)). [Table 4 on page 16](#) shows the port mappings for QFX3600 Node devices.

Table 4: QFX3600 Node Device Port Mappings

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
Q0	Not supported on this port	xle-0/1/0	fte-0/1/0
Q1	Not supported on this port	xle-0/1/1	fte-0/1/1
Q2	xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11	xle-0/1/2	fte-0/1/2
Q3	xe-0/0/12 xe-0/0/13 xe-0/0/14 xe-0/0/15	xle-0/1/3	fte-0/1/3
Q4	xe-0/0/16 xe-0/0/17 xe-0/0/18 xe-0/0/19	xle-0/1/4	fte-0/1/4
Q5	xe-0/0/20 xe-0/0/21 xe-0/0/22 xe-0/0/23	xle-0/1/5	fte-0/1/5

Table 4: QFX3600 Node Device Port Mappings (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
Q6	xe-0/0/24	xle-0/1/6	fte-0/1/6
	xe-0/0/25		
	xe-0/0/26		
	xe-0/0/27		
Q7	xe-0/0/28	xle-0/1/7	fte-0/1/7
	xe-0/0/29		
	xe-0/0/30		
	xe-0/0/31		
Q8	xe-0/0/32	xle-0/1/8	Not supported on this port
	xe-0/0/33		
	xe-0/0/34		
	xe-0/0/35		
Q9	xe-0/0/36	xle-0/1/9	Not supported on this port
	xe-0/0/37		
	xe-0/0/38		
	xe-0/0/39		
Q10	xe-0/0/40	xle-0/1/10	Not supported on this port
	xe-0/0/41		
	xe-0/0/42		
	xe-0/0/43		

Table 4: QFX3600 Node Device Port Mappings (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
Q11	xe-0/0/44 xe-0/0/45 xe-0/0/46 xe-0/0/47	xle-0/1/11	Not supported on this port
Q12	xe-0/0/48 xe-0/0/49 xe-0/0/50 xe-0/0/51	xle-0/1/12	Not supported on this port
Q13	xe-0/0/52 xe-0/0/53 xe-0/0/54 xe-0/0/55	xle-0/1/13	Not supported on this port
Q14	xe-0/0/56 xe-0/0/57 xe-0/0/58 xe-0/0/59	xle-0/1/14	Not supported on this port
Q15	xe-0/0/60 xe-0/0/61 xe-0/0/62 xe-0/0/63	xle-0/1/15	Not supported on this port

The QFX5100-48S Node device provides 48 10-Gigabit Ethernet interfaces and 6 40-Gbps QSFP+ interfaces. By default, 4 interfaces (labeled **48** through **51**) are configured for 40-Gbps uplink connections between your Node device and your Interconnect devices, and 2 interfaces (labeled **52** and **53**) support 40-Gigabit Ethernet connections to either endpoint systems (such as servers and storage devices) or

external networks. Optionally, you can choose to configure the middle two interfaces (**50** and **51**) for 40-Gigabit Ethernet connections to either endpoint systems or external networks, and you can choose to configure the last two interfaces (**52** and **53**) for uplink connections between your Node device and your Interconnect devices (see [“Configuring the QSFP+ Port Type on QFX5100 Devices” on page 414](#)).

[Table 5 on page 19](#) shows the port mappings for QFX5100-48S Node devices.

Table 5: QFX5100-48S Node Device Port Mappings

Port Number	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
48	Not supported on this PIC	fte-0/1/0
49	Not supported on this PIC	fte-0/1/1
50	xle-0/1/2	fte-0/1/2
51	xle-0/1/3	fte-0/1/3
52	xle-0/1/4	fte-0/1/4
53	xle-0/1/5	fte-0/1/5

Link Aggregation

Link aggregation enables you to create link aggregation groups across Node devices within a network Node group or redundant server Node group. You can include up to eight Ethernet interfaces in a LAG. You can have up to 48 LAGs within a redundant server Node group, and 128 LAGs in a network Node group. To configure a LAG, include the **aggregated-devices** statement at the **[edit chassis node-group node-group-name]** hierarchy level and the **device-count** statement at the **[edit chassis node-group node-group-name aggregated-devices ethernet]** hierarchy level. Additionally, include any aggregated Ethernet options (**minimum-links** and **link-speed**) at the **[edit interfaces interface-name aggregated-ether-options]** hierarchy level and the **802.3ad** statement at the **[edit interfaces interface-name ether-options]** hierarchy level. To configure the Link Aggregation Control Protocol (LACP), include the **lacp** statement at the **[edit interfaces aggregated-ether-options]** hierarchy level.

RELATED DOCUMENTATION

[Configuring the Port Type on QFX3600 Node Devices | 408](#)

[Configuring the QSFP+ Port Type on QFX5100 Devices | 414](#)

Hardware Architecture Overview

IN THIS CHAPTER

- [Understanding the QFabric System Hardware Architecture | 21](#)
- [Understanding the Director Group | 26](#)
- [Understanding Routing Engines in the QFabric System | 27](#)
- [Understanding Interconnect Devices | 29](#)
- [Understanding Node Devices | 33](#)
- [Understanding Node Groups | 38](#)
- [Understanding Port Oversubscription on Node Devices | 40](#)

Understanding the QFabric System Hardware Architecture

IN THIS SECTION

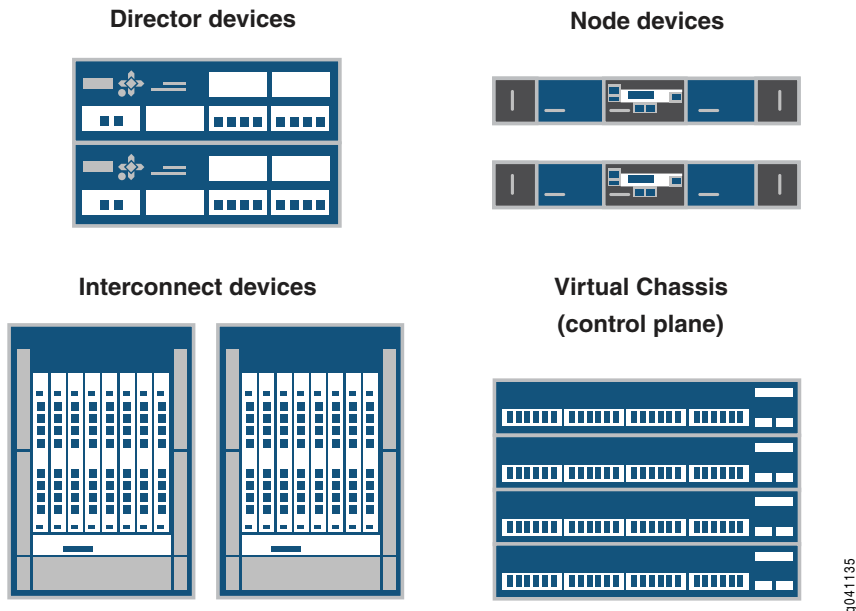
- [QFabric System Hardware Architecture Overview | 21](#)
- [QFX3000-G QFabric System Features | 25](#)
- [QFX3000-M QFabric System Features | 25](#)

QFabric System Hardware Architecture Overview

The QFabric system is a single-layer networking tier that connects servers and storage devices to one another across a high-speed, unified core fabric. You can view the QFabric system as a single, extremely large, nonblocking, high-performance Layer 2 and Layer 3 switching system. The reason you can consider the QFabric system as a single system is that the Director software running on the Director group allows the main QFabric system administrator to access and configure every device and port in the QFabric system from a single location. Although you configure the system as a single entity, the fabric contains four major hardware components. The hardware components can be chassis-based, group-based, or a hybrid of the two. As a result, it is important to understand the four types of generic QFabric system

components and their functions, regardless of which hardware environment you decide to implement. A representation of these components is shown in [Figure 3 on page 22](#).

Figure 3: QFabric System Hardware Architecture



The four major QFabric system components include the following:

- Director group**—The *Director group* is a management platform that establishes, monitors, and maintains all components in the QFabric system. It is a set of Director devices that run the Junos operating system (Junos OS) on top of a CentOS foundation. The Director group handles tasks such as QFabric system network topology discovery, Node and Interconnect device configuration and startup, and Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), and Network File System (NFS) services. The Director group also runs the software for management applications, hosts and load-balances internal processes for the QFabric system, and starts additional QFabric system processes as requested.
- Node devices**—A *Node device* is a hardware system located on the ingress of the QFabric system that connects to endpoints (such as servers or storage devices) or external networks, and is connected to the heart of the QFabric system through an Interconnect device. A Node device can be used in a manner similar to how a top-of-rack switch is implemented. By default, Node devices connect to servers or storage devices. However, when you group Node devices together to connect to a network that is external to the QFabric system, the formation is known as a *network Node group*.
- Interconnect devices**—An *Interconnect device* acts as the primary fabric for data plane traffic traversing the QFabric system between Node devices. To reduce latency to a minimum, the Interconnect device

implements multistage Clos switching to provide nonblocking interconnections between any of the Node devices in the system.

- **Control plane network**—The *control plane network* is an out-of-band Gigabit Ethernet management network that connects all QFabric system components. For example, you can use a group of EX4200 Ethernet switches or a group of EX4300 Ethernet switches configured as a Virtual Chassis to enable the control plane network. The control plane network connects the Director group to the management ports of the Node and Interconnect devices. By keeping the control plane network separate from the data plane, the QFabric system can scale to support thousands of servers and storage devices.

The four major QFabric system components can be assembled from a variety of hardware options. Currently supported hardware configurations are shown in [Table 6 on page 23](#).

Table 6: Supported QFabric System Hardware Configurations

QFabric System Configuration	Director Group	Node Device	Interconnect Device	Control Plane Device
QFX3000-G QFabric system	QFX3100 Director group	QFX3500, QFX3600, and QFX5100-48S, QFX5100-48T, and QFX5100-24Q Node devices NOTE: There can be a maximum of 128 Node devices in the QFX3000-G QFabric system.	QFX3008-I Interconnect device NOTE: There can be a maximum of four Interconnect devices in the QFX3000-G QFabric system.	Two Virtual Chassis composed of either four EX4200-48T or four EX4300-48T switches each (for a copper-based control plane) or eight EX4200-24F or four EX4300-48P switches each (for a fiber-based control plane)

Table 6: Supported QFabric System Hardware Configurations (*continued*)

QFabric System Configuration	Director Group	Node Device	Interconnect Device	Control Plane Device
QFX3000-M QFabric system	QFX3100 Director group NOTE: For a copper-based QFX3000-M QFabric system control plane network, use QFX3100 Director devices with RJ-45 network modules installed. For a fiber-based control plane network, use QFX3100 Director devices with SFP network modules installed.	QFX3500, QFX3600, and QFX5100-48S, QFX5100-48T, and QFX5100-24Q Node devices NOTE: <ul style="list-style-type: none"> There can be a maximum of 16 Node devices in the QFX3000-M QFabric system using QFX3600-I as Interconnect devices and 32 Node devices using the QFX5100-24Q as Interconnect devices. NOTE: QFX5100-24Q Interconnect devices and QFX3600-I Interconnect devices cannot be mixed on the same QFabric system. <ul style="list-style-type: none"> For a copper-based QFX3000-M QFabric system control plane network, use QFX3500 Node devices with a 1000BASE-T management board installed. For a fiber-based control plane network, use QFX3500 Node devices with an SFP management board installed. 	QFX5100-24Q or QFX3600-I Interconnect devices NOTE: There can be a maximum of four Interconnect devices in the QFX3000-M QFabric system.	Two EX4200 Ethernet or two EX4300 switches NOTE: For a copper-based QFX3000-M QFabric system control plane network, use EX4200-24T or EX4300-48T switches with an SFP+ uplink module installed. For a fiber-based control plane network, use EX4200-24F or EX4300-48P switches with an SFP+ uplink module installed.

To complete the system, external Routing Engines (such as the fabric manager Routing Engines, network Node group Routing Engines, and fabric control Routing Engines) run on the Director group and implement QFabric system control plane functions. The control plane network provides the control plane connections between the Node devices, the Interconnect devices, and the Routing Engines running on the Director group.

QFX3000-G QFabric System Features

A QFX3000-G QFabric system provides the following key features:

- Support for up to 128 Node devices and 4 Interconnect devices, which provides a maximum of 6144 10-Gigabit Ethernet ports.
- Low port-to-port latencies that scale as the system size grows from 48 to 6144 10-Gigabit Ethernet ports.
- Support for up to 384,000 total ingress queues at each Node device to the QFabric system Interconnect backplane.
- Support for Converged Enhanced Ethernet (CEE) traffic.

QFX3000-M QFabric System Features

A QFX3000-M QFabric system provides the following key features:

- Support for up to 32 Node devices and 4 QFX5100-24Q Interconnect devices or 16 Node device and 4 QFX3600-I Interconnect devices.

NOTE: You may not mix QFX5100-24Q Interconnect devices with QFX3600-I Interconnect devices on the same QFX3000-M QFabric system.

- Low port-to-port latencies that scale as the system size grows from 48 to 768 10-Gigabit Ethernet ports.

RELATED DOCUMENTATION

[Understanding QFabric System Terminology | 7](#)

[Understanding the QFabric System Software Architecture | 45](#)

[Understanding the Director Group | 26](#)

[Understanding Routing Engines in the QFabric System | 27](#)

[Understanding Interconnect Devices | 29](#)

[Understanding Node Devices | 33](#)

[Understanding Node Groups | 38](#)

[Understanding Partitions | 47](#)

Understanding the Director Group

IN THIS SECTION

- [Director Group Components | 26](#)
- [Director Group Services | 26](#)

Because the Director group provides management services for the QFabric system, it is important to understand the components of the cluster and how the Director group supports the needs of the greater fabric.

Director Group Components

When you build a Director group, consider the following elements and concepts.

- **Director device**—A single management device for the QFabric system. Director devices with a hard drive provide full processing services and are used to build the Director group.
- **Director group**—A set of Director devices. The Director group is essential to the QFabric system, which cannot operate properly without it. The Director group shares and load-balances processing tasks for the QFabric system, performs topology discovery, assigns identifiers to QFabric system components, and manages interfabric communication. The primary devices in a Director group are Director devices that contain hard drives. The Director devices run dual processes in active or standby mode for maximum redundancy.

When you add additional Director devices to the group, the Director group coordinates their activities and distributes processing loads across all available Director devices. The additional Director devices provide the Director group with additional memory and processing power. Supplementing the Director group with extra Director devices allows the group to scale efficiently and serve the needs of the entire QFabric system as it grows.

Director Group Services

The Director group is a management platform that establishes, monitors, and maintains all components in the QFabric system. It is a set of Director devices that run the Junos operating system (Junos OS) on top of a CentOS foundation. The Director group handles tasks such as QFabric system network topology discovery, Node and Interconnect device configuration and startup, and Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), and Network File System (NFS) services. The Director group also runs the software for management applications, hosts and load-balances internal processes for the

QFabric system, maintains configuration and topology databases, and starts additional QFabric system processes as requested.

Another critical role provided by the Director group is the hosting of the virtual Junos Routing Engines. These Routing Engines provide services for the QFabric system to keep it operating smoothly.

RELATED DOCUMENTATION

[Performing the QFabric System Initial Setup on a QFX3100 Director Group | 363](#)

[Understanding Routing Engines in the QFabric System | 27](#)

[Understanding the QFabric System Hardware Architecture | 21](#)

Understanding Routing Engines in the QFabric System

IN THIS SECTION

- [Hardware-Based Routing Engines | 27](#)
- [Software-Based External Routing Engines | 27](#)

Routing Engines perform many important processing tasks in the QFabric system. Knowing where the Routing Engines are located and what services they provide enables you to troubleshoot the QFabric system and ensure that it is running the way it should.

Hardware-Based Routing Engines


A traditional Juniper Networks Routing Engine is a hardware field-replaceable unit that runs routing protocols, builds the routing and switching tables, sends routing information to the Packet Forwarding Engine, and handles several software processes for the device (such as interface control, chassis component monitoring, system management, and user access). Node devices that are part of server Node groups in the QFabric system that connect to servers or storage devices implement Routing Engine functions locally using this traditional hardware method.

Software-Based External Routing Engines

The QFabric system also uses external Routing Engines that run in software on the Director group. In contrast with traditional Routing Engines, the functions and processes provided by software-based Routing

Engines are segmented, specialized, and distributed across multiple Routing Engine instances running on the Director group. Such separation provides redundancy for these functions and enables the QFabric system to scale. [Figure 4 on page 28](#) shows the external Routing Engine types.

Figure 4: External Routing Engine Types



These special-purpose external Routing Engine instances running on the Director group provide the following major services for the QFabric system:

- **Fabric manager Routing Engine**—Provides services to all devices in the QFabric system, such as system initialization, topology discovery, internal IP address and ID assignment, and interdevice communication. The fabric manager Routing Engine authenticates Interconnect and Node devices, and maintains a database for system components. A single fabric manager Routing Engine instance is generated to manage the entire QFabric system.
- **Fabric control Routing Engine**—Runs the fabric control protocol to share route information between available devices in a partition. A pair of redundant route distribution Routing Engine instances is generated for every partition in the QFabric system, and both instances are active.
- **Diagnostic Routing Engine**—Gathers operational information that allows QFabric system administrators to monitor the health of the QFabric system. A single Routing Engine instance is generated for the entire QFabric system.
- **Network Node group Routing Engine**—Provides Routing Engine functionality for groups of Node devices bundled together as a single Layer 3 routing device, which is used to connect to external networks. A pair of redundant Routing Engine instances is generated for every network Node group in the QFabric system.

RELATED DOCUMENTATION

[Understanding the Director Group | 26](#)

[Understanding the QFabric System Control Plane | 49](#)

[Understanding the QFabric System Hardware Architecture | 21](#)

Understanding Interconnect Devices

IN THIS SECTION


- [Interconnect Device Introduction | 29](#)
- [QFX3008-I Interconnect Devices | 30](#)
- [QFX3600-I Interconnect Devices | 31](#)
- [QFX5100-24Q Interconnect Devices | 32](#)

Interconnect devices in a QFabric system provide a way for the Node devices to connect with one another over a high-speed backplane. By understanding the role of Interconnect devices, you can harness the benefits of low latency, superb scalability, and minimum packet processing offered by a single-tier data center architecture.

Interconnect Device Introduction

Interconnect devices act as the primary fabric for data plane traffic traversing the QFabric system between Node devices. The main task for the Interconnect devices is to transfer traffic between the Node devices as quickly as possible across a high-speed, available path backplane. To reduce latency to a minimum, larger Interconnect devices (such as the QFX3008-I Interconnect device) implement multistage Clos switching to provide nonblocking connections between any of the Node devices in the system. [Figure 5 on page 30](#) shows an example of how Clos switching works in the QFX3008-I Interconnect device.

Figure 5: Clos Switching for QFX3008-I Interconnect Devices

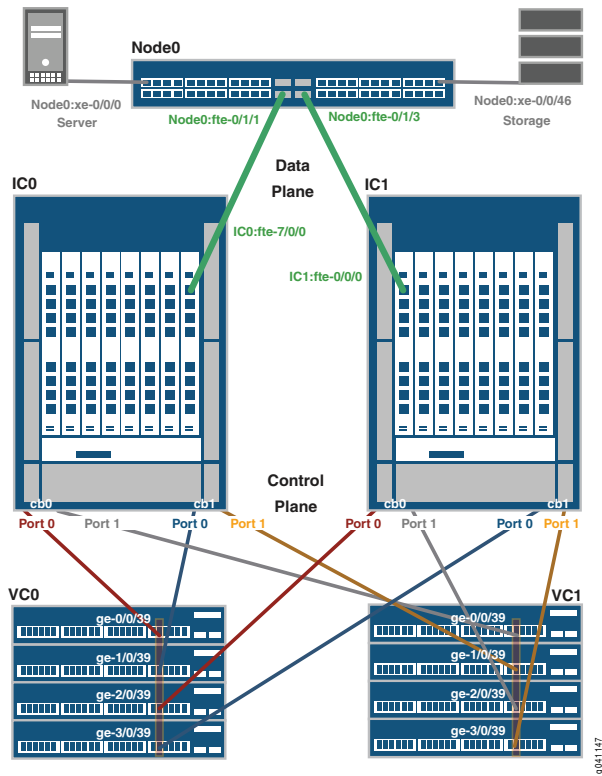


Traffic enters a QSFP+ port from a Node device, and an ingress chipset provides stage F1 processing. For the F2 stage, the frame is sent to a rear card and processed by a midplane chipset. Lastly, an egress chipset on the front card QSFP+ port handles processing tasks for the F3 stage. At each of the three Clos stages, a switching table chooses the best path and determines where to send the frame to reach the next stage. The F1 and F3 stages can be handled by the same front card or different front cards, depending on the best path selected by the fabric. After the frame traverses the Interconnect device backplane, the Interconnect device sends the frame to the egress Node device.

QFX3008-I Interconnect Devices

The QFX3008-I Interconnect device contains eight slots in the front of the chassis. In each slot, you can install a front card containing 16 40-Gbps quad small form-factor pluggable plus (QSFP+) ports. A fully configured system offers a total capacity of 128 QSFP+ connections. These front card ports attach to the high-speed backplane to reach the eight slots in the rear of the chassis, which provide the heavy-duty interconnections for the entire QFX3000-G QFabric system. In addition, four interfaces (two per Control Board) provide Gigabit Ethernet access to the control plane management network. [Figure 6 on page 31](#) shows an example of the data plane and control plane connections for QFX3008-I Interconnect devices.

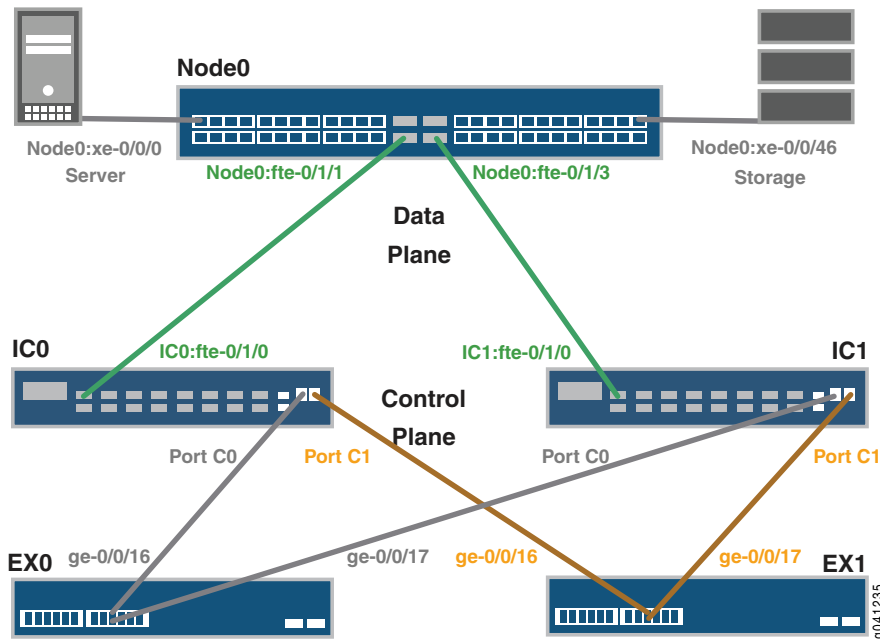
Figure 6: QFX3008-I Data Plane and Control Plane Connections



QFX3600-I Interconnect Devices

The QFX3600-I Interconnect device has 16 40-Gbps quad small form-factor pluggable plus (QSFP+) ports that provide interconnections for the entire QFX3000-M QFabric system. In addition, two management ports provide Gigabit Ethernet access to the control plane management network. [Figure 7 on page 32](#) shows an example of the data plane and control plane connections for a QFX3600-I Interconnect device.


Figure 7: QFX3600-I Data Plane and Control Plane Connections



QFX5100-24Q Interconnect Devices

The QFX5100-24Q switch can be configured to operate either as an Interconnect device or as a Node device in a QFX3000-M QFabric system. The QFX5100-24Q has 24 40-Gbps QSFP+ ports and can hold an additional 8 40-Gbps QSFP+ ports. These QSFP+ ports provide interconnects for the entire QFabric system. QFX5100 devices have a minimum of two management ports: a 10/100/1000BASE-T RJ-45 port and a 1-Gbps SFP ports. If you plan to use all fiber connections, be sure to order the product SKUs have an additional 1-Gbps SFP port that can be used either for fiber or copper connections.

Figure 8: QFX5100-24Q Data Plane and Control Plane Connections



RELATED DOCUMENTATION

[Understanding Node Devices | 33](#)

[Understanding the QFabric System Data Plane | 53](#)

[Understanding the QFabric System Control Plane | 49](#)

[Understanding the QFabric System Hardware Architecture | 21](#)

Understanding Node Devices

IN THIS SECTION

- [Node Device Introduction | 34](#)
- [QFX3500 Node Devices | 34](#)
- [QFX3600 Node Devices | 35](#)
- [QFX5100 Node Devices | 36](#)

Node devices in a QFabric system provide a way for servers, storage devices, and external networks to connect to the QFabric system. By understanding the role of Node devices, you can design your QFabric system topology to take advantage of the unique benefits offered by a single-tier data center architecture.

Node Device Introduction

A *Node device* in the QFabric system connects either endpoint systems (such as application servers and storage devices) or external networks to Interconnect devices. It can be used similarly to the way a top-of-rack switch is implemented in a data center. Node devices provide an access point to the QFabric system, allowing data to flow into and out of the QFabric system. Because all Node devices in the QFabric system connect through a backplane of Interconnect devices, in essence all Node devices are connected to one another. This directly connected design model eliminates multiple tiers of aggregation and core devices and provides minimum latency, maximum scalability, and rapid transport of server-to-server traffic and QFabric system-to-external network traffic.

Sets of Node devices can be bundled together into *Node groups*, in which each group operates as a single virtual entity. Node groups that connect to servers and storage devices are known as *server Node groups*, and Node groups that connect to external networks are known as *network Node groups*.

QFX3500 Node Devices

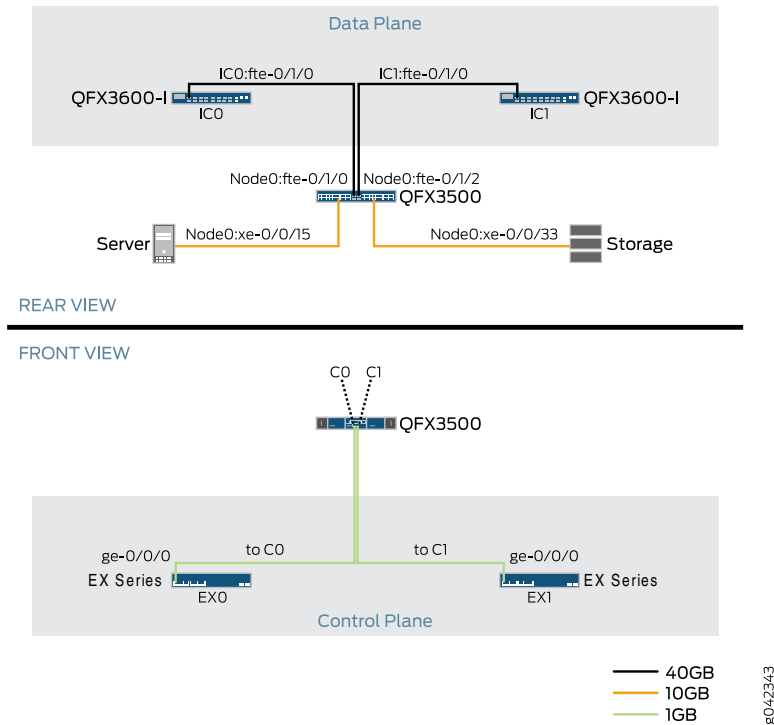
A QFX3500 Node device provides up to 48 10-Gigabit Ethernet interfaces to connect to endpoints or external networks. You can configure 12 of these 48 interfaces to support 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel. You can also configure the remaining 36 interfaces with Gigabit Ethernet.

NOTE: You can configure interface ports **0** through **47** as 10-Gigabit Ethernet ports, **0** through **5** and **42** through **47** as Fibre Channel over Ethernet ports, and **6** through **41** as Gigabit Ethernet ports. However, you cannot configure any Fibre Channel over Ethernet ports as Gigabit Ethernet ports or vice versa.

In addition to these server and network interfaces, there are four uplink interfaces to connect the QFX3500 Node device to Interconnect devices in a QFabric system. These uplinks use 40-Gbps quad small form-factor pluggable plus (QSFP+) interfaces.

The control plane requires two management ports on the QFX3500 chassis to connect the Node device to the control plane network. [Figure 9 on page 35](#) shows an example of the data plane and control plane connections for a QFX3500 Node device.

Figure 9: QFX3500 Data Plane and Control Plane Connections




QFX3600 Node Devices

A QFX3600 Node device provides 16 40-Gbps QSFP+ interfaces. By default, 4 interfaces (labeled **Q0** through **Q3**) are configured for 40-Gbps uplink connections between your QFX3600 Node device and your Interconnect device, and 12 interfaces (labeled **Q4** through **Q15**) use QSFP+ direct-attach copper (DAC) breakout cables or QSFP+ transceivers with fiber breakout cables to support 48 10-Gigabit Ethernet interfaces for connections to either endpoint systems or external networks. Optionally, you can choose to configure the first eight interfaces (**Q0** through **Q7**) for uplink connections between your Node device and your Interconnect devices, and interfaces **Q2** through **Q15** for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks.

The control plane requires two management ports on the QFX3600 chassis to connect the Node device to the control plane network. [Figure 10 on page 36](#) shows an example of the data plane and control plane connections for a QFX3600 Node device.

Figure 10: QFX3600 Data Plane and Control Plane Connections



QFX5100 Node Devices

Three models of the QFX5100 line of switches are supported as Node devices on a QFabric system:

- QFX5100-48S
- QFX5100-48T
- QFX5100-24Q

Figure 11 on page 37 shows an example of the data plane and control plane connections for a QFX5100 Node device.

Figure 11: QFX5100 Data Plane and Control Plane Connections



- QFX5100-48S

A QFX5100-48S Node device provides 48 10-Gigabit Ethernet interfaces to connect to endpoints or external networks and 6 40-Gbps QSFP+ interfaces. By default, 4 of the QSFP+ interfaces (labeled **fte-0/1/0** through **fte-0/1/1**) are configured for 40-Gbps uplink connections between your Node device and your Interconnect devices, and 2 QSFP+ interfaces (labeled **xle-0/1/4** and **xle-0/1/5**) provide 40-Gigabit Ethernet connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the middle two interfaces (**xle-0/1/8** and **xle-0/1/15**) for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks, and you can choose to configure the last two interfaces (**fte-0/1/4** and **fte-0/1/5**) for uplink connections between your Node device and your Interconnect devices.

- QFX5100-48T

A QFX5100-48T Node device provides 48 10GBASE-T interfaces to connect to endpoints or external networks and 6 40-Gbps QSFP+ interfaces.

By default, 4 of the QSFP+ interfaces (labeled **fte-0/1/0** through **fte-0/1/3**) are configured for 40-Gbps uplink connections between your Node device and your Interconnect devices, and 2 QSFP+ interfaces (labeled **xle-0/1/4** and **xle-0/1/5**) provide 40-Gigabit Ethernet connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the middle two interfaces (**xle-0/1/8** and **xle-0/1/15**) for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks, and you can choose to configure the last

two interfaces (**fte-0/1/4** and **fte-0/1/5**) for uplink connections between your Node device and your Interconnect devices.

- QFX5100-24Q

A QFX5100-24Q Node device provides 24 QSFP+ interfaces **fte-0/1/0** through **fte-0/1/7** as uplinks and **xle-0/1/8** to **xle-0/1/23** as endpoint systems or external networks.

The QFX5100-24Q has two expansion bays. With the optional QFX-EM-4Q expansion modules, the QFX5100-24Q can provide an additional 8 40-Gbps interfaces that are channelized into 10G ports. The QFX-EM-4Q expansion modules only support 40 Gigabit XLE interfaces. Valid interfaces are in the range of **xle-0/2/0** to **xle-0/2/3** for the first expansion module and from **xle-0/3/0** to **xle-0/3/3** in the second expansion module.

The control plane requires two management ports on the QFX5100 chassis to connect the Node device to the control plane network.

RELATED DOCUMENTATION

Converting the Device Mode for a QFabric System Component 322
Configuring Aliases for the QFabric System 388
Configuring Node Groups for the QFabric System 402
Configuring the Port Type on QFX3600 Node Devices 408
Understanding Node Groups 38
Understanding Interconnect Devices 29
Understanding the QFabric System Data Plane 53
Understanding the QFabric System Control Plane 49
Understanding the QFabric System Hardware Architecture 21

Understanding Node Groups

IN THIS SECTION

- [Network Node Groups | 39](#)
- [Server Node Groups | 39](#)

Node groups help you combine multiple Node devices into a single virtual entity within the QFabric system to enable redundancy and scalability at the edge of the data center.

Network Node Groups

A set of one or more Node devices that connect to an external network is called a *network Node group*. The network Node group also relies on two external Routing Engines running on the Director group. These redundant *network Node group Routing Engines* run the routing protocols required to support the connections from the network Node group to external networks.

When configured, the Node devices within a network Node group and the network Node group Routing Engines work together in tandem as a single entity. By default, network Node group Routing Engines are part of the **NW-NG-0** network Node group but no Node devices are included in the group. As a result, you must configure Node devices to be part of a network Node group.

In a QFabric system deployment that requires connectivity to external networks, you can modify the automatically generated network Node group by including its preset name **NW-NG-0** in the Node group configuration. Within a network Node group, you can include a minimum of one Node device up to a maximum of eight Node devices. By adding more Node devices to the group, you provide enhanced scalability and redundancy for your network Node group.

NOTE: The QFabric system creates a single **NW-NG-0** network Node group for the default partition. You cannot configure a second network Node group inside the default partition. The remaining Node devices within the default partition are reserved to connect to servers, storage, or other endpoints internal to the QFabric system. These Node devices either can be retained in the automatically generated server Node groups or can be configured as part of a redundant server Node group.

Server Node Groups

A *server Node group* is a set of one or more Node devices that connect to servers or storage devices. Unlike Node devices that are part of a network Node group and rely on an external Routing Engine, a Node device within a server Node group connects directly to endpoints and implements the Routing Engine functions locally, using the local CPU built into the Node device itself.

There are two different server Node group types:

- **Autogenerated server Node group**—By default, each Node device is placed in its own self-named Node group to connect to servers and storage. No configuration is necessary, but the QFabric system provides no redundancy for this type of Node group.

- **Redundant server Node group**—You can override the default, autogenerated server Node group assignment by manually configuring a *redundant server Node group* that contains a maximum of two Node devices. You can use a redundant server Node group to provide multihoming services to servers and storage, as well as configure aggregated LAG connections that span the two Node devices.

NOTE: The Node devices in a redundant server Node group must be of the same type, either two QFX3500 Node devices, two QFX3600 Node devices, or two QFX5100 Node devices. You cannot mix and match different Node device models in the same redundant server Node group.

RELATED DOCUMENTATION

[Configuring Node Groups for the QFabric System | 402](#)

[Understanding Node Devices | 33](#)

[Understanding Routing Engines in the QFabric System | 27](#)

[Understanding the QFabric System Hardware Architecture | 21](#)

Understanding Port Oversubscription on Node Devices

Each Node device in a QFabric system can have a different port oversubscription configuration. For example, you can have a Node device with 3:1 port oversubscription, another with 6:1 oversubscription, and yet another with 1:1 oversubscription.

The port oversubscription ratio on a Node device is based on the total amount of bandwidth from the server-facing connections and the total number of uplink connections to the Interconnect devices.. To determine your oversubscription ratio, multiply the number of server ports by the server-port speed, multiply the number of uplink ports by the uplink-port speed, and divide the total server-facing bandwidth by the total uplink-facing bandwidth. For example, If you use 32 10-Gigabit Ethernet server ports (320 gigabits) and 8 40-Gigabit Ethernet uplink ports (320 gigabits) on a QFX3600 Node device, you can configure 1:1 port oversubscription by connecting the eight uplink ports (labeled Q0 through Q7) from the Node device to the Interconnect devices and splitting the remaining 8 40-Gigabit Ethernet ports into 32 10-Gigabit Ethernet server ports.

Other important factors to consider are:

- When you connect more than one port on a Node device to an Interconnect device, you reduce the overall number of Node devices supported on the QFabric system.

- When you have two Interconnect devices and a single connection from each Node device, you can connect any remaining ports as access ports not server ports.

[Table 7 on page 41](#) shows the oversubscription ratio for ports on QFX3500 Node devices in default mode based on the number of Interconnect devices and the number of connections from each Node device to each Interconnect device.

Table 7: Oversubscription Ratio on QFX3500 Node Devices

Number of Interconnect Devices	Number of Connections from Each Node Device to Each Interconnect Device	Oversubscription Ratio on Node Device
2	1	6:1
2	2	3:1
4	1	3:1

[Table 8 on page 41](#) shows the oversubscription ratio for ports on QFX3600 Node devices in default mode based on the number of Interconnect devices and the number of connections from each Node device to each Interconnect device.

Table 8: Oversubscription Ratio on QFX3600 Node Devices

Number of Interconnect Devices	Number of Connections from Each Node Device to Each Interconnect Device	Oversubscription Ratio on Node Device
2	1	6:1
2	2	3:1
2	4	1:1
4	1	3:1
4	2	1:1

[Table 9 on page 42](#) shows the oversubscription ratio for ports on QFX5100-48S Node devices and QFX5100-48T Node devices based on the number of Interconnect devices and the number of connections from each Node device to each Interconnect device.

Table 9: Oversubscription Ratio on QFX5100-48S and QFX5100-48T Node Devices

Number of Interconnect Devices	Number of Connections from Each Node Device to Each Interconnect Device	Oversubscription Ratio on Node Device
2	2	3.5:1
2	1	8:1
4	1	3.5:1

Unlike other Node Devices, the QFX5100-24Q supports the QFX-EM-4Q expansion module, that impact the oversubscription ratio. [Table 10 on page 42](#) shows the oversubscription ratio for the base model without expansion modules inserted into the device and with two QFX-EM-4Q modules installed. Other factors that impact the oversubscription ratio include the system mode selected for the device. See [Table 11 on page 43](#) for details on system mode.

Table 10: Oversubscription Ratio on QFX5100-24Q Node Devices

Number of Interconnect Devices	Number of Connections from Each Node Device to Each Interconnect Device	Oversubscription Ratio on Node Device	Oversubscription Ratio on Node Device with Two QFX-EM-4Q Installed
2	1	8:1	12:1
2	2	4:1	6:1
2	4	2:1	3:1
2	8	N/A	1:1
4	1	4:1	6:1
4	2	2:1	3:1
4	4	N/A	1:1

[Table 11 on page 43](#) shows the ports the default port configuration and identifies the ports that can be converted on Node devices. Because the QFX5100-24Q has expansion bays that allow you to add additional types of ports, see [Table 12 on page 44](#) for the supported system modes and default configuration of that Node device.

Table 11: Port Configurations on Node Devices

Model	Total Ports	Fixed FTE Ports	Ports that can be configured FTE or XLE	Ports that can be configured XLE or XE	Default Configuration
QFX3500	<ul style="list-style-type: none"> • 48 x 10G • 4 x 40G 	Q0 through Q3	NA	NA	<ul style="list-style-type: none"> • 48 x 10G (xe-0/0/0 through xe-0/0/47) • 4 x 40G (fte-0/1/0 through fte-0/1/3)
QFX3600	16 x 40G	Q0 and Q1	Q2 through Q7	Q2 through Q15	<ul style="list-style-type: none"> • 12 x 40G xle-0/1/4 through xle-0/1/15) • 4 x 40G (fte-0/1/0 through fte-0/1/3)
QFX5100-48S and QFX5100-48T	<ul style="list-style-type: none"> • 48 x 10G • 6 x 40G 	0 and 1	2 through 5	2 through 5	<ul style="list-style-type: none"> • 2 x 40G (xle-0/1/4 through xle-0/1/5)) • 4 x 40G fte-0/1/0 through fte-0/1/3) • 48 x 10G (xe-0/0/0 through xe-0/0/47)
			NOTE: Port 2 and port 3 can be configured together as either as XE or XLE. Port 4 and port 5 can be independently configured as XE or XLE.		

The QFX5100-24Q can be configured up to 104 10G ports using different system modes to achieve varying levels of port density. See [Table 12 on page 44](#) for the default configurations for each system mode.

Table 12: QFX5100-24Q System Mode Default Port Configuration

System Mode	Total Ports	Fixed FTE Ports	Ports that can be configured FTE or XLE	Ports that can be configured XLE or XE	Disabled	Default Configuration	QFX-EM-4Q Expansion Modules
Default	24 x 40G (base configuration) optional expansion model configuration of two modules: <ul style="list-style-type: none"> 4 x 4 x 10G (QFX-EM-4Q) 	0 through 7	8 through 15	8 through 23 NOTE: Of these 16 available ports, only 12 may be channelized.	NA	<ul style="list-style-type: none"> 16 x 40G (xle-0/1/8 through xle-0/1/23) 8 x 40G (fte-0/1/0 through fte-0/1/7) 	Ports in the expansion modules are xe by default. The expansion module ports are supported, but cannot be channelized. These ports also cannot be converted to XLE or FTE ports (Xe-0/2/0 through xe-0/2/15 for 1 expansion module and x-0/3/0 and 0/3/15 for a second expansion module)

RELATED DOCUMENTATION

[Connecting a QFX3500 Node Device to a QFX3008-I Interconnect Device](#)

[Connecting a QFX3600 Node Device to a QFX3008-I Interconnect Device](#)

[Connecting a QFX5100 Node Device to a QFX3008-I Interconnect Device](#)

[Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device | 311](#)

[Connecting a QFX3600 Node Device to a QFX3600-I Interconnect Device | 306](#)

[Connecting a QFX5100 Node Device to a QFX3600-I Interconnect Device | 301](#)

Software Architecture Overview

IN THIS CHAPTER

- [Understanding the QFabric System Software Architecture | 45](#)
- [Understanding the Director Software | 46](#)
- [Understanding Partitions | 47](#)
- [Understanding the QFabric System Control Plane | 49](#)
- [Understanding the QFabric System Data Plane | 53](#)
- [Understanding QFabric Multicast Flow Groups | 55](#)

Understanding the QFabric System Software Architecture

The software architecture for the QFabric system environment has been designed to provide a high-speed, low-latency, nonblocking fabric for data center traffic. This topic explores how the software architecture for a QFabric system supports these goals.

Key components of the QFabric system software architecture include:

- A single administrative view of all QFabric system components provides unified management, configuration, monitoring, and troubleshooting of the QFabric system. This view is provided by the QFX Series Director software running on the Director group. A primary administrator can access the unified view through the default partition.
- A fabric control protocol enables rapid transport of data traffic between QFabric system components. This unique feature of the software architecture distributes route information for each device within the QFabric system, and removes the need to run spanning-tree protocols inside the QFabric system network.
- A fabric management protocol provides rapid transport of control traffic between QFabric system components. This protocol helps identify and initialize QFabric system resources, supports device redundancy, and supports management communication throughout the QFabric system.
- A control plane network that is separate from the data plane network provides high availability for the QFabric system.

The software also provides access to relevant features in the Junos operating system (Junos OS) that support QFabric system functionality. Support is available for most switching features available on EX Series Ethernet switches and many routing features available on M Series, MX Series, and T Series routing platforms.

RELATED DOCUMENTATION

Understanding QFabric System Terminology 7
Understanding the QFabric System Hardware Architecture 21
Understanding the Director Software 46
Understanding Partitions 47
Understanding the QFabric System Control Plane 49
Understanding the QFabric System Data Plane 53
Understanding QFabric Multicast Flow Groups 55

Understanding the Director Software

The Director software provides a single view into the QFabric system so that it can be managed as a single entity. This topic explains how the Director software interacts with the components of the QFabric system to maintain operations from a central location.

Because the QFabric system consists of multiple Director, Node, and Interconnect devices, the architects of the QFabric system determined that it would be useful to manage the entire system as a single logical entity. As a result, the Director software handles administration tasks for the entire QFabric system, such as fabric management and configuration. The Director software runs on the *Director group*, provides a single consolidated view of the QFabric system, and enables the main QFabric system administrator to configure, manage, monitor, and troubleshoot QFabric system components from a centralized location. In the Junos operating system (Junos OS) command-line interface (CLI), you can access the Director software by logging in to the default partition.

The Director software handles the following major tasks for the QFabric system:

- Provides command-line interface (CLI) access to all QFabric system components that you have permission to manage or view.
- Evaluates configuration statements and operational mode commands for their scope and sends requests to the applicable Director, Node, and Interconnect devices. (This operation is sometimes referred to as *scattering*.)

- Consolidates responses from Director, Node, and Interconnect devices, and displays output from the devices in a unified, centralized manner. (This operation is sometimes referred to as *gathering*.)
- Coordinates configuration and operational efforts with a database housed in the Director group to store and retrieve configurations, software images, event logs, and system log messages.
- Facilitates control plane communication between the Node devices, the Routing Engine services running on the Director group, and the Interconnect devices.
- Runs parallel processes on the Director group devices to provide high availability for the QFabric system.
- Coordinates interactions with QFabric system components to provide load balancing of processing tasks across the Director group devices.
- Manages user access and privileges.
- Enables you to configure, manage, monitor, and troubleshoot QFabric system components that are assigned to you.
- Gathers QFabric system inventory and topology details.
- Offers a way to manage Director group devices, including the ability to add and delete Director devices in the group, set and switch mastership in the Director group, and monitor Director group status.
- Provides a centralized way to coordinate software upgrades for QFabric system components.

The Director software provides a backbone of functionality that supports the entire QFabric system. It is an essential component of the QFabric system that enables you to implement the system in a logical and efficient way.

RELATED DOCUMENTATION

[Gaining Access to the QFabric System Through the Default Partition | 375](#)

[Understanding the Director Group | 26](#)

[Understanding the QFabric System Software Architecture | 45](#)

Understanding Partitions

IN THIS SECTION


- [QFabric System Default Partition | 48](#)

Partitions provide a way to allocate specified virtual and physical resources within your QFabric system. This topic covers:

QFabric System Default Partition

By default, all equipment and virtual resources in the QFabric system belong to the *default partition*. As a result, the QFabric system in its initial state has a single broadcast domain that is administered by a single main administrator. [Figure 12 on page 48](#) shows a topology with the default settings—a single collection that contains all the devices in the QFabric system.

Figure 12: QFabric System Topology - Default Partition



NOTE: The initial release of the QFabric system supports a single default partition. All equipment and resources belong to the default partition.

A partition provides the following functions:

- Fault isolation and separation from other partitions at the control plane level.
- A separate configuration domain for the Node devices within the partition.
- A Layer 2 domain in which MAC learning takes place, and members of the same VLAN can communicate with each other. To provide network connectivity between partitions, you need to enable Layer 3 routing by way of a routed VLAN interface (RVI).

RELATED DOCUMENTATION

[Gaining Access to the QFabric System Through the Default Partition | 375](#)

[Understanding the QFabric System Software Architecture | 45](#)

[Understanding the QFabric System Hardware Architecture | 21](#)

Understanding the QFabric System Control Plane

IN THIS SECTION


- [Control Plane Elements | 50](#)
- [Control Plane Services | 52](#)

The control plane in the QFabric system transports management traffic between QFabric system components to facilitate system operations, configuration, and maintenance. This topic covers:

Control Plane Elements

Control traffic within a QFabric system is carried across a redundant, scalable, out-of-band, Ethernet switching network called the *control plane* network. To maintain high availability, the QFabric system control plane is separated from the QFabric system data plane. [Figure 13 on page 50](#) shows a diagram of the QFabric system devices that compose the control plane network.

Figure 13: QFabric System Control Plane Network



The control plane consists of the following elements:

- **Control plane switches**—Provide connectivity to the management interfaces of all QFabric system components in the control plane network, including the Node devices, the Interconnect devices, and the Director group. When you interconnect all QFabric system devices to the control plane switches, the Director group can manage the entire system. Depending on the size and scale of your QFabric system, the control plane switches might be standalone switches or might be groups of switches bundled into a Virtual Chassis (See the Example topics in the Related Documentation section of this topic to learn more about the control plane switch configuration required for your QFabric system.)

For example, the control plane for the QFX3000-G QFabric system requires two Virtual Chassis containing either four EX4200 switch members or four EX4300 switch members each. The two Virtual Chassis connect to each other across a 10-Gigabit Ethernet LAG link to provide maximum resiliency for the QFabric system control plane. The control plane for the QFX3000-M QFabric system requires either two EX4200 switches or two EX4300 switches, interconnected and configured in a LAG for redundancy.

NOTE: You cannot mix EX4200 switches and EX4300 switches in the control plane in a QFabric system.

- **Connections between the management interfaces of the Node devices and the control plane switches**—Enable control plane connectivity from the Node devices to the rest of the QFabric system. You must connect two management interfaces from each Node device to the control plane switches. Connect each interface to a different control plane switch to provide system resiliency.
- **Connections between the management interfaces of the Interconnect devices and the control plane switches**—Enable control plane connectivity from the Interconnect devices to the rest of the QFabric system. You must connect the interfaces in each Interconnect device to the control plane switches. Connect each interface to a different control plane switch to provide system resiliency.

For example, on QFX3008-I Interconnect devices in a QFX3000-G QFabric system, there are two Control Boards and two interfaces per Control Board, for a total of four connections per Interconnect device. To provide system resiliency, connect one interface from each Control Board to the first control plane Virtual Chassis, and connect the second interface from each Control Board to the second control plane Virtual Chassis.

- **Connections between the network module interfaces of the Director group and the control plane switches**—Enable control plane connectivity from the Director group to the rest of the QFabric system. You must connect some interfaces from the first network module in a Director device to one control plane switch, and connect some interfaces from the second network module in a Director device to the second control plane switch. Also, you must connect the ports from the first network module to the primary control plane switch for each Director device (which may vary depending on the configuration of your Director group).

For guidance on the QFabric control plane configuration and cabling recommendations, see:

- *Example: Configuring the Virtual Chassis for a Copper-Based QFX3000-G QFabric System Control Plane*

- *Example: Configuring a Fiber-Based Control Plane for the QFX3000-G QFabric System*
- [Example: Configuring EX Series Switches for the QFX3000-M QFabric System Control Plane on page 329](#)
- **Routing Engines**—Although they are automatically provisioned, specialized Routing Engines implement services such as default QFabric system infrastructure, device management, route sharing, and diagnostics to support the QFabric system. Routing Engines for control plane functions are virtual entities that run on the Director group.
- **Fabric management protocol**—A link-state protocol runs on the control plane network to identify and initialize QFabric system resources, support device redundancy, and support management communication throughout the QFabric system. The protocol is enabled by default.

Control Plane Services

The QFabric system control plane provides the infrastructure to support the following services for the QFabric system:

- System initialization
- Topology discovery
- Internal IP address and unique ID assignment
- Route information sharing
- Configuration delivery to Node devices
- Interdevice communication between Node devices, Interconnect devices, and the Director group

Many of these services are provided by the external Routing Engines that run in software on the Director group.

RELATED DOCUMENTATION

Example: Configuring the Virtual Chassis for a Copper-Based QFX3000-G QFabric System Control Plane

Example: Configuring a Fiber-Based Control Plane for the QFX3000-G QFabric System

[Example: Configuring EX Series Switches for the QFX3000-M QFabric System Control Plane | 329](#)

[Understanding the QFabric System Data Plane | 53](#)

[Understanding Routing Engines in the QFabric System | 27](#)

[Understanding the QFabric System Hardware Architecture | 21](#)

Understanding the QFabric System Data Plane

IN THIS SECTION

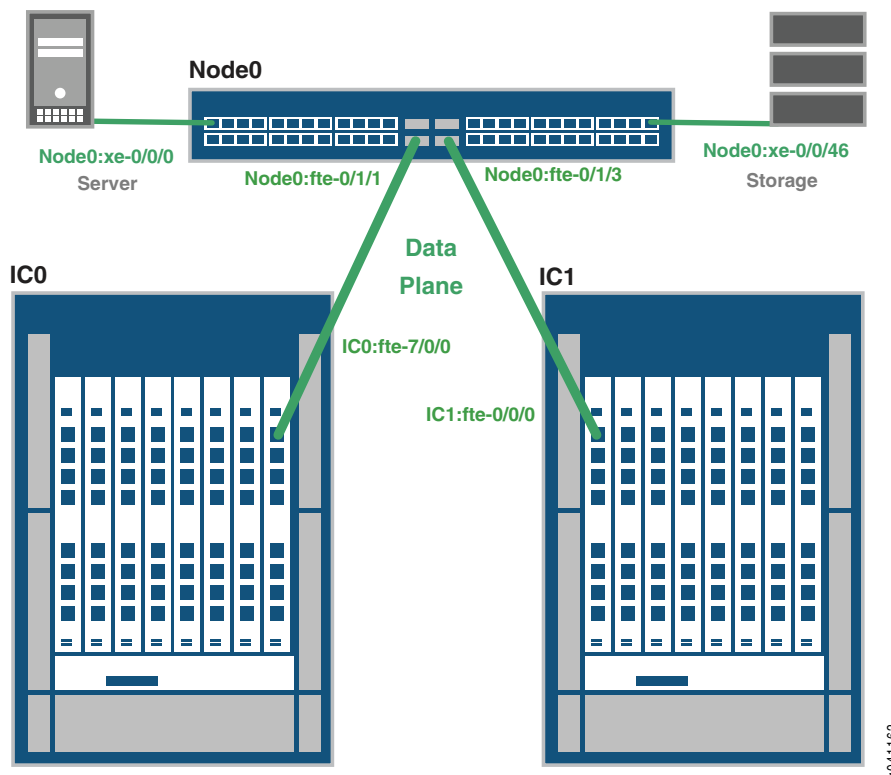
- Data Plane Components | 53
- QFabric System Fabric | 54

The data plane in the QFabric system transfers application traffic between QFabric system components rapidly and efficiently. This topic covers:

Data Plane Components

Data traffic within a QFabric system is carried across a redundant, high-performance, and scalable *data plane*. To maintain high availability, the QFabric system data plane is separated physically from the QFabric system control plane and uses a different network. [Figure 14 on page 53](#) shows an example diagram of the QFabric system data plane network.

Figure 14: QFabric System Data Plane Network




The QFabric system data plane includes the following high-speed data connections and elements:

- 10-Gigabit Ethernet or 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel connections between QFabric system endpoints (such as servers or storage devices) and the Node devices.
- 40-Gbps quad, small form-factor pluggable plus (QSFP+) connections between the Node devices and the Interconnect devices.
- 10-Gigabit Ethernet connections between external networks and the Node devices contained in the network Node group.
- A fabric control protocol, used to distribute route information to all devices connected to the QFabric system data plane.

QFabric System Fabric

Unlike traditional data centers that employ a multi-tiered hierarchy of switches, a QFabric system contains a single tier of Node devices connected to one another across a backplane of Interconnect devices. The QFabric system fabric is a distributed, multistage network that consists of a fabric queuing and scheduling system implemented in the Node devices, and a distributed cross-connect system implemented in the Interconnect devices. The cross-connect system for the QFX3008-I Interconnect device is shown as an example in [Figure 15 on page 54](#).

Figure 15: QFX3008-I Interconnect Device Cross-Connect System



The design of the cross-connect system provides multistage Clos switching, which results in nonblocking paths for data traffic and any-to-any connectivity for the Node devices. Because all Node devices are

connected through the Interconnect device, the QFabric system offers very low port-to-port latencies. In addition, dynamic load balancing and low-latency packet flows provide for scaling the port count and bandwidth capacity of a QFabric system.

RELATED DOCUMENTATION

[Understanding the QFabric System Control Plane | 49](#)

[Understanding the QFabric System Hardware Architecture | 21](#)

Understanding QFabric Multicast Flow Groups

IN THIS SECTION

- [How Do I Implement a Flow Group in a Fabric? | 55](#)
- [Do Flow Groups Impact QFabric System Performance? | 56](#)
- [What Platforms Support Flow Groups? | 56](#)
- [How Do Multicast Flow Groups Work? | 56](#)
- [What are QFabric Multicast Flow Group Limitations? | 57](#)

In a QFabric system, a hash function is used to select an Interconnect device to forward traffic between two Node devices. Since this hash function is performed on all Interconnect devices, it is possible for redundant multicast streams to flow through one Interconnect device, making that Interconnect device a potential single point of failure for the redundant flows. Some applications, such as financial transactions, require that the redundant multicast streams flow through different Interconnect devices to prevent a single Interconnect device from potentially dropping both streams of multicast traffic during a failure. You can enforce this use of diverse Interconnect devices by using the QFabric flow groups feature.

For an explanation of QFabric system, see [“QFabric System Overview” on page 3](#).

How Do I Implement a Flow Group in a Fabric?

There are two ways to configure a QFabric system, from the CLI and from Network Director. Flow control using flow groups for QFabric systems, however, can be configured from the CLI only, using the [flow-groups](#) statement.

You create a flow group from the CLI by designating the name of the flow group, the Node devices used by the flow group, and the Interconnect devices used by the nodes in the flow group.

When you configure a flow group, the software assigns a priority to each used Interconnect device. You can override the software settings using the CLI.

Do Flow Groups Impact QFabric System Performance?

Flow group configuration does not affect L3 multicast scaling numbers or performance.

What Platforms Support Flow Groups?

32 bit QFabric-G and QFabric-M platforms support flow groups. For more information about these platforms, see *Understanding QFX3000-G QFabric System Hardware Configurations* and [“Understanding QFX3000-M QFabric System Hardware Configurations” on page 112](#).

How Do Multicast Flow Groups Work?

Link preferences, which restrict the traffic from ingress Node devices to Interconnect devices, are set to **normal** by default when they are not yet a part of a flow group. When you create a flow group, the link preferences between indicated Interconnect devices and Node devices are automatically changed from **normal** to either **high** or **never**, meaning that any preference determined by the normal algorithm is ignored and the flow is determined by the flow group link settings.

NOTE: You can override the link preference by using the **preference** option of the **flow-groups** CLI statement.

The three link preferences used by the fabric software are:

- **high**—Use the indicated Interconnect device in the hashing algorithm for the specified flow.
- **normal**—Use the indicated Interconnect device for backup when Interconnect devices set to **high** fail.
- **never**—Do not use the indicated Interconnect device in the hashing algorithm for the specified flow.

Interconnect device preference settings are assigned by the software. When you configure a flow group, preferences for some flows are set to **high** and some are set to **never**. This is how flows are enforced. Unassigned Interconnect devices (and Node devices) all belong to a default group and are set to **normal**. When an assigned Interconnect device fails, the corresponding flow looks for a **normal** Interconnect device to use as backup if one is available. This helps prevent loss of data in the event of an Interconnect device failure.

What are QFabric Multicast Flow Group Limitations?

- An Interconnect device can belong to only one flow group.
- A QFabric system can contain up to four flow groups.
- By default, Node devices in a flow group cannot use Interconnect devices in other flow groups. You can override this setting using the **preference** option of the **flow-groups** CLI statement.

RELATED DOCUMENTATION

[Segregating QFabric Traffic Flows With Flow Groups | 429](#)

[Understanding QFabric System Terminology | 7](#)

[Understanding the QFabric System Hardware Architecture | 21](#)

[Understanding the QFabric System Software Architecture | 45](#)

[Example: Creating a QFabric Flow Group | 431](#)

Software Features

IN THIS CHAPTER

- Understanding Software Upgrade on the QFabric System | 59
- Understanding Nonstop Software Upgrade for QFabric Systems | 61
- Understanding Statements and Commands on the QFabric System | 66
- Understanding NTP on the QFabric System | 68
- Understanding Network Management Implementation on the QFabric System | 68
- Understanding the Implementation of SNMP on the QFabric System | 69
- Understanding the Implementation of System Log Messages on the QFabric System | 72
- Understanding User and Access Management Features on the QFabric System | 74
- Understanding QFabric System Login Classes | 75
- Understanding Interfaces on the QFabric System | 76
- Understanding Layer 3 Features on the QFabric System | 81
- Understanding Telnet on the QFabric System | 82
- Understanding Security Features on the QFabric System | 83
- Understanding Port Mirroring | 84
- Understanding Fibre Channel Fabrics on the QFabric System | 91
- Understanding CoS Fabric Forwarding Class Sets | 93

Understanding Software Upgrade on the QFabric System

IN THIS SECTION

- Operational Software Commands | 60
- Operational Reboot Commands | 60

The QFabric system software package contains software for the QFabric system infrastructure and for all of the different component devices in the QFabric system: Director group, Interconnect devices, and Node devices.

Operational Software Commands

The **request system software download** CLI command enables you to download the software package to various locations: for example, USB device, remote server, or FTP site.

The following CLI commands enable you to install the software for the Director group, Interconnect devices, Node devices, and the QFabric system infrastructure. You may need to specify the **reboot** option depending on which devices or QFabric infrastructure you are installing the software. The **reboot** option works differently depending on whether you install the software on the QFabric system infrastructure or on a particular device in the QFabric system.

- **request system software add component all**

This command installs software for the Director group, fabric control Routing Engine, fabric manager Routing Engine, Interconnect devices, and network and server Node groups.

- **request system software add component director-group**

This command installs software for the Director group and the default partition, which is where you access the QFabric system CLI.

- **request system software add component fabric**

This command installs the software for the fabric control Routing Engines and the Interconnect devices.

- **request system software add component *node-group-name***

This command installs software for a server Node group or a network Node group.

Additionally, you can back up your current QFabric configuration file and installation-specific parameters using the **request system software configuration-backup** command. We recommend that you save this file to an external location, like an FTP site or USB device, but you can save it locally.

Operational Reboot Commands

The following commands enable you to reboot the entire QFabric system, various Node devices, or the QFabric system infrastructure:

- **request system reboot all**

This command reboots the Director group, fabric control Routing Engines, fabric manager Routing Engine, Interconnect devices, and network and server Node groups.

- **request system reboot director-group**

This command reboots the Director group and the default partition, which is where you access the QFabric system CLI.

- **request system reboot fabric**

This command reboots the fabric control Routing Engines and the Interconnect devices.

- **request system reboot node-group**

This command reboots a server Node group or a network node group.

RELATED DOCUMENTATION

| [Upgrading Software on a QFabric System | 566](#)

Understanding Nonstop Software Upgrade for QFabric Systems

The framework that underlies a nonstop software upgrade in a QFabric system enables you to upgrade the system in a step-by-step manner and minimize the impact to the continuous operation of the system. This topic explains how a nonstop software upgrade works in a QFabric system, the steps that are involved, and the procedures that you need to implement to experience the benefits of this style of software upgrade.

Nonstop software upgrade enables some QFabric system components to continue operating while similar components in the system are being upgraded. In general, the QFabric system upgrades redundant components in stages so that some components remain operational and continue forwarding traffic while their equivalent counterparts upgrade to a new version of software.

TIP: Use the following guidelines to decide when to implement a nonstop software upgrade:

TIP: Before you perform a nonstop software upgrade, contact JTAC to perform a pre-upgrade health check on the QFabric system.

- If you need to upgrade all components of the system in the shortest amount of time (approximately one hour) and you do not need to retain the forwarding resiliency of the data plane, issue the **request system software add component all** command to perform a standard software upgrade. All components of the QFabric system upgrade simultaneously and expediently, but this type of upgrade does not provide resiliency or switchover capabilities.
- If you need to minimize service impact, preserve the forwarding operations of the data plane during the upgrade, and are willing to take the extra time required for component switchovers (in many cases, several hours), issue the three nonstop software upgrade commands (**request system software nonstop-upgrade (director-group | fabric | node-group)** described in this topic in the correct order.

NOTE:

- Before you begin a nonstop software upgrade, issue the **request system software download** command to copy the software to the QFabric system.
- Each of the 3 nonstop software upgrade steps must be considered parts of the whole process. You must complete all 3 steps of a nonstop software upgrade in the correct order to ensure the proper operation of the QFabric system.
- Open two SSH sessions to the QFabric CLI. Use one session to monitor the upgrade itself and use a second session to verify that the QFabric system components respond to operational mode commands as expected. For more information on verification of the upgrade, see [“Verifying Nonstop Software Upgrade for QFabric Systems” on page 540](#).
- Issue the **show fabric administration inventory** command to verify that all upgraded components are operational at the end of a step before beginning the next step.
- Once you start the nonstop software upgrade process, we strongly recommend that you complete all 3 steps within 12 hours.

The three steps to a successful nonstop software upgrade must be performed in the following order:

- **Director group**—The first step upgrades the Director devices, the fabric manager Routing Engine, and the diagnostic Routing Engine. To perform the first step, issue the **request system software**

nonstop-upgrade director-group command. The key actions that occur during a Director group upgrade are:

1. Connecting to the QFabric system by way of an SSH connection. This action establishes a load-balanced CLI session on one of the Director devices in the Director group.
2. The QFabric system downloads and installs the new software in both Director devices.
3. The Director device hosting the CLI session becomes the master for all QFabric system processes running on the Director group, such as the fabric manager and network Node group Routing Engines.
4. The QFabric system installs the new software for the backup fabric manager Routing Engine on the backup Director device.
5. The backup Director device reboots to activate the new software.
6. The master Director device begins a 15 minute sequence that includes a temporary suspension of QFabric services and a QFabric database transfer. You cannot issue operational mode commands in the QFabric CLI during this period.
7. The QFabric system installs the new software for the fabric manager and diagnostic Routing Engines on the Director group master.
8. The QFabric system switches mastership of all QFabric processes from the master Director device to the backup Director device.
9. The master Director device reboots to activate the new software.
10. The CLI session terminates, and logging back in to the QFabric system with a new SSH connection establishes the session on the new master Director device (the original backup).
11. The previous master Director device resumes operation as a backup and the associated processes (such as the fabric manager and network Node group Routing Engines) become backup as well. The fabric control Routing Engine associated with this Director device returns to active status.

NOTE: After the Director group nonstop software upgrade completes, any Interconnect device or Node device that reboots will automatically download the new software, install it, and reboot again. As a result, try not to restart any QFabric system devices before you complete the rest of the nonstop software upgrade steps.

TIP:

- To enable BGP and OSPF to continue operating on the network Node group during a Director group nonstop service upgrade, we recommend that you configure graceful restart for these routing protocols. For more information on graceful restart, see [“Configuring Graceful Restart for QFabric Systems” on page 424](#).
 - Wait 15 minutes after the second Director device returns to service and hosts Routing Engine processes before proceeding to step 2—the fabric upgrade. You can verify the operational status of both Director devices by issuing the **show fabric administration inventory director-group status** command. Also, issue the **show fabric administration inventory infrastructure** command to verify when the Routing Engine processes become load balanced (typically, there will be three to four Routing Engines running on each Director device).
- Fabric—The second step upgrades the Interconnect devices and the fabric control Routing Engines. To perform the second step, issue the **request system software nonstop-upgrade fabric** command. The key actions that occur during a fabric upgrade are:
 1. The QFabric system downloads, validates, and installs the new software in all Interconnect devices and fabric control Routing Engines (FC-0 and FC-1).
 2. One fabric control Routing Engine reboots and comes back online.
 3. The other fabric control Routing Engine reboots and comes back online.
 4. The first Interconnect device reboots, comes back online, and resumes the forwarding of traffic.
 5. Subsequent Interconnect devices reboot one at a time, come back online, and return to service.

NOTE:

- If the software does not load properly on any one of the fabric components, all components revert back to the original software version.
 - If one of the components in a fabric upgrade does not reboot successfully, issue the **request system reboot fabric** command to reattempt the rebooting process for this fabric component and activate the new software.
- Node group—The third and final step upgrades Node groups. You can choose to upgrade a network Node group, a redundant server Node group, or individual server Node groups. You can upgrade the Node groups one at a time or in groups (known as upgrade groups). However, you must upgrade all Node groups in your QFabric system before you can complete the nonstop software upgrade process. To perform the third step, issue the **request system software nonstop-upgrade node-group** command. The key actions that occur during a network Node group upgrade are:

1. The QFabric system copies the new software to each Node device one at a time.
2. The QFabric system validates and then installs the new software in all Node devices simultaneously.
3. The system copies the software to the network Node group Routing Engines.
4. The QFabric system validates and then installs the software in the network Node group Routing Engines one at a time -- first the backup, then the master.
5. The backup network Node group Routing Engine reboots and comes back online.
6. The supporting Node devices reboot and come back online one at a time.

NOTE: To reduce the total upgrade duration, configure an upgrade group. All Node devices within the upgrade group reboot at the same time.

7. The master network Node group Routing Engine relinquishes mastership to the backup, reboots, and comes back online.

The key actions that occur during a redundant server Node group upgrade are:

1. The QFabric system copies the new software to the backup Node device, then the master Node device.
2. The QFabric system validates and then installs the new software on the backup Node device, then the master Node device.
3. The backup Node device reboots, comes back online, and becomes the master Node device.
4. The previous master Node device reboots and comes back online as a backup Node device.

NOTE: For redundant server Node groups, both Node devices must be online before the upgrade will proceed. If one of the devices is no longer available, remove the Node device from the Node group configuration before you issue the nonstop software upgrade command.

The key actions that occur during a server Node group upgrade for a Node group that contains one member are:

1. The Node device downloads the software package and validates the software.
2. The Node device installs the software and reboots.

NOTE: Because there is no redundancy for Node groups containing a single Node device, traffic loss occurs when the device reboots during the upgrade.

RELATED DOCUMENTATION

[Nonstop Software Upgrade Checklist for QFabric Systems | 529](#)

[Performing a Nonstop Software Upgrade on the QFabric System | 533](#)

[Verifying Nonstop Software Upgrade for QFabric Systems | 540](#)

[request system software nonstop-upgrade | 639](#)

[request system software add](#)

[Configuring Graceful Restart for QFabric Systems | 424](#)

Understanding Statements and Commands on the QFabric System

IN THIS SECTION

- [Chassis Statements | 66](#)
- [Chassis Commands | 66](#)

Chassis Statements

The following chassis statements enable you to configure various options for your Interconnect devices, Node groups (network and server), and Node devices:

- **interconnect-device**
- **node-group**
- **node-device**

Chassis Commands

The Junos OS CLI contains additions to the existing chassis commands. These additions reflect new options as a result of adding the **interconnect-device**, **node-group**, and **node-device** chassis statements at the **[edit chassis]** hierarchy level.

The following chassis commands enable you to monitor and configure the QFabric system hardware and software options at various hierarchy levels:

- **clear chassis display message**
- **request chassis beacon**

- request chassis cb (QFX3000-G QFabric systems only)
- request chassis fabric (QFX3000-G QFabric systems only)
- request chassis fpc
- request chassis routing-engine master
- set chassis aggregated-devices
- set chassis alarm
- set chassis container-devices
- set chassis craft-lockout
- set chassis display
- set chassis fpc
- set chassis routing-engine
- show chassis alarms
- show chassis beacon
- show chassis environment
- show chassis fan (QFX3000-G QFabric systems only)
- show chassis fabric
- show chassis firmware
- show chassis fpc
- show chassis hardware
- show chassis lcd
- show chassis led
- show chassis location
- show chassis mac-addresses
- show chassis nonstop-upgrade
- show chassis pic
- show chassis routing-engine
- show chassis temperature-thresholds
- show chassis zones

RELATED DOCUMENTATION

[QFabric System Initial and Default Configuration Information | 319](#)

[Understanding User and Access Management Features on the QFabric System | 74](#)

[Generating the MAC Address Range for a QFabric System | 362](#)

[Performing the QFabric System Initial Setup on a QFX3100 Director Group | 363](#)

Understanding NTP on the QFabric System

Network Time Protocol (NTP) enables you to synchronize the time across the network. This is especially helpful for correlating log events and replicating databases and file systems. The QFabric system synchronizes time with servers that are external to the system and operates in client mode only.

To configure NTP, include the **server address** and **authentication-key** statements at the **[edit system ntp]** hierarchy level.

RELATED DOCUMENTATION

[NTP Time Server and Time Services Overview \(QFabric System\)](#)

[Synchronizing and Coordinating Time Distribution Using NTP](#)

[Configuring NTP Authentication Keys \(QFabric System\)](#)

[Configuring the NTP Time Server and Time Services \(QFabric System\)](#)

[Configuring the Switch to Listen for Broadcast Messages Using NTP](#)

[Configuring the Switch to Listen for Multicast Messages Using NTP](#)

[Example: Configuring NTP](#)

[Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization](#)

Understanding Network Management Implementation on the QFabric System

This topic describes network management features on the QFabric system that are implemented differently than on other devices running Junos OS.

The following network management features are supported on the QFabric system:

- **System log messages**—The QFabric system monitors events that occur on its component devices, distributes system log messages about those events to all external system log message servers (hosts) that are configured, and archives the messages. Component devices include Node devices, Interconnect

devices, Director devices, and the Virtual Chassis. You configure system log messages at the **[edit system syslog]** hierarchy level. Use the **show log filename** operational mode command to view messages.

- **Simple Network Management Protocol (SNMP) Version 1 (v1) and v2c**—SNMP monitors network devices from a central location. The SNMP implementation on the QFabric system supports the basic SNMP architecture of Junos OS with some limitations, including a reduced set of MIB objects, read-only access for SNMP communities, and limited support for SNMP requests. You configure SNMP at the **[edit snmp]** hierarchy level. Only the **show snmp statistics** operational mode command is supported, but you can issue SNMP requests using external SNMP client applications.
- **Advanced Insight Solutions (AIS)**—AIS provides tools and processes to automate the delivery of support services for the QFabric system. AIS components include Advanced Insight Scripts (AI-Scripts) and Advanced Insight Manager (AIM). You install AI-Scripts using the **request system scripts add** operational mode command. However, the **jais-activate-scripts.slax** file used during installation is preconfigured for the QFabric system and cannot be changed.

NOTE: Do not install Junos Space and AIS on the control plane network EX4200 switches or EX4200 Virtual Chassis in a QFX3000 QFabric system

RELATED DOCUMENTATION

[Advanced Insight Scripts \(AI-Scripts\) Release Notes](#)

[Understanding Device and Network Management Features](#)

[Overview of Junos OS System Log Messages](#)

[Understanding the Implementation of SNMP on the QFabric System | 69](#)

[SNMP MIBs Support](#)

Understanding the Implementation of SNMP on the QFabric System

SNMP monitors network devices from a central location. The QFabric system supports the basic SNMP architecture of Junos OS, but its implementation of SNMP differs from that of other devices running Junos OS. This topic provides an overview of the SNMP implementation on the QFabric system.

As in other SNMP systems, the SNMP manager resides on the network management system (NMS) of the network to which the QFabric system belongs. The SNMP agent resides in the QFabric Director software and is responsible for receiving and distributing all traps as well as responding to all the queries of the SNMP manager. For example, traps that are generated by a Node device are sent to the SNMP agent in

the Director software, which in turn processes and sends them to the target IP addresses that are defined in the SNMP configuration.

NOTE: In its SNMP implementation, the QFabric system acts as an SNMP proxy server, and requires more time to process SNMP requests than a typical Junos OS device does. The default timeout setting on most SNMP client applications is 3 seconds, which is not enough time for the QFabric system to respond to SNMP requests, so the results of your `mibwalk` command may be incomplete. For this reason, we recommend that you change the SNMP timeout setting to 5 seconds or longer for the QFabric system to complete the responses to your requests.

Support for SNMP on the QFabric system includes:

- Support for the SNMP Version 1 (v1) and v2.

NOTE: Only SNMPv2 traps are supported on the QFabric system.

- Support for the following standard MIBs:
 - RFC 1155, *Structure and Identification of Management Information for TCP/IP-based Internets*
 - RFC 1157, *A Simple Network Management Protocol (SNMP)*
 - RFC 1212, *Concise MIB Definitions*
 - RFC 1213, *Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II* (partial support, including the system group and interfaces group)
 - RFC 1215, *A Convention for Defining Traps for use with the SNMP*
 - RFC 1901, *Introduction to Community-based SNMPv2*
 - RFC 1905, *Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)*
 - RFC 1907, *Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)*
 - RFC 2011, *SNMPv2 Management Information Base for the Internet Protocol Using SMIv2*
 - RFC 2012, *SNMPv2 Management Information Base for the Transmission Control Protocol Using SMIv2*
 - RFC 2013, *SNMPv2 Management Information Base for the User Datagram Protocol Using SMIv2*
 - RFC 2233, *The Interfaces Group MIB Using SMIv2*
 - RFC 2571, *An Architecture for Describing SNMP Management Frameworks* (read-only access) (excluding SNMPv3)

- RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)* (read-only access) (excluding SNMPv3)
- RFC 2576, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework* (excluding SNMPv3)
- RFC 2578, *Structure of Management Information Version 2 (SMIv2)*
- RFC 2579, *Textual Conventions for SMIv2*
- RFC 2580, *Conformance Statements for SMIv2*
- RFC 2665, *Definitions of Managed Objects for the Ethernet-like Interface Types*
- RFC 2863, *The Interfaces Group MIB*
- RFC 3410, *Introduction and Applicability Statements for Internet Standard Management Framework* (excluding SNMPv3)
- RFC 3411, *An Architecture for Describing Simple Network Management Protocol (SNMP) Management Framework* (excluding SNMPv3)
- RFC 3412, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)* (excluding SNMPv3)
- RFC 3413, *Simple Network Management Protocol (SNMP) Applications* (excluding SNMPv3)
- RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*
- RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP)*
- RFC 3418, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*
- RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework* (excluding SNMPv3)
- RFC 4188, *Definitions of Managed Objects for Bridges*
- RFC 4293, *Management Information Base for the Internet Protocol (IP)*
- RFC 4363b, *Q-Bridge VLAN MIB*
- Support for the following Juniper Networks enterprise-specific MIBs:
 - Chassis MIB (mib-jnx-chassis.txt)
 - Class-of-Service MIB (mib-jnx-cos.txt)
 - Configuration Management MIB (mib-jnx-cfgmgmt.txt)
 - Fabric Chassis MIB (mib-jnx-fabric-chassis.txt)
 - Interface MIB Extensions (mib-jnx-if-extensions.txt)
 - Power Supply Unit MIB (mib-jnx-power-supply-unit.txt)

- QFabric MIB (mib-jnx-qf-smi.txt)
- Utility MIB (mib-jnx-util.txt)
- Support for operational mode commands—Limited to the **show snmp statistics** command. You may issue other SNMP requests, including **get**, **get next**, and **walk** requests, by using external SNMP client applications.

RELATED DOCUMENTATION

SNMP MIBs Support

SNMP Traps Support

Understanding the Implementation of System Log Messages on the QFabric System

This topic provides an overview of system log (syslog) messages as implemented on the QFabric system.

The QFabric system monitors events that occur on its component devices and distributes system log messages about those events to all external system log message servers (hosts) that are configured. Component devices may include Node devices, Interconnect devices, Director devices, and the Virtual Chassis. Messages are stored for viewing only in the QFabric system database. To view the messages, issue the **show log** command.

You configure system log messages by using the **host** and **file** statements at the **[edit system syslog]** hierarchy level. Use the **show log filename** operational mode command to view the messages.

NOTE: On the QFabric system, a syslog file named **messages** with a size of 100 MB is configured by default. If you do not configure a filename, you can use the default filename **messages** with the **show log filename** command.

All messages with a severity level of **notice** or higher are logged. Messages with a facility level of **interactive-commands** on Node devices are not logged.

The QFabric system supports the following system log message features:

- The **file filename** and **host hostname** statements at the **[edit system syslog]** hierarchy level are supported. Other statements at that hierarchy level are not supported.
- You can specify the maximum amount of data that is displayed when you issue the **show log filename** command by configuring the **file filename archive maximum-file-size** statement.
- You can specify that one or more system log message servers receive messages, which are sent to each server that is configured.
- If you configured an alias for a device or interface, the alias is displayed in the message for the device or interface.
- The level of detail that is included in a message depends on the facility and severity levels that are configured. Messages include the highest level of detail available for the configured facility and severity levels.
- The unit of time is measured and displayed in seconds, and not milliseconds. If you attempt to configure the **time-format** option in milliseconds, the log output displays **000**.

Starting in Junos OS Release 13.1, the QFabric system supports these additional syslog features:

- You can filter the output of the **show log filename** operational mode command by device type and device ID or device alias when you specify the **device-type (device-id | device-alias)** optional parameters. Device types include **director-device**, **infrastructure-device**, **interconnect-device**, and **node-device**.
- You can specify the syslog structured data output format when you configure the **structured-data** statement at the **[edit system syslog file filename]** and **[edit system syslog host hostname]** hierarchy levels.

NOTE: Information displayed in the structured data output for system logs originating from the Director software may not be complete.

- You can filter the types of logs that the Director group collects from a component device when you configure the **filter all facility severity** or **filter all match "regular-expression"** statements at the **[edit system syslog]** hierarchy level.

Unsupported syslog features include:

- File access to syslog messages
- Monitoring of syslog messages

RELATED DOCUMENTATION

[Example: Configuring System Log Messages](#) | 421

Understanding User and Access Management Features on the QFabric System

The QFabric system supports the following user and access management features:

- User authentication
- RADIUS
- Link Layer Discovery Protocol (LLDP)
- SSH
- TACACS+
- Access privilege management

The specific functionality, features, options, syntax, and hierarchy levels of some of the user and access management commands and configuration statements implemented on the QFabric system may differ somewhat from the same commands and configuration statements on standard Junos OS. See the configuration statement or command topic in the documentation set for additional information, and use the help (?) command-line function to display specific information as needed.

Some user and access management features are not yet fully supported in the full QFabric architecture, although full support is planned for future releases. The user and access management features currently unsupported on the QFabric system include:

- Full RADIUS server support, including RADIUS accounting
- **accounting-options** configuration statement hierarchy
- **tacplus-options** configuration statement

Understanding QFabric System Login Classes

In some cases (such as device-level troubleshooting), it is useful to log in to individual QFabric system components so you can view and manage issues on a per-device basis. This topic explains the login classes that provide individual component access within a QFabric system.

NOTE: Under normal operating conditions, you should manage the QFabric system as a single entity by using the QFabric system default partition command-line interface (CLI). The default partition CLI provides you with the ability to configure and monitor your entire QFabric system from a central location and should be used as the primary way to manage the system.

The QFabric system offers three special preset login classes that provide different levels of access to individual components within a QFabric system:

- **qfabric-admin**—Provides the ability to log in to individual QFabric system components and manage them. This class is equivalent to setting the following permissions: **access**, **admin**, **clear**, **firewall**, **interface**, **maintenance**, **network**, **reset**, **routing**, **secret**, **security**, **snmp**, **system**, **trace**, and **view**. The *qfabric-admin* class also enables you issue all operational mode commands except **configure**. To provide QFabric system component-level login and management privileges, include the **qfabric-admin** statement at the **[edit system login user *username* authentication remote-debug-permission]** hierarchy level.
- **qfabric-operator**—Provides the privilege to log in to individual QFabric system components and view component operations and configurations. This class is equivalent to setting the following permissions: **trace** and **view**. The *qfabric-operator* class also enables you issue the **monitor** and **show log messages** operational mode commands. To provide limited QFabric system component-level access, include the **qfabric-operator** statement at the **[edit system login user *username* authentication remote-debug-permission]** hierarchy level.
- **qfabric-user**—Prevents access to individual QFabric system components. This class is the default setting for all QFabric system users and is equivalent to the preset Junos OS class of **unauthorized**. To prevent a user from accessing individual QFabric system components, include the **qfabric-user** statement at the **[edit system login user *username* authentication remote-debug-permission]** hierarchy level.

When you perform the initial setup for the Director group, you must specify a username and password for QFabric components. Once configured, this information is stored in the QFabric system and mapped to the QFabric system login classes. Such mapping allows users with the proper login class (**qfabric-admin** or **qfabric-operator**) to log in automatically to a component without being prompted for the username and password.

After you assign the **qfabric-admin** or **qfabric-operator** class to a user, the user can log in to an individual QFabric system component by issuing the **request component login *component-name*** command. You can access Node devices, Interconnect devices, and virtual Junos Routing Engines (diagnostics, fabric control, and fabric manager) one at a time when you issue this command. To leave the CLI prompt of a component

and return to the QFabric system default partition CLI, issue the **exit** command from the component's operational mode CLI prompt.

RELATED DOCUMENTATION

| [Junos OS Login Classes Overview](#)

Understanding Interfaces on the QFabric System

IN THIS SECTION

- [Four-Level Interface Naming Convention | 76](#)
- [QSFP+ Interfaces | 76](#)
- [Link Aggregation | 80](#)

This topic describes:

Four-Level Interface Naming Convention

When you configure an interface on the QFabric system, the interface name needs to follow a four-level naming convention that enables you to identify an interface as part of either a Node device or a Node group. Include the name of the network or server Node group at the beginning of the interface name.

The four-level interface naming convention is:

device-name:type-fpc/pic/port

where *device-name* is the name of the Node device or Node group. The remainder of the naming convention elements are the same as those in the QFX3500 switch interface naming convention.

An example of a four-level interface name is:

node2:xe-0/0/2

QSFP+ Interfaces

The QFX3500 Node device provides four 40-Gbps QSFP+ (quad small form-factor pluggable plus) interfaces (labeled **Q0** through **Q3**) for uplink connections between your Node device and your Interconnect devices.

The QFX3600 Node device provides 16 40-Gbps QSFP+ interfaces. By default, 4 interfaces (labeled **Q0** through **Q3**) are configured for 40-Gbps uplink connections between your Node device and your Interconnect devices, and 12 interfaces (labeled **Q4** through **Q15**) use QSFP+ direct-attach copper (DAC) breakout cables or QSFP+ transceivers with fiber breakout cables to support 48 10-Gigabit Ethernet interfaces for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the first eight interfaces (Q0 through Q7) for uplink connections between your Node device and your Interconnect devices, and interfaces Q2 through Q15 for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks (see [“Configuring the Port Type on QFX3600 Node Devices” on page 408](#)). [Table 4 on page 16](#) shows the port mappings for QFX3600 Node devices.

Table 13: QFX3600 Node Device Port Mappings

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
Q0	Not supported on this port	xle-0/1/0	fte-0/1/0
Q1	Not supported on this port	xle-0/1/1	fte-0/1/1
Q2	xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11	xle-0/1/2	fte-0/1/2
Q3	xe-0/0/12 xe-0/0/13 xe-0/0/14 xe-0/0/15	xle-0/1/3	fte-0/1/3
Q4	xe-0/0/16 xe-0/0/17 xe-0/0/18 xe-0/0/19	xle-0/1/4	fte-0/1/4

Table 13: QFX3600 Node Device Port Mappings (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
Q5	xe-0/0/20	xle-0/1/5	fte-0/1/5
	xe-0/0/21		
	xe-0/0/22		
	xe-0/0/23		
Q6	xe-0/0/24	xle-0/1/6	fte-0/1/6
	xe-0/0/25		
	xe-0/0/26		
	xe-0/0/27		
Q7	xe-0/0/28	xle-0/1/7	fte-0/1/7
	xe-0/0/29		
	xe-0/0/30		
	xe-0/0/31		
Q8	xe-0/0/32	xle-0/1/8	Not supported on this port
	xe-0/0/33		
	xe-0/0/34		
	xe-0/0/35		
Q9	xe-0/0/36	xle-0/1/9	Not supported on this port
	xe-0/0/37		
	xe-0/0/38		
	xe-0/0/39		

Table 13: QFX3600 Node Device Port Mappings (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
Q10	xe-0/0/40	xle-0/1/10	Not supported on this port
	xe-0/0/41		
	xe-0/0/42		
	xe-0/0/43		
Q11	xe-0/0/44	xle-0/1/11	Not supported on this port
	xe-0/0/45		
	xe-0/0/46		
	xe-0/0/47		
Q12	xe-0/0/48	xle-0/1/12	Not supported on this port
	xe-0/0/49		
	xe-0/0/50		
	xe-0/0/51		
Q13	xe-0/0/52	xle-0/1/13	Not supported on this port
	xe-0/0/53		
	xe-0/0/54		
	xe-0/0/55		
Q14	xe-0/0/56	xle-0/1/14	Not supported on this port
	xe-0/0/57		
	xe-0/0/58		
	xe-0/0/59		

Table 13: QFX3600 Node Device Port Mappings (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
Q15	xe-0/0/60	xle-0/1/15	Not supported on this port
	xe-0/0/61		
	xe-0/0/62		
	xe-0/0/63		

The QFX5100-48S Node device provides 48 10-Gigabit Ethernet interfaces and 6 40-Gbps QSFP+ interfaces. By default, 4 interfaces (labeled **48** through **51**) are configured for 40-Gbps uplink connections between your Node device and your Interconnect devices, and 2 interfaces (labeled **52** and **53**) support 40-Gigabit Ethernet connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the middle two interfaces (**50** and **51**) for 40-Gigabit Ethernet connections to either endpoint systems or external networks, and you can choose to configure the last two interfaces (**52** and **53**) for uplink connections between your Node device and your Interconnect devices (see [“Configuring the QSFP+ Port Type on QFX5100 Devices” on page 414](#)).

[Table 5 on page 19](#) shows the port mappings for QFX5100-48S Node devices.

Table 14: QFX5100-48S Node Device Port Mappings

Port Number	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
48	Not supported on this PIC	fte-0/1/0
49	Not supported on this PIC	fte-0/1/1
50	xle-0/1/2	fte-0/1/2
51	xle-0/1/3	fte-0/1/3
52	xle-0/1/4	fte-0/1/4
53	xle-0/1/5	fte-0/1/5

Link Aggregation

Link aggregation enables you to create link aggregation groups across Node devices within a network Node group or redundant server Node group. You can include up to eight Ethernet interfaces in a LAG.

You can have up to 48 LAGs within a redundant server Node group, and 128 LAGs in a network Node group. To configure a LAG, include the **aggregated-devices** statement at the **[edit chassis node-group node-group-name]** hierarchy level and the **device-count** statement at the **[edit chassis node-group node-group-name aggregated-devices ethernet]** hierarchy level. Additionally, include any aggregated Ethernet options (**minimum-links** and **link-speed**) at the **[edit interfaces interface-name aggregated-ether-options]** hierarchy level and the **802.3ad** statement at the **[edit interfaces interface-name ether-options]** hierarchy level. To configure the Link Aggregation Control Protocol (LACP), include the **lacp** statement at the **[edit interfaces aggregated-ether-options]** hierarchy level.

RELATED DOCUMENTATION

[Configuring the Port Type on QFX3600 Node Devices | 408](#)

[Configuring the QSFP+ Port Type on QFX5100 Devices | 414](#)

Understanding Layer 3 Features on the QFabric System

The QFabric system supports the following Layer 3 features:

- Static routes, which enable you to manually configure and enter routes directly into the routing table.
- Routed VLAN Interfaces, which are a special type of Layer 3 virtual interface that enable you to forward packets between VLANs without using a router to connect the VLANs. Using this approach to connect VLANs reduces complexity and avoids the costs associated with purchasing, installing, managing, powering, and cooling another device.
- Routing protocols for routing traffic. The following routing protocols are supported on QFabric systems:
 - Border Gateway Protocol (BGP), which is an exterior gateway protocol (EGP) for routing traffic between autonomous systems (ASs).
 - Open Shortest Path First (OSPF) protocol, which is an interior gateway protocol (IGP) for routing traffic within an autonomous system (AS). QFabric systems support OSPFv1 and OSPFv2.

NOTE:

- When you configure routing protocols on the QFabric system, you must use interfaces from the Node devices assigned to the network Node group. If you try to configure routing protocols on interfaces from the Node devices assigned to server Node groups, the configuration commit operation fails.
- You can configure routing protocols by including statements at the **[edit protocols]** hierarchy level. If you want to isolate customer traffic on your network, you can configure virtual router routing instances at the **[edit routing-instances]** hierarchy level, and configure routing protocols for each virtual router routing instance by including statements at the **[edit routing-instances routing-instance-name protocols]** hierarchy level.

RELATED DOCUMENTATION

Understanding Virtual Router Routing Instances

Understanding Telnet on the QFabric System

This topic describes the support for the Telnet protocol on QFabric systems.

Telnet service is available for devices running Junos OS, including QFX Series devices. However, on QFabric systems, Telnet support is limited and the following conditions apply:

- You can telnet from a QFabric system to external devices that are connected to the QFabric system by way of the network Node group. To connect to these external devices, issue the **telnet** command from the QFabric default partition CLI.
- You cannot use the Telnet protocol to connect from the QFabric system default partition CLI to individual components. To access system components, you must issue the **request component login** command instead.

RELATED DOCUMENTATION

[request component login](#) | 605

telnet

Understanding Security Features on the QFabric System

The QFabric system supports the following security features:

- Firewall filters provide rules that define whether to accept or discard packets that are transiting an interface. If a packet is accepted, you can configure additional actions to perform on the packet, such as class-of-service (CoS) marking (grouping similar types of traffic together and treating each type of traffic as a class with its own level of service priority) and traffic policing (controlling the maximum rate of traffic sent or received).
- Policing (rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service. You use policers to apply limits to traffic flow and set consequences for packets that exceed these limits—usually applying a higher loss priority—so that if packets encounter downstream congestion, they can be discarded first. Policers apply only to unicast packets.
- MAC limiting protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You enable this feature on Layer 2 interfaces (ports). MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface or on all the Layer 2 access interfaces on the switch. Junos OS provides two MAC limiting methods:
 - Maximum number of MAC addresses—You configure the maximum number of dynamic MAC addresses allowed per interface. When the limit is exceeded, incoming packets with new MAC addresses can be ignored, dropped, or logged. You can also specify that the interface be shut down or temporarily disabled.
 - Allowed MAC—You configure specific “allowed” MAC addresses for the access interface. Any MAC address that is not in the list of configured addresses is not learned, and the switch logs an appropriate message. Allowed MAC binds MAC addresses to a VLAN so that the address does not get registered outside the VLAN. If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.
- Storm control causes a switch to monitor traffic levels and take a specified action when a specified traffic level—called the storm control level—is exceeded, thus preventing packets from proliferating and degrading service. You can configure switches to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when the storm control level is exceeded.

To understand more about QFX Series devices Licensing, see [Software Features That Require Licenses for QFX Series](#). Please refer to the [Juniper Licensing Guide](#) for general information about License Management. Please refer to the product [Data Sheets](#) for details, or contact your Juniper Account Team or Juniper Partner.

RELATED DOCUMENTATION

Overview of Policers

Overview of Firewall Filters

Understanding MAC Limiting and MAC Move Limiting for Port Security

Understanding Storm Control

Understanding Port Mirroring

IN THIS SECTION

- [Port Mirroring Overview | 84](#)
- [Port Mirroring Instance Types | 85](#)
- [Port-Mirroring Terminology | 85](#)
- [Port Mirroring and STP | 87](#)
- [Port Mirroring Constraints and Limitations | 88](#)

Port Mirroring Overview

Port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for local monitoring or to a VLAN for remote monitoring. Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

Port mirroring is needed for traffic analysis on a switch because a switch normally sends packets only to the port to which the destination device is connected. You configure port mirroring on the switch to send copies of unicast traffic to a local interface or a VLAN and run an analyzer application on a device connected to the interface or VLAN. You configure port mirroring by using the **analyzer** statement.

Keep performance in mind when configuring port mirroring. For example, If you mirror traffic from multiple ports, the mirrored traffic may exceed the capacity of the output interface. We recommend that you limit the amount of copied traffic by selecting specific interfaces instead of using the **all** keyword. You can also limit the amount of mirrored traffic by using a firewall filter to send specific traffic to a port mirroring instance. Mirroring only the necessary packets reduces the possibility of a performance impact.

You can use port mirroring to copy any of the following:

- All packets entering or exiting an interface (in any combination)—For example, you can send copies of the packets entering some interfaces and the packets exiting other interfaces to the same local interface or VLAN. If you configure port mirroring to copy packets exiting an interface, traffic that originates on

that switch or Node device (in a QFabric system) is not copied when it egresses. Only switched traffic is copied on egress. (See the limitation on egress mirroring below.)

- All packets entering a VLAN—You cannot use port mirroring to copy packets exiting a VLAN.
- Firewall-filtered sample—Sample of packets entering a port or VLAN. Configure a firewall filter to select certain packets for mirroring.

NOTE: Firewall filters are not supported on egress ports; therefore, you cannot specify policy-based sampling of packets exiting an interface.

Port Mirroring Instance Types

To configure port mirroring, you configure an instance of one of the following types:

- Analyzer instance: You must specify the input and output for the instance. This instance type is useful for ensuring that all traffic transiting an interface or VLAN is mirrored and sent to the analyzer device.
- Port-mirroring instance: You do not specify an input for this instance type. Instead, you, create a firewall filter that specifies the required traffic and directs it to the mirror. This instance type is useful for controlling which types of traffic should be mirrored. When you use a port-mirroring instance, you can direct traffic to it in the following ways:
 - Specify the name of the port-mirroring instance in the firewall filter using the **port-mirror-instance instance-name** action. You should use this approach if there are multiple port-mirroring instances defined.
 - Configure the filter to send the mirrored packets to the output interface defined in the instance using the **port-mirror** action. You can use this approach if there is only one port-mirroring instance defined.

Port-Mirroring Terminology

[Table 15 on page 85](#) lists the terms used in the documentation about port mirroring and provides definitions.

Table 15: Port Mirroring Terms and Definitions

Term	Description
Analyzer instance	Port-mirroring configuration that includes a name, source interfaces or source VLAN, and a destination for mirrored packets (either a local access interface or a VLAN).

Table 15: Port Mirroring Terms and Definitions (*continued*)

<p>Port mirroring instance</p> <p>NOTE: Port mirroring instance feature is not supported on NFX150 devices.</p>	<p>A port-mirroring configuration that does not specify an input.. A firewall filter must be used to send traffic to the port mirror. Use the action port-mirror-instance <i>instance-name</i> in the firewall filter configuration to send packets to the port mirror.</p>
<p>Output interface (also known as monitor interface)</p>	<p>Access interface to which packet copies are sent and to which a device running an analyzer application is connected.</p> <p>The following limitations apply to an output interface:</p> <ul style="list-style-type: none"> • Cannot also be a source port. • Cannot be used for switching. • Cannot be an aggregated Ethernet interface (LAG). • Does not participate in Layer 2 protocols, such as Spanning Tree Protocol (STP). • Loses any existing VLAN associations when you configure it as an analyzer output interface. <p>If the capacity of the output interface is insufficient to handle the traffic from the source ports, overflow packets are dropped.</p>
<p>Output IP address</p>	<p>IP address of the device running an analyzer application. The device can be on a remote network. When you use this feature, the mirrored packets are GRE-encapsulated. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)</p> <ul style="list-style-type: none"> • An output IP address cannot be in the same subnetwork as any of the switch's management interfaces. • If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).

Table 15: Port Mirroring Terms and Definitions (*continued*)

Output VLAN (also known as monitor or analyzer VLAN)	<p>VLAN to which copies are sent and to which a device running an analyzer application is connected. The analyzer VLAN can span multiple switches.</p> <p>The following limitations apply to an output VLAN:</p> <ul style="list-style-type: none"> • Cannot be a private VLAN or VLAN range. • Cannot be shared by multiple analyzer statements. • An output VLAN interface cannot be a member of any other VLAN. • An output VLAN interface cannot be an aggregated Ethernet interface (LAG). • On some switches, only one interface can be a member of the analyzer VLAN. This limitation does not apply on the QFX10000 switch if traffic is mirrored on ingress. In this case, multiple QFX10000 interfaces can belong to the output VLAN, and traffic is mirrored to all of those interfaces. If traffic is mirrored on egress on a QFX10000 switch, only one interface can be a member of the analyzer VLAN.
Input interface (also known as mirrored or monitored interface)	Interface that provides traffic to be mirrored. This traffic can be entering or exiting the interface. (Ingress or egress traffic can be mirrored.) An input interface cannot also be an output interface for an analyzer.
Monitoring station	Computer running an analyzer application.
Local port mirroring	Port-mirroring configuration in which the mirrored packets are sent to an interface on the same switch.
Remote port mirroring	Flooding mirrored packets to an output (analyzer) VLAN that you create to receive mirror traffic or sending the mirrored packets to a remote IP address. (You cannot send mirrored packets to a remote IP address on a QFabric system.)
Policy-based mirroring	Mirroring of packets that match the match a firewall filter term. The action analyzer analyzer-name is used in the firewall filter to send the packets to the analyzer.

Port Mirroring and STP

The behavior of STP in a port-mirroring configuration depends on the version of Junos OS you are using:

- Junos OS 13.2X50, Junos OS 13.2X51-D25 or earlier, Junos OS 13.2X52: If you enable STP, port mirroring might not work because STP might block the mirrored packets.
- Junos OS 13.2X51-D30, Junos OS 14.1X53: STP is disabled for mirrored traffic. You must ensure that your topology prevents loops for this traffic.

Port Mirroring Constraints and Limitations

IN THIS SECTION

- [Local and Remote Port Mirroring | 88](#)
- [Remote Port Mirroring Only | 90](#)
- [Port Mirroring Constraints on OCX Series Switches | 90](#)

Local and Remote Port Mirroring

The following constraints and limitations apply to local and remote port mirroring:

- You can create a total of four port-mirroring configurations.
- You can create a total of four port-mirroring configurations on each Node group in a QFabric system, subject to the following constraints:
 - As many as four of the configurations can be for local port mirroring.
 - As many as three of the configurations can be for remote port mirroring.
- Regardless of whether you are configuring a standalone switch or a Node group, the following limits apply:
 - There can be no more than two configurations that mirror ingress traffic. (If you configure a firewall filter to send traffic to a port mirror—that is, you use the **analyzer** action modifier in a filter term—this counts as an ingress mirroring configuration for switch or Node group on which the filter is applied.)
 - There can be no more than two configurations that mirror egress traffic.

NOTE: On QFabric systems, there is no system-wide limit on the total number of mirror sessions.

- You can configure no more than one type of output in one port-mirroring configuration. That is, you can use no more than one of the following to complete a **set analyzer name output** statement:
 - **interface**
 - **ip-address**
 - **vlan**
- If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs on a standalone switch or QFabric system. If you do so, some VLAN packets might contain incorrect VLAN IDs. This applies to any VLAN packets—not only the mirrored copies.

- The **ratio** and **loss-priority** options are not supported.
- Packets with physical layer errors are filtered out and are not sent to the output port or VLAN.
- If you use sFlow monitoring to sample traffic, it does not sample the mirror copies when they exit from the output interface.
- You cannot mirror packets exiting or entering the following ports:
 - Dedicated Virtual Chassis interfaces
 - Management interfaces (me0 or vme0)
 - Fibre Channel interfaces
 - Integrated routing and bridging (IRB) interfaces (also known as routed VLAN interfaces, or RVIs)
- An aggregated Ethernet interface cannot be an output interface if the input is a VLAN or if traffic is sent to the analyzer by a firewall filter.
- When packet copies are sent out the output interface, they are not modified for any changes that are normally applied on egress, such as CoS rewriting.
- An interface can be the input interface for only one mirroring configuration. Do not use the same interface as the input interface for multiple mirroring configurations.
- CPU-generated packets (such as ARP, ICMP, BPDU, and LACP packets) cannot be mirrored on egress.
- VLAN-based mirroring is not supported for STP traffic.
- (QFabric systems only) If you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on different Node devices, the mirrored copies have incorrect VLAN IDs. This limitation does not apply if you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on the *same* Node device. In this case the mirrored copies have the correct VLAN IDs (as long as you do not configure more than 2000 VLANs on the QFabric system).
- True egress mirroring is defined as mirroring the exact number of copies and the exact packet modifications that went out the egress switched port. Because the processor on QFX5xxx (including QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210) and EX4600 (including EX4600 and EX4650) switches implements egress mirroring in the ingress pipeline, those switches do not provide accurate egress packet modifications, so egress mirrored traffic can carry incorrect VLAN tags that differ from the tags in the original traffic.
- If you configure a port-mirroring instance to mirror traffic exiting from an interface that performs VXLAN encapsulation, the source and destination MAC addresses of the mirrored packets will not be the same as those of the original traffic.
- Mirroring on member interfaces of a LAG is not supported.
- Egress VLAN mirroring is not supported.

Remote Port Mirroring Only

The following constraints and limitations apply to remote port mirroring:

- If you configure an output IP address, the address cannot be in the same subnetwork as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).
- An output VLAN cannot be a private VLAN or VLAN range.
- An output VLAN cannot be shared by multiple **analyzer** statements.
- An output VLAN interface cannot be a member of any other VLAN.
- An output VLAN interface cannot be an aggregated Ethernet interface.
- If the output VLAN has more than one member interface, then traffic is mirrored only to the first member of the VLAN, and other members of the same VLAN do not carry any mirrored traffic.
- If you attempt to configure more than one analyzer session for remote port mirroring to an IP address (GRE encapsulation) and the IP addresses of the analyzers are reachable through the same interface, then only one analyzer session is configured.

Port Mirroring Constraints on OCX Series Switches

The following constraints and limitations apply to port mirroring on OCX Series switches:

- You can create a total of four port-mirroring configurations. The following constraints also apply:
 - There can be no more than two configurations that mirror ingress traffic.
 - There can be no more than two configurations that mirror egress traffic.
- If you use sFlow monitoring to sample traffic, it does not sample the mirror copies when they exit from the output interface.
- You can create only one port-mirroring session.
- You cannot mirror packets exiting or entering the following ports:
 - Dedicated Virtual Chassis interfaces
 - Management interfaces (me0 or vme0)
 - Fibre Channel interfaces
 - Routed VLAN interfaces or IRB interfaces
- An aggregated Ethernet interface cannot be an output interface.
- Do not include an 802.1Q subinterface that has a unit number other than 0 in a port mirroring configuration. Port mirroring does not work with subinterfaces if their unit number is not 0. (You configure 802.1Q subinterfaces using the **vlan-tagging** statement.)

- When packet copies are sent out the output interface, they are not modified for any changes that are normally applied on egress, such as CoS rewriting.
- An interface can be the input interface for only one mirroring configuration. Do not use the same interface as the input interface for multiple mirroring configurations.
- CPU-generated packets (such as ARP, ICMP, BPDU, and LACP packets) cannot be mirrored on egress.
- VLAN-based mirroring is not supported for STP traffic.

SEE ALSO


Understanding Port Mirroring
Example: Mirroring Employee Web Traffic with a Firewall Filter
Configuring Port Mirroring

RELATED DOCUMENTATION

Configuring Port Mirroring
Examples: Configuring Port Mirroring for Local Analysis
Examples: Configuring Port Mirroring for Local Analysis
Example: Configuring Port Mirroring for Remote Analysis
Troubleshooting Port Mirroring

Understanding Fibre Channel Fabrics on the QFabric System

A Fibre Channel (FC) fabric on a QFabric system is a construct that you configure on a QFX3500 Node device when the Node device is in FCoE-FC gateway mode. The FC fabric on a QFabric Node device is not the same as an FC fabric on a storage area network (SAN). The FC fabric on a QFabric Node device is local to that particular node device. We call the FC fabric on a QFabric Node device a *local FC fabric* to differentiate it from an FC fabric on the SAN.

**NOTE:** The QFX3600 Node device does not support FC or FCoE features.

A local FC fabric does not span Node devices and does not span the fabric Interconnect device. Local FC fabrics are entirely contained on a single Node device. A local FC fabric creates associations that connect

FCoE devices that have converged network adapters (CNAs) on the Ethernet network to an FC switch or FCoE forwarder (FCF) on the FC network. A local FC fabric consists of:

- A unique fabric name.
- A unique fabric ID.
- One or more FCoE VLAN interfaces that include one or more 10-Gigabit Ethernet interfaces connected to FCoE devices. The FCoE VLANs transport traffic between the FCoE servers and the FCoE-FC gateway. Each FCoE VLAN must carry only FCoE traffic. You cannot mix FCoE traffic and standard Ethernet traffic on the same VLAN.

The 10-Gigabit Ethernet interfaces that connect to FCoE devices must include a native VLAN to transport FIP traffic because FIP VLAN discovery and notification frames are exchanged as untagged packets.

Each FCoE VLAN interface can present multiple VF_Port interfaces to the FCoE network.

- One or more native FC interfaces. The native FC interfaces transport traffic between the gateway and the FC switch or FCF.

TIP: If the network does not use a dual-rail architecture for redundancy, configure more than one native FC interface for each local FC fabric to create redundant connections between the FCoE devices and the FC network. If one physical link goes down, any sessions it carried can log in again and connect to the FC network on a different interface.

All of the FC and FCoE traffic that belongs to a local FC fabric on a Node device must enter and exit that Node device. This means that the FC switch or FCF and the FCoE devices in the Ethernet network must be connected to the same Node device. The interfaces that connect to the FC switch and the interfaces that connect to the FCoE devices must be included in the local FC fabric. You cannot configure a local FC fabric that spans more than one Node device.

Traffic flows from FC and FCoE devices that are not in the same local FC fabric remain separate and cannot communicate with each other through the FCoE-FC gateway.

NOTE: The QFabric system enforces commit checks to ensure that local FC fabrics and FCoE VLANs on FCoE-FC gateways do not span more than one Node device.

RELATED DOCUMENTATION

Overview of Fibre Channel

Understanding an FCoE-FC Gateway

Understanding FCoE-FC Gateway Functions

Understanding Interfaces on an FCoE-FC Gateway

Understanding CoS Fabric Forwarding Class Sets

IN THIS SECTION

- [Default Fabric Forwarding Class Sets | 95](#)
- [Fabric Forwarding Class Set Configuration and Implementation | 98](#)
- [QFabric System CoS | 100](#)
- [Support for Flow Control and Lossless Transport Across the Fabric | 100](#)
- [Viewing Fabric Forwarding Class Set Information | 103](#)
- [Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences | 105](#)

Fabric forwarding class sets (fabric fc-sets) are similar to the fc-sets (priority groups) you configure on Node devices. The major differences are:

1. Fabric fc-sets group traffic for transport across the QFX3008-I or QFX3600-I Interconnect device (the fabric). Node device fc-sets group traffic on a Node device for transport across that Node device.
2. Fabric fc-sets are global. They apply to the entire fabric. Node device fc-sets apply only to the Node device on which they are configured.
3. Fabric fc-sets are mapped directly to Interconnect device output queues; in this way, they behave similarly to forwarding classes on a Node device.

Fabric fc-sets map to Interconnect device fabric output queues statically—you cannot configure the mapping of fabric fc-sets to fabric output queues. All traffic in a fabric fc-set maps to the same output queue.

Node device fc-sets include forwarding classes that map to Node device output queues, and you can configure the mapping of forwarding classes to output queues (or you can use the default mapping). Because output queues are mapped to forwarding classes, different classes of traffic in a Node device fc-set can be mapped to different output queues.

Node device fc-sets consist of forwarding classes containing traffic that requires similar CoS treatment. (Forwarding classes are default forwarding classes or user-defined forwarding classes.) You can configure CoS for each fc-set to determine how the traffic of its forwarding classes is scheduled on a Node device.

When traffic exits a Node device interface and enters an Interconnect device fabric interface, the Interconnect device uses the same forwarding classes to group traffic. The forwarding classes are mapped to global fabric fc-sets for transport across the fabric. Like fc-sets on a Node device, fabric fc-sets also contain traffic that requires similar CoS treatment. Also like fc-sets on a Node device, you can configure CoS on fabric fc-sets.

Fabric fc-sets reside on the Interconnect device and are global to the QFabric system. Fabric fc-sets apply to all traffic that traverses the fabric. The mapping of forwarding classes to fabric fc-sets is global and applies to all forwarding classes with traffic that traverses the fabric from all connected Node devices. You can change the mapping of forwarding classes to fabric fc-sets. All mapping changes you make are global. For example, if you change the fabric fc-set to forwarding class mapping of the default best-effort forwarding class, then every Node device's best-effort forwarding class traffic that traverses the fabric is mapped to that fabric fc-set. The CoS you configure on a fabric fc-set applies to all the traffic that belongs to that fabric fc-set, from all connected Node devices.

This topic describes:

Default Fabric Forwarding Class Sets

Interconnect devices have 12 default fabric fc-sets, including five visible default fabric fc-sets, four for unicast traffic and one for multideestination (multicast, broadcast, and destination lookup failure) traffic.

There are also seven hidden default fabric fc-sets. There are three hidden default fabric fc-sets for multideestination traffic that you can use if you want to map different multideestination forwarding classes to different multideestination fabric fc-sets. There are four hidden default fabric fc-sets for lossless traffic that you can use to map different lossless forwarding classes (priorities) to different lossless fabric fc-sets.

[Table 16 on page 95](#) shows the default fabric fc-sets:

Table 16: Default Fabric Forwarding Class Sets

Fabric Forwarding Class Set Name	Characteristics
fabric_fcset_be	Transports best-effort unicast traffic across the fabric.
fabric_fcset_strict_high	Transports unicast traffic that has been configured with strict-high priority and in the network-control forwarding class across the fabric. This fabric fc-set receives as much bandwidth across the fabric as it needs to service the traffic in the group up to the entire fabric interface bandwidth. For this reason, exercise caution when mapping traffic to this fabric fc-set to avoid starving other traffic.
fabric_fcset_noloss1	Transports unicast traffic in the default fcoe forwarding class across the fabric.
fabric_fcset_noloss2	Transports unicast traffic in the default no-loss forwarding class across the fabric.
fabric_fcset_noloss3	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.
fabric_fcset_noloss4	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.
fabric_fcset_noloss5	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.
fabric_fcset_noloss6	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.

Table 16: Default Fabric Forwarding Class Sets (*continued*)

Fabric Forwarding Class Set Name	Characteristics
fabric_fcset_multicast1	Transports multdestination traffic in the mcast forwarding class across the fabric. This fabric fc-set is valid only for multdestination forwarding classes.
fabric_fcset_multicast2	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes.
fabric_fcset_multicast3	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes.
fabric_fcset_multicast4	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes.

The five default forwarding classes (**best-effort**, **fcoe**, **no-loss**, **network-control**, and **mcast**) are mapped to the fabric fc-sets by default as shown in [Table 17 on page 96](#).

Table 17: Default Forwarding Class to Fabric Forwarding Class Set Mapping

Forwarding Class	Fabric Forwarding Class Set	Fabric Output Queue	Maximum MTU Supported for Lossless Operation
best-effort	fabric_fcset_be	0	NA
network-control	fabric_fcset_strict_high	7	NA
fcoe	fabric_fcset_noloss1	1	9K
no-loss	fabric_fcset_noloss2	2	9K
mcast	fabric_fcset_multicast1	8	NA
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss3	3	9k

Table 17: Default Forwarding Class to Fabric Forwarding Class Set Mapping (*continued*)

Forwarding Class	Fabric Forwarding Class Set	Fabric Output Queue	Maximum MTU Supported for Lossless Operation
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss4	4	9k
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss5	5	9k
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss6	6	9k
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_multicast2	9	NA
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_multicast3	10	NA
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_multicast4	11	NA

The maximum fiber cable length between the QFabric system Node device and the QFabric system Interconnect device is 150 meters.

TIP: If you explicitly configure lossless forwarding classes, we recommend that you map each user-configured lossless forwarding class to an unused fabric fc-set (fabric_fcset_noloss3 through fabric_fcset_noloss6) on a one-to-one basis: one lossless forwarding class mapped to one lossless fabric fc-set.

The reason for one-to-one mapping is to avoid fate sharing of lossless flows. Because each fabric fc-set is mapped statically to an output queue, when you map more than one forwarding class to a fabric fc-set, all of the traffic in all of the forwarding classes that belong to the fabric fc-set uses the same output queue. If that output queue becomes congested due to congestion caused by one of the flows, the other flows are also affected. (They share fate because the flow that congests the output queue affects flows that are not experiencing congestion.)

If you want to map different multdestination forwarding classes to different multdestination fabric fc-sets, use one or more of the hidden multdestination fabric fc-sets.

NOTE: The global mapping of forwarding classes to fabric fc-sets is independent of the mapping of forwarding classes to Node device fc-sets. Global mapping of forwarding classes to fabric fc-sets occurs only on the Interconnect device. The Node device mapping of forwarding classes to fc-sets does not affect the global mapping of forwarding classes to fabric fc-sets on the Interconnect device, and vice versa.

When you define new forwarding classes on a Node device, you explicitly map those forwarding classes to Node device fc-sets. However, new (user-created) forwarding classes are mapped by default to fabric fc-sets. (You can override the default mapping if you want to configure the forwarding class to fabric fc-set mapping explicitly, as described in the next section.)

By default:

- All best-effort traffic forwarding classes that you create are mapped to the **fabric_fcset_be** fabric fc-set.
- All lossless traffic forwarding classes that you create are mapped to the **fabric_fcset_noloss1** or **fabric_fcset_noloss2** fabric fc-set.

NOTE: To avoid fate sharing, we recommend that you configure one-to-one mapping of user-configured lossless forwarding classes to lossless fabric fc-sets instead of using the default mapping. You can also use firewall filters to mitigate fate sharing by separating flows that belong to the same forwarding class as the traffic traverses the Interconnect device (see *Understanding How to Mitigate Fate Sharing on a QFabric System Interconnect Device by Remapping Traffic Flows (Forwarding Classes)* for more information.)

- All multidestination traffic forwarding classes that you create are mapped to the **fabric_fcset_multicast1** fabric fc-set.
- All **strict-high** priority traffic and **network-control** forwarding classes that you create are mapped to the **fabric_fcset_strict_high** fabric fc-set.

Fabric Forwarding Class Set Configuration and Implementation

IN THIS SECTION

- Mapping Forwarding Classes to Fabric Forwarding Class Sets | 99
- Fabric Forwarding Class Set Implementation | 99

You can map forwarding classes to fabric fc-sets and configure CoS scheduling for fabric fc-sets. This section describes:

Mapping Forwarding Classes to Fabric Forwarding Class Sets

If you do not want to use the default mapping of forwarding classes to fabric fc-sets, you can map forwarding classes to fabric fc-sets in the same way as you map forwarding classes to Node device fc-sets. To do this, use exactly the same statement that you use to map forwarding classes to fc-sets, but instead of specifying a Node device fc-set name, specify a fabric fc-set name.

NOTE: The global mapping of forwarding classes to fabric fc-sets does not affect the mapping of forwarding classes to Node device fc-sets. The global forwarding class mapping to fabric fc-sets pertains to the traffic only when it enters, traverses, and exits the fabric. The forwarding class mapping to fc-sets on a Node device is valid within that Node device.

Mapping forwarding classes to fabric fc-sets does not affect the scheduling configuration of the forwarding classes or fc-sets on Node devices. Fabric fc-set scheduling pertains to traffic only when it enters, traverses, and exits the Interconnect device fabric.

If you change the mapping of a forwarding class to a fabric fc-set, the new mapping is global and applies to all traffic in that forwarding class, regardless of which Node device forwards the traffic to the Interconnect device.

- To assign one or more forwarding classes to a fabric fc-set:

```
[edit class-of-service]
user@switch# set forwarding-class-sets fabric-forwarding-class-set-name class forwarding-class-name
```

For example, to map a user-defined forwarding class named **best-effort-2** to the fabric fc-set **fabric_fcset_be**:

```
[edit class-of-service]
user@switch# set forwarding-class-sets fabric_fcset_be class best-effort-2
```

NOTE: Because fabric fc-set configuration is global, in this example all forwarding classes with the name **best-effort-2** on all of the Node devices connected to the fabric use the **fabric_fcset_be** fabric fc-set to transport traffic across the fabric.

Fabric Forwarding Class Set Implementation

The following rules apply to fabric fc-sets:

- You cannot create new fabric fc-sets. Only the twelve default fabric fc-sets are available.
- You cannot delete a default fabric fc-set.
- You cannot attach a fabric fc-set to a Node device interface. Fabric fc-sets are used only on the Interconnect device fabric, not on Node devices.
- You can map only multdestination forwarding classes to multdestination fabric fc-sets.
- You cannot map multdestination forwarding classes to unicast fabric fc-sets.
- You cannot map unicast forwarding classes to multdestination fabric fc-sets.

QFabric System CoS

When traffic enters and exits the same QFabric system Node device, CoS works the same as it works on a standalone switch.

However, when traffic enters a QFabric system Node device, crosses the Interconnect device, and then exits a different Node device, CoS is applied differently:

1. Traffic entering the ingress Node device receives the CoS configured at the Node ingress (packet classification and congestion notification profile for PFC).
2. When traffic goes from the ingress Node device to the Interconnect device, the fabric fc-set CoS is applied to the traffic.
3. When traffic goes from the Interconnect device to the egress Node device, the egress Node device applies CoS at the egress port (egress queue scheduling, WRED, and IEEE 802.1p or DSCP code-point rewrite).

Traffic that traverses the Interconnect device can use the default CoS fabric scheduling or you can configure two-tier hierarchical CoS scheduling explicitly on fabric fc-sets as described in *Understanding CoS Scheduling Across the QFabric System*.

Support for Flow Control and Lossless Transport Across the Fabric

The Interconnect device incorporates flow control mechanisms to support lossless transport during periods of congestion on the fabric. To support the priority-based flow control (PFC) feature on the Node devices, the fabric interfaces use LLFC to support lossless transport for up to six IEEE 802.1p priorities when the following two configuration constraints are met:

1. The IEEE 802.1p priority used for the traffic that requires lossless transport is mapped to a lossless forwarding class on the Node devices.

2. The lossless forwarding class must be mapped to a lossless fabric fc-set on the Interconnect device (**fabric_fcset_noloss1**, **fabric_fcset_noloss2**, **fabric_fcset_noloss3**, **fabric_fcset_noloss4**, **fabric_fcset_noloss5**, or **fabric_fcset_noloss6**).

When traffic meets the two configuration constraints, the fabric propagates the back pressure from the egress Node device across the fabric to the ingress Node device during periods of congestion. However, to achieve end-to-end lossless transport across the switch, you must also configure a congestion notification profile to enable PFC on the Node device ingress ports.

For all other combinations of IEEE 802.1p priority to forwarding class mapping and all other combinations of forwarding class to fabric fc-set mapping, the congestion control mechanism is normal packet drop. For example:

- **Case 1**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric_fcset_noloss1** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 2**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 3**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric_fcset_noloss2** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 4**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 5**—If the IEEE 802.1p priority 5 is mapped to the **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 6**—If the IEEE 802.1p priority 5 is mapped to the **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric_fcset_noloss1** fabric fc-set, then the congestion control mechanism is packet drop.

NOTE: Lossless transport across the fabric also must meet the following two conditions:

1. The maximum cable length between the Node device and the Interconnect device is a 150 meters of fiber cable.
2. The maximum frame size is 9216 bytes.

If the MTU is 9216 KB, in some cases the QFabric system supports only five lossless forwarding classes instead of six lossless forwarding classes because of headroom buffer limitations.

The number of IEEE 802.1p priorities (forwarding classes) the QFabric system can support for lossless transport across the Interconnect device fabric depends on several factors:

- Approximate fiber cable length—The longer the fiber cable that connects Node device fabric (FTE) ports to the Interconnect device fabric ports, the more data the connected ports need to buffer when a pause is asserted. (The longer the fiber cable, the more frames are traversing the cable when a pause is asserted. Each port must be able to store all of the “in transit” frames in the buffer to preserve lossless behavior and avoid dropping frames.)
- MTU size—The larger the maximum frame sizes the buffer must hold, the fewer frames the buffer can hold. The larger the MTU size, the more buffer space each frame consumes.
- Total number of Node device fabric ports connected to the Interconnect device—The higher the number of connected fabric ports, the more headroom buffer space the Node device needs on those fabric ports to support the lossless flows that traverse the Interconnect device. Because more buffer space is used on the Node device fabric ports, less buffer space is available for the Node device access ports, and a lower total number of lossless flows are supported.

The QFabric system supports six lossless priorities (forwarding classes) under most conditions. The priority group headroom that remains after allocating headroom to lossless flows is sufficient to support best-effort and multidestination traffic.

[Table 18 on page 103](#) shows how many lossless priorities the QFabric system supports under different conditions (fiber cable lengths and MTUs) in cases when the QFabric system supports fewer than six lossless priorities. The number of lossless priorities is the same regardless of how many Node device FTE ports are connected to the Interconnect device. However, the higher the number of FTE ports connected to the Interconnect device, the lower the number of total lossless flows supported. In all cases that are not shown in [Table 18 on page 103](#), the QFabric system supports six lossless priorities.

NOTE: The system does not perform a configuration commit check that compares available system resources with the number of lossless forwarding classes configured. If you commit a configuration with more lossless forwarding classes than the system resources can support, frames in lossless forwarding classes might be dropped.

Table 18: Lossless Priority (Forwarding Class) Support for Node Devices When Fewer than Six Lossless Priorities Are Supported

MTU in Bytes	Fiber Cable Length in Meters (Approximate)	Maximum Number of Lossless Priorities (Forwarding Classes) on the Node Device
9216 (9K)	100	5
9216 (9K)	150	5

NOTE: The total number of lossless flows decreases as resource consumption increases. For a Node device, the higher the number of FTE ports connected to the Interconnect device, the larger the MTU, and the longer the fiber cable length, the fewer total lossless flows the QFabric system can support.

Viewing Fabric Forwarding Class Set Information

You can display information about fabric fc-sets using the same CLI command you use to display information about Node device fc-sets:

```
user@switch> show class-of-service forwarding-class-set
```

```
Forwarding class set: fabric_fcset_be, Type: fabric-type, Forwarding class set
index: 1
  Forwarding class      Index
  best-effort           0

Forwarding class set: fabric_fcset_mcast1, Type: fabric-type, Forwarding class set
index: 5
  Forwarding class      Index
  mcast                 8

Forwarding class set: fabric_fcset_mcast2, Type: fabric-type, Forwarding class set
```



```

index: 6

Forwarding class set: fabric_fcset_mcast3, Type: fabric-type, Forwarding class set
index: 7

Forwarding class set: fabric_fcset_mcast4, Type: fabric-type, Forwarding class set
index: 8

Forwarding class set: fabric_fcset_noloss1, Type: fabric-type, Forwarding class
set index: 2
  Forwarding class          Index
  fcoe                     1

Forwarding class set: fabric_fcset_noloss2, Type: fabric-type, Forwarding class
set index: 3
  Forwarding class          Index
  no-loss                   2

Forwarding class set: fabric_fcset_noloss3, Type: fabric-type, Forwarding class
set index: 9

Forwarding class set: fabric_fcset_noloss4, Type: fabric-type, Forwarding class
set index: 10

Forwarding class set: fabric_fcset_noloss5, Type: fabric-type, Forwarding class
set index: 11

Forwarding class set: fabric_fcset_noloss6, Type: fabric-type, Forwarding class
set index: 12

Forwarding class set: fabric_fcset_strict_high, Type: fabric-type, Forwarding class
set index: 4
  Forwarding class          Index
  network-control          3

```

[Table 19 on page 104](#) describes the meaning of the **show class-of-service forwarding-class-set** output fields when you display fabric fc-set information.

Table 19: show class-of-service forwarding-class-set Command Output Fields

Field Name	Field Description
Forwarding class set	Name of the fabric forwarding class set.

Table 19: show class-of-service forwarding-class-set Command Output Fields (*continued*)

Field Name	Field Description
Type	Type of forwarding class set: <ul style="list-style-type: none"> • Fabric-type—Fabric fc-set • Normal-type—Node device fc-set
Forwarding class set index	Index of this forwarding class set.
Forwarding class	Name of a forwarding class.
Index	Index of the forwarding class.

Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences

Table 20 on page 105 summarizes the differences between fabric fc-sets and fc-sets:

Table 20: Summary of Differences Between Fabric fc-sets and Local fc-sets

Characteristic	Fabric fc-set	Local fc-set
Location	QFX3008-I or QFX3600-I Interconnect device (the fabric).	QFabric Node device.
Global or local	Global, valid for the entire fabric.	Local to the Node device on which the fc-set is configured.
Ability to create (define) a new fc-set	No. Use the 12 default fabric fc-sets provided.	Yes.
Ability to configure CoS	User-configurable using fabric fc-set scheduler maps.	User-configurable using traffic control profiles.
Ability to map forwarding classes to an fc-set	Yes. Mapping is global and applies to all forwarding classes across the Interconnect device fabric (traffic from all connected Node devices).	Yes. Mapping is local to a Node device and applies only to the forwarding classes on the Node device.

RELATED DOCUMENTATION

Understanding CoS Forwarding Class Sets (Priority Groups)

Understanding CoS Scheduling Across the QFabric System

Understanding CoS Scheduling on QFabric System Node Device Fabric (fte) Ports

Understanding Default CoS Scheduling on QFabric System Interconnect Devices (Junos OS Release 13.1 and Later Releases)

Understanding How to Mitigate Fate Sharing on a QFabric System Interconnect Device by Remapping Traffic Flows (Forwarding Classes)

Defining CoS Forwarding Class Sets

Example: Configuring Forwarding Class Sets

Example: Configuring CoS Scheduling Across the QFabric System

show class-of-service forwarding-class-set

2

PART

Installation

Before You Begin | **109**

Ports and Connectors | **137**

Power | **147**

Installing a QFX3100 Director Device | **165**

Installing a QFX3600-I Interconnect Device or QFX3600 Node Device | **177**

Installing a QFX3500 Node Device | **199**

Installing a QFX5100 Node Device | **219**

Cabling a Copper-Based Control Plane for the QFX3000-M QFabric System | **239**

Cabling a Fiber-Based Control Plane for the QFX3000-M QFabric System | **267**

Cabling the Data Plane for the QFX3000-M QFabric System | **301**

Before You Begin

IN THIS CHAPTER

- QFX3000-M QFabric System Installation Overview | 109
- Understanding QFX3000-M QFabric System Hardware Configurations | 112
- Planning a QFX3000-M QFabric System Deployment | 118
- General Site Guidelines | 124
- Site Electrical Wiring Guidelines | 124
- Environmental Requirements and Specifications for a QFX3100 Director Device | 125
- Environmental Requirements and Specifications for QFX3600 and QFX3600-I Devices | 126
- Environmental Requirements and Specifications for a QFX3500 Device | 128
- Environmental Requirements and Specifications for a QFX5100 Device | 129
- Environmental Requirements and Specifications for EX Series Switches | 130

QFX3000-M QFabric System Installation Overview

A QFX3000-M QFabric system is formed by interconnecting QFX3500 or QFX3600 Node devices, QFX3600-I Interconnect devices, and QFX3100 Director devices. Two EX4200 or EX4300 switches are used to interconnect the control plane and management network. For more information about the role of each device in the QFX3000-M QFabric system, see [“Understanding QFX3000-M QFabric System Hardware Configurations” on page 112](#).

Before you begin to install the QFX3000-M QFabric system:

- Read *General Safety Guidelines and Warnings*.
- Review [“Planning a QFX3000-M QFabric System Deployment” on page 118](#) and the topics it references. Do not start the installation until you have completed the site preparation checklists for each device type:
 - *Site Preparation Checklist for a QFX3100 Director Device*
 - *Site Preparation Checklist for a QFX5100 Device*
 - *Site Preparation Checklist for a QFX3600 or QFX3600-I Device*

- *Site Preparation Checklist for a QFX3500 Device*
- *Site Preparation Checklist for EX4200 Switches* or *Site Preparation Checklist for EX4300 Switches*.

NOTE: Do not mix control plane EX4200 and EX4300 switches in the same QFabric system.

To install a QFX3000-M QFabric system:

1. Install all the devices in their permanent location, connect the devices to earth ground, and connect power to the devices. See:
 - [Installing and Connecting a QFX3100 Director Device on page 165](#)
 - [Installing and Connecting a QFX5100 Device on page 219](#)
 - [Installing and Connecting a QFX3600 or QFX3600-I Device on page 177](#)
 - [Installing and Connecting a QFX3500 Device on page 199](#)
 - [Installing and Connecting an EX4200 Switch](#) or [Installing and Connecting an EX4300 Switch](#)
2. Ensure that each Node device is set to Node device mode. By default, the devices work as standalone switches. You perform this step using the console (**CON**) port on each Node device. Leave the Node devices powered on. See [“Converting the Device Mode for a QFabric System Component” on page 322](#).
3. Interconnect the two EX Series switches using the SFP+ uplink module ports. The EX4200 uses 1-Gigabit Ethernet ports; the EX4300 uses 10-Gigabit Ethernet ports. These ports will later be configured in a LAG. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#).
4. Interconnect the two QFX3100 Director devices for control plane redundancy. See [“Connecting QFX3100 Director Devices in a Director Group” on page 242](#).
5. Based on your QFX3000-M QFabric system control plane type, connect each QFX Series device to each EX Series switch for control plane interconnection.
 - For a copper-based QFabric system control plane, see:
 - [Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network on page 245](#)
 - [Connecting a QFX3600-I Interconnect Device to a Copper-Based QFX3000-M QFabric System Control Plane Network on page 252](#)
 - [Connecting a QFX3500 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network on page 263](#)

- [Connecting a QFX3600 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network on page 259](#)
 - [Connecting a QFX5100 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network on page 256](#)
 - For a fiber-based QFabric system control plane, see:
 - [Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-M QFabric System Control Plane Network on page 273](#)
 - [Connecting a QFX3600-I Interconnect Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network on page 281](#)
 - [Connecting a QFX3500 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network on page 290](#)
 - [Connecting a QFX3600 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network on page 295](#)
 - [Connecting a QFX5100 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network on page 286](#)
6. Connect each Node device to each QFX3600-I Interconnect device for data plane interconnection. See:
- [Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device on page 311](#)
 - [Connecting a QFX3600 Node Device to a QFX3600-I Interconnect Device on page 306](#)
 - [Connecting a QFX5100 Node Device to a QFX3600-I Interconnect Device on page 301](#)
- Leave the Node and Interconnect devices powered on.
7. Configure each EX4200 or EX4300 switch using the recommended configuration as described in this example: [“Example: Configuring EX Series Switches for the QFX3000-M QFabric System Control Plane” on page 329](#).
- Leave the EX Series switches powered on.
8. Power on the QFX3100 Director devices and complete the initial configuration for the QFX3000-M QFabric system described in [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 363](#).

NOTE: When you power on the QFX3100 Director devices in the Director group for the first time the first device to be powered on assumes the Director Group 0 (*dg0*) role. The second device to be powered on assumes the Director Group 1 (*dg1*) role.

Understanding QFX3000-M QFabric System Hardware Configurations

The QFX3000-M QFabric system is made up of multiple hardware components:

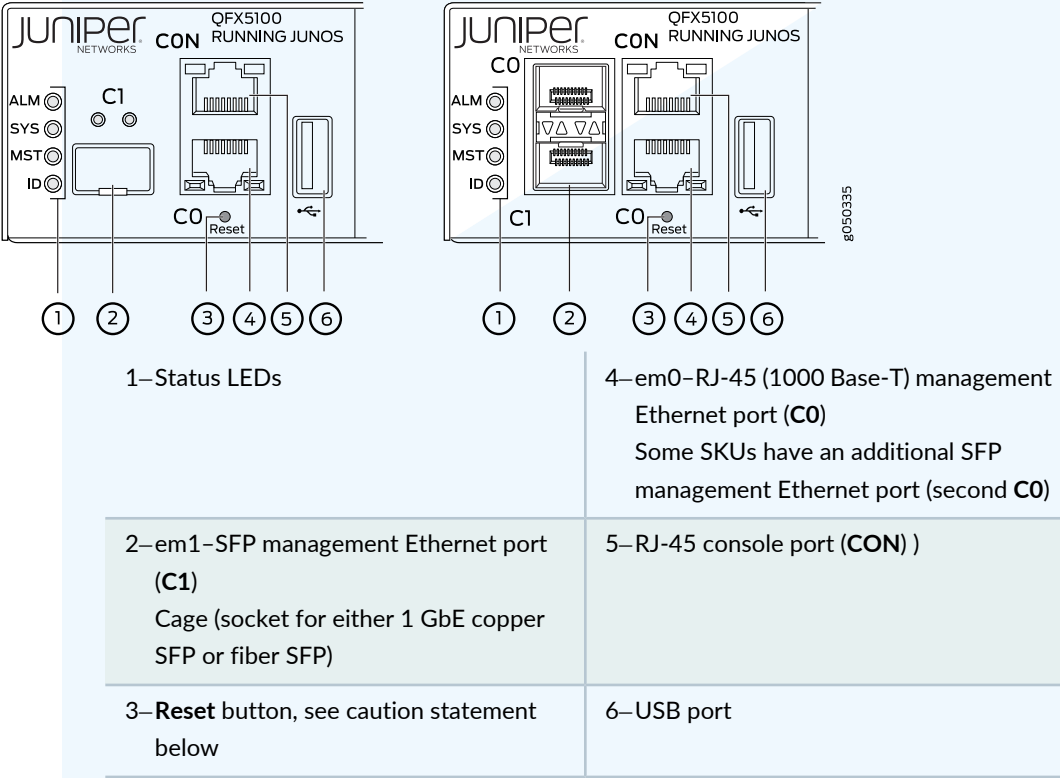
- **EX Series switches**—Two EX4200 switches or two EX4300 switches are required for a QFX3000-M QFabric system control plane and management network interconnection. You may not mix EX4200 and EX4300 switches on the same QFabric system.
 - **EX4200**—Use EX4200-24T switches with an SFP+ uplink module installed for copper-based control planes and EX4200-24F for fiber-based control planes. Two small form-factor pluggable (SFP) Gigabit Ethernet uplink ports on each EX4200 switch interconnect the two EX4200 switches.
 - **EX4300**—Use EX4300-48T switches with an SFP+ uplink module installed for copper-based control planes and EX4300-48P for fiber-based control planes. Four quad small form-factor pluggable plus (QSFP+) 40-Gigabit Ethernet uplink ports on each EX4300 interconnect the two EX Series switches.
- **QFX3100 Director devices**—Two QFX3100 Director devices are required for a QFX3000-M QFabric system. Together, the two Director devices are called a *Director group*.
 - Four Gigabit Ethernet ports on each QFX3100 Director device provide connection to the control plane and management network through the EX4200 switches.
 - Two Gigabit Ethernet ports on each QFX3100 Director device interconnect two Director devices in a Director group.
 - One Gigabit Ethernet RJ-45 management port on each QFX3100 Director device provides connection to the management network through your out-of-band management network.

NOTE: For a copper-based QFX3000-M QFabric system control plane network, use QFX3100 Director devices with RJ-45 network modules installed. For a fiber-based control plane network, use QFX3100 Director devices with SFP network modules installed.

- **QFX5100-24Q or QFX3600-I Interconnect devices**—Two interconnect devices are required for a QFX3000-M QFabric system. Up to four interconnect devices can be used in a QFX3000-M QFabric system.
 - **QFX5100-24Q Interconnect device**
 - Up to 24 40-Gbps QSFP+ ports on each Interconnect device connect the QFX5100, QFX3600, or QFX3500 Node devices to the data plane network across fiber-optic cables and a high-speed backplane.
 - Two Gigabit Ethernet management ports connect each QFX5100-24Q Interconnect device to the control plane and management network through the EX4200 switches.

NOTE: For a copper-based QFX3000-M QFabric system control plane network, use the 1000BASE-T management port labeled **C0** and the SFP management port labeled **C1** on QFX5100-24Q Interconnect devices. For a fiber-based control plane network, use the SFP management ports labeled **C0** and **C1** on QFX5100-24Q Interconnect devices.

Figure 16: QFX5100-24Q Management Panel



CAUTION: Do not use the **Reset** button to restart the power sequence unless under the direction of Juniper Networks Technical Assistance Center (JTAC).

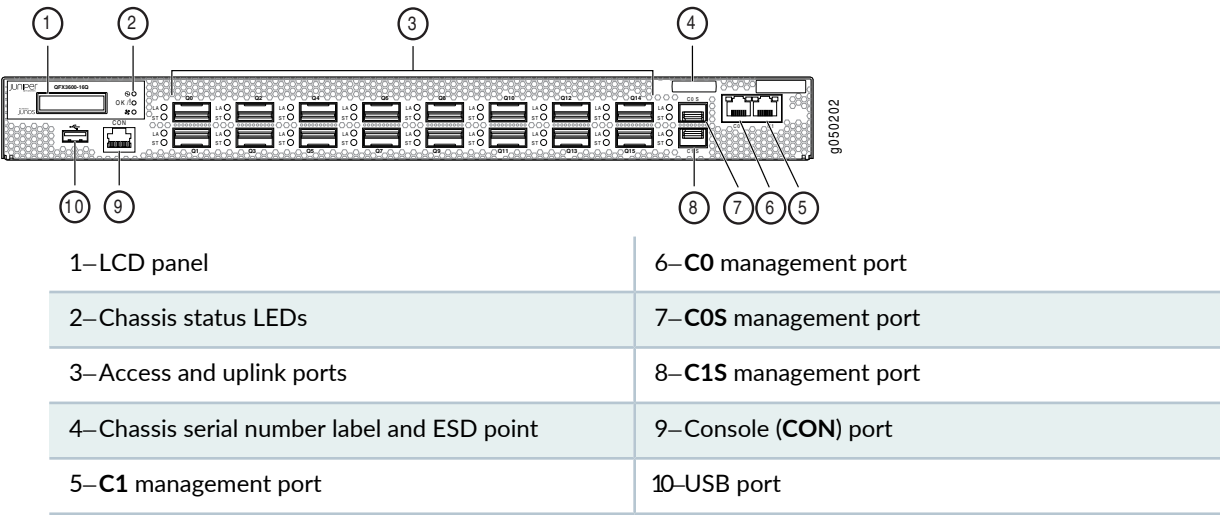
- QFX3600-I Interconnect device
- Up to 16 40-Gbps quad form-factor-pluggable plus (QSFP+) ports on each QFX3600-I Interconnect device connect the QFX3500 or QFX3600 Node devices to the data plane network across fiber-optic cables and a high-speed backplane.

- Two Gigabit Ethernet management ports connect each QFX3600-I Interconnect device to the control plane and management network through the EX Series switches.

NOTE: For a copper-based QFX3000-M QFabric system control plane network, use the 1000BASE-T management ports labeled **C0** and **C1** on QFX3600-I Interconnect devices. For a fiber-based control plane network, use the SFP management ports labeled **C0S** and **C1S** on QFX3600-I Interconnect devices.

Figure 17 on page 114 shows the front of a QFX3600-I Interconnect device.

Figure 17: QFX3600-I Interconnect Device Front Panel



- QFX3500, QFX3600, QFX5100-24Q, QFX5100-48S, or QFX5100-48T Node devices—Up to 16 Node devices can be connected to the QFX3000-M QFabric system.

Table 21: Connecting Node Devices to Servers, Storage, or External Networks

Node Device	Description
QFX3500	<p>Total of 48 10-Gbps SFP+ ports.</p> <ul style="list-style-type: none"> • 12 (ports 0 through 5 and 42 through 47) can be configured as 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel ports. • 36 (ports 6 through 41) can be configured as Gigabit Ethernet ports. • No limitations on which ports can be configured as 10-Gigabit Ethernet ports. • However, if a Fibre Channel transceiver is in ports 0 through 5 or 42 through 47, <i>all</i> of the ports in that group must be configured as Fibre Channel, even though only one port in a group might have a Fibre Channel transceiver, and no port in the group can use a 10-Gigabit Ethernet transceiver.

Table 21: Connecting Node Devices to Servers, Storage, or External Networks (*continued*)

Node Device	Description
QFX3600	<p>Up to 14 40-Gbps QSFP+ ports (Q2 through Q15) on each Node device use QSFP+ DAC breakout cables or QSFP+ transceivers with fiber breakout cables to support 48 10-Gigabit Ethernet interfaces.</p> <ul style="list-style-type: none"> • By default, 12 ports (Q4 through Q15) are configured for 10-Gigabit Ethernet connections between the Node device and servers, storage, or external networks. • Optionally, you can configure ports Q2 through Q15 for 10-Gigabit Ethernet connections.
QFX5100-24Q	<p>The base model supports 24 40-Gbps Ethernet QSFP+ ports. With two QFX-EM-4Q expansion modules inserted, eight additional QSFP+ ports provide a total of 32 40-Gbps ports. The availability of these ports depends on the system mode you configure.</p> <p>Optionally, you can configure these ports on each Node device use QSFP+ direct-attach copper (DAC) breakout cables or QSFP+ transceivers with fiber breakout cables to support 96 10-Gigabit Ethernet interfaces. With two QFX-EM-4Q expansion modules inserted, eight additional QSFP+ ports provide a total of 104 10-Gbps ports.</p>
QFX5100-48S	<p>Total of 48 10-Gbps SFP+ ports and 6 40-Gbps QSFP+ ports.</p> <ul style="list-style-type: none"> • 48 (ports 0 through 47) can be configured as either 1-Gigabit Ethernet, using a breakout cable, or 10-Gigabit Ethernet interfaces. • By default, 4 of the QSFP+ interfaces (labeled fte-0/1/0 through fte-0/1/3) are configured for 40-Gbps uplink connections between your Node device and your Interconnect devices, and 2 QSFP+ interfaces (labeled xle-0/1/4 and xle-0/1/5) provide 40-Gigabit Ethernet connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the middle two interfaces (xle-0/1/8 and xle-0/1/15) for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks, and you can choose to configure the last two interfaces (fte-0/1/4 and fte-0/1/5) for uplink connections between your Node device and your Interconnect devices.
QFX5100-48T	<p>The base supports 48 10GBASE-T ports and 6 QSFP+ ports.</p> <ul style="list-style-type: none"> • The forty-eight copper physical ports (0 through 47) are tri-speed (10GbE/1GbE/100 Mbps) and can be configured as access ports. • By default, 4 of the QSFP+ interfaces (labeled fte-0/1/0 through fte-0/1/3) are configured for 40-Gbps uplink connections between your Node device and your Interconnect devices, and 2 QSFP+ interfaces (labeled xle-0/1/4 and xle-0/1/5) provide 40-Gigabit Ethernet connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the middle two interfaces (xle-0/1/8 and xle-0/1/15) for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks, and you can choose to configure the last two interfaces (fte-0/1/4 and fte-0/1/5) for uplink connections between your Node device and your Interconnect devices.

Table 22: Connecting Node Devices to the Data Plane

Node Device	Interconnect Device	Description
QFX3500	QFX3600-I or QFX5100-24Q	Up to four 40-Gbps QSFP+ uplink ports (Q0 through Q3) connect to the data plane network.
QFX3600	QFX3600-I or QFX5100-24Q	<p>Up to eight 40-Gbps QSFP+ uplink ports (Q0 through Q7) connect to the data plane network.</p> <ul style="list-style-type: none"> By default, four ports (Q0 through Q3) are configured for connection between the Node device and the data plane network. Optionally, you can configure the first eight ports (Q0 through Q7) for this connection.
QFX5100-24Q	QFX3600-I	When introducing a QFX5100-24Q into an existing QFX3000-M QFabric System, install the QFX5100-24Q as the interconnect device and convert the QFX3600-I to a QFX3600 for use as a node device. See "Converting the Device Mode for a QFabric System Component" on page 322.
	QFX5100-24Q	Any of the 24 40-QSFP+ uplink ports (0 through 23) may be used to connect to the data plane network.
QFX5100-48S	QFX3600-I or QFX5100-24Q	Up to six 40-Gbps QSFP+ uplink ports (48 through 53) connect to the data plane network.
QFX5100-48T	QFX3600-I or QFX5100-24Q	Up to six 40-Gbps QSFP+ uplink ports (48 through 53) connect to the data plane network.

Table 23: Connecting Node Devices to the Control Plane and Management Network through EX4200 Switches

Node Device	Description
QFX3500	<p>Two Gigabit Ethernet management ports on each Node device provide connection to the control plane and management network.</p> <ul style="list-style-type: none"> For a copper-based QFX3000-M QFabric system control plane network, use the management ports on Node devices with a 1000BASE-T management board installed. For a fiber-based control plane network, use the management ports on Node devices with an SFP management board installed.

Table 23: Connecting Node Devices to the Control Plane and Management Network through EX4200 Switches (continued)

Node Device	Description
QFX3600	<p>Two Gigabit Ethernet management ports on each Node device provide connection to the control plane and management network.</p> <ul style="list-style-type: none"> For a copper-based control plane network, use the 1000BASE-T management ports (C0 and C1) on Node devices. For a fiber-based control plane network, use the SFP management ports (C0S and C1S) on Node devices.
QFX5100-24Q	<p>Two Gigabit Ethernet management ports on each Node device provide connection to the control plane and management network.</p> <ul style="list-style-type: none"> For a copper-based QFX3000-M QFabric system control plane network, use the RJ-45 management port (C0) and the SFP management port (C1) on Node devices with a 1 Gbps copper SFP. For a fiber-based control plane network, use the SFP management port (C0) and (C1) on Node devices with a 1 Gbps fiber SFP. <p>NOTE: Not all QFX5100-24Q product SKUs support an all-fiber connection.</p>
QFX5100-48S	<p>Two Gigabit Ethernet management ports on each Node device provide connection to the control plane and management network.</p> <ul style="list-style-type: none"> For a copper-based QFX3000-M QFabric system control plane network, use the RJ-45 management port (C0) and the SFP management port (C1) on Node devices with a 1 Gbps copper SFP. For a fiber-based control plane network, use the SFP management port (C0) and (C1) on Node devices with a 1 Gbps fiber SFP. <p>NOTE: Not all QFX5100-48S product SKUs support an all-fiber connection.</p>
QFX5100-48S	<p>Two Gigabit Ethernet management ports on each Node device provide connection to the control plane and management network.</p> <ul style="list-style-type: none"> For a copper-based QFX3000-M QFabric system control plane network, use the RJ-45 management port (C0) and the SFP management port (C1) on Node devices with a 1 Gbps copper SFP. For a fiber-based control plane network, use the SFP management port (C0) and (C1) on Node devices with a 1 Gbps fiber SFP. <p>NOTE: Not all QFX5100-48S product SKUs support an all-fiber connection.</p>

Table 23: Connecting Node Devices to the Control Plane and Management Network through EX4200 Switches (continued)

Node Device	Description
QFX5100-48T	<p>Two Gigabit Ethernet management ports on each Node device provide connection to the control plane and management network.</p> <ul style="list-style-type: none"> For a copper-based QFX3000-M QFabric system control plane network, use the RJ-45 management port (C0) and the SFP management port (C1) on Node devices with a 1 Gbps copper SFP. For a fiber-based control plane network, use the SFP management port (C0) and (C1) on Node devices with a 1 Gbps fiber SFP.

RELATED DOCUMENTATION

[QFX3000-M QFabric System Installation Overview | 109](#)

Management Panel of a QFX5100 Device

Planning a QFX3000-M QFabric System Deployment

A QFX3000-M QFabric system is formed by interconnecting QFX3500, QFX3600, or QFX5100 Node devices, QFX3600-I or QFX5100-24Q Interconnect devices, and QFX3100 Director devices. Two Juniper Networks EX Series switches are used to interconnect the control plane and management network.

Before installing a QFX3000-M QFabric system, you must consider the following factors:

- The number of devices in the QFabric system and their location.
 - You must have two QFX3100 Director devices operating in a Director group.
 - You can have up to four Interconnect devices, either QFX5100-24Q or QF3600-I, but not intermixed in the same QFabric Node group.
 - You can interconnect up to 16 QFX Series Switches as Node devices using either QFX5100-24Q or QFX3600-I as the Interconnect Devices. You may not mix QFX5100-24Q Interconnect devices with QFX3600-I Interconnect devices on the same QFX3000-M QFabric system. Supported models are:
 - QFX5100-24Q
 - QFX5100-48S
 - QFX5100-48T

- QFX3600
- QFX3500
- The number of Node devices you require depends on the following factors:
 - The number of access ports you need for connections to either endpoint systems (such as servers and storage devices) or external networks.
 - The oversubscription ratio you need on the access ports.
 - The number of access ports supported on each Node device based on that oversubscription ratio.

Table 24 on page 119 shows the number of 10-Gigabit Ethernet access ports supported on Node devices based on the oversubscription ratio you need on the access ports.

Table 24: Number of 10-Gigabit Ethernet Access Ports Supported on Node Devices Based on Oversubscription Ratio

Oversubscription Ratio on Access Ports	Number of 10-Gigabit Ethernet Access Ports Supported on a QFX3500 Node Device	Number of 10-Gigabit Ethernet Access Ports Supported on a QFX3600 Node Device	Number of 10-Gigabit Ethernet Access Ports Supported on a QFX5100-24Q Node Device	Number of 10-Gigabit Ethernet Access Ports Supported on a QFX5100-48S and QFX5100-48T Node Device
1:1 (no oversubscription)	1:1 oversubscription not supported	32	104	1:1 oversubscription not supported
3:1	48	48	104 ports with 2 QFX-EM-4Q expansion modules	3:1 oversubscription not supported
3.5:1	3.5:1 oversubscription not supported	3.5:1 oversubscription not supported	3.5:1 oversubscription not supported	56
6:1	48	48	96 ports with 2 QFX-EM-4Q expansion modules	3:1 oversubscription not supported

To calculate the required number of Node devices (N_n) for your QFabric system, divide the number of access ports you need for connections to either endpoint systems or external networks (N_p) by the number of access ports supported on each Node device (N_a) based on the required oversubscription ratio, and round up the resulting value. For example, if you need 300 10-Gigabit Ethernet access ports at 3:1 oversubscription, and your Node device supports 48 10-Gigabit Ethernet access ports at 3:1 oversubscription, you require 7 Node devices as shown below:

$$N_n = N_p / N_a$$

$$N_n = 300/48$$

$$N_n = 6.25 \text{ (rounded up to 7)}$$

- The number of Interconnect devices you require depends on the number of Node devices and the oversubscription ratio required on the access ports of the Node devices. See [Table 25 on page 120](#) to determine the number of Interconnect devices you require. For example, if you plan to install 14 Node devices and require 3:1 oversubscription ratio on the Node devices, you must install 4 Interconnect devices.

Table 25: Maximum Number of Node Devices Supported Based on Oversubscription Ratio and Number of Interconnect Devices

Oversubscription Ratio	QFX Series Switch Models	Number of Interconnect Devices	Maximum Number of Node Devices
1:1	QFX3600 or	2	4
		4	8
	QFX5100-24Q	8	32 connecting to QFX5100-24Q Interconnect devices and 16 connecting to QFX3600-I Interconnect devices
2:1	QFX5100-24Q	4	32 connecting to QFX5100-24Q Interconnect devices and 16 connecting to QFX3600-I Interconnect devices
3:1	<ul style="list-style-type: none"> • QFX3600 • QFX3500 	2	32 connecting to QFX5100-24Q Interconnect devices and 16 connecting to QFX3600-I Interconnect devices
		4	32 connecting to QFX5100-24Q Interconnect devices and 16 connecting to QFX3600-I Interconnect devices

Table 25: Maximum Number of Node Devices Supported Based on Oversubscription Ratio and Number of Interconnect Devices (*continued*)

Oversubscription Ratio	QFX Series Switch Models	Number of Interconnect Devices	Maximum Number of Node Devices
3.5:1	QFX5100-48S or QFX5100-48T	2	32 connecting to QFX5100-24Q Interconnect devices and 16 connecting to QFX3600-I Interconnect devices
		4	32 connecting to QFX5100-24Q Interconnect devices and 16 connecting to QFX3600-I Interconnect devices
6:1	<ul style="list-style-type: none"> • QFX5100-24Q • QFX5100-48S • QFX5100-48T • QFX3600 • QFX3500 	2	32 connecting to QFX5100-24Q Interconnect devices and 16 connecting to QFX3600-I Interconnect devices

For information about the size and strength of racks for the devices, see the following topics:

- *Rack Requirements for a QFX3100 Director Device*
- *Rack Requirements for a QFX5100 Device*
- *Rack Requirements for a QFX3600 or QFX3600-I Device*
- *Rack Requirements for a QFX3500 Device*
- *Rack Requirements for EX4200 Switches*

For the dimensions and weights of the devices, see the following topics:

- *Chassis Physical Specifications for a QFX3100 Director Device*
- *Chassis Physical Specifications for a QFX5100 Device*
- *Chassis Physical Specifications for QFX3600 and QFX3600-I Devices*
- *Chassis Physical Specifications for a QFX3500 Chassis*
- *Chassis Physical Specifications for EX4200 Switches*
- The type of control plane and management network—copper-based or fiber-based.
 - For a copper-based control plane and management network, use:

- QFX3100 Director devices with RJ-45 network modules installed
- QFX5100 Node devices using the RJ-45 network port (C0 and a 1 Gbps SFP installed in network port (C1).
- QFX3500 Node devices with a 1000BASE-T management board installed
- EX4200-24T switches with an SFP+ uplink module installed
- For a fiber-based control plane and management network, use:
 - QFX3100 Director devices with SFP network modules installed
 - QFX5100 Node devices using 1 Gbps SFP installed in network ports (C0) and (C1).
 - QFX3500 Node devices with an SFP management board installed
 - EX4200-24F switches with an SFP+ uplink module installed
- Cabling requirements for the control plane and management network
 - The copper-based control plane and management network are interconnected using standard 1000BASE-T Ethernet over copper wiring. Each network segment can be a maximum length of 100 m (328 ft). See [“Cable Specifications for Copper-Based Control Plane Connections for the QFabric System” on page 144.](#)
 - The fiber-based control plane and management network are interconnected using small form-factor pluggable (SFP) transceivers over fiber-optic cabling. Each network segment can be a maximum length of 10 km (6.2 miles). However, keep in mind that each network segment in the data plane is limited to a maximum length of 150 m (492 ft). See [The Hardware Compatibility Tool.](#)
- Cabling requirements for the data plane—The data plane is interconnected using standard 40GBASE-SR 40-Gigabit Ethernet QSFP+ optical transceivers over fiber-optic wiring. If you use OM3 optical fiber, each network segment can be a maximum of 100 m (328 ft). If you use OM4 optical fiber, each network segment can be a maximum of 150 m (492 ft). However, keep in mind that if you are using a copper-based control plane and management network, each copper network segment is limited to a maximum length of 100 m (328 ft). This limitation of the copper control plane network segment effectively limits the network segments in your data plane to 100 m (328 ft). See [The Hardware Compatibility Tool.](#)

The data plane cabling requirements depend on the number of connections required from each Node device to Interconnect devices based on the oversubscription ratio you need on the access ports of the Node device. [Table 26 on page 123](#) shows the number of connections required from each Node device to Interconnect devices based on the oversubscription ratio you need on the Node device.

Table 26: Number of Connections Required Between Node and Interconnect Devices Based on Oversubscription Ratio

Oversubscription Ratio	QFX Series Switch Model	Number of Connections from Each Node Device to Interconnect Devices
1:1	<ul style="list-style-type: none"> • QFX3600 • QFX5100-24Q 	8
3:1	<ul style="list-style-type: none"> • QFX3600 • QFX3500 	4
3.5:1	<ul style="list-style-type: none"> • QFX5100-48S • QFX5100-48T 	4
6:1	<ul style="list-style-type: none"> • QFX5100-24Q • QFX5100-48S • QFX5100-48T • QFX3600 • QFX3500 	2

- **Power supply**—You must plan the installation site to meet the power requirements of all the devices in the QFabric system. For information on power requirements and configuration options for each device, see the following topics:
 - [AC Power Specifications for a QFX3100 Director Device on page 147](#)
 - [AC Power Specifications for a QFX3600 or QFX3600-I Device on page 149](#)
 - [DC Power Specifications for a QFX3600 or QFX3600-I Device on page 161](#)
 - [AC Power Specifications for a QFX5100 Device on page 151](#)
 - [DC Power Specifications for a QFX5100 Device on page 162](#)
 - [AC Power Specifications for a QFX3500 Device on page 150](#)
 - [DC Power Specifications for a QFX3500 Device on page 162](#)
 - [Power Specifications for EX4200 Switches](#)
- For more information about the site preparation requirements for each device, see the following topics:
 - [Site Preparation Checklist for a QFX3100 Director Device](#)
 - [Site Preparation Checklist for a QFX3600 or QFX3600-I Device](#)
 - [Site Preparation Checklist for a QFX5100 Device](#)

- *Site Preparation Checklist for a QFX3500 Device*
- *Site Preparation Checklist for EX4200 Switches*

RELATED DOCUMENTATION

QFX3000-M QFabric System Installation Overview 109
Understanding QFX3000-M QFabric System Hardware Configurations 112
Clearance Requirements for Airflow and Hardware Maintenance for a QFX3100 Director Device
Clearance Requirements for Airflow and Hardware Maintenance for a QFX5100 Device
Clearance Requirements for Airflow and Hardware Maintenance for a QFX3600 or QFX3600-I Device
Clearance Requirements for Airflow and Hardware Maintenance for a QFX3500 Device
Clearance Requirements for Airflow and Hardware Maintenance for EX4200 Switches

General Site Guidelines

Efficient device operation requires proper site planning and maintenance and proper layout of the equipment, rack or cabinet (if used), and wiring closet.

To plan and create an acceptable operating environment for your device and prevent environmentally caused equipment failures:

- Keep the area around the chassis free from dust and conductive material, such as metal flakes.
- Follow prescribed airflow guidelines to ensure that the cooling system functions properly and that exhaust from other equipment does not blow into the intake vents of the device.
- Follow the prescribed electrostatic discharge (ESD) prevention procedures to prevent damaging the equipment. Static discharge can cause components to fail completely or intermittently over time.
- Install the device in a secure area, so that only authorized personnel can access the device.

Site Electrical Wiring Guidelines

[Table 27 on page 125](#) describes the factors you must consider while planning the electrical wiring at your site.



WARNING: You must provide a properly grounded and shielded environment and use electrical surge-suppression devices.

Table 27: Site Electrical Wiring Guidelines

Site Wiring Factor	Guidelines
Signaling limitations	<p>If your site experiences any of the following problems, consult experts in electrical surge suppression and shielding:</p> <ul style="list-style-type: none"> • Improperly installed wires cause radio frequency interference (RFI). • Damage from lightning strikes occurs when wires exceed recommended distances or pass between buildings. • Electromagnetic pulses (EMPs) caused by lightning damage unshielded conductors and electronic devices.
Radio frequency interference	<p>To reduce or eliminate RFI from your site wiring, do the following:</p> <ul style="list-style-type: none"> • Use a twisted-pair cable with a good distribution of grounding conductors. • If you must exceed the recommended distances, use a high-quality twisted-pair cable with one ground conductor for each data signal when applicable.
Electromagnetic compatibility	<p>If your site is susceptible to problems with electromagnetic compatibility (EMC), particularly from lightning or radio transmitters, seek expert advice.</p> <p>Some of the problems caused by strong sources of electromagnetic interference (EMI) are:</p> <ul style="list-style-type: none"> • Destruction of the signal drivers and receivers in the device • Electrical hazards as a result of power surges conducted over the lines into the equipment

Environmental Requirements and Specifications for a QFX3100 Director Device

The device must be installed in a rack or cabinet housed in a dry, clean, well-ventilated, and temperature-controlled environment.

Follow these environmental guidelines:

- The site must be as dust-free as possible, because dust can clog air intake vents and filters, reducing the efficiency of the device cooling system.

- Maintain ambient airflow for normal device operation. If the airflow is blocked or restricted, or if the intake air is too warm, the device might overheat, leading to the device temperature monitor shutting down the device to protect the hardware components.

Table 28 on page 126 provides the required environmental conditions for normal device operation.

Table 28: QFX3100 Director Device Environmental Tolerances

Description	Tolerance
Altitude	No performance degradation to 10,000 feet (3048 meters)
Relative humidity	Normal operation ensured in relative humidity range of 5% through 85%, noncondensing
Temperature	<ul style="list-style-type: none"> • Normal operation ensured in temperature range of 32° F through 104° F (0° C through 40° C) • Short-term operation ensured in temperature range of 23° F through 131° F (–5° C through 55° C) <p>NOTE: As defined in NEBS GR-63-CORE, Issue 3, short-term events can be up to 96 hours in duration but not more than 15 days per year.</p> <ul style="list-style-type: none"> • Nonoperating storage temperature in shipping container: –40° F through 158° F (–40° C through 70° C)

NOTE: Install QFX Series devices only in restricted areas, such as dedicated equipment rooms and equipment closets, in accordance with Articles 110-16, 110-17, and 110-18 of the National Electrical Code, ANSI/NFPA 70.

RELATED DOCUMENTATION

Clearance Requirements for Airflow and Hardware Maintenance for a QFX3100 Director Device

[Installing and Connecting a QFX3100 Director Device | 165](#)

Environmental Requirements and Specifications for QFX3600 and QFX3600-I Devices

The QFX3600 and QFX3600-I devices must be installed in a rack or cabinet housed in a dry, clean, well-ventilated, and temperature-controlled environment.

Follow these environmental guidelines:

- The site must be as dust-free as possible, because dust can clog air intake vents and filters, reducing the efficiency of the device cooling system.
- Maintain ambient airflow for normal device operation. If the airflow is blocked or restricted, or if the intake air is too warm, the device might overheat, leading to the device temperature monitor shutting down the device to protect the hardware components.

[Table 29 on page 127](#) provides the required environmental conditions for normal device operation.

Table 29: QFX3600 and QFX3600-I Device Environmental Tolerances

Description	Tolerance
Altitude	No performance degradation to 10,000 feet (3048 meters)
Relative humidity	Normal operation ensured in relative humidity range of 5% through 85%, noncondensing
Temperature	<ul style="list-style-type: none"> • Normal operation ensured in temperature range of 32° F through 104° F (0° C through 40° C) • Short-term operation ensured in temperature range of 23° F through 131° F (-5° C through 55° C) <p>NOTE: As defined in NEBS GR-63-CORE, Issue 3, short-term events can be up to 96 hours in duration but not more than 15 days per year.</p> <ul style="list-style-type: none"> • Nonoperating storage temperature in shipping container: -40° F through 158° F (-40° C through 70° C)
Seismic	Designed to comply with Zone 4 earthquake requirements per NEBS GR-63-CORE, Issue 3.

NOTE: Install QFX Series devices only in restricted areas, such as dedicated equipment rooms and equipment closets, in accordance with Articles 110-16, 110-17, and 110-18 of the National Electrical Code, ANSI/NFPA 70.

RELATED DOCUMENTATION

Clearance Requirements for Airflow and Hardware Maintenance for a QFX3600 or QFX3600-I Device

[Installing and Connecting a QFX3600 or QFX3600-I Device](#) | 177

Environmental Requirements and Specifications for a QFX3500 Device

The device must be installed in a rack or cabinet housed in a dry, clean, well-ventilated, and temperature-controlled environment.

Follow these environmental guidelines:

- The site must be as dust-free as possible, because dust can clog air intake vents and filters, reducing the efficiency of the device cooling system.
- Maintain ambient airflow for normal device operation. If the airflow is blocked or restricted, or if the intake air is too warm, the device might overheat, leading to the device temperature monitor shutting down the device to protect the hardware components.

Table 30 on page 128 provides the required environmental conditions for normal device operation.

Table 30: QFX3500 Device Environmental Tolerances

Description	Tolerance
Altitude	No performance degradation to 10,000 feet (3048 meters)
Relative humidity	Normal operation ensured in relative humidity range of 5% through 85%, noncondensing
Temperature	<ul style="list-style-type: none"> • Normal operation ensured in temperature range of 32° F through 104° F (0° C through 40° C) • Short-term operation ensured in temperature range of 23° F through 131° F (–5° C through 55° C) <p>NOTE: As defined in NEBS GR-63-CORE, Issue 3, short-term events can be up to 96 hours in duration but not more than 15 days per year.</p> <ul style="list-style-type: none"> • Nonoperating storage temperature in shipping container: –40° F through 158° F (–40° C through 70° C)
Seismic	Designed to comply with Zone 4 earthquake requirements per NEBS GR-63-CORE, Issue 3.

NOTE: Install QFX Series devices only in restricted areas, such as dedicated equipment rooms and equipment closets, in accordance with Articles 110-16, 110-17, and 110-18 of the National Electrical Code, ANSI/NFPA 70.

RELATED DOCUMENTATION

Clearance Requirements for Airflow and Hardware Maintenance for a QFX3500 Device

[Installing and Connecting a QFX3500 Device](#) | 199

Environmental Requirements and Specifications for a QFX5100 Device

The switch must be installed in a rack or cabinet. It must be housed in a dry, clean, well-ventilated, and temperature-controlled environment.

Follow these environmental guidelines:

- The site must be as dust-free as possible, because dust can clog air intake vents and filters, reducing the efficiency of the switch cooling system.
- Maintain ambient airflow for normal switch operation. If the airflow is blocked or restricted, or if the intake air is too warm, the switch might overheat, leading to the switch temperature monitor shutting down the device to protect the hardware components.

[Table 31 on page 129](#) provides the required environmental conditions for normal switch operation.

Table 31: QFX5100 Switch Environmental Tolerances

Description	Tolerance
Altitude	No performance degradation to 6,562 feet (2000 meters)
Relative humidity	<p>Normal operation ensured in relative humidity range of 5% through 90%, noncondensing</p> <ul style="list-style-type: none"> • Short-term operation ensured in relative humidity range of 5% through 93%, noncondensing <p>NOTE: As defined in NEBS GR-63-CORE, Issue 3, short-term events can be up to 96 hours in duration but not more than 15 days per year.</p>
Temperature	<ul style="list-style-type: none"> • Normal operation ensured in temperature range of 32° F through 104° F (0° C through 40° C) <p>NOTE: Customers with QFX5100-48T switches should ensure the room temperature does not exceed a 2° C increase or decrease per minute.</p> <ul style="list-style-type: none"> • Nonoperating storage temperature in shipping container: -40° F through 158° F (-40° C through 70° C)
Seismic	Designed to comply with Zone 4 earthquake requirements per NEBS GR-63-CORE, Issue 3.

NOTE: Install QFX Series devices only in restricted areas, such as dedicated equipment rooms and equipment closets, in accordance with Articles 110-16, 110-17, and 110-18 of the National Electrical Code, ANSI/NFPA 70.

RELATED DOCUMENTATION

Clearance Requirements for Airflow and Hardware Maintenance for a QFX5100 Device
[Installing and Connecting a QFX5100 Device | 219](#)

Environmental Requirements and Specifications for EX Series Switches

The switch must be installed in a rack or cabinet housed in a dry, clean, well-ventilated, and temperature-controlled environment.

Ensure that these environmental guidelines are followed:

- The site must be as dust-free as possible, because dust can clog air intake vents and filters, reducing the efficiency of the switch cooling system.
- Maintain ambient airflow for normal switch operation. If the airflow is blocked or restricted, or if the intake air is too warm, the switch might overheat, leading to the switch temperature monitor shutting down the switch to protect the hardware components.

[Table 32 on page 131](#) provides the required environmental conditions for normal switch operation.

Table 32: EX Series Switch Environmental Tolerances

Switch or device	Environment Tolerance			
	Altitude	Relative Humidity	Temperature	Seismic
EX2200-C	No performance degradation up to 5,000 feet (1524 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation ensured in the temperature range 32° F (0° C) through 104° F (40° C) at altitudes up to 5,000 ft (1,524 m). For information about extended temperature SFP transceivers supported on EX2200 switches, see <i>Pluggable Transceivers Supported on EX2200 Switches</i> .	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.
EX2200 (except EX2200-C switches)	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation ensured in the temperature range 32° F (0° C) through 113° F (45° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.
EX2300-C	No performance degradation up to 5,000 feet (1524 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation ensured in the temperature range 32° F (0° C) through 104° F (40° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.
EX2300 (except EX2300-C switches)	No performance degradation up to 13,000 feet (3962 meters) at 104° F (40° C) as per GR-63	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation ensured in the temperature range 32° F (0° C) through 113° F (45° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.
EX3200	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation ensured in the temperature range 32° F (0° C) through 113° F (45° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.
EX3300	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation ensured in the temperature range 32° F (0° C) through 113° F (45° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.

Table 32: EX Series Switch Environmental Tolerances (*continued*)

Switch or device	Environment Tolerance			
	Altitude	Relative Humidity	Temperature	Seismic
EX3400	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation ensured in the temperature range 32° F (0° C) through 113° F (45° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.
EX4200	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation ensured in the temperature range 32° F (0° C) through 113° F (45° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.
EX4300 The maximum thermal output for EX4300-48T is 423 BTU/hour and for EX4300-48P is 5844 BTU/hour.	EX4300 switches except the EX4300-48MP model— No performance degradation up to 10,000 feet (3048 meters) EX4300-48MP model— No performance degradation up to 6,000 feet (1829 meters)	EX4300 switches except the EX4300-48MP model— Normal operation ensured in the relative humidity range 10% through 85% (noncondensing) EX4300-48MP model— Normal operation ensured in the relative humidity range 5% through 90% (noncondensing)	Normal operation ensured in the temperature range 32° F (0° C) through 113° F (45° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.
EX4500	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation ensured in the temperature range 32° F (0° C) through 113° F (45° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.

Table 32: EX Series Switch Environmental Tolerances (*continued*)

Switch or device	Environment Tolerance			
	Altitude	Relative Humidity	Temperature	Seismic
EX4550	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	<ul style="list-style-type: none"> EX4550-32F switches— Normal operation ensured in the temperature range 32° F (0° C) through 113° F (45° C) EX4550-32T switches— Normal operation is ensured in the temperature range 32° F through 104° F (40° C) 	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.
EX4600	No performance degradation to 6,562 feet (2000 meters)	<p>Normal operation ensured in the relative humidity range 5% through 90%, noncondensing</p> <ul style="list-style-type: none"> Short-term operation ensured in the relative humidity range 5% through 93%, noncondensing <p>NOTE: As defined in NEBS GR-63-CORE, Issue 4, short-term events can be up to 96 hours in duration but not more than 15 days per year.</p>	<ul style="list-style-type: none"> Normal operation ensured in the temperature range 32° F (0° C) through 113° F (45° C) Nonoperating storage temperature in shipping container: – 40° F (–40° C) through 158° F (70° C) 	Designed to comply with Zone 4 earthquake requirements per NEBS GR-63-CORE, Issue 4.
EX4650	No performance degradation to 6,000 feet (1829 meters)	Normal operation ensured in the relative humidity range 10% through 85% (condensing)	Normal operation is ensured in the temperature range 32° F (0° C) through 104° F (40° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.
EX6210	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation is ensured in the temperature range 32° F (0° C) through 104° F (40° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.

Table 32: EX Series Switch Environmental Tolerances (*continued*)

Switch or device	Environment Tolerance			
	Altitude	Relative Humidity	Temperature	Seismic
EX8208	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation is ensured in the temperature range 32° F (0° C) through 104° F (40° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.
EX8216	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation is ensured in the temperature range 32° F (0° C) through 104° F (40° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.
EX9204	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 5% through 90% (noncondensing)	Normal operation is ensured in the temperature range 32° F (0° C) through 104° F (40° C) Nonoperating storage temperature in shipping container: - 40° F (-40° C) to 158° F (70° C)	Complies with Zone 4 earthquake requirements as per GR-63.
EX9208	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 5% through 90% (noncondensing)	Normal operation is ensured in the temperature range 32° F (0° C) through 104° F (40° C) Nonoperating storage temperature in shipping container: - 40° F (-40° C) to 158° F (70° C)	Complies with Zone 4 earthquake requirements as per GR-63.
EX9214	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 5% through 90% (noncondensing)	Normal operation is ensured in the temperature range 32° F (0° C) through 104° F (40° C) Nonoperating storage temperature in shipping container: - 40° F (-40° C) through 158° F (70° C)	Complies with Zone 4 earthquake requirements as per GR-63.

Table 32: EX Series Switch Environmental Tolerances (*continued*)

Switch or device	Environment Tolerance			
	Altitude	Relative Humidity	Temperature	Seismic
EX9251 The maximum thermal output is 1705 BTU/hour (500 W).	No performance degradation up to 10,000 ft (3048 m)	Normal operation ensured in relative humidity range of 5% to 90%, noncondensing	Normal operation ensured in temperature range of 32° F (0° C) to 104° F (40° C) Nonoperating storage temperature in shipping container: - 40° F (-40° C) to 158° F (70° C)	Complies with Telcordia Technologies Zone 4 earthquake requirements
XRE200	No performance degradation up to 10,000 feet (3048 meters)	Normal operation ensured in the relative humidity range 10% through 85% (noncondensing)	Normal operation ensured in the temperature range 41° F (5° C) through 104° F (40° C)	Complies with Zone 4 earthquake requirements as per GR-63, Issue 4.

NOTE: Install EX Series switches only in restricted areas, such as dedicated equipment rooms and equipment closets, in accordance with Articles 110- 16, 110- 17, and 110- 18 of the National Electrical Code, ANSI/NFPA 70.

Ports and Connectors

IN THIS CHAPTER

- Determining Interface Support for the QFX3600 Device | 137
- Determining Interface Support for the QFX3500 Device | 138
- Determining Interface Support for the QFX5100 Device | 140
- Determining Transceiver Support for QFabric Systems | 143
- Cable Specifications for Copper-Based Control Plane Connections for the QFabric System | 144

Determining Interface Support for the QFX3600 Device

The QFX3600 device provides 16 QSFP+ ports, which support QSFP+ transceivers and QSFP+ DAC or DAC breakout cables.

On a QFX3600 Node device in a QFabric system, four ports (labeled **Q0** through **Q3**) operate as 40-Gbps data uplink (**fte**) ports for uplink connections between your QFX3600 Node device and your Interconnect device. Twelve ports (labeled **Q4** through **Q15**) operate in 10-Gigabit Ethernet (**xe**) mode to support 48 10-Gigabit Ethernet interfaces for connections to either endpoint systems or external networks. Optionally, you can choose to configure ports **Q0** through **Q7** to operate as 40-Gbps data uplink (**fte**) ports, and ports **Q2** through **Q15** to operate in 10-Gigabit Ethernet (**xe**) or 40-Gigabit Ethernet (**xle**) mode. See ["Configuring the Port Type on QFX3600 Node Devices" on page 408](#) for more information.

On a QFX3600 standalone switch, ports **Q0** through **Q15** operate as 40-Gigabit Ethernet (**xle**) ports. Optionally, you can choose to configure ports **Q0** through **Q15** to operate as 10-Gigabit Ethernet (**xe**) ports. See *Configuring the Port Type on QFX3600 Standalone Switches* for more information.

NOTE: On a QFX3600 standalone switch, you can either configure up to 63 or 64 10-Gigabit Ethernet ports on ports **Q0** through **Q15**, depending on the Junos OS release running on the switch. See the Channelizing Interfaces topic for your specific Junos OS release for further details.

You can use SFP transceivers to connect the QFX3600 device to a management network, or the control plane and management network of a QFabric system. The 1000BASE-SX Gigabit Ethernet SFP module

(QFX-SFP-1GE-SX) is supported in the SFP management ports labeled **C0S** and **C1S**. The QFX3600 device also has two 1000BASE-T RJ-45 management ports (labeled **C0** and **C1**), which can be used to connect the QFX3600 device to a management network, or the control plane and management network of a QFabric system.

NOTE: The QSFP+ DAC cables consist of a cable assembly terminated with QSFP+ transceivers on either end. If you use the QSFP+ DAC cable as the data plane connection between a QFX3600-I Interconnect device and a QFX3600 or QFX3500 Node device, the interface is automatically configured to operate at 40 Gbps. If you use the QSFP+ DAC cable to interconnect a QFX3600 Node device with another device, the interface is automatically configured to operate as four 10-Gigabit Ethernet interfaces over one cable.

You can find information about the optical transceivers supported on your Juniper device by using the Hardware Compatibility Tool. In addition to transceiver and connection type, the optical and cable characteristics—where applicable—are documented for each transceiver. The Hardware Compatibility Tool enables you to search by product, displaying all the transceivers supported on that device, or category, by interface speed or type. The list of supported transceivers for the QFX3600 and QFX3600-I Interconnect device is located at <https://pathfinder.juniper.net/hct/product/#prd=QFX3600>.



CAUTION: If you face a problem running a Juniper Networks device that uses a third-party optic or cable, the Juniper Networks Technical Assistance Center (JTAC) can help you diagnose the source of the problem. Your JTAC engineer might recommend that you check the third-party optic or cable and potentially replace it with an equivalent Juniper Networks optic or cable that is qualified for the device.

RELATED DOCUMENTATION

Front Panel of a QFX3600 Device

Rear Panel of QFX3600 and QFX3600-I Devices

Determining Interface Support for the QFX3500 Device


The 48 small form-factor pluggable plus (SFP+) access ports in the QFX3500 device support SFP and SFP+ transceivers, as well as SFP+ direct-attach copper (DAC) cables.

The four quad small form-factor pluggable plus (QSFP+) uplink ports in the QFX3500 device support QSFP+ transceivers, as well as QSFP+ DAC and DAC breakout cables. When the QFX3500 device is operating as a standalone switch, each QSFP+ port can be configured to operate as 10-Gigabit Ethernet interfaces or a single 40-Gigabit Ethernet interface. By default, the uplink ports on a standalone switch are configured as 10-Gigabit Ethernet interfaces.

NOTE: Ports **Q1** through **Q3** support four of the 10-Gigabit Ethernet interfaces. Together, the three QSFP+ ports provide up to 12 10-Gigabit Ethernet interfaces. Port **Q0** also supports 10-Gigabit Ethernet interfaces, but the number of interfaces supported depends on the release of Junos OS running on the switch. Some Junos OS releases support three 10-Gigabit Ethernet interfaces, others support four. See the topic on Channelizing Interface for your specific release.

Figure 18 on page 139 shows the uplink ports and lists the interfaces created on each port when they are configured as 10-Gigabit Ethernet interfaces.

Figure 18: QSFP+ Uplink Port Locations



*Port availability is release dependent.
See the topic on Channelizing Interfaces for your Junos Release.

If your QFX3500 device has an SFP management board instead of the 1000BASE-T RJ-45 management board, you can use the SFP transceivers to connect the QFX3500 device to a management network, or the control plane and management network of a QFabric system. The 1000BASE-SX Gigabit Ethernet SFP module (QFX-SFP-1GE-SX) is supported in the SFP management ports in the following software releases:

- Junos OS 12.1X49-D1 or later for the QFX3500 device as a standalone switch
- Junos OS 12.2X50-D10 or later for the QFX3500 Node device in a QFabric system

You can find information about the optical transceivers supported on your Juniper device by using the Hardware Compatibility Tool. In addition to transceiver and connection type, the optical and cable characteristics—where applicable—are documented for each transceiver. The Hardware Compatibility Tool enables you to search by product, displaying all the transceivers supported on that device, or category, by interface speed or type. The list of supported transceivers for the QFX3500 is located at

<https://pathfinder.juniper.net/hct/product/#prd=QFX3500>.



CAUTION: If you face a problem running a Juniper Networks device that uses a third-party optic or cable, the Juniper Networks Technical Assistance Center (JTAC) can help you diagnose the source of the problem. Your JTAC engineer might recommend that you check the third-party optic or cable and potentially replace it with an equivalent Juniper Networks optic or cable that is qualified for the device.

RELATED DOCUMENTATION

| *Front Panel of a QFX3500 Device*

Determining Interface Support for the QFX5100 Device

All product SKUs of the QFX5100 supply quad small form-factor pluggable plus (QSFP+) ports for use as uplinks, as access ports, or as Virtual Chassis ports (VCPs). These 40 GbE ports support QSFP+ transceivers, QSFP+ direct-attach copper (DAC) cables, and DAC breakout cables (DACBO). The QFX5100-48S has 6 QSFP+ ports; the QFX5100-96S has 8 QSFP+ ports; the QFX5100-24Q has 24 built-in QSFP+ ports that can all be used as uplinks. You can also add two QFX-EM-4Q expansion modules to the QFX5100-24Q for additional QSFP+ uplink ports. Each QSFP+ port on a QFX5100-24Q can be configured to operate as 10-Gigabit Ethernet interface by using a breakout cable or as a single 40-Gigabit Ethernet interface. See [“Configuring the QSFP+ Port Type on QFX5100 Devices” on page 414](#) for more information.

On all QFX5100 product SKUs, the ports are enabled by default and the default config adds the ports to the default VLAN.

Downlink ports are product SKU-specific:

- QFX5100-96S—has 96 small form-factor pluggable plus (SFP+) ports that support SFP and SFP+ transceivers, as well as DAC cables.
- QFX5100-48S—has 48 SFP+ ports that support SFP and SFP+ transceivers, as well as DAC cables.
- QFX5100-48T—has 6 QSFP+ uplink ports.
- QFX5100-24Q—has 24 QSFP+ access ports that can be configured to operate as 10-Gigabit Ethernet interfaces or as a single 40-Gigabit Ethernet interface.
- QFX5100-24Q-AA—has 24 QSFP+ access ports that can be configured to operate as 10-Gigabit Ethernet interfaces or as a single 40-Gigabit Ethernet interface.

[Figure 19 on page 141](#) shows the location of SFP+ and QSFP+ ports for the QFX5100-96S, [Figure 20 on page 141](#) shows these ports for the QFX5100-48S device, [Figure 21 on page 141](#) shows the

RJ45 and QSFP+ ports for the QFX5100-48T device, and [Figure 22 on page 142](#) shows the location of QSFP+ ports for the QFX5100-24Q device.

Figure 19: Port Panel QFX5100-96S Device




Figure 20: Port Panel QFX5100-48S Device

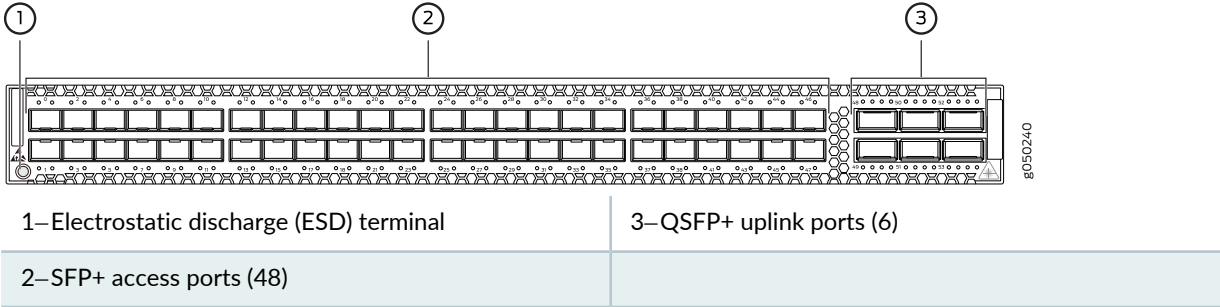


Figure 21: Port Panel QFX5100-48T Device




Figure 22: Port Panel QFX5100-24Q Device





Figure 23: Port Panel QFX5100-24Q-AA Device



You can find information about the optical transceivers supported on your Juniper device by using the Hardware Compatibility Tool. In addition to transceiver and connection type, the optical and cable characteristics—where applicable—are documented for each transceiver. The Hardware Compatibility Tool enables you to search by product, displaying all the transceivers supported on that device, or category, by interface speed or type. The list of supported transceivers for the QFX5100 is located at <https://pathfinder.juniper.net/hct/product/#prd=QFX5100>.



CAUTION: If you face a problem running a Juniper Networks device that uses a third-party optic or cable, the Juniper Networks Technical Assistance Center (JTAC) can help you diagnose the source of the problem. Your JTAC engineer might recommend that you check the third-party optic or cable and potentially replace it with an equivalent Juniper Networks optic or cable that is qualified for the device.

RELATED DOCUMENTATION

Port Panel of a QFX5100-24Q Device
Port Panel of a QFX5100-24Q-AA Device

Port Panel of QFX5100-48S and QFX5100-48SH Devices

Port Panel of QFX5100-48T and QFX5100-48TH Devices

Port Panel of a QFX5100-96S Device

Determining Transceiver Support for QFabric Systems

You can find information about the optical transceivers supported on your Juniper device by using the Hardware Compatibility Tool. In addition to transceiver and connection type, the optical and cable characteristics—where applicable—are documented for each transceiver. The Hardware Compatibility Tool enables you to search by product, displaying all the transceivers supported on that device, or category, by interface speed or type. The list of supported transceivers for QFabric Systems is located at <https://pathfinder.juniper.net/hct/product/>.



CAUTION: If you face a problem running a Juniper Networks device that uses a third-party optic or cable, the Juniper Networks Technical Assistance Center (JTAC) can help you diagnose the source of the problem. Your JTAC engineer might recommend that you check the third-party optic or cable and potentially replace it with an equivalent Juniper Networks optic or cable that is qualified for the device.

RELATED DOCUMENTATION

Front Panel of a QFX3100 Director Device

Rear Cards in a QFX3008-I Interconnect Device

Port Panel of QFX5100-48S and QFX5100-48SH Devices

Port Panel of QFX5100-48T and QFX5100-48TH Devices

Port Panel of a QFX5100-24Q Device

Front Panel of a QFX3600 Device

Front Panel of a QFX3500 Device

Cable Specifications for Copper-Based Control Plane Connections for the QFabric System

QFX Series devices support using RJ-45 patch cables to interconnect the copper-based QFabric system control plane. The RJ-45 patch cables connect to 1000BASE-T ports on the network modules in the QFX3100 Director device and the management ports in the QFX3500 Node device, QFX3600 Node device, and QFX3600-I Interconnect device.

QFX5100 Node devices and the QFX5100-24Q Interconnect devices connect using a 1000BASE-T port **C1** and a 1000BASE-T SFP module (Juniper model number QFX-SFP-1GE-T) installed in the the SFP management Ethernet cage **C0**. Some QFX5100 SKUs have an additional SFP management port (second **C0**).

In the QFX3008-I Interconnect device Control Board, the RJ-45 patch cables are used with 1000BASE-T SFP modules (Juniper model number QFX-SFP-1GE-T) installed in the SFP+ ports.

NOTE: For information about the QFX-SFP-1GE-T SFP module, see the [Hardware Compatibility Tool](#).

Table 33 on page 144 lists the specifications for the cables that connect the QFabric system control plane.

Table 33: Cable Specifications for Copper-Based Control Plane Connections for the QFabric System

Port on QFX Series Device	Cable Specification	Maximum Length	Device Receptacle
<ul style="list-style-type: none"> QFX3100 Director device network module ports QFX3008-I Interconnect device Control Board management ports (with 1000BASE-T SFP module) QFX3600-I Interconnect device management (C0 and C1) ports QFX5100 Interconnect device management C0 and C1 (with 1000BASE-T SFP module) QFX3600 Node device management (C0 and C1) ports QFX3500 Node device management (C0 and C1) ports QFX5100 Node device management C0 and C1 (with 1000BASE-T SFP module) 	Category 5 cable or equivalent suitable for 1000BASE-T operation with RJ-45 connectors	328 feet (100 meters)	RJ-45

RELATED DOCUMENTATION

Connecting QFX3100 Director Devices to a Copper-Based QFX3000-G QFabric System Control Plane Network

[Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network | 245](#)

[Connecting a QFX3600-I Interconnect Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 252](#)

Connecting a QFX3500 Node Device to a Copper-Based QFX3000-G QFabric System Control Plane Network

Connecting a QFX3600 Node Device to a Copper-Based QFX3000-G QFabric System Control Plane Network

[Connecting a QFX3500 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 263](#)

[Connecting a QFX3600 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 259](#)

Connecting a QFX5100 Node Device to a Copper-Based QFX3000-G QFabric System Control Plane Network

[Connecting a QFX5100 Node Device to a QFX3600-I Interconnect Device | 301](#)

Power

IN THIS CHAPTER

- AC Power Specifications for a QFX3100 Director Device | 147
- AC Power Cord Specifications for a QFX3100 Director Device | 148
- AC Power Specifications for a QFX3600 or QFX3600-I Device | 149
- AC Power Specifications for a QFX3500 Device | 150
- AC Power Specifications for a QFX5100 Device | 151
- AC Power Cord Specifications for a QFX Series Device | 152
- AC Power Supply Specifications for EX4300 Switches | 154
- AC Power Cord Specifications for an EX4300 Switch | 155
- AC Power Cord Specifications for EX4200 Switches | 159
- DC Power Specifications for a QFX3600 or QFX3600-I Device | 161
- DC Power Specifications for a QFX3500 Device | 162
- DC Power Specifications for a QFX5100 Device | 162
- DC Power Supply Specifications for EX4300 Switches | 164

AC Power Specifications for a QFX3100 Director Device

Table 34 on page 147 describes the AC power specifications for a QFX3100 Director device.

Table 34: AC Power Specifications for a QFX3100 Director Device

Item	Specifications
AC input voltage	Operating range: 100–240 VAC
AC input line frequency	50–60 Hz
AC input current rating	<ul style="list-style-type: none">● 5 A at 100 VAC● 2 A at 240 VAC
Typical power consumption	476 W

Table 34: AC Power Specifications for a QFX3100 Director Device (continued)

Item	Specifications
Maximum power consumption	220 W

RELATED DOCUMENTATION

[AC Power Supply in a QFX3100 Director Device](#)

[AC Power Cord Specifications for a QFX3100 Director Device](#) | 148

AC Power Cord Specifications for a QFX3100 Director Device

Detachable AC power cords are supplied with the QFX3100 Director device. The coupler is type C13 as described by International Electrotechnical Commission (IEC) standard 60320. The plug at the male end of the power cord fits into the power source outlet that is standard for your geographical location.



CAUTION: The supplied AC power cord for the switches is intended for use with the QFX3100 Director device only and not for any other use.

NOTE: In North America, AC power cords must not exceed 4.5 meters (approximately 14.75 feet) in length, to comply with National Electrical Code (NEC) Sections 400-8 (NFPA 75, 5-2.2) and 210-52 and Canadian Electrical Code (CEC) Section 4-010(3). The cords supplied with the QFX3100 Director device are in compliance.

Table 35 on page 148 lists AC power cord specifications provided for each country or region.

Table 35: AC Power Cord Specifications for a QFX3100 Director Device


Country or Region	Electrical Specifications	Plug Standards
Australia	250 VAC, 10 A, 50 Hz	AS/NZ 3112-1993
China	250 VAC, 10 A, 50 Hz	GB2099.1 1996 and GB1002 1996 (CH1-10P)

Table 35: AC Power Cord Specifications for a QFX3100 Director Device (continued)

Country or Region	Electrical Specifications	Plug Standards
Europe (except Italy, Switzerland, and United Kingdom)	250 VAC, 10 A, 50 Hz	CEE (7) VII
Italy	250 VAC, 10 A, 50 Hz	CEI 23-16/VII
Japan	125 VAC, 12 A, 50 Hz or 60 Hz	JIS 8303
Korea	250 VAC, 10A, 50 Hz	CEE 7/4
North America	125 VAC, 13 A, 60 Hz	NEMA 5-15
Switzerland	250 VAC, 10A, 50 Hz	SEV 1011 SEV 6534/2
United Kingdom	250 VAC, 10 A, 50 Hz	BS 1363A

Figure 24 on page 149 illustrates the plug on the power cord for some of the countries or regions listed in Table 35 on page 148.

Figure 24: AC Plug Types



RELATED DOCUMENTATION

- AC Power Supply in a QFX3100 Director Device
- General Safety Guidelines and Warnings
- General Electrical Safety Guidelines and Warnings
- Prevention of Electrostatic Discharge Damage

AC Power Specifications for a QFX3600 or QFX3600-I Device

Table 36 on page 150 describes the AC power specifications for a QFX3600 or QFX3600-I device.

Table 36: AC Power Specifications for a QFX3600 or QFX3600-I Device

Item	Specification
AC input voltage	Operating range: <ul style="list-style-type: none"> • 100–240 VAC
AC input line frequency	50–60 Hz
AC input current rating	<ul style="list-style-type: none"> • 4 A at 100VAC • 2 A at 240 VAC
Typical power consumption	255 W
Maximum power consumption	345 W

RELATED DOCUMENTATION

[AC Power Cord Specifications for a QFX Series Device | 152](#)
[AC Power Supply for a QFX3500, QFX3600, or QFX3600-I Device](#)
[General Safety Guidelines and Warnings](#)
[General Electrical Safety Guidelines and Warnings](#)

AC Power Specifications for a QFX3500 Device

Table 37 on page 150 describes the AC power specifications for a QFX3500 device.

Table 37: AC Power Specifications for a QFX3500 Device

Item	Specification
AC input voltage	Operating range: <ul style="list-style-type: none"> • 100–127 VAC • 200–240 VAC
AC input line frequency	50–60 Hz
AC input current rating	<ul style="list-style-type: none"> • 7.8 A at 100–127 VAC • 3.8 A at 200–240 VAC

Table 37: AC Power Specifications for a QFX3500 Device (*continued*)

Item	Specification
Typical power consumption	230 W
Maximum power consumption	365 W

RELATED DOCUMENTATION

[AC Power Cord Specifications for a QFX Series Device | 152](#)
[AC Power Supply for a QFX3500, QFX3600, or QFX3600-I Device](#)
[General Safety Guidelines and Warnings](#)
[General Electrical Safety Guidelines and Warnings](#)

AC Power Specifications for a QFX5100 Device

Table 38 on page 151 describes the AC power specifications for a QFX5100 device.

Table 38: AC Power Specifications for a QFX5100 Device

Item	Specification
AC input voltage	Operating range: <ul style="list-style-type: none"> • 100 / 240 VAC
AC input line frequency	50–60 Hz (all product SKUs)
AC input current rating	<ul style="list-style-type: none"> • 4.5 A at 100–120 VAC • 2.0 A at 200–240 VAC
Typical power consumption	
QFX5100-24Q	230 W
QFX5100-48S and QFX5100-48SH	230 W
QFX5100-48T and QFX5100-48TH	322 W
QFX5100-96S	315 W

Table 38: AC Power Specifications for a QFX5100 Device *(continued)*

Item	Specification
Maximum power consumption	
QFX5100-24Q	365 W
QFX5100-48S and QFX5100-48SH	365 W
QFX5100-48T and QFX5100-48TH	395 W
QFX5100-96S	470 W

RELATED DOCUMENTATION

[AC Power Cord Specifications for a QFX Series Device | 152](#)

AC Power Supply for a QFX5100 Device

General Safety Guidelines and Warnings

General Electrical Safety Guidelines and Warnings

AC Power Cord Specifications for a QFX Series Device

Detachable AC power cords are shipped with the chassis, if you include them as part of your order. The coupler is type C13 as described by International Electrotechnical Commission (IEC) standard 60320. The plug at the male end of the power cord fits into the power source outlet that is standard for your geographical location.

NOTE: In North America, AC power cords must not exceed 14.75 feet (approximately 4.5 meters) in length, to comply with National Electrical Code (NEC) Sections 400-8 (NFPA 75, 5-2.2) and 210-52, and Canadian Electrical Code (CEC) Section 4-010(3). The cords that can be ordered for the QFX Series switches are in compliance.

Table 39 on page 153 lists AC power cord specifications provided for each country or region.

Table 39: AC Power Cord Specifications

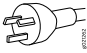
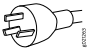







Country/Region	Electrical Specifications	Plug Standards	Shipped Juniper Model Number	Spare Juniper Model Number	Graphic
Australia	250 VAC, 10 A, 50 Hz	AS/NZ 3109-1996	CG_CBL-C13-06-AU	CBL-EX-PWR-C13-AU	
China	250 VAC, 10 A, 50 Hz	GB 1002-1996	CG_CBL-C13-06-CH	CBL-EX-PWR-C13-CH	
Europe (except Italy, Switzerland, and United Kingdom)	250 VAC, 10 A, 50 Hz	CEE (7) VII	CG_CBL-C13-06-EU	CBL-EX-PWR-C13-EU	
Italy	250 VAC, 10 A, 50 Hz	CEI 23-16/VII	CG_CBL-C13-06-IT	CBL-EX-PWR-C13-IT	
Japan	125 VAC, 12 A, 50 Hz or 60 Hz	JIS C8303	CG_CBL-C13-06-JP	CBL-EX-PWR-C13-JP	
North America	(QFX10002-36Q, QFX10002-72Q only) 125 VAC, 13 A, 60 Hz (all models) 250 VAC, 13 A, 60 Hz (all models) 250 VAC, 13 A, 60 Hz	CAN/CSA No. 49-92 NEMA L6-15 NEMA 6-15	CG_CBL-C13-06-US	CBL-EX-PWR-C13-US CBL-PW-C13-250-US CBL-PWR-C13-250-US	
South Korea	250 VAC, 10 A, 60 Hz 250 VAC, 13 A, 60 Hz	KSC 8305; K60884-1	CG_CBL-C13-06-KR	CBL-EX-PWR-C13-KR	
Switzerland	250 VAC, 10 A, 50 Hz	SEV 1011 SEV 1991; EN 60320 C13	CG_CBL-C13-06-SZ	CBL-EX-PWR-C13-SZ	
United Kingdom		BS 1363/A	CG_CBL-C13-06-UK	CBL-EX-PWR-C13-UK	

Table 39: AC Power Cord Specifications (*continued*)

Country/Region	Electrical Specifications	Plug Standards	Shipped Juniper Model Number	Spare Juniper Model Number	Graphic
	250 VAC, 10 A, 50 Hz				

AC Power Supply Specifications for EX4300 Switches

EX4300 switches support 350 W, 715 W, and 1100 W AC power supplies.

The tables in this topic provides power supply specification of AC power supplies used in an EX4300 switch:

- [Table 40 on page 154](#): 350 W AC power supply specifications
- [Table 41 on page 155](#): 715 W AC power supply specifications
- [Table 42 on page 155](#): 1100 W AC power supply specifications

Table 40: 350 W AC Power Supply Specifications for an EX4300 Switch

Item	Specification
AC input voltage	<ul style="list-style-type: none"> • Low-voltage line: 100–120 VAC • High-voltage line: 200–240 VAC
AC input line frequency	50–60 Hz
AC input current rating	<ul style="list-style-type: none"> • Low-voltage line: 4 A • High-voltage line: 2 A
Output power	350 W

NOTE: The 32-port EX4300 switches support only 350 W AC power supplies with front-to-back airflow direction.

Table 41: 715 W AC Power Supply Specifications for an EX4300 Switch

Item	Specification
AC input voltage	<ul style="list-style-type: none"> • Low-voltage line: 100–120 VAC • High-voltage line: 200–240 VAC
AC input line frequency	50–60 Hz
AC input current rating	<ul style="list-style-type: none"> • Low-voltage line: 11 A • High-voltage line: 5 A
Output power	715 W

Table 42: 1100 W AC Power Supply Specifications for an EX4300 Switch

Item	Specification
AC input voltage	<ul style="list-style-type: none"> • Low-voltage line: 115–120 VAC • High-voltage line: 200–240 VAC
AC input line frequency	50–60 Hz
AC input current rating	<ul style="list-style-type: none"> • Low-voltage line: 12 A • High-voltage line: 6 A
Output power	1100 W

RELATED DOCUMENTATION

AC Power Supply in EX4300 Switches

AC Power Supply LEDs in EX4300 Switches

AC Power Cord Specifications for an EX4300 Switch

Each AC power supply has a single AC appliance inlet that requires a dedicated AC power feed. A detachable AC power cord is supplied with each AC power supply. The 350 W AC and the 715 W AC power supplies are shipped with AC power cords with the C13 coupler type and the 1100 W AC power supplies and 1400 W AC power supplies are shipped with AC power cord with the C15 coupler type as described by International Electrotechnical Commission (IEC) standard 60320. The plug at the male end of the power cord fits into the power source outlet that is standard for your geographical location.

NOTE: In North America, AC power cords must not exceed 14.75 ft (4.5 m) in length, to comply with National Electrical Code (NEC) Section 400-8 (NFPA 75, 5-2.2) and Canadian Electrical Code (CEC) Section 4-010(3).

The tables in this topic list the AC power cords specifications provided for different power supplies for each country or region.

- [Table 43 on page 156](#)—Power cords for 350 W AC for EX4300 switches except EX4300-48MP and EX4300-48MP-S switches and 715 W AC power supplies for EX4300 switches
- [Table 44 on page 158](#)—Power cords for 1100 W AC power supplies for EX4300 switches and 1400 W AC power supplies for EX4300-48MP and EX4300-48MP-S Switches

Table 43: AC Power Cord Specifications for 350 W Power Supplies for EX4300 Switches Except EX4300-48MP and EX4300-48MP-S Switches and 715 W AC Power Supplies for EX4300 Switches

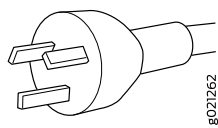
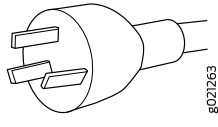
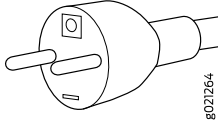
Country/ Region	Electrical Specifications	Plug Standards	Juniper Model Number	Graphic
Argentina	250 VAC, 10 A, 50 Hz	IRAM 2073 Type RA/3	CBL-EX-PWR-C13-AR	No graphic available
Australia	250 VAC, 10 A, 50 Hz	AS/NZS 3112 Type SAA/3	CBL-EX-PWR-C13-AU	
Brazil	250 VAC, 10 A, 50 Hz	NBR 14136 Type BR/3	CBL-EX-PWR-C13-BR	No graphic available
China	250 VAC, 10 A, 50 Hz	GB 1002-1996 Type PRC/3	CBL-EX-PWR-C13-CH	
Europe (except Italy, Switzerland, and United Kingdom)	250 VAC, 10 A, 50 Hz	CEE (7) VII Type VIIG	CBL-EX-PWR-C13-EU	
India	250 VAC, 10 A, 50 Hz	IS 1293 Type IND/3	CBL-EX-PWR-C13-IN	No graphic available

Table 43: AC Power Cord Specifications for 350 W Power Supplies for EX4300 Switches Except EX4300-48MP and EX4300-48MP-S Switches and 715 W AC Power Supplies for EX4300 Switches (continued)

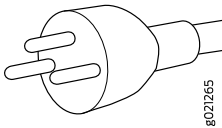
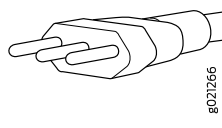
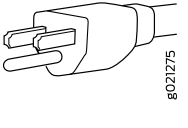
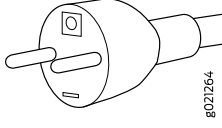
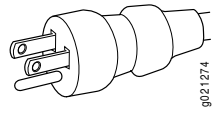
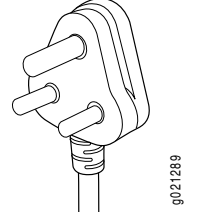
Country/ Region	Electrical Specifications	Plug Standards	Juniper Model Number	Graphic
Israel	250 VAC, 10 A, 50 Hz	SI 32/1971 Type IL/3G	CBL-EX-PWR-C13-IL	
Italy	250 VAC, 10 A, 50 Hz	CEI 23-16 Type I/3G	CBL-EX-PWR-C13-IT	
Japan	125 VAC, 12 A, 50 Hz or 60 Hz	SS-00259 Type VCTF	CBL-EX-PWR-C13-JP	
Korea	250 VAC, 10 A, 50 Hz or 60 Hz	CEE (7) VII Type VIIGK	CBL-EX-PWR-C13-KR	
North America	125 VAC, 13 A, 60 Hz	NEMA 5-15 Type N5-15	CBL-EX-PWR-C13-US	
South Africa	250 VAC, 10 A, 50 Hz	SABS 164/1:1992 Type ZA/3	CBL-EX-PWR-C13-SA	
Switzerland	250 VAC, 10 A, 50 Hz	SEV 6534-2 Type 12G	CBL-EX-PWR-C13-SZ	No graphic available

Table 43: AC Power Cord Specifications for 350 W Power Supplies for EX4300 Switches Except EX4300-48MP and EX4300-48MP-S Switches and 715 W AC Power Supplies for EX4300 Switches (continued)

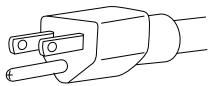
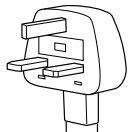
Country/ Region	Electrical Specifications	Plug Standards	Juniper Model Number	Graphic
Taiwan	125 VAC, 10 A, 50 Hz	NEMA 5-15P Type N5-15P	CBL-EX-PWR-C13-TW	 9021288
United Kingdom	250 VAC, 10 A, 50 Hz	BS 1363/A Type BS89/13	CBL-EX-PWR-C13-UK	 8021271

Table 44: AC Power Cord Specifications for 1100 W AC Power Supplies for EX4300 Switches and 1400 W AC Power Supplies for EX4300-48MP and EX4300-48MP-S Switches

Country/Region	Electrical Specifications	Plug Standards	Juniper Model Number
Argentina	250 VAC, 10 A, 50 Hz	IRAM 2073 Type RA/3	CBL-PWR-C15M-HITEMP-AR
Australia	250 VAC, 10 A, 50 Hz	AS/NZS 3112-2000 Type SAA/3	CBL-PWR-C15M-HITEMP-AU
Brazil	250 VAC, 10 A, 50 Hz	NBR 14136 Type BR/3	CBL-PWR-C15M-HITEMP-BR
China	250 VAC, 10 A, 50 Hz	GB2099, GB1002 Type PRC/3	CBL-PWR-C15M-HITEMP-CH
Europe (except Italy, Switzerland, and United Kingdom)	250 VAC, 10 A, 50 Hz	CEE (7) VII Type VIIG	CBL-PWR-C15M-HITEMP-EU
Israel	250 VAC, 10 A, 50 Hz	SI 32 Type IL/3G	CBL-PWR-C15M-HITEMP-IL
India	250 VAC, 10 A, 50 Hz	SABS 164/1:1992 Type ZA/3	CBL-PWR-C15M-HITEMP-IN
Italy	250 VAC, 10 A, 50 Hz	CEI 23-16 Type I/3G	CBL-PWR-C15M-HITEMP-IT
Japan	125 VAC, 15 A, 50 Hz or 60 Hz	JIS 8303 Type 498GJ	CBL-PWR-C15M-HITEMP-JP
Korea	250 VAC, 10 A, 50 Hz	CEE (7) VII Type VIIG	CBL-PWR-C15M-HITEMP-KR
South Africa	250 VAC, 10 A, 50 Hz	SABS 164/1:1992 Type ZA/3	CBL-PWR-C15M-HITEMP-SA

Table 44: AC Power Cord Specifications for 1100 W AC Power Supplies for EX4300 Switches and 1400 W AC Power Supplies for EX4300-48MP and EX4300-48MP-S Switches (continued)

Country/Region	Electrical Specifications	Plug Standards	Juniper Model Number
North America	125 VAC, 15 A, 60 Hz	NEMA 5-15 Type N5/15	CBL-PWR-C15M-HITEMP-US
Switzerland	250 VAC, 10 A, 50 Hz	SEV 1011 / 6534-2 Type 12G	CBL-PWR-C15M-HITEMP-SZ
United Kingdom	250 VAC, 10 A, 50 Hz	BS 1363/A Type BS89/13	CBL-PWR-C15M-HITEMP-UK



CAUTION: The AC power cord for the EX4300 switch is intended for use with this switch only. Do not use the cord with any other product.



CAUTION: Power cords must not block access to switch components.

RELATED DOCUMENTATION

Connecting AC Power to an EX4300 Switch

AC Power Cord Specifications for EX4200 Switches

A detachable AC power cord is supplied with the AC power supplies. The coupler is type C13 as described by International Electrotechnical Commission (IEC) standard 60320. The plug at the male end of the power cord fits into the power source outlet that is standard for your geographical location.



CAUTION: The AC power cord provided with each power supply is intended for use with that power supply only and not for any other use.

NOTE: In North America, AC power cords must not exceed 4.5 meters in length, to comply with National Electrical Code (NEC) Sections 400-8 (NFPA 75, 5-2.2) and 210-52 and Canadian Electrical Code (CEC) Section 4-010(3). The cords supplied with the switch are in compliance.

Table 45 on page 160 gives the AC power cord specifications for the countries and regions listed in the table.

Table 45: AC Power Cord Specifications


Country/Region	Electrical Specifications	Plug Standards	Juniper Model Number
Argentina	250 VAC, 10 A, 50 Hz	IRAM 2073 Type RA/3	CBL-EX-PWR-C13-AR
Australia	250 VAC, 10 A, 50 Hz	AS/NZS 3112 Type SAA/3	CBL-EX-PWR-C13-AU
Brazil	250 VAC, 10 A, 50 Hz	NBR 14136 Type BR/3	CBL-EX-PWR-C13-BR
China	250 VAC, 10 A, 50 Hz	GB 1002-1996 Type PRC/3	CBL-EX-PWR-C13-CH
Europe (except Italy, Switzerland, and United Kingdom)	250 VAC, 10 A, 50 Hz	CEE (7) VII Type VIIG	CBL-EX-PWR-C13-EU
India	250 VAC, 10 A, 50 Hz	IS 1293 Type IND/3	CBL-EX-PWR-C13-IN
Israel	250 VAC, 10 A, 50 Hz	SI 32/1971 Type IL/3G	CBL-EX-PWR-C13-IL
Italy	250 VAC, 10 A, 50 Hz	CEI 23-16 Type I/3G	CBL-EX-PWR-C13-IT
Japan	125 VAC, 12 A, 50 Hz or 60 Hz	SS-00259 Type VCTF	CBL-EX-PWR-C13-JP
Korea	250 VAC, 10 A, 50 Hz or 60 Hz	CEE (7) VII Type VIIGK	CBL-EX-PWR-C13-KR
North America	125 VAC, 13 A, 60 Hz	NEMA 5-15 Type N5-15	CBL-EX-PWR-C13-US
South Africa	250 VAC, 10 A, 50 Hz	SABS 164/1:1992 Type ZA/13	CBL-EX-PWR-C13-SA
Switzerland	250 VAC, 10 A, 50 Hz	SEV 6534-2 Type 12G	CBL-EX-PWR-C13-SZ

Table 45: AC Power Cord Specifications (continued)

Country/Region	Electrical Specifications	Plug Standards	Juniper Model Number
Taiwan	125 VAC, 11 A and 15 A, 50 Hz	NEMA 5-15P Type N5-15P	CBL-EX-PWR-C13-TW
United Kingdom	250 VAC, 10 A, 50 Hz	BS 1363/A Type BS89/13	CBL-EX-PWR-C13-UK

Figure 25 on page 161 illustrates the plug on the power cord for some of the countries or regions listed in Table 45 on page 160.

Figure 25: AC Plug Types



DC Power Specifications for a QFX3600 or QFX3600-I Device

Table 46 on page 161 describes the DC power specifications for a QFX3600 or QFX3600-I device.

Table 46: DC Power Specifications for a QFX3600 or QFX3600-I Device

Item	Specifications
DC input voltage	<ul style="list-style-type: none"> • Minimum operating voltage: -40 VDC • Nominal operating voltage: -48 VDC • Operating voltage range: -40 VDC through -72 VDC
DC input current rating	8 A maximum at nominal operating voltage (-48 VDC)
Typical power consumption	341 W
Maximum power consumption	252 W

RELATED DOCUMENTATION

DC Power Supply for a QFX3500, QFX3600, or QFX3600-I Device

DC Power Specifications for a QFX3500 Device

Table 47 on page 162 describes the DC power specifications for a QFX3500 device.

Table 47: DC Power Specifications for a QFX3500 Device

Item	Specifications
DC input voltage	<ul style="list-style-type: none">• Minimum operating voltage: -40 VDC• Nominal operating voltage: -48 VDC• Operating voltage range: -40 VDC through -72 VDC
DC input current rating	7 A maximum at nominal operating voltage (-48 VDC)
Typical power consumption	250 W
Maximum power consumption	385 W

RELATED DOCUMENTATION

DC Power Supply for a QFX3500, QFX3600, or QFX3600-I Device
DC Power Supply LEDs on a QFX3500, QFX3600, or QFX3600-I Device

DC Power Specifications for a QFX5100 Device

Table 48 on page 163 describes the DC power specifications for DC product SKUs of the QFX5100 device.

Table 48: DC Power Specifications for a QFX5100 Device

Item	Product SKUs	Specifications
DC input voltage	QFX5100-24Q QFX5100-48S QFX5100-48T	<ul style="list-style-type: none"> Rated operating voltage: -48 VDC to -60 VDC Operating voltage range: -40 VDC through -72 VDC
	QFX5100-96S	<ul style="list-style-type: none"> Rated operating voltage: VDC -48 VDC to -60 VDC Operating voltage range: -40 VDC through -72 VDC
DC input current rating	QFX5100-24Q QFX5100-48S QFX5100-48T QFX5100-96S	10 A maximum
Typical power consumption	QFX5100-48S QFX5100-48T QFX5100-24Q	300 W
	QFX5100-96S	315 W
Maximum power consumption	QFX5100-24Q QFX5100-48S QFX5100-48T	385 W
	QFX5100-96S	470 W

RELATED DOCUMENTATION

DC Power Supply in a QFX5100 Device

DC Power Supply LEDs on a QFX5100 Device

DC Power Supply Specifications for EX4300 Switches

Table 47 on page 162 lists the power supply specifications for a DC power supply used in an EX4300 switch.

Table 49: DC Power Supply Specifications for an EX4300 Switch

Item	Specifications
DC input voltage	<ul style="list-style-type: none"> Nominal operating voltage: -48 VDC Operating voltage range: -48 VDC through -60 VDC
DC input current rating	4 A maximum at nominal operating voltage (-48 VDC)
Output power	550 W

RELATED DOCUMENTATION

DC Power Supply in EX4300 Switches

DC Power Supply LEDs in EX4300 Switches

Installing a QFX3100 Director Device

IN THIS CHAPTER

- [Installing and Connecting a QFX3100 Director Device | 165](#)
- [Unpacking a QFX3100 Director Device | 166](#)
- [Mounting a QFX3100 Director Device on Two Posts in a Rack or Cabinet | 168](#)
- [Mounting a QFX3100 Director Device on Four Posts in a Rack or Cabinet | 170](#)
- [Connecting AC Power to a QFX3100 Director Device | 172](#)
- [Connecting a QFX Series Device to a Management Console | 173](#)
- [Powering On a QFX3100 Director Device | 175](#)

Installing and Connecting a QFX3100 Director Device

To install and connect a QFX3100 Director device:

1. Follow instructions in [“Unpacking a QFX3100 Director Device” on page 166](#).
2. Mount the QFX3100 Director device by following instructions appropriate for your site:
 - [“Mounting a QFX3100 Director Device on Two Posts in a Rack or Cabinet” on page 168](#) (using the mounting brackets provided)
 - [“Mounting a QFX3100 Director Device on Four Posts in a Rack or Cabinet” on page 170](#) (using the mounting brackets provided)
3. Follow instructions in [“Connecting AC Power to a QFX3100 Director Device” on page 172](#) to connect power.
4. See *QFX3000-G QFabric System Installation Overview* for information about the steps to install and configure your QFX3000-G QFabric system. See [“QFX3000-M QFabric System Installation Overview” on page 109](#) for information about the steps to install and configure your QFX3000-M QFabric system.

RELATED DOCUMENTATION

Rack Requirements for a QFX3100 Director Device

Cabinet Requirements for a QFX3100 Director Device

Clearance Requirements for Airflow and Hardware Maintenance for a QFX3100 Director Device

Unpacking a QFX3100 Director Device

The QFX3100 Director devices are shipped in a cardboard carton, secured with foam packing material. The carton also contains an accessory box and quick start instructions.



CAUTION: QFX3100 Director devices are maximally protected inside the shipping carton. Do not unpack the Director devices until you are ready to begin installation.

To unpack a QFX3100 Director device (see [Figure 26 on page 167](#)):

1. Move the shipping carton to a staging area as close to the installation site as possible but where you have enough room to remove the system components.
2. Position the carton so that the arrows are pointing up.
3. Open the top flaps on the shipping carton.
4. Remove the accessory box and verify the contents against the parts inventory.
5. Pull out the packing material holding the QFX3100 Director device in place.
6. Verify the components received against the inventory provided in [Table 50 on page 167](#).
7. Save the shipping carton and packing materials in case you need to move or ship the device later.

Figure 26: Unpacking a QFX3100 Director Device




Table 50: Inventory of Components Provided with a QFX3100 Director Device

Component	Quantity
QFX3100 Director device	1
Fan module (installed)	3
AC power supply module (installed)	2
4-port Ethernet network module (installed)	2
Hard disk drive (HDD) module (installed)	2
Mounting screws	8
Two-post rack-mount kit	1
Four-post rack-mount kit	1

RELATED DOCUMENTATION

Mounting a QFX3100 Director Device on Two Posts in a Rack or Cabinet

You can mount a QFX3100 Director device on two posts of a 19-in. rack or cabinet by using the mounting brackets provided with the device. (The remainder of this topic uses “rack” to mean “rack or cabinet.”)

You can mount the QFX3100 Director device on four posts of a four-post rack by using the side rail brackets provided with the device. See [“Mounting a QFX3100 Director Device on Four Posts in a Rack or Cabinet” on page 170.](#)

Before mounting the device on two posts in a rack:

- Verify that the site meets the requirements described in *Site Preparation Checklist for a QFX3100 Director Device*.
- Place the rack in its permanent location, allowing adequate clearance for airflow and maintenance, and secure it to the building structure.
- Read *General Safety Guidelines and Warnings*.
- Remove the device from the shipping carton (see [“Unpacking a QFX3100 Director Device” on page 166.](#))

Ensure that you have the following parts and tools available:

- Electrostatic discharge (ESD) grounding strap (provided)
- Phillips (+) screwdriver, number 2
- Four mid-mount mounting brackets and mounting screws (provided)
- Screws to secure the chassis to the rack (not provided)

NOTE: One person must be available to lift the QFX3100 Director device while another secures it to the rack.




CAUTION: If you are mounting multiple units on the rack, mount the heaviest unit at the bottom and mount the others from bottom to top in order of decreasing weight.

To mount the QFX3100 Director device on two posts in a rack:

1. Place the QFX3100 Director device on a flat, stable surface.
2. Align one mid-mount bracket to the mid-mount bracket holes near the center of the side of the QFX3100 Director device. Ensure that the bracket is aligned with the mounting holes and that the bracket face is facing the rack post.
3. Attach the mounting bracket to the QFX3100 Director device.
4. Mount the attached mounting bracket to the rack post. Tighten all screws.
5. Attach the mounting bracket on the opposite end of the same side to the QFX3100 Director device and mount it to the rack post. Tighten all screws.
6. Repeat this procedure for the mounting brackets on the opposite side of the QFX3100 Director device. Tighten all screws.
7. Ensure that the QFX3100 Director device chassis is level by verifying that all screws on one side of the rack are aligned with the screws on the other side. See [Figure 27 on page 169](#).

Figure 27: Mounting the QFX3100 Director Device on Two Posts in a Rack



Attach the front bracket to the chassis, and secure the chassis to the post. Attach the rear bracket to the other side of the post, and secure the chassis to the rear bracket, adjusting the bracket width as needed.

RELATED DOCUMENTATION

Rack-Mounting and Cabinet-Mounting Warnings

[Installing and Connecting a QFX3100 Director Device | 165](#)

[Connecting AC Power to a QFX3100 Director Device | 172](#)

Mounting a QFX3100 Director Device on Four Posts in a Rack or Cabinet

You can mount a QFX3100 Director device on four posts of a 19-in. rack or cabinet by using the adjustable rear mounting brackets provided. (The remainder of this topic uses “rack” to mean “rack or cabinet.”)

Before mounting the QFX3100 Director device on four posts in a rack:

- Verify that the site meets the requirements described in *Site Preparation Checklist for a QFX3100 Director Device*.
- Place the rack in its permanent location, allowing adequate clearance for airflow and maintenance, and secure it to the building structure.
- Read *General Safety Guidelines and Warnings*.
- Remove the QFX3100 Director device from the shipping carton (see [“Unpacking a QFX3100 Director Device” on page 166](#)).

Ensure that you have the following parts and tools available:

- Electrostatic discharge (ESD) grounding strap (provided).
- Phillips (+) screwdriver, number 2.
- Screws to secure the chassis and mounting brackets to the rack (not provided).
- One pair of adjustable rear mounting brackets (provided). These mounting brackets support the rear of the chassis, and must be installed.
- Screws to attach the mounting brackets to the chassis (provided).



CAUTION: If you are mounting multiple units on a rack, mount the heaviest unit at the bottom of the rack and mount the other units from the bottom of the rack to the top in decreasing order of the weight of the units.


To mount the QFX3100 Director device on four posts in a rack:

1. Place the QFX3100 Director device on a flat, stable surface.
2. Measure the distance between the front and rear rack rails. Using this measurement, attach the adjustable rear mounting brackets on the chassis using the provided screws.
3. Flip the hinged rack-mounting plates at the end of the brackets outward.

NOTE: The device weighs approximately 41.2 lb (18.73 kg). Installing the QFX3100 Director device in a rack or cabinet requires one person to lift it and a second person to secure it to the rack.

4. Have one person grasp both sides of the device, lift it, and position it in the rack, aligning the bracket holes with the holes in the rack.
5. Have a second person install a mounting screw—and cage nut and washer if your rack requires them—in each of the four bracket holes to secure the device to the front rack rails.
6. While still supporting the chassis, have the second person install a mounting screw—and cage nut and washer if your rack requires them—in each of the four bracket holes on the adjustable rear mounting brackets to secure the device to the rear rack rails.
7. Ensure that the chassis is level by verifying that all the screws on the front of the rack are aligned with the screws at the back of the rack. See [Figure 28 on page 171](#).

Figure 28: Mounting a QFX3100 Director Device on Four Posts in a Rack or Cabinet



RELATED DOCUMENTATION

Connecting Earth Ground to an EX Series Switch

[Connecting AC Power to a QFX3100 Director Device | 172](#)

[Installing and Connecting a QFX3100 Director Device | 165](#)

Connecting AC Power to a QFX3100 Director Device

The power supply in a QFX3100 Director device is a hot-removable and hot-insertable field-replaceable unit (FRU) located on the far right side of the rear panel. You can remove and replace a single power supply without powering off the QFX3100 Director device or disrupting QFX3100 Director device functions.

Before you begin connecting AC power to a QFX3100 Director device:

- Install the power supply in the chassis. See *Installing a Power Supply in a QFX3100 Director Device (includes video)*.

NOTE: Each power supply must be connected to a dedicated power source outlet to ensure power supply redundancy.

Ensure that you have the following parts and tools available:

- A power cord appropriate for your geographical location

To connect AC power to a QFX3100 Director device (see [Figure 29 on page 173](#)):


1. Ensure that the power supplies are fully inserted in the QFX3100 Director device.
2. Locate the power cords shipped with the QFX3100 Director device; the cords have plugs appropriate for your geographical location. See [“AC Power Cord Specifications for a QFX3100 Director Device” on page 148](#).



WARNING: Ensure that the power cord does not block access to QFX3100 Director device components or drape where people can trip on it.

3. Insert the coupler end of the power cord into the AC power cord inlet on the AC power supply faceplate (see [Figure 29 on page 173](#)).

Figure 29: Connecting an AC Power Cord to an AC Power Supply in a QFX3100 Director Device



4. If the AC power source outlet has a power switch, set it to the OFF (O) position.
5. Insert the power cord plug into an AC power source outlet.
6. If the AC power source outlet has a power switch, set it to the ON (I) position.
7. Repeat these steps for the second AC power supply.
8. Press the power switch on the rear panel of the QFX3100 Director device to power on the device.

NOTE: Momentarily pressing the power switch causes the system to power on or causes a power event to the operating system, which causes a graceful shutdown. Pressing the power switch for 4 seconds or longer causes an abrupt power shutdown.

9. Verify that the power LED on the power supply is lit and is on steadily.

RELATED DOCUMENTATION

[Installing and Connecting a QFX3100 Director Device | 165](#)

AC Power Supply in a QFX3100 Director Device

Connecting a QFX Series Device to a Management Console

The QFX Series has a console port with an RJ-45 connector. Use the console port to connect the device to a management console or to a console server.

Ensure that you have an RJ-45 to DB-9 rollover cable available. An RJ-45 cable with an RJ-45 to DB-9 adapter is provided with the device.

NOTE: If your laptop or PC does not have a DB-9 male connector pin and you want to connect your laptop or PC directly to the QFX Series, use a combination of the RJ-45 cable and RJ-45 to DB-9 adapter supplied with the device and a USB to DB-9 male adapter. You must provide the USB to DB-9 male adapter.


To connect the QFX Series to a management console (see [Figure 30 on page 174](#) and [Figure 31 on page 174](#)):

1. Connect one end of the Ethernet cable to the console port (labeled **CON**).
2. Connect the other end of the Ethernet cable into the console server (see [Figure 30 on page 174](#)) or management console (see [Figure 31 on page 174](#)).

Figure 30: Connecting the QFX Series to a Management Console Through a Console Server



Figure 31: Connecting the QFX Series Directly to a Management Console



RELATED DOCUMENTATION

Console Port Connector Pinout Information

Configuring Junos OS to Set Console and Auxiliary Port Properties

Powering On a QFX3100 Director Device

Before you power on the QFX3100 Director device, ensure that:

- All required QFX3100 Director device components are installed.
- You understand how to protect the QFX3100 Director device from electrostatic damage. See *Prevention of Electrostatic Discharge Damage*.

Ensure that you have the following parts and tools available to power on the QFX3100 Director device:

- An external management device such as a PC to monitor the startup process—For connecting a management device to the console port, see [“Connecting a QFX Series Device to a Management Console” on page 173](#). For connecting a management device to the management port, see [“Connecting a QFX3100 Director Device to a Network for Out-of-Band Management” on page 247](#).

NOTE: You cannot use the management (**MGMT**) port to perform the initial configuration of the QFX3100 Director device. You must configure the management ports before you can successfully connect to the QFX3100 Director device using these ports. See [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 363](#).

When you power on the QFX3100 Director devices in the Director group for the first time the first device to be powered on assumes the Director Group 0 (*dg0*) role. The second device to be powered on assumes the Director Group 1 (*dg1*) role.

To power on the QFX3100 Director device:

1. Ensure that the power supplies are fully inserted in the QFX3100 Director device and that each of their handles is flush against the faceplate.
2. Ensure that the source power cord is inserted securely into the appliance inlet for each AC power supply.
3. Switch on the site circuit breakers.
4. Press the power switch on the rear panel of the QFX3100 Director device to power on the device.

NOTE: Pressing the power switch momentarily either causes the system to power on or causes a graceful shutdown. Pressing the power switch for 4 seconds or longer causes an abrupt power shutdown.

5. Observe the power supply faceplate LEDs. If the power supply is installed correctly and functioning normally, the AC power supply LED is green.
6. On the external management device, monitor the startup process to ensure that the system boots properly.

NOTE: After you power on a power supply, wait for at least 60 seconds before you turn it off. After you power off a power supply, wait for at least 60 seconds before you turn it back on.

RELATED DOCUMENTATION

Powering Off a QFX3100 Director Device

AC Power Supply in a QFX3100 Director Device

[AC Power Cord Specifications for a QFX3100 Director Device](#) | 148

Installing a QFX3600-I Interconnect Device or QFX3600 Node Device

IN THIS CHAPTER

- [Installing and Connecting a QFX3600 or QFX3600-I Device | 177](#)
- [Unpacking a QFX3600 or QFX3600-I Device | 178](#)
- [Mounting a QFX3600 or QFX3600-I Device on Two Posts in a Rack or Cabinet | 180](#)
- [Mounting a QFX3600 or QFX3600-I Device on Four Posts in a Rack or Cabinet | 182](#)
- [Connecting Earth Ground to QFX3600 or QFX3600-I Devices | 187](#)
- [Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device | 188](#)
- [Connecting DC Power to a QFX3500, QFX3600, or QFX3600-I Device | 191](#)
- [Connecting a QFX Series Device to a Management Console | 196](#)

Installing and Connecting a QFX3600 or QFX3600-I Device

To install and connect a QFX3600 or QFX3600-I device:

1. Follow the instructions in [“Unpacking a QFX3600 or QFX3600-I Device” on page 178](#).
2. Mount the device by following the instructions appropriate for your site:
 - [Mounting a QFX3600 or QFX3600-I Device on Two Posts in a Rack or Cabinet on page 180](#)
 - [Mounting a QFX3600 or QFX3600-I Device on Four Posts in a Rack or Cabinet on page 182](#)
3. Follow the instructions in [“Connecting Earth Ground to QFX3600 or QFX3600-I Devices” on page 187](#).
4. Follow the instructions for connecting power as appropriate for your site:
 - [Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device on page 188](#)
 - [Connecting DC Power to a QFX3500, QFX3600, or QFX3600-I Device on page 191](#)
5. Depending on how you will be using the QFX3600 or QFX3600-I device, take one of the following actions:

- If you are using the QFX3600 device as a standalone switch, follow the instructions in *Configuring a QFX3600 Device as a Standalone Switch*.
- If you are using the QFX3600 device as a Node device in a QFX3000-G QFabric system, see *QFX3000-G QFabric System Installation Overview* for information about the steps to install and configure your QFX3000-G QFabric system.
- If you are using the QFX3600 device as a Node device in a QFX3000-M QFabric system, see [“QFX3000-M QFabric System Installation Overview” on page 109](#) for information about the steps to install and configure your QFX3000-M QFabric system.
- If you are using the QFX3600-I device as an Interconnect device in a QFX3000-M QFabric system, see [“QFX3000-M QFabric System Installation Overview” on page 109](#) for information about the steps to install and configure your QFX3000-M QFabric system.

RELATED DOCUMENTATION

Rack Requirements for a QFX3600 or QFX3600-I Device

Cabinet Requirements for a QFX3600 or QFX3600-I Device

Clearance Requirements for Airflow and Hardware Maintenance for a QFX3600 or QFX3600-I Device

Unpacking a QFX3600 or QFX3600-I Device

The QFX3600 or QFX3600-I device chassis is a rigid sheet-metal structure that houses the hardware components. QFX3600 and QFX3600-I devices are shipped in a cardboard carton, secured with foam packing material. The carton also contains an accessory box and quick start instructions.



CAUTION: QFX3600 and QFX3600-I devices are maximally protected inside the shipping carton. Do not unpack the device until you are ready to begin installation.

To unpack a QFX3600 or QFX3600-I device:

1. Move the shipping carton to a staging area as close to the installation site as possible, but where you have enough room to remove the system components.
2. Position the carton so that the arrows are pointing up.
3. Open the top flaps on the shipping carton.

4. Remove the accessory box and verify the contents against the inventory included in the box.
[Table 51 on page 179](#) lists the inventory of components supplied with a QFX3600 or QFX3600-I device.
5. Pull out the packing material holding the device in place.
6. Verify the chassis components received:
 - Three fan trays
 - Two power supplies
7. Save the shipping carton and packing materials in case you need to move or ship the device later.

Table 51: Accessory Kit Part Contents

Parts	Quantity
Chassis grounding lug	1
M5 screws to attach the chassis grounding lug to the protective earth terminal on the chassis	2
Electrostatic discharge (ESD) grounding strap	1
NOTE: Use only clip-style ESD grounding straps with the chassis grounding lug.	
SFP/SFP+ port dust covers	2
QSFP+ port dust covers	16
RJ-45 cable and RJ-45 to DB-9 adapter for console port connection	1
Mounting brackets for front-mounting in a four-post rack or cabinet	2
M4 flat-head screws to attach the brackets for front-mounting in a rack or cabinet	6
Rear installation blades for front-mounting in a four-post rack or cabinet	2
Mounting brackets for front-mounting in a two-post rack or cabinet	2
Mounting brackets for mid-mounting in a two-post rack or cabinet	2
M4 pan-head screws to attach the brackets for front-mounting or mid-mounting in a rack or cabinet	6

RELATED DOCUMENTATION

Mounting a QFX3600 or QFX3600-I Device on Two Posts in a Rack or Cabinet

You can mid-mount a QFX3600 or QFX3600-I device on two posts of a 19-in. rack or cabinet by using the mounting brackets provided with the device. (The remainder of this topic uses “rack” to mean “rack or cabinet.”)

You can also mount the device on four posts of a four-post rack by using the mounting brackets provided with the device. See [“Mounting a QFX3600 or QFX3600-I Device on Four Posts in a Rack or Cabinet” on page 182.](#)

The holes in the mounting brackets are placed at 1 U (1.75 in., or 4.45 cm.) apart so that the device can be mounted in any rack that provides holes spaced at that distance.

Before mounting a QFX3600 or QFX3600-I device on two posts in a rack:

- Ensure you understand how to prevent electrostatic discharge (ESD) damage. See *Prevention of Electrostatic Discharge Damage*.
- Verify that the site meets the requirements described in *Site Preparation Checklist for a QFX3600 or QFX3600-I Device*.
- Place the rack in its permanent location, allowing adequate clearance for airflow and maintenance, and secure it to the building structure.
- Read *General Safety Guidelines and Warnings*, with particular attention to *Chassis Lifting Guidelines for a QFX3600 or QFX3600-I Device*.
- Remove the device from the shipping carton (see [“Unpacking a QFX3600 or QFX3600-I Device” on page 178.](#))

Ensure that you have the following parts and tools available:

- ESD grounding strap (provided)
- One pair of mounting brackets depending on how you want to mount the device (provided)
 - Use the front/rear mounting brackets (part number 540-038579) to front-mount or rear-mount the device.
 - Use the mid-mounting brackets (part number 540-038665) to mid-mount the device.
- 6 Phillips 4x6-mm pan-head mounting screws (provided)
- Four screws to secure the chassis to the rack (not provided)
- Appropriate screwdriver for the mounting screws (not provided)

NOTE: One person must be available to lift the device while another secures the device to the rack.




CAUTION: If you are mounting multiple device on a rack, mount a device in the bottom of the rack first and proceed to mount the rest of the devices from bottom to top.

To mount the device on two posts in a rack:

1. Attach the ESD grounding strap to your bare wrist and to a site ESD point.
2. Place the device on a flat, stable surface.
3. Align the mid-mounting brackets (part number 540-038665) with the holes on the side panels of the device chassis, such that the mounting ears are in the center of the side panels. See [Figure 32 on page 182](#).
4. Insert mounting screws into the aligned holes. Tighten the screws.
5. Have one person grasp both sides of the device, lift the device, and position it in the rack, aligning the mounting ear holes with the threaded holes in the rack or cabinet rail. Align the bottom mounting ear hole in both the mounting brackets with a hole in each rack rail, making sure the chassis is level. See [Figure 32 on page 182](#).
6. Have a second person secure the device to the rack by using the appropriate screws. Tighten the screws.
7. Ensure that the device chassis is level by verifying that all screws on one side of the rack are aligned with the screws on the other side.

Figure 32: Mounting the Device on Two Posts in a Rack



RELATED DOCUMENTATION

Rack-Mounting and Cabinet-Mounting Warnings

[Connecting Earth Ground to QFX3600 or QFX3600-I Devices | 187](#)

[Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device | 188](#)

[Connecting DC Power to a QFX3500, QFX3600, or QFX3600-I Device | 191](#)

Mounting a QFX3600 or QFX3600-I Device on Four Posts in a Rack or Cabinet

You can front-mount or rear-mount a QFX3600 or QFX3600-I device on four posts in a 19-in. rack or cabinet by using the mounting brackets and installation blades provided with the device. (The remainder of this topic uses “rack” to mean “rack or cabinet.”) The front and rear rack rails must be spaced between 19.3 in. (49 cm) and 36 in. (91.4 cm) front to back.

You can also mount the device on two posts of a 19-in. rack or cabinet by using the mounting brackets provided with the device. See [“Mounting a QFX3600 or QFX3600-I Device on Two Posts in a Rack or Cabinet” on page 180.](#)

The holes in the mounting brackets and installation blades are placed at 1 U (1.75 in., or 4.45 cm.) apart so that the device can be mounted in any rack that provides holes spaced at that distance.

Before you begin mounting a QFX3600 or QFX3600-I device on the rack or cabinet:

- Ensure you understand how to prevent electrostatic discharge (ESD) damage. See *Prevention of Electrostatic Discharge Damage*.
- Verify that the site meets the requirements described in *Site Preparation Checklist for a QFX3600 or QFX3600-I Device*.
- Place the rack in its permanent location, allowing adequate clearance for airflow and maintenance, and secure it to the building structure.
- Read “[General Site Guidelines](#)” on [page 124](#), with particular attention to *Chassis Lifting Guidelines for a QFX3600 or QFX3600-I Device*.
- Remove the device from the shipping carton (see “[Unpacking a QFX3600 or QFX3600-I Device](#)” on [page 178](#)).

Ensure that you have the following parts and tools available to mount the device on four posts in a rack:

- ESD grounding strap (provided).
- One pair of mounting brackets (part number 540-038596) (provided).
- One pair of rear installation blades (part number 540-038598). These installation blades support the rear of the chassis, and must be installed (provided).
- 6 Phillips 4x5-mm flat-head mounting screws (provided).
- Eight screws to secure the chassis and rear installation blades to the rack (not provided).
- Appropriate screwdriver for the mounting screws (not provided).



WARNING: QFX3600 and QFX3600-I devices must be supported at all four corners. Mounting the chassis using only the front brackets damages the chassis and can result in serious bodily injury.



CAUTION: If you are installing the QFX3600 or QFX3600-I device above 60 in. (152.4 cm) from the floor, you must remove the power supplies and fan trays before attempting to install the device, or ask someone to assist you during the installation.




CAUTION: If you are mounting multiple devices on a rack, mount the device in the lowest position of the rack first, and proceed to mount the rest of the devices from bottom to top.

To mount the device on four posts in a rack:

1. Attach the ESD grounding strap to your bare wrist and to a site ESD point.
2. Decide where to position the device in the rack.
3. Install the rear installation blades. See [Figure 33 on page 184](#).
 - a. With two mounting screws—and cage nuts and washers if your rack requires them—attach one of the rear installation blades to the left rear of the rack at the point where you want to mount the device. Tighten the screws.
 - b. Position the second rear installation blade at the desired position in the right rear of the rack, so that it is on the same rack level as the left rear installation blade. If the right and left rear installation blades are not on the same level, the chassis will rest at an angle in the rack instead of resting flat and level.
 - c. With two mounting screws—and cage nuts and washers if your rack requires them—attach the second rear installation blade to the right rear of the rack at the point where you want to mount the device. Tighten the screws.

Figure 33: Attaching the Installation Blades to the Rear of the Rack




4. Prepare the device for mounting.
 - a. Place the device on a flat, stable surface.
 - b. Align the mounting brackets along the front or rear of the side panels of the device chassis depending on how you want to mount the device. For example, if you want to front-mount the device, align the brackets along the front of the side panel such that the mounting ears are in the front of the device chassis. See [Figure 34 on page 185](#).
 - c. Align the holes in the mounting brackets with holes on the side panels of the device chassis.
 - d. Insert mounting screws into the aligned holes. Tighten the screws.
5. Mount the device.

- a. Grasp both sides of the device, lift it, and position it in the rack so that the rear of the chassis slides onto the installation blade. See [Figure 34 on page 185](#).

TIP: If someone is assisting you, have one person stand at the rear of the rack where the installation blade is installed, to help guide the device onto the installation blade.

- b. Align the holes in the front brackets on the chassis with the holes in the rack. Ensure that the chassis is level.
- c. With four mounting screws—and cage nuts and washers if your rack requires them—secure the front of the device to the rack. Tighten the screws.
- d. Ensure that the device chassis is level by verifying that all screws on one side of the rack are aligned with the screws on the other side.

Figure 34: Mounting the Device on Four Posts



g050214

RELATED DOCUMENTATION

Rack-Mounting and Cabinet-Mounting Warnings

[Connecting Earth Ground to QFX3600 or QFX3600-I Devices | 187](#)

[Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device | 188](#)

[Connecting DC Power to a QFX3500, QFX3600, or QFX3600-I Device | 191](#)

Configuring a QFX3600 Device as a Standalone Switch

Connecting Earth Ground to QFX3600 or QFX3600-I Devices

To meet safety and electromagnetic interference (EMI) requirements and to ensure proper operation, you must connect the QFX3600 and QFX3600-I devices to earth ground before you connect it to power.

For installations that require a separate grounding conductor to the chassis, use the protective earthing terminal on the left rear of the chassis to connect to the earth ground (see [Figure 35 on page 188](#)).

Before you connect earth ground to the protective earthing terminal of a QFX3600 or QFX3600-I device, ensure that a licensed electrician has attached an appropriate grounding lug to the grounding cable.



CAUTION: Using a grounding cable with an incorrectly attached lug can damage the switch.

NOTE: Mount your device in the rack or cabinet before attaching the grounding lug to the switch. See [“Mounting a QFX3600 or QFX3600-I Device on Two Posts in a Rack or Cabinet” on page 180](#) and [“Mounting a QFX3600 or QFX3600-I Device on Four Posts in a Rack or Cabinet” on page 182](#).

Ensure that you have the following parts and tools available:

- Grounding cable for your QFX3600 or QFX3600-I device—The grounding cable must be 14 AWG (2 mm²), minimum 90° C wire, or as permitted by the local code.
- Grounding lug for your grounding cable—The grounding lug required is a Panduit LCD10-10A-L or equivalent. This grounding lug is provided in the accessory kit.
- Two M5 screws with integrated washers—The screws are used to secure the grounding lug to the protective earthing terminal. The screws are provided in the accessory kit.
- Phillips (+) screwdriver, number 2.


An AC-powered QFX3600 or QFX3600-I device chassis gains additional grounding when you plug the power supply in the switch into a grounded AC power outlet by using an AC power cord appropriate for your geographical location. See [“AC Power Cord Specifications for a QFX Series Device” on page 152](#).

To connect earth ground to a QFX3600 or QFX3600-I device:

1. Connect one end of the grounding cable to a proper earth ground, such as the rack in which the device is mounted.
2. Place the grounding lug attached to the grounding cable over the protective earthing terminal.

3. Secure the grounding lug to the protective earthing terminal with screws.
4. Dress the grounding cable and ensure that it does not touch or block access to other switch components and that it does not drape where people could trip over it.

Figure 35: Connecting a Grounding Cable to a QFX3600 or QFX3600-I Device



RELATED DOCUMENTATION

General Safety Guidelines and Warnings

Grounded Equipment Warning

[Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device | 188](#)

[Connecting DC Power to a QFX3500, QFX3600, or QFX3600-I Device | 191](#)

Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device

The QFX3500, QFX3600, and QFX3600-I devices are shipped from the factory with two 650 W power supplies pre-installed. Each power supply is a hot-removable and hot-insertable field-replaceable unit (FRU) when the second power supply is installed and running. You can install replacement power supplies without powering off the device or disrupting the switching function.

Ensure that you have a power cord appropriate for your geographical location available to connect AC power to the device.

Before you begin connecting AC power to the device:

- Ensure that you have taken the necessary precautions to prevent electrostatic discharge (ESD) damage (see *Prevention of Electrostatic Discharge Damage*).
- Ensure that you have connected the device chassis to earth ground.



CAUTION: Before you connect power to the device, a licensed electrician must attach a cable lug to the grounding and power cables that you supply. A cable with an incorrectly attached lug can damage the device (for example, by causing a short circuit).

To meet safety and electromagnetic interference (EMI) requirements and to ensure proper operation, you must connect the chassis to earth ground before you connect it to power. For installations that require a separate grounding conductor to the chassis, use the protective earthing terminal on the device chassis to connect to the earth ground. For instructions on connecting earth ground, see [“Connecting Earth Ground to a QFX3500 Device” on page 207](#) or [“Connecting Earth Ground to QFX3600 or QFX3600-I Devices” on page 187](#). The device gains additional grounding when you plug the power supply in the device into a grounded AC power outlet by using the AC power cord appropriate for your geographical location (see [“AC Power Cord Specifications for a QFX Series Device” on page 152](#)).

- Install the power supply in the chassis. For instructions on installing a power supply in a QFX3500 device, see *Installing a Power Supply in a QFX3500 Device*. For instructions on installing a power supply in a QFX3600 or QFX3600-I device, see *Installing a Power Supply in a QFX3600 or QFX3600-I Device*.

NOTE: Each power supply must be connected to a dedicated power source outlet.

To connect AC power to a QFX3500, QFX3600, or QFX3600-I device:

1. Attach the grounding strap to your bare wrist and to a site ESD point.
2. Ensure that the power supplies are fully inserted in the chassis and the latches are secure. If only one power supply is installed, ensure a that blank cover panel is installed over the second power supply slot.
3. Locate the power cord or cords shipped with the device; the cords have plugs appropriate for your geographical location. See [“AC Power Cord Specifications for a QFX Series Device” on page 152](#).



WARNING: Ensure that the power cord does not block access to device components or drape where people can trip on it.

4. Connect each power supply to the power sources. Insert the coupler end of the power cord into the AC power cord inlet on the AC power supply faceplate.

5. Push the power cord retainer onto the power cord (see [Figure 36 on page 190](#) or [Figure 37 on page 190](#)).

Figure 36: Connecting an AC Power Cord to an AC Power Supply in a QFX3500 Device





Figure 37: Connecting an AC Power Cord to an AC Power Supply in a QFX3600 or QFX3600-I Device



6. If the AC power source outlet has a power switch, set it to the OFF (O) position.

NOTE: The device powers on as soon as power is provided to the power supply. There is no power switch on the device.

7. Insert the power cord plug into an AC power source outlet.
8. If the AC power source outlet has a power switch, set it to the ON (I) position.
9. Verify that the AC and DC LEDs on each power supply are lit green.

If the amber fault LED is lit, remove power from the power supply, and replace the power supply (see *Removing a Power Supply from a QFX3500 Device* or *Removing a Power Supply from a QFX3600 or QFX3600-I Device*). Do not remove the power supply until you have a replacement power supply ready: the power supplies or a blank cover panel must be installed in the device to ensure proper airflow.



CAUTION: Replace a failed power supply with a blank panel or new power supply within 1 minute of removal to prevent chassis overheating.

RELATED DOCUMENTATION

AC Power Supply for a QFX3500, QFX3600, or QFX3600-I Device

AC Power Supply LEDs on a QFX3500, QFX3600, or QFX3600-I Device

Connecting DC Power to a QFX3500, QFX3600, or QFX3600-I Device

The QFX3500, QFX3600, and QFX3600-I devices are shipped from the factory with two 650 W power supplies pre-installed. Each power supply is a hot-removable and hot-insertable field-replaceable unit (FRU) when the second power supply is installed and running. You can install replacement power supplies without powering off the device or disrupting the switching function.



WARNING: DC-powered QFX3500, QFX3600 and QFX3600-I devices are intended for installation only in a restricted access location.

NOTE: The battery returns of the DC power supply should be connected as an isolated DC return (DC-I).

Before you begin connecting DC power to the device:

- Ensure that you have taken the necessary precautions to prevent electrostatic discharge (ESD) damage (see *Prevention of Electrostatic Discharge Damage*).
- Ensure that you have connected the device chassis to earth ground.



CAUTION: Before you connect power to the device, a licensed electrician must attach a cable lug to the grounding and power cables that you supply. A cable with an incorrectly attached lug can damage the device (for example, by causing a short circuit).

To meet safety and electromagnetic interference (EMI) requirements and to ensure proper operation, you must connect the chassis to earth ground before you connect it to power. For installations that require a separate grounding conductor to the chassis, use the protective earthing terminal on the device chassis to connect to the earth ground. For instructions on connecting earth ground, see [“Connecting Earth Ground to a QFX3500 Device” on page 207](#) or [“Connecting Earth Ground to QFX3600 or QFX3600-I Devices” on page 187](#).

- Install the power supply in the chassis. For instructions on installing a power supply in a QFX3500 device, see *Installing a Power Supply in a QFX3500 Device*. For instructions on installing a power supply in a QFX3600 or QFX3600-I device, see *Installing a Power Supply in a QFX3600 or QFX3600-I Device*

Ensure that you have the following parts and tools available:

- DC power source cables (14–16 AWG) with ring lug (Molex 190700069 or equivalent) (not provided)
- Phillips (+) screwdriver, number 2 (not provided)
- Multimeter (not provided)

To connect DC power to a QFX3500, QFX3600 or QFX3600-I device:

1. Attach the grounding strap to your bare wrist and to a site ESD point.
2. Verify that the DC power cables are correctly labeled before making connections to the power supply. In a typical power distribution scheme where the return is connected to chassis ground at the battery plant, you can use a multimeter to verify the resistance of the –48V and RTN DC cables to chassis ground:
 - The cable with very low resistance (indicating a closed circuit) to chassis ground is positive (+) and will be installed on the V+ (return) DC power input terminal.
 - The cable with very high resistance (indicating an open circuit) to chassis ground is negative (–) and will be installed on the V– (input) DC power input terminal.



CAUTION: You must ensure that power connections maintain the proper polarity. The power source cables might be labeled (+) and (–) to indicate their polarity. There is no standard color coding for DC power cables. The color coding used by the external DC power source at your site determines the color coding for the leads on the power cables that attach to the DC power input terminals on each power supply.

3. Ensure that the input circuit breaker is open so that the voltage across the DC power source cable leads is 0 V and that the cable leads do not become active while you are connecting DC power.

NOTE: The V+ terminals are referred to as +RTN, and V– terminals are referred to as –48 V in *DC Power Wiring Sequence Warning* and *DC Power Electrical Safety Guidelines*.

4. Ensure that the power supplies are fully inserted in the chassis.
5. Remove the terminal block cover. The terminal block cover is a piece of clear plastic that snaps into place over the terminal block (see [Figure 38 on page 195](#)).
6. Remove the screws on the terminals using the screwdriver. Save the screws.



WARNING: Ensure that the power cables do not block access to device components or drape where people can trip on them.


7. Connect each power supply to the power sources. Secure power source cables to the power supplies by screwing the ring lugs attached to the cables to the appropriate terminals by using the screw from the terminals (see [Figure 39 on page 196](#) and [Figure 38 on page 195](#)).



CAUTION: The DC power supply has four terminals labeled V+, V+, V-, and V- for connecting DC power source cables labeled positive (+) and negative (-). The V+ terminals are shunted internally together, as are the V- terminals. The same polarity terminal can be wired together from the same source to provide an additional current path in a higher power chassis. Do not connect the terminals to different sources. For example, connect -48 V from DC source feed A to the input terminals of one power supply and connect -48 V from feed B to the input terminals of the second power supply on the other side of the chassis. This configuration provides the commonly deployed A/B feed redundancy for the system.

- a. Secure the ring lug of the positive (+) DC power source cable to the V+ terminal on the DC power supply.
- b. Secure the ring lug of the negative (-) DC power source cable to the V- terminal on the DC power supply.
- c. Tighten the screws on the power supply terminals until snug using the screwdriver. Do not overtighten—apply between 5 in-lb (0.56 Nm) and 6 in-lb (0.68 Nm) of torque to the screws.

Figure 38: DC Power Supply Faceplate for a QFX3500, QFX3600 or QFX3600-I Device




1—Shunt positive input terminals	5—ESD grounding point
2—Shunt negative input terminals	6—Fault LED
3—Terminal block	7—Output LED
4—Ejector lever	8—Input LED



CAUTION: The V+ terminals are shunted internally together, as are the V- terminals. The same polarity terminal can be wired together from the same source to provide an additional current path in a higher power chassis. Do not connect the terminals to different sources.

Figure 39: Securing Ring Lugs to the Terminals on the QFX3500, QFX3600 or QFX3600-I DC Power Supply



8. Replace the terminal block cover.

9. Close the input circuit breaker.

NOTE: The device powers on as soon as power is provided to the power supply. There is no power switch on the device.

10. Verify that the **IN** and **OUT** LEDs on the power supply are lit green and are on steadily.

RELATED DOCUMENTATION

DC Power Supply for a QFX3500, QFX3600, or QFX3600-I Device

DC Power Supply LEDs on a QFX3500, QFX3600, or QFX3600-I Device

Connecting a QFX Series Device to a Management Console

The QFX Series has a console port with an RJ-45 connector. Use the console port to connect the device to a management console or to a console server.

Ensure that you have an RJ-45 to DB-9 rollover cable available. An RJ-45 cable with an RJ-45 to DB-9 adapter is provided with the device.

NOTE: If your laptop or PC does not have a DB-9 male connector pin and you want to connect your laptop or PC directly to the QFX Series, use a combination of the RJ-45 cable and RJ-45 to DB-9 adapter supplied with the device and a USB to DB-9 male adapter. You must provide the USB to DB-9 male adapter.


To connect the QFX Series to a management console (see [Figure 30 on page 174](#) and [Figure 31 on page 174](#)):

1. Connect one end of the Ethernet cable to the console port (labeled **CON**).
2. Connect the other end of the Ethernet cable into the console server (see [Figure 30 on page 174](#)) or management console (see [Figure 31 on page 174](#)).

Figure 40: Connecting the QFX Series to a Management Console Through a Console Server



Figure 41: Connecting the QFX Series Directly to a Management Console



RELATED DOCUMENTATION

Console Port Connector Pinout Information

Configuring Junos OS to Set Console and Auxiliary Port Properties

Installing a QFX3500 Node Device

IN THIS CHAPTER

- [Installing and Connecting a QFX3500 Device | 199](#)
- [Unpacking a QFX3500 Device | 200](#)
- [Mounting a QFX3500 Device in a Rack or Cabinet | 201](#)
- [Connecting Earth Ground to a QFX3500 Device | 207](#)
- [Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device | 209](#)
- [Connecting DC Power to a QFX3500, QFX3600, or QFX3600-I Device | 211](#)
- [Connecting a QFX Series Device to a Management Console | 216](#)

Installing and Connecting a QFX3500 Device

To install and connect a QFX3500 device:

1. Follow the instructions in [“Unpacking a QFX3500 Device” on page 200](#).
2. Mount the device by following the instructions in [“Mounting a QFX3500 Device in a Rack or Cabinet” on page 201](#).
3. Follow the instructions in [“Connecting Earth Ground to a QFX3500 Device” on page 207](#).
4. Follow the instructions in [“Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device” on page 188](#).
5. Depending on how you will be using the QFX3500 device, do one of the following:
 - If you are using the QFX3500 device as a standalone switch, follow the instructions in *Configuring a QFX3500 Device as a Standalone Switch*.
 - If you are using the QFX3500 device as a Node device in a QFX3000-G QFabric system, see *QFX3000-G QFabric System Installation Overview* for information about the steps to install and configure your QFX3000-G QFabric system.

- If you are using the QFX3500 device as a Node device in a QFX3000-M QFabric system, see [“QFX3000-M QFabric System Installation Overview” on page 109](#) for information about the steps to install and configure your QFX3000-M QFabric system.

RELATED DOCUMENTATION

Rack Requirements for a QFX3500 Device

Cabinet Requirements for a QFX3500 Device

Clearance Requirements for Airflow and Hardware Maintenance for a QFX3500 Device

Unpacking a QFX3500 Device

The QFX3500 device chassis is a rigid sheet-metal structure that houses the hardware components. A QFX3500 device is shipped in a cardboard carton, secured with foam packing material. The carton also contains an accessory box and quick start instructions.



CAUTION: QFX3500 devices are maximally protected inside the shipping carton. Do not unpack the device until you are ready to begin installation.

To unpack a QFX3500 device:

1. Move the shipping carton to a staging area as close to the installation site as possible, but where you have enough room to remove the system components.
2. Position the carton so that the arrows are pointing up.
3. Open the top flaps on the shipping carton.
4. Remove the accessory box and verify the contents against the inventory included in the box. [Table 52 on page 201](#) lists the inventory of components supplied with a QFX3500 device.
5. Pull out the packing material holding the device in place.
6. Verify the chassis components received:
 - Management board
 - Two fan trays

- One or two power supplies, depending on your order. If only one power supply is installed, a blank panel should be installed on the second power supply slot.

7. Save the shipping carton and packing materials in case you need to move or ship the device later.

Table 52: Inventory of Components Supplied with a QFX3500 Device

Component	Quantity
Chassis with management board, two fan trays, and one or two power supplies	1
Rear installation blades	2
RJ-45 cable and RJ-45 to DB-9 adapter	1
SFP/SFP+ port dust covers	48
QSFP+ port dust covers	4
Electrostatic discharge (ESD) grounding strap	1

RELATED DOCUMENTATION

[Mounting a QFX3500 Device in a Rack or Cabinet | 201](#)

[Installing and Connecting a QFX3500 Device | 199](#)

Mounting a QFX3500 Device in a Rack or Cabinet

IN THIS SECTION

- [Before You Begin Rack Installation | 202](#)
- [Two Mounting Rails Procedure | 203](#)
- [Four Mounting Rails Procedure | 205](#)

You can mount a QFX3500 device on four posts in a 19-in. rack or cabinet by using the mounting kits provided with the device. Choose one of the following two mounting kits provided for the different QFX3500 chassis configurations.

- If your installation kit has two rails and your QFX3500 has mounting holes integrated as part of the chassis, use [“Two Mounting Rails Procedure” on page 203](#). This configuration aligns the management end of the device flush with the rack. The adjustable rails allow for installation into racks having different depths.
- If your installation kit has four rails and the QFX3500 does not have mounting holes as part of the chassis faceplate, use [“Four Mounting Rails Procedure” on page 205](#). This configuration allows either end of the device to be mounted flush with the rack and still be adjustable for racks with different depths.

(The remainder of this topic uses “rack” to mean “rack or cabinet.”) The front and rear rack rails must be spaced between 28 in. (71.1 cm) and 36 in. (91.4 cm) front to back.

Before You Begin Rack Installation

Before you begin mounting a QFX3500 device in the rack or cabinet:

- If replacing an existing QFX3500, remove previous rack-mount hardware. The mounting bracket and mounting blade in this procedure is not compatible with other Juniper mounting kits.
- Ensure that you understand how to prevent electrostatic discharge (ESD) damage. See *Prevention of Electrostatic Discharge Damage*.
- Verify that the site meets the requirements described in *Site Preparation Checklist for a QFX3500 Device*.
- Place the rack in its permanent location, allowing adequate clearance for airflow and maintenance, and secure it to the building structure.
- Read [“General Site Guidelines” on page 124](#), with particular attention to *Chassis Lifting Guidelines for a QFX3500 Device*.
- Remove the device from the shipping carton (see [“Unpacking a QFX3500 Device” on page 200](#)).

Ensure that you have the following parts and tools available to mount the device on four posts in a rack:

- ESD grounding strap (provided).
- One pair of rear installation blades. These installation blades support the rear of the chassis and must be installed (provided).
- Eight screws to secure the chassis and rear installation blades to the rack (not provided).
- Appropriate screwdriver for the mounting screws (not provided).



WARNING: The QFX3500 device must be supported at all four corners. Mounting the chassis using only the front brackets will damage the chassis and can result in serious bodily injury.



CAUTION: If you are installing the QFX3500 device above 60 in. (152.4 cm) from the floor, you must remove the power supplies, fan trays, and management board before attempting to install the device, or ask someone to assist you during the installation.




CAUTION: If you are mounting multiple devices on a rack, mount the device in the lowest position of the rack first and proceed to mount the rest of the devices from bottom to top.

Two Mounting Rails Procedure

To mount the device on four posts in a rack using a two-rail kit:

1. Attach the ESD grounding strap to your bare wrist and to a site ESD point.
2. With two mounting screws—and cage nuts and washers if your rack requires them—attach one of the rear installation blades to the rear of the rack at the point where you want to mount the device. Tighten the screws. The blade helps support the rear of the chassis. You install the second rear installation blade after securing both front mounting brackets. See [Figure 42 on page 203](#).


Figure 42: Installing an Installation Blade in a Rack



3. Grasp both sides of the device, lift it, and position it in the rack so that the blade receptacle at the rear of the chassis catches and slides onto the installation blade. See [Figure 43 on page 204](#).

TIP: If someone is assisting you, have one person stand at the rear of the rack where the installation blade is installed, to help guide the device onto the installation blade.

Figure 43: Mounting the QFX3500 Device on Four Posts in a Rack Using a Two-Rail Kit




4. Align the holes in the front brackets on the chassis with the holes in the rack. Ensure that the chassis is level.
5. With four mounting screws—and cage nuts and washers if your rack requires them—secure the front of the device to the rack. Insert the first screw on the opposite corner from the rear installation blade you installed. Tighten the screws.
6. Ensure that the device chassis is level by verifying that all the screws on the front of the rack are aligned with the screws at the back of the rack.
7. With two mounting screws—and cage nuts and washers if your rack requires them—slide the second rear mounting blade into the blade receptacle on the chassis, and secure it to the rear of the rack by tightening the screws. You might need to loosen and adjust the first mounting blade to install the second blade.

Four Mounting Rails Procedure

To mount the device on four posts in a rack using a four -rail kit:


1. Attach the ESD grounding strap to your bare wrist and to a site ESD point.
2. Decide whether the management end of the device or the ports are to be placed at the front of the rack.
3. Align the holes in the mounting rail with the button fasteners on the side of the device and slide the holes over the fasteners to click into place. See [Figure 44 on page 205](#) to see the proper alignment.

Figure 44: Attaching a Mounting Rail to the QFX3500



4. With three mounting screws—and cage nuts and washers if your rack requires them—attach the mounting rail to the device. Tighten the screws.
5. Repeats steps 3 and 4 on the opposite side of the device. One end of the device now has front facing mounting holes, the other end none.
6. With two mounting screws—and cage nuts and washers if your rack requires them—attach one of the rear installation blades to the rear of the rack at the point where you want to mount the device. Tighten the screws. The blade helps support the rear of the chassis. You install the second rear installation blade after securing both front mounting brackets. See [Figure 42 on page 203](#) for detail on installing the rear blade.
7. Grasp both sides of the device, lift it, and position it in the rack so that the blade receptacle at the rear of the chassis catches and slides onto the installation blade. See [Figure 45 on page 206](#).

Figure 45: Slide Mounting Rail onto the Rear Mounting Blade



8. Align the holes in the mounting brackets with the holes in the rack. Ensure that the chassis is level.
9. With four front mounting screws—and cage nuts and washers if your rack requires them—attach the mounting bracket to the rack. Insert the first screw on the opposite corner from the mounting blade. Tighten the screws.
10. Ensure that the switch chassis is level by verifying that the screws on the front of the rack are aligned with the screws at the back of the rack.
11. With two mounting screws—and cage nuts and washers if your rack requires them—slide the second rear mounting blade into the blade receptacle on the mounting blade, and secure it to the rear of the rack by tightening the screws. You might need to loosen and adjust the other mounting blade to install the second blade.

RELATED DOCUMENTATION

Rack-Mounting and Cabinet-Mounting Warnings

[Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device](#) | 188

Configuring a QFX3500 Device as a Standalone Switch

Connecting Earth Ground to a QFX3500 Device

To meet safety and electromagnetic interference (EMI) requirements and to ensure proper operation, you must connect the QFX3500 device to earth ground before you connect it to power.

For installations that require a separate grounding conductor to the chassis, you must attach a protective earthing terminal bracket on the QFX3500 device left front mounting bracket to connect to the earth ground (see [Figure 46 on page 208](#)).

Before you connect earth ground to the protective earthing terminal of a QFX3500 device, ensure that a licensed electrician has attached an appropriate grounding lug to the grounding cable.



CAUTION: Using a grounding cable with an incorrectly attached lug can damage the device.

NOTE: Mount your device in the rack or cabinet before attaching the grounding lug to the device. See [“Mounting a QFX3500 Device in a Rack or Cabinet” on page 201](#).

Ensure that you have the following parts and tools available:


- Protective earthing terminal bracket—This L-shaped bracket attaches to a post on the QFX3500 device left front mounting bracket, providing a protective earthing terminal for the device. This bracket is provided in the accessory kit.
- Grounding cable for your QFX3500 device—The grounding cable must be 14 AWG (2 mm²), minimum 90° C wire, or as permitted by the local code.
- Grounding lug for your grounding cable—The grounding lug required is a Panduit LCD10-10A-L or equivalent. This grounding lug is provided in the accessory kit.
- Three M4 hex nuts with integrated washers—One nut and washer are required to secure the grounding lug bracket to the left front mounting bracket, and two nuts and washers are used to secure the grounding lug to the grounding lug bracket protective earthing terminal. Four nuts are provided in the accessory kit.
- 7-mm wrench or socket with driver to attach all three nuts.

An AC-powered QFX3500 device chassis gains additional grounding when you plug the power supply in the device into a grounded AC power outlet by using an AC power cord appropriate for your geographical location. See [“AC Power Cord Specifications for a QFX Series Device” on page 152](#).

To connect earth ground to a QFX3500 device:

1. Secure the provided protective earthing terminal bracket to the threaded post on the QFX3500 device left front mounting bracket with the nut provided. The posts on the protective earthing terminal bracket should point to the left. See [Figure 46 on page 208](#).

Figure 46: Connecting a Grounding Cable to a QFX3500 Device



2. Connect one end of the grounding cable to a proper earth ground, such as the rack in which the device is mounted.
3. Place the grounding lug attached to the grounding cable over the protective earthing terminal on the protective earthing terminal bracket.
4. Secure the grounding lug to the protective earthing terminal with two nuts.
5. Dress the grounding cable and ensure that it does not touch or block access to other device components and that it does not drape where people could trip over it.

RELATED DOCUMENTATION

General Safety Guidelines and Warnings

Grounded Equipment Warning

[Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device | 188](#)

[Connecting DC Power to a QFX3500, QFX3600, or QFX3600-I Device | 191](#)

Connecting AC Power to a QFX3500, QFX3600, or QFX3600-I Device

The QFX3500, QFX3600, and QFX3600-I devices are shipped from the factory with two 650 W power supplies pre-installed. Each power supply is a hot-removable and hot-insertable field-replaceable unit (FRU) when the second power supply is installed and running. You can install replacement power supplies without powering off the device or disrupting the switching function.

Ensure that you have a power cord appropriate for your geographical location available to connect AC power to the device.

Before you begin connecting AC power to the device:

- Ensure that you have taken the necessary precautions to prevent electrostatic discharge (ESD) damage (see *Prevention of Electrostatic Discharge Damage*).
- Ensure that you have connected the device chassis to earth ground.



CAUTION: Before you connect power to the device, a licensed electrician must attach a cable lug to the grounding and power cables that you supply. A cable with an incorrectly attached lug can damage the device (for example, by causing a short circuit).

To meet safety and electromagnetic interference (EMI) requirements and to ensure proper operation, you must connect the chassis to earth ground before you connect it to power. For installations that require a separate grounding conductor to the chassis, use the protective earthing terminal on the device chassis to connect to the earth ground. For instructions on connecting earth ground, see [“Connecting Earth Ground to a QFX3500 Device” on page 207](#) or [“Connecting Earth Ground to QFX3600 or QFX3600-I Devices” on page 187](#). The device gains additional grounding when you plug the power supply in the device into a grounded AC power outlet by using the AC power cord appropriate for your geographical location (see [“AC Power Cord Specifications for a QFX Series Device” on page 152](#)).

- Install the power supply in the chassis. For instructions on installing a power supply in a QFX3500 device, see *Installing a Power Supply in a QFX3500 Device*. For instructions on installing a power supply in a QFX3600 or QFX3600-I device, see *Installing a Power Supply in a QFX3600 or QFX3600-I Device*.

NOTE: Each power supply must be connected to a dedicated power source outlet.

To connect AC power to a QFX3500, QFX3600, or QFX3600-I device:

1. Attach the grounding strap to your bare wrist and to a site ESD point.
2. Ensure that the power supplies are fully inserted in the chassis and the latches are secure. If only one power supply is installed, ensure a that blank cover panel is installed over the second power supply slot.
3. Locate the power cord or cords shipped with the device; the cords have plugs appropriate for your geographical location. See [“AC Power Cord Specifications for a QFX Series Device” on page 152](#).



WARNING: Ensure that the power cord does not block access to device components or drape where people can trip on it.

4. Connect each power supply to the power sources. Insert the coupler end of the power cord into the AC power cord inlet on the AC power supply faceplate.
5. Push the power cord retainer onto the power cord (see [Figure 36 on page 190](#) or [Figure 37 on page 190](#)).

Figure 47: Connecting an AC Power Cord to an AC Power Supply in a QFX3500 Device





Figure 48: Connecting an AC Power Cord to an AC Power Supply in a QFX3600 or QFX3600-I Device



6. If the AC power source outlet has a power switch, set it to the OFF (O) position.

NOTE: The device powers on as soon as power is provided to the power supply. There is no power switch on the device.

7. Insert the power cord plug into an AC power source outlet.
8. If the AC power source outlet has a power switch, set it to the ON (I) position.
9. Verify that the AC and DC LEDs on each power supply are lit green.

If the amber fault LED is lit, remove power from the power supply, and replace the power supply (see *Removing a Power Supply from a QFX3500 Device* or *Removing a Power Supply from a QFX3600 or QFX3600-I Device*). Do not remove the power supply until you have a replacement power supply ready: the power supplies or a blank cover panel must be installed in the device to ensure proper airflow.



CAUTION: Replace a failed power supply with a blank panel or new power supply within 1 minute of removal to prevent chassis overheating.

RELATED DOCUMENTATION

AC Power Supply for a QFX3500, QFX3600, or QFX3600-I Device

AC Power Supply LEDs on a QFX3500, QFX3600, or QFX3600-I Device

Connecting DC Power to a QFX3500, QFX3600, or QFX3600-I Device

The QFX3500, QFX3600, and QFX3600-I devices are shipped from the factory with two 650 W power supplies pre-installed. Each power supply is a hot-removable and hot-insertable field-replaceable unit (FRU) when the second power supply is installed and running. You can install replacement power supplies without powering off the device or disrupting the switching function.



WARNING: DC-powered QFX3500, QFX3600 and QFX3600-I devices are intended for installation only in a restricted access location.

NOTE: The battery returns of the DC power supply should be connected as an isolated DC return (DC-I).

Before you begin connecting DC power to the device:

- Ensure that you have taken the necessary precautions to prevent electrostatic discharge (ESD) damage (see *Prevention of Electrostatic Discharge Damage*).
- Ensure that you have connected the device chassis to earth ground.



CAUTION: Before you connect power to the device, a licensed electrician must attach a cable lug to the grounding and power cables that you supply. A cable with an incorrectly attached lug can damage the device (for example, by causing a short circuit).

To meet safety and electromagnetic interference (EMI) requirements and to ensure proper operation, you must connect the chassis to earth ground before you connect it to power. For installations that require a separate grounding conductor to the chassis, use the protective earthing terminal on the device chassis to connect to the earth ground. For instructions on connecting earth ground, see [“Connecting Earth Ground to a QFX3500 Device” on page 207](#) or [“Connecting Earth Ground to QFX3600 or QFX3600-I Devices” on page 187](#).

- Install the power supply in the chassis. For instructions on installing a power supply in a QFX3500 device, see *Installing a Power Supply in a QFX3500 Device*. For instructions on installing a power supply in a QFX3600 or QFX3600-I device, see *Installing a Power Supply in a QFX3600 or QFX3600-I Device*

Ensure that you have the following parts and tools available:

- DC power source cables (14–16 AWG) with ring lug (Molex 190700069 or equivalent) (not provided)
- Phillips (+) screwdriver, number 2 (not provided)
- Multimeter (not provided)

To connect DC power to a QFX3500, QFX3600 or QFX3600-I device:

1. Attach the grounding strap to your bare wrist and to a site ESD point.
2. Verify that the DC power cables are correctly labeled before making connections to the power supply. In a typical power distribution scheme where the return is connected to chassis ground at the battery plant, you can use a multimeter to verify the resistance of the –48V and RTN DC cables to chassis ground:

- The cable with very low resistance (indicating a closed circuit) to chassis ground is positive (+) and will be installed on the V+ (return) DC power input terminal.
- The cable with very high resistance (indicating an open circuit) to chassis ground is negative (-) and will be installed on the V- (input) DC power input terminal.



CAUTION: You must ensure that power connections maintain the proper polarity. The power source cables might be labeled (+) and (-) to indicate their polarity. There is no standard color coding for DC power cables. The color coding used by the external DC power source at your site determines the color coding for the leads on the power cables that attach to the DC power input terminals on each power supply.

3. Ensure that the input circuit breaker is open so that the voltage across the DC power source cable leads is 0 V and that the cable leads do not become active while you are connecting DC power.

NOTE: The V+ terminals are referred to as +RTN, and V- terminals are referred to as -48 V in *DC Power Wiring Sequence Warning* and *DC Power Electrical Safety Guidelines*.

4. Ensure that the power supplies are fully inserted in the chassis.
5. Remove the terminal block cover. The terminal block cover is a piece of clear plastic that snaps into place over the terminal block (see [Figure 38 on page 195](#)).
6. Remove the screws on the terminals using the screwdriver. Save the screws.



WARNING: Ensure that the power cables do not block access to device components or drape where people can trip on them.


7. Connect each power supply to the power sources. Secure power source cables to the power supplies by screwing the ring lugs attached to the cables to the appropriate terminals by using the screw from the terminals (see [Figure 39 on page 196](#) and [Figure 38 on page 195](#)).



CAUTION: The DC power supply has four terminals labeled V+, V+, V-, and V- for connecting DC power source cables labeled positive (+) and negative (-). The V+ terminals are shunted internally together, as are the V- terminals. The same polarity terminal can be wired together from the same source to provide an additional current path in a higher power chassis. Do not connect the terminals to different sources. For example, connect -48 V from DC source feed A to the input terminals of one power supply and connect -48 V from feed B to the input terminals of the second power supply on the other side of the chassis. This configuration provides the commonly deployed A/B feed redundancy for the system.

- a. Secure the ring lug of the positive (+) DC power source cable to the V+ terminal on the DC power supply.
- b. Secure the ring lug of the negative (-) DC power source cable to the V- terminal on the DC power supply.
- c. Tighten the screws on the power supply terminals until snug using the screwdriver. Do not overtighten—apply between 5 in-lb (0.56 Nm) and 6 in-lb (0.68 Nm) of torque to the screws.

Figure 49: DC Power Supply Faceplate for a QFX3500, QFX3600 or QFX3600-I Device




1—Shunt positive input terminals	5—ESD grounding point
2—Shunt negative input terminals	6—Fault LED
3—Terminal block	7—Output LED
4—Ejector lever	8—Input LED



CAUTION: The V+ terminals are shunted internally together, as are the V- terminals. The same polarity terminal can be wired together from the same source to provide an additional current path in a higher power chassis. Do not connect the terminals to different sources.

Figure 50: Securing Ring Lugs to the Terminals on the QFX3500, QFX3600 or QFX3600-I DC Power Supply



8. Replace the terminal block cover.

9. Close the input circuit breaker.

NOTE: The device powers on as soon as power is provided to the power supply. There is no power switch on the device.

10. Verify that the **IN** and **OUT** LEDs on the power supply are lit green and are on steadily.

RELATED DOCUMENTATION

DC Power Supply for a QFX3500, QFX3600, or QFX3600-I Device

DC Power Supply LEDs on a QFX3500, QFX3600, or QFX3600-I Device

Connecting a QFX Series Device to a Management Console

The QFX Series has a console port with an RJ-45 connector. Use the console port to connect the device to a management console or to a console server.

Ensure that you have an RJ-45 to DB-9 rollover cable available. An RJ-45 cable with an RJ-45 to DB-9 adapter is provided with the device.

NOTE: If your laptop or PC does not have a DB-9 male connector pin and you want to connect your laptop or PC directly to the QFX Series, use a combination of the RJ-45 cable and RJ-45 to DB-9 adapter supplied with the device and a USB to DB-9 male adapter. You must provide the USB to DB-9 male adapter.


To connect the QFX Series to a management console (see [Figure 30 on page 174](#) and [Figure 31 on page 174](#)):

1. Connect one end of the Ethernet cable to the console port (labeled **CON**).
2. Connect the other end of the Ethernet cable into the console server (see [Figure 30 on page 174](#)) or management console (see [Figure 31 on page 174](#)).

Figure 51: Connecting the QFX Series to a Management Console Through a Console Server



Figure 52: Connecting the QFX Series Directly to a Management Console



RELATED DOCUMENTATION

Console Port Connector Pinout Information

Configuring Junos OS to Set Console and Auxiliary Port Properties

Installing a QFX5100 Node Device

IN THIS CHAPTER

- [Installing and Connecting a QFX5100 Device | 219](#)
- [Unpacking a QFX5100 Device | 220](#)
- [Mounting a QFX5100 Device in a Rack or Cabinet | 222](#)
- [Connecting Earth Ground to a QFX5100 Device | 227](#)
- [Connecting AC Power to a QFX5100 Device | 228](#)
- [Connecting DC Power to a QFX5100 Device | 231](#)
- [Connecting a QFX Series Device to a Management Console | 236](#)

Installing and Connecting a QFX5100 Device

You can mount a QFX5100 device:

- Flush with the front of a 19-in. four-post rack. Use the standard mounting brackets provided with the switch for this configuration.
- Recessed 2 in. (5 cm) from the front of a 19-in. four-post rack. Use the extension bracket provided in the standard mounting kit for this configuration. Recessed mounting is primarily used in enclosed cabinets.

To install and connect a QFX5100 device:

1. Follow the instructions in [“Unpacking a QFX5100 Device” on page 220](#).
2. Determine how the device is to be mounted.

Flush or recessed mounted in a rack or cabinet, see [“Mounting a QFX5100 Device in a Rack or Cabinet” on page 222](#).

3. Follow the instructions in:
 - a. [Connecting Earth Ground to a QFX5100 Device on page 227](#)

- b. [“Connecting AC Power to a QFX5100 Device” on page 228](#) or [Connecting DC Power to a QFX5100 Device on page 231](#)
 - c. *Registering Products—Mandatory for Validating SLAs*
4. Depending on how you plan to use the QFX5100 device, do one of the following:
- If you are using the QFX5100 device as a standalone switch, follow the instructions in *Configuring a QFX5100 Device*.
 - If you are using the QFX5100 device as a Node device in a QFX3000-G QFabric system, see *QFX3000-G QFabric System Installation Overview* for information about the steps to install and configure your QFX3000-G QFabric system.
 - If you are using the QFX5100 device as a Node device in a QFX3000-M QFabric system, see [“QFX3000-M QFabric System Installation Overview” on page 109](#) for information about the steps to install and configure your QFX3000-M QFabric system.
 - If you are using a QFX5100-24Q as an Interconnect device in a QFX3000-M QFabric system, see [“Connecting a QFX5100 Node Device to a QFX5100-24Q Interconnect Device” on page 304](#).
 - If you are using the QFX5100 device in a Virtual Chassis Fabric, see *Connecting a QFX5100 Device in a Virtual Chassis Fabric*.

RELATED DOCUMENTATION

Rack Requirements for a QFX5100 Device

Cabinet Requirements for a QFX5100 Device

Clearance Requirements for Airflow and Hardware Maintenance for a QFX5100 Device

Unpacking a QFX5100 Device

The QFX5100 switch chassis is a rigid sheet-metal structure that houses the hardware components. A QFX5100 device is shipped in a cardboard carton, secured with foam packing material. The carton also contains an accessory box and quick start instructions.



CAUTION: QFX5100 devices are maximally protected inside the shipping carton. Do not unpack the switch until you are ready to begin installation.

To unpack a QFX5100 device:

1. Move the shipping carton to a staging area as close to the installation site as possible, but where you have enough room to remove the system components.
2. Position the carton so that the arrows are pointing up.
3. Open the top flaps on the shipping carton.
4. Remove the accessory box and verify the contents against the inventory included in the box.
[Table 53 on page 221](#) lists the inventory of components supplied with a QFX5100 device.
5. Pull out the packing material holding the switch in place.
6. Verify the chassis components received:
 - Two power supplies
 - Fan modules
 - Five fan modules for 1 U devices
 - Three fan modules for 2 U devices

NOTE: Product SKU QFX5100-24Q: If you ordered the optional high-speed uplink modules, they are packaged as components and must be installed in the switch

7. Save the shipping carton and packing materials in case you need to move or ship the switch later.

Table 53: Inventory of Components Supplied with a QFX5100 Device

Component	Quantity
Chassis with five fan modules and two power supplies. The QFX5100-96S has three fan modules.	1
Rear mounting blades	2
Front mounting brackets	2
Extension brackets	2
RJ-45 cable and RJ-45 to DB-9 adapter	1
Power cords (AC systems only)	2

RELATED DOCUMENTATION

[Mounting a QFX5100 Device in a Rack or Cabinet | 222](#)

[Installing and Connecting a QFX5100 Device | 219](#)

Mounting a QFX5100 Device in a Rack or Cabinet

IN THIS SECTION

- [Before You Begin Rack Installation | 222](#)
- [Four-Post Procedure | 224](#)

You can mount all QFX5100 switches on a four post 19-in. rack or cabinet using the mounting kit provided with the device.

For four post rack or cabinet installations, the mounting kit contains two front mounting rails with two matching rear mounting blades. This configuration allows either end of the switch to be mounted flush with the rack and still be adjustable for racks with different depths.

(The remainder of this topic uses “rack” to mean “rack or cabinet.”) The front and rear rack rails must be spaced between 28 in. (71.1 cm) and 36 in. (91.4 cm) front to back.

Before You Begin Rack Installation

Before you begin mounting a QFX5100 switch in the rack or cabinet:

1. Ensure that you understand how to prevent electrostatic discharge (ESD) damage. See *Prevention of Electrostatic Discharge Damage*.
2. Verify that the site meets the requirements described in *Site Preparation Checklist for a QFX5100 Device*.
3. Place the rack in its permanent location, allowing adequate clearance for airflow and maintenance, and secure it to the building structure.
4. Read “[General Site Guidelines](#)” on page 124, with particular attention to *Chassis Lifting Guidelines for a QFX5100 Device*.

5. Remove the switch from the shipping carton (see [“Unpacking a QFX5100 Device” on page 220](#)).
6. Ensure that you have the following parts and tools available to mount the switch in a rack:
 - ESD grounding strap (not provided).
 - Blades, rails, or brackets (provided).
 - For four-post installations:
 - One pair of rear mounting blades. These mounting blades support the rear of the chassis and must be installed (provided).
 - One pair of front mounting rails. The mounting blades slide into the mounting rails to support the switch (provided).
 - Screws to secure the mounting rails to the chassis (provided).
 - Twelve screws for 1 U chassis
 - Twenty-four screws for QFX5100-96S
 - Eight screws to secure the chassis and rear installation blades to the rack (not provided).
 - Appropriate screwdriver for the mounting screws (not provided).
 - Two power cords with plugs appropriate to your geographical location (provided).
 - RJ-45 cable and RJ-45 to DB-9 serial port adapter (provided).
 - Management host, such as a PC laptop, with a serial port (not provided).

Optional equipment: Grounding cable kit with bracket, lug, and three nuts with integrated washers.



WARNING: The 1 U versions of QFX5100 switches must be supported at all four corners. Mounting the chassis using only the front brackets will damage the chassis and can result in serious bodily injury.



CAUTION: All QFX5100 switches require two people for installation, one person to lift the switch into place and another person to attach the switch to the rack. If you are installing the QFX5100 switch above 60 in. (152.4 cm) from the floor, you can remove the power supplies and fan modules to minimize the weight before attempting to install the switch.



CAUTION: If you are mounting multiple switches on a rack, mount the switch in the lowest position of the rack first. Proceed to mount the rest of the switches from the bottom to the top of the rack to minimize the risk of the rack toppling.

Four-Post Procedure

To mount the switch on four posts in a rack using the provided mounting kit:

1. Attach the ESD grounding strap to your bare wrist and to a site ESD point.
2. Decide whether the Field Replaceable Unit (FRU) end of the switch or the port end is to be placed at the front of the rack. Position the switch in such a manner that the **AIR IN** labels on components are next to the cold aisle and **AIR OUT** labels on components are next to the hot aisle.
3. Align the holes in the mounting rail with the holes on the side of the chassis. See [Figure 44 on page 205](#) through [Figure 55 on page 225](#) for examples the proper alignment of 1 U and 2 U chassis systems.

Figure 53: Attaching Mounting Rails to the QFX5100-24Q




Figure 54: Attaching Mounting Rails to the QFX5100-48S





Figure 55: Attaching Mounting Rails to the QFX5100-96S



4. Attach the mounting rail to the switch using the mounting screws (and cage nuts and washers if your rack requires them). Tighten the screws.
5. Repeats steps 4 and 5 on the opposite side of the switch.
6. Have one person grasp both sides of the switch, lift it, and position it in the rack so that the front bracket is aligned with the rack holes.
7. Have a second person secure the front of the switch to the rack using four mounting screws (and cage nuts and washers if your rack requires them.) Tighten the screws. See [Figure 45 on page 206](#) and [Figure 57 on page 226](#) for examples of connecting the mounting rails and blades.

Figure 56: Attach 1 U Switch to Rack






Figure 57: Slide Mounting Rail onto the QFX5100-96S Rear Mounting Blade



8. Continue to support the switch while sliding the rear mounting-blades into the channel of the side mounting-rails and securing the blades to the rack. Use the four mounting screws (and cage nuts and washers if your rack requires them) to attach each blade to the rack. (Use eight front-mounting screws for the QFX5100-96S.) Tighten the screws. See [Figure 58 on page 226](#).

Figure 58: Slide Mounting Blade into 1 U Mounting Rail



9. Ensure that the switch chassis is level by verifying that all the screws on the front of the rack are aligned with the screws at the back of the rack.

RELATED DOCUMENTATION

Rack-Mounting and Cabinet-Mounting Warnings

[Connecting AC Power to a QFX5100 Device | 228](#)

[Connecting Earth Ground to a QFX5100 Device | 227](#)

Connecting Earth Ground to a QFX5100 Device

To meet safety and electromagnetic interference (EMI) requirements and to ensure proper operation, you must connect the QFX5100 device to earth ground before you connect it to power.

For installations that require a separate grounding conductor to the chassis, you must attach a protective earthing terminal bracket on the QFX5100 device left front mounting bracket to connect to the earth ground (see [Figure 59 on page 228](#)).

Before you connect earth ground to the protective earthing terminal of a QFX5100 device, ensure that a licensed electrician has attached an appropriate grounding lug to the grounding cable.



CAUTION: Using a grounding cable with an incorrectly attached lug can damage the switch.

NOTE: Mount your switch in the rack or cabinet before attaching the grounding lug to the switch. See [“Mounting a QFX5100 Device in a Rack or Cabinet” on page 222](#).

Ensure that you have the following parts and tools available:

- Protective earthing terminal bracket—This bracket attaches to the QFX5100 switch chassis through the left front mounting bracket, providing a protective earthing terminal for the switch.
- Grounding cable for your QFX5100 device—The grounding cable must be 14 AWG (2 mm²), minimum 90° C wire, or as permitted by the local code.
- Grounding lug for your grounding cable—The grounding lug required is a Panduit LCD10-10A-L or equivalent.
- Two SAE 10-32 washers and screws—To attach the grounding lug to the protective earthing terminal.
- Screwdriver to attach the screws.

An AC-powered QFX5100 switch chassis gains additional grounding when you plug the power supply in the switch into a grounded AC power outlet by using an AC power cord appropriate for your geographical location. See [“AC Power Cord Specifications for a QFX Series Device” on page 152](#).

To connect earth ground to a QFX5100 device:

1. Attach one end of the grounding cable to an appropriate earth ground site, such as the mounting rack.
2. Position the grounding lug over the protective earthing terminal on the side of the chassis, which is visible through the mounting bracket.

3. Secure the grounding lug to the protective earthing terminal with the washers and screws. See [Figure 59 on page 228](#) and [Figure 60 on page 228](#).

Figure 59: Connecting a Grounding Cable to a 1 U QFX5100 Device





Figure 60: Connecting a Grounding Cable to the 2 U QFX5100-96S Device



4. Dress the grounding cable and ensure that it does not touch or block access to other device components and that it does not drape where people could trip over it.

RELATED DOCUMENTATION

General Safety Guidelines and Warnings

Grounded Equipment Warning

[Connecting AC Power to a QFX5100 Device | 228](#)

[Connecting DC Power to a QFX5100 Device | 231](#)

Connecting AC Power to a QFX5100 Device

The QFX5100 is shipped from the factory with two power supplies. Each power supply is a hot-removable and hot-insertable field-replaceable unit (FRU) when the second power supply is installed and running.

You can install replacement power supplies in the two slots next to the fan modules without powering off the switch or disrupting the switching function.

Ensure that you have a power cord appropriate for your geographical location available to connect AC power to the switch.

Before you begin connecting AC power to the switch:

- Ensure that you have taken the necessary precautions to prevent electrostatic discharge (ESD) damage (see *Prevention of Electrostatic Discharge Damage*).
- Ensure that you have connected the switch chassis to earth ground.



CAUTION: Before you connect power to the switch, a licensed electrician must attach a cable lug to the grounding and power cables that you supply. A cable with an incorrectly attached lug can damage the switch (for example, by causing a short circuit).

To meet safety and electromagnetic interference (EMI) requirements and to ensure proper operation, you must connect the chassis to earth ground before you connect it to power. For installations that require a separate grounding conductor to the chassis, use the protective earthing terminal on the switch chassis to connect to the earth ground. For instructions on connecting earth ground, see [“Connecting Earth Ground to a QFX5100 Device” on page 227](#). The switch gains additional grounding when you plug the power supply in the switch into a grounded AC power outlet by using the AC power cord appropriate for your geographical location (see *AC Power Supply for a QFX5100 Device*).

- Install the power supply in the chassis. For instructions on installing a power supply in a QFX5100 device, see *Installing a Power Supply in a QFX5100 Device*.

NOTE: Each power supply must be connected to a dedicated power source outlet.

To connect AC power to a QFX5100 device:

1. Attach the grounding strap to your bare wrist and to a site ESD point.
2. Ensure that the power supplies are fully inserted in the chassis and the latches are secure. If only one power supply is installed, ensure a that blank cover panel is installed over the second power supply slot.


3. Locate the power cord or cords shipped with the switch; the cords have plugs appropriate for your geographical location. See [“AC Power Cord Specifications for a QFX Series Device” on page 152](#).



WARNING: Ensure that the power cord does not block access to device components or drape where people can trip on it.


4. Connect each power supply to the power sources. Insert the coupler end of the power cord into the AC power cord inlet on the AC power supply faceplate.
5. Push the power cord retainer onto the power cord (see [Figure 61 on page 230](#) and [Figure 62 on page 230](#)).

Figure 61: Connecting an AC Power Cord to an AC Power Supply in a 1 U QFX5100 Device



1—Power cord retainer

Figure 62: Connecting an AC Power Cord to an AC Power Supply in a 2 U QFX5100 Device



6. If the AC power source outlet has a power switch, set it to the OFF (O) position.

NOTE: The switch powers on as soon as power is provided to the power supply. There is no power switch on the device.

7. Insert the power cord plug into an AC power source outlet.
8. If the AC power source outlet has a power switch, set it to the ON (I) position.
9. Verify that the AC and DC LEDs on each power supply are lit green.

If the amber fault LED is lit, remove power from the power supply, and replace the power supply (see *Removing a Power Supply from a QFX5100 Device*). Do not remove the power supply until you have a replacement power supply ready: the power supplies or a blank cover panel must be installed in the switch to ensure proper airflow.



CAUTION: Replace a failed power supply with a blank panel or new power supply within 1 minute of removal to prevent chassis overheating.



CAUTION: A system reboot with Routing Engine FPGA version 7.1 might not successfully boot the Junos OS software. In case of a system reboot failure, you need to power cycle the switch. To check the current FPGA version, issue the **show chassis firmware** command.

RELATED DOCUMENTATION

[AC Power Supply for a QFX5100 Device](#)

[AC Power Supply LEDs on a QFX5100 Device](#)

Connecting DC Power to a QFX5100 Device

The QFX5100 is shipped from the factory with two power supplies. Each power supply is a hot-removable and hot-insertable field-replaceable unit (FRU) when the second power supply is installed and running. You can install replacement power supplies in the two slots next to the fan modules without powering off the switch or disrupting the switching function.



WARNING: DC-powered QFX5100 devices are intended for installation only in a restricted access location.

NOTE: The battery returns of the DC power supply must be connected as an isolated DC return (DC-I).

Before you begin connecting DC power to the switch:

- Ensure that you have taken the necessary precautions to prevent electrostatic discharge (ESD) damage (see *Prevention of Electrostatic Discharge Damage*).
- Ensure that you have connected the switch chassis to earth ground.



CAUTION: Before you connect power to the switch, a licensed electrician must attach a cable lug to the grounding and power cables that you supply. A cable with an incorrectly attached lug can damage the switch (for example, by causing a short circuit).

To meet safety and electromagnetic interference (EMI) requirements and to ensure proper operation, you must connect the chassis to earth ground before you connect it to power. For installations that require a separate grounding conductor to the chassis, use the protective earthing terminal on the switch chassis to connect to the earth ground. For instructions on connecting earth ground, see [“Connecting Earth Ground to a QFX5100 Device” on page 227](#).

- Install the power supply in the chassis. For instructions on installing a power supply in a QFX5100 device, see *Installing a Power Supply in a QFX5100 Device*.

Ensure that you have the following parts and tools available:

- DC power source cables (14–16 AWG) with ring lug (Molex 190700069 or equivalent) (not provided)
- Phillips (+) screwdriver, number 2 (not provided)
- Multimeter (not provided)

To connect DC power to a QFX5100 device:

1. Attach the grounding strap to your bare wrist and to a site ESD point.
2. Verify that the DC power cables are correctly labeled before making connections to the power supply. In a typical power distribution scheme where the return is connected to chassis ground at the battery plant, you can use a multimeter to verify the resistance of the –48V and RTN DC cables to chassis ground:
 - The cable with very low resistance (indicating a closed circuit) to chassis ground is positive (+) and will be installed on the V+ (return) DC power input terminal.

- The cable with very high resistance (indicating an open circuit) to chassis ground is negative (–) and will be installed on the V– (input) DC power input terminal.



CAUTION: You must ensure that power connections maintain the proper polarity. The power source cables might be labeled (+) and (–) to indicate their polarity. There is no standard color coding for DC power cables. The color coding used by the external DC power source at your site determines the color coding for the leads on the power cables that attach to the DC power input terminals on each power supply.

3. Ensure that the input circuit breaker is open so that the voltage across the DC power source cable leads is 0 V and that the cable leads do not become active while you are connecting DC power.

NOTE: The V+ terminals are referred to as +RTN, and V– terminals are referred to as –48 V in *DC Power Wiring Sequence Warning* and *DC Power Electrical Safety Guidelines*.

4. Ensure that the power supplies are fully inserted in the chassis.
5. Remove the terminal block cover. The terminal block cover is a piece of clear plastic that snaps into place over the terminal block (see [Figure 63 on page 235](#)).
6. Remove the screws on the terminals using the screwdriver. Save the screws.



WARNING: Ensure that the power cables do not block access to device components or drape where people can trip on them.

7. Connect each power supply to the power sources. Secure power source cables to the power supplies by screwing the ring lugs attached to the cables to the appropriate terminals by using the screw from the terminals (see [Figure 63 on page 235](#) and [Figure 64 on page 236](#)).

The QFX5100 is designed to operate with a DC power supply that has a single, non-redundant, feed input. For source redundancy, two DC power supplies must be installed in QFX5100; connect source (A) to one power supply and connect source (B) to the second power supply. This configuration provides the commonly deployed A/B feed redundancy for the system.

The terminal block of the power supply has four terminals labeled V+, V+, V–, and V– for connecting DC power source cables labeled positive (+) and negative (–). The V+ terminals are shunted internally together, as are the V– terminals.




CAUTION: The connection between each power source and power supply must include a circuit breaker.

Do not connect two sources to a single power supply because doing so can potentially cause circulating current in feed wires whenever there is any difference in the voltage of the two sources.

NOTE: For QFX5100-96S installations using battery backup, a single 15 A circuit breaker is recommended.

- a. Secure the ring lug of the positive (+) DC power source cable to the V+ terminal on the DC power supply.
- b. Secure the ring lug of the negative (–) DC power source cable to the V– terminal on the DC power supply.
- c. Tighten the screws on the power supply terminals until snug using the screwdriver. Do not overtighten—apply between 5 in-lb (0.56 Nm) and 6 in-lb (0.68 Nm) of torque to the screws.

Figure 63: DC Power Supply Faceplate for a QFX5100 Device




1—Shunt negative input terminals (-48V)	5—ESD grounding point
2—Shunt positive input terminals (+RTN)	6—Fault LED
3—Terminal block	7—Output LED
4—Ejector lever	8—Input LED



CAUTION: The V+ terminals are shunted internally together, as are the V- terminals. The same polarity terminal can be wired together from the same source to provide an additional current path in a higher power chassis. Do not connect the terminals to different sources.

Figure 64: Securing Ring Lugs to the Terminals on the QFX5100 DC Power Supply



8. Replace the terminal block cover.

9. Close the input circuit breaker.

NOTE: The switch powers on as soon as power is provided to the power supply. There is no power switch on the device.

10. Verify that the **IN** and **OUT** LEDs on the power supply are lit green and are on steadily.



CAUTION: A system reboot with Routing Engine FPGA version 7.1 might not successfully boot the Junos OS software. In case of a system reboot failure, you need to power cycle the switch. To check the current FPGA version, issue the **show chassis firmware** command.

RELATED DOCUMENTATION

DC Power Supply in a QFX5100 Device

DC Power Supply LEDs on a QFX5100 Device

Connecting a QFX Series Device to a Management Console

The QFX Series has a console port with an RJ-45 connector. Use the console port to connect the device to a management console or to a console server.

Ensure that you have an RJ-45 to DB-9 rollover cable available. An RJ-45 cable with an RJ-45 to DB-9 adapter is provided with the device.

NOTE: If your laptop or PC does not have a DB-9 male connector pin and you want to connect your laptop or PC directly to the QFX Series, use a combination of the RJ-45 cable and RJ-45 to DB-9 adapter supplied with the device and a USB to DB-9 male adapter. You must provide the USB to DB-9 male adapter.


To connect the QFX Series to a management console (see [Figure 30 on page 174](#) and [Figure 31 on page 174](#)):

1. Connect one end of the Ethernet cable to the console port (labeled **CON**).
2. Connect the other end of the Ethernet cable into the console server (see [Figure 30 on page 174](#)) or management console (see [Figure 31 on page 174](#)).

Figure 65: Connecting the QFX Series to a Management Console Through a Console Server



Figure 66: Connecting the QFX Series Directly to a Management Console



RELATED DOCUMENTATION

Console Port Connector Pinout Information

Configuring Junos OS to Set Console and Auxiliary Port Properties

Cabling a Copper-Based Control Plane for the QFX3000-M QFabric System


IN THIS CHAPTER

- Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy | 239
- Connecting QFX3100 Director Devices in a Director Group | 242
- Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network | 245
- Connecting a QFX3100 Director Device to a Network for Out-of-Band Management | 247
- Connecting a QFX5100-24Q Interconnect Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 248
- Connecting a QFX3600-I Interconnect Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 252
- Connecting a QFX5100 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 256
- Connecting a QFX3600 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 259
- Connecting a QFX3500 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 263

Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 or to two EX4300 switches. For redundancy and communication, you must connect the two EX Series switches using the Ethernet SFP+ uplink module ports configured as a link aggregation group (LAG) (see [Figure 67 on page 240](#)).

Figure 67: QFX3000-M QFabric System Control Plane—Inter-EX Series Switch LAG Connections



For an EX4200 control plane, use the 1-Gigabit Ethernet SFP+ uplink module ports; for an EX4300 control plane, use the 10-Gigabit Ethernet SFP+ uplink module ports. Do not mix EX4200 and EX4300 switches in the same control plane.

Specific ports have been reserved on the EX Series switches to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

Before you begin to interconnect two EX Series switches for QFX3000-M QFabric system control plane redundancy:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). See [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).
- Install your EX4200 or EX4300 switches. See *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.

NOTE: For a copper-based control plane network, either use EX4200-24T switches with an SFP+ uplink module installed or use an EX4300-48T with an SFP+ uplink module installed. For a fiber-based control plane network, use EX4200-24F switches with an SFP+ uplink module installed or an EX4300-48P with an SFP+ uplink module installed.

- Install transceivers in ports **0** and **1** of the uplink module (see *Installing a Transceiver*).
 - For an EX4200 Switch control plane, ensure that you have installed 1-Gigabit Ethernet SFP transceivers in each EX4200 switch uplink module. For a list of supported transceivers, see *Pluggable Transceivers Supported on EX4200 Switches*.
 - For an EX4300 Switch control plane, ensure that you have installed 10-Gigabit Ethernet SFP+ transceivers in each EX4300 switch uplink module. For a list of supported transceivers, see *Pluggable Transceivers Supported on EX4300 Switches*.
- Ensure that you have appropriate fiber-optic cables (see *Pluggable Transceivers Supported on EX4200 Switches* and *Pluggable Transceivers Supported on EX4300 Switches*).
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings*).


To interconnect two EX Series switches for QFabric system control plane redundancy (see [Figure 67 on page 240](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
2. Remove the rubber safety cap from the SFP optical transceiver in port **0** of the first EX Series switch uplink module. Save the cap.
3. Insert the cable connector into the optical transceiver (see [Figure 68 on page 241](#)).

Figure 68: Connecting a Fiber-Optic Cable to an Optical Transceiver Installed in an EX Series Switch



4. If the connector at the other end of the fiber-optic cable is covered by a rubber safety cap, remove the cap. Save the cap.
5. Remove the rubber safety cap from the SFP optical transceiver in port **0** of the *second* EX Series switch uplink module. Save the cap.
6. Insert the cable connector into the optical transceiver.
7. Repeat Step [1](#) through Step [6](#) to interconnect uplink module port **1** on both the EX Series switches.
8. Secure the cables so that they are not supporting their own weight. Place excess cable out of the way in a neatly coiled loop. Placing fasteners on a loop helps cables maintain their shape.



CAUTION: Do not bend fiber-optic cables beyond their minimum bend radius. An arc smaller than a few inches in diameter can damage the cables and cause problems that are difficult to diagnose.

Do not let fiber-optic cables hang free from the connector. Do not allow fastened loops of cables to dangle, which stresses the cables at the fastening point.

RELATED DOCUMENTATION

[Installing and Connecting a QFX3100 Director Device | 165](#)

[Installing and Connecting a QFX3600 or QFX3600-I Device | 177](#)

[Installing and Connecting a QFX3500 Device | 199](#)

Connecting QFX3100 Director Devices in a Director Group

A QFabric system requires two QFX3100 Director devices interconnected as a QFX3100 Director *group* (see [Figure 69 on page 242](#) through [Figure 71 on page 243](#)).

Figure 69: QFX3100 Director Group Control Plane Connections for QFX3000-G QFabric System Using Copper-Based Control Plane




Figure 70: QFX3100 Director Group Control Plane Connections for QFX3000-G QFabric System Using Fiber-Based Control Plane





Figure 71: QFX3100 Director Group Control Plane Connections for QFX3000-M QFabric System



The second QFX3100 Director device provides redundancy for the control plane and management network.

Before you begin to connect QFX3100 Director devices in a Director group:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [Installing and Connecting a QFX3008-I Interconnect Device](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).

NOTE: For a copper-based QFX3000-M QFabric system control plane network, use QFX3100 Director devices with RJ-45 network modules installed. For a fiber-based control plane network, use QFX3100 Director devices with SFP network modules installed.

- Ensure that you have appropriate transceivers and cables available. For cable specifications, see [“Cable Specifications for Copper-Based Control Plane Connections for the QFabric System” on page 144](#) and [“Determining Transceiver Support for QFabric Systems” on page 143](#).



CAUTION: 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. Although a single inter-Director device can lose a link and regain its connection, the loss of both inter-Director device links causes one of the Director devices to become isolated from the Director group. In [Figure 69 on page 242](#) through [Figure 71 on page 243](#), these redundant patch cables are shown in red and connect port 3 to port 3.

To connect QFX3100 Director devices in a Director group (see [Figure 69 on page 242](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. Connect the port labeled **3** on the first network module on one of the Director devices to the corresponding port (labeled **3**) on the first network module on the second Director device.
2. Connect the port labeled **3** on the *second* network module on one of the Director devices to the corresponding port (labeled **3**) on the *second* network module on the second Director device.

RELATED DOCUMENTATION

Connecting QFX3100 Director Devices to a Copper-Based QFX3000-G QFabric System Control Plane Network


Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-G QFabric System Control Plane Network

[Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network | 245](#)

Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 or two EX4300 switches. QFX3100 Director devices have two 4-port network modules. Use the network module ports to connect the QFX3100 Director group to the EX Series switches (see [Figure 72 on page 245](#)).

Figure 72: QFX3100 Director Group to EX Series Switch Connections for QFX3000-M QFabric System



Use the following QFX3100 Director devices and EX Series switches for a copper-based QFX3000-M QFabric system control plane network:

- QFX3100 Director devices with 1000BASE-T network modules installed. Each 1000BASE-T network module provides four 1000BASE-T ports (labeled **0** through **3**) with RJ-45 connectors.
- EX4200-24T or EX4300-48T switches with an SFP+ uplink module installed.

Specific ports have been reserved on the EX Series switches to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

Before you begin to connect a QFX3100 Director device to the copper-based QFX3000-M QFabric system control plane network:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). See [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).
- Install your EX4200 or EX4300 switches. See *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.

NOTE: Use either EX4200-24T switches with an SFP+ uplink module installed or use an EX4300-48T with an SFP+ uplink module installed.

- Install transceivers in ports **0** and **1** of the uplink module (see *Installing a Transceiver*).
 - For an EX4200 Switch control plane, ensure that you have installed 1-Gigabit Ethernet SFP transceivers in each EX4200 switch uplink module. For a list of supported transceivers, see *Pluggable Transceivers Supported on EX4200 Switches*.
 - For an EX4300 Switch control plane, ensure that you have installed 10-Gigabit Ethernet SFP+ transceivers in each EX4300 switch uplink module. For a list of supported transceivers, see *Pluggable Transceivers Supported on EX4300 Switches*.
- Interconnect the two EX Series switches. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#).
- Connect the two QFX3100 Director devices to create a Director group. See [“Connecting QFX3100 Director Devices in a Director Group” on page 242](#).
- Ensure that you have four RJ-45 patch cables available for each Director device. For cable specifications, see [“Cable Specifications for Copper-Based Control Plane Connections for the QFabric System” on page 144](#).

To connect a QFX3100 Director device to the copper-based QFX3000-M QFabric control plane network (see [Figure 72 on page 245](#)):

1. Connect both network modules on the first Director device (labeled **DG0** in [Figure 72 on page 245](#)) to the two EX4200 or EX4300 switches (labeled **EX0** and **EX1** in [Figure 72 on page 245](#)). You connect the first two ports (labeled **0** and **1**) on the first network module to the first EX Series switch (**EX0**). You connect the first two ports on the second network module (also labeled **0** and **1**) to the second EX Series switch (**EX1**). The ports used are the same on each EX Series switch.

Table 54: QFX3100 Director Device-to-EX Series Switch Control Plane Port Assignments for DG0

Network Module Port 0	Network Module Port 1	Network Module Port 3
EX Series switch port ge-0/0/20	EX Series switch port ge-0/0/21	Connect this port to the identical port on the other Director device. See “Connecting QFX3100 Director Devices in a Director Group” on page 242 .

2. Connect both network modules on the *second* Director device (labeled **DG1** in [Figure 72 on page 245](#)) to the two EX4200 or EX4300 switches (labeled **EX0** and **EX1** in [Figure 72 on page 245](#)). You connect the first two ports on the first network module to the first EX Series switch (**EX0**). You connect the first two ports on the second network module to the second EX Series switch (**EX1**). The ports used are the same on each EX Series switch.

Table 55: Second QFX3100 Director Device-to-EX Series Switch Control Plane Port Assignments for DG1

Network Module Port 0	Network Module Port 1	Network Module Port 3
EX Series switch port ge-0/0/22	EX Series switch port ge-0/0/23	Connect this port to the identical port on the other Director device. See "Connecting QFX3100 Director Devices in a Director Group" on page 242.

RELATED DOCUMENTATION

| [Installing and Connecting a QFX3100 Director Device](#) | 165

Connecting a QFX3100 Director Device to a Network for Out-of-Band Management

Use the management port on your QFX3100 Director device to connect each Director device in your Director group to your out-of-band management network.

NOTE: You cannot use the management port to perform the initial configuration of the QFX3100 Director device. You must configure the management port before you can successfully connect to the QFX3100 Director device using this port. See ["Performing the QFabric System Initial Setup on a QFX3100 Director Group"](#) on page 363.

Ensure that you have an RJ-45 patch cable available.

To connect a QFX3100 Director device to a network for out-of-band management:

1. Connect one end of the Ethernet cable to the management port (labeled **MGMT** on the Director device front panel).
2. Connect the other end of the Ethernet cable to your management device or management network.

3. Repeat these steps for the second Director device.

RELATED DOCUMENTATION

Management Port Connector Pinouts for the QFX Series

Cable Specifications for Console and Management Connections for the QFX Series

Connecting QFX3100 Director Devices to a Copper-Based QFX3000-G QFabric System Control Plane Network

[Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network | 245](#)

[Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 273](#)


[Connecting a QFX Series Device to a Management Console | 173](#)

Connecting a QFX5100-24Q Interconnect Device to a Copper-Based QFX3000-M QFabric System Control Plane Network

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 or two EX4300 switches. QFX5100-24Q Interconnect devices have two management ports (labeled **C0** and **C1**), to connect the QFX5100-24Q Interconnect device to each EX Series switch (see [Figure 73 on page 248](#)).

NOTE: You may not mix EX4200 and EX4300 switches in the same QFabric system.

Figure 73: QFX5100-24Q Interconnect Device Copper-Based Control Plane Connections for QFX3000-M QFabric System



Specific ports have been reserved on the EX Series switch to connect to the Interconnect devices, Node devices, and QFX3100 Director devices in your QFabric system. Such design simplifies installation and facilitates timely deployment of a QFabric system.

Before you begin to connect a QFX5100-24Q Interconnect device to the copper-based QFX3000-M QFabric system control plane and management network:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), [“Installing and Connecting a QFX3500 Device” on page 199](#), and [“Installing and Connecting a QFX5100 Device” on page 219](#).

NOTE: There can be a maximum of four Interconnect devices in the QFX3000-M QFabric system.

- Install your EX4200 or EX4300 switches. See *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.
- Interconnect the two EX Series switches using the SFP+ uplink module ports. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#).

NOTE: Use either EX4200-24T switches with an SFP+ uplink module installed or use an EX4300-48T with an SFP+ uplink module installed.

- Ensure that you have two RJ-45 patch cables available for each Interconnect device. For cable specifications, see [“Cable Specifications for Copper-Based Control Plane Connections for the QFabric System” on page 144](#).
- Use [Table 56 on page 250](#) to determine the QFX5100-24Q Interconnect device-to-EX Series switch port mappings. Specific ports have been reserved on the EX Series switch to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

NOTE: The numerical identifiers for each Interconnect device below are not preassigned to the Interconnect devices that are shipped to you. They represent the order in which you connect the Node devices. For example, the first Interconnect device (IC0) is connected to port 16 (ge-0/0/16) on the first EX Series switch **EX0** and the second EX Series switch **EX1**.

Table 56: QFX5100-24Q Interconnect Device-to-EX Series Switch Port Mappings for Copper-Based Control Plane

Interconnect Device	EX Series Switch EX0	EX Series Switch EX1
IC0	IC0, management port C0 on the QFX5100-24Q to port 16 (ge-0/0/16) on the EX Series switch	IC0, management port C1 on the QFX5100-24Q to port 16 (ge-0/0/16) on the EX Series switch
IC1	IC1, management port C0 on the QFX5100-24 to port 17 (ge-0/0/17) on the EX Series switch	IC1, management port C1 on the QFX5100-24 to port 17 (ge-0/0/17) on the EX Series switch
IC2	IC2, management port C0 on the QFX5100-24 to port 18 (ge-0/0/18) on the EX Series switch	IC2, management port C1 on the QFX5100-24 to port 18 (ge-0/0/18) on the EX Series switch
IC3	IC3, management port C0 on the QFX5100-24 to port 19 (ge-0/0/19) on the EX Series switch	IC3, management port C1 on the QFX5100-24 to port 19 (ge-0/0/19) on the EX Series switch

To connect each QFX5100-24Q Interconnect device to the copper-based control plane network:

1. Connect the first QFX5100-24Q Interconnect device.

- a. Connect one end of the first RJ-45 patch cable to the management port labeled **C0**.

NOTE: Some product SKUs have two ports labeled **C0**. The second **C0** port is for fiber connections.

- b. Connect the other end of that cable to port ge-0/0/16 on the first EX Series switch.
- c. Connect one end of the second RJ-45 patch cable to the management port labeled **C1**.
- d. Connect the other end of that cable to port ge-0/0/16 on the *second* EX Series switch.

2. Connect the second QFX5100-24Q Interconnect device.

- a. Connect one end of the first RJ-45 patch cable to the management port labeled **C0**.
- b. Connect the other end of that cable to port ge-0/0/17 on the first EX Series switch.
- c. Connect one end of the second RJ-45 patch cable to the management port labeled **C1**.
- d. Connect the other end of that cable to port ge-0/0/17 on the *second* EX Series switch.

3. Connect the third QFX5100-24Q Interconnect device if used.

- a. Connect one end of the first RJ-45 patch cable to the management port labeled **C0**.
- b. Connect the other end of that cable to port ge-0/0/18 on the first EX Series switch.
- c. Connect one end of the second RJ-45 patch cable to the management port labeled **C1**.
- d. Connect the other end of that cable to port ge-0/0/18 on the *second* EX Series switch.

4. Connect the fourth QFX5100-24Q Interconnect device if used.

- a. Connect one end of the first RJ-45 patch cable to the management port labeled **C0**.
- b. Connect the other end of that cable to port ge-0/0/19 on the first EX Series switch.

- c. Connect one end of the second RJ-45 patch cable to the management port labeled **C1**.
- d. Connect the other end of that cable to port ge-0/0/19 on the *second* EX Series switch.


RELATED DOCUMENTATION

Connecting a QFX3500 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network 263
Connecting a QFX3500 Node Device to an QFX5100-24Q Interconnect Device 313
Connecting a QFX3600 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network 259
Connecting a QFX3600 Node Device to a QFX5100-24Q Interconnect Device 309
Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network 245
Connecting QFX3100 Director Devices in a Director Group 242

Connecting a QFX3600-I Interconnect Device to a Copper-Based QFX3000-M QFabric System Control Plane Network

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 switches. QFX3600-I Interconnect devices have two 1000BASE-T management ports (labeled **C0** and **C1**) with RJ-45 connectors. For a copper-based control plane network, use the 1000BASE-T management ports to connect the QFX3600-I Interconnect device to each EX4200 switch (see [Figure 74 on page 252](#)).

Figure 74: QFX3600-I Interconnect Device Copper-Based Control Plane Connections for QFX3000-M QFabric System



Specific ports have been reserved on the EX4200 switch to connect to the Interconnect devices, Node devices, and QFX3100 Director devices in your QFabric system. Such design simplifies installation and facilitates timely deployment of a QFabric system.

Before you begin to connect a QFX3600-I Interconnect device to the copper-based QFX3000-M QFabric system control plane and management network:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).

NOTE: There can be a maximum of four Interconnect devices in the QFX3000-M QFabric system.

- Install your EX4200 switches. See *Installing and Connecting an EX4200 Switch*.

NOTE: For a copper-based control plane network, use EX4200-24T switches with an SFP+ uplink module installed.

- Interconnect the two EX4200 switches using the 1-Gigabit Ethernet SFP+ uplink module ports. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#).
- Ensure that you have two RJ-45 patch cables available for each Interconnect device. For cable specifications, see [“Cable Specifications for Copper-Based Control Plane Connections for the QFabric System” on page 144](#).
- Use [Table 57 on page 254](#) to determine the QFX3600-I Interconnect device-to-EX4200 switch port mappings. Specific ports have been reserved on the EX4200 switch to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

NOTE: The numerical identifiers for each Interconnect device below are not preassigned to the Interconnect devices that are shipped to you. They represent the order in which you connect the Node devices. For example, the first Interconnect device (IC0) is connected to port 16 (ge-0/0/16) on the first EX4200 switch **EX0** and the second EX4200 switch **EX1**.

Table 57: QFX3600-I Interconnect Device-to-EX4200 Switch Port Mappings for Copper-Based Control Plane

Interconnect Device	EX4200 Switch EX0	EX4200 Switch EX1
IC0	IC0, management port C0 to port 16 (ge-0/0/16)	IC0, management port C1 to port 16 (ge-0/0/16)
IC1	IC1, management port C0 to port 17 (ge-0/0/17)	IC1, management port C1 to port 17 (ge-0/0/17)
IC2	IC2, management port C0 to port 18 (ge-0/0/18)	IC2, management port C1 to port 18 (ge-0/0/18)
IC3	IC3, management port C0 to port 19 (ge-0/0/19)	IC3, management port C1 to port 19 (ge-0/0/19)

To connect each QFX3600-I Interconnect device to the copper-based control plane network:

1. Connect the first QFX3600-I Interconnect device.
 - a. Connect one end of the first RJ-45 patch cable to the management port labeled **C0**.
 - b. Connect the other end of that cable to port ge-0/0/16 on the first EX4200 switch.
 - c. Connect one end of the second RJ-45 patch cable to the management port labeled **C1**.
 - d. Connect the other end of that cable to port ge-0/0/16 on the *second* EX4200 switch.
2. Connect the second QFX3600-I Interconnect device.
 - a. Connect one end of the first RJ-45 patch cable to the management port labeled **C0**.
 - b. Connect the other end of that cable to port ge-0/0/17 on the first EX4200 switch.
 - c. Connect one end of the second RJ-45 patch cable to the management port labeled **C1**.
 - d. Connect the other end of that cable to port ge-0/0/17 on the *second* EX4200 switch.
3. Connect the third QFX3600-I Interconnect device if used.
 - a. Connect one end of the first RJ-45 patch cable to the management port labeled **C0**.
 - b. Connect the other end of that cable to port ge-0/0/18 on the first EX4200 switch.
 - c. Connect one end of the second RJ-45 patch cable to the management port labeled **C1**.
 - d. Connect the other end of that cable to port ge-0/0/18 on the *second* EX4200 switch.
4. Connect the fourth QFX3600-I Interconnect device if used.
 - a. Connect one end of the first RJ-45 patch cable to the management port labeled **C0**.
 - b. Connect the other end of that cable to port ge-0/0/19 on the first EX4200 switch.
 - c. Connect one end of the second RJ-45 patch cable to the management port labeled **C1**.
 - d. Connect the other end of that cable to port ge-0/0/19 on the *second* EX4200 switch.

RELATED DOCUMENTATION

[Connecting a QFX3500 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 263](#)

[Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device | 311](#)

[Connecting a QFX3600 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 259](#)

[Connecting a QFX3600 Node Device to a QFX3600-I Interconnect Device | 306](#)


[Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network | 245](#)

[Connecting QFX3100 Director Devices in a Director Group | 242](#)

Connecting a QFX5100 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 switches. QFX5100 Node devices have an RJ-45 **C0** located under the **CON** port, and an SFP cage **C1** that can be used as a copper management port. Use the RJ-45 management port (labeled **C0** and the SFP management port **C1**) to connect the QFX5100 Node device to each Virtual Chassis. See [Figure 75 on page 256](#) for an example of control plane cabling.

Figure 75: QFX5100 Node Device Copper-Based Control Plane Connections for QFX3000-M QFabric System



CAUTION: For redundancy, each QFX5100 Node device must be connected to two EX4200 switches.

Use the following QFX5100 Node devices and EX4200 switches for a copper-based QFX3000-M QFabric system control plane network:

- QFX5100 Node devices with:
 - em0-RJ-45 (1000BASE-T) management Ethernet port (**C0**)

- em1-SFP management Ethernet port (**C1**) cage (socket for 1 GbE copper SFP)
- EX4200-24T switches with an SFP+ uplink module installed.

Before you begin to connect a QFX3500 Node device to the copper-based QFX3000-M QFabric system control plane network:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#) and [“Installing and Connecting a QFX5100 Device” on page 219](#).
- Install your EX4200 switches. See *Installing and Connecting an EX4200 Switch*.
- Interconnect the two EX4200 switches using the 1-Gigabit Ethernet SFP+ uplink module ports. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#).
- Ensure that you have two RJ-45 patch cables available for each Node device. For cable specifications, see [“Cable Specifications for Copper-Based Control Plane Connections for the QFabric System” on page 144](#).
- Use [Table 58 on page 257](#) to determine the QFX5100 Node device-to-EX4200 switch port mappings. Specific ports have been reserved on the EX4200 switch to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

NOTE: The numerical identifiers for each Node device below are not preassigned to the Node devices that are shipped to you. They represent the order in which you connect the Node devices. For example, the first Node device (Node0) is connected to port 0 (ge-0/0/0) on the first EX4200 switch **EX0** and the second EX4200 switch **EX1**.

Table 58: QFX5100 Node Device-to-EX4200 Switch Port Mappings for Copper-Based Control Plane

Node Device	EX4200 Switch EX0	EX4200 Switch EX1
Node 0	Node0, management port C0 to port 0 (ge-0/0/0)	Node0, management port C1 to port 0 (ge-0/0/0)
Node 1	Node1, management port C0 to port 1 (ge-0/0/1)	Node1, management port C1 to port 1 (ge-0/0/1)
Node 2	Node2, management port C0 to port 2 (ge-0/0/2)	Node2, management port C1 to port 2 (ge-0/0/2)
Node 3	Node3, management port C0 to port 3 (ge-0/0/3)	Node3, management port C1 to port 3 (ge-0/0/3)

Table 58: QFX5100 Node Device-to-EX4200 Switch Port Mappings for Copper-Based Control Plane (continued)

Node Device	EX4200 Switch EX0	EX4200 Switch EX1
Node 4	Node4, management port C0 to port 4 (ge-0/0/4)	Node4, management port C1 to port 4 (ge-0/0/4)
Node 5	Node5, management port C0 to port 5 (ge-0/0/5)	Node5, management port C1 to port 5 (ge-0/0/5)
Node 6	Node6, management port C0 to port 6 (ge-0/0/6)	Node6, management port C1 to port 6 (ge-0/0/6)
Node 7	Node7, management port C0 to port 7 (ge-0/0/7)	Node7, management port C1 to port 7 (ge-0/0/7)
Node 8	Node8, management port C0 to port 8 (ge-0/0/8)	Node8, management port C1 to port 8 (ge-0/0/8)
Node 9	Node9, management port C0 to port 9 (ge-0/0/9)	Node9, management port C1 to port 9 (ge-0/0/9)
Node 10	Node10, management port C0 to port 10 (ge-0/0/10)	Node10, management port C1 to port 10 (ge-0/0/10)
Node 11	Node11, management port C0 to port 11 (ge-0/0/11)	Node11, management port C1 to port 11 (ge-0/0/11)
Node 12	Node12, management port C0 to port 12 (ge-0/0/12)	Node12, management port C1 to port 12 (ge-0/0/12)
Node 13	Node13, management port C0 to port 13 (ge-0/0/13)	Node13, management port C1 to port 13 (ge-0/0/13)
Node 14	Node14, management port C0 to port 14 (ge-0/0/14)	Node14, management port C1 to port 14 (ge-0/0/14)
Node 15	Node15, management port C0 to port 15 (ge-0/0/15)	Node15, management port C1 to port 15 (ge-0/0/15)

To connect a QFX5100 Node device to the copper-based QFX3000-M QFabric system control plane network:

1. Connect one end of the first RJ-45 patch cable to the first management port (labeled **C0**) on the Node device management panel.
2. Connect the other end of that cable to the appropriate port on the EX4200 switch. See [Table 58 on page 257](#).
3. Connect one end of the second RJ-45 patch cable to the second management port (labeled **C1**) on the Node device management panel.
4. Connect the other end of that cable to the appropriate port on the *second* EX4200 switch. This should be the same port number that you connected to in Step 2. For example, if you connected the first cable to ge-0/0/0 on the first EX4200 switch, you connect the second cable to ge-0/0/0 on the second EX4200 switch.
5. Repeat this procedure for each QFX5100 Node device.

RELATED DOCUMENTATION


Connecting a QFX3600-I Interconnect Device to a Copper-Based QFX3000-M QFabric System Control Plane Network 252
Connecting a QFX5100 Node Device to a QFX3600-I Interconnect Device 301
Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device 311
Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network 245
Connecting QFX3100 Director Devices in a Director Group 242

Connecting a QFX3600 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 or two EX4300 switches. QFX3600 Node devices have two 1000BASE-T management ports (labeled **C0** and **C1**) with RJ-45 connectors. For a copper-based control plane network, use the 1000BASE-T management ports to connect the QFX3600 Node device to each EX Series switch (see [Figure 76 on page 260](#)).

NOTE: You may not mix EX4200 and EX4300 switches in the same QFabric system.

Figure 76: QFX3600 Node Device Copper-Based Control Plane Connections for QFX3000-M QFabric System



CAUTION: For redundancy, each QFX3600 Node device must be connected to two EX Series switches.

Before you begin to connect a QFX3600 Node device to the copper-based QFX3000-M QFabric system control plane network:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).
- Install your EX4200 or EX4300 switches. See *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.
- Interconnect the two EX Series switches. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#).

NOTE: Use either an EX4200-24T switch with an SFP+ uplink module installed or use an EX4300-48T switch with an SFP+ uplink module installed.

- Ensure that you have two RJ-45 patch cables available for each Node device. For cable specifications, see [“Cable Specifications for Copper-Based Control Plane Connections for the QFabric System” on page 144](#).
- Use [Table 59 on page 261](#) to determine the QFX3600 Node device-to-EX Series switch port mappings. Specific ports have been reserved on the EX Series switch to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

NOTE: The numerical identifiers for each Node device below are not preassigned to the Node devices that are shipped to you. They represent the order in which you connect the Node devices. For example, the first Node device (Node0) is connected to port 0 (ge-0/0/0) on the first EX Series switch **EX0** and port 0 (ge-0/0/0) on the second EX Series switch **EX1**.

Table 59: QFX3600 Node Device-to-EX Series Switch Port Mappings for Copper-Based Control Plane

Node Device	EX Series Switch EX0	EX Series Switch EX1
Node 0	Node0, management port C0 to port 0 (ge-0/0/0)	Node0, management port C1 to port 0 (ge-0/0/0)
Node 1	Node1, management port C0 to port 1 (ge-0/0/1)	Node1, management port C1 to port 1 (ge-0/0/1)
Node 2	Node2, management port C0 to port 2 (ge-0/0/2)	Node2, management port C1 to port 2 (ge-0/0/2)
Node 3	Node3, management port C0 to port 3 (ge-0/0/3)	Node3, management port C1 to port 3 (ge-0/0/3)
Node 4	Node4, management port C0 to port 4 (ge-0/0/4)	Node4, management port C1 to port 4 (ge-0/0/4)
Node 5	Node5, management port C0 to port 5 (ge-0/0/5)	Node5, management port C1 to port 5 (ge-0/0/5)
Node 6	Node6, management port C0 to port 6 (ge-0/0/6)	Node6, management port C1 to port 6 (ge-0/0/6)
Node 7	Node7, management port C0 to port 7 (ge-0/0/7)	Node7, management port C1 to port 7 (ge-0/0/7)
Node 8	Node8, management port C0 to port 8 (ge-0/0/8)	Node8, management port C1 to port 8 (ge-0/0/8)
Node 9	Node9, management port C0 to port 9 (ge-0/0/9)	Node9, management port C1 to port 9 (ge-0/0/9)
Node 10	Node10, management port C0 to port 10 (ge-0/0/10)	Node10, management port C1 to port 10 (ge-0/0/10)

Table 59: QFX3600 Node Device-to-EX Series Switch Port Mappings for Copper-Based Control Plane (continued)

Node Device	EX Series Switch EX0	EX Series Switch EX1
Node 11	Node11, management port C0 to port 11 (ge-0/0/11)	Node11, management port C1 to port 11 (ge-0/0/11)
Node 12	Node12, management port C0 to port 12 (ge-0/0/12)	Node12, management port C1 to port 12 (ge-0/0/12)
Node 13	Node13, management port C0 to port 13 (ge-0/0/13)	Node13, management port C1 to port 13 (ge-0/0/13)
Node 14	Node14, management port C0 to port 14 (ge-0/0/14)	Node14, management port C1 to port 14 (ge-0/0/14)
Node 15	Node15, management port C0 to port 15 (ge-0/0/15)	Node15, management port C1 to port 15 (ge-0/0/15)

To connect a QFX3600 Node device to the copper-based QFX3000-M QFabric system control plane network (see [Figure 76 on page 260](#)):

1. Connect one end of the first RJ-45 patch cable to the first 1000BASE-T management port (labeled **C0**) on the Node device front panel.
2. Connect the other end of that cable to the appropriate port on the EX Series switch. See [Table 59 on page 261](#).
3. Connect one end of the second RJ-45 patch cable to the second 1000BASE-T management port (labeled **C1**) on the Node device front panel.
4. Connect the other end of that cable to the appropriate port on the *second* EX Series switch. This should be the same port number that you connected to in Step 2. For example, if you connected the first cable to ge-0/0/0 on the first EX Series switch, you connect the second cable to ge-0/0/0 on the second EX Series switch.
5. Repeat this procedure for each QFX3600 Node device.

RELATED DOCUMENTATION

[Connecting a QFX3600-I Interconnect Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 252](#)

[Connecting a QFX3600 Node Device to a QFX3600-I Interconnect Device | 306](#)

[Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network | 245](#)


[Connecting QFX3100 Director Devices in a Director Group | 242](#)

Connecting a QFX3500 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 or two EX4300 switches. QFX3500 Node devices have two management ports labeled **C0** and **C1**. Use the management ports to connect the QFX3500 Node device to each EX Series switch (see [Figure 77 on page 263](#)).

NOTE: You may not mix EX4200 and EX4300 switches in the same QFabric system.

Figure 77: QFX3500 Node Device Copper-Based Control Plane Connections for QFX3000-M QFabric System



CAUTION: For redundancy, each QFX3500 Node device must be connected to two EX Series switches.

Use the following QFX3500 Node devices and EX Series switches for a copper-based QFX3000-M QFabric system control plane network:

- QFX3500 Node devices with a 1000BASE-T management board installed. The 1000BASE-T management board provides two 1000BASE-T management ports (labeled **C0** and **C1**) with RJ-45 connectors.

- Either an EX4200-24T switch with an SFP+ uplink module installed or an EX4300-48T switch with an SFP+ uplink module installed.

Before you begin to connect a QFX3500 Node device to the copper-based QFX3000-M QFabric system control plane network:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).
- Install your EX4200 or EX4300 switches. See *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.
- Interconnect the two EX Series switches using the SFP+ uplink module. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#).
- Ensure that you have two RJ-45 patch cables available for each Node device. For cable specifications, see [“Cable Specifications for Copper-Based Control Plane Connections for the QFabric System” on page 144](#).
- Use [Table 60 on page 264](#) to determine the QFX3500 Node device-to-EX Series switch port mappings. Specific ports have been reserved on the EX Series switch to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

NOTE: The numerical identifiers for each Node device below are not preassigned to the Node devices that are shipped to you. They represent the order in which you connect the Node devices. For example, the first Node device (Node0) is connected to port 0 (ge-0/0/0) on the first EX Series switch **EX0** and the second EX Series switch **EX1**.

Table 60: QFX3500 Node Device-to-EX Series Switch Port Mappings for Copper-Based Control Plane

Node Device	EX Series Switch EX0	EX Series Switch EX1
Node 0	Node0, management port C0 to port 0 (ge-0/0/0)	Node0, management port C1 to port 0 (ge-0/0/0)
Node 1	Node1, management port C0 to port 1 (ge-0/0/1)	Node1, management port C1 to port 1 (ge-0/0/1)
Node 2	Node2, management port C0 to port 2 (ge-0/0/2)	Node2, management port C1 to port 2 (ge-0/0/2)
Node 3	Node3, management port C0 to port 3 (ge-0/0/3)	Node3, management port C1 to port 3 (ge-0/0/3)

Table 60: QFX3500 Node Device-to-EX Series Switch Port Mappings for Copper-Based Control Plane (continued)

Node Device	EX Series Switch EX0	EX Series Switch EX1
Node 4	Node4, management port C0 to port 4 (ge-0/0/4)	Node4, management port C1 to port 4 (ge-0/0/4)
Node 5	Node5, management port C0 to port 5 (ge-0/0/5)	Node5, management port C1 to port 5 (ge-0/0/5)
Node 6	Node6, management port C0 to port 6 (ge-0/0/6)	Node6, management port C1 to port 6 (ge-0/0/6)
Node 7	Node7, management port C0 to port 7 (ge-0/0/7)	Node7, management port C1 to port 7 (ge-0/0/7)
Node 8	Node8, management port C0 to port 8 (ge-0/0/8)	Node8, management port C1 to port 8 (ge-0/0/8)
Node 9	Node9, management port C0 to port 9 (ge-0/0/9)	Node9, management port C1 to port 9 (ge-0/0/9)
Node 10	Node10, management port C0 to port 10 (ge-0/0/10)	Node10, management port C1 to port 10 (ge-0/0/10)
Node 11	Node11, management port C0 to port 11 (ge-0/0/11)	Node11, management port C1 to port 11 (ge-0/0/11)
Node 12	Node12, management port C0 to port 12 (ge-0/0/12)	Node12, management port C1 to port 12 (ge-0/0/12)
Node 13	Node13, management port C0 to port 13 (ge-0/0/13)	Node13, management port C1 to port 13 (ge-0/0/13)
Node 14	Node14, management port C0 to port 14 (ge-0/0/14)	Node14, management port C1 to port 14 (ge-0/0/14)
Node 15	Node15, management port C0 to port 15 (ge-0/0/15)	Node15, management port C1 to port 15 (ge-0/0/15)

To connect a QFX3500 Node device to the copper-based QFX3000-M QFabric system control plane network (see [Figure 77 on page 263](#)):

1. Connect one end of the first RJ-45 patch cable to the first management port (labeled **C0**) on the Node device management board.
2. Connect the other end of that cable to the appropriate port on the EX Series switch. See [Table 60 on page 264](#).
3. Connect one end of the second RJ-45 patch cable to the second management port (labeled **C1**) on the Node device management board.
4. Connect the other end of that cable to the appropriate port on the *second* EX Series switch. This should be the same port number that you connected to in [Step 2](#). For example, if you connected the first cable to ge-0/0/0 on the first EX Series switch, you connect the second cable to ge-0/0/0 on the second EX Series switch.
5. Repeat this procedure for each QFX3500 Node device.

RELATED DOCUMENTATION

[Connecting a QFX3600-I Interconnect Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 252](#)

[Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device | 311](#)

[Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network | 245](#)

[Connecting QFX3100 Director Devices in a Director Group | 242](#)

Cabling a Fiber-Based Control Plane for the QFX3000-M QFabric System


IN THIS CHAPTER

- Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy | 267
- Connecting QFX3100 Director Devices in a Director Group | 270
- Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 273
- Connecting a QFX3100 Director Device to a Network for Out-of-Band Management | 275
- Connecting a QFX5100-24Q Interconnect Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 276
- Connecting a QFX3600-I Interconnect Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 281
- Connecting a QFX5100 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 286
- Connecting a QFX3500 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 290
- Connecting a QFX3600 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 295

Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 or to two EX4300 switches. For redundancy and communication, you must connect the two EX Series switches using the Ethernet SFP+ uplink module ports configured as a link aggregation group (LAG) (see [Figure 67 on page 240](#)).

Figure 78: QFX3000-M QFabric System Control Plane—Inter-EX Series Switch LAG Connections



For an EX4200 control plane, use the 1-Gigabit Ethernet SFP+ uplink module ports; for an EX4300 control plane, use the 10-Gigabit Ethernet SFP+ uplink module ports. Do not mix EX4200 and EX4300 switches in the same control plane.

Specific ports have been reserved on the EX Series switches to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

Before you begin to interconnect two EX Series switches for QFX3000-M QFabric system control plane redundancy:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). See [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).
- Install your EX4200 or EX4300 switches. See *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.

NOTE: For a copper-based control plane network, either use EX4200-24T switches with an SFP+ uplink module installed or use an EX4300-48T with an SFP+ uplink module installed. For a fiber-based control plane network, use EX4200-24F switches with an SFP+ uplink module installed or an EX4300-48P with an SFP+ uplink module installed.

- Install transceivers in ports **0** and **1** of the uplink module (see *Installing a Transceiver*).
 - For an EX4200 Switch control plane, ensure that you have installed 1-Gigabit Ethernet SFP transceivers in each EX4200 switch uplink module. For a list of supported transceivers, see *Pluggable Transceivers Supported on EX4200 Switches*.
 - For an EX4300 Switch control plane, ensure that you have installed 10-Gigabit Ethernet SFP+ transceivers in each EX4300 switch uplink module. For a list of supported transceivers, see *Pluggable Transceivers Supported on EX4300 Switches*.
- Ensure that you have appropriate fiber-optic cables (see *Pluggable Transceivers Supported on EX4200 Switches* and *Pluggable Transceivers Supported on EX4300 Switches*).
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings*).


To interconnect two EX Series switches for QFabric system control plane redundancy (see [Figure 67 on page 240](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
2. Remove the rubber safety cap from the SFP optical transceiver in port **0** of the first EX Series switch uplink module. Save the cap.
3. Insert the cable connector into the optical transceiver (see [Figure 68 on page 241](#)).

Figure 79: Connecting a Fiber-Optic Cable to an Optical Transceiver Installed in an EX Series Switch



4. If the connector at the other end of the fiber-optic cable is covered by a rubber safety cap, remove the cap. Save the cap.
5. Remove the rubber safety cap from the SFP optical transceiver in port **0** of the *second* EX Series switch uplink module. Save the cap.
6. Insert the cable connector into the optical transceiver.
7. Repeat Step [1](#) through Step [6](#) to interconnect uplink module port **1** on both the EX Series switches.
8. Secure the cables so that they are not supporting their own weight. Place excess cable out of the way in a neatly coiled loop. Placing fasteners on a loop helps cables maintain their shape.



CAUTION: Do not bend fiber-optic cables beyond their minimum bend radius. An arc smaller than a few inches in diameter can damage the cables and cause problems that are difficult to diagnose.

Do not let fiber-optic cables hang free from the connector. Do not allow fastened loops of cables to dangle, which stresses the cables at the fastening point.

RELATED DOCUMENTATION

[Installing and Connecting a QFX3100 Director Device | 165](#)

[Installing and Connecting a QFX3600 or QFX3600-I Device | 177](#)

[Installing and Connecting a QFX3500 Device | 199](#)

Connecting QFX3100 Director Devices in a Director Group

A QFabric system requires two QFX3100 Director devices interconnected as a QFX3100 Director *group* (see [Figure 69 on page 242](#) through [Figure 71 on page 243](#)).

Figure 80: QFX3100 Director Group Control Plane Connections for QFX3000-G QFabric System Using Copper-Based Control Plane




Figure 81: QFX3100 Director Group Control Plane Connections for QFX3000-G QFabric System Using Fiber-Based Control Plane





Figure 82: QFX3100 Director Group Control Plane Connections for QFX3000-M QFabric System



The second QFX3100 Director device provides redundancy for the control plane and management network.

Before you begin to connect QFX3100 Director devices in a Director group:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [Installing and Connecting a QFX3008-I Interconnect Device](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).

NOTE: For a copper-based QFX3000-M QFabric system control plane network, use QFX3100 Director devices with RJ-45 network modules installed. For a fiber-based control plane network, use QFX3100 Director devices with SFP network modules installed.

- Ensure that you have appropriate transceivers and cables available. For cable specifications, see [“Cable Specifications for Copper-Based Control Plane Connections for the QFabric System” on page 144](#) and [“Determining Transceiver Support for QFabric Systems” on page 143](#).



CAUTION: 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. Although a single inter-Director device can lose a link and regain its connection, the loss of both inter-Director device links causes one of the Director devices to become isolated from the Director group. In [Figure 69 on page 242](#) through [Figure 71 on page 243](#), these redundant patch cables are shown in red and connect port 3 to port 3.

To connect QFX3100 Director devices in a Director group (see [Figure 69 on page 242](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. Connect the port labeled **3** on the first network module on one of the Director devices to the corresponding port (labeled **3**) on the first network module on the second Director device.
2. Connect the port labeled **3** on the *second* network module on one of the Director devices to the corresponding port (labeled **3**) on the *second* network module on the second Director device.

RELATED DOCUMENTATION

Connecting QFX3100 Director Devices to a Copper-Based QFX3000-G QFabric System Control Plane Network


Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-G QFabric System Control Plane Network

[Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network | 245](#)

Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-M QFabric System Control Plane Network

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 switches or to two EX4300 switches. QFX3100 Director devices have two 4-port network modules. Use the network module ports to connect the QFX3100 Director group to the EX Series switches (see [Figure 83 on page 273](#)).

Figure 83: QFX3100 Director Group to EX Series Switch Connections for QFX3000-M QFabric System



Use the following QFX3100 Director devices and EX Series switches for a fiber-based QFX3000-M QFabric system control plane network:

- QFX3100 Director devices with SFP network modules installed. Each SFP network module provides four SFP ports labeled 0 through 3.
- Either an EX4200-24F switch with an SFP+ uplink module installed or an EX4300-48P with an SFP+ uplink module installed.

Specific ports have been reserved on the EX Series switches to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

Before you begin to connect a QFX3100 Director device to the fiber-based QFX3000-M QFabric system control plane network:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). See [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).
- Install your EX4200 or EX4300 switches. See *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.
- Interconnect the two EX Series switches using the SFP+ uplink module ports. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#)
- Connect the two QFX3100 Director devices to create a Director group. See [“Connecting QFX3100 Director Devices in a Director Group” on page 242](#).
- Ensure that you have installed transceivers in the ports labeled **0** and **1** on both the network modules on each QFX3100 Director device (see *Installing a Transceiver in a QFX Series Device*). For a list of supported transceivers and required cables, see [The Hardware Compatibility Tool](#).
- Install transceivers in ports **0** and **1** of the uplink module (see *Installing a Transceiver*).
 - For an EX4200 Switch control plane, ensure that you have installed 1-Gigabit Ethernet SFP transceivers in each EX4200 switch uplink module. For a list of supported transceivers, see [The Hardware Compatibility Tool](#).
 - For an EX4300 Switch control plane, ensure that you have installed 10-Gigabit Ethernet SFP+ transceivers in each EX4300 switch uplink module. For a list of supported transceivers, see [The Hardware Compatibility Tool](#).
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings for the QFX Series*).

To connect a QFX3100 Director device to the fiber-based QFX3000-M QFabric control plane network (see [Figure 83 on page 273](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. Connect both network modules on the first Director device (labeled **DG0** in [Figure 83 on page 273](#)) to the two EX Series switches (labeled **EX0** and **EX1** in [Figure 83 on page 273](#)). You connect the first two ports (labeled **0** and **1**) on the first network module to the first EX Series switch (**EX0**). You connect the first two ports on the second network module (also labeled **0** and **1**) to the second EX Series switch (**EX1**). The ports used are the same on each switch.

Table 61: QFX3100 Director Device-to-EX Series Switch Control Plane Port Assignments for DG0

Network Module Port 0	Network Module Port 1	Network Module Port 3
EX Series switch port ge-0/0/20	EX Series switch port ge-0/0/21	Connect this port to the identical port on the other Director device. See “Connecting QFX3100 Director Devices in a Director Group” on page 242.

2. Connect both network modules on the *second* Director device (labeled **DG1** in [Figure 83 on page 273](#)) to the two EX Series switches (labeled **EX0** and **EX1** in [Figure 83 on page 273](#)). You connect the first two ports on the first network module to the first EX Series switch (**EX0**). You connect the first two ports on the second network module to the second EX Series switch (**EX1**). The ports used are the same on each switch.

Table 62: Second QFX3100 Director Device-to-EX Series Switch Control Plane Port Assignments for DG1

Network Module Port 0	Network Module Port 1	Network Module Port 3
EX Series switch port ge-0/0/22	EX Series switch port ge-0/0/23	Connect this port to the identical port on the other Director device. See “Connecting QFX3100 Director Devices in a Director Group” on page 242.

RELATED DOCUMENTATION

| [Installing and Connecting a QFX3100 Director Device](#) | 165

Connecting a QFX3100 Director Device to a Network for Out-of-Band Management

Use the management port on your QFX3100 Director device to connect each Director device in your Director group to your out-of-band management network.

NOTE: You cannot use the management port to perform the initial configuration of the QFX3100 Director device. You must configure the management port before you can successfully connect to the QFX3100 Director device using this port. See [“Performing the QFabric System Initial Setup on a QFX3100 Director Group”](#) on page 363.

Ensure that you have an RJ-45 patch cable available.

To connect a QFX3100 Director device to a network for out-of-band management:

1. Connect one end of the Ethernet cable to the management port (labeled **MGMT** on the Director device front panel).
2. Connect the other end of the Ethernet cable to your management device or management network.
3. Repeat these steps for the second Director device.

RELATED DOCUMENTATION


Management Port Connector Pinouts for the QFX Series
Cable Specifications for Console and Management Connections for the QFX Series
Connecting QFX3100 Director Devices to a Copper-Based QFX3000-G QFabric System Control Plane Network
Connecting QFX3100 Director Devices to a Copper-Based QFX3000-M QFabric System Control Plane Network 245
Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-M QFabric System Control Plane Network 273
Connecting a QFX Series Device to a Management Console 173

Connecting a QFX5100-24Q Interconnect Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 or two EX4300 switches. QFX5100-24Q Interconnect devices have two management ports (labeled **C0** and **C1**), to connect the QFX5100-24Q Interconnect device to each EX Series switch (see [Figure 84 on page 277](#)).

NOTE: You may not mix EX4200 and EX4300 switches in the same QFabric system.

Figure 84: QFX5100-24Q Interconnect Device Fiber-Based Control Plane Connections for QFX3000-M QFabric System



Specific ports have been reserved on the EX Series switch to connect to the Interconnect devices, Node devices, and QFX3100 Director devices in your QFabric system. Such design simplifies installation and facilitates timely deployment of a QFabric system.

Before you begin to connect a QFX5100-24Q Interconnect device to the fiber-based QFX3000-M QFabric system control plane and management network:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#), and [“Installing and Connecting a QFX5100 Device” on page 219](#).

NOTE: There can be a maximum of four Interconnect devices in the QFX3000-M QFabric system.

- Install your EX4200 or EX4300 switches. See *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.

NOTE: For a fiber-based control plane network, use an EX4200-24F switch or an EX4300-48P switch with an SFP+ uplink module installed.

- Interconnect the two EX Series switches using the SFP+ uplink module ports. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#).
- Ensure that you have installed 1-Gigabit Ethernet SFP transceivers in management ports **C0** and **C1** on each QFX5100-24Q Interconnect device (see *Installing a Transceiver in a QFX Series Device*). For a list of supported transceivers and required cables, see [“Determining Transceiver Support for QFabric Systems” on page 143](#).

- Ensure that you have installed 1-Gigabit Ethernet SFP transceivers in the ports (labeled **16** through **19**) you are using on each EX Series switch (see *Installing a Transceiver*). For a list of supported transceivers, see [“Determining Transceiver Support for QFabric Systems” on page 143](#).
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings for the QFX Series*).
- Use [Table 63 on page 278](#) to determine the QFX5100-24Q Interconnect device-to-EX Series switch port mappings. Specific ports have been reserved on the EX Series switch to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

NOTE: The numerical identifiers for each Interconnect device below are not preassigned to the Interconnect devices that are shipped to you. They represent the order in which you connect the Node devices. For example, the first Interconnect device (IC0) is connected to port 16 (ge-0/0/16) on the first EX Series switch **EX0** and the second EX Series switch **EX1**.

Table 63: QFX3600-I Interconnect Device-to-EX Series Switch Port Mappings for Fiber-Based Control Plane

Interconnect Device	EX Series Switch EX0	EX Series Switch EX1
IC0	IC0, management port C0 to port 16 (ge-0/0/16)	IC0, management port C1 to port 16 (ge-0/0/16)
IC1	IC1, management port C0 to port 17 (ge-0/0/17)	IC1, management port C1 to port 17 (ge-0/0/17)
IC2	IC2, management port C0 to port 18 (ge-0/0/18)	IC2, management port C1 to port 18 (ge-0/0/18)
IC3	IC3, management port C0 to port 19 (ge-0/0/19)	IC3, management port C1 to port 19 (ge-0/0/19)

To connect each QFX5100-24Q Interconnect device to the fiber-based control plane network:



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. Connect the first QFX5100-24Q Interconnect device.
 - a. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
 - b. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C0**. Save the cap.
 - c. Insert the cable connector into the optical transceiver.
 - d. Connect the other end of that cable to port ge-0/0/16 on the first EX Series switch.
 - e. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C1**. Save the cap.
 - f. Insert the cable connector into the optical transceiver.
 - g. Connect the other end of that cable to port ge-0/0/16 on the *second* EX Series switch.
2. Connect the second QFX5100-24Q Interconnect device.
 - a. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
 - b. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C0**. Save the cap.
 - c. Insert the cable connector into the optical transceiver.
 - d. Connect the other end of that cable to port ge-0/0/17 on the first EX Series switch.
 - e. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C1**. Save the cap.
 - f. Insert the cable connector into the optical transceiver.

- g. Connect the other end of that cable to port ge-0/0/17 on the *second* EX Series switch.
3. Connect the third QFX5100-24Q Interconnect device if used.
- a. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
 - b. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C0**. Save the cap.
 - c. Insert the cable connector into the optical transceiver.
 - d. Connect the other end of that cable to port ge-0/0/18 on the first EX Series switch.
 - e. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C1**. Save the cap.
 - f. Insert the cable connector into the optical transceiver.
 - g. Connect the other end of that cable to port ge-0/0/18 on the *second* EX Series switch.
4. Connect the fourth QFX5100-24Q Interconnect device if used.
- a. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
 - b. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C0**. Save the cap.
 - c. Insert the cable connector into the optical transceiver.
 - d. Connect the other end of that cable to port ge-0/0/19 on the first EX Series switch.
 - e. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C1**. Save the cap.
 - f. Insert the cable connector into the optical transceiver.
 - g. Connect the other end of that cable to port ge-0/0/19 on the *second* EX Series switch.

RELATED DOCUMENTATION

[Connecting a QFX3500 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 290](#)

[Connecting a QFX3500 Node Device to an QFX5100-24Q Interconnect Device | 313](#)

[Connecting a QFX3600 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 295](#)

[Connecting a QFX3600 Node Device to a QFX5100-24Q Interconnect Device | 309](#)

[Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 273](#)

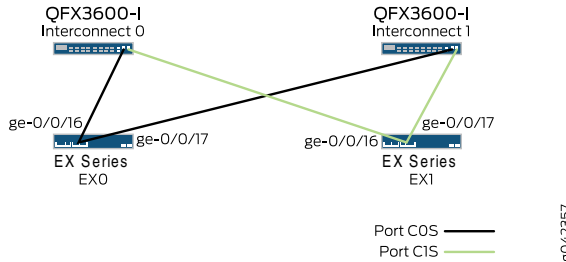
[Connecting QFX3100 Director Devices in a Director Group | 242](#)

Connecting a QFX3600-I Interconnect Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 or two EX4300 switches. QFX3600-I Interconnect devices have two SFP management ports labeled **C0S** and **C1S**. For a fiber-based control plane network, use the SFP management ports to connect the QFX3600-I Interconnect device to each EX Series switch (see [Figure 85 on page 281](#)).

NOTE: You may not mix EX4200 and EX4300 switches in the same QFabric system.

Figure 85: QFX3600-I Interconnect Device Fiber-Based Control Plane Connections for QFX3000-M QFabric System



Specific ports have been reserved on the EX Series switch to connect to the Interconnect devices, Node devices, and QFX3100 Director devices in your QFabric system. Such design simplifies installation and facilitates timely deployment of a QFabric system.

Before you begin to connect a QFX3600-I Interconnect device to the fiber-based QFX3000-M QFabric system control plane and management network:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).

NOTE: There can be a maximum of four Interconnect devices in the QFX3000-M QFabric system.

- Install your EX4200 or EX4300 switches. See *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.

NOTE: For a fiber-based control plane network, use an EX4200-24F switch or an EX4300-48P switch with an SFP+ uplink module installed.

- Interconnect the two EX Series switches using the SFP+ uplink module ports. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#).
- Ensure that you have installed 1-Gigabit Ethernet SFP transceivers in management ports **C0S** and **C1S** on each QFX3600-I Interconnect device (see *Installing a Transceiver in a QFX Series Device*). For a list of supported transceivers and required cables, see [The Hardware Compatibility Tool](#).
- Ensure that you have installed 1-Gigabit Ethernet SFP transceivers in the ports (labeled **16** through **19**) you are using on each EX Series switch (see *Installing a Transceiver*). For a list of supported transceivers, see [The Hardware Compatibility Tool](#).
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings for the QFX Series*).
- Use [Table 64 on page 283](#) to determine the QFX3600-I Interconnect device-to-EX Series switch port mappings. Specific ports have been reserved on the EX Series switch to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

NOTE: The numerical identifiers for each Interconnect device below are not preassigned to the Interconnect devices that are shipped to you. They represent the order in which you connect the Node devices. For example, the first Interconnect device (IC0) is connected to port 16 (ge-0/0/16) on the first EX Series switch **EX0** and the second EX Series switch **EX1**.

Table 64: QFX3600-I Interconnect Device-to-EX Series Switch Port Mappings for Fiber-Based Control Plane

Interconnect Device	EX Series Switch EX0	EX Series Switch EX1
IC0	IC0, management port C0S to port 16 (ge-0/0/16)	IC0, management port C1S to port 16 (ge-0/0/16)
IC1	IC1, management port C0S to port 17 (ge-0/0/17)	IC1, management port C1S to port 17 (ge-0/0/17)
IC2	IC2, management port C0S to port 18 (ge-0/0/18)	IC2, management port C1S to port 18 (ge-0/0/18)
IC3	IC3, management port C0S to port 19 (ge-0/0/19)	IC3, management port C1S to port 19 (ge-0/0/19)

To connect each QFX3600-I Interconnect device to the fiber-based control plane network:



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. Connect the first QFX3600-I Interconnect device.
 - a. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
 - b. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C0S**. Save the cap.
 - c. Insert the cable connector into the optical transceiver.
 - d. Connect the other end of that cable to port ge-0/0/16 on the first EX Series switch.
 - e. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C1S**. Save the cap.
 - f. Insert the cable connector into the optical transceiver.
 - g. Connect the other end of that cable to port ge-0/0/16 on the *second* EX Series switch.
2. Connect the second QFX3600-I Interconnect device.
 - a. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
 - b. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C0S**. Save the cap.
 - c. Insert the cable connector into the optical transceiver.
 - d. Connect the other end of that cable to port ge-0/0/17 on the first EX Series switch.
 - e. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C1S**. Save the cap.
 - f. Insert the cable connector into the optical transceiver.

- g. Connect the other end of that cable to port ge-0/0/17 on the *second* EX Series switch.
3. Connect the third QFX3600-I Interconnect device if used.
- a. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
 - b. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C0S**. Save the cap.
 - c. Insert the cable connector into the optical transceiver.
 - d. Connect the other end of that cable to port ge-0/0/18 on the first EX Series switch.
 - e. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C1S**. Save the cap.
 - f. Insert the cable connector into the optical transceiver.
 - g. Connect the other end of that cable to port ge-0/0/18 on the *second* EX Series switch.
4. Connect the fourth QFX3600-I Interconnect device if used.
- a. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
 - b. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C0S**. Save the cap.
 - c. Insert the cable connector into the optical transceiver.
 - d. Connect the other end of that cable to port ge-0/0/19 on the first EX Series switch.
 - e. Remove the rubber safety cap from the SFP optical transceiver in the management port labeled **C1S**. Save the cap.
 - f. Insert the cable connector into the optical transceiver.
 - g. Connect the other end of that cable to port ge-0/0/19 on the *second* EX Series switch.

RELATED DOCUMENTATION

[Connecting a QFX3500 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 290](#)

[Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device | 311](#)

[Connecting a QFX3600 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 295](#)

[Connecting a QFX3600 Node Device to a QFX3600-I Interconnect Device | 306](#)

[Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 273](#)


[Connecting QFX3100 Director Devices in a Director Group | 242](#)

Connecting a QFX5100 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 or two EX4300 switches. QFX5100 Node devices have an RJ-45 **C0** located under the **CON** port, and a small form-factor pluggable (SFP) cage **C1** that can be used as a copper management port. Use the RJ-45 management port (labeled **C0** and the SFP management port **C1**) to connect the QFX5100 Node device to each EX Series switch. See [Figure 86 on page 286](#) for an example of this cabling.

NOTE: You may not mix EX4200 and EX4300 switches in the same QFabric system.

Figure 86: QFX5100 Node Device Fiber-Based Control Plane Connections for QFX3000-M QFabric System



CAUTION: For redundancy, each QFX5100 Node device must be connected to two EX Series switches.

Use the following QFX5100 Node devices and EX Series switches for a fiber-based QFX3000-M QFabric system control plane network:

- QFX5100-48S Node devices with two 1 Gbps SFP modules installed in the management ports labeled **C0** and **C1**.
- Either an EX4200-24F switches with an SFP+ uplink module installed or an EX4300-48P with an SFP+ uplink module installed.

Before you begin to connect a QFX5100 Node device to the fiber-based QFX3000-M QFabric system control plane network:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#), and [“Installing and Connecting a QFX5100 Device” on page 219](#).
- Install your EX4200 or EX4300 switches. See *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.
- Interconnect the two EX Series switches. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#).

NOTE: Use either EX4200-24F switches with an SFP+ uplink module installed or use an EX4300-48P with an SFP+ uplink module installed.

- Ensure that you have installed 1-Gigabit Ethernet SFP transceivers in management ports **C0** and **C1** on each QFX5100 Node device (see *Installing a Transceiver in a QFX Series Device*). For a list of supported transceivers and required cables, see [The Hardware Compatibility Tool](#).
- Follow the port assignments in [Table 63 on page 278](#) and ensure that you have installed 1-Gigabit Ethernet SFP transceivers in the ports (labeled **0** through **15**) you are using on each EX4200 switch (see *Installing a Transceiver*). For a list of supported transceivers, see [The Hardware Compatibility Tool](#).
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings for the QFX Series*).
- Use [Table 63 on page 278](#) to determine the QFX5100 Node device-to-EX Series switch port mappings. Specific ports have been reserved on the EX Series switch to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

NOTE: The numerical identifiers for each Node device below are not preassigned to the Node devices that are shipped to you. They represent the order in which you connect the Node devices. For example, the first Node device (Node0) is connected to port 0 (ge-0/0/0) on the first EX Series switch **EX0** and port 0 (ge-0/0/0) on the second EX Series switch **EX1**.

Table 65: QFX5100 Node Device-to-EX Series Switch Port Mappings for Fiber-Based Control Plane

Node Device	EX Series Switch EX0	EX Series Switch EX1
Node 0	Node0, management port C0 to port 0 (ge-0/0/0)	Node0, management port C1 to port 0 (ge-0/0/0)
Node 1	Node1, management port C0 to port 1 (ge-0/0/1)	Node1, management port C1 to port 1 (ge-0/0/1)
Node 2	Node2, management port C0 to port 2 (ge-0/0/2)	Node2, management port C1 to port 2 (ge-0/0/2)
Node 3	Node3, management port C0 to port 3 (ge-0/0/3)	Node3, management port C1 to port 3 (ge-0/0/3)
Node 4	Node4, management port C0 to port 4 (ge-0/0/4)	Node4, management port C1 to port 4 (ge-0/0/4)
Node 5	Node5, management port C0 to port 5 (ge-0/0/5)	Node5, management port C1 to port 5 (ge-0/0/5)
Node 6	Node6, management port C0 to port 6 (ge-0/0/6)	Node6, management port C1 to port 6 (ge-0/0/6)
Node 7	Node7, management port C0 to port 7 (ge-0/0/7)	Node7, management port C1 to port 7 (ge-0/0/7)
Node 8	Node8, management port C0 to port 8 (ge-0/0/8)	Node8, management port C1 to port 8 (ge-0/0/8)
Node 9	Node9, management port C0 to port 9 (ge-0/0/9)	Node9, management port C1 to port 9 (ge-0/0/9)
Node 10	Node10, management port C0 to port 10 (ge-0/0/10)	Node10, management port C1 to port 10 (ge-0/0/10)

Table 65: QFX5100 Node Device-to-EX Series Switch Port Mappings for Fiber-Based Control Plane (continued)

Node Device	EX Series Switch EX0	EX Series Switch EX1
Node 11	Node11, management port C0 to port 11 (ge-0/0/11)	Node11, management port C1 to port 11 (ge-0/0/11)
Node 12	Node12, management port C0 to port 12 (ge-0/0/12)	Node12, management port C1 to port 12 (ge-0/0/12)
Node 13	Node13, management port C0 to port 13 (ge-0/0/13)	Node13, management port C1 to port 13 (ge-0/0/13)
Node 14	Node14, management port C0 to port 14 (ge-0/0/14)	Node14, management port C1 to port 14 (ge-0/0/14)
Node 15	Node15, management port C0 to port 15 (ge-0/0/15)	Node15, management port C1 to port 15 (ge-0/0/15)

To connect a QFX5100 Node device to the fiber-based QFX3000-M QFabric system control plane network (see [Figure 86 on page 286](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
2. Remove the rubber safety cap from the SFP optical transceiver in the first management port (labeled **C0**) on the Node device management board. Save the cap.
3. Insert the cable connector into the optical transceiver.
4. Connect the other end of that cable to the appropriate port on the EX Series switch. See [Table 63 on page 278](#).
5. Remove the rubber safety cap from the SFP optical transceiver in the second management port (labeled **C1**) on the Node device management board. Save the cap.
6. Insert the cable connector into the optical transceiver.

7. Connect the other end of that cable to the appropriate port on the *second* EX4200 switch. This should be the same port number that you connected to in Step 4. For example, if you connected the first cable to ge-0/0/0 on the first EX Series switch, you connect the second cable to ge-0/0/0 on the second EX Series switch.
8. Repeat this procedure for each QFX5100 Node device.
9. Secure the cables so that they are not supporting their own weight. Place excess cable out of the way in a neatly coiled loop. Placing fasteners on a loop helps cables maintain their shape.



CAUTION: Do not bend fiber-optic cables beyond their minimum bend radius. An arc smaller than a few inches in diameter can damage the cables and cause problems that are difficult to diagnose.

Do not let fiber-optic cables hang free from the connector. Do not allow fastened loops of cables to dangle, which stresses the cables at the fastening point.

RELATED DOCUMENTATION

[Connecting a QFX3600-I Interconnect Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 281](#)

[Connecting a QFX5100 Node Device to a QFX3600-I Interconnect Device | 301](#)

[Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 273](#)


[Connecting QFX3100 Director Devices in a Director Group | 242](#)

Connecting a QFX3500 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 or two EX4300 switches. QFX3500 Node devices have two management ports labeled **C0** and **C1**. Use the management ports to connect the QFX3500 Node device to each EX Series switch (see [Figure 87 on page 291](#)).

NOTE: You may not mix EX4200 and EX4300 switches in the same QFabric system.

Figure 87: QFX3500 Node Device Fiber-Based Control Plane Connections for QFX3000-M QFabric System



CAUTION: For redundancy, each QFX3500 Node device must be connected to two EX Series switches.

Use the following QFX3500 Node devices and EX Series switches for a fiber-based QFX3000-M QFabric system control plane network:

- QFX3500 Node devices with an SFP management board installed. The SFP management board provides two small form-factor pluggable (SFP) 1-Gbps management ports labeled **C0** and **C1**.
- Either an EX4200-24F switch with an SFP+ uplink module installed or an EX4300-48P switch with an SFP+ uplink module installed.

Before you begin to connect a QFX3500 Node device to the fiber-based QFX3000-M QFabric system control plane network:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).
- Install your EX4200 or EX4300 switches. See *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.
- Interconnect the two EX Series switches. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#).
- Ensure that you have installed 1-Gigabit Ethernet SFP transceivers in management ports **C0** and **C1** on each QFX3500 Node device (see *Installing a Transceiver in a QFX Series Device*). For a list of supported transceivers and required cables, see [The Hardware Compatibility Tool](#).

- Follow the port assignments in [Table 66 on page 292](#) and ensure that you have installed 1-Gigabit Ethernet SFP transceivers in the ports (labeled **0** through **15**) you are using on each EX Series switch (see *Installing a Transceiver*). For a list of supported transceivers, see [The Hardware Compatibility Tool](#).
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings for the QFX Series*).
- Use [Table 66 on page 292](#) to determine the QFX3500 Node device-to-EX4200 switch port mappings. Specific ports have been reserved on the EX Series switch to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

NOTE: The numerical identifiers for each Node device below are not preassigned to the Node devices that are shipped to you. They represent the order in which you connect the Node devices. For example, the first Node device (Node0) is connected to port 0 (ge-0/0/0) on the first EX Series switch **EX0** and port 0 (ge-0/0/0) on the second EX Series switch **EX1**.

Table 66: QFX3500 Node Device-to-EX Series Switch Port Mappings for Fiber-Based Control Plane

Node Device	EX Series Switch EX0	EX Series Switch EX1
Node0	Node0, management port C0 to port 0 (ge-0/0/0)	Node0, management port C1 to port 0 (ge-0/0/0)
Node1	Node1, management port C0 to port 1 (ge-0/0/1)	Node1, management port C1 to port 1 (ge-0/0/1)
Node2	Node2, management port C0 to port 2 (ge-0/0/2)	Node2, management port C1 to port 2 (ge-0/0/2)
Node3	Node3, management port C0 to port 3 (ge-0/0/3)	Node3, management port C1 to port 3 (ge-0/0/3)
Node4	Node4, management port C0 to port 4 (ge-0/0/4)	Node4, management port C1 to port 4 (ge-0/0/4)
Node5	Node5, management port C0 to port 5 (ge-0/0/5)	Node5, management port C1 to port 5 (ge-0/0/5)
Node6	Node6, management port C0 to port 6 (ge-0/0/6)	Node6, management port C1 to port 6 (ge-0/0/6)
Node7	Node7, management port C0 to port 7 (ge-0/0/7)	Node7, management port C1 to port 7 (ge-0/0/7)

Table 66: QFX3500 Node Device-to-EX Series Switch Port Mappings for Fiber-Based Control Plane (continued)

Node Device	EX Series Switch EX0	EX Series Switch EX1
Node8	Node8, management port C0 to port 8 (ge-0/0/8)	Node8, management port C1 to port 8 (ge-0/0/8)
Node9	Node9, management port C0 to port 9 (ge-0/0/9)	Node9, management port C1 to port 9 (ge-0/0/9)
Node10	Node10, management port C0 to port 10 (ge-0/0/10)	Node10, management port C1 to port 10 (ge-0/0/10)
Node11	Node11, management port C0 to port 11 (ge-0/0/11)	Node11, management port C1 to port 11 (ge-0/0/11)
Node12	Node12, management port C0 to port 12 (ge-0/0/12)	Node12, management port C1 to port 12 (ge-0/0/12)
Node13	Node13, management port C0 to port 13 (ge-0/0/13)	Node13, management port C1 to port 13 (ge-0/0/13)
Node14	Node14, management port C0 to port 14 (ge-0/0/14)	Node14, management port C1 to port 14 (ge-0/0/14)
Node15	Node15, management port C0 to port 15 (ge-0/0/15)	Node15, management port C1 to port 15 (ge-0/0/15)

To connect a QFX3500 Node device to the fiber-based QFX3000-M QFabric system control plane network (see [Figure 87 on page 291](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
2. Remove the rubber safety cap from the SFP optical transceiver in the first management port (labeled **C0**) on the Node device management board. Save the cap.
3. Insert the cable connector into the optical transceiver.

4. Connect the other end of that cable to the appropriate port on the EX Series switch. See [Table 66 on page 292](#).
5. Remove the rubber safety cap from the SFP optical transceiver in the second management port (labeled **C1**) on the Node device management board. Save the cap.
6. Insert the cable connector into the optical transceiver.
7. Connect the other end of that cable to the appropriate port on the *second* EX Series switch. This should be the same port number that you connected to in Step 4. For example, if you connected the first cable to ge-0/0/0 on the first EX Series switch, you connect the second cable to ge-0/0/0 on the second EX Series switch.
8. Repeat this procedure for each QFX3500 Node device.
9. Secure the cables so that they are not supporting their own weight. Place excess cable out of the way in a neatly coiled loop. Placing fasteners on a loop helps cables maintain their shape.



CAUTION: Do not bend fiber-optic cables beyond their minimum bend radius. An arc smaller than a few inches in diameter can damage the cables and cause problems that are difficult to diagnose.

Do not let fiber-optic cables hang free from the connector. Do not allow fastened loops of cables to dangle, which stresses the cables at the fastening point.

RELATED DOCUMENTATION

[Connecting a QFX3600-I Interconnect Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 281](#)

[Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device | 311](#)

[Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 273](#)


[Connecting QFX3100 Director Devices in a Director Group | 242](#)

Connecting a QFX3600 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network

A QFX3000-M QFabric system control plane and management network is formed by connecting the QFX Series devices in your network to two EX4200 or two EX4300 switches. QFX3600 Node devices have two SFP management ports labeled **C0S** and **C1S**. For a fiber-based control plane network, use the SFP management ports to connect the QFX3600 Node device to each EX Series switch (see [Figure 88 on page 295](#)).

NOTE: You may not mix EX4200 and EX4300 switches in the same QFabric system.

Figure 88: QFX3600 Node Device Fiber-Based Control Plane Connections for QFX3000-M QFabric System



CAUTION: For redundancy, each QFX3600 Node device must be connected to two EX Series switches.

Before you begin to connect a QFX3600 Node device to the fiber-based QFX3000-M QFabric system control plane network:

- Install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).
- Install your EX4200 or EX4300 switches. See *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.
- Interconnect the two EX Series switches. See [“Interconnecting Two EX Series Switches for QFX3000-M QFabric System Control Plane Redundancy” on page 239](#).

NOTE: Use either an EX4200-24F switch with an SFP+ uplink module installed or use an EX4300-48P switch with an SFP+ uplink module installed.

- Ensure that you have installed 1-Gigabit Ethernet SFP transceivers in management ports **C0S** and **C1S** on each QFX3600 Node device (see *Installing a Transceiver in a QFX Series Device*). For a list of supported transceivers and required cables, see [The Hardware Compatibility Tool](#).
- Follow the port assignments in [Table 67 on page 296](#) and ensure that you have installed 1-Gigabit Ethernet SFP transceivers in the ports (labeled **0** through **15**) you are using on each EX Series switch (see *Installing a Transceiver*). For a list of supported transceivers, see [The Hardware Compatibility Tool](#).
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings for the QFX Series*).
- Use [Table 67 on page 296](#) to determine the QFX3600 Node device-to-EX4200 switch port mappings. Specific ports have been reserved on the EX Series switch to connect to each of the QFX Series device types. Such design simplifies installation and facilitates timely deployment of a QFabric system.

NOTE: The numerical identifiers for each Node device below are not preassigned to the Node devices that are shipped to you. They represent the order in which you connect the Node devices. For example, the first Node device (Node0) is connected to port 0 (ge-0/0/0) on the first EX Series switch **EX0** and port 0 (ge-0/0/0) on the second EX Series switch **EX1**.

Table 67: QFX3600 Node Device-to-EX Series Switch Port Mappings for Fiber-Based Control Plane

Node Device	EX Series Switch EX0	EX Series Switch EX1
Node0	Node0, management port C0S to port 0 (ge-0/0/0)	Node0, management port C1S to port 0 (ge-0/0/0)
Node1	Node1, management port C0S to port 1 (ge-0/0/1)	Node1, management port C1S to port 1 (ge-0/0/1)
Node2	Node2, management port C0S to port 2 (ge-0/0/2)	Node2, management port C1S to port 2 (ge-0/0/2)
Node3	Node3, management port C0S to port 3 (ge-0/0/3)	Node3, management port C1S to port 3 (ge-0/0/3)
Node4	Node4, management port C0S to port 4 (ge-0/0/4)	Node4, management port C1S to port 4 (ge-0/0/4)

Table 67: QFX3600 Node Device-to-EX Series Switch Port Mappings for Fiber-Based Control Plane (continued)

Node Device	EX Series Switch EX0	EX Series Switch EX1
Node5	Node5, management port C0S to port 5 (ge-0/0/5)	Node5, management port C1S to port 5 (ge-0/0/5)
Node6	Node6, management port C0S to port 6 (ge-0/0/6)	Node6, management port C1S to port 6 (ge-0/0/6)
Node7	Node7, management port C0S to port 7 (ge-0/0/7)	Node7, management port C1S to port 7 (ge-0/0/7)
Node8	Node8, management port C0S to port 8 (ge-0/0/8)	Node8, management port C1S to port 8 (ge-0/0/8)
Node9	Node9, management port C0S to port 9 (ge-0/0/9)	Node9, management port C1S to port 9 (ge-0/0/9)
Node10	Node10, management port C0S to port 10 (ge-0/0/10)	Node10, management port C1S to port 10 (ge-0/0/10)
Node11	Node11, management port C0S to port 11 (ge-0/0/11)	Node11, management port C1S to port 11 (ge-0/0/11)
Node12	Node12, management port C0S to port 12 (ge-0/0/12)	Node12, management port C1S to port 12 (ge-0/0/12)
Node13	Node13, management port C0S to port 13 (ge-0/0/13)	Node13, management port C1S to port 13 (ge-0/0/13)
Node14	Node14, management port C0S to port 14 (ge-0/0/14)	Node14, management port C1S to port 14 (ge-0/0/14)
Node15	Node15, management port C0S to port 15 (ge-0/0/15)	Node15, management port C1S to port 15 (ge-0/0/15)

To connect a QFX3600 Node device to the fiber-based QFX3000-M QFabric system control plane network (see [Figure 88 on page 295](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
2. Remove the rubber safety cap from the SFP optical transceiver in the first SFP management port (labeled **C0S**) on the Node device front panel. Save the cap.
3. Insert the cable connector into the optical transceiver.
4. Connect the other end of that cable to the appropriate port on the EX Series switch. See [Table 67 on page 296](#).
5. Remove the rubber safety cap from the SFP optical transceiver in the second SFP management port (labeled **C1S**) on the Node device front panel. Save the cap.
6. Insert the cable connector into the optical transceiver.
7. Connect the other end of that cable to the appropriate port on the *second* EX Series switch. This should be the same port number that you connected to in [Step 4](#). For example, if you connected the first cable to ge-0/0/0 on the first EX Series switch, you connect the second cable to ge-0/0/0 on the second EX Series switch.
8. Repeat this procedure for each QFX3600 Node device.
9. Secure the cables so that they are not supporting their own weight. Place excess cable out of the way in a neatly coiled loop. Placing fasteners on a loop helps cables maintain their shape.



CAUTION: Do not bend fiber-optic cables beyond their minimum bend radius. An arc smaller than a few inches in diameter can damage the cables and cause problems that are difficult to diagnose.

Do not let fiber-optic cables hang free from the connector. Do not allow fastened loops of cables to dangle, which stresses the cables at the fastening point.

RELATED DOCUMENTATION

[Connecting a QFX3600-I Interconnect Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 281](#)

[Connecting a QFX3600 Node Device to a QFX3600-I Interconnect Device | 306](#)

[Connecting QFX3100 Director Devices to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 273](#)

[Connecting QFX3100 Director Devices in a Director Group | 242](#)

Cabling the Data Plane for the QFX3000-M QFabric System

IN THIS CHAPTER

- Connecting a QFX5100 Node Device to a QFX3600-I Interconnect Device | 301
- Connecting a QFX5100 Node Device to a QFX5100-24Q Interconnect Device | 304
- Connecting a QFX3600 Node Device to a QFX3600-I Interconnect Device | 306
- Connecting a QFX3600 Node Device to a QFX5100-24Q Interconnect Device | 309
- Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device | 311
- Connecting a QFX3500 Node Device to an QFX5100-24Q Interconnect Device | 313


Connecting a QFX5100 Node Device to a QFX3600-I Interconnect Device

To form the data plane in a QFX3000-M QFabric system using the QFX5100 line of switches, you connect the QSFP+ uplink ports on the QFX5100 Node device to the QSFP+ ports on the QFX3600-I Interconnect device. See [Table 68 on page 301](#) for the available ports for connecting to the QFX5100-24Q and [Figure 89 on page 302](#) for a diagram of the cabling. For a description of the port ranges and interface numbering for data plane and control plane, see *Understanding Port Ranges and System Modes*.

Table 68: Port Numbering for Data Plane Connections

QFX5100 Node Device Model	Port Numbers
QFX5100-24Q	0 through 23
QFX5100-48S	48 through 53
QFX5100-48T	48 through 53

Figure 89: QFX5100 Node Device Data Plane Connects for QFX3000-M QFabric System



CAUTION: For redundancy, each QFX5100 Node device must be connected to each QFX3600-I Interconnect device. For example, if you have two QFX3600-I Interconnect devices, then at least one uplink port on each QFX5100 Node device must be connected to each QFX3600-I Interconnect device. If you are connecting all four uplink ports to two QFX3600-I Interconnect devices, we recommend connecting two uplink ports to each Interconnect device.

Before you begin to cable the QFX3000-M QFabric data plane:

- Review [The Hardware Compatibility Tool](#) for information about the optical interface characteristics.
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings for the QFX Series*).
- Ensure you have installed QSFP+ transceivers in each port you are using. See *Installing a Transceiver in a QFX Series Device*.
- Ensure that you have appropriate fiber-optic cables (see *Cable Specifications for QSFP+ and QSFP28 Transceivers*).

To connect a QFX5100 Node device to a QFX3600-I Interconnect device (see [Figure 89 on page 302](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
2. Remove the rubber safety cap from the QSFP+ optical transceiver on the uplink port of the QFX5100 Node device. Save the cap.
3. Insert the cable connector into the optical transceiver.
4. If the connector at the other end of the fiber-optic cable is covered by a rubber safety cap, remove the cap. Save the cap.
5. Remove the rubber safety cap from the QSFP+ optical transceiver on the QFX3600-I Interconnect device. Save the cap.
6. Insert the cable connector into the optical transceiver.
7. Secure the cables so that they are not supporting their own weight. Place excess cable out of the way in a neatly coiled loop. Placing fasteners on a loop helps cables maintain their shape.



CAUTION: Do not bend fiber-optic cables beyond their minimum bend radius. An arc smaller than a few inches in diameter can damage the cables and cause problems that are difficult to diagnose.

Do not let fiber-optic cables hang free from the connector. Do not allow fastened loops of cables to dangle, which stresses the cables at the fastening point.

RELATED DOCUMENTATION

| [Determining Transceiver Support for QFabric Systems](#) | 143

Connecting a QFX5100 Node Device to a QFX5100-24Q Interconnect Device

To form the data plane in a QFX3000-M QFabric system using the QFX5100 line of switches, you connect the QSFP+ uplink ports on the QFX5100 Node device to the QSFP+ ports on the QFX5100-24Q Interconnect device. See [Table 68 on page 301](#) for the available ports for connecting to the QFX5100-24Q and [Figure 90 on page 304](#) for a diagram of the cabling. For a description of the port ranges and interface numbering for data plane and control plane, see *Understanding Port Ranges and System Modes*.

Table 69: Port Numbering for Data Plane Connections

QFX5100 Node Device Model	Port Numbers
QFX5100-24Q	0 through 23
QFX5100-48S	48 through 53
QFX5100-48T	48 through 53

Figure 90: QFX5100 Node Device Data Plane Connections to an QFX5100-24Q Interconnect Device for QFX3000-M QFabric System





CAUTION: For redundancy, each QFX5100 Node device must be connected to each QFX5100-24Q Interconnect device. For example, if you have two QFX5100-24Q Interconnect devices, then at least one uplink port on each QFX5100 Node device must be connected to each QFX5100-24Q Interconnect device. If you are connecting four uplink ports to two QFX5100-24Q Interconnect devices, we recommend connecting two uplink ports to each Interconnect device.

Before you begin to cable the QFX3000-M QFabric data plane:

- Review [The Hardware Compatibility Tool](#) for information about the optical interface characteristics of the QFX5100.
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings for the QFX Series*).
- Ensure you have installed QSFP+ transceivers in each port you are using. See *Installing a Transceiver in a QFX Series Device*.
- Ensure that you have appropriate fiber-optic cables (see *Cable Specifications for QSFP+ and QSFP28 Transceivers*).

To connect a QFX5100 Node device to a QFX5100-24Q Interconnect device (see [Figure 90 on page 304](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
2. Remove the rubber safety cap from the QSFP+ optical transceiver on the uplink port of the QFX5100 Node device. Save the cap.
3. Insert the cable connector into the optical transceiver.
4. If the connector at the other end of the fiber-optic cable is covered by a rubber safety cap, remove the cap. Save the cap.
5. Remove the rubber safety cap from the QSFP+ optical transceiver on the QFX5100-24Q Interconnect device. Save the cap.

6. Insert the cable connector into the optical transceiver.
7. Secure the cables so that they are not supporting their own weight. Place excess cable out of the way in a neatly coiled loop. Placing fasteners on a loop helps cables maintain their shape.



CAUTION: Do not bend fiber-optic cables beyond their minimum bend radius. An arc smaller than a few inches in diameter can damage the cables and cause problems that are difficult to diagnose.

Do not let fiber-optic cables hang free from the connector. Do not allow fastened loops of cables to dangle, which stresses the cables at the fastening point.

RELATED DOCUMENTATION

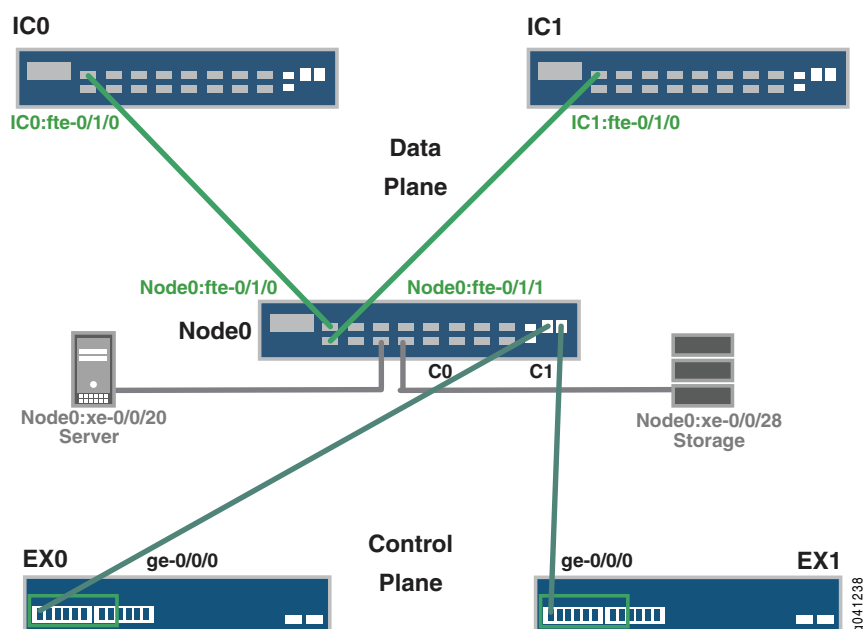
[Determining Transceiver Support for QFabric Systems](#) | 143

Connecting a QFX3600 Node Device to a QFX3600-I Interconnect Device

To form the data plane in a QFX3000-M QFabric system, you connect the QSFP+ uplink ports (labeled **Q0** through **Q7**) on the QFX3600 Node device to the QSFP+ ports (labeled **Q0** through **Q15**) on the QFX3600-I Interconnect device (see [Figure 91 on page 307](#)).

NOTE: By default, four ports (labeled **Q0** through **Q3**) are configured for 40-Gbps uplink connections between your QFX3600 Node device and your Interconnect device. Optionally, you can choose to configure the first eight ports (labeled **Q0** through **Q7**) for the uplink connections (see [“Configuring the Port Type on QFX3600 Node Devices” on page 408](#)).

Figure 91: QFX3600 Node Device Data Plane Connections for QFX3000-M QFabric System



CAUTION: For redundancy, each QFX3600 Node device must be connected to each QFX3600-I Interconnect device. For example, if you have two QFX3600-I Interconnect devices, then at least one uplink port on each QFX3600 Node device must be connected to each QFX3600-I Interconnect device. If you are connecting all four uplink ports to two QFX3600-I Interconnect devices, we recommend connecting two uplink ports to each Interconnect device.

Before you begin to cable the QFX3000-M QFabric system data plane:

- Review [The Hardware Compatibility Tool](#) for information about the optical interface characteristics.
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings for the QFX Series*).
- Ensure you have installed QSFP+ transceivers in each port you are using. See *Installing a Transceiver in a QFX Series Device*.
- Ensure that you have appropriate fiber-optic cables (see *Cable Specifications for QSFP+ and QSFP28 Transceivers*).

To connect a QFX3600 Node device to a QFX3600-I Interconnect device (see [Figure 91 on page 307](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
2. Remove the rubber safety cap from the QSFP+ optical transceiver on the uplink port of the QFX3600 Node device. Save the cap.
3. Insert the cable connector into the optical transceiver.
4. If the connector at the other end of the fiber-optic cable is covered by a rubber safety cap, remove the cap. Save the cap.
5. Remove the rubber safety cap from the QSFP+ optical transceiver on the QFX3600-I Interconnect device. Save the cap.
6. Insert the cable connector into the optical transceiver.
7. Secure the cables so that they are not supporting their own weight. Place excess cable out of the way in a neatly coiled loop. Placing fasteners on a loop helps cables maintain their shape.



CAUTION: Do not bend fiber-optic cables beyond their minimum bend radius. Bending the cables beyond their minimum bend radius can damage the cables and cause problems that are difficult to diagnose.

Do not let fiber-optic cables hang free from the connector. Do not allow fastened loops of cables to dangle, which stresses the cables at the fastening point.

RELATED DOCUMENTATION

[Determining Transceiver Support for QFabric Systems | 143](#)

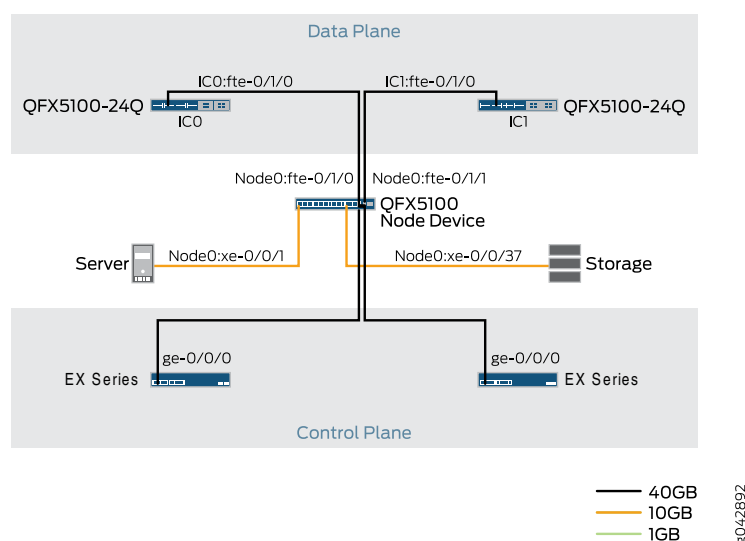
[Configuring the Port Type on QFX3600 Node Devices | 408](#)

Connecting a QFX3600 Node Device to a QFX5100-24Q Interconnect Device

To form the data plane in a QFX3000-M QFabric system, you connect the QSFP+ uplink ports (labeled **Q0** through **Q7**) on the QFX3600 Node device to the QSFP+ ports (labeled **Q** through **23**) on the QFX5100-24Q Interconnect device (see [Figure 92 on page 309](#)).

NOTE: By default, four ports (labeled **Q0** through **Q3**) are configured for 40-Gbps uplink connections between your QFX3600 Node device and your Interconnect device. Optionally, you can choose to configure the first eight ports (labeled **Q0** through **Q7**) for the uplink connections (see [“Configuring the Port Type on QFX3600 Node Devices” on page 408](#)).

Figure 92: QFX3600 Node Device Data Plane Connections to QFX5100-24Q for QFX3000-M QFabric System



CAUTION: For redundancy, each QFX3600 Node device must be connected to each QFX5100-24Q Interconnect device. For example, if you have two QFX5100-24Q Interconnect devices, then at least one uplink port on each QFX3600 Node device must be connected to each QFX5100-24Q Interconnect device. If you are connecting all four uplink ports to two QFX5100-24Q Interconnect devices, we recommend connecting two uplink ports to each Interconnect device.

Before you begin to cable the QFX3000-M QFabric system data plane:

- Review [The Hardware Compatibility Tool](#) for information about the optical interface characteristics.
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings for the QFX Series*).
- Ensure you have installed QSFP+ transceivers in each port you are using. See *Installing a Transceiver in a QFX Series Device*.
- Ensure that you have appropriate fiber-optic cables (see *Cable Specifications for QSFP+ and QSFP28 Transceivers*).

To connect a QFX3600 Node device to a QFX5100-24Q Interconnect device :



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
2. Remove the rubber safety cap from the QSFP+ optical transceiver on the uplink port of the QFX3600 Node device. Save the cap.
3. Insert the cable connector into the optical transceiver.
4. If the connector at the other end of the fiber-optic cable is covered by a rubber safety cap, remove the cap. Save the cap.
5. Remove the rubber safety cap from the QSFP+ optical transceiver on the QFX5100-24Q Interconnect device. Save the cap.
6. Insert the cable connector into the optical transceiver.
7. Secure the cables so that they are not supporting their own weight. Place excess cable out of the way in a neatly coiled loop. Placing fasteners on a loop helps cables maintain their shape.



CAUTION: Do not bend fiber-optic cables beyond their minimum bend radius. Bending the cables beyond their minimum bend radius can damage the cables and cause problems that are difficult to diagnose.

Do not let fiber-optic cables hang free from the connector. Do not allow fastened loops of cables to dangle, which stresses the cables at the fastening point.

RELATED DOCUMENTATION

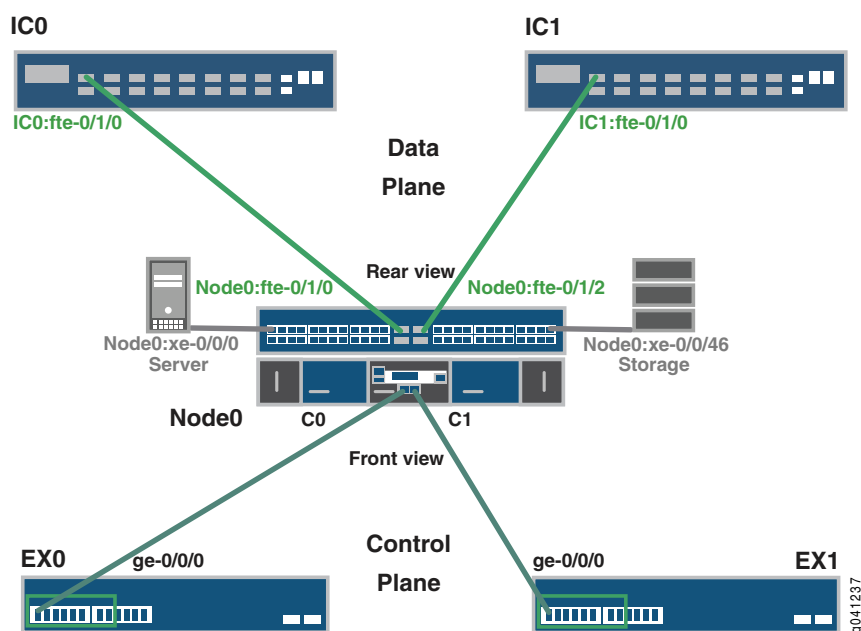
[Determining Transceiver Support for QFabric Systems | 143](#)

[Configuring the Port Type on QFX3600 Node Devices | 408](#)

Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device

To form the data plane in a QFX3000-M QFabric system, you connect the QSFP+ uplink ports (labeled **Q0** through **Q3**) on the QFX3500 Node device to the QSFP+ ports (labeled **Q0** through **Q15**) on the QFX3600-I Interconnect device (see [Figure 93 on page 311](#)).

Figure 93: QFX3500 Node Device Data Plane Connections for QFX3000-M QFabric System





CAUTION: For redundancy, each QFX3500 Node device must be connected to each QFX3600-I Interconnect device. For example, if you have two QFX3600-I Interconnect devices, then at least one uplink port on each QFX3500 Node device must be connected to each QFX3600-I Interconnect device. If you are connecting all four uplink ports to two QFX3600-I Interconnect devices, we recommend connecting two uplink ports to each Interconnect device.

Before you begin to cable the QFX3000-M QFabric data plane:

- Review [The Hardware Compatibility Tool](#) for information about the optical interface characteristics.
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings for the QFX Series*).
- Ensure you have installed QSFP+ transceivers in each port you are using. See *Installing a Transceiver in a QFX Series Device*.
- Ensure that you have appropriate fiber-optic cables (see *Cable Specifications for QSFP+ and QSFP28 Transceivers*).

To connect a QFX3500 Node device to a QFX3600 Interconnect device (see [Figure 93 on page 311](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
2. Remove the rubber safety cap from the QSFP+ optical transceiver on the uplink port of the QFX3500 Node device. Save the cap.
3. Insert the cable connector into the optical transceiver.
4. If the connector at the other end of the fiber-optic cable is covered by a rubber safety cap, remove the cap. Save the cap.
5. Remove the rubber safety cap from the QSFP+ optical transceiver on the QFX3600-I Interconnect device. Save the cap.

6. Insert the cable connector into the optical transceiver.
7. Secure the cables so that they are not supporting their own weight. Place excess cable out of the way in a neatly coiled loop. Placing fasteners on a loop helps cables maintain their shape.



CAUTION: Do not bend fiber-optic cables beyond their minimum bend radius. An arc smaller than a few inches in diameter can damage the cables and cause problems that are difficult to diagnose.

Do not let fiber-optic cables hang free from the connector. Do not allow fastened loops of cables to dangle, which stresses the cables at the fastening point.


RELATED DOCUMENTATION

[Determining Transceiver Support for QFabric Systems](#) | 143

Connecting a QFX3500 Node Device to an QFX5100-24Q Interconnect Device

To form the data plane in a QFX3000-M QFabric system, you connect the QSFP+ uplink ports (labeled **Q0** through **Q3**) on the QFX3500 Node device to the QSFP+ ports (labeled **Q** through **23**) on the QFX5100-24Q Interconnect device (see [Figure 93 on page 311](#)).

Figure 94: QFX3500 Node Device Data Plane Connections for QFX3000-M QFabric System



CAUTION: For redundancy, each QFX3500 Node device must be connected to each QFX5100-24Q Interconnect device. For example, if you have two QFX5100-24Q Interconnect devices, then at least one uplink port on each QFX3500 Node device must be connected to each QFX5100-24Q Interconnect device. If you are connecting all four uplink ports to two QFX5100-24Q Interconnect devices, we recommend connecting two uplink ports to each Interconnect device.

Before you begin to cable the QFX3000-M QFabric data plane:

- Review [The Hardware Compatibility Tool](#) for information about the optical interface characteristics.
- Ensure that you have taken the necessary precautions for safe handling of lasers (see *Laser and LED Safety Guidelines and Warnings for the QFX Series*).
- Ensure you have installed QSFP+ transceivers in each port you are using. See *Installing a Transceiver in a QFX Series Device*.
- Ensure that you have appropriate fiber-optic cables (see *Cable Specifications for QSFP+ and QSFP28 Transceivers*).

To connect a QFX3500 Node device to a QFX5100-24Q Interconnect device (see [Figure 93 on page 311](#)):



WARNING: Do not look directly into a fiber-optic transceiver or into the ends of fiber-optic cables. Fiber-optic transceivers and fiber-optic cables connected to transceivers emit laser light that can damage your eyes.

1. If the fiber-optic cable connector is covered by a rubber safety cap, remove the cap. Save the cap.
2. Remove the rubber safety cap from the QSFP+ optical transceiver on the uplink port of the QFX3500 Node device. Save the cap.
3. Insert the cable connector into the optical transceiver.
4. If the connector at the other end of the fiber-optic cable is covered by a rubber safety cap, remove the cap. Save the cap.
5. Remove the rubber safety cap from the QSFP+ optical transceiver on the QFX5100-24Q Interconnect device. Save the cap.
6. Insert the cable connector into the optical transceiver.
7. Secure the cables so that they are not supporting their own weight. Place excess cable out of the way in a neatly coiled loop. Placing fasteners on a loop helps cables maintain their shape.



CAUTION: Do not bend fiber-optic cables beyond their minimum bend radius. An arc smaller than a few inches in diameter can damage the cables and cause problems that are difficult to diagnose.

Do not let fiber-optic cables hang free from the connector. Do not allow fastened loops of cables to dangle, which stresses the cables at the fastening point.

RELATED DOCUMENTATION

| [Determining Transceiver Support for QFabric Systems](#) | 143

3

PART

Configuration

Initial Setup | **319**

QFabric System Configuration | **371**

QFabric System Fabric OAM Configuration | **439**

Configuration Statements | **457**

Initial Setup

IN THIS CHAPTER

- QFabric System Initial and Default Configuration Information | 319
- Converting the Device Mode for a QFabric System Component | 322
- Example: Configuring EX Series Switches for the QFX3000-M QFabric System Control Plane | 329
- Importing a QFX3000-M QFabric System Control Plane EX Series Switch Configuration with a USB Flash Drive | 360
- Generating the MAC Address Range for a QFabric System | 362
- Performing the QFabric System Initial Setup on a QFX3100 Director Group | 363

QFabric System Initial and Default Configuration Information

Once you install the hardware for the QFabric system, you can configure the Junos operating system (Junos OS) to begin using the system. This topic discusses which setup activities you need to perform and which activities are handled automatically by the QFabric system.

The fabric manager Routing Engine in the Director group automatically handles some of the initial setup activities, including:

- Assignment of IP addresses and unique identifiers to each QFabric system component by way of the management control plane
- Inclusion of all QFabric system devices within the default partition
- Establishment of interdevice communication and connectivity through the use of a fabric provisioning protocol and a fabric management protocol

The initial configuration tasks you need to perform to bring up the QFabric system and make it operational include:

- Converting any standalone devices, such as QFX3500 and QFX3600 devices, to Node device mode (see [“Converting the Device Mode for a QFabric System Component” on page 322](#))
- Setting up the QFabric system control plane cabling, topology, and configuration

- To set up the control plane cabling, topology, and configuration for the QFX3000-G QFabric system, see *Example: Configuring the Virtual Chassis for a Copper-Based QFX3000-G QFabric System Control Plane*.
- To set up a copper or fiber-based control plane cabling, topology, and configuration for the QFX3000-M QFabric system, see [“Example: Configuring EX Series Switches for the QFX3000-M QFabric System Control Plane” on page 329](#).
- Accessing the Director group through a console connection, turning on the devices, and running through the initial setup script (see [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 363](#)), which prompts you to:
 - Set IP addresses for the Director devices in the Director group.
 - Set an IP address for the default partition.
 - Add the software serial number for your QFabric system. (Review the e-mail containing the software serial number that you received from Juniper Networks when you purchased your QFabric system.)
 - Set the starting MAC address and the range of MAC addresses for the QFabric system. (See [“Generating the MAC Address Range for a QFabric System” on page 362](#) for this information.)
 - Set a root password for the Director devices.
 - Set a root password for the QFabric system components, such as Node devices, Interconnect devices, and infrastructure.
- Logging into the default partition by using the IP address you configured when you ran the Director group initial setup script (See [“Gaining Access to the QFabric System Through the Default Partition” on page 375](#))
- Configuring basic system settings for the default partition, such as time, location, and default gateways

NOTE: Unlike other Juniper Networks devices that run Junos OS, a QFabric system does not have a default factory configuration (containing the basic configuration settings for system logging, interfaces, protocols, and so on) that is loaded when you first install and power on the Director devices. Therefore, you must configure all the settings required for your QFabric system through the default partition CLI.

- Configuring aliases for Node devices (see [“Configuring Aliases for the QFabric System” on page 388](#))
- Configuring VLANs and interfaces for the QFabric system devices
- Configuring redundant server Node groups to provide resiliency for server and storage connections (see [“Configuring Node Groups for the QFabric System” on page 402](#))
- Configuring a network Node group to connect the QFabric system to external networks (see [“Configuring Node Groups for the QFabric System” on page 402](#))

- Configuring the port type on QFX3600 Node devices (see [“Configuring the Port Type on QFX3600 Node Devices”](#) on page 408)
- Configuring routing protocols to run on the network Node group interfaces and reach external networks

NOTE: When you configure routing protocols on the QFabric system, you must use interfaces from the Node devices assigned to the network Node group. If you try to configure routing protocols on interfaces from the Node devices assigned to server Node groups, the configuration commit operation fails.

- Generating and adding the license keys for the QFabric system (see *Generating the License Keys for a QFabric System* and *Adding New Licenses (CLI Procedure)*)

RELATED DOCUMENTATION

QFX3000-G QFabric System Installation Overview

[QFX3000-M QFabric System Installation Overview | 109](#)

[Converting the Device Mode for a QFabric System Component | 322](#)

Example: Configuring the Virtual Chassis for a Copper-Based QFX3000-G QFabric System Control Plane

[Example: Configuring EX Series Switches for the QFX3000-M QFabric System Control Plane | 329](#)

[Generating the MAC Address Range for a QFabric System | 362](#)

[Performing the QFabric System Initial Setup on a QFX3100 Director Group | 363](#)

[Understanding QFabric System Administration Tasks and Utilities | 371](#)

[Gaining Access to the QFabric System Through the Default Partition | 375](#)

[Configuring Aliases for the QFabric System | 388](#)

[Configuring Node Groups for the QFabric System | 402](#)

[Configuring the Port Type on QFX3600 Node Devices | 408](#)

Generating the License Keys for a QFabric System

Adding New Licenses (CLI Procedure)

Converting the Device Mode for a QFabric System Component

You can configure some devices to act as a standalone switch or participate in a QFabric system in a particular role. To change the role of your device, you must set the device mode. [Table 70 on page 322](#) shows the device modes available for various devices.

Table 70: Support for device mode options

Device mode	QFX3500	QFX3600	QFX5100
Interconnect device	N/A	Supported	Supported for QFX3000-M
Node device	Supported	Supported	Supported
Standalone	Supported	Supported	N/A

To convert a device to a different mode, issue the **request chassis device-mode** command and specify the desired device mode. You verify the current and future device mode with the **show chassis device-mode** command.

When you convert a device from standalone mode to either Node device or Interconnect device mode, the software prepares the device to be configured automatically by the QFabric system. However, changing the device mode erases all configuration data on the device.

NOTE: The QFX3600 switch requires Jloader Release 1.1.8 before you can convert the switch to Interconnect device mode. For more information, see: [Jloader 1.1.8 Release for QFX-Series Platforms](#).



CAUTION: We recommend that you back up your device configuration to an external location before converting a device to a different device mode.

The following procedures illustrate the conversion options available when you modify a device mode:

- Convert from standalone switch mode to Node device mode
- Convert from Node device mode to Interconnect device mode
- Convert from Interconnect device mode to Node device mode
- Convert from Node device mode or Interconnect device mode to standalone switch mode

Standalone Switch to Node Device

To convert your device from standalone mode to Node device mode, follow these steps:

1. Connect to your standalone device through the console port and log in as the root user.
2. Back up your device configuration to an external location.

```
root@switch# save configuration-name external-path
```

3. Upgrade the software on your device to a QFabric system **Node and Interconnect device** software package that matches the QFabric system complete software package used by your QFabric system. If the complete software package for your QFabric system is named **jinstall-qfabric-13.2X52-D10.2.rpm**, you need to install the **jinstall-qfabric-5-13.2X52-D10.2-domestic-signed.tgz** package on your QFX5100 device and the **jinstall-qfx-13.2X52-D10.2-domestic-signed.tgz** package on your QFX3500 or QFX3600 device. Matching the two software packages ensures a smooth and successful addition of the device to the QFabric system inventory.

```
root@switch# request system software add software-package-name reboot
```

NOTE: After you install the correct software, the QFX5100 device is placed into Node device mode by default and cannot be converted to any other mode in Junos OS Release 13.2X52-D10.

4. Check the current device mode by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
```

```
Current device-mode : Standalone
Future device-mode after reboot : Standalone
```

5. Issue the **request chassis device-mode** command and select the desired device mode.

```
root@switch> request chassis device-mode node-device
```

```
Device mode set to 'node-device' mode.
Please reboot the system to complete the process.
```

6. Verify the future device mode by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
```



```
Current device-mode : Standalone
Future device-mode after reboot : Node-device
```

7. Reboot the device.

```
root@switch> request system reboot
```

```
Reboot the system ? [yes,no] (no) yes

Shutdown NOW!
[pid 34992]

root@switch>

*** FINAL System shutdown message from root@switch ***
System going down IMMEDIATELY
```

8. Verify that the new device mode has been enabled by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
```

```
Current device-mode : Node-device
Future device-mode after reboot : Node-device
```

9. To enable a converted device to participate in the QFabric system, locate the applicable network cables for your device and connect the device ports to the control plane and data plane.

10. (Optional) If you change the device back from Node device mode to standalone mode, restore the saved backup configuration from your external location.

```
root@switch# load configuration-name external-path
```

Node Device to Interconnect Device

To convert your device from Node device mode to Interconnect device mode, follow these steps:

1. From the default partition CLI prompt, back up your QFabric system configuration to an external location.

```
user@qfabric# save configuration-name external-path
```

2. Connect to your device through the console port and log in as the root user.

3. Check the current device mode by issuing the **show chassis device-mode** command.


```
root@switch> show chassis device-mode
```

```
Current device-mode : Node-device
Future device-mode after reboot : Node-device
```

4. Issue the **request chassis device-mode** command and select the desired device mode.

```
root@switch> request chassis device-mode interconnect-device
```

```
Device mode set to 'interconnect-device' mode.
Please reboot the system to complete the process.
```

5. Verify the future device mode by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
```

```
Current device-mode : Node-device
Future device-mode after reboot : Interconnect-device
```

6. Reboot the device.

```
root@switch> request system reboot
```

```
Reboot the system ? [yes,no] (no) yes

Shutdown NOW!
[pid 34992]

root@switch>

*** FINAL System shutdown message from root@switch ***
System going down IMMEDIATELY
```

7. Verify that the new device mode has been enabled by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
```

```
Current device-mode : Interconnect-device
Future device-mode after reboot : Interconnect-device
```

8. To enable a converted device to participate in the QFabric system in its new role, move the device to a different rack (as needed), locate the applicable network cables for your device, connect the device ports to the control plane and data plane per the design for your specific QFabric system, and reconfigure any aliases for the device at the QFabric default partition CLI prompt.

Interconnect Device to Node Device

To convert your device from Interconnect device mode to Node device mode, follow these steps:

1. From the default partition CLI prompt, back up your QFabric system configuration to an external location.

```
user@qfabric# save configuration-name external-path
```

2. Connect to your device through the console port and log in as the root user.
3. Check the current device mode by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
```

```
Current device-mode : Interconnect-device
Future device-mode after reboot : Interconnect-device
```

4. Issue the **request chassis device-mode** command and select the desired device mode.

```
root@switch> request chassis device-mode node-device
```

```
Device mode set to 'node-device' mode.
Please reboot the system to complete the process.
```

5. Verify the future device mode by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
```

```
Current device-mode : Interconnect-device
Future device-mode after reboot : Node-device
```

6. Reboot the device.

```
root@switch> request system reboot
```

```
Reboot the system ? [yes,no] (no) yes

Shutdown NOW!
[pid 34992]

root@switch>

*** FINAL System shutdown message from root@switch ***
System going down IMMEDIATELY
```


7. Verify that the new device mode has been enabled by issuing the **show chassis device-mode** command.

```
root@switch> show chassis device-mode
```

```
Current device-mode : Node-device
Future device-mode after reboot : Node-device
```

8. To enable a converted device to participate in the QFabric system in its new role, move the device to a different rack (as needed), locate the applicable network cables for your device, connect the device ports to the control plane and data plane per the design for your specific QFabric system, and reconfigure any aliases for the device at the QFabric default partition CLI prompt.

QFabric Component (Interconnect or Node Device) to Standalone Switch

To convert your QFabric component from either Interconnect device mode or Node device mode to standalone switch mode, follow these steps:

1. From the default partition CLI prompt, back up your QFabric system configuration to an external location.

```
user@qfabric# save configuration-name external-path
```

2. Connect to the desired QFabric component through the console port of the device and log in as the root user.

3. Check the current device mode by issuing the **show chassis device-mode** command.

```
root@node1> show chassis device-mode
```

```
Current device-mode : Node-device
Future device-mode after reboot : Node-device
```

4. Issue the **request chassis device-mode standalone** command to convert the component to standalone switch mode, while the component is still connected to the QFabric system.

```
root@node1> request chassis device-mode standalone
```

```
Device mode set to 'standalone' mode.
Please reboot the system to complete the process.
```


NOTE: Always convert the device mode to **standalone** before you remove the component from the QFabric system. If you remove the component from the QFabric system before converting the device mode to **standalone**, the switch might not operate properly. For example, the output of the **show chassis hardware** command might display no FPCs or interfaces for the switch.

5. Verify the future device mode by issuing the **show chassis device-mode** command.

```
root@node1> show chassis device-mode
```

```
Current device-mode : Node-device
Future device-mode after reboot : Standalone
```

6. Reboot the component to complete the conversion process.

```
root@node1> request system reboot
```

```
Reboot the system ? [yes,no] (no) yes

Shutdown NOW!
[pid 34992]

root@node1>

*** FINAL System shutdown message from root@node1 ***
System going down IMMEDIATELY
```

7. Disconnect and remove the component from the QFabric system. You may now operate the device as a standalone switch.

RELATED DOCUMENTATION

[request chassis device-mode | 600](#)

[show chassis device-mode | 653](#)

Software Installation and Upgrade Overview

Connecting a QFX3500 Node Device to a Copper-Based QFX3000-G QFabric System Control Plane Network

Connecting a QFX3600 Node Device to a Copper-Based QFX3000-G QFabric System Control Plane Network

Connecting a QFX5100 Node Device to a Copper-Based QFX3000-G QFabric System Control Plane Network

Connecting a QFX3500 Node Device to a Fiber-Based QFX3000-G QFabric System Control Plane Network

Connecting a QFX3600 Node Device to a Fiber-Based QFX3000-G QFabric System Control Plane Network

Connecting a QFX5100 Node Device to a Fiber-Based QFX3000-G QFabric System Control Plane Network

Connecting a QFX3500 Node Device to a QFX3008-I Interconnect Device

Connecting a QFX3600 Node Device to a QFX3008-I Interconnect Device

Connecting a QFX5100 Node Device to a QFX3008-I Interconnect Device

[Connecting a QFX3500 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 263](#)

[Connecting a QFX3600 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 259](#)

[Connecting a QFX5100 Node Device to a Copper-Based QFX3000-M QFabric System Control Plane Network | 256](#)

[Connecting a QFX3500 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 290](#)

[Connecting a QFX3600 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 295](#)

[Connecting a QFX5100 Node Device to a Fiber-Based QFX3000-M QFabric System Control Plane Network | 286](#)

[Connecting a QFX3500 Node Device to a QFX3600-I Interconnect Device | 311](#)

[Connecting a QFX3600 Node Device to a QFX3600-I Interconnect Device | 306](#)

[Connecting a QFX5100 Node Device to a QFX3600-I Interconnect Device | 301](#)

Example: Configuring EX Series Switches for the QFX3000-M QFabric System Control Plane

IN THIS SECTION

- [Requirements | 330](#)
- [Overview | 331](#)
- [Configuration | 340](#)
- [Verification | 354](#)

This example shows you how to connect QFabric system components and configure EX4200 switches used for the QFX3000-M QFabric system control plane network. Proper wiring of Director devices, Interconnect devices, and Node devices to the EX Series switches, combined with a standard configuration, enables you to bring up the internal QFabric system management network and prepare your QFabric system for full operation.

NOTE: EX4200 or EX4300 switch configuration is the same for both the copper-based and fiber-based QFX3000-M QFabric system control plane networks. Hence, a separate example for configuring EX4300 switches or the fiber-based control plane network is not provided.

However, because you cannot mix and match fiber and copper in the same control plane network, you must select only one type of control plane for each QFX3000-M QFabric system you install. The primary focus of this example is a copper-based control plane network with EX4200-24T switches. Before you use this example to configure a fiber-based control plane network with EX4200 switches, ensure that you have installed and wired the QFabric system hardware and EX4200 switches as required for a fiber-based control plane network (see [“QFX3000-M QFabric System Installation Overview” on page 109](#)). Before you use this example to configure a copper or fiber-based control plane network with EX4300 switches, ensure that you have installed and wired the QFabric system hardware and EX4300 switches as required in *Installing and Connecting an EX4300 Switch*.

Requirements

This example uses the following hardware and software components:

- One QFX3000-M QFabric system containing:
 - Two QFX3100 Director devices with 1000BASE-T network modules installed
 - Two QFX3600-I Interconnect devices
 - Eight QFX3500 Node devices with a 1000BASE-T management board installed
- Two EX4200-24T switches with SFP+ uplink module installed

NOTE: Instead of two EX4200-24T switches, you can use two EX4300-48T switches for a copper-based control plane network. For a fiber-based control plane network, you can use two EX4200-24F switches or two EX4300-48P switches.

- Junos OS Release 13.2X52-D10 for the QFabric system components
- Junos OS Release 12.3R6.6 for the EX Series switches

Before you begin:

- Rack, mount, and install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see [“Installing and Connecting a QFX3100 Director Device” on page 165](#), [“Installing and Connecting a QFX3600 or QFX3600-I Device” on page 177](#), and [“Installing and Connecting a QFX3500 Device” on page 199](#).
- Rack, mount, and install your control plane EX Series switches. For more information, see *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.

Overview


The QFX3000-M QFabric system control plane network connects the Director group, Interconnect devices, and Node devices in a QFabric system across a pair of redundant EX Series switches. By separating the management control plane from the data plane, the QFabric system can scale efficiently. The copper-based control plane network uses Gigabit Ethernet cabling and connections between components, and two 1-Gigabit Ethernet connections configured in a link aggregation group (LAG) between the redundant EX Series switches.

Specific ports have been reserved on the EX Series switches to connect to each of the QFabric system device types. Such design simplifies installation and facilitates timely deployment of a QFabric system. It also permits the use of a standard EX Series switch configuration included as part of this example. The standard configuration can scale from the 8 Node devices shown in this example to a maximum of 16 Node devices.

Topology

[Figure 95 on page 331](#) shows the general port ranges where QFabric system devices must be connected to the EX Series switches. For each EX Series switch, connect ports **0** through **15** to Node devices, ports **16** through **19** to Interconnect devices, ports **20** through **23** to Director devices, and uplink ports **0** and **1** to the other control plane switch as an inter-switch LAG. [Table 71 on page 332](#) shows the details of the QFabric system component-to-EX Series switch port mappings.

Figure 95: QFX3000-M QFabric System Control Plane—EX Series Switch Port Ranges



**CAUTION:**

- The control plane network within a QFabric system is a critical component of the system that should not be shared with other network traffic. In order to scale efficiently, the control plane network must be reserved for the QFabric system and its components. As a result, the ports of the QFabric system control plane must never be used for any purpose other than to transport QFabric system control plane traffic, and we neither recommend nor support the connection of other devices to the QFabric system control plane network.
- Do not install Junos Space and AI-Scripts (AIS) on the control plane network EX Series switches in a QFX3000-M QFabric system.

Table 71 on page 332 shows the specific mappings of QFabric system control plane network ports from the QFabric system components to the EX Series switches.

NOTE: The uplink ports 2 and 3 on the EX Series switches are reserved for future use.

Table 71: QFX3000-M QFabric System Copper-Based Control Plane—QFabric Component-to-EX Series Switch Port Mappings

EX Series Switch 1 (EX0)	EX Series Switch 2 (EX1)	QFabric System Component
Node0, management port C0 to port 0 (ge-0/0/0)	Node0, management port C1 to port 0 (ge-0/0/0)	Node device 0
Node1, management port C0 to port 1 (ge-0/0/1)	Node1, management port C1 to port 1 (ge-0/0/1)	Node device 1
Node2, management port C0 to port 2 (ge-0/0/2)	Node2, management port C1 to port 2 (ge-0/0/2)	Node device 2
Node3, management port C0 to port 3 (ge-0/0/3)	Node3, management port C1 to port 3 (ge-0/0/3)	Node device 3
Node4, management port C0 to port 4 (ge-0/0/4)	Node4, management port C1 to port 4 (ge-0/0/4)	Node device 4
Node5, management port C0 to port 5 (ge-0/0/5)	Node5, management port C1 to port 5 (ge-0/0/5)	Node device 5

Table 71: QFX3000-M QFabric System Copper-Based Control Plane—QFabric Component-to-EX Series Switch Port Mappings (continued)

EX Series Switch 1 (EX0)	EX Series Switch 2 (EX1)	QFabric System Component
Node6, management port C0 to port 6 (ge-0/0/6)	Node6, management port C1 to port 6 (ge-0/0/6)	Node device 6
Node7, management port C0 to port 7 (ge-0/0/7)	Node7, management port C1 to port 7 (ge-0/0/7)	Node device 7
Node8, management port C0 to port 8 (ge-0/0/8)	Node8, management port C1 to port 8 (ge-0/0/8)	Node device 8
Node9, management port C0 to port 9 (ge-0/0/9)	Node9, management port C1 to port 9 (ge-0/0/9)	Node device 9
Node10, management port C0 to port 10 (ge-0/0/10)	Node10, management port C1 to port 10 (ge-0/0/10)	Node device 10
Node11, management port C0 to port 11 (ge-0/0/11)	Node11, management port C1 to port 11 (ge-0/0/11)	Node device 11
Node12, management port C0 to port 12 (ge-0/0/12)	Node12, management port C1 to port 12 (ge-0/0/12)	Node device 12
Node13, management port C0 to port 13 (ge-0/0/13)	Node13, management port C1 to port 13 (ge-0/0/13)	Node device 13
Node14, management port C0 to port 14 (ge-0/0/14)	Node14, management port C1 to port 14 (ge-0/0/14)	Node device 14
Node15, management port C0 to port 15 (ge-0/0/15)	Node15, management port C1 to port 15 (ge-0/0/15)	Node device 15
IC0, management port C0 to port 16 (ge-0/0/16)	IC0, management port C1 to port 16 (ge-0/0/16)	Interconnect device 0
IC1, management port C0 to port 17 (ge-0/0/17)	IC1, management port C1 to port 17 (ge-0/0/17)	Interconnect device 1
IC2, management port C0 to port 18 (ge-0/0/18)	IC2, management port C1 to port 18 (ge-0/0/18)	Interconnect device 2

Table 71: QFX3000-M QFabric System Copper-Based Control Plane—QFabric Component-to-EX Series Switch Port Mappings (*continued*)


EX Series Switch 1 (EX0)	EX Series Switch 2 (EX1)	QFabric System Component
IC3, management port C0 to port 19 (ge-0/0/19)	IC3, management port C1 to port 19 (ge-0/0/19)	Interconnect device 3
DG0 module 0, port 0 to port 20 (ge-0/0/20)	DG0 module 1, port 0 to port 20 (ge-0/0/20)	Director device 0
DG0 module 0, port 1 to port 21 (ge-0/0/21)	DG0 module 1, port 1 to port 21 (ge-0/0/21)	Director device 0
DG1 module 0, port 0 to port 22 (ge-0/0/22)	DG1 module 1, port 0 to port 22 (ge-0/0/22)	Director device 1
DG1 module 0, port 1 to port 23 (ge-0/0/23)	DG1 module 1, port 1 to port 23 (ge-0/0/23)	Director device 1
EX0, uplink port 0 to EX1, uplink port 0 (ge-0/1/0)	EX1, uplink port 0 to EX0, uplink port 0 (ge-0/1/0)	Inter-EX Series switch LAG
EX0, uplink port 1 to EX1, uplink port 1 (ge-0/1/1)	EX1, uplink port 1 to EX0, uplink port 1 (ge-0/1/1)	Inter-EX Series switch LAG
Reserved Uplink port 2 (ge-0/1/2)	Reserved Uplink port 2 (ge-0/1/2)	Future use
Reserved Uplink port 3 (ge-0/1/3)	Reserved Uplink port 3 (ge-0/1/3)	Future use

Next, connect the Director devices to the EX Series switches. In general, you want to accomplish the following:

- Connect two ports from one network module in a Director device to the first EX Series switch, and two ports from the second network module to the second EX Series switch.
- Connect the Director devices to each other and create a Director group. You can use either straight-through RJ-45 patch cables or crossover cables, because the Director devices contain autosensing modules. Connect one port from each network module on the first Director device to one port in each network module on the second Director device.

Figure 96 on page 335 shows the specific ports on the Director group that you must connect to the EX Series switches and interconnect between the Director devices.

Figure 96: QFX3000-M QFabric System Control Plane—Director Group to EX Series Switch Connections



In this specific example, connect ports **0** and **1** from module **0** on Director device DG0 to ports **20** and **21** on EX Series switch EX0 (ge-0/0/20 and ge-0/0/21), and connect ports **0** and **1** from module **1** to ports **20** and **21** on the *second* EX Series switch EX1 (ge-0/0/20 and ge-0/0/21).

For Director device DG1, connect ports **0** and **1** from module **0** to ports **22** and **23** on EX Series switch EX0 (ge-0/0/22 and ge-0/0/23), and connect ports **0** and **1** from module **1** to ports **22** and **23** on the *second* EX Series switch EX1 (ge-0/0/22 and ge-0/0/23).

To form the Director group, connect port **3** on module **0** on Director device DG0 to port **3** on module **0** on Director device DG1. Similarly, connect port **3** on module **1** on Director device DG0 to port **3** on module **1** on Director device DG1. [Table 72 on page 335](#) shows the port mappings for the Director group in this example.

Table 72: Director Group Port Mappings

Director Device	EX Series Switch EX0	EX Series Switch EX1
DG0	<ul style="list-style-type: none"> DG0 module 0, port 0 to port 20 on EX0 (ge-0/0/20) DG0 module 0, port 1 to port 21 on EX0 (ge-0/0/21) DG0 module 0, port 3 to module 0, port 3 on DG1 	<ul style="list-style-type: none"> DG0 module 1, port 0 to port 20 on EX1 (ge-0/0/20) DG0 module 1, port 1 to port 21 on EX1 (ge-0/0/21) DG0 module 1, port 3 to module 1, port 3 on DG1

Table 72: Director Group Port Mappings (continued)

Director Device	EX Series Switch EX0	EX Series Switch EX1
DG1	<ul style="list-style-type: none"> DG1 module 0, port 0 to port 22 on EX0 (ge-0/0/22) DG1 module 0, port 1 to port 23 on EX0 (ge-0/0/23) DG1 module 0, port 3 to module 0, port 3 on DG0 	<ul style="list-style-type: none"> DG1 module 1, port 0 to port 22 on EX1 (ge-0/0/22) DG1 module 1, port 1 to port 23 on EX1 (ge-0/0/23) DG1 module 1, port 3 to module 1, port 3 on DG0


In the software, the ports of each network module on a Director device are reversed, numbered from right to left, and incremented sequentially across modules. If you issue interface operational commands directly on the Director device, note the following port mappings as shown in [Table 73 on page 336](#):

Table 73: Hardware to Software Port Mappings for Director Device Network Modules

Network Module	Port 0	Port 1	Port 2	Port 3
Module 0	eth5	eth4	eth3	eth2
Module 1	eth9	eth8	eth7	eth6

[Figure 97 on page 336](#) shows the specific ports on the QFX3600-I Interconnect devices that you must connect to the EX Series switches. In general, connect the first management port in an Interconnect device to the first EX Series switch, and the second management port to the second EX Series switch.

Figure 97: QFX3000-M QFabric System Control Plane—Interconnect Device to EX Series Switch Connections



In this specific example, for both Interconnect devices IC0 and IC1, connect management port **C0** to EX Series switches EX0 and EX1 and management port **C1** to EX Series switches EX0 and EX1. Connect the management port **C0** cables to port 16 on EX Series switches EX0 and EX1 (ge-0/0/16), and connect the management port **C1** cables to port 17 on EX Series switches EX0 and EX1 (ge-0/0/17). [Table 74 on page 337](#) shows the port mappings for the Node devices in this example.

Table 74: Interconnect Device Port Mappings

Interconnect Device	EX Series Switch EX0	EX Series Switch EX1
IC0	IC0, management port C0 to port 16 (ge-0/0/16)	IC0, management port C1 to port 16 (ge-0/0/16)
IC1	IC1, management port C0 to port 17 (ge-0/0/17)	IC1, management port C1 to port 17 (ge-0/0/17)

Figure 98 on page 337, Figure 99 on page 337, and Figure 100 on page 337 show the specific ports on the Node devices that you must connect to the EX Series switches when using a copper-based control plane. In general, connect the first management port from a Node device to the first EX Series switch, and the second management port to the second EX Series switch.

Figure 98: QFX3000-M QFabric System Control Plane—QFX3500 Node Device to EX Series Switch Connections




Figure 99: QFX3000-M QFabric System Control Plane—QFX3600 Node Device to EX Series Switch Connections





Figure 100: QFX3000-M QFabric System Control Plane—QFX5100 Node Device to EX Series Switch Connections



When implementing a fiber-based control plane, refer to [Figure 101 on page 338](#), [Figure 102 on page 338](#), and [Figure 103 on page 338](#) for the proper control plane connections.

Figure 101: QFX3000-M QFabric System Fiber-Based Control Plane—QFX3500 Node Device to EX Series Switch Connections




Figure 102: QFX3000-M QFabric System Fiber-Based Control Plane—QFX3600 Node Device to EX Series Switch Connections





Figure 103: QFX3000-M QFabric System Fiber-Based Control Plane—QFX5100 Node Device to EX Series Switch Connections



In this specific example, for Node device Node0, connect management port **C0** (also known as me5) to EX Series switch EX0 port **0** (ge-0/0/0), and connect management port **C1** (also known as me6) to the *second* EX Series switch EX1 port **0** (ge-0/0/0).


For the remaining seven Node devices, connect management port **C0** to the ge-0/0/X port on EX Series switch EX0 that matches the Node device number. Similarly, connect management port **C1** to the port on the *second* EX Series switch EX1 that matches the Node device number. For example, you would connect Node device Node5 to port **5** (ge-0/0/5). [Table 75 on page 339](#) shows the full set of port mappings for the Node devices in this example.

Table 75: Node Device to EX Series Switch Port Mappings

Node Device	EX Series Switch EX0	EX Series Switch EX1
Node0	Node0, management port C0 to port 0 (ge-0/0/0)	Node0, management port C1 to port 0 (ge-0/0/0)
Node1	Node1, management port C0 to port 1 (ge-0/0/1)	Node1, management port C1 to port 1 (ge-0/0/1)
Node2	Node2, management port C0 to port 2 (ge-0/0/2)	Node2, management port C1 to port 2 (ge-0/0/2)
Node3	Node3, management port C0 to port 3 (ge-0/0/3)	Node3, management port C1 to port 3 (ge-0/0/3)
Node4	Node4, management port C0 to port 4 (ge-0/0/4)	Node4, management port C1 to port 4 (ge-0/0/4)
Node5	Node5, management port C0 to port 5 (ge-0/0/5)	Node5, management port C1 to port 5 (ge-0/0/5)
Node6	Node6, management port C0 to port 6 (ge-0/0/6)	Node6, management port C1 to port 6 (ge-0/0/6)
Node7	Node7, management port C0 to port 7 (ge-0/0/7)	Node7, management port C1 to port 7 (ge-0/0/7)

Figure 104 on page 339 shows the specific uplink ports on the first EX Series switch that you must connect to the second EX Series switch. These connections create a link aggregation group (LAG) that provides redundancy and resiliency for the EX Series switch portion of the control plane. In general, connect each 1-Gigabit Ethernet uplink port from the first EX Series switch to the corresponding 1-Gigabit Ethernet uplink port on the second EX Series switch.

Figure 104: QFX3000-M QFabric System Control Plane—Inter-EX Series Switch LAG Connections



In this specific example, for EX Series switch EX0, connect uplink port **0** (ge-0/1/0) to EX Series switch EX1 uplink port **0** (ge-0/1/0). Then connect uplink port **1** (ge-0/1/1) on EX Series switch EX0 to uplink port **1** (ge-0/1/1) on EX Series switch EX1.

Table 76 on page 340 shows the port mappings for the EX Series switch LAG connections in this example.

Table 76: EX Series Switch LAG Port Mappings

EX0 and EX1	EX0	EX1
Uplink port 0	ge-0/1/0 to ge-0/1/0	ge-0/1/0 to ge-0/1/0
Uplink port 1	ge-0/1/1 to ge-0/1/1	ge-0/1/1 to ge-0/1/1

Configuration

IN THIS SECTION

- [xref target has no title]

CLI Quick Configuration

To configure the QFX3000-M QFabric system control plane EX Series switches quickly, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network, and then copy and paste the commands into the EX Series switch CLI at the **[edit]** hierarchy level.

```
set groups qfabric chassis aggregated-devices ethernet device-count 3
set groups qfabric chassis alarm management-ethernet link-down ignore
set groups qfabric chassis lcd-menu fpc 0 menu-item maintenance-menu disable
set groups qfabric protocols rstp interface ae2.0 mode point-to-point
set groups qfabric protocols rstp interface all edge
set groups qfabric protocols rstp interface all no-root-port
set groups qfabric protocols rstp bpdu-block-on-edge
set groups qfabric protocols lldp interface all
set groups qfabric ethernet-switching-options storm-control interface all bandwidth 10000
set groups qfabric vlans qfabric vlan-id 100
set groups qfabric vlans qfabric dot1q-tunneling
set groups qfabric-int interfaces <*> mtu 9216
set groups qfabric-int interfaces <*> unit 0 family ethernet-switching port-mode access
set groups qfabric-int interfaces <*> unit 0 family ethernet-switching vlan members qfabric
set groups qfabric-ae interfaces <*> aggregated-ether-options link-speed 1g
set groups qfabric-ae interfaces <*> aggregated-ether-options lacp active
set apply-groups qfabric
set chassis fpc 0 pic 1 sfpplus pic-mode 1g
set interfaces interface-range Node_Device_Interfaces member "ge-0/0/[0-15]"
set interfaces interface-range Node_Device_Interfaces description "QFabric Node Device"
```



```

set interfaces interface-range Node_Device_Interfaces mtu 9216
set interfaces interface-range Node_Device_Interfaces unit 0 family ethernet-switching port-mode access
set interfaces interface-range Node_Device_Interfaces unit 0 family ethernet-switching vlan members qfabric
set interfaces interface-range Interconnect_Device_Interfaces member "ge-0/0/[16-17]"
set interfaces interface-range Interconnect_Device_Interfaces description "QFabric Interconnect Device"
set interfaces interface-range Interconnect_Device_Interfaces mtu 9216
set interfaces interface-range Interconnect_Device_Interfaces unit 0 family ethernet-switching port-mode
    access
set interfaces interface-range Interconnect_Device_Interfaces unit 0 family ethernet-switching vlan members
    qfabric
set interfaces interface-range Director_Device_DG0_LAG_Interfaces member "ge-0/0/[20-21]"
set interfaces interface-range Director_Device_DG0_LAG_Interfaces description "QFabric Director Device -
    DG0"
set interfaces interface-range Director_Device_DG0_LAG_Interfaces ether-options speed 1g
set interfaces interface-range Director_Device_DG0_LAG_Interfaces ether-options 802.3ad ae0
set interfaces interface-range Director_Device_DG1_LAG_Interfaces member "ge-0/0/[22-23]"
set interfaces interface-range Director_Device_DG1_LAG_Interfaces description "QFabric Director Device -
    DG1"
set interfaces interface-range Director_Device_DG1_LAG_Interfaces ether-options speed 1g
set interfaces interface-range Director_Device_DG1_LAG_Interfaces ether-options 802.3ad ae1
set interfaces interface-range Control_Plane_Inter_LAG_Interfaces member "ge-0/1/[0-1]"
set interfaces interface-range Control_Plane_Inter_LAG_Interfaces description "QFabric Control Plane (Inter -
    Switch LAG)"
set interfaces interface-range Control_Plane_Inter_LAG_Interfaces ether-options 802.3ad ae2
set interfaces ae0 apply-groups qfabric-int
set interfaces ae0 apply-groups qfabric-ae
set interfaces ae0 description "QFabric Director Device - DG0"
set interfaces ae1 apply-groups qfabric-int
set interfaces ae1 apply-groups qfabric-ae
set interfaces ae1 description "QFabric Director Device - DG1"
set interfaces ae2 description "QFabric Control Plane (Inter-Switch LAG)"
set interfaces ae2 mtu 9216
set interfaces ae2 aggregated-ether-options link-speed 1g
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 unit 0 family ethernet-switching vlan members qfabric
set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_3 loss-priority low
    code-points 110
set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_3 loss-priority low
    code-points 111
set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_2 loss-priority low
    code-points 100
set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_2 loss-priority high
    code-points 101

```



```

set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_0 loss-priority low
code-points 010
set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_0 loss-priority high
code-points 001
set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class class_3 loss-priority low
code-points 110
set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class class_3 loss-priority low
code-points 111
set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class class_2 loss-priority low
code-points 100
set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class class_2 loss-priority high
code-points 101
set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class class_0 loss-priority low
code-points 010
set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class class_0 loss-priority high
code-points 001
set class-of-service forwarding-classes class class_3 queue-num 7
set class-of-service forwarding-classes class class_2 queue-num 2
set class-of-service forwarding-classes class class_0 queue-num 0
set class-of-service interfaces ge-*/0/* scheduler-map cpe_network_smap
set class-of-service interfaces ge-*/0/* unit 0 classifiers ieee-802.1 onep_qfabric_classifier
set class-of-service interfaces ge-*/0/* unit 0 classifiers inet-precedence IP_qfabric_classifier
set class-of-service interfaces ae* scheduler-map cpe_network_smap
set class-of-service interfaces ae* unit 0 classifiers ieee-802.1 onep_qfabric_classifier
set class-of-service interfaces ae* unit 0 classifiers inet-precedence IP_qfabric_classifier
set class-of-service scheduler-maps cpe_network_smap forwarding-class class_3 scheduler scheduler_3
set class-of-service scheduler-maps cpe_network_smap forwarding-class class_2 scheduler scheduler_2
set class-of-service scheduler-maps cpe_network_smap forwarding-class class_0 scheduler scheduler_0
set class-of-service schedulers scheduler_3 buffer-size percent 30
set class-of-service schedulers scheduler_3 priority strict-high
set class-of-service schedulers scheduler_2 transmit-rate percent 75
set class-of-service schedulers scheduler_2 buffer-size percent 30
set class-of-service schedulers scheduler_2 priority low
set class-of-service schedulers scheduler_0 transmit-rate percent 25
set class-of-service schedulers scheduler_0 buffer-size percent 40
set class-of-service schedulers scheduler_0 priority low
set system host-name qfabric-control-plane
set system services ssh
set system services telnet
set system services web-management http
set system syslog user * any emergency
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file messages archive world-readable

```



```

set system syslog file messages explicit-priority
set system syslog file interactive-commands interactive-commands any
set system syslog file secure authorization info
set system syslog file default-log-messages any any
set system syslog file default-log-messages structured-data
set system syslog file console any error
set system syslog time-format millisecond
set interfaces me0 unit 0 family inet address 192.168.157.26/24
set routing-options static route 0.0.0.0/0 next-hop 192.168.157.1

```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an EX Series switch for the QFX3000-M QFabric system control plane network:

1. Create a configuration group called `qfabric` to define global QFabric system control plane properties. Set up the number of aggregated Ethernet devices, configure alarm and LCD management, activate loop prevention and storm control, specify a global VLAN (VLAN 100) and 802.1q tunneling, define options for aggregated Ethernet interfaces, and apply the `qfabric` group settings to the configuration.

```

[edit]
user@switch# set groups qfabric chassis aggregated-devices ethernet device-count 3
user@switch# set groups qfabric chassis alarm management-ethernet link-down ignore
user@switch# set groups qfabric chassis lcd-menu fpc 0 menu-item maintenance-menu disable
user@switch# set groups qfabric protocols rstp interface ae2.0 mode point-to-point
user@switch# set groups qfabric protocols rstp interface all edge
user@switch# set groups qfabric protocols rstp interface all no-root-port
user@switch# set groups qfabric protocols rstp bpdu-block-on-edge
user@switch# set groups qfabric protocols lldp interface all
user@switch# set groups qfabric ethernet-switching-options storm-control interface all bandwidth 10000
user@switch# set groups qfabric vlans qfabric vlan-id 100
user@switch# set groups qfabric vlans qfabric dot1q-tunneling
user@switch# set groups qfabric-int interfaces <*> mtu 9216
user@switch# set groups qfabric-int interfaces <*> unit 0 family ethernet-switching port-mode access
user@switch# set groups qfabric-int interfaces <*> unit 0 family ethernet-switching vlan members qfabric
user@switch# set groups qfabric-ae interfaces <*> aggregated-ether-options link-speed 1g
user@switch# set groups qfabric-ae interfaces <*> aggregated-ether-options lacp active
user@switch# set apply-groups qfabric

```

2. Configure interfaces for the QFabric system control plane network. Enable the EX Series switch SFP+ uplink module for 1-Gigabit Ethernet operation. Set the interface ranges where Node devices (0 through 15), Interconnect devices (16 and 17), and Director devices (20 through 23) connect to the control

plane network through the EX Series switches. Configure the inter-EX Series switch LAG connections for the ae2 interface and apply the qfabric-int and qfabric-ae configuration groups to the aggregated Ethernet interfaces (ae0 and ae1) for the Director devices.

```
[edit]
user@switch# set chassis fpc 0 pic 1 sfppplus pic-mode 1g
user@switch# set interfaces interface-range Node_Device_Interfaces member "ge-0/0/[0-15]"
user@switch# set interfaces interface-range Node_Device_Interfaces description "QFabric Node Device"
user@switch# set interfaces interface-range Node_Device_Interfaces mtu 9216
user@switch# set interfaces interface-range Node_Device_Interfaces unit 0 family ethernet-switching
    port-mode access
user@switch# set interfaces interface-range Node_Device_Interfaces unit 0 family ethernet-switching vlan
    members qfabric
user@switch# set interfaces interface-range Interconnect_Device_Interfaces member "ge-0/0/[16-17]"
user@switch# set interfaces interface-range Interconnect_Device_Interfaces description "QFabric Interconnect
    Device"
user@switch# set interfaces interface-range Interconnect_Device_Interfaces mtu 9216
user@switch# set interfaces interface-range Interconnect_Device_Interfaces unit 0 family ethernet-switching
    port-mode access
user@switch# set interfaces interface-range Interconnect_Device_Interfaces unit 0 family ethernet-switching
    vlan members qfabric
user@switch# set interfaces interface-range Director_Device_DG0_LAG_Interfaces member "ge-0/0/[20-21]"
user@switch# set interfaces interface-range Director_Device_DG0_LAG_Interfaces description "QFabric
    Director Device - DG0"
user@switch# set interfaces interface-range Director_Device_DG0_LAG_Interfaces ether-options speed 1g
user@switch# set interfaces interface-range Director_Device_DG0_LAG_Interfaces ether-options 802.3ad
    ae0
user@switch# set interfaces interface-range Director_Device_DG1_LAG_Interfaces member "ge-0/0/[22-23]"
user@switch# set interfaces interface-range Director_Device_DG1_LAG_Interfaces description "QFabric
    Director Device - DG1"
user@switch# set interfaces interface-range Director_Device_DG1_LAG_Interfaces ether-options speed 1g
user@switch# set interfaces interface-range Director_Device_DG1_LAG_Interfaces ether-options 802.3ad
    ae1
user@switch# set interfaces interface-range Control_Plane_Inter_LAG_Interfaces member "ge-0/1/[0-1]"
user@switch# set interfaces interface-range Control_Plane_Inter_LAG_Interfaces description "QFabric
    Control Plane (Inter - Switch LAG)"
user@switch# set interfaces interface-range Control_Plane_Inter_LAG_Interfaces ether-options 802.3ad
    ae2
user@switch# set interfaces ae0 apply-groups qfabric-int
user@switch# set interfaces ae0 apply-groups qfabric-ae
user@switch# set interfaces ae0 description "QFabric Director Device - DG0"
user@switch# set interfaces ae1 apply-groups qfabric-int
user@switch# set interfaces ae1 apply-groups qfabric-ae
user@switch# set interfaces ae1 description "QFabric Director Device - DG1"
user@switch# set interfaces ae2 description "QFabric Control Plane (Inter-Switch LAG)"
```



```

user@switch# set interfaces ae2 mtu 9216
user@switch# set interfaces ae2 aggregated-ether-options link-speed 1g
user@switch# set interfaces ae2 aggregated-ether-options lacp active
user@switch# set interfaces ae2 unit 0 family ethernet-switching vlan members qfabric

```

3. Enable class of service (CoS) for the QFabric system control plane network. Establish forwarding classes, priorities, scheduler maps, classifiers, and queues for three types of traffic: control traffic, interdevice traffic, and best-effort traffic.

```

[edit]
user@switch# set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_3
    loss-priority low code-points 110
user@switch# set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_3
    loss-priority low code-points 111
user@switch# set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_2
    loss-priority low code-points 100
user@switch# set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_2
    loss-priority high code-points 101
user@switch# set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_0
    loss-priority low code-points 010
user@switch# set class-of-service classifiers ieee-802.1 onep_qfabric_classifier forwarding-class class_0
    loss-priority high code-points 001
user@switch# set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class class_3
    loss-priority low code-points 110
user@switch# set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class class_3
    loss-priority low code-points 111
user@switch# set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class class_2
    loss-priority low code-points 100
user@switch# set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class class_2
    loss-priority high code-points 101
user@switch# set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class class_0
    loss-priority low code-points 010
user@switch# set class-of-service classifiers inet-precedence IP_qfabric_classifier forwarding-class class_0
    loss-priority high code-points 001
user@switch# set class-of-service forwarding-classes class class_3 queue-num 7
user@switch# set class-of-service forwarding-classes class class_2 queue-num 2
user@switch# set class-of-service forwarding-classes class class_0 queue-num 0
user@switch# set class-of-service interfaces ge-*/0/* scheduler-map cpe_network_smap
user@switch# set class-of-service interfaces ge-*/0/* unit 0 classifiers ieee-802.1 onep_qfabric_classifier
user@switch# set class-of-service interfaces ge-*/0/* unit 0 classifiers inet-precedence IP_qfabric_classifier
user@switch# set class-of-service interfaces ae* scheduler-map cpe_network_smap
user@switch# set class-of-service interfaces ae* unit 0 classifiers ieee-802.1 onep_qfabric_classifier
user@switch# set class-of-service interfaces ae* unit 0 classifiers inet-precedence IP_qfabric_classifier

```



```

user@switch# set class-of-service scheduler-maps cpe_network_smap forwarding-class class_3 scheduler
scheduler_3
user@switch# set class-of-service scheduler-maps cpe_network_smap forwarding-class class_2 scheduler
scheduler_2
user@switch# set class-of-service scheduler-maps cpe_network_smap forwarding-class class_0 scheduler
scheduler_0
user@switch# set class-of-service schedulers scheduler_3 buffer-size percent 30
user@switch# set class-of-service schedulers scheduler_3 priority strict-high
user@switch# set class-of-service schedulers scheduler_2 transmit-rate percent 75
user@switch# set class-of-service schedulers scheduler_2 buffer-size percent 30
user@switch# set class-of-service schedulers scheduler_2 priority low
user@switch# set class-of-service schedulers scheduler_0 transmit-rate percent 25
user@switch# set class-of-service schedulers scheduler_0 buffer-size percent 40
user@switch# set class-of-service schedulers scheduler_0 priority low

```

4. Configure settings to enable the EX Series switches to interoperate with your management network. Set a hostname, system services (such as Telnet), system log thresholds, management interface parameters, default routes, and any additional preferences you might have.

```

[edit]
user@switch# set system host-name qfabric-control-plane
user@switch# set system services ssh
user@switch# set system services telnet
user@switch# set system services web-management http
user@switch# set system syslog user * any emergency
user@switch# set system syslog file messages any notice
user@switch# set system syslog file messages authorization info
user@switch# set system syslog file messages archive world-readable
user@switch# set system syslog file messages explicit-priority
user@switch# set system syslog file interactive-commands interactive-commands any
user@switch# set system syslog file secure authorization info
user@switch# set system syslog file default-log-messages any any
user@switch# set system syslog file default-log-messages structured-data
user@switch# set system syslog file console any error
user@switch# set system syslog time-format millisecond
user@switch# set interfaces me0 unit 0 family inet address 192.168.157.26/24
user@switch# set routing-options static route 0.0.0.0/0 next-hop 192.168.157.1

```

Results

To view the configuration, issue the **show** command in configuration mode or the **show configuration** command in operational mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

The following configuration is the standard configuration that applies universally to both EX Series switches in your QFabric system control plane network.

```
[edit]
groups {
  qfabric {
    chassis {
      aggregated-devices {
        ethernet {
          device-count 3;
        }
      }
      alarm {
        management-ethernet {
          link-down ignore;
        }
      }
      lcd-menu {
        fpc 0 {
          maintenance-menu disable;
        }
      }
    }
  }
  protocols {
    rstp {
      interface ae2.0 {
        mode point-to-point;
      }
      interface all {
        edge;
        no-root-port;
      }
      bpdu-block-on-edge;
    }
    lldp {
      interface all;
    }
  }
  ethernet-switching-options {
    storm-control {
      interface all {
        bandwidth 10000;
      }
    }
  }
}
```



```

vllans {
  qlfabric {
    vllan-id 100;
    dot1q-tunneling;
  }
}
qlfabric-int {
  interfaces {
    <*> {
      mtu 9216;
      unit 0 {
        family ethernet-switching {
          port-mode access;
          vllan {
            members qlfabric;
          }
        }
      }
    }
  }
}
qlfabric-ae {
  interfaces {
    <*> {
      aggregated-ether-options {
        link-speed 1g;
        lacp {
          active;
        }
      }
    }
  }
}
apply-groups [qlfabric];
chassis {
  fpc 0 {
    pic 1 {
      sfpplus {
        pic-mode 10g;
      }
    }
  }
}

```



```

}
interfaces {
  interface-range Node_Device_Interfaces {
    member "ge-0/0/[0-15]";
    description "QFabric Node Device";
    mtu 9216;
    unit 0 {
      family ethernet-switching {
        port-mode access;
        vlan {
          members qfabric;
        }
      }
    }
  }
  interface-range Interconnect_Device_Interfaces {
    member "ge-0/0/[16-17]";
    description "QFabric Interconnect Device";
    mtu 9216;
    unit 0 {
      family ethernet-switching {
        port-mode access;
        vlan {
          members qfabric;
        }
      }
    }
  }
  interface-range Director_Device_DG0_LAG_Interfaces {
    member "ge-0/0/[20-21]";
    description "QFabric Director Device - DG0";
    ether-options {
      speed {
        1g;
      }
      802.3ad ae0;
    }
  }
  interface-range Director_Device_DG1_LAG_Interfaces {
    member "ge-0/0/[22-23]";
    description "QFabric Director Device - DG1";
    ether-options {
      speed {
        1g;
      }
    }
  }
}

```



```

    }
    802.3ad ae1;
  }
}
interface-range Control_Plane_Inter_LAG_Interfaces {
  member "ge-0/1/[0-1]";
  description "QFabric Control Plane (Inter-Switch LAG)";
  ether-options {
    802.3ad ae2;
  }
}
ae0 {
  apply-groups [ qfabric-int qfabric-ae ];
  description "QFabric Director Device - DG0";
}
ae1 {
  apply-groups [ qfabric-int qfabric-ae ];
  description "QFabric Director Device - DG1";
}
ae2 {
  description "QFabric Control Plane (Inter-Switch LAG)";
  mtu 9216;
  aggregated-ether-options {
    link-speed 1g;
    lacp {
      active;
    }
  }
  unit 0 {
    family ethernet-switching {
      vlan {
        members qfabric;
      }
    }
  }
}
}
class-of-service {
  classifiers {
    ieee-802.1 onep_qfabric_classifier {
      forwarding-class class_3 {
        loss-priority low code-points [ 110 111 ];
      }
      forwarding-class class_2 {

```



```

        loss-priority low code-points 100;
        loss-priority high code-points 101;
    }
    forwarding-class class_0 {
        loss-priority low code-points 010;
        loss-priority high code-points 001;
    }
}
inet-precedence IP_qfabric_classifier {
    forwarding-class class_3 {
        loss-priority low code-points [ 110 111 ];
    }
    forwarding-class class_2 {
        loss-priority low code-points 100;
        loss-priority high code-points 101;
    }
    forwarding-class class_0 {
        loss-priority low code-points 010;
        loss-priority high code-points 001;
    }
}
}
forwarding-classes {
    class class_3 queue-num 7;
    class class_2 queue-num 2;
    class class_0 queue-num 0;
}
interfaces {
    ge-*/0/* {
        scheduler-map cpe_network_smap;
        unit 0 {
            classifiers {
                ieee-802.1 onep_qfabric_classifier;
                inet-precedence IP_qfabric_classifier;
            }
        }
    }
}
ae* {
    scheduler-map cpe_network_smap;
    unit 0 {
        classifiers {
            ieee-802.1 onep_qfabric_classifier;
            inet-precedence IP_qfabric_classifier;
        }
    }
}

```



```

    }
  }
}
scheduler-maps {
  cpe_network_smap {
    forwarding-class class_3 scheduler scheduler_3;
    forwarding-class class_2 scheduler scheduler_2;
    forwarding-class class_0 scheduler scheduler_0;
  }
}
schedulers {
  scheduler_3 {
    buffer-size percent 30;
    priority strict-high;
  }
  scheduler_2 {
    transmit-rate percent 75;
    buffer-size percent 30;
    priority low;
  }
  scheduler_0 {
    transmit-rate percent 25;
    buffer-size percent 40;
    priority low;
  }
}
}

```

The following portion of the configuration applies to the specific requirements of your management network. Modify this section to meet the needs of your network.

```

[edit]
system {
  host-name qfabric-control-plane;
  services {
    ssh;
    telnet;
    web-management {
      http;
    }
  }
}
syslog {
  user * {
    any emergency;
  }
}

```



```

    }
    file messages {
        any notice;
        authorization info;
        archive world-readable;
        explicit-priority;
    }
    file interactive-commands {
        interactive-commands any;
    }
    file secure {
        authorization info;
    }
    file default-log-messages {
        any any;
        structured-data;
    }
    file console {
        any error;
    }
    time-format millisecond;
}
}
interfaces {
    me0 {
        unit 0 {
            family inet {
                address 192.168.157.26/24;
            }
        }
    }
}
routing-options {
    static {
        route 0.0.0.0/0 next-hop 192.168.157.1;
    }
}
}

```

To verify the syntax of your configuration prior to committing it, enter **commit check** from configuration mode. If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying the QFX3000-M QFabric System Control Plane—EX Series Switch EX0 | 354](#)
- [Verifying the QFX3000-M QFabric System Control Plane—EX Series Switch EX1 | 356](#)

Confirm that the EX Series switch configuration is working properly.

Verifying the QFX3000-M QFabric System Control Plane—EX Series Switch EX0

Purpose

Verify that the control plane is properly connected on your first EX Series switch.

Action

Connect to the Junos OS CLI of EX Series switch EX0, either from your management network or from the console port of the switch. In operational mode, enter the **show interfaces terse** command.

Sample Output

user@ex0> **show interfaces terse**

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	eth-switch		
ge-0/0/1	up	up			
ge-0/0/1.0	up	up	eth-switch		
ge-0/0/2	up	up			
ge-0/0/2.0	up	up	eth-switch		
ge-0/0/3	up	up			
ge-0/0/3.0	up	up	eth-switch		
ge-0/0/4	up	up			
ge-0/0/4.0	up	up	eth-switch		
ge-0/0/5	up	up			
ge-0/0/5.0	up	up	eth-switch		
ge-0/0/6	up	up			
ge-0/0/6.0	up	up	eth-switch		
ge-0/0/7	up	up			
ge-0/0/7.0	up	up	eth-switch		

ge-0/0/8	up	down	
ge-0/0/8.0	up	down	eth-switch
ge-0/0/9	up	down	
ge-0/0/9.0	up	down	eth-switch
ge-0/0/10	up	down	
ge-0/0/10.0	up	down	eth-switch
ge-0/0/11	up	down	
ge-0/0/11.0	up	down	eth-switch
ge-0/0/12	up	down	
ge-0/0/12.0	up	down	eth-switch
ge-0/0/13	up	down	
ge-0/0/13.0	up	down	eth-switch
ge-0/0/14	up	down	
ge-0/0/14.0	up	down	eth-switch
ge-0/0/15	up	down	
ge-0/0/15.0	up	down	eth-switch
ge-0/0/16	up	up	
ge-0/0/16.0	up	up	eth-switch
ge-0/0/17	up	up	
ge-0/0/17.0	up	up	eth-switch
ge-0/0/18	up	down	
ge-0/0/18.0	up	down	eth-switch
ge-0/0/19	up	down	
ge-0/0/19.0	up	down	eth-switch
ge-0/0/20	up	up	
ge-0/0/20.0	up	up	aenet --> ae0.0
ge-0/0/21	up	up	
ge-0/0/21.0	up	up	aenet --> ae0.0
ge-0/0/22	up	up	
ge-0/0/22.0	up	up	aenet --> ae1.0
ge-0/0/23	up	up	
ge-0/0/23.0	up	up	aenet --> ae1.0
ge-0/1/0	up	up	
ge-0/1/0.0	up	up	aenet --> ae2.0
ge-0/1/1	up	up	
ge-0/1/1.0	up	up	aenet --> ae2.0
vcp-0	up	down	
vcp-0.32768	up	down	
vcp-1	up	down	
vcp-1.32768	up	down	
ae0	up	up	
ae0.0	up	up	eth-switch
ae1	up	up	
ae1.0	up	up	eth-switch


```

ae2                up    up
ae2.0              up    up    eth-switch
bme0               up    up
bme0.32768         up    up    inet    128.0.0.1/2
                                   128.0.0.16/2
                                   128.0.0.32/2
                                   tnp    0x10
dsc                up    up
gre                up    up
ipip               up    up
lo0               up    up
lo0.0             up    up    inet    127.0.0.1    --> 0/0
lsi               up    up
me0               up    up
me0.0             up    up    inet    192.168.157.26/24
mtun              up    up
pimd              up    up
pime              up    up
tap               up    up
vlan              up    up
vme               up    down

```

Meaning

In the output of the **show interfaces terse** command, if all interfaces that connect to the QFabric system devices are listed as **up** (such as ge-0/0/16 and ge-0/0/17 for the Interconnect devices; ge-0/0/20 through ge-0/0/23 for the Director devices; ge-0/0/0 through ge-0/0/7 for the Node devices; and ge-0/1/0 and ge-0/1/1 for the inter-EX Series switch connections), the control plane is properly connected.

Verifying the QFX3000-M QFabric System Control Plane—EX Series Switch EX1

Purpose

Verify that the control plane is properly connected on your second EX Series switch.

Action

Connect to the Junos OS CLI of EX Series switch EX1, either from your management network or from the console port of the switch. In operational mode, enter the **show interfaces terse** command.

Sample Output

```
user@ex1> show interfaces terse
```


Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	eth-switch		
ge-0/0/1	up	up			
ge-0/0/1.0	up	up	eth-switch		
ge-0/0/2	up	up			
ge-0/0/2.0	up	up	eth-switch		
ge-0/0/3	up	up			
ge-0/0/3.0	up	up	eth-switch		
ge-0/0/4	up	up			
ge-0/0/4.0	up	up	eth-switch		
ge-0/0/5	up	up			
ge-0/0/5.0	up	up	eth-switch		
ge-0/0/6	up	up			
ge-0/0/6.0	up	up	eth-switch		
ge-0/0/7	up	up			
ge-0/0/7.0	up	up	eth-switch		
ge-0/0/8	up	down			
ge-0/0/8.0	up	down	eth-switch		
ge-0/0/9	up	down			
ge-0/0/9.0	up	down	eth-switch		
ge-0/0/10	up	down			
ge-0/0/10.0	up	down	eth-switch		
ge-0/0/11	up	down			
ge-0/0/11.0	up	down	eth-switch		
ge-0/0/12	up	down			
ge-0/0/12.0	up	down	eth-switch		
ge-0/0/13	up	down			
ge-0/0/13.0	up	down	eth-switch		
ge-0/0/14	up	down			
ge-0/0/14.0	up	down	eth-switch		
ge-0/0/15	up	down			
ge-0/0/15.0	up	down	eth-switch		
ge-0/0/16	up	up			
ge-0/0/16.0	up	up	eth-switch		
ge-0/0/17	up	up			
ge-0/0/17.0	up	up	eth-switch		
ge-0/0/18	up	down			
ge-0/0/18.0	up	down	eth-switch		
ge-0/0/19	up	down			
ge-0/0/19.0	up	down	eth-switch		
ge-0/0/20	up	up			
ge-0/0/20.0	up	up	aenet	--> ae0.0	
ge-0/0/21	up	up			


```

ge-0/0/21.0      up    up    aenet    --> ae0.0
ge-0/0/22        up    up
ge-0/0/22.0      up    up    aenet    --> ae1.0
ge-0/0/23        up    up
ge-0/0/23.0      up    up    aenet    --> ae1.0
ge-0/1/0         up    up
ge-0/1/0.0       up    up    aenet    --> ae2.0
ge-0/1/1         up    up
ge-0/1/1.0       up    up    aenet    --> ae2.0
vcp-0            up    down
vcp-0.32768      up    down
vcp-1            up    down
vcp-1.32768      up    down
ae0              up    down
ae0.0            up    down    eth-switch
ae1              up    down
ae1.0            up    down    eth-switch
ae2              up    up
ae2.0            up    up    eth-switch
bme0             up    up
bme0.32768       up    up    inet      128.0.0.1/2
                                   128.0.0.16/2
                                   128.0.0.32/2
                                   tnp      0x10
dsc              up    up
gre              up    up
ipip             up    up
lo0              up    up
lo0.0            up    up    inet      127.0.0.1      --> 0/0
lsi              up    up
me0              up    up
me0.0            up    up    inet      192.168.157.26/24
mtun             up    up
pimd             up    up
pime             up    up
tap              up    up
vlan             up    up
vme              up    down

```

Meaning

In the output of the **show interfaces terse** command, if all interfaces that connect to the QFabric system devices are listed as **up** (such as ge-0/0/16 and ge-0/0/17 for the Interconnect devices; ge-0/0/20 through ge-0/0/23 for the Director devices; ge-0/0/0 through ge-0/0/7 for the Node devices; and ge-0/1/0 and ge-0/1/1 for the inter-EX Series switch connections), the control plane is properly connected.

RELATED DOCUMENTATION

QFX3000-M QFabric System Installation Overview 109
Installing and Connecting a QFX3100 Director Device 165
Installing and Connecting a QFX3600 or QFX3600-I Device 177
Installing and Connecting a QFX3500 Device 199
<i>Installing and Connecting an EX4200 Switch</i>
Understanding the QFabric System Control Plane 49

Importing a QFX3000-M QFabric System Control Plane EX Series Switch Configuration with a USB Flash Drive

There are two methods of importing the configuration file to the QFX3000-M QFabric system control plane EX Series switches.

- Download the configuration file onto a USB flash drive from the Juniper Networks software download site before inserting the USB flash drive into the EX Series switch USB port
- Copy and paste the configuration from “[Example: Configuring EX Series Switches for the QFX3000-M QFabric System Control Plane](#)” on page 329.

NOTE: The EX Series switch configuration is the same for both the copper-based and fiber-based QFX3000-M QFabric system control plane networks.

Before you begin:

- Rack, mount, and install your QFabric system hardware (Director group, Interconnect devices, and Node devices). For more information, see “[Installing and Connecting a QFX3100 Director Device](#)” on page 165, “[Installing and Connecting a QFX3600 or QFX3600-I Device](#)” on page 177, and “[Installing and Connecting a QFX3500 Device](#)” on page 199.
- Rack, mount, and install your EX Series switches for the QFabric system control plane. For more information, see *Installing and Connecting an EX4200 Switch* or *Installing and Connecting an EX4300 Switch*.
- Select a USB flash drive that meets the EX Series switch USB port specifications. See *USB Port Specifications for an EX Series Switch*.
- Use a computer or other device to download the configuration file from the Internet and copy it to the USB flash drive.

To import the control plane EX Series switch configuration file onto a USB flash drive:

1. In a browser, go to <https://www.juniper.net/support/downloads/junos.html> .

The Junos Platforms - Download Software page appears.

NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp> .

2. In the QFX Series box, select **QFX3000-M QFabric System**.

The QFX3000-M QFabric System - Download Software page appears.

3. Click the **Software** tab and select the software release number from the **Release** list that appears to the right of the Software tab.

A login screen appears.

4. In the QFabric System Control Plane Network section, select **QFX3000-M Control Plane Network Configuration**.

A login screen appears.

5. Enter your user ID and password and click **Login**.
6. Read the End User License Agreement, select the **I agree** option button, and then click **Proceed**.
7. Save the configuration file onto the USB flash drive using your computer or other device.
8. Insert the USB flash drive into the USB port on the EX Series switch.
9. Save the file to the `/var/home/username` directory on the EX Series switch.
10. Load the configuration file into the switch.

```
user@switch# load override filename
```

11. Commit the configuration.

```
user@switch# commit  
Load complete
```

12. Remove the USB flash drive from the switch.

RELATED DOCUMENTATION

| [Example: Configuring EX Series Switches for the QFX3000-M QFabric System Control Plane](#) | 329

Generating the MAC Address Range for a QFabric System

Each QFabric system requires a range of reserved MAC addresses that is assigned by Juniper Networks. You must specify the MAC address range when you perform the initial setup of the QFX3100 Director group (see “[Performing the QFabric System Initial Setup on a QFX3100 Director Group](#)” on page 363). Additionally, refer to [Activate Your QFabric System](#) for more information.

When you purchase a QFabric system, you receive an e-mail containing a software serial number from Juniper Networks. You can use the software serial number to generate the MAC address range for your QFabric system.

To generate the MAC address range for a QFabric system:

1. In a browser, log in to the Juniper Networks License Management System at <https://www.juniper.net/lcrs/license.do>.

The Manage Product Licenses page appears.

NOTE: To access the licensing site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

2. On the Generate Licenses tab, select **QFX Series Product** from the drop-down list, and click **Go**.

The Generate Licenses - QFX Series Product page appears.

3. Select the **QFX Series Product Fabric** option button, and click **Continue**.

The Generate Licenses - QFX Series Product Fabrics page appears.

4. In the **Software Serial No** field, enter the software serial number for your QFabric system, and press the Tab key.

The starting MAC address and number of MAC addresses for your QFabric system are displayed.

5. (Optional) Click **Download/Email MAC Address** to download or e-mail the MAC address range.

The Download/Email MAC Address page appears.

To download the MAC address range:

- Select the **Download to this computer** option button, and click **OK**.

To e-mail the MAC address range:

- Select the **Send e-mail to e-mail ID** option button, and click **OK**.

RELATED DOCUMENTATION

| [Performing the QFabric System Initial Setup on a QFX3100 Director Group](#) | 363

Performing the QFabric System Initial Setup on a QFX3100 Director Group

You must perform the initial setup of the QFX3100 Director group through the console port. (Before configuring the QFX3100 Director group, see [“Installing and Connecting a QFX3100 Director Device” on page 165.](#))

Before you begin connecting and configuring a QFX3100 Director group, set the following parameter values on the console server or PC:

- Baud Rate—9600
- Flow Control—None
- Data—8
- Parity—None
- Stop Bits—1
- DCD State—Disregard

NOTE: When you use the SecureCRT client to connect to a Director device for the initial setup of a QFabric system, the backspace key does not work. As a workaround, use the Shift+Delete key combination in SecureCRT as a backspace key equivalent or use a different UNIX client to support the backspace key natively.

The initial setup requires that you specify certain values for your QFabric system. These include:

- Software serial number for your QFabric system (found in the e-mail containing the software serial number that you received from Juniper Networks when you purchased your QFabric system)
- IP addresses and a default gateway IP address for your QFabric system default partition

- IP addresses for your Director group device management ports
 - Range of reserved MAC addresses for your QFabric system (see [“Generating the MAC Address Range for a QFabric System”](#) on page 362 or [Activate Your QFabric System](#) for this information)
 - Root password for your Director group
 - Root password for the QFabric system components such as the Node devices, Interconnect devices, and infrastructure
- [Performing an Initial Setup | 364](#)
 - [Restoring a Backup Configuration | 368](#)

Performing an Initial Setup

The initial setup can be performed either manually or by using a previously saved backup configuration.

To connect and configure the QFX3100 Director group manually from the console:

1. Connect the console port of one of the Director devices to a laptop or PC using an RJ-45 to DB-9 rollover cable. An RJ-45 to DB-9 rollover cable is supplied with each QFX3100 Director device. The console (**CONSOLE**) port is located on the front panel of the device.
2. Log in as **root**. If the software booted before you connected to the console port, you might need to press the Enter key for the prompt to appear.

```
dg0 login: root
```

NOTE: The prompt is either **dg0 login** or **dg1 login** depending on the Director device to which you connected your cable.

3. For manual configuration or for initial installation, enter **no** when prompted to specify the backup file. The current Director device configuration is displayed.

```
Initial Configuration
```

```
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
```

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

```
Specify a backup file? [y/n]: n
```

```
Existing local configuration:
```


4. Enter the IP addresses and prefixes for both Director devices.

NOTE: The Director group devices and QFabric system default partition IP addresses must be on the same subnet as your management network.

Please enter the Director Group 0 IP address and prefix: ***ip address/prefix***

Please enter the Director Group 1 IP address and prefix: ***ip address/prefix***

Please enter the Director Group Subnet Mask: ***subnet mask***

5. Enter the gateway IP address for the Director group.

Please enter the Director Group gateway IP address: ***gateway ip address***

6. Enter the default partition IP address. (You will use this address to log in to the QFabric system on subsequent connections.)

Please enter the QFabric default partition IP address: ***ip address***

7. (Optional) Enter the IPv6 addresses for both Director devices and the gateway IPv6 address for the Director group.

Would you like to input IPv6 addresses for Director Group nodes? (y/n): ***y***

Please enter the Director Group 0 IPv6 address or 'y' to use /0: ***IPv6 address***

Please enter the Director Group 1 IPv6 address or 'y' to use /0: ***IPv6 address***

Please enter the Director Group gateway IPv6 address or 'y' to use /0 : ***IPv6 address***

8. Enter the MAC address information.

Please enter the starting MAC address: ***mac address***

Please enter the number of MAC addresses: ***number of mac addresses***

NOTE: The minimum number of MAC addresses accepted is 4000.

9. Enter the QFabric system software serial number.

Please enter the QFabric serial ID: ***serial id***

10. Create the Director device root password.

Please enter a Director device root password: **director-device-password**

Please re-enter password: **director-device password**

11. Create a password for the QFabric system components.

NOTE: If you need to change the component password after the QFabric system is operational, issue the **device-authentication** statement at the **[edit system]** hierarchy level in the QFabric default partition CLI.

Please enter a password for QFabric components (Node devices, Interconnect devices, and infrastructure): **component-password**

Please re-enter password: **component-password**

Note: please record your passwords for recovery purposes.



CAUTION: Carefully save your passwords for future reference, because some cannot be recovered on a QFabric system.

12. Enter the QFabric system platform type.

Supported platform types:

1. QFX3000-G

2. QFX3000-M

Please select product type: **number corresponding to platform type**

13. Confirm the initial configuration. Ensure that the information is accurate before proceeding.

Does the following configuration appear correct?

```

Director Group 0 IPv4/Prefix          [10.94.200.9/24]
Director Group 1 IPv4/Prefix          [10.94.200.10/24]
Director Group IPv4 Gateway           [10.94.200.250]
Director Group 0 IPv6/Prefix          [2000:1:2:3::a5e:c809/64]
Director Group 1 IPv6/Prefix          [2000:1:2:3::a5e:c80a/64]
Director Group IPv6 Gateway           [2000:0001:0002:0003:0226:88ff:fe7b:e880]
QFabric Default Partition (IPv6 address) [2000:1:2:3::0a5e:c802/64]
```


QFabric Serial ID	[qfsn-0123456789]
Director Device Password	[*****]

NOTE: Only addresses of the IP version(s) you entered will appear in the configuration.

14. Confirm the initial setup.

[y/n]: y



CAUTION: Resetting this initial configuration requires assistance from Juniper Networks customer support or [“Performing a QFabric System Recovery Installation on the Director Group” on page 835](#). As a result, make sure you are certain the values you entered are correct before you enter **yes**.

15. The director device displays the configuration.

```

Saving temporary configuration...
Configuring peer...
Configuring local interfaces...
Configuring interface eth0 with [10.49.214.74/24:10.49.214.254]
Configured interface eth0 with [10.49.214.74/24:10.49.214.254]
Configuring QFabric software with an initial pool of 4000 MAC addresses
[00:11:00:00:00:00 - 00:11:00:00:0f:3b]
Configuring QFabric address [10.49.214.150]
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.
Configuration complete. Director Group services will auto start within 30
seconds.
```


Restoring a Backup Configuration

Before you restore a backup configuration for the Director group:

- You must have a backup configuration file. You create the backup file with the **request system software configuration-backup** command and save it on an external USB flash drive.
- If you need to reinstall the system software, perform that operation first (see [“Performing a QFabric System Recovery Installation on the Director Group” on page 835](#)).

To connect and configure the Director group with a backup configuration:

1. Log in as **root**. If the software booted before you connected to the console port, you might need to press the Enter key for the prompt to appear.

```
dg0 login: root
```

NOTE: The prompt is either **dg0 login** or **dg1 login** depending on the Director device to which you connected your cable.

2. To use a previously saved backup configuration, enter **yes** when prompted to specify the backup file and then enter the path and filename of the backup configuration.

```
Specify a back up file? [y/n]: y
```

```
Please specify the full path of the configuration backup file: path/filename
```

3. Confirm the restoration of the configuration from the backup. Ensure that the information is accurate before proceeding.

```
Does the following configuration appear correct?
```

```
Director Group 0 IP/Prefix          [10.49.214.74/24]
Director Group 1 IP/Prefix          [10.49.214.75/24]
Director Group Gateway               [10.49.214.254]
Starting MAC address                 [00:11:00:00:00:00]
Number of MAC addresses               [4000]
QFabric Default Partition IP         [10.49.214.150]
QFabric serial ID                    [qfsn-123456789]
Director Device Password             [*****]
```


QFabric component Password	[*****]
Product Type:	[QFX3000-G]

4. Confirm the backup restoration.

[y/n]: **y**

The Director device displays the configuration.

```

Saving temporary configuration...
Configuring peer...
Configuring local interfaces...
Configuring interface eth0 with [10.49.214.74/24:10.49.214.254]
Configured interface eth0 with [10.49.214.74/24:10.49.214.254]
Configuring QFabric software with an initial pool of 4000 MAC addresses
[00:11:00:00:00:00 - 00:11:00:00:0f:3b]
Configuring QFabric address [10.49.214.150]
Reconfiguring QFabric software static configuration
Applying the new Director device password
Applying the QFabric component password
Configuration complete. Director Group services will auto start within 30
seconds.
```

RELATED DOCUMENTATION

[Generating the MAC Address Range for a QFabric System | 362](#)

[Gaining Access to the QFabric System Through the Default Partition | 375](#)

[QFabric System Initial and Default Configuration Information | 319](#)

[Installing and Connecting a QFX3100 Director Device | 165](#)

[Performing a QFabric System Recovery Installation on the Director Group | 835](#)

[*request system software configuration-backup*](#)

[device-authentication | 464](#)

QFabric System Configuration

IN THIS CHAPTER

- Understanding QFabric System Administration Tasks and Utilities | 371
- Gaining Access to the QFabric System Through the Default Partition | 375
- Example: Configuring QFabric System Login Classes | 377
- Example: Configuring RADIUS Authentication on a QFabric System | 387
- Configuring Aliases for the QFabric System | 388
- Configuring Node Groups for the QFabric System | 402
- Configuring the Port Type on QFX3600 Node Devices | 408
- Configuring the QSFP+ Port Type on QFX5100 Devices | 414
- Example: Configuring SNMP | 417
- Example: Configuring System Log Messages | 421
- Configuring Graceful Restart for QFabric Systems | 424
- Optimizing the Number of Multicast Flows on QFabric Systems | 428
- Segregating QFabric Traffic Flows With Flow Groups | 429
- Example: Creating a QFabric Flow Group | 431

Understanding QFabric System Administration Tasks and Utilities

The following items describe QFabric system components, common administration tasks that you perform on the QFabric system, or utilities that help you to manage the QFabric system and its components.

- **Converting the device mode (QFX3500 and QFX3600 devices)**—Enables you to convert a QFX3500, QFX3600, or QFX5100 device into a Node device so it can be deployed within a QFabric system. By default, QFX3500, QFX3600, and QFX5100 devices operate in *standalone* mode. Before the devices can participate within a QFabric system environment, you must change the device mode for the switch to *node-device* mode. To convert a QFX3500, QFX3600, or QFX5100 device from standalone mode to Node device mode, connect to the console port of the device, issue the **request chassis device-mode node-device** command, verify the future device mode with the **show chassis device-mode** command, connect the management port of the device to the QFabric system control plane, and reboot the device.

NOTE:

- Before you convert the device mode, you must upgrade the software on your standalone device to a QFabric system Node and Interconnect device software package that matches the QFabric system complete software package used by your QFabric system. For example, if the complete software package for your QFabric system is named **jinstall-qfabric-11.3X30.6.rpm**, you need to install the **jinstall-qfx-11.3X30.6-domestic-signed.tgz** package on your standalone device. Matching the two software packages ensures a smooth and successful addition of the device to the QFabric system inventory.
- Converting the device mode erases the switch configuration. We recommend that you save your configuration to an external server or USB flash drive before executing the device mode conversion commands and rebooting the switch.

- **QFabric system control plane Ethernet network (EX4200 or EX4300 switches to support the QFabric system)**—Provides a separate control plane network within the QFabric system to handle management traffic. This design enables the data plane network to focus on efficient, low-latency delivery of data, voice, and video traffic.
 - The QFX3000-G QFabric system control plane uses two sets of four EX4200 or EX4300 switches each, configured as a pair of Virtual Chassis to connect all components within the QFabric system. The dual Virtual Chassis architecture provides redundancy and high availability to ensure reliable QFabric system operation for the Director group, the Interconnect devices, and the Node devices.
 - The QFX3000-M QFabric system control plane uses two EX4200 or EX4300 switches to connect all components within the QFabric system. The two EX4200 or EX4300 switches provide redundancy and high availability to ensure reliable QFabric system operation for the Director group, the Interconnect devices, and the Node devices.

Because the level of detail necessary to fully understand the control plane connections, cabling, topology, and configuration is beyond the scope of this topic, see:

- *Example: Configuring the Virtual Chassis for a Copper-Based QFX3000-G QFabric System Control Plane* for information about a QFX3000-G QFabric system with a copper-based control plane

- [“Example: Configuring EX Series Switches for the QFX3000-M QFabric System Control Plane” on page 329](#) for information about a QFX3000-M QFabric system with a copper or fiber-based control plane
- **QFabric system data plane network**—Provides a separate network to handle rapid delivery of data plane traffic. The data plane uses QSFP+ interfaces and fiber-optic cabling to connect QFabric system components at speeds of 40 Gbps. By creating a redundant set of connections between the Node devices and the backplane-like Interconnect devices, the data plane enables the Node devices to appear as if they are directly connected to one another in a single tier. To view the connection status of the QFabric system data plane, issue the **show chassis fabric connectivity** command.
- **Director group (QFX3100 Director devices within a QFabric system)**—Provides a redundant, resilient platform that manages the QFabric system components. Two QFX3100 Director devices work together to ensure high availability of the system and load-balance system processes, such as the command-line interface (CLI) and shared storage. To configure the Director group for operation, install and cable two Director devices as a Director group, connect to the console port of one of the Director devices, and perform the initial setup. The setup script starts automatically the first time you power on the Director device. For more information, see [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 363](#). To monitor the status of the Director group, log in to the QFabric system default partition and issue the **show fabric administration inventory director-group status** command.
- **Automatic detection and configuration of QFabric system components**—Enables QFabric system components to join the QFabric system automatically. When you install the QFabric system, activate the control plane and Director group, and power on the Node and Interconnect devices, the Director group recognizes these devices, sends each device its own portion of the Junos OS configuration, and adds them to the QFabric system inventory. By default, each individual Node device is placed into a unique server Node group that contains only that single Node device. No configuration is required for the default assignments. The default settings can be overridden when you add Node devices into a redundant server Node group (containing a pair of Node devices) or a network Node group (that can contain up to eight Node devices, run routing protocols, and connect to external networks).
- **QFabric system Routing Engines**—Support the QFabric system by providing virtual, redundant instances of Junos OS that run on the Director group. The Routing Engines perform fabric management tasks, maintain control of the fabric, and host the operation of routing protocols for network Node groups. Because they are generated in pairs, the Routing Engines provide additional high availability for the QFabric system. No configuration is required. To view the status of the QFabric system Routing Engines, issue the **show fabric administration inventory infrastructure** command.
- **QFabric system command-line interface**—Enables you to configure all components of the QFabric system from a single location by using the Junos OS CLI. To access this central location, you need to log in to the QFabric system default partition (an IP address you specify during the initial setup of the Director group). For more information, see [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 363](#).

Most existing Junos OS configuration statements and operational mode commands are supported (for example, interfaces, VLANs, protocols, and firewall filters).

To view QFabric system components and check connectivity of the system, issue the **show fabric administration inventory** commands.

- **Alias configuration for Director devices, Interconnect devices, and Node devices**—Enables you to set user-defined aliases for QFabric system Director devices, Interconnect devices, and Node devices to facilitate usability of the QFabric system as it scales. Aliased names appear in the output of many QFabric system operational commands, such as **show fabric administration inventory**. To map the hardware serial number of a Director device, Interconnect device or Node device to a user-defined name, see [“Configuring Aliases for the QFabric System” on page 388](#).
- **Node group configuration**—Enables you to cluster several Node devices together to provide redundancy, resiliency, and high availability at the ingress and egress points of the QFabric system. There are two types of Node groups you can configure:
 - **Redundant server Node group**—Enables the grouped Node devices to connect the QFabric system to local servers and storage devices. A redundant server Node group can contain a maximum of two Node devices and supports LAG connections that can span both devices.

NOTE: The Node devices in a redundant server Node group must be of the same type, such as a QFX3500 Node, a QFX3600 Node, or a QFX5100 Node. For example, you cannot add a QFX3500 and a QFX3600 Node device to the same redundant server Node group.

- **Network Node group**—Enables the grouped Node devices to connect the QFabric system to external networks and run routing protocols such as BGP and OSPF. A network Node group can contain up to eight Node devices and supports LAG connections.

NOTE:

- The name of the network Node group in the default partition, NW-NG-0, is preset. You must use this name when adding Node devices to the network Node group. You cannot specify a different name.
- When you configure routing protocols on the QFabric system, you must use interfaces from the Node devices assigned to the network Node group. If you try to configure routing protocols on interfaces from the Node devices assigned to server Node groups, the configuration commit operation fails.

To configure a redundant server Node group, include two Node devices with the **node-device node-device-name** statement at the **[edit fabric resources node-group node-group-name]** hierarchy level.

To configure a network Node group, include the **network-domain** statement at the **[edit fabric resources node-group NW-NG-0]** hierarchy level. In addition, include between two and eight Node devices with

the **node-device** *node-device-name* statement at the [edit fabric resources node-group NW-NG-0] hierarchy level.

RELATED DOCUMENTATION

[Converting the Device Mode for a QFabric System Component | 322](#)

Example: Configuring the Virtual Chassis for a Copper-Based QFX3000-G QFabric System Control Plane

[Example: Configuring EX Series Switches for the QFX3000-M QFabric System Control Plane | 329](#)

[show chassis fabric connectivity | 749](#)

[Performing the QFabric System Initial Setup on a QFX3100 Director Group | 363](#)

[show fabric administration inventory director-group status | 786](#)

[show fabric administration inventory infrastructure | 793](#)

[show fabric administration inventory | 780](#)

[Configuring Aliases for the QFabric System | 388](#)

[Configuring Node Groups for the QFabric System | 402](#)

Gaining Access to the QFabric System Through the Default Partition

This topic explains how to log in to the QFabric system default partition so you can access the Junos OS command-line interface (CLI) and configure the system.

Before you access the QFabric system default partition:

- Install the QFabric system hardware components, including connecting the network and power cables.
- Convert any QFX3500, QFX3600, or QFX5100 standalone devices to *node-device* mode.
- Connect all components to the control plane Ethernet network.
- Turn on the Director group and run the initial setup script. Remember to write down the IP address of the default partition, which must be on the same subnet as your management network.

To access the default partition:

1. Open an SSH connection to the QFabric default partition. Use the IP address you set for the default partition as part of the QFabric initial setup procedure. In your network, you can simplify access to the QFabric system by mapping the default partition IP address to a name.

```
[root@customer ~]# ssh root@192.168.1.49
```



```
Last login: Fri Sep  2 21:34:54 2011 from customer
Juniper QFabric Director 11.3.5043 2011-08-26 18:05:21 UTC
```

```
RUNNING ON DIRECTOR DEVICE : dg1
root@qfabric>
```

NOTE: The QFabric system is load balanced, so the CLI session might be hosted on either Director device **DG0** or **DG1**.

2. Enter configuration mode (the default mode in the QFabric system is **configure private**), configure a root password and hostname for the default partition, and assign QFabric administrator privileges to the root user.

```
root@qfabric> configure
```

```
warning: Using private edit on QF/Director
warning: uncommitted changes will be discarded on exit
Entering configuration mode

[edit]
```

```
root@qfabric# set system root-authentication plain-text-password
```

```
New password: My-Password
Retype new password: My-Password
```

```
root@qfabric# set system root-authentication remote-debug-permission qfabric-admin
```

```
root@qfabric# set system host-name my-qfabric
```

```
[edit]
```

```
root@qfabric# commit
```

```
commit complete
```

```
[edit]
```

```
root@my-qfabric#
```


3. Configure your QFabric system as needed. You can configure routing protocols, interfaces, VLANs, and other features as needed. Keep in mind that interfaces require the four-level interface naming convention (device-name:*fpc/pic/port*).

RELATED DOCUMENTATION

[Performing the QFabric System Initial Setup on a QFX3100 Director Group | 363](#)

[QFabric System Initial and Default Configuration Information | 319](#)

[Understanding Interfaces on the QFabric System | 15](#)

Example: Configuring QFabric System Login Classes

IN THIS SECTION

- [Requirements | 377](#)
- [Overview | 378](#)
- [Configuration | 379](#)
- [Verification | 381](#)

This example shows you how to assign the correct login class to users so they can access components within a QFabric system.

Requirements

This example uses the following hardware and software components:

- One QFX3000-G QFabric system containing:
 - Two QFX3100 Director devices
 - Two QFX3008-I Interconnect devices
 - Eight QFX3500 Node devices
 - Junos OS Release 12.2 for these QFX Series components
- Eight EX4200 switches, used to make two redundant Virtual Chassis with four members apiece

- Junos OS Release 12.1R1.9 for the EX Series switches used in the Virtual Chassis

Before you begin:

- Perform the initial setup of the QFabric system on the Director group, which includes the creation of a username and password for the QFabric system components. See [“Performing the QFabric System Initial Setup on a QFX3100 Director Group” on page 363](#).

Overview

The QFabric system offers three special preset login classes that provide different levels of access to individual components within a QFabric system (such as Node devices and Interconnect devices). The *qfabric-admin* class provides the ability to log in to individual QFabric system components and manage them. The *qfabric-operator* class enables the user to log in to individual components and view component-level operations and configurations. The *qfabric-user* class prevents access to individual QFabric system components.

You include these classes in your configuration at the **[edit system login user username authentication remote-debug-permission]** hierarchy level. The key task is to decide which class you should apply to users based on their need to access QFabric system components.

NOTE: To set QFabric system login classes for a root user, include the **remote-debug-permission** statement at the **[edit system root-authentication]** hierarchy level and specify the *qfabric-admin* class.

If you assign the *qfabric-admin* or the *qfabric-operator* class to a user, the QFabric system maps the user to a list of authorized users who are permitted to access components. To facilitate ease of use, the QFabric system uses the component password you specified during the initial setup of the Director group. When users assigned the *qfabric-admin* or the *qfabric-operator* class log in to a component by issuing the **request component login** operational mode command, the QFabric system verifies the class and sends the username and password to the component. The component accepts these credentials and permits access.

NOTE:

- The three QFabric system login classes give access to the components only. To provide access to the QFabric system as a whole through the default partition command-line interface (CLI), you must configure the usual Junos OS login classes or permissions (such as the *super-user* class). For more information about login classes, see *Junos OS Login Classes Overview*.
- If you have completed the QFabric system initial setup and the system is operational, you can change the component password by issuing the **device-authentication** statement at the **[edit system]** hierarchy level in the QFabric default partition CLI.

Topology

This example defines three users: Adam, Oscar, and Ulf. Adam needs to manage QFabric system components, Oscar needs limited access, and Ulf should not have any access to the components. As a result, assign the qfabric-admin class to Adam, the qfabric-operator class to Oscar, and the qfabric-user class to Ulf. However, all three users should have all permissions to access the QFabric system CLI.

Configuration

IN THIS SECTION

- [\[xref target has no title\]](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system login class all-qfabric permissions all
set system login user Adam class all-qfabric
set system login user Adam authentication encrypted-password "$1$aoYSFkvE$G/dYqsTV5iSvVW2sND69U."
set system login user Adam authentication remote-debug-permission qfabric-admin
set system login user Oscar class all-qfabric
set system login user Oscar authentication encrypted-password "$1$3e.3wJQ8$31SrZV0.efdBk.ZJncKm0"
set system login user Oscar authentication remote-debug-permission qfabric-operator
set system login user Ulf class all-qfabric
set system login user Ulf authentication encrypted-password "$1$qt9Ncm0o$okNYSN8O4fVITE/SHBdYj0"
set system login user Ulf authentication remote-debug-permission qfabric-user
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To provide the same access to the QFabric system CLI for all users, but different QFabric system component-level access to different users:

1. Define and provide all-qfabric access and passwords to all three users. This administrator-defined class provides full permissions, enabling the users to log in to the QFabric system default partition and use the CLI. Alternatively, you can assign the super-user class to these users to accomplish the same goal.

[edit]


```

user@qfabric# set system login class all-qfabric permissions all
user@qfabric# set system login user Adam class all-qfabric
user@qfabric# set system login user Adam authentication encrypted-password
"$1$aoYSFkvE$G/dYqsTV5iSvVW2sND69U."
user@qfabric# set system login user Oscar class all-qfabric
user@qfabric# set system login user Oscar authentication encrypted-password
"$1$3e.3wJQ8$31SrZV0.efdBk.ZJncKm0"
user@qfabric# set system login user Ulf class all-qfabric
user@qfabric# set system login user Ulf authentication encrypted-password
"$1$qt9Ncm0o$okNYSN8O4fVITE/SHBdYj0"

```

2. Provide qfabric-admin component access to Adam so he can manage QFabric system components.

```

[edit]
user@qfabric# set system login user Adam authentication remote-debug-permission qfabric-admin

```

3. Provide qfabric-operator component access to Oscar so he can view the CLI at the QFabric system components.

```

[edit]
user@qfabric# set system login user Oscar authentication remote-debug-permission qfabric-operator

```

4. Assign qfabric-user component restrictions to Ulf to prevent him from accessing the QFabric system components.

```

[edit]
user@qfabric# set system login user Ulf authentication remote-debug-permission qfabric-user

```

Results

From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example.

```

[edit]
system {
  login {
    class all-qfabric {
      permissions all;
    }
  }
}

```



```

    }
    user Adam {
        class all-qfabric;
        authentication {
            encrypted-password "$1$aoYSFkvE$G/dYqsTV5iSvVW2sND69U."; ## SECRET-DATA
            remote-debug-permission qfabric-admin;
        }
    }
    user Oscar {
        class all-qfabric;
        authentication {
            encrypted-password "$1$3e.3wJQ8$31SrZV0.efdRbk.ZJncKm0"; ## SECRET-DATA
            remote-debug-permission qfabric-operator;
        }
    }
    user Ulf {
        class all-qfabric;
        authentication {
            encrypted-password "$1$qt9Ncm0o$okNYSN8O4fVITE/SHBdYj0"; ## SECRET-DATA
            remote-debug-permission qfabric-user;
        }
    }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying qfabric-admin Access | 382](#)
- [Verifying qfabric-operator Access | 384](#)
- [Verifying qfabric-user Access | 386](#)

Confirm that the QFabric system and component-level access configuration is working properly for all three users. Adam, Oscar, and Ulf should have equivalent, full-permission access to the QFabric system CLI. Adam should have management-level access to components. Oscar should have read-only access to components. Ulf should have no component-level access.

Verifying qfabric-admin Access

Purpose

Verify that Adam can access the QFabric system CLI at the default partition and manage QFabric system components.

Action

From a management station on your network, issue the **ssh user@qfabric** command and enter the password to open an SSH session for Adam to the QFabric system. Issue the **?** command to view the CLI operational mode commands that Adam has permission to use on the QFabric system default partition.

```
> ssh Adam@qfabric.network.net
Warning: Permanently added 'qfabric.network.net' (RSA) to the list of known hosts.
Adam@qfabric.network.net's password:
Last login: Sun Nov 20 14:12:29 2011 from 192.168.28.19
Juniper QFabric Director 11.3.5510 2011-10-21 16:31:44 UTC

RUNNING ON DIRECTOR DEVICE : dg0
Adam@qfabric>

Adam@qfabric> ?
Possible completions:
  clear          Clear information in the system
  configure      Manipulate software configuration information
  file           Perform file operations
  help           Provide help information
  load           Load information from file
  op             Invoke an operation script
  ping           Ping remote target
  quit           Exit the management session
  request        Make system-level requests
  restart        Restart software process
  save           Save information to file
  set            Set CLI properties, date/time, craft interface message
  show           Show system information
  telnet         Telnet to another host
  test           Perform diagnostic debugging
  traceroute     Trace route to remote host
```


Issue the **request component login ?** command to view the components that Adam can access. Next, issue the **request component login *component-name*** command to log in to a Node device without being prompted for a username or password.

```
Adam@qfabric> request component login ?
Possible completions:
  <[Enter]>          Execute this command
  <node-name>        Inventory name for the remote node
  BBAK0372           Node device
  BBAK0394           Node device
  DRE-0              Diagnostic routing engine
  EE3093             Node device
  FC-0              Fabric control
  FC-1              Fabric control
  FM-0              Fabric manager
  NW-NG-0           Node group
  WS001/RE0          Interconnect device control board
  WS001/RE1          Interconnect device control board
  |                 Pipe through a command

Adam@qfabric> request component login EE3093
Warning: Permanently added 'qfnode-ee3093,169.254.128.14' (RSA) to the list of
known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
```

Finally, issue the **?** command to view the CLI operational mode commands that Adam has the permission to use on the Node device. Notice that the CLI prompt now indicates Adam's component access level (**qfabric-admin**) as the username and the Node device identifier (**EE3093**) as the host.

```
qfabric-admin@EE3093> ?
Possible completions:
  clear              Clear information in the system
  file               Perform file operations
  help               Provide help information
  load               Load information from file
  monitor            Show real-time debugging information
  mtrace             Trace multicast path from source to receiver
  op                 Invoke an operation script
  ping               Ping remote target
  quit               Exit the management session
  request            Make system-level requests
  restart            Restart software process
```


save	Save information to file
set	Set CLI properties, date/time, craft interface message
show	Show system information
ssh	Start secure shell on another host
start	Start shell
telnet	Telnet to another host
test	Perform diagnostic debugging
traceroute	Trace route to remote host

Meaning

The output shows that Adam has received the proper permissions to access the QFabric system CLI and log in to individual components with management-level access.

Verifying qfabric-operator Access

Purpose

Verify that Oscar can access the QFabric system CLI at the default partition and view the CLI on the QFabric system components.

Action

From a management station on your network, issue the **ssh user@qfabric** command and enter the password to open an SSH session for Oscar to the QFabric system. Issue the **?** command to view the CLI operational mode commands that Oscar has permission to use on the QFabric system default partition. Notice that these permissions are the same as those given to Adam.

```
> ssh Oscar@qfabric.network.net
Warning: Permanently added 'qfabric.network.net' (RSA) to the list of known hosts.
Oscar@qfabric.network.net's password:
Last login: Sun Nov 19 19:21:29 2011 from 192.168.28.14
Juniper QFabric Director 11.3.5510 2011-10-22 18:33:41 UTC
```

```
RUNNING ON DIRECTOR DEVICE : dgl
```

```
Oscar@qfabric>
```

```
Oscar@qfabric> ?
```

```
Possible completions:
```

clear	Clear information in the system
configure	Manipulate software configuration information
file	Perform file operations
help	Provide help information
load	Load information from file
op	Invoke an operation script

ping	Ping remote target
quit	Exit the management session
request	Make system-level requests
restart	Restart software process
save	Save information to file
set	Set CLI properties, date/time, craft interface message
show	Show system information
telnet	Telnet to another host
test	Perform diagnostic debugging
traceroute	Trace route to remote host

Issue the **request component login *component-name*** command to log in to a Node device without being prompted for a username or password.

```
Oscar@qfabric> request component login EE3093
Warning: Permanently added 'qfnode-ee3093,169.254.128.14' (RSA) to the list of
known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
```

Finally, issue the ? command to view the CLI operational mode commands that Oscar has permission to use on the Node device. Notice that the CLI prompt now indicates Oscar's component access level (**qfabric-operator**) as the username and the Node device identifier (**EE3093**) as the host. Additionally, Oscar has fewer CLI commands available than Adam because of Oscar's read-only qfabric-operator login class.

```
qfabric-operator@EE3093> ?
Possible completions:
  file          Perform file operations
  help          Provide help information
  load          Load information from file
  op            Invoke an operation script
  quit          Exit the management session
  request       Make system-level requests
  save          Save information to file
  set           Set CLI properties, date/time, craft interface message
  show          Show system information
  start         Start shell
  test          Perform diagnostic debugging
```


Meaning

The output shows that Oscar has full permissions to access the QFabric system CLI, but only read-only access when he logs in to individual components. Oscar's permissions on the QFabric system are the same as Adam's, but Oscar has fewer permissions than Adam on the Node device.

Verifying qfabric-user Access

Purpose

Verify that Ulf has full access to the QFabric system CLI at the default partition but cannot access the QFabric system components.

Action

From a management station on your network, issue the **ssh user@qfabric** command and enter the password to open an SSH session for Ulf to the QFabric system. Issue the ? command to view the CLI operational mode commands that Ulf has permission to use on the QFabric system default partition. Notice that these permissions are the same as those given to Adam and Oscar.

```
> ssh Ulf@qfabric.network.net
Warning: Permanently added 'qfabric.network.net' (RSA) to the list of known hosts.
Ulf@qfabric.network.net's password:
Last login: Sun Nov 17 17:12:24 2011 from 192.168.28.22
Juniper QFabric Director 11.3.5510 2011-10-23 19:23:31 UTC

RUNNING ON DIRECTOR DEVICE : dg0
Ulf@qfabric>

Ulf@qfabric> ?
Possible completions:
  clear          Clear information in the system
  configure      Manipulate software configuration information
  file           Perform file operations
  help           Provide help information
  load           Load information from file
  op             Invoke an operation script
  ping           Ping remote target
  quit           Exit the management session
  request        Make system-level requests
  restart        Restart software process
  save           Save information to file
  set            Set CLI properties, date/time, craft interface message
  show           Show system information
  telnet         Telnet to another host
```


test	Perform diagnostic debugging
traceroute	Trace route to remote host

When Ulf issues the **request component login *component-name*** command, the Node device denies his access attempt.

```
Ulf@qfabric> request component login EE3093
error: User Ulf does not have sufficient permissions to login to device EE3093
```

Meaning

The output shows that Ulf has full permissions to access the QFabric system CLI in the same way as Adam and Oscar. However, unlike Adam and Oscar, Ulf cannot access individual components because of the qfabric-user login class assigned to him.

RELATED DOCUMENTATION

[Understanding QFabric System Login Classes | 75](#)

[Performing the QFabric System Initial Setup on a QFX3100 Director Group | 363](#)

[Junos OS Login Classes Overview](#)

Example: Configuring RADIUS Authentication on a QFabric System

RADIUS authentication is a method of authenticating users who are attempting to access a network device. On a QFabric system, users are load balanced on each of the Director devices. Each Director device needs to be able to communicate with the RADIUS server. Packets sent to the RADIUS server originate from the Director device IP addresses.

The following example shows how to configure RADIUS authentication on the QFabric system:

Perform the following steps to configure RADIUS authentication on the QFabric system:

1. Configure the order in which the authentication methods are used.

For example:

```
user@switch # set system authentication-order [radius password]
```

In this example, RADIUS authentication is the first authentication method that Junos OS will use when a user logs into the system.

2. Configure the IP address of the RADIUS server and the secret password. The secret password on the switch must match the secret password on the RADIUS server.

For example:

```
user@switch # set system radius-server 172.28.36.108 secret testing123
```

3. Assign the login class and the template account for the user.

For example:

```
user@switch # set system login user remote class super-user
```

Here are the results of your configuration:

```
[edit]
system {
  authentication-order [ radius password ];
  login {
    user remote {
      class super-ruser;
    }
  }
}
radius-server {
  172.28.36.108 {
    secret test123
  }
}
```

RELATED DOCUMENTATION

| *Configuring RADIUS Authentication (QFX Series or OCX Series)*

Configuring Aliases for the QFabric System

This topic explains how to configure aliases for components of the QFabric system, such as Director devices, Interconnect devices, and Node devices. Aliases replace the hardware serial numbers of components, making it easier to identify system devices and simplify configuration tasks.

Before you create aliases in a QFabric system:

- Issue one of the **show fabric administration inventory** commands to view the components that are available for aliasing and their hardware serial numbers.

NOTE: The following rules apply to QFabric component alias naming:

- Alias names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.
- The maximum length of an alias name is 30 characters.
- Alias names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.
- You cannot use the reserved names **all**, **fabric**, or **director-group** as an alias name.

To create an alias for a Node device:

1. Discover the serial number of the Node device you wish to rename by issuing the **set fabric aliases node-device ?** context-sensitive help command.

```
root@qfabric# set fabric aliases node-device ?
```

```
Possible completions:
  <aliasable-item-name>  The name of the item to be aliased
  BBAK8309               Node device
  BBAK8283               Node device
  BBAK8891               Node device
  BBAK8868               Node device
  BBAK8276               Node device
  BBAK8273               Node device
[edit]
```

As an alternate way to discover the serial number for a Node device, issue the **show fabric administration inventory node-devices** command. In this case, the serial numbers for the Node devices are **BBAK8309BBAK8283BBAK8891BBAK8868BBAK8276** and **BBAK8273**.

```
root@qfabric> show fabric administration inventory node-devices
```

Item	Identifier	Connection	Configuration
Node device			
BBAK8309		Connected	
BBAK8283		Connected	
BBAK8891		Connected	
BBAK8868		Connected	
BBAK8276		Connected	
BBAK8273		Connected	

2. Specify the serial number of the Node device and the desired alias name by including the **node-device** statement at the **[edit fabric aliases]** hierarchy level.

```
[edit fabric aliases]
root@qfabric# set node-device BBAK8309 Node0
root@qfabric# set node-device BBAK8283 Node1
root@qfabric# set node-device BBAK8891 Node2
root@qfabric# set node-device BBAK8868 Node3
root@qfabric# set node-device BBAK8276 Node4
root@qfabric# set node-device BBAK8273 Node5
```

3. Review your configuration and issue the **commit** command.


```
[edit]
root@qfabric# show fabric
```

```
aliases {
  node-device BBAK8309 {
    Node0;
  }
  node-device BBAK8283 {
    Node1;
  }
  node-device BBAK8891 {
    Node2;
  }
  node-device BBAK8868 {
    Node3;
  }
  node-device BBAK8276 {
    Node4;
  }
  node-device BBAK8273 {
    Node5;
  }
}
```

```
[edit]
root@qfabric# commit
```

```
commit complete
```

4. To view that your aliases are operational, issue the **show fabric administration inventory node-devices** command.

```
root@qfabric> show fabric administration inventory node-devices
```

Item	Identifier	Connection	Configuration
Node device			
node0	BBAK8309	Connected	
node1	BBAK8283	Connected	
node2	BBAK8891	Connected	
node3	BBAK8868	Connected	
node4	BBAK8276	Connected	
node5	BBAK8273	Connected	

NOTE: If you attempt to commit all configuration settings for a new Node group (such as the Node group itself, aliasing, and other features) at the same time, the commit operation might appear to succeed when it actually has failed. For this reason, we recommend configuring and verifying Node groups and aliases first, followed by configuring and verifying other features. Establishing the Node groups and aliases first enables the QFabric system to reject any potentially unsupported configuration. The resulting commit errors indicate where the configuration problem lies. To verify the establishment of Node groups and aliases before configuring other features, issue the **show fabric administration inventory** command.

To create an alias for a Node group:

- Specify a name for the Node group when you include the **node-group** statement at the **[edit fabric resources]** hierarchy level.

NOTE: You cannot use the **aliases** statement at the **[edit fabric]** hierarchy level to create an aliased name for a Node group.

To create an alias for a Director device:

1. Discover the serial number of the Director device you wish to rename by issuing the **set fabric aliases director-device ?** context-sensitive help command.

```
root@qfabric# set fabric aliases director-device ?
```

```
Possible completions:
  <aliasable-item-name>  The name of the item to be aliased
    0281052011000001      Director device
    0281052011000032      Director device
[edit]
```

As an alternate way to discover the serial number for a Director device, issue the **show fabric administration inventory director-group status** command. In this case, the serial number for Director device DG0 is **0281052011000001** and the serial number for Director device DG1 is **0281052011000032**.

```
root@qfabric> show fabric administration inventory director-group status
```

```
Director Group Status Tue Jun  5 15:11:26 UTC 2012
```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	8%	17363152k	4	3 days, 20:55 hrs
dg1	online	backup	10.49.215.39	6%	20157440k	3	3 days, 20:55 hrs

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:4	optimal
-------------	---------

Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size	Used	Avail	Used%	Mounted on
423G	5.4G	395G	2%	/
99M	16M	79M	17%	/boot
93G	7.3G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online	
Partition Manager	online	
Software Mirroring	online	
Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status

Management Interface	up
Control Plane Bridge	up
Control Plane LAG	up
CP Link [0/2]	up
CP Link [0/1]	up
CP Link [0/0]	up
CP Link [1/2]	down
CP Link [1/1]	down
CP Link [1/0]	down
Crossover LAG	up
CP Link [0/3]	up
CP Link [1/3]	up

Member	Device Id/Alias	Status	Role
--------	-----------------	--------	------


```

-----
dgl      0281052011000032 online backup

Director Group Managed Services
-----

Shared File System           online
Network File System          online
Virtual Machine Server        online
Load Balancer/DHCP            online

Hard Drive Status
-----

Volume ID:8                   optimal
Physical ID:1                 online
Physical ID:0                 online
SCSI ID:1                     100%
SCSI ID:0                     100%

Size  Used Avail Used% Mounted on
----  -
423G  5.5G 395G  2%   /
99M   16M  79M   17%  /boot
93G   7.3G 86G   8%   /pbdata

Director Group Processes
-----

Director Group Manager        online
Partition Manager             online
Software Mirroring             online
Shared File System master      online
Secure Shell Process           online
Network File System            online
DHCP Server master             online    backup
FTP Server                     online
Syslog                         online
Distributed Management          online
SNMP Trap Forwarder            online
SNMP Process                   online
Platform Management            online

Interface Link Status
-----

Management Interface          up

```



```

Control Plane Bridge          up
Control Plane LAG             up
CP Link [0/2]                 up
CP Link [0/1]                 up
CP Link [0/0]                 up
CP Link [1/2]                 down
CP Link [1/1]                 down
CP Link [1/0]                 down
Crossover LAG                 up
CP Link [0/3]                 up
CP Link [1/3]                 up

```

2. Specify the serial number of the Director device and the desired alias name by including the **director-device** statement at the **[edit fabric aliases]** hierarchy level.

```

[edit fabric aliases]
root@qfabric# set director-device 0281052011000001 Director0
root@qfabric# set director-device 0281052011000032 Director1

```

3. Review your configuration and issue the **commit** command.

```

[edit]
root@qfabric# show fabric

```

```

aliases {
  director-device 0281052011000001 {
    Director0;
  }
  director-device 0281052011000032 {
    Director1;
  }
}

```

```

[edit]
root@qfabric# commit

```

```

commit complete

```

4. To view that your aliases are operational, issue the **show fabric administration inventory director-group status** command. In this case, the serial numbers in the **Device Id/Alias** field have been replaced with the **Director0** and **Director1** aliased names.

```

root@qfabric> show fabric administration inventory director-group status

```


Director Group Status Tue Jun 5 15:11:26 UTC 2012

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	8%	17363152k	4	3 days, 20:55 hrs
dg1	online	backup	10.49.215.39	6%	20157440k	3	3 days, 20:55 hrs

Member	Device Id/Alias	Status	Role
dg0	Director0	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:4	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size	Used	Avail	Used%	Mounted on
423G	5.4G	395G	2%	/
99M	16M	79M	17%	/boot
93G	7.3G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online
Partition Manager	online
Software Mirroring	online

Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status

Management Interface	up
Control Plane Bridge	up
Control Plane LAG	up
CP Link [0/2]	up
CP Link [0/1]	up
CP Link [0/0]	up
CP Link [1/2]	down
CP Link [1/1]	down
CP Link [1/0]	down
Crossover LAG	up
CP Link [0/3]	up
CP Link [1/3]	up

Member	Device Id/Alias	Status	Role
-----	-----	-----	-----
dgl	Director1	online	backup

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:8	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size	Used	Avail	Used%	Mounted on
423G	5.5G	395G	2%	/
99M	16M	79M	17%	/boot
93G	7.3G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online	
Partition Manager	online	
Software Mirroring	online	
Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	backup
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status

Management Interface	up
Control Plane Bridge	up
Control Plane LAG	up
CP Link [0/2]	up
CP Link [0/1]	up
CP Link [0/0]	up
CP Link [1/2]	down
CP Link [1/1]	down
CP Link [1/0]	down
Crossover LAG	up
CP Link [0/3]	up
CP Link [1/3]	up

To create an alias for an Interconnect device:

1. Discover the serial number of the Interconnect device you wish to rename by issuing the **set fabric aliases interconnect-device ?** context-sensitive help command.

```
root@qfabric# set fabric aliases interconnect-device ?
```

```
Possible completions:
  <aliasable-item-name>  The name of the item to be aliased
  IC-F1249               Interconnect device
  IC-F4912               Interconnect device
[edit]
```

As an alternate way to discover the serial number for an Interconnect device, issue the **show fabric administration inventory interconnect-devices** command. In this case, the serial numbers for the Interconnect devices are **IC-F1249** and **IC-F4912**.

```
root@qfabric> show fabric administration inventory interconnect-devices
```

Item	Identifier	Connection	Configuration
Interconnect device			
IC-F1249		Connected	Configured
F1249/RE0		Connected	
IC-F4912		Connected	Configured
F4912/RE0		Connected	

2. Specify the serial number of the Interconnect device and the desired alias name by including the **interconnect-device** statement at the **[edit fabric aliases]** hierarchy level.

```
[edit fabric aliases]
```

```
root@qfabric# set interconnect-device IC-F1249 Interconnect0
```

```
root@qfabric# set interconnect-device IC-F4912 Interconnect1
```

3. Review your configuration and issue the **commit** command.

```
[edit]
```

```
root@qfabric# show fabric
```

```
aliases {
  interconnect-device IC-F1249 {
    Interconnect0;
  }
  interconnect-device IC-F4912 {
    Interconnect1;
  }
}
```



```
}  
}  
  
[edit]  
root@qfabric# commit  
  
commit complete
```

4. To view that your aliases are operational, issue the **show fabric administration inventory interconnect-devices** command.

```
root@qfabric> show fabric administration inventory interconnect-devices
```

Item	Identifier	Connection	Configuration
Interconnect device			
Interconnect0	IC-F1249	Connected	Configured
F1249/RE0		Connected	
Interconnect1	IC-F4912	Connected	Configured
F4912/RE0		Connected	

RELATED DOCUMENTATION

aliases	 459
show fabric administration inventory	 780
show fabric administration inventory director-group status	 786
show fabric administration inventory interconnect-devices	 797
show fabric administration inventory node-devices	 800
Understanding the Director Group	 26
Understanding Interconnect Devices	 29
Understanding Node Devices	 33

Configuring Node Groups for the QFabric System

This topic explains how to configure Node groups for Node devices within the QFabric system. Node groups provide redundancy for Node devices and make your QFabric system more resilient.

There are three types of Node groups in a QFabric system:

- **Automatically generated server Node groups**—By default, every Node device that joins the QFabric system is placed within an automatically generated server Node group that contains one Node device (the device itself). Server Node groups connect to servers and storage devices.
- **Network Node groups**—You can assign up to eight Node devices to a network Node group. When grouped together, the Node devices within a network Node group connect to other routers running routing protocols such as OSPF and BGP.
- **Redundant server Node groups**—You can assign two Node devices to a redundant server Node group. When grouped together, you can create link aggregation groups (LAGs) that span the interfaces on both Node devices to provide resiliency and redundancy.

Before you create Node groups in a QFabric system:

- Make sure your QFabric system is operational.
- Issue the **show fabric administration inventory node-devices** command to display the Node devices that are available to add to a Node group.
- Issue the **show fabric administration inventory node-groups** command to display the existing Node groups.

NOTE: The following rules apply to QFabric Node group naming:

- Node group names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.
- The maximum length of a Node group name is 30 characters.
- Node group names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.
- You cannot use the reserved names **all**, **fabric**, or **director-group** as a Node group name.

NOTE: If you attempt to commit all configuration settings for a new Node group (such as the Node group itself, aliasing, and other features) at the same time, the commit operation might appear to succeed when it actually has failed. For this reason, we recommend configuring and verifying Node groups and aliases first, followed by configuring and verifying other features. Establishing the Node groups and aliases first enables the QFabric system to reject any potentially unsupported configuration. The resulting commit errors indicate where the configuration problem lies. To verify the establishment of Node groups and aliases before configuring other features, issue the **show fabric administration inventory** command.

To display an automatically generated server Node group:

- Issue the **show fabric administration inventory node-groups** command and look for Node groups containing a single Node device that has the same name or serial number as the server Node group.

```
root@qfabric> show fabric administration inventory node-groups
```

Item	Identifier	Connection	Configuration
Node group			
BBAK8281		Connected	Configured
BBAK8281		Connected	
BBAK8835		Connected	Configured
BBAK8835		Connected	
NW-NG-0		Connected	Configured
Node0	BBAK8309	Connected	
Node1	BBAK8283	Connected	
S1		Connected	Configured
Node2	BBAK8891	Connected	
Node3	BBAK8868	Connected	

To create a network Node group:

1. Specify the Node devices you wish to add to the network Node group by including the **node-device** statement at the **[edit fabric resources node-group NW-NG-0]** hierarchy level.

NOTE:

- The network Node group must use the predefined name NW-NG-0. You must use this name when adding Node devices to the network Node group. You cannot specify a different name. Also, you can configure only one network Node group per partition.
- When you configure routing protocols on the QFabric system, you must use interfaces from the Node devices assigned to the network Node group. If you try to configure routing protocols on interfaces from the Node devices assigned to server Node groups, the configuration commit operation fails.

```
[edit]
```

```
root@qfabric# set fabric resources node-group NW-NG-0 node-device Node0
```

```
root@qfabric# set fabric resources node-group NW-NG-0 node-device Node1
```

2. To designate the Node group as a network Node group, include the **network-domain** statement at the **[edit fabric resources node-group NW-NG-0]** hierarchy level.


```
[edit]
```

```
root@qfabric# set fabric resources node-group NW-NG-0 network-domain
```

3. Review your configuration and issue the **commit** command.

```
[edit]
```

```
root@qfabric# show fabric
```

```
resources {  
  node-group NW-NG-0 {  
    network-domain;  
    node-device Node0;  
    node-device Node1;  
  }  
}
```

```
[edit]
```

```
root@qfabric# commit
```

```
commit complete
```

NOTE: When you add or delete Node devices from a Node group configuration, the corresponding Node devices reboot when you commit the configuration change.

4. To determine if your network Node group is operational, issue the **show fabric administration inventory node-groups** command in operational mode.

```
root@qfabric>show fabric administration inventory node-groups NW-NG-0
```

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
Node0	BBAK8309	Connected	
Node1	BBAK8283	Connected	

To create a redundant server Node group:

1. Specify the two Node devices you wish to add to the redundant server Node group by including the **node-device** statement at the **[edit fabric resources node-group node-group-name]** hierarchy level.

NOTE: Ensure that the two Node devices are of the same type, either two QFX3500 Node devices, two QFX3600 Node devices, or two QFX5100 Node devices. You cannot add different Node device types to the same redundant server Node group.

```
[edit]
```

```
root@qfabric# set fabric resources node-group S1 node-device Node2
```

```
root@qfabric# set fabric resources node-group S1 node-device Node3
```

2. Review your configuration and issue the **commit** command.

```
[edit]
```

```
root@qfabric# show fabric
```

```
resources {
  node-group S1 {
    node-device Node2;
    node-device Node3;
  }
}
```

```
[edit]
```

```
root@qfabric# commit
```

```
commit complete
```

NOTE: When you add or delete Node devices from a Node group configuration, the corresponding Node devices reboot when you commit the configuration change.

3. To determine if your redundant server Node groups are operational, issue the **show fabric administration inventory node-groups *redundant-server-node-group-name*** command in operational mode.

```
root@qfabric> show fabric administration inventory node-groups S1
```

Item	Identifier	Connection	Configuration
Node group			
S1		Connected	Configured
Node2	BBAK8891	Connected	
Node3	BBAK8868	Connected	

RELATED DOCUMENTATION

show fabric administration inventory node-groups 802
show fabric administration inventory node-devices 800
Understanding Node Groups 38
node-group (Resources) 512

Configuring the Port Type on QFX3600 Node Devices

The QFX3600 Node device provides 16 40-Gbps QSFP+ ports. By default, four ports (labeled **Q0** through **Q3**) operate as 40-gigabit data plane (*fte*) uplink ports for uplink connections between your Node device and your Interconnect devices. Twelve ports (labeled **Q4** through **Q15**) operate as 10-Gigabit Ethernet (*xe*) ports to support 48 10-Gigabit Ethernet interfaces for connections to either endpoint systems or external networks. Optionally, you can choose to configure ports Q0 through Q7 to operate as 40-gigabit data plane uplink ports, and ports Q2 through Q15 to operate as 10-Gigabit Ethernet or 40-Gigabit Ethernet (*xle*) ports.

NOTE: You can use QSFP+ to four SFP+ breakout cables or QSFP+ transceivers with fiber breakout cables to connect the 10-Gigabit Ethernet ports to other devices.

NOTE: When you delete the port type configuration for an individual port or a block of ports, the ports return to operating in their default port type. For example, when you delete the 40-Gigabit Ethernet (*xle*) port configuration for port Q4, the port returns to operating as a 10-Gigabit Ethernet (*xe*) port.

NOTE: When the 40-Gigabit Ethernet (*xle*) ports of a QFX3600 Node device carry traffic at the full line rate, loss of untagged Layer 2 or Layer 3 traffic going across the fabric might occur, as well as increased latency on the Node device. Such effects result from the addition of a 4-byte header to packets traversing the uplink ports on the Node device. The percentage of traffic loss depends on the size of the packets: the greater the packet size, the lower the traffic loss and vice versa. This problem does not affect tagged traffic.

This topic explains how to configure the port type on QFX3600 Node devices.

Before you configure the port type on QFX3600 Node devices:

- Make sure your QFabric system is operational.
- Issue the **show fabric administration inventory node-groups** command to display the existing Node groups and the Node devices in each Node group.

NOTE:

- Only ports Q0 through Q7 can be configured to operate as 40-gigabit data plane (fte) uplink ports.
- Only ports Q2 through Q15 can be configured to operate as 10-Gigabit Ethernet (xe) or 40-Gigabit Ethernet (xle) ports.



CAUTION: The Packet Forwarding Engine on the QFX3600 Node device is restarted when you commit the port type configuration changes. As a result, you might experience packet loss on the Node device.

The following message may be displayed in the system log file when the Packet Forwarding Engine is restarted. You can ignore this message.

Pipe write error: Broken pipe

flush operation failed

The following steps describe how to configure either a block of ports or an individual port to operate as 40-gigabit data plane uplink (fte) ports, as well as how to delete a 40-gigabit data plane uplink (fte) port configuration.

1. To configure a block of ports to operate as 40-gigabit data plane uplink (fte) ports, specify a port range:

```
[edit chassis node-group name node-device name pic 1]
root@qfabric# set fte port-range port-range-low port-range-high
```

For example, to configure ports Q4 through Q7 to operate as 40-gigabit data plane uplink ports:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 1]
root@qfabric# set fte port-range 4 7
```

2. To configure an individual port to operate as a 40-gigabit data plane uplink (fte) port, specify a port number:

```
[edit chassis node-group name node-device name pic 1]
root@qfabric# set fte port port-number
```


For example, to configure port Q4 to operate as a 40-gigabit data plane uplink port:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 1]
root@qfabric# set fte port 4
```

3. Review your configuration and issue the **commit** command.

```
[edit]
root@qfabric# commit
commit complete
```

4. To delete the 40-gigabit data plane uplink (fte) port configuration for a block of ports, specify a port range:

```
[edit chassis node-group name node-device name pic 1]
root@qfabric# delete fte port-range port-range-low port-range-high
```

For example, to delete the 40-gigabit data plane uplink port configuration for ports Q4 through Q7:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 1]
root@qfabric# delete fte port-range 4 7
```

5. To delete the 40-gigabit data plane uplink (fte) port configuration for an individual port, specify a port number:

```
[edit chassis node-group name node-device name pic 1]
root@qfabric# delete fte port port-number
```

For example, to delete the 40-gigabit data plane uplink port configuration for port Q4:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 1]
root@qfabric# delete fte port 4
```

The following steps describe how to configure either a block of ports or an individual port to operate as 10-Gigabit Ethernet (xe) ports, as well as how to delete a 10-Gigabit Ethernet (xe) port configuration.

1. To configure a block of ports to operate as 10-Gigabit Ethernet (xe) ports, specify a port range:

```
[edit chassis node-group name node-device name pic 0]
```



```
root@qfabric# set xe port-range port-range-low port-range-high
```

For example, to configure ports Q4 through Q7 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 0]
root@qfabric# set xe port-range 4 7
```

2. To configure an individual port to operate as a 10-Gigabit Ethernet port, specify a port number:

```
[edit chassis node-group name node-device name pic 0]
root@qfabric# set xe port port-number
```

For example, to configure port Q4 to operate as a 10-Gigabit Ethernet port:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 0]
root@qfabric# set xe port 4
```

3. Review your configuration and issue the **commit** command.

```
[edit]
root@qfabric# commit
commit complete
```

4. To delete the 10-Gigabit Ethernet (xe) port configuration for a block of ports, specify a port range:

```
[edit chassis node-group name node-device name pic 0]
root@qfabric# delete xe port-range port-range-low port-range-high
```

For example, to delete the 10-Gigabit Ethernet port configuration for ports Q4 through Q7:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 0]
root@qfabric# delete xe port-range 4 7
```

5. To delete the 10-Gigabit Ethernet (xe) port configuration for an individual port, specify a port number:

```
[edit chassis node-group name node-device name pic 0]
root@qfabric# delete xe port port-number
```


For example, to delete the 10-Gigabit Ethernet port configuration for port Q4:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 0]
root@qfabric# delete xe port 4
```

The following steps describe how to configure either a block of ports or an individual port to operate as 40-Gigabit Ethernet (xle) ports, as well as how to delete a 40-Gigabit Ethernet (xle) port configuration.

1. To configure a block of ports to operate as 40-Gigabit Ethernet (xle) ports, specify a port range:

```
[edit chassis node-group name node-device name pic 1]
root@qfabric# set xle port-range port-range-low port-range-high
```

For example, to configure ports Q4 through Q7 to operate as 40-Gigabit Ethernet ports:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 1]
root@qfabric# set xle port-range 4 7
```

2. To configure an individual port to operate as a 40-Gigabit Ethernet (xle) port, specify a port number:

```
[edit chassis node-group name node-device name pic 1]
root@qfabric# set xle port port-number
```

For example, to configure port Q4 to operate as a 40-Gigabit Ethernet port:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 1]
root@qfabric# set xle port 4
```

3. Review your configuration and issue the **commit** command.

```
[edit]
root@qfabric# commit
commit complete
```

4. To delete the 40-Gigabit Ethernet (xle) port configuration for block of ports, specify a port range:

```
[edit chassis node-group name node-device name pic 1]
root@qfabric# delete xle port-range port-range-low port-range-high
```


For example, to delete the 40-Gigabit Ethernet port configuration for ports Q4 through Q7:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 1]  
root@qfabric# delete xle port-range 4 7
```

5. To delete the 40-Gigabit Ethernet (xle) port configuration for an individual port, specify a port number:

```
[edit chassis node-group name node-device name pic 1]  
root@qfabric# delete xle port port-number
```

For example, to delete the 40-Gigabit Ethernet port configuration for port Q4:

```
[edit chassis node-group BBAK8281 node-device BBAK8309 pic 1]  
root@qfabric# delete xle port 4
```

RELATED DOCUMENTATION

[Understanding Node Devices | 33](#)

[Understanding Interfaces on the QFabric System | 15](#)

[pic | 514](#)

Configuring the QSFP+ Port Type on QFX5100 Devices

You can convert default 40-Gigabit Ethernet data plane uplink interfaces (fte) to 40-Gigabit Ethernet access interfaces (xle) ports, and default 40-Gigabit Ethernet interfaces (xle) to 40-Gigabit Ethernet data plane uplink interfaces (fte). Ports Q0 and Q1 are fixed fte ports and cannot be changed. Ports Q2 and Q3 are fte ports by default but can be changed to xle ports. Ports Q4 and Q5 are xle ports by default but can be changed to fte ports.

NOTE: On QFX5100-24Q switches, ports Q1 through Q7 are fixed FTE ports and cannot be changed.

NOTE: You must configure xle ports in pairs, not individually, otherwise functionality is not guaranteed.



CAUTION: The Packet Forwarding Engine on a QFX5100 switch is restarted when you commit port type configuration changes (for example, configuring or deleting an fte or xle port). As a result, you might experience packet loss on the device.

The following steps describe how to configure either a block of ports or an individual port, as well as how to delete these configurations.

1. To configure a block of ports to operate as 40-Gigabit Ethernet interfaces (xle) , specify a port range:

```
[edit chassis (QFX Series) node-group name node-device name pic 1]
user@switch# set xle port-range port-range-low port-range-high
```

For example, to configure ports Q4 through Q5 to operate as 40-Gigabit Ethernet interfaces (xle):

```
[edit chassis node-group name node-device name pic 1]
user@switch# set xle port-range 4 5
```

2. To configure a block of ports to operate as 40-Gigabit Ethernet data plane uplink interfaces (fte), specify a port range:

```
[edit chassis (QFX Series) node-group name node-device name pic 1]
```



```
user@switch# set fte port-range port-range-low port-range-high
```

For example, to configure ports Q4 through Q5 to operate as 40-Gigabit Ethernet data plane uplink interfaces (fte):

```
[edit chassis node-group name node-device name pic 1]
user@switch# set fte port-range 4 5
```

3. To configure an individual port to operate as a 40-Gigabit Ethernet data plane uplink interfaces (fte), specify a port number:

```
[edit chassis node-group name node-device name pic 1]
user@switch# set fte port port-number
```

For example, to configure port Q4 to operate as a 40-Gigabit Ethernet data plane uplink interfaces (fte):

```
[edit chassis node-group name node-device name pic 1]
user@switch# set fte port 4
```

4. Review your configuration and issue the **commit** command.

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

5. To delete a block of ports configured as 40-Gigabit Ethernet (xle) ports, specify a port range:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete xle port-range port-range-low port-range-high
```

For example, to delete the 40-Gigabit Ethernet access interface (xle) port configuration for ports Q2 through Q3:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete xle port-range 2 3
```

6. To delete an individual port configured as a 40-Gigabit Ethernet (xle) interface:


```
[edit chassis node-group name node-device name pic 1]
user@switch# delete xle port port-number
```

For example, to delete the 40-Gigabit Ethernet interface (xle) for port Q2:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete xle port 2
```

7. To delete a block of ports configured as 40-Gigabit Ethernet data plane uplink interfaces (fte), specify a port range:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete fte port-range port-range-low port-range-high
```

For example, to delete the block of ports configured as 40-Gigabit Ethernet data plane uplink interfaces (fte) for ports Q4 through Q5:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete fte port-range 4 5
```

8. To delete an individual port configured as a 40-Gigabit Ethernet data plane uplink interfaces (fte):

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete fte port port-number
```

For example, to delete the 40-Gigabit Ethernet data plane uplink interfaces (fte) for port Q4:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete fte port 4
```

9. Review your configuration and issue the **commit** command.

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

RELATED DOCUMENTATION

| *pic*

Example: Configuring SNMP

IN THIS SECTION

- [Requirements | 417](#)
- [Overview | 417](#)
- [Configuration | 417](#)

By default, SNMP is disabled on devices running Junos OS. This example describes the steps for configuring SNMP on the QFabric system.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2
- Network management system (NMS) (running the SNMP manager)
- QFabric system (running the SNMP agent) with multiple Node devices

Overview

Because SNMP is disabled by default on devices running Junos OS, you must enable SNMP on your device by including configuration statements at the **[edit snmp]** hierarchy level. At a minimum, you must configure the **community public** statement. The community defined as public grants read-only access to MIB data to any client.

If no **clients** statement is configured, all clients are allowed. We recommend that you always include the **restrict** option to limit SNMP client access to the switch.

The network topology in this example includes an NMS, a QFabric system with four Node devices, and external SNMP servers that are configured for receiving traps.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set snmp name "snmp qfabric" description "qfabric0 switch"
set snmp location "Lab 4 Row 11" contact "qfabric-admin@qfabric0"
set snmp community public authorization read-only
set snmp client-list list0 192.168.0.0/24
set snmp community public client-list-name list0
set snmp community public clients 192.170.0.0/24 restrict
set snmp trap-group "qf-traps" destination-port 155 targets 192.168.0.100
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *Junos OS CLI User Guide*.

To configure SNMP on the QFabric system:

NOTE: If the name, description, location, contact, or community name contains spaces, enclose the text in quotation marks (" ").

1. Configure the SNMP system name:

```
[edit snmp]
user@switch# set name "snmp qfabric"
```

NOTE: The above configured SNMP system name can be accessed:

- By doing a query with the SNMPGet on policy object identifier (OID) sysName.0.
- From the generic jnxSyslogTrap. To send the jnxSyslogTrap, configure the trap events at **[edit event-options policy]** hierarchy.

2. Specify a description.

```
[edit snmp]
user@switch# set description "qfabric0 system"
```


This string is placed into the MIB II sysDescription object.

3. Specify the physical location of the QFabric system.

```
[edit snmp]
user@switch# set location "Lab 4 Row 11"
```

This string is placed into the MIB II sysLocation object.

4. Specify an administrative contact for the SNMP system.

```
[edit snmp]
user@switch# set contact "qfabric-admin@qfabric0"
```

This name is placed into the MIB II sysContact object.

5. Specify a unique SNMP community name and the read-only authorization level.

NOTE: The **read-write** option is not supported on the QFabric system.

```
[edit snmp]
user@switch# set community public authorization read-only
```

6. Create a client list with a set of IP addresses that can use the SNMP community.

```
[edit snmp]
user@switch# set client-list list0 192.168.0.0/24
user@switch# set community public client-list-name list0
```

7. Specify IP addresses of clients that are restricted from using the community.

```
[edit snmp]
user@switch# set community public clients 198.51.100.0/24 restrict
```

8. Configure a trap group, destination port, and a target to receive the SNMP traps in the trap group.

```
[edit snmp]
```



```
user@switch# set trap-group "qf-traps" destination-port 155 targets 192.168.0.100
```

NOTE: You do not need to include the **destination-port** statement if you use the default port 162.

The trap group qf-traps is configured to send traps to 192.168.0.100.

Results

From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@switch# show
snmp {
  name "snmp qfabric";
  description "qfabric0 system";
  location "Lab 4 Row 11";
  contact "qfabric-admin@qfabric0";
  client-list list0 {
    192.168.0.0/24;
  }
  community public {
    authorization read-only;
    clients {
      198.51.100.0/24 restrict;
    }
  }
  trap-group qf-traps {
    destination-port 155;
    targets {
      192.168.0.100;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

RELATED DOCUMENTATION

[Understanding the Implementation of SNMP on the QFabric System](#) | 69

Example: Configuring System Log Messages

IN THIS SECTION

- [Requirements | 421](#)
- [Overview | 421](#)
- [Configuration | 422](#)

The QFabric system monitors events that occur on its component devices and distributes system log messages about those events to all external system log message servers (hosts) that are configured. Component devices may include Node devices, Interconnect devices, Director devices, and the Virtual Chassis. Messages are stored for viewing only in the QFabric system database. To view the messages, issue the **show log** command.

This example describes how to configure system log messages on the QFabric system.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2
- QFabric system
- External servers that can be configured as system log message hosts

Overview

Component devices that generate system log message events may include Node devices, Interconnect devices, Director devices, and the control plane switches. The following configuration example includes these components in the QFabric system:

- Director software running on the Director group
- Control plane switches

- Interconnect device
- Multiple Node devices

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system syslog host 10.1.1.12 any error
set system syslog file qflogs
set system syslog file qflogs structured-data brief
set system syslog file qflogs archive size 1g
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *Junos OS CLI User Guide*.

To configure system messages from the QFabric Director device:

1. Specify a host, any facility, and the **error** severity level.

```
[edit system syslog]
user@switch# set host 10.1.1.12 any error
```

NOTE: You can configure more than one system log message server (host). The QFabric system sends the messages to each server configured.

2. (Optional) Specify a filename to capture log messages.

NOTE: On the QFabric system, a syslog file named **messages** is configured implicitly with facility and severity levels of **any any** and a file size of 100 MBs. Therefore, you cannot specify the filename **messages** in your configuration, and automatic command completion does not work for that filename.


```
[edit system syslog]
user@switch# set file qflogs structured-data brief
user@switch# set file qflogs
```

3. (Optional) Configure the maximum size of your system log message archive file. This example specifies an archive size of 1 GB.

```
[edit system syslog]
user@switch# set file qflogs archive size 1g
```

Results

From configuration mode, confirm your configuration by entering the **show system** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@switch# show system
syslog {
  file qflogs {
  }
  host 10.1.1.12 {
    any error;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

RELATED DOCUMENTATION

[Understanding the Implementation of System Log Messages on the QFabric System | 72](#)

[syslog \(QFabric System\) | 520](#)

[show log | 813](#)

Configuring Graceful Restart for QFabric Systems

IN THIS SECTION

- [Enabling Graceful Restart | 424](#)
- [Configuring Graceful Restart Options for BGP | 425](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 | 426](#)
- [Tracking Graceful Restart Events | 428](#)

When you configure graceful restart in the QFabric CLI, the QFabric system applies the configuration to the network Node group to participate in graceful restart operations with devices external to the QFabric system. Such configuration preserves routing table state and helps neighboring routing devices to resume routing operations more quickly after a system restart. This also enables the network Node group to resume routing operations rapidly if there is a restart in the QFabric system (such as a software upgrade). As a result, we recommend enabling graceful restart for routing protocols in the QFabric CLI.

NOTE: The QFabric system also uses graceful restart internally within the fabric to facilitate interfabric resiliency and recovery. This internal feature is enabled by default with no configuration required.

Enabling Graceful Restart

By default, graceful restart is disabled. To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level.

For example:

```
routing-options {  
  graceful-restart;  
}
```

To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level.

NOTE: Helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode on a per-protocol basis.

```
[edit]
routing-options {
  graceful-restart {
    disable;
    restart-duration seconds;
  }
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

When graceful restart is enabled for all routing protocols at the **[edit routing-options graceful-restart]** hierarchy level, you can disable graceful restart on a per-protocol basis.

NOTE: If you configure graceful restart after a BGP or LDP session has been established, the BGP or LDP session restarts and the peers negotiate graceful restart capabilities. Also, the BGP peer routing statistics are reset to zero.

Configuring Graceful Restart Options for BGP

To configure the duration of the BGP graceful restart period, include the **restart-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the **stale-routes-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.

```
[edit]
protocols {
  bgp {
    graceful-restart {
      disable;
      restart-time seconds;
      stale-routes-time seconds;
    }
  }
}
```



```

    }
}
routing-options {
    graceful-restart;
}

```

To disable BGP graceful restart capability for all BGP sessions, include the **disable** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.

NOTE: To set BGP graceful restart properties or disable them for a group, include the desired statements at the **[edit protocols bgp group group-name graceful-restart]** hierarchy level.

To set BGP graceful restart properties or disable them for a specific neighbor in a group, include the desired statements at the **[edit protocols bgp group group-name neighbor ip-address graceful-restart]** hierarchy level.

NOTE: Configuring graceful restart for BGP resets the BGP peer routing statistics to zero. Also, existing BGP sessions restart, and the peers negotiate graceful restart capabilities.

Configuring Graceful Restart Options for OSPF and OSPFv3

To configure the duration of the OSPF/OSPFv3 graceful restart period, include the **restart-duration** statement at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level. To specify the length of time for which the router notifies helper routers that it has completed graceful restart, include the **notify-duration** at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level. Strict OSPF link-state advertisement (LSA) checking results in the termination of graceful restart by a helping router. To disable strict LSA checking, include the **no-strict-lsa-checking** statement at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level.

```

[edit]
protocols {
    ospf | ospfv3 {
        graceful-restart {
            disable;
            helper-disable
            no-strict-lsa-checking;
            notify-duration seconds;
            restart-duration seconds;
        }
    }
}

```



```

    }
  }
}
routing-options {
  graceful-restart;
}

```

To disable OSPF/OSPFv3 graceful restart, include the **disable** statement at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level.

Starting with Release 11.3, the Junos OS supports both the standard (based on RFC 3623, *Graceful OSPF Restart*) and the restart signaling-based (as specified in RFC 4811, RFC 4812, and RFC 4813) helper modes for OSPF version 2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default. To disable the helper mode for OSPF version 2 graceful restart configurations, include the **helper-disable <both | restart-signaling | standard>** statement at the **[edit protocols ospf graceful-restart]** hierarchy level. Note that the last committed statement always takes precedence over the previous one.

```

[edit protocols ospf]
  graceful-restart {
    helper-disable <both | restart-signaling | standard>
  }

```

To reenabling the helper mode, delete the **helper-disable** statement from the configuration by using the **delete protocols ospf graceful-restart helper-disable <restart-signaling | standard | both>** command. In this case also, the last executed command takes precedence over the previous ones.

NOTE:

Restart signaling-based helper mode is not supported for OSPFv3 configurations. To disable helper mode for OSPFv3 configurations, include the **helper-disable** statement at the **[edit protocols ospfv3 graceful-restart]** hierarchy level.

TIP: You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols (ospf | ospfv3)]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 428](#).

NOTE: If you configure BFD and graceful restart for OSPF, graceful restart might not work as expected.

Tracking Graceful Restart Events

To track the progress of a graceful restart event, you can configure graceful restart trace options flags for IS-IS and OSPF/OSPFv3. To configure graceful restart trace options, include the **graceful-restart** statement at the **[edit protocols *protocol* traceoptions flag]** hierarchy level:

```
[edit protocols]
isis {
  traceoptions {
    flag graceful-restart;
  }
}
(ospf | ospf3) {
  traceoptions {
    flag graceful-restart;
  }
}
```

RELATED DOCUMENTATION

Graceful Restart Concepts

Verifying Graceful Restart Operation

Optimizing the Number of Multicast Flows on QFabric Systems

Because of the distributed nature of QFabric systems, the default configuration does not allow the maximum number of supported Layer 3 multicast flows to be created. To allow a QFabric system to create the maximum number of supported flows, configure the following statement:

```
set fabric routing-options multicast fabric-optimized-distribution
```

After configuring this statement, you must reboot the QFabric Director group to make the change take effect.

RELATED DOCUMENTATION

Segregating QFabric Traffic Flows With Flow Groups

In a QFabric system, a hash function is used to select an Interconnect device to forward traffic between two ingress switches, which are called Node devices. Since this hash function is performed on all Interconnect devices, it is possible for redundant multicast streams to flow through one Interconnect device, making that Interconnect device a potential single point of failure for the redundant flows. Some operations, such as financial transactions, require redundant data flows to use different Interconnect devices to avoid point-of-failure data loss. To enforce this in a fabric, you can create a flow group for multicast or unicast traffic from a particular ingress switch. In a flow group, switches are forced to use different, specified Interconnect devices to forward redundant traffic.

NOTE: An Interconnect device can belong to only one flow group but can serve as a backup to other flow groups.

To direct redundant QFabric system traffic flows to different Interconnect devices:

1. Create a flow group and indicate the Node device where the redundant data originates:

```
[edit fabric flow-groups]
set flow-group-name node-device node-device-name
```

2. Indicate at least one Interconnect device connected to the switch to be used for redundancy:

```
[edit fabric flow-groups flow-group-name]
set interconnect-device interconnect-device-name
```

If all redundant data from the named switch will use the flow group, you are finished with the configuration.

3. You can alter a flow group's configuration by overriding the software's Node device preferences for Interconnect devices. Interconnect devices that reside in the default group (those that are not assigned to a user-configured flow group) have a link preference value set to **normal** in the software. When you create a flow group, the **flow-groups** configuration statement resets some Interconnect devices' preference to **high** (for those assigned to a flow group) and others to **never** (for those assigned to a different flow group) to create the flow.

- To add a Node device to multiple flow groups, set the link between the Node device in the original flow group and a different Interconnect device to **high**:

```
[edit fabric flow-groups flow-group-name]
set node-device node-device-name interconnect-device new-IC-name preference
high
```

- To configure a Node device to use an Interconnect device from a different flow group if none in the same flow group are available, set the link between the Node device and the other group's Interconnect device to **normal**:

```
[edit fabric flow-groups flow-group-name]
set node-device node-device-name interconnect-device IC-name-from-other-flow-group
preference normal
```

- To increase the priority of a switch for Interconnect devices, set the switch's Interconnect device link preference to **high**.

```
[edit fabric flow-groups flow-group-name]
set node-device node-device-name interconnect-device interconnect-device-name
preference high
```

- To prevent a flow group's switches from using an Interconnect device, set the Interconnect device link preference to **never**.

```
[edit fabric flow-groups flow-group-name]
set node-device node-device-name interconnect-device interconnect-device-name
preference never
```

For an example of flow group configuration, see [“Example: Creating a QFabric Flow Group”](#) on page 431.

TIP: To add a TOR to multiple flow groups:

set fabric resources node-device TOR-B interconnect-device IC-A preference HIGH

To configure a TOR to use an Interconnect device from a different flow group if none in the same flow group are available:

set fabric resources node-device TOR-B interconnect-device IC-A preference NORMAL

To increase the priority of default-flowgroup Interconnect devices for a specific TOR:

set fabric resources node-device TOR-A interconnect-device ALL_ICS preference HIGH

To configure default-flowgroup TORs to never use a particular Interconnect device:

set fabric resources node-device ALL_TORS interconnect-device IC-A preference NEVER

RELATED DOCUMENTATION

[flow-groups | 479](#)[Understanding QFabric Multicast Flow Groups | 55](#)[Example: Creating a QFabric Flow Group | 431](#)

Example: Creating a QFabric Flow Group

IN THIS SECTION

- [Requirements | 431](#)
- [Overview | 432](#)
- [Configuration | 433](#)
- [Verification | 437](#)

In a QFabric system, a hash function is used to select an Interconnect device to forward traffic between two Node devices. Since this hash function is performed on all Interconnect devices, it is possible for redundant multicast streams to flow through one Interconnect device, making that Interconnect device a potential single point of failure for the redundant flows. Some applications require that the redundant multicast streams flow through different Interconnect devices to prevent a single Interconnect device from potentially dropping both streams of multicast traffic during a failure. You can enforce this use of separate Interconnect devices by using the QFabric system flow groups feature. This example creates two flow groups.

TIP: A single point of failure can occur because the hash function can pick the same Interconnect device for two different sets of Node devices. For example, a set of Node devices (N1, N2) can be hashed to Interconnect device IC1 and another set of Node devices (N3, N4) are also hashed to IC1, so IC1 can then be a potential single point of failure.

Requirements

This example uses the following hardware and software components:

- QFX5100-F switches

- QFabric system
- Junos 14.1X53-D15

Before you create a flow group on your QFX Series switches, be sure you have installed and set up a QFabric system. See *QFX3000-G QFabric System Installation Overview*.

Overview

In this example, you configure two **flow-groups** flows on QFX5100-F switches in a QFabric system. Each flow will use its configured Interconnect device. You can optionally configure a backup option for one of the flows to use the Interconnect device in the other flow in case of failure.

Topology

This example creates two flow groups, flowA and flowB, in a QFabric system.

You always have a default flow group that contains any unassigned Interconnect devices and unassigned Node devices—in this case, the only member of the default flow group will be the Interconnect device IC3 after your configurations. The flow groups, Node devices (TORs—Top-of-Rack devices), and Interconnect devices that you configure are listed in [Table 77 on page 432](#).

Table 77: Flow-group Example Devices

Flow Group Name	Node Devices	Interconnect Device Names
flowA	P2025-C	IC-4892
	P1360-C	IC-4
flowB	BBAK8822	IC-F4908
	P1523-C	
default flow-group		IC3

The software enforces flows by setting each Node device to one of three preferences for Interconnect devices:

- **high**—Use this Interconnect device in the hashing algorithm for the specified flow.
- **normal**—Use this Interconnect device for backup when no Interconnect devices with high settings are available.
- **never**—Do not use this Interconnect device in the hashing algorithm for the specified flow.

NOTE: Unassigned Interconnect devices in the default group are set to **normal**, meaning when no Interconnect device set to **high** is available in a flow group, the flow will look for unused Interconnect devices set to **normal** and use one that is available. This helps prevent loss of data in the event of an Interconnect device failure. The **high** and **never** settings for Interconnect devices are set by the software when you configure flow groups.

Configuration

IN THIS SECTION

- [Configure Flow Group flowA | 434](#)
- [Configure Flow Group flowB | 434](#)
- [Configure a Backup Interconnect-Device For FlowA \(Optional\) | 435](#)
- [Results | 437](#)

In this example, you configure two flow groups: flowA and flowB. FlowA has two Interconnect devices set to **high**. If both Interconnect devices in flowA fail, there is one automatic backup because the default group contains an Interconnect device. You can also, optionally, provide a specific backup Interconnect device for flowA—in this example, we specifically configure the Interconnect device for flowB as backup for flowA. Therefore, if both Interconnect devices in flowA (IC-4892 and IC-4) fail, an Interconnect device with preference **normal** will take over. IC3 is available at this point.

CLI Quick Configuration

To quickly configure flow groups flowA and flowB with backup for the Interconnect device in flowA, copy and paste these commands into your CLI:

```
set fabric flow-groups flowA node-device P2025-C
set fabric flow-groups flowA node-device P1360-C
set fabric flow-groups flowA interconnect-device IC-4892
set fabric flow-groups flowA interconnect-device IC4

set fabric flow-groups flowB node-device BBAK8822
set fabric flow-groups flowB node-device P1523-C
set fabric flow-groups flowB interconnect-device IC-F4908
```



```
set fabric flow-groups node-device P2025-C interconnect-device IC-4908 preference
normal
```

Configure Flow Group flowA

Step-by-Step Procedure

1. Set up two Node devices for flowA:

```
set fabric flow-groups flowA node-device P2025-C
set fabric flow-groups flowA node-device P1360-C
```

2. Set up two Interconnect devices for flowA:

```
set fabric flow-groups flowA interconnect-device IC-4892
set fabric flow-groups flowA interconnect-device IC4
```

Results

At this point, the software's Node devices preference settings for flow A are P2025-C and P1360-C, with the following Interconnect device preferences for flowA:

- IC-4892: **high** (assigned to flowA)
- IC4: **high** (assigned to flowA)
- IC-F4908: **normal** (not assigned to any specific flow group and is available for backup)
- IC3: **normal** (not assigned to any specific flow group and available for backup)

NOTE: Neither IC-F4908 or IC3 have been assigned yet, so they are both set to **normal** in flowA and are available for automatic backup. That will change with the next configuration—when flowB is configured, the IC-F4908 preference for flowA will automatically change.

Configure Flow Group flowB

Step-by-Step Procedure

1. Set up two Node devices for flowB:

```
set fabric flow-groups flowB node-device BBAK8822
set fabric flow-groups flowB node-device P1523-C
```

2. Set up one Interconnect device for flowB:

```
set fabric flow-groups flowB interconnect-device IC-F4908
```


Results

At this point, the software's Node device preference settings for flowB are BBAK8822 and P1523-C, with the following Interconnect device preferences for flowB:

- IC-F4908: **high** (assigned to flowB)
- IC4: **never** (assigned to another flow group, flowA)
- IC-4892: **never** (assigned to another flow group, flowA)
- IC3: **normal** (not assigned to any specific flow group and available for backup)

Also, now the software's Interconnect device preference settings for flowA have changed automatically as a result of assigning Interconnect device IC-F4908 to another flow group, flowB. At this point, the Interconnect device preference settings for flowA are as follows:

- IC-4892: **high** (assigned to flowA)
- IC4: **high** (assigned to flowA)
- IC-F4908: **never** (assigned to another flow group, flowB)
- IC3: **normal** (not assigned to any specific flow group and is available for backup)

TIP: After an Interconnect device is explicitly assigned to a flow group, all other flow groups mark its preference for their flow as **never**.

Configure a Backup Interconnect-Device For FlowA (Optional)

Step-by-Step Procedure

The software sets Interconnect device preferences automatically—however, you can override this setting using the **preference** option of the **flow-groups** CLI statement. You might want to do this for backup reasons.

In this example, the TOR switch P2025-C will always use either Interconnect device IC-F4908 or IC4 because these Interconnect devices are configured for flow group flowA. After adding the optional preference override setting described here, if both Interconnect devices IC4 and IC-F4908 subsequently fail, the TOR switches in flowA will start using an available Interconnect device that is set to preference **normal**, instead of dropping traffic. This step overrides the TOR switch P1523-C preference for the Interconnect device IC-4892, changing it from **never** to **normal**, resulting in two Interconnect devices with preference setting **normal**—IC-4892 in flowB and IC3 in the default group.

1. Override the TOR switch P2025-C preference for the Interconnect device IC-4908, changing it from **never** to **normal**:

```
set fabric flow-groups node-device P2025-C interconnect-device IC-4908
preference normal
```

Results

At this point, the Interconnect device preference settings for the Node device P2025-C are as follows:

- IC-4892: **high** (assigned to flowA)
- IC4: **high** (assigned to flowA)
- IC-F4908: **normal** (by explicit preference configuration for Node device P2025-C)
- IC3: **normal** (not assigned to any specific flow group and available for backup)

The preference settings for the Node devices in flowB are the same as they were before you configured the Interconnect device preference override for the Node device P2025-C in flowA:

- IC-F4908: **high** (assigned to flowB)
- IC4: **never** (assigned to flowA)
- IC-4892: **never** (assigned to flowA)
- IC3: **normal** not assigned to any specific flow group and available for backup)

TIP: To see more reasons you might want to alter the software's preference settings, see [“Segregating QFabric Traffic Flows With Flow Groups” on page 429](#).

Results

Type **show fabric flow-groups** in configuration mode to see your configuration.

```
root@qfabric# show fabric flow-groups
```

```
flowA {
  node-device {
    P2025-C;
    P1360-C;
  }
  interconnect-device {
    IC-F4892;
    IC-4;
  }
}
flowB {
  node-device {
    P1523-C;
    BBAK8822;
  }
  interconnect-device {
    IC-F4908;
  }
}
node-device P2025-C {
  interconnect-device IC-F4908 {
    preference normal;
  }
}
```

Verification

Verifying Flow Groups

Purpose

Check to see if the INEs in a QFabric system recognize the flow groups.

Action

After configuring the flow groups, you can see if the preference configuration has taken effect by running the following command on any of the INEs in a QFabric system:

```
root@P2025-C> show fabric multicast vccpdf-adjacency
```

Flags: S - Stale

Src Dev id	Src INE	Src Dev type	Dest Dev id	Interface	Flags	Src Port	Dest Port	link
Preference								
2	260	TOR	256	n/a		-1	-1	
NORMAL								
2	260	TOR	768	n/a		-1	-1	
NORMAL								
3	262	TOR	256	n/a		-1	-1	
NORMAL								
3	262	TOR	774	n/a		-1	-1	
NORMAL								
5	265(s)	TOR	257	fte-0/1/2.32768		2	9	HIGH
256	256	F1_F3	2	n/a		-1	-1	
NORMAL								

RELATED DOCUMENTATION

Understanding QFabric Multicast Flow Groups 55
Segregating QFabric Traffic Flows With Flow Groups 429
QFabric System Overview 3

QFabric System Fabric OAM Configuration

IN THIS CHAPTER

- Overview of Internal Fabric Monitoring | 439
- Configuring a Fabric Maintenance Association | 441
- Configuring Flow Specifications | 442
- Example: Configuring Internal Fabric OAM Monitoring | 448

Overview of Internal Fabric Monitoring

Internal fabric monitoring is a feature of the Operation, Administration, and Maintenance (OAM) of the QFabric system. This feature enables you to validate the flow path of protocol data units (PDUs) across a given VLAN on the QFabric system using the unicast ping, multicast ping, and traceroute operations.

Internal fabric monitoring is useful for fault detection on the QFabric system. For example, if a PDU reaches a destination that is not part of the VLAN configuration, the operation (unicast ping, multicast ping, or traceroute) displays the exception on the console at runtime.

The unicast and multicast ping operations send PDUs from a source interface (called the source fabric maintenance endpoint [FMEP]) to a destination FMEP. The destination FMEP sends a response to the source FMEP when the PDUs are received.

The traceroute operation, also called a flow linktrace, traces the path taken by a specific, learned unicast flow from a source FMEP to a destination FMEP in a VLAN. The source and destination FMEPs may be on the same Node device, different Node devices connected to the same Interconnect device, or different Node devices connected to different Interconnect devices. The flow path is the sequence of Packet Forwarding Engine forwarding hops along which the PDU travels. The hop-by-hop sequence and number of hops are reported in terms of fabric maintenance intermediate points (FMIPs), which are interfaces on the Packet Forwarding Engine of the Interconnect device. An FMIP sends a response to the source FMEP when the traceroute PDU is received.

The following internal fabric monitoring commands are supported:

- **show oam fabric flow specification**
- **show oam fabric interfaces**

- **ping fabric unicast-flow**
- **ping fabric multicast-flow**
- **traceroute fabric unicast-flow**

The following are the main components of the internal fabric monitoring feature:

- **FMEP**—Represents the source or destination point (endpoint) in the monitoring operations. An FMEP is an interface through which PDUs are sent (source FMEP) or received (destination FMEP). Upon receipt of the PDUs, the destination FMEP sends a response to the source FMEP to validate the monitoring flow. FMEPs are associated with a VLAN in the fabric maintenance association (FMA) configuration, and the source and destination FMEP addresses are configured in the flow specification.
- **FMA**—Associates a set of FMEPs with a VLAN. The FMA defines the VLAN and FMEP parameters, including the VLAN name, FMEP identifiers, FMEP names, and the interface names of the FMEPs. The FMEPs defined in the FMA are the source and destination FMEPs used in the monitoring commands.

NOTE: A default FMA is automatically created for each Node group in the QFabric system. The default FMA is used to send error response PDUs (for example, in the case of a VLAN leak) and responses to PDUs that are not mapped to a specific interface in the QFabric system.

- **Flow specification**—Configures the flow type and FMEP addressing parameters. Unicast flow types include an Ethernet type (other than IPv4) and Ethernet IPv4. Multicast flow types include the Ethernet IPv4 and VLAN flood type. The flow specification also defines parameters within each flow type, such as MAC or IPv4 addresses of the source and destination FMEPs. The names and identifiers of these FMEPs are configured in the FMA parameters.

To enable internal fabric monitoring, configure the **fabric-maintenance-associations** and **flow-specs** statements at the **[edit protocols oam fabric]** hierarchy level.

The **ping fabric unicast-flow**, **ping fabric multicast-flow**, and **traceroute fabric unicast-flow** commands require that you specify the flow specification and FMA names, as well as the source and destination FMEP names.

RELATED DOCUMENTATION

[Configuring a Fabric Maintenance Association | 441](#)

[Configuring Flow Specifications | 442](#)

[fabric \(OAM\) | 470](#)

[ping fabric multicast-flow | 591](#)

[ping fabric unicast-flow | 593](#)

[traceroute fabric unicast-flow | 828](#)

Configuring a Fabric Maintenance Association

On a QFabric system, enabling internal fabric monitoring using the ping and traceroute operations requires that a fabric maintenance association (FMA) be configured in addition to a flow specification. The FMA associates a particular VLAN with a set of source and destination interfaces called fabric maintenance endpoints (FMEPs) that are used for tracing the flow path between the endpoints. Configuration of the VLAN name is mandatory.

Before configuring the FMA, configure a VLAN name (and not a VLAN ID) for the VLAN in which the internal fabric monitoring occurs. If you already configured a VLAN ID instead of a VLAN name, you must first delete the VLAN ID and then specify a VLAN name before you configure the FMA.

NOTE: Because each FMA and VLAN have a one-to-one correspondence, an FMA in each QFabric system must be given a unique name.

To configure an FMA:

1. Configure an FMA name that is unique within the QFabric system:

```
[edit protocols oam fabric]
user@host# set fabric-maintenance-associations fma-name
```

2. (Optional) Configure the FMA description:

```
[edit protocols oam fabric]
user@host# set fabric-maintenance-associations fma-name description string
```

3. Configure the VLAN name:

```
[edit protocols oam fabric]
user@host# set fabric-maintenance-associations fma-name vlan-name vlan-name
```

4. Configure an FMEP, including the identifier, name, description, and associated interface:


```
[edit protocols oam fabric]
user@host# set fabric-maintenance-associations fma-name fabric-maintenance-end-points fmem-id fmem-name
name description string fmem-interface interface-name
```

5. Repeat Step 4 to configure additional FMEPs.

RELATED DOCUMENTATION

[Overview of Internal Fabric Monitoring | 439](#)

[Configuring Flow Specifications | 442](#)

[fabric \(OAM\) | 470](#)

[ping fabric multicast-flow | 591](#)

[ping fabric unicast-flow | 593](#)

[traceroute fabric unicast-flow | 828](#)

Configuring Flow Specifications

IN THIS SECTION

- [Configuring a Unicast Ethernet Flow Specification | 443](#)
- [Configuring a Unicast Ethernet IPv4 Flow Specification | 445](#)
- [Configuring a Multicast IPv4 Flow Specification | 446](#)
- [Configuring a Multicast VLAN Flood Flow Specification | 447](#)

On a QFabric system, enabling internal fabric monitoring using the ping and traceroute operations requires that a flow specification be configured in addition to a fabric maintenance association (FMA). A flow specification is used to determine the contents of the packet that is used for a particular operation. For example, an Ethernet type flow specification describes the contents of a pure Layer 2 packet, whereas an Ethernet IPv4 type flow specification describes an Ethernet IPv4 packet.

NOTE: In the flow specification configuration, some parameters are optional. If you do not configure the optional parameters, the system generates random values that are used as PDU contents.

In order to replicate actual traffic behavior, internal fabric monitoring requires that the following flow specification fields be filled with the same values as the traffic packets:

- Unicast Ethernet—Source MAC address, destination MAC address, Ethertype
- Unicast IPv4—Source IP address, destination IP address, source L4 port, destination L4 port, destination MAC address

The configuration for each flow specification type includes parameters specific to that particular type. Use one of the following procedures to configure the flow specification for your flow type:

Configuring a Unicast Ethernet Flow Specification

This procedure describes how to configure a flow specification for a unicast Ethernet flow type. This configuration is applicable for the unicast ping and traceroute operations.

NOTE: Some parameters are optional. If you do not configure the optional parameters, the system generates random values that are used as PDU contents.

In order to replicate actual traffic behavior, internal fabric monitoring requires that the following flow specification fields be filled with the same values as the traffic packets: source MAC address, destination MAC address, and Ethertype.

To configure a unicast Ethernet flow specification:

1. Configure a flow specification name:

```
[edit protocols oam fabric]
user@host# set flow-specs flow-specification-name
```

2. (Optional) Set the Ethernet frame size:

```
[edit protocols oam fabric]
user@host# set ethernet-frame-size ethernet-frame-size
```

3. Set the flow specification type as unicast Ethernet and specify the EtherType:


```
[edit protocols oam fabric flow-specs flow-specification-name]
user@host# set unicast-ethernet ethertype ethertype
```

4. (Optional) Set the MAC address and the MAC address range of the source FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet]
user@host# set source-mac source-mac-address source-mac-mask source-mac-mask
```

NOTE: Configuration of the **source-mac** and **source-mac-mask** parameters is optional, but the **source-mac-mask** parameter can be specified only if the **source-mac** parameter is also configured.

5. (Optional) Set the MAC address or address range mask for the destination FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet]
user@host# set destination-mac destination-mac-address destination-mac-mask destination-mac-mask
```

NOTE: Configuration of the **destination-mac** and **destination-mac-mask** parameters is optional, but the **destination-mac-mask** parameter can be specified only if the **destination-mac** parameter is also configured.

Configuring a Unicast Ethernet IPv4 Flow Specification

This procedure describes how to configure a flow specification for a unicast Ethernet IPv4 flow type. This configuration is applicable for the unicast ping and traceroute operations.

NOTE: All parameters except the flow specification name are optional. If you do not configure the optional parameters, the system generates random values that are used as PDU contents.

In order to replicate actual traffic behavior, internal fabric monitoring requires that the following flow specification fields be filled with the same values as the traffic packets: source IP address, destination IP address, source L4 port, destination L4 port, destination MAC address.

To configure a unicast Ethernet IPv4 flow specification:

1. Configure a flow specification name:

```
[edit protocols oam fabric]
user@host# set flow-specs flow-specification-name
```

2. (Optional) Set the Ethernet frame size:

```
[edit protocols oam fabric flow-specs flow-specification-name]
user@host# set ethernet-frame-size size
```

3. Set the flow specification type as unicast Ethernet IPv4:

```
[edit protocols oam fabric flow-specs flow-specification-name]
user@host# set unicast-ethernet-ipv4
```

4. (Optional) Set the IPv4 address or the address range of the source FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# set source-ip source-ip-address
```

or

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# set source-ip-mask source-ip-mask
```


5. (Optional) Set the IPv4 address or the address range of the destination FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# set destination-ip destination-ip-address
```

or

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# set destination-ip-mask destination-ip-mask
```

6. (Optional) Set the IPv4 protocol type:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# ip-proto protocol
```

7. (Optional) Set the L4 port (TCP or UDP port number) of the source FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# source-l4-port source-l4-port-number
```

8. (Optional) Set the L4 port of the destination FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name unicast-ethernet-ipv4]
user@host# destination-l4-port destination-l4-port-number
```

Configuring a Multicast IPv4 Flow Specification

This procedure describes how to configure a flow specification for a multicast IPv4 flow type.

NOTE: Some parameters are optional. If you do not configure the optional parameters, the system generates random values that are used as PDU contents.

To configure a multicast IPv4 flow specification:

1. Configure a flow specification name:

```
[edit protocols oam fabric]
```



```
user@host# set flow-specs flow-specification-name
```

2. Set the flow specification type as multicast IPv4:

```
[edit protocols oam fabric flow-specs flow-specification-name]  
user@host# set multicast-ipv4
```

3. (Optional) Set the IPv4 address of the source FMEP:

```
[edit protocols oam fabric flow-specs flow-specification-name multicast-ipv4]  
user@host# set source-ip ipv4-address
```

4. Set the IPv4 multicast group address:

```
[edit protocols oam fabric flow-specs flow-specification-name multicast-ipv4]  
user@host# set dest-ip-multicast-group ipv4-mcast-address
```

Configuring a Multicast VLAN Flood Flow Specification

This procedure describes how to configure a flow specification for a multicast VLAN flood flow type.

To configure a multicast VLAN flood flow specification:

1. Configure a flow specification name:

```
[edit protocols oam fabric]  
user@host# set flow-specs flow-specification-name
```

2. Set the flow specification type as multicast VLAN flood:

```
[edit protocols oam fabric flow-specs flow-specification-name]  
user@host# set multicast-vlan-flood
```

RELATED DOCUMENTATION

[Overview of Internal Fabric Monitoring](#) | 439

[Configuring a Fabric Maintenance Association | 441](#)[fabric \(OAM\) | 470](#)[ping fabric multicast-flow | 591](#)[ping fabric unicast-flow | 593](#)[traceroute fabric unicast-flow | 828](#)

Example: Configuring Internal Fabric OAM Monitoring

IN THIS SECTION

- [Requirements | 448](#)
- [Overview | 449](#)
- [Configuration | 451](#)
- [Verification | 453](#)

This example shows how to configure the internal fabric Operation, Administration, Maintenance (OAM) monitoring feature on the QFabric system, including the fabric maintenance association (FMA) and flow specifications for unicast Ethernet traffic.

Internal fabric monitoring enables you to validate the flow path of protocol data units (PDUs) across a given VLAN on the QFabric system using the ping operation.

Internal fabric monitoring is useful for fault detection on the QFabric system. For example, if a PDU reaches a destination that is not part of the VLAN configuration, the ping operation displays the exception on the console at runtime.

Requirements

This example uses the following hardware and software components:

- One QFabric system with the following components:
 - A Director group (two Director devices)
 - Two QFabric system Interconnect devices
 - Node devices, including:
 - One Node device (Node-81) configured in a default network Node group (NNW-NG-0)

- One Node device (Node-80) in an autogenerated server Node group (SNG-80)
 - Two Node devices (Node-83 and Node-84) configured in a redundant server Node group (RSNG-8384)
- Junos OS Release 12.2 or later

Before you configure the FMA, first configure a VLAN in which the internal fabric OAM monitoring occurs. You can usually configure a VLAN by specifying a VLAN ID or VLAN name. However, in the case of the VLAN used for internal fabric OAM monitoring, you must specify a VLAN name (not a VLAN ID). The FMA configuration requires a VLAN name and not a VLAN ID.

If you have already configured a VLAN ID for the VLAN, you must delete the VLAN ID from the associated interfaces and then add a VLAN name.

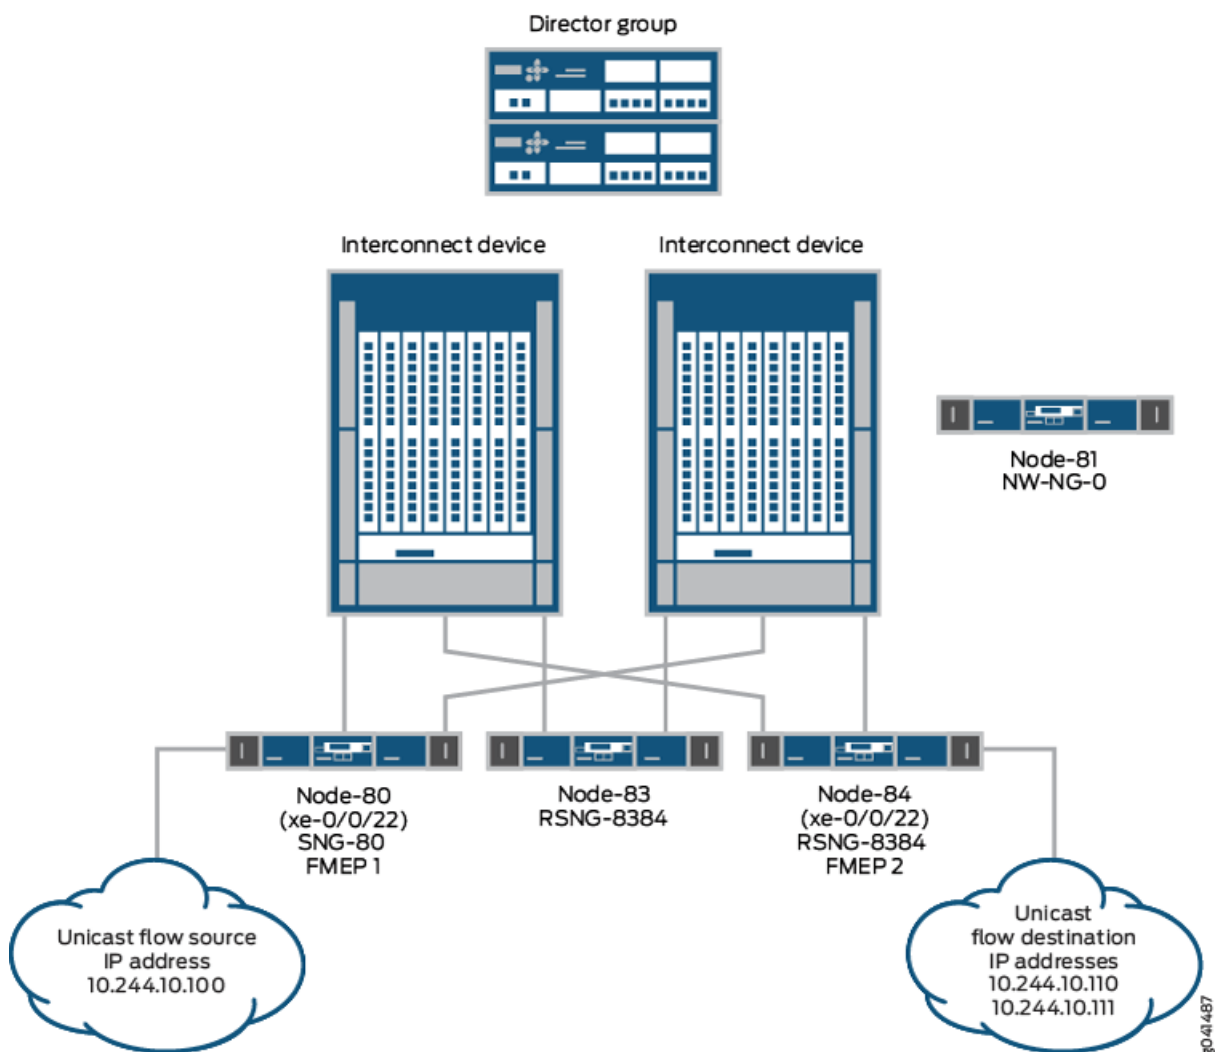
Overview

On a QFabric system, enabling internal fabric OAM monitoring requires that you configure an FMA and a flow specification.

- The *FMA* associates a particular VLAN with a set of source and destination interfaces called fabric maintenance endpoints (FMEPs) that are used for tracing the data path between the endpoints. Configuration of the VLAN name is mandatory when you configure an FMA.
- A *flow specification* is used to determine the contents of a packet that is used for a particular operation. For example, an Ethernet type flow specification describes the contents of a pure Layer 2 packet, whereas an Ethernet IPv4 type flow specification describes an Ethernet IPv4 packet.

Topology

This example uses the following topology:



- The QFabric system Node devices are connected to the Interconnect devices.
- Node device Node-80 is configured in the server Node group SNG-80.
- Node devices Node-83 and Node-84 are configured in a redundant server Node group RSNG-8384.
- The device with the unicast flow source IP address is connected to Node-80.
- The device with the unicast flow destination IP addresses is connected to Node-84.

[Table 78 on page 451](#) maps the association of the fabric OAM elements configured in this example.

Table 78: Fabric OAM Configuration Elements

Interface	FMA	VLAN Name	FMEP ID	Flow Specifications
Node-80:xe-0/0/22	fma1	v10	1	<ul style="list-style-type: none"> flowspec1 flowspec2
Node-84:xe-0/0/22	fma1	v10	2	<ul style="list-style-type: none"> flowspec1 flowspec2

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

set vlans v10 vlan-id 10
set interfaces Node-80:xe-0/0/22 unit 0 family ethernet-switching vlan members v10
set interfaces Node-84:xe-0/0/22 unit 0 family ethernet-switching vlan members v10
set protocols oam fabric fabric-maintenance-associations fma1 vlan-name v10
set protocols oam fabric fabric-maintenance-associations fma1 fabric-maintenance-end-points 1 fmep-interface
Node-80:xe-0/0/22.0
set protocols oam fabric fabric-maintenance-associations fma1 fabric-maintenance-end-points 2 fmep-interface
Node-84:xe-0/0/22.0
set protocols oam fabric flow-specs flowspec1 unicast-ethernet-ipv4 source-ip 10.244.10.100
set protocols oam fabric flow-specs flowspec1 unicast-ethernet-ipv4 destination-ip 10.244.10.110
set protocols oam fabric flow-specs flowspec1 unicast-ethernet-ipv4 destination-mac 00:00:80:8d:2d:15
set protocols oam fabric flow-specs flowspec1 unicast-ethernet-ipv4 source-l4-port 60
set protocols oam fabric flow-specs flowspec1 unicast-ethernet-ipv4 destination-l4-port 80
set protocols oam fabric flow-specs flowspec2 unicast-ethernet-ipv4 source-ip 10.244.10.100
set protocols oam fabric flow-specs flowspec2 unicast-ethernet-ipv4 destination-ip 10.244.10.111
set protocols oam fabric flow-specs flowspec2 unicast-ethernet-ipv4 destination-mac 00:00:80:8d:2d:17
set protocols oam fabric flow-specs flowspec2 unicast-ethernet-ipv4 source-l4-port 60
set protocols oam fabric flow-specs flowspec2 unicast-ethernet-ipv4 destination-l4-port 80

```

Step-by-Step Procedure

To configure internal fabric OAM monitoring:

1. Configure a name and ID for the VLAN.

```

[edit]
set vlans v10 vlan-id 10

```


2. Assign a set of interfaces to the VLAN.

```
[edit]
set interfaces Node-80:xe-0/0/22 unit 0 family ethernet-switching vlan members v10
set interfaces Node-84:xe-0/0/22 unit 0 family ethernet-switching vlan members v10
```

3. Define the FMA and associate it with a VLAN name.

```
[edit protocols oam fabric]
set fabric-maintenance-associations fma1 vlan-name v10
```

4. Configure FMEPs for the FMA.

```
[edit protocols oam fabric]
set fabric-maintenance-associations fma1 fabric-maintenance-end-points 1 fmep-interface
  Node-80:xe-0/0/22.0
set fabric-maintenance-associations fma1 fabric-maintenance-end-points 2 fmep-interface
  Node-84:xe-0/0/22.0
```

5. Configure two IPv4 flow specifications for unicast traffic from FMEP 1 to FMEP 2.

```
[edit protocols oam fabric]
set flow-specs flowspec1 unicast-ethernet-ipv4 source-ip 10.244.10.100
set flow-specs flowspec1 unicast-ethernet-ipv4 destination-ip 10.244.10.110
set flow-specs flowspec1 unicast-ethernet-ipv4 destination-mac 00:00:80:8d:2d:15
set flow-specs flowspec1 unicast-ethernet-ipv4 source-l4-port 60
set flow-specs flowspec1 unicast-ethernet-ipv4 destination-l4-port 80
set flow-specs flowspec2 unicast-ethernet-ipv4 source-ip 10.244.10.100
set flow-specs flowspec2 unicast-ethernet-ipv4 destination-ip 10.244.10.111
set flow-specs flowspec2 unicast-ethernet-ipv4 destination-mac 00:00:80:8d:2d:17
set flow-specs flowspec2 unicast-ethernet-ipv4 source-l4-port 60
set flow-specs flowspec2 unicast-ethernet-ipv4 destination-l4-port 80
```

Results

Display the results of the configuration.

```
root@qfabric> show protocols oam fabric
```

```
fabric-maintenance-associations {
  fma1 {
```



```

    vlan-name v10;
    fabric-maintenance-end-points {
        1 {
            fmep-interface Node-80:xe-0/0/22.0;
        }
        2 {
            fmep-interface Node-84:xe-0/0/22.0;
        }
    }
}
}
flow-specs {
    flowspec1 {
        unicast-ethernet-ipv4 {
            source-ip 10.244.10.100;
            destination-ip 10.244.10.110;
            destination-mac 0:0:80:8d:2d:15;
            source-l4-port 60;
            destination-l4-port 80;
        }
    }
    flowspec2 {
        unicast-ethernet-ipv4 {
            source-ip 10.244.10.100;
            destination-ip 10.244.10.111;
            destination-mac 0:0:80:8d:2d:17;
            source-l4-port 60;
            destination-l4-port 80;
        }
    }
}
}

```

Verification

IN THIS SECTION

- [Verifying Configuration of Fabric OAM Interfaces | 454](#)
- [Verifying Configuration of the Fabric OAM Flow Specification | 454](#)
- [Verifying Operation of the Unicast ping Command | 455](#)

Confirm that the configuration is working properly.

Verifying Configuration of Fabric OAM Interfaces

Purpose

Verify that the interfaces are associated with the FMA.

Action

Display the fabric OAM interfaces.

```
root@qfabric> show oam fabric interfaces
```

Interface-name	Fabric-Maintenance Association	VLAN	Interface state	MEP Identifier	MEP Name
NW-NG-0:NULL fmep-default-BBAK8793	fma-default	*	up	32790	
NW-NG-0:NULL fmep-default-BBAK8798	fma-default	*	up	32791	
RSNG-8384:NULL fmep-default-BBAK8748	fma-default	*	up	32794	
RSNG-8384:NULL fmep-default-BBAK0465	fma-default	*	up	32795	
Node-84:xe-0/0/22.0	fma1	V10	up	2	
SNG-60:NULL fmep-default-BBAK9603	fma-default	*	up	32793	
SNG-80:NULL fmep-default-BBAK8728	fma-default	*	up	32792	
Node-80:xe-0/0/22.0	fma1	V10	up	1	

Meaning

Interfaces Node-84:xe-0/0/22.0 and Node-80:xe-0/0/22.0 are associated with the FMA, FMEP, and VLAN as configured.

Interfaces that are not configured for fabric OAM monitoring (NW-NG-0:NULL, RSNG-8384:NULL, SNG-60:NULL, and SNG-80:NULL) display an asterisk (*) in the VLAN column, and the default FMA and FMEP names in the Fabric-Maintenance Association and MEP Name columns.

Verifying Configuration of the Fabric OAM Flow Specification

Purpose

Verify the configuration of the flow specifications.

Action

Display the fabric OAM flow specification configuration.

```
root@qfabric> show oam fabric flow-specification
```

```
Flow specification name : flowspec1 Type : Ethernet Unicast IPV4
  Ethernet frame size : Unspecified
  Source-IP : 10.244.10.100
  Source-IP Mask : Unspecified
  Destination-IP : 10.244.10.110
  Destination-IP Mask : Unspecified
  Destination-mac : 0:0:80:8d:2d:15
  IP-protocol : Unspecified
  Source-L4-port : 60
  Destination-L4-port : 80
Flow specification name : flowspec2 Type : Ethernet Unicast IPV4
  Ethernet frame size : Unspecified
  Source-IP : 10.244.10.100
  Source-IP Mask : Unspecified
  Destination-IP : 10.244.10.111
  Destination-IP Mask : Unspecified
  Destination-mac : 0:0:80:8d:2d:17
  IP-protocol : Unspecified
  Source-L4-port : 60
  Destination-L4-port : 80
```

Meaning

The output shows that flow specifications were configured as the user intended.

Verifying Operation of the Unicast ping Command

Purpose

Verify that the fabric OAM unicast **ping** command works.

Action

Issue the **ping fabric unicast-flow** command.

```
root@qfabric> ping fabric unicast-flow source-fmep-id 1 dest-fmep-id 2 flow-spec-name flowspec1
fma-name fma1 count 3
```

```
Fabric flow ping between source Node-80 destination Node-84
received response from fmep-id 2...
received response from fmep-id 2...
received response from fmep-id 2...
```



```
sent 3 requests, received 3 responses

root@qfabric> ping fabric unicast-flow source-fmep-id 1 dest-fmep-id 2 fma-name
fma1 flow-spec-name flowspec2
Fabric flow ping between source Node-80 destination Node-84
received response from fmep-id 2...
sent 1 requests, received 1 responses
```

Meaning

The fabric OAM unicast ping request from the source Node device (Node-80) to the destination Node device (Node-84) has been successful.

RELATED DOCUMENTATION

[Overview of Internal Fabric Monitoring | 439](#)

[fabric-maintenance-associations | 473](#)

[fabric-maintenance-end-points | 475](#)

[flow-specs | 486](#)

[show oam fabric flow-specification | 820](#)

[show oam fabric interfaces | 824](#)

Configuration Statements

IN THIS CHAPTER

- aliases | 459
- archive (QFabric System) | 460
- chassis (QFabric System) | 461
- device-authentication | 464
- director-device (Aliases) | 466
- ethernet-frame-size | 467
- fabric | 468
- fabric (OAM) | 470
- fabric-control | 472
- fabric-maintenance-associations | 473
- fabric-maintenance-end-points | 475
- file (QFabric System) | 477
- flow-groups | 479
- flow-specs | 486
- graceful-restart (Enabling Globally) | 488
- graceful-restart (Fabric Control) | 490
- graceful-restart (Protocols BGP) | 491
- graceful-restart (Protocols OSPF) | 493
- interconnect-device (Chassis) | 495
- interconnect-device (Aliases) | 497
- interconnect-device (Flow Groups) | 499
- multicast (QFabric Routing Options) | 500
- multicast-ipv4 | 501
- multicast-vlan-flood | 502
- network-domain | 503
- no-make-before-break | 504
- node-device (Aliases) | 505
- node-device (Chassis) | 507

- node-device (Flow Groups) | 508
- node-device (Resources) | 509
- node-group (Chassis) | 510
- node-group (Resources) | 512
- pic (Port) | 514
- protocols (Fabric) | 516
- resources | 517
- restart-time (Fabric Control) | 518
- routing-options (QFabric System) | 519
- syslog (QFabric System) | 520
- stale-routes-time (Fabric Control) | 521
- unicast-ethernet | 522
- unicast-ethernet-ipv4 | 524

aliases

Syntax

```
aliases {
  director-device director-device-name {
    assigned-director-device-name;
  }
  interconnect-device interconnect-device-name {
    assigned-interconnect-device-name;
  }
  node-device node-device-name {
    assigned-node-device-name;
  }
}
```

Hierarchy Level

[edit fabric]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Options **director-group** and **interconnect-device** introduced in Junos OS Release 12.2 for the QFX Series.

Description

(QFabric systems only) Specify the mapping of user defined names to QFabric system components. You can reassign names for Director devices, Interconnect devices, and Node devices.

The remaining statements are explained separately. See [CLI Explorer](#).

NOTE: The following rules apply to QFabric component alias naming:

- Alias names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.
- The maximum length of an alias name is 30 characters.
- Alias names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.
- You cannot use the reserved names **all**, **fabric**, or **director-group** as an alias name.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Aliases for the QFabric System | 388](#)

[Understanding Node Devices | 33](#)

archive (QFabric System)

Syntax

```
archive {  
    size size;  
}
```

Hierarchy Level

```
[edit system syslog file filename]
```

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure the archiving properties for the system message log file.

Options

size *size*—Maximum amount of system log message data that the QFabric system stores in the log file.

Syntax: xk to specify the number of kilobytes, xm for the number of megabytes, or xg for the number of gigabytes

Range: 65 KB through 1 GB

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[syslog | 520](#)

chassis (QFabric System)

Syntax

```
chassis {
  interconnect-device {
    alarm {
      (ethernet (Alarm) | management-ethernet) {
        link-down (red | yellow | ignore);
      }
    }
  }
  container-devices {
    device-count number;
  }
  craft-lockout {
    alarm {
      interface-type {
        link-down (red | yellow | ignore);
      }
    }
    container-devices {
      device-count number;
    }
    fpc slot {
      power (on | off);
    }
    routing-engine {
      on-disk-failure {
        disk-failure-action (halt | reboot);
      }
    }
  }
  fpc slot {
    power (on | off);
  }
  routing-engine {
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
}
node-group name {
  aggregated-devices {
    ethernet {
      device-count number;
    }
  }
}
```



```

    }
  }
  alarm {
    interface-type {
      link-down (ignore | red | yellow);
    }
  }
  container-devices {
    device-count number;
  }
  node-device name {
    fibre-channel {
      port-range {
        port-range-low port-range-high;
      }
    }
    pic pic-number {
      fte {
        port port-number;
        port-range port-range-low port-range-high;
      }
      xe {
        port port-number;
        port-range port-range-low port-range-high;
      }
      xle {
        port port-number;
        port-range port-range-low port-range-high;
      }
    }
  }
  routing-engine {
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
}

```

Hierarchy Level

[edit]

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Configure chassis-specific properties for the switch.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

device-authentication

Syntax

```
device-authentication {
  (encrypted-password "password" | plain-text-password);
}
```

Hierarchy Level

```
[edit system]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure the authentication password used when accessing individual QFabric system components with the **request component login** command.

NOTE: Configuring this statement overrides the password for individual QFabric system components that was set during the initial QFabric system setup procedure.

Options

encrypted-password "password"— Specify the MD5 or other encrypted authentication password. You can specify only one encrypted password.

You cannot configure a blank password for the **encrypted-password** option using blank quotation marks (" "). You must configure a password of 1 through 128 characters and enclose the password in quotation marks.

plain-text-password—Plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring QFabric System Login Classes | 377

Performing the QFabric System Initial Setup on a QFX3100 Director Group | 363

director-device (Aliases)

Syntax

```
director-device director-device-name {
    assigned-director-device-name;
}
```

Hierarchy Level

```
[edit fabric aliases]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

(QFabric systems only) Specify the mapping of user-defined names to QFabric system Director devices.

NOTE: The following rules apply to QFabric component alias naming:

- Alias names must use alphabetic (**A** through **Z** and **a** through **z**), numeric (**0** through **9**), or dash (-) characters.
- The maximum length of an alias name is 30 characters.
- Alias names are case sensitive. For example, **MY-NG-1** and **my-ng-1** refer to different components.
- You cannot use the reserved names **all**, **fabric**, or **director-group** as an alias name.

Options

director-device-name—Specify a user-defined name for a QFabric system Director device.

assigned-director-device-name—Specify the Director device identifier or name that has been provided. Identifiers are usually auto-generated by the Director software, and names are usually provided by the administrator of the default partition.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Aliases for the QFabric System | 388](#)
[Understanding the Director Group | 26](#)

ethernet-frame-size

Syntax

```
ethernet-frame-size ethernet-frame-size;
```

Hierarchy Level

```
[edit protocols oam fabric flow-specs flow-specification-name]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure the Ethernet frame size in the flow specification parameters for internal fabric monitoring.

Options

ethernet-frame-size—Integer defining the size (in bytes) of the Ethernet frame.

Range: 256 to 9116 bytes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Overview of Internal Fabric Monitoring | 439](#)
[Configuring a Fabric Maintenance Association | 441](#)
[Configuring Flow Specifications | 442](#)
[fabric \(OAM\) | 470](#)
[ping fabric multicast-flow | 591](#)
[ping fabric unicast-flow | 593](#)
[traceroute fabric unicast-flow | 828](#)

fabric

Syntax

```

fabric
  aliases {
    director-device director-device-name {
      assigned-director-device-name;
    }
    interconnect-device interconnect-device-name {
      assigned-interconnect-device-name;
    }
    node-device node-device-name {
      assigned-node-device-name;
    }
  }
  flow-groups {
    flow-group-name {
      node-device (node-device-name | serial-ID | alias);
      interconnect-device interconnect-device-name;
    }
    node-device (node-device-name | serial-ID | alias) {
      interconnect-device interconnect-device-name {
        preference (high | normal | never);
      }
    }
  }
  protocols {
    fabric-control {
      graceful-restart {
        restart-time seconds;
        stale-routes-time seconds;
      }
    }
  }
  resources {
    node-group node-group-name {
      node-device node-device-name;
      network-domain;
    }
  }
}

```

Hierarchy Level

[edit]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Define resources, routing options, fabric control protocol settings, and the mapping of user-defined names to QFabric system components.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[show fabric administration inventory](#) | **780**

[Understanding QFabric System Administration Tasks and Utilities](#) | **371**

fabric (OAM)

Syntax

```

fabric {
  fabric-maintenance-associations {
    fma-name {
      description string;
      fabric-maintenance-end-points {
        fmep-id {
          description string;
          fmep-interface interface-name;
          fmep-name name;
        }
      }
      vlan-name vlan-name;
    }
  }
  flow-specs {
    flow-specification-name {
      ethernet-frame-size size;
      multicast-ipv4 {
        dest-ip-multicast-group ipv4-mcast-address;
        source-ip ipv4-address;
      }
      multicast-vlan-flood;
      unicast-ethernet {
        destination-mac destination-mac-address destination-mac-mask destination-mac-mask;
        ethertype ethertype;
        source-mac mac-address source-mac-mask source-mac-mask;
      }
      unicast-ethernet-ipv4 {
        destination-ip destination-ip-address destination-ip-mask destination-ip-mask;
        destination-l4-port destination-l4-port-number;
        destination-mac destination-mac-address;
        ip-proto protocol;
        source-ip source-ip-address source-ip-mask source-ip-mask;
        source-l4-port source-l4-port-number;
      }
    }
  }
  faboam-trace-options {
    file filename{
      files number;
      no-stamp;
    }
  }
}

```



```

        no-world-readable;
        replace;
        size size;
        world-readable;
    }
    flag {
        all;
        debug-all;
        ffping;
        generic;
        netio;
        packets;
        trace-route;
    }
}
}
}

```

Hierarchy Level

[edit protocols oam]

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure the fabric maintenance association (FMA) and flow specification parameters for internal fabric monitoring.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring 439
Configuring a Fabric Maintenance Association 441
Configuring Flow Specifications 442
ping fabric multicast-flow 591
ping fabric unicast-flow 593
traceroute fabric unicast-flow 828

fabric-control

Syntax

```
fabric-control {
  graceful-restart {
    restart-time seconds;
    stale-routes-time seconds;
  }
}
```

Hierarchy Level

[edit fabric [protocols](#)]

Release Information

Statement introduced in Junos OS Release 13.2X52-D10 for the QFX Series.

Description

Specify attributes for the fabric control protocol.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Routing Engines in the QFabric System 27
--

fabric-maintenance-associations

Syntax

```
fabric-maintenance-associations {
  fma-name {
    description string;
    fabric-maintenance-end-points {
      fmep-id {
        description string;
        fmep-interface interface-name;
        fmep-name name;
      }
    }
    vlan-name vlan-name;
  }
}
```

Hierarchy Level

```
[edit protocols oam fabric]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure the fabric maintenance association (FMA) parameters for internal fabric monitoring. The FMA defines a set of QFabric system interfaces (fabric maintenance endpoints [FMEPs]) that are members of a given VLAN which are used in internal fabric monitoring operations. The FMA associates the set of FMEPs with a VLAN. Configuration of the VLAN is mandatory.

The remaining statements under this hierarchy level are explained separately.

Options

description *string*—Description of the FMA.

Syntax: 1 to 128 alphanumeric characters

fma-name—Name of the FMA.

Syntax: 1 to 20 alphanumeric characters

vlan-name *vlan-name*—Name of the VLAN for which the FMA is defined.

NOTE: Configuring the VLAN option is mandatory.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Overview of Internal Fabric Monitoring | 439](#)

[Configuring a Fabric Maintenance Association | 441](#)

[Configuring Flow Specifications | 442](#)

[fabric \(OAM\) | 470](#)

[ping fabric multicast-flow | 591](#)

[ping fabric unicast-flow | 593](#)

[traceroute fabric unicast-flow | 828](#)

fabric-maintenance-end-points

Syntax

```
fabric-maintenance-end-points {
  fmep-id {
    description string;
    fmep-interface interface-name;
    fmep-name name;
  }
}
```

Hierarchy Level

```
[edit protocols oam fabric fabric-maintenance-associations fma-name]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure the fabric maintenance association endpoint (FMEP) parameters within a fabric maintenance association (FMA). The FMA associates the FMEP with a VLAN for purposes of internal fabric monitoring.

Options

description *string*—Description of the FMEP.

Syntax: 1 to 128 alphanumeric characters

fmep-id—32-bit integer that is used to identify an FMEP.

Range: 1 to 64,000

fmep-name name—Name of the fabric maintenance association endpoint.

Syntax: 1 to 20 alphanumeric characters

fmep-interface interface-name—Interface on a QFabric system Node device.

Syntax: *node-device-name:interface-type-fpc/pic/slot.unit*

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Overview of Internal Fabric Monitoring | 439](#)

[Configuring a Fabric Maintenance Association | 441](#)

[Configuring Flow Specifications | 442](#)

[fabric \(OAM\) | 470](#)

[ping fabric multicast-flow | 591](#)

[ping fabric unicast-flow | 593](#)

[traceroute fabric unicast-flow | 828](#)

file (QFabric System)

Syntax

```
file filename {  
  archive {  
    size maximum-file-size;  
  }  
  explicit-priority;  
  facility severity;  
  match "regular-expression";  
  structured-data {  
    brief;  
  }  
}
```

Hierarchy Level

[edit system [syslog](#)]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure the logging of system messages to a file.

Options

facility—Class of messages to log. To specify multiple classes, include multiple **facility severity** statements.

filename—Filename that you specify with the **show log** command.

Default: Filename **messages**

severity—Severity of the messages that belong to the facility specified by the paired **facility** name. Messages with severities at the specified level and higher are logged.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

flow-groups

Syntax

```

flow-groups {
  flow-group-name {
    node-device (node-device-name | serial-ID | alias);
    interconnect-device interconnect-device-name;
  }
  node-device (node-device-name | serial-ID | alias) {
    interconnect-device interconnect-device-name {
      preference (high | normal | never);
    }
  }
}

```

Hierarchy Level

[edit fabric]

Release Information

Statement introduced in Junos OS Release 14.1X53-D15 for the QFX Series.

Description

In a QFabric system, a hash function is used to select an Interconnect device to forward traffic between two ingress switches, which are called Node devices. Since this hash function is performed on all Interconnect devices, it is possible for redundant multicast streams to flow through one Interconnect device, making that Interconnect device a potential single point of failure for the redundant flows. Some operations, such as financial transactions, require redundant data flows to use different Interconnect devices to avoid point-of-failure data loss. To enforce this in a fabric, you can create a flow group for multicast or unicast traffic from a particular ingress switch. In a flow group, switches are forced to use different, specified Interconnect devices to forward redundant traffic.

NOTE: The following rules apply to QFabric system flow control:

- A Node device can belong to only one flow group.
- A fabric can contain up to four flow groups.
- Node devices in flow groups can only use the defined Interconnect devices in their flow groups unless all defined Interconnect devices fail. In case of failure, the flow group checks for Interconnect devices with preference set to **normal**—this includes all Interconnect devices in the default group and any Interconnect devices you specifically set to preference **normal**.

The following example creates a flow group. The name of the flow group, test1, is configured first. Next, the Node device TA3713480203 is configured, and the Interconnect device IC-P6603-C is configured last.

```
[edit fabric]
root@qfabric# set flow-groups ?
Possible completions:
<flow group name>          Flow Group name
+ apply-groups             Groups from which to inherit configuration data
+ apply-groups-except      Don't inherit configuration data from these groups

> node-device              Override the link preference for Node device(s)

root@qfabric# set flow-groups test1 ?
Possible completions:
+ apply-groups             Groups from which to inherit configuration data
+ apply-groups-except      Don't inherit configuration data from these groups
> interconnect-device      Interconnect device(s) associated with this flow
group
> node-device              Node device(s) associated with this flow group

root@qfabric# set flow-groups test1 node-device ?

Possible completions:
<node-device-name> Serial ID or alias of node-device
P3686-C               Node device
TA3713480203          Node device
TA3713500165          Node device
TA3713500185          Node device
TA3713500307          Node device

root@qfabric# set flow-groups test1 node-device TA3713480203 ?
Possible completions:
<Enter>               Execute this command
+ apply-group          Groups from which to inherit configuration data
+ apply-groups-except  Don't inherit configuration data from these groups
|                     Pipe through a command

root@qfabric# set fabric flow-groups test1 node-device TA3713480203

root@qfabric# set flow-groups test1 interconnect-device ?
Possible completions:
<interconnect-device-name> Serial ID or alias of interconnect-device
IC-P6601-C             Interconnect device
IC-P6603-C             Interconnect device
```



```

root@qfabric# set flow-groups test1 interconnect-device IC-P6603-C
root@qfabric# show flow-groups
test1 {
    node-device {
        TA3713480203;
    }
    interconnect-device {
        IC-P6603-C;
    }
}

```

The following example overrides the software settings the flow group test1 Node device TA3713500307 was assigned by the software for Interconnect device IC-P6601-C. Node device TA3713500307 has its preference for Interconnect device IC-P6601-C set to **never** by the software, meaning that Node device TA3713500307 can never use Interconnect device IC-P6601-C. You can change this preference so that Interconnect device IC-P6601-C can be used by Node device TA3713500307 if the specifically-configured Interconnect devices all fail. You use the preference option of the **flow-groups** configuration statement to set Interconnect device IC-P6601-C preference to **normal** for Node device TA3713500307.

To configure the preference setting for a Node device to use a particular Interconnect device, you first select the Node device, then specify the Interconnect device, and finally specify the preference option:


```

[edit fabric]
root@qfabric# set flow-groups ?

Possible completions:
<flow-group name>      Flow Group name
+ apply-groups          Groups from which to inherit configuration data
+ apply-groups-except   Don't inherit configuration data from these groups >
node-device            Override the link preference for Node device(s)
test1                  Flow Group name

root@qfabric# set flow-groups node-device ?

Possible completions:
<name> serial ID or alias of node-device
P3686-C                Node device
TA3713480203           Node device
TA3713500185           Node device
TA3713500307           Node device

root@qfabric# set fabric flow-groups node-device TA3713500307 ?
Possible completions:
+ apply-groups          Groups from which to inherit configuration data
+ apply-groups-except   Don't inherit configuration data from these groups
interconnect device

root@qfabric# set fabric flow-groups node-device TA3713500307 interconnect-device ?
Possible completions:
<name>                  Serial ID or alias of interconnect-device
IC-P6601-C              Interconnect device
IC-P6603-C              Interconnect device

root@qfabric# set fabric flow-groups node-device TA3713500307 interconnect-device
IC-P6601-C
preference ?
Possible completions:
high                    High link preference
never                   Never use this link
normal                  Normal link reference

root@qfabric# set fabric flow-groups node-device TA3713500307 interconnect-device
IC-P6601-C preference never

root@qfabric# show fabric flow-groups
test1 {

```



```

node-device {
    TA3713480203;
}
interconnect-device {
    IC-P6603-C;
}
}
node-device TA3713500307 {
    interconnect-device IC-P6601-C {
        preference never;
    }
}

```

TIP: To add a TOR to multiple flow groups:

set fabric resources node-device TOR-B interconnect-device IC-A preference HIGH

To configure a TOR to use an Interconnect device from a different flow group if none in the same flow group are available:

set fabric resources node-device TOR-B interconnect-device IC-A preference NORMAL

To increase the priority of default-flowgroup Interconnect devices for a specific TOR:

set fabric resources node-device TOR-A interconnect-device ALL_ICS preference HIGH

To configure default-flowgroup TORs to never use a particular Interconnect device:

Set fabric resources node-device ALL_TORS interconnect-device IC-A preference NEVER

Options

node-device *node-device-name* —Specify a QFabric system ingress Node switch by device name or alias in the QFabric system.

interconnect-device *interconnect-device-name*—Specify a QFabric system Interconnect device name.

preference *preference*—Overrides the algorithm results that assign preference to Node devices to use particular Interconnect devices. You can set preference values for these devices to **normal**, **high**, or **never**. (See [“Example: Creating a QFabric Flow Group” on page 431](#) for more details on setting **preference** values and their meanings.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Creating a QFabric Flow Group | 431](#)

[Understanding QFabric Multicast Flow Groups | 55](#)

[Segregating QFabric Traffic Flows With Flow Groups | 429](#)

flow-specs

Syntax

```

flow-specs {
  flow-specification-name {
    ethernet-frame-size ethernet-frame-size;
    multicast-ipv4 {
      dest-ip-multicast-group ipv4-mcast-address;
      source-ip ipv4-address;
    }
    multicast-vlan-flood;
    unicast-ethernet {
      destination-mac destination-mac-address destination-mac-mask destination-mac-mask;
      ethertype ethertype;
      source-mac mac-address source-mac-mask source-mac-mask;
    }
    unicast-ethernet-ipv4 {
      destination-ip destination-ip-address destination-ip-mask destination-ip-mask;
      destination-l4-port destination-l4-port-number;
      destination-mac destination-mac-address;
      ip-proto protocol;
      source-ip source-ip-address source-ip-mask source-ip-mask;
      source-l4-port source-l4-port-number;
    }
  }
}

```

Hierarchy Level

[edit protocols oam fabric]

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure the fabric flow specification parameters for internal fabric monitoring. The following conditions apply:

- You must specify a protocol type for each flow specification: **unicast-ethernet**, **unicast-ethernet-ipv4**, **multicast-ipv4**, or **multicast-vlan-flood**.
- You can specify only one type for each flow specification.
- All other parameters are optional.

The remaining statements under this hierarchy level are explained separately.

Options

flow-specification-name—Name of the flow specification.

Syntax: 1 to 20 alphanumeric characters

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring 439
Configuring a Fabric Maintenance Association 441
Configuring Flow Specifications 442
fabric (OAM) 470
ping fabric multicast-flow 591
ping fabric unicast-flow 593
traceroute fabric unicast-flow 828

graceful-restart (Enabling Globally)

Syntax

```
graceful-restart {  
    disable;  
    helper-disable;  
    maximum-helper-recovery-time seconds;  
    maximum-helper-restart-time seconds;  
    notify-duration seconds;  
    recovery-time seconds;  
    restart-duration seconds;  
    stale-routes-time seconds;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],  
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options],  
[edit routing-options],  
[edit routing-instances routing-instance-name routing-options]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

You configure the graceful restart routing option globally to enable the feature, but not to enable graceful restart for all routing protocols in a routing instance. To enable graceful restart globally, include the graceful-restart statement under the **[edit routing options]** hierarchy level. This enables graceful restart globally for all routing protocols. You can, optionally, modify the global settings at the individual protocol level.

NOTE:

- For VPNs, the **graceful-restart** statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.
- For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.
- LDP sessions flap when **graceful-restart** configurations change.

Default

Graceful restart is disabled by default.

Options

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Enabling Graceful Restart

Configuring Routing Protocols Graceful Restart

Configuring Graceful Restart for MPLS-Related Protocols

Configuring VPN Graceful Restart

Configuring Logical System Graceful Restart

[Configuring Graceful Restart for QFabric Systems | 424](#)

graceful-restart (Fabric Control)

Syntax

```
graceful-restart {  
    restart-time seconds;  
    stale-routes-time seconds;  
}
```

Hierarchy Level

[edit fabric protocols [fabric-control](#)]

Release Information

Statement introduced in Junos OS Release 13.2X52-D10 for the QFX Series.

Description

Configure graceful restart parameters for the fabric control in a QFabric system.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Routing Engines in the QFabric System](#) | 27

graceful-restart (Protocols BGP)

Syntax

```
graceful-restart {  
    disable;  
    restart-time seconds;  
    stale-routes-time seconds;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],  
[edit logical-systems logical-system-name protocols bgp group group-name],  
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],  
[edit protocols bgp],  
[edit protocols bgp group group-name],  
[edit protocols bgp group group-name neighbor address]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default. However, helper mode, the ability to assist a neighboring router attempting a graceful restart, is enabled by default.

To configure the duration of the BGP graceful restart period, include the **restart-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the **stale-routes-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.

NOTE: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.

Enable graceful restart mode for BGP (and other protocols) by configuring graceful-restart at the routing-options level. Note that you cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally.

For example, this configuration is required to enable graceful restart:

```
routing-options {  
    graceful-restart  
}
```

If you want to disable graceful restart for some protocols, you can do this at the protocol's graceful-restart command. The following configuration along with the configuration above will keep graceful restart for all protocols but BGP.

```
protocols{  
    bgp{  
        graceful-restart; {  
            disable;  
        }  
    }  
}
```

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

[Configuring Graceful Restart for QFabric Systems](#) | 424

High Availability User Guide

graceful-restart (Protocols OSPF)

Syntax

```
graceful-restart {
  disable;
  helper-disable (standard | restart-signaling | both);
  no-strict-lsa-checking;
  notify-duration seconds;
  restart-duration seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the **no-strict-lsa-checking** statement introduced in Junos OS Release 8.5.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for the helper mode **standard**, **restart-signaling**, and **both** options introduced in Junos OS Release 11.4.

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure graceful restart for OSPF.

Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the **[edit routing-options]** hierarchy level.

Options

disable—Disable graceful restart for OSPF.

helper-disable (standard | restart-signaling | both)—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The **standard**, **restart-signaling**, and **both** options are only supported for OSPFv2. Specify **standard** to disable helper mode for standard graceful restart (based on RFC 3623). Specify **restart-signaling** to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and

RFC 4813). Specify **both** to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.

Default: Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.

no-strict-lsa-checking—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.

NOTE: The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols (ospf | ospf3)** command.

notify-duration seconds—Estimated time needed to send out purged grace LSAs over all the interfaces.

Range: 1 through 3600 seconds

Default: 30 seconds

restart-duration seconds—Estimated time needed to reacquire a full OSPF neighbor from each area.

Range: 1 through 3600 seconds

Default: 180 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Graceful Restart for OSPF

Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart

Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart

Example: Disabling Strict LSA Checking for OSPF Graceful Restart

interconnect-device (Chassis)

Syntax

```
interconnect-device {
  alarm {
    (ethernet (Alarm) | management-ethernet) {
      link-down (red | yellow | ignore);
    }
  }
  container-devices {
    device-count number;
  }
  craft-lockout {
    alarm {
      interface-type {
        link-down (red | yellow | ignore);
      }
    }
    container-devices {
      device-count number;
    }
    fpc slot {
      power (on | off);
    }
    routing-engine {
      on-disk-failure {
        disk-failure-action (halt | reboot);
      }
    }
  }
  fpc slot {
    power (on | off);
  }
  routing-engine {
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
}
```

Hierarchy Level

[edit chassis]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure properties specific to a QFabric system Interconnect device.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Understanding Interconnect Devices](#) | 29

interconnect-device (Aliases)

Syntax

```
interconnect-device interconnect-device-name {
    assigned-interconnect-device-name;
}
```

Hierarchy Level

[edit fabric aliases]

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

(QFabric systems only) Specify the mapping of user-defined names to QFabric system Interconnect devices.

NOTE: The following rules apply to QFabric component alias naming:

- Alias names must use alphabetic (**A** through **Z** and **a** through **z**), numeric (**0** through **9**), or dash (-) characters.
- The maximum length of an alias name is 30 characters.
- Alias names are case sensitive. For example, **MY-NG-1** and **my-ng-1** refer to different components.
- You cannot use the reserved names **all**, **fabric**, or **director-group** as an alias name.

Options

interconnect-device-name—Specify a user-defined name for a QFabric system Interconnect device.

assigned-interconnect-device-name—Specify the Interconnect device identifier or name that has been provided. Identifiers are usually auto-generated by the Interconnect software, and names are usually provided by the administrator of the default partition.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Aliases for the QFabric System](#) | 388

[Understanding Interconnect Devices](#) | 29

interconnect-device (Flow Groups)

Syntax

```
interconnect-device interconnect-device-name;
```

Hierarchy Level

```
[edit fabric flow-groups flow-group-name]
```

Release Information

Statement introduced in Junos OS Release 14.1X53-D15 for the QFX Series.

Description

In a QFabric system, a hash function is used to select an interconnect device to forward traffic between two Node devices. Since this hash function is performed on all available Interconnect devices, it is possible for redundant multicast streams to flow through one Interconnect device, making that Interconnect device a potential single point of failure for the redundant flows. You can enforce the use of specific Interconnect devices by using the QFabric system **flow-groups** statement.

```
[edit fabric]
root@qfabric# set flow-groups test1 interconnect-device ?
Possible completions:
  <interconnect-device-name> Serial ID or alias of interconnect-device
  IC-P6601-C   Interconnect-device
  IC-P6602-C   Interconnect device
  IC-P6603-C   Interconnect device
```

Options

interconnect-device-name—assigns Interconnect devices to the named flow group

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[flow-groups](#) | 479

[Example: Creating a QFabric Flow Group](#) | 431

multicast (QFabric Routing Options)

Syntax

```
multicast {  
    fabric-optimized-distribution  
    no-make-before-break;  
}
```

Hierarchy Level

```
[edit fabric routing-options]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

(QFabric systems only) Set multicast routing options in a QFabric system, such as [no-make-before-break](#).

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding the QFabric System Data Plane](#) | 53

multicast-ipv4

Syntax

```
multicast-ipv4 {
  source-ip source-ip-address;
  dest-ip-multicast-group ipv4-mcast-address;
}
```

Hierarchy Level

```
[edit protocols oam fabric flow-specs flow-specification-name]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure multicast parameters for the IPv4 Ethernet protocol type in the internal fabric monitoring flow specifications.

Options

dest-ip-multicast-group *ipv4-mcast-address*—IPv4 multicast group address of the destination fabric maintenance endpoint (FMEP) that is configured in the flow specification. Configuration of this parameter is mandatory.

Syntax: Dotted decimal notation without the prefix

source-ip *source-ip-address*—(Optional) IPv4 multicast address of the source fabric maintenance endpoint (FMEP) that is configured in the flow specification.

Syntax: Dotted decimal notation without the prefix

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Overview of Internal Fabric Monitoring | 439](#)

[Configuring a Fabric Maintenance Association | 441](#)

[Configuring Flow Specifications | 442](#)

[fabric \(OAM\) | 470](#)

[ping fabric multicast-flow | 591](#)

[ping fabric unicast-flow | 593](#)

[traceroute fabric unicast-flow | 828](#)

multicast-vlan-flood

Syntax

```
multicast-vlan-flood
```

Hierarchy Level

```
[edit protocols oam fabric flow-specs flow-specification-name]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure parameters of the multicast VLAN flood protocol type in the internal fabric monitoring flow specification.

Options

There are no configuration options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Overview of Internal Fabric Monitoring | 439](#)

[Configuring a Fabric Maintenance Association | 441](#)

[Configuring Flow Specifications | 442](#)

[fabric \(OAM\) | 470](#)

[ping fabric multicast-flow | 591](#)

[ping fabric unicast-flow | 593](#)

[traceroute fabric unicast-flow | 828](#)

network-domain

Syntax

```
network-domain;
```

Hierarchy Level

```
[edit fabric resources node-group node-group-name]
```

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Designate a Node group to become a *network Node group*, which is used to route traffic between a QFabric system and external networks. The absence of the **network-domain** configuration statement implies that the Node group is a *server Node group*, which is used to group sets of Node devices that are connected to servers or storage devices.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Node Groups for the QFabric System | 402](#)

[Understanding Node Groups | 38](#)

no-make-before-break

Syntax

```
no-make-before-break;
```

Hierarchy Level

```
[edit fabric routing-options multicast]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

(QFabric systems only) Disable the default *make-before-break* multicast feature on QFabric systems to prevent the duplication of system traffic. This feature increases the speed of data plane traffic, but carries the risk of minor traffic losses when compared with the default make-before-break method of creating new multicast fabric paths before tearing down the old paths. The absence of the **no-make-before-break** configuration statement implies that the make-before-break default system behavior is in effect.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding the QFabric System Data Plane](#) | 53

node-device (Aliases)

Syntax

```
node-device node-device-name {
    assigned-node-device-name;
}
```

Hierarchy Level

[edit fabric aliases]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Specify the mapping of user-defined names to QFabric system Node devices.

NOTE: The following rules apply to QFabric component alias naming:

- Alias names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.
- The maximum length of an alias name is 30 characters.
- Alias names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.
- You cannot use the reserved names **all**, **fabric**, or **director-group** as an alias name.

Options

node-device-name—Specify a user-defined name for a QFabric system Node device.

assigned-node-device-name—Specify the Node device identifier or name that has been provided. Identifiers are usually autogenerated by the Director software, and names are usually provided by the administrator of the default partition.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Aliases for the QFabric System | 388](#)

[Understanding Node Devices | 33](#)

node-device (Chassis)

Syntax

```
node-device name {  
    fibre-channel {  
        port-range {  
            port-range-low port-range-high;  
        }  
    }  
    pic pic-number {  
        fte {  
            port port-number;  
            port-range port-range-low port-range-high;  
        }  
        xe {  
            port port-number;  
            port-range port-range-low port-range-high;  
        }  
    }  
}
```

Hierarchy Level

[edit chassis [node-group](#)]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure properties specific to a Node device in a QFabric system.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Link Aggregation](#)

node-device (Flow Groups)

Syntax

```
node-device (node-device-name | serial-ID | alias );
```

Hierarchy Level

```
[edit fabric flow-groups flow-group-name]
```

Release Information

Statement introduced in Junos OS Release 14.1X53-D15 for the QFX Series.

Description

In a QFabric system, a hash function is used to select an Interconnect device to forward traffic between two Node devices. Since this hash function is performed on all available Interconnect devices, it is possible for redundant multicast streams to flow through one Interconnect device, making that Interconnect device a potential single point of failure for the redundant flows. You can enforce the use of specific Interconnect devices by using the QFabric system **flow-groups** statement.

NOTE: A Node device can belong to only one flow group.

```
root@qfabric# set fabric flow-groups test1 node-device ?
Possible completions:
<node-device-name> Serial ID or alias of node-device
P3686-C          Node device
TA3713480203     Node device
TA3713500165     Node device
TA3713500185     Node device
TA3713500307     Node device
```

Options

node-device-name (or serial-ID or alias)—selects the indicated ingress Node device switch for a flow group

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[flow-groups](#) | 479

[Segregating QFabric Traffic Flows With Flow Groups](#) | 429

[Example: Creating a QFabric Flow Group](#) | 431

[Understanding QFabric Multicast Flow Groups](#) | 55

node-device (Resources)

Syntax

```
node-device node-device-name;
```

Hierarchy Level

```
[edit fabric resources node-group node-group-name]
```

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Assign Node devices to Node groups within a QFabric system.

NOTE: Ensure that the Node devices you assign to a redundant server Node group are of the same type, either two QFX3500 Node devices, two QFX3600 Node devices, or two QFX5100 Node devices. You cannot mix and match different Node device types in the same redundant server Node group.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Node Groups for the QFabric System](#) | 402

[Understanding Node Groups](#) | 38

node-group (Chassis)

Syntax

```
node-group name {
  aggregated-devices {
    ethernet {
      device-count number;
    }
  }
  alarm {
    interface-type {
      link-down (ignore | red | yellow);
    }
  }
  container-devices {
    device-count number;
  }
  node-device name {
    fibre-channel {
      port-range {
        port-range-low port-range-high;
      }
    }
  }
  pic pic-number {
    fte {
      port port-number;
      port-range port-range-low port-range-high;
    }
    xe {
      port port-number;
      port-range port-range-low port-range-high;
    }
  }
  routing-engine {
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
}
```

Hierarchy Level

[edit chassis]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure properties specific to a Node group.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring Link Aggregation*

node-group (Resources)

Syntax

```
node-group node-group-name {
  network-domain;
  node-device node-device-name;
}
```

Hierarchy Level

[edit fabric resources]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Define Node groups within a QFabric system.

NOTE: The following rules apply to QFabric Node group naming:

- The network Node group must use the predefined name *NW-NG-0*. You must use this name when adding Node devices to the network Node group. You cannot specify a different name. Also, you can configure only one network Node group per partition.
- Node group names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.
- The maximum length of a Node group name is 30 characters.
- Node group names are case sensitive. For example, *MY-NG-1* and *my-ng-1* refer to different components.
- You cannot use the reserved names **all**, **fabric**, or **director-group** as a Node group name.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Node Groups for the QFabric System | 402

Understanding Node Groups | 38

pic (Port)

Syntax

```
pic pic-number {
    fte {
        (port port-number | port-range port-range-low port-range-high);
    }
    xe {
        (port port-number | port-range port-range-low port-range-high);
    }
}
```

Hierarchy Level

[edit [chassis interconnect-device](#) *name* fpc slot]

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

(QFX3600 Node device only) Configure port types on QFX3600 Node devices.



CAUTION: The Packet Forwarding Engine on the QFX3600 Node device is restarted when you commit the port type configuration changes. As a result, you might experience packet loss on the Node device.

Options

pic *pic-number*—Number of the physical interface card (PIC) on which you want to configure port types. Specify **0** to configure **xe** (10-Gigabit Ethernet) type ports. Specify **1** to configure **fte** (40-gigabit data plane uplink) type ports.

fte—Configure a specific port or a range of ports to operate as 40-gigabit data plane uplink ports.

xe—Configure a specific port or a range of ports to operate as four 10-Gigabit Ethernet ports.

port-number—Port number on which you want to configure the port type. Valid values are 0 through 7 if the port type is **fte**, and 2 through 15 if the port type is **xe**.

port-range-low—Lowest-numbered port in the range of ports. The lowest possible value is 0 if the port type is **fte**. The lowest possible value is 2 if the port type is **xe**.

port-range-high—Highest-numbered port in the range of ports. The highest possible value is 7 if the port type is **fte**. The highest possible value is 15 if the port type is **xe**.

NOTE:

- By default, ports Q0 through Q3 operate as fte type ports, and ports Q4 through Q15 operate as xe type ports.
- Only ports Q0 through Q7 can be configured as fte type ports.
- Only ports Q2 through Q15 can be configured as xe type ports.

NOTE: When you delete the port type configuration for an individual port or a block of ports, the ports return to operating in their default port type.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Port Type on QFX3600 Node Devices](#) | 408

protocols (Fabric)

Syntax

```
protocols {  
  fabric-control {  
    graceful-restart {  
      restart-time seconds;  
      stale-routes-time seconds;  
    }  
  }  
}
```

Hierarchy Level

[edit fabric]

Release Information

Statement introduced in Junos OS Release 13.2X52-D10 for the QFX Series.

Description

Specify attributes for the fabric control protocol.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Routing Engines in the QFabric System](#) | 27

resources

Syntax

```
resources {  
  node-group node-group-name {  
    network-domain;  
    node-device node-device-name;  
  }  
}
```

Hierarchy Level

```
[edit fabric]
```

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Define resources within a QFabric system and handle a variety of component assignments.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Node Groups for the QFabric System | 402](#)

[Understanding Node Groups | 38](#)

restart-time (Fabric Control)

Syntax

```
restart-time seconds;
```

Hierarchy Level

```
[edit fabric protocols fabric-control graceful-restart]
```

Release Information

Statement introduced in Junos OS Release 13.2X52-D10 for the QFX Series.

Description

Configure the duration of the graceful restart period for the fabric control Routing Engine.

The graceful restart resynchronization process takes longer when the QFabric system contains node groups that have a large number of VLANs. The graceful-restart duration should, therefore, be set higher when the QFabric system contains at least one node group with a large number of VLANs.

Configure a restart time of 600 seconds if the number of VLAN members (vmembers) exceeds 32k.



CAUTION: Configuring the restart time restarts the session between the fabric control Routing Engine and the Node groups. Traffic is dropped as a result of this restart. Normal QFabric system operations should resume once the session has restarted without any further user actions.

Options

seconds—Duration of the graceful restart period.

Default: 300 seconds

Range: 300 to 900 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Bridging and VLANs on Switches

routing-options (QFabric System)

Syntax

```
routing-options {  
  multicast {  
    no-make-before-break;  
  }  
}
```

Hierarchy Level

```
[edit fabric]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

(QFabric systems only) Set routing options in a QFabric system.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

syslog (QFabric System)

Syntax

```
syslog {  
  file filename {  
    archive {  
      size maximum-file-size;  
    }  
    explicit-priority;  
    facility severity;  
    match "regular-expression";  
    structured-data;  
  }  
  filter all {  
    facility severity;  
    match "regular-expression";  
  }  
  host hostname {  
    explicit-priority;  
    facility severity;  
    facility-override facility;  
    log-prefix string;  
    match "regular-expression";  
    structured-data;  
  }  
}
```

Hierarchy Level

[edit system]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure system log messages for the QFabric system.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding the Implementation of System Log Messages on the QFabric System | 72](#)

[Directing System Log Messages to a Remote Machine](#)

stale-routes-time (Fabric Control)

Syntax

```
stale-routes-time seconds;
```

Hierarchy Level

```
[edit fabric protocols fabric-control graceful-restart]
```

Release Information

Statement introduced in Junos OS Release 13.2X52-D10 for the QFX Series.

Description

Set the length of time that the fabric control Routing Engine waits to receive messages from devices before declaring them down. Configure a stale routes time of 1800 seconds if the number of VLAN members (vmembers) exceeds 32k.

Options

seconds—Amount of time that the fabric control Routing Engine waits to receive messages from other devices before declaring them down.

Default: 900 seconds

Range: 900 to 1800 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Bridging and VLANs on Switches](#)

[Understanding Routing Engines in the QFabric System | 27](#)

unicast-ethernet

Syntax

```
unicast-ethernet {
  destination-mac destination-mac-address destination-mac-mask destination-mac-mask;
  ethertype ethertype;
  source-mac source-mac-address source-mac-mask source-mac-mask;
}
```

Hierarchy Level

```
[edit protocols oam fabric flow-specs flow-specification-name]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure unicast Ethernet parameters for a flow specification used in internal fabric monitoring.

Options

destination-mac *destination-mac-address*—(Optional) MAC address of the destination fabric maintenance endpoint (FMEP) that is configured in the flow specification.

destination-mac-mask *destination-mac-mask*—(Optional) MAC address range mask for the destination FMEP that is configured in the flow specification. You cannot configure the **destination-mac-mask** parameter unless you also configure the **destination-mac** parameter.

The masked bits denote the range that needs to be generated. The other bits obtain their values from the corresponding source or destination MAC address parameter. The mask must be contiguous.

Values: Maximum of 6 contiguous masked bits, or a mask range of 64

ethertype *ethertype*—Valid Ethernet type other than the IPv4 Ethertype.

Values: 2 bytes in decimal or hexadecimal notation

source-mac *source-mac-address*—(Optional) MAC address of the source FMEP or interface that is configured in the flow specification.

source-mac-mask *source-mac-mask*—(Optional) MAC address range mask for the source FMEP that is configured in the flow specification. You cannot configure the **source-mac-mask** parameter unless you also configure the **source-mac** parameter.

The masked bits denote the range that needs to be generated. The other bits obtain their values from the corresponding source or destination MAC address parameter. The mask must be contiguous.

Values: Maximum of 6 contiguous masked bits, or a mask range of 64

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring 439
Configuring a Fabric Maintenance Association 441
Configuring Flow Specifications 442
fabric (OAM) 470
ping fabric multicast-flow 591
ping fabric unicast-flow 593
traceroute fabric unicast-flow 828

unicast-ethernet-ipv4

Syntax

```
unicast-ethernet-ipv4 {
  source-ip source-ip-address source-ip-mask source-ip-mask;
  destination-ip destination-ip-address destination-ip-mask destination-ip-mask;
  ip-proto protocol;
  source-l4-port source-l4-port-number;
  destination-l4-port destination-l4-port-number;
}
```

Hierarchy Level

```
[edit protocols oam fabric flow-specs flow-specification-name]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure unicast parameters for the Ethernet IPv4 protocol type in the internal fabric monitoring flow specifications.

Options

destination-ip *destination-ip-address*—Unicast IPv4 address of the destination fabric maintenance endpoint (FMEP) that is configured in the fabric maintenance association (FMA).

Syntax: Dotted decimal notation without the prefix

destination-l4-port *destination-l4-port-number*—L4 TCP or UDP port for the destination fabric maintenance endpoint (FMEP).

Values: 2-byte decimal or hexadecimal notation

ip-proto *protocol*—IPv4 protocol type.

Values: 1-byte value in decimal or hexadecimal notation

destination-ip-mask *destination-ip-mask*—IPv4 address range mask for the destination FMEP that is configured in the flow specification.

The masked bits denote the range that needs to be generated. The other bits obtain their values from the corresponding source or destination IPv4 address parameter. The mask must be contiguous.

Syntax: Dotted decimal notation without the prefix

Values: Maximum of 6 contiguous masked bits, or a mask range of 64

source-ip *source-ip-address*—Unicast IPv4 address of the source fabric maintenance endpoint (FMEP) that is configured in the fabric maintenance association (FMA).

Syntax: Dotted decimal notation without the prefix

source-l4-port *source-l4-port-number*—TCP or UDP port for the source fabric maintenance endpoint (FMEP).

Values: 2-byte decimal or hexadecimal notation

source-ip-mask *source-ip-mask*—IPv4 address range mask for the source FMEP that is configured in the flow specification.

The masked bits denote the range that needs to be generated. The other bits obtain their values from the corresponding source or destination IPv4 address parameter. The mask must be contiguous.

Syntax: Dotted decimal notation without the prefix

Values: Maximum of 6 contiguous masked bits, or a mask range of 64

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring 439
Configuring a Fabric Maintenance Association 441
Configuring Flow Specifications 442
fabric (OAM) 470
ping fabric multicast-flow 591
ping fabric unicast-flow 593
traceroute fabric unicast-flow 828

4

PART

Administration

Software Upgrade, Downgrade, and Recovery | **529**

QFabric System Administration and Maintenance | **583**

Operational Mode Commands | **589**

Software Upgrade, Downgrade, and Recovery

IN THIS CHAPTER

- Nonstop Software Upgrade Checklist for QFabric Systems | 529
- Performing a Nonstop Software Upgrade on the QFabric System | 533
- Verifying Nonstop Software Upgrade for QFabric Systems | 540
- Upgrading Software on a QFabric System | 566
- Downgrading Software on a QFabric System | 572
- Performing System Backup and Recovery for a QFabric System | 581

Nonstop Software Upgrade Checklist for QFabric Systems

IN THIS SECTION

- Preparing for an NSSU | 529
- Upgrading Director Groups | 531
- Upgrading QFabric | 532
- Upgrading Network Node Groups | 532

This topic lists the precautions that must be taken before you begin a nonstop software upgrade (NSSU) on QFabric systems. This topic includes:

Preparing for an NSSU

NOTE: Before you perform a nonstop software upgrade, contact JTAC to perform a pre-upgrade health check on the QFabric system.

In preparation for an NSSU, ensure the following:

- The NSSU is being done during a maintenance window.
- Console connectivity is available for both Director groups.
- The existing configuration has been backed up.

Next, complete the following steps:

1. Collect the output of the **lsvm** and **dns.dump** commands. **lsvm** command provides the universal unique identifiers (UUIDs) of all the virtual machines, and the **dns.dump** provides the IP addresses of all the components that can be used in case of failure.
2. Run the following commands from the Director device shell to perform sanity checks of the entire QFabric system:

```
/opt/DCF/scripts/director_upgrade_sanity_check.pl
/opt/DCF/scripts/fabric_upgrade_sanity_check.pl
/opt/DCF/scripts/nodegroup_upgrade_sanity_check.pl
```

NOTE: In case of any errors, contact Juniper Networks Technical Assistance Center (JTAC).

3. Enter the **request component login DRE-0** command to connect to DRE-0. Check the available storage using the **show system storage** command. There should be about 500 MB of space available in the /var partition. If not, delete the files under /var/log. Also, check the /var/log/kdm and /var/tmp directories for any large files that can be deleted. Issue the **clear log filename-with-path** command to clear the large files in the /var/log/ and in /var/tmp/ directories.
4. Modify the multicast discovery advertisements (MDA) timer on all Fibre Channel (FC) and Fibre Channel over Ethernet (FCoE) gateways connected to QFabric, if applicable. We recommend this step to prevent FC session drops during NSSU. Enter the following command on FC and FCoE gateways: **set fc-fabrics fabric-name protocols fip fka-adv-period 45000**.
5. Enter the following CLI commands to update the system log configuration. This will reduce the amount of system log messages added to the database during upgrade. Logging to an external server does not have any impact during upgrade.
 - **set system syslog host syslog server address**
 - For remote servers:


```
set system syslog file cli_commands interactive-commands any
```


set system syslog file *for all files configured* any critical

set system syslog file messages any critical

Upgrading Director Groups

The NSSU process individually upgrades members of a Director group, so that one device in the group is always operational. It switches the mastership of the Routing Engine processes to the backup Director device, before upgrading the master Director device.

Before starting the NSSU for the Director group:

- Run the sanity check scripts from the Director device shell to ensure that the Director device is ready to be upgraded.

```
root@dg0 ~]# /opt/dcf/scripts/director_upgrade_sanity_check.pl
Checking if director-device upgrade is currently in progress.
Checking VM status.
Checking for communication between director devices.
Checking inventory status of VMs.
Checking VM versions.
Checking VM passwords.
Checking chassis alarms.
```

- Verify whether both the Director devices are online.

```
root@Qfabric_device> show fabric administration inventory director-group status
```

Enter the following CLI commands to clean the storage and system log.

```
request system storage cleanup director-group
request system storage cleanup infrastructure FM-0
request system storage cleanup qfabric component all
clear log messages
```

After these processes are complete, you can run the NSSU on the master Director group. After the upgrade is complete, verify the process using the following CLI commands:

```
show version component director-group
show version component DRE-0
show fabric admin inventory
show fabric admin inventory infrastructure
show fabric admin inventory director-group status
```


Restore the system log configuration that you updated in step 5.

Upgrading QFabric

Before starting the NSSU on QFabric:

- Run the sanity check scripts from the Director device shell to ensure that the Director device is ready to be upgraded.

```
root@dg0 ~]# /opt/def/scripts/fabric_upgrade_sanity_check.pl
Checking if director-device upgrade is currently in progress.
Checking VM status.
Checking for communication between director devices.
Checking inventory status of ICs/FCs.
Checking FC passwords.
Checking IC passwords.
Checking chassis alarms.
The system is not ready for upgrade, please check the following:
Check password for P8959-C/RE0 cannot complete. Node P8959-C/RE0 not reachable.
```

Enter the following CLI commands to clean the storage.

```
request system storage cleanup infrastructure FC-0
request system storage cleanup infrastructure FC-1
request system storage cleanup interconnect-device ic-name
```

After these processes are complete, you can run the NSSU on QFabric. After the upgrade is complete, verify the process using the following CLI commands:

```
show version component fabric
show fabric admin inventory
show fabric admin inventory infrastructure
show chassis fabric connectivity
```

Upgrading Network Node Groups

Before starting the NSSU on the network Node groups:

- Run the sanity check scripts from the Director device shell to ensure that the Node groups are ready to be upgraded.


```

root@dg0 ~]# /opt/dcf/scripts/nodegroup_upgrade_sanity_check.pl
Checking if director-device upgrade is currently in progress.
Checking VM status.
Checking for communication between director devices.
Checking inventory status of NINEs/SINEs/RNGs.
Checking Server INE passwords.
Checking NW-NG-0 passwords.
Checking chassis alarms.

```

Enter the following CLI commands to clean the storage.

```
request system storage cleanup node-group NW-NG-0
```

After these processes are complete, you can run the NSSU on the network Node groups. After the upgrade is complete, verify the process using the following CLI commands:

```

show version component NW-NG-0
show fabric admin inventory NW-NG-0
show fabric admin inventory infrastructure
show chassis fabric connectivity

```

NOTE: Follow the same procedure for redundant server Node groups (RSNGs) and server Node groups (SNGs).

RELATED DOCUMENTATION

[Understanding Nonstop Software Upgrade for QFabric Systems | 61](#)

[Performing a Nonstop Software Upgrade on the QFabric System | 533](#)

Performing a Nonstop Software Upgrade on the QFabric System

NOTE: Before you perform a nonstop software upgrade, contact JTAC to perform a pre-upgrade health check on the QFabric system.

NOTE: Before you can perform a nonstop software upgrade to Junos OS Release 13.1X50-D10, you must have Junos OS Release 12.2X50-D42 or later installed. You cannot perform a nonstop software upgrade with Junos OS Release 12.2X50-D41 or earlier. Contact the Juniper Technical Assistance Center for information on how to download Junos OS Release 12.2X50-D42. Performing a standard software upgrade (that is, issuing the **request system software add component all** command) does not require that you that you upgrade to an intermediate Junos OS software release.

To perform a nonstop software upgrade to Junos OS Release 13.1X50-D10:

1. First perform a nonstop software upgrade to Junos OS Release 12.2X50-D42.
2. Then perform a nonstop software upgrade to Junos OS Release 13.1X50-D10.

Nonstop software upgrade enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. This feature introduces several high availability improvements to the QFabric system software upgrade process, including:

- Upgrading members of a Director group or Node group one at a time so that one device in the group is always operational
- Switching mastership of Routing Engine processes to the backup Director device before upgrading the master Director device
- Rebooting Interconnect devices and fabric control Routing Engines one at a time, so that one Interconnect device or one fabric control Routing Engine is always operational
- Switching mastership of a Node group to the backup Node device before upgrading the master Node device
- Specifying an upgrade group if you want all Node devices in a Node group to be upgraded in parallel (which shortens the time of the upgrade)
- Rebooting devices automatically as part of the nonstop upgrade process

When performing a nonstop upgrade, start with the Director group upgrade, then issue the fabric upgrade, and end with the Node group upgrades.

NOTE: Because there is no redundancy for Node groups containing a single Node device, traffic loss occurs when the device reboots during the upgrade. For node-groups defined with two node-devices, both must be online in order for upgrade to succeed.

NOTE: Before you install the software, we recommend that you back up your current configuration files by issuing the **request system software configuration-backup** command.

NOTE: Before you can perform a nonstop software upgrade in your QFabric system, you must first upgrade your system to Junos OS Release 12.2 by using a conventional upgrade method such as issuing the **request system software add component all** command.

This topic describes the following tasks:

- [Backing Up the Current Configuration Files | 535](#)
- [Downloading Software Files Using a Browser | 536](#)
- [Retrieving Software Files for Download | 537](#)
- [Performing a Nonstop Software Upgrade for Director Devices in a Director Group | 537](#)
- [Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components | 537](#)
- [\(Optional\) Creating Upgrade Groups for Node Groups | 538](#)
- [Performing a Nonstop Software Upgrade on a Node Group | 539](#)

Backing Up the Current Configuration Files

To back up your current configuration files:

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

Back up the configuration files to a local directory, remote server, or removable drive (for example, an external USB flash drive).

For example:

```
user@qfabric> request system software configuration-backup/media/USB/
```


Downloading Software Files Using a Browser

NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

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 package you want to download in the **QFabric System Install Package** section:
 - If you want to upgrade the entire QFabric system, select **QFabric System - Complete Install Package**.
 - If you want to upgrade either a single Node or Interconnect device for recovery purposes, select **Node and Interconnect Device Install Package**. For information on how to perform a recovery installation on either a Node or Interconnect device, see [“Performing a Recovery Installation” on page 845](#).

A login screen appears.
7. Enter your user ID and password and press **Enter**.
8. Read the End User License Agreement, select the **I agree** option button, and then click **Proceed**.
9. Save the **jinstall-qfabric-version.rpm** file on your computer.

Retrieving Software Files for Download

Retrieve the software from the location in which you downloaded it. To do this, issue the **request system software download** command. The software package is copied from where you downloaded it and is placed locally on the QFabric system.

- To retrieve the software:

```
user@qfabric> request system software download /path/package-name
```

For example:

```
user@qfabric> request system software download
ftp://server/files/jinstall-qfabric-12.2X50-D10.3.rpm
```

Performing a Nonstop Software Upgrade for Director Devices in a Director Group

NOTE: If you reboot any Node groups or Interconnect devices after you perform a nonstop upgrade on the Director group, these devices are upgraded to the same version of software that is running on the Director group.

To upgrade the software on the Director devices in a Director group:

- Issue the **request system software nonstop-upgrade director-group package-name** command.

For example:

```
user@qfabric> request system software nonstop-upgrade director-group
jinstall-qfabric-12.2X50-D10.3.rpm
```

When the Director group is rebooting the first Director device during the upgrade and transferring mastership to the peer director device, you might not be able to log in and access the QFabric system CLI for several minutes while user interface services come back online.

Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components

Before you perform a nonstop upgrade on the Interconnect devices and other fabric-related components, verify that both Director devices in the Director group are online. Both Director devices must be online before you attempt to perform a nonstop upgrade. To do verify that both Director devices are online, issue the **show fabric administration inventory director-group status** command.

To install the software on the Interconnect device and other components in the fabric:

- Issue the **request system software nonstop-upgrade fabric *package-name*** command.

For example:

```
user@qfabric> request system software nonstop-upgrade fabric jinstall-qfabric-12.2X50-D10.3.rpm
```

(Optional) Creating Upgrade Groups for Node Groups

Upgrade groups enable two or more Node devices in a Node group, or an entire Node group, to be rebooted at the same time. If you do not create an upgrade group, the Node devices are upgraded one at a time. Before performing a nonstop upgrade on a Node group, create an upgrade group and include the devices you want to reboot at the same time.

NOTE: If you add Node devices that have links to the same link aggregation group (LAG), there might be traffic loss.

- Create the upgrade group by issuing the **set chassis node-group *node-group-name* nssu upgrade-group *upgrade-group-name* node-devices** command at the [edit chassis] hierarchy.

For example:

```
user@qfabric# set chassis node-group nodegroup1 nssu upgrade-group upgrade1 node-devices [
node1 node2 ]
```


Performing a Nonstop Software Upgrade on a Node Group

When you perform a nonstop software upgrade on a network Node group, the Node devices in the network Node group are upgraded in a serial fashion except when upgrade groups are configured. If you perform a nonstop upgrade on a redundant server Node group, both Node devices must be online for a successful upgrade. If one of the Node devices is no longer available, remove it from the configuration before you perform the nonstop software upgrade. If you perform a nonstop upgrade on a Node group with only one Node device, traffic loss occurs while the Node device is rebooting.

NOTE: You can upgrade multiple Node groups with this command. However, if more than one Node group is specified, there may be traffic loss depending on the topology of the network.

To install software on a Node group:

- Issue the **request system software nonstop-upgrade node-group *node-group-name* *package-name*** command.

To perform a nonstop upgrade on one Node group:

```
user@qfabric> request system software nonstop-upgrade node-group nodegroup1
jinstall-qfabric-12.2X50-D10.3.rpm
```

To perform a nonstop upgrade on more than one Node group:

```
user@qfabric> request system software nonstop-upgrade node-group [nodegroup1 nodegroup2
nodegroup3] jinstall-qfabric-12.2X50-D10.3.rpm
```

RELATED DOCUMENTATION

[Nonstop Software Upgrade Checklist for QFabric Systems | 529](#)

[Configuring Graceful Restart for QFabric Systems | 424](#)

[Understanding Nonstop Software Upgrade for QFabric Systems | 61](#)

Verifying Nonstop Software Upgrade for QFabric Systems

IN THIS SECTION

- [Verifying a Director Group Nonstop Software Upgrade | 540](#)
- [Verifying a Fabric Nonstop Software Upgrade | 558](#)
- [Verifying a Redundant Server Node Group Nonstop Software Upgrade | 560](#)
- [Verifying a Network Node Group Nonstop Software Upgrade | 564](#)

NOTE: Before you perform a nonstop software upgrade, contact JTAC to perform a pre-upgrade health check on the QFabric system.

This topic discusses how you can monitor the progress of each of the three steps in a nonstop software upgrade. By identifying the key actions and events that define this process, you can track the status of the upgrade with confidence.

TIP: When performing a nonstop software upgrade, open two SSH sessions to the QFabric CLI. Use one session to monitor the upgrade itself and use a second session to verify that the QFabric system components respond to operational mode commands as expected.

Verifying a Director Group Nonstop Software Upgrade

Purpose

During the Director group portion of a nonstop software upgrade, you should expect to see the Director device that hosts the CLI session selected as the master device. When mastership of all processes moves to the master, the QFabric system upgrades the backup Director device and this Director device reboots. After the backup Director device comes back online, the master Director device suspends CLI operations for 15 minutes, upgrades itself, and reboots. At this point, the backup becomes the new master Director device and you can issue CLI operational commands. Finally, the former master comes back online as a backup and both devices are operational once again. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

Action


```

-----
Shared File System           online
Network File System          online
Virtual Machine Server       online
Load Balancer/DHCP           online

Hard Drive Status
-----
Volume ID:4                  optimal
Physical ID:1                online
Physical ID:0                online
SCSI ID:1                    100%
SCSI ID:0                    100%

Size  Used Avail Used% Mounted on
-----
423G  5.4G 395G   2%   /
99M   16M  79M   17% /boot
93G   7.3G 86G    8%  /pbdata

Director Group Processes
-----
Director Group Manager       online
Partition Manager            online
Software Mirroring            online
Shared File System master    online
Secure Shell Process          online
Network File System           online
DHCP Server master           online      master
FTP Server                    online
Syslog                        online
Distributed Management         online
SNMP Trap Forwarder           online
SNMP Process                  online
Platform Management           online

Interface Link Status
-----
Management Interface         up
Control Plane Bridge          up
Control Plane LAG             up
CP Link [0/2]                 up
CP Link [0/1]                 up

```



```

CP Link [0/0]          up
CP Link [1/2]          down
CP Link [1/1]          down
CP Link [1/0]          down
Crossover LAG          up
CP Link [0/3]          up
CP Link [1/3]          up

```

```

Member Device Id/Alias  Status  Role
-----
dgl      0281052011000032 online  backup

```

Director Group Managed Services

```

Shared File System      online
Network File System     online
Virtual Machine Server  online
Load Balancer/DHCP      online

```

Hard Drive Status

```

Volume ID:8             optimal
Physical ID:1           online
Physical ID:0           online
SCSI ID:1               100%
SCSI ID:0               100%

```

```

Size  Used Avail Used% Mounted on
----  -
423G  5.5G 395G   2%   /
99M   16M  79M   17%  /boot
93G   7.3G 86G    8%   /pbdata

```

Director Group Processes

```

Director Group Manager      online
Partition Manager           online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         online
DHCP Server master          online    backup
FTP Server                   online

```



```

Syslog                online
Distributed Management online
SNMP Trap Forwarder   online
SNMP Process          online
Platform Management   online

```

Interface Link Status

```
-----
```

```

Management Interface    up
Control Plane Bridge     up
Control Plane LAG        up
CP Link [0/2]            up
CP Link [0/1]            up
CP Link [0/0]            up
CP Link [1/2]            down
CP Link [1/1]            down
CP Link [1/0]            down
Crossover LAG           up
CP Link [0/3]            up
CP Link [1/3]            up

```

session1@qfabric> show fabric session-host

```
Identifier: 0281052011000032
```

2. In a second SSH session to the QFabric CLI, issue the request for the Director group nonstop software upgrade.

```

root@qfabric> request system software nonstop-upgrade director-group
jinstall-qfabric-12.2X50-D10.3.rpm

```

3. If the CLI session is being hosted by the master Director device, skip to step 4. However, if the CLI session is hosted by the backup Director device, the Director group mastership switches to the backup device after you issue the nonstop software upgrade command. In this example, mastership switches to Director device DG1.

session1@qfabric> show fabric administration inventory director-group status

```
Director Group Status Tue Jun  5 15:12:20 UTC 2012
```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	backup	10.49.215.38	8%	31905924k	0	3 days, 21:16 hrs


```

dg1    online master    10.49.215.39    6%  18010368k    3    3 days, 21:16 hrs

```

```

Member Device Id/Alias  Status  Role
-----

```

```

dg0    0281052011000001 online  backup

```

```

Director Group Managed Services
-----

```

```

Shared File System      offline
Network File System     offline
Virtual Machine Server  offline
Load Balancer/DHCP      offline

```

```

Hard Drive Status
-----

```

```

Volume ID:4             optimal
Physical ID:1           online
Physical ID:0           online
SCSI ID:1               100%
SCSI ID:0               100%

```

```

Size  Used Avail Used% Mounted on
-----

```

```

423G 5.4G 395G  2%  /
99M  16M  79M   17% /boot

```

```

Director Group Processes
-----

```

```

Director Group Manager      online
Partition Manager           online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         offline
DHCP Server master          offline  backup
FTP Server                  online
Syslog                      online
Distributed Management       offline
SNMP Trap Forwarder         offline
SNMP Process                offline
Platform Management         online

```

```

Interface Link Status

```



```

-----
Management Interface          up
Control Plane Bridge          up
Control Plane LAG             up
CP Link [0/2]                 up
CP Link [0/1]                 up
CP Link [0/0]                 up
CP Link [1/2]                 down
CP Link [1/1]                 down
CP Link [1/0]                 down
Crossover LAG                 up
CP Link [0/3]                 up
CP Link [1/3]                 up

Member Device Id/Alias  Status  Role
-----
dgl      0281052011000032 online  master

Master Services
-----
Database Server          online
Load Balancer Director   online
QFabric Partition Address online

Director Group Managed Services
-----
Shared File System       online
Network File System      online
Virtual Machine Server   online
Load Balancer/DHCP       online

Hard Drive Status
-----
Volume ID:8              optimal
Physical ID:1            online
Physical ID:0            online
SCSI ID:1                100%
SCSI ID:0                100%

Size  Used Avail Used% Mounted on
-----
423G  6.0G 395G   2%   /
99M   16M  79M   17%  /boot

```



```

93G  7.3G 86G   8%  /pbdata

Director Group Processes
-----
Director Group Manager      online
Partition Manager          online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         online
DHCP Server master          online      master
FTP Server                  online
Syslog                     online
Distributed Management      online
SNMP Trap Forwarder         online
SNMP Process                online
Platform Management         online

Interface Link Status
-----
Management Interface        up
Control Plane Bridge        up
Control Plane LAG           up
CP Link [0/2]               up
CP Link [0/1]               up
CP Link [0/0]               up
CP Link [1/2]               down
CP Link [1/1]               down
CP Link [1/0]               down
Crossover LAG               up
CP Link [0/3]               up
CP Link [1/3]               up

```

session1@qfabric> show fabric session-host

```
Identifier: 0281052011000032
```

4. The Director group nonstop software upgrade process continues by downloading and installing software for the fabric manager Routing Engines and the Director devices.

root@qfabric>

```
Validating update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm
```



```

upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to complete phase one of rolling upgrade
[Peer Update Status]: Peer completed phase one of rolling upgrade

```

5. When the system upgrades and reboots the backup Director device DG0, notice how this device is not displayed in the output of the **show fabric administration inventory director-group status** command. Because Director device DG1 appears, this means that the DG1 is operational and acts as the master device.

NOTE: If your second SSH session is being hosted by the rebooting Director device, your session terminates and you need to log back in to establish a new session running on the active Director device.

session1@qfabric> show fabric administration inventory director-group status

```

Director Group Status Tue Jun  5 15:41:14 UTC 2012

Member Status Role      Mgmt Address      CPU Free Memory VMs Up Time
-----
dg1      online master    10.49.215.39      6%  8372272k       4   3 days, 21:25 hrs

```



```

Member Device Id/Alias  Status  Role
-----
dg1      0281052011000032 online  master

Master Services
-----
Database Server          online
Load Balancer Director   online
QFabric Partition Address online

Director Group Managed Services
-----
Shared File System       online
Network File System      online
Virtual Machine Server   online
Load Balancer/DHCP       online

Hard Drive Status
-----
Volume ID:8              optimal
Physical ID:1            online
Physical ID:0            online
SCSI ID:1                100%
SCSI ID:0                100%

Size  Used Avail Used% Mounted on
----  -
423G  6.0G 395G   2%   /
99M   16M  79M   17%  /boot
93G   7.3G 86G    8%  /pbdata

Director Group Processes
-----
Director Group Manager   online
Partition Manager        online
Software Mirroring        online
Shared File System master online
Secure Shell Process     online
Network File System      online
DHCP Server master       online    master
FTP Server               online
Syslog                   online
Distributed Management    online

```



```

SNMP Trap Forwarder      online
SNMP Process             online
Platform Management      online

```

Interface Link Status

```
-----
```

```

Management Interface      up
Control Plane Bridge      up
Control Plane LAG         up
CP Link [0/2]             up
CP Link [0/1]             up
CP Link [0/0]             up
CP Link [1/2]             down
CP Link [1/1]             down
CP Link [1/0]             down
Crossover LAG            up
CP Link [0/3]             up
CP Link [1/3]             up

```

6. The upgrade continues with master Director device DG1 suspending CLI services for 15 minutes, transferring mastership to Director device DG0, and then rebooting Director device DG1 (which terminates the CLI session).

root@qfabric>

```

[Peer Update Status]: Setting peer DG node as the master SFC

Delaying start of local upgrade to allow peer services time to initialize [15
minutes]
Delaying start of local upgrade to allow peer services time to initialize [15
minutes]
Delaying start of local upgrade to allow peer services time to initialize [12
minutes]
Delaying start of local upgrade to allow peer services time to initialize [9
minutes]
Delaying start of local upgrade to allow peer services time to initialize [6
minutes]
Delaying start of local upgrade to allow peer services time to initialize [3
minutes]
[Peer Update Status]: Check for VMs on dg0
Triggering Final Stage of Fabric Manager Upgrade:
Updating FM-0 to Junos version 12.2X50-D10.3
[Status 2012-06-05 16:10:12]: Fabric Manager: Upgrade Final Stage started
[NW-NG-0 2012-06-05 16:10:22]: Transferring NW-NG-0 Mastership to REMOTE DG

```



```

[NW-NG-0 2012-06-05 16:11:44]: Finished NW-NG-0 Mastership switch
[Status 2012-06-05 16:11:45]: Upgrading FM-0 VM on worker DG to 12.2X50-D10.3
[DRE-0 2012-06-05 16:12:43]: Retrieving package
[DRE-0 2012-06-05 16:13:46]: ----- re0: -----
[Status 2012-06-05 16:15:17]: Load completed with 0 errors...
[Status 2012-06-05 16:15:17]: Reboot is required to complete upgrade ...
[DRE-0 2012-06-05 16:15:22]: Waiting for DRE-0 to terminate ...
[DRE-0 2012-06-05 16:15:34]: Waiting for DRE-0 to come back ...
[DRE-0 2012-06-05 16:18:44]: Running Uptime Test for DRE-0
[DRE-0 2012-06-05 16:18:51]: Uptime Test for DRE-0 Passed ...
[Status 2012-06-05 16:18:51]: DRE-0 booted successfully ...
Performing post install shutdown and cleanup

Broadcast message from root (Tue Jun 5 16:18:51 2012):

The system is going down for reboot NOW!
Director group upgrade complete

root@qfabric> Read from remote host qfabric-partition0: Connection reset by peer
Connection to qfabric-partition0 closed.

```

7. Upon reopening the SSH session, notice that Director device DG0 is now the master device hosting the session and Director device DG1 does not appear in the QFabric system inventory while it is rebooting.

session1@qfabric> show fabric session-host

```
Identifier: 0281052011000001
```

session1@qfabric> show fabric administration inventory director-group status

```

Director Group Status Tue Jun 5 16:21:23 UTC 2012

Member Status Role      Mgmt Address      CPU Free Memory VMs Up Time
-----
dg0      online master    10.49.215.38      13% 20739560k      3   36:29 mins

Member Device Id/Alias  Status  Role
-----
dg0      0281052011000001 online  master

Master Services
-----

```



```

Database Server          online
Load Balancer Director   online
QFabric Partition Address online

```

Director Group Managed Services

```
-----
```

```

Shared File System       online
Network File System      online
Virtual Machine Server   online
Load Balancer/DHCP       online

```

Hard Drive Status

```
-----
```

```

Volume ID:4              optimal
Physical ID:1            online
Physical ID:0            online
SCSI ID:1                100%
SCSI ID:0                100%

```

Size Used Avail Used% Mounted on

```
-----
```

```

423G 5.3G 396G 2% /
99M 16M 79M 17% /boot
93G 7.4G 86G 8% /pbdata

```

Director Group Processes

```
-----
```

```

Director Group Manager   online
Partition Manager        online
Software Mirroring       online
Shared File System master online
Secure Shell Process     online
Network File System      online
DHCP Server master       online    master
FTP Server               online
Syslog                   online
Distributed Management   online
SNMP Trap Forwarder      online
SNMP Process             online
Platform Management      online

```

Interface Link Status

```
-----
```



```

Management Interface      up
Control Plane Bridge      up
Control Plane LAG         up
CP Link [0/2]             up
CP Link [0/1]             up
CP Link [0/0]             up
CP Link [1/2]             down
CP Link [1/1]             down
CP Link [1/0]             down
Crossover LAG             up
CP Link [0/3]             up
CP Link [1/3]             up

```

8. When Director device DG1 comes back online, it returns to the QFabric system inventory as a backup Director device and hosts some of the Routing Engine processes (which should appear load balanced between the master and backup Director devices).

session1@qfabric> show fabric administration inventory director-group status

```

root@qfabric> show fabric administration inventory director-group status
Director Group Status Tue Jun  5 16:41:02 UTC 2012

```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	15%	14759920k	6	56:09 mins
dg1	online	backup	10.49.215.39	8%	31486680k	0	07:51 mins

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:4	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size Used Avail Used% Mounted on

423G	5.3G	396G	2%	/
99M	16M	79M	17%	/boot
93G	7.4G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online	
Partition Manager	online	
Software Mirroring	online	
Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status

Management Interface	up
Control Plane Bridge	up
Control Plane LAG	up
CP Link [0/2]	up
CP Link [0/1]	up
CP Link [0/0]	up
CP Link [1/2]	down
CP Link [1/1]	down
CP Link [1/0]	down
Crossover LAG	up


```

CP Link [0/3]                up
CP Link [1/3]                up

```

```

Member Device Id/Alias  Status  Role
-----
dg1      0281052011000032 online  backup

```

Director Group Managed Services

```

-----
Shared File System          online
Network File System         online
Virtual Machine Server      online
Load Balancer/DHCP          online

```

Hard Drive Status

```

-----
Volume ID:8                 optimal
Physical ID:1               online
Physical ID:0               online
SCSI ID:1                   100%
SCSI ID:0                   100%

```

```

Size  Used Avail Used% Mounted on
----  -
423G  5.3G 396G   2%   /
99M   16M  79M   17%  /boot
93G   7.4G 86G    8%  /pbdata

```

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager           online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         online
DHCP Server master          online    backup
FTP Server                  online
Syslog                      online
Distributed Management      online
SNMP Trap Forwarder         online
SNMP Process                online
Platform Management         online

```


Interface Link Status

```

-----
Management Interface          up
Control Plane Bridge          up
Control Plane LAG             up
CP Link [0/2]                 up
CP Link [0/1]                 up
CP Link [0/0]                 up
CP Link [1/2]                 down
CP Link [1/1]                 down
CP Link [1/0]                 down
Crossover LAG                 up
CP Link [0/3]                 up
CP Link [1/3]                 up

```

session1@qfabric> show fabric administration inventory infrastructure

dg0:

Routing Engine Type	Hostname	PID	CPU-Use(%)
Fabric control	QFabric_default_FC-1_RE0	27906	2.5
Network Node group	QFabric_default_NW-NG-1_RE1	20421	1.8
Fabric manager	FM-0	4211	1.8
Debug Routing Engine	QFabric_DRE	1575	3.3

dg1:

Routing Engine Type	Hostname	PID	CPU-Use(%)
Fabric control	QFabric_default_FC-0_RE0	5686	2.3
Network Node group	QFabric_default_NW-NG-0_RE0	5866	1.9
Fabric manager	FM-1	572	1.6

Verifying a Fabric Nonstop Software Upgrade

Purpose

During the fabric portion of a nonstop software upgrade, you should expect to see both fabric control Routing Engines upgrade first, followed by the upgrade of each Interconnect device one at a time. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

Action

1. In an SSH session to the QFabric CLI, issue the request for the fabric nonstop software upgrade.

```
root@qfabric> request system software nonstop-upgrade fabric jinstall-qfabric-12.2X50-D10.3.rpm
```

```
[FC-0      2012-06-05 16:48:53]: Retrieving package
[FC-1      2012-06-05 16:48:53]: Retrieving package
[IC-F4912  2012-06-05 16:48:59]: Retrieving package
[FC-0      2012-06-05 16:49:51]: ----- re0: -----
[FC-1      2012-06-05 16:49:52]: ----- re0: -----
[IC-F4912  2012-06-05 16:49:54]: ----- re0: -----
[IC-F4912  2012-06-05 16:50:42]: Step 1 of 20 Creating temporary file system
[IC-F4912  2012-06-05 16:50:42]: Step 2 of 20 Determining installation source
[IC-F4912  2012-06-05 16:50:43]: Step 3 of 20 Processing format options
[IC-F4912  2012-06-05 16:50:43]: Step 4 of 20 Determining installation slice
[IC-F4912  2012-06-05 16:50:43]: Step 5 of 20 Creating and labeling new slices
[IC-F4912  2012-06-05 16:50:44]: Step 6 of 20 Create and mount new file system
[IC-F4912  2012-06-05 16:50:53]: Step 7 of 20 Getting OS bundles
[IC-F4912  2012-06-05 16:50:53]: Step 8 of 20 Updating recovery media
[IC-F4912  2012-06-05 16:51:17]: Step 9 of 20 Extracting incoming image
[IC-F4912  2012-06-05 16:52:56]: Step 10 of 20 Unpacking OS packages
[IC-F4912  2012-06-05 16:52:59]: Step 11 of 20 Mounting jbase package
[IC-F4912  2012-06-05 16:53:28]: Step 12 of 20 Creating base OS symbolic links
[IC-F4912  2012-06-05 16:54:45]: Step 13 of 20 Creating fstab
[IC-F4912  2012-06-05 16:54:45]: Step 14 of 20 Creating new system files
[IC-F4912  2012-06-05 16:54:46]: Step 15 of 20 Adding jbundle package
[IC-F4912  2012-06-05 16:58:15]: Step 16 of 20 Backing up system data
[IC-F4912  2012-06-05 16:58:18]: Step 17 of 20 Setting up shared partition data
[IC-F4912  2012-06-05 16:58:18]: Step 18 of 20 Checking package sanity in
installation
[IC-F4912  2012-06-05 16:58:18]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[IC-F4912  2012-06-05 16:58:22]: Step 20 of 20 Setting da0s1 as new active
partition
[Status    2012-06-05 16:58:34]: Load completed with 0 errors...
[Status    2012-06-05 16:58:34]: Reboot is required to complete upgrade ...
[Status    2012-06-05 16:58:34]: Trying to Connect to Node: FC-0
```



```
[Status 2012-06-05 16:58:39]: Rebooting FC-0
[Status 2012-06-05 16:58:39]: Trying to Connect to Node: FC-1
[Status 2012-06-05 16:58:44]: Rebooting FC-1
[Status 2012-06-05 16:58:44]: Trying to Connect to Node: IC-F4912
[Status 2012-06-05 16:58:50]: Rebooting IC-F4912
Success
```

2. When the fabric components reboot, they appear as **Disconnected** in the output of the **show fabric administration inventory infrastructure fabric-controls** and **show fabric administration inventory interconnect-devices** commands.

session1@qfabric> show fabric administration inventory infrastructure fabric-controls

Item	Identifier	Connection	Configuration
Fabric control			
FC-0		Disconnected	
FC-1		Disconnected	

session1@qfabric> show fabric administration inventory interconnect-devices IC-F4912

Item	Identifier	Connection	Configuration
Interconnect device			
IC-F4912		Disconnected	
F4912/RE0		Disconnected	

3. When the fabric components return to full service, they appear as **Connected** in the output of the **show fabric administration inventory** command.

session1@qfabric> show fabric administration inventory

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
P1507-C		Connected	
RSNG		Connected	Configured
P1550-C		Connected	
P1571-C		Connected	
Interconnect device			
IC-F4912		Connected	Configured
F4912/RE0		Connected	
Fabric manager			
FM-0		Connected	Configured
Fabric control			

FC-0	Connected	Configured
FC-1	Connected	Configured
Diagnostic routing engine		
DRE-0	Connected	Configured

Verifying a Redundant Server Node Group Nonstop Software Upgrade

Purpose

During the redundant server Node group portion of a nonstop software upgrade, you should expect to see the backup Node device upgrade first, followed by the upgrade of the master Node device. Server Node groups with a single device upgrade the device in the same way as a standalone switch. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

Action

1. In an SSH session to the QFabric CLI, issue the request for the redundant server Node group nonstop software upgrade.

```
root@qfabric> request system software nonstop-upgrade node-group RSNG
jinstall-qfabric-12.2X50-D10.3.rpm
```

```
Upgrading target(s): RSNG

[RSNG 2012-06-05 17:26:44]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
[RSNG 2012-06-05 17:26:44]: Retrieving package
[RSNG 2012-06-05 17:28:56]: Pushing bundle to fpc1
[RSNG 2012-06-05 17:29:26]: fpc1: Validate package...
[RSNG 2012-06-05 17:35:22]: fpc0: Validate package...
[RSNG 2012-06-05 17:35:49]: ----- fpc1 -----
[RSNG 2012-06-05 17:36:25]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:36:26]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:36:26]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:36:26]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:36:27]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:36:27]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:36:35]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:36:35]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:36:56]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:38:07]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:38:16]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:38:41]: Step 12 of 20 Creating base OS symbolic links
```



```

[RSNG      2012-06-05 17:39:41]: Step 13 of 20 Creating fstab
[RSNG      2012-06-05 17:39:42]: Step 14 of 20 Creating new system files
[RSNG      2012-06-05 17:39:42]: Step 15 of 20 Adding jbundle package
[RSNG      2012-06-05 17:42:16]: Step 16 of 20 Backing up system data
[RSNG      2012-06-05 17:42:32]: Step 17 of 20 Setting up shared partition data
[RSNG      2012-06-05 17:42:33]: Step 18 of 20 Checking package sanity in
installation
[RSNG      2012-06-05 17:42:33]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[RSNG      2012-06-05 17:42:36]: Step 20 of 20 Setting da0s2 as new active
partition
[RSNG      2012-06-05 17:42:51]: ----- fpc0 - master -----
[RSNG      2012-06-05 17:42:51]: Step 1 of 20 Creating temporary file system
[RSNG      2012-06-05 17:42:51]: Step 2 of 20 Determining installation source
[RSNG      2012-06-05 17:42:51]: Step 3 of 20 Processing format options
[RSNG      2012-06-05 17:42:51]: Step 4 of 20 Determining installation slice
[RSNG      2012-06-05 17:42:51]: Step 5 of 20 Creating and labeling new slices
[RSNG      2012-06-05 17:42:51]: Step 6 of 20 Create and mount new file system
[RSNG      2012-06-05 17:42:51]: Step 7 of 20 Getting OS bundles
[RSNG      2012-06-05 17:42:51]: Step 8 of 20 Updating recovery media
[RSNG      2012-06-05 17:42:51]: Step 9 of 20 Extracting incoming image
[RSNG      2012-06-05 17:42:51]: Step 10 of 20 Unpacking OS packages
[RSNG      2012-06-05 17:42:51]: Step 11 of 20 Mounting jbase package
[RSNG      2012-06-05 17:42:51]: Step 12 of 20 Creating base OS symbolic links
[RSNG      2012-06-05 17:42:51]: Step 13 of 20 Creating fstab
[RSNG      2012-06-05 17:42:51]: Step 14 of 20 Creating new system files
[RSNG      2012-06-05 17:42:51]: Step 15 of 20 Adding jbundle package
[RSNG      2012-06-05 17:42:51]: Step 16 of 20 Backing up system data
[RSNG      2012-06-05 17:42:51]: Step 17 of 20 Setting up shared partition data
[RSNG      2012-06-05 17:42:51]: Step 18 of 20 Checking package sanity in
installation
[RSNG      2012-06-05 17:42:51]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[RSNG      2012-06-05 17:42:51]: Step 20 of 20 Setting da0s2 as new active
partition
[RSNG      2012-06-05 17:43:36]: Rebooting Backup RE
[RSNG      2012-06-05 17:43:36]: ----- Rebooting fpc1 -----
[RSNG      2012-06-05 17:50:12]: Initiating Chassis In-Service-Upgrade
[RSNG      2012-06-05 17:50:33]: Upgrading group: 0 fpc: 0
[RSNG      2012-06-05 17:52:38]: Upgrade complete for group:0
[RSNG      2012-06-05 17:52:38]: Upgrading group: 1 fpc: 1
[RSNG      2012-06-05 17:54:42]: Upgrade complete for group:1
[RSNG      2012-06-05 17:54:42]: Finished processing all upgrade groups, last
group :1

```



```
[RSNG      2012-06-05 17:54:48]: Preparing for Switchover
[RSNG      2012-06-05 17:55:38]: Switchover Completed
[Status    2012-06-05 17:55:41]: Upgrade completed with 0 errors
Success
```

2. Issue the **show system software upgrade status** command to view the status of the upgrade.

```
root@qfabric> show system software upgrade status
```

```
Wed Jan 16 22:06:02 2013 Software nonstop upgrade on:
RSNG in progress
```

3. During the redundant server Node group upgrade, the backup Node device (in this case, P1571-C) is upgraded first and appears in the **Disconnected** state in the output of the **show fabric administration inventory** command.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
P1507-C		Connected	
RSNG		Connected	Configured
P1550-C		Connected	
P1571-C		Disconnected	
Interconnect device			
IC-F4912		Connected	Configured
F4912/RE0		Connected	
Fabric manager			
FM-0		Connected	Configured
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured
Diagnostic routing engine			
DRE-0		Connected	Configured

4. After the backup Node device comes back online, the master Node device (in this case, P1550-C) appears in the **Disconnected** state in the output of the **show fabric administration inventory** command while the master Node device upgrades its software.

```
session1@qfabric> show fabric administration inventory
```


Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
P1507-C		Connected	
RSNG		Connected	Configured
P1550-C		Disconnected	
P1571-C		Connected	
Interconnect device			
IC-F4912		Connected	Configured
F4912/RE0		Connected	
Fabric manager			
FM-0		Connected	Configured
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured
Diagnostic routing engine			
DRE-0		Connected	Configured

5. After both Node devices in the redundant server Node group come back online, both Node devices appear as **Connected** to indicate the successful completion of the Node group nonstop software upgrade step.

session1@qfabric> show fabric administration inventory

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
P1507-C		Connected	
RSNG		Connected	Configured
P1550-C		Connected	
P1571-C		Connected	
Interconnect device			
IC-F4912		Connected	Configured
F4912/RE0		Connected	
Fabric manager			
FM-0		Connected	Configured
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured
Diagnostic routing engine			
DRE-0		Connected	Configured

Verifying a Network Node Group Nonstop Software Upgrade

Purpose

During the network Node group portion of a nonstop software upgrade, you should expect to see the backup network Node group Routing Engine upgrade first, followed by the Node devices within the network Node group upgrading one at a time, and ending with the upgrade of the master network Node group Routing Engine. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

NOTE: If you configure an upgrade group for Node groups containing 2 or more Node devices, all Node devices within the upgrade group reboot at the same time.

Action

1. In an SSH session to the QFabric CLI, issue the request for the network Node group nonstop software upgrade.

```
root@qfabric> request system software nonstop-upgrade node-group NW-NG-0
jinstall-qfabric-12.2X50-D10.3.rpm
```

```
Upgrading target(s): NW-NG-0

[NW-NG-0 2012-06-01 09:45:06]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
[NW-NG-0 2012-06-01 09:45:06]: Retrieving package
[NW-NG-0 2012-06-01 09:46:18]: Pushing bundle to fpc0
[NW-NG-0 2012-06-01 09:46:52]: fpc0: Validate package...
[NW-NG-0 2012-06-01 09:53:26]: ----- fpc0 -----
[NW-NG-0 2012-06-01 09:54:01]: Step 1 of 20 Creating temporary file system
[NW-NG-0 2012-06-01 09:54:01]: Step 2 of 20 Determining installation source
[NW-NG-0 2012-06-01 09:54:02]: Step 3 of 20 Processing format options
[NW-NG-0 2012-06-01 09:54:02]: Step 4 of 20 Determining installation slice
[NW-NG-0 2012-06-01 09:54:02]: Step 5 of 20 Creating and labeling new slices
[NW-NG-0 2012-06-01 09:54:03]: Step 6 of 20 Create and mount new file system
[NW-NG-0 2012-06-01 09:54:10]: Step 7 of 20 Getting OS bundles
[NW-NG-0 2012-06-01 09:54:10]: Step 8 of 20 Updating recovery media
[NW-NG-0 2012-06-01 09:54:31]: Step 9 of 20 Extracting incoming image
[NW-NG-0 2012-06-01 09:55:43]: Step 10 of 20 Unpacking OS packages
[NW-NG-0 2012-06-01 09:55:46]: Step 11 of 20 Mounting jbase package
[NW-NG-0 2012-06-01 09:56:09]: Step 12 of 20 Creating base OS symbolic links
[NW-NG-0 2012-06-01 09:57:05]: Step 13 of 20 Creating fstab
[NW-NG-0 2012-06-01 09:57:05]: Step 14 of 20 Creating new system files
[NW-NG-0 2012-06-01 09:57:05]: Step 15 of 20 Adding jbundle package
```



```

[NW-NG-0 2012-06-01 09:59:30]: Step 16 of 20 Backing up system data
[NW-NG-0 2012-06-01 09:59:44]: Step 17 of 20 Setting up shared partition data
[NW-NG-0 2012-06-01 09:59:44]: Step 18 of 20 Checking package sanity in
installation
[NW-NG-0 2012-06-01 09:59:44]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[NW-NG-0 2012-06-01 09:59:47]: Step 20 of 20 Setting da0s1 as new active
partition
[NW-NG-0 2012-06-01 09:59:55]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-dc-re.tgz
[NW-NG-0 2012-06-01 09:59:55]: Retrieving package
[NW-NG-0 2012-06-01 10:01:04]: Pushing bundle to rel
[NW-NG-0 2012-06-01 10:01:35]: rel: Validate package...
[NW-NG-0 2012-06-01 10:02:56]: re0: Validate package...
[NW-NG-0 2012-06-01 10:04:45]: Rebooting Backup RE
[NW-NG-0 2012-06-01 10:08:31]: Initiating Chassis In-Service-Upgrade
[NW-NG-0 2012-06-01 10:08:52]: Upgrading group: 0 fpc: 0
[NW-NG-0 2012-06-01 10:18:33]: Upgrade complete for group:0
[NW-NG-0 2012-06-01 10:18:33]: Finished processing all upgrade groups, last
group :0
[NW-NG-0 2012-06-01 10:18:37]: Preparing for Switchover
[NW-NG-0 2012-06-01 10:18:55]: Switchover Completed
[Status 2012-06-01 10:18:58]: Upgrade completed with 0 errors
Success

```

2. Issue the **show system software upgrade status** command to view the status of the upgrade.

```
root@qfabric> show system software upgrade status
```

```

Wed Jan 16 22:06:02 2013 Software nonstop upgrade on:
NW-NG-0 in progress

```

3. Verify the progress of the upgrade by issuing the **show chassis nonstop-upgrade node-group**, **show fabric administration inventory**, **show fabric administration inventory infrastructure**, and **show fabric administration inventory node-groups NW-NG-0** commands. You should see the backup network Node group Routing Engine reboot first, followed by each Node device within the network Node group, and ending with the reboot of master network Node group Routing Engine. Restarting devices appear as **Disconnected** in the output of the **show fabric administration inventory** command and restarting Routing Engines do not appear in output of the **show fabric administration inventory infrastructure** command until they return to service.

RELATED DOCUMENTATION

Nonstop Software Upgrade Checklist for QFabric Systems	529
Performing a Nonstop Software Upgrade on the QFabric System	533
Understanding Nonstop Software Upgrade for QFabric Systems	61
show chassis nonstop-upgrade node-group	778
show fabric administration inventory	780
show fabric administration inventory director-group status	786
show fabric administration inventory infrastructure	793
show fabric administration inventory interconnect-devices	797
show fabric administration inventory node-groups	802

Upgrading Software on a QFabric System

The QFabric system software package contains software for all of the different components in the QFabric system, such as the Director group, Interconnect devices, Node devices, and other QFabric system components. You can upgrade the software on all of the QFabric components at the same time using the **request system software add *package-name* component all reboot** command.

NOTE: Downgrading software on a QFabric system is not supported.

This topic describes the following tasks:

- [Backing Up the Current Configuration Files](#) | [566](#)
- [Downloading Software Files Using a Browser](#) | [567](#)
- [Retrieving Software Files for Download](#) | [568](#)
- [Installing the Software Package on the Entire QFabric System](#) | [568](#)

Backing Up the Current Configuration Files

To back up your current configuration files:

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

Back up the configuration files to a local directory, remote server, or removable drive (for example, an external USB flash drive).

For example:

```
user@switch> request system software configuration-backup /media/USB/
```

Downloading Software Files Using a Browser

NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website

<https://www.juniper.net/registration/Register.jsp> .

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 package you want to download in the **QFabric System Install Package** section:
 - If you want to upgrade the entire QFabric system, select **QFabric System - Complete Install Package**.
 - If you want to upgrade either a single Node or Interconnect device for recovery purposes, select **Node and Interconnect Device Install Package**. For information on how to perform a recovery installation on either a Node or Interconnect device, see [“Performing a Recovery Installation” on page 845](#).
- A login screen appears.
7. Enter your user ID and password and click **Login**.
 8. Read the End User License Agreement, select the **I agree** option button, and then click **Proceed**.
 9. Save the **jinstall-qfabric-version.rpm** file on your computer.

Retrieving Software Files for Download

Retrieve the software from the location in which you downloaded it. To do this, issue the **request system software download** command. The software package is copied from where you downloaded it and is placed locally on the QFabric system.

- To retrieve the software:

```
user@switch> request system software download /path/package-name
```

For example:

```
user@switch> request system software download ftp://server/files/jinstall-qfabric-11.3X30.6.rpm
```

Installing the Software Package on the Entire QFabric System

NOTE: On a QFabric system, a QFX3500 Node device or QFX3600 Node device might not be able to participate as a Node device in the QFabric system if the Node device is running a different version of software from that of the Director group. This mismatch of software versions between the Node device and the Director group can occur when the Node device is introduced into the setup, and both Director devices go offline before the Node device completes its auto-upgrade process to upgrade its software version to the same software version running on the Director group. The workaround is to reboot the QFX3500 or QFX3600 Node device once the Director group comes back online. The QFX3500 or QFX3600 Node device will initiate auto-upgrade and upgrade its software version from the Director group.

1. Issue the **request system software add package-name component all reboot** command.

For example:

```
user@switch> request system software add jinstall-qfabric-11.3X30.6.rpm component all reboot
```

NOTE: If you receive an error message after issuing the **request system software add package-name component all reboot** command that says that the configuration file cannot be loaded as is, you will need to enter configuration mode, make any necessary changes to the configuration file, and then commit the changes.

NOTE: The default value for a QFabric system software upgrade is **validate**. The validation step adds up to 10 minutes to the overall software upgrade. If the validation fails, the upgrade does not proceed and the QFabric system automatically issues the **request system software rollback** command to restore the current software image. If you upgrade more than one component (for example, by issuing the **component all** option), validation failure on one device stops the upgrade process for the other devices. If you do not want to validate the software package against the current configuration, issue the **no-validate** option.

2. After the reboot has finished, verify that the new version of software has been properly installed by issuing the **show version component all** command.

```
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]

```

RELATED DOCUMENTATION

Software Installation and Upgrade Overview

[Performing a QFabric System Recovery Installation on the Director Group | 835](#)

Upgrading Jloader Software on QFX Series Devices

request system software add

Software Installation and Upgrade Guide

Downgrading Software on a QFabric System

IN THIS SECTION

- (Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive | 573
- Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software | 575

Starting in Junos OS Release 14.1, if a software upgrade or configuration changes has made the QFabric system unstable or inoperable, you can rollback or downgrade to a previous version of software and configuration. The software and configuration that you rollback to is called a restore-point. The restore-point is stored in a dedicated partition. You can create a checksum (MD5 hash) for the partition in which the restore-partition is stored and verify the integrity of the restore-point partition.

NOTE: The ability to downgrade the software does not replace the existing back up and restore functionality.

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 *Switchingbox*, 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. Burn the software image you just downloaded to your UNIX workstation onto your external USB flash drive using the **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
```

18. Perform the steps in [“Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software”](#) on page 575 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**.

The Director device reboots from the local disk on which the software was just installed.

5. Log in as root on the Director device.

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]

6. 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.

7. Issue the **ls /mnt** command to list the *mount* directory.

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

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

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

9. 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/
```

10. 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
```

11. 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
```

12. Enter **y** to continue.

13. 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.

14. 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.

15. 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
```

16. 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
```

17. 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

```

Release History Table

Release	Description
14.1	Starting in Junos OS Release 14.1, if a software upgrade or configuration changes has made the QFabric system unstable or inoperable, you can rollback or downgrade to a previous version of software and configuration.

Performing System Backup and Recovery for a QFabric System

Many routers and switches require an administrator to recover the software package and the configuration file for the device separately. In the case of a device failure, this means the administrator might need to perform two separate tasks (if neither the software package nor the configuration file can be recovered).

In contrast, the QFabric system uses a unique mechanism that saves the backup and recovery files for both the Junos OS software and the system configuration into a single collection. The following QFabric system backup and recovery mechanism simplifies and streamlines the recovery process so you can return to normal operations as quickly as possible.

To backup and recover your QFabric system:

1. (First time only) Implement the following one-time procedure to prepare your QFabric system to use the system backup and recovery feature:
 - Insert a Juniper Networks software installation USB flash drive into the master Director device. (This drive was provided to you as one of the components of your QFabric system shipment.)
 - Issue the **request system software format-qfabric-backup** command. The contents and format of the USB flash drive are copied to the Director group shared directory and are used as the basis for all future backup and recovery operations.

```
user@qfabric> request system software format-qfabric-backup
```

```
Copying QFabric USB template image from /dev/sdb(Unigen,PQS4000,4009 MB).....
```

- Remove the Juniper Networks software installation USB drive from the master Director device.
2. Issue the **request system software system-backup** command to backup the software package and configuration file. This command saves the current files necessary to recover the QFabric system. The files are saved to a shared memory directory in the Director group.

NOTE: As you upgrade your system with new software and change the system configuration over time, remember to reissue this command periodically to save the newest files for recovery purposes.

```
user@qfabric> request system software system-backup
```

```
user@qfabric>
```


3. Insert a 4 GB or larger USB flash drive into the master Director device for your Director group, and issue the **request system software system-backup usb-create** command. This command copies the recovery files that have been backed up in the Director group and transfers them to the USB flash drive to create a recovery USB drive.

NOTE: Issuing this command overwrites the contents of the USB flash drive with the QFabric system recovery files.

```
user@qfabric> request system software system-backup usb-create /dev/sdb
```

```
Issuing this command will overwrite the contents of the USB drive.  
Continue? [yes,no] (no) yes
```

```
This operation will access the USB drive on 0281042010000013.  
Are you sure you want to continue? [yes,no] (no) yes
```

```
Copying QFabric recovery media to /dev/sdb...  
Successfully copied QFabric recovery media to /dev/sdb
```

4. Remove the recovery USB drive from the Director device, and store it securely in a known location that you will remember when you need to use the recovery USB drive.
5. If the QFabric system fails, power off the Director group, insert the recovery USB drive into the master Director device of your Director group, turn on power to the Director device, and follow the prompts to recover your system. This step restores the software package and the configuration file for your QFabric system.

RELATED DOCUMENTATION

[request system software format-qfabric-backup | 638](#)

[request system software system-backup | 647](#)

QFabric System Administration and Maintenance

IN THIS CHAPTER

- [Overview of Internal Fabric Monitoring | 583](#)
- [QFabric System Operational Mode Commands | 585](#)
- [Filtering Operational Mode Command Output in a QFabric System | 587](#)

Overview of Internal Fabric Monitoring

Internal fabric monitoring is a feature of the Operation, Administration, and Maintenance (OAM) of the QFabric system. This feature enables you to validate the flow path of protocol data units (PDUs) across a given VLAN on the QFabric system using the unicast ping, multicast ping, and traceroute operations.

Internal fabric monitoring is useful for fault detection on the QFabric system. For example, if a PDU reaches a destination that is not part of the VLAN configuration, the operation (unicast ping, multicast ping, or traceroute) displays the exception on the console at runtime.

The unicast and multicast ping operations send PDUs from a source interface (called the source fabric maintenance endpoint [FMEP]) to a destination FMEP. The destination FMEP sends a response to the source FMEP when the PDUs are received.

The traceroute operation, also called a flow linktrace, traces the path taken by a specific, learned unicast flow from a source FMEP to a destination FMEP in a VLAN. The source and destination FMEPs may be on the same Node device, different Node devices connected to the same Interconnect device, or different Node devices connected to different Interconnect devices. The flow path is the sequence of Packet Forwarding Engine forwarding hops along which the PDU travels. The hop-by-hop sequence and number of hops are reported in terms of fabric maintenance intermediate points (FMIPs), which are interfaces on the Packet Forwarding Engine of the Interconnect device. An FMIP sends a response to the source FMEP when the traceroute PDU is received.

The following internal fabric monitoring commands are supported:

- **show oam fabric flow specification**
- **show oam fabric interfaces**
- **ping fabric unicast-flow**

- **ping fabric multicast-flow**
- **traceroute fabric unicast-flow**

The following are the main components of the internal fabric monitoring feature:

- **FMEP**—Represents the source or destination point (endpoint) in the monitoring operations. An FMEP is an interface through which PDUs are sent (source FMEP) or received (destination FMEP). Upon receipt of the PDUs, the destination FMEP sends a response to the source FMEP to validate the monitoring flow. FMEPs are associated with a VLAN in the fabric maintenance association (FMA) configuration, and the source and destination FMEP addresses are configured in the flow specification.
- **FMA**—Associates a set of FMEPs with a VLAN. The FMA defines the VLAN and FMEP parameters, including the VLAN name, FMEP identifiers, FMEP names, and the interface names of the FMEPs. The FMEPs defined in the FMA are the source and destination FMEPs used in the monitoring commands.

NOTE: A default FMA is automatically created for each Node group in the QFabric system. The default FMA is used to send error response PDUs (for example, in the case of a VLAN leak) and responses to PDUs that are not mapped to a specific interface in the QFabric system.

- **Flow specification**—Configures the flow type and FMEP addressing parameters. Unicast flow types include an Ethernet type (other than IPv4) and Ethernet IPv4. Multicast flow types include the Ethernet IPv4 and VLAN flood type. The flow specification also defines parameters within each flow type, such as MAC or IPv4 addresses of the source and destination FMEPs. The names and identifiers of these FMEPs are configured in the FMA parameters.

To enable internal fabric monitoring, configure the **fabric-maintenance-associations** and **flow-specs** statements at the **[edit protocols oam fabric]** hierarchy level.

The **ping fabric unicast-flow**, **ping fabric multicast-flow**, and **traceroute fabric unicast-flow** commands require that you specify the flow specification and FMA names, as well as the source and destination FMEP names.

RELATED DOCUMENTATION

[Configuring a Fabric Maintenance Association | 441](#)

[Configuring Flow Specifications | 442](#)

[fabric \(OAM\) | 470](#)

[ping fabric multicast-flow | 591](#)

[ping fabric unicast-flow | 593](#)

[traceroute fabric unicast-flow | 828](#)

QFabric System Operational Mode Commands

Table 79 on page 585 summarizes the command line interface (CLI) commands that you can use to monitor and troubleshoot the QFabric system operations.

Table 79: QFabric System Operational Mode Commands

Task	Command
Bring an Interconnect device Control Board offline or online.	<code>request chassis cb</code>
Select the operating mode for the device.	<code>request chassis device-mode</code>
Set the Interconnect device Flexible PIC Concentrator (FPC) offline or online for the QFabric system.	<code>request chassis fabric fpc</code>
Log in to individual QFabric system components for device level troubleshooting.	<code>request component login</code>
Select a Director device to become the new primary device within a Director group.	<code>request fabric administration director-group change-master</code>
Power off an Interconnect device in a QFX3000-G QFabric system (take the device offline and gracefully shut it down).	<code>request fabric administration power-off interconnect-device</code>
Power off an Interconnect device FPC card in a QFX3000-G QFabric system.	<code>request fabric administration power-off interconnect-device fpc</code>
Power off node devices in server Node groups, redundant server Node groups, and network Node groups In a QFX3000-G Qfabric system, or power off an Interconnect device in a QFX3000-M QFabric system.	<code>request fabric administration power-off node-device</code>
Remove a disconnected component from the QFabric system inventory.	<code>request fabric administration remove</code>
Add a MAC range to the MAC pool assigned to the QFabric system.	<code>request fabric administration system mac-pool add</code>
Delete a MAC range from the MAC block assigned to the QFabric system.	<code>request fabric administration system mac-pool delete</code>
Halt a Director device.	<code>request system halt director-device</code>
Reboot QFabric system components.	<code>request system reboot</code>

Table 79: QFabric System Operational Mode Commands (*continued*)

Task	Command
Upgrade the software version of the QFabric system by using the nonstop software upgrade method (which preserves forwarding functionality during the upgrade and enables components to upgrade on a rotating basis).	request system software nonstop-upgrade
Save the software package and system configuration files to be able to recover the QFabric system in the case of a system failure.	request system software system-backup
Specify information to be displayed on the LCD panel of a QFabric system device.	set chassis display message
Display information about the operating mode of the device.	show chassis device-mode
Display information about the Control Board of an Interconnect device.	show chassis environment cb
Display the status of Ethernet switching in the Control Board of an Interconnect device.	show chassis ethernet-switch interconnect-device cb
Display the status of Ethernet switching in the Flexible PIC Controller (FPC) of an Interconnect device.	show chassis ethernet-switch interconnect-device fpc
Display the status of the data plane connections in the QFabric system.	show chassis fabric connectivity
Display the fabric management status of devices in your QFabric system.	show chassis fabric device
Display information shown on the LCD screen of a QFabric system device.	show chassis lcd
Display the status of a nonstop software upgrade for a Node group.	show chassis nonstop-upgrade node-group
Display all devices that belong to the QFabric system.	show fabric administration inventory
Display the Director devices that belong to a QFabric system Director group.	show fabric administration inventory director-group status
Display the services running on the Director group for the QFabric system.	show fabric administration inventory infrastructure
Display the Interconnect devices that belong to a QFabric system.	show fabric administration inventory interconnect-devices
Display the Node devices that belong to the QFabric system.	show fabric administration inventory node-devices

Table 79: QFabric System Operational Mode Commands (*continued*)

Task	Command
Display the Node groups and the corresponding Node devices that belong to the QFabric system.	show fabric administration inventory node-groups
Display all devices that belong to the QFabric system.	show fabric inventory
Display the MAC addresses that belong to a QFabric system Director group.	show fabric administration system mac-pool
Display the Director device that hosts the QFabric CLI session.	show fabric session-host
Display the system log messages in the specified file.	show log
Display the status of a QFabric system software upgrade.	show system software upgrade status

Filtering Operational Mode Command Output in a QFabric System

When you issue an operational mode command in a QFabric system, the output generated can be fairly extensive because of the number of components contained within the system. To make the output more accessible, you can filter the output by appending the **| filter** option to the end of most Junos OS commands.

1. To filter operational mode command output and limit it to a Node group, include the **| filter node-group node-group-name** option at the end of your Junos OS operational mode command.

```
root@qfabric> show interfaces terse | filter node-group NW-NG-0
```

```

Interface           Admin Link Proto  Local           Remote
NW-NG-0:dsc         up   up
NW-NG-0:em0         up   up
NW-NG-0:em1         up   up
NW-NG-0:gre         up   up
NW-NG-0:ipip        up   up
NW-NG-0:lo0         up   up
NW-NG-0:lo0.16384   up   up   inet    127.0.0.1       --> 0/0
NW-NG-0:lo0.16385   up   up   inet
NW-NG-0:lsi         up   up
NW-NG-0:mtun        up   up
NW-NG-0:pimd        up   up

```


NW-NG-0:pime	up	up
NW-NG-0:tap	up	up
Node01:ge-0/0/10	up	up
Node01:ge-0/0/40	up	up
Node01:ge-0/0/41	up	up
vlan	up	up

2. To filter operational mode command output and limit it to a set of Node groups, include the | **filter node-group** option at the end of your Junos OS operational mode command and specify the list of Node group names in brackets.

root@qfabric> show ethernet-switching interfaces | filter node-group [NW-NG-0 RSNG-1]

Interface	State	VLAN members	Tag	Tagging	Blocking
NW-NG-0:ae0.0	up	v200	200	tagged	unblocked
		v50	50	tagged	unblocked
		v51	51	tagged	unblocked
		v52	52	tagged	unblocked
		v53	53	tagged	unblocked
RSNG-1:ae0.0	up	v200	200	untagged	unblocked
RSNG-1:ae47.0	up	v50	50	tagged	unblocked
		v51	51	tagged	unblocked
		v52	52	tagged	unblocked
		v53	53	tagged	unblocked

Operational Mode Commands

IN THIS CHAPTER

- ping fabric multicast-flow | 591
- ping fabric unicast-flow | 593
- request chassis cb | 596
- request chassis device-mode | 600
- request chassis fabric fpc | 603
- request component login | 605
- request fabric administration director-group change-master | 608
- request fabric administration power-off interconnect-device | 610
- request fabric administration power-off interconnect-device fpc | 612
- request fabric administration power-off node-device | 614
- request fabric administration remove | 616
- request fabric administration system mac-pool add | 618
- request fabric administration system mac-pool delete | 619
- request system halt | 621
- request system reboot | 628
- request system software format-qfabric-backup | 638
- request system software nonstop-upgrade | 639
- request system software system-backup | 647
- set chassis display message | 649
- show chassis device-mode | 653
- show chassis environment cb | 656
- show chassis ethernet-switch interconnect-device cb | 692
- show chassis ethernet-switch interconnect-device fpc | 715
- show chassis fabric connectivity | 749
- show chassis fabric device | 758
- show chassis lcd | 760
- show chassis nonstop-upgrade node-group | 778
- show fabric administration inventory | 780

- [show fabric administration inventory director-group status | 786](#)
- [show fabric administration inventory infrastructure | 793](#)
- [show fabric administration inventory interconnect-devices | 797](#)
- [show fabric administration inventory node-devices | 800](#)
- [show fabric administration inventory node-groups | 802](#)
- [show fabric administration system mac-pool | 805](#)
- [show fabric inventory | 807](#)
- [show fabric session-host | 811](#)
- [show log | 813](#)
- [show oam fabric flow-specification | 820](#)
- [show oam fabric interfaces | 824](#)
- [show system software upgrade status | 826](#)
- [traceroute fabric unicast-flow | 828](#)

ping fabric multicast-flow

Syntax

```
ping fabric multicast-flow source-fmep-id source-fmep-id fma-name fma-name flow-spec-name flow-specification-name
<verbose>
```

Release Information

Command introduced in Junos OS Release 12.2 for the QFX Series.

Description

Send a fabric multicast ping command to the multicast group destination fabric maintenance endpoints (FMEPs) configured in the flow specification. Send a single PDU with a timeout after 1 second.

Options

flow-spec-name *flow-specification-name*—Name of the flow specification that defines the protocol parameters (unicast or multicast) for the fabric ping operation.

fma-name *fma-name*—Name of the FMA.

source-fmep-id *source-fmep-id*—FMEP ID that is the source of the fabric ping operation.

verbose—(Optional) Detailed version of the output display.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring	439
Configuring a Fabric Maintenance Association	441
Configuring Flow Specifications	442
fabric (OAM)	470
ping fabric unicast-flow	593
traceroute fabric unicast-flow	828

List of Sample Output

[ping fabric multicast-flow on page 592](#)

Output Fields

Table 80 on page 592 lists the output fields for the **ping fabric-multicast-flow** command. Output fields are listed in the approximate order in which they appear.

Table 80: ping fabric multicast-flow Output Fields

Field Name	Field Description
Using fabric flow-specification	Name of the flow specification.
Ethernet frame-size	Ethernet frame size in bytes.
Source-IP	MAC address of the source FMEP interface.
Destination-IP	MAC address of the destination FMEP interface.
Protocol	IP protocol.
received response from fmep-id	FMEP ID from which responses are received.

Sample Output

ping fabric multicast-flow

```
user@host> ping fabric multicast-flow source-fmep-id 3 fma-name fma1 flow-spec-name fspec2
verbose
```

```
Using fabric flow-specification: fspec2
Ethernet frame-size: 256
Source-IP: 87.238.205.21
Destination-IP: 225.0.0.1 IP
Protocol: 17
received response from fmep-id 32775...
received response from fmep-id 32794...
received response from fmep-id 32795...
sent 1 requests, received 3 responses
```


ping fabric unicast-flow

Syntax

```
ping fabric unicast-flow fma-name fma-name source-fmep-id source-fmep-id dest-fmep-id dest-fmep-id
    flow-spec-name flow-specification-name
    <count count>
    <forced-fte-interface fte-interface>
    <verbose>
```

Release Information

Command introduced in Junos OS Release 12.2 for the QFX Series.

Description

Send a unicast fabric ping command to the destination fabric maintenance endpoints (FMEPs) specified.

Options

count *count*—(Optional) Number of fabric ping PDUs to send.

The maximum PDU count is 5. The default PDU count is 1.

dest-fmep-id *destination-fmep-id*—ID of the destination FMEP of the fabric ping operation.

flow-spec-name *flow-specification-name*—Name of the flow specification that defines the protocol parameters (unicast or multicast) for the fabric ping operation.

fma-name *fma-name*—Name of the FMA.

forced-fte-interface *fte-interface*—(Optional) Forces the fabric ping operation to use the specified FTE interface to inject the PDU into the fabric instead of using the internal forwarding path lookup table to determine the FTE interface.

source-fmep-id *source-fmep-id*—ID of the FMEP that is the source of the fabric ping operation.

verbose—(Optional) Detailed version of the output display.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Overview of Internal Fabric Monitoring | 439](#)

[Configuring a Fabric Maintenance Association | 441](#)

[Configuring Flow Specifications | 442](#)

[fabric \(OAM\) | 470](#)

[ping fabric multicast-flow | 591](#)

[traceroute fabric unicast-flow | 828](#)

List of Sample Output

[ping fabric unicast-flow on page 594](#)

[ping fabric unicast-flow \(With Masking Enabled on the Flow Specification\) on page 595](#)

Output Fields

Table 81 on page 594 lists the output fields for the **ping fabric-unicast-flow** command. Output fields are listed in the approximate order in which they appear.

Table 81: ping fabric unicast-flow Output Fields

Field Name	Field Description
Fabric flow ping between source <i>source-fmep-name</i> destination <i>destination-fmep-name</i>	Path of the unicast-flow ping between the source and destination FMEPs.
Using fabric flow-specification	Name of the flow specification.
Ethernet frame-size	Ethernet frame size in bytes.
Source-MAC	MAC address of the source FMEP interface.
Destination-MAC	MAC address of the destination FMEP interface.
Ethertype	EtherType protocol.
received response from fmep-id	ID of the FMEP sending the response.

Sample Output

ping fabric unicast-flow

```
user@host> ping fabric unicast-flow fma-name fma1 source-fmep-id 1 dest-fmep-id 2 flow-spec-name
fspec1 verbose
```

```
Fabric flow ping between source ED1494 destination ED1497
Using fabric flow-specification: fspec1
```



```

Ethernet frame-size: 256
Source-MAC: 0:2:B1:61:3B:96
Destination-MAC: 0:26:78:90:70:21
Ethertype: 8295
received response from fmep-id 2...
sent 1 requests, received 1 responses

```

ping fabric unicast-flow (With Masking Enabled on the Flow Specification)

user@host> **ping fabric unicast-flow fma-name fma1 source-fmep-id 1 dest-fmep-id 2 flow-spec-name fspec1 verbose**

```

Fabric flow ping between source ED1494 destination ED1497
Using fabric flow-specification: fspec1
Ethernet frame-size: 256
Source-MAC: 0:1:20:A0:0:1
Destination-MAC: 0:8A:3A:B3:7D:67
Source-MAC Mask: FF:F8:FF:FF:FF:FF
Ethertype: 26785
Current Source-MAC value: 0:1:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:2:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:3:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:4:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:5:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:6:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:7:20:a0:0:1
received response from fmep-id 2...
sent 7 requests, received 7 responses

```


request chassis cb

List of Syntax

[Syntax on page 596](#)

[Syntax \(TX Matrix Router\) on page 596](#)

[Syntax \(TX Matrix Plus Router\) on page 596](#)

[Syntax \(QFabric System\) on page 596](#)

[Syntax \(EX9253 Switches\) on page 596](#)

Syntax

```
request chassis cb (offline | online) slot slot-number
```

Syntax (TX Matrix Router)

```
request chassis cb (offline | online) <slot slot-number | lcc number slot cb-slot-number | scc number slot cb-slot-number>
```

Syntax (TX Matrix Plus Router)

```
request chassis cb (offline | online) <slot slot-number | lcc number slot cb-slot-number | sfc number slot cb-slot-number>
```

Syntax (QFabric System)

```
request chassis cb (offline | online) interconnect-device name slot slot-number
<interconnect-device name slot slot-number (offline | online)>
```

Syntax (EX9253 Switches)

```
request chassis cb (offline | online) name slot slot-number
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS 9.4 for EX Series switches.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

Command introduced in Junos OS 11.3 for QFX Series.

Command introduced in Junos OS Release 12.3 for MX2010 and MX2020 Universal Routing Platforms.

Command introduced in Junos OS Release 17.2 for MX2008 Universal Routing Platforms.

Command introduced in Junos OS Release 17.3 for MX10003 Universal Routing Platforms.

Command introduced in Junos OS Release 18.2 for EX9253 Series Switches.

Description

(M120, M320, and MX Series routers and T Series routers, QFabric systems, and EX8200 switches only)
Control the operation of the Control Board (CB).

Options

offline—Take the Control Board offline.

NOTE: On a QFabric system, to bring the backup Control Board on a QFX3008-I Interconnect device offline, issue the **request chassis cb slot *backup-slot-number* offline** command.

NOTE: Only backup Control Board can be turned offline or online. To turn a Control Board offline or to bring it back online, the Routing Engine should be turned offline first.

online—Bring the Control Board online.

interconnect-device *name*—(QFabric systems only) (Optional) Bring the QFX3008-I Interconnect device Control Board either offline or online:

slot *slot-number*—Control Board slot number:

- (TX Matrix and TX Matrix Plus routers only) On a TX Matrix router, if you specify the number of the T640 router by using the *lcc number* option (the recommended method), replace *cb-slot-number* with a value from 0 through 1.

Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 or T4000 router by using the *lcc number* option (the recommended method), replace *cb-slot-number* with a value from 0 through 1.

- M320 router—Replace *slot-number* with a value from 0 through 1.
- MX480/MX240 routers—Replace *slot-number* with a value from 0 through 1.
- MX960 router—Replace *slot-number* with a value from 0 through 2.
- MX2020, MX2010, and MX2008 routers—Replace *slot-number* with 0 or 1.
- EX8208 switch—Replace *slot-number* with a value from 0 through 2.
- EX8216 switch—Replace *slot-number* with a value from 0 through 1.
- QFabric System—Replace *slot-number* with a value from 0 through 1.

lcc *number*—(TX Matrix, TX Matrix Plus routers only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

sfc *number*—(TX Matrix Plus routers only) (Optional) Change the CB status for the TX Matrix Plus router (switch-fabric chassis). Replace *number* with 0.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[show chassis environment cb | 656](#)

Understanding Switching Control Board Redundancy

List of Sample Output

[request chassis cb on page 598](#)

[request chassis cb interconnect-device \(QFabric System\) on page 599](#)

[request chassis cb \(MX2020 Router\) on page 599](#)

[request chassis cb \(MX2010 Router\) on page 599](#)

[request chassis cb \(MX2008 Router\) on page 599](#)

[request chassis cb \(MX10003 Router\) on page 599](#)

[request chassis cb \(EX9253 Switch\) on page 599](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis cb

user@host> **request chassis cb offline slot 1**

```
Backup CB 1 cannot be set offline, backup RE is online
```


request chassis cb interconnect-device (QFabric System)

```
user@switch> request chassis cb interconnect-device interconnect1 offline slot 1
```

```
Backup CB 1 cannot be set offline, backup RE is online
```

request chassis cb (MX2020 Router)

```
user@host> request chassis cb offline slot 1
```

```
Backup CB 1 cannot be set offline, backup RE is online
```

request chassis cb (MX2010 Router)

```
user@host> request chassis cb offline slot 1
```

```
Backup CB 1 cannot be set offline, backup RE is online
```

request chassis cb (MX2008 Router)

```
user@host>request chassis cb offline slot 1
```

```
Backup CB 1 cannot be set offline, backup RE is online
```

request chassis cb (MX10003 Router)

```
user@host>request chassis cb online slot 1
```

```
CB 1 appears to be online already
```

request chassis cb (EX9253 Switch)

```
user@switch>request chassis cb offline slot 1
```

```
Offline initiated, use "show chassis environment cb" to verify
```


request chassis device-mode

Syntax

```
request chassis device-mode (interconnect-device | node-device | standalone)
```

Release Information

Command introduced in Junos OS Release 11.2 for the QFX Series.

interconnect-device option introduced in Junos OS Release 13.1 for the QFX Series.

Description

Select the operating mode for the device, which acts either as a device within a QFabric system or as a standalone switch.

NOTE:

- Issue the **request chassis device-mode** command only when your management station is connected directly to the device over a console port connection.
- Changing the device mode erases all configuration data on the device. When you convert a device to a different device mode, we recommend that you back up your device configuration to an external location before issuing the **request chassis device-mode** command.

Options

interconnect-device—Set the device to operate as an Interconnect device within a QFabric system. To complete the Interconnect device mode conversion process, you must connect the device to the QFabric system management control plane and reboot the device.

node-device—Set the device to operate as a Node device within a QFabric system. To complete the Node device mode conversion process, you must connect the device to the QFabric system management control plane and reboot the device.

standalone—Set the device to operate as a standalone switch. If the device starts in Node device or Interconnect device mode, you must reboot the device to return to standalone mode. Standalone mode is the factory default setting.

Required Privilege Level

admin

RELATED DOCUMENTATION

[Converting the Device Mode for a QFabric System Component](#) | 322

[show chassis device-mode | 653](#)

[Understanding Interconnect Devices | 29](#)

[Understanding Node Devices | 33](#)

[Understanding the QFabric System Hardware Architecture | 21](#)

List of Sample Output

[request chassis device-mode interconnect-device \(Starting in Node Device or Standalone Mode\) on page 601](#)

[request chassis device-mode node-device \(Starting in Interconnect Device or Standalone Mode\) on page 601](#)

[request chassis device-mode standalone \(Starting in Interconnect Device or Node Device Mode\) on page 601](#)

[request chassis device-mode standalone \(Starting in Standalone Mode\) on page 601](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis device-mode interconnect-device (Starting in Node Device or Standalone Mode)

```
user@switch> request chassis device-mode interconnect-device
```

```
Device mode set to 'interconnect-device' mode.
Please reboot the system to complete the process.
```

request chassis device-mode node-device (Starting in Interconnect Device or Standalone Mode)

```
user@switch> request chassis device-mode node-device
```

```
Device mode set to 'node-device' mode.
Please reboot the system to complete the process.
```

request chassis device-mode standalone (Starting in Interconnect Device or Node Device Mode)

```
user@switch> request chassis device-mode standalone
```

```
Device mode set to 'standalone' mode.
Please reboot the system to complete the process.
```

request chassis device-mode standalone (Starting in Standalone Mode)

```
user@switch> request chassis device-mode standalone
```



```
Device mode set to 'standalone' mode.  
No reboot required.
```


request chassis fabric fpc

Syntax

```
request chassis fabric fpc interconnect-device interconnect-device-name slot slot-number (offline | online)
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Set the Interconnect device Flexible PIC Concentrator (FPC) offline or online for the QFabric system. When the FPC is offline, traffic is redirected to other FPCs and is not lost while you remove or install an FPC. After issuing this command, you must issue the **request chassis fpc** command.

Options

interconnect-device *interconnect-device-name*—Set the Interconnect device containing the FPC you want to bring either offline or online.

slot *slot-number*—Set the specific FPC slot on the Interconnect device.

offline—Set the Interconnect device FPC to offline for removal.

online—Set the Interconnect device FPC to online after installation.

Required Privilege Level

admin

RELATED DOCUMENTATION

| [request chassis fpc](#)

List of Sample Output

[request chassis fabric fpc online on page 604](#)

[request chassis fabric fpc offline on page 604](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis fabric fpc online

user@qfabric> **request chassis fabric fpc interconnect-device IC-YW3781 offline slot 15**

Graceful offline of the fabric card has been initiated. Please wait 20 seconds before offlining or removing the card.

request chassis fabric fpc offline

user@qfabric> **request chassis fabric fpc interconnect-device IC-YW3781 online slot 15**

Bring the FPC online by issuing the "request chassis fpc online" command.

request component login

Syntax

```
request component login component-name
```

Release Information

Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Log in to a QFabric system component. To gain access to individual components by way of the **request component login** command, you must first provide the **qfabric-admin** or **qfabric-operator** class privilege to your user (for more information, see: *remote-debug-permission*).

Options

component-name—Specify the QFabric system component to which you wish to log in.

Required Privilege Level

admin

RELATED DOCUMENTATION

List of Sample Output

[request component login \(with qfabric-admin Privileges\) on page 605](#)

[request component login \(with qfabric-operator Privileges\) on page 606](#)

[request component login \(with qfabric-user Privileges\) on page 607](#)

Sample Output

The three sample output displays show the results of attempts to log in to Node device EE3093. The results differ depending on the privilege level assigned to the user.

request component login (with qfabric-admin Privileges)

```
admin@qfabric> request component login EE3093
```

```
Warning: Permanently added 'qfabric-node-ee3093,192.0.2.0' (RSA) to the list of
known hosts.
```

```
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
```



```
{master}
qfabric-admin@node-ee3093> ?
Possible completions:
clear          Clear information in the system
file           Perform file operations
help           Provide help information
load           Load information from file
monitor        Show real-time debugging information
mtrace         Trace multicast path from source to receiver
op             Invoke an operation script
ping           Ping remote target
quit           Exit the management session
request        Make system-level requests
restart        Restart software process
save           Save information to file
set            Set CLI properties, date/time, craft interface message
show           Show system information
ssh            Start secure shell on another host
start          Start shell
telnet         Telnet to another host
test           Perform diagnostic debugging
traceroute     Trace route to remote host{master}
qfabric-admin@node-ee3093>
```

request component login (with qfabric-operator Privileges)

operator@qfabric> **request component login EE3093**

```
Warning: Permanently added 'qfabric-node-EE3093,192.0.2.0' (RSA) to the list of
known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
qfabric-operator@node-EE3093> ?
Possible completions:
file           Perform file operations
help           Provide help information
load           Load information from file
op             Invoke an operation script
quit           Exit the management session
request        Make system-level requests
save           Save information to file
set            Set CLI properties, date/time, craft interface message
show           Show system information
start          Start shell
```



```
test                Perform diagnostic debugging
{master}
qfabric-operator@node-ee3093>
```

request component login (with qfabric-user Privileges)

```
user0@qfabric> request component login EE3093
```

```
error: User user0 does not have sufficient permissions to login to device ee3093
```


request fabric administration director-group change-master

Syntax

```
request fabric administration director-group change-master (director-device director-device-name)
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Select a Director device to become the new primary device within a Director group. The specified device becomes the new master Director device, and the previous master Director device becomes a backup Director device.

Options

none—Change the device that controls the Director group. Assign the current backup Director device as the new master and the current master Director device as the backup.

director-device *director-device-name*—Specify which Director device should become the primary device within the Director group.

Required Privilege Level

admin

RELATED DOCUMENTATION

[show fabric administration inventory director-group status | 786](#)

[Performing the QFabric System Initial Setup on a QFX3100 Director Group | 363](#)

[Understanding the Director Group | 26](#)

List of Sample Output

[request fabric administration director-group change-master on page 608](#)

Sample Output

```
request fabric administration director-group change-master
```

```
user@qfabric> request fabric administration director-group change-master
```



```
Do you intend to switchover mastership? [yes,no] (no) yes
```

```
Cluster master successfully switched
```


request fabric administration power-off interconnect-device

Syntax

```
request fabric administration power-off interconnect--device (serial-id | alias-name)
```

Release Information

Command introduced in Junos OS Release 14.1X53-D15 for the QFX Series.

Description

In QFX3000-G QFabric systems, you can use this command to power off QFX3008-I Interconnect devices.

This command systematically takes the device offline and gracefully shuts down the device while preserving system state information. A message appears on the console or console log, confirming that the operating system has stopped on the device.

Options

serial-id | alias-name—Provide the serial ID or alias name of the device.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[request fabric administration power-off interconnect-device fpc](#) | [612](#)

[request fabric administration power-off node-device](#) | [614](#)

[show fabric administration inventory](#) | [780](#)

List of Sample Output

[request fabric administration power-off interconnect-device on page 610](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request fabric administration power-off interconnect-device

```
user@qfabric request fabric administration power-off interconnect-device IC-12345
```



```
STEP 1 of 7 (Acquiring lock):  
Acquiring lock to perform this operation  
Acquired lock to perform this operation  
STEP 2 of 7 (Performing pre-checks):  
interconnect-device IC-12345 is online  
STEP 3 of 7 (Gracefully offlining the interconnect):  
Disabling fabric protocol on interconnect-device  
Disabled fabric protocol on interconnect-device  
STEP 4 of 7 (Waiting for convergence):  
Waiting for convergence(this will take few minutes)  
Convergence complete  
STEP 5 of 7 (Powering-off the interconnect):  
Now, powering-off the interconnect-device (this may take some time)  
Powered-off the interconnect-device  
STEP 6 of 7 (Updating inventory):  
Removing the interconnect-device - IC-A0004 from inventory  
interconnect-device - IC-12345 is removed from inventory  
STEP 7 of 7 (Releasing lock):  
Releasing the lock
```


request fabric administration power-off interconnect-device fpc

Syntax

```
request fabric administration power-off interconnect--device (serial-id | alias-name) fpc slot-number
```

Release Information

Command introduced in Junos OS Release 14.1X53-D15 for the QFX Series.

Description

In QFX3000-G QFabric systems, you can use this command to power off a 16-Port QSFP+ front card in a QFX3008-I Interconnect device.

A message appears on the console or console log, confirming that the operating system has stopped on the device.

Options

serial-id | alias-name—Provide the serial ID or alias name of the QFX3008-I Interconnect device.

slot slot-number—Provide the slot number of the 16-Port QSFP+ front card in the QFX3008-I Interconnect device. The range of slot numbers is 0 through 15

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[request fabric administration power-off interconnect-device | 610](#)

[request fabric administration power-off node-device | 614](#)

[show fabric administration inventory | 780](#)

List of Sample Output

[request fabric administration power-off interconnect-device fpc 13 on page 613](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request fabric administration power-off interconnect-device fpc 13

user@qfabric **request fabric administration power-off interconnect-device IC-12345 fpc 13**

```
STEP 1 of 6 (Acquiring lock):  
Acquiring lock to perform this operation  
Acquired lock to perform this operation  
STEP 2 of 6 (Performing pre-checks):  
interconnect-device IC-12345 is online  
STEP 3 of 6 (Gracefully offlining the interconnect linecard):  
Gracefully offlining the fpc  
fpc is successfully offlined  
STEP 4 of 6 (Waiting for convergence):  
Waiting for convergence(this will take few minutes)  
request fabric administration power-off interconnect-device IC-12345  
Convergence complete  
STEP 5 of 6 (Powering-off the interconnect linecard):  
Now, powering-off the fpc on the interconnect-device  
Powered-off the fpc on the interconnect-device  
STEP 6 of 6 (Releasing lock):  
Releasing the lock
```


request fabric administration power-off node-device

Syntax

```
request fabric administration power-off node-device (serial-id | alias-name)
```

Release Information

Command introduced in Junos OS Release 14.1X53-D15 for the QFX Series.

Description

In QFX3000-G QFabric systems, you can use this command to power off QFX3500, QFX3600, QFX5100-48S, QFX5100-48T, and QFX5100-24Q devices in server Node groups, redundant server Node groups, and network Node groups. In QFX3000-M QFabric systems, you can use this command to power off QFX5100-24Q Interconnect devices and QFX3600-I Interconnect devices.

This command systematically takes the device offline and gracefully shuts down the device while preserving system state information. A message appears on the console or console log, confirming that the operating system has stopped on the device.

Options

serial-id | alias-name—Provide the serial ID or alias name of the device.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[request fabric administration power-off interconnect-device | 610](#)

[request fabric administration power-off interconnect-device fpc | 612](#)

[show fabric administration inventory | 780](#)

List of Sample Output

[request fabric administration power-off node-device on page 615](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request fabric administration power-off node-device

user@qfabric **request fabric administration power-off node-device ED1234**

```
STEP 1 of 8 (Acquiring lock):
Acquiring lock to perform this operation
Acquired lock to perform this operation
STEP 2 of 8 (Performing pre-checks):
node-device ED1491 is online
STEP 3 of 8 (Mastership switch-over):
node-device ED1491 is Server Node-group
STEP 4 of 8 (Gracefully offlining the node-device):
node-device is being offlined
node-device is successfully offlined
STEP 5 of 8 (Waiting for convergence):
Waiting for convergence(this will take few minutes)
Convergence complete
STEP 6 of 8 (Powering-off the node-device):
Now, powering-off the node-device
Powered-off the node-device
STEP 7 of 8 (Updating inventory):
Removing the node-device - ED1491 from inventory
node-device - ED1491 is removed from inventory
STEP 8 of 8 (Releasing lock):
Releasing the lock
```


request fabric administration remove

Syntax

```
request fabric administration remove (interconnect-device interconnect-device-name | node-device node-device-name)
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Remove a disconnected Interconnect or Node device from the QFabric system inventory so that it does not appear in the output of the **show fabric administration inventory** command.

NOTE:

- You cannot remove any devices that appear in the **Connected** state in the output of the **show fabric administration inventory** command.
- For Node devices, you can only remove a device if it belongs to an autogenerated server Node group that contains a single Node device. Node devices contained within redundant server Node groups or network Node groups cannot be removed directly. To remove a Node device that is part of a group, delete the device from the Node group configuration first before attempting to remove the device from the inventory.

Options

interconnect-device *interconnect-device-name*—Remove a disconnected Interconnect device from the QFabric system inventory.

node-device *node-device-name*—Remove a disconnected Node device from the QFabric system inventory.

Required Privilege Level

admin

RELATED DOCUMENTATION

[show fabric administration inventory | 780](#)

[show fabric administration inventory interconnect-devices | 797](#)

[show fabric administration inventory node-devices | 800](#)

[show fabric administration inventory node-groups | 802](#)

List of Sample Output

[request fabric administration remove interconnect-device on page 617](#)

[request fabric administration remove node-device on page 617](#)

Sample Output

request fabric administration remove interconnect-device

user@qfabric> **request fabric administration remove interconnect-device IC1**

```
Device successfully removed
```

Sample Output

request fabric administration remove node-device

user@qfabric> **request fabric administration remove node-device node5**

```
Device successfully removed
```


request fabric administration system mac-pool add

Syntax

```
request fabric administration system mac-pool add mac-base starting-mac-address count number-of-mac-address
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Add a MAC address pool to expand the initial set of MAC addresses assigned to the QFabric system.

Options

mac-base *starting-mac-address*—Set the starting MAC address for a pool of addresses assigned to the QFabric system.

count *number-of-mac-address*—Set the total number of MAC addresses in the specified address pool assigned to the QFabric system.

Required Privilege Level

admin

RELATED DOCUMENTATION

[request fabric administration system mac-pool delete](#) | 619

[show fabric administration system mac-pool](#) | 805

List of Sample Output

[request fabric administration system mac-pool add mac-base starting-mac-address count](#) on page 618

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
request fabric administration system mac-pool add mac-base starting-mac-address count
```

```
user@switch> request fabric administration system mac-pool add mac-base 02:00:00:11:22:00 count
10
```


request fabric administration system mac-pool delete

Syntax

```
request fabric administration system mac-pool delete mac-base starting-mac-address
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Delete a range of MAC addresses assigned manually to the QFabric system.

NOTE: You cannot delete the MAC address range assigned during the initial setup of the QFabric system. Also, you cannot delete a MAC address range if the MAC address block is still in use.

Options

mac-base *starting-mac-address*—Specify the starting MAC address for a pool of addresses you wish to remove from the QFabric system.

Additional Information

After you issue the **request fabric administration system mac-pool delete** command, issue the [show fabric administration system mac-pool](#) command to verify that the MAC address range has been deleted.

Required Privilege Level

admin

RELATED DOCUMENTATION

[request fabric administration system mac-pool add](#) | 618

[show fabric administration system mac-pool](#) | 805

List of Sample Output

[request fabric administration system mac-pool delete mac-base](#) on page 620

Sample Output

```
request fabric administration system mac-pool delete mac-base
```

```
user@switch> request fabric administration system mac-pool delete mac-base 02:00:00:11:22:00
```


request system halt

List of Syntax

[Syntax on page 621](#)

[Syntax \(EX Series Switches\) on page 621](#)

[Syntax \(PTX Series\) on page 621](#)

[Syntax \(TX Matrix Router\) on page 622](#)

[Syntax \(TX Matrix Plus Router\) on page 622](#)

[Syntax \(MX Series Router\) on page 622](#)

[Syntax \(QFX Series\) on page 623](#)

Syntax

```
request system halt
<at time>
<backup-routing-engine>
<both-routing-engines>
<other-routing-engine>
<in minutes>
<media (compact-flash | disk | removable-compact-flash | usb)>
<message "text">
```

Syntax (EX Series Switches)

```
request system halt
<all-members>
<at time>
<backup-routing-engine>
<both-routing-engines>
<in minutes>
<local>
<media (external | internal)>
<member member-id>
<message "text">
<other-routing-engine>
<slice slice>
```

Syntax (PTX Series)

```
request system halt
<at time>
<backup-routing-engine>
<both-routing-engines>
```



```

<other-routing-engine>
<in minutes>
<media (compact-flash | disk)>
<message "text">

```

Syntax (TX Matrix Router)

```

request system halt
<all-lcc | lcc number | scc>
<at time>
<backup-routing-engine>
<both-routing-engines>
<other-routing-engine>
<in minutes>
<media (compact-flash | disk)>
<message "text">

```

Syntax (TX Matrix Plus Router)

```

request system halt
<all-chassis | all-lcc | lcc number | sfc number>
<at time>
<backup-routing-engine>
<both-routing-engines>
<other-routing-engine>
<in minutes>
<media (compact-flash | disk)>
<message "text">

```

Syntax (MX Series Router)

```

request system halt
<all-members>
<at time>
<backup-routing-engine>
<both-routing-engines>
<in minutes>
<local>
<media (external | internal)>
<member member-id>
<message "text">
<other-routing-engine>

```


Syntax (QFX Series)

```
request system halt
<all-members>
<at time>
<both-routing-engines>
<director-device director-device-id>
<in minutes>
<local>
<media >
<member member-id>
<message "text">
<other-routing-engine>
<slice slice>
```

Release Information

Command introduced before Junos OS Release 7.4.

other-routing-engine option introduced in Junos OS Release 8.0.

Command introduced in Junos OS Release 9.0 for EX Series switches.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

Command introduced in Junos OS Release 11.1 for the QFX Series.

director-device option introduced for QFabric systems in Junos OS Release 12.2.

backup-routing-engine option introduced in Junos OS Release 13.1.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Stop the router or switch software.

NOTE: When you issue this command on an individual component—for example, a Node device—in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos Routing Engines will reboot.” If you want to halt only one member of a Node group, issue this command with the **member** option on the Node device CLI, because you cannot issue this command from the QFabric CLI. Also, issuing this command might cause traffic loss on an individual component.

When you issue this command on a QFX5100 switch, you are not prompted to reboot. You must power cycle the switch to reboot.

NOTE: For the routers with the Routing Engines RE-S-2x00x6, RE-PTX-2x00x8, and RE-S-2x00x8, this command is deprecated and might be removed completely in a future release.

On these routers, this command is replaced with the **request vmhost halt** command which provides similar functionality.

Options

none—Stop the router or switch software immediately.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Halt all chassis.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, halt all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, halt all T1600 routers (or line-card chassis) connected to the TX Matrix Plus router.

all-members—(Optional) Halt all members of the Virtual Chassis configuration.

at time —(Optional) Time at which to stop the software, specified in one of the following ways:

- **now**—Stop the software immediately. This is the default.
- **+minutes**—Number of minutes from now to stop the software.
- **yymmddhhmm**—Absolute time at which to stop the software, specified as year, month, day, hour, and minute.
- **hh:mm**—Absolute time on the current day at which to stop the software.

backup-routing-engine—(Optional) Halt the backup Routing Engine. This command halts the backup Routing Engine, regardless from which Routing Engine the command is executed. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. If you issue the command from the backup Routing Engine, the backup Routing Engine is halted.

both-routing-engines—(Optional) Halt both Routing Engines at the same time.

director-device *director-device-id*—(QFabric systems only) Halt a specific Director device.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, halt a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, halt a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(Optional) Halt the local Virtual Chassis member.

in *minutes*—(Optional) Number of minutes from now to stop the software. This option is an alias for the **at +*minutes*** option.

media (compact-flash | disk)—(Optional) Boot medium for the next boot.

media (external | internal)—(EX Series and QFX Series switches and MX Series routers only) (Optional) Halt the boot media:

- **external**—Halt the external mass storage device.
- **internal**—Halt the internal flash device.

member *member-id*—(Optional) Halt the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, ***member-id*** can only be 0 or 1.

message "text"—(Optional) Message to display to all system users before stopping the software.

other-routing-engine—(Optional) Halt the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is halted.

scc—(TX Matrix routers only) (Optional) Halt the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Halt the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

slice *slice*—(EX Series and QFX Series switches only) (Optional) Halt a partition on the boot media. This option has the following suboptions:

- 1—Halt partition 1.
- 2—Halt partition 2.
- **alternate**—Reboot from the alternate partition.

Additional Information

On the M7i router, the **request system halt** command does not immediately power down the Packet Forwarding Engine. The power-down process can take as long as 5 minutes.

On a TX Matrix router and TX Matrix Plus router if you issue the **request system halt** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are halted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are halted.

NOTE: If you have a router or switch with two Routing Engines and you want to shut the power off to the router or switch or remove a Routing Engine, you must first halt the backup Routing Engine (if it has been upgraded), and then halt the master Routing Engine. To halt a Routing Engine, issue the **request system halt** command. You can also halt both Routing Engines at the same time by issuing the **request system halt both-routing-engines** command.

Required Privilege Level
maintenance

RELATED DOCUMENTATION

- clear system reboot*
- request system power-off*
- request vmhost halt*
- show virtual-chassis*
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

- List of Sample Output**
- [request system halt on page 626](#)
 - [request system halt \(In 2 Hours\) on page 627](#)
 - [request system halt \(Immediately\) on page 627](#)
 - [request system halt \(At 1:20 AM\) on page 627](#)

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system halt
user@host> **request system halt**


```
Halt the system ? [yes,no] (no) yes

*** FINAL System shutdown message from root@section2 ***
System going down IMMEDIATELY
Terminated
...
syncing disks... 11 8 done
The operating system has halted.
Please press any key to reboot.
```

request system halt (In 2 Hours)

The following example, which assumes that the time is 5 PM (1700), illustrates three different ways to request that the system stop 2 hours from now:

```
user@host> request system halt at +120

user@host> request system halt in 120

user@host> request system halt at 19:00
```

request system halt (Immediately)

```
user@host> request system halt at now
```

request system halt (At 1:20 AM)

To stop the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system halt at yymmdd120
```

```
request system halt at 120
Halt the system at 120? [yes,no] (no) yes
```


request system reboot

List of Syntax

[Syntax on page 628](#)

[Syntax \(EX Series Switches and EX Series Virtual Chassis\) on page 628](#)

[Syntax \(MX Series Routers and MX Series Virtual Chassis, EX9200 Switches and EX9200 Virtual Chassis\) on page 628](#)

[Syntax \(QFabric Systems\) on page 629](#)

[Syntax \(QFX Series Switches and QFX Series Virtual Chassis, Virtual Chassis Fabric\) on page 629](#)

[Syntax \(TX Matrix Router\) on page 629](#)

[Syntax \(TX Matrix Plus Router\) on page 629](#)

Syntax

```
request system reboot
  <at time>
  <both-routing-engines>
  <in minutes>
  <media (compact-flash | disk | removable-compact-flash | usb)>
  <message "text">
  <other-routing-engine>
```

Syntax (EX Series Switches and EX Series Virtual Chassis)

```
request system reboot
  <all-members | local | member member-id>
  <at time>
  <in minutes>
  <media (external | internal)> | <media (compact-flash | disk | removable-compact-flash | usb)>
  <message "text">
  <slice slice>
```

Syntax (MX Series Routers and MX Series Virtual Chassis, EX9200 Switches and EX9200 Virtual Chassis)

```
request system reboot
  <all-members | local | member member-id>
  <at time>
  <both-routing-engines>
  <in minutes>
  <media (external | internal)> | <media (compact-flash | disk | usb)> | <junos | network | oam | usb>
  <message "text">
  <other-routing-engine>
```


Syntax (QFabric Systems)

```
request system reboot
<all <graceful>>
<at time>
<director-device name>
<director-group <graceful>>
<fabric <graceful>>
<in minutes>
<in-service>
<media>
<message "text">
<node-group name>
<slice slice>
```

Syntax (QFX Series Switches and QFX Series Virtual Chassis, Virtual Chassis Fabric)

```
request system reboot
<all-members | local | member member-id>
<at time>
<in minutes>
<in-service>
<hypervisor>
<junos | network | oam | usb>
<message "text">
<slice slice>
```

Syntax (TX Matrix Router)

```
request system reboot
<all-chassis | all-lcc | lcc number | scc>
<at time>
<both-routing-engines>
<in minutes>
<media (compact-flash | disk)>
<message "text">
<other-routing-engine>
```

Syntax (TX Matrix Plus Router)

```
request system reboot
<all-chassis | all-lcc | lcc number | sfc number>
<at time>
```



```
<both-routing-engines>
<in minutes>
<media (compact-flash | disk)>
<message "text">
<other-routing-engine>
<partition (1 | 2 | alternate)>
```

Release Information

Command introduced before Junos OS Release 7.4.

Option **other-routing-engine** introduced in Junos OS Release 8.0.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Option **sfc** introduced for the TX Matrix Plus router in Junos OS Release 9.6.

Option **partition** changed to **slice** in Junos OS Release 10.0 for EX Series switches.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Option **both-routing-engines** introduced in Junos OS Release 12.1.

Description

Reboot the software.

This command can be used on standalone devices and on devices supported in a Virtual Chassis, Virtual Chassis Fabric, or QFabric system.

NOTE: Starting with Junos OS Release 15.1F3, the statement **request system reboot** reboots only the guest operating system on the PTX5000 with RE-PTX-X8-64G and, MX240, MX480, and MX960 with RE-S-X6-64G.

Starting with Junos OS Release 15.1F5, the statement **request system reboot** reboots only the guest operating system on the MX2010, and MX2020 with REMX2K-X8-64G.

NOTE: Starting from Junos OS Release 17.2R1, PTX10008 routers do not support the **request system reboot** command. Starting from Junos OS Release 17.4R1, PTX10016 routers do not support the **request system reboot** command. Use the **request vmhost reboot** command instead of the **request system reboot** command on the PTX10008 and PTX10016 routers to reboot the Junos OS software package or bundle on the router. See *request vmhost reboot*.

NOTE: On a QFabric system, to avoid traffic loss on the network Node group, switch mastership of the Routing Engine to the backup Routing Engine, and then reboot.

Options

The options described here are not all supported on every platform or release of Junos OS. Refer to the Syntax sections for the options commonly available on each type of platform.

none—Reboot the software immediately.

all-chassis—(Optional) On a TX Matrix router or TX Matrix Plus router, reboot all routers connected to the TX Matrix or TX Matrix Plus router, respectively.

all-lcc—(Optional) On a TX Matrix router or TX Matrix Plus router, reboot all line card chassis connected to the TX Matrix or TX Matrix Plus router, respectively.

all-members | local | member *member-id*—(Optional) Specify which member of the Virtual Chassis to reboot:

- **all-members**—Reboots each switch that is a member of the Virtual Chassis.
- **local**—Reboots only the local switch (switch where you are logged in).
- **member *member-id***—Reboots the specified member switch of the Virtual Chassis

at *time*—(Optional) Time at which to reboot the software, specified in one of the following ways:

- **now**—Stop or reboot the software immediately. This is the default.
- **+*minutes***—Number of minutes from now to reboot the software.
- ***yyymmddhhmm***—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.
- ***hh:mm***—Absolute time on the current day at which to stop the software, specified in 24-hour time.

both-routing-engines—(Optional) Reboot both Routing Engines at the same time.

hypervisor—(Optional) Reboot Junos OS, host OS, and any installed guest VMs.

in *minutes*—(Optional) Number of minutes from now to reboot the software. The minimum value is 1. This option is an alias for the **at +*minutes*** option.

in-service—(Optional) Enables you to reset the software state (no software version change) of the system with minimal disruption in data and control traffic.

junos—(Optional) Reboot from the Junos OS (main) volume.

lcc *number*—(Optional) Line-card chassis (LCC) number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

media (compact-flash | disk | removable-compact-flash | usb)—(Optional) Use the indicated boot medium for the next boot.

media (external | internal)—(Optional) Use the indicated boot medium for the next boot:

- **external**—Reboot the device using a software package stored on an external boot source, such as a USB flash drive.
- **internal**—Reboot the device using a software package stored in an internal memory source.

message "text"—(Optional) Message to display to all system users before stopping or rebooting the software.

network—(Optional) Reboot using the Preboot Execution Environment (PXE) boot method over the network.

oam—(Optional) Reboot from the maintenance volume (OAM volume, usually the compact flash drive).

other-routing-engine—(Optional) Reboot the other Routing Engine from which the command is issued.

For example, if you issue the command from the master Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is rebooted.

partition *partition*—(Optional) Reboot using the specified partition on the boot media. This option is equivalent to the **slice** option that is supported on some devices. Specify one of the following *partition* values:

- **1**—Reboot from partition 1.
- **2**—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition.

scc—(Optional) Reboot the Routing Engine on the TX Matrix switch-card chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted.

sfc *number*—(Optional) Reboot the Routing Engine on the TX Matrix Plus switch-fabric chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted. Replace *number* with 0.

slice *slice*—(Optional) Reboot using the specified partition on the boot media. This option was originally the **partitiion** option but was renamed to **slice** on EX Series and QFX Series switches. Specify one of the following *slice* values:

- **1**—Reboot from partition 1.

- **2**—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition (which did not boot the switch at the last bootup).

NOTE: The **slice** option is not supported on QFX Series switches that have no alternate slice when Junos OS boots as a Virtual Machine (VM). To switch to the previous version of Junos OS, issue the **request system software rollback** command.

usb—(Optional) Reboot from a USB device.

The following options are available only on QFabric Systems:

all—(Optional) Reboots the software on the Director group, fabric control Routing Engines, fabric manager Routing Engines, Interconnect devices, and network and server Node groups.

director-device name—(Optional) Reboots the software on the Director device and the default partition (QFabric CLI).

director-group—(Optional) Reboots the software on the Director group and the default partition (QFabric CLI).

fabric—(Optional) Reboots the fabric control Routing Engines and the Interconnect devices.

node-group name—(Optional) Reboots the software on a server Node group or a network Node group.

graceful—(Optional) Enables the QFabric component to reboot with minimal impact to network traffic. This sub-option is only available for the **all**, **fabric**, and **director-group** options.

Additional Information

Reboot requests are recorded in the system log files, which you can view with the **show log** command (see [show log](#)). Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the **show system processes** command (see *show system processes*).

On a TX Matrix or TX Matrix Plus router, if you issue the **request system reboot** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are rebooted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are rebooted.

NOTE: Before issuing the **request system reboot** command on a TX Matrix Plus router with no options or the **all-chassis**, **all-lcc**, **lcc number**, or **sfc** options, verify that master Routing Engine for all routers in the routing matrix are in the same slot number. If the master Routing Engine for a line-card chassis is in a different slot number than the master Routing Engine for a TX Matrix Plus router, the line-card chassis might become logically disconnected from the routing matrix after the **request system reboot** command.

NOTE: To reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) first, and then reboot the master Routing Engine.

Required Privilege Level
maintenance

RELATED DOCUMENTATION

clear system reboot

[request system halt](#) | [621](#)

[Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

request vmhost reboot

List of Sample Output

[request system reboot on page 634](#)
[request system reboot \(at 2300\) on page 635](#)
[request system reboot \(in 2 Hours\) on page 635](#)
[request system reboot \(Immediately\) on page 635](#)
[request system reboot \(at 1:20 AM\) on page 635](#)
[request system reboot in-service on page 635](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system reboot

```
user@host> request system reboot
```



```
Reboot the system ? [yes,no] (no)
```

request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
```

```
Reboot the system ? [yes,no] (no) yes
```

```
shutdown: [pid 186]
```

```
*** System shutdown message from root@test.example.net ***
```

```
System going down at 23:00
```

request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
```

```
user@host> request system reboot in 120
```

```
user@host> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@host> request system reboot at now
```

request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
```

```
request system reboot at 120
```

```
Reboot the system at 120? [yes,no] (no) yes
```

request system reboot in-service

```
user@switch> request system reboot in-service
```



```

Reboot the system ? [yes,no]
[Feb 22 02:37:04]:ISSU: Validating Image

PRE ISSR CHECK:
-----
PFE Status                : Online
Member Id zero            : Valid
VC not in mixed or fabric mode : Valid
Member is single node vc   : Valid
BFD minimum-interval check done : Valid
GRES enabled              : Valid
NSR enabled               : Valid
drop-all-tcp not configured : Valid
Ready for ISSR            : Valid

warning: Do NOT use /user during ISSR. Changes to /user during ISSR may get lost!
Current image is jinstall-jcp-i386-flex-18.1.img
[Feb 22 02:37:14]:ISSU: Preparing Backup RE
Prepare for ISSR
[Feb 22 02:37:19]:ISSU: Backup RE Prepare Done
Spawning the backup RE
Spawn backup RE, index 1 successful
Starting secondary dataplane
Second dataplane container started
GRES in progress
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade for ISSR
Chassis ISSU Started
[Feb 22 02:42:55]:ISSU: Preparing Daemons
[Feb 22 02:43:00]:ISSU: Daemons Ready for ISSU
[Feb 22 02:43:05]:ISSU: Starting Upgrade for FRUs
[Feb 22 02:43:15]:ISSU: FPC Warm Booting
[Feb 22 02:44:16]:ISSU: FPC Warm Booted
[Feb 22 02:44:27]:ISSU: Preparing for Switchover
[Feb 22 02:44:31]:ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item                Status                Reason
  FPC 0                Online (ISSU)
Send ISSR done to chassisd on backup RE
Chassis ISSU Completed
Removing dcpfe0 eth1 128.168.0.16 IP

```



```
Bringing down bme00
Post Chassis ISSU processing done
[Feb 22 02:44:33]:ISSU: IDLE
Stopping primary dataplane
Clearing ISSU states
Console and management sessions will be disconnected. Please login again.
device_handoff successful ret: 0
Shutdown NOW!
[pid 14305]

*** FINAL System shutdown message from root@sw-duckhorn-01 ***

System going down IMMEDIATELY
```


request system software format-qfabric-backup

Syntax

```
request system software format-qfabric-backup
```

Release Information

Command introduced in Junos OS Release 13.1 for the QFX Series.

Description

(QFabric systems only) Copy the install media files from a USB flash drive to your QFabric system recovery directory on the Director group. You must issue this command before you can use the **request system software system-backup** and **request system software system-backup copy-to-usb** commands.

Options

none—Copy the install media files from a USB flash drive to a Director group recovery directory.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[Performing System Backup and Recovery for a QFabric System | 581](#)

[request system software system-backup | 647](#)

[Performing a QFabric System Recovery Installation on the Director Group | 835](#)

List of Sample Output

[request system software format-qfabric-backup on page 638](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
request system software format-qfabric-backup
```

```
user@qfabric> request system software format-qfabric-backup
```

```
Copying QFabric USB template image from /dev/sdb(Unigen,PQS4000,4009 MB).....
```


request system software nonstop-upgrade

Syntax

```
request system software nonstop-upgrade package-name
<fabric >
<director-group>
<node-group name>
```

Release Information

Command introduced in Junos OS Release 12.2 for the QFX Series.

Description

Nonstop software upgrade enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. You should upgrade the devices in the following order: Director group, fabric controls and Interconnect devices, and network and server Node groups.

NOTE: Before you perform a nonstop software upgrade, contact JTAC to perform a pre-upgrade health check on the QFabric system.

Options

package-name—Location from which the software is to be installed. For example:

- ***protocol://hostname/pathname/package-name***—For a software package or bundle that is to be downloaded and installed from a remote location. Replace ***protocol*** with one of the following:
 - **ftp**—File Transfer Protocol.
Use **ftp://hostname/pathname/package-name**. To specify authentication credentials, use **ftp://<username>:<password>@hostname/pathname/package-name**. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
 - **http**—Hypertext Transfer Protocol.
Use **http://hostname/pathname/package-name**. To specify authentication credentials, use **http://<username>:<password>@hostname/pathname/package-name**. If a password is required and you omit it, you are prompted for it.
 - **scp**—Secure copy (available only for Canada and U.S. version).

Use `scp://hostname/pathname/package-name`. To specify authentication credentials, use `scp://<username>:<password>@hostname/pathname/package-name`.

NOTE:

- The **pathname** in the protocol is the relative path to the user's home directory on the remote system and not the root directory.

director-group—Install software package on the Director group and Fabric managers.

fabric—Install software package on the Interconnect devices and Fabric controls.

node-group name —Install software package on the redundant server Node group, server Node group, or network Node group.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[Nonstop Software Upgrade Checklist for QFabric Systems | 529](#)

[Performing a Nonstop Software Upgrade on the QFabric System | 533](#)

[Verifying Nonstop Software Upgrade for QFabric Systems | 540](#)

[show chassis nonstop-upgrade node-group | 778](#)

List of Sample Output

[request system software nonstop-upgrade director-group on page 640](#)

[request system software nonstop-upgrade fabric on page 643](#)

[request system software nonstop-upgrade node-group \(Redundant Server Node Group\) on page 644](#)

[request system software nonstop-upgrade node-group \(Server Node Group\) on page 646](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
request system software nonstop-upgrade director-group
```

```
user@qfabric> request system software nonstop-upgrade director-group
```

```
jinstall-qfabric-12.2X50-D10.3.rpm
```



```

[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one
of rolling upgrade
[Peer Update Status]: Waiting for peer to complete phase one of rolling upgrade
[Peer Update Status]: Peer completed phase one of rolling upgrade
Setting peer DG node as the master SFC

```

```

Delaying start of local upgrade to allow peer services time to initialize [15
minutes]
Delaying start of local upgrade to allow peer services time to initialize [15
minutes]
Delaying start of local upgrade to allow peer services time to initialize [12
minutes]
Delaying start of local upgrade to allow peer services time to initialize [9
minutes]
Delaying start of local upgrade to allow peer services time to initialize [6
minutes]
Delaying start of local upgrade to allow peer services time to initialize [3
minutes]
[Peer Update Status]: Check for VMs on dg0
Triggering Final Stage of Fabric Manager Upgrade:
Updating FM-0 to Junos version 12.2X50-D10.3
[Status 2012-06-05 16:10:12]: Fabric Manager: Upgrade Final Stage started
[NW-NG-0 2012-06-05 16:10:22]: Transferring NW-NG-0 Mastership to REMOTE DG
[NW-NG-0 2012-06-05 16:11:44]: Finished NW-NG-0 Mastership switch

```



```
[Status 2012-06-05 16:11:45]: Upgrading FM-0 VM on worker DG to 12.2X50-D10.3
[DRE-0 2012-06-05 16:12:43]: Retrieving package
[DRE-0 2012-06-05 16:13:46]: ----- re0: -----
[Status 2012-06-05 16:15:17]: Load completed with 0 errors...
[Status 2012-06-05 16:15:17]: Reboot is required to complete upgrade ...
[DRE-0 2012-06-05 16:15:22]: Waiting for DRE-0 to terminate ...
[DRE-0 2012-06-05 16:15:34]: Waiting for DRE-0 to come back ...
[DRE-0 2012-06-05 16:18:44]: Running Uptime Test for DRE-0
[DRE-0 2012-06-05 16:18:51]: Uptime Test for DRE-0 Passed ...
[Status 2012-06-05 16:18:51]: DRE-0 booted successfully ...
Performing post install shutdown and cleanup
```

```
Broadcast message from root (Tue Jun 5 16:18:51 2012):
```

```
The system is going down for reboot NOW!
Director group upgrade complete
```

```
root@qfabric> Read from remote host qfabric-partition0: Connection reset by peer
Connection to qfabric-partition0 closed.
```

request system software nonstop-upgrade fabric

```
user@qfabric> request system software nonstop-upgrade fabric jinstall-qfabric-12.2X50-D10.3.rpm
```

```
[FC-0 2012-06-05 16:48:53]: Retrieving package
[FC-1 2012-06-05 16:48:53]: Retrieving package
[IC-F4912 2012-06-05 16:48:59]: Retrieving package
[FC-0 2012-06-05 16:49:51]: ----- re0: -----
[FC-1 2012-06-05 16:49:52]: ----- re0: -----
[IC-F4912 2012-06-05 16:49:54]: ----- re0: -----
[IC-F4912 2012-06-05 16:50:42]: Step 1 of 20 Creating temporary file system
[IC-F4912 2012-06-05 16:50:42]: Step 2 of 20 Determining installation source
[IC-F4912 2012-06-05 16:50:43]: Step 3 of 20 Processing format options
[IC-F4912 2012-06-05 16:50:43]: Step 4 of 20 Determining installation slice
[IC-F4912 2012-06-05 16:50:43]: Step 5 of 20 Creating and labeling new slices
[IC-F4912 2012-06-05 16:50:44]: Step 6 of 20 Create and mount new file system
[IC-F4912 2012-06-05 16:50:53]: Step 7 of 20 Getting OS bundles
[IC-F4912 2012-06-05 16:50:53]: Step 8 of 20 Updating recovery media
[IC-F4912 2012-06-05 16:51:17]: Step 9 of 20 Extracting incoming image
[IC-F4912 2012-06-05 16:52:56]: Step 10 of 20 Unpacking OS packages
[IC-F4912 2012-06-05 16:52:59]: Step 11 of 20 Mounting jbase package
[IC-F4912 2012-06-05 16:53:28]: Step 12 of 20 Creating base OS symbolic links
[IC-F4912 2012-06-05 16:54:45]: Step 13 of 20 Creating fstab
[IC-F4912 2012-06-05 16:54:45]: Step 14 of 20 Creating new system files
```



```

[IC-F4912 2012-06-05 16:54:46]: Step 15 of 20 Adding jbundle package
[IC-F4912 2012-06-05 16:58:15]: Step 16 of 20 Backing up system data
[IC-F4912 2012-06-05 16:58:18]: Step 17 of 20 Setting up shared partition data
[IC-F4912 2012-06-05 16:58:18]: Step 18 of 20 Checking package sanity in
installation
[IC-F4912 2012-06-05 16:58:18]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[IC-F4912 2012-06-05 16:58:22]: Step 20 of 20 Setting da0s1 as new active partition
[Status 2012-06-05 16:58:34]: Load completed with 0 errors...
[Status 2012-06-05 16:58:34]: Reboot is required to complete upgrade ...
[Status 2012-06-05 16:58:34]: Trying to Connect to Node: FC-0
[Status 2012-06-05 16:58:39]: Rebooting FC-0
[Status 2012-06-05 16:58:39]: Trying to Connect to Node: FC-1
[Status 2012-06-05 16:58:44]: Rebooting FC-1
[Status 2012-06-05 16:58:44]: Trying to Connect to Node: IC-F4912
[Status 2012-06-05 16:58:50]: Rebooting IC-F4912
Success

```

request system software nonstop-upgrade node-group (Redundant Server Node Group)

user@qfabric> request system software nonstop-upgrade node-group RSNG

jinstall-qfabric-12.2X50-D10.3.rpm

Upgrading target(s): RSNG

```

[RSNG 2012-06-05 17:26:44]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
[RSNG 2012-06-05 17:26:44]: Retrieving package
[RSNG 2012-06-05 17:28:56]: Pushing bundle to fpc1
[RSNG 2012-06-05 17:29:26]: fpc1: Validate package...
[RSNG 2012-06-05 17:35:22]: fpc0: Validate package...
[RSNG 2012-06-05 17:35:49]: ----- fpc1 -----
[RSNG 2012-06-05 17:36:25]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:36:26]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:36:26]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:36:26]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:36:27]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:36:27]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:36:35]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:36:35]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:36:56]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:38:07]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:38:16]: Step 11 of 20 Mounting jbase package

```



```

[RSNG      2012-06-05 17:38:41]: Step 12 of 20 Creating base OS symbolic links
[RSNG      2012-06-05 17:39:41]: Step 13 of 20 Creating fstab
[RSNG      2012-06-05 17:39:42]: Step 14 of 20 Creating new system files
[RSNG      2012-06-05 17:39:42]: Step 15 of 20 Adding jbundle package
[RSNG      2012-06-05 17:42:16]: Step 16 of 20 Backing up system data
[RSNG      2012-06-05 17:42:32]: Step 17 of 20 Setting up shared partition data
[RSNG      2012-06-05 17:42:33]: Step 18 of 20 Checking package sanity in
installation
[RSNG      2012-06-05 17:42:33]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[RSNG      2012-06-05 17:42:36]: Step 20 of 20 Setting da0s2 as new active partition
[RSNG      2012-06-05 17:42:51]: ----- fpc0 - master -----
[RSNG      2012-06-05 17:42:51]: Step 1 of 20 Creating temporary file system
[RSNG      2012-06-05 17:42:51]: Step 2 of 20 Determining installation source
[RSNG      2012-06-05 17:42:51]: Step 3 of 20 Processing format options
[RSNG      2012-06-05 17:42:51]: Step 4 of 20 Determining installation slice
[RSNG      2012-06-05 17:42:51]: Step 5 of 20 Creating and labeling new slices
[RSNG      2012-06-05 17:42:51]: Step 6 of 20 Create and mount new file system
[RSNG      2012-06-05 17:42:51]: Step 7 of 20 Getting OS bundles
[RSNG      2012-06-05 17:42:51]: Step 8 of 20 Updating recovery media
[RSNG      2012-06-05 17:42:51]: Step 9 of 20 Extracting incoming image
[RSNG      2012-06-05 17:42:51]: Step 10 of 20 Unpacking OS packages
[RSNG      2012-06-05 17:42:51]: Step 11 of 20 Mounting jbase package
[RSNG      2012-06-05 17:42:51]: Step 12 of 20 Creating base OS symbolic links
[RSNG      2012-06-05 17:42:51]: Step 13 of 20 Creating fstab
[RSNG      2012-06-05 17:42:51]: Step 14 of 20 Creating new system files
[RSNG      2012-06-05 17:42:51]: Step 15 of 20 Adding jbundle package
[RSNG      2012-06-05 17:42:51]: Step 16 of 20 Backing up system data
[RSNG      2012-06-05 17:42:51]: Step 17 of 20 Setting up shared partition data
[RSNG      2012-06-05 17:42:51]: Step 18 of 20 Checking package sanity in
installation
[RSNG      2012-06-05 17:42:51]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[RSNG      2012-06-05 17:42:51]: Step 20 of 20 Setting da0s2 as new active partition
[RSNG      2012-06-05 17:43:36]: Rebooting Backup RE
[RSNG      2012-06-05 17:43:36]: ----- Rebooting fpc1 -----
[RSNG      2012-06-05 17:50:12]: Initiating Chassis In-Service-Upgrade
[RSNG      2012-06-05 17:50:33]: Upgrading group: 0 fpc: 0
[RSNG      2012-06-05 17:52:38]: Upgrade complete for group:0
[RSNG      2012-06-05 17:52:38]: Upgrading group: 1 fpc: 1
[RSNG      2012-06-05 17:54:42]: Upgrade complete for group:1
[RSNG      2012-06-05 17:54:42]: Finished processing all upgrade groups, last group
:1
[RSNG      2012-06-05 17:54:48]: Preparing for Switchover

```



```
[RSNG      2012-06-05 17:55:38]: Switchover Completed
[Status    2012-06-05 17:55:41]: Upgrade completed with 0 errors
Success
```

request system software nonstop-upgrade node-group (Server Node Group)

```
user@qfabric> request system software nonstop-upgrade node-group P1507-C
jinstall-qfabric-12.2X50-D10.3.rpm
```

```
Upgrading target(s): P1507-C
```

```
[P1507-C  2012-06-26 14:02:44]: Retrieving package
[P1507-C  2012-06-26 14:03:21]: ----- P1507-C: -----
[P1507-C  2012-06-26 14:03:59]: Step 1 of 20 Creating temporary file system
[P1507-C  2012-06-26 14:03:59]: Step 2 of 20 Determining installation source
[P1507-C  2012-06-26 14:03:59]: Step 3 of 20 Processing format options
[P1507-C  2012-06-26 14:03:59]: Step 4 of 20 Determining installation slice
[P1507-C  2012-06-26 14:04:00]: Step 5 of 20 Creating and labeling new slices
[P1507-C  2012-06-26 14:04:00]: Step 6 of 20 Create and mount new file system
[P1507-C  2012-06-26 14:04:08]: Step 7 of 20 Getting OS bundles
[P1507-C  2012-06-26 14:04:09]: Step 8 of 20 Updating recovery media
[P1507-C  2012-06-26 14:04:29]: Step 9 of 20 Extracting incoming image
[P1507-C  2012-06-26 14:05:42]: Step 10 of 20 Unpacking OS packages
[P1507-C  2012-06-26 14:05:49]: Step 11 of 20 Mounting jbase package
[P1507-C  2012-06-26 14:06:14]: Step 12 of 20 Creating base OS symbolic links
[P1507-C  2012-06-26 14:07:15]: Step 13 of 20 Creating fstab
[P1507-C  2012-06-26 14:07:15]: Step 14 of 20 Creating new system files
[P1507-C  2012-06-26 14:07:16]: Step 15 of 20 Adding jbundle package
[P1507-C  2012-06-26 14:09:52]: Step 16 of 20 Backing up system data
[P1507-C  2012-06-26 14:10:07]: Step 17 of 20 Setting up shared partition data
[P1507-C  2012-06-26 14:10:07]: Step 18 of 20 Checking package sanity in
installation
[P1507-C  2012-06-26 14:10:08]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[P1507-C  2012-06-26 14:10:11]: Step 20 of 20 Setting da0s2 as new active partition
[Status    2012-06-26 14:10:25]: Trying to Connect to Node: P1507-C
[Status    2012-06-26 14:10:32]: Rebooting P1507-C
[Status    2012-06-26 14:10:32]: Upgrade completed with 0 errors
Success
```


request system software system-backup

Syntax

```
request system software system-backup  
<usb-create>
```

Release Information

Command introduced in Junos OS Release 13.1 for the QFX Series.

Description

(QFabric systems only) Save a copy of the current QFabric system configuration file and the current software package for recovery purposes. You can use these saved files to restore your QFabric system to full operation after a system failure or shutdown.

Options

none—Copy the QFabric system software package and system configuration file to a Director group recovery directory.

NOTE: If this command fails, insert a Juniper Networks software installation USB flash drive into the master Director device and issue the **request system software format-qfabric-backup** command. For more details about this prerequisite procedure that is required before you can use the QFabric system backup and recovery feature, see [“Performing System Backup and Recovery for a QFabric System” on page 581](#).

usb-create—Copy the QFabric system software package and system configuration file from the Director group recovery directory to a USB flash drive. When the files have been copied, you can use the USB flash drive to help your QFabric system recover from a failure condition.

NOTE: You must issue the **request system software system-backup** command (which saves the files to the Director group) before you can issue the **request system software system-backup usb-create** command.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[Performing System Backup and Recovery for a QFabric System | 581](#)

[request system software format-qfabric-backup | 638](#)

save

request system software configuration-backup

[Performing a QFabric System Recovery Installation on the Director Group | 835](#)

List of Sample Output

[request system software system-backup on page 648](#)

[request system software system-backup usb-create on page 648](#)

Output Fields

When you enter these commands, you are provided feedback on the status of your request.

Sample Output

request system software system-backup

user@qfabric> **request system software system-backup**

```
user@qfabric>
```

request system software system-backup usb-create

user@qfabric> **request system software system-backup usb-create /dev/sdb**

```
Issuing this command will overwrite the contents of the USB drive.
Continue? [yes,no] (no) yes
```

```
This operation will access the USB drive on 0281042010000013.
Are you sure you want to continue? [yes,no] (no) yes
```

```
Copying QFabric recovery media to /dev/sdb...
Successfully copied QFabric recovery media to /dev/sdb
```


set chassis display message

List of Syntax

[Syntax on page 649](#)

[Syntax \(TX Matrix Router\) on page 649](#)

[Syntax \(TX Matrix Plus Router\) on page 649](#)

Syntax

```
set chassis display message "message"
<permanent>
```

Syntax (TX Matrix Router)

```
set chassis display message "message" (lcc number | scc)
<permanent>
```

Syntax (TX Matrix Plus Router)

```
set chassis display message "message " (fpc-slot slot-number | lcc number | sfc number)
<permanent>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

sfc option for TX Matrix Plus router introduced in Junos OS Release 9.6.

Description

Display or stop a text message on the craft interface display, which is on the front of the router, or on the LCD panel display on the switch. The craft interface alternates the display of text messages with standard craft interface messages three times, switching between messages every 60 seconds.

NOTE: On T Series routers, when this command is executed with the **permanent** option, the display of the text message alternates with that of the standard craft interface message continuously every 60 seconds.

By default, on both the router and the switch, the text message is displayed for 5 minutes. The craft interface display has four 20-character lines. The LCD panel display has two 16-character lines, and text messages appear only on the second line.

Options

"message"—Message to display. On the craft interface display, if the message is longer than 20 characters, it wraps onto the next line. If a word does not fit on one line, the entire word moves down to the next line. Any portion of the message that does not fit on the display is truncated. An empty pair of quotation marks (" ") deletes the text message from the craft interface display. On the LCD panel display, the message is limited to 16 characters.

fpc-slot slot-number—(TX Matrix Plus routers and EX4200 and QFX Series only) On the router or switch, display the text message on the craft interface for a specific Flexible PIC Concentrator (FPC). Replace **slot-number** with a value from **0** through **31**. On the switch, display the text message for a specific member of a Virtual Chassis, where **fpc-slot slot-number** corresponds to the member ID. Replace **slot-number** with a value from **0** through **9**. On the QFX Series, the **slot-number** is always **0**. On a TX Matrix Plus router with 3D SIBs replace **slot-number** with a value from **0** through **63**.

lcc number—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

permanent—(Optional) Display a text message on the craft interface display or LCD panel display permanently.

scc—(TX Matrix routers only) Display the text message on the craft interface display of the TX Matrix router (switch-card chassis).

sfc number—(TX Matrix Plus routers only) Display the text message on the craft interface display of the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level

clear

RELATED DOCUMENTATION

Configuring the LCD Panel on EX Series Switches (CLI Procedure)

clear chassis display message

show chassis craft-interface

List of Sample Output

[set chassis display message \(Creating\) on page 651](#)

[set chassis display message \(Deleting\) on page 651](#)

Output Fields

See *show chassis craft-interface* for an explanation of output fields.

Sample Output

set chassis display message (Creating)

The following example shows how to set the display message and verify the result:

```
user@host> set chassis display message "NOC contact Dusty (888) 555-1234"
```

```
message sent
```

```
user@host> show chassis craft-interface
```

```
Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
Host OK LED:    On
Host fail LED:  Off
FPCs           0  1  2  3  4  5  6  7
-----
Green  ..  *..  *  *.
Red    .....
LCD screen:
+-----+
|NOC contact Dusty  |
| (888) 555-1234    |
+-----+
```

set chassis display message (Deleting)

The following example shows how to delete the display message and verify that the message is removed:

```
user@host> set chassis display message ""
```

```
message sent
```

```
user@host> show chassis craft-interface
```



```

Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
Host OK LED:    On
Host fail LED:  Off
FPCs           0  1  2  3  4  5  6  7

```

```
-----
```

```
Green  ..  *..  *  *.
```

```
Red    .....
```

```
LCD screen:
```

```

+-----+
|host           |
|Up: 0+17:05:47 |
|               |
|Temperature OK |
+-----+

```


show chassis device-mode

Syntax

```
show chassis device-mode
```

Release Information

Command introduced in Junos OS Release 11.2 for the QFX Series.

Description

Display information about the operating mode of the device. For example, QFX3500 devices operate either as a single switch in standalone mode or as a QFabric system Node device in node-device mode.

NOTE: Issue the **show chassis device-mode** command only when your management station is connected directly to the device over a console port connection.

Options

There are no options for this command.

Required Privilege Level

admin

RELATED DOCUMENTATION

Converting the Device Mode for a QFabric System Component 322
request chassis device-mode 600
Understanding Interconnect Devices 29
Understanding Node Devices 33

List of Sample Output

- [show chassis device-mode \(Interconnect Device Mode\) on page 654](#)
- [show chassis device-mode \(Interconnect Device Mode, but Node Device-Ready\) on page 654](#)
- [show chassis device-mode \(Interconnect Device Mode, but Standalone-Ready\) on page 654](#)
- [show chassis device-mode \(Node Device Mode\) on page 655](#)
- [show chassis device-mode \(Node Device Mode, but Interconnect Device-Ready\) on page 655](#)
- [show chassis device-mode \(Node Device Mode, but Standalone-Ready\) on page 655](#)
- [show chassis device-mode \(Standalone Mode\) on page 655](#)
- [show chassis device-mode \(Standalone Mode, but Interconnect Device-Ready\) on page 655](#)

[show chassis device-mode \(Standalone Mode, but Node Device-Ready\) on page 655](#)

Output Fields

[Table 82 on page 654](#) lists the output fields for the **show chassis device-mode** command. Output fields are listed in the approximate order in which they appear.

Table 82: show chassis device-mode Output Fields

Field Name	Field Description
Current device-mode	Existing operational mode for the device. The device can be in Interconnect device mode, Node device mode, or standalone mode.
Future device-mode after reboot	Future operational mode for the device after you reboot it. The device can be set to enter Interconnect device mode, Node device mode, or standalone mode. NOTE: To set the future mode of the device, issue the request chassis device-mode command.

Sample Output

show chassis device-mode (Interconnect Device Mode)

```
user@switch> show chassis device-mode
```

```
Current device-mode : Interconnect-device
Future device-mode after reboot : Interconnect-device
```

show chassis device-mode (Interconnect Device Mode, but Node Device-Ready)

```
user@switch> show chassis device-mode
```

```
Current device-mode : Interconnect-device
Future device-mode after reboot : Node-device
```

show chassis device-mode (Interconnect Device Mode, but Standalone-Ready)

```
user@switch> show chassis device-mode
```

```
Current device-mode : Interconnect-device
Future device-mode after reboot : Standalone
```


show chassis device-mode (Node Device Mode)

```
user@switch> show chassis device-mode
```

```
Current device-mode : Node-device  
Future device-mode after reboot : Node-device
```

show chassis device-mode (Node Device Mode, but Interconnect Device-Ready)

```
user@switch> show chassis device-mode
```

```
Current device-mode : Node-device  
Future device-mode after reboot : Interconnect-device
```

show chassis device-mode (Node Device Mode, but Standalone-Ready)

```
user@switch> show chassis device-mode
```

```
Current device-mode : Node-device  
Future device-mode after reboot : Standalone
```

show chassis device-mode (Standalone Mode)

```
user@switch> show chassis device-mode
```

```
Current device-mode : Standalone  
Future device-mode after reboot : Standalone
```

show chassis device-mode (Standalone Mode, but Interconnect Device-Ready)

```
user@switch> show chassis device-mode
```

```
Current device-mode : Standalone  
Future device-mode after reboot : Interconnect-device
```

show chassis device-mode (Standalone Mode, but Node Device-Ready)

```
user@switch> show chassis device-mode
```

```
Current device-mode : Standalone  
Future device-mode after reboot : Node-device
```


show chassis environment cb

List of Syntax

[Syntax on page 656](#)

[Syntax \(TX Matrix Routers\) on page 656](#)

[Syntax \(TX Matrix Plus Routers\) on page 656](#)

[Syntax \(MX Series Routers\) on page 656](#)

[Syntax \(MX104 Universal Routing Platforms\) on page 656](#)

[Syntax \(MX2010, MX2020, MX10003, MX204, MX2008, and MX10008 Universal Routing Platforms; EX9251 and EX9253 Switches\) on page 657](#)

[Syntax \(QFabric System\) on page 657](#)

Syntax

```
show chassis environment cb  
<slot>
```

Syntax (TX Matrix Routers)

```
show chassis environment cb  
<lcc number | scc>  
<slot>
```

Syntax (TX Matrix Plus Routers)

```
show chassis environment cb  
<lcc number | sfc number >  
<slot>
```

Syntax (MX Series Routers)

```
show chassis environment cb  
<slot>  
<all-members>  
<local>  
<member member-id>
```

Syntax (MX104 Universal Routing Platforms)

```
show chassis environment cb
```


Syntax (MX2010, MX2020, MX10003, MX204, MX2008, and MX10008 Universal Routing Platforms; EX9251 and EX9253 Switches)

```
show chassis environment cb
<slot>
```

Syntax (QFabric System)

```
show chassis environment cb
<slot interconnect-device interconnect-device-name>
< interconnect-device interconnect-device-name slot>
```

Release Information

Command introduced before Junos Release 7.4.

Command introduced in Junos OS Release 9.4 for EX Series switches.

Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

Command introduced in Junos OS Release 12.1 for T4000 Core Routers.

option introduced for the TX Matrix Plus router in Junos Release 9.6.

Command introduced in Junos OS Release 11.3 for the QFX Series.

Command introduced in Junos OS Release 12.3 for MX2010 and MX2020 Universal Routing Platforms.

Command introduced in Junos OS Release 13.2 for MX104 Universal Routing Platforms.

Command introduced in Junos OS Release 17.2 for MX2008 Universal Routing Platforms and PTX10008 Routers.

Command introduced in Junos OS Release 17.3 for MX10003 Universal Routing Platforms.

Command introduced in Junos OS Release 17.4 for MX204 Universal Routing Platforms.

Command introduced in Junos OS Release 18.1R1 for EX9251 switches.

Command introduced in Junos OS Release 18.2R1 for MX10008 Universal Routing Platforms, and EX9253 switches.

Description

(M120, M320, MX Series, and T Series routers, EX8200 switches, and PTX Series Packet Transport Routers only) Display environmental information about the Control Boards (CBs).

Options

none—Display environmental information about all CBs. For a TX Matrix router, display environmental information about all CBs on the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display environmental information about all CBs on the TX Matrix Plus router and its attached T1600 or T4000 routers.

all-members—(MX Series routers only) (Optional) Display environmental information about the CBs on all the members of the Virtual Chassis configuration.

interconnect-device—(QFabric systems only) Display environmental information about CBs on the Interconnect device.

lcc *number*—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display environmental information about the CBs on the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display environmental information about the CBs on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

scc—(TX Matrix router only) (Optional) Display environmental information about the CBs in the TX Matrix router (switch-card chassis).

sfc *number*—(TX Matrix Plus router only) (Optional) Display environmental information about the CBs in the TX Matrix Plus router (or switch-fabric chassis).

slot—(Optional) Display environmental information about the specified CB. On routers and PTX Series Packet Transport Routers, replace *slot* with **0** or **1**. On EX Series switches replace *slot* with **0**, **1**, or **2**. On QFX Series switches, replace *slot* with **0** or **1**.

Required Privilege Level

view

RELATED DOCUMENTATION

[request chassis cb](#) | 596

Understanding Switching Control Board Redundancy

List of Sample Output

[show chassis environment cb \(M120 Router\) on page 660](#)

[show chassis environment cb \(M320 Router\) on page 661](#)

[show chassis environment cb \(MX80 Router\) on page 662](#)

[show chassis environment cb \(MX104 Router\) on page 662](#)
[show chassis environment cb \(MX240 Router\) on page 663](#)
[show chassis environment cb \(MX240 Router with Enhanced MX SCB\) on page 663](#)
[show chassis environment cb \(MX480 Router\) on page 664](#)
[show chassis environment cb \(MX480 Router with Enhanced MX SCB\) on page 665](#)
[show chassis environment cb \(MX960 Router\) on page 665](#)
[show chassis environment cb \(MX960 Router with Enhanced MX SCB\) on page 666](#)
[show chassis environment cb \(MX2020 Router\) on page 667](#)
[show chassis environment cb \(MX2010 Router\) on page 668](#)
[show chassis environment cb \(MX2008 Router\) on page 669](#)
[show chassis environment cb \(MX10003 Router\) on page 670](#)
[show chassis environment cb \(MX204 Router\) on page 672](#)
[show chassis environment cb \(MX10008 Router\) on page 673](#)
[show chassis environment cb \(T4000 Core Router\) on page 674](#)
[show chassis environment cb \(TX Matrix Router\) on page 675](#)
[show chassis environment cb \(TX Matrix Plus Router\) on page 676](#)
[show chassis environment cb \(EX8200 Switch\) on page 681](#)
[show chassis environment cb \(EX8208 Switch\) on page 683](#)
[show chassis environment cb \(EX9251 Switch\) on page 684](#)
[show chassis environment cb \(EX9253 Switch\) on page 686](#)
[show chassis environment cb \(PTX5000 Packet Transport Router\) on page 687](#)
[show chassis environment cb \(PTX10008 Router\) on page 687](#)
[show chassis environment cb \(PTX10016 Router\) on page 689](#)
[show chassis environment cb \(QFabric System\) on page 690](#)
[show chassis environment cb 0 \(PTX10003-80C and PTX10003-160C Routers\) on page 691](#)

Output Fields

Table 83 on page 659 lists the output fields for the **show chassis environment cb** command. Output fields are listed in the approximate order in which they appear.

Table 83: show chassis environment cb Output Fields

Field Name	Field Description
State	<p>Status of the CB. If two CBs are installed and online, one is functioning as the master, and the other is the standby.</p> <ul style="list-style-type: none"> • Online—CB is online and running. • Offline— CB is powered down. <p>NOTE: On the EX8208 switch, the installation can include three CBs.</p>

Table 83: show chassis environment cb Output Fields (*continued*)

Field Name	Field Description
Temperature	<p>Temperature in Celsius (C) and Fahrenheit (F) of the air flowing past the CB.</p> <ul style="list-style-type: none"> • Temperature Intake—Measures the temperature of the air intake to cool the power supplies. • Temperature Exhaust—Measures the temperature of the hot air exhaust. <p>NOTE: On the MX2010, MX2020, and MX2008 routers, the intake temperature measures the temperature of the air intake to cool the Control Board (CB). The MX2010, MX2020, and MX2008 routers include intake and exhaust temperatures for multiple zones (Intake A, Intake B, Intake C, Exhaust A, Exhaust B, and TCBC).</p>
Power	Power required and measured on the CB. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
BUS Revision	Revision level of the generic bus device. (Not on switches.)
FPGA Revision	Revision level of the field-programmable gate array (FPGA). (Not on switches.)
PMBus device (on MX240, MX480, and MX960 routers with Enhanced MX SCB)	<p>Enhanced SCB on MX 240, MX480, and MX960 routers allows the system to save power by supplying only the amount of voltage that is required. Configurable PMBus devices are used to provide the voltage for each individual device. There is one PMBus device for each XF ASIC so that the output can be customized to each device. The following PMBus device information is displayed for routers with Enhanced MX SCB:</p> <ul style="list-style-type: none"> • Expected voltage • Measured voltage • Measured current • Calculated power

Sample Output

show chassis environment cb (M120 Router)

user@host> **show chassis environment cb**

```

CB 0 status:
  State                Online Master
  Temperature          33 degrees C / 91 degrees F
  Power
    1.2 V              1214 mV
    1.5 V              1495 mV

```



```

    2.5 V                2494 mV
    3.3 V                3319 mV
    5.0 V                5085 mV
    3.3 V bias          3296 mV
    Bus Revision         12
    FPGA Revision        17
CB 1 status:
  State                  Online Standby
  Temperature            34 degrees C / 93 degrees F
  Power:
    1.2 V                1195 mV
    1.5 V                1495 mV
    2.5 V                2504 mV
    3.3 V                3312 mV
    5.0 V                5111 mV
    3.3 V bias          3296 mV
    Bus Revision         12
    FPGA Revision        17

```

show chassis environment cb (M320 Router)

user@host> **show chassis environment cb**

```

CB 0 status:
  State                  Online Master
  Temperature            29 degrees C / 84 degrees F
  Power:
    1.8 V                1805 mV
    2.5 V                2501 mV
    3.3 V                3293 mV
    4.6 V                4725 mV
    5.0 V                5032 mV
    12.0 V               11975 mV
    3.3 V bias          3286 mV
    8.0 V bias          7589 mV
    BUS Revision         40
    FPGA Revision        7
CB 1 status:
  State                  Online Standby
  Temperature            32 degrees C / 89 degrees F
  Power:
    1.8 V                1802 mV
    2.5 V                2482 mV
    3.3 V                3289 mV

```


4.6 V	4720 mV
5.0 V	5001 mV
12.0 V	11946 mV
3.3 V bias	3274 mV
8.0 V bias	7562 mV
BUS Revision	40
FPGA Revision	7

show chassis environment cb (MX80 Router)

user@host> show chassis environment cb

```

CB 0 status:
  State                Online Master
  Temperature          36 degrees C / 96 degrees F
  Power 1
    1.0 V              1034 mV
    1.0 V MQ           1037 mV
    1.0 V LU           1005 mV
    1.2 V              1218 mV
    1.5 V              1524 mV
    1.8 V              1814 mV
    2.5 V              2558 mV
    3.3 V              3296 mV
    5.0 V              5233 mV
    5.0 V bias         5207 mV
    12.0 V             12162 mV

```

show chassis environment cb (MX104 Router)

user@host > show chassis environment cb

```

CB 0 status:
  State                Online Master
  Temperature          33 degrees C / 91 degrees F
  Power 1
    0.75 V             751 mV
    1.0 V              1005 mV
    1.1 V              1113 mV
    1.5 V              1494 mV
    2.5 V              2518 mV
    3.3 V              3338 mV
    5.0 V              4960 mV

```



```

    12.0 V                12006 mV
    FPGA Revision         25
    CB 1 status:
    State                  Empty

```

show chassis environment cb (MX240 Router)

```
user@host> show chassis environment cb
```

```

CB 0 status:
  State                  Online Standby
  Temperature             37 degrees C / 98 degrees F
  Power 1
    1.2 V                1208 mV
    1.5 V                1521 mV
    1.8 V                1811 mV
    2.5 V                2513 mV
    3.3 V                3332 mV
    5.0 V                5059 mV
    12.0 V               12162 mV
    1.25 V               1260 mV
    3.3 V SM3            3306 mV
    5.0 V RE             5085 mV
    12.0 V RE            11872 mV
  Power 2
    11.3 V bias PEM      11272 mV
    4.6 V bias MidPlane  4827 mV
    11.3 V bias FPD      11272 mV
    11.3 V bias POE 0    11292 mV
    11.3 V bias POE 1    11253 mV
  Bus Revision           42
  FPGA Revision          1

```

show chassis environment cb (MX240 Router with Enhanced MX SCB)

```
user@host> show chassis environment cb
```

```

CB 0 status:
  State                  Online Standby
  Temperature             37 degrees C / 98 degrees F
  Power 1
    1.2 V                1208 mV
    1.5 V                1521 mV

```



```

1.8 V          1811 mV
2.5 V          2513 mV
3.3 V          3332 mV
5.0 V          5059 mV
12.0 V         12162 mV
1.25 V         1260 mV
3.3 V SM3      3306 mV
5.0 V RE       5085 mV
12.0 V RE      11872 mV
Power 2
11.3 V bias PEM 11272 mV
4.6 V bias MidPlane 4827 mV
11.3 V bias FPD 11272 mV
11.3 V bias POE 0 11292 mV
11.3 V bias POE 1 11253 mV
Bus Revision    42
FPGA Revision   1
PMBus           Expected Measured Measured Calculated
device          voltage  voltage  current  power
XF ASIC A       1000 mV   997 mV   11031 mA 10997 mW
XF ASIC B       1000 mV   996 mV   12125 mA 12076 mW

```

show chassis environment cb (MX480 Router)

```
user@host> show chassis environment cb
```

```

CB 0 status:
State          Online Master
Temperature    41 degrees C / 105 degrees F
Power 1
1.2 V          1202 mV
1.5 V          1511 mV
1.8 V          1798 mV
2.5 V          2507 mV
3.3 V          3312 mV
5.0 V          5027 mV
12.0 V         12200 mV
1.25 V         1260 mV
3.3 V SM3      3293 mV
5 V RE         5040 mV
12 V RE        11910 mV
Power 2
11.3 V bias PEM 11156 mV
4.6 V bias MidPlane 4801 mV

```



```

11.3 V bias FPD          11214 mV
11.3 V bias POE 0        11098 mV
11.3 V bias POE 1        11330 mV
Bus Revision              42
FPGA Revision             1

```

show chassis environment cb (MX480 Router with Enhanced MX SCB)

```
user@host> show chassis environment cb
```

```

CB 0 status:
State                      Online Master
Temperature                 41 degrees C / 105 degrees F
Power 1
  1.2 V                     1202 mV
  1.5 V                     1511 mV
  1.8 V                     1798 mV
  2.5 V                     2507 mV
  3.3 V                     3312 mV
  5.0 V                     5027 mV
  12.0 V                    12200 mV
  1.25 V                    1260 mV
  3.3 V SM3                 3293 mV
  5 V RE                    5040 mV
  12 V RE                   11910 mV
Power 2
  11.3 V bias PEM           11156 mV
  4.6 V bias MidPlane       4801 mV
  11.3 V bias FPD           11214 mV
  11.3 V bias POE 0         11098 mV
  11.3 V bias POE 1         11330 mV
Bus Revision                42
FPGA Revision               1
PMBus                       Expected  Measured  Measured  Calculated
device                      voltage   voltage   current   power
XF ASIC A                   1000 mV   997 mV   11031 mA  10997 mW
XF ASIC B                   1000 mV   996 mV   12125 mA  12076 mW

```

show chassis environment cb (MX960 Router)

```
user@host> show chassis environment cb
```



```
CB 0 status:
  State                Online Master
  Temperature          24 degrees C / 75 degrees F
  Power 1
    1.2 V              1965 mV
    1.5 V              2465 mV
    1.8 V              2990 mV
    2.5 V              3296 mV
    3.3 V              3296 mV
    5.0 V              6593 mV
    12.0 V             13187 mV
    3.3 V bias         3296 mV
    1.25 V             1994 mV
    3.3 V SM3          3296 mV
    5 V RE             6593 mV
    12 V RE            13174 mV
  Power 2              Sensor failure
  Bus Revision         4
  FPGA Revision        3
```

show chassis environment cb (MX960 Router with Enhanced MX SCB)

user@host> **show chassis environment cb**

```
CB 0 status:
  State                Online Master
  Temperature          24 degrees C / 75 degrees F
  Power 1
    1.2 V              1965 mV
    1.5 V              2465 mV
    1.8 V              2990 mV
    2.5 V              3296 mV
    3.3 V              3296 mV
    5.0 V              6593 mV
    12.0 V             13187 mV
    3.3 V bias         3296 mV
    1.25 V             1994 mV
    3.3 V SM3          3296 mV
    5 V RE             6593 mV
    12 V RE            13174 mV
  Power 2              Sensor failure
  Bus Revision         4
  FPGA Revision        3
  PMBus                Expected   Measured   Measured   Calculated
```


device	voltage	voltage	current	power
XF ASIC A	1000 mV	997 mV	11031 mA	10997 mW
XF ASIC B	1000 mV	996 mV	12125 mA	12076 mW

show chassis environment cb (MX2020 Router)

user@host> show chassis environment cb

CB 0 status:

```

State                               Online Master
IntakeA-Zone0 Temperature           44 degrees C / 111 degrees F
IntakeB-Zone1 Temperature           34 degrees C / 93 degrees F
IntakeC-Zone0 Temperature           45 degrees C / 113 degrees F
ExhaustA-Zone0 Temperature          43 degrees C / 109 degrees F
ExhaustB-Zone1 Temperature          36 degrees C / 96 degrees F
TCBC-Zone0 Temperature              39 degrees C / 102 degrees F
Power 1
  1.0 V                             1011 mV
  1.2 V                             1208 mV
  1.8 V                             1801 mV
  2.5 V                             2552 mV
  3.3 V                             3312 mV
  5.0 V                             5040 mV
  5.0 V RE                           4988 mV
  12.0 V                             12065 mV
  12.0 V RE                           12046 mV
Bus Revision                         99
FPGA Revision                        270

```

CB 1 status:

```

State                               Online Standby
IntakeA-Zone0 Temperature           45 degrees C / 113 degrees F
IntakeB-Zone1 Temperature           41 degrees C / 105 degrees F
IntakeC-Zone0 Temperature           46 degrees C / 114 degrees F
ExhaustA-Zone0 Temperature          44 degrees C / 111 degrees F
ExhaustB-Zone1 Temperature          41 degrees C / 105 degrees F
TCBC-Zone0 Temperature              45 degrees C / 113 degrees F
Power 1
  1.0 V                             1008 mV
  1.2 V                             1208 mV
  1.8 V                             1798 mV
  2.5 V                             2539 mV
  3.3 V                             3325 mV
  5.0 V                             5033 mV
  5.0 V RE                           4950 mV

```


12.0 V	12046 mV
12.0 V RE	11968 mV
Bus Revision	99
FPGA Revision	0

show chassis environment cb (MX2010 Router)

user@host> show chassis environment cb

CB 0 status:

State	Online Master
IntakeA-Zone0 Temperature	36 degrees C / 96 degrees F
IntakeB-Zone1 Temperature	30 degrees C / 86 degrees F
IntakeC-Zone0 Temperature	38 degrees C / 100 degrees F
ExhaustA-Zone0 Temperature	36 degrees C / 96 degrees F
ExhaustB-Zone1 Temperature	32 degrees C / 89 degrees F
TCBC-Zone0 Temperature	34 degrees C / 93 degrees F
Power 1	
1.0 V	1015 mV
1.2 V	1205 mV
1.8 V	1804 mV
2.5 V	2552 mV
3.3 V	3325 mV
5.0 V	5020 mV
5.0 V RE	4988 mV
12.0 V	12104 mV
12.0 V RE	12026 mV
Bus Revision	100
FPGA Revision	270

CB 1 status:

State	Online
IntakeA-Zone0 Temperature	35 degrees C / 95 degrees F
IntakeB-Zone1 Temperature	28 degrees C / 82 degrees F
IntakeC-Zone0 Temperature	37 degrees C / 98 degrees F
ExhaustA-Zone0 Temperature	34 degrees C / 93 degrees F
ExhaustB-Zone1 Temperature	29 degrees C / 84 degrees F
TCBC-Zone0 Temperature	32 degrees C / 89 degrees F
Power 1	
1.0 V	1011 mV
1.2 V	1208 mV
1.8 V	1788 mV
2.5 V	2526 mV
3.3 V	3319 mV
5.0 V	5046 mV


```

5.0 V RE          4975 mV
12.0 V            12046 mV
12.0 V RE        12007 mV
Bus Revision      100
FPGA Revision     0

```

show chassis environment cb (MX2008 Router)

```
user@host> show chassis environment cb
```

```

CB 0 status:
State                Online Master
Inlet1 Temperature   37 degrees C / 98 degrees F
Inlet2 Temperature   45 degrees C / 113 degrees F
Inlet3 Temperature   44 degrees C / 111 degrees F
Inlet4 Temperature   42 degrees C / 107 degrees F
Exhaust1 Temperature 30 degrees C / 86 degrees F
Exhaust2 Temperature 40 degrees C / 104 degrees F
Exhaust3 Temperature 48 degrees C / 118 degrees F
Exhaust4 Temperature 46 degrees C / 114 degrees F
Power 1
  1.0 V PHY          989 mV
  1.15 V             1150 mV
  1.2 V bias         1189 mV
  1.5 V              1488 mV
  1.8 V              1772 mV
  2.5 V              2462 mV
  3.3 V bias         3296 mV
  VCCIO              1028 mV
Power 2
  1.1 V              1099 mV
  3.3 V              3300 mV
Power 3
  0.95 V XL710       949 mV
  1.05 V             1050 mV
Power 4
  1.2 V              1199 mV
  5.0 V              4999 mV
Power 5
  1.0 V              1000 mV
  1.2 V PHY          1199 mV
Bus Revision         114
FPGA Revision        1
CB 1 status:

```



```

State                               Online Standby
Inlet1 Temperature                  30 degrees C / 86 degrees F
Inlet2 Temperature                  31 degrees C / 87 degrees F
Inlet3 Temperature                  29 degrees C / 84 degrees F
Inlet4 Temperature                  32 degrees C / 89 degrees F
Exhaust1 Temperature                30 degrees C / 86 degrees F
Exhaust2 Temperature                33 degrees C / 91 degrees F
Exhaust3 Temperature                34 degrees C / 93 degrees F
Exhaust4 Temperature                35 degrees C / 95 degrees F
Power 1
  1.0 V PHY                         986 mV
  1.15 V                           1153 mV
  1.2 V bias                        1195 mV
  1.5 V                             1498 mV
  1.8 V                             1798 mV
  2.5 V                             2494 mV
  3.3 V bias                        3296 mV
  VCCIO                             1034 mV
Power 2
  1.1 V                             1100 mV
  3.3 V                             3300 mV
Power 3
  0.95 V XL710                      949 mV
  1.05 V                            1050 mV
Power 4
  1.2 V                             1199 mV
  5.0 V                             5000 mV
Power 5
  1.0 V                             1000 mV
  1.2 V PHY                         1199 mV
Bus Revision                        114
FPGA Revision                       1

```

show chassis environment cb (MX10003 Router)

```
user@host> show chassis environment cb
```

```

CB 0 status:
  State                               Online Master
CB 0 Exhaust Temp Sensor 0x49 35 degrees C / 95 degrees F
CB 0 Inlet Temp Sensor 0x49 28 degrees C / 82 degrees F
Power
  VDD1V5_PCH                         1489 mV

```


VDDIO	940 mV
VDD3V3_PCH	3332 mV
VDD2V5_AB	2508 mV
VDD1V8_CLC	1764 mV
VDD3V3	3292 mV
VDD2V5_CD	2508 mV
VDD1V2_CBC_GTX	0 mV
VDD1V8_GLS_GTX	0 mV
VDD1V2_CBC	0 mV
VDD1V8_GLS	0 mV
BIAS3V3_BP	0 mV
VDD1V2_GH	1200 mV
VDD3V3_CBC	3299 mV
VDD1V2_CD	1199 mV
BIAS3V3	3340 mV
VDD1V2_AB	1200 mV
VDD5V0	5000 mV
VDD1V05	1049 mV
VDD1V05	1050 mV
VCORE	1780 mV
12V	12272 mV
	3952 mA
	48984 mW

CB 1 status:

State Online Standby

CB 1 Exhaust Temp Sensor 0x49 35 degrees C / 95 degrees F

CB 1 Inlet Temp Sensor 0x49 31 degrees C / 87 degrees F

Power

VDD1V5_PCH	1489 mV
VDDIO	940 mV
VDD3V3_PCH	3351 mV
VDD2V5_AB	2508 mV
VDD1V8_CLC	1764 mV
VDD3V3	3312 mV
VDD2V5_CD	2508 mV
VDD1V2_CBC_GTX	1195 mV
VDD1V8_GLS_GTX	1764 mV
VDD1V2_CBC	1195 mV
VDD1V8_GLS	1783 mV
BIAS3V3_BP	4096 mV
VDD1V2_GH	1200 mV
VDD3V3_CBC	3300 mV
VDD1V2_CD	1200 mV
BIAS3V3	3339 mV
VDD1V2_AB	1200 mV
VDD5V0	5000 mV

VDD1V05	1050 mV		
VDD1V05	1050 mV		
VCORE	1780 mV		
12V	12351 mV	3823 mA	45007 mW

show chassis environment cb (MX204 Router)

user@host> show chassis environment cb

CB 0 status:

State	Online Master
CB 0 Top Right Inlet Sensor	35 degrees C / 95 degrees F
CB 0 Top Left Inlet Sensor	37 degrees C / 98 degrees F
CB 0 Top Right Exhaust Sensor	43 degrees C / 109 degrees F
CB 0 Top Left Exhaust Sensor	50 degrees C / 122 degrees F
CB 0 CPU Core-0 Temp	48 degrees C / 118 degrees F
CB 0 CPU Core-1 Temp	48 degrees C / 118 degrees F
CB 0 CPU Core-2 Temp	48 degrees C / 118 degrees F
CB 0 CPU Core-3 Temp	47 degrees C / 116 degrees F
CB 0 CPU Core-4 Temp	47 degrees C / 116 degrees F
CB 0 CPU Core-5 Temp	47 degrees C / 116 degrees F
CB 0 CPU Core-6 Temp	47 degrees C / 116 degrees F
CB 0 CPU Core-7 Temp	47 degrees C / 116 degrees F

Power

VDD1V5_PCH	1509 mV
VDDIO	950 mV
VDD3V3_PCH	3312 mV
VDD2V5_AB	2508 mV
VDD1V8_FRMR	1813 mV
VDD3V3	3312 mV
VDD2V5_CD	2508 mV
VDD1V8_PLL	1813 mV
VDD1V5	1499 mV
EA0_1V5	1499 mV
EA0_1V04	1038 mV
EA0_PLL_1V0	999 mV
EA0_2V5	2508 mV
BIAS3V	3332 mV
VDD1V2_CD	1214 mV
VDD1V2_AB	1215 mV
VDD1V05	1050 mV
BIAS3V3	3309 mV
VDD1V0	1015 mV

VDD1V8	1804 mV		
VDD1V2	1199 mV		
VDD2V5	2504 mV		
EA0_VDD0V9	949 mV		
EA0_HM1_VDD0V9	899 mV		
EA0_VDD0V9R2	952 mV		
EA0_VDD1V0	1000 mV		
VDD3V3	3304 mV		
EA0_XR_VDD1V2	1199 mV		
EA0_XR_VDD0V9	903 mV		
EA0_HM_VDDM1V2	1199 mV		
EA0_HM_VDD1V2	1199 mV		
VDDCPU0	1770 mV		
12V Hotswap A	11968 mV	4696 mA	55466 mW
12V Hotswap B	12048 mV	14936 mA	180652 mW

show chassis environment cb (MX10008 Router)

user@host> show chassis environment cb

```

CB 0 status:
  State                               Online Master
CB 0 Intake A Temp Sensor             24 degrees C / 75 degrees F
CB 0 Intake B Temp Sensor             24 degrees C / 75 degrees F
CB 0 Exhaust A Temp Sensor            28 degrees C / 82 degrees F
CB 0 Exhaust B Temp Sensor            30 degrees C / 86 degrees F
CB 0 Middle Temp Sensor               28 degrees C / 82 degrees F
Power
  GESW_VDD1V0                         1000 mV
  VDD1V0                              1000 mV
  VDD1V2                              1199 mV
  VDD3V3                              3299 mV
  XL710_VCCD                          950 mV
  VDD1V05                             1050 mV
  VDD2V5                              2500 mV
  FPGA_VDD1V2                         1200 mV
  VDD1V8                              1800 mV
  VDD1V15                             1150 mV
  VDD1V1                              1099 mV
  VCCIO                               950 mV
  PHY_VDD1V0                          1000 mV
  VDD5V0                              4998 mV
  FPGA_VDD1V5                         1496 mV
  VDD1V5                              1496 mV

```



```

12V                                12281 mV    7700 mA    92400 mW
CB 1 status:
  State                            Online Standby
CB 1 Intake A Temp Sensor          24 degrees C / 75 degrees F
CB 1 Intake B Temp Sensor          23 degrees C / 73 degrees F
CB 1 Exhaust A Temp Sensor         27 degrees C / 80 degrees F
CB 1 Exhaust B Temp Sensor         30 degrees C / 86 degrees F
CB 1 Middle Temp Sensor            28 degrees C / 82 degrees F
Power
  GESW_VDD1V0                      999 mV
  VDD1V0                           1000 mV
  VDD1V2                           1199 mV
  VDD3V3                           3299 mV
  XL710_VCCD                       950 mV
  VDD1V05                          1050 mV
  VDD2V5                           2499 mV
  FPGA_VDD1V2                      1200 mV
  VDD1V8                           1799 mV
  VDD1V15                          1150 mV
  VDD1V1                           1100 mV
  VCCIO                             949 mV
  PHY_VDD1V0                       999 mV
  VDD5V0                           5000 mV
  FPGA_VDD1V5                      1502 mV
  VDD1V5                           1496 mV
12V                                12281 mV    8002 mA    96024 mW
```

show chassis environment cb (T4000 Core Router)

user@host> **show chassis environment cb**

```

CB 0 status:
  State                            Online Master
  Temperature                      33 degrees C / 91 degrees F
  Power 1
    1.8 V                          1805 mV
    2.5 V                          2523 mV
    3.3 V                          3324 mV
    3.3 V bias                     3296 mV
    4.6 V                          4680 mV
    5.0 V                          4893 mV
    8.0 V bias                     7572 mV
    12.0 V                         11916 mV
  Power 2
```



```

    1.0 V          993 mV
    1.2 V          1210 mV
    3.3 V RE       3330 mV
  Bus Revision    51
  FPGA Revision   5
CB 1 status:
  State           Online Standby
  Temperature      33 degrees C / 91 degrees F
  Power 1
    1.8 V          1810 mV
    2.5 V          2496 mV
    3.3 V          3308 mV
    3.3 V bias     3286 mV
    4.6 V          4692 mV
    5.0 V          4954 mV
    8.0 V bias     7282 mV
    12.0 V         11926 mV
  Power 2
    1.0 V          993 mV
    1.2 V          1185 mV
    3.3 V RE       3316 mV
  Bus Revision    51
  FPGA Revision   5

```

show chassis environment cb (TX Matrix Router)

```
user@host> show chassis environment cb
```

```

-----
CB 0 status:
  State           Online Master
  Temperature      32 degrees C / 89 degrees F
  Power:
    1.8 V          1797 mV
    2.5 V          2477 mV
    3.3 V          3311 mV
    4.6 V          4727 mV
    5.0 V          5015 mV
    12.0 V         12185 mV
    3.3 V bias     3304 mV
    8.0 V bias     7870 mV
  BUS Revision     40
  FPGA Revision    1
CB 1 status:

```



```

    State                               Online Standby
...

lcc0-re0:
-----

CB 0 status:
    State                               Online Master
    Temperature                         32 degrees C / 89 degrees F
    Power:
        1.8 V                           1787 mV
        2.5 V                           2473 mV
        3.3 V                           3306 mV
        4.6 V                           4793 mV
        5.0 V                           5025 mV
        12.0 V                          12156 mV
        3.3 V bias                       3289 mV
        8.0 V bias                       7609 mV
    BUS Revision                        40
    FPGA Revision                       5
CB 1 status:
    State                               Online Standby
....
    BUS Revision                        40
    FPGA Revision                       5

lcc2-re0:
-----

CB 0 status:
    State                               Online Master
...
CB 1 status:
    State                               Online Standby
...

```

show chassis environment cb (TX Matrix Plus Router)

user@host> show chassis environment cb

```

sfc0-re0:
-----

CB 0 status:
    State                               Online Master
    Temperature                         38 degrees C / 100 degrees F
    Power 1

```



```

1.0 V          1005 mV
1.1 V          1108 mV
1.2 V          1205 mV
1.25 V         1269 mV
1.5 V          1508 mV
1.8 V          1814 mV
2.5 V          2507 mV
3.3 V          3306 mV
3.3 V bias     3300 mV
9.0 V          9058 mV
9.0 V RE       9107 mV
Power 2
3.9 V          3963 mV
5.0 V          5020 mV
9.0 V          9087 mV
Bus Revision   79
FPGA Revision  23
CB 1 status:
State          Online Standby
Temperature     39 degrees C / 102 degrees F
Power 1
1.0 V          1002 mV
1.1 V          1105 mV
1.2 V          1198 mV
1.25 V         1276 mV
1.5 V          1504 mV
1.8 V          1804 mV
2.5 V          2507 mV
3.3 V          3300 mV
3.3 V bias     3293 mV
9.0 V          9039 mV
9.0 V RE       9049 mV
Power 2
3.9 V          3892 mV
5.0 V          5040 mV
9.0 V          9058 mV
Bus Revision   79
FPGA Revision  23

lcc0-re0:
-----
CB 0 status:
State          Online Master
Temperature     39 degrees C / 102 degrees F

```



```

Power 1
  1.8 V          1799 mV
  2.5 V          2499 mV
  3.3 V          3327 mV
  3.3 V bias     3299 mV
  4.6 V          4673 mV
  5.0 V          4918 mV
  8.0 V bias     7308 mV
  12.0 V         11887 mV
Power 2
  1.0 V          996 mV
  1.2 V          1199 mV
  3.3 V RE       3319 mV
Bus Revision     51
FPGA Revision    3
CB 1 status:
  State          Online Standby
  Temperature     40 degrees C / 104 degrees F
Power 1
  1.8 V          1800 mV
  2.5 V          2496 mV
  3.3 V          3322 mV
  3.3 V bias     3284 mV
  4.6 V          4680 mV
  5.0 V          4954 mV
  8.0 V bias     7284 mV
  12.0 V         11902 mV
Power 2
  1.0 V          998 mV
  1.2 V          1205 mV
  3.3 V RE       3327 mV
Bus Revision     51
FPGA Revision    3

lcc1-re0:
-----
CB 0 status:
  State          Online Master
  Temperature     41 degrees C / 105 degrees F
Power 1
  1.8 V          1804 mV
  2.5 V          2517 mV
  3.3 V          3300 mV
  3.3 V bias     3284 mV

```


4.6 V	4681 mV
5.0 V	4927 mV
8.0 V bias	7357 mV
12.0 V	11907 mV
Power 2	
1.0 V	991 mV
1.2 V	1202 mV
3.3 V RE	3301 mV
Bus Revision	51
FPGA Revision	3
CB 1 status:	
State	Online Standby
Temperature	40 degrees C / 104 degrees F
Power 1	
1.8 V	1805 mV
2.5 V	2528 mV
3.3 V	3324 mV
3.3 V bias	3289 mV
4.6 V	4694 mV
5.0 V	4959 mV
8.0 V bias	7311 mV
12.0 V	11926 mV
Power 2	
1.0 V	998 mV
1.2 V	1200 mV
3.3 V RE	3313 mV
Bus Revision	51
FPGA Revision	3

lcc2-re0:

CB 0 status:	
State	Online Master
Temperature	41 degrees C / 105 degrees F
Power 1	
1.8 V	1805 mV
2.5 V	2494 mV
3.3 V	3333 mV
3.3 V bias	3296 mV
4.6 V	4673 mV
5.0 V	4901 mV
8.0 V bias	7343 mV
12.0 V	11916 mV
Power 2	


```

    1.0 V          993 mV
    1.2 V          1213 mV
    3.3 V RE       3328 mV
Bus Revision      51
FPGA Revision     3
CB 1 status:
State             Online Standby
Temperature        41 degrees C / 105 degrees F
Power 1
    1.8 V          1804 mV
    2.5 V          2523 mV
    3.3 V          3334 mV
    3.3 V bias     3291 mV
    4.6 V          4697 mV
    5.0 V          4969 mV
    8.0 V bias     7308 mV
    12.0 V         11936 mV
Power 2
    1.0 V          996 mV
    1.2 V          1200 mV
    3.3 V RE       3328 mV
Bus Revision      51
FPGA Revision     3

```

```
lcc3-re0:
```

```

-----
CB 0 status:
State             Online Master
Temperature        37 degrees C / 98 degrees F
Power 1
    1.8 V          1809 mV
    2.5 V          2510 mV
    3.3 V          3296 mV
    3.3 V bias     3291 mV
    4.6 V          4670 mV
    5.0 V          4905 mV
    8.0 V bias     7211 mV
    12.0 V         11882 mV
Power 2
    1.0 V          996 mV
    1.2 V          1188 mV
    3.3 V RE       3326 mV
Bus Revision      51
FPGA Revision     5

```



```

CB 1 status:
  State                Online Standby
  Temperature          38 degrees C / 100 degrees F
  Power 1
    1.8 V              1813 mV
    2.5 V              2510 mV
    3.3 V              3322 mV
    3.3 V bias         3289 mV
    4.6 V              4692 mV
    5.0 V              4967 mV
    8.0 V bias         7194 mV
    12.0 V             11916 mV
  Power 2
    1.0 V              996 mV
    1.2 V              1205 mV
    3.3 V RE           3273 mV
  Bus Revision         51
  FPGA Revision        5

```

show chassis environment cb (EX8200 Switch)

```
user@host> show chassis environment cb
```

```

CB 0 status:
  State                Online Master
  Temperature Intake    20 degrees C / 68 degrees F
  Temperature Exhaust   24 degrees C / 75 degrees F
  Power 1
    1.1 V              1086 mV
    1.2 V              1179 mV
    1.2 V *            1182 mV
    1.2 V *            1182 mV
    1.25 V             1211 mV
    1.5 V              1472 mV
    1.8 V              1756 mV
    2.5 V              2449 mV
    3.3 V              3254 mV
    3.3 V bias         3300 mV
    5.0 V              4911 mV
    12.0 V             11891 mV
  Power 2
    3.3 V bias *       3615 mV
    3.3 V bias *       3615 mV

```



```

3.3 V bias *      3567 mV
3.3 V bias *      3664 mV
4.3 V bias *      4224 mV
4.3 V bias *      4215 mV
4.3 V bias *      4224 mV
4.3 V bias *      4205 mV
4.3 V bias *      4195 mV
4.3 V bias *      4215 mV
5.0 V bias        4920 mV

CB 1 status:
State              Online Standby
Temperature Intake  19 degrees C / 66 degrees F
Temperature Exhaust 23 degrees C / 73 degrees F
Power 1
1.1 V              1082 mV
1.2 V              1169 mV
1.2 V *            1179 mV
1.2 V *            1179 mV
1.25 V             1214 mV
1.5 V              1482 mV
1.8 V              1759 mV
2.5 V              2481 mV
3.3 V              3248 mV
3.3 V bias         3306 mV
5.0 V              4911 mV
12.0 V             11910 mV

Power 2
3.3 V bias *      3644 mV
3.3 V bias *      3664 mV
3.3 V bias *      3586 mV
3.3 V bias *      3654 mV
4.3 V bias *      4224 mV
4.3 V bias *      4215 mV
4.3 V bias *      4224 mV
4.3 V bias *      4205 mV
4.3 V bias *      4244 mV
4.3 V bias *      4215 mV
5.0 V bias        4930 mV

CB 2 status:
State              Online
Temperature Intake  19 degrees C / 66 degrees F
Temperature Exhaust 23 degrees C / 73 degrees F
Power 1
1.2 V              1195 mV

```


1.5 V	1511 mV
1.8 V	1804 mV
2.5 V	2526 mV
3.3 V	3300 mV
3.3 V bias	3306 mV
12.0 V	12220 mV

show chassis environment cb (EX8208 Switch)

user@host> show chassis environment cb

```

CB 0 status:
  State                               Online Master
  Temperature Intake                   20 degrees C / 68 degrees F
  Temperature Exhaust                  24 degrees C / 75 degrees F
  Power 1
    1.1 V                             1086 mV
    1.2 V                             1179 mV
    1.2 V *                           1182 mV
    1.2 V *                           1182 mV
    1.25 V                            1211 mV
    1.5 V                             1466 mV
    1.8 V                             1759 mV
    2.5 V                             2455 mV
    3.3 V                             3261 mV
    3.3 V bias                        3300 mV
    5.0 V                             4930 mV
    12.0 V                            11891 mV
  Power 2
    3.3 V bias *                      3606 mV
    3.3 V bias *                      3615 mV
    3.3 V bias *                      3567 mV
    3.3 V bias *                      3673 mV
    4.3 V bias *                      4224 mV
    4.3 V bias *                      4215 mV
    4.3 V bias *                      4234 mV
    4.3 V bias *                      4205 mV
    4.3 V bias *                      4186 mV
    4.3 V bias *                      4215 mV
    5.0 V bias                        4940 mV
CB 1 status:
  State                               Online Standby
  Temperature Intake                   19 degrees C / 66 degrees F
  Temperature Exhaust                  23 degrees C / 73 degrees F

```



```

Power 1
  1.1 V          1086 mV
  1.2 V          1169 mV
  1.2 V *        1179 mV
  1.2 V *        1179 mV
  1.25 V         1211 mV
  1.5 V          1479 mV
  1.8 V          1759 mV
  2.5 V          2475 mV
  3.3 V          3235 mV
  3.3 V bias     3306 mV
  5.0 V          4930 mV
  12.0 V         11891 mV
Power 2
  3.3 V bias *   3644 mV
  3.3 V bias *   3664 mV
  3.3 V bias *   3586 mV
  3.3 V bias *   3654 mV
  4.3 V bias *   4215 mV
  4.3 V bias *   4224 mV
  4.3 V bias *   4215 mV
  4.3 V bias *   4215 mV
  4.3 V bias *   4234 mV
  4.3 V bias *   4224 mV
  5.0 V bias     4920 mV
CB 2 status:
State           Online
Temperature Intake  20 degrees C / 68 degrees F
Temperature Exhaust 24 degrees C / 75 degrees F
Power 1
  1.2 V          1202 mV
  1.5 V          1508 mV
  1.8 V          1804 mV
  2.5 V          2520 mV
  3.3 V          3300 mV
  3.3 V bias     3300 mV
  12.0 V         12200 mV

```

show chassis environment cb (EX9251 Switch)

```
user@switch> show chassis environment cb
```

```

CB 0 status:
State           Online Master

```



```

CB 0 Top Right Inlet Sensor29 degrees C / 84 degrees F
CB 0 Top Left Inlet Sensor 28 degrees C / 82 degrees F
CB 0 Top Right Exhaust Sensor40 degrees C / 104 degrees F
CB 0 Top Left Exhaust Sensor59 degrees C / 138 degrees F
CB 0 CPU Core-0 Temp          45 degrees C / 113 degrees F
CB 0 CPU Core-1 Temp          44 degrees C / 111 degrees F
CB 0 CPU Core-2 Temp          44 degrees C / 111 degrees F
CB 0 CPU Core-3 Temp          44 degrees C / 111 degrees F
CB 0 CPU Core-4 Temp          45 degrees C / 113 degrees F
CB 0 CPU Core-5 Temp          44 degrees C / 111 degrees F
CB 0 CPU Core-6 Temp          44 degrees C / 111 degrees F
CB 0 CPU Core-7 Temp          43 degrees C / 109 degrees F

```

Power

VDD1V5_PCH	1499 mV
VDDIO	950 mV
VDD3V3_PCH	3312 mV
VDD2V5_AB	2489 mV
VDD1V8_FRMR	1793 mV
VDD3V3	3292 mV
VDD2V5_CD	2508 mV
VDD1V8_PLL	1793 mV
VDD1V5	1499 mV
EA0_1V5	1499 mV
EA0_1V04	999 mV
EA0_PLL_1V0	999 mV
EA0_2V5	2508 mV
BIAS3V	3292 mV
VDD1V2_CD	1215 mV
VDD1V2_AB	1214 mV
VDD1V05	1050 mV
BIAS3V3	3309 mV
VDD1V0	1014 mV
VDD1V8	1805 mV
VDD1V2	1200 mV
VDD2V5	2504 mV
EA0_VDD0V9	949 mV
EA0_HM1_VDD0V9	899 mV
EA0_VDD0V9R2	952 mV
EA0_VDD1V0	999 mV
VDD3V3	3305 mV
EA0_XR_VDD1V2	1199 mV
EA0_XR_VDD0V9	903 mV
EA0_HM_VDDM1V2	1199 mV
EA0_HM_VDD1V2	1199 mV

VDDCPU0	1770 mV			
12V Hotswap A	11955 mV	4861 mA	59347 mW	
12V Hotswap B	11916 mV	15046 mA	180887 mW	

show chassis environment cb (EX9253 Switch)

user@switch> show chassis environment cb

```

CB 0 status:
  State                               Online Master
CB 0 Exhaust Temp Sensor             38 degrees C / 100 degrees F
CB 0 Inlet Temp Sensor                32 degrees C / 89 degrees F
CB 0 CPU DIE Temp Sensor              43 degrees C / 109 degrees F
Power
  VDD1V5_PCH                         1489 mV
  VDDIO                              940 mV
  VDD3V3_PCH                         3332 mV
  VDD2V5_AB                          2508 mV
  VDD1V8_CLC                         1783 mV
  VDD3V3                             3312 mV
  VDD2V5_CD                          2508 mV
  VDD1V2_CBC_GTX                     1195 mV
  VDD1V8_GLS_GTX                     1783 mV
  VDD1V2_CBC                         1176 mV
  VDD1V8_GLS                         1783 mV
  BIAS3V3_BP                         3978 mV
  VDD1V2_GH                          1200 mV
  VDD3V3_CBC                         3299 mV
  VDD1V2_CD                          1200 mV
  BIAS3V3                            3340 mV
  VDD1V2_AB                          1199 mV
  VDD5V0                             5000 mV
  VDD1V05                            1050 mV
  VDD1V05                            1050 mV
  VCORE                             1770 mV
  12V                                12061 mV    4806 mA    57841 mW

CB 1 status:
  State                               Offline
CB 1 Exhaust Temp Sensor             32 degrees C / 89 degrees F
CB 1 Inlet Temp Sensor                29 degrees C / 84 degrees F
CB 1 CPU DIE Temp Sensor              43 degrees C / 109 degrees F
Power                                Disabled

```


show chassis environment cb (PTX5000 Packet Transport Router)

```
user@host> show chassis environment cb
```

```
CB 0 status:
  State                               Online Master
  Intake Temperature                  38 degrees C / 100 degrees F
  Exhaust A Temperature               45 degrees C / 113 degrees F
  Exhaust B Temperature               42 degrees C / 107 degrees F
  Power 1
    1.2 V                             1200 mV
    1.25 V                            1250 mV
    2.5 V                             2500 mV
    3.3 V                             3300 mV
  Power 2
    1.0 V                             1000 mV
    3.3 V bias                        3293 mV
    3.9 V                             3921 mV
  Bus Revision                        132
  FPGA Revision                       27
CB 1 status:
  State                               Online Standby
  Intake Temperature                  34 degrees C / 93 degrees F
  Exhaust A Temperature               39 degrees C / 102 degrees F
  Exhaust B Temperature               36 degrees C / 96 degrees F
  Power 1
    1.2 V                             1199 mV
    1.25 V                            1250 mV
    2.5 V                             2499 mV
    3.3 V                             3299 mV
  Power 2
    1.0 V                             1000 mV
    3.3 V bias                        3312 mV
    3.9 V                             3961 mV
  Bus Revision                        132
  FPGA Revision                       28
```

show chassis environment cb (PTX10008 Router)

```
user@host> show chassis environment cb
```

```
CB 0 status:
  State                               Online Master
  CB 0 Intake Temp Sensor             28 degrees C / 82 degrees F
  CB 0 Exhaust Temp Sensor            32 degrees C / 89 degrees F
```


Power

VDD 2.5V	2489 mV		
Bias 3.3V	3332 mV		
VDD 3.3V	3292 mV		
VCC 1.8V	1822 mV		
VDD 1.2V	1205 mV		
VCC 1V	999 mV		
VCC CPU 1.8V	1803 mV		
VDD 2.5V	2489 mV		
VCC Aux 5V	5115 mV		
VDD DDR 1.5V	1499 mV		
VTT SA CPU 0.8V	803 mV		
VTT CPU 1.05V	1048 mV		
VCC Core CPU	901 mV		
VCC PCH 1.5V	1519 mV		
VDD 1.05V	1048 mV		
VCC 2.5V	2508 mV		
FORT VCCA 1V	960 mV		
VDD .85V	862 mV		
VTT DDRA .75V	744 mV		
VTT DDRB .75V	744 mV		
12V	12285 mV	3779 mA	46339 mW

CB 1 status:

State	Online Standby
CB 1 Intake Temp Sensor	27 degrees C / 80 degrees F
CB 1 Exhaust Temp Sensor	32 degrees C / 89 degrees F

Power

VDD 2.5V	2489 mV
Bias 3.3V	3332 mV
VDD 3.3V	3273 mV
VCC 1.8V	1822 mV
VDD 1.2V	1195 mV
VCC 1V	999 mV
VCC CPU 1.8V	1783 mV
VDD 2.5V	2489 mV
VCC Aux 5V	5056 mV
VDD DDR 1.5V	1499 mV
VTT SA CPU 0.8V	793 mV
VTT CPU 1.05V	1048 mV
VCC Core CPU	882 mV
VCC PCH 1.5V	1509 mV
VDD 1.05V	1048 mV
VCC 2.5V	2489 mV
FORT VCCA 1V	960 mV

VDD .85V	862 mV		
VTT DDRA .75V	744 mV		
VTT DDRB .75V	744 mV		
12V	12391 mV	3779 mA	46727 mW

show chassis environment cb (PTX10016 Router)

user@host> show chassis environment cb

```

CB 0 status:
  State                               Online Master
CB 0 Intake Temp Sensor               20 degrees C / 68 degrees F
CB 0 Exhaust Temp Sensor              24 degrees C / 75 degrees F
Power
  VDD 2.5V                            2508 mV
  Bias 3.3V                           3351 mV
  VDD 3.3V                             3292 mV
  VCC 1.8V                             1832 mV
  VDD 1.2V                             1205 mV
  VCC 1V                               999 mV
  VCC CPU 1.8V                         1793 mV
  VDD 2.5V                             2508 mV
  VCC Aux 5V                           5056 mV
  VDD DDR 1.5V                         1509 mV
  VTT SA CPU 0.8V                      803 mV
  VTT CPU 1.05V                        1048 mV
  VCC Core CPU                         960 mV
  VCC PCH 1.5V                         1519 mV
  VDD 1.05V                            1058 mV
  VCC 2.5V                             2528 mV
  FORT VCCA 1V                         960 mV
  VDD .85V                             852 mV
  VTT DDRA .75V                        744 mV
  VTT DDRB .75V                        744 mV
  12V                                  12259 mV    3649 mA    45173 mW

CB 1 status:
  State                               Online Standby
CB 1 Intake Temp Sensor               20 degrees C / 68 degrees F
CB 1 Exhaust Temp Sensor              23 degrees C / 73 degrees F
Power
  VDD 2.5V                            2508 mV
  Bias 3.3V                           3312 mV
  VDD 3.3V                             3273 mV
  VCC 1.8V                             1822 mV

```


VDD 1.2V	1195 mV		
VCC 1V	989 mV		
VCC CPU 1.8V	1783 mV		
VDD 2.5V	2489 mV		
VCC Aux 5V	5086 mV		
VDD DDR 1.5V	1499 mV		
VTT SA CPU 0.8V	803 mV		
VTT CPU 1.05V	1048 mV		
VCC Core CPU	1029 mV		
VCC PCH 1.5V	1519 mV		
VDD 1.05V	1048 mV		
VCC 2.5V	2528 mV		
FORT VCCA 1V	960 mV		
VDD .85V	862 mV		
VTT DDRA .75V	744 mV		
VTT DDRB .75V	744 mV		
12V	12285 mV	3952 mA	48447 mW

show chassis environment cb (QFabric System)

user@switch> show chassis environment cb interconnect-device IC-123 0

CB 0 status:

State	Online Master
Left Intake Temperature	33 degrees C / 91 degrees F
Right Intake Temperature	33 degrees C / 91 degrees F
Left Exhaust Temperature	36 degrees C / 96 degrees F
Right Exhaust Temperature	35 degrees C / 95 degrees F
Power	OK
VDD 3V3	3294 mV
VDD 2V5	2436 mV
VDD 1V8	1746 mV
VDD 1V5	1460 mV
VDD 1V25	1210 mV
VDD 1V2	1164 mV
CPU CORE 1V2	1120 mV
VDD 1V0	968 mV
VDD 5V0	5088 mV
CPU MP BIAS 4V3	4050 mV
BIAS 3V3	3180 mV
VTT 0V9	866 mV

show chassis environment cb 0 (PTX10003-80C and PTX10003-160C Routers)

```
user@router> show chassis environment cb 0
```

```

CB 0 status:
  State                               Online Master
  Power 1
    PCIe SW0 power, CORE 0.9 V      894 mV
  Power 2
    CLOCK MAIN, 2.5 V                2495 mV
    PCIe SW0 power, SERDES 0.9 V     897 mV
  Power 3
    PCIe SW1 power, CORE 0.9 V      897 mV
  Power 4
    PCIe SW1 power, SERDES 0.9 V     894 mV
  Power 5
    Ethernet SW0 power, 1.0 V        998 mV
  Power 6
    Ethernet SW1 power, CORE 1.0 V   998 mV
  Power 7
    12 V                             12034 mV
  Bus Revision                        0
  FPGA Revision                      0

```

Starting in Junos OS Evolved Release 19.1R1, on PTX10003-80C and PTX10003-160C devices, the **show chassis environment cb** command does not show the Bus and FPGA revision information. If you want to view the FPGA revision or version information for the CB, use the **show system firmware** command.

show chassis ethernet-switch interconnect-device cb

Syntax

```
show chassis ethernet-switch interconnect-device name cb  
<detail>  
<port number>  
<slotnumber>
```

Release Information

Command introduced in Junos OS Release 12.2 for the QFX Series.

Description

(QFX3000-G QFabric systems only) Display Ethernet switch information for the Control Board (CB) ports in an Interconnect device.

Options

none—Display Ethernet switch information about each connected port on each online CB in the Interconnect device.

detail—(Optional) Display detailed status information for all CBs or for the CB in the specified slot in the Interconnect device.

port *number*—(Optional) Display Ethernet switch information about a specific port on a CB in the Interconnect device.

slot *number*—(Optional) Display Ethernet switch information about a CB in a specific slot in the Interconnect device.

Required Privilege Level

view

RELATED DOCUMENTATION

[chassis | 461](#)

[show chassis environment cb | 656](#)

[show chassis ethernet-switch interconnect-device fpc | 715](#)

List of Sample Output

[show chassis ethernet-switch interconnect-device cb on page 697](#)

[show chassis ethernet-switch interconnect-device cb detail on page 699](#)

[show chassis ethernet-switch interconnect-device cb detail slot port on page 706](#)

Output Fields

Table 84 on page 693 lists the output fields for the **show chassis ethernet-switch interconnect-device cb** command. Output fields are listed in the approximate order in which they appear.

Table 84: show chassis ethernet-switch interconnect-device fpc Output Fields

Field Name	Field Description
Link is good on port n connected to device	Information about the link between each port on the FPC's Ethernet switch and one of the following devices: <ul style="list-style-type: none"> • FWD-SWITCH-0 • FWD-SWITCH-1 • CB0 • CB1
Speed is	Speed at which the Ethernet link is running: 10 Mb When the device is RE or Other RE on the TX Matrix router, the speed is 1000 Mb .
Duplex is	Duplex type of the Ethernet link: full or half .
Autonegotiate is Enabled (or Disabled)	By default, built-in Fast Ethernet ports on a PIC autonegotiate whether to operate at 10 Mbps or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the no-concatenate statement at the [edit chassis] hierarchy level, as described in the <i>Junos OS System Basics Configuration Guide</i>).
Flow Control TX is Enabled (or Disabled)	Flow control in the transmit direction is enabled (or disabled). Flow control regulates the flow of packets from the switch to the remote side of the connection.
Flow Control RX is Enabled (or Disabled)	Flow control in the receive direction is enabled (or disabled). Flow control regulates the flow of packets from the remote side of the connection to the switch.
TX Octets	Number of octets sent.
TX Packets 64 Octets	Number of transmitted packets of size 64 octets.
TX Packets 65-127 Octets	Number of transmitted frames of size 65 through 127 octets.
TX Packets 128-255 Octets	Number of transmitted frames of size 128 through 255 octets.
TX Packets 256-511 Octets	Number of transmitted frames of size 256 through 511 octets.

Table 84: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
TX Packets 512-1023 Octets	Number of transmitted frames of size 512 through 1023 octets.
TX Packets 1024-1518 Octets	Number of transmitted frames of size 1024 through 1518 octets.
TX Packets 1519-2047 Octets	Number of transmitted frames of size 1519 through 2047 octets.
TX Packets 2048-4095 Octets	Number of transmitted frames of size 2048 through 4095 octets.
TX Packets 4096-9216 Octets	Number of transmitted frames of size 4096 through 9216 octets.
TX Packets 9217-16383 Octets	Number of transmitted frames of size 9217 through 16383 octets.
TX Multicast packets	Number of multicast packets sent.
TX Broadcast packets	Number of broadcast packets sent.
TX Single Collision frames	Number of packets sent after one collision.
TX Mult. Collision frames	Number of packets sent after multiple collisions.
TX Late Collision Frames	Number of packets aborted during sending because of collisions after 64 bytes.
TX Excessive collisions	Number of packets not sent because of too many collisions.
TX Collision frames	Number of collision packets sent.
TX PAUSEMAC Ctrl Frames	Number of Media Access Control (MAC) frames containing PAUSE commands sent.
TX MAC ctrl frames	Number of MAC control packets sent.

Table 84: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
TX Frame deferred Xmns	Number of frames deferred in x milliseconds.
TX Oversize Packets	Number of oversized packets sent.
TX Jabbers	Total number of frames sent that exceed the maximum byte count and contain CRC errors .
TX FCS Error Counter	Number of packets discarded because of frame check sequence errors.
TX Fragment Counter	Number of fragmented packets sent.
TX Byte Counter	Number of bytes sent.
RX Octets	Number of octets received.
RX Packets 64 Octets	Number of received packets of size 64 octets.
RX Packets 65-127 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 128-255 Octets	Number of received packets of size 128 through 255 octets.
RX Packets 256-511 Octets	Number of received packets of size 256 through 511 octets.
RX Packets 512-1023 Octets	Number of received packets of size 512 through 1023 octets.
RX Packets 1024-1518 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 1519-2047 Octets	Number of received packets of size 1519 through 2047 octets.
RX Packets 2048-4095 Octets	Number of received packets of size 2048 through 4095 octets.
RX Packets 4096-9216 Octets	Number of received packets of size 4096 through 9216 octets.

Table 84: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
RX Multicast Packets	Number of multicast packets received.
RX Broadcast Packets	Number of broadcast packets received.
RX FCS Errors	Number of packets discarded because of frame check sequence errors.
RX Align Errors	Number of incomplete octets received.
RX Fragments	Number of fragmented packets received.
RX Symbol errors	Number of symbols received that the router did not correctly decode.
RX Unsupported opcodes	Number of packets received with unsupported op codes.
RX Out of Range Length	Number of packets received with an out of range length.
RX False Carrier Errors	Number of packets received with false carrier errors.
RX Undersize Packets	Number of undersized packets received.
RX Oversize Packets	Number of oversized packets received.
RX Jabbers	Total number of frames received that exceed the maximum byte count and contain CRC errors .
RX 1519-1522 Good Vlan frms	
RX MTU Exceed Counter	Number of packets received that exceed the MTU.
RX Control Frame Counter	Number of control frames received.
RX Pause Frame Counter	Number of pause frames received.

Table 84: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
RX Byte Counter	Number of bytes received.

Sample Output

show chassis ethernet-switch interconnect-device cb

user@switch> **show chassis ethernet-switch interconnect-device IC-WS001 cb**

```

Displaying summary for switch 0
Link is down on XE port 1 connected to device: FPC7
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 2 connected to device: FPC6
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 3 connected to device: FPC5
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 5 connected to device: FPC4
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 7 connected to device: FPC3
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 9 connected to device: FPC2
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 10 connected to device: FPC1
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  TX Octets          326358
  RX Octets          237947

```


Link is good on XE port 11 connected to device: FPC0

Speed is 10000Mb

Duplex is full

Autonegotiate is Enabled

TX Octets 548249

RX Octets 386013

Link is down on XE port 20 connected to device: SFP3

Flow Control TX is Disabled

Flow Control RX is Disabled

Link is down on XE port 21 connected to device: SFP2

Flow Control TX is Disabled

Flow Control RX is Disabled

Link is good on XE port 22 connected to device: SFP1

Speed is 1000Mb

Duplex is full

Autonegotiate is Enabled

TX Octets 1

RX Octets 11704758

Link is good on XE port 23 connected to device: SFP0

Speed is 1000Mb

Duplex is full

Autonegotiate is Enabled

TX Octets 1500022

RX Octets 11629453

Link is good on XE port 24 connected to device: VCCPD

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

TX Octets 23332467

RX Octets 1500023

Link is good on GE port 25 connected to device: SFI

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

TX Octets 643918

RX Octets 894548

show chassis ethernet-switch interconnect-device cb detail

```
user@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 cb detail
```

```
Port statistics for CB switch

Link is down on XE port 1 connected to device: FPC7

Link is down on XE port 2 connected to device: FPC6

Link is down on XE port 3 connected to device: FPC5

Link is down on XE port 5 connected to device: FPC4

Link is down on XE port 7 connected to device: FPC3

Link is down on XE port 9 connected to device: FPC2
Statistics for port 10 connected to device FPC1:
  TX Packets 64 Octets      0
  TX Packets 65-127 Octets  1
  TX Packets 128-255 Octets 319293
  TX Packets 256-511 Octets 5043
  TX Packets 512-1023 Octets 2072
  TX Packets 1024-1518 Octets 6
  TX Packets 1519-2047 Octets 0
  TX Packets 2048-4095 Octets 0
  TX Packets 4096-9216 Octets 0
  TX Packets 9217-16383 Octets 0
  TX Octets      326415
  TX Multicast Packets 0
  TX Broadcast Packets 1
  TX PAUSEMAC Ctrl Frames 0
  TX Oversize Packets 0
  TX FCS Error Counter 0
  TX Fragment Counter 0
  TX Byte Counter 71659246
  TX Packet OK Counter 326415
  TX Pause Packet Counter 0
  TX Unicast Counter 326414
  RX Packets 64 Octets 0
  RX Packets 65-127 Octets 0
  RX Packets 128-255 Octets 235428
  RX Packets 256-511 Octets 2134
  RX Packets 512-1023 Octets 420
  RX Packets 1024-1518 Octets 6
```



```

RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Packets 9217-16383 Octets  0
RX Octets                237988
RX Multicast Packets      0
RX Broadcast Packets      0
RX FCS Errors             0
RX Fragments              0
RX MAC Control Packets    0
RX Out of Range Length    0
RX Undersize Packets      0
RX Oversize Packets       0
RX Jabbers                 0
RX Control Frame Counter  0
RX Pause Frame Counter    0
RX Byte Counter           55821504
RX Unicast Frame Count    237988
RX Packet OK Count        237988

```

Statistics for port 11 connected to device FPC0:

```

TX Packets 64 Octets      0
TX Packets 65-127 Octets  1
TX Packets 128-255 Octets 535483
TX Packets 256-511 Octets 9289
TX Packets 512-1023 Octets 3564
TX Packets 1024-1518 Octets 5
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                548342
TX Multicast Packets      0
TX Broadcast Packets      1
TX PAUSEMAC Ctrl Frames   0
TX Oversize Packets       0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           120498414
TX Packet OK Counter      548342
TX Pause Packet Counter   0
TX Unicast Counter        548341
RX Packets 64 Octets      0
RX Packets 65-127 Octets  0
RX Packets 128-255 Octets 382931

```



```

RX Packets 256-511 Octets    2762
RX Packets 512-1023 Octets   386
RX Packets 1024-1518 Octets   3
RX Packets 1519-2047 Octets   0
RX Packets 2048-4095 Octets   0
RX Packets 4096-9216 Octets   0
RX Packets 9217-16383 Octets  0
RX Octets                    386082
RX Multicast Packets          0
RX Broadcast Packets          0
RX FCS Errors                 0
RX Fragments                  0
RX MAC Control Packets        0
RX Out of Range Length        0
RX Undersize Packets          0
RX Oversize Packets           0
RX Jabbers                    0
RX Control Frame Counter      0
RX Pause Frame Counter        0
RX Byte Counter               90717369
RX Unicast Frame Count        386082
RX Packet OK Count            386082

```

Link is down on XE port 20 connected to device: SFP3

Link is down on XE port 21 connected to device: SFP2

Statistics for port 22 connected to device SFP1:

```

TX Packets 64 Octets          0
TX Packets 65-127 Octets      0
TX Packets 128-255 Octets     0
TX Packets 256-511 Octets     0
TX Packets 512-1023 Octets    0
TX Packets 1024-1518 Octets    1
TX Packets 1519-2047 Octets    0
TX Packets 2048-4095 Octets    0
TX Packets 4096-9216 Octets    0
TX 1519-1522 Good Vlan Frms   0
TX Octets                     1
TX Multicast Packets          1
TX Broadcast Packets          0
TX Single Collision Frames     0
TX Mult. Collision Frames      0
TX Late Collisions             0
TX Excessive Collisions        0

```


TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xtns	0
TX Frame Excessive Deferral	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	1422
RX Packet OK Count	1
RX Packets 64 Octets	230013
RX Packets 65-127 Octets	174529
RX Packets 128-255 Octets	286735
RX Packets 256-511 Octets	343412
RX Packets 512-1023 Octets	172152
RX Packets 1024-1518 Octets	10500065
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	11706906
RX Multicast Packets	11672320
RX Broadcast Packets	34460
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	2379464164
RX Packet OK Count	11706906

Statistics for port 23 connected to device SFP0:

TX Packets 64 Octets	3
TX Packets 65-127 Octets	484733
TX Packets 128-255 Octets	219112
TX Packets 256-511 Octets	129014

TX Packets 512-1023 Octets	503
TX Packets 1024-1518 Octets	666958
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	0
TX Octets	1500323
TX Multicast Packets	794098
TX Broadcast Packets	1040
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xmns	0
TX Frame Excessive Deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	1065466891
RX Packet OK Count	1500323
RX Packets 64 Octets	341563
RX Packets 65-127 Octets	430810
RX Packets 128-255 Octets	318279
RX Packets 256-511 Octets	347147
RX Packets 512-1023 Octets	184798
RX Packets 1024-1518 Octets	10008993
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	11631590
RX Multicast Packets	10878484
RX Broadcast Packets	33420
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0

RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	1720484325
RX Packet OK Count	11631591

Statistics for port 24 connected to device VCCPD:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1176546
TX Packets 128-255 Octets	604988
TX Packets 256-511 Octets	690561
TX Packets 512-1023 Octets	356942
TX Packets 1024-1518 Octets	20507438
TX Packets 1519-2047 Octets	278
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	278
TX Octets	23336753
TX Multicast Packets	22549383
TX Broadcast Packets	67862
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xtns	0
TX Frame Excessive Deferral	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	4191296788
RX Packet OK Count	23336753
RX Packets 64 Octets	3
RX Packets 65-127 Octets	484673
RX Packets 128-255 Octets	219074
RX Packets 256-511 Octets	129100
RX Packets 512-1023 Octets	516
RX Packets 1024-1518 Octets	666959
RX Packets 1519-2047 Octets	0


```

RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Octets                    1500325
RX Multicast Packets         794099
RX Broadcast Packets         1040
RX FCS Errors                0
RX Align Errors              0
RX Fragments                 0
RX Symbol Errors             0
RX Unsupported Opcodes       0
RX Out of Range Length       0
RX False Carrier Errors      0
RX Undersize Packets         0
RX Oversize Packets          0
RX Jabbers                   0
RX 1519-1522 Good Vlan Frms 0
RX MTU Exceed Counter        0
RX Control Frame Counter     0
RX Pause Frame Counter       0
RX Byte Counter              1071469739
RX Packet OK Count           1500325

```

Statistics for port 25 connected to device SFI:

```

TX Packets 64 Octets         12
TX Packets 65-127 Octets     1
TX Packets 128-255 Octets    618363
TX Packets 256-511 Octets    4896
TX Packets 512-1023 Octets   806
TX Packets 1024-1518 Octets  19950
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX 1519-1522 Good Vlan Frms 0
TX Octets                    644028
TX Multicast Packets         4
TX Broadcast Packets         19954
TX Single Collision Frames   0
TX Mult. Collision Frames    0
TX Late Collisions           0
TX Excessive Collisions      0
TX Collision Frames          0
TX PAUSEMAC Ctrl Frames      0
TX MAC Ctrl Frames           0
TX Frame Deferred Xtns       0
TX Frame Excessive Deferl    0

```


TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	167039705
RX Packet OK Count	644028
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	854776
RX Packets 256-511 Octets	14332
RX Packets 512-1023 Octets	5636
RX Packets 1024-1518 Octets	19954
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	894698
RX Multicast Packets	0
RX Broadcast Packets	19943
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	212658920
RX Packet OK Count	894698

show chassis ethernet-switch interconnect-device cb detail slot port

user@qfabric> **show chassis ethernet-switch interconnect-device IC-WS001 cb slot 1 port 1**

re0:

Port statistics for CB switch

Link is down on XE port 1 connected to device: FPC7

Link is down on XE port 2 connected to device: FPC6

Link is down on XE port 3 connected to device: FPC5

Link is down on XE port 5 connected to device: FPC4

Link is down on XE port 7 connected to device: FPC3

Link is down on XE port 9 connected to device: FPC2

Statistics for port 10 connected to device FPC1:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	319366
TX Packets 256-511 Octets	5043
TX Packets 512-1023 Octets	2072
TX Packets 1024-1518 Octets	6
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX Packets 9217-16383 Octets	0
TX Octets	326488
TX Multicast Packets	0
TX Broadcast Packets	1
TX PAUSEMAC Ctrl Frames	0
TX Oversize Packets	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	71675330
TX Packet OK Counter	326488
TX Pause Packet Counter	0
TX Unicast Counter	326487
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	235481
RX Packets 256-511 Octets	2134
RX Packets 512-1023 Octets	420
RX Packets 1024-1518 Octets	6
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	238041
RX Multicast Packets	0

RX Broadcast Packets	0
RX FCS Errors	0
RX Fragments	0
RX MAC Control Packets	0
RX Out of Range Length	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	55834224
RX Unicast Frame Count	238041
RX Packet OK Count	238041

Statistics for port 11 connected to device FPC0:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	535606
TX Packets 256-511 Octets	9289
TX Packets 512-1023 Octets	3564
TX Packets 1024-1518 Octets	5
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX Packets 9217-16383 Octets	0
TX Octets	548465
TX Multicast Packets	0
TX Broadcast Packets	1
TX PAUSEMAC Ctrl Frames	0
TX Oversize Packets	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	120525524
TX Packet OK Counter	548465
TX Pause Packet Counter	0
TX Unicast Counter	548464
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	383018
RX Packets 256-511 Octets	2762
RX Packets 512-1023 Octets	386
RX Packets 1024-1518 Octets	3
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0


```

RX Packets 9217-16383 Octets  0
RX Octets                      386169
RX Multicast Packets          0
RX Broadcast Packets          0
RX FCS Errors                  0
RX Fragments                  0
RX MAC Control Packets        0
RX Out of Range Length        0
RX Undersize Packets          0
RX Oversize Packets           0
RX Jabbers                     0
RX Control Frame Counter      0
RX Pause Frame Counter        0
RX Byte Counter                90738249
RX Unicast Frame Count         386169
RX Packet OK Count             386169

```

Link is down on XE port 20 connected to device: SFP3

Link is down on XE port 21 connected to device: SFP2

Statistics for port 22 connected to device SFP1:

```

TX Packets 64 Octets          0
TX Packets 65-127 Octets      0
TX Packets 128-255 Octets     0
TX Packets 256-511 Octets     0
TX Packets 512-1023 Octets    0
TX Packets 1024-1518 Octets   1
TX Packets 1519-2047 Octets   0
TX Packets 2048-4095 Octets   0
TX Packets 4096-9216 Octets   0
TX 1519-1522 Good Vlan Frms  0
TX Octets                      1
TX Multicast Packets          1
TX Broadcast Packets          0
TX Single Collision Frames     0
TX Mult. Collision Frames     0
TX Late Collisions             0
TX Excessive Collisions        0
TX Collision Frames            0
TX PAUSEMAC Ctrl Frames       0
TX MAC Ctrl Frames            0
TX Frame Deferred Xtns        0
TX Frame Excessive Deferral   0
TX Oversize Packets           0

```


TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	1422
RX Packet OK Count	1
RX Packets 64 Octets	230071
RX Packets 65-127 Octets	174571
RX Packets 128-255 Octets	286812
RX Packets 256-511 Octets	343500
RX Packets 512-1023 Octets	172203
RX Packets 1024-1518 Octets	10502544
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	11709701
RX Multicast Packets	11675110
RX Broadcast Packets	34465
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	2383079858
RX Packet OK Count	11709701

Statistics for port 23 connected to device SFP0:

TX Packets 64 Octets	3
TX Packets 65-127 Octets	485048
TX Packets 128-255 Octets	219200
TX Packets 256-511 Octets	129053
TX Packets 512-1023 Octets	503
TX Packets 1024-1518 Octets	667127
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	0

TX Octets	1500934
TX Multicast Packets	794300
TX Broadcast Packets	1040
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xmns	0
TX Frame Excessive Deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	1065764997
RX Packet OK Count	1500934
RX Packets 64 Octets	341648
RX Packets 65-127 Octets	431183
RX Packets 128-255 Octets	318367
RX Packets 256-511 Octets	347225
RX Packets 512-1023 Octets	184849
RX Packets 1024-1518 Octets	10011311
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	11634583
RX Multicast Packets	10881071
RX Broadcast Packets	33425
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0


```

RX Byte Counter          1723893006
RX Packet OK Count       11634583
Statistics for port 24 connected to device VCCPD:
TX Packets 64 Octets      0
TX Packets 65-127 Octets  1177102
TX Packets 128-255 Octets 605153
TX Packets 256-511 Octets 690727
TX Packets 512-1023 Octets 357044
TX Packets 1024-1518 Octets 20512235
TX Packets 1519-2047 Octets 278
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan Frms 278
TX Octets                 23342539
TX Multicast Packets      22554760
TX Broadcast Packets      67872
TX Single Collision Frames 0
TX Mult. Collision Frames 0
TX Late Collisions        0
TX Excessive Collisions   0
TX Collision Frames       0
TX PAUSEMAC Ctrl Frames   0
TX MAC Ctrl Frames        0
TX Frame Deferred Xmns    0
TX Frame Excessive Deferl 0
TX Oversize Packets       0
TX Jabbers                0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           4198344167
RX Packet OK Count        23342539
RX Packets 64 Octets      3
RX Packets 65-127 Octets  484985
RX Packets 128-255 Octets 219164
RX Packets 256-511 Octets 129139
RX Packets 512-1023 Octets 516
RX Packets 1024-1518 Octets 667128
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                 1500935
RX Multicast Packets      794301
RX Broadcast Packets      1040
RX FCS Errors             0

```


RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	1071770147
RX Packet OK Count	1500935

Statistics for port 25 connected to device SFI:

TX Packets 64 Octets	12
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	618503
TX Packets 256-511 Octets	4896
TX Packets 512-1023 Octets	806
TX Packets 1024-1518 Octets	19950
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan Frms	0
TX Octets	644168
TX Multicast Packets	4
TX Broadcast Packets	19954
TX Single Collision Frames	0
TX Mult. Collision Frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision Frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC Ctrl Frames	0
TX Frame Deferred Xtns	0
TX Frame Excessive Deferral	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	167073305
RX Packet OK Count	644168

RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	854972
RX Packets 256-511 Octets	14332
RX Packets 512-1023 Octets	5636
RX Packets 1024-1518 Octets	19954
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	894894
RX Multicast Packets	0
RX Broadcast Packets	19943
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol Errors	0
RX Unsupported Opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan Frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	212702114
RX Packet OK Count	894894

show chassis ethernet-switch interconnect-device fpc

Syntax

```
show chassis ethernet-switch interconnect-device name fpc
<detail>
<port number>
<slot number>
```

Release Information

Command introduced in Junos OS Release 12.2 for the QFX Series.

Description

(QFX3000-G QFabric systems only) Display Ethernet switch information for the front card Flexible Port Concentrators (FPCs) in an Interconnect device.

Options

none—Display Ethernet switch information about each connected port on each online FPC in the Interconnect device.

detail—(Optional) Display detailed status information for all FPCs or for the FPC in the specified slot in the Interconnect device.

port *number*—(Optional) Display Ethernet switch information about a specific port on an FPC in the Interconnect device.

slot *number*—(Optional) Display Ethernet switch information about an FPC in a specific slot in the Interconnect device.

Required Privilege Level

view

RELATED DOCUMENTATION

[chassis | 461](#)

[show chassis environment fpc](#)

[show chassis ethernet-switch interconnect-device cb | 692](#)

List of Sample Output

[show chassis ethernet-switch interconnect-device fpc on page 720](#)

[show chassis ethernet-switch interconnect-device fpc detail on page 723](#)

[show chassis ethernet-switch fpc detail slot on page 732](#)

[show chassis ethernet-switch fpc interconnect-device port on page 742](#)

[show chassis ethernet-switch fpc interconnect-device detail port on page 744](#)

Output Fields

[Table 84 on page 693](#) lists the output fields for the **show chassis ethernet-switch interconnect-device fpc** command. Output fields are listed in the approximate order in which they appear.

Table 85: show chassis ethernet-switch interconnect-device fpc Output Fields

Field Name	Field Description
Link is good on port n connected to device	Information about the link between each port on the FPC's Ethernet switch and one of the following devices: <ul style="list-style-type: none"> • FWD-SWITCH-0 • FWD-SWITCH-1 • CB0 • CB1
Speed is	Speed at which the Ethernet link is running: 10 Mb When the device is RE or Other RE on the TX Matrix router, the speed is 1000 Mb .
Duplex is	Duplex type of the Ethernet link: full or half .
Autonegotiate is Enabled (or Disabled)	By default, built-in Fast Ethernet ports on a PIC autonegotiate whether to operate at 10 Mbps or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the no-concatenate statement at the [edit chassis] hierarchy level, as described in the <i>Junos OS System Basics Configuration Guide</i>).
TX Octets	Number of octets sent.
TX Packets 64 Octets	Number of transmitted frames of size 64 octets.
TX Packets 65-127 Octets	Number of transmitted frames of size 65 through 127 octets.
TX Packets 128-255 Octets	Number of transmitted frames of size 128 through 255 octets.
TX Packets 256-511 Octets	Number of transmitted frames of size 256 through 511 octets.
TX Packets 512-1023 Octets	Number of transmitted frames of size 512 through 1023 octets.

Table 85: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
TX Packets 1024-1518 Octets	Number of transmitted frames of size 1024 through 1518 octets.
TX Packets 1519-2047 Octets	Number of transmitted frames of size 1519 through 2047 octets.
TX Packets 2048-4095 Octets	Number of transmitted frames of size 2048 through 4095 octets.
TX Packets 4096-9216 Octets	Number of transmitted frames of size 4096 through 9216 octets.
TX Packets 9217-16383 Octets	Number of transmitted frames of size 9217 through 16383 octets.
TX Multicast packets	Number of multicast packets sent.
TX Broadcast packets	Number of broadcast packets sent.
TX Single Collision frames	Number of packets sent after one collision.
TX Mult. Collision frames	Number of packets sent after multiple collisions.
TX Late Collision Frames	Number of packets aborted during sending because of collisions after 64 bytes.
TX Excessive collisions	Number of packets not sent because of too many collisions.
TX Collision frames	Number of collision packets sent.
TX PAUSEMAC Ctrl Frames	Number of Media Access Control (MAC) frames containing PAUSE commands sent.
TX MAC ctrl frames	Number of MAC control packets sent.
TX Frame deferred Xmns	Number of frames deferred in x milliseconds.

Table 85: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
TX Oversize Packets	Number of oversized packets sent.
TX Jabbers	Total number of frames sent that exceed the maximum byte count and contain CRC errors .
TX FCS Error Counter	Number of packets discarded because of frame check sequence errors.
TX Fragment Counter	Number of fragmented packets sent.
TX Byte Counter	Number of bytes sent.
RX Octets	Number of octets received.
RX Packets 64 Octets	Number of received packets of size 64 octets.
RX Packets 65-127 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 128-255 Octets	Number of received packets of size 128 through 255 octets.
RX Packets 256-511 Octets	Number of received packets of size 256 through 511 octets.
RX Packets 512-1023 Octets	Number of received packets of size 512 through 1023 octets.
RX Packets 1024-1518 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 1519-2047 Octets	Number of received packets of size 1519 through 2047 octets.
RX Packets 2048-4095 Octets	Number of received packets of size 2048 through 4095 octets.
RX Packets 4096-9216 Octets	Number of received packets of size 4096 through 9216 octets.
RX Multicast Packets	Number of multicast packets received.

Table 85: show chassis ethernet-switch interconnect-device fpc Output Fields (*continued*)

Field Name	Field Description
RX Broadcast Packets	Number of broadcast packets received.
RX FCS Errors	Number of packets discarded because of frame check sequence errors.
RX Align Errors	Number of incomplete octets received.
RX Fragments	Number of fragmented packets received.
RX Symbol errors	Number of symbols received that the router did not correctly decode.
RX Unsupported opcodes	Number of packets received with unsupported op codes.
RX Out of Range Length	Number of packets received with an out of range length.
RX False Carrier Errors	Number of packets received with false carrier errors.
RX Undersize Packets	Number of undersized packets received.
RX Oversize Packets	Number of oversized packets received.
RX Jabbers	Total number of frames received that exceed the maximum byte count and contain CRC errors .
RX 1519-1522 Good Vlan frms	Number of transmitted frames of size 1519 through 1522 octets that are good VLAN frames.
RX MTU Exceed Counter	Number of packets received that exceed the MTU.
RX Control Frame Counter	Number of control frames received.
RX Pause Frame Counter	Number of pause frames received.
RX Byte Counter	Number of bytes received.

Sample Output

show chassis ethernet-switch interconnect-device fpc

user@qfabric> **show chassis ethernet-switch interconnect-device IC-WS001 fpc**

```
Summary for switch on FC0
Link is good on GE port 2 connected to device: FWD-SWITCH-0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                124638
  RX Octets                86496

Link is good on GE port 4 connected to device: FWD-SWITCH-1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                82191
  RX Octets                58979

Link is good on XE port 28 connected to device: CB0
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                145475
  RX Octets                206828

Link is good on XE port 29 connected to device: CB1
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                1
  RX Octets                0

Summary for switch on FC1
Link is good on GE port 2 connected to device: FWD-SWITCH-0
```



```

Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets                82290
RX Octets                 59443

```

Link is good on GE port 4 connected to device: FWD-SWITCH-1

```

Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets                40900
RX Octets                 30013

```

Link is good on XE port 28 connected to device: CB0

```

Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets                89456
RX Octets                123189

```

Link is good on XE port 29 connected to device: CB1

```

Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets                1
RX Octets                 0

```

root@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 fpc

Summary for switch on FC0

Link is good on GE port 2 connected to device: FWD-SWITCH-0

```

Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets                124697

```


RX Octets 86535

Link is good on GE port 4 connected to device: FWD-SWITCH-1

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 82229

RX Octets 59009

Link is good on XE port 28 connected to device: CB0

Speed is 10000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 145544

RX Octets 206925

Link is good on XE port 29 connected to device: CB1

Speed is 10000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 1

RX Octets 0

Summary for switch on FC1

Link is good on GE port 2 connected to device: FWD-SWITCH-0

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

Flow Control TX is Disabled

Flow Control RX is Disabled

TX Octets 82327

RX Octets 59472

Link is good on GE port 4 connected to device: FWD-SWITCH-1

Speed is 1000Mb

Duplex is full

Autonegotiate is Disabled

Flow Control TX is Disabled


```

Flow Control RX is Disabled
TX Octets          40918
RX Octets          30028

```

Link is good on XE port 28 connected to device: CB0

```

Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          89500
RX Octets          123244

```

Link is good on XE port 29 connected to device: CB1

```

Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          1
RX Octets          0

```

show chassis ethernet-switch interconnect-device fpc detail

user@host> **show chassis ethernet-switch interconnect-device IC-WS001 fpc detail**

```

Port statistics for FC0 switch
Statistics for port 2 connected to device FWD-SWITCH-0:
TX Packets 64 Octets          0
TX Packets 65-127 Octets      1
TX Packets 128-255 Octets     121716
TX Packets 256-511 Octets     2200
TX Packets 512-1023 Octets    823
TX Packets 1024-1518 Octets   2
TX Packets 1519-2047 Octets   0
TX Packets 2048-4095 Octets   0
TX Packets 4096-9216 Octets   0
TX 1519-1522 Good Vlan frms  0
TX Octets                    124742
TX Multicast Packets          0
TX Broadcast Packets          1
TX Single Collision frames    0
TX Mult. Collision frames     0
TX Late Collisions            0

```


TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xmns	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	27391588
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	85924
RX Packets 256-511 Octets	555
RX Packets 512-1023 Octets	86
RX Packets 1024-1518 Octets	1
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	86566
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	20380581

Statistics for port 4 connected to device FWD-SWITCH-1:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	80374
TX Packets 256-511 Octets	1347
TX Packets 512-1023 Octets	532

TX Packets 1024-1518 Octets	3
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	82257
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xmns	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	18146746
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	58410
RX Packets 256-511 Octets	522
RX Packets 512-1023 Octets	96
RX Packets 1024-1518 Octets	2
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	59030
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0


```

RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter      0
RX Control Frame Counter   0
RX Pause Frame Counter     0
RX Byte Counter            13882179

```

Statistics for port 28 connected to device CB0:

```

TX Packets 64 Octets      0
TX Packets 65-127 Octets  0
TX Packets 128-255 Octets 144334
TX Packets 256-511 Octets 1077
TX Packets 512-1023 Octets 182
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                  145596
TX Multicast Packets      0
TX Broadcast Packets      0
TX PAUSEMAC Ctrl Frames   0
TX Oversize Packets       0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           34262760
RX Packets 64 Octets      0
RX Packets 65-127 Octets  1
RX Packets 128-255 Octets 202090
RX Packets 256-511 Octets 3547
RX Packets 512-1023 Octets 1355
RX Packets 1024-1518 Octets 5
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                  206998
RX Multicast Packets      0
RX Broadcast Packets      1
RX FCS Errors             0
RX Fragments              0
RX MAC Control Packets    0
RX Out of Range Length    0
RX Undersize Packets      0
RX Oversize Packets       0
RX Jabbers                0

```



```

RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             45538262

```

Statistics for port 29 connected to device CB1:

```

TX Packets 64 Octets        0
TX Packets 65-127 Octets    1
TX Packets 128-255 Octets   0
TX Packets 256-511 Octets   0
TX Packets 512-1023 Octets  0
TX Packets 1024-1518 Octets 0
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                   1
TX Multicast Packets        0
TX Broadcast Packets        1
TX PAUSEMAC Ctrl Frames     0
TX Oversize Packets         0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             72
RX Packets 64 Octets        0
RX Packets 65-127 Octets    0
RX Packets 128-255 Octets   0
RX Packets 256-511 Octets   0
RX Packets 512-1023 Octets  0
RX Packets 1024-1518 Octets 0
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                   0
RX Multicast Packets        0
RX Broadcast Packets        0
RX FCS Errors               0
RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    0
RX Pause Frame Counter      0

```


RX Byte Counter	0
-----------------	---

Port statistics for FC1 switch

Statistics for port 2 connected to device FWD-SWITCH-0:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	80560
TX Packets 256-511 Octets	1279
TX Packets 512-1023 Octets	514
TX Packets 1024-1518 Octets	3
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	82357
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xmns	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	18059906
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	58733
RX Packets 256-511 Octets	639
RX Packets 512-1023 Octets	119
RX Packets 1024-1518 Octets	3
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	59494
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0

RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	13994432

Statistics for port 4 connected to device FWD-SWITCH-1:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	39971
TX Packets 256-511 Octets	668
TX Packets 512-1023 Octets	290
TX Packets 1024-1518 Octets	3
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	40933
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xmns	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	9050841
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0


```

RX Packets 128-255 Octets    29767
RX Packets 256-511 Octets    225
RX Packets 512-1023 Octets   44
RX Packets 1024-1518 Octets   3
RX Packets 1519-2047 Octets   0
RX Packets 2048-4095 Octets   0
RX Packets 4096-9216 Octets   0
RX Octets                     30039
RX Multicast Packets          0
RX Broadcast Packets          0
RX FCS Errors                 0
RX Align Errors               0
RX Fragments                  0
RX Symbol errors              0
RX Unsupported opcodes        0
RX Out of Range Length        0
RX False Carrier Errors       0
RX Undersize Packets          0
RX Oversize Packets           0
RX Jabbers                    0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter         0
RX Control Frame Counter      0
RX Pause Frame Counter        0
RX Byte Counter               7043738

```

Statistics for port 28 connected to device CB0:

```

TX Packets 64 Octets         0
TX Packets 65-127 Octets     0
TX Packets 128-255 Octets    88500
TX Packets 256-511 Octets    864
TX Packets 512-1023 Octets   163
TX Packets 1024-1518 Octets   6
TX Packets 1519-2047 Octets   0
TX Packets 2048-4095 Octets   0
TX Packets 4096-9216 Octets   0
TX Packets 9217-16383 Octets  0
TX Octets                     89533
TX Multicast Packets          0
TX Broadcast Packets          0
TX PAUSEMAC Ctrl Frames      0
TX Oversize Packets           0
TX FCS Error Counter          0
TX Fragment Counter           0
TX Byte Counter               21038170

```



```

RX Packets 64 Octets      0
RX Packets 65-127 Octets  1
RX Packets 128-255 Octets 120531
RX Packets 256-511 Octets 1947
RX Packets 512-1023 Octets 804
RX Packets 1024-1518 Octets 6
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                  123289
RX Multicast Packets      0
RX Broadcast Packets      1
RX FCS Errors             0
RX Fragments              0
RX MAC Control Packets    0
RX Out of Range Length    0
RX Undersize Packets      0
RX Oversize Packets       0
RX Jabbers                0
RX Control Frame Counter  0
RX Pause Frame Counter    0
RX Byte Counter           27110675

```

Statistics for port 29 connected to device CB1:

```

TX Packets 64 Octets      0
TX Packets 65-127 Octets  1
TX Packets 128-255 Octets  0
TX Packets 256-511 Octets  0
TX Packets 512-1023 Octets  0
TX Packets 1024-1518 Octets 0
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                  1
TX Multicast Packets      0
TX Broadcast Packets      1
TX PAUSEMAC Ctrl Frames   0
TX Oversize Packets       0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           72
RX Packets 64 Octets      0
RX Packets 65-127 Octets  0

```



```

RX Packets 128-255 Octets    0
RX Packets 256-511 Octets   0
RX Packets 512-1023 Octets  0
RX Packets 1024-1518 Octets 0
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                    0
RX Multicast Packets         0
RX Broadcast Packets         0
RX FCS Errors                0
RX Fragments                 0
RX MAC Control Packets       0
RX Out of Range Length       0
RX Undersize Packets         0
RX Oversize Packets          0
RX Jabbers                   0
RX Control Frame Counter     0
RX Pause Frame Counter       0
RX Byte Counter              0

```

```

Port statistics for FC2 switch
Empty fpc slot number 2

```

```

Port statistics for FC3 switch
Empty fpc slot number 3

```

```

Port statistics for FC4 switch
Empty fpc slot number 4

```

```

Port statistics for FC5 switch
Empty fpc slot number 5

```

```

Port statistics for FC6 switch
Empty fpc slot number 6

```

```

Port statistics for FC7 switch
Empty fpc slot number 7

```

show chassis ethernet-switch fpc detail slot

user@qfabric> **show chassis ethernet-switch fpc detail 0**


```
re0:
```

```
-----
Port statistics for FC0 switch
```

```
Statistics for port 2 connected to device FWD-SWITCH-0:
```

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	121823
TX Packets 256-511 Octets	2200
TX Packets 512-1023 Octets	823
TX Packets 1024-1518 Octets	2
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	124849
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xmns	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	27414524
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	85998
RX Packets 256-511 Octets	557
RX Packets 512-1023 Octets	86
RX Packets 1024-1518 Octets	1
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	86642
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0

RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	20398564

Statistics for port 4 connected to device FWD-SWITCH-1:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	80443
TX Packets 256-511 Octets	1347
TX Packets 512-1023 Octets	532
TX Packets 1024-1518 Octets	3
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	82326
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xmns	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	18161734
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0


```

RX Packets 128-255 Octets    58460
RX Packets 256-511 Octets    523
RX Packets 512-1023 Octets   96
RX Packets 1024-1518 Octets   2
RX Packets 1519-2047 Octets   0
RX Packets 2048-4095 Octets   0
RX Packets 4096-9216 Octets   0
RX Octets                    59081
RX Multicast Packets         0
RX Broadcast Packets         0
RX FCS Errors                0
RX Align Errors              0
RX Fragments                 0
RX Symbol errors             0
RX Unsupported opcodes       0
RX Out of Range Length       0
RX False Carrier Errors      0
RX Undersize Packets         0
RX Oversize Packets          0
RX Jabbers                   0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter        0
RX Control Frame Counter     0
RX Pause Frame Counter       0
RX Byte Counter              13894171
Statistics for port 28 connected to device CB0:
TX Packets 64 Octets         0
TX Packets 65-127 Octets     0
TX Packets 128-255 Octets    144458
TX Packets 256-511 Octets    1080
TX Packets 512-1023 Octets   182
TX Packets 1024-1518 Octets   3
TX Packets 1519-2047 Octets   0
TX Packets 2048-4095 Octets   0
TX Packets 4096-9216 Octets   0
TX Packets 9217-16383 Octets  0
TX Octets                    145723
TX Multicast Packets         0
TX Broadcast Packets         0
TX PAUSEMAC Ctrl Frames     0
TX Oversize Packets          0
TX FCS Error Counter         0
TX Fragment Counter          0
TX Byte Counter              34292735

```



```

RX Packets 64 Octets      0
RX Packets 65-127 Octets  1
RX Packets 128-255 Octets 202266
RX Packets 256-511 Octets 3547
RX Packets 512-1023 Octets 1355
RX Packets 1024-1518 Octets 5
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                  207174
RX Multicast Packets      0
RX Broadcast Packets      1
RX FCS Errors             0
RX Fragments              0
RX MAC Control Packets    0
RX Out of Range Length    0
RX Undersize Packets      0
RX Oversize Packets       0
RX Jabbers                0
RX Control Frame Counter  0
RX Pause Frame Counter    0
RX Byte Counter           45576186

```

Statistics for port 29 connected to device CB1:

```

TX Packets 64 Octets      0
TX Packets 65-127 Octets  1
TX Packets 128-255 Octets  0
TX Packets 256-511 Octets  0
TX Packets 512-1023 Octets  0
TX Packets 1024-1518 Octets 0
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                  1
TX Multicast Packets      0
TX Broadcast Packets      1
TX PAUSEMAC Ctrl Frames   0
TX Oversize Packets       0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           72
RX Packets 64 Octets      0
RX Packets 65-127 Octets  0

```


RX Packets 128-255 Octets	0
RX Packets 256-511 Octets	0
RX Packets 512-1023 Octets	0
RX Packets 1024-1518 Octets	0
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	0
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Fragments	0
RX MAC Control Packets	0
RX Out of Range Length	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	0

Port statistics for Fc1 switch

Statistics for port 2 connected to device FWD-SWITCH-0:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	80629
TX Packets 256-511 Octets	1279
TX Packets 512-1023 Octets	514
TX Packets 1024-1518 Octets	3
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	82426
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0

TX Frame deferred Xtns	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	18074790
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	58785
RX Packets 256-511 Octets	640
RX Packets 512-1023 Octets	119
RX Packets 1024-1518 Octets	3
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	59547
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	14006842

Statistics for port 4 connected to device FWD-SWITCH-1:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	40004
TX Packets 256-511 Octets	668
TX Packets 512-1023 Octets	290
TX Packets 1024-1518 Octets	3
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0

TX 1519-1522 Good Vlan frms	0
TX Octets	40966
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xmns	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	9058102
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	29794
RX Packets 256-511 Octets	225
RX Packets 512-1023 Octets	44
RX Packets 1024-1518 Octets	3
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	30066
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0


```

RX Byte Counter          7050000
Statistics for port 28 connected to device CB0:
TX Packets 64 Octets      0
TX Packets 65-127 Octets  0
TX Packets 128-255 Octets 88579
TX Packets 256-511 Octets 865
TX Packets 512-1023 Octets 163
TX Packets 1024-1518 Octets 6
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                  89613
TX Multicast Packets       0
TX Broadcast Packets       0
TX PAUSEMAC Ctrl Frames    0
TX Oversize Packets        0
TX FCS Error Counter       0
TX Fragment Counter        0
TX Byte Counter            21056842
RX Packets 64 Octets      0
RX Packets 65-127 Octets  1
RX Packets 128-255 Octets 120633
RX Packets 256-511 Octets 1947
RX Packets 512-1023 Octets 804
RX Packets 1024-1518 Octets 6
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                  123391
RX Multicast Packets       0
RX Broadcast Packets       1
RX FCS Errors              0
RX Fragments               0
RX MAC Control Packets     0
RX Out of Range Length     0
RX Undersize Packets       0
RX Oversize Packets        0
RX Jabbers                 0
RX Control Frame Counter   0
RX Pause Frame Counter     0
RX Byte Counter            27132820
Statistics for port 29 connected to device CB1:

```


TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	0
TX Packets 256-511 Octets	0
TX Packets 512-1023 Octets	0
TX Packets 1024-1518 Octets	0
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX Packets 9217-16383 Octets	0
TX Octets	1
TX Multicast Packets	0
TX Broadcast Packets	1
TX PAUSEMAC Ctrl Frames	0
TX Oversize Packets	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	72
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	0
RX Packets 256-511 Octets	0
RX Packets 512-1023 Octets	0
RX Packets 1024-1518 Octets	0
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	0
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Fragments	0
RX MAC Control Packets	0
RX Out of Range Length	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	0

Port statistics for FC2 switch
Empty fpc slot number 2


```
Port statistics for FC3 switch
Empty fpc slot number 3
```

```
Port statistics for FC4 switch
Empty fpc slot number 4
```

```
Port statistics for FC5 switch
Empty fpc slot number 5
```

```
Port statistics for FC6 switch
Empty fpc slot number 6
```

```
Port statistics for FC7 switch
Empty fpc slot number 7
```

show chassis ethernet-switch fpc interconnect-device port

user@qfabric> **show chassis ethernet-switch fpc interconnect-device IC-WS001 port 2**

```
Summary for switch on FC0
Link is good on GE port 2 connected to device: FWD-SWITCH-0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                319466
  RX Octets                221869
```

```
Link is good on GE port 4 connected to device: FWD-SWITCH-1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                210295
  RX Octets                151164
```

```
Link is good on XE port 28 connected to device: CB0
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
```



```
Flow Control RX is Disabled
TX Octets                373033
RX Octets                 529760
```

Link is good on XE port 29 connected to device: CB1

```
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets                1
RX Octets                 0
```

Summary for switch on FC1

Link is good on GE port 2 connected to device: FWD-SWITCH-0

```
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets                210760
RX Octets                 152617
```

Link is good on GE port 4 connected to device: FWD-SWITCH-1

```
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets                104587
RX Octets                 77315
```

Link is good on XE port 28 connected to device: CB0

```
Speed is 10000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets                229932
RX Octets                 315346
```

Link is good on XE port 29 connected to device: CB1

```
Speed is 10000Mb
Duplex is full
```



```

Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled
TX Octets          1
RX Octets          0

```

show chassis ethernet-switch fpc interconnect-device detail port

user@qfabric> show chassis ethernet-switch fpc interconnect-device IC-WS001 detail port 2

```

Port statistics for FC0 switch
Statistics for port 2 connected to device FWD-SWITCH-0:
TX Packets 64 Octets          0
TX Packets 65-127 Octets     1
TX Packets 128-255 Octets    311974
TX Packets 256-511 Octets    5552
TX Packets 512-1023 Octets   2084
TX Packets 1024-1518 Octets  2
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX 1519-1522 Good Vlan frms  0
TX Octets                    319613
TX Multicast Packets         0
TX Broadcast Packets         1
TX Single Collision frames   0
TX Mult. Collision frames    0
TX Late Collisions           0
TX Excessive Collisions      0
TX Collision frames          0
TX PAUSEMAC Ctrl Frames      0
TX MAC ctrl frames           0
TX Frame deferred Xmns       0
TX Frame excessive deferl    0
TX Oversize Packets          0
TX Jabbers                   0
TX FCS Error Counter         0
TX Fragment Counter          0
TX Byte Counter              70091196
RX Packets 64 Octets          0
RX Packets 65-127 Octets     0
RX Packets 128-255 Octets    220284
RX Packets 256-511 Octets    1486
RX Packets 512-1023 Octets   198

```


RX Packets 1024-1518 Octets	1
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	221969
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	52192002

Statistics for port 4 connected to device FWD-SWITCH-1:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	205595
TX Packets 256-511 Octets	3426
TX Packets 512-1023 Octets	1366
TX Packets 1024-1518 Octets	3
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	210391
TX Multicast Packets	0
TX Broadcast Packets	1
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xmns	0


```

TX Frame excessive deferl 0
TX Oversize Packets 0
TX Jabbers 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 46380018
RX Packets 64 Octets 0
RX Packets 65-127 Octets 0
RX Packets 128-255 Octets 149866
RX Packets 256-511 Octets 1194
RX Packets 512-1023 Octets 173
RX Packets 1024-1518 Octets 2
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 151235
RX Multicast Packets 0
RX Broadcast Packets 0
RX FCS Errors 0
RX Align Errors 0
RX Fragments 0
RX Symbol errors 0
RX Unsupported opcodes 0
RX Out of Range Length 0
RX False Carrier Errors 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 35496911

```

Statistics for port 28 connected to device CB0:

```

TX Packets 64 Octets 0
TX Packets 65-127 Octets 0
TX Packets 128-255 Octets 370150
TX Packets 256-511 Octets 2680
TX Packets 512-1023 Octets 371
TX Packets 1024-1518 Octets 3
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0

```


TX Octets	373204
TX Multicast Packets	0
TX Broadcast Packets	0
TX PAUSEMAC Ctrl Frames	0
TX Oversize Packets	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	87688913
RX Packets 64 Octets	0
RX Packets 65-127 Octets	1
RX Packets 128-255 Octets	517569
RX Packets 256-511 Octets	8978
RX Packets 512-1023 Octets	3450
RX Packets 1024-1518 Octets	5
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	530003
RX Multicast Packets	0
RX Broadcast Packets	1
RX FCS Errors	0
RX Fragments	0
RX MAC Control Packets	0
RX Out of Range Length	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	116471142

Statistics for port 29 connected to device CB1:

TX Packets 64 Octets	0
TX Packets 65-127 Octets	1
TX Packets 128-255 Octets	0
TX Packets 256-511 Octets	0
TX Packets 512-1023 Octets	0
TX Packets 1024-1518 Octets	0
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX Packets 9217-16383 Octets	0
TX Octets	1
TX Multicast Packets	0

TX Broadcast Packets	1
TX PAUSEMAC Ctrl Frames	0
TX Oversize Packets	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	72
RX Packets 64 Octets	0
RX Packets 65-127 Octets	0
RX Packets 128-255 Octets	0
RX Packets 256-511 Octets	0
RX Packets 512-1023 Octets	0
RX Packets 1024-1518 Octets	0
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	0
RX Multicast Packets	0
RX Broadcast Packets	0
RX FCS Errors	0
RX Fragments	0
RX MAC Control Packets	0
RX Out of Range Length	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	0

show chassis fabric connectivity

Syntax

```
show chassis fabric connectivity (device | slot)
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Display the status of the data plane connections in your QFabric system.

Options

- none**—Display the status of all data plane connections in your QFabric system.
- device**—Display the status of the data plane connections for a specific device.
- slot**—Display the status of the data plane connections for a specific Flexible PIC Concentrator (FPC) slot on a specific device.

Required Privilege Level

admin

RELATED DOCUMENTATION

- [request chassis fabric fpc | 603](#)
- [show chassis fabric device | 758](#)
- [Understanding Interconnect Devices | 29](#)

List of Sample Output

- [show chassis fabric connectivity on page 750](#)
- [show chassis fabric connectivity device on page 754](#)
- [show chassis fabric connectivity device device-name slot on page 756](#)

Output Fields

[Table 86 on page 749](#) lists the output fields for the **show chassis fabric connectivity** command. Output fields are listed in the approximate order in which they appear.

Table 86: show chassis fabric connectivity Output Fields

Field Name	Field Description
Device ID	Hardware serial identifier of the QFabric system component.

Table 86: show chassis fabric connectivity Output Fields (*continued*)

Field Name	Field Description
Type	Model number of the QFabric system component. Values include qfxc08-3008 (QFX3008-I Interconnect device), qfx3600-I (QFX3600-I Interconnect device), and supported Node device names such as qfx3500 (QFX3500 Node device) and qfx3600-16q (QFX3600 Node device).
Fabric: Incoming links	Displays inbound data plane (fte-) connections between Node devices and Interconnect devices, and their status (such as Ok).
Fabric: Outgoing links	Displays outbound data plane (fte-) connections between Node devices and Interconnect devices, and their status (such as Ok).

Sample Output

show chassis fabric connectivity

user@qfabric> **show chassis fabric connectivity**

```

Device ID: ED1487, Type: qfx3500
  Fabric: Incoming links:
    A0010:fte-0/0/1          -> ED1487:fte-0/1/2          Ok
    A0010:fte-3/0/1          -> ED1487:fte-0/1/3          Ok
  Fabric: Outgoing links:
    ED1487:fte-0/1/2          -> A0010:fte-0/0/1          Ok
    ED1487:fte-0/1/3          -> A0010:fte-3/0/1          Ok
Device ID: ED3683, Type: qfx3500
  Fabric: Incoming links:
    A0010:fte-0/0/2          -> ED3683:fte-0/1/2          Ok
    A0010:fte-3/0/2          -> ED3683:fte-0/1/3          Ok
  Fabric: Outgoing links:
    ED3683:fte-0/1/2          -> A0010:fte-0/0/2          Ok
    ED3683:fte-0/1/3          -> A0010:fte-3/0/2          Ok
Device ID: ED3705, Type: qfx3500
  Fabric: Incoming links:
    A0010:fte-0/0/0          -> ED3705:fte-0/1/2          Ok
    A0010:fte-3/0/0          -> ED3705:fte-0/1/3          Ok
  Fabric: Outgoing links:
    ED3705:fte-0/1/2          -> A0010:fte-0/0/0          Ok
    ED3705:fte-0/1/3          -> A0010:fte-3/0/0          Ok
Device ID: ED3707, Type: qfx3500
  Fabric: Incoming links:

```



```

A0010:fte-0/0/8          -> ED3707:fte-0/1/2          Ok
A0010:fte-3/0/8          -> ED3707:fte-0/1/3          Ok
Fabric: Outgoing links:
ED3707:fte-0/1/2          -> A0010:fte-0/0/8          Ok
ED3707:fte-0/1/3          -> A0010:fte-3/0/8          Ok
Device ID: ED3711, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/9          -> ED3711:fte-0/1/2          Ok
A0010:fte-3/0/9          -> ED3711:fte-0/1/3          Ok
Fabric: Outgoing links:
ED3711:fte-0/1/2          -> A0010:fte-0/0/9          Ok
ED3711:fte-0/1/3          -> A0010:fte-3/0/9          Ok
Device ID: ED3702, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/10         -> ED3702:fte-0/1/2          Ok
A0010:fte-3/0/10         -> ED3702:fte-0/1/3          Ok
Fabric: Outgoing links:
ED3702:fte-0/1/2         -> A0010:fte-0/0/10         Ok
ED3702:fte-0/1/3         -> A0010:fte-3/0/10         Ok
Device ID: BBAK8737, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/11         -> BBAK8737:fte-0/1/1          Ok
Fabric: Outgoing links:
BBAK8737:fte-0/1/1       -> A0010:fte-0/0/11         Ok
Device ID: BBAK8777, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/5          -> BBAK8777:fte-0/1/0          Ok
A0010:fte-0/0/6          -> BBAK8777:fte-0/1/1          Ok
Fabric: Outgoing links:
BBAK8777:fte-0/1/0       -> A0010:fte-0/0/5          Ok
BBAK8777:fte-0/1/1       -> A0010:fte-0/0/6          Ok
Device ID: BBAK8866, Type: qfx3500
Fabric: Incoming links:
A0010:fte-0/0/3          -> BBAK8866:fte-0/1/0          Ok
A0010:fte-0/0/4          -> BBAK8866:fte-0/1/1          Ok
Fabric: Outgoing links:
BBAK8866:fte-0/1/0       -> A0010:fte-0/0/3          Ok
BBAK8866:fte-0/1/1       -> A0010:fte-0/0/4          Ok
Device ID: BBAK8810, Type: qfx3500
Fabric: Incoming links:
A0010:fte-3/0/5          -> BBAK8810:fte-0/1/0          Ok
A0010:fte-3/0/6          -> BBAK8810:fte-0/1/1          Ok
Fabric: Outgoing links:
BBAK8810:fte-0/1/0       -> A0010:fte-3/0/5          Ok

```



```

    BBAK8810:fte-0/1/1          -> A0010:fte-3/0/6          Ok
Device ID: BBAK8854, Type: qfx3500
    Fabric: Incoming links:
    A0010:fte-3/0/3             -> BBAK8854:fte-0/1/0          Ok
    A0010:fte-3/0/4             -> BBAK8854:fte-0/1/1          Ok
    Fabric: Outgoing links:
    BBAK8854:fte-0/1/0          -> A0010:fte-3/0/3          Ok
    BBAK8854:fte-0/1/1          -> A0010:fte-3/0/4          Ok
Device ID: BBAK8885, Type: qfx3500
    Fabric: Incoming links:
    A0010:fte-0/0/14           -> BBAK8885:fte-0/1/0          Ok
    A0010:fte-0/0/15           -> BBAK8885:fte-0/1/1          Ok
    Fabric: Outgoing links:
    BBAK8885:fte-0/1/0          -> A0010:fte-0/0/14          Ok
    BBAK8885:fte-0/1/1          -> A0010:fte-0/0/15          Ok
Device ID: BBAK8864, Type: qfx3500
    Fabric: Incoming links:
    A0010:fte-3/0/7             -> BBAK8864:fte-0/1/0          Ok
    A0010:fte-3/0/11           -> BBAK8864:fte-0/1/1          Ok
    Fabric: Outgoing links:
    BBAK8864:fte-0/1/0          -> A0010:fte-3/0/7          Ok
    BBAK8864:fte-0/1/1          -> A0010:fte-3/0/11          Ok
Device ID: BBAK8759, Type: qfx3500
    Fabric: Incoming links:
    A0010:fte-3/0/12           -> BBAK8759:fte-0/1/0          Ok
    A0010:fte-3/0/13           -> BBAK8759:fte-0/1/1          Ok
    Fabric: Outgoing links:
    BBAK8759:fte-0/1/0          -> A0010:fte-3/0/12          Ok
    BBAK8759:fte-0/1/1          -> A0010:fte-3/0/13          Ok
Device ID: BBAK8704, Type: qfx3500
    Fabric: Incoming links:
    A0010:fte-0/0/12           -> BBAK8704:fte-0/1/0          Ok
    A0010:fte-0/0/13           -> BBAK8704:fte-0/1/1          Ok
    Fabric: Outgoing links:
    BBAK8704:fte-0/1/0          -> A0010:fte-0/0/12          Ok
    BBAK8704:fte-0/1/1          -> A0010:fte-0/0/13          Ok
Device ID: BBAK8714, Type: qfx3500
    Fabric: Incoming links:
    A0010:fte-3/0/14           -> BBAK8714:fte-0/1/0          Ok
    A0010:fte-3/0/15           -> BBAK8714:fte-0/1/1          Ok
    Fabric: Outgoing links:
    BBAK8714:fte-0/1/0          -> A0010:fte-3/0/14          Ok
    BBAK8714:fte-0/1/1          -> A0010:fte-3/0/15          Ok
Device ID: A0010, Type: qfxc08-3008

```


Front Card 0 : Incoming links:

ED3705:fte-0/1/2	-> A0010:fte-0/0/0	Ok
ED1487:fte-0/1/2	-> A0010:fte-0/0/1	Ok
ED3683:fte-0/1/2	-> A0010:fte-0/0/2	Ok
BBAK8866:fte-0/1/0	-> A0010:fte-0/0/3	Ok
BBAK8866:fte-0/1/1	-> A0010:fte-0/0/4	Ok
BBAK8777:fte-0/1/0	-> A0010:fte-0/0/5	Ok
BBAK8777:fte-0/1/1	-> A0010:fte-0/0/6	Ok
ED3707:fte-0/1/2	-> A0010:fte-0/0/8	Ok
ED3711:fte-0/1/2	-> A0010:fte-0/0/9	Ok
ED3702:fte-0/1/2	-> A0010:fte-0/0/10	Ok
BBAK8737:fte-0/1/1	-> A0010:fte-0/0/11	Ok
BBAK8704:fte-0/1/0	-> A0010:fte-0/0/12	Ok
BBAK8704:fte-0/1/1	-> A0010:fte-0/0/13	Ok
BBAK8885:fte-0/1/0	-> A0010:fte-0/0/14	Ok
BBAK8885:fte-0/1/1	-> A0010:fte-0/0/15	Ok

Front Card 0 : Outgoing links:

A0010:fte-0/0/8	-> ED3707:fte-0/1/2	Ok
A0010:fte-0/0/9	-> ED3711:fte-0/1/2	Ok
A0010:fte-0/0/10	-> ED3702:fte-0/1/2	Ok
A0010:fte-0/0/11	-> BBAK8737:fte-0/1/1	Ok
A0010:fte-0/0/12	-> BBAK8704:fte-0/1/0	Ok
A0010:fte-0/0/13	-> BBAK8704:fte-0/1/1	Ok
A0010:fte-0/0/14	-> BBAK8885:fte-0/1/0	Ok
A0010:fte-0/0/15	-> BBAK8885:fte-0/1/1	Ok
A0010:fte-0/0/0	-> ED3705:fte-0/1/2	Ok
A0010:fte-0/0/1	-> ED1487:fte-0/1/2	Ok
A0010:fte-0/0/2	-> ED3683:fte-0/1/2	Ok
A0010:fte-0/0/3	-> BBAK8866:fte-0/1/0	Ok
A0010:fte-0/0/4	-> BBAK8866:fte-0/1/1	Ok
A0010:fte-0/0/5	-> BBAK8777:fte-0/1/0	Ok
A0010:fte-0/0/6	-> BBAK8777:fte-0/1/1	Ok

Front Card 3 : Incoming links:

ED3705:fte-0/1/3	-> A0010:fte-3/0/0	Ok
ED1487:fte-0/1/3	-> A0010:fte-3/0/1	Ok
ED3683:fte-0/1/3	-> A0010:fte-3/0/2	Ok
BBAK8854:fte-0/1/0	-> A0010:fte-3/0/3	Ok
BBAK8854:fte-0/1/1	-> A0010:fte-3/0/4	Ok
BBAK8810:fte-0/1/0	-> A0010:fte-3/0/5	Ok
BBAK8810:fte-0/1/1	-> A0010:fte-3/0/6	Ok
BBAK8864:fte-0/1/0	-> A0010:fte-3/0/7	Ok
ED3707:fte-0/1/3	-> A0010:fte-3/0/8	Ok
ED3711:fte-0/1/3	-> A0010:fte-3/0/9	Ok
ED3702:fte-0/1/3	-> A0010:fte-3/0/10	Ok


```

BBAK8864:fte-0/1/1      -> A0010:fte-3/0/11      Ok
BBAK8759:fte-0/1/0      -> A0010:fte-3/0/12      Ok
BBAK8759:fte-0/1/1      -> A0010:fte-3/0/13      Ok
BBAK8714:fte-0/1/0      -> A0010:fte-3/0/14      Ok
BBAK8714:fte-0/1/1      -> A0010:fte-3/0/15      Ok
Front Card 3 : Outgoing links:
A0010:fte-3/0/8         -> ED3707:fte-0/1/3      Ok
A0010:fte-3/0/9         -> ED3711:fte-0/1/3      Ok
A0010:fte-3/0/10        -> ED3702:fte-0/1/3      Ok
A0010:fte-3/0/11        -> BBAK8864:fte-0/1/1      Ok
A0010:fte-3/0/12        -> BBAK8759:fte-0/1/0      Ok
A0010:fte-3/0/13        -> BBAK8759:fte-0/1/1      Ok
A0010:fte-3/0/14        -> BBAK8714:fte-0/1/0      Ok
A0010:fte-3/0/15        -> BBAK8714:fte-0/1/1      Ok
A0010:fte-3/0/0         -> ED3705:fte-0/1/3      Ok
A0010:fte-3/0/1         -> ED1487:fte-0/1/3      Ok
A0010:fte-3/0/2         -> ED3683:fte-0/1/3      Ok
A0010:fte-3/0/3         -> BBAK8854:fte-0/1/0      Ok
A0010:fte-3/0/4         -> BBAK8854:fte-0/1/1      Ok
A0010:fte-3/0/5         -> BBAK8810:fte-0/1/0      Ok
A0010:fte-3/0/6         -> BBAK8810:fte-0/1/1      Ok
A0010:fte-3/0/7         -> BBAK8864:fte-0/1/0      Ok

```

show chassis fabric connectivity device

user@qfabric> **show chassis fabric connectivity device BBAK8714**

```

Device ID: BBAK8714, Type: qfx3500
Fabric: Incoming links:
A0010:fte-3/0/14        -> BBAK8714:fte-0/1/0      Ok
A0010:fte-3/0/15        -> BBAK8714:fte-0/1/1      Ok
Fabric: Outgoing links:
BBAK8714:fte-0/1/0      -> A0010:fte-3/0/14      Ok
BBAK8714:fte-0/1/1      -> A0010:fte-3/0/15      Ok

```

user@qfabric> **show chassis fabric connectivity device A0010**

```

Device ID: A0010, Type: qfxc08-3008
Front Card 0 : Incoming links:
ED3705:fte-0/1/2        -> A0010:fte-0/0/0      Ok
ED1487:fte-0/1/2        -> A0010:fte-0/0/1      Ok
ED3683:fte-0/1/2        -> A0010:fte-0/0/2      Ok

```


BBAK8866:fte-0/1/0	-> A0010:fte-0/0/3	Ok
BBAK8866:fte-0/1/1	-> A0010:fte-0/0/4	Ok
BBAK8777:fte-0/1/0	-> A0010:fte-0/0/5	Ok
BBAK8777:fte-0/1/1	-> A0010:fte-0/0/6	Ok
ED3707:fte-0/1/2	-> A0010:fte-0/0/8	Ok
ED3711:fte-0/1/2	-> A0010:fte-0/0/9	Ok
ED3702:fte-0/1/2	-> A0010:fte-0/0/10	Ok
BBAK8737:fte-0/1/1	-> A0010:fte-0/0/11	Ok
BBAK8704:fte-0/1/0	-> A0010:fte-0/0/12	Ok
BBAK8704:fte-0/1/1	-> A0010:fte-0/0/13	Ok
BBAK8885:fte-0/1/0	-> A0010:fte-0/0/14	Ok
BBAK8885:fte-0/1/1	-> A0010:fte-0/0/15	Ok

Front Card 0 : Outgoing links:

A0010:fte-0/0/8	-> ED3707:fte-0/1/2	Ok
A0010:fte-0/0/9	-> ED3711:fte-0/1/2	Ok
A0010:fte-0/0/10	-> ED3702:fte-0/1/2	Ok
A0010:fte-0/0/11	-> BBAK8737:fte-0/1/1	Ok
A0010:fte-0/0/12	-> BBAK8704:fte-0/1/0	Ok
A0010:fte-0/0/13	-> BBAK8704:fte-0/1/1	Ok
A0010:fte-0/0/14	-> BBAK8885:fte-0/1/0	Ok
A0010:fte-0/0/15	-> BBAK8885:fte-0/1/1	Ok
A0010:fte-0/0/0	-> ED3705:fte-0/1/2	Ok
A0010:fte-0/0/1	-> ED1487:fte-0/1/2	Ok
A0010:fte-0/0/2	-> ED3683:fte-0/1/2	Ok
A0010:fte-0/0/3	-> BBAK8866:fte-0/1/0	Ok
A0010:fte-0/0/4	-> BBAK8866:fte-0/1/1	Ok
A0010:fte-0/0/5	-> BBAK8777:fte-0/1/0	Ok
A0010:fte-0/0/6	-> BBAK8777:fte-0/1/1	Ok

Front Card 3 : Incoming links:

ED3705:fte-0/1/3	-> A0010:fte-3/0/0	Ok
ED1487:fte-0/1/3	-> A0010:fte-3/0/1	Ok
ED3683:fte-0/1/3	-> A0010:fte-3/0/2	Ok
BBAK8854:fte-0/1/0	-> A0010:fte-3/0/3	Ok
BBAK8854:fte-0/1/1	-> A0010:fte-3/0/4	Ok
BBAK8810:fte-0/1/0	-> A0010:fte-3/0/5	Ok
BBAK8810:fte-0/1/1	-> A0010:fte-3/0/6	Ok
BBAK8864:fte-0/1/0	-> A0010:fte-3/0/7	Ok
ED3707:fte-0/1/3	-> A0010:fte-3/0/8	Ok
ED3711:fte-0/1/3	-> A0010:fte-3/0/9	Ok
ED3702:fte-0/1/3	-> A0010:fte-3/0/10	Ok
BBAK8864:fte-0/1/1	-> A0010:fte-3/0/11	Ok
BBAK8759:fte-0/1/0	-> A0010:fte-3/0/12	Ok
BBAK8759:fte-0/1/1	-> A0010:fte-3/0/13	Ok
BBAK8714:fte-0/1/0	-> A0010:fte-3/0/14	Ok


```

BBAK8714:fte-0/1/1      -> A0010:fte-3/0/15      Ok
Front Card 3 : Outgoing links:
A0010:fte-3/0/8         -> ED3707:fte-0/1/3        Ok
A0010:fte-3/0/9         -> ED3711:fte-0/1/3        Ok
A0010:fte-3/0/10        -> ED3702:fte-0/1/3        Ok
A0010:fte-3/0/11        -> BBAK8864:fte-0/1/1      Ok
A0010:fte-3/0/12        -> BBAK8759:fte-0/1/0      Ok
A0010:fte-3/0/13        -> BBAK8759:fte-0/1/1      Ok
A0010:fte-3/0/14        -> BBAK8714:fte-0/1/0      Ok
A0010:fte-3/0/15        -> BBAK8714:fte-0/1/1      Ok
A0010:fte-3/0/0         -> ED3705:fte-0/1/3        Ok
A0010:fte-3/0/1         -> ED1487:fte-0/1/3        Ok
A0010:fte-3/0/2         -> ED3683:fte-0/1/3        Ok
A0010:fte-3/0/3         -> BBAK8854:fte-0/1/0      Ok
A0010:fte-3/0/4         -> BBAK8854:fte-0/1/1      Ok
A0010:fte-3/0/5         -> BBAK8810:fte-0/1/0      Ok
A0010:fte-3/0/6         -> BBAK8810:fte-0/1/1      Ok
A0010:fte-3/0/7         -> BBAK8864:fte-0/1/0      Ok
Ok
A0010:fte-3/0/3         -> BBAK8854:fte-0/1/0      Ok
A0010:fte-3/0/4         -> BBAK8854:fte-0/1/1      Ok
A0010:fte-3/0/5         -> BBAK8810:fte-0/1/0      Ok
A0010:fte-3/0/6         -> BBAK8810:fte-0/1/1      Ok
A0010:fte-3/0/7         -> BBAK8864:fte-0/1/0      Ok

```

show chassis fabric connectivity device device-name slot

user@qfabric> **show chassis fabric connectivity device A0010 slot 3**

```

Device ID: A0010, Type: qfxc08-3008
Front Card 3 : Incoming links:
ED3705:fte-0/1/3        -> A0010:fte-3/0/0        Ok
ED1487:fte-0/1/3        -> A0010:fte-3/0/1        Ok
ED3683:fte-0/1/3        -> A0010:fte-3/0/2        Ok
BBAK8854:fte-0/1/0       -> A0010:fte-3/0/3        Ok
BBAK8854:fte-0/1/1       -> A0010:fte-3/0/4        Ok
BBAK8810:fte-0/1/0       -> A0010:fte-3/0/5        Ok
BBAK8810:fte-0/1/1       -> A0010:fte-3/0/6        Ok
BBAK8864:fte-0/1/0       -> A0010:fte-3/0/7        Ok
ED3707:fte-0/1/3        -> A0010:fte-3/0/8        Ok
ED3711:fte-0/1/3        -> A0010:fte-3/0/9        Ok
ED3702:fte-0/1/3        -> A0010:fte-3/0/10       Ok
BBAK8864:fte-0/1/1       -> A0010:fte-3/0/11       Ok
BBAK8759:fte-0/1/0       -> A0010:fte-3/0/12       Ok

```


BBAK8759:fte-0/1/1	-> A0010:fte-3/0/13	Ok
BBAK8714:fte-0/1/0	-> A0010:fte-3/0/14	Ok
BBAK8714:fte-0/1/1	-> A0010:fte-3/0/15	Ok
Front Card 3 : Outgoing links:		
A0010:fte-3/0/8	-> ED3707:fte-0/1/3	Ok
A0010:fte-3/0/9	-> ED3711:fte-0/1/3	Ok
A0010:fte-3/0/10	-> ED3702:fte-0/1/3	Ok
A0010:fte-3/0/11	-> BBAK8864:fte-0/1/1	Ok
A0010:fte-3/0/12	-> BBAK8759:fte-0/1/0	Ok
A0010:fte-3/0/13	-> BBAK8759:fte-0/1/1	Ok
A0010:fte-3/0/14	-> BBAK8714:fte-0/1/0	Ok
A0010:fte-3/0/15	-> BBAK8714:fte-0/1/1	Ok
A0010:fte-3/0/0	-> ED3705:fte-0/1/3	Ok
A0010:fte-3/0/1	-> ED1487:fte-0/1/3	Ok
A0010:fte-3/0/2	-> ED3683:fte-0/1/3	Ok
A0010:fte-3/0/3	-> BBAK8854:fte-0/1/0	Ok
A0010:fte-3/0/4	-> BBAK8854:fte-0/1/1	Ok
A0010:fte-3/0/5	-> BBAK8810:fte-0/1/0	Ok
A0010:fte-3/0/6	-> BBAK8810:fte-0/1/1	Ok
A0010:fte-3/0/7	-> BBAK8864:fte-0/1/0	Ok

show chassis fabric device

Syntax

```
show chassis fabric device device-name
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Display the fabric management status of devices in your QFabric system.

Options

none—Display the fabric management status for all devices in your QFabric system.

device-name—Display the fabric management status for a specific device in your QFabric system. You can enter either the alias name or the serial number of the device.

Required Privilege Level

admin

RELATED DOCUMENTATION

request chassis fabric fpc	 603
show chassis fabric connectivity	 749
show fabric administration inventory	 780
Understanding Interconnect Devices	 29

List of Sample Output

[show chassis fabric device on page 759](#)

Output Fields

[Table 87 on page 758](#) lists the output fields for the **show chassis fabric device** command. Output fields are listed in the approximate order in which they appear.

Table 87: show chassis fabric device Output Fields

Field Name	Field Description
Device ID	Hardware serial identifier of the QFabric system component.

Table 87: show chassis fabric device Output Fields (*continued*)

Field Name	Field Description
Type	Model number of the QFabric system component. Values include qfxc08-3008 (QFX3008-I Interconnect device), qfx3600-I (QFX3600-I Interconnect device), and supported Node device names such as qfx3500 (QFX3500 Node device) and qfx3600-16q (QFX3600 Node device).
Management status	Displays keepalive status for fabric management processes on a specific device. Values include On and Off .
Hardware status	Displays operational status of the device participating in fabric management (such as Ok).

Sample Output

show chassis fabric device

user@qfabric> **show chassis fabric device**

```
Device ID: node2, Type: qfx3500
Management status: On, Hardware status: Ok

Device ID: node3, Type: qfx3500
Management status: On, Hardware status: Ok

Device ID: node0, Type: qfx3500
Management status: On, Hardware status: Ok

Device ID: node1, Type: qfx3500
Management status: On, Hardware status: Ok
```


show chassis lcd

List of Syntax

[show chassis lcd \(EX Series\) on page 760](#)

[show chassis lcd \(QFX Series\) on page 760](#)

[show chassis lcd \(OCX Series\) on page 760](#)

show chassis lcd (EX Series)

```
show chassis lcd
<fpc-slot fpc-slot-number>
<menu <(all-members | local | member member-id)>>
```

show chassis lcd (QFX Series)

```
show chassis lcd
<fpc-slot fpc-slot-number>
<interconnect-device device-id>
<node-device device-id>
```

show chassis lcd (OCX Series)

```
show chassis lcd
<fpc-slot fpc-slot-number>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.

menu option introduced in Junos OS Release 10.2 for EX Series switches.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Command introduced in Junos OS Release 13.1 for QFabric systems.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display the information that appears on the LCD panel of EX3200, EX3300, EX4200, EX4500, EX6200, and EX8200 switches, XRE200 External Routing Engines, QFX Series standalone switches, OCX Series switches, and Interconnect devices and Node devices within a QFabric system. Display the status of the currently selected port parameter of the Status LED for each network port on the device.

Options

none—Display the information that appears on the LCD panel (for any EX Series member switch in a Virtual Chassis or for XRE200 External Routing Engines, display the information for all Virtual Chassis members). Display the status of the currently selected port parameter of the Status LED for each network port.

fpc-slot <fpc-slot-number>—(Optional) Display the information as follows:

- (EX3200, EX3300, EX4200, and EX4500 switches, QFX Series, or OCX Series) Display the information that appears on the LCD panel for either an FPC slot with no **fpc-slot-number** value specified or for the FPC slot specified by **fpc-slot 0**. **fpc-slot** refers to the switch itself and **0** is the only valid value for **fpc-slot-number**. Output for these options is the same as for the **none** option.

Also display the status of the currently selected port parameter of the Status LED for each network port.

- (EX Series Virtual Chassis member switches or XRE200 External Routing Engines) If no **fpc-slot-number** value is specified, display the information that appears on the LCD panel for all members of the Virtual Chassis. Output for this option is the same as for the **none** option. If the **fpc-slot-number** value is specified (it equals the **member-id** value), display the information for the specified member.

Also display the status of the currently selected port parameter of the Status LED for each network port.

- (EX6200 or EX8200 switches)—Display the information that appears on the LCD panel for the line card in the line-card slot specified by the **fpc-slot-number** value.

Also display the status of the currently selected port parameter of the Status LED for each network port.

interconnect-device device-id—(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Interconnect device.

menu—(Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel.

menu all-members—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for all Virtual Chassis members.

menu local—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the Virtual Chassis member from which you issued the command.

menu member member-id—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the specified Virtual Chassis member.

node-device device-id—(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Node device.

Required Privilege Level

view

RELATED DOCUMENTATION

[LCD Panel in EX3200 Switches](#)

[LCD Panel in EX4200 Switches](#)

[LCD Panel in EX4500 Switches](#)

[LCD Panel in an EX8200 Switch](#)

[LCD Panel in an XRE200 External Routing Engine](#)

[Configuring the LCD Panel on EX Series Switches \(CLI Procedure\)](#)

[set chassis display message](#) | [649](#)

List of Sample Output

[show chassis lcd \(Two-Member EX4200 Virtual Chassis\) on page 763](#)

[show chassis lcd fpc-slot 1 \(EX4200 Virtual Chassis\) on page 765](#)

[show chassis lcd \(EX8200 Switch\) on page 766](#)

[show chassis lcd fpc-slot 2 \(EX8200 Switch\) on page 768](#)

[show chassis lcd menu \(EX4200 Switch\) on page 768](#)

[show chassis lcd menu \(EX8200 Switch\) on page 769](#)

[show chassis lcd \(QFX3500 Switches\) on page 769](#)

[show chassis lcd \(XRE200 External Routing Engine in EX8200 Virtual Chassis\) on page 770](#)

[show chassis lcd interconnect-device \(QFabric Systems\) on page 774](#)

[show chassis lcd node-device \(QFabric Systems\) on page 777](#)

Output Fields

[Table 88 on page 762](#) lists the output fields for the **show chassis lcd** command. Output fields are listed in the approximate order in which they appear.

Table 88: show chassis lcd Output Fields

Field Name	Field Description
membernumber (XRE200 External Routing Engine)	Member ID of the device whose content is being displayed.
Front panel contents for slot Front panel contents (EX6200, EX8200 switch, XRE200 External Routing Engine, and QFX Series)	<p>FPC slot number of the switch whose content is being displayed. The number is always 0, except for EX4200 switches in a Virtual Chassis, where it is the member ID value.</p> <p>On EX6200 switches, EX8200 switches, and XRE200 External Routing Engines, no slot number is displayed.</p> <p>On XRE200 External Routing Engines, this field appears under the member number field for each member device in the EX8200 Virtual Chassis.</p>

Table 88: show chassis lcd Output Fields (*continued*)

Field Name	Field Description
LCD screen	<p>The first line displays the hostname (for Virtual Chassis members, displays the member ID, the current role, and hostname; for EX8200 switches, displays RE and the hostname). The second line displays the currently selected port parameter of the Status LED and the alarms counter. The Status LED port parameters are:</p> <ul style="list-style-type: none"> • ADM—Administrative • SPD—Speed • DPX—Duplex • POE—Power over Ethernet (EX3200 and EX4200 switches only)
LEDs status	Current state of the Alarms, System, and Master LEDs (chassis status LEDs).
Interface	Names of the interfaces on the switch.
LED (ADM/SPD/DPX/POE)	<p>State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are:</p> <p>NOTE: The XRE200 External Routing Engine always displays the NA parameter. The QFX Series products do not have any of the port parameters listed below.</p> <ul style="list-style-type: none"> • ADM—Administrative • SPD—Speed • DPX—Duplex • NA—Not applicable. • POE—Power over Ethernet
fpcx	On standalone EX Series and QFX Series switches, always 0 . On EX Series Virtual Chassis member switches, member ID of the Virtual Chassis member whose LCD menu is displayed.

Sample Output

show chassis lcd (Two-Member EX4200 Virtual Chassis)

```
user@switch> show chassis lcd
```

```
Front panel contents for slot: 0
-----
LCD screen:
  00:BK switch1
```



```

    LED:SPD ALARM 00
LEDs status:
    Alarms LED: Off
    System LED: Green
    Master LED: Off
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0       Off
ge-0/0/1       Off
ge-0/0/2       Off
ge-0/0/3       Off
ge-0/0/4       Off
ge-0/0/5       Off
ge-0/0/6       Off
ge-0/0/7       Off
ge-0/0/8       Off
ge-0/0/9       Off
ge-0/0/10      Off
ge-0/0/11      Off
ge-0/0/12      Off
ge-0/0/13      Off
ge-0/0/14      Off
ge-0/0/15      Off
ge-0/0/16      Off
ge-0/0/17      Off
ge-0/0/18      Off
ge-0/0/19      Off
ge-0/0/20      Off
ge-0/0/21      Off
ge-0/0/22      Off
ge-0/0/23      Off
Front panel contents for slot: 1
-----
LCD screen:
    01:RE switch2
    LED:SPD ALARM 01
LEDs status:
    Alarms LED: Yellow
    System LED: Green
    Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-1/0/0       Off
ge-1/0/1       Off

```



```

ge-1/0/2      Off
ge-1/0/3      Off
ge-1/0/4      Off
ge-1/0/5      Off
ge-1/0/6      Off
ge-1/0/7      Off
ge-1/0/8      Off
ge-1/0/9      Off
ge-1/0/10     Off
ge-1/0/11     Off
ge-1/0/12     Off
ge-1/0/13     Off
ge-1/0/14     Off
ge-1/0/15     Off
ge-1/0/16     Off
ge-1/0/17     Off
ge-1/0/18     Off
ge-1/0/19     Off
ge-1/0/20     Off
ge-1/0/21     Off
ge-1/0/22     Off
ge-1/0/23     Off

```

The output for the **show chassis lcd fpc-slot** command is the same as the output for the **show chassis lcd** command.

show chassis lcd fpc-slot 1 (EX4200 Virtual Chassis)

```
user@switch> show chassis lcd fpc-slot 1
```

```

Front panel contents for slot: 1
-----
LCD screen:
  01:RE switch2
  LED:SPD ALARM 01
LEDs status:
  Alarms LED: Yellow
  System LED: Green
  Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-1/0/0      Off
ge-1/0/1      Off

```



```

ge-1/0/2      Off
ge-1/0/3      Off
ge-1/0/4      Off
ge-1/0/5      Off
ge-1/0/6      Off
ge-1/0/7      Off
ge-1/0/8      Off
ge-1/0/9      Off
ge-1/0/10     Off
ge-1/0/11     Off
ge-1/0/12     Off
ge-1/0/13     Off
ge-1/0/14     Off
ge-1/0/15     Off
ge-1/0/16     Off
ge-1/0/17     Off
ge-1/0/18     Off
ge-1/0/19     Off
ge-1/0/20     Off
ge-1/0/21     Off
ge-1/0/22     Off
ge-1/0/23     Off

```

show chassis lcd (EX8200 Switch)

```
user@switch> show chassis lcd
```

```

Front panel contents:
-----
LCD screen:
  RE st-8200-r
  LED:ADM ALARM 01
LEDs status:
  Alarms LED: Yellow
  System LED: Yellow
  Master LED: Green
Interface      LED(ADM/SPD/DPX)
-----
ge-0/0/0      Off
ge-0/0/1      Off
ge-0/0/2      Off
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Off

```


ge-0/0/6	Off
ge-0/0/7	Off
ge-0/0/8	Off
ge-0/0/9	Off
ge-0/0/10	Off
ge-0/0/11	Off
ge-0/0/12	Off
ge-0/0/13	Off
ge-0/0/14	Off
ge-0/0/15	Off
ge-0/0/16	Off
ge-0/0/17	Off
ge-0/0/18	Off
ge-0/0/19	Off
ge-0/0/20	Off
ge-0/0/21	Off
ge-0/0/22	Off
ge-0/0/23	Off
ge-0/0/24	Off
ge-0/0/25	Off
ge-0/0/26	Off
ge-0/0/27	Off
ge-0/0/28	Off
ge-0/0/29	Off
ge-0/0/30	Off
ge-0/0/31	Off
ge-0/0/32	Off
ge-0/0/33	Off
ge-0/0/34	Off
ge-0/0/35	Off
ge-0/0/36	Off
ge-0/0/37	Off
ge-0/0/38	Off
ge-0/0/39	Off
ge-0/0/40	Off
ge-0/0/41	Off
ge-0/0/42	Off
ge-0/0/43	Off
ge-0/0/44	Off
ge-0/0/45	Off
ge-0/0/46	Off
ge-0/0/47	Off
xe-2/0/0	Off
xe-2/0/1	Off

xe-2/0/2	Off
xe-2/0/3	Off
xe-2/0/4	Off
xe-2/0/5	Off
xe-2/0/6	Off
xe-2/0/7	Off
xe-3/0/0	Off
xe-3/0/1	Off
xe-3/0/2	Off
xe-3/0/3	Off
xe-3/0/4	Off
xe-3/0/5	Off
xe-3/0/6	Off
xe-3/0/7	Off
xe-5/0/0	Off
xe-5/0/1	Off
xe-5/0/2	Off
xe-5/0/3	Off
xe-5/0/4	Off
xe-5/0/5	Off
xe-5/0/6	On
xe-5/0/7	On
xe-7/0/5	Off

show chassis lcd fpc-slot 2 (EX8200 Switch)

show chassis lcd fpc-slot 2

Interface	LED (ADM/SPD/DPX)

xe-2/0/0	Off
xe-2/0/1	Off
xe-2/0/2	Off
xe-2/0/3	Off
xe-2/0/4	Off
xe-2/0/5	Off
xe-2/0/6	Off
xe-2/0/7	Off

show chassis lcd menu (EX4200 Switch)

user@switch> **show chassis lcd menu**


```
fpc0:
-----
status-menu
status-menu vcp-status
status-menu power-status
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu vc-uplink-config
maintenance-menu factory-default
```

On an EX4200 switch in a Virtual Chassis, the output for the **show chassis lcd menu all-members** command is the same as the output for the **show chassis lcd menu** command.

show chassis lcd menu (EX8200 Switch)

```
user@switch> show chassis lcd menu
```

```
status-menu
status-menu sf-status1-menu
status-menu sf-status2-menu
status-menu psu-status1-menu
status-menu psu-status2-menu
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu factory-default
```

show chassis lcd (QFX3500 Switches)

```
user@switch> show chassis lcd
```

```
Front panel contents for slot: 0
-----
LCD screen:
00:RE switch
ALARM 01
```



```

LEDs status:
Status/Beacon LED: Yellow Blinking
Interface STATUS LED ACTIVITY LED
-----
fte-0/1/0 Off Off

```

show chassis lcd (XRE200 External Routing Engine in EX8200 Virtual Chassis)

user@external-routing-engine> show chassis lcd

```

member0:
-----
Front panel contents:
-----
LCD screen:
    RE ex8200-member0
    LED:ADM ALARM 04
LEDs status:
    Alarms LED: Red
    System LED: Yellow
    Master LED: Green

member1:
-----

member8:
-----
Front panel contents:
-----
LCD screen:
    BACKUP

member9:
-----
Front panel contents:
-----
LCD screen:
    09:RE xre200-member9
    LED: NA ALARM 01
Interface          LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0           On
ge-0/0/1           On
ge-0/0/2           On

```


ge-0/0/3	On
ge-0/0/4	Off
ge-0/0/5	Off
ge-0/0/6	Off
ge-0/0/7	Off
ge-0/0/8	Off
ge-0/0/9	Off
ge-0/0/10	On
ge-0/0/11	Off
ge-0/0/12	Off
ge-0/0/13	Off
ge-0/0/14	Off
ge-0/0/15	Off
ge-0/0/16	Off
ge-0/0/17	Off
ge-0/0/18	Off
ge-0/0/19	Off
ge-0/0/20	Off
ge-0/0/21	Off
ge-0/0/22	Off
ge-0/0/23	Off
ge-0/0/24	Off
ge-0/0/25	Off
ge-0/0/26	Off
ge-0/0/27	Off
ge-0/0/28	Off
ge-0/0/29	Off
ge-0/0/30	Off
ge-0/0/31	Off
ge-0/0/32	Off
ge-0/0/33	Off
ge-0/0/34	Off
ge-0/0/35	Off
ge-0/0/36	Off
ge-0/0/37	Off
ge-0/0/38	Off
ge-0/0/39	Off
ge-0/0/40	On
ge-0/0/41	On
ge-0/0/42	On
ge-0/0/43	On
ge-0/0/44	On
ge-0/0/45	On
ge-0/0/46	On

ge-0/0/47	On
ge-16/0/0	On
ge-16/0/1	Off
ge-16/0/2	On
ge-16/0/3	Off
ge-16/0/4	On
ge-16/0/5	Off
ge-16/0/6	On
ge-16/0/7	Off
ge-16/0/8	Off
ge-16/0/9	Off
ge-16/0/10	Off
ge-16/0/11	Off
ge-16/0/12	Off
ge-16/0/13	On
ge-16/0/14	Off
ge-16/0/15	On
ge-16/0/16	Off
ge-16/0/17	On
ge-16/0/18	On
ge-16/0/19	On
ge-16/0/20	On
ge-16/0/21	Off
ge-16/0/22	On
ge-16/0/23	Off
ge-16/0/24	Off
ge-16/0/25	Off
ge-16/0/26	On
ge-16/0/27	Off
ge-16/0/28	Off
ge-16/0/29	Off
ge-16/0/30	On
ge-16/0/31	Off
ge-16/0/32	On
ge-16/0/33	On
ge-16/0/34	On
ge-16/0/35	Off
ge-16/0/36	On
ge-16/0/37	Off
ge-16/0/38	Off
ge-16/0/39	Off
ge-16/0/40	Off
ge-16/0/41	Off
ge-16/0/42	On

ge-16/0/43	Off
ge-16/0/44	Off
ge-16/0/45	Off
ge-16/0/46	Off
ge-16/0/47	Off
xe-19/0/0	Off
xe-19/0/1	On
xe-19/0/2	On
xe-19/0/3	On
xe-19/0/4	On
xe-19/0/5	On
ge-22/0/0	Off
ge-22/0/1	Off
ge-22/0/2	On
ge-22/0/3	Off
ge-22/0/4	On
ge-22/0/5	On
ge-22/0/6	On
ge-22/0/7	On
ge-22/0/8	Off
ge-22/0/9	Off
ge-22/0/10	Off
ge-22/0/11	Off
ge-22/0/12	Off
ge-22/0/13	Off
ge-22/0/14	Off
ge-22/0/15	Off
ge-22/0/16	On
ge-22/0/17	Off
ge-22/0/18	On
ge-22/0/19	Off
ge-22/0/20	On
ge-22/0/21	Off
ge-22/0/22	On
ge-22/0/23	Off
ge-22/0/24	On
ge-22/0/25	Off
ge-22/0/26	Off
ge-22/0/27	Off
ge-22/0/28	Off
ge-22/0/29	Off
ge-22/0/30	Off
ge-22/0/31	Off
ge-22/0/32	On


```

ge-22/0/33      Off
ge-22/0/34      On
ge-22/0/35      Off
ge-22/0/36      Off
ge-22/0/37      Off
ge-22/0/38      Off
ge-22/0/39      Off
ge-22/0/40      Off
ge-22/0/41      Off
ge-22/0/42      Off
ge-22/0/43      Off
ge-22/0/44      Off
ge-22/0/45      Off
ge-22/0/46      Off
ge-22/0/47      Off

```

show chassis lcd interconnect-device (QFabric Systems)

show chassis lcd interconnect-device IC-F1012

```

                                Front Panel Module Information
                                -----
                                LCD screen:
                                IC-F1012          3 Alarms active

LEDs status:
  Status LED: Green
  Power LED : Green
  Major Alarm LED: off
  Minor Alarm LED: Yellow
  Fan 0 LED : Green
  Fan 1 LED : Green
  Fan 2 LED : Green
  Fan 3 LED : Green
  Fan 4 LED : Green
  Fan 5 LED : Green
  Fan 6 LED : Green
  Fan 7 LED : Green
  Fan 8 LED : Green
  Fan 9 LED : Green
  PEM 0 LED : Green
  PEM 1 LED : Green
  PEM 2 LED : Green
  PEM 3 LED : off
  PEM 4 LED : off

```


PEM 5 LED : off

LED info for: CB - 0

LEDs status:

Status LED: Green

Mastership LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:pme0 :	Green	N/A
IC-F1012:pme1 :	Green	N/A
IC-F1012:pme2 :	off	N/A
IC-F1012:pme3 :	off	N/A

LED info for: CB - 1

LEDs status:

Status LED: Green

Mastership LED: Amber

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:pme0 :	Green	N/A
IC-F1012:pme1 :	Green	N/A
IC-F1012:pme2 :	off	N/A
IC-F1012:pme3 :	off	N/A

LED info for: FC 0 FPC - 0

LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-0/0/0	Green	N/A
IC-F1012:fte-0/0/1	Green	N/A
IC-F1012:fte-0/0/2	Green	N/A
IC-F1012:fte-0/0/3	Green	N/A
IC-F1012:fte-0/0/4	Green	N/A

LED info for: FC 1 FPC - 1

LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-1/0/0	Green	N/A
IC-F1012:fte-1/0/1	Green	N/A
IC-F1012:fte-1/0/2	Green	N/A
IC-F1012:fte-1/0/3	Green	N/A
IC-F1012:fte-1/0/4	Green	N/A

LED info for: RC 0 FPC - 8

LEDs status:

Status LED: Green

LED info for: RC 1 FPC - 9

LEDs status:

Status LED: Green

LED info for: RC 2 FPC - 10

LEDs status:

Status LED: Green

LED info for: RC 3 FPC - 11

LEDs status:

Status LED: Green

LED info for: RC 4 FPC - 12

LEDs status:

Status LED: Green

LED info for: RC 5 FPC - 13

LEDs status:

Status LED: Green

LED info for: RC 6 FPC - 14

LEDs status:

Status LED: Green


```
LED info for: RC 7 FPC - 15
```

```
-----
```

```
LEDs status:
```

```
Status LED: Green
```

show chassis lcd node-device (QFabric Systems)

show chassis lcd node-device P3774-C

```
Front panel contents for: P3774-C
```

```
-----
```

```
LCD screen:
```

```
P3774-C
```

```
LEDs status:
```

```
Status/Beacon LED: Yellow Blinking
```

Interface	STATUS LED	LINK/ACTIVITY LED

P3774-C:xe-0/0/6	Green	Green
P3774-C:xe-0/0/7	Green	Green
P3774-C:ge-0/0/10	Green	Green
P3774-C:ge-0/0/11	Green	Green Blinking
P3774-C:ge-0/0/12	Green	Off
P3774-C:ge-0/0/13	Green	Green Blinking
P3774-C:ge-0/0/20	Green	Green
P3774-C:ge-0/0/21	Green	Green
P3774-C:ge-0/0/22	Green	Green Blinking
P3774-C:ge-0/0/23	Green	Off
P3774-C:ge-0/0/30	Green	Green
P3774-C:ge-0/0/31	Green	Green
P3774-C:ge-0/0/32	Green	Green Blinking
P3774-C:ge-0/0/33	Green	Green Blinking
P3774-C:fte-0/1/0	Green	Green
P3774-C:fte-0/1/1	Green	Green Blinking
P3774-C:fte-0/1/2	Green	Green Blinking
P3774-C:fte-0/1/3	Green	Green

show chassis nonstop-upgrade node-group

Syntax

```
show chassis nonstop-upgrade node-group node-group-name
```

Release Information

Command introduced in Junos OS Release 12.2 for the QFX Series.

Description

Display the status of the Node group after the most recent nonstop software upgrade (NSSU).

Required Privilege Level

view

RELATED DOCUMENTATION

- [Performing a Nonstop Software Upgrade on the QFabric System | 533](#)
- [request system software nonstop-upgrade | 639](#)

List of Sample Output

[show chassis nonstop-upgrade node-group on page 779](#)

Output Fields

[Table 89 on page 778](#) lists the output fields for the **show chassis nonstop-upgrade node-group** command. Output fields are listed in the approximate order in which they appear.

Table 89: show chassis nonstop-upgrade node-group Output Fields

Field Name	Field Description
Item	Node device slot number.
Status	State of Node device: <ul style="list-style-type: none"> • Error—Node device is in an error state. • Offline—Node device is powered down. • Online—Node device is online and running.
Reason	Reason for the state (if the line card is offline).

Sample Output

```
show chassis nonstop-upgrade node-group
```

```
user@qfabric> show chassis nonstop-upgrade node-group NW-NG-0
```

Item	Status	Reason
P1550-C	Online	

show fabric administration inventory

Syntax

```
show fabric administration inventory
<brief | detail | summary | terse>
<director-group (status)>
<infrastructure (fabric-controls | fabric-managers | diagnostic-routing-engines)>
<interconnect-devices interconnect-device-name>
<node-devices node-device-name>
<node-groups node-group-name>
<summary>
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Display all devices that belong to the QFabric system. You can narrow the level of output by specifying a device type.

NOTE: If your Node devices do not appear in the output of the **show fabric administration inventory** command, check the cabling of your system.

Options

none—Display all devices within a QFabric system.

brief | detail | summary | terse—(Optional) Display the specified level of output.

director-group (status)—(Optional) Display the status for the Director group within a QFabric system.

infrastructure (fabric-controls | fabric-managers | diagnostic-routing-engines)—(Optional) Display information for the fabric control Routing Engine, fabric manager Routing Engine, and diagnostic Routing Engine running on the Director group for the QFabric system.

interconnect-devices *interconnect-device-name*—(Optional) Display a specific Interconnect device within a QFabric system.

node-devices *node-device-name*—(Optional) Display a specific Node device within a QFabric system.

node-groups *node-group-name*—(Optional) Display a specific Node group within a QFabric system.

Required Privilege Level

admin

RELATED DOCUMENTATION

Performing the QFabric System Initial Setup on a QFX3100 Director Group	363
Configuring Aliases for the QFabric System	388
Configuring Node Groups for the QFabric System	402
show fabric administration inventory infrastructure	793
show fabric administration inventory director-group status	786
show fabric administration inventory node-devices	800
show fabric administration inventory node-groups	802
show fabric administration inventory interconnect-devices	797
show fabric inventory	807

List of Sample Output

- [show fabric administration inventory on page 783](#)
- [show fabric administration inventory detail on page 784](#)
- [show fabric administration inventory summary on page 785](#)

Output Fields

Table 90 on page 781 lists the output fields for the **show fabric administration inventory** command. Output fields are listed in the approximate order in which they appear.

Table 90: show fabric administration inventory Output Fields

Field Name	Field Description	Level of Out
Item	Type of QFabric system component being viewed. Possible values include Node group , Interconnect device , Fabric control , Fabric manager , Diagnostic routing engine , Director group , and Ungrouped Node device .	detail none
Identifier	Hardware serial identifier of a QFabric system component. When you configure an alias name for a component, the ID is displayed.	detail none
Connection	Status of a QFabric system component: either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the listed component.	detail none
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured , Failed (with details about the failure), Pending (in the process of being written or retried), or Unknown .	detail none

Table 90: show fabric administration inventory Output Fields (continued)

Field Name	Field Description	Level of Out
Node group	Name of the Node groups associated with the QFabric system, and the Node devices assigned to each Node group. The group can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the devices in the group. This field also displays the serial ID for the Node group and the status for the Node group.	detail none
Interconnect device	Name of the Interconnect devices associated with the Node group. The device can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the device. This field also displays the serial ID and configuration status for the Interconnect device.	detail none
Fabric manager	Name of the primary virtual Junos Routing Engine associated with the QFabric system. The fabric manager Routing Engine can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the fabric manager Routing Engine.	detail none
Fabric control	Name of the virtual Junos Routing Engines responsible for route selection within a QFabric system partition. The fabric control Routing Engine can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the fabric control Routing Engine.	detail none
Diagnostic routing engine	Name of the virtual Junos Routing Engine responsible for troubleshooting and diagnostic utilities within a QFabric system partition. The diagnostic Routing Engine can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the diagnostic Routing Engine.	detail none
Director group	Identifier for the Director devices that are part of the Director group in a QFabric system. Each Director device can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the device.	detail none
Connection uptime	Length of time the component has been operational. The time is listed in days, hours, minutes, and seconds (D+HH:MM:SS).	detail
Network domain	Indicates whether a Node group is a network Node group (Yes) or a server Node group (No).	detail
Member id	Member identification number for a Node device within a Node group.	detail
Node group master	Indicates whether or not a Node device acts as a master device within a Node group. Values for this field are Yes or No .	detail

Table 90: show fabric administration inventory Output Fields (*continued*)

Field Name	Field Description	Level of Out
Configuration checkout failed	Provides troubleshooting details about a failed configuration on a component. This field includes the following output: <ul style="list-style-type: none"> • Edit path—Displays the configuration hierarchy level where the problem occurs. • Statement—Displays the configuration statement that causes the problem. • Message—Provides a suggested workaround to resolve the problem. 	detail
Member	Name of a Routing Engine allocated to an Interconnect device.	detail
Master	Indicates whether or not a Routing Engine acts as a master device within an Interconnect device. Values are Yes or No .	detail
Total Node devices	Number of connected and disconnected Node devices in the QFabric system.	summary
Total connected Node devices	Number of available Node devices in the QFabric system.	summary
Total disconnected Node devices	Number of unavailable Node devices in the QFabric system.	summary
Total Interconnect devices	Number of Interconnect devices in the QFabric system.	summary

Sample Output

show fabric administration inventory

```
user@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Ungrouped Node device			
Node6	BBAK8979	Disconnected	
Node group			
P3359-C		Connected	Configured
P3359-C		Connected	
P3865-C		Connected	Configured
P3865-C		Connected	
RSNG-1		Connected	Configured

Node-3	BBAK8276	Connected	
Node-4	BBAK8273	Connected	
NW-NG-0		Connected	Configured
Node-0	BBAK8309	Connected	
Node-1	BBAK8283	Connected	
Interconnect device			
IC-F1032		Connected	Configured
F1032/RE0		Connected	
F1032/RE1		Connected	
IC-F1092		Connected	Configured
F1092/RE0		Connected	
F1092/RE1		Connected	
Fabric manager			
FM-0		Connected	Configured
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured
Diagnostic routing engine			
DRE-0		Connected	Configured
Director group			
0281112011000023		Connected	
0281112011000082		Connected	

show fabric administration inventory detail

user@qfabric> show fabric administration inventory detail

```
Node group: NW-NG-0, Connected, Configured
Connection uptime: 2+22:40:46
Network domain: Yes
```

```
Node device: node01, Connected,
Member id: 0
```

```
Node group: RSNG, Connected, Configured
Connection uptime: 1:20:22
```

```
Node device: node02, Connected,
Member id: 0
```

```
Node device: node03, Connected,
Member id: 1
Node group master: Yes
```



```

Node group: BBAK0423, Connected, Failed (invalid configuration)
  Connection uptime: 0:01:06

Configuration checkout failed:
  Edit path: edit ethernet-switching-options
  Statement: analyzer
    Message: Vlan vlan1 which is configured as output vlan for analyzer session
    an1 contains untagged interface.
      The analyzer output vlan should only contain tagged interface.

Node device: node0, Connected,
  Member id: 0
  Node group master: Yes

Interconnect device: IC-F4912, Connected, Configured
  Connection uptime: 0:08:05

Member: F4912/RE0, Connected,
  Connection uptime: 0:08:05
  Master: Yes

Fabric manager: FM-0, Connected, Configured
  Connection uptime: 2+22:40:47

Fabric control: FC-0, Connected, Configured
  Connection uptime: 2+22:17:38

Fabric control: FC-1, Connected, Configured
  Connection uptime: 2+22:17:57

Diagnostic routing engine: DRE-0, Connected, Configured
  Connection uptime: 2+22:39:56

```

show fabric administration inventory summary

user@qfabric> **show fabric administration inventory summary**

```

Total Node devices: 3
Total connected Node devices: 3
Total disconnected Node devices: 0
Total Interconnect devices: 1

```


show fabric administration inventory director-group status

Syntax

```
show fabric administration inventory director-group status (target director-device-name | all)
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Display the status of Director devices that belong to a QFabric system Director group.

Options

none—Display the status of all Director devices within a QFabric system.

all—Display the status of all Director devices within a QFabric system.

target *director-device-name*—Display the status of a specific Director device within a QFabric system.

Required Privilege Level

admin

RELATED DOCUMENTATION

- [request fabric administration director-group change-master | 608](#)
- [Performing the QFabric System Initial Setup on a QFX3100 Director Group | 363](#)
- [Understanding the Director Group | 26](#)

List of Sample Output

[show fabric administration inventory director-group status on page 788](#)

[show fabric administration inventory director-group status target on page 791](#)

Output Fields

[Table 91 on page 786](#) lists the output fields for the **show fabric administration inventory director-group status** command. Output fields are listed in the approximate order in which they appear.

Table 91: show fabric administration inventory director-group status Output Fields

Field Name	Field Description
Director Group Status	Timestamp for the Director group status report.

Table 91: show fabric administration inventory director-group status Output Fields (*continued*)

Field Name	Field Description
Member	Name of the Director device.
Status	Current operational mode of the Director device: online or offline .
Role	High availability operational role of the Director device: master or standby .
Mgmt Address	Management IP address of the Director device.
CPU	Percentage of CPU processing memory being used by the Director device.
Free Memory	Available storage memory on the Director device.
VMs	Number of virtual machines operating on the Director device. A Routing Engine issue exists on a particular Director device when the system displays a star (*) in the VMs column along with the following message: Error in retrieving VM Count in dgX .
Up Time	Length of time the Director device has been operating.
Device Id/Alias	Name or identifier of the Director device.
Master Services	Operational status of the database server, load balancer, and QFabric partition address.
Director Group Managed Services	Operational status of the shared file system, network file system, virtual machine server, and DHCP load balancer.
Hard Drive Status	Operational status of the Director device hard drive, including information about the volume identifier, physical identifiers, and SCSI identifiers. There is also status information for the drive partitions, including directory size, available and used drive space, utilization, directory locations, and drive resynchronization progress. A hard drive issue exists when the system displays one of the following messages: Error in retrieving Hard disk status or Error in retrieving Hard disk storage status .
Director Group Processes	Operational status of the Director group processes, such as device managers, SSH, NFS, FTP, system log messages, and SNMP.
Interface Link Status	Operational status of the Director device interfaces, such as the management interface, the control plane bridge interface, the control plane LAG, the control plane links (where the first number represents the Ethernet module [0 or 1] and the second number represents the port [0 - 3]), and the inter-Director crossover LAG. The state of the interfaces can be up or down . The Active designation indicates that a control plane link is active.

Sample Output

show fabric administration inventory director-group status

user@qfabric> **show fabric administration inventory director-group status**

Director Group Status Thu Aug 2 17:36:34 UTC 2012

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.94.215.38	8%	15191684k	4	6 days, 06:24 hrs
dg1	online	backup	10.94.215.39	7%	17733160k	3	6 days, 06:24 hrs

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:0C72369323C3837B	optimal
Physical ID:0	online
Physical ID:1	online
Resync Progress Remaining:0	0%
Resync Progress Remaining:1	0%

Size	Used	Avail	Used%	Mounted on
423G	9.4G	391G	3%	/
99M	16M	79M	17%	/boot
93G	11G	83G	12%	/pbdata

Director Group Processes

Director Group Manager	online	
Partition Manager	online	
Software Mirroring	online	
Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status

Management Interface	up	
Control Plane Bridge	up	
Control Plane LAG	up	
CP Link [0/2]	up	Active
CP Link [0/1]	up	Active
CP Link [0/0]	up	Active
CP Link [1/2]	down	
CP Link [1/1]	down	
CP Link [1/0]	down	
Crossover LAG	up	
CP Link [0/3]	up	
CP Link [1/3]	up	

Member	Device Id/Alias	Status	Role
dg1	0281052011000032	online	backup

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

```

Volume ID:0C9BA4D5F21511B9      optimal
Physical ID:0                      online
Physical ID:1                      online
Resync Progress Remaining:0        0%
Resync Progress Remaining:1        0%

```

```

Size  Used Avail Used% Mounted on
----  -
423G  9.8G 391G   3%   /
99M   16M  79M   17%  /boot
93G   11G  83G   12%  /pbdata

```

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager          online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         online
DHCP Server master          online    backup
FTP Server                  online
Syslog                      online
Distributed Management      online
SNMP Trap Forwarder         online
SNMP Process                offline
Platform Management         online

```

Interface Link Status

```

-----
Management Interface        up
Control Plane Bridge        up
Control Plane LAG           up
CP Link [0/2]               up      Active
CP Link [0/1]               up      Active
CP Link [0/0]               up      Active
CP Link [1/2]               down
CP Link [1/1]               down
CP Link [1/0]               down
Crossover LAG               up
CP Link [0/3]               up
CP Link [1/3]               up

```


show fabric administration inventory director-group status target

```
user@qfabric> show fabric administration inventory director-group status target 0281052011000004
```

```
Director Group Status Thu Aug 16 02:25:37 UTC 2012
```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	backup	10.94.195.109	7%	10009364k	3	23:44 hrs
dg1	online	master	10.94.195.110	9%	6120712k	4	1 day, 33:17 mins

Member	Device Id/Alias	Status	Role
dg1	0281052011000004	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Physical ID:0	online
Physical ID:1	online

Size	Used	Avail	Used%	Mounted on
423G	7.2G	394G	2%	/
99M	20M	75M	21%	/boot
93G	8.8G	85G	10%	/pbdata

Director Group Processes

Director Group Manager	online
Partition Manager	online
Software Mirroring	online

Shared File System master	online	
Secure Shell Process	online	
Network File System	online	
DHCP Server master	online	master
FTP Server	online	
Syslog	online	
Distributed Management	online	
SNMP Trap Forwarder	online	
SNMP Process	online	
Platform Management	online	

Interface Link Status

Management Interface	up	
Control Plane Bridge	up	
Control Plane LAG	up	
CP Link [0/2]	up	Active
CP Link [0/1]	up	Active
CP Link [0/0]	up	Active
CP Link [1/2]	up	
CP Link [1/1]	up	
CP Link [1/0]	up	
Crossover LAG	up	
CP Link [0/3]	up	
CP Link [1/3]	up	

show fabric administration inventory infrastructure

Syntax

```
show administrator inventory infrastructure
<brief | detail>
(fabric-controls | fabric-managers | diagnostic-routing-engines)
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Display the services running on the Director group for the QFabric system. These services can include external Routing Engines that are used to support QFabric system operations, such as partitioning and fabric control.

Options

none—Display all services running on the Director group, which are used to support the QFabric system.

brief | detail—(Optional) Display the specified level of output.

fabric-managers—Display information for the fabric manager Routing Engine running on the Director group, which is used to support all partitions in the QFabric system.

fabric-controls—Display information for the fabric control Routing Engine running on the Director group, which is used to support route information in the QFabric system.

diagnostic-routing-engines—Display information for the diagnostic Routing Engine running on the Director group, which is responsible for troubleshooting and diagnostic utilities within a QFabric system partition.

Required Privilege Level

admin

RELATED DOCUMENTATION

[show fabric administration inventory](#) | [780](#)

List of Sample Output

[show fabric administration inventory infrastructure on page 795](#)

[show fabric administration inventory infrastructure fabric-controls on page 795](#)

[show fabric administration inventory infrastructure fabric-managers on page 796](#)

[show fabric administration inventory infrastructure diagnostic-routing-engines on page 796](#)

Output Fields

Table 92 on page 794 lists the output fields for the **show fabric administration inventory infrastructure** command. Output fields are listed in the approximate order in which they appear.

Table 92: show fabric administration inventory infrastructure Output Fields

Field Name	Field Description
Routing Engine Type	Type of virtual Junos Routing Engine being viewed. Examples include the network Node group, fabric control, fabric manager, and diagnostic Routing Engines.
Hostname	Name of the QFabric system component.
PID	Process identifier for the component.
CPU-Use (%)	Percentage of CPU processing memory being used by the component.
Fabric control	<p>Name of the virtual Junos Routing Engines responsible for route selection within a QFabric system partition.</p> <p>With the fabric-controls option, the fabric control Routing Engine can be either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for this virtual device. This field also displays the identifier and configuration status for the fabric control Routing Engine.</p>
Fabric manager	<p>Name of the virtual Junos Routing Engine that manages the QFabric system.</p> <p>With the fabric-managers option, the fabric manager Routing Engine can be either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for this virtual device. This field also displays the identifier and configuration status for the fabric manager Routing Engine.</p>
Network Node group	Name of the virtual Junos Routing Engine instance that handles routing processes for a network Node group.
Diagnostic	<p>Name of the virtual Junos Routing Engine responsible for troubleshooting and diagnostic utilities within a QFabric system partition.</p> <p>With the diagnostic-routing-engines option, the diagnostic Routing Engine can be either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the diagnostic Routing Engine.</p>
Item	Type of QFabric system component being viewed.
Identifier	Hardware serial identifier of a QFabric system component.

Table 92: show fabric administration inventory infrastructure Output Fields (*continued*)

Field Name	Field Description
Connection	Status of a QFabric system component: either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the listed component.
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured , Failed (unsuccessful), Pending (in the process of being written or retried), or Unknown .

Sample Output

show fabric administration inventory infrastructure

```
user@qfabric> show fabric administration inventory infrastructure
```

```
dg0:
Routing Engine Type      Hostname                  PID      CPU-Use(%)
-----
Fabric manager           FM-0                     9832     1.0
Network Node group       QFabric_default_NW-NG-1_RE1 24633    4.2
Fabric control           QFabric_default_FC-1_RE0   25374    1.8
Diagnostic               QFabric_DRE               6789     1.3

dg1:
Routing Engine Type      Hostname                  PID      CPU-Use(%)
-----
Fabric manager           FM-1                     572      1.6
Network Node group       QFabric_default_NW-NG-0_RE0 19217    7.8
Fabric control           QFabric_default_FC-0_RE0   20071    1.9
```

show fabric administration inventory infrastructure fabric-controls

```
user@qfabric> show fabric administration inventory infrastructure fabric-controls fabric-controls
```

Item	Identifier	Connection	Configuration
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured

show fabric administration inventory infrastructure fabric-managers

```
user@qfabric> show fabric administration inventory infrastructure fabric-managers
```

Item	Identifier	Connection	Configuration
Fabric manager			
FM-0		Connected	Configured

show fabric administration inventory infrastructure diagnostic-routing-engines

```
user@qfabric> show fabric administration inventory infrastructure diagnostic-routing-engines
```

Item	Identifier	Connection	Configuration
Diagnostic routing engine			
DRE-0		Connected	Configured

show fabric administration inventory interconnect-devices

Syntax

```
show fabric administration inventory interconnect-devices interconnect-device-name
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Display the Interconnect devices that belong to a QFabric system.

Options

none—Display all Interconnect devices within a QFabric system.

interconnect-device-name—Display a specific Interconnect device within a QFabric system.

Required Privilege Level

admin

RELATED DOCUMENTATION

request chassis fabric fpc 603
show chassis fabric connectivity 749
show chassis fabric device 758
show fabric administration inventory 780
Understanding Interconnect Devices 29

List of Sample Output

[show fabric administration inventory interconnect-devices on page 798](#)

[show fabric administration inventory interconnect-devices device-name on page 798](#)

Output Fields

[Table 93 on page 797](#) lists the output fields for the **show fabric administration inventory interconnect-devices** command. Output fields are listed in the approximate order in which they appear.

Table 93: show fabric administration inventory interconnect-devices Output Fields

Field Name	Field Description
Interconnect device	Name of the Interconnect devices associated with the partition.

Table 93: show fabric administration inventory interconnect-devices Output Fields (continued)

Field Name	Field Description
Item	Type of QFabric system component being viewed. Interconnect devices either display the alias name (if configured) or the hardware serial identifier and Control Board Routing Engine numbers.
Identifier	Hardware serial identifier of a QFabric system component. When you configure an alias name for a component, the ID is displayed.
Connection	Status of a QFabric system component: either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the listed component.
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured , Failed (unsuccessful), Pending (in the process of being written or retried), or Unknown .

Sample Output

show fabric administration inventory interconnect-devices

```
user@qfabric> show fabric administration inventory interconnect-devices
```

Item	Identifier	Connection	Configuration
Interconnect device			
IC-YW3781		Connected	Configured
YW3781/RE0		Connected	
YW3781/RE1		Connected	
IC-YW3798		Connected	Configured
YW3798/RE0		Connected	
YW3798/RE1		Connected	

show fabric administration inventory interconnect-devices device-name

```
user@qfabric> show fabric administration inventory interconnect-devices IC-YW3781
```

Item	Identifier	Connection	Configuration
Interconnect device			
IC-YW3781		Connected	Configured
YW3781/RE0		Connected	

YW3781/RE1

Connected

show fabric administration inventory node-devices

Syntax


```
show fabric administration inventory node-devices node-device-name
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Display the Node devices that belong to the QFabric system.

**NOTE:** If your Node devices do not appear in the output of the **show fabric administration inventory node-devices** command, check the cabling of your system.

Options

none—Display all Node devices within the QFabric system.

node-device-name—Display a specific Node device within the QFabric system.

Required Privilege Level

admin

RELATED DOCUMENTATION

Configuring Aliases for the QFabric System	388
Configuring Node Groups for the QFabric System	402
show fabric administration inventory node-groups	802
show fabric administration inventory	780
Understanding Node Devices	33

List of Sample Output

- [show fabric administration inventory node-devices on page 801](#)
- [show fabric administration inventory node-devices device-name on page 801](#)

Output Fields

[Table 94 on page 801](#) lists the output fields for the **show fabric administration inventory node-devices** command. Output fields are listed in the approximate order in which they appear.

Table 94: show fabric administration inventory node-devices Output Fields

Field Name	Field Description
Item	Type of QFabric system component being viewed.
Identifier	Hardware serial identifier of a QFabric system component.
Connection	Status of a QFabric system component: either Connected or Not Connected , depending on whether or not the Director software has detected keepalive messages for the listed component.
Node device	Name of the Node devices associated with the Node group. The device can be either Connected or Not Connected , depending on whether or not the Director software has detected keepalive messages for the device. This field also displays the serial ID and configuration status for the Node device.

Sample Output

show fabric administration inventory node-devices

```
user@qfabric> show fabric administration inventory node-devices
```

Item	Identifier	Connection
Node device		
node1	P3749-C	Connected
node2	P3767-C	Connected
node3	P3850-C	Connected
node4	P3947-C	Connected

show fabric administration inventory node-devices device-name

```
user@qfabric> show fabric administration inventory node-devices node0
```

Item	Identifier	Connection
Node device		
node1	P3749-C	Connected

show fabric administration inventory node-groups

Syntax

```
show fabric administration inventory node-groups node-group-name
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Display the Node groups and the corresponding Node devices that belong to the QFabric system.

Options

none—Display all Node groups within the QFabric system.

node-group-name—Display information for a specific Node group within the QFabric system.

Required Privilege Level

admin

RELATED DOCUMENTATION

- [Configuring Node Groups for the QFabric System | 402](#)
- [show fabric administration inventory node-devices | 800](#)
- [show fabric administration inventory | 780](#)

List of Sample Output

- [show fabric administration inventory node-groups on page 803](#)
- [show fabric administration inventory node-groups node-group-name on page 803](#)

Output Fields

[Table 95 on page 802](#) lists the output fields for the **show fabric administration inventory node-groups** command. Output fields are listed in the approximate order in which they appear.

Table 95: show fabric administration inventory node-groups Output Fields

Field Name	Field Description
Node group	Name of the Node groups associated with the partition.

Table 95: show fabric administration inventory node-groups Output Fields (*continued*)

Field Name	Field Description
Item	Type of QFabric system component being viewed. <ul style="list-style-type: none"> Autogenerated Node groups display the hardware serial identifier for both the name of the Node group and the name of the included Node device. User-configured Node groups either display the alias name (if configured) or the hardware serial identifier for each Node device contained in the Node group.
Identifier	Hardware serial identifier of a QFabric system component. When you configure an alias name for a component, the ID is displayed.
Connection	Status of a QFabric system component: either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the listed component.
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured , Failed (unsuccessful), Pending (in the process of being written or retried), or Unknown .

Sample Output

show fabric administration inventory node-groups

```
user@qfabric> show fabric administration inventory node-groups
```

Item	Identifier	Connection	Configuration
Node group			
BBAK8891		Connected	Configured
BBAK8891		Connected	
BBAK8868		Connected	Configured
BBAK8868		Connected	
RSNG-1		Connected	Configured
Node-3	BBAK8276	Connected	
Node-4	BBAK8273	Connected	
NW-NG-0		Connected	Configured
Node-0	BBAK8309	Connected	
Node-1	BBAK8283	Connected	

show fabric administration inventory node-groups node-group-name

```
user@qfabric> show fabric administration inventory node-groups RSNG-1
```


Item	Identifier	Connection	Configuration
Node group			
RSNG-1		Connected	Configured
Node-3	BBAK8276	Connected	
Node-4	BBAK8273	Connected	

show fabric administration system mac-pool

Syntax

```
show fabric administration system mac-pool
```

Release Information

Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

(QFabric systems only) Display the MAC addresses that belong to a QFabric Director group.

Options

There are no options for this command.

Required Privilege Level

admin

RELATED DOCUMENTATION

- [request fabric administration system mac-pool add | 618](#)
- [request fabric administration system mac-pool delete | 619](#)

List of Sample Output

[show fabric administration system mac-pool on page 806](#)

Output Fields

[Table 96 on page 805](#) lists the output fields for the **show fabric administration system mac-pool** command. Output fields are listed in the approximate order in which they appear.

Table 96: show fabric administration system mac-pool Output Fields

Field Name	Field Description
MAC Block Base	Starting MAC address for the pool assigned to the QFabric system.
Total MACs	Total number of MAC addresses assigned to the QFabric system.
Available MACs	Number of available MAC addresses from the total.

Sample Output

show fabric administration system mac-pool

user@qfabric> show fabric administration system mac-pool

Mac Block Base	Total MACs	Available MACs
00:11:00:00:00:00	4096	4084
02:00:00:11:22:00	10	10

show fabric inventory

Syntax

```
show fabric inventory
<brief | detail | terse>
<infrastructure fabric-controls <FC-0 | FC-1>>
<node-devices node-device-name>
<node-groups node-group-name>
```

Release Information

Command introduced in Junos OS Release 12.2 for the QFX Series.

Description

(QFabric systems only) Display Node devices, Node groups, and fabric control Routing Engines that belong to the QFabric system. You can narrow the level of output by specifying a device type.

NOTE:

- If you have administrator privileges, issue the **show fabric administration inventory** command to view all devices in your QFabric system (including Interconnect devices and Director devices).
- If your Node devices do not appear in the output of the **show fabric inventory** command, check the cabling of your system.

Options

none—Display all devices within a QFabric system.

brief | detail | terse—(Optional) Display the specified level of output.

infrastructure fabric-controls <FC-0 | FC-1>—(Optional) Display information for all fabric control Routing Engines running on the Director group within the QFabric system, or the individual fabric control Routing Engine you specify (either FC-0 or FC-1).

node-devices *node-device-name*—(Optional) Display a specific Node device within a QFabric system.

node-groups *node-group-name*—(Optional) Display a specific Node group within a QFabric system.

Required Privilege Level

admin

RELATED DOCUMENTATION

[Performing the QFabric System Initial Setup on a QFX3100 Director Group | 363](#)

[Configuring Aliases for the QFabric System | 388](#)

[Configuring Node Groups for the QFabric System | 402](#)

[show fabric administration inventory | 780](#)

List of Sample Output

[show fabric inventory on page 809](#)

[show fabric inventory infrastructure fabric-controls on page 809](#)

[show fabric inventory node-devices on page 809](#)

[show fabric inventory node-groups on page 810](#)

Output Fields

[Table 97 on page 808](#) lists the output fields for the **show fabric inventory** command. Output fields are listed in the approximate order in which they appear.

Table 97: show fabric inventory Output Fields

Field Name	Field Description
Item	Type of QFabric system component being viewed. Possible values include Node device , Node group , Fabric control , and Ungrouped Node device .
Identifier	Hardware serial identifier of a QFabric system component.
Connection	Status of a QFabric system component: either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the listed component.
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured , Failed (unsuccessful), Pending (in the process of being written or retried), or Unknown .
Node group	Name of the Node groups associated with the QFabric system, and the Node devices assigned to each Node group. The group can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the devices in the group. This field also displays the serial ID for the Node group and the status for the Node group.
Node device	Name of the Node devices associated with the Node group. The device can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for the device. This field also displays the serial ID and configuration status for the Node device.

Table 97: show fabric inventory Output Fields (*continued*)

Field Name	Field Description
Fabric control	Name of the virtual Junos Routing Engines responsible for route selection within a QFabric system partition. The fabric control Routing Engine can be either Connected or Disconnected , depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the fabric control Routing Engine.

Sample Output

show fabric inventory

```
user@qfabric> show fabric inventory
```

Item	Identifier	Connection	Configuration
Ungrouped Node device			
P3747-C		Disconnected	
Node group			
NW-NG-0		Connected	Configured
Node-1	P4093-C	Connected	
RSNG-1		Connected	Configured
Node-2	P4514-C	Connected	
Node-3	P3917-C	Connected	
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured

show fabric inventory infrastructure fabric-controls

```
user@qfabric> show fabric inventory infrastructure fabric-controls
```

Item	Identifier	Connection	Configuration
Fabric control			
FC-0		Connected	Configured
FC-1		Connected	Configured

show fabric inventory node-devices

```
user@qfabric> show fabric inventory node-devices
```


Item	Identifier	Connection	Configuration
Node device			
Node-1	P4093-C	Connected	
Node-2	P4514-C	Connected	
Node-3	P3917-C	Connected	

show fabric inventory node-groups

user@qfabric> **show fabric inventory node-groups**

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
Node-1	P4093-C	Connected	
RSNG-1		Connected	Configured
Node-2	P4514-C	Connected	
Node-3	P3917-C	Connected	

show fabric session-host

Syntax

```
show fabric session-host
```

Release Information

Command introduced in Junos OS Release 12.2 for the QFX Series.

Description

(QFabric systems only) Display the Director device within the Director group that hosts the QFabric CLI session.

Options

none—Display the Director device hosting the QFabric CLI session.

Required Privilege Level

admin

RELATED DOCUMENTATION

Understanding the Director Group 26
show fabric administration inventory director-group status 786

List of Sample Output

[show fabric session-host on page 812](#)

Output Fields

[Table 98 on page 811](#) lists the output fields for the **show fabric session-host** command. Output fields are listed in the approximate order in which they appear.

Table 98: show fabric session-host Output Fields

Field Name	Field Description
Identifier	Hardware serial identifier of the Director device that hosts the SSH QFabric CLI session.

Sample Output

```
show fabric session-host
```

```
user@qfabric> show fabric session-host
```

```
Identifier: 0281052011000032
```


show log

List of Syntax

[Syntax on page 813](#)

[Syntax \(QFX Series and OCX Series\) on page 813](#)

[Syntax \(TX Matrix Router\) on page 813](#)

Syntax

```
show log
<filename | user <username>>
```

Syntax (QFX Series and OCX Series)

```
show log filename
<device-type (device-id | device-alias)>
```

Syntax (TX Matrix Router)

```
show log
<all-lcc | lcc number | scc>
<filename | user <username>>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Option *device-type (device-id | device-alias)* is introduced in Junos OS Release 13.1 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

List log files, display log file contents, or display information about users who have logged in to the router or switch.

NOTE: On MX Series routers, modifying a configuration to replace a service interface with another service interface is treated as a catastrophic event. When you modify a configuration, the entire configuration associated with the service interface—including NAT pools, rules, and service sets—is deleted and then re-created for the newly specified service interface. If there are active sessions associated with the service interface that is being replaced, these sessions are deleted and the NAT pools are then released, which leads to the generation of the NAT_POOL_RELEASE system log messages. However, because NAT pools are already deleted as a result of the catastrophic configuration change and no longer exist, the NAT_POOL_RELEASE system log messages are not generated for the changed configuration.

Options

none—List all log files.

<all-lcc | lcc *number* | scc>—(Routing matrix only)(Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace ***number*** with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).

device-type—(QFabric system only) (Optional) Display log messages for only one of the following device types:

- **director-device**—Display logs for Director devices.
- **infrastructure-device**—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).
- **interconnect-device**—Display logs for Interconnect devices.
- **node-device**—Display logs for Node devices.

NOTE: If you specify the **device-type** optional parameter, you must also specify either the **device-id** or **device-alias** optional parameter.

(*device-id* | *device-alias*)—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.

NOTE: The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of **messages**.

user <username>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include **username**, display logging information about the specified user.

Required Privilege Level

trace

RELATED DOCUMENTATION

| [syslog \(System\)](#)

List of Sample Output

[show log on page 815](#)

[show log filename on page 816](#)

[show log filename \(QFabric System\) on page 818](#)

[show log user on page 819](#)

Sample Output

show log

user@host> **show log**

```
total 57518
-rw-r--r-- 1 root bin      211663 Oct  1 19:44 dcd
-rw-r--r-- 1 root bin      999947 Oct  1 19:41 dcd.0
-rw-r--r-- 1 root bin      999994 Oct  1 17:48 dcd.1
-rw-r--r-- 1 root bin      238815 Oct  1 19:44 rpd
-rw-r--r-- 1 root bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r-- 1 root bin      1061095 Oct  1 12:13 rpd.1
-rw-r--r-- 1 root bin      1052026 Oct  1 06:08 rpd.2
-rw-r--r-- 1 root bin      1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin      1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin      1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin      1056350 Sep 30 07:04 rpd.6
```



```
-rw-r--r--  1 root  bin      1048876 Sep 30 03:21 rpd.7
-rw-rw-r--  1 root  bin      19656 Oct  1 19:37 wttmp
```

show log filename

user@host> show log rpd

```
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr 192.0.2.21
nhop type local nhop 192.0.2.21
Oct  1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr 192.0.2.22
nhop type unicast nhop 192.0.2.22
Oct  1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...
```

user@host:LSYS1> show log flow_lsys1.log

```
Nov  7 07:34:09 07:34:09.491800:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table
lock

Nov  7 07:34:09 07:34:09.491809:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route
table lock

Nov  7 07:34:09 07:34:09.491840:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table
lock

Nov  7 07:34:09 07:34:09.491841:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route
table lock

Nov  7 07:34:09 07:34:09.491854:CID-0:THREAD_ID-00:LSYS_ID-01:RT:cache final sw_nh
0x0

Nov  7 07:34:09 07:34:09.491868:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table
lock

Nov  7 07:34:09 07:34:09.491869:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route
```



```
table lock
```

```
Nov  7 07:34:09 07:34:09.491881:CID-0:THREAD_ID-00:LSYS_ID-01:RT:cache final sw_nh
0x0
```

```
user@host:TSYS1> show log flow_tsys1.log
```

```
Nov  7 13:21:47
13:21:47.217744:CID-0:THREAD_ID-05:LSYS_ID-32:RT:<192.0.2.0/0->198.51.100.0/9011;1,0x0>
:

Nov  7 13:21:47 13:21:47.217747:CID-0:THREAD_ID-05:LSYS_ID-32:RT:packet [84] ipid
= 39281, @0x7f490ae56d52

Nov  7 13:21:47 13:21:47.217749:CID-0:THREAD_ID-05:LSYS_ID-32:RT:----
flow_process_pkt: (thd 5): flow_ctxt type 0, common flag 0x0, mbuf 0x4882b600,
rtbl7

Nov  7 13:21:47 13:21:47.217752:CID-0:THREAD_ID-05:LSYS_ID-32:RT: flow process pak
fast ifl 88 in_ifp lt-0/0/0.101

Nov  7 13:21:47 13:21:47.217753:CID-0:THREAD_ID-05:LSYS_ID-32:RT:
lt-0/0/0.101:192.0.2.0->198.51.100.0, icmp, (0/0)

Nov  7 13:21:47 13:21:47.217756:CID-0:THREAD_ID-05:LSYS_ID-32:RT: find flow: table
0x11d0a2680, hash 20069(0xffff), sa 192.0.2.0, da 198.51.100.0, sp 0, d0

Nov  7 13:21:47 13:21:47.217760:CID-0:THREAD_ID-05:LSYS_ID-32:RT:Found: session
id 0x12. sess tok 28685

Nov  7 13:21:47 13:21:47.217761:CID-0:THREAD_ID-05:LSYS_ID-32:RT: flow got session.

Nov  7 13:21:47 13:21:47.217761:CID-0:THREAD_ID-05:LSYS_ID-32:RT: flow session
id 18

Nov  7 13:21:47 13:21:47.217763:CID-0:THREAD_ID-05:LSYS_ID-32:RT: vector bits 0x200
vector 0x84ae85f0

Nov  7 13:21:47 13:21:47.217764:CID-0:THREAD_ID-05:LSYS_ID-32:RT:set nat
0x11e463550(18) timeout const to 2

Nov  7 13:21:47 13:21:47.217765:CID-0:THREAD_ID-05:LSYS_ID-32:RT: set_nat_timeout
2 on session 18
```



```

Nov  7 13:21:47 13:21:47.217765:CID-0:THREAD_ID-05:LSYS_ID-32:RT:refresh nat
0x11e463550(18) timeout to 2

Nov  7 13:21:47 13:21:47.217767:CID-0:THREAD_ID-05:LSYS_ID-32:RT:insert usp tag
for apps

Nov  7 13:21:47 13:21:47.217768:CID-0:THREAD_ID-05:LSYS_ID-32:RT:mbuf 0x4882b600,
exit nh 0xffffb0006

```

show log filename (QFabric System)

user@qfabric> show log messages

```

Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486 file:
UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486 file:
UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)

```



```
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492 file:
  UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492 file:
  UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)
```

show log user

user@host> **show log user**

usera	mg2546		Thu Oct 1 19:37	still logged in
usera	mg2529		Thu Oct 1 19:08 - 19:36	(00:28)
usera	mg2518		Thu Oct 1 18:53 - 18:58	(00:04)
root	mg1575		Wed Sep 30 18:39 - 18:41	(00:02)
root	ttyp2	aaa.bbbb.com	Wed Sep 30 18:39 - 18:41	(00:02)
userb	ttyp1	192.0.2.0	Wed Sep 30 01:03 - 01:22	(00:19)

show oam fabric flow-specification

Syntax

```
show oam fabric flow-specification
<flow-specification-name>
```

Release Information

Command introduced in Junos OS Release 12.2 for the QFX Series.

Description

Display the fabric flow specifications that are configured in a QFabric system Node group.

Options

flow-specification-name—(Optional) Name of a flow specification.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring 439
Configuring a Fabric Maintenance Association 441
Configuring Flow Specifications 442
fabric (OAM) 470
ping fabric multicast-flow 591
ping fabric unicast-flow 593
traceroute fabric unicast-flow 828

List of Sample Output

[show oam fabric flow-specification \(All\) on page 822](#)

[show oam fabric flow-specification \(Flow Specification Specified\) on page 822](#)

Output Fields

[Table 99 on page 821](#) lists the output fields for the **show oam fabric flow-specification** command. Output fields are listed in the approximate order in which they appear.

Table 99: show oam fabric flow-specification Output Fields

Field Name	Field Description
Flow specification name	Name of the flow specification.
Type	Flow specification type: <ul style="list-style-type: none"> • Ethernet • Ethernet Unicast IPV4 • Multicast IPV4 • Multicast VLAN flood
Ethernet frame size	Ethernet frame size in bytes.
Ether type	EtherType protocol.
Source-MAC	MAC address of the source FMEP interface.
Source-MAC Mask	MAC address mask of the source FMEP interface.
Destination-MAC	MAC address of the destination FMEP interface.
Destination-MAC Mask	MAC address mask of the destination FMEP interface.
Source-IP	IP address of the source FMEP interface.
Source-IP Mask	IP address mask of the source FMEP interface.
Destination-IP	IP address of the destination FMEP interface.
Destination-IP mask	IP address mask of the destination FMEP interface.
IP-Protocol	IPv4 protocol configured.
Source-L4-Port	Source TCP or UDP port.
Destination-L4-Port	Destination TCP or UDP port.
Destination multicast group IP	IP address of the destination multicast group.

Sample Output

show oam fabric flow-specification (All)

user@host> show oam fabric flow-specification

```
Flow specification name : fspec1 Type : Ethernet
  Ethernet frame size : Unspecified
  Ether type : Unspecified
  Source-MAC : Unspecified
  Source-MAC Mask : Unspecified
  Destination-MAC : Unspecified
  Destination-MAC Mask : Unspecified
Flow specification name : fspec2 Type : Ethernet
  Ethernet frame size : Unspecified
  Ether type : 1792
  Source-MAC : 0:0:a0:f0:cc:22
  Source-MAC Mask : Unspecified
  Destination-MAC : 0:11:22:33:a0:e2
  Destination-MAC Mask : Unspecified
Flow specification name : fspec3 Type : Ethernet Unicast IPV4
  Ethernet frame size : Unspecified
  Source-IP : 121.0.0.1
  Source-IP Mask : Unspecified
  Destination-IP : 132.1.0.1
  Destination-IP Mask : Unspecified
  Destination-mac : Unspecified
  IP-protocol : TCP
  Source-L4-port : 270
  Destination-L4-port : Unspecified
Flow specification name : fspec4 Type : Multicast IPV4
  Ethernet frame size : Unspecified
  Source IP : 100.0.0.23
  Destination multicast group IP: 225.0.0.1
Flow specification name : fspec5 Type : Multicast VLAN flood
  Ethernet frame size : Unspecified
```

show oam fabric flow-specification (Flow Specification Specified)

user@host> show oam fabric flow-specification fspec3

```
Flow specification name : fspec3 Type : Ethernet Unicast IPV4
```



```
Ethernet frame size : Unspecified
Source-IP : 121.0.0.1
Source-IP Mask : Unspecified
Destination-IP : 132.1.0.1
Destination-IP Mask : Unspecified
Destination-mac : Unspecified
IP-protocol : TCP
Source-L4-port : 270
Destination-L4-port : Unspecified
```


show oam fabric interfaces

Syntax

```
show oam fabric interfaces
<interface-name>
<fabric-maintenance-association fma-name>
<brief>
```

Release Information

Command introduced in Junos OS Release 12.2 for the QFX Series.

Description

Display the fabric maintenance associations (FMAs) and fabric maintenance endpoints (FMEPs) that are configured in a QFabric system Node group.

If you do not specify an interface, display information for an FMA (if one is specified) or all FMAs configured in the Node group.

Options

brief—(Optional) Summarized version of the output.

fabric-maintenance-association *fma-name*—(Optional) Name of a specific FMA.

interface-name—(Optional) Name of a specific interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Overview of Internal Fabric Monitoring | 439](#)

[Configuring a Fabric Maintenance Association | 441](#)

[Configuring Flow Specifications | 442](#)

[fabric \(OAM\) | 470](#)

[ping fabric multicast-flow | 591](#)

[ping fabric unicast-flow | 593](#)

[traceroute fabric unicast-flow | 828](#)

List of Sample Output

[show oam fabric interfaces on page 825](#)

Output Fields

Table 100 on page 825 lists the output fields for the **show oam fabric interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 100: show oam fabric interfaces Output Fields

Field Name	Field Description
Interface-name	Name of the interface.
Fabric-Maintenance Association	FMA name.
VLAN	VLAN name.
Interface state	State of the interface: Up or down.
MEP Identifier	FMEP identifier.
MEP Name	FMEP name.

Sample Output

show oam fabric interfaces

```
user@host> show oam fabric interfaces
```

Interface-name	Fabric-Maintenance Association	VLAN	Interface state	MEP Identifier	MEP Name
ED1479:NULL	fma-default	*	up	32774	
fmeop-default-ED1479					
ED1494:NULL	fma-default	*	up	32773	
fmeop-default-ED1494					
ED1494:xe-0/0/7.0	fma1	v100	up	1	
ED1497:NULL	fma-default	*	up	32775	
fmeop-default-ED1497					
ED1497:xe-0/0/8.0	fma1	v100	up	2	
P3613-C:NULL	fma-default	*	up	32777	
fmeop-default-P3613-C					

show system software upgrade status

Syntax

```
show system software upgrade status
```

Release Information

Command introduced in Junos OS Release 13.1 for the QFX Series.

Description

(QFabric systems only) Display the status of a software upgrade, including details for both nonstop software upgrades and *component all* style upgrades.

Required Privilege Level

view

RELATED DOCUMENTATION

- [Performing a Nonstop Software Upgrade on the QFabric System | 533](#)
- [Verifying Nonstop Software Upgrade for QFabric Systems | 540](#)
- [Upgrading Software on a QFabric System | 566](#)
- [request system software nonstop-upgrade | 639](#)

List of Sample Output

- [show system software upgrade status \(Nonstop Software Upgrade\) on page 827](#)
- [show system software upgrade status \(Component All Upgrade\) on page 827](#)

Output Fields

Table 101 on page 826 lists the output fields for the **show system software upgrade status** command. Output fields are listed in the approximate order in which they appear.

Table 101: show system software upgrade status Output Fields

Field Name	Field Description
Timestamp	Displays the day of the week, month, date, hour, minute, second, and year when you issue the show system software upgrade status command. An example of the timestamp format is as follows: Wed Jan 16 22:06:02 2013.

Table 101: show system software upgrade status Output Fields (*continued*)

Field Name	Field Description
Software nonstop upgrade on:	<p>Status of the upgrade:</p> <ul style="list-style-type: none"> • FM-0 in progress—A Director group nonstop software upgrade is in process for the fabric manager Routing Engine. • NW-NG-0 in progress—A Node group nonstop software upgrade is in process for the network Node group. • RSNG in progress—A Node group nonstop software upgrade is in process for the redundant server Node group. • all in progress—A <i>component all</i> style upgrade is in process for the entire QFabric system.

Sample Output

show system software upgrade status (Nonstop Software Upgrade)

```
user@qfabric> show system software upgrade status
```

```
Wed Jan 16 22:06:02 2013 Software nonstop upgrade on:
      NW-NG-0 in progress
      RSNG in progress
```

show system software upgrade status (Component All Upgrade)

```
user@qfabric> show system software upgrade status
```

```
Wed Jan 16 22:37:48 2013 Software component upgrade on:
      all in progress
```


traceroute fabric unicast-flow

Syntax

```
traceroute fabric unicast-flow flow-spec-name flow-specification-name fma fma-name
<source-fmep fmep-name>
<dest-fmep fmep-name>
<forced-fte-interface fte-interface>
<count count>
<verbose>
```

Release Information

Command introduced in Junos OS Release 12.2 for the QFX Series.

Description

Trace the path taken by a specific unicast flow (ping) operation across a VLAN on the QFabric system from a source maintenance endpoint (FMEP) to a destination FMEP. The source and destination FMEPs may be on the same Node device, different Node devices connected to the same Interconnect device, or different Node devices connected to different Interconnect devices.

The flow path is the sequence of Packet Forwarding Engine forwarding hops along which the protocol data unit (PDU) travels. The hop-by-hop sequence and number of hops are reported in terms of fabric maintenance intermediate points (FMIPs), which are interfaces on the Packet Forwarding Engine of the Interconnect device. An FMIP sends a response to the source FMEP when the traceroute PDU is received.

Options

dest-fmep-id *source-fmep-id*—(Optional) ID of the destination FMEP of the traceroute operation.

flow-spec-name *flow-specification-name*—Name of the flow specification that defines the parameters of the traceroute operation.

fma-name *fma-name*—Name of the FMA that associates a VLAN with the FMEPs.

forced-fte-interface *fte-interface*—(Optional) Option to force the fabric ping operation to use the specified FTE interface to inject the PDU into the fabric instead of using the internal forwarding path lookup table to determine the FTE interface.

source-fmep-id *source-fmep-id*—(Optional) ID of the FMEP that is the source of the traceroute operation.

verbose—(Optional) Detailed version of the output display.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- [Overview of Internal Fabric Monitoring | 439](#)
- [Configuring a Fabric Maintenance Association | 441](#)
- [Configuring Flow Specifications | 442](#)
- [fabric \(OAM\) | 470](#)
- [ping fabric multicast-flow | 591](#)
- [ping fabric unicast-flow | 593](#)

List of Sample Output

[traceroute fabric unicast-flow on page 830](#)

[traceroute fabric unicast-flow \(Verbose\) on page 830](#)

Output Fields

[Table 102 on page 829](#) lists the output fields for the **traceroute fabric-unicast-flow** command. Output fields are listed in the approximate order in which they appear.

Table 102: traceroute fabric unicast-flow Output Fields

Field Name	Field Description
Fabric flow traceroute between source <i>source-fmep-name</i> destination <i>destination-fmep-name</i>	Source and destination FMEP names.
Using fabric flow-specification	Name of the flow specification.
Ethernet frame-size	Ethernet frame size.
Source-MAC	MAC address of the source FMEP interface.
Destination-MAC	MAC address of the destination FMEP interface.
Ethertype	EtherType protocol.
Received response from fmip-id, interface	FMIP identifier and interface name of the FMIP that received the PDU and responded to the source FMEP.
Received total <i>number</i> responses	Number of responses that are received after the traceroute operation has finished.

Sample Output

traceroute fabric unicast-flow

```
user@host> traceroute fabric unicast-flow fma-name fma1 source-fmep-id 1 dest-fmep-id 2
flow-spec-name fspec1
```

```
Received response from fmip-id 128.40.0.8, interface A0003:0/1/2
Received response from fmip-id 128.128.0.8, interface A0003:0/0/8
Received response from fmep-id 2
Received total 3 responses
```

Sample Output

traceroute fabric unicast-flow (Verbose)

```
user@host> traceroute fabric unicast-flow fma-name fma1 source-fmep-id 1 dest-fmep-id 2
flow-spec-name fspec1 verbose
```

```
Fabric flow traceroute between source ED1494 destination ED1497
Using fabric flow-specification: fspec1
Ethernet frame-size: 256
Source-MAC: 0:EF:7:99:85:10
Destination-MAC: 0:D3:A6:11:39:21
Ethertype: 11564
Received response from fmip-id 128.40.0.8, interface A0003:0/1/2
Received response from fmip-id 128.128.0.8, interface A0003:0/0/8
Received response from fmep-id 2
Received total 3 responses
```


5

PART

Troubleshooting

QFabric System Troubleshooting | **833**

QFabric System Troubleshooting

IN THIS CHAPTER

- [Performing System Backup and Recovery for a QFabric System | 833](#)
- [Performing a QFabric System Recovery Installation on the Director Group | 835](#)
- [Performing a Recovery Installation | 845](#)
- [Creating an Emergency Boot Device for QFX Series Switches | 847](#)

Performing System Backup and Recovery for a QFabric System

Many routers and switches require an administrator to recover the software package and the configuration file for the device separately. In the case of a device failure, this means the administrator might need to perform two separate tasks (if neither the software package nor the configuration file can be recovered).

In contrast, the QFabric system uses a unique mechanism that saves the backup and recovery files for both the Junos OS software and the system configuration into a single collection. The following QFabric system backup and recovery mechanism simplifies and streamlines the recovery process so you can return to normal operations as quickly as possible.

To backup and recover your QFabric system:

1. (First time only) Implement the following one-time procedure to prepare your QFabric system to use the system backup and recovery feature:
 - Insert a Juniper Networks software installation USB flash drive into the master Director device. (This drive was provided to you as one of the components of your QFabric system shipment.)
 - Issue the **request system software format-qfabric-backup** command. The contents and format of the USB flash drive are copied to the Director group shared directory and are used as the basis for all future backup and recovery operations.

```
user@qfabric> request system software format-qfabric-backup
```

```
Copying QFabric USB template image from /dev/sdb(Unigen,PQS4000,4009 MB).....
```


- Remove the Juniper Networks software installation USB drive from the master Director device.
2. Issue the **request system software system-backup** command to backup the software package and configuration file. This command saves the current files necessary to recover the QFabric system. The files are saved to a shared memory directory in the Director group.

NOTE: As you upgrade your system with new software and change the system configuration over time, remember to reissue this command periodically to save the newest files for recovery purposes.

```
user@qfabric> request system software system-backup
```

```
user@qfabric>
```

3. Insert a 4 GB or larger USB flash drive into the master Director device for your Director group, and issue the **request system software system-backup usb-create** command. This command copies the recovery files that have been backed up in the Director group and transfers them to the USB flash drive to create a recovery USB drive.

NOTE: Issuing this command overwrites the contents of the USB flash drive with the QFabric system recovery files.

```
user@qfabric> request system software system-backup usb-create /dev/sdb
```

```
Issuing this command will overwrite the contents of the USB drive.
Continue? [yes,no] (no) yes
```

```
This operation will access the USB drive on 0281042010000013.
Are you sure you want to continue? [yes,no] (no) yes
```

```
Copying QFabric recovery media to /dev/sdb...
Successfully copied QFabric recovery media to /dev/sdb
```

4. Remove the recovery USB drive from the Director device, and store it securely in a known location that you will remember when you need to use the recovery USB drive.
5. If the QFabric system fails, power off the Director group, insert the recovery USB drive into the master Director device of your Director group, turn on power to the Director device, and follow the prompts

to recover your system. This step restores the software package and the configuration file for your QFabric system.

RELATED DOCUMENTATION

[request system software format-qfabric-backup | 638](#)

[request system software system-backup | 647](#)

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 | 836](#)
- [Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software | 838](#)

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 *Switchingbox*, 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 575](#) 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 | 363](#)

[Upgrading Software on a QFabric System | 566](#)

request system software configuration-backup

request system software configuration-restore

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” on page 847](#) 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 not 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](#) | 847

Creating an Emergency Boot Device for QFX Series Switches

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.

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

NOTE: You can create the emergency boot device on another Juniper Networks switch or router, or any PC or laptop 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. Use **gunzip** to unzip the image file.
3. Insert a USB device into the USB port.
4. From the Junos OS command-line interface (CLI), start the shell:

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

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 root, 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 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>
```