

DAY ONE: CONFIGURING EX SERIES ETHERNET SWITCHES

**3rd
Edition**

You need to configure your EX Series Ethernet switch and you need to get it done today. This practical, best-selling book, now in its third edition, shows you what to do and exactly how to do it.

By Yong Kim

DAY ONE: CONFIGURING EX SERIES ETHERNET SWITCHES, 3rd. Edition

The Juniper Networks EX Series Ethernet Switches deliver a high-performance, scalable solution for campus, branch office, and data center environments. You can deploy cost-effective Junos switching solutions that deliver carrier-class reliability, security risk management, network virtualization, application control, and reduced total cost of ownership. This book gives you both configuration background and key samples so you can get your switch up and optimally running in your network. No theory, no long introductions, just straightforward configurational how-to's.

"This *Day One* book does an excellent job of providing you with the necessary information to get the EX Switches in your environment up and running correctly without trying to reteach you the history or basics of Ethernet switching."

Brandon Bennett, Senior IT Engineer

JNCIE-ER #46, JNCIP-M, JNCIA-EX, CCIE R&S #19406

IT'S DAY ONE AND YOU HAVE A JOB TO DO, SO LEARN HOW TO:

- Manage an EX Series switch using the Junos command line interface (CLI).
- Set key Virtual Chassis configurations using various interconnection methods, as well as important design considerations for your Virtual Chassis configuration.
- Configure Link Aggregation Group (LAG).
- Configure Layer 2 Switching and Layer 3 Routing.
- Configure basic IP connectivity and elements to enable remote access.
- Configure basic static routing.
- Set various Ethernet-switching options such as voice VLAN, Layer 2 security (DHCP snooping, Dynamic ARP Inspection, etc.), or other Layer 2-specific features.
- Configure key EX Series switch features such as Ethernet OAM, MVRP, Multicast, EZQOS-Voice, and Port mirroring.

Juniper Networks Books are singularly focused on network productivity and efficiency. Peruse the complete library at www.juniper.net/books.

Published by Juniper Networks Books

ISBN 978-1-936779-14-7



JUNIPER
NETWORKS

Day One: Configuring EX Series Ethernet Switches, 3rd Edition

By Yong Kim

<i>Chapter 1: EX Series Overview</i>	<i>7</i>
<i>Chapter 2: Virtual Chassis Physical Connections</i>	<i>15</i>
<i>Chapter 3: Network Topology (Logical Topology)</i>	<i>37</i>
<i>Chapter 4: Ethernet Switching</i>	<i>55</i>
<i>Chapter 5: EX Series Features</i>	<i>69</i>

© 2015 by Juniper Networks, Inc. All rights reserved. Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, 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..

Published by Juniper Networks Books
Writers: Yong Kim
Editor in Chief: Patrick Ames
Copyediting and Proofing: Nancy Koerbel
Third Edition Technical Reviewer: Steve Puluka

ISBN: 978-1-936779-14-7 (print)
Printed in the USA by Vervante Corporation.
ISBN: 978-1-936779-15-4 (ebook)

Version History: v5 (Third Edition) March 2015
5 6 7 8 9 10 #7100127

About the Author

Yong Kim was a Senior Technical Marketing Engineer for Juniper Networks in the Fabric and Switching Technology Business Group. Yong has over a decade of experience in network and security solution design, implementation, and troubleshooting. Prior to joining Juniper Networks he worked at Cisco Systems in a variety of roles including the Proof of Concept (POC) lab and TAC.

This book is available in a variety of formats at: www.juniper.net/dayone.

What You Need to Know Before Reading this Book

Before reading this booklet you should have a basic understanding of the Junos operating system. Specifically, you should be able to change configurations, and to navigate through the command line hierarchy. You should reference other *Day One* book in the *Junos Fundamentals Series* (<http://www.juniper.net/dayone>), any of the excellent books in the *Juniper Networks Technical Library* (<http://www.juniper.net/books>), and any material about Junos and its operation at www.juniper.net, to help you acquire this background.

Other knowledge that you'll need as you read this book:

- Understanding of TCP/IP.
- Knowing basic switching concepts including bridging and Spanning Tree Protocol(s).
- Familiarity with interface naming in devices running the Junos operating system.
- Although it is not mandatory to complete the reading of this book, access to EX Series devices can help you practice configuring the various scenarios covered in the following pages, increasing the speed of implementing the EX Series devices in your network.

After Reading this Book, You'll Be Able To

- Manage an EX Series switch using the Junos command-line interface (CLI).
- Set key Virtual Chassis configurations using various interconnection methods, as well as important design considerations for your Virtual Chassis configuration.
- Configure Link Aggregation Group (LAG).
- Configure Layer 2 Switching and Layer 3 Routing.
- Configure basic IP connectivity and elements to enable remote access.
- Configure basic static routing.
- Set various Ethernet-switching options such as voice VLAN, Layer 2 security (DHCP snooping, Dynamic ARP Inspection, etc.), or other Layer 2-specific features.
- Configure key EX Series switch features such as Ethernet OAM, MVRP, Multicast, EZQOS-Voice, and Port mirroring.

The EX Series Ethernet Switches

The *EX Series Ethernet Switches* is a mouthful to pronounce. And the Junos device comes in several different platforms designed for a variety of networking usage. There are many types of EX Series Ethernet switches for a variety of deployment scenarios from small branch office to data center core.

This book simplifies terminology by using the term *EX*, or *the EX*.

NOTE Some features of the EX Series Ethernet Switches are configured differently on different platforms and this book attempts to point that out.

Chapter 1

EX Series Overview

Exploring the EX4200 Ethernet Switch 8

Managing an EX Series Ethernet Switch 11

The Juniper Networks EX Series Ethernet Switches deliver a high-performance, scalable solution for campus, branch office, and data center environments. With the EX Series switches, you can deploy cost effective Junos switching solutions that deliver carrier-class reliability, security risk management, network virtualization, application control, and reduced total cost of ownership.

If you have administered or operated other Ethernet switches, the Juniper Networks EX Series Ethernet Switches should appear familiar to you. However, if this is your first time setting up an Ethernet switch, this booklet guides you through the process.

The EX Series consists of several switch product families:

- the *entry-level* EX2200, and EX2200-C line of Ethernet switches;
- the EX3200, EX3300, EX4200, and EX4300 line of fixed-configuration Ethernet switches;
- the EX4500, EX4550, and the EX4600 10GbE top-of-rack/Aggregation Ethernet switches;
- the EX6200 and EX8200 *modular* switches;
- and, the chassis-based EX9200 programmable switch.

The EX2200, EX2200-C, EX3300, EX4200, EX4300, EX4500, EX4550, EX4600, EX8200, and EX9200 switches feature Juniper's Virtual Chassis technology (more about that in Chapter 2). This book focuses on the steps for configuring an EX4200 switch.

MORE? For more information about each specific line of EX Series switch, see the product literature at <http://www.juniper.net/us/en/products-services/switching/ex-series/>.

Exploring the EX4200 Ethernet Switch

When configuring an Ethernet switch the first step is becoming familiar with the physical layout of the device. The rear panel of the EX4200 switch (see Figure 1.1) includes a number of ports.

- The Console port: The switch can be configured via a rear-panel RS-232 serial interface that uses an RJ-45 connector. A computer can be directly attached to the switch console port and configured using a terminal-emulation program. If consoled this way the terminal emulation software should be configured with the

following parameters: 9600 baud rate; 8 data bits; No Parity: 1 stop bit; and, No Flow Control.

- The Management port: A dedicated rear-panel Ethernet RJ-45 port, located to the left of the console port, is available for performing out-of-band switch management. The port uses an auto-sensing RJ-45 connector to support a 10/100/1000 BASE-T connection. Two LEDs located next to the port indicate link activity and port status. The management port requires an IP address and a subnet mask to be configured for switch management and administration.
- USB port: Storage devices such as flash drives can be connected directly to the EX4200 switch via a rear-panel USB port. USB flash drives can be used to store and upload configuration files or Junos software releases.
- Virtual Chassis port (VCP): The dual rear-panel VCPs enable EX4200 switches to be interconnected over a dedicated 128 gigabit-per-second (Gbps) high-speed virtual backplane. Switches deployed in close proximity, such as in wiring closets, or in top-of-rack data center applications, can be easily connected using a Virtual Chassis cable, which is covered in Chapter 2.

NOTE The VCP uses a specific Virtual Chassis cable (that is included) to interconnect EX4200 Ethernet switches. For more information, see the *Connecting a Virtual Chassis Cable to an EX4200 Switch Guide* at <http://www.juniper.net/techpubs>.




Figure 1.1

Rear Panel of EX4200 Ethernet Switch

The front panel of the EX4200 switch (see Figure 1.2) includes an LCD panel, an optional uplink module bay, and up to 48 host network ports.

- **LCD panel:** The backlit LCD panel displays various types of information about the switch, including key stages of the boot process, the host name of the switch, the switch's role in a Virtual Chassis configuration in an abbreviated form, member ID in a Virtual Chassis, and current operations such as initial switch setup and reboot.
- **LCD buttons and status LEDs:** Located next to the LCD panel, the LEDs and buttons allow you to quickly determine switch status and perform basic operations. The top button, labeled Menu, enables you to cycle through various LCD panel menus. The bottom button, labeled Enter, allows you to confirm the selection. The Enter button also works as confirmation when used in the LCD panel's maintenance mode.

MORE? The LCD panel and buttons also serve other useful purposes, such as returning the switch to factory default settings or rebooting the switch without requiring a computer for management. See the *LCD Panel in EX3200 and EX4200 Switches* documentation at the EX Switches section at <http://www.juniper.net/techpubs>.

- **Status LEDs,** located next to the LCD buttons, illuminate in various colors to report the status of the switch.
- **Uplink module:** An optional, field-replaceable unit (FRU) optical interface uplink module can be installed in the slot located on the lower-right corner of the EX4200 switch. The optional front-panel uplink modules can support either four gigabit Ethernet (GbE) ports with SFP optical transceivers, two 10GbE ports with XFP optical transceivers, or a user-configurable option offering either two 10GbE or four GbE ports with SFP+ optical transceivers for high-speed backbone or link-aggregation connections between wiring closets and upstream aggregation switches.

NOTE The uplink module that can be configured for either two 10GbE ports with SFP+ optical transceivers, or four GbE ports, has SFP fixed port numbering. Therefore, when the uplink is configured in 10GbE mode, the ports that should be configured are the first (0) and third (2) ports

on the module. For example, if the SFP+ uplink module is present in the first Virtual Chassis member switch (0), the 10GbE ports that need to be configured are xe-0/1/0 and xe-0/1/2.

In addition to the SFP+ optical transceivers and fiber cables, another physical cable that has transceivers directly attached to the cable itself (copper), or Direct Attach Cable (DAC), can be used in environments where interconnected devices are located in relatively close proximity such as an adjacent rack or within the rack. Such DAC, also known as , are suitable for short distances, typically up to 7 meters (23 ft), on most models of EX Series Switches. The primary advantage of the DACs is the high cost-effectiveness for providing network connectivity in short distance. Please see *SFP+ Direct Attach Cables for EX Series Switches* documentation at <http://www.juniper.net/techpubs>.

- Network port: An EX4200 switch offers either 24 or 48 10/100/1000BASE-T Ethernet ports located on the front panel where hosts are typically connected. A model offering 24 100BASE-FX/1000BASE-X SFP optic ports is also available with the EX4200 line of switches.




Figure 1.2 EX4200-48T Ethernet Switch Front Panel

Managing an EX Series Ethernet Switch

An EX Series switch can be managed by either the Junos command-line interface (CLI), or by a web-based interface such as Juniper Web Device Manager or J-Web. The CLI can be accessed two ways: in-band or out-of-band. Neither method is necessarily better than the other, and the choice is really a personal preference. Whichever method is used, however, the first step is to connect to the switch and log in. (This book assumes that the switch has been powered on and the boot process has been completed.)

MORE? For more information on getting started with CLI configuration and commands, see *Day One: Exploring the Junos CLI* for step-by-step instructions for logging in to a network device: <http://www.juniper.net/techpubs>.

In-Band Management

It's possible to manage and configure the switch in-band by using the front-panel network ports. Whether this method is selected for convenience, or to comply with corporate policy, in-band management requires minimal up-front configuration.

This method does not require a separate network subnet to be created or utilized; simply use the IP address that has been allocated and configured for the network ports, and connect a computer for management. In-band management is available only when the switch is booted, initialized, and configured properly.

Out-of-Band Management

The rear-panel console or management Ethernet ports can be used for out-of-band switch management. When using the console port, the only requirement is that the computer has terminal emulation software installed that is properly configured for console access.

If you would like to use the management port instead, a minimal configuration requiring a valid IP address and subnet mask, similar to in-band management, is needed. When using the management port, the switch is accessed via an out-of-band port rather than through the in-band network ports in the front panel. Whichever out-of-band management method is used, the switch needs to be booted and initialized properly with minimal configuration for management port.

Management During the Initial Installation

One of the common methods of the initial installation and set up of the switch is to use out-of-band management method via console. During this time, the red alarm LED on the front panel status LEDs is lit. This tends to cause an alarm since the LED will remain lit as long as the management Ethernet port is not connected and status is down as shown:

```
user@switch> show chassis alarms
1 alarms currently active
Alarm time           Class  Description
2014-xx-xx 10:37:04 PST Major  Management Ethernet Link Down
```

This alarm LED will not be turned off even when the management Ethernet interface is configured to be shut (administratively disabled), and will turn off once the management Ethernet port is connected and the status changes to up, possibly during the actual deployment of the switch after the completion of initial configuration. However, if there is no plan to utilize the management Ethernet port, this can potentially mask a real issue later where such alerts can be present. This behavior of setting off a major alarm when the management Ethernet port is down can be turned off by the user:

```
user@switch# set chassis alarm management-ethernet link-down ignore
```

TIP By default, the EX Series switch has a user login credential of `root` as the username and no password. See *Day One: Configuring Junos Basics* for how to change the Junos password for your device: <http://www.juniper.net/techpubs>.

J-Web Management

Juniper Web Device Manager (J-Web) is a graphical user interface (GUI) that you can use to manage the switch. With J-Web, it is possible to navigate the interface, scroll pages, and expand and collapse elements just like a typical Web browser.

The J-Web interface provides GUI tools for performing all the same tasks available via the Junos CLI, including a CLI Viewer to observe the current configuration, a CLI Editor for viewing and modifying the configuration, and a Point & Click CLI editor for navigating through all of the available CLI statements.

MORE? To learn more about the Junos Web Device Manager, see the *Connecting and Configuring an EX Series Switch J-Web Guide* at <http://www.juniper.net/techpubs>.

Summary

This chapter discussed the different ways of consoling to your EX switch. Again, there is no right or wrong way to console, there is only the way that you might prefer. Junos provides multiple methods for the initial configurations and deployment of your EX Series Ethernet Switch.

You'll use this information throughout this book as it helps you place your EX Switch within your network and configure it.

Now that you know what one switch looks like, let's turn to how to set up multiple EX switches together in a Virtual Chassis, interconnecting and operating as a single, high-bandwidth device.

Chapter 2

Virtual Chassis Physical Connections

<i>Virtual Chassis Configuration</i>	16
<i>Virtual Chassis Port Numbering</i>	21
<i>Virtual Chassis Implementation</i>	24
<i>Network Role</i>	31
<i>Link Aggregation Group (LAG)</i>	34

Most of Juniper Networks EX Series Ethernet switches offer Virtual Chassis technology, which allows a various number (as many as ten on EX3300, EX4200, EX4300, EX4500, EX4550, and EX4600 models on supported Junos releases) of switches to be interconnected and operated as a single, high-bandwidth device. The EX4200 line of Ethernet switches (or Virtual Chassis members) in particular, which is covered in this book, can be interconnected via the dedicated Virtual Chassis ports on the rear panel of each switch, through optional uplink module ports, or via front-panel optical SFP network ports configured as Virtual Chassis ports on an EX4200-24F switch.

Supported EX Series Ethernet switches deployed in a Virtual Chassis configuration are managed and monitored by Virtual Chassis Control Protocol (VCCP) as a single, logical device. This approach greatly simplifies network operations, allows the logical grouping of physical devices even if they reside in different locations, and provides efficient utilization of resources.

This chapter covers how Virtual Chassis configurations are formed using various interconnection methods, along with design considerations for Virtual Chassis configuration.

Virtual Chassis Configuration

EX4200 switches can be deployed as part of a Virtual Chassis configuration in a variety of ways: in a single rack, across several racks, in a single wiring closet, or spanning multiple wiring closets on different floors or in different buildings.

There are two types of physical Virtual Chassis configurations. One, called a “dedicated configuration,” consists of adjacent switches interconnected with special Virtual Chassis port cables connected to the rear-panel Virtual Chassis ports on each switch as shown in Figure 2.1.

A Virtual Chassis configuration may be extended by using optional uplink ports, or by configuring front-panel optical SFP network ports on EX4200-24F switches as Virtual Chassis ports to allow a greater distance between two directly-connected member switches. A Virtual Chassis configuration interconnected via GbE or 10GbE uplink ports or front-panel optical SFP network ports is called an “extended configuration” and is shown in Figure 2.2.




Figure 2.1 Dedicated Virtual Chassis Configuration




Figure 2.2 Extended Virtual Chassis Configuration

There are three basic cabling options for interconnecting switches in a Virtual Chassis configuration: daisy-chained ring, braided ring, and extended Virtual Chassis configuration.

BEST PRACTICE

Virtual Chassis technology does not require cable connections to be in the form of a ring. However, it is highly recommended that you close the loop with a ring configuration to provide resiliency.

Daisy-chained Ring Configuration

In a daisy-chained ring configuration, each member in a Virtual Chassis configuration is connected to the member immediately adjacent to it. Members at the end of the Virtual Chassis configuration are connected to each other using a long Virtual Chassis cable to complete the ring topology. As shown in Figure 2.3, the daisy-chained ring configuration provides a simple and intuitive method for interconnecting devices.




Figure 2.3 EX4200 Virtual Chassis Configuration in a Ring Topology Using the Daisy-chained Ring Method

Braided Ring Configuration

You can use the braided-ring cabling method to support a Virtual Chassis configuration with Virtual Chassis port cables, as shown in Figure 2.4. In a braided-ring cabling configuration, alternating members of a Virtual Chassis configuration are connected. The two member pairs at each end are directly connected to each other to complete the ring topology.




Figure 2.4 EX4200 Virtual Chassis Using the Braided ring Configuration

Extended Configuration

For extended configurations where Virtual Chassis members are spread across a geographic region, Virtual Chassis members can be interconnected via optional GbE or 10GbE uplink modules, or via the front-panel optical SFP network ports on an EX4200-24F. Ports can be configured to function as Virtual Chassis ports so that interconnected switches are recognized as members of the same Virtual Chassis configuration. Multiple uplinks may also be used to interconnect extended Virtual Chassis configurations for increased bandwidth and path redundancy.

NOTE The extended Virtual Chassis connections can be bundled into a single logical group to provide more Virtual Chassis bandwidth and resiliency on supported Junos releases. If two or more optical ports are configured as Virtual Chassis ports (VCPs) connecting the same member switches, the optical ports configured as VCPs will form a link aggregation group (LAG) automatically when such ports are configured to operate at the same link speeds. For more information on the VCP configurable types of interfaces on supported models of EX Series switches and the maximum number of interfaces as VCPs in a single VCP LAG, see *Understanding EX Series Virtual Chassis Components* at <http://www.juniper.net/techpubs/>.

Use the following CLI command to configure optional GbE or 10GbE uplink ports as extended Virtual Chassis ports:

```
user@switch> request virtual-chassis vc-port set pic-slot <pic-slot> port <port> member <member-id>
```

To provide greater flexibility for various environments, Virtual Chassis configurations can be formed using a combination of both dedicated and extended Virtual Chassis connections.

MORE? Table 2.1 lists the configurable interfaces for Virtual Chassis on supported releases that are currently documented. It is highly recommended to reference the release information for the exact maximum number of interfaces as well as the supported interface on a particular model and Junos release. In addition, there are instances where additional configuration restriction may apply, such as the non-mix-and-match of interfaces with different speeds like VCP in VCP LAG, or the restrictions of 1000BASE-T copper SFP transceiver (EX-SFP-1GE-T) on certain models such as EX3300 or EX4200. For more information go to the following: <http://www.juniper.net/techpubs/> and <http://www.juniper.net/us/en/products-services/switching/ex-series/>.

Table 2.1 A List of Configurable Interfaces for Virtual Chassis

Model	Configurable interfaces for Virtual Chassis on supported Junos releases	Maximum number of interfaces as VCPs in a single VCP LAG
EX2200	Built-in or uplink GbE interface (SFP or RJ-45)	8
EX3300	Uplink dual-speed 10GbE (SFP+)/GbE (SFP) interfaces (last two ports (2,3) are configured as VCPs as default)	-
EX4200	Dedicated 32 Gbps VCP ports (64 Gbps bidirectional/port), uplink GbE (SFP) or 10GbE (SFP+/XFP) interface, front-panel GbE (SFP) interface of EX4200-24F model	8
EX4300	Built-in or uplink 10GbE (SFP+) or 40GbE (QSFP+) interface (built-in 40GbE (SFP+) ports are configured as VCPs as default)	4 for 40GbE (QSFP+) for highest backplane capacity
EX4500	Built-in or uplink 10GbE (SFP+) interface, 32 Gbps VCP ports (64 Gbps bidirectional/port) on Virtual Chassis module	8
EX4550	Built-in or uplink 10GbE (SFP+) interface, uplink 40GbE (QSFP+) interface module, 32 Gbps VCP ports (64 Gbps bidirectional/port) on Virtual Chassis module	8
EX4600	10GbE (SFP+) or 40GbE (QSFP+) interface	16
EX8200	10GbE (SFP+) interface	12

Virtual Chassis Port Numbering

There are two dedicated Virtual Chassis ports on the rear panel of each EX4200 switch, designated VCP 0, and VCP 1. The interfaces for these dedicated ports are operational by default when the ports are cabled with dedicated Virtual Chassis port cables. Virtual Chassis ports do not have port-number dependencies; for example, VCP 0 may be interconnected to VCP 0 or VCP 1 on another Virtual Chassis switch member.

Each switch network port on a Virtual Chassis member is numbered $x/y/z$, where:

- x is the member ID of the switch
- y is the port interface controller (PIC) ID. Network ports are always on PIC 0 and uplink module ports are always on PIC 1
- z is the port number on the uplink or network port PIC

For example, port number 0/1/3 indicates the fourth port (because port numbering starts at 0) on the uplink module (PIC ID 1) on the first member switch (0) in a Virtual Chassis configuration:

```
user@switch> show interfaces ge-0/1/3
Physical interface: ge-0/1/3, Enabled, Physical link is Up
...
```

MORE? If you need more information on getting started with CLI configuration and commands, go get *Day One: Exploring the Junos CLI* for step-by-step instructions on logging in to a network device at: <http://www.juniper.net/dayone>.

Virtual Chassis Member Roles

Each member in a Virtual Chassis configuration is assigned a specific role that determines the functions it performs.

In a Virtual Chassis configuration, one member is assigned the *master* or Routing Engine role (RE), and is responsible for managing other members in the Virtual Chassis configuration. A second member is assigned the *backup* role (BK) and takes over the master role if the master switch should fail. All other members are assigned a *linecard* role (LC). The system executes a mastership election algorithm to determine member roles.

MORE? For more information about the Virtual Chassis mastership election algorithm, see the *Understanding Virtual Chassis Components Guide* at <http://www.juniper.net/techpubs>.

Master Role (RE)

The Master switch in a Virtual Chassis configuration performs the following functions:

- Operates as the active Routing Engine for the Virtual Chassis configuration.
- Manages all member switches in the Virtual Chassis configuration.
- Runs Junos for the Virtual Chassis configuration.
- Runs the chassis management processes and network control protocols.
- Receives and transmits routing information.
- Represents all member switches (the hostname and other properties that are assigned to the master switch apply to all members of the Virtual Chassis configuration).
- Holds the active and master copy of the entire Virtual Chassis configuration.

Backup Role (BK)

The member switch that serves as the backup in a Virtual Chassis configuration performs the following functions:

- Operates as the backup Routing Engine for the Virtual Chassis configuration.
- Maintains synchronization with the master switch so that it can take over the master role in the event of a master switch failure.
- Runs Junos for the Virtual Chassis configuration in a backup role.
- Synchronizes with the master switch protocol states, forwarding table and other configurations, so that it is prepared to maintain network connectivity with no or minimal disruption in case the master switch becomes unavailable.

Linecard Role (LC)

Line card member switches perform the following functions:

- Run Junos for Virtual Chassis configuration in linecard role.
- Detect switch error conditions, such as an unplugged cable, on any interfaces that have been configured through the master switch and relay this information to the master switch.
- Receive updates about forwarding information from the master switch and program these updates into the local Packet Forwarding Engine to forward traffic.
- A line card member in a Virtual Chassis configuration does not run full network control protocols while in that role. However, if a master or backup switch fails, one of the line card switches takes over the backup role.

Member Switch and Member ID

Potentially, each EX4200 switch is eligible to become a member of a Virtual Chassis configuration in a dynamic installation scenario. When an EX4200 switch is powered on, it receives a member ID. If the switch is powered on as a standalone switch, its member ID is always 0. When the switch is interconnected with other member switches in a Virtual Chassis configuration, its member ID (0 through 9) is assigned by the master based on various factors, such as the order in which the switch was added to the Virtual Chassis configuration. As each switch is added and powered on, it receives the next available (unused) member ID, and that member ID is displayed on the front-panel LCD.

If the Virtual Chassis configuration previously included a member switch and that member was physically disconnected or removed from the Virtual Chassis configuration, its member ID is not automatically available for assignment as part of the standard sequential assignment by the master. For example, you might have a Virtual Chassis configuration composed of member 0, member 2, and member 3, because member 1 was removed from the Virtual Chassis configuration. When you add another member switch and power it on, the master assigns it as member 4.

However, you can use a command to explicitly change the member ID of the new member switch to ID 1:

```
user@switch> request virtual-chassis renumber member-id 4 new-member-id 1
```

Virtual Chassis Implementation

There are two methods for implementing Virtual Chassis technology: dynamic and preprovisioning.

The dynamic method offers a simple plug-and-play option for building a Virtual Chassis configuration. While the dynamic method does not require any manual configuration, it does not allow you to select the master and backup switches, and it does not prevent certain user errors, such as adding the wrong switch into a Virtual Chassis configuration.

The preprovisioning method requires prior planning and manual configuration before installing the Virtual Chassis configuration. Since all member switches and their roles in a given Virtual Chassis must be configured manually, this method minimizes user error and provides consistent and deterministic results if a member switch fails.

BEST PRACTICE

Dynamic method is the default setting when the switch is powered up for the first time. However, the preprovisioning method is recommended to minimize potential user errors and maximize operational consistency.

Dynamic Installation

The dynamic installation method can be used to build a Virtual Chassis configuration or to add new members to an existing Virtual Chassis configuration without prior user configuration.

In a dynamic installation, the role (master, backup, or linecard), which a member switch assumes within the Virtual Chassis configuration, can be designated by configuring its mastership priority from 1 to 255. The mastership priority value is the factor with the highest precedence for selecting the master of the Virtual Chassis configuration. When an EX4200 switch powers on, it receives the default mastership priority value of 128. Although it is not required, it is recommended that the master and backup switches be designated by configuring the mastership priority of these switches to be the highest value of all members.

The following CLI configuration can be used to change the priority of each member. For example, member 0 is configured for higher priority of 250 for master Routing Engine role, while member 1 is configured for lower (but still higher than the default value of 128) for backup role:


```
user@switch# set virtual-chassis member 0 mastership-priority 250
user@switch# set virtual-chassis member 1 mastership-priority 200
```

NOTE The Virtual Chassis mastership priority value ranges from 0 to 255.

When assigning mastership priority, it is also recommended that you configure the highest possible mastership priority value (255) for the master and backup switches. This configuration ensures that these members continue to function as the master and backup switches when new members are added to the Virtual Chassis configuration. In addition, doing so helps to ensure a smooth transition from master to backup if the master switch becomes unavailable. This configuration also prevents the original master switch from retaking control from the backup switch when the original master switch comes back online, a situation sometimes referred to as *flapping*, or *pre-emption*, that can reduce the efficiency of system operation.

Factory Defaults

It is recommended that factory defaults be loaded on *all* Virtual Chassis switch members before adding these switches to the Virtual Chassis configuration if the switch is not out of the box. This procedure prevents unexpected behavior during the addition of the new member, such as new master reelection and wiping out the current configuration.

Factory defaults can be loaded in either of the following ways:

1. Use the following configuration mode CLI commands:

```
user@switch# load factory-default
user@switch# set system root-authentication plain-password
```

- Then follow the prompts to configure a root password to apply the change:

```
user@switch# commit
```

2. Use the LCD menus on the switch:

- Press the Menu button next to the LCD panel until Maintenance Menu appears.
- Press the Enter button to select Maintenance Menu.
- Press the Menu button until Load Factory menu appears.
- Press Enter to select.
- Press Enter again to confirm when prompted.

Preprovisioned Installation

A preprovisioned configuration allows you to deterministically control the member ID and role assigned to a member switch by associating the switch to its serial number. A preprovisioned configuration file links the serial number of each EX4200 switch to a designated member ID and role. The serial number must be specified in the configuration file for the member to be recognized as part of the Virtual Chassis configuration.

In this configuration, two members must be configured in the role of routing-engine to become eligible for election as the master and backup switches. When these two members are listed in the preprovisioned configuration, one functions as the master switch of the Virtual Chassis configuration while the other functions as the backup switch. In pre-provision configuration, these two member switches can only have the role of Routing Engine and cannot be manually configured as either master or backup.

Any additional members that are not eligible for election as the master or backup switch can be specified as *line cards* in the preprovisioned configuration.

Explicitly configuring a member switch with the role of line card makes it ineligible for functioning as a master or backup switch.

In addition, the preprovisioned configuration provides the option of not explicitly assigning a role to a member switch, making it eligible for election as the backup if the master or the backup switch fails. It can also become the master switch if both the master and backup switches fail.

The mastership priority value is assigned by Junos based on the specified role:

- The master and backup switches (members in Routing Engine role) are assigned a mastership priority of 129.
- A line card switch is assigned a mastership priority of 0, making it ineligible to participate in the master election.
- A switch that is not explicitly assigned a role is configured with a mastership priority of 128 (default), making it eligible to participate in the master election.

To Configure a Preprovisioned Virtual Chassis

1. Set the Virtual Chassis configuration mode to preprovisioned method on the master switch:

```
user@switch# set virtual-chassis preprovisioned
```

2. Define the members of the Virtual Chassis with serial numbers and their roles:

```
user@switch# set virtual-chassis member 0 serial-number AB0123456789 role routing-engine
user@switch# set virtual-chassis member 1 serial-number CD0123456789 role routing-engine
user@switch# set virtual-chassis member 2 serial-number EF0123456789 role line-card
```

Preprovisioned Virtual Chassis also allows switches that are not currently connected and powered on. Therefore, all relevant configuration can be configured on the master switch during initial staging and installation, then additional member switches can be added at a later time during actual deployment.

To View Virtual Chassis Configuration for All Member Switches

1. Enter the `show virtual-chassis status` command. Note that the Virtual Chassis is shown as preprovisioned with specific priority values assigned to different roles:

```
user@switch> show virtual-chassis status
```

```
Preprovisioned Virtual Chassis
Virtual Chassis ID: 1234.5678.90ab
Virtual Chassis Mode: Enabled
```

Member ID	Status	Serial No	Model	Mstr prio	Role	Mixed Mode	Neighbor ID	List Interface
0 (FPC 0)	Prsnt	AB0123456789	ex4200-48t	129	Master*	N	1	vcp-1
							2	vcp-0
1 (FPC 1)	Prsnt	CD0123456789	ex4200-48t	129	Backup	N	2	vcp-1
							0	vcp-0
2 (FPC 2)	Prsnt	EF0123456789	ex4200-48t	0	Linecard	N	0	vcp-1
							1	vcp-0

2. Enter the `show virtual-chassis vc-port all-members` command to view VCP status on all members:

```
user@switch> show virtual-chassis vc-port all-members
fpc0:
```

Interface or PIC / Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0	Dedicated	1	Up	32000	1	vcp-1
vcp-1	Dedicated	2	Up	32000	2	vcp-0

fpc1:

Interface or PIC / Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0	Dedicated	1	Up	32000	2	vcp-1
vcp-1	Dedicated	2	Up	32000	0	vcp-0

fpc2:

Interface or PIC / Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0	Dedicated	1	Up	32000	0	vcp-1
vcp-1	Dedicated	2	Up	32000	1	vcp-0

TIP Preprovisioned Virtual Chassis offers additional benefit of allowing Nonstop Software Upgrade (NSSU) on supported models. For more information, please see the *Understanding Nonstop Software Upgrade on EX Series Switches Guide* at <http://www.juniper.net/techpubs/>.

BEST PRACTICE One of the most common deployments of EX4200 switches is to deploy them in pairs. In such case, two switches will form a Virtual Chassis for redundancy purposes. One of the important behaviors to understand in this two-member Virtual Chassis configuration is that there is a Virtual Chassis split detection, which is enabled by default. It is generally recommended to disable this particular behavior specifically for this two-member Virtual Chassis configuration to ensure continuous operation in the event of a Virtual Chassis split.

To disable split-detection, enter the following configuration:

```
user@switch# set virtual-chassis no-split-detection
```

Assigning an IP Address to a Virtual Chassis Configuration

A Virtual Chassis configuration is managed as a single logical network element. As such, it has only one management IP address, which is configured on the virtual management Ethernet (VME) interface. This VME interface is a logical IP interface associated with the Virtual

Chassis internal management virtual LAN (VLAN) that connects the management Ethernet interfaces of all member switches in a Virtual Chassis configuration. To assign an IP address, the following CLI configuration can be used:

```
user@switch> configure
[edit]
user@switch# set interfaces vme unit 0 family inet address <ip-address>/<subnet-mask>
```

BEST PRACTICE For better resiliency, it is recommended that the VME interface be configured for IP address management rather than individual Management Ethernet (me0).

Synchronizing Virtual Chassis Members

Whenever the configuration settings on the master switch are changed, propagating changes to all other switches in the Virtual Chassis configuration is recommended. To do this, use the following configuration-mode CLI command:

```
user@switch> configure
[edit]
user@switch# commit synchronize
```

The default behavior of commit can be set to synchronize. This way, it is not necessary to remember to use the above commit synchronize command every time.

The default can be changed by following:

```
user@switch# set system commit synchronize
```

MORE? To learn more about implementing Virtual Chassis technology, see the *Virtual Chassis Technology Best Practices Guide* at <http://www.juniper.net/techpubs/>.

Monitoring Operation with CLI Commands

Virtual Chassis configurations can be monitored with CLI commands. Information can be displayed for all members in a Virtual Chassis configuration or for one specific member.

To view member details for all members in a Virtual Chassis configuration, enter the `show virtual-chassis status` command. The following output is from a Virtual Chassis with dynamic configuration:

```
user@switch> show virtual-chassis
```

```
Virtual Chassis ID: 1234.5678.90ab
```

```
Virtual Chassis Mode: Enabled
```

Member ID	Status	Serial No	Model	Mstr prio	Role	Mixed Mode	Neighbor List ID Interface
0 (FPC 0)	Prsnt	AB0123456789	ex4200-48t	250	Master*	N 1	vcp-0
						1	vcp-1
1 (FPC 1)	Prsnt	CD0123456789	ex4200-48t	200	Backup	N 0	vcp-0
						0	vcp-1

```
Member ID for next new member: 2 (FPC 2)
```

MORE? In addition, different models of the EX Series Ethernet switches, such as EX4200, EX4500, and EX4550, can be mixed together to form a single Virtual Chassis in mixed mode using various interfaces including Virtual Chassis module for EX4500/EX4550. For more information on mixed-mode configuration, please see the *Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches* at <http://www.juniper.net/techpubs>.

Table 2.2 Maximum Number of Switches (Members) in a Virtual Chassis on Supported Junos Releases

Model	Maximum number of switches in Virtual Chassis	Supports Mixed models in Virtual Chassis
EX2200	4	-
EX3300	10	-
EX4200	10	With EX4500 and/or EX4550
EX4300	10	EX4300 as linecard with EX4600
EX4500	10	With EX4200 and/or EX4550
EX4550	10	With EX4200 and/or EX4500
EX4600	10	With EX4300 as linecard
EX8200	4 with XRE200 External Routing Engine	-
EX9200	2 with dual Routing Engines installed on both switches	-

NOTE It is highly recommended that you verify the exact number of switches (members) in a Virtual Chassis by referencing documents such as *EX Series Virtual Chassis Overview*, as well as *Junos for the EX Series* release notes at <http://www.juniper.net/techpubs> since the maximum number may vary depending on the model and the Junos release.

MORE? Unlike other EX Series switches such as EX4200, the EX8200 modular switches require XRE200 external routing engines for Virtual Chassis configuration. The Virtual Chassis Control Interface (VCCI) module on the XRE200 external routing engine will be connecting to the GbE management (MGMT) interface of the routing engine of EX8216 or Switch Fabric and Routing Engine of EX8208 (up to four EX8200 switches, which can be mixed). For redundancy purposes, a backup XRE200 external routing engine with full-mesh topology is recommended, where the redundant XRE200 will also be interconnected to the master XRE200. In addition, the member EX8200 switches also require VCP links between them. For more information on EX8200 Virtual Chassis, see *EX8200 Virtual Chassis Overview* at <http://www.juniper.net/techpubs>.

MORE? The EX9200 Series switches can also be configured as Virtual Chassis on supported Junos releases. Up to two EX9200 switches can be configured as a Virtual Chassis, and both EX9200 switches must have dual routing engines installed. For more information on the EX9200 Virtual Chassis, see *Understanding EX9200 Virtual Chassis* at <http://www.juniper.net/techpubs>.

Network Role

With the details of Virtual Chassis technology covered, you might wonder where you would actually deploy a Virtual Chassis configuration. First, however, some fundamentals of network roles should be covered.

An enterprise LAN architecture may span up to three layers, from end-user computers and devices connected to wiring closet switches at the access layer to the core layer at the center of a large enterprise LAN. This hierarchical topology segments the network into physical building blocks, simplifying operation and increasing availability. Each layer within the hierarchical infrastructure has a specific role:

- Access layer: provides an access control boundary and network connectivity to end users in a LAN.
- Aggregation layer: aggregates connections and traffic flows from multiple access-layer switches delivering traffic to core-layer switches.
- Core layer: provides connectivity between aggregation-layer switches and the routers connecting to the WAN or the Internet to enable network collaboration.

NOTE This book focuses primarily on three-layered LAN designs, although you can implement a two-layered design with a converged aggregation and a core layer that is prevalent in either small campuses or branches as well as in data centers with a collapsed core/aggregation design.

MORE? To learn more about designing an Enterprise network, see the *Campus LAN Design Guide* at <http://www.juniper.net/techpubs>.

Access Layer

The access layer provides network connectivity to the network's users by connecting devices such as PCs, network printers, IP phones, and Power over Ethernet/Power over Ethernet+ (PoE/PoE+) cameras to the LAN. Access-layer switches are typically deployed in the wiring closets of each floor in each building or facility.

Typical LANs use VLANs to logically group sets of users, devices, or data, which reside in the access layer, into logical networks through software configuration instead of physically relocating devices on the LAN. VLANs help address issues such as scalability, security, and network management, covered in detail in Chapter 4.

The EX4200 Ethernet switch with Virtual Chassis technology would be an access-layer solution with either 24 or 48 10/100/1000BASE-T ports or 24 100BASE-FX/1000BASE-X ports. One of the unique advantages of the EX4200 Ethernet switches is their pay-as-you-grow design; you can start with a single EX4200 switch and incrementally add up to nine more switches to the Virtual Chassis configuration.

Each EX4200/EX4300 Ethernet switch supports optional uplinks that can be used to interconnect the switches from the access layer to the aggregation layer. In case a high port density is needed in a wiring closet or end-of-row for GbE ports, EX6200 can be another great switch for such an environment. For a single box solution, where hardware redundancy isn't required and the port count is 48 or less, the EX3200/EX3300 or EX2200 are ideal switches for this type of deployment.

Aggregation Layer

The aggregation layer, sometimes referred to as the *distribution layer* in Enterprise networks, aggregates connections and traffic flows from multiple access-layer switches to provide high-density connectivity to the core layer. The primary function of switches at the aggregation

layer is to provide scalability, high density, and high availability.

The EX4200/EX4300 switches in a Virtual Chassis configuration, EX4500/EX4550, EX4600, or the EX8200 line of modular Ethernet switches can provide the required performance and services needed at the aggregation layer. The EX6200 can also be deployed as end-of-row GbE aggregation in the data center. The EX4500/EX4550 is a 10GbE (or GbE) switch with modular uplink slots while EX4600 is a 10GbE switch with 40GbE uplink capability.

The EX8200 line of Ethernet switches offers up to 64 (8-slot chassis) or 128 (16-slot chassis) 10GbE ports at line rate. When greater 10GbE densities are needed, line cards such as 40-port 10GbE line card or combination line cards are available for environment where port density is the most important requirement.

The EX4200-24F 24-port 100BASE-FX/1000BASE-X switch with optional two-port 10GbE uplink module in a Virtual Chassis configuration is a solution for low-to-medium density GbE aggregation layers.

MORE? For more information about the EX4500/EX4550 and EX6200/EX8200 line of modular Ethernet switches, see the product information at <http://www.juniper.net/techpubs>.

Core Layer

The core layer, sometimes referred to as the *backbone*, provides a fabric for high-speed packet switching between multiple aggregation layers or the access layer in a collapsed network, serving as the foundation for reliability and efficiency.

The core layer typically utilizes a 10GbE interface to handle the high amount of throughput and performance. High availability is also an important aspect; the core layer typically incorporates multiple core layer switches to provide system and network redundancy.

The EX8200 line of modular Ethernet switches offers a core-layer solution as it provides redundant Routing Engines and switch fabrics, as well as redundant power supplies and fans. In addition, redundant links to each core layer device are provided in the event of a device or link failure.

Further, the EX9200 line of modular Ethernet switches can provide programmable, flexible, and scalable modular core solutions, especially for cloud applications and server virtualization across campus and data centers.

MORE? For data center networks requiring scalability beyond Virtual Chassis, a Virtual Chassis Fabric architecture that utilizes spine-leaf architecture, or a QFabric system can be considered. See the *Juniper Switches Provide Scale And Simplicity For Any Size Data Center* solution brief at <http://www.juniper.net/us/en/products-services/switching/qfabric-system> for additional information.

As for providing link redundancy, connecting multiple redundant links between network devices would be the first step, and another solution is to group the multiple links as if they are a single high-capacity link between the network devices by using a link aggregation group.

Link Aggregation Group (LAG)

A link aggregation group (LAG) is a group of multiple physical links combined in a single logical bundle. The LAG balances traffic across the member links within an aggregated Ethernet bundle and effectively increases the link bandwidth as shown in Figure 2.9. Another advantage of link aggregation is increased availability, because the LAG is composed of multiple member links. If one member link fails, the LAG continues to carry traffic over the remaining links.

A LAG is typically configured on the EX Series Ethernet switch uplink where uplink ports are connected to other network devices upstream, providing the benefit of LAG for hosts downstream.

A LAG can be either a Layer 2 port or Layer 3 port (port-layer mode is covered in Chapter 4). You can configure LAG by either static or dynamic methods, and when configuring using dynamic methods, Link Aggregation Control Protocol (LACP) can be used.



Figure 2.5 Two EX Series Ethernet Switches Connected via LAG

Link Aggregation Group Guidelines

Some guidelines to keep in mind when configuring a LAG on an EX Series Ethernet switch:

- A LAG is configured as an aggregate Ethernet interface.
- All link speeds and duplex settings need to be identical.

- The maximum number of physical links in a LAG is eight for the EX2200, EX2200-C, EX3200, EX3300, EX4200, EX4500, EX4550, and EX6200 switches, or twelve for EX8200 switches.
- Up to 32 LAGs are supported for EX2200, EX2200-C, and EX3200.
- Up to 111 LAGs are supported for EX3300 EX4200, EX4500, EX4550, and EX6200 from Junos 12.3.
- Up to 255 LAGs are supported on EX8200 Ethernet Switches.
- The LAG must be configured on both sides of the link.

ALERT The maximum number of LAGs supported may change through different releases of Junos. It is recommended to check the maximum number of LAGs for a particular release at <http://www.juniper.net/techpubs>.

NOTE It is not necessary to make the ports in LAG contiguous; in case of a Virtual Chassis configuration, LAG can be across switch members.

Link Aggregation Control Protocol (LACP)

Per IEEE 802.3ad specifications, LACP defines the bundling of multiple physical ports. LACP provides basic error checking for misconfiguration, ensuring that LAG is properly configured on both ends of the LAG. Should there be a misconfiguration, the LAG would not become active.

As a part of the protocol definition, LACP exchanges are made between *actors* (the transmitting link) and *partners* (the receiving link). The LACP mode can be either active or passive.

ALERT! If both ends are both in passive mode, they do not exchange LACP packets, which results in the LAG not coming up. By default, LACP is in passive mode. To initiate transmission of LACP packets and thus bring up the LAG, LACP must be enabled in active mode on at least one side of the LAG.

To Configure a Dynamic LAG with LACP

1. Define the number of LAGs in the switch (or in Virtual Chassis configuration):

```
user@switch# set chassis aggregated-devices ethernet device-count 1
```

2. Delete existing interface configuration (using ge-0/0/10 and ge-0/0/11 in this example):

```
user@switch# delete interfaces ge-0/0/10
user@switch# delete interfaces ge-0/0/11
```

3. Configure interfaces to be a part of a LAG:

```
user@switch# set interfaces ge-0/0/10 ether-options 802.3ad ae0
user@switch# set interfaces ge-0/0/11 ether-options 802.3ad ae0
```

4. Configure the LACP (using active mode):

```
user@switch# set interfaces ae0 aggregated-ether-options lacp active periodic fast
```

5. Configure the LAG interface as a Layer 2 trunk port to transport all VLANs. Port modes such as access and trunk are covered in Chapter 4:

```
user@switch# set interfaces ae0 unit 0 family ethernet-switching port-mode trunk vlan members all
```

NOTE By default, the actor and partner send LACP packets every second (fast mode). The interval can be fast (every second) or slow (every 30 seconds).

To View LAG Details for All Members In a LAG

1. Enter the show lacp interfaces ae0 command:

```
user@switch> show lacp interfaces ae0
```

```
Aggregated interface: ae0
```

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
ge-0/0/10	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-0/0/10	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-0/0/11	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-0/0/11	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
ACP protocol:	Receive State	Transmit State	Mux State						
ge-0/0/10	Current	Fast periodic	Collecting distributing						
ge-0/0/11	Current	Fast periodic	Collecting distributing						

MORE? To learn more about Link Aggregation Groups, see *Understanding Aggregated Ethernet Interface and LACP* at <http://www.juniper.net/techpubs>.

Chapter 3

Network Topology (Logical Topology)

<i>Layer 3 (Routing)</i>	39
<i>Layer 2 (Switching)</i>	41
<i>Bridge Protocol Data Unit (BPDU) Protection</i>	46
<i>Redundant Trunk Group</i>	51

Chapter 2 discussed the physical topology (Layer 1 of the OSI Model) and where the EX Series switches can be deployed in the network – the EX8200 or EX9200 at the core/aggregation layers; the EX4200, EX4300, EX4500/EX4550, EX4600, EX6200, or EX8200 in a Virtual Chassis at the aggregation/access layers; and, the EX2200, EX3200, EX3300, EX4200/EX4300 standalone or in a Virtual Chassis configuration, as well as EX6200 at the access layer.

Let's move the layers of the OSI Model up to the data link (Layer 2), and network layer (Layer 3), to discuss where the EX switches fit in the overall network topology. Generally speaking, the data link layer or Layer 2 is responsible for data transfer between entities within the same network. The Layer 2 domain can be confined to a single networking device or it can expand to multiple networking devices (across multiple wiring closets), as shown in Figure 3.1. The network layer, or Layer 3, is responsible for transferring data between networks. It facilitates communication between devices that are in different networks.

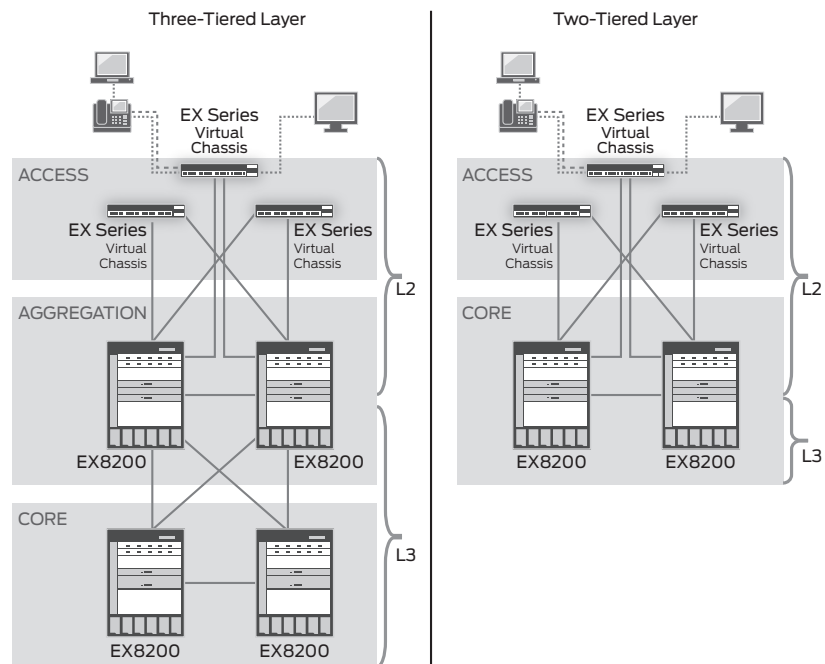


Figure 3.1 Routing and Switching Domains for a Three- and Two-Tiered Network

Layer 3 (Routing)

Routing typically starts at the aggregation layer for the majority of enterprise campus deployments, although there are some deployments that move the Layer 3 boundary from the aggregation to the access. The benefits of routing at the access layer include eliminating spanning-tree and having multipath active-active links.

An IP address defines a host and gives it a “location” within the network. All data that passes through the network starts at an IP host (source) and ends at another host (destination). IP configuration on the EX Series switches follows the same command syntax as the other Junos-based platforms, including the T, M, MX, SRX, and J-series devices.

Layer 3 Interface (IPv4 or IPv6)

EX Series switches support *single stack* (IPv4 or IPv6 only), *dual IP stack* (IPv4 and IPv6), or any combination of single- and dual-stack configurations. The EX4200 includes IPv4 routing and switching, and IPv6 switching and IPv6 routing in the base license.

TIP It is highly recommended that you verify the availability of any particular feature since it may vary depending on the model and the Junos release.

MORE? For detailed information about feature support by Junos version and device, use the Feature Explorer on the Juniper support site at <http://pathfinder.juniper.net/feature-explorer/>.

The following command is an example of an IPv4 address configuration:

```
user@switch# set interfaces ge-0/0/0 unit 0 family inet address x.x.x.x/yy
```

The following command is an example of an IPv6 address configuration:

```
user@switch# set interface ge-0/0/0 unit 0 family inet6 address xxxx::xxxx/yy
```

An IP address can be configured at the physical port or a virtual VLAN interface, also known as the routed VLAN interface (RVI).

Routed VLAN Interface (RVI)

An RVI is a logical Layer 3 interface that provides routing functionality for a given VLAN. Configuring an RVI is a two-step process. The first step is to configure an IP address on the RVI (similar to configuring an IP address on a physical port except that it is for a VLAN interface):

```
user@switch# set interfaces vlan unit 1 family inet address x.x.x.x/yy
```

NOTE For additional RVIs, just increase the unit number. The unit number can be arbitrary and does not have to be sequential. However, it is recommended that the RVI unit number match the VLAN-ID.

The second step is to bind the RVI to a VLAN with the following command:

```
user@switch# set vlans vlan-name 13-interface vlan.1
```

Here is another example, where two RVIs are created for two different VLANs:

```
user@switch# set interfaces vlan unit 1 family inet address 10.0.1.1/24
user@switch# set interfaces vlan unit 2 family inet address 10.0.2.1/24
```

```
user@switch# set vlans vlan-1 13-interface vlan.1
user@switch# set vlans vlan-2 13-interface vlan.2
```

NOTE To configure an IPv6 address, use family inet6.

Routing Protocols (OSPF)

The next step is to enable a routing protocol. Similar to other Junos-based platforms, routing protocol configuration is performed under the protocols stanza in Junos. The EX4200 Series switches support RIP, OSPF, IS-IS, and BGP. RIP and OSPF are part of the base license, whereas IS-IS, MPLS, and BGP require the Advanced Feature License.

NOTE This book focuses on basic OSPF configuration and does not go into detail about the OSPF protocol itself. For more advanced configurations on OSPF, or for configuring other routing protocols, please reference the *Junos Routing Protocols Configuration Guide* at <http://www.juniper.net/techpubs>.

OSPF is a two-tier hierarchical link-state routing protocol. Each router builds a routing database based on the OSPF link-state advertisement (LSA). The following command enables OSPF on the EX Series switches:

```
user@switch# set protocols ospf area 0.0.0.0 interface vlan.1
```

The `show ospf neighbor` command provides a good OSPF summary between adjacencies, such as the local interface, the IP address OSPF is enabled on, the respective adjacency state, and the neighbor's information:

```
user@switch> show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
172.16.31.2	ge-0/0/23.0	Full	10.0.0.2	128	32
172.16.3.2	vlan.1	Full	10.0.0.3	1	16

Use the `show ospf route` command to view the OSPF routes learned from other OSPF-enabled routers or the `show route` command to view all of the routing tables:

```
user@switch> show ospf route
```

Topology default Route Table:

Prefix	Path	Route	NH	Metric	NextHop	NextHop
	Type	Type	Type		Interface	Address/LSP
1.0.0.1	Intra	Area/AS	BR IP	2	ge-0/0/0.0	192.168.150.2
1.0.0.2	Intra	Area/AS	BR IP	2	ge-0/0/0.0	192.168.150.2
172.16.3.2	Intra	Router	IP	1	vlan.1	172.16.3.2
192.0.0.1	Intra	Router	IP	1	ge-0/0/0.0	192.168.150.2
10.0.0.1/32	Intra	Network	IP	0	lo0.0	
172.16.3.0/24	Intra	Network	IP	1	vlan.1	
172.16.31.0/24	Intra	Network	IP	1	ge-0/0/23.0	
172.16.81.0/24	Intra	Network	IP	3	ge-0/0/0.0	192.168.150.2
172.16.82.0/24	Intra	Network	IP	3	ge-0/0/0.0	192.168.150.2
192.168.150.0/24	Intra	Network	IP	1	ge-0/0/0.0	

Layer 2 (Switching)

The Layer 2 (switching) domain is typically at the access layer and can span multiple switches. With Layer 2 loops and the nature of Layer 2 domains, traffic can be broadcast across the domain, creating the possibility of traffic from a source returning to that source endlessly (see Figure 3.2) – thus the need for a protocol such as Spanning Tree to manage Layer 2 loops. If the loops are not prevented, then the network is susceptible to outages due to broadcast storms.

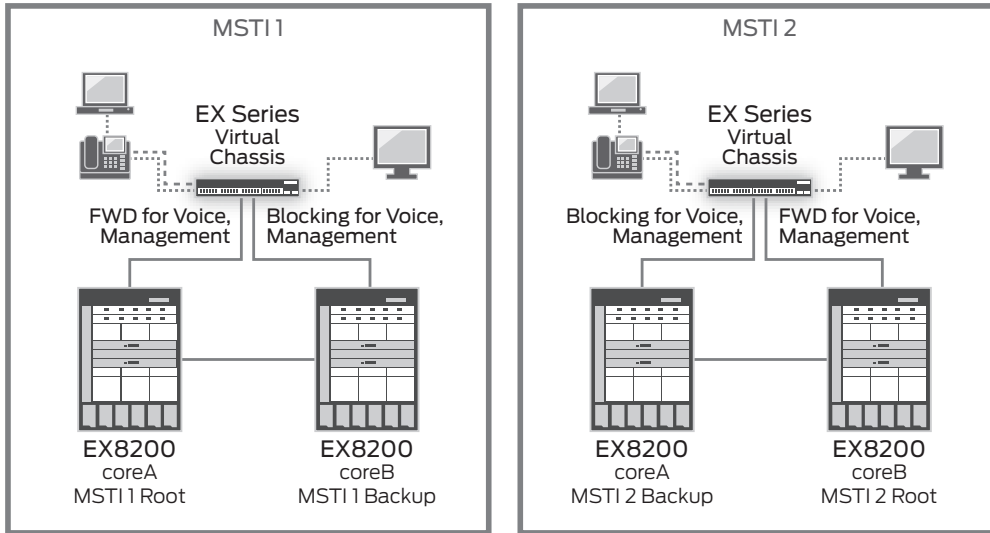


Figure 3.2 Example of MSTP which Provides Active-Active Uplink While Maintaining a Loop-free L2 Topology

Spanning Tree is a Layer 2 protocol that ensures a loop-free network by blocking redundant Layer 2 paths. Bridge protocol data units (BPDUs) are exchanged between switches, which contain Bridge-ID and path-costs. Bridge-ID is composed of bridge-priority and MAC-addresses, which allow switches to elect a root-bridge. Once a root-bridge is elected (lowest Bridge-ID), non-root builds a shortest path to the root bridge and blocks any redundant paths.

EX Series switches support four different flavors of the Spanning Tree Protocol (STP):

- 802.1D (STP): Supports a single instance of the STP (supports one spanning-tree (Layer 2) forwarding topology).
- 802.1w (Rapid Spanning Tree Protocol, or RSTP): Same as STP, but improves the convergence time through the enhancement of bridge communications/interactions. It is backward compatible to STP.
- 802.1s (Multiple Spanning Tree Protocol, or MSTP): MSTP is an extension of RSTP (supports rapid convergence) and increases the number of Layer 2 topology instances in STP. Therefore, each

instance has a different spanning-tree forwarding topology. MSTP supports up to 64 instances, which allows STP to forward traffic on all links but still maintain a loop-free topology. It is backward compatible to STP/RSTP.

- **VLAN Spanning Tree Protocol (VSTP):** VSTP is a per-VLAN STP. Each VLAN has its own spanning-tree instance. VSTP supports rapid convergence as defined by RSTP/MSTP. The EX Series switches support up to 253 VLAN STP instances.

All the spanning-tree protocols are configured under the Junos protocol stanza. This book will cover the basic configurations for RSTP, MSTP, and VSTP.

MORE? To learn more about STPs, please reference the *Spanning Tree Protocols for EX Series Switches*, available at <http://www.juniper.net>.

Rapid Spanning Tree Protocol (RSTP)

RSTP is enabled on the EX Series switches by default. Therefore, one can plug an EX Series switch into the network and, through RSTP, create a loop-free network. However, it is recommended that the bridge priority be configured based on where the switch is placed in the network; bridge priority either increases or decreases the likelihood that the switch will become a root bridge. A lower bridge priority increases the chance of the switch becoming a root bridge. Root bridges influence the Layer 2 forwarding topology as each bridge will forward or block links based on the lowest-cost path to the root bridge.

By default, switch bridge priority is 32768. The command to change the priority is:

```
user@switch# set protocols rstp bridge-priority bridge-priority-value
```

The spanning-tree bridge priority value is between 0 and 65535.

Multiple Spanning Tree Protocol (MSTP)

Besides being an extension of RSTP and supporting the rapid convergence defined by that protocol, MSTP increases the number of supported spanning-tree instances from 1 (STP/RSTP) to 64. This allows VLAN load balancing between a pair of redundant uplinks (active-active uplinks), providing a better link usage in comparison to STP/

RSTP (active-standby uplinks).

NOTE MSTP cannot be enabled with other STPs; therefore, you must “delete” or “deactivate” any other running STPs.

To take advantage of these features, all MSTP-enabled switches must be part of the same *region*. A region is a group of MSTP switches that all have the same MSTP parameters — configuration name, revision level, and MSTI (the number of MSTIs and VLAN mapping must be identical). If any of these parameters are different, then the switches will be in different regions, eliminating the ability to support multiple spanning-tree instances between the switches:

```
user@switch# set protocols mstp configuration-name configuration-name
user@switch# set protocols mstp revision-level revision-level-number
```

NOTE Common spanning tree (CST) bridge priorities and spanning-tree timers are configured under the main MSTP context.

Multiple Spanning-tree Instances (MSTI)

MSTI is a mapping of VLAN(s) to a spanning-tree instance. A group of VLANs mapped to the same MSTI implies those VLANs share the same spanning-tree forwarding topology. This is because each MSTI builds the shortest path to the MSTI root bridge of which it is a part. MSTI Bridge-ID is locally significant to that instance.

The following is a mapping of a VLAN to the instance:

```
user@switch# set protocols mstp msti msti-number vlan vlan-ids
```

The MSTI number can be any number between 1 to 64. VLAN-IDs can be configured as a name, or `vlan-id`, or as a range (1-100, [1 3 5 7-10]).

The following command is used to configure the bridge-priority (0 to 65535) for the MSTI:

```
user@switch# set protocols mstp msti msti-number bridge-priority bridge-priority-value
```

VLAN Spanning Tree Protocol (VSTP)

VSTP provides multiple spanning-tree instances, but there is just one spanning-tree instance for each VLAN. This is in contrast to MSTP, which allows the mapping of many VLANs to one instance. However, it has some similarities to RSTP/MSTP in terms of functionality: it follows the same port states and roles; and, it also utilizes the rapid convergence that is commonly seen with RSTP/MSTP.

Each VLAN can be configured with unique bridge-priority and spanning-tree parameters. The following command is used to enable VSTP on a VLAN:

```
user@switch # set protocols vstp vlan vlan-id
```

The following command is used to configure bridge-priority for a given VLAN:

```
user@switch# set protocols vstp vlan vlan-id bridge-priority bridge-priority-value
```

NOTE RSTP can be configured with VSTP. This allows interoperability with Cisco PVST+/R-PVST+.

The following show commands are available for all STPs. The show spanning-tree bridge command can be used to obtain basic spanning tree-information such as protocol, Bridge ID, and timers.

```
user@switch> show spanning-tree bridge
STP bridge parameters
Context ID                : 0
Enabled protocol          : RSTP
Root ID                   : 4096.00:19:e2:50:86:60
Hello time                 : 2 seconds
Maximum age               : 20 seconds
Forward delay              : 15 seconds
Message age               : 0
Number of topology changes : 10
Time since last topology change : 7642 seconds
Local parameters
Bridge ID                 : 4096.00:19:e2:50:86:60
Extended system ID        : 0
Internal instance ID      : 0
```

Another useful command is the show spanning-tree interface, which shows the interface spanning-tree port states and port roles:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae0.0	128:1	128:1	4096.0019e2508660	10000	FWD	DESG
ge-0/0/0.0	128:513	128:513	4096.0019e2508660	20000	FWD	DESG
ge-0/0/3.0	128:516	128:516	32768.0019e2508660	20000	BLK	DIS
ge-0/0/4.0	128:517	128:517	32768.0019e2508660	20000	BLK	DIS
ge-0/0/5.0	128:518	128:518	32768.0019e2508660	20000	BLK	DIS

The following command is specific to MSTP. It provides a summary of MSTP configuration, such as configuration name, revision level, and MSTI-VLAN mappings. It is a good validation command to see whether a switch is part of the desired MSTP region:

```
user@switch> show spanning-tree mstp configuration
MSTP information
Context identifier      : 0
Region name            : MST-Region-1
Revision               : 2
Configuration digest    : 0x57c9f50482c9c9ae3c404a5d3212715d
```

MSTI	Member VLANs
0	0,401-4094
1	1-100
2	101-200
3	201-300
4	301-400

Bridge Protocol Data Unit (BPDU) Protection

Layer 2 loops can be prevented using various Spanning Tree Protocols (STP, RSTP, MSTP, and VSTP) covered previously. All the Spanning Tree Protocols use bridge protocol data unit (BPDU) to communicate in order to provide the loop-free network. It is essential to protect such communication in order to prevent STP miscalculation, which can potentially cause network instability and outage.

There are several methods to provide BPDU protection. Depending on the type of configuration being used, the EX Series Ethernet switches provide various methods to accommodate such different kinds of environments.

BPDU Protection For Edge Interfaces

Typically, at the access layer, switches will be connecting to many end-hosts including PCs, IP phones, or servers. Such hosts are not generally expected to participate in STP, and are usually expected to be able to communicate as soon as the devices are physically connected. In such instances, the switch interfaces that are connected to the end-hosts should be configured as edge interface, indicating that those interfaces are the "edge" of the STP boundary. Because these interfaces will transition into forwarding state immediately, it is very important to prevent any BPDUs from entering the network through those edge interfaces.

To Configure BPDU Protection for Edge Interfaces:

1. Define the interface where end-host is connected as edge interface (using RSTP as well as ge-0/0/12 and ge-0/0/13 in this example):

```
user@switch# set protocols rstp interface ge-0/0/12.0 edge
user@switch# set protocols rstp interface ge-0/0/13.0 edge
```

2. Enable BPDU protection on the edge interfaces:

```
user@switch# set protocols rstp bpdu-block-on-edge
```

When a BPDU is received on the edge interface when the BPDU protection is enabled for edge interface, such interface will shut down, thus preventing any traffic from passing through the interface. This can be viewed using the `show spanning-tree interface` command where the interface will indicate the role of “DIS (BPDU-INCON)”:

```
user@switch> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
Interface    Port ID    Designated    Designated    Port    State    Role
              port ID    port ID      bridge ID     Cost
<output truncated>
ge-0/0/12.0  128:518    128:518      32768.0123456789a0  20000  BLK     DIS (Bpdu-
Incon)
ge-0/0/13.0  128:519    128:519      32768.0123456789a0  20000  BLK     DIS (Bpdu-
Incon)
<output truncated>
```

There are two methods to re-enable such interface.

1. The first method is to manually clear the interface:

```
user@switch> clear ethernet-switching bpdu-error interface <interface>
```

2. The second method is to allow interface to re-enable after a certain time period. The timer value can be configured:

```
user@switch# set ethernet-switching-options bpdu-block disable-timeout <time>
```

The time can be defined from 10 seconds to 1 hour (3600 seconds).

BPDU Protection for Interface in Port s Shutdown Mode

There can be instances where BPDU protection is desired on interfaces that are not explicitly configured as an edge interface. As with the BPDU protection on edge interfaces, similar functionality can be provided (using ge-0/0/14 in this example):

```
user@switch# set ethernet-switching-options bpdu-block interface ge-0/0/14.0 shutdown
```

To View BPDU Protection for Interface in Port Shutdown Mode

The interface will be shut down when BPDU is received. The interface state can be viewed by the following CLI command:

```
user@switch> show ethernet-switching interfaces
Interface  State  VLAN members  Tag  Tagging  Blocking
<output truncated>
ge-0/0/14.0 down  default                               untagged Disabled by bpdu-control
<output truncated>
```

The interface can be re-enabled by the same methods as the BPDU protection for edge interface where interface can be manually cleared using `clear ethernet-switching bpdu-error interface` command or using the `disable-timeout` value.

BPDU Protection for Interface in BPDU Drop Mode

If it is not desirable to shutdown the interface when BPDU is received, it is also possible to simply have the BPDUs to be dropped while retaining the interface status as Up to allow the end-host to continue to communicate (all other traffic will pass through on that interface). In this case, it is important to note that STP must not be configured on such interfaces. The following CLI configuration can be used:

1. Disable STP on the interface (in this example, RSTP and ge-0/0/15 is used):

```
user@switch# set protocols rstp interface ge-0/0/15.0 disable
```

2. Set the BPDU protection to drop mode:

```
user@switch# set ethernet-switching-options bpdu-block interface ge-0/0/15.0 drop
```

To View BPDU Protection for Interface in BPDU Drop Mode

The interface will remain in up state when BPDU is received while the BPDUs received continue to drop. All other traffic will continue to pass normally. The interface state can be viewed by the following CLI command:

```
user@switch> show ethernet-switching interfaces
Interface  State  VLAN members  Tag  Tagging  Blocking
<output truncated>
ge-0/0/15.0 up    default                               untagged unblocked-xSTP bpdu
                                                filter enabled
<output truncated>
```


MORE? For more information on BPDU protection, see *Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches* at <http://www.juniper.net/techpubs>.

Root Protection

One of the most important aspects of STP is the root bridge/switch. As all switches participating in STP will be using BPDUs to communicate the information regarding the root switch and subsequent configurations of STP to provide loop-free Layer 2 network, it is important to protect this communication regarding the root switch to prevent any instability or sub-optimal forwarding at Layer 2.

The root protection can be enabled on interfaces that should not be receiving lower-priority (superior) BPDU from the root bridge and should not be elected as the root/alternate port. Typically, such interfaces are designated interface (downstream) on the administrative boundary of STP. The root protection will block the interface from becoming the root/alternate port, and is enabled for all instances of STP. The root protection will be participating in the STP, however, and will block the interface when it receives superior BPDUs.

NOTE An interface cannot be configured for both root protection and loop protection at the same time.

To configure root protection, the following CLI configuration can be used (using RSTP and interface ge-0/0/16 in example):

```
user@switch# set protocols rstp interface ge-0/0/16.0 no-root-port
```

When superior BPDU is received on the interface where root protection is configured, it can be viewed using the following CLI command, where the role of the interface will be indicated DIS (Root-Incon):

```
user@switch> show spanning-tree interface
Spanning tree interface parameters for instance 0
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
<output truncated>						
ge-0/0/16.0	128:522	128:522	32768.0123456789a0	20000	BLK	DIS (Root-Incon)
<output truncated>						

The interface will transition back to normal (forwarding) state from blocking without any manual intervention once it stops receiving superior BPDUs.

MORE? For more information on root protection, please reference *Understanding Root Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches* at <http://www.juniper.net/techpubs>.

Loop Protection

Switches in STP communicate using BPDUs to provide a loop-free network. Such communication using BPDUs allows switches to correctly determine which interface should be forwarding and which interface should be blocking to prevent any loops.

However, in a case where a blocked interface stops receiving expected superior BPDUs while it is still at up state, a loop can result as such interface can eventually transition to forwarding state. Such error can be caused by hardware error, software configuration, as well as problem with the physical medium between the switch and its neighbor. The loop protection will prevent this behavior and block the interface instead of allowing the interface to transition into forward state.

NOTE As stated in root protection section, an interface cannot be configured for both root protection and loop protection at the same time.

To configure loop protection, the following CLI configuration can be used (using RSTP and interface ge-0/0/17 in the example):

```
user@switch# set protocols rstp interface ge-0/0/17.0 bpdutimeout-action block
```

When superior BPDUs are no longer received on the blocked interface with loop protection configured, the interface state can be viewed using the following CLI command, where the role of the interface will be indicated DIS (Loop-Incon):

```
user@switch> show spanning-tree interface
Spanning tree interface parameters for instance 0
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
<output truncated> ge-0/0/17.0	128:523	128:523	32768.0123456789a0	20000	BLK	DIS (Loop-Incon)
<output truncated>						

The interface will transition out of inconsistent state without any manual intervention once it starts to receive superior BPDUs.

MORE? For more information on loop protection, please reference *Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches* at <http://www.juniper.net/techpubs>.

Redundant Trunk Group (RTG)

Redundant Trunk Group (RTG) is an alternative feature on the EX Series switches, that provides a loop-free Layer 2 topology without requiring STP to be running on the access-layer switch. RTG accomplishes this by making one link active and the other link a standby. For the links that are enabled for RTG, they do not transmit/forward BPDUs and drop BPDUs if received on RTG-enabled ports. Switchover occurs when the physical link is down as shown in Figure 3.3. RTG should only be configured on the access switches.

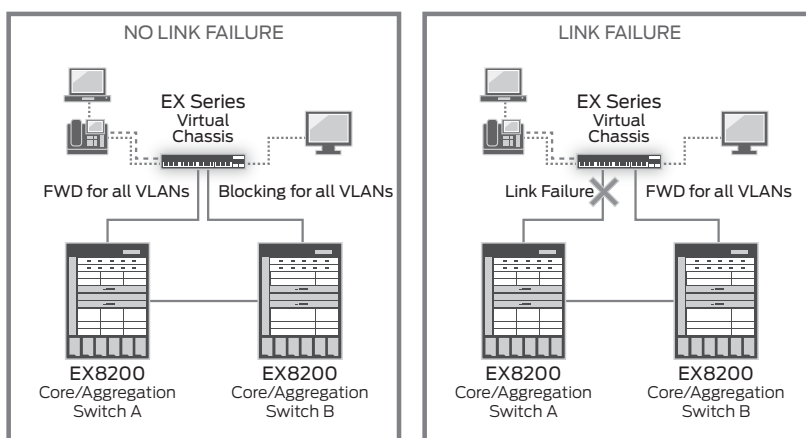


Figure 3.3 RTG Before and After a Primary Link Failure

Up to 16 RTG groups are supported for the EX Series switches, and the EX8200 series and EX8200 Virtual Chassis support up to 254 RTGs. A maximum of two links can be configured in an RTG group; one will be active and forwarding traffic while the other remains in standby mode. The highest numbered interface in an RTG group is the active link, regardless of the order in which the command was entered.

NOTE RTG and STP are mutually exclusive. STP needs to be disabled for interfaces configured for RTG since Rapid Spanning Tree Protocol (RSTP) is enabled by default.

The following command is to disable STP globally:

```
user@switch# delete protocols [stp|rstp|mstp|vstp]
```

The other alternative is to disable STP on an interface:

```
user@switch# set protocols [stp|rstp|mstp|vstp] interface interface-name disable
```

TIP A second option is recommended by keeping STP enabled for other ports that are not enabled for RTG to help prevent any user error that may induce a Layer 2 loop.

RTG is configured under the Junos ethernet-switching-options stanza:

```
user@switch# set ethernet-switching-options redundant-trunk-group RTG-1 interface ge-0/1/0.0
user@switch# set ethernet-switching-options redundant-trunk-group RTG-1 interface ge-0/1/1.0
```

The show redundant-trunk-group command is used to view the RTG link states. Notice that the interface ge-0/1/1.0 is active:

```
user@switch> show redundant-trunk-group
```

Group name	Interface	State	Time of last flap	Flap count
RTG-1	ge-0/1/1.0	Up/Act	Never	0
	ge-0/1/0.0	Up	Never	0

NOTE It is recommended that you keep STP enabled on the core/aggregation switches to protect against any configuration or physical error that can lead to Layer 2 loop.

Primary Keyword

The *primary* keyword serves two functions. First, the link that is configured as primary is active and forwarding during normal operation. Second, it preempts the other link when it is able to take over. Anytime the primary link is up, it will always become active and forwarding:

```
user@switch# set ethernet-switching-options redundant-trunk-group group RTG-1 interface ge-0/1/1.0 primary
```

Notice interface ge-0/1/0.0 is active and has the label “Pri” next to it to indicating that primary was configured on that port:

```
user@switch# run show redundant-trunk-group
```

Group name	Interface	State	Time of last flap	Flap count
RTG-1	ge-0/1/0.0	Up/Pri/Act	Never	0
	ge-0/1/1.0	Up	Never	0

Optionally, the length of time that a re-enabled primary link waits prior to takeover from an active secondary link can be changed (default 120 seconds, ranges from 1 to 600 seconds):

```
user@switch# set ethernet-switching-options redundant-trunk-  
group group <RTG NAME> preempt-cutover-timer <seconds>
```

MORE? To learn more about RTG, please refer to *Understanding Redundant Trunk Links on EX Series Switches* at <http://www.juniper.net/tech-pubs>.

Chapter 4

Ethernet Switching

<i>Virtual LAN (VLAN)</i>	56
<i>Link Layer Discovery Protocol (LLDP)</i>	61
<i>Voice VLAN</i>	65
<i>Interface Range</i>	67

The Ethernet switching process (eswd) is a new process for Junos that is responsible for managing and controlling all Layer 2 functionality for the EX Series switches. Its responsibilities include MAC address table, VLANs, and Layer 2 protocols (i.e., STP, LLDP, etc.). With the introduction of eswd, a few additions were made to the Junos CLI:

- A new family, **ethernet-switching**, has been added. Family **ethernet-switching** transitions a logical unit into a Layer 2 port, and is discussed further under the *Port Mode* section later in this chapter.
- Two new configuration stanzas were introduced in Junos:
 - **VLAN**: Manages VLAN database, membership, and functionality.
 - **Ethernet-switching-options**: Configures Layer 2-specific features, such as voice VLAN, access security (DHCP snooping, Dynamic ARP Inspection, etc.), or other Layer 2-specific features. Access security features are covered in Chapter 5.

Virtual LAN (VLAN)

A local area network (LAN) is a collection of devices that belong to the same Layer 2 broadcast domain – similar to devices connecting to a hub. A virtual LAN (VLAN) extends that concept to multiple logical LANs existing on the same Layer 2 device such as a switch, or essentially a group of switch ports that share the same Layer 2 broadcast domain, as shown in Figure 4.1.




Figure 4.1 EX Series Switch Divided into Multiple Logical VLANs

The EX Series switches support up to 4,094 VLANs, for which any VLAN-ID can be used. By default, all ports are part of VLAN “default” with a null VLAN-ID (as shown below).


```
user@switch> show vlans
```

Name	Tag	Interfaces
default		ge-0/0/0.0*, ge-0/0/3.0, ge-0/0/4.0, ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0, ge-0/0/9.0, ge-0/0/10.0, ge-0/0/11.0, ge-0/0/12.0, ge-0/0/13.0, ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0, ge-0/0/17.0, ge-0/0/18.0, ge-0/0/19.0, ge-0/0/20.0, ge-0/0/21.0, ge-0/0/22.0, ge-0/0/23.0

NOTE The above output may vary depending on the EX Series switch model. The asterisk (*) denotes the port is active (link up).

Adding or deleting a VLAN is done under the VLAN stanza. The minimum VLAN configuration is defining a VLAN name:

```
user@switch# set vlans faculty
```

To delete a VLAN, replace the command `set` with `delete`.

Within the same command line, an 802.1Q VLAN-ID – a numerical value between 1 and 4094 – can be assigned. A VLAN-ID is only required when the switches are connected by a trunk link and extended across the switch. For example:

```
user@switch# set vlans faculty vlan-id 10
```

VLAN Range

A VLAN range allows users to define a range of VLANs with a single command such as:

```
user@switch# set vlans vlan-name vlan-range low-high
```

The `vlan-range` does not support discontinuous-numbered VLAN-IDs. In addition, any attributes configured under the `vlan-range` are inherited by all VLANs in the VLAN range.

For example, the sample configuration below shows a system-generated VLAN name `Bldg_A` with a VLAN-ID range from 20 to 30. The MAC table aging-time has been changed from 300 seconds (default) to 60 seconds. This change will apply for the VLANs in the VLAN range, VLANs 20 to 30:

```
user@switch# show vlans
Bldg_A {
    vlan-range 20-30;
    mac-table-aging-time 60;
}
```

VLANs that are created using VLAN range have the naming structure of __VLAN-name_VLAN-ID__ to provide each VLAN a unique name as shown here:

```
user@switch> show vlans
Name           Tag    Interfaces
__Bldg_A_20__  20
               None
__Bldg_A_21__  21
               None
__Bldg_A_22__  22
               None
__Bldg_A_23__  23
<output truncated>
```

VLAN Membership

Placing a port into a VLAN can be done in one of two ways, either VLAN-centric or port-centric. Neither method offers any advantage over the other, as the results will be the same.

Membership: VLAN-centric

Use the following command to configure the VLAN membership under the VLAN:

```
user@switch# set vlans faculty interface ge-0/0/0.0
```

Membership: Port-centric

Use one of the following commands to configure the VLAN membership under the interface:

```
user@switch# set interfaces ge-0/0/0.0 family ethernet-switching vlan members faculty
```

Or:

```
user@switch# set interfaces ge-0/0/0.0 family ethernet-switching vlan members 10
```

TIP For easier CLI management, centralization of the VLAN membership configuration can be used. For access port, configure all the VLAN membership under the VLAN stanza. For trunk ports, configure all the VLAN membership under the interface (port-centric method). See also the Interface Range section later in this chapter.

VLAN list is supported under the port-centric method. The following configuration, which is very useful for trunk port, is acceptable:

```
user@switch# set interfaces ge-0/1/0.0 family ethernet-
switching vlan members [1 5 7-100]
```

Besides `show vlan` another useful command is `show ethernet-switching interfaces <interface-name>`. This command details the VLAN membership, 802.1Q tag, and forwarding state:

```
user@switch> show ethernet-switching interfaces ge-0/1/0
Interface      State  VLAN members      Tag  Tagging  Blocking
ge-0/1/0.0     up    default           1    untagged unblocked
               faculty          10    tagged  unblocked
               student          30    tagged  unblocked
               voice            5    tagged  unblocked
```




Figure 4.2 VLAN Editing Using J-web

Port Roles (Port Mode)

Endpoints typically dictate the port mode for which the switch is configured. For example, if the end point is host (PC), then typically the port will be configured as an access port. If there is a phone plus a PC, then most likely it is an access-port plus voice VLAN support. The three switch port types are *access* (for end-hosts, IP phones, or servers), *trunk* (for servers, switches, or service devices such as firewall, IDP, etc.), or *routed* (for routers, switches, or wireless access points). Table 4.1 shows a matrix of device and port type.

Table 4.1 Switch Ports Commonly Configured for Endpoints

Port Type \ Device	Access	Trunk	Routed
Host	✓		
Host + IP Telephony	✓	✓	
Server	✓	✓	
Network Devices	✓	✓	✓
Service Devices	✓	✓	




Figure 4.3 Port Configuration Using J-web

An access interface is a Layer 2 port that is a member of one VLAN. It is commonly connected to hosts or servers. To configure use the following:

```
user@switch# set interfaces ge-0/0/0.0 family ethernet-switching port-mode access
```

A trunk interface is a Layer 2 port and a member of multiple VLANs. Common connections are servers, switches, service devices, or any devices that need to extend multiple VLANs over a single link. To configure:

```
user@switch# set interfaces ge-0/1/0.0 family ethernet-switching port-mode trunk
```

A routed interface is an interface with an IP address, usually configured between two routed nodes. Use something akin to the following:

```
user@switch# set interfaces ge-0/1/1.0 family inet address 10.1.3.1/30
```

Access interface with Voice VLAN support is an access port with additional support where the end-host and an IP telephone are connected to the same switch port in a daisy-chain connection (see Figure 4.4). Physically, voice and data traffic are connected to the same port, but logically they are in separate VLANs. The data traffic is sent and received as untagged, whereas the voice traffic is tagged. See Voice VLAN section for configuration.



Figure 4.4 Switch Port Configured as an Access Port With Voice VLAN, IP Telephony, and PC Sharing the Same Switch Port

Link Layer Discovery Protocol (LLDP)

The Link Layer Discovery Protocol (LLDP), defined as the IEEE 802.1AB standard, allows network devices to advertise their identity and capabilities on the LAN. In particular, this advertised information allows EX Series switches to identify a variety of devices that can interoperate efficiently in a LAN.

LLDP-capable devices, called *agent* per standard, transmit information in the form of type, length, and value (TLV) messages, called *Link Layer Discovery Protocol Data Units* (LLD PDUs), to neighboring devices. These messages can include device-specific information such as chassis and port identification, and system name and capabilities. The LLD PDU is sent from each agent, and is stored on the receiving agent. It must be refreshed periodically to remain valid.

By default, EX4200 Ethernet switches have LLDP enabled, but should you need to re-enable them or on other models, use the following CLI configuration:

```
user@switch# set protocols lldp interface all
```

If more granular control is required, LLDP can also be enabled on a per-interface basis by specifying the interface rather than the use of the `all` keyword:

```
user@switch# set protocols lldp interface ge-0/0/0
```

MORE? For additional LLDP configuration information such as LLDP TLV, start timer, and advertise interval settings, please see *Configuring LLDP* at <http://www.juniper.net/techpubs>.

LLDP-MED

LLDP-Media Endpoint Discovery (LLDP-MED) is an extension of the LLDP (IEEE 802.1AB) standard that supports interoperability between voice over IP (VoIP) endpoint devices and other networking end devices. LLDP-MED is commonly used for discovering VoIP phones connected to networked devices such as switches.

In addition to the TLV information that is transmitted on the LLDP agents, LLDP-MED includes additional information such as network policy discovery and Power over Ethernet (PoE) management.

The network policy TLV advertises the VLAN information (see voice VLAN section) for which the interface is configured, as well as associ-

ated Layer 2 and Layer 3 attributes such as 802.1Q tagging, and QoS information such as DSCP. The switch uses this TLV to ensure that voice traffic gets treated with appropriate priority by advertising this information to the IP phone.

For supported devices, the PoE management TLV lets the switch advertise the power level and PoE priority required. For example, the switch can compare the power required by an IP telephone connected to a PoE interface with available resources. If the switch cannot deliver the resources required by the IP phone, the switch could negotiate with the IP phone until a compromise on power is reached.

And the location information advertises the configured physical location of the endpoint. This can be determined either by physical location or by Emergency Line Identification Number (ELIN).

MORE? For additional information about LLDP-MED TLVs, see the documentation on *Configuring LLDP-MED* at <http://www.juniper.net/tech-pubs>.

EX4200 Ethernet switches have LLDP-MED enabled by default, but should you need to re-enable it on other switch models, use the following configuration:

```
user@switch# set protocols lldp-med interface all
```

Similar to LLDP, if more granular control is required, LLDP-MED can also be enabled on a per-interface basis by specifying the interface rather than the use of the `all` keyword:

```
user@switch# set protocols lldp-med interface ge-0/0/0
```

MORE? For additional LLDP-MED configuration information, such as location information and fast start settings that are simply beyond the scope of this book, please see <http://www.juniper.net/techpubs>.




Figure 4.5 LLDP/LLDP-MED Configuration Using J-web

LLDP and LLDP-MED Interaction

By default, interfaces configured with both LLDP and LLDP-MED will only advertise TLVs defined in LLDP. Once the interface detects an LLDP-MED-capable device by receiving LLDP-MED TLVs, the interface will toggle to send LLDP-MED TLVs out on the interface.

For verifying LLDP status on EX4200 Ethernet switches, use the `show lldp` command:

```
user@switch> show lldp
```

```
LLDP                               : Enabled
Advertisement interval              : 30 seconds
Transmit delay                     : 2 seconds
Hold timer                         : 4 seconds
Notification interval              : 0 Second(s)
Config Trap Interval               : 0 seconds
Connection Hold timer              : 300 seconds
```

```
LLDP MED                           : Enabled
MED fast start count                : 3 Packets
```

Interface	Parent Interface	LLDP	LLDP-MED
all	-	Enabled	Enabled

MORE? For more information on the LLDP/LLDP-MED show CLI command output, please see <http://www.juniper.net/techpubs>.

One of the most useful aspects of LLDP information is the list of neighbors on the database of the EX4200 Ethernet switch. Use the `show lldp neighbors` command:

```
user@switch> show lldp neighbors
```

Local Interface	Parent Interface	Chassis Id	Port info	System Name
ge-0/0/0.0	-	00:11:22:33:44:00	ge-0/0/10.0	L2-Switch
ge-0/0/1.0	-	00:55:66:77:88:00	ge-0/0/5.0	L2-Switch
ge-0/0/2.0	-	00:99:aa:bb:cc:00	ge-0/0/12.0	L2-Switch

In the event an existing LLDP neighbor list needs to be cleared, you can clear it using the following:

```
user@switch> clear lldp neighbors
```

Individual interfaces can be specified if it is not desirable to clear the entire database:

```
user@switch> clear lldp neighbors interface ge-0/0/0
```

And it is also useful to see what information is being advertised to the neighbors, as shown here with the `show lldp local-information` command:

```
user@switch> show lldp local-information
```

LLDP Local Information details

```
Chassis ID   : 12:34:56:78:9a:bc
System name  : switch
System descr : Juniper Networks, Inc. ex4200-
448t , version 12.3R6.6 Build date: 2014-xx-xx 08:38:30 UTC
```

System Capabilities

```
Supported   : Bridge Router
Enabled     : Bridge Router
```

Interface name	Parent Interface	SNMP Index	Interface description	Status
Tunneling				
vme.0	-	36	Up	Disabled
me0.0	-	34	Up	Disabled
ge-0/0/0.0	-	503	Up	Disabled
ge-0/0/1.0	-	505	Up	Disabled

<output truncated>

Collected statistics on EX4200 Ethernet switches can be viewed by using the `statistics` keyword:

```
user@switch> show lldp statistics
```

Interface	Parent Interface	Received	Unknown TLVs	With Errors
ge-0/0/0.0	-	158502	0	0
ge-0/0/1.0	-	158510	0	0
ge-0/0/2.0	-	158517	0	0

Discarded TLVs	Transmitted	Untransmitted
0	158502	1
0	158510	1
0	158517	1

Finally, use the `clear` keyword to clear the collected LLDP statistics on the EX4200 switch:

```
user@switch> clear lldp statistics
```

TIP Individual interfaces can also be specified if necessary (similar to the `clear lldp neighbors interface ge-0/0/0` CLI command).

Voice VLAN

Voice VLAN allows 802.1Q-tagged packets into access ports, which is very useful when multiple devices such as an end-host and a VoIP phone are connected to a single port. The EX4200 Ethernet switch can advertise the voice VLAN-ID and QoS information to the VoIP phone through LLDP-MED to ease deployment. Remember that LLDP and LLDP-MED are enabled on EX4200 Ethernet switches by default, so if the VoIP phone supports LLDP-MED, then it will utilize the LLDP-MED information provided by the EX4200 Ethernet switch to tag the voice packets with the appropriate VLAN-ID as well as any QoS markings.

To configure the Voice VLAN feature, first configure the interface as a normal access interface as part of end-host VLAN (please see the earlier section for VLAN membership for configuration examples). Next, enable the Voice VLAN feature with the following command to enable the access port to accept both tagged and untagged packets (where VOIP-VLAN is the VLAN name):

```
user@switch# set ethernet-switching-options voip interface ge-0/0/0.0 vlan VOIP-VLAN
```

An optional command allows LLDP-MED to advertise the QoS code-point associated with the configured forwarding-class when enabled:

```
user@switch# set ethernet-switching-options voip interface <interface_name> forwarding-class <forwarding_class_name>
```

NOTE To advertise the proper QoS code point, a behavior aggregate (BA) must be bound to the interface. See the EZCoS-Voice section in Chapter 5.

MORE? For more information on IP telephony with the EX Series Ethernet Switches, see the *Deploying IP Telephony with Juniper Networks EX Series Ethernet Switches* application note at <http://www.juniper.net/products-services/switching/ex-series>.

Validating or Determining Port States

The following show commands are helpful for validating or determining the port state. The `show interface interface_name` command is useful to see what the port type is:

```

user@switch> show interfaces ge-0/0/0.0
Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 119)
Flags: Device-Down SNMP-Traps Encapsulation: ENET2
Input packets : 0
Output packets: 710
Protocol eth-switch    <-- L2 port
Flags: Is-Primary      <-- no flags, therefore access-port

```

```

user@switch> show interfaces ge-0/0/0.0
Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 119)
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Input packets : 0
Output packets: 710
Protocol eth-switch    <-- L2 port
Flags: Trunk-Mode      <-- trunk port

```

```

user@switch> show interfaces ge-0/0/0.0
Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 119)
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Input packets : 0
Output packets: 711
Protocol inet          <-- L3 port
Flags: None
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 192.168.32/24, Local: 192.168.32.1,
Broadcast: 192.168.32.255

```

Another useful command is `show ethernet-switching interface <interface_name> detail`. This Layer 2 show command provides information on Layer 2 port state, VLAN membership, port forwarding states, and number of learned MAC address:

```

user@switch> show ethernet-switching interfaces ge-0/0/22 detail
Interface: ge-0/0/23.0, Index: 68, State: up, Port mode: Access
VLAN membership:
    student, 802.1Q Tag: 30, untagged, msti-id: 0, unblocked
    voip-vlan, 802.1Q Tag: 5, tagged, msti-id: 0, unblocked
Number of MACs learned on IFL: 2

```

```

user@switch> show ethernet-switching interfaces ge-0/1/0 detail
Interface: ge-0/1/0.0, Index: 69, State: up, Port mode: Trunk
VLAN membership:
    faculty, 802.1Q Tag: 10, tagged, msti-id: 0, unblocked
    student, 802.1Q Tag: 30, tagged, msti-id: 0, unblocked
    voip-vlan, 802.1Q Tag: 5, tagged, msti-id: 0, unblocked
Number of MACs learned on IFL: 1000

```

Interface Range

The interface range function allows users to apply a common set of configurations across a group of interfaces within a given range, simplifying EX Series switch configuration and reducing the number of lines in the configuration file. Interface range is a very useful feature when deploying EX4200 switches in a Virtual Chassis configuration, or when deploying chassis-based modular switches in cases where every interface is not explicitly defined within the default configurations. Interface range is configured under the interface stanza:

```
user@switch# set interfaces interface-range interface-range-name [member|member-range]
```

Use the member-range to add range of interfaces either within or across members/linecards. Note that regular expression is not supported under this statement. An example would be:

- member-range ge-0/0/0 to ge-2/0/47;
- member-range ge-3/0/0 to 3/0/23;

To add individual interfaces, or multiple interfaces using limited regular expression, use an asterisk (*) or a range in form of square brackets [start-end]:

- member ge-0/0/0;
- member ge-0/*/*;
- member ge-0/0/[0-23];

NOTE Multiple member ranges, members, or a combination of both can be configured under the same interface-range group.

Let's use an interface range example: half the ports on an access switch are assigned to the faculty and the other half are assigned to students. Instead of configuring the VLAN membership on a per-interface basis, one can use an interface range command to collectively apply one set of configuration options to the faculty group and another set to the student group:

```
user@switch# set interfaces interface-range faculty-ports member ge-0/0/[0-23]
user@switch# set interfaces interface-range faculty-ports unit 0 family ethernet-switching vlan members faculty
```

```
user@switch# set interfaces interface-range student-ports member ge-0/0/[24-47]
user@switch# set interfaces interface-range student-ports unit 0 family ethernet-switching vlan members student
```

An alternative method of assigning the VLAN membership under the VLAN stanza is to first create two interface groups—one for faculty and one for students—as access groups, and then reference the interface-range group name under the VLAN stanza:

```
user@switch# set interfaces interface-range faculty-ports member ge-0/0/[0-23]
user@switch# set interfaces interface-range faculty-ports unit 0 family ethernet-switching
```

```
user@switch# set interfaces interface-range student-ports member ge-0/0/[24-47]
user@switch# set interfaces interface-range student-ports unit 0 family ethernet-switching
```

```
user@switch# set vlans faculty interface faculty-ports
user@switch# set vlans student interface student-ports
```

Each interface-range group name becomes an interface entity and can be referenced in other parts of the Junos CLI, such as features under ethernet-switching-options or features under protocols like spanning-tree, OSPF, and 802.1X to name a few. By referencing the interface-range group, features will be applied uniformly to all ports within that group. Conversely, a feature can be applied to single port, or subset, of the interface-range group:

```
user@switch# set interfaces interface-range faculty-ports member ge-0/0/[0-23]
user@switch# set protocols rstp interface ge-0/0/0 edge
```

Chapter 5

EX Series Features

<i>OAM Link-Fault Management (802.3ah)</i>	70
<i>MVRP (802.1ak)</i>	71
<i>Multicast and Multicast Routing</i>	73
<i>EZQoS-Voice</i>	75
<i>Access Port Security</i>	79
<i>Power over Ethernet (PoE/PoE+)</i>	84
<i>Port Mirroring</i>	88
<i>sFlow</i>	91

Let's go through some of the EX Series switch features that are commonly used in both campus and branch deployment:

- Ethernet Operation, Administration, and Maintenance (OAM) (802.3ah), which protects against a uni-directional link;
- MVRP (802.1ak) helps VLAN management across switched networks;
- Multicast for delivery options to a subset or group of users;
- EZQoS-Voice that takes the guessing out of CoS configuration;
- Access port security to help protect LAN from man-in-the-middle or denial of service (DoS) attacks;
- Power over Ethernet (PoE) to provide power to connected devices;
- Port mirroring for network policy enforcement or identifying problems such as abnormal or excessive bandwidth during troubleshooting.
- And sFlow for network monitoring and traffic flow visibility.

There are, of course, lots of other features in the EX Ethernet Switch platform that may be taken advantage of in your network. Seek out the documentation, *Junos Enterprise Switching* by Reynolds & Marschke (O'Reilly Media, 2009), and especially the feature overview of each new Junos operating system release available at <http://www.juniper.net/us/en/community/junos/releases/>.

OAM Link-Fault Management (802.3ah)

IEEE 802.3ah is a standards-based feature that encompasses operation, administration, and maintenance (OAM) to help increase reliability and streamline administration and maintenance for Ethernet. The 802.3ah standard is a link-layer and point-to-point protocol; thus, it does not extend beyond the local link. While the 802.3ah standard provides remote failure indication, remote loopback, link monitoring, and discovery, we will focus on how it can be used to detect a unidirectional link, which occurs when a link between two devices is still up, but one device is no longer receiving traffic because of a hardware or software error.

The 802.3ah standard needs to be supported and enabled on the interfaces of both devices. Through discovery – OAM protocol data unit (PDU) – the two endpoints will establish adjacencies and learn

each other's capabilities. If one end loses adjacency at any time, then the interface can be forced down.

802.3ah is configured under the oam stanza in Junos. The first step is to configure the OAM action profile for loss-adjacency; when adjacency is lost, it brings the link down:

```
user@switch# set protocols oam ethernet link-fault-management action-profile action-
profile-name event link-adjacency-loss
user@switch# set protocols oam ethernet link-fault-management action-profile action-
profile-name action link-down
```

Next, enable 802.3ah on the interfaces:

```
user@switch# set protocols oam ethernet link-fault-management interface ge-
0/1/0.0 link-discovery active
```

And the last step is to bind the action profile to the interface:

```
user@switch# set interface ge-0/1/0.0 apply-action-profile action-profile-name
```

You can use the `show oam ethernet link-fault-management` command to validate 802.3ah. The output provides information on the neighboring capabilities as well as the action-profile that has been invoked. When the output displays a MAC address for the peer address, and the discovery state is Send Any, then OAM link-fault-management is configured correctly.

```
user@switch> show oam ethernet link-fault-management
Interface: ge-0/0/23.0
Status: Running, Discovery state: Send Any
Peer address: 00:11:22:33:44:55
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote entity information:
  Remote MUX action: forwarding, Remote parser action: forwarding
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: unsupported, Link events: supported
  Variable requests: unsupported
Application profile statistics:
  Profile Name          Invoked    Executed
  down-link             0          0
```

MVRP (802.1ak)

MVRP is a standards-based protocol that supersedes the Generic VLAN Registration Protocol (GVRP). It is used to dynamically manage VLANs across a Layer 2 network to reduce the management overhead for a switched network and improve the bandwidth efficien-

cy by pruning VLANs on trunk ports. Through join and leave messaging, MVRP allows switches to register or withdraw VLAN information with other switches in the same Layer 2 domain. The joins and leaves are sent across on trunk ports and will follow the active spanning-tree topology as shown in Figure 5.1.




Figure 5.1 VLAN Information Being Propagated Across the Layer 2 Network

MVRP is disabled by default on EX switches and needs to be configured on the trunk ports as follows:

```
user@switch# set protocols mvrp interface <interface-name>
```

VLAN and VLAN membership are configured on the edge switches (*both* edge switches) of the network. MVRP will propagate and build the Layer 2 path between the edge devices.

NOTE To manually configure a port to be part of a MVRP-learned VLAN, the corresponding VLAN-ID needs to be manually configured on the switch.

Similar to VLAN range configuration in Chapter 4, VLANs that are learned by MVRP have the following naming structure: `__mvrp_vlan-id__`. The `show mvrp dynamic-vlan-memberships` MVRP command can be used to view VLAN membership learned from MVRP, (the standard `show vlan` will also display the VLAN learned by MVRP):

```
user@switch> show mvrp dynamic-vlan-memberships
VLAN Name      Interfaces
-----
__mvrp_10__    ge-0/0/0.0
                ge-0/0/1.0
```

Using the `statistics` keyword, you can view MVRP statistics, such as joins and leaves:

```
user@switch> show mvrp statistics interface ge-0/1/0
MVRP statistics
```



```
Interface name           : ge-0/1/0.0
MRPDU received          : 162
Invalid PDU received    : 0
New received            : 0
Join Empty received     : 380
Join In received        : 106
```

<output truncated>

Multicast and Multicast Routing

Multicast allows delivery of packets from a single source to a specific subset of users or many destination members.

Multicast routing is supported in the base license of the EX Series switches. The EX Series switches support three different Protocol Independent Multicast (PIM) modes (PIM is a family of multicast routing protocols for IP networks):

- PIM DM (dense mode, flood and prune): Multicast join requests are initially flooded to all PIM DM-enabled routers. If there are no downstream members, then the router prunes towards the source.
- PIM SM (sparse mode, explicit join): The destination/receiver member must send an explicit “join” request to the rendezvous point (RP) router.
- PIM SSM (source-specific): One-to-many model; receiving hosts must join with either Internet Group Management Protocol version 3 (IGMPv3) or Multicast Listener Discovery version 2 (MLDv2).

NOTE This book only provides configuration syntax for PIM SM and static rendezvous point (RP).

All multicast routing configuration is done under PIM stanza in Junos.

In shared tree, RP is the root of the multicast distribution tree. Initially, the source of the multicast and PIM join requests from the last hop router, first converge at the RP. The RP needs to be reachable by all multicast routers. The following command should be configured on the router designated as the RP:

```
user@switch# set protocols pim rp local address <ip_address>
```

TIP The loopback 0 interface is recommended to be the RP interface.

For all other routers, configure:

```
user@switch# set protocols pim rp static address <ip_address>
```

Any routed interface, including RP interace, that will be routing multicast traffic, needs to be enabled for PIM-SM:

```
user@switch# set protocols pim interface <interface_name> mode sparse
```

The show pim rps command is to verify the RP. Its output provides a RP address, how the RP is learned, number of active multicast groups, and the multicast group the RP can forward:

```
user@swtich> show pim rps
```

Instance: PIM.master

Address family INET

RP address	Type	Holdtime	Timeout	Groups	Group prefixes
10.1.1.1	static	0	None	1	224.0.0.0/4

This show pim neighbors command is used to validate PIM neighbors:

```
user@switch> show pim neighbors
```

Instance: PIM.master

B = Bidirectional Capable, G = Generation Identifier,

H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,

P = Hello Option DR Priority

Interface	IP V Mode	Option	Uptime	Neighbor addr
ge-1/0/23.0	4 2	HPLG	02:18:42	10.1.2.2

The show multicast route command displays the multicast route for a given multicast group, as well as the multicast source and the upstream and downstream multicast path:

```
user@switch> show multicast route
```

Family: INET

Group: 224.0.1.39

Source: 1.1.1.2/32

Upstream interface: ge-0/1/0.0

Downstream interface list:

local ge-1/0/23.0

Multicast Switching

By default, a switch treats a multicast packet much like a broadcast packet – it floods to all ports within the VLAN with the exception of

the source port. IGMP snooping regulates multicast traffic by monitoring the IGMP transmission between the router and host to build a table, associating the Layer 3 multicast group and the switch port on a per-VLAN basis. The switch knows which port to forward the multicast packet. IGMP snooping is enabled by default.

For hosts that do not support IGMP, the group can be manually configured using:

```
user@switch# set protocols igmp-snooping vlan <vlan_name> interface <interface_name> static group <multicast_ip_group_address>
```

The `show igmp-snooping membership` command is to view the IGMP snooping table that was built by the switch. The output provides all the multicast groups on a per-VLAN basis:

```
user@switch> show igmp-snooping membership
VLAN: v2
225.1.1.1      *              199 secs
Interfaces: xe-0/0/1.0, xe-0/0/2.0, xe-0/0/3.0
```

EZQoS-Voice

The EX lines of switches support Class of Service (CoS), which can help meet business applications requirements while ensuring that specialized traffic does not exceed the latency and jitter requirements of the network. The EX Series switches support up to eight CoS queues per port and each queue can be uniquely molded to best serve business needs. In order to ensure that applications meet the required service level, it is recommended to enable CoS end-to-end.

The basic CoS building blocks for the EX Series switches are classification, policing, queuing, scheduling, and remarking, shown in Figure 5.2. Configuring CoS can be a daunting task, as it requires proper knowledge and QoS configuration. *How do I classify traffic? How much bandwidth should I allocate? How much buffer should be allocated between the queues?* These are all questions that users face when deploying QoS.



Figure 5.2 General EX Series Switch QoS Stages

EZQoS-Voice removes the complexity and helps streamline the CoS

configuration for best-effort, video, voice, and network control type traffic. It provides a base configuration that addresses traffic classification, traffic queuing, and traffic scheduling.

NOTE EZQoS-Voice does not implement all of the QoS stages, but it is available if required. For more information on CoS on the EX Series switches, please reference the EX Series switch technical documentation at <http://www.juniper.net/techpub/>.

Classifying Traffic

Classifying traffic is the first QoS process, which is done when the switch first receives traffic. By separating traffic flows, the switch can handle traffic based on its prioritization. Traffic differentiation can be accomplished by using any of the numerous port classification methods:

- Behavior Aggregate (BA): Classify traffic base on 802.1p, DSCP, or IP Precedence.
- Multifield (MF) Classifier: Classifying traffic base on Layer 2, Layer3, and Layer4 information.
- Port-based: Although this isn't differentiating traffic, but rather characterizing all incoming traffic to a specified forwarding-class.

EZQoS-Voice uses BA and classifies traffic based on the DSCP values, some of which are listed in Table 5.1. Based on DSCP, the packet will be associated with a particular class of service servicing level, forwarding class. The forwarding class is mapped to a given egress queue.

Table 5.1 Default Settings for the EZQoS-VOICE Template

Forwarding Class	Queue	DSCP	Scheduler
Best-Effort	0	0-23, 25, 26-33, 35-45, 46-47, 49-55, 57-63	SDWRR
Video	4	34	SDWRR
Voice	5	46	Strict-Priority
Network-Control	7	24, 26, 48, 56	Strict-Priority

Queuing Traffic

The important factors for queuing traffic are the number of queues, the queue depth, and queue management. EX Series switches support up to

eight queues per port, and EZQoS-Voice uses four of the eight queues. Each queue is responsible for certain traffic classes (forwarding-class); EZQoS-Voice uses queues 0, 4, 5, 7, which are associated to best-effort, video, voice, and network-control respectively. Each queue is configured with a different buffer size based on the traffic type and platform.

Scheduling Traffic

There are two different types of queue schedulers that can be configured for the queue – *strict-priority* (strict-high) or *shaped deficit weighted round robin* (SDWRR) (low). If the queue is configured for the strict-high, then any time packets are in this queue they are always serviced. When queues are configured for SDWRR, queues are serviced in round-robin fashion (from high queue to low queue) while preserving the overall bandwidth distribution base on weight.

The bandwidth distribution on the EX4200 switch for best effort and video is 30/70; on the EX8200 it is 20/50. Voice and network control are treated as strict priority, thus anytime voice or network control packets are in queue, they are serviced immediately.

The EZQoS-Voice template is saved as a file, *ezqos-voice.conf* in the */etc/config* directory. Use the `load merge` command to load and merge the EZQoS-Voice template into the configuration:

```
user@switch# load merge /etc/config/ezqos-voip.conf
```

NOTE EZQoS-Voice is an editable template. Administrators can edit or build off of the template to better meet their business or network requirements.

The template is loaded under the Junos group stanza as *ezqos-voip*. Even though the template is part of the configuration, the EZQoS-Voice configuration is not active.

The next step is to activate it by applying the group (*ezqos-voip*) under the Junos CoS stanza:

```
user@switch# set class-of-service apply-groups ezqos-voip
```

Finally, bind the classifier and scheduler to the interface(s):

```
user@switch# set class-of-service interfaces ge-0/0/0 unit 0 classifier dscp ezqos-  
dscp-classifier  
user@switch# set class-of-service interfaces ge-0/0/0 scheduler-map ezqos-voip-sched-  
maps
```

TIP An asterisk can be used to simplify and reduce repetitive configuration for like interfaces (ge or xe). Asterisks will apply the same classifier and/or scheduler to all of the same type interfaces, so: `set class-of-service interfaces ge-* unit 0 classifier dscp ezqos-dscp-classifier`.

The majority of the Junos show commands for CoS are under the `show class-of-service stanza`. The `show interface <interface-name> extensive | find`, and `<Cos Information>` or `show class-of-service interface <interface-name>`, are good summary commands:

```
user@switch> show class-of-service interface ge-0/0/0
```

```
Physical interface: ge-0/0/0, Index: 129
```

```
Queues supported: 8, Queues in use: 5
```

```
Scheduler map: ezqos-voip-sched-maps, Index: 37585
```

```
Logical interface: ge-0/0/0.0, Index: 2684275700
```

Object	Name	Type	Index
Classifier	ezqos-dscp-classifier	dscp	57624

In the sample output shown here, the `show` command provides the number of configured egress queues, the configured scheduler, and the configured and type of classifier.

To view specific classifier or scheduler-map configuration, use:

```
user@switch> show class-of-service classifier name classifier-name
```

```
user@switch> show class-of-service scheduler-map scheduler-map-name
```

Another useful command to validate proper traffic queuing and/or to see any packet drops is the `show interface interface-name [detail|extensive] | find <Queue counters>` or `show interface queue <interface-name>` commands:

```
user@switch> show interfaces queue ge-0/0/0
```

```
Physical interface: ge-0/0/0, Enabled, Physical link is Down
```

```
Interface index: 129, SNMP ifIndex: 501
```

```
Forwarding classes: 16 supported, 5 in use
```

```
Egress queues: 8 supported, 5 in use
```

```
Queue: 0, Forwarding classes: ezqos-best-effort
```

```
Queued:
```

Packets	:	Not Available
Bytes	:	Not Available
Packets	:	41570904
Bytes	:	5320940436
Tail-dropped packets	:	0

```
<output truncated>
```

Access Port Security

Like any other network device on an Ethernet LAN, Ethernet switches are vulnerable to malicious attacks such as address spoofing and man-in-the-middle attacks (shown in Figure 5.3). The EX Series Ethernet switches include many access security features to protect access ports against such attacks, which can disrupt network access and negatively impact productivity. While there are various categories of attacks, the EX Series Ethernet switches allow you to selectively configure the appropriate access security protection features with minimal configurations.




Figure 5.3 Hacker Posing as the Gateway (man-in-the-middle-attack)

DHCP Snooping

Dynamic Host Configuration Protocol (DHCP) dynamically allocates IP addresses for (DHCP) clients, leasing addresses to devices on a temporary basis so that the addresses can be reused. End devices that require IP addresses obtained through DHCP must communicate with a DHCP server across the LAN.

DHCP snooping prevents rogue, non-legitimate DHCP servers by allowing the switch to become aware of DHCP packets. The switch actively filters and blocks incoming DHCP server-type messages on ports that are not defined as DHCP server ports (untrusted ports). On the other hand, the switch builds and maintains a DHCP snooping binding database consisting of DHCP snooping entries where client MAC addresses, obtained IP addresses via DHCP processes, port information, VLAN information, and additional information regarding DHCP leases are stored. Once a DHCP client releases an IP address or a DHCP lease expires, the associated DHCP snooping binding entry is removed from the database.

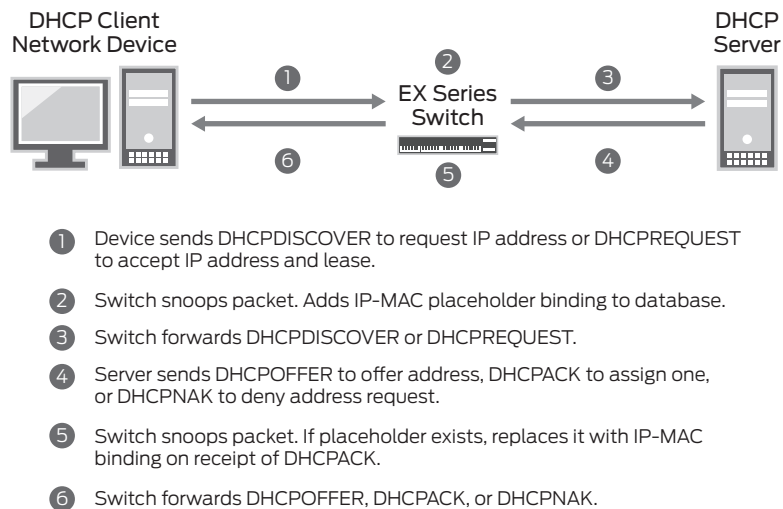


Figure 5.4 DHCP Snooping Process

TIP DHCP snooping is a requirement for other access port security features such as Dynamic ARP Inspection (DAI) and IP source guard.

When enabling DHCP snooping on an EX Series switch, the following guidelines should be kept in mind:

1. All access ports clients are typically expected to be connected to are untrusted, while trunk ports that network infrastructures are connected are trusted by default.

2. On untrusted ports, only DHCP client-type messages such as *discoveries/requests* are allowed; all other DHCP packets are dropped. The switch also builds a DHCP snooping database on these ports where MAC addresses, port locations, VLAN, and IP-binding from DHCP exchanges between the client and server are stored in the database.

3. If you move a network device from one VLAN to another, where typically the device has to acquire a new IP address, its entry in the DHCP snooping binding database including the VLAN ID is updated.

DHCP snooping is most effective in cases where a rogue DHCP server is impersonating a legitimate DHCP server on a LAN segment, providing lease offers to DHCP clients that disrupt their network access. The rogue server might also assign itself as the network's default gateway within the DHCP lease offer packets, enabling the attacker to receive packets from clients to “sniff” network traffic and launch a man-in-the-middle attack, misdirecting network traffic intended for legitimate devices and resources.

The DHCP snooping feature is enabled on a per-VLAN basis. You can use the following configuration to enable DHCP snooping feature on EX Series Ethernet switches:

```
user@switch# set ethernet-switching-options secure-access-port vlan vlan_
name examine-dhcp
```

If there is a local DHCP server connected to the switch on an access port rather than a trunk port, the port characteristics need to be changed from “untrusted” to “trusted.”

It is also important to ensure that the DHCP server interface is physically secure. It is recommended that access to the DHCP server be monitored and controlled at the site before configuring the port as trusted:

```
user@switch# set ethernet-switching-options secure-access-port interface interface_
name dhcp-trusted
```

Use the following command to configure static entry for the DHCP snooping database, for devices that have static IP addresses and do not rely on DHCP.

```
user@switch# set ethernet-switching-options secure-access-port interface <interface_
name> static-ip <ip_address> mac <mac_address> vlan <vlan_name>
```

NOTE By default, the IP-MAC bindings are lost when the switch is rebooted and DHCP clients (the network devices or hosts) must reacquire bindings. However, you can configure the bindings to persist by setting the `dhcp-snooping-file` statement to store the database file either locally or remotely.

This command shows the DHCP snooping binding database:

```
user@switch> show dhcp snooping binding
```

DHCP Snooping Information:

MAC address	IP address	Lease (seconds)	Type	VLAN	Interface
00:01:23:45:67:89	192.168.1.10	-	static	corp-access	ge-0/0/10.0
00:01:23:45:67:90	192.168.2.11	653	dynamic	corp-access	ge-0/0/11.0
00:01:23:45:67:91	192.168.2.12	720	dynamic	corp-access	ge-0/0/12.0

Dynamic ARP Inspection (DAI)

In order to send IP packets on a network (such as an Ethernet network), mapping an IP address (Layer 3) to an Ethernet media access control (MAC) address (Layer 2) is required. Address Resolution Protocol (ARP) is used to map MAC addresses to IP addresses on an Ethernet LAN.

Network devices maintain this mapping in an ARP cache that they consult when forwarding packets to other network devices. If the ARP cache does not contain an existing entry for the destination device, the device broadcasts an ARP request for the destination device's address and stores the response in the cache.

Dynamic ARP Inspection (DAI) validates ARP packets on the network. The switch intercepts ARP packets from access ports and checks them against the IP-MAC database (DHCP snooping binding database) populated through DHCP snooping. Therefore, this feature is dependent on DHCP snooping in order to make filtering decisions upon receiving ARP packets from untrusted ports as defined in DHCP snooping. If a mismatch is found, then the ARP packet is dropped, preventing any man-in-the-middle attacks such as ARP spoofing/poisoning.

ALERT! It is important to remember that DAI is entirely dependent on DHCP snooping, specifically the DHCP snooping binding database. If there is no corresponding DHCP snooping entry in the binding database, any ARP packets received on the untrusted port are dropped.

NOTE The concept of untrusted and trusted ports on DAI and IP source guard is the same as with the DHCP snooping feature.

In an ARP spoofing attack, an attacker generates an ARP packet and sends it to the network, typically to start a man-in-the-middle attack. The attacker associates its own MAC address with the IP address of a network device connected to the switch by sending an ARP packet that spoofs the MAC address of another device (target) on the LAN. A common type of ARP spoofing uses gratuitous ARP; this is a type of ARP packet used when a network device, such as an end host, sends an ARP request to resolve its own IP address. In a normal LAN, this gratuitous ARP message would indicate that there are two devices with the same MAC address. The gratuitous ARP message is also sent when an end host's network interface card is changed, or a device is rebooted, so other network devices on the LAN update their ARP caches.

However, in an ARP spoofing attack, an attacker maliciously poisons the device's ARP cache by announcing itself as the targeted device. Any traffic sent to that IP address is instead sent to the attacker impersonating a legitimate device. Once the attacker is receiving traffic intended for a legitimate device, he can create various types of mischief, including sniffing the packets and launching man-in-the-middle attacks. (In a man-in-the-middle attack, the attacker intercepts messages between two hosts, reads them, and perhaps alters them, all without the original hosts knowing that their communications have been compromised.)

The DAI feature is also enabled on a per-VLAN basis, and you can use the following configuration to enable the DAI feature on EX Series Ethernet switches:

```
user@switch# set ethernet-switching-options secure-access-port vlan vlan_name arp-
inspection
```

Use this show command to show the DAI statistics:

```
user@switch> show arp inspection statistics
```

ARP inspection statistics:

Interface	Packets received	ARP inspection pass	ARP inspection failed
ge-0/0/10.0	9	9	0
ge-0/0/11.0	30	30	0
ge-0/0/12.0	25	24	1

IP Source Guard

IP source guard is effective against IP spoofing attacks on Ethernet LANs. IP spoofing is typically used by attackers to prevent LAN administrators from identifying the actual source of attacks. The IP source guard feature is similar to DAI, although the feature is applicable to IP packets rather than ARP packets from devices on untrusted ports.

TIP A typical form of IP spoofing is a Denial of Service (DoS) attack, where the attacker floods a target with TCP SYN packets in an attempt to overwhelm the device while hiding the actual source of the attack.

The IP source guard feature is dependent on the EX Series DHCP snooping feature because it requires the DHCP snooping binding database to make filtering decisions when inspecting IP packets from devices on untrusted ports. IP source guard cross-checks the IP source address and the port upon which it was received; if the packet does not match the DHCP snooping binding database, then the packet is discarded. The IP source guard feature is configured on a per-VLAN basis:

```
user@switch# set ethernet-switching-options secure-access-port <vlan_name> ip-source-guard
```

The `show ip-source-guard` command shows the IP source guard information:

```
user@switch> show ip-source-guard
```

IP source guard information:

Interface	Tag	IP Address	MAC Address	VLAN
ge-0/0/11.0	0	192.168.2.11	00:01:23:45:67:90	corp-access
ge-0/0/12.0	0	192.168.2.12	00:01:23:45:67:91	corp-access

MORE? For more information about access port security CLI configuration, see the *Port Security on EX Series Switches Guide* at <http://www.juniper.net/techpubs>.

Power over Ethernet (PoE/PoE+)

Power over Ethernet (PoE) refers to the ability to pass electric power over a copper Ethernet LAN cable. PoE is a standard defined as IEEE 802.3af, which specifies the delivery of a regulated 15.4 watts of power at the output from power sourcing equipment (PSE). PoE+ is another standard defined as IEEE 802.3at, which specifies enhance-

ment for higher power level of 30 watts of power at the output from PSE. This power is utilized by a connected powered device (PD) such as VoIP phones, wireless access points, and IP-based video cameras as shown in Figure 5.5.




Figure 5.5 Powered Devices (PD) Connected to an EX4200 Switch

The ability to deliver power over the same Ethernet LAN cables used to transmit data has eliminated the need to attach PDs to electrical outlets. Additional benefits include simplified device deployment, lower cost of deployment, greater flexibility, and remote management.

The EX2200, EX3200, EX 3300, EX4200, EX4300, EX6200, and EX8200 switches all provide support for PoE/PoE+, wherein the switch acts as the PSE.

NOTE “P” models on EX3200 and EX4200 lines provide support for enhanced PoE (up to 18.6 watts at PSE) with supported Junos releases. “PX” models on EX3200 and E4200 support PoE+.

NOTE The EX4200 switch provides either full or partial PoE on all models (with the exception of the fiber-based EX4200-24F model). The full PoE models provide power on all 24 or 48 ports, while the partial PoE models provide power on first eight ports only.

NOTE PoE is enabled by default on the fixed-configuration EX Series switches that support PoE. You can activate PoE simply by connecting PDs to the powered ports.

Use the following CLI command to configure PoE:

```
user@switch# set poe interface all
```

For PoE management, there are two modes available on the EX Series switches:

- Static mode: as the name suggests, this mode allocates a specified amount of power from the switch's available power budget to the individual interface.
- Class mode: allocates power for interfaces based on the class of PD connected to the port. The amount of power allocated will be the maximum of the class of the PD. Refer to Table 5.2 for each PoE class and corresponding power allocation range.

Table 5.2 PoE Class and Power Allocation

PoE Class	Max Power at Output Port of PSE
0	15.4 watts reserved
1	4 watts
2	7 watts
3	15.4 watts
4 (PoE+ only, PSE type 2)	30.0 watts

ALERT! The default PoE management mode is static. For the EX2200, it is recommended that the mode be changed from static to class. For more information, please refer to *Understanding PoE on EX Series Switches* at <http://www.juniper.net/techpubs/>.

NOTE Although the amount of output power on the PSE is listed in Table 5.2, the actual power received on the PD must take line loss into account. For example, in case of Class 3 PoE, the specified 15.4 watts would need to subtract 16% to account for power loss, which would guarantee 12.95 watts on the PD. IEEE 802.3af compliant PDs require up to 12.95 watts. For Class 4 PoE+, available power at PD is 25.50 watts.

The set poe management class command can be used to change the PoE power management mode:

```
user@switch# set poe management class
```

For the purposes of verifying PoE status on EX Series switches, use the show poe interface command:

```
user@switch> show poe interface
```

Interface	Admin status	Oper status	Max power	Priority	Power consumption	Class
ge-0/0/0	Enabled	ON	15.4W	Low	12.95W	0
ge-0/0/1	Enabled	ON	15.4W	Low	12.95W	0
ge-0/0/2	Enabled	ON	15.4W	Low	12.95W	0
ge-0/0/3	Enabled	ON	15.4W	Low	12.95W	0
ge-0/0/4	Enabled	ON	15.4W	Low	12.95W	0
ge-0/0/5	Enabled	ON	15.4W	Low	12.95W	0
ge-0/0/6	Enabled	ON	15.4W	Low	12.95W	0
ge-0/0/7	Enabled	ON	15.4W	Low	12.95W	0

```
user@switch> show poe interface ge-0/0/0
```

```
PoE interface status:
```

```
PoE interface           : ge-0/0/0
Administrative status   : Enabled
Operational status     : ON
Power limit on the interface : 15.4W
Priority                 : Low
Power consumed          : 12.95W
Class of power device   : 0
```

```
user@switch> show poe controller
```

Controller index	Maximum power	Power consumption	Guard band	Management
0	305 W	0W	0W	Static

Additional methods are available on the EX Series switches to track PoE power consumption and distribution through interfaces:

- EX Series switches can reserve a limited amount of power (maximum 19 watts) for handling a power spike. This can be configured using guard-band:

```
user@switch# set poe guard-band 15
```
- In case of an insufficient PoE power budget for connected PDs, interfaces can be set with a PoE priority of either *high* or *low* so that interfaces designated as high priority would be guaranteed power. In situations where the power budget is limited, low priority interfaces would not be supplied with power in deference to the high priority interfaces.

NOTE It is recommended that you place more business-critical PoE PDs on high-priority interfaces so they continue to be powered in case the switch's power budget drops.

Use the following CLI command to change the PoE priority on an interface:

```
user@switch# set poe interface ge-0/0/0 priority high
```

And per-interface PoE power consumption can be monitored using telemetries:

```
user@switch# set poe interface all telemetries
```

NOTE For information on configuration of additional support of PoE see *Configuring PoE* at <http://www.juniper.net/techpubs>.

Port Mirroring

An Ethernet switch such as the EX4200 normally does not flood out every packet when the destination MAC address is known. However, there are times when it is necessary to receive copies of packets for traffic analysis on interfaces that are different than the originally intended destination interface. Port mirroring can be used to analyze traffic on EX Series Ethernet Switches at Layer 2. It can be used for business and network policy enforcement regarding proper network usage and for identifying problems such as abnormal or excessive bandwidth usage from nodes or applications during troubleshooting.

Port mirroring copies packets from a source to a destination. This source and destination pairing is considered a *session* of port mirroring. Mirrored packets can in turn be analyzed using a protocol analyzer application. The protocol analyzer can be run on a host directly connected to the destination port locally (see Figure 5.6), or on a remotely located monitoring station, which can be on a different Ethernet switch with a VLAN configured as the destination (as in Figure 5.7).

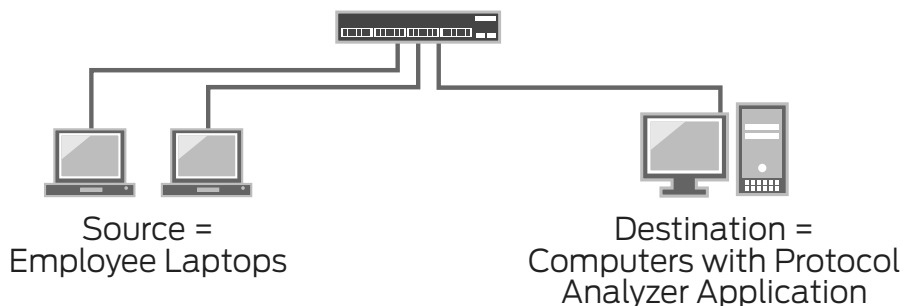


Figure 5.6 Local Port Mirroring




Figure 5.7 Remote Port Mirroring

ALERT! Port mirroring is implemented at the hardware level on EX Series Ethernet switches. As such, the hardware capabilities are different, depending on the EX Series Ethernet switch model. For example, the EX4200 supports one session per system, while the EX8200 supports seven sessions per system. See the *Understanding Port Mirroring on EX Series Switches* at <http://www.juniper.net/techpubs> for detailed guidelines.

There are a number of ways that packets can be mirrored:

- Packets entering (ingress) and/or exiting (egress) the port.
- Multiple ports can also be the source for mirroring session.
- Packets entering (ingress) or exiting (egress) the VLAN.

ALERT! Several limitations must be considered when configuring port mirroring. A source port of the port mirroring session cannot also be a destination port, and the destination port does not participate in Layer 2 protocols such as STP. For more information about these limitations, please see <http://www.juniper.net/techpubs>.

To Configure the Source of Port Mirroring:

1. Set the ingress packets on an interface to become the source of mirroring:

```
user@switch# set ethernet-switching-options analyzer LOCAL-  
MIRROR input ingress interface ge-0/0/0.0
```

2. Set the egress packets on an interface to become the source of mirroring:

```
user@switch# set ethernet-switching-options analyzer LOCAL-  
MIRROR input egress interface ge-0/0/1.0
```

3. Set the ingress packets on a VLAN to become the source of mirroring:

```
user@switch# set ethernet-switching-options analyzer LOCAL-  
MIRROR input ingress vlan Employee_VLAN
```

To configure the Destination of Port Mirroring:

1. Set a port as the destination:

```
user@switch# set ethernet-switching-options analyzer LOCAL-  
MIRROR output interface ge-0/0/10.0
```

To Transport Mirrored Packets to a Remotely Located Monitoring Station that is Running a Protocol Analyzer Application:

1. Set VLAN can be configured as the destination:

```
user@switch# set ethernet-switching-options analyzer REMOTE-  
MIRROR output vlan Mirror_VLAN
```

Configuration of port-mirror session can be verified by using the show analyzer command:

```
user@switch> show analyzer  
Analyzer name           : LOCAL-MIRROR  
Output interface        : ge-0/0/10.0  
Mirror ratio            : 1  
Loss priority           : Low  
Ingress monitored interfaces : ge-0/0/0.0  
Egress monitored interfaces : ge-0/0/1.0
```

The EX Series Ethernet switches support statistical sampling of mirroring. This allows mirroring a packet out of a configured ratio such as 1:x. By default the ratio is 1, which is every packet (1:1 ratio). This can be incremented up to a maximum value of 2047, which would mirror one packet out of every 2047 packets on the given source. To change the mirror ratio from default value (1):

```
user@switch# set ethernet-switching-options analyzer MIRRORING ratio 1000
```

By default, mirrored packets have a loss priority of low, which means mirrored packets would have a lower priority than regular traffic, and in case there is congestion, packets with lower priority are dropped. This setting can be changed to *high* if necessary.

To set the loss priority to high:

```
user@switch# set ethernet-switching-options analyzer MIRRORING loss-priority high
```

In addition, there are often times when specifically selected packets, rather than entire packets, must traverse the mirroring source. The EX Series Ethernet Switches allow policy-based port mirroring where a firewall filter can be configured to select certain packets to be mirrored to the analyzer. For more information on policy-based mirroring using firewall filters, please see <http://www.juniper.net/techpubs>.

sFlow

Network monitoring and traffic flow visibility are important aspects of network device operation. The EX Series Ethernet switches provide support for sFlow monitoring technology, described in RFC 3176, for switched or routed networks. sFlow monitoring consists of an sFlow agent, which is embedded in the switch itself, and the sFlow collector which is typically centralized.

sFlow agent samples network packets at a set pace and sends the samples to a defined collector via UDP. Such pace can be set to be based on either packet (one packet out of a specified number of packets from an interface that has sFlow enabled) or time (interface statistics sample at a specified interval from an interface that has sFlow enabled).

Each datagram sent from the sFlow agent to the collector consists of IP address of the sFlow agent, number of samples, interface through which the packets traversed (ingress and egress), and source and destination interface/VLAN for the packets.

EX Series Ethernet switches implemented the distributed sFlow architecture. The sFlow contains two separate sampling sub-agents, with its unique ID for data source identification, with each Packet Forwarding Engine. A sub-agent follows its own state independently and forwards the sample messages to the sFlow agent. The sFlow agent in turn gathers the samples into datagrams to be sent to the collector. This allows the protocol overhead to be reduced significantly at the centralized collector.

NOTE Only raw packet headers are sampled by sFlow on the switches. A raw packet is the entire Layer 2 Ethernet frame.

To Configure sFlow:

1. Specify the collector IP address and UDP port to be used (up to four collectors can be configured):

```
user@switch# set protocols sflow collector <ip-address> udp-port <port-number>
```

2. Enable sFlow on desired interface (using ge-0/0/19 in this example):

```
user@switch# set protocols sflow interfaces ge-0/0/19
```

NOTE sFlow cannot be enabled on Layer 3 VLAN-tagged interfaces or LAG interfaces such as ae0, but the individual member interfaces that belong to the LAG can be configured for sFlow.

3. Specify the sFlow agent polling interval (0 means disable):

```
user@switch# set protocols sflow interfaces ge-0/0/19 polling-interval <seconds>
```

4. Specify the ingress/egress sampling rate (can also be configured at interface level):

```
user@switch# set protocols sflow sample-rate egress <number>
user@switch# set protocols sflow sample-rate ingress <number>
```

5. Specify agent ID for sFlow agent:

```
user@switch# set protocols sflow agent-id <ip-address>
```

NOTE The agent IP address information is essential for the sFlow collector to determine the source of the sFlow information. It is generally recommended that you configure the agent IP address for consistency. However, if it is not specified, then an IP Address is automatically assigned to the agent by using the IP address of Virtual management Ethernet (VME) interface, or if not present, the management Ethernet interface. If neither of the interfaces are configured with IP address, then any Layer 3 interface or routed VLAN interface (RVI) IP address is used.

6. Specify the source IP address for sFlow datagrams to be sent to the collector:

```
user@switch# set protocols sflow source-ip <ip-address>
```

To Review How sFlow is Configured:

sFlow configuration and system status can be viewed by the following CLI command:

```
user@switch> show sflow
sFlow                : Enabled
Sample limit         : 300 packets/second
Polling interval     : 30 second
Sample rate egress   : 1:1000: Enabled
Sample rate ingress  : 1:1000: Enabled
Agent ID             : 10.0.0.1
Source IP address    : 10.0.0.1
```

NOTE The sample limit of 300 packets/second is defined by the switch and is not user-configurable.

sFlow can be viewed at interface level for interface specific details:

```
user@switch> show sflow interface
```

Interface	Status		Sample rate		Adapted sample rate		Polling interval
	Egress	Ingress	Egress	Ingress	Egress	Ingress	
ge-0/0/0.0	Enabled	Enabled	1000	1000	1000	1000	30
ge-0/0/19.0	Enabled	Enabled	1000	1000	1000	1000	30

sFlow collector information can be viewed by the following CLI command:

```
user@switch> show sflow collector
```

Collector address	Udp-port	No. of samples
10.1.1.100	6343	10000000

MORE? For more details on sFlow including its adaptive sampling using binary backup algorithm to reduce the load on the system, please reference *Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch* at <http://www.juniper.net/techpubs>.

