

# Configuring Juniper Networks Routers

---

Revision 7.a

*Student Guide  
Volume 1*

**Juniper Networks, Inc.**  
1194 North Mathilda Avenue  
Sunnyvale, CA 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

Course Number: EDU-JUN-CJNR-M

Juniper Networks, the Juniper Networks logo, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOS 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.

*Configuring Juniper Networks Routers Volume 1 Student Guide, Revision 7.a*

Copyright © 2006, Juniper Networks, Inc.

All rights reserved. Printed in USA.

**Revision History**

Revision 6.a—May 2004

Revision 6.b—April 2005

Revision 7.a—September 2006

The information in this document is current as of the date listed above.

The information in this document has been carefully verified and is believed to be accurate for software Release 7.6. Juniper Networks assumes no responsibilities for any inaccuracies that may appear in this document. In no event will Juniper Networks be liable for direct, indirect, special, exemplary, incidental or consequential damages resulting from any defect or omission in this document, even if advised of the possibility of such damages.

The information in this document is current as of the date listed above.

The information in this document has been carefully verified and is believed to be accurate for software Release 6.3. Juniper Networks assumes no responsibilities for any inaccuracies that may appear in this document. In no event will Juniper Networks be liable for direct, indirect, special, exemplary, incidental or consequential damages resulting from any defect or omission in this document, even if advised of the possibility of such damages.

Juniper Networks reserves the right to change, modify, transfer or otherwise revise this publication without notice.

**YEAR 2000 NOTICE**

Juniper Networks hardware and software products do not suffer from Year 2000 problems and hence are Year 2000 compliant. The JUNOS software has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

**SOFTWARE LICENSE**

The terms and conditions for using this software are described in the software license contained in the acknowledgment to your purchase order or, to the extent applicable, to any reseller agreement or end-user purchase agreement executed between you and Juniper Networks. By using this software, you indicate that you understand and agree to be bound by those terms and conditions.

Generally speaking, the software license restricts the manner in which you are permitted to use the software and may contain prohibitions against certain uses. The software license may state conditions under which the license is automatically terminated. You should consult the license for further details.

# Contents

---

- Module 1: JUNOS Software CLI Review**
- Module 2: Initial System Configuration**
- Module 3: Protocol Independent Routing Properties**
- Module 4: RIP**
- Module 5: Routing Policy**
- Module 6: OSPF Operation, Configuration, and Troubleshooting**
- Module 7: IS-IS Operation, Configuration, and Troubleshooting**
- Module 8: BGP Operation, Configuration, and Troubleshooting**
- Appendix A: Overview of M-series and T-series Hardware**
- Appendix B: Platform Troubleshooting**
- Appendix C: Additional Features**

---

## Course Overview

---

*Configuring Juniper Networks Routers (CJNR-M) Volume 1* is an instructor-led course that covers the configuration of system and protocol-related features for the Juniper Networks M-series and T-series platforms. This class is a combination of lecture and lab to allow ample time for *hands-on* exposure to the JUNOS software configuration-mode and operational-mode analysis of protocol operation.

### Objectives

After successfully completing this volume, you will be able to:

- Perform initial configuration on M-series and T-series platforms;
- Describe the basic functionality of the RIP routing protocol and how to configure it on a Juniper Networks M-series or T-series platform;
- Use the routing policy within JUNOS software to control routes within the routing and forwarding tables;
- Describe the basic functionality of the OSPF routing protocol and how to configure it on a Juniper Networks M-series or T-series platform;
- Describe the basic functionality of the IS-IS routing protocol and how to configure it on a Juniper Networks M-series or T-series platform; and
- Describe the basic functionality of the BGP routing protocol and how to configure it on a Juniper Networks M-series or T-series platform.

### Intended Audience

The primary audiences for this course include the following:

- Personnel who are unfamiliar with Juniper Networks M-series and T-series platform configuration;
- Internet engineers; and
- Network operations center engineers.

The secondary audiences for this course include the following:

- Juniper Networks and partner sales representatives;
- Juniper Networks and partner systems engineers; and
- Juniper Networks employees (such as hardware engineers, software engineers, and TAC engineers).

### Course Level

*Configuring Juniper Networks Routers Volume 1* is an intermediate-level course designed to provide an in-depth exposure to general configuration and routing protocol operation to prepare students for the more advanced courses available in the Juniper Networks training curriculum.



## Prerequisites

The CJNR-M Volume 1 prerequisites are:

- A basic familiarity with the JUNOS software CLI
- TCP/IP basics;
- Link-state routing protocols, including a familiarity with either OSPF or IS-IS;
- Operation of BGP4 and knowledge of the BGP4 mandatory attributes; and
- General interdomain routing issues.

While not required, familiarity with some type of UNIX system is helpful.

## Course Agenda

---

### Day 1

**Module 0: Introduction and Overview**

**Module 1: JUNOS Software CLI Review**

Gaining Access to the CLI  
CLI Modes and Feature Overview  
Configuration Mode

**Module 2: Initial Configuration**

Getting the Router up and Running  
Configuration Groups  
Software Upgrades and Installation  
Configuring Interfaces

### Day 2

**Module 3: Protocol Independent Routing Properties**

Static Routes  
Aggregate Routes  
Generated Routes  
JUNOS Software Routing Tables  
Route Preferences and Active Route Selection

**Module 4: RIP**

Routing Information Protocol Overview  
RIP Version 2 Message Types and Limitations  
Configuring and Monitoring RIP

**Module 5: Routing Policy**

Policy Overview  
Match Conditions  
Applying Policy  
Route Filters

**Module 6: OSPF Operation, Configuration, and Troubleshooting**

OSPF Overview  
OSPF Scalability  
Adjacency Formation and DR Election  
Configuring OSPF  
OSPF Monitoring and Troubleshooting

**Day 3**

**Module 7: IS-IS Operation, Configuration, and Troubleshooting**

IS-IS Overview  
IS-IS PDUs  
Monitoring IS-IS Operation

**Module 8: BGP Operation, Configuration, and Troubleshooting**

BGP Overview  
BGP Route Selection  
BGP Attributes  
Configuring and Monitoring BGP in JUNOS Software

## Document Conventions

---

The following table lists the syntax-related style conventions used throughout this document:

Style	Description	Usage Example
Arial	Lab instructions and descriptive text.	If told to do so by your instructor, enter the following commands to restore the factory default configuration.
Courier New	Operational displays and noncommand-related syntax.	<code>commit complete</code> Exiting configuration mode
Courier New bold	Command syntax is displayed in bold to differentiate commands from descriptive text.  Please note that the Courier New bold style can be combined with other styles as needed, for example, to indicate a command that involves the use of a locally defined named variable.	<code>erx1:isp-1#configure terminal</code> Or The user can display interface status with the <code>show interfaces</code> command and may make use of the extensive switch, as needed, to obtain additional information.
<i>Courier New italic</i>	Predefined syntax variables such as named policies or passwords.	You will now apply the <i>ospf-test-policy</i> to the OSPF routing instance as an export policy.
<u><i>Courier New italic underline</i></u>	A syntax variable that the reader is expected to define locally.	You will now apply your <u><i>ospf-export-policy</i></u> to the OSPF routing instance as an export policy.

## Additional Information

---

### Education Services Offerings

You can obtain information on the latest Education Services offerings, course dates, and class locations from the World Wide Web by pointing your Web browser to:  
<http://www.juniper.net/training/>.

### About this Publication

The *Configuring Juniper Networks Routers Volume 1 Student Guide* was developed and tested using JUNOS software Release 7.6. Previous and later versions of software may behave differently so you should always consult the documentation and release notes for the version of code you are running before reporting errors.

This document is written and maintained by the Juniper Networks Education Services development team. Please send questions and suggestions for improvement to [training@juniper.net](mailto:training@juniper.net).

### Technical Publications

You can print technical manuals and release notes directly from the Internet in a variety of formats:

1. Go to <http://www.juniper.net/techpubs/>.
2. Locate the specific software or hardware release and title you need, and choose the format in which you want to view or print the document.

Documentation sets and CDs are available through your local Juniper Networks sales office or account representative.

### Juniper Networks Support

For technical support, contact Juniper Networks at <http://www.juniper.net/customers/support/>, or at 1-888-314-JTAC (within the United States) or 408-745-2121 (from outside the United States).





**Configuring Juniper Networks Routers**

***Module 0: Introduction and Overview***

Copyright © 2006, Juniper Networks, Inc.

CJNR-M 7.a.7.6.1

## Configuring Juniper Networks Routers

### Module Objectives

---

- After successfully completing this module, you will be able to:
  - Get to know one another
  - Identify the objectives, prerequisites, facilities, and materials used during this course
  - Identify additional Juniper Networks courses
  - Describe the Juniper Networks Technical Certification Program (JNTCP)

Copyright © 2006, Juniper Networks, Inc.

#### **This Module Discusses:**

- Objectives and course content information;
- Additional Juniper Networks courses; and
- Juniper Networks Technical Certification Program.



## Configuring Juniper Networks Routers

### Introductions

---

- What is your name?
- Where do you work?
- What is your primary role in your organization?
- What kind of network experience do you have?
- What is the most important thing for you to learn in this training session?



Copyright © 2006, Juniper Networks, Inc.

### Introductions

This slide serves to break the ice by having class attendees introduce themselves and state their reasons for attending the class.

# Configuring Juniper Networks Routers

## Course Contents (Volume 1)

---

- **Volume 1 Contents: Days 1–3**
  - Review of the JUNOS software command-line interface and configuration basics
  - Initial system configuration
  - Protocol-independent properties
  - RIP
  - Routing policy
  - OSPF operation, configuration, and troubleshooting
  - IS-IS operation, configuration, and troubleshooting
  - BGP operation, configuration, and troubleshooting

Copyright © 2006, Juniper Networks, Inc.

### Course Contents

The Configuring Juniper Networks Routers (CJNR-M) course is an intermediate-level, instructor-led course that focuses on protocol operation and configuration on M-series and T-series platforms. Nearly all of this information is applicable to J-series platforms as well, although this course does not discuss the J-Web graphic user interface. This course provides a substantial review of the operational characteristics of common protocols, including OSPF, IS-IS, BGP, RIP, multicast, and MPLS to facilitate effective configuration and operational analysis of these protocols.

The course combines both lecture and labs, with significant time allocated for hands-on experience with JUNOS Internet software configuration, operation, and protocol and platform troubleshooting. The complete five-day CJNR-M class is the recommended way to prepare students for attending the advanced courses in the Juniper Networks training curriculum. Juniper Networks offers the five-day course in two separate volumes, Configuring Juniper Networks Routers—Part One and Configuring Juniper Networks Routers—Part Two, if you are unable to take the recommended five-day course at one time.

The slide lists topics for the first three days of the CJNR-M course.

## Configuring Juniper Networks Routers

### Course Contents (Volume 2)

- Volume 2 Contents: Days 4–5
  - MPLS concepts
  - LSP signaling with RSVP
  - Named paths explicit route objects
  - Firewall filtering
  - Multicast

Copyright © 2006, Juniper Networks, Inc.

### Course Contents

The topics shown on the slide are addressed on days four and five of the CJNR-M course.

## Configuring Juniper Networks Routers

### Prerequisites

---

- Basic familiarity with the JUNOS software CLI
- TCP/IP basics
- Link-state routing protocols, including a familiarity with either OSPF or IS-IS
- Operation of BGP4 and knowledge of the BGP4 mandatory attributes
- General interdomain routing issues

Copyright © 2006, Juniper Networks, Inc.

### Prerequisites


The slide lists the prerequisites for the CJNR-M course. While not absolutely necessary, a basic familiarity with the JUNOS software command-line interface is suggested.

## Configuring Juniper Networks Routers

### Course Administration

---

- **Sign-in sheet**
- **Schedule**
  - Class times
  - Breaks
  - Lunch
- **Break and restroom facilities**
- **Communications**
  - Telephones
  - Cellular phones and pagers
  - Internet access



Copyright © 2006, Juniper Networks, Inc.

#### General Course Administration

This slide documents general aspects of classroom administration.

## Configuring Juniper Networks Routers

### Education Materials

---

- **Available in class:**
  - Lecture material
  - Lab guide
  - Lab equipment
- **Available outside of class:**
  - Online documentation at [www.juniper.net](http://www.juniper.net)
  - Juniper Networks Technical Assistance Center (JTAC)
- **Available through your account representative:**
  - Documentation CD
  - Printed documentation



Copyright © 2006, Juniper Networks, Inc.

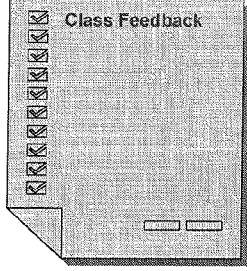

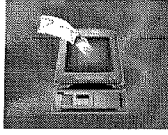
### Training and Study Materials

This slide describes several options for obtaining study and preparation materials.

## Configuring Juniper Networks Routers

### Satisfaction Feedback

---



- **Please be sure to tell us how we did!**
  - You will receive a survey to complete either at the end of class or sent to you via e-mail within two weeks
- **Completed surveys:**
  - Help us serve you better
  - Ensure that you receive a certificate of completion

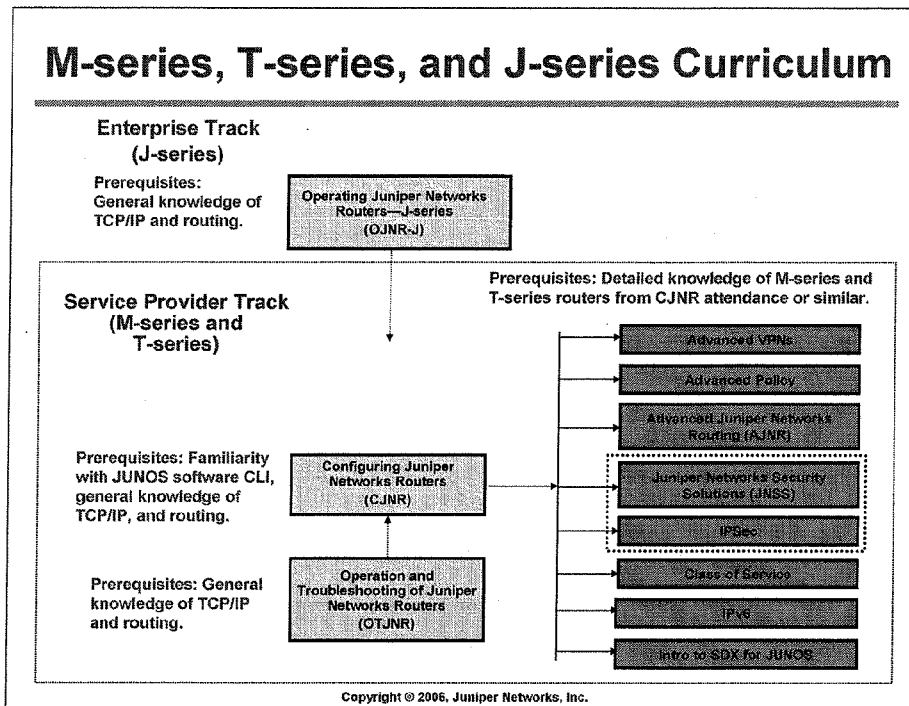
Copyright © 2008, Juniper Networks, Inc.

### Satisfaction Feedback

Juniper Networks uses an electronic survey system to collect and analyze your comments and feedback. Depending on the class you are taking, please complete the survey at the end of the class, or be sure to look for an e-mail about two weeks from class completion that directs you to complete an on-line survey form (be sure to provide us with your current e-mail address).

Submitting your feedback entitles you to a certificate of class completion. We thank you in advance for taking the time to help us improve our educational offerings.

# Configuring Juniper Networks Routers

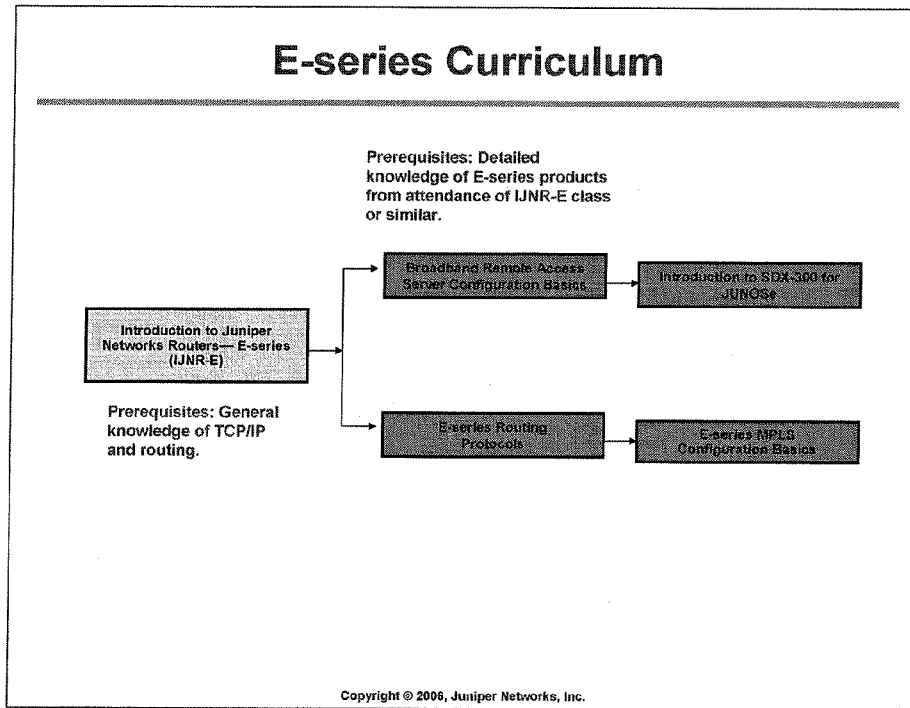


## M-series, T-series, and J-series Curriculum

This graphic displays the primary Education Services offerings that support Juniper Networks M-series, T-series, and J-series router technologies.



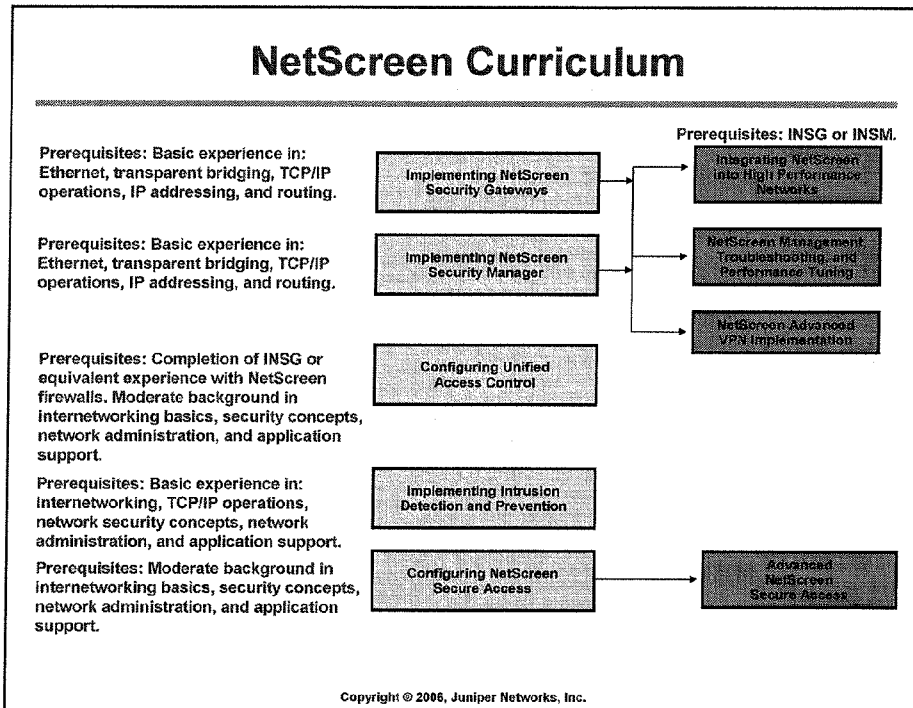
## Configuring Juniper Networks Routers



### E-series Curriculum

This graphic displays the primary Education Services offerings that support Juniper Networks E-series router technologies.

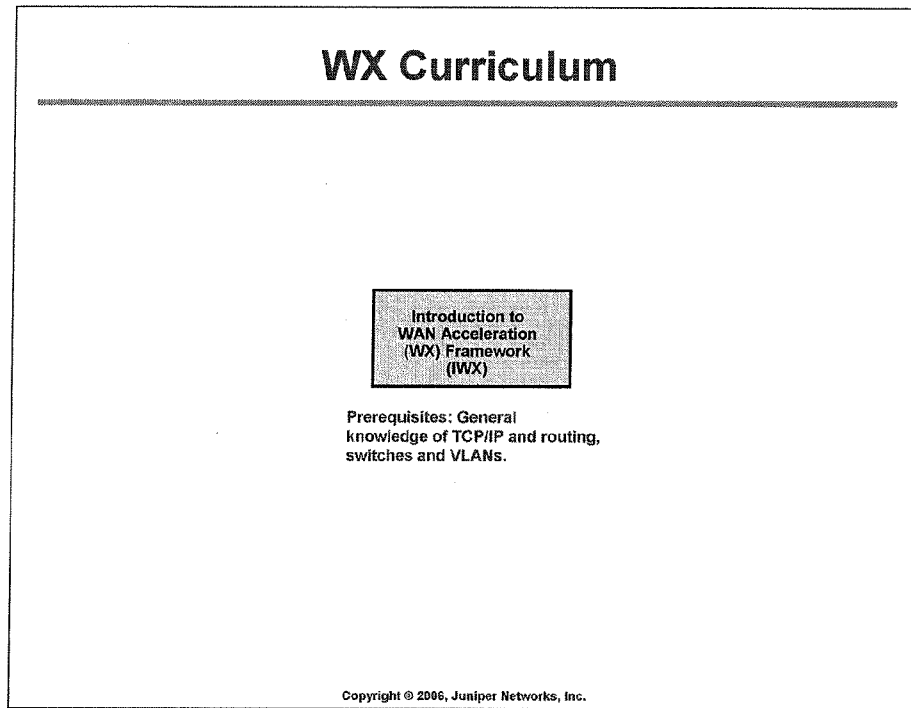
## Configuring Juniper Networks Routers



### NetScreen Curriculum

This graphic displays the primary Education Services offerings that support Juniper Networks NetScreen security technologies.

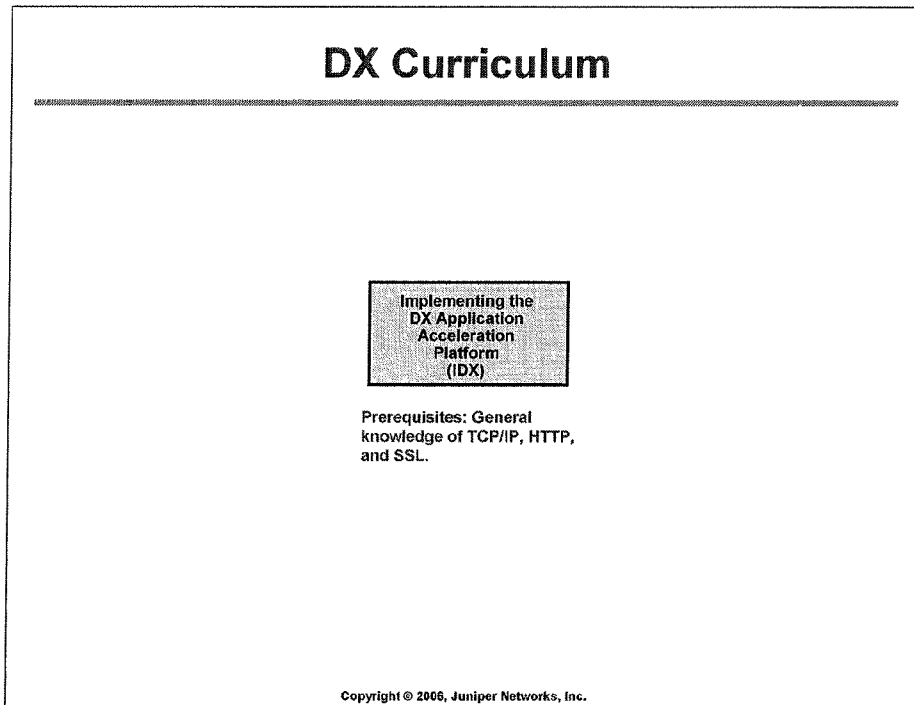
## Configuring Juniper Networks Routers



### **WX Curriculum**

This graphic displays the primary Education Services offerings that support Juniper Networks WX Framework technologies.

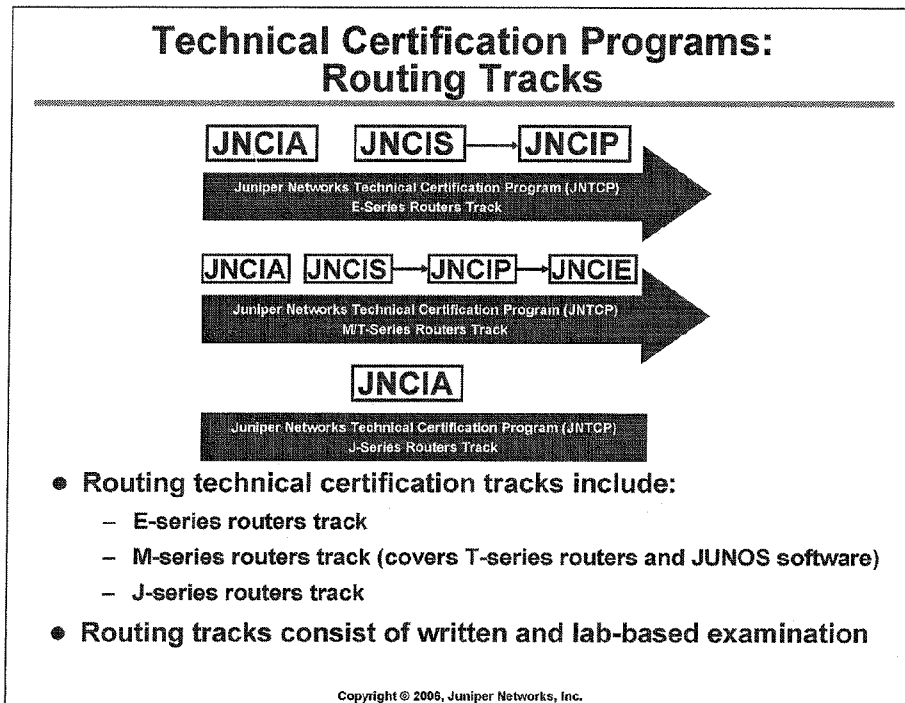
## Configuring Juniper Networks Routers



### **DX Curriculum**

This graphic displays the primary Education Services offerings that support Juniper Networks DX Application Acceleration Platform technologies.

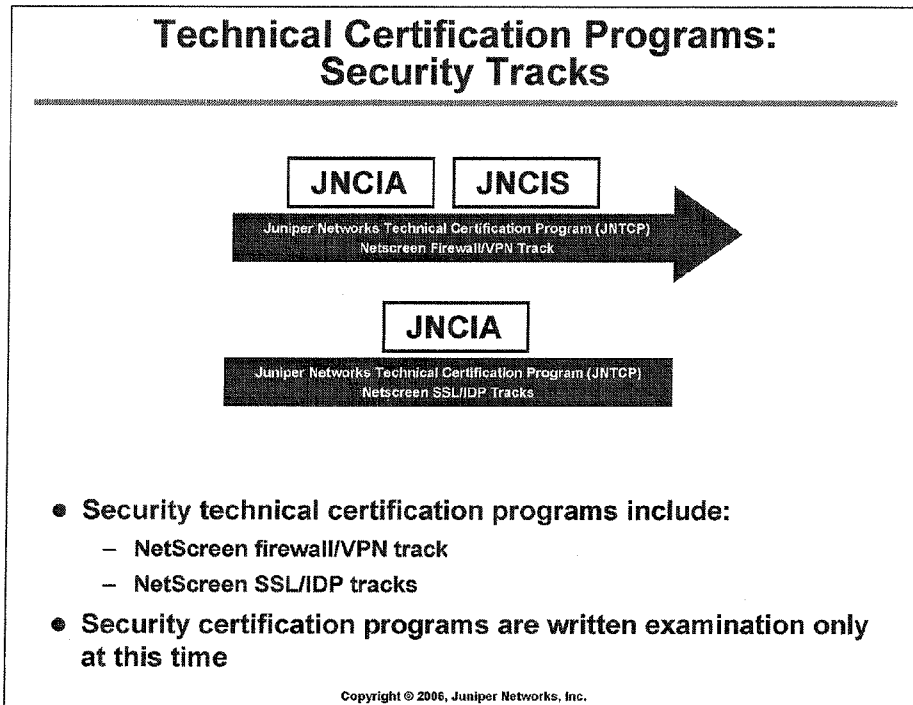
## Configuring Juniper Networks Routers



### Technical Certification Programs: Routing Tracks

This slide outlines the current levels of technical certification offered by Juniper Networks.

## Configuring Juniper Networks Routers



### Technical Certification Programs: Security Tracks

This slide outlines the current levels of technical certification offered by Juniper Networks.

## Configuring Juniper Networks Routers

### Juniper Networks Certified Internet Associate

---

- **JNCIA**

- Computer-based, written exam
- Delivered at Prometric testing centers worldwide
- 60 questions, 60 minutes
- Passing Score: 70%
- \$125 USD
- Prerequisite certification: none
- Benefits provided to JNCIAs:
  - Certificate
  - Logo usage
  - Industry recognition
- Validates candidate's general knowledge of IP technologies, platform operating system, and hardware



Copyright © 2006, Juniper Networks, Inc.

### The JNCIA Certification

This slide details the JNCIA certification level.

## Configuring Juniper Networks Routers

### Juniper Networks Certified Internet Specialist

---

- JNCIS

- Computer-based, written exam
- Delivered at Prometric testing centers worldwide
- Prerequisite for the JNCIP lab exam
- 75 questions, 90 minutes
- Passing Score: 70%
- \$125 USD
- Prerequisite certification: none
- Benefits provided to JNCISs:
  - Certificate
  - Logo usage
  - Provides ability to take JNCIP exam
  - Industry recognition as an IP and routing platform specialist
- Validates candidate's advanced knowledge of platform operating system, hardware, and IP technologies



Copyright © 2006, Juniper Networks, Inc.

### The JNCIS Certification

This slide details the JNCIS certification level.



## Configuring Juniper Networks Routers

### Juniper Networks Certified Internet Professional

---

- **JNCIP**

- One-day, lab-based exam
- Tests candidate's configuration and design skills for essential technologies
- Testing centers: Sunnyvale, Amsterdam, Herndon, Westford, Remote
- Prerequisite for the JNCIE lab exam
- \$1,250 USD
- Prerequisite certification: JNCIS
- Benefits provided to JNCIPs:
  - Certificate
  - Logo usage
  - Provides ability to take JNCIE exam
  - Industry recognition as an IP and routing platform professional
- Validates candidate's practical platform configuration skills



Copyright © 2006, Juniper Networks, Inc.

### The JNCIP Certification

This slide details the JNCIP certification level.

## Configuring Juniper Networks Routers

### Juniper Networks Certified Internet Expert

---

- JNCIE

- One-day, lab-based exam
- Tests candidate's advanced configuration and design skills for essential and specialized technologies
- Testing centers: Sunnyvale, Amsterdam, Herndon, Remote
- \$1,250 USD
- Prerequisite certification: JNCIP
- Currently only available in the M-series routers track
- Benefits provided to JNCIEs:
  - Crystal plaque and certificate
  - Logo usage
  - Worldwide recognition as an Internet Expert
- The most challenging and respected exam of its type in the industry



Copyright © 2006, Juniper Networks, Inc.



### The JNCIE Certification

This slide details the JNCIE certification level.

## Configuring Juniper Networks Routers

### Certification Preparation

- **Training and study resources**
  - JNTCP Website
    - [www.juniper.net/certification](http://www.juniper.net/certification)
  - Education Services training classes
    - <http://www.juniper.net/training>
  - Juniper Networks documents and white papers
    - <http://www.juniper.net/techpubs/>
    - <http://www.juniper.net/techcenter/>
  - Sybex JNTCP preparation guides
    - JNCIA and JNCIP available in bookstores now
  - *Juniper Networks Routers: The Complete Reference*
    - Available in bookstores now
    - Covers M-series and T-series platforms
- **Practical exams: Lots of hands-on practice**
  - On-the-job experience
  - Education Services training classes
  - Equipment access




Copyright © 2008, Juniper Networks, Inc.

### Prepping and Studying

This slide lists some options for those interested in prepping for Juniper Networks certification.

## Configuring Juniper Networks Routers



### **Any Questions?**

If you have any questions or concerns about the class you are attending, we suggest that you voice them now so that your instructor can best address your needs during class.



**Configuring Juniper Networks Routers**

***Module 1: JUNOS Software CLI Review***

Copyright © 2006, Juniper Networks, Inc.

CJNR-M-7.a.7.6.1

### Module Objectives

---

- **After successfully completing this module, you will be able to:**
  - **Log in to a Juniper Networks J-series, M-series or T-series router**
  - **Issue operational-mode commands**
  - **Enter the configuration mode**
  - **Navigate the candidate configuration**
  - **Modify the candidate configuration**
  - **Commit a new active configuration**
  - **Compare configuration files**
  - **Save and manipulate configuration files**
  - **Run operational-mode commands while in configuration mode**

Copyright © 2006, Juniper Networks, Inc.

#### **This Module Discusses:**

- Logging in to a Juniper Networks J-series, M-series or T-series platform;
- Operational-mode commands;
- Navigating the configuration hierarchy;
- Committing a new configuration;
- Comparing configuration files;
- Saving and manipulating configuration files; and
- Running operational-mode commands while in the configuration mode.

## **Agenda: CLI Overview**

---

- **Gaining Access to the CLI**
  - CLI Modes and Feature Overview
  - Configuration Mode

Copyright © 2006, Juniper Networks, Inc.

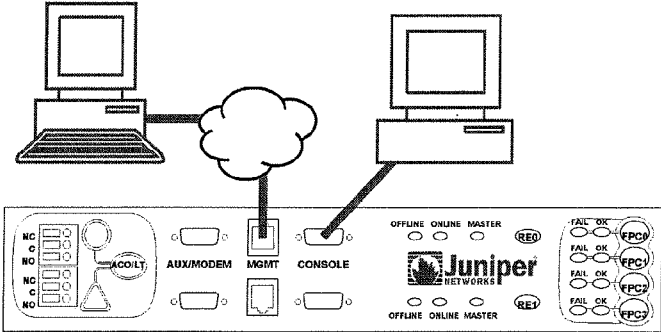
### **Module Agenda**

- ☛ This slide outlines the topics for this chapter and highlights the section we cover first.

## Configuring Juniper Networks Routers

### Access Router's Management Ports

- **Console**
  - Db9 EIA-232 @ 9600 Bps, 8/N/1 (preconfigured)
- **Management port, using Telnet, SSH**
  - Requires configuration
- **J-series, M-series and T-series use same interface**



Copyright © 2006, Juniper Networks, Inc.

### Console Port

The console port is the only preconfigured port on the router. You use the console port to access the CLI.

The JUNOS Internet software CLI is the interface to the software that you use whenever you access the router, either from the console or through a remote network connection. The CLI starts automatically when you log in as a nonroot user and provides commands to perform various tasks, including configuring the JUNOS software, and monitoring and troubleshooting the software, network connectivity, and the router hardware.

The CLI is a straightforward command-line interface. You type a command on a single line, and the command is executed when you press the Enter key.

### Alternative Access

You can also access the CLI using the management interface (`fxp0`) or auxiliary port. This access requires configuration, however. Also, Telnet and SSH access is available.

### Equivalent Interface

This course focuses primarily on M-series and T-series platforms. However, nearly all commands that we discuss are also applicable to J-series platforms. JUNOS software includes a Web-based interface that is installed by default on J-series platforms and that you can install as an option on M-series and T-series platforms. Specific differences of the J-series platform and the Web interface are discussed in the Operation of Juniper Networks Routers, J-series (OJNR-J) course. In general, in this course, when you see mention of software features that pertain to M-series and T-series platforms, it is probably safe to assume that they apply to J-series platforms as well.



## Agenda: CLI Overview

---

- Gaining Access to the CLI
- CLI Modes and Feature Overview
- Configuration Mode

Copyright © 2006, Juniper Networks, Inc.

### CLI Modes and Feature Overview

The next set of slides describe the CLI operational mode and provides a general overview of CLI features.

### CLI Modes and Feature Overview

---

- **CLI operational mode**
  - Editing command lines
  - Command completion/history
  - Context-sensitive and documentation-based help
  - UNIX-style pipes
- **CLI configuration mode**
  - Object-oriented hierarchy
  - Configuration groups
  - Jumping between levels
  - Candidate configuration with sanity checking
  - Automatic rollback capability
  - Showing portions of configuration while configuring
  - Running operational-mode commands from within configuration
  - Saving, loading, and deleting configuration files
  - Wildcard deletes

Copyright © 2006, Juniper Networks, Inc.

#### **CLI Operational Mode**

Use the CLI operational mode to monitor and troubleshoot the operation of the router.

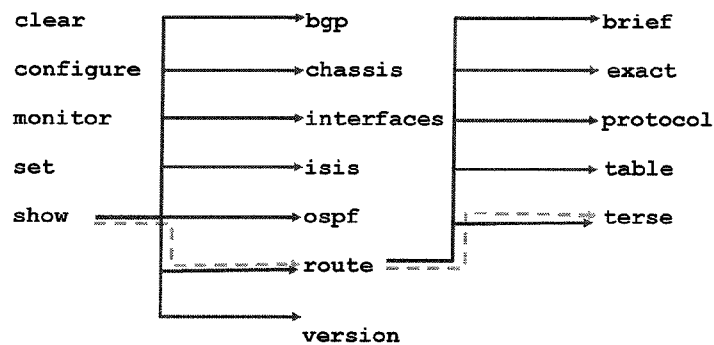
#### **CLI Configuration Mode**

Use the CLI configuration mode when actually modifying the router's configuration..

### CLI Operational Mode

- Commands are executed (mainly) from the default CLI level (`user@host>`)

- Can be executed from configuration mode with the `run` command
- Hierarchy of commands



Copyright © 2006, Juniper Networks, Inc.

### Operational Mode

Operational-mode CLI commands are used to monitor and control the operation of the router. The operational-mode commands are hierarchically structured, as shown on the slide. For example, the `show` command displays various types of information about the system and its environment. One of the possible options for the `show` command is `route`, which displays information about routing tables. Adding the `terse` switch, as in `show route terse`, results in a reduction of output in line with a terse reply.

Key operational-mode capabilities include:

- Entering configuration mode;
- Controlling the CLI environment;
- Exiting the CLI;
- Monitoring and troubleshooting:
  - `clear`
  - `monitor`
  - `ping`
  - `show`
  - `test`
  - `traceroute`;
- Connecting to other network systems;
- Copying files;
- Restarting software processes; and
- Performing system-level operations.

## Configuring Juniper Networks Routers

### Editing Command Lines

---

- **EMACS-style editing sequences are supported**

```
lab@omaha> show interfaces ▲
Ctrl-b
lab@omaha> show interfaces ▲
Ctrl-a
lab@omaha> show interfaces ▲
Ctrl-f
lab@omaha> show interfaces ▲
Ctrl-e
lab@omaha> show interfaces ▲
```

Keyboard sequence →

Cursor position

- **Configure/set a VT-100 terminal type to use arrow keys in addition to EMACS-based control sequences**

Copyright © 2006, Juniper Networks, Inc.

### EMACS-Style Control Keys

The CLI supports EMACS style keyboard sequences that allow you to move around on a command line and delete specific characters or words. The following sequences are supported:

- Ctrl-B: Moves cursor left one character
- Ctrl-A: Moves cursor to the beginning of the command line
- Ctrl-F: Moves cursor right one character
- Ctrl-E: Moves cursor to the end of the command line
- Delete/BS: Deletes character before cursor
- Ctrl-D: Deletes character over the cursor
- Ctrl-K: Deletes from cursor to end of line
- Ctrl-U: Deletes all characters/negates current command
- Ctrl-W: Deletes entire word to left of cursor
- Ctrl-L: Redraws the current line
- Ctrl-P/Ctrl-N: Repeats previous and next command in command history

### Configure/Set a VT-100 Terminal Type

You can enable the use of both EMACS sequences *and* your keyboard's arrow keys by either setting a per-session VT-100 terminal type or by configuring the system's console port to run as a VT-100 terminal type:

Per CLI session (no configuration):

```
user@host> set cli terminal vt100
```

Configure the system console port:

```
[edit system ports]
```

```
lab@host# set console type vt100
```

### Command Completion

#### ●Space bar completes a command

```
lab@HongKong> sh<space>ow i<space>
                ^'i' is ambiguous.
Possible completions:
  igmp           Show Internet Group Management Protocol information
  ike            Show Internet Key Exchange information
  ilmi           Show interim local management interface information
  interfaces     Show interface information
  ipsec          Show IP Security information
  ipv6           Show IP version 6 information
  isis           Show Intermediate System-to-Intermediate System
  information
lab@HongKong> show i
```

#### ●Tab key completes a variable

Copyright © 2006, Juniper Networks, Inc.

#### Space Completion

The CLI provides a completion function. Therefore, you do not always have to type the full command or command option name for the CLI to recognize it.

To complete a command or option that you have typed partially, press the Space bar. If the partially typed letters begin a string that uniquely identifies a command, the CLI displays the complete command name. Otherwise, the CLI beeps to indicate that you have entered an ambiguous command, and it displays the possible completions.

The command completion option is on by default, but you can turn it off.

#### Tab Completion

You can also use the Tab key to complete variables. Examples of variables include policy names, AS paths, community names, and IP addresses.

## Configuring Juniper Networks Routers

### CLI Modes

---

- **Operational mode**
  - Monitor and troubleshoot the software, network connectivity, and router hardware

```
lab@host>
```

The > character identifies operational mode
  
- **Configuration mode**
  - Configure the router, including interfaces, general routing information, routing protocols, user access, and system hardware properties

```
[edit]  
lab@host#
```

The # character identifies configuration mode

Copyright © 2006, Juniper Networks, Inc.

### Operational Mode

In operational mode, you use the CLI to monitor and troubleshoot the router. The `monitor`, `ping`, `show`, `test`, and `traceroute` commands let you display information and statistics about the software running on the router, such as routing table entries, and let you test network connectivity.

### Configuration Mode

You configure JUNOS software by entering configuration mode and creating a hierarchy of configuration statements. You can configure all properties of JUNOS software, including interfaces, general routing information, routing protocols, and user access, as well as several system hardware properties.

## Context-Sensitive Help

Type a question mark (?) anywhere on command line

```
lab@host> ?  
Possible completions:  
clear          Clear information in the system  
configure      Manipulate software configuration information  
file           Perform file operations  
help           Provide help information  
...  
lab@host> clear ?  
Possible completions:  
arp            Clear address resolution information  
bfd            Clear Bidirectional Forwarding Detection  
               information  
bgp            Clear Border Gateway Protocol information  
cli            Clear command-line interface settings  
firewall       Clear firewall counters  
...
```

Copyright © 2006, Juniper Networks, Inc.

### Need Help?

The CLI provides context-sensitive help at any point in a command line. Help tells you which options are acceptable at the current point in the command and provides a brief description of each command or command option.

To get help at any time while in the Juniper Networks CLI, type a question mark (?). You do not need to press Enter. If you type the question mark at the command-line prompt, the CLI lists the available commands and options. If you type the question mark after entering the complete name of a command or an option, the CLI lists the available commands and options and then redisplay the command name and options that you typed. If you type the question mark in the middle of a command name, the CLI lists possible command completions that match the letters you have entered so far, then redisplay the letters that you typed.

## Configuring Juniper Networks Routers

### Topical Help

#### The `help topic` command provides information on concepts

```
lab@Sydney> help topic groups ?
```

Possible completions:

<code>apply-groups</code>	Specify where configuration group is inherited
<code>apply-groups-except</code>	Specify where configuration group is not inherited
<code>display-inheritance</code>	Show statements inherited from configuration group
<code>examples</code>	Overview of configuration group examples
<code>groups</code>	Create configuration group
<code>junos-defaults</code>	Default junos-defaults configuration group
<code>overview</code>	Configuration groups overview
<code>wildcards</code>	Wildcard use in configuration groups

```
lab@host> help topic icmp lifetime
```

Using Wildcards with Configuration Groups

You can use wildcards to identify names and allow one statement to provide data for a variety of statements. For example, grouping the configuration of the `sonet-options` statement over all SONET/SDH interfaces or the `dead interval` for Open Shortest Path First (OSPF) over all Asynchronous

...

Copyright © 2006, Juniper Networks, Inc.

### Help on General Concepts

There are various ways to use the `help` command. The `help topic` command displays usage guidelines for the statement. In the example on the slide, we are receiving information on using wildcards with configuration groups.



## Getting Help on Configuration Syntax

### The `help reference` command provides configuration-related information

```
lab@host> help reference groups apply-groups  
apply-groups
```

#### Syntax

```
apply-groups [ group-names ];
```

#### Hierarchy Level

All hierarchy levels

#### Release Information

Statement introduced before JUNOS Release 7.4.

#### Description

Apply a configuration group to a specific hierarchy level in a configuration, to have a configuration inherit the statements in the configuration group.

You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes

Copyright © 2006, Juniper Networks, Inc.

### Help on JUNOS Software Configuration

- The `help reference` command displays summary information for the statement. In other words, it contains JUNOS software-specific, configuration-related information. In the example on the slide we are using the `help` command for information on the `apply-groups` configuration command. Notice the difference between the `help reference` command shown here and the `help topic` command from the previous slide.

### Using | (Pipe)

- **The pipe function is used to filter output**

- Available in all modes and context

```
user@host> show route | ?
Possible completions:
count          Count occurrences
display        Display additional information
except         Show only text that does not match a pattern
find           Search for the first occurrence of a pattern
hold           Hold text without exiting the --More-- prompt
last           Display the last screen of lines in the output
match          Show only text that matches a pattern
no-more        Don't paginate output
request        Make system-level requests
resolve        Resolve IP addresses
save           Save output text to a file
trim           Trim specified number of columns from start of line
```

Copyright © 2006, Juniper Networks, Inc.

### Using Pipe

For operational and configuration commands that display output, such as the `show` commands, you can filter the output. When help is displayed for these commands, one of the options listed is `|`, called a pipe, which allows the command output to be filtered. To filter the output of an operational-mode or a configuration-mode command, add a pipe and option to the end of the command. The options are:

- `compare ( filename | rollback n )`: Available in configuration mode only using the `show` command. Compares configuration changes with another configuration file.
- `count`: Displays the number of lines in the output.
- `display (changed | commit-scripts | detail | inheritance | set | xml )`: Available in configuration mode only. Displays configuration in various formats. For example, `set` displays the set commands required to produce the displayed configuration, and `inheritance` displays the configuration after inheriting from configuration groups.
- `except regular-expression`: Ignores a text matching a regular expression when searching the output. If the regular expression contains spaces, operators, or wildcard characters, you must enclose it in quotation marks.
- `find regular-expression`: Displays the output starting at the first occurrence of text matching a regular expression. If the regular expression contains spaces, operators, or wildcard characters, you must enclose it in quotation marks.

*Continued on next page.*

## Configuring Juniper Networks Routers

### The Pipe Commands (contd.)

- hold: Holds text without exiting the --(more)-- prompt.
- last: Displays the last screen of information.
- match regular-expression: Searches for text matching a regular expression. If the regular expression contains spaces, operators, or wildcard characters, you must enclose it in quotation marks.
- no-more: Displays output all at once rather than one screen at a time.
- resolve: Converts IP addresses to DNS names. Truncates to fit original size unless you specify full-names.
- save filename: Saves the output to a file or URL.
- trim: Trims specified number of columns from the start line.

### Key Operational-Mode Commands

---

- Where we are going...
  - Rebooting and shutting down
  - Analyzing log and trace files
  - Miscellaneous log file commands

Copyright © 2008, Juniper Networks, Inc.

### Key Operational-Mode Commands

The slide shows the topics examined in the following section.

### Rebooting and Shutting Down

- You should always gracefully shut down JUNOS software before removing power

- Rebooting the system:

```
user@host> request system reboot ?
Possible completions:
<[Enter]>      Execute this command
at            Time at which to perform the operation
in           Number of minutes to delay before operation
media        Boot media for next boot
message      Message to display to all users
|           Pipe through a command
```

- Shutting down the system:

```
lab@San_Jose-3> request system halt ?
Possible completions:
<[Enter]>      Execute this command
at            Time at which to perform the operation
both-routing-engines  Halt both Routing Engines
in           Number of minutes to delay before operation
media        Boot media for next boot
message      Message to display to all users
|           Pipe through a command
```

Copyright © 2006, Juniper Networks, Inc.

### Graceful Shutdowns and Rebooting

JUNOS software runs on a multitasking, multiuser operating system based on FreeBSD. As with any UNIX system, you should always gracefully shut the system down before removing power. Failing to shut down in a graceful manner can result in file system corruption that, in the best of cases, simply prolongs the next boot, and in the worst of cases, might actually prevent a successful boot.

The `request system reboot` command causes the router to reboot. Reboot requests are recorded to the system log files, which you can view with the `show log` command. The reboot command takes a variety of arguments that you can use to schedule the reboot, generate a system message, or specify the medium for which the router should use to boot. Media options include `compact-flash` and `disk`. The router always boots from removable medium when the medium is installed, and a cold boot is performed.

The `request system halt` command gracefully stops the router software and prepares the router to be shut down. It is critical to note that you must either cycle power or hit the Enter key on the terminal attached to the router's console port to effect a reboot of a router that has executed a shutdown command.

Both the reboot and shutdown command require user confirmation of the action:

```
user@host> request system reboot
Reboot the system ? [yes,no] (no) yes
```

```
*** FINAL System shutdown message from root@host ***
System going down IMMEDIATELY
```

```
shutdown: [pid 15049]
Shutdown NOW!
```

### Analyzing Log and Trace Files

- Log and trace files are stored in `/var/log`

- Use the `show log file-name` command to display contents
  - Hint: Get help on available options at the `more` prompt by entering an `h`
  - Be sure to make use of the CLI's pipe functionality!

```
lab@host> show log messages | match fail
Jan 29 12:40:47 Montreal-3 rpd[2228]: RPD_ISIS_ADJDOWN: IS-IS lost L2
adjacency to Amsterdam-3 on so-0/3/1.0, reason: 3-Way Handshake Failed
```

- Cascade instances of the CLI's pipe function to evoke a logical AND type search:

```
lab@host> show log messages | match so-0/3/1 | match TRAP
Feb 18 18:51:28 Montreal-3 mib2d[2227]: SNMP_TRAP_LINK_DOWN: ifIndex
34, ifAdminStatus up(1), ifOperStatus down(2), ifName so-0/3/1.0
```

- Use quotes and the pipe (`|`) character to evoke a logical OR:

```
show log messages | match "fpc | sfm | kernel | panic"
show log messages | match "-0|-1|-2|-3|-4"
```

Search by message priority or keywords

Copyright © 2006, Juniper Networks, Inc.

### Viewing Logs and Traces

By default, log and trace files are stored in `/var/log`. To view stored log files, use the command `show log`. Recall that the CLI automatically pauses when there is more than one screen's worth of information, and that at this `more` prompt, you can enter a forward slash (`/`) character to conduct a forward search. As a hint, enter `h` when at a `more` prompt for a context help screen of available commands:

```
Jan  7 18:22:40 Parsing config file
---(Help for CLI automore)---
Clear all match and except strings:                c or C
Display all line matching a regexp:                m or M <string>
Display all lines except those matching a regexp:  e or E <string>
Display this help text:                            h
Don't hold in automore at bottom of output:       N
Hold in automore at bottom of output:             H
Move down half display:                            TAB, d, or ^D
Move down one line:                               Enter, j, ^N, ^X, ^Z, or Down-Arrow
. . .
```

Being able to cascade multiple instances of the CLI's pipe functionality is a real benefit when you must search a long file for associated entries. In the example, the match function is cascaded so that only lines containing the words `so-0/3/1` and `TRAP` are displayed; this creates a logical AND type matching function. Being able to search for multiple criteria in a logical OR fashion is extremely handy, especially when you are not quite sure what you are looking for. The slide provides two examples of a logical OR type search. The former is based on human-readable keywords while the latter makes use of explicit message priority codes to display all messages ranging from Level 0 (emergency) to Level 4 (warning). Note that searching by message priority is predicated on enabling syslog priority with the `explicit-priority` keyword.

### Miscellaneous Log File Commands

- Monitor a log/trace in real-time with the CLI's **monitor** command

```
user@host> monitor start filename
```

- Shows updates to monitored file(s) asynchronously
- `monitor start filename | match pattern` filters output
- Use Esc-Q to enable/disable real-time output to screen
- Issue a `monitor stop` command to cease all monitoring

- To stop a tracing operation, delete a trace flag or the entire stanza:

```
[edit protocols bgp traceoptions]
```

```
user@host# delete flag open
```

- Log/trace file manipulation:

- Use the `clear` command to truncate (clear) log/trace files

```
user@host> clear log filename
```

- Use the `file delete` command to delete log/trace files

```
user@host> file delete filename
```

Copyright © 2006, Juniper Networks, Inc.

### Monitoring Logs and Trace Files

Use the `monitor` CLI command to view real-time log information. You can monitor several log files at one time. The messages from each log are identified by filename, where filename is the name of the file from which entries are being displayed. This line is displayed initially and when the CLI switches between log files.

Using Esc-Q enables and disables syslog output to screen; using `monitor stop` ceases all monitoring. Note that you can use the CLI's `match` functionality to monitor a file in real time, while only displaying entries that match your search criteria. To make use of the functionality, use a command in the form of:

```
lab@San_Jose> monitor start messages | match fail
```

```
lab@San_Jose>
```

### Stop Tracing through Configuration

If you do not delete or disable all trace flags, tracing continues in the background, and the output continues to be written to the specified file. The file remains on the Routing Engine hard drive until either it is deleted manually or overwritten according to the `traceoptions` file parameters. To disable all tracing at a particular hierarchy, issue a `delete traceoptions` command at that hierarchy, and commit the changes.

### Log/Trace File Manipulation

To truncate files used for logging, use the `clear log filename` command.

To delete a file, use the `file delete` command. If you want, you can use wildcards with `delete`, `compare`, `copy`, `list`, and `rename` operations.

### Agenda: CLI Overview

---

- Gaining Access to the CLI
- CLI Modes and Feature Overview
- **Configuration Mode**

Copyright © 2006, Juniper Networks, Inc.

### Configuration Mode

The tasks you can perform in configuration mode, which are covered in the following pages, are:

- *Entering configuration:* Type `configure` to enter configuration mode.
- *Moving within the configuration hierarchy:* Use the `edit`, `up`, `top` and `exit` commands to move between levels.
- *Viewing the candidate configuration and identifying differences:* Use `show` commands while in configuration mode along with the `|` operator.
- *Activating the candidate configuration:* Use the `commit` command to activate the configuration.
- *Backing out of configuration changes:* Use the `rollback` command to restore a previous configuration; better yet, try `commit-confirmed`.
- *Manipulating configuration files:* Use the `save` and `load` commands to save and load configuration files in a variety of situations.



### Entering Configuration Mode

- Type `configure` or `edit` at the CLI operational-mode prompt

```
root@lab2> configure
Entering configuration mode
[edit]
root@lab2#
```

- To allow a single user to edit the configuration, type `configure exclusive`
- `configure private` allows the user to edit a private copy of the candidate configuration
  - Multiple users can edit private candidate configurations simultaneously
  - At commit time, the user's private changes are merged back into the global configuration

Copyright © 2006, Juniper Networks, Inc.

### Starting Configuration Mode

You enter configuration mode by issuing the `configure` command or the `edit` command from the CLI operational mode. If, when you enter configuration mode, another user is also in configuration mode, a message indicates who the user is and what portion of the configuration the user is viewing or editing.

In configuration mode, the prompt changes from the angle bracket (>) of operational mode to the pound sign (#), preceded by the name of the user and the name of the router.

The portion of the prompt in brackets, such as `[edit]`, is a banner indicating that you are in configuration mode and specifying your location within the statement hierarchy.

### Exclusive Configuration

By default, multiple users can enter configuration mode and commit changes. To allow only a single user to edit the configuration, use the `configure exclusive` command.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Private Configuration

Entering configuration mode using `configure private` allows multiple users to edit the configuration while only committing their private changes (you must issue a `commit` command from the `[edit]` hierarchy). If private users issue a `rollback 0` command, only their changes are discarded. If two users are in private mode and both make the same change (`user_1` changes the system hostname to `foo` while `user_2` sets the name to `faa`) then the second `commit` will fail with an error message to avoid configuration conflicts. The second user's changes are placed into effect if a second `commit` is issued, however.

When a user is in private mode, other users must enter private mode or use `configure exclusive` to become the master, or they cannot modify the candidate configuration. Exiting private configuration without committing changes results in the loss of any modifications made to the private candidate configuration.

### Configuration Hierarchy

---

- **Create a hierarchy of configuration statements**

- Enter commands in CLI configuration mode

```
root@lab2# set chassis alarm sonet lol red
```

- And the resulting configuration hierarchy is created...

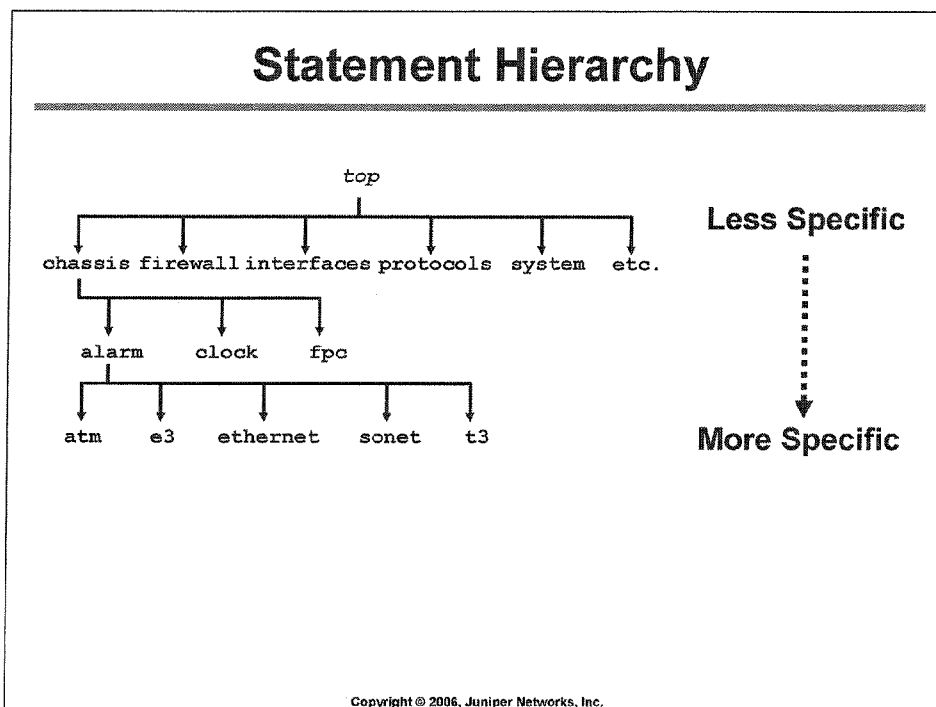
```
chassis {  
  alarm {  
    sonet {  
      lol red;  
    }  
  }  
}
```

Copyright © 2008, Juniper Networks, Inc.

### Enter Commands and Display

To configure the Juniper Networks J-series, M-series, or T-series platforms, including the routing protocols, the router interfaces, network management, and user access, you enter CLI commands in configuration mode. In configuration mode, the CLI provides commands that let you configure the system, load an ASCII text file that contains the system configuration, activate a configuration, and save the configuration to a text file.

## Configuring Juniper Networks Routers



### Statement Hierarchy

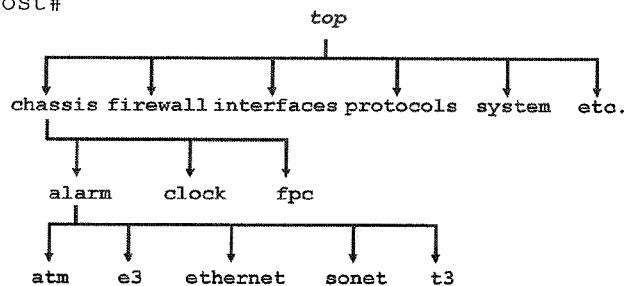
In configuration mode, you enter commands that affect the statement hierarchy. The statement hierarchy stores configuration information and is independent of the CLI operational-mode command hierarchy. The commands available in configuration mode are also independent of the commands available in operational mode. For example, CLI operational mode includes a `show` command to display specific information, while CLI configuration mode provides a `show` command to display the statement hierarchy. The two commands are independent of each other.

The statement hierarchy is organized in a tree structure similar to Windows folders or UNIX directories, grouping related information into a particular branch of the tree.

### Moving between Levels (1 of 2)

- Moving between levels of the statement hierarchy
  - edit functions like a change directory (CD) command

```
[edit]
user@host# edit chassis alarm ethernet
[edit chassis alarm ethernet]
user@host#
```



Copyright © 2008, Juniper Networks, Inc.

### Changing Directories

- To move down through an existing configuration statement hierarchy or to create a hierarchy and move down to that level, use the `edit` command, specifying your desired hierarchy level. After you issue an `edit` command, the configuration mode banner changes to indicate your current level in the hierarchy.

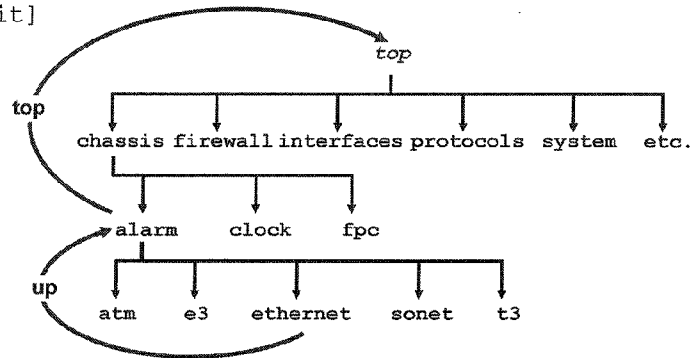
## Configuring Juniper Networks Routers

### Moving between Levels (2 of 2)

```
[edit chassis alarm ethernet]  
user@host# up
```

```
[edit chassis alarm]  
user@host# top
```

```
[edit]
```



Copyright © 2006, Juniper Networks, Inc.

### Level Navigation

To return to your previous location in the statement hierarchy, use the `exit` command. This command is, in effect, the opposite of the `edit` command. Entering `exit` at the top level of the hierarchy exits configuration mode.

To move up in the configuration statement hierarchy one level at a time, use the `up` command. To move to the top of the statement hierarchy from any location, use the `top` command.

### CLI Enhancements

- **Relative configuration commands**

- **Display/edit any portion of the hierarchy**

```
[edit interfaces so-5/1/0 unit 0 family inet]
lab@host# top show system login
class superuser-local {
    permissions all;
}
[edit interfaces so-5/1/0 unit 0 family inet]
lab@host# top edit protocols ospf
[edit protocols ospf]
lab@host#
```

- **Operational-mode show configuration command supports a configuration path**

```
lab@host> show configuration interfaces fxp0
unit 0 {
    family inet {
        address 10.250.0.134/16;
    }
}
```

Copyright © 2006, Juniper Networks, Inc.

#### Relative Configuration Statements

- Starting with JUNOS software Release 5.3, you can enter commands from any level in the hierarchy by issuing the `top` command. As seen on the slide, the use of this command allows you to view every portion of the configuration, regardless of which directory you are located in. It also allows you to change directories without having to jump to the top of the directory. Thus, in the example on the slide, the user went from the `[edit interfaces]` hierarchy to the `[protocols ospf]` hierarchy by simply issuing a single command.

#### Viewing the Configuration in Operational Mode

- Starting in JUNOS software Release 5.3, the `show configuration` command takes a configuration path. Thus, instead of viewing the entire configuration, you can view a portion of the configuration by specifying the configuration hierarchy (previous to Release 5.3, similar functionality could be achieved using pipe commands). We see this feature on the slide, where the user is in operational mode and is viewing only the `fxp0` interface portion of the configuration.

## Configuring Juniper Networks Routers

### Viewing Candidate Configuration

---

```
[edit]
user@host# show chassis alarm
sonet {
    los red;
    pll yellow;
}
[edit]
```

You can display just the portions that concern you from the root of the hierarchy...

```
user@host# edit chassis alarm
[edit chassis alarm]
user@host# show
sonet {
    los red;
    pll yellow;
}
[edit chassis alarm]
```

...or use edit to park yourself at a specific sub-hierarchy

Copyright © 2006, Juniper Networks, Inc.

### Displaying the Configuration

To display the candidate configuration, use the configuration-mode `show` command. This command displays the configuration at the current hierarchy level or at the specified level below the current location.

The `show` command has the following syntax: `show statement-path`. When displaying the configuration, the CLI indents each subordinate hierarchy level, inserts braces to indicate the beginning and end of each hierarchy level, and places a semicolon at the end of statements that are at the lowest level of the hierarchy. The display format is the same format you use when creating an ASCII configuration file, and it is also the same format that the CLI uses when saving a configuration to an ASCII file.

In cases where an empty statement leads to an invalid configuration because it is incomplete or meaningless, the `show` command does not display any of the statement path.



### Identifying Configuration File Differences

- Change the candidate configuration

```
[edit chassis]
user@host# set alarm sonet lol red
[edit chassis]
user@host# delete alarm sonet pll
```

- Display differences between the candidate and active configurations

```
[edit chassis]
user@host# show | compare
[edit chassis alarm sonet]
+ lol red;
- pll yellow;
```

- Compare arbitrary files

```
user@host> file compare files filename 1 filename 2

user@host> show configuration | compare rollback number
```

Copyright © 2006, Juniper Networks, Inc.

### Modifying a Candidate Configuration

The example on the slide modifies a candidate configuration by setting a loss-of-light (LOL) SONET/SDH alarm and removing a phase-locked loop (PLL) alarm that was previously committed.

### Viewing Differences

Piping the output of a `show` command to the CLI `compare` function displays the differences between the candidate configuration file and the active configuration. Starting with JUNOS software Release 5.3, configuration comparison is now *patch*-like. Thus, instead of showing the entire configuration and where changes were made, only the actual changes are shown. By using the pipe switch you can save the configuration differences to file name of your choosing. Once saved, you can issue a `load patch filename` command to merge the contents of the patch file into the candidate configuration where they can be viewed, edited, and ultimately committed.

### Viewing Differences in Other Files

The `compare` function allows you to view differences in any text files, including log files and any of the 49 rollback configurations. The slide shows the `lab` user at an operational-mode prompt displaying the differences between two files in the first syntax example (the resulting output is not shown for the sake of brevity). In the second syntax example, the user is comparing the current configuration to the contents of a given rollback file.

### Removing Statements

- Statements added with `set` are removed with the `delete` command

- Removes everything from the specified hierarchy down
- Use wildcard deletes to save time

```
[edit chassis alarm sonet]
lab@host# show
los red;
```

Note that the final argument (red) is not specified in the delete statement

```
[edit chassis alarm sonet]
lab@host# delete los
```

- Pop Quiz: You have just disabled an interface with a `set interface interface-name disable` statement. How do you re-enable this interface?

Copyright © 2006, Juniper Networks, Inc.

### Removing Configuration Statements

Use the configuration-mode `delete` command to remove statements that were added to the configuration with a `set` command. This command deletes the statement and all its subordinate statements and identifiers. Deleting a statement or an identifier effectively *unconfigures* the functionality associated with that statement or identifier, returning that functionality to its default condition.

You cannot delete the final leaf on a particular configuration hierarchy. For example, trying to delete `red`, which is a leaf to a `set chassis alarms sonet lol` statement, result in an error message.

Consider using the wildcard `delete` function when deleting individual statements is too arduous and deleting an entire configuration sub-hierarchy lacks the granularity that is needed. Sample syntax for a wildcard `delete` is shown:

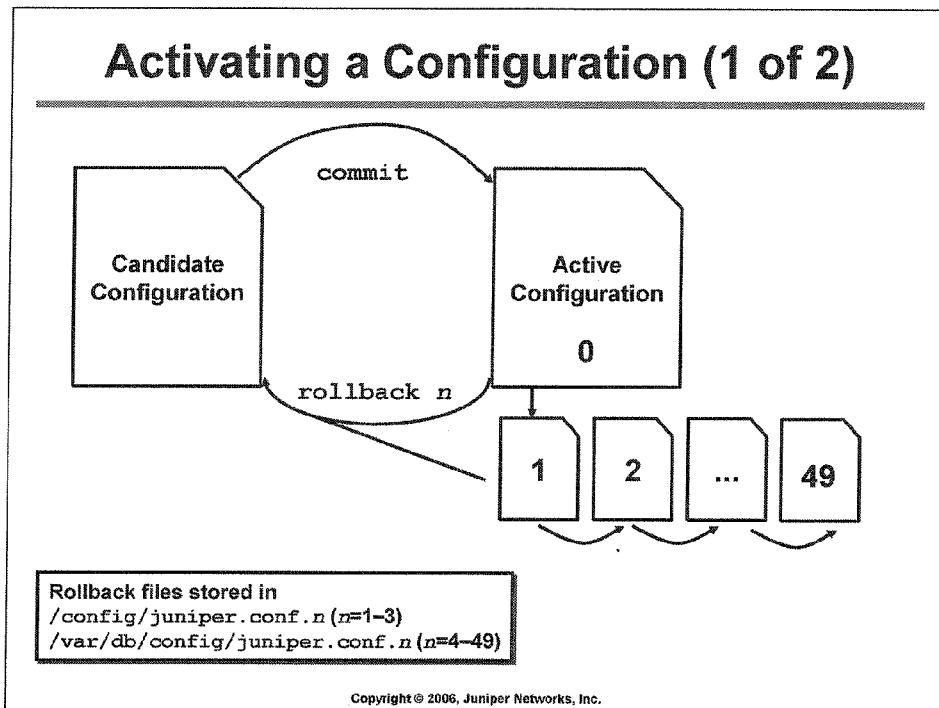
```
[edit]
lab@Sao_Paulo# wildcard delete interfaces fe-*
  matched: fe-0/0/2
Delete 1 objects? [yes,no] (no) yes
```

In addition to deleting configuration statements, you should also consider the use of `deactivate` to cause the specified portion of the configuration hierarchy to be ignored while still retaining the original configuration. Issue an `activate` command to place the configuration back into effect. Also consider the use of `disable` for interfaces. Use the `set` command to add a `disable` statement to flag a given interface as being administratively disabled.

### Pop-Quiz!

Issue a `delete interface interface-name disable` to delete the `disable` statement placed into effect with a `set` command. This syntax has been known to strike some folks as being more than a bit on the double-negative side; then again, these same folks tend to agree that a `no shutdown` statement, as used for similar functionality on other vendor's equipment, is equally counter-intuitive!

## Configuring Juniper Networks Routers



### Active versus Candidate Configuration

When you edit a configuration, you work in a copy of the current configuration to create a *candidate* configuration. The changes you make to the candidate configuration are visible in the CLI immediately, so if multiple users are editing the configuration at the same time, all users can see all changes.

To have a candidate configuration take effect, you must `commit` the changes. At this time, the candidate file is checked for proper syntax, activated, and marked as the current, operational software configuration file. If the syntax is not correct, an error message indicates the location of the error, and no part of the configuration is activated. You must correct the errors before recommitting the configuration. When you commit a configuration (which you can do from any hierarchy level), you commit the entire configuration in its current form. If more than one user is modifying the configuration, committing it saves and activates the changes of all the users. Use the `commit check` command to validate a candidate configuration without actually placing it into effect.

When a new configuration is committed, the old, active configuration is saved as `/config/juniper.conf.1`. The previous configuration can be easily recovered with a `rollback 1` command. Each existing backup is renumbered and pushed further out, storing the oldest copy as number 49.

JUNOS software stores a maximum of 50 previously committed configurations (including the current active configuration known as `rollback 0`). The first three rollbacks (1-3) are stored in the `/config` directory, which resides on the solid-state flash disk. The remainder are stored in the `/var/db/config` directory, which resides on the hard disk.

### Activating a Configuration (2 of 2)

- Remote configuration changes require caution
  - Might disrupt remote connectivity to router
    - Use `commit confirmed` to temporarily activate a configuration (default is 10 minutes)
    - If configuration is not confirmed, router returns to previous configuration automatically; a second `commit` confirms the changes
- Use the `synchronize` switch to mirror the new configuration to a backup RE
- Support for scheduled and commented commits
  - Use the `commit at` time option (Release 5.5)

```
[edit]
user@host# commit at 20:01:00
configuration check succeeds
commit at will be executed at 2003-08-08 20:01:00 UTC
The configuration has been changed but not committed
Exiting configuration mode
```
  - Comments can be added to the `commits` log with the `comment` switch (Release 6.1)

Copyright © 2006, Juniper Networks, Inc.

### Remote Configuration Is Risky

The system never commits a candidate configuration on its own. As previously discussed, when you load or merge a configuration file, you must commit the changes before they can take effect. Special caution should be taken when making configuration changes remotely. This is because a simple configuration mistake might leave the router inaccessible to remote connections, which means that you might find yourself locked out with no way to issue the `rollback 1` command.

The `commit` command supports a `confirmed` switch that is specifically designed to avoid the pitfalls associated with remote configuration. When you issue a `commit confirmed time-out` command, the system starts a timer, during which it expects to see another `commit`. If a second `commit` does not occur within the time-out value specified (a range of 1 to 65,535 is supported, with 10 minutes being the default), the system performs a `rollback 1` `commit` sequence on your behalf. After the automatic rollback you can load the `rollback 1` file to look for your mistake.

### Synchronizing Changes between REs

When working on a system with redundant REs, you generally want all changes made to the active RE to be mirrored to the backup RE. By synchronizing the two configurations, you ensure the expected behavior in the event of an RE mastership change. While the active configuration file can be manually copied to the backup RE, most users opt to use the `commit` command's `synchronize` switch. By adding the `synchronize` switch, the candidate configuration is committed on the master RE and is then internally copied and committed on the backup RE automatically.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Scheduled and Commented Commits

Starting with JUNOS software Release 5.5, you can schedule a commit that occurs at a specific time using the `commit at time` command. To cancel a pending commit, issue a `clear system commit` command. In Release 6.1 you can add a comment to the `commits` log using the `commit comment "comment-string"` option.

### Backing out of Configuration Changes

---

- Use the `rollback` command to restore one of the last 50 previously committed configurations
- Use `rollback` (or `rollback 0`) to reset the candidate configuration to the configuration currently running (which is the last version committed)
  - `rollback 1` loads the configuration before that
  - `rollback n` loads *n* configurations before that

Copyright © 2006, Juniper Networks, Inc.

#### Backing out of Changes

The software saves the last 50 committed versions of the configuration. To return to one of these versions previously committed and load it into configuration mode without activating it, use the CLI configuration `rollback` command. By default, the system returns to the most recently committed configuration:

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

To activate the configuration that you loaded, issue the `commit` command:

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

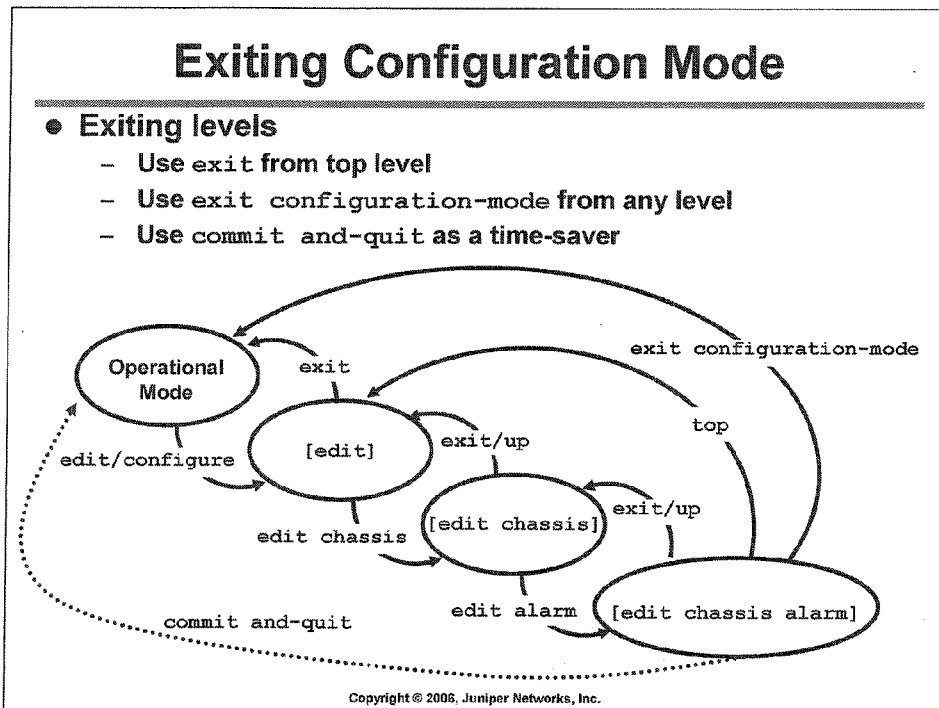
#### Specifying Rollback Files

To return to a version prior to the configuration most recently committed, include the version number in the `rollback` command:

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

The version argument can be a number in the range 0 through 49. The most recently saved configuration is version 0, which is a copy of the currently active configuration. The oldest committed configuration automatically saved is version 49.

## Configuring Juniper Networks Routers



### Exiting Levels

To exit CLI configuration mode and return to CLI operational mode, enter the `exit` command at the top level, or enter the `exit configuration-mode` command at any level. The slide illustrates the various methods of moving within the statement hierarchy. Note that `up` moves you up one level in the hierarchy while `exit` returns you to your previous location in the hierarchy.

### Saving Configuration Files

- Save current candidate configuration using `save` command

```
[edit]
user@router# save filename
```

- File saved to user's home directory unless full path name is specified
  - Only saves from the current hierarchy down!
- File name can specify:
    - A URL
    - A target on redundant Routing Engine
    - SSH `user@host:filename` notation
  - Additional capabilities:
    - `terminal` option for save commands
      - Simplifies load operations from terminal buffers
    - Pipe option for `display set`
      - Displays the `set` statements used to create a configuration
    - Periodic saves to a remote host

Copyright © 2006, Juniper Networks, Inc.

### Saving Files

You can save the software configuration from your current configuration session to an ASCII file. Doing this saves the configuration in its current form, including any uncommitted changes. If more than one user is modifying the configuration, saving it saves the changes made by all the users.

Note that only configuration statements at the current hierarchy level and below are saved. To save the entire candidate configuration, you must be at the top level of the configuration hierarchy. By default, the CLI saves the configuration to the specified file in your home directory. For example, user *Doug* would store files in `/var/home/Doug`. You can change this default by specifying a path name.

### Specifying File Names

You can specify a filename in one of the following ways:

- `ftp://user@host/path/filename`: Puts file in location explicitly described by this URL.
- `re0:filename` or `rel:filename`: Puts file on redundant Routing Engine 0 or Routing Engine 1, if present.
- `system:filename`, `system:path/filename`, `username@system:filename`, or `username@system:path/filename`: Puts file on a remote system using the SSH protocol. The default path is the user's home directory on the remote system.
- `a:filename` or `a:path/filename` (M40 router only): Puts file on the router's LS-120 floppy drive. The default path is `/` (the root-level directory). The floppy can be in either MS-DOS or UNIX (UFS) format.

*Continued on next page.*



## Configuring Juniper Networks Routers

### Additional Capabilities

JUNOS software supports saving configuration data to a terminal device. With this option the appropriate configuration hierarchy name, curly braces, and `replace` tag are added to readily accommodate pasting into another router's configuration using some form of load-terminal operation. You can also save the output to a file for later use in a file load operation. An example of `load terminal` at work is provided here:

```
[edit protocols ospf]
lab@router# save terminal
protocols {
  replace:
    ospf {
      area 0.0.0.0 {
        interface fe-0/0/0.0;
        interface fe-0/0/1.0;
        interface so-0/3/0.0;
      }
    }
}
```

Wrote 10 lines of configuration to 'terminal'

You can pipe output to `display set`. This feature converts a configuration into the actual `set` statements used to create the configuration; this option is intended to simplify the editing of configuration data being cut and pasted between routers:

```
[edit protocols ospf]
user@host# show
area 0.0.0.0 {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
[edit protocols ospf]
user@host# # show | display set
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

You can configure either a periodic or commit-driven upload of the router's configuration to a particular host using FTP. A typical configuration is shown:

```
[edit system archival]
user@host# show
configuration {
  transfer-on-commit;
  archive-sites {
    "ftp://lab:lab@10.250.0.139";
  }
}
```

Note that because a destination file name is not specified in the FTP URL, the file written to the archive host takes the form of `routername juniper.conf date time`.

### Loading Configuration Files

- Configuration information can come from an ASCII file or terminal emulation capture buffer
- The `load` command supports various arguments:
  - Override an existing configuration:
    - `load override filename`
  - Merge new statements into current configuration:
    - `load merge filename`
  - Replace existing statements in current configuration:
    - `load replace filename`
  - Take input from terminal capture buffer:
    - `load (replace | merge | override) terminal`
  - Load relative to current configuration hierarchy:
    - `load (replace | merge) (filename | terminal) relative`
- Changes candidate configuration only
  - You must issue a `commit` to activate

Copyright © 2006, Juniper Networks, Inc.

#### Loading a Configuration

You can use the configuration-mode `load` command to load a complete or partial configuration from a local file, a file on a remote machine, or from a terminal emulation program's capture buffer. The `load` command supports several arguments that determine the specifics of the operation.

#### Load Options

- `merge`: Combines the current configuration with the configuration being loaded.
- `override`: Completely overwrites the current configuration with the configuration being loaded. You must perform override operations at the root of the configuration hierarchy.
- `replace`: Looks for a `replace` tag in the configuration being loaded. Existing statements of the same name are replaced with the those in the loaded configuration for stanzas marked with the `replace` tag.
- `terminal`: Uses the text you type at the terminal as input to the configuration. Type Ctrl D to end terminal input. Usually this option is used in conjunction with a terminal emulation program's copy/paste functionality to copy and paste configuration data from one system to another.
- `relative`: Normally, a `load merge` or `load replace` operation requires that the data being loaded contain a full path to the related configuration hierarchy. The `relative` option negates this need by telling the router to *assume* that the data being loaded should be added *relative* to the current configuration hierarchy.

#### Changes Candidate Configuration Only

In all cases, after the `load` operation is complete, you must issue a `commit` to activate the changes made to the configuration.

## Configuring Juniper Networks Routers

### run is Cool

- Use the run command to execute operational-mode CLI commands from within configuration

- Can be a real time-saver when testing the effect of a recent change

```
[edit interfaces so-0/1/1]
lab@Amsterdam# set unit 0 family inet address 10.0.24.2/24
```

```
[edit interfaces so-0/1/1]
lab@Amsterdam# commit
commit complete
```

Test configuration changes without leaving configuration mode with run

```
[edit interfaces so-0/1/1]
lab@Amsterdam# run ping 10.0.24.1 count 1
PING 10.0.24.1 (10.0.24.1): 56 data bytes
64 bytes from 10.0.24.1: icmp_seq=0 ttl=255 time=0.967 ms

--- 10.0.24.1 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.967/0.967/0.967/0.000 ms
```

Copyright © 2006, Juniper Networks, Inc.

### Running with the Big Dogs (in the tall grass)

The `run` command allows you to execute operational-mode commands while in the configuration mode. It is similar to the `do` command on other vendor's equipment. This extremely handy time-saver works for all operational-mode commands and is supported at all configuration hierarchies. In the example on the slide, the operator is editing the configuration for the router's `so-0/1/1` interface. After assigning what is hoped to be the correct IP address, the change is committed (without the confirmed switch), and the `run` command is invoked to execute a quick ping test.

### Review Questions

---

1. What are the two types of CLI modes?
2. How can you navigate up two levels in the configuration hierarchy?
3. What is the purpose of using the `confirmed` switch when committing changes?
4. What command restores the candidate configuration to the currently active configuration?
5. How can you display differences between an active and candidate configuration?
6. When loading configuration files, what is the difference between the `merge`, `override`, and `replace` arguments?
7. How can you display the status of an interface while in configuration mode?

Copyright © 2006, Juniper Networks, Inc.

#### This Module Discussed:

- Logging into a Juniper Networks M-series or T-series platform;
- Operational-mode commands;
- Navigating the configuration hierarchy;
- Committing a new configuration;
- Comparing configuration files;
- Saving and manipulating configuration files; and
- Running operational-mode commands while in the configuration mode.

## **Lab 1: The JUNOS Software CLI**

---

### **Lab Objective:**

**Familiarization with the JUNOS software CLI**

Copyright © 2006, Juniper Networks, Inc.

6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100



## Configuring Juniper Networks Routers

### *Module 2: Initial Configuration*

Copyright © 2006, Juniper Networks, Inc.

CJNR-M-7.a.7.6.1

### Module Objectives

---

- **After successfully completing this module, you will be able to:**
  - Explain user authentication and authorization options
  - Describe the use of configuration groups
  - Configure system logging and tracing
  - Configure interfaces
  - Perform typical initial system configuration according to a checklist

Copyright © 2006, Juniper Networks, Inc.

#### **This Module Discusses:**

- User authentication and authorization;
- Configuration groups;
- System logging and tracing;
- Interface configuration; and
- A typical initial configuration checklist.



## **Agenda: Initial System Configuration**

### **→ User Authentication and Authorization**

- Configuration Groups
- System Logging and Tracing
- Interface Configuration
- Initial Configuration Checklist and Examples

Copyright © 2006, Juniper Networks, Inc.

### **User Authentication and Authorization**

- This slide outlines the module agenda and highlights the area we cover next. The next section describes how users gain access to the CLI and how they can be authorized to perform specific functions.

### User Authentication

---

- **Local**
  - Name and password
  - Individual accounts and home directories
  - Per-user command *class* permissions
- **RADIUS/TACACS+**
  - Supports authentication, per-class authorization, and extended regular expressions that alter the permissions associated with the user's login class
- **Authentication order can be specified**
  - By default, fall back to local authentication when RADIUS or TACACS+ fails

Copyright © 2006, Juniper Networks, Inc.

#### Local

With local password authentication, you can configure a password for each user to log into the router. After successfully logging in, the router displays the CLI prompt (>), which is preceded by the name of the user and the name of the router.

#### RADIUS/TACACS

RADIUS and TACACS+ are authentication methods used for validating users who attempt to access the router. They are both distributed client-server systems. The RADIUS and TACACS+ clients run on the Juniper Networks router; the server runs on a host connected to a remote network. Both protocols allow for user authentication, and with the appropriate Juniper Networks extensions loaded on the server, both approaches can provide for user authorization using extended regular expressions.

#### Authentication Order

You can configure the router to be both a RADIUS and TACACS+ client, and you can prioritize the order in which the software tries the different authentication methods when verifying that a user can access the router. For each login attempt, JUNOS software tries the authentication methods in order, until the password matches.

### Login Class Permissions

- Each nonroot user is associated with a login class
  - Each login class can be associated with one or more permission flags
    - Sample permissions include `access`, `configuration`, and `clear`
    - Individual commands can be allowed or denied with regular expressions
- Default login classes and permissions
  - `operator`
    - `Clear`, `network`, `reset`, `trace`, `view`
  - `read-only`
    - `View`
  - `super-user` (also known as `superuser`)
    - `All`
  - `unauthorized`
    - `None`

Copyright © 2006, Juniper Networks, Inc.

### Login Classes and Permissions

When a nonroot user logs in, JUNOS software determines that user's permissions based on the flags specified in the associated login class. For example, the `network` flag allows the operator CLI access to network utilities like `ping` and `traceroute` while the `clear` flag allows the user to issue `clear` commands like `clear interfaces statistics`. Regular expressions for commands that should be specifically allowed or denied, despite the permissions flags that might or might not be set, are also supported.

### Default Login Classes

In addition to user-definable classes, JUNOS software supports the default login classes shown on the slide. Note that you cannot override the permissions associated with a default login class. Use the CLI help function to display a truncated listing of the supported permission flags for use in user-defined login classes:

```
[edit system login]
lab@host# set class test permissions ?
Possible completions:
[
access          Open a set of values
access-control  Can view network access configuration
admin           Can modify network access configuration
admin-control   Can view user accounts
all             Can modify user accounts
clear          Can modify any configuration values
configure       Can enter configuration mode
control        Can modify any configuration values
field          Special for field (debug) support
firewall       Can view firewall configuration
...
```

## Configuring Juniper Networks Routers

### Login Class Configuration Example

- This configuration defines two nonroot users in the local database
  - The *ops* user has limited permissions, while the *lab* user has all possible permissions

```
[edit system login]
lab@Sao Paulo# show
class ops {
  permissions [ clear network view view-configuration ];
}
user lab {
  uid 2000;
  class superuser;
  authentication {
    encrypted-password "$1$EcLbIfpB$wzX7xVMo9ou8zmzdm4gHy/"; # SECRET-DATA
  }
}
user ops {
  uid 2004;
  class ops;
  authentication {
    encrypted-password "$1$b.a0nccU$kxy6uliTADLzObenDV0jq."; # SECRET-DATA
  }
}
```

Definition of permissions for the ops login class

Predefined login class with all permissions granted

A custom login class

Copyright © 2006, Juniper Networks, Inc.

### Login Class Example

This slide provides an example of [edit system login] stanza that defines two nonroot users. The *lab* user is associated with the predefined superuser class that awards all possible permissions. Note that the *superuser* class is not explicitly defined as you cannot alter the permissions associated with a predefined login class. In contrast, the *ops* user is associated with a custom login class, which is also called *ops* in this example. The *ops* class grants a limited set of permissions, as shown on the slide.

### RADIUS Authentication Example

- Use `authentication-order` to specify the sequence in which a user should be authenticated
  - Local password is the default
- Pop quiz: Based on this configuration, can the `lab` user log in if the RADIUS server is unreachable?

```
[edit system]
root@Sao_Paulo# show
host-name Sao Paulo;
authentication-order [ radius password ];
root-authentication {
    encrypted-password "$1$5Jkjbxwx$UT2e1FhTb0yVgRfGjN8IE1"; # SECRET-DATA
}
radius-server {
    10.0.1.201 secret "$9$3tRO/CuvMXwYo"; # SECRET-DATA
}
login {
    user lab {
        uid 2000;
        class superuser;
    }
}
. . .
```

Copyright © 2006, Juniper Networks, Inc.

#### Specifying the Authentication Order

By default users are authenticated according to entries in the router's local password database. To alter this behavior use the `authentication-order` statement, as shown in the example on the slide. This configuration forces the system to attempt authorization through RADIUS first, followed by an attempt to authenticate against the local password database (which always occurs by default). The `radius-server` stanza defines the properties associated with the RADIUS server, including its address, shared secret, retry attempts, etc.

#### Pop Quiz

In this example the `lab` user does not have a local password defined. As a result the `lab` user cannot authenticate locally, so the failure of the RADIUS server will prevent this user from logging into the router.

## Configuring Juniper Networks Routers

### **Agenda: Initial System Configuration**

---

- User Authentication and Authorization
- **Configuration Groups**
- System Logging and Tracing
- Interface Configuration
- Initial Configuration Checklist and Examples

Copyright © 2006, Juniper Networks, Inc.

### **Configuration Groups**

The next section describes how configuration groups can simplify repetitive configuration tasks and how these groups support Routing Engine (RE) redundancy.

## Configuration Groups

- **Groups of statements that you can apply to different sections of a configuration**
  - Shortcut method of applying the same parameters to many parts of a configuration
  - Required for redundant RE support
- **Target area of configuration inherits information from source of configuration data**

```
groups {  
    group-name {  
        configuration-data;  
    }  
}
```

Copyright © 2006, Juniper Networks, Inc.

### Configuration Groups

Configuration groups are configuration statements that you can use to direct the inheritance of that group's statements in the rest of the configuration. You can apply the same group to different sections of the configuration, and different sections of one group's configuration statements can be inherited in different places in the configuration.

Configuration groups create smaller, more logically constructed configuration files. For example, you can group statements together that are repeated in many places in the configuration, such as when configuring the encapsulation type for point-to-point interfaces, which limits all updates to that configuration group. Wildcards allow configuration data to be inherited by any object that matches a wildcard expression.

The configuration group mechanism is separate from the grouping mechanisms used elsewhere in the configuration, such as BGP groups. Configuration groups provide a generic mechanism that you can use throughout the configuration, but that are known only to JUNOS software. The individual software processes that perform the actions directed by the configuration receive the expanded form of the configuration; they do not have any knowledge of configuration groups.

Configuration groups are required when supporting redundant Routing Engines because of the need to share a common configuration file between the two REs. By using configuration groups you ensure that each RE acts only upon the relevant portions of the shared configuration file.

*Continued on next page.*

---

## Configuring Juniper Networks Routers

### Group Inheritance

Configuration groups use true inheritance, which involves a dynamic, ongoing relationship between the source of the configuration data and the target of that data. The target automatically inherits data values changed in the configuration group. The inherited values can be overridden in the target without affecting the source from which they were inherited. To have a configuration inherit the statements in a configuration group, include the `apply-groups` statement: `set apply-groups group-names;`



## Configuring Juniper Networks Routers

### Configuration Group Example

```
[edit]
lab@SanJose-re0# show groups re0
re0 {
  system {
    host-name SanJose-re0;
  }
  interfaces {
    fxp0 {
      unit 0 {
        family inet {
          address 192.168.200.51/24;
        }
      }
    }
  }
}
```

```
[edit]
lab@SanJose-re0# show groups re1
re1 {
  system {
    host-name SanJose-re1;
  }
  interfaces {
    fxp0 {
      unit 0 {
        family inet {
          address 192.168.200.52/24;
        }
      }
    }
  }
}
```

```
[edit]
lab@SanJose-re0# set apply-groups [ re0 re1 ];
```

Copyright © 2006, Juniper Networks, Inc.

### Predefined Groups

You can use two special configuration group names in a chassis with redundant Routing Engines. When defined, you can use groups `re0` and `re1` to apply Routing Engine-specific configuration data for the active Routing Engine. For example, you can use these groups to define a unique system name and/or a unique `fxp0` IP address for each Routing Engine.

For this specific example, note that both group names are applied at the global configuration level with a `set apply-groups` command.

## Configuring Juniper Networks Routers

### Interface Group Example

<pre>[edit] lab@SanJose# show groups all-atm {   interfaces {     &lt;at-*&gt; {       encapsulation atm-pvc;       atm-options {         vpi 0 maximum-vcbs 200;       }       unit 100 {         point-to-point;         vci 0.100;       }     }   } }  [edit] lab@SanJose# set interfaces apply- groups all-atm</pre>	<pre>[edit] [edit interfaces] lab@San_Jose-3# show apply-groups all-atm; at-0/0/1 {   unit 100 {     family inet {       address 1.1.1.1/24;     }   } }</pre>
---	--

Copyright © 2006, Juniper Networks, Inc.

### User-Defined Group

The example on the slide shows the configuration of a group called *all-atm*, which is designed to set common parameters for ATM interfaces, such as the indication of a point-to-point logical interface type and the VPI/VCI values associated with logical unit 100.

In this example, the *all-atm* group is applied to all ATM interfaces in the chassis by virtue of the `apply-groups all-atm` statement being specified at the `[edit interfaces]` hierarchy and because the *all-atm* group makes use of wildcards that match on all possible FPC, PIC and port numbers. Notice when looking at the interface configuration we only see the group applied and not the attributes from that group.

So, how do you display configuration attributes that are inherited from a group as opposed to being explicitly configured?

### Displaying Inherited Configuration

```
[edit]
lab@San_Jose# show interfaces | display inheritance
at-0/0/1 {
  ##
  ## 'atm-pvc' was inherited from group 'all-atm'
  ##
  encapsulation atm-pvc;
  ##
  ## 'atm-options' was inherited from group 'all-atm'
  ##
  atm-options {
    ##
    ## '0' was inherited from group 'all-atm'
    ##
    vpi 0 {
      ##
      ## '200' was inherited from group 'all-atm'
      ##
      maximum-vcs 200;
    }
  }
  unit 100 {
    ##
    ## 'point-to-point' was inherited from group 'all-atm'
    ##
    point-to-point;
    ##
  }
  . . .
}
```

Hint: Pipe results to `except #` to remove lines beginning with # from the display

Copyright © 2006, Juniper Networks, Inc.

### Viewing Inheritance

- Pipe configuration output to `display inheritance` to have the CLI call out those portions of a configuration that were inherited from a configuration group. The example on the slide shows how ## characters are used to clarify what is inherited versus what is explicitly configured. You can use an additional pipe operation to remove these characters if you want:

```
[edit]
lab@San_Jose-3# show interfaces | display inheritance | except #
at-0/0/1 {
  encapsulation atm-pvc;
  atm-options {
    vpi 0 {
      maximum-vcs 200;
    }
  }
  unit 100 {
    point-to-point;
    vci 0.100;
    family inet {
      address 1.1.1.1/24;
    }
  }
}
```

## Configuring Juniper Networks Routers

### **Agenda: Initial System Configuration**

---

- User Authentication and Authorization
- Configuration Groups
- **System Logging and Tracing**
- Interface Configuration
- Initial Configuration Checklist and Examples

Copyright © 2006, Juniper Networks, Inc.

### **System Logging and Tracing**

The next section describes system logging and tracing capabilities and configuration.

### System Logging and Tracing

- Logging and tracing allows you to monitor system and protocol events
  - System logging
    - Standard UNIX syslog syntax and options
    - Primary destination is `/var/log/messages`
  - Tracing operations
    - Protocol-specific information, for example, BGP or OSPF
    - General routing and interface behavior

```
lab@Sao_Paulo> show log messages | match fail
May 10 20:38:20 Sao Paulo chassisd[2269]: CHASSISD_SNMP_TRAP6: SNMP
trap: Power Supply failed, jnxContentsContainerIndex 2,
jnxContentsL1Index 1, jnxContentsL2Index 0, jnxContentsL3Index 0,
jnxContentsDescr Power Supply A, jnxOperatingState/Temp 6
```

This log entry indicates a power supply failure

Copyright © 2008, Juniper Networks, Inc.

### System Logging and Tracing

System logging and tracing operations allow you to track events that occur in the router—both normal router operations and error conditions—and to monitor protocol exchanges stemming from the routing protocol process running on the local router. The results of tracing and logging operations are placed in files in the `/var/log` directory on the router.

System logging (syslog) operations use a UNIX syslog-style mechanism to record system-wide, high-level operations, such as interface state changes or users logging in or out of the router. You configure system logging functionality by using the `syslog` statement at the `[edit system]` hierarchy.

Tracing operations are generally performed to obtain detailed information about the operation of one or more specific routing protocols. For example, with tracing you can record the various type of routing protocol packets sent and received, along with any protocol error events that might transpire. You configure tracing operations by using the `traceoptions` statement under the specific protocol that is to be monitored. For example, to trace OSPF protocol operation, you include the `traceoptions` statement at the `[edit protocol ospf]` hierarchy.

Tracing is often referred to as *debug* on equipment made by other vendors. Like *their* debug functionality, tracing is normally deployed *only* when there is a need to perform failure analysis. In contrast, you should enable system logging at all times, whether or not you believe that problems might be present.

### System Logging Facilities

- The facility determines the type/class of events that should be logged

– Facilities available in Release 6.3:

any	All facilities
authorization	Authorization system
change-log	Configuration change log
conflict-log	Configuration conflict log
daemon	Various system processes
dfc	Dynamic flow capture
firewall	Firewall filtering system
ftp	File Transfer Protocol process
interactive-commands	Commands executed by the UI
kernel	Kernel
pfe	Packet Forwarding Engine
user	User processes

Copyright © 2006, Juniper Networks, Inc.


### Syslog Facilities and Severity Levels

You can think of syslog facilities as categories of logging. By specifying one or more facilities, you achieve varying levels of logging granularity. This slide shows the facility classes support in JUNOS software Release 6.3. Specifying `all` is a shortcut for explicitly listing all facilities individually.

In addition to specifying what type of facility should be logged, you must also indicate the severity level at which you want logging to begin. Messages at the specific severity level, as well as those that are considered more severe, are logged. In most cases specifying a severity level of `debug` is not a good idea because this results in a potentially mind-numbing degree of chatter in your logs.

## Syslog Severity Levels

- **Setting a severity level causes router to log all messages at or above the specified priority**
  - Logging at the **critical** level also causes **alert** and **emergency** messages to appear



Copyright © 2006, Juniper Networks, Inc.

### Setting a Severity Level

Severity levels, in order of decreasing severity, are as follows:

- **emergency**: Panic or other conditions that make the system unusable;
- **alert**: Conditions that should be corrected immediately, such as a corrupted system database;
- **critical**: Critical conditions, such as hard-drive errors;
- **error**: Standard error conditions;
- **warning**: System warning messages;
- **notice**: Conditions that are not error conditions but that might warrant special handling;
- **info**: Informational messages (the default);
- **any**: All severity levels; and
- **none**: No messages.

Use the **any** severity level as a shortcut for explicitly naming all possible severities, and use **none** to disable logging of a specific facility.

## Configuring Juniper Networks Routers

### Writing to a Local File

- Use the `file` keyword to write entries to the named file on the local hard drive
  - Log and trace files are housed in `/var/log`
- Use the `archive` keyword to set system-wide defaults

```
file filename {  
    facility severity-level;  
    archive {  
        files number;  
        size size;  
        (world-readable | no-world-readable);  
    }  
}  
archive size 1m files 5;
```

The file to which the entries are written

What should be logged

Archive settings for log history

Copyright © 2006, Juniper Networks, Inc.

### Writing to a Local System Log File

In most cases you write syslog entries to one or more files that reside on the local hard drive. Use the `file` keyword to configure logging to a local file housed on the router's hard disk in the `/var/log` directory.

By default, the software retains up to ten copies of the syslog file, with each such file limited to a platform-dependant size. J-series routers default to 128 KB, TX platforms default to 10 MB, and all other M-series and T-series platforms default to 1 MB. You can alter these default archive settings by specifying that from 1 to 1000 files should be retained, and by specifying their maximum size in the range of 64 KB through 1 GB.

By default, log files are owned by the root user and have read permission only for users in the wheel (superuser) class. Users with the `trace` bit set in their login class can view the contents of log and trace files using the CLI command `show log log-file`. Use the `world-readable` flag to provide all users with read permission if access from the shell is required.

### Setting System-Wide File Defaults

If you want the same archive settings for most of your syslog files, use the `archive` keyword at the `edit system syslog` level of the hierarchy. This setting determines the default archive setting for all syslog files.



### Other Syslog Output Options

---

- Write to:

- A host

```
host hostname {  
    facility level;  
}
```

- A user

```
user (username | *) {  
    facility level;  
}
```

- The console

```
console {  
    facility level;  
}
```

Copyright © 2006, Juniper Networks, Inc.

#### Writing to Other Destinations

- Besides writing to a local file, you can also configure the router to write entries to nonfile destinations like a remote syslog host, a particular user (when logged in), or to the system console. With appropriate settings, you can have a given syslog entry written to multiple destinations simultaneously.

## Configuring Juniper Networks Routers

### Syslog Configuration Example

```
[edit system syslog]
lab@host# show
/* send all error messages to file "errors" with explicit priority */
file errors {
  any error;
  explicit-priority;
}
/* send all daemon at level info and above, and anything, */
/* warning and above, to host hot-dog.juniper.net */
host hot-dog.juniper.net {
  any warning;
  daemon info;
}
/* send all security-related information to file "security" */
file security {
  authorization info;
  interactive-commands info;
}
/* send generic messages (authorization at level notice and above, */
/* the rest at level warning and above) to file "messages" */
file messages {
  any warning;
  authorization notice;
  archive size 10m files 20 no-world-readable;
}
}
```

Comments

The log file name

The level at which to begin logging

The syslog facility

Archive and permission settings for the messages file

Copyright © 2006, Juniper Networks, Inc.

### System Logging Options Example

The example on the slide shows various syslog configurations that result in messages being written to local log files and to a remote host. General syslog configuration options include:

- archive: Configures how to archive system logging files;
- console: Configures the types of syslog messages to log to the system console;
- facility: Displays the class of log messages;
- file *filename*: Configures the types of syslog messages to log to the specified file; and
- files *number*: Displays the maximum number of system log files.

You can configure support for explicit priority in syslog messages. This configuration alters the normal syslog message format by adding a numeric priority value. The explicit priority value can simplify the task of parsing log files for important messages. For example, you can search for all messages at priority 7. The presence of explicit priority also accommodates the use of tools developed to parse the logs generated by other vendors' equipment.

*Continued on next page.*

## Configuring Juniper Networks Routers

### System Logging Options Example (contd.)

The mapping of numeric codes to message severity is as follows:

0	emergency	System panic or other condition that causes the routing platform to stop functioning
1	alert	Conditions that require immediate correction, such as a corrupted system database
2	critical	Critical conditions, such as hard drive errors
3	error	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
4	warning	Conditions that warrant monitoring
5	notice	Conditions that are not errors but might warrant special handling
6	info	Events or nonerror conditions of interest
7	debug	Software debugging messages; specify this level only when so directed by a technical support representative

Examples of syslog message, both with and without an explicit priority are provided here.

With explicit-priority:

```
Aug 21 12:36:30 router1 chassisd[522]: %DAEMON-6  
CHASSISD_PARSE_COMPLETE:
```

And the same message using the default format:

```
Aug 21 12:36:30 router1 chassisd[522]: CHASSISD_PARSE_COMPLETE: Using  
new configuration
```

## Configuring Juniper Networks Routers

### Tracing Example

- Tracing is normally used to troubleshoot routing protocol operation
  - Configure with the `traceoptions` statement under the protocol to be traced

```
[edit protocols isis]
lab@Sao_Paulo# show
Traceoptions {
  file isis-trace;
  flag error detail;
  flag hello detail;
}
level 1 disable;
interface fe-0/0/2.0;
interface lo0.0;
```

The trace file used to store trace output

What is being traced, and at what level of detail

Copyright © 2006, Juniper Networks, Inc.

### Tracing Example

This slide shows a typical protocol tracing configuration. In this case the IS-IS protocol is being traced to a file called `isis-trace`, and the operator has configured detailed tracing for all errors and for IS-IS hello packets. While the available trace flags vary by protocol, the overall concept remains the same. You specify a trace file name and one or more flags that identify what should be traced. Here are IS-IS flag options:

```
[edit protocols isis]
lab@Sao_Paulo# set traceoptions flag ?
Possible completions:
  all                Trace everything
  cs                 Trace complete sequence number (CSN) packets
  error              Trace errored packets
  general            Trace general events
  graceful-restart   Trace graceful restart events
  hello              Trace hello packets
  ldp-synchronization Trace synchronization between IS-IS and LDP
  lsp                Trace link-state packets
  lsp-generation     Trace LSP generation
  normal             Trace normal events
  packets            Trace IS-IS packets
  policy             Trace policy processing
  psn                Trace partial sequence number (PSN) packets
  route              Trace routing information
  spf                Trace SPF events
  state              Trace state transitions
  task               Trace routing protocol task processing
  timer              Trace routing protocol timer processing
```

Continued on next page.

## Configuring Juniper Networks Routers

### Tracing Example (contd.)

Sample output from the IS-IS tracing configuration shown on the previous slide is provided here:

```
[edit protocols isis]
lab@Sao_Paulo# run show log isis-trace
May 12 13:16:23 ISIS L2 periodic xmit to 01:80:c2:00:00:15 interface fe-0/0/2.0
May 12 13:16:31 ISIS L2 periodic xmit to 01:80:c2:00:00:15 interface fe-0/0/2.0
May 12 13:16:39 ISIS L2 periodic xmit to 01:80:c2:00:00:15 interface fe-0/0/2.0
May 12 13:23:50 Sending L2 LAN IIH on fe-0/0/2.0
May 12 13:23:50     max area 0, circuit type l2
May 12 13:23:50     hold time 27, priority 64, circuit id Sao_Paulo.02
May 12 13:23:50     neighbor 0:90:69:68:44:2
May 12 13:23:50     speaks IP
May 12 13:23:50     speaks IPv6
May 12 13:23:50     IP address 10.0.13.2
May 12 13:23:50     area address 49.0001 (3)
May 12 13:23:50     restart RR reset RA reset holdtime 0
May 12 13:23:50     1424 bytes of total padding
May 12 13:23:50 Received L2 LAN IIH, source id 0808.0808.0808 on fe-0/0/2.0
May 12 13:23:50     intf index 67, snpa 0:90:69:68:44:2
May 12 13:23:50     max area 0, circuit type l2, packet length 48
May 12 13:23:50     hold time 27, priority 64, circuit id 0808.0808.0808.02
May 12 13:23:50     speaks IP
May 12 13:23:50     speaks IPV6
May 12 13:23:50     IP address 10.0.13.1
May 12 13:23:50     area address 49.0001 (3)
May 12 13:23:50     restart RR reset RA reset holdtime 0
May 12 13:23:50     updating neighbor 0808.0808.0808
```

## Configuring Juniper Networks Routers

### **Agenda: Initial System Configuration**

---

- User Authentication and Authorization
- Configuration Groups
- System Logging and Tracing
- **Interface Configuration**
- Initial Configuration Checklist and Examples

Copyright © 2006, Juniper Networks, Inc.

### **Interface Configuration**

The next section describes the configuration of permanent and transient interfaces.

## Configuring Interfaces

---

- **Where we are going...**
  - Permanent and transient interfaces
  - Interface naming and selected media types
  - Logical units
  - Physical and logical interface properties
  - Configuration examples

Copyright © 2006, Juniper Networks, Inc.

### Configuring Interfaces

The items covered in the following pages include:

- *Standard interfaces*: Juniper Networks M-series and T-series platforms carry the full range of standard interfaces.
- *Interface names*: Interfaces are named by type and location in the chassis.
- *Permanent interfaces*: Two permanent interfaces, `fxp0` and `fxp1`, exist on Juniper Networks M-series and T-series platforms.
- *Interface properties*: You can configure both physical and logical properties in a given interface.

## Configuring Juniper Networks Routers

### Permanent Interfaces

- Router has several permanent interfaces
  - Out-of-band management interface is called `fxp0`
    - Requires configuration
  - Internal Routing Engine to Packet Forwarding Engine connection is called `fxp1/bcm0`
  - Internal RE-to-RE connection is `fxp2` or `em0`
    - Internal interfaces do not require any configuration; do not attempt to modify these interfaces!

Copyright © 2006, Juniper Networks, Inc.

### Permanent Interfaces

Each Juniper Networks M-series and T-series platform has several permanent interfaces. One—the management Ethernet interface—provides an out-of-band method for connecting to the router. You can connect to the management interface over a network using utilities such as SSH and Telnet, and SNMP can also use the management interface to gather statistics from the router. The `fxp0` management interface requires configuration to operate.

A second permanent interface provides internal Ethernet-based connectivity between the RE and the Packet Forwarding Engine (PFE). This interface is named `fxp1` on most platforms; in the case of the M320, the interface operates at 1 Gbps and is called `bcm0`. A proprietary protocol known as the Trivial Network Protocol (TNP) operates over these interfaces. Note that traffic arriving on the out-of-band interface is not permitted to egress on a PFE port, and vice versa. This design is intended to isolate the out-of-band network from transient traffic and to preserve bandwidth on the internal `fxp1` or `bcm0` control interface.

On platforms with redundant REs, a third interface (`fxp2` or `em0`) is used to interconnect REs for heartbeat and internal communications exchanges.

Internal interfaces like `fxp1` and `fxp2` do not require any configuration. You should not attempt to configure or modify these interfaces.




### Transient Interfaces

- PICs support transient interfaces
  - PICs plug into FPCs
  - FPC plugs into chassis
- Transient interfaces are named according to:
  - Interface media type
  - FPC slot number
  - PIC slot number within FPC
  - PIC port number
  - Channel number where applicable
- Naming example:

**at-0/2/3**

= port 3 of an ATM PIC in slot 2 on FPC 0



Copyright © 2006, Juniper Networks, Inc.

### Transient Interfaces

Each FPC can support from two to four PICs, depending on the platform. PICs provide the actual physical interfaces to the network. These physical interfaces are the router's transient interfaces. We refer to them to as transient because you can hot-swap FPCs and PICs on most platforms at any time. From the point of view of the Packet Forwarding Engine, you can place any FPC into any slot, and you can generally place any combination of PICs in any location on an FPC (consult the release notes for the PICs in question for the specific version of JUNOS software you will be using). These characteristics makes PFE interfaces transient.

You can use a transient interface for networking or to provide hardware-assisted services within the router. Service examples include IPSec, stateful firewall, and GRE tunneling. You must configure each of the transient interfaces based on which slot the FPC is installed, in which location in the FPC the PIC is installed, and to which port you are connecting.

### Transient Interface Naming

JUNOS software uses a standard naming convention when naming interfaces. You must configure each of the standard interfaces based on the slot in which the FPC is installed, the location in which the PIC is installed, and for some PICs, the port to which you are connecting. When dealing with a channelized PIC, you must also reference the correct channel/time-slot value using a `:n` form of syntax.

### Interface Naming Example

The slide shows an example of an ATM interface. Note that Juniper Networks also offers channelized interfaces; these require a colon (:) to distinguish one channel from another on the same port. For example, a channel from a channelized T-1 interface might be called `ct1-0/3/0:10`.

### Selected Interface Media Types

- **Media types:**
  - **at:** ATM over SONET/SDH ports
  - **e1:** E1 ports
  - **e3:** E3 ports
  - **fe:** Fast Ethernet ports
  - **so:** SONET/SDH ports
  - **t1:** T1 ports
  - **t3:** DS-3 ports
  - **ge:** Gigabit Ethernet ports
  - **ae:** Aggregated Ethernet ports
- **Various IP services and internal interface types**
  - **No media or ports associated with IP services or internally generated interfaces**
    - Examples include Adaptive Services and passive monitoring PICs

Copyright © 2006, Juniper Networks, Inc.

### Interface Media Types

The slide shows the list of interface media types.

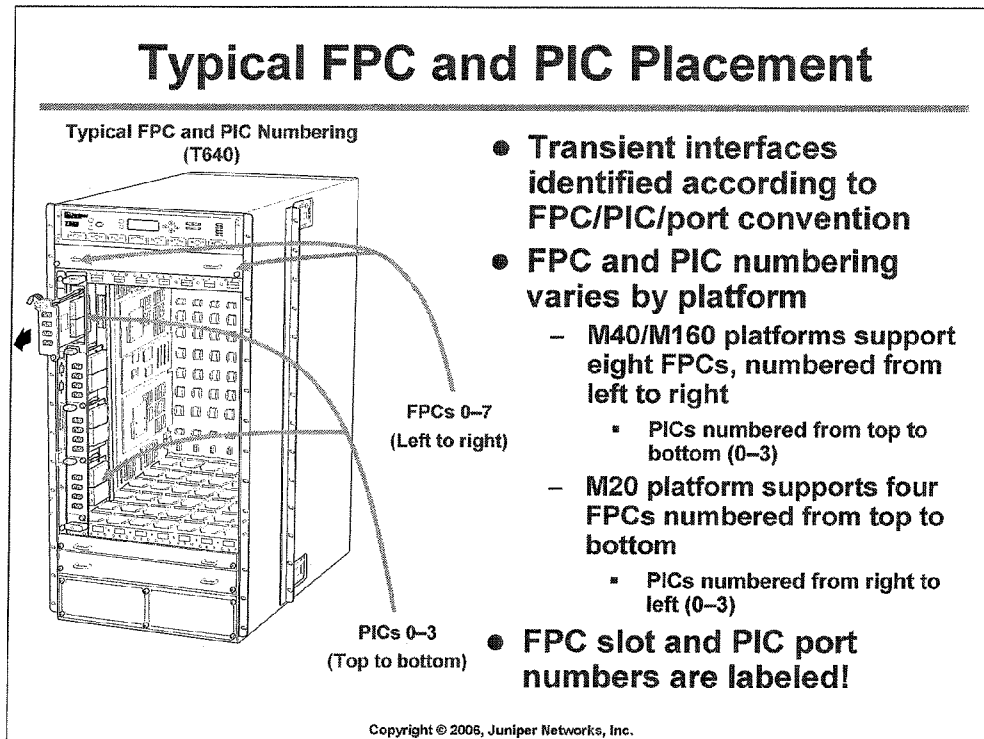
### IP Services PICs

IP Services PICs provide value-added services such as tunneling or the management of multilink bundles. IP Services PICs do not have ports or media associated with them, but they do have two-letter interface type designations as shown below. Actual coverage of the services provided by these PICs is beyond the scope of this class.

- **es:** Encryption interface;
- **gr:** Generic route encapsulation tunnel interface;
- **ip:** IP-over-IP encapsulation tunnel interface;
- **ls:** Link services interface;
- **ml:** Multilink interface;
- **mo:** Passive monitoring interface;
- **mt:** Multicast tunnel interface;
- **sp:** Adaptive services interfaces; and
- **vt:** Virtual loopback tunnel interface.

Internally generated and nonconfigurable interfaces include:

- **gre;**
- **mtun;**
- **ipip;** and
- **tap.**



### Identifying Transient Interfaces

Transient interfaces are identified by the interface's FPC slot number, the PIC slot number, and the PIC's physical port number in the form of media-type-fpc-slot/pic-slot/port-number. Channelized interfaces identify a particular subchannel with the addition of a suffix in the form of :sub-channel-number. A logical unit (aka, a subinterface) is identified with a suffix in the form of .logical-interface-number.

### FPC and PIC Slot Numbering Varies

The FPC and PIC slot numbering varies by platform due to some platforms using vertically aligned FPC slots while other platforms use a horizontal FPC arrangement. The slide details the differences in FPC and PIC slot numbering for the M40/M160 platforms versus the M20 platform. The graphic shows typical FPC and PIC numbering in the content of the T640 platform, which makes use of vertically aligned FPCs.

### The Upside?

The upside to this story is that each platform has labels that clearly identify the FPC slot number and PIC number. Further, each PIC has a label to identify the number associated with that PIC's physical ports.

### Logical Units

---

so-5/2/3.43

- **Logical units are like sub-interfaces in other equipment**
  - In JUNOS software, a logical unit is always required
    - Also used to support multipoint technologies like Frame Relay, ATM, or VLANs
- **Interface unit number is separate in meaning from the actual circuit identifier; can be any arbitrary value**
  - Suggested convention is to keep them the same
- **PPP/HDLC encapsulations support only one logical unit**
  - Must configure unit number as zero for these encapsulations
- **Multiple protocol addresses are supported on a single logical unit**
  - Typing in additional addresses does not override previous address
    - Watch for multiple addresses when correcting addressing mistakes!

Copyright © 2006, Juniper Networks, Inc.

### Logical Interfaces

Each physical interface descriptor can contain one or more logical interface descriptors. These descriptors allow you to map one or more logical (sometimes called virtual) interfaces to a single physical device. Creating multiple logical interfaces is useful for ATM and Frame Relay networks, in which you can associate multiple virtual circuits or data link layer connections with a single physical interface.

### Circuit Identifier versus Unit Number

The unit number and the circuit identifier are different in meaning. The circuit identifier identifies the logical tunnel or circuit, while the unit is used to identify a logical partition of the physical interface.

Although not required, it is generally considered best practice to keep the unit number and circuit identifier the same. This practice can greatly aid in troubleshooting when you have many logical circuits.

### Point-to-Point Encapsulations

PPP and Cisco HDLC encapsulations support only a single logical interface, and its logical unit number must be zero. Frame Relay and ATM encapsulations support multiple logical interfaces, so you can configure one or more logical unit numbers.

### Addressing Issues

A Juniper Networks M-series or T-series platform can have more than one address on a single logical interface. Issuing a second `set` command does not overwrite the previous address but simply adds to that address. Use of the CLI's `rename` command is an excellent way to correct addressing mistakes.

Also note that JUNOS software forms IGP adjacencies over all addresses when the IGP is configured on these interfaces; this behavior is worth noting because some vendors form an adjacency *only* over the primary address of an interface.

### Interface Properties

---

- **Physical properties**
  - **Clocking**
  - **Scrambling**
  - **FCS**
  - **MTU**
  - **Data link layer protocol, keepalives**
  - **Diagnostic characteristics**
    - **Local, remote, and facility loopback**
    - **BERT**
- **Logical properties**
  - **Protocol family (inet, inet6, iso, mpls)**
  - **Addresses (IP address, ISO NET address)**
  - **Virtual circuits (VCI/VPI, DLCI)**
  - **Other characteristics**

Copyright © 2006, Juniper Networks, Inc.

#### Physical Properties

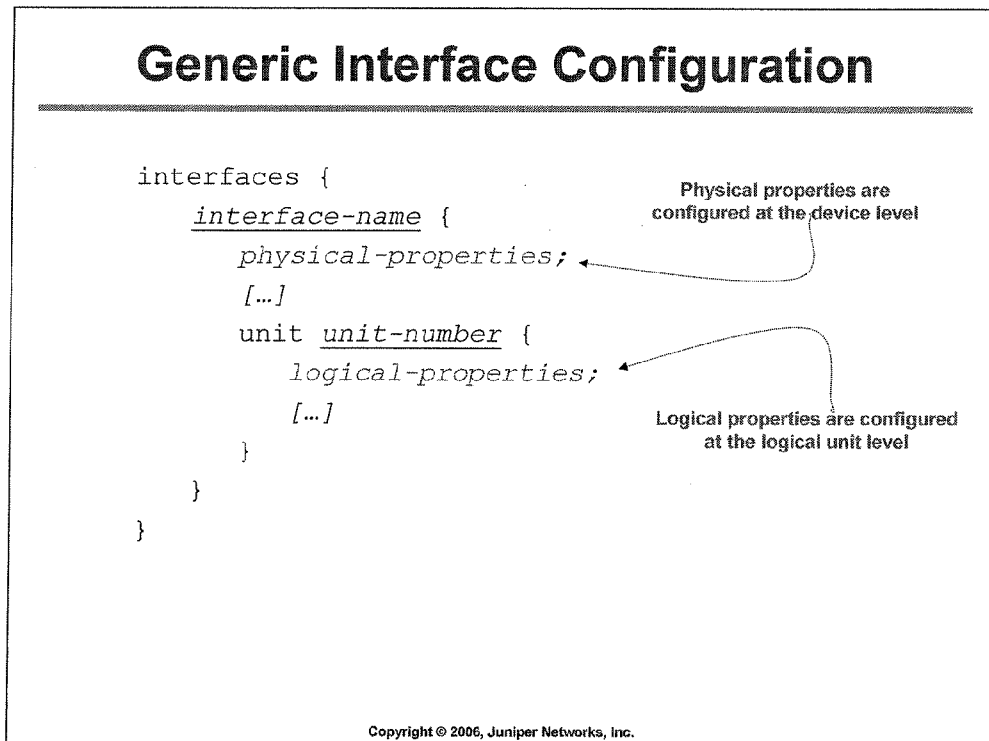
The following list provides details of the interface's physical properties:

- *Clocking*: Refers to the interface clock source, either internal or external.
- *Scrambling*: Refers to payload scrambling, which can be on or off.
- *Frame check sequence (FCS)*: You can modify to 32-bit mode (the default is 16-bit mode).
- *Maximum transmission unit (MTU)*: You can vary the size from 256 to 9192 bytes.
- *Data link layer protocol, keepalives*: You can change the data link layer protocol for the particular media type (for example, PPP to Cisco HDLC), and you can turn keepalives on or off.
- *Diagnostic characteristics*: You can enable local or remote loopbacks or set up at BERT test (see Module 5).

#### Logical Properties

The following list provides details of the interface's logical properties:

- *Protocol family*: Refers to the protocol family you want to use, such as family iso, inet, or mpls.
- *Addresses*: Refers to the address associated with the particular family (for example, IP address using family inet).
- *Virtual circuits*: Refers to the virtual circuit identifier, such as a DLCI, VPI/VCI, or VLAN tag.
- *Other characteristics*: Some other configurable options include Inverse ARP, traps, and accounting profiles.



### Statement Hierarchy

All interfaces have the same configuration hierarchy organization. JUNOS software considers all properties defined directly under the interface name to be the physical properties of that interface. The unit number represents a particular logical interface or subinterface. JUNOS software considers all properties defined directly under the unit number to be the logical properties of each particular subinterface.

### Configuring Physical Properties

- **Configure physical properties of the interface using the `set` command from the `[edit]` hierarchy:**

```
[edit]
lab@omaha# set interfaces so-1/0/3 no-keepalives
```

- **Or, park yourself at a sub-hierarchy**

```
lab@omaha> configure
[edit]
lab@omaha# edit interfaces so-1/0/3
[edit interfaces so-1/0/3]
lab@omaha# set no-keepalives
```

Copyright © 2006, Juniper Networks, Inc.

#### Configure from the Top

- ☛ In the example on the slide, we are setting no keepalives on the SONET/SDH interface. Notice that we execute this command at the top, or root, level.

#### Or, Park Yourself

Alternately, you can accomplish the same goal by navigating into the `[edit interfaces]` directory. In this case, we are parking ourselves in the `[edit interfaces so-1/0/3]` directory and then issuing the `set` command. Both examples accomplish the same goal of setting no keepalives on that interface.

### Logical Interface Settings

---

- **Logical settings**
  - Protocol family (`inet`, `inet6`, `iso`, `mpls`)
    - Protocol MTU
    - Protocol addressing
    - Other protocol options
  - Virtual circuit identifiers (VPI/VCI, DLCI)
  - Other properties according to circuit characteristics

Copyright © 2006, Juniper Networks, Inc.

### Logical Settings

For a physical interface device to function, you must configure at least one logical interface on that device. For each logical interface, you must specify, at a minimum, the protocol family that the interface supports. You also can configure other logical interface properties. These properties vary by PIC and encapsulation type but include the IP address of the interface, even if the interface does not support multicast traffic, DLCIs, VCIs and VPIs, and traffic shaping.



### Configuring Logical Interfaces

- Use the `set` command to configure a logical interface using the unit number

– For example:

```
lab@omaha> configure
[edit]
lab@omaha# set interfaces so-1/0/3 unit 40 dlcI 40
```

- Or park yourself at the unit level:

```
lab@omaha> configure
[edit]
lab@omaha# edit interfaces so-1/0/3 unit 40
[edit interfaces so-1/0/3 unit 40]
lab@omaha# set dlcI 40
```

Copyright © 2006, Juniper Networks, Inc.

#### Configuration of Logical Interfaces: Method 1

- ☛ Use the `set` command at the top of the configuration tree. You must specify logical properties after the unit number.

#### Configuration of Logical Interfaces: Method 2

You also can park yourself under that logical unit and configure properties using the normal `set` commands.

### Configuring Protocol Families

---

- **Each major protocol is called a family**
  - Multiple families can live on the same logical interface
  - Family encompasses entire protocol suite
    - Internet protocol has TCP, UDP, and ICMP as family members
- **Supported protocol families are:**
  - IP (`inet`)
  - IPv6 (`inet6`)
  - International Standards Organization (`iso`)
  - Traffic engineering (`mpls`)

Copyright © 2006, Juniper Networks, Inc.

### One Happy Family

You can specify more than one family on a logical interface. Specifying a family encompasses the entire protocol suite. For example, when you enable family `inet` on an interface, this family enables TCP, UDP, and ICMP to run.

### Supported Protocol Families

The following protocol families are supported:

- `inet` (Internet Protocol): You must configure this protocol family for the logical interface to support IP traffic, including OSPF, BGP, and ICMP.
- `inet6` (IP version 6): You must configure this protocol family for the logical interface to support IPv6 traffic.
- `iso` (International Standards Organization): You must configure this protocol family for the logical interface to support IS-IS traffic.
- `mpls` (Multiprotocol Label Switching): You must configure this protocol family for the logical interface to participate in an MPLS path.

### Internet Protocol Family (`inet`)

---

- **Allows you to set:**
  - IP address: address *A.B.C.D/prefix length*
  - Remote address on point-to-point links: destination *A.B.C.D*
  - Broadcast address: broadcast *A.B.C.D*
  - Primary address: primary
  - Preferred address: preferred
  - MTU size: mtu *bytes*
  - ICMP redirect control: no-redirects
  - Multicasts only: multicast-only

Copyright © 2006, Juniper Networks, Inc.

#### Family `inet` Options

The following list provides the details of the options shown on the slide:

- *Interface address*: This address is defined as the interface address/destination prefix.
- *Address of the remote side of the connection (for point-to-point interfaces only)*: Specify this address in the destination statement.
- *Broadcast address for the interface's subnet*: Specify this address in the broadcast statement.
- *Primary address*: Each interface has a primary local address. If an interface has more than one address, the primary local address is used by default as the source address when you originate packets out the interface where the destination gives no hint about the subnet (for example, some ping commands). By default, the primary address on an interface is the lowest numbered non-127 preferred address on the interface. To override the default and explicitly configure the preferred address, include the `primary` statement when configuring the address.
- *Preferred address*: Each subnet on an interface has a preferred local address. If you configure more than one address on the same subnet, the preferred local address is chosen by default as the source address when you originate packets to destinations on the subnet. By default, the preferred address is the lowest numbered address on the subnet. To override the default and explicitly configure the preferred address, include the `preferred` statement when configuring the address.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Family `inet` Options (contd.)

- *MTU size:* For each interface, you can configure an interface-specific MTU. If you increase the size of the protocol MTU, you must ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead.
- *ICMP redirect control:* Do not send protocol redirect messages on the interface. To disable the sending of protocol redirect messages for the entire router, include the `no-redirects` statement at the `[edit system]` hierarchy level.
- *Multicast only:* Configure the unit and family so that it can transmit and receive multicast traffic only. You can configure this property on the IP family only.

### Protocol Family Example

- **Sample configuration for the inet family:**

```
lab@omaha> configure
[edit]
lab@omaha# edit interfaces so-1/0/3
[edit interfaces so-1/0/3]
lab@omaha# set unit 0 family inet address 10.0.20.1/24
```

Note the use of CIDR notation for mask length

- **Displayed as:**

```
[edit interfaces so-0/1/3]
lab@omaha# show
unit 0 {
    family inet {
        address 10.0.20.1/24;
    }
}
```

- **Use display set to convert configuration stanza to set commands**

```
[edit interfaces so-0/1/3]
lab@omaha# show | display set
set interfaces so-0/1/3 unit 0 family inet address 10.0.20.1/24
```

Copyright © 2008, Juniper Networks, Inc.

### Sample Configuration

This slide shows an example of configuring a family inet address.

### Show Interface

When viewing the configuration, you see that the protocol family and address has a logical property configured after the unit number. You can conveniently show all the logical properties associated with a specific logical unit while in configuration mode by adding the logical unit number to the interface name, as shown:

```
[edit]
lab@San_Jose-3# show interfaces at-0/1/0.100
vci 100;
family inet {
    address 10.0.0.1/24;
}
```

This command syntax is a form of shorthand when compared to the syntax previously required:

```
[edit]
lab@San_Jose-3# show interfaces at-0/1/0 unit 100
vci 100;
family inet {
    address 10.0.0.1/24;
}
```

## Configuring Juniper Networks Routers

### Interface Configuration Examples

```
[edit interfaces]
lab@Sydney# show at-0/2/1
description "SY to HK and DE";
atm-options {
  vpi 0 {
    maximum-vc 200;
  }
}
unit 0 {
  description "to HK";
  vci 100;
  family inet {
    address 10.0.15.1/24;
  }
}
unit 101 {
  description "to DE";
  vci 101;
  family inet {
    address 172.16.0.1/24;
  }
}
```

**An ATM interface with multiple units**

```
[edit interfaces]
lab@Sydney# show fe-0/0/2
unit 0 {
  family inet {
    address 10.0.13.1/24;
  }
  family mpls;
}
```

**Fast Ethernet with inet and mpls support**

```
[edit interfaces]
lab@Sydney# show so-0/1/3
no-keepalives;
encapsulation frame-relay;
unit 100 {
  dlci 100;
  family inet {
    address 4.4.4.4/24;
  }
}
```

**A SONET interface running Frame Relay with keepalives (LMI) disabled**

Copyright © 2006, Juniper Networks, Inc.

### Interface Configuration Examples

This slide shows three configuration examples for common interface types. You can use cut and paste in conjunction with the `load merge terminal` command to modify these configurations for use in your router. Piping the output of a `show` command to `display set` is an excellent way to see the commands that were used to create a given configuration stanza.

Note that each configuration example makes use of at least one logical unit, and that a protocol family and related logical properties are specified at the unit level. The commands used to configure the ATM interface shown on the slide are shown here using the CLI's support for piped output to `display set`. In this example, the `relative switch` is added, revealing commands that are appropriate for the current edit level:

```
[edit interfaces at-0/2/1]
lab@Sydney# show | display set relative
set description "SY to HK and DE"
set atm-options vpi 0 maximum-vc 200
set unit 0 description "to HK"
set unit 0 vci 100
set unit 0 family inet address 10.0.15.1/24
set unit 101 description "to DE"
set unit 101 vci 101
set unit 101 family inet address 172.16.0.1/24
```

*Continued on next page.*

## Configuring Juniper Networks Routers

### Interface Configuration Examples (contd.)

The configuration below is somewhat complicated because it reflects a channelized DS3 Q-PIC interface that is broken down into channelized and unchannelized DS1s. In the former case, two DS0 channel bundles are defined, along with the related ds interfaces:

```
[edit interfaces]
lab@Sydney# show | no-more
ct3-0/1/0 {
    description "Q-PIC based CT3 to CT1 and non-channelized T1s.";
    t3-options {
        loopback remote;
        loop-timing;
    }
    partition 1 interface-type ct1;
    partition 2-28 interface-type t1;
}
ct1-0/1/0:1 {
    description "CT1 to NxDS0s.";
    t1-options {
        line-encoding ami;
        framing sf;
        bert-algorithm all-ones-repeating;
    }
    partition 1 timeslots 1-10 interface-type ds;
    partition 2 timeslots 11-23 interface-type ds;
}
ds-0/1/0:1:1 {
    description "first DS0 channel bundle of ct1-0/1/0:1";
    unit 0 {
        family inet {
            address 1.1.1.1/24;
        }
    }
}
ds-0/1/0:1:2 {
    description "Second DS0 channel bundle of ct1-0/1/0:1";
    unit 0 {
        family inet {
            address 2.2.2.2/24;
        }
    }
}
t1-0/1/0:2 {
    description "First full T1 from ct3-0/1/0, range is t1-0/1/0:[2-28]";
    encapsulation cisco-hdlc;
    unit 0 {
        family inet {
            address 3.3.3.3/24;
        }
    }
}
```

**A Q-PIC CT3 interface**

**Device-level options and partition definitions**

**First partition is a channelized DS1; DS1 properties for this channel are configured**

**Definition of DS0 bundles on the CT1**

**ds interface types relating to ct1-0/1/0:1**

**First unchannelized DS1**

## Configuring Juniper Networks Routers

### Disabling and Deactivating

- Use `deactivate` to cause the related stanza to be ignored

```
[edit interfaces]
lab@San_Jose# deactivate so-0/1/0

[edit interfaces]
lab@San_Jose# show so-0/1/0
##
## inactive: interfaces so-0/1/0
##
unit 0 {
  family inet {
    address 10.0.1.2/24;
  }
}
```

- Setting an interface or logical unit to disable signals JUNOS software to treat that interface as administratively down

```
[edit interfaces]
lab@San_Jose# set so-0/1/0 disable

[edit interfaces]
lab@San_Jose# show so-0/1/0
disable
unit 0 {
  family inet {
    address 10.0.1.2/24;
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

### Deactivating an Interface

You can deactivate statements and identifiers in a configuration so that they do not take effect when you issue the `commit` command. Any deactivated statements and identifiers are marked with the `inactive:` tag. They remain in the configuration but are not activated when you issue a `commit` command. In essence, the deactivated portion of the configuration is returned to the factory defaults.

To deactivate a statement or identifier, use the `deactivate` configuration mode command: `deactivate (statement | identifier)`. To reactivate a statement or identifier, use the `activate` configuration mode command: `activate (statement | identifier)`.

Deactivating the `so-0/1/0` interface as shown on the slide results in the related configuration being ignored while the interface remains administratively enabled:

```
lab@San_Jose> show interfaces so-0/1/0 terse
Interface           Admin Link Proto Local           Remote
so-0/1/0            up      up
```

### Disabling an Interface

For interfaces you can include a `disable` statement to administratively disable the related functionality. When you deactivate a statement, JUNOS software completely ignores that specific object or property and does not apply it at all when you issue a `commit` command. When you disable a functionality, it is activated when you issue a `commit` command but is treated as being down or administratively disabled, as shown here:

```
lab@San_Jose> show interfaces so-0/1/0 terse
Interface           Admin Link Proto Local           Remote
so-0/1/0            down   up
so-0/1/0.0         up     down   inet  10.0.1.2/24
```



### **Agenda: Initial System Configuration**

---

- User Authentication and Authorization
- Configuration Groups
- System Logging and Tracing
- Interface Configuration
- **Initial Configuration Checklist and Examples**

Copyright © 2006, Juniper Networks, Inc.

#### **Initial Configuration Checklist**

- The next section provides a checklist of typical configuration tasks for a newly installed M-series or T-series platform along with typical configuration examples.

### Initial Configuration Checklist

- **The following items are normally configured at initial system installation:**
  - Root password
  - Host name
  - Domain name and DNS server address
  - Configuration file compression (no longer necessary)
  - System logging
  - Out-of-band management interface
  - Default and backup routers for management network
  - Configure system services for remote access
  - User accounts
  - System time
  - Loopback and transient interfaces
  - Remaining configuration needed to place the router into service (protocols, firewall filters, etc.)

Copyright © 2006, Juniper Networks, Inc.

### Initial Configuration

When you receive a Juniper Networks M-series or T-series router, JUNOS Internet software is preinstalled. Once you power on the router, it is ready for you to configure it. You can configure the router from a console connected to the router's console port. We recommend you configure the following items at installation time:

- Root password (By default, the only user that can access a router is *root*. The root user must log in through the console port. There is no root password specified in the initial active configuration, so we recommend setting this password immediately.);
- Hostname or name of the router;
- IP address of a domain name system (DNS) server;
- Configuration file compression is automatic (you must configure it explicitly on older versions);
- System logging (syslog);
- Management Ethernet interface (`fxp0`) IP address and prefix length (you could use this IP address for an out-of-band management network);
- IP address of a default router;
- System services for remote access (Telnet, SSH, and FTP);
- Local user accounts;
- Time of day/Network Time Protocol;
- Loopback and transient interfaces; and
- Any remaining functionality needed to place the router into service, for example, routing protocols, routing policy, firewall filters, etc.

## Configuring Juniper Networks Routers

### Initial Configuration (1 of 10)

- Log in as root

```
.....  
starting local daemons:..  
Fri Jan 17 22:23:32 UTC 1997
```

```
Amnesiac (ttyd0)
```

Amnesiac indicates a factory default configuration

```
login: root  
Last login: Fri Jan 17 22:21:55 on ttyd0
```

```
--- JUNOS 5.2R2.3 built 2002-03-23 02:44:36 UTC
```

```
Terminal type? [vt100] <enter>
```

BSD shell prompt

```
root%
```

- Start CLI

```
root% cli  
root>
```

Copyright © 2006, Juniper Networks, Inc.

### Log in as Root

- ❑ When you receive an M-series or T-series platform from the factory, the root password is not set. To log into the router for the first time, you must login through the console port using the `root` username with no password. The *Amnesiac* indication shows the lack of a configured hostname in a factory default configuration.

### Start the CLI

When logging in with the root username, you must start the CLI process on the Routing Engine using the `cli` command, as shown on the slide.

## Configuring Juniper Networks Routers

### Initial Configuration (2 of 10)

- Enter configuration mode

```
root> configure
[edit]
root#
```

- Configure root password

- Plain text

```
root# set system root-authentication plain-text-
password
```

- Pre-encrypted password

```
root# set system root-authentication encrypted-
password encrypted-password
```

Do not enter a clear text password in this mode!

- Secure Shell (SSH) key

```
root# set system root-authentication ssh-rsa key
```

Copyright © 2006, Juniper Networks, Inc.

### Enter Configuration Mode

After starting the CLI, you enter operational mode. You can make changes to the configuration only in configuration mode. Enter configuration mode by entering the command `configure` at the operational-mode prompt, as shown on the slide.

### Set the Root Password

You should always set the root password, which you can do using the following methods:

- *Plain-text password:* When setting the plain-text password, the system prompts for the password on a separate line and does not log the plain-text password if you enable command logging.
- *Pre-encrypted password:* Using this method, JUNOS software expects a pre-encrypted password to be pasted at the end of the configuration statement, `set system root-authentication encrypted-password encrypted password`
- *SSH key:* The Secure Shell (SSH) package is only available on domestic systems.

## Configuring Juniper Networks Routers

### Initial Configuration (3 of 10)

- **Configure router name**

```
[edit]
root# set system host-name lab2
```

- **Configure router domain name**

```
[edit]
root# set system domain-name domain-name.tld
```

- **Configure name server address**

```
[edit]
root@# set system name-server ns-address
```

- **Configure configuration file compression**

- Is the default for recent versions
- For older versions:

```
[edit]
root@# set system compress-configuration-files
```

Copyright © 2008, Juniper Networks, Inc.

#### Set the Router Name

Set the router name or hostname by typing the command `set system host-name hostname`.

#### Set the Domain Name for the Router

If you are using DNS, set the router's domain name by typing the `set system domain-name domain.tld` command.

#### Set the Name Server Address

To use DNS you must also specify the location of name server's IP address to enable the generation of DNS queries.

#### Enable Configuration File Compression (Optional)

Because the active configuration file (`juniper.conf`) and the first three rollback copies are stored on the router's flash memory, large configurations might result in a shortage of flash memory storage space. By enabling compression you can reduce the size of your configuration files by as much as 90%. We recommend that you enable compression when the size of your configuration file exceeds 3 MB. Issue a `file list /config detail` command to display the size of your configuration files. To enable compression, issue a `set system compress-configuration-files` command from the `[edit]` hierarchy. Note that you must issue a `commit` twice to activate the compression and that the name of the configuration file changes to `juniper.conf.gz` as a result.

## Configuring Juniper Networks Routers

### Initial Configuration (4 of 10)

- **Adjust syslog parameters as needed**

- Interactive command and configuration change logging is a good idea
- Adjusting archive settings for more history also recommended

```
[edit system syslog]
root@lab2# show
user * {
    any emergency;
}
file messages {
    any notice;
    authorization info;
    archive size 1m files 20;
}
file cli-commands {
    interactive-commands any;
    archive size 1m files 10;
}
file config-changes {
    change-log info;
    archive size 1m files 10;
}
```

Archive settings adjusted on default syslog file

Interactive commands and configuration changes

Copyright © 2006, Juniper Networks, Inc.

### Adjust Syslog Settings as Needed

Although the default syslog settings are quite workable, it is common to see a configuration that tweaks system logging as shown on the slide. In this case all interactive CLI commands (to include both configuration-mode and operational-mode commands) are logged to a file named `cli-commands`. A separate log file named `config-changes` is used to store all long entries relating to configuration changes.

All the syslog files are using modified archive settings to provide additional history over the defaults. These syslog settings will help you by identifying what users issued which commands, and when.

### Initial Configuration (5 of 10)

- Commit changes so far

```
[edit]
root# commit
commit complete
```

Note host name takes effect after the commit

```
[edit]
root@lab2#
```

- Configure management interface IP address and prefix

```
[edit]
root@lab2# set interfaces fxp0 unit 0 family inet address
ip-address/prefix-length
```

- Define a backup router

- Used when routing daemon is not running
  - Required when using redundant Routing Engines

```
[edit]
root@lab2# set system backup-router gateway-address
```

Copyright © 2006, Juniper Networks, Inc.

### Commit Changes

To make these changes part of the active configuration, issue the `commit` command. This slide shows how the newly assigned host name takes effect after the operator issues a `commit`.

### Management Interface IP Address and Prefix

Set the management Ethernet port (`fxp0`) using the command `set fxp0 unit 0 family inet address ip-address/prefix-length`.

### Define a Backup Router

The backup router statement is used to define a default route that supports `fxp0` out-of-band management traffic while the system boots and the routing protocol daemon (`rpd`) is not yet started. This feature is particularly useful when you want to route an SNMP system-up trap to a remote monitoring station. The backup router also acts as a safety net should the routing process (`rpd`) fail to start. In this case, the backup router remains installed in the routing and forwarding tables, and access to the router through the management interface remains enabled. The backup default route is removed when the routing process starts successfully.

Note that the backup-router statement is necessary when your platform has redundant REs and you must be able to access/communicate with either RE from a remote network location. This is because `rpd` is not active on a backup RE.

### Initial Configuration (6 of 10)

- Define static route for OoB management network

```
[edit]
root@lab2# edit routing-options

[edit routing-options]
root@lab2# set static route ip-address/prefix-length
next-hop OoB-next-hop-address no-readvertise
```

- Configure system services for remote access

```
[edit]
root@lab2# set system services ssh

[edit]
root@lab2# set system services telnet

[edit]
root@lab2# set system services ftp
```

Copyright © 2006, Juniper Networks, Inc.

### Define a Static Route for OoB Management Network

You should define a static route used to reach destinations on the OoB management network because dynamic routing protocols are normally not run over the OoB network. In most cases this is a default route that points to the first-hop router on the `fxp0` management network. Note that the static route is used only when `rpd` is running. The `backup-router` configuration is used when `rpd` is not running, such as for a backup RE. The `no-readvertise` tag ensures that the static route intended for use by the OoB network cannot be advertised in a routing protocol. The `retain` flag ensures that the static route stays active during RE transitions.

### Configure Remote Access Protocols

In most cases you will want to configure one or more remote access protocols to support Telnet or SSH logins from users and to support FTP-based file transfers. The factory defaults do not enable these services, which forces local console access during initial system configurations. Note that the root user can login remotely using SSH, but not Telnet. If you want, you can prevent remote root logins using SSH. A user can use Telnet to access a router and then become root at a shell prompt using the `su` command, but the clear-text nature of Telnet makes this a bad idea.

You can configure aspects of some services, as shown here for the SSH protocol:

```
[edit]
lab@lab2# set system services ssh ?
Possible completions:
 <[Enter]>          Execute this command
+ apply-groups      Groups from which to inherit configuration data
+ apply-groups-except Don't inherit configuration data from these groups
  connection-limit  Maximum number of allowed connections (1..250)
+ protocol-version  Specify ssh protocol versions supported
  rate-limit        Maximum number of connections per minute (1..250)
  root-login        Configure root access via ssh
```



## Initial Configuration (7 of 10)

### ● Configure user accounts

- Use predefined login classes, or create your own

```
[edit system login]
root@lab2# show
user dr-data {
  full-name "The Doctor 'O Data";
  uid 2003;
  class superuser;
  authentication {
    encrypted-password "$1$B78jkPLd$8VVjFv6D.ZQQev/5rstET0"; #
SECRET-DATA
  }
}

[edit system login]
root@lab2# show | display set
set system login user dr-data full-name "The Doctor 'O Data"
set system login user dr-data uid 2003
set system login user dr-data class superuser
set system login user dr-data authentication encrypted-password
"$1$B78jkPLd$8VVjFv6D.ZQQev/5rstET0"
```

The user ID is created automatically when not explicitly configured

The commands used to create the *dr-data* user account, courtesy of display set

Copyright © 2006, Juniper Networks, Inc.

### User Accounts

You must configure user accounts so that users other than root can log into the router. Each user defined in the configuration receives a home directory on the hard drive at `/var/home/username`.

For each account, you define the login name for the user and, optionally, information that identifies the user. To create user accounts, include the `user` statement at the `[edit system login]` hierarchy level.

For each user account, you can define the following items:

- **User name:** Name that identifies the user. It must be unique within the router. Do not include spaces, colons, or commas in the user name.
- **User's full name:** If the full name contains spaces, enclose it in quotation marks. Do not include colons or commas.
- **User identifier (UID):** Numeric identifier associated with the user account name. The identifier must be in the range 100 through 64000 and must be unique within the router. If you do not assign a UID to a user name, the software assigns one when you commit the configuration, preferring the lowest available number. You must ensure that the UID is unique. However, it is possible to assign the same UID to different users. If you do this, the CLI displays a warning configuration, then assigns the duplicate UID.

Continued on next page.

## Configuring Juniper Networks Routers

### User Accounts (contd.)

- *User's access privilege:* One of the login classes you defined in the class statement at the [edit system login] hierarchy level or one of the default classes.
- *Authentication method or methods and passwords that the user can use to access the router.* You can use SSH or an MD5 password, or you can enter a plain-text password that JUNOS software encrypts using MD5-style encryption before entering it in the password database.

All users who can log into the router must be in a login class. With login classes, you define the following:

- Access privileges users have when they are logged into the router;
- Commands and statements that users can and cannot specify; and
- How long a login session can be idle before it times out and the user is logged off.

You can define any number of login classes. The software contains a few predefined login classes, which are operator, view-only, super-user, and unauthorized.

### Initial Configuration (8 of 10)

- **Configure time zone and manually set the time of day**

- **Configure time zone:**

```
[edit]
```

```
root@lab2# set system time-zone America/Los_Angeles
```

- **Set date and time manually**

```
root@lab2> set date ?
```

```
Possible completions:
```

```
<time> New date and time (YYYYMMDDhhmm.ss)
```

```
ntp Set date/time using Network Time Protocol
```

```
servers
```

```
root@lab2> set date 200405141017.20
```

```
Fri May 14 10:17:20 PDT 2004
```

- **Or, configure NTP**

Copyright © 2006, Juniper Networks, Inc.

### Time

To set the current date and time on the router, use the `set date YYYYMMDDHHMM.ss` command. YYYY is the four-digit year, MM is the two-digit month, DD is the two-digit date, hh is the two-digit hour, mm is the two-digit minute, and ss is the two-digit second. At a minimum, you must specify the two-digit minute. All other parts of the date and time are optional.

The default local time zone on the router is UTC (Coordinated Universal Time, formerly known as Greenwich Mean Time). To modify the local time zone, include the `time-zone` statement at the `[edit system]` hierarchy level. You specify the time zone using the continent/country/zone primary name. For the time zone change to take effect for all processes running on the router, you must reboot the router.

### Use NTP

The Network Time Protocol (NTP) provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network. NTP uses a returnable-time design in which a distributed subnet of time servers operating in a self-organizing, hierarchical, master-slave configuration synchronizes local clocks within the subnet and to national time standards by means of wire or radio. The servers also can redistribute reference time using local routing algorithms and time daemons. RFC 1305, *Network Time Protocol (Version 3) Specification, Implementation, and Analysis*, defines NTP.

When configuring NTP, you do not actively configure time servers. Rather, all clients are also servers. An NTP server is not believed unless it, in turn, is synchronized to another NTP server—which itself must be synchronized to something upstream, eventually terminating in a high-precision clock.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Use NTP (contd.)

If the time difference between the local router clock and the NTP server clock is more than 128 milliseconds but less than 128 seconds, the clocks are slowly stepped into synchronization. However, if the difference is more than 128 seconds, the clocks are not synchronized. You must set the time on the local router so that the difference is less than 128 seconds to start the synchronization process. On the local router, you set the date and time using the `set date` command. To set the time automatically, use the `boot-server` statement at the `[edit system ntp]` hierarchy level, specifying the address of an NTP server along with the desired authentication mechanisms and keys. A sample NTP stanza is shown here:

```
[edit system]
root@lab2# show ntp
boot-server 10.0.1.201;
authentication-key 1 type md5 value "$9$AWTtt1hW87wYohS"; # SECRET-DATA
server 10.0.1.201 key 1; # SECRET-DATA
trusted-key 1;
```

## Configuring Juniper Networks Routers

### Initial Configuration (9 of 10)

#### Configure loopback and transient interfaces

```
[edit interfaces]
root@lab2# set lo0 unit 0 family inet address 192.168.12.1

[edit interfaces]
root@lab2# set fe-0/0/2 unit 0 family inet address 10.0.13.2/24

[edit interfaces]
root@lab2# show lo0
unit 0 {
    family inet {
        address 192.168.12.1/32;
    }
}

[edit interfaces]
root@lab2# show fe-0/0/2
unit 0 {
    family inet {
        address 10.0.13.2/24;
    }
}
```

Loopback interface  
must use a /32

Copyright © 2006, Juniper Networks, Inc.

#### Configure Loopback and Transient Interfaces

In most cases you will want to assign a unique IP address to the router's loopback interface for use as a router ID and to accommodate loopback peering for IBGP sessions. Because no one buys a Juniper Networks router just to ping their loopback address, you will likely also need to configure one or more transient interfaces.

This slide shows sample command syntax used to configure the `inet` protocol family and a related IPv4 address, on both a loopback and transient interface. Note that omitting the mask length results in an assumed /32, which is required for the loopback interface; a commit error is returned when a /32 is not specified on the `lo0` interface:

```
[edit interfaces lo0]
lab@Sao_Paulo-3# set unit 0 family inet address 192.168.12.1/24

[edit interfaces lo0]
lab@Sao_Paulo-3# commit check

[edit interfaces lo0 unit 0 family inet]
'address 192.168.12.1/24'
  Loopback addresses' prefix must be 32 bits
error: configuration check-out failed

[edit interfaces lo0]
lab@Sao_Paulo-3# rename unit 0 family inet address 192.168.12.1/24 to address
192.168.12.1/32

[edit interfaces lo0]
lab@Sao_Paulo-3# commit check
configuration check succeeds
```

CLI's rename command is used  
to correct an addressing  
mistake

### Initial Configuration (10 of 10)

---

- **Configure remaining items required to place the router into service**
  - Routing protocols (OSPF, IS-IS, BGP, PIM, etc)
  - Routing policies
  - Firewall filters to secure the local router and possible attached devices
  - MPLS traffic engineering
- **These tasks are detailed in subsequent modules**

Copyright © 2006, Juniper Networks, Inc.

#### **Configure Remaining Functionality**

At this stage you should have a baseline system that supports an out-of-band management network, system logging, access for authorized users, and loopback/transient interface configuration. The specifics of your installation environment generally determine what is configured next. In most cases you move onto the configuration of routing protocols, routing policy, and firewall filters that are designed to harden the local router or to provide protection to attached devices.

#### **All in Due Time**

The remaining functionality is detailed in subsequent modules. You will be adding additional functionality to your initial system configuration as you proceed through this course.

## Review Questions

---

1. What is the default root password?
2. Describe at least three parameters normally configured as part of initial system configuration.
3. Explain when a backup router is needed.
4. Describe how a router's permanent interfaces are used.
5. List three examples of physical interface parameters.
6. List two examples of logical interface settings.
7. What FPC is associated with interface at-0/3/2.135?
8. In the previous question, what does the .135 represent?

Copyright © 2006, Juniper Networks, Inc.

### This Module Discussed:

- User authentication and authorization;
- Configuration groups;
- System logging and tracing;
- Interface configuration; and
- Initial configuration checklist.

## **Lab 2: Initial Configuration**

---

**Lab Objectives:**  
**Perform initial system configuration and  
monitor the router's operation**

Copyright © 2006, Juniper Networks, Inc.





**Configuring Juniper Networks Routers**

***Module 3: Protocol-Independent  
Routing Properties***

Copyright © 2006, Juniper Networks, Inc.

CJNR-M-7.a.7.6.1

### Module Objectives

---

- After successfully completing this module, you will be able to:
  - Create static routes
  - Create aggregate routes
  - Create generated routes
  - Describe JUNOS software routing tables
  - Describe the JUNOS software route selection process
  - Explain load-balancing options

Copyright © 2006, Juniper Networks, Inc.

#### **This Module Discusses:**

- Static routes;
- Aggregate routes;
- Generated routes;
- The JUNOS software routing tables;
- The active route selection process; and
- JUNOS software load-balancing options.

## Protocol-Independent Routing

---

- Where we are going...
  - Static, aggregate, generated, and martian routes
  - Routing tables and route preferences
  - Configure routing protocol process system logging
  - Load balancing

Copyright © 2006, Juniper Networks, Inc.

### Protocol-Independent Routing

The slide shows the topics examined on the following pages.

### Static Routes

- Manually configured routes added to the routing table
- Once active, routes remain in the routing table until deleted
- Route configured at the `routing-options` hierarchy level

```
[edit]
routing-options {
  static {
    defaults {
      static-options;
    }
    route destination-prefix/
      next-hop;
      static-options;
    }
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

### Static Routes

You can use static routes in a networking environment for multiple purposes, including a default route for the autonomous system (AS) as well as routes to customer networks. Unlike dynamic routing protocols, you manually configure the routing information provided by static routes on each router in the network.

### Active Static Routes

In addition, these routes are permanently in the routing table until a network administrator removes them or they become *nonactive*. One possible way for a static route to be nonactive is for the IP address of the next hop to be unreachable across a directly connected interface.

### Configuration Hierarchy

All configuration for static routes occurs at the `[edit routing-options]` level of the hierarchy. The possible options to assign to static routes are:

- `as-path`: Used if this route is intended to be redistributed into BGP and you want to add values manually to the AS path attribute.
- `community`: Used if this route is intended for BGP, and you want to add community values to the route for use in your AS.
- `metric`: If multiple routes share the same preference value, the route with the best metric becomes active in the routing table. Use this value to prefer one route over another in this case.
- `preference`: The default preference value of static routes is 5. This preference makes them more likely to be active than OSPF, IS-IS, or BGP for matching prefixes. Use this option to increase the value of the static routes to prefer other sources of routing information.

### Static Route Configuration

- **Static routes require the configuration of a next hop**
  - Valid options are IP address, discard, and reject
- **Defaults section affects all static routes**
- **Qualified next-hop option allows independent preference for static routes to the same destination**
- **Recursive static routes allow you to configure a route to an IP address that is not connected directly to the router**

```
routing-options {
  static {
    defaults {
      preference 250;
    }
    route 192.168.20.0/24 next-hop 10.0.0.1;
    route 192.168.21.0/24 discard;
    route 192.168.22.0/24 reject;
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

### Static Route Requirements

To be committed to the active configuration, static routes must have a next-hop value configured. Often that next-hop value is the IP address of the neighboring router headed toward the ultimate destination. Another possibility is that the next-hop value is the *bit bucket*. This phrase is analogous to dropping the packet off the network. Within JUNOS software, the way to represent the dropping of packets is with the keywords `reject` or `discard`. Both options drop the packet from the network. The difference between them is in the action the router takes after the drop action. If you specify `reject` as the next-hop value, the router sends an ICMP message (that is, the network unreachable message) back to the source of the IP packet. If you specify `discard` as the next-hop value, the router does *not* send back an ICMP message; the router drops the packet silently.

Once in the configuration, static routes appear in the routing table if they are active. Active static routes have a valid next-hop option. Routes with `reject` or `discard` options as next hops always are active and present in the routing table. Routes with an IP address as a next hop are present only if that address is reachable across a directly connected interface on the router.

### Defaults Section

Within the `[edit routing-options static]` configuration hierarchy, the `defaults` section can contain static route options. Any options configured within this section are applied to all static routes on the router. In the example on the slide, we changed the preference value to 250. This value means that all static routes on the router have that preference value.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Qualified Next-Hop Option

The qualified-next-hop option allows independent preference for static routes to the same destination.

### Recursive Static Routes

Recursive static routes allow you to configure a route to an IP address that is not connected to the router directly but can be resolved through the `inet.0` and `inet.3` routing tables. To configure these routes, include the `resolve` statement at the `[edit routing-options static route address]` hierarchy level.

### Aggregate Routes

- Route prefixes in the network can be combined into a single entry in the routing table
- Aggregate routes become active once one or more contributing routes are active
- You configure aggregate routes at the `routing-options` hierarchy level

```
[edit]
routing-options {
  aggregate {
    defaults {
      aggregate-options;
    }
    route destination-prefix {
      policy policy-name;
      aggregate-options;
    }
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

### Combining Aggregate Routes

Aggregate routes are useful for reducing the size of routing tables within an AS. Like static routes, you configure aggregate routes manually on each router in the network.

### Activating Aggregate Routes

Aggregate routes become active in the routing table when at least one of the contributing routes (more specific routes) for the aggregate is also active in the routing table.

### Configuring Aggregate Routes

You configure aggregate routes at the `[edit routing-options]` hierarchy level. The possible options to assign to aggregate routes are:

- `as-path`: Used if this route is intended to be redistributed into BGP and you want to add values manually to the AS path attribute.
- `community`: Used if this route is intended for BGP, and you want to add community values to the route for use in your AS.
- `metric`: If multiple routes share the same preference value, the route with the best metric becomes active in the routing table. Use this value to prefer one route over another in this case.
- `policy`: By default, you can use all possible, more-specific contributing routes to activate an aggregate route. To alter this default, you can use a policy to accept or reject certain routes that should or should not be used.
- `preference`: The default preference value of aggregate routes is 130. Use this option to alter the value of aggregate routes.

### Aggregate Route Configuration

- The default next hop for an aggregate is `reject`
  - `discard` is also a valid option
- Defaults section affects all aggregate routes

```
routing-options {
  aggregate {
    defaults {
      community 1:888;
    }
    route 192.168.16.0/21;
    route 192.168.24.0/21 discard;
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

#### Default Next Hop

The default next-hop value for aggregate routes is `reject`. Much as with static routes, this default means that the router drops the packet from the network and sends an ICMP network unreachable message back to the source of the IP packet. The other possible next-hop value for aggregate routes is `discard`. As before, the router does *not* send back an ICMP message, and it drops the packet silently.

#### Defaults Sections

Within the `[edit routing-options aggregate]` configuration hierarchy, the defaults section can contain aggregate route options. Any options configured within this section are applied to all aggregate routes on the router. In the example on the slide, we set the community value to `1:888`. This value means that all aggregate routes on the router have that community value. Note that you can configure only one aggregate route per prefix.



## Configuring Juniper Networks Routers

### Generated Routes

- Similar to aggregate routes
  - Become active once one or more contributing routes are active
- Often used as the *route of last resort* or *floating static*
- Configured at the `routing-options` hierarchy level

```
[edit]
routing-options {
  generate {
    defaults {
      generate-options;
    }
    route destination-prefix {
      policy policy-name;
      generate-options;
    }
  }
}
```

Copyright © 2008, Juniper Networks, Inc.

### Generated Routes

Generated routes are similar in nature and configuration to aggregate routes. They are configured manually on each router in the network. Generated routes become active in the routing table when at least one of the contributing routes (more specific routes) for the generated route is also active in the routing table. The next hop associated with a generated route is the same next hop as that of the contributing route with the lowest number prefix, not prefix length.

### Route of Last Resort

A generated route is often referred to as a *route of last resort*. This reference is due to one of the uses of a generated route, which is to source a default route.

### Generated Routes Configuration

You configure generated routes at the `[edit routing-options]` hierarchy level.

The possible options to assign to generated routes are:

- `as-path`: Used if this route is intended to be redistributed into BGP and you want to explicitly set the AS path attribute.
- `community`: Used if this route is intended for BGP, and you want to add community values to the route for use in your AS.
- `metric`: If multiple routes share the same preference value, the route with the best metric becomes active in the routing table. Use this value to prefer one route over another in this case.
- `policy`: By default, you can use all possible, more-specific contributing routes to activate a generated route. To alter this default, you can use a policy to accept or reject certain routes that should or should not be used.
- `preference`: The default preference value of generated routes is 130. Use this option to alter the value of the generated routes.

### Generated Route Configuration

- The default next hop is the next hop of the primary contributing route
  - Discard is also a valid option
- Defaults section affects all generated routes

```
routing-options {
  generate {
    defaults {
      metric 5;
    }
    route 172.16.64.0/20;
    route 172.16.80.0/20 discard;
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

#### Default Next Hop

Unlike aggregate routes, the default next-hop value for a generated route is the IP next-hop address of the primary contributing route. The primary route is considered the contributing route with the smallest numerical prefix value (not prefix length). The other possible next-hop value for generated routes is discard. As before, the router does *not* send back an ICMP message, and it drops the packet silently.

Because generated routes are so similar to aggregate routes, they appear in the routing table as an aggregate route with a preference value of 130.

#### Default Section

Within the [edit routing-options generate] configuration hierarchy, the defaults section can contain generated route options. Any options configured within this section are applied to all generated routes on the router. In the example on the slide, we set the metric value to 5. Thus, all generated routes on the router have that metric value. Note that you can configure only one generated route per prefix.

## Contributing Routes

- All routes that fall within a prefix defined for an aggregate or generated route
- Only active routes with a valid forwarding next hop can contribute to a generated route
- A given route can only contribute to one summary
- View contributing routes with `show route protocol aggregate extensive` command

```
lab@London> show route protocol aggregate detail
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
172.16.64.0/20 (1 entry, 1 announced) ← A generated route
  *Aggregate Preference: 130
    Next hop: 10.0.22.1 via so-0/1/1.0, selected
    State: <Active Int Ext>
    Age: 8:37
    Task: Aggregate
    Announcement bits (1): 0-KRT
    AS path: I
    Flags: Generate Depth: 0 Active
    Contributing Routes (1):
      172.16.65.0/24 proto Static ← The contributing route
__juniper_private1__ .inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Copyright © 2006, Juniper Networks, Inc.
```

### Contributing Routes

Contributing routes are important to both aggregate and generated routes (also called *summary routes*). By definition, all routes that share the most common bits with the summary and have a longer bit-mask length are eligible to contribute to the summary.

### Active Routes

Only active routes with a valid forwarding next-hop can contribute to a generated route, because a generated route must be associated with a forwarding next hop. This means that a static route pointing to a discard or reject next hop cannot serve as a contributor to an aggregate or generated route. Any active route can contribute to an aggregate route because an aggregate route always points to a discard (or reject) next hop under the assumption that traffic will match one of the aggregate route's more-specific contributing routes for forwarding purposes.

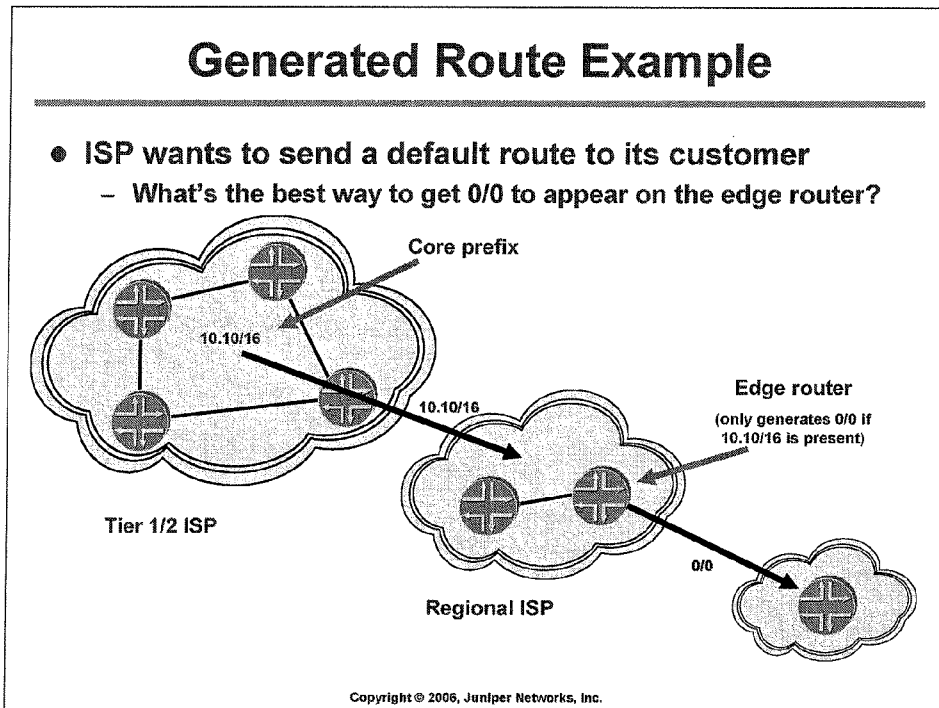
### Summary Route Contribution Limit

While all active routes in the routing table are eligible to contribute to a summary, a single route can contribute to only one summary at a time.

### Viewing Contributing Routes

The `show route extensive` CLI command displays the summary route and all currently contributing routes. In addition to listing the prefixes contributing actively, you can also see which protocol placed that contributing route into the routing table. In the example on the slide, the aggregate route 172.16.64/20 is active in the routing table and has one current contributing route.

## Configuring Juniper Networks Routers



### One Way to Use a Generated Route

One example of using a generated route is an ISP that wants to send a default route to one of its customers. The use of a generated route here is better than a simple static route to establish this default route.

In this example, the regional ISP *could* configure a static route for the 0/0 network on the edge router leading to the customer and redistribute that route to the customer. The problem with using a static route, however, is that the static route is always active on the router, even if no upstream connection exists to which the ISP actually can route the traffic. This method gets the customer traffic to the ISP even if there is no real place to which the ISP can pass off the 0/0 traffic.

Therefore, what would be the easiest way to associate the 0/0 route with upstream connectivity? A generated route with a policy applied might work. In this case, the generated route becomes a type of *dynamic* static route with an IP next hop. By default, a generated route of 0/0 uses all routes in the routing table as potential contributing routes. Thus, the generated route gives you the same net effect as a static route to 0/0, without the drawbacks of the 0/0 static route.

We can use a policy to control which routes from the routing table can contribute to a specific generated route. We can match only that one route in a policy called *only-certain-routes*. Then, we can apply that policy to the generated route within the [edit routing-options generate] section of the configuration hierarchy.

Once we have applied this policy, the 0/0 route only appears in the `inet.0` routing table when the 10.10/16 upstream route is also present. Should the 10.10/16 route disappear, the 0/0 route will also disappear. We have created a sort of *dynamic* static route.

### Martian Addresses

- Address prefixes for which the routers ignore all associated routing information
- Martians are not installed into the routing table
- In JUNOS software, the default martian addresses are:
  - 0.0.0.0/8 orlonger
  - 127.0.0.0/8 orlonger
  - 128.0.0.0/16 orlonger
  - 191.255.0.0/16 orlonger
  - 192.0.0.0/24 orlonger
  - 223.255.255.0/24 orlonger
  - 240.0.0.0/4 orlonger

Copyright © 2008, Juniper Networks, Inc.

#### Martian Addresses

Within the networking environment, *martian addresses* are prefixes that should never be used to route data packets. The router ignores all routing information associated with them. Typically, these are addresses reserved by an Internet authority or are prefixes set aside for special uses (loopback addresses).

#### Martians and the Routing Table

JUNOS software does not allow martian addresses to be placed into the routing table. JUNOS software has seven default martian addresses, which are listed on the slide. The phrase *orlonger* is a policy match term that means that all routes sharing the common prefix bits and having a longer bit-mask length are considered martians also.

#### Default Martian Addresses

Most noticeably missing from this martian list are the default route (0.0.0.0/0) and the three RFC 1918 address ranges: 10/8, 172.16/12, and 192.168/16. This absence is due to the operation of a martian address—it is never allowed into the routing table. It is more likely that a network will want to use these prefixes internally than to never use them. Therefore, they are not included by default in the martian list. Of course, a good Internet service provider (ISP) would not want to receive those prefixes from customers or peers, so they can be filtered out at the edge of the network by using a policy (which is covered in another module).

### Adding Martian Addresses

- Additional prefixes can be added to the martian list
- Configured at the `routing-options` hierarchy level

```
routing-options {
  martians {
    destination-prefix match-type;
  }
}

[edit]
routing-options {
  martians {
    10.0.0.0/8 orlonger;
    172.16.0.0/12 orlonger;
    192.168.0.0/16 orlonger;
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

### Adding Additional Prefixes

You can add or remove prefixes from the default martian list. To see the current list of martian addresses, use the `show route martians` command. On the slide, the martian address list includes the three RFC 1918 private address ranges.

### Configuring Additional Martian Addresses

Within the `[edit routing-options martians]` configuration hierarchy, you can specify additional prefixes. In addition to the prefix and bit mask, you must also include a match type, which is a policy control that determines which exact routes are covered by the configuration statement. The details of policy match types are covered in another module; however, the six default match types are:

- exact;
- longer;
- orlonger;
- prefix-length-range;
- through; and
- upto.

### Routing Tables

- Juniper Networks M-series and T-series platforms have eight predefined routing tables:
  - `inet.0` for unicast routes
  - `inet.1` for the multicast forwarding cache
  - `inet.2` for MBGP routes to provide reverse path forwarding (RPF) checks
  - `inet.3` for MPLS path information
  - `inet6.0` for IPv6 routes
  - `mpls.0` for MPLS next hops
  - `__juniper_private1__.inet.0`
  - `__juniper_private1__.inet6.0`

Copyright © 2006, Juniper Networks, Inc.

#### The Predefined Routing Tables

Juniper Networks M-series and T-series platforms have eight predefined routing tables, which are listed on the slide. You can display routing table information by using the operational CLI command `show route`.

The default routing tables are created and displayed when needed. For example, if a router has no IP information of any kind, then no tables are displayed. As you configure IP addresses and activate routing protocols, then the `inet.0` table is created. If you configure traffic engineering protocols, such as RSVP and MPLS, then the `inet.3` and `mpls.0` tables are created. The table `inet6.0` is not populated until you configure and use IPv6 addresses.

Normal IP packet forwarding is based on information stored in the `inet.0` unicast routing table. You can create additional routing tables for various purposes, such as when configuring multiple routing instances. Using RIB groups, you can copy routes from one table into another. A common example is the use of a RIB group to place interface routes into the `inet.2` table to accommodate DVMRP RPF checks.

JUNOS software creates two private routing instances by default. The software uses these routing instances to communicate between the RE and the FPC, as well as some service PICs, if present. These instances have their own routing tables—one for IPv4 communication, and one for IPv6 communication. As with VPNs, the routes in these tables are completely independent from routes in any other tables. You can safely ignore them when planning your network. If you find the output for these tables in the `show route` command to be annoying, you can add the `table table` switch as in the following example:

```
lab@HongKong> show route table inet.0
```

### Routing Table Protocols

- Within JUNOS software, many sources of routing information exist
  - Referred to as *protocols* in the routing table
- Default protocols:
  - Direct
  - Local
  - Static
  - RSVP
  - LDP
  - OSPF
  - IS-IS
  - RIP
  - Aggregate
  - BGP

Copyright © 2006, Juniper Networks, Inc.

### Sources of Routing Information

As information is added to the routing tables, it is associated with the source of that information. In Juniper Networks language, this source of information is called a *protocol*. (Note: The term *protocol*, as used here, is actually a superset of routing protocols because things like static and aggregate routes are NOT dynamic routing protocols.)

### Default Protocols

The Juniper Networks default protocols for M-series and T-series platforms are:

- *Direct*: The subnet address assigned to a router's interface;
- *Local*: The /32 interface address on a directly connected interface;
- *Static*: Static routes;
- *RSVP*: Resource Reservation Protocol—used for traffic engineering path setup;
- *LDP*: Label Distribution Protocol—used for traffic engineering path setup;
- *OSPF*: Open Shortest Path First—an IGP link-state routing protocol;
- *IS-IS*: Intermediate-System-to-Intermediate-System—an IGP link-state routing protocol;
- *RIP*: Routing Information Protocol—an IGP distance-vector routing protocol;
- *Aggregate*: Aggregate and generated routes; and
- *BGP*: Border Gateway Protocol—an EGP path-vector routing protocol.



### Protocol Preference

- Each protocol has a default preference value
  - Preference is a measure of desirability
    - Used as a tiebreaker when the same prefix is learned through multiple sources; protocols with a lower preference are preferred
- Selected default preference values circa release 6.0:

Direct/Local: 0	...
Static: 5	PIM: 105
RSVP: 7	DVMRP: 110
LDP: 9	Aggregate routes: 130
OSPF internal route: 10	OSPF AS external routes: 150
IS-IS Level 1 internal route: 15	IS-IS Level 1 external route: 160
IS-IS Level 2 internal route: 18	IS-IS Level 2 external route: 165
RIP: 100	BGP: 170
RIPng: 100	MSDP: 175
...	

Copyright © 2006, Juniper Networks, Inc.

### Global Preference Varies by Protocol

It is possible for multiple routing information sources (that is, protocols) to supply the routing table with the exact same prefix. Because only one active route per prefix can exist, the routing table must have a way to differentiate among those choices. This differentiation is the function of the preference value.

Each protocol within JUNOS software is assigned a preference value. This value is a measure of a protocol's *believability*. Preference values can range from 0 through 255, with a lower value being preferred over a higher one.

### Default Preference Values

The slide details the default preference values in JUNOS software as of Release 6.0. Generally speaking, you can change the preference value for a given protocol globally or on a per-route basis through the use of routing policy. The preference of the following route types cannot be modified, however:

- Directly connected networks;
- System routes;
- Redirects;
- Kernel;
- SNMP;
- Router discovery; and
- Down interface routes.

## Configuring Juniper Networks Routers

### The Main Routing Table: `inet.0`

#### Sample `inet.0` routing table for unicast routes:

```
user@host> show route

inet.0: 49 destinations, 49 routes (49 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.11.0/24      *[Direct/0] 1d 08:19:20
                  > via at-0/1/0.100
10.0.11.1/32     *[Local/0] 1d 08:19:20
                  Local
192.168.1.0/24   *[BGP/170] 00:06:08, localpref 100
                  AS path: 1 I
                  > to 10.0.11.2 via at-0/1/0.100
192.168.16.0/21  *[Static/5] 00:02:40
                  Discard
                  [Aggregate/130] 00:36:17
                  Reject
192.168.20.0/24 *[Static/5] 00:06:12
                  Reject
```

Copyright © 2006, Juniper Networks, Inc.

#### Unicast Routing Table

Remember the following points when working with the `inet.0` routing table:

- The routing table is identified, and the total number of routes is displayed (49 routes in this example).
- Prefixes are listed in numerical order.
- Each prefix is added with its protocol assigned, and that protocol's preference value displayed.
- Active routes are noted with an asterisk (\*). Only active routes are placed into the forwarding table.
- If applicable, next-hop information is displayed.

### Load Balancing

---

- **Default is to randomly choose from multiple equal-cost paths on a per-prefix basis**
  - You can change this default to allow multiple next-hop addresses into the forwarding table with a `per-packet` configuration
- **Load-balancing characteristics vary according to Internet Processor version**
  - **Internet Processor I: traffic is balanced according to prefix among up to eight next hops**
    - Per-packet behavior when `per-packet` is configured
  - **Internet Processor II: traffic is balanced according to prefix among up to sixteen next hops**
    - Per-flow balancing when `per-packet` is configured; packets for individual flows are forwarded to a common next hop

Copyright © 2006, Juniper Networks, Inc.

#### Default Behavior

When a single routing protocol (OSPF, for example) provides the routing table with multiple equal-cost paths to a single prefix, the default behavior of the routing table is to pick one of these paths for use in forwarding to that prefix. This results in per-prefix load balancing.

You can change this default behavior using a `per-packet` export policy, which is applied as `export` to the forwarding table at the `[edit routing-options forwarding-table]` hierarchy, to allow multiple equal-cost paths into the forwarding table. The number of paths eligible for balancing and the balancing behavior are specific to the type of Internet Processor ASIC installed in the router. The Internet Processor can handle eight equal-cost paths with per-prefix or per-packet load balancing, while the Internet Processor II can handle up to 16 paths with per-prefix or per-flow balancing.

#### Load-Balancing Behavior Varies

The Internet Processor II ASIC provides load-balancing enhancements over the original Internet Processor. Besides supporting twice as many equal-cost paths (16 vs. 8), the Internet Processor II also supports balancing based on individual traffic flows. A traffic flow is classified as a combination of the following parameters by default: source IP address, destination IP address, source port, destination port, transport protocol, and the incoming router interface.

You can influence the hash algorithm by including a `hash` configuration at the `[edit forwarding-options]` hierarchy. The use of this option is beyond the scope of this course.

### Review Questions

---

1. What is the difference between a static and an aggregate route?
2. What is the purpose of the martian table?
3. Describe the JUNOS software route selection process.
4. List and describe three or more standard JUNOS software routing tables.
5. How would you create multiple static routes to the same destination with independent next-hop preference?

Copyright © 2006, Juniper Networks, Inc.

#### This Module Discussed:

- Static routes;
- Aggregate routes;
- Generated routes;
- The JUNOS software routing tables;
- Active route selection process; and
- JUNOS software load-balancing options.



**Configuring Juniper Networks Routers**

***Module 4: Routing Information Protocol***

Copyright © 2006, Juniper Networks, Inc.

CJNR-M-7.a.7.6.1

## **Module Objectives**

---

- **After successfully completing this module, you will be able to:**
  - **Describe the basic operation of RIP**
  - **Configure RIP**
  - **Monitor and troubleshoot the operation of RIP**

Copyright © 2006, Juniper Networks, Inc.

### **This Module Discusses:**

- The operation of the Routing Information Protocol (RIP);
- RIP configuration in JUNOS software; and
- Using operational-mode commands to monitor and troubleshoot RIP operation.

## Routing Information Protocol

---

- Where we are going...
  - What is RIP?
  - RIP characteristics
  - RIP message types
  - RIPv2 features
  - RIP limitations
  - JUNOS software RIP support
  - Configuring RIP—incoming
  - Configuring RIP— outgoing
  - Useful commands

Copyright © 2006, Juniper Networks, Inc.

### RIP

The slide shows the topics discussed in the following pages.

### What Is RIP?

---

- **RIP is an IGP**
  - Used within an autonomous system
- **Two versions:**
  - RIPv1 (RFC 1058)
  - RIPv2 (RFC 2453)

Copyright © 2006, Juniper Networks, Inc.

#### **RIP Is an Interior Gateway Protocol**

RIP is an interior gateway protocol (IGP) used *within* an autonomous system. RIP advertises routes between devices within the autonomous system.

#### **Two Versions**

Two versions of RIP exist: RIPv1 and RIPv2. RIPv2 did not change the protocol; it expanded the capabilities of it. RFC 1058 defines RIPv1; RFC 2453 defines RIPv2. RIPv2 is the newer of the two protocols. RIPv1 and RIPv2 can interoperate if RIPv1 ignores all fields that must be zero. RIPv2 allows more information to be included in RIP packets and provides a simple authentication mechanism. It also supports VLSM.

RIP is based on the ROUTED program, originally distributed with version 4.3 of the Berkeley Software Division UNIX software. In most UNIX systems, the ROUTED routing daemon dynamically builds the routing table based on information it receives through RIP updates. When routing starts, it issues a request for routing updates and then listens for responses to its request. When a system configured to supply RIP information hears the request, it responds with an update packet based on the information in its routing table. The update packet contains the destination addresses from the routing table and the routing metric associated with each destination. Update packets are not just issued in response to requests, they are also issued periodically to keep routing information accurate.



### RIP Characteristics

---

- Distance-vector routing protocol
- Hop count is used as the metric for path selection, based on Bellman-Ford distance-vector routing algorithm
  - Maximum allowable hop count is 15
- Routing updates are broadcast every 30 seconds

Copyright © 2008, Juniper Networks, Inc.

#### Distance-Vector Routing Protocol

When determining the best path to a destination, RIP uses a combination of hop count (that is, distance) and the next hop (that is, vector).

#### Hop Count

The longest network path in an RIP network is 15 hops between the source and the destination. The assumption here is that the metric count for each network or hop has a cost of one. The 15-hop limitation exists to prevent the creation of an infinitely long network path. With an upper limit of 15 hops, the protocol treats a metric of 16, referred to as *infinity*, to mean that the destination network is unreachable, referred to as *network unreachable*.

When the routing daemon starts on a newly initialized router, it first determines all interfaces that are initialized and then sends them a request packet to each interface asking the router on the other end of the link to send its complete routing table.

#### Routing Updates

Upon receiving an update from another router, the requesting router validates the response and might or might not update its routing table. If updating is required, the update can take the form of adding a route to the table, modifying an existing entry, or deleting an existing entry. Upon receipt of all replies from connected routers, the requesting router builds and updates its routing table.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Routing Updates (contd.)

Each entry in the routing table consists of:

- Network reachability information, the network ID, and the metric;
- Next-hop information;
- The interface through which a packet must pass; and
- A timer indicating the age of a routing entry.

Every 30 seconds, RIP sends all or part of the router's routing table to each of its neighbor's directly connected routers. The routing table is either broadcast to its neighbors on an Ethernet segment or sent to the other end of a point-to-point link. These periodic updates allow a router running RIP to respond to network changes.

RIP also supports triggered updates. A triggered update occurs when a metric changes on a route and can include only the changed entry or entries.

### RIP Message Types

---

- **Two message types**
  - **Request message**
    - Asks neighbors to send routes
  - **Response message**
    - Carries route updates
    - Advertises 25 routes per update
- **Router decides how to handle routes in update**
  - Add, modify, or delete

Copyright © 2006, Juniper Networks, Inc.

### Update Process

A request message asks neighboring routers to send an update, and a response message carries the update from the neighboring routers. When a router receives an update from a neighbor, RIP adds the cost of the network over which the update is received to the advertised metric. The new value is used when comparing routes. RIP stores unknown routes immediately. If a router must advertise more than 25 routes, it must send out an additional response message.

### Route Updates

RIP evaluates known routes by comparing the metric, or cost, of the route presently in the table to the metric of the received route with the following decisions:

- If the cost is lower, RIP adds the new route to the table.
- Where the router advertising the network is the same as that which originally provided it, RIP adopts the route, even where the metric is larger.
- If the advertised hop count is higher than the recorded hop count and the recorded next-hop router originated the update, RIP marks the route as unreachable for a specific hold-down period. At the end of the hold-down period, if the same neighbor is still advertising the higher hop count, RIP accepts the new metric.
- The router can receive both RIPv1 and RIPv2 update messages, with 25 route entries per message. RIP uses timers to enable the router to make the decisions described above.

### RIPv2 Features

---

- **Backward compatible with RIPv1**
- **Authentication on a per-message basis**
  - Simple password or MD5 authentication
- **Multicast updates**
  - Multicast address of 224.0.0.9
  - You can enable broadcast
- **Update includes prefix length**
  - Allows VLSM
- **Update includes next-hop address**
  - Similar function as ICMP redirect
- **RIPv1/v2 interoperability**

Copyright © 2006, Juniper Networks, Inc.

#### Backward Compatibility

RIPv2 is totally *backward compatible* with RIPv1. If a RIPv2 router receives a RIPv1 request message, it should respond with a RIPv1 response message. If you configure the router to send only RIPv2 messages, it should not respond to a RIPv1 request message.

#### Authentication per Message

*Authentication* is possible with RIPv2. The authentication scheme for RIPv2 uses the space of an entire RIP entry. If the address family identifier of the first—and only the first—entry in the message is 0xFFFF, the remainder of the entry contains the authentication. Thus, at most, 24 RIP entries in the remainder of the message can exist. If authentication is not in use, then no entries in the message should have an address family identifier of 0xFFFF. Currently, the only *authentication type* is a simple password, and it is type 2. The remaining 16 octets contain the plain-text password. If the password is under 16 octets, it must be left-justified and padded to the right with nulls (0x00).

#### Multicast Updates

*Multicasting* was added to reduce unnecessary processing of RIP updates by hosts who are not involved in RIPv2 processing. The IP multicast address is 224.0.0.9. On nonbroadcast multi-access networks, like Frame Relay or ATM, you can use unicast addressing.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Prefix Length

RIPv2 can perform *classless routing*, where the prefix length is included in the RIP updates. Another benefit of having a destination prefix length associated with an update is that you can use *variable-length destination prefixes*, thus eliminating the requirement that all destination prefixes in the Internet have the same length.

### Next-Hop Address

With RIPv2, updates also include the *next-hop address*, providing functionality similar to an ICMP redirect. The purpose of the next-hop field is to eliminate packets being routed through extra hops in the system. This field is particularly useful when RIP is not running on all the routers on a network.

### RIPv1/v2 Interoperability

RFC 1723 defines a compatibility mode switch with the following four settings, which allow versions 1 and 2 to interoperate:

- *RIP-1*: Only RIPv1 messages transmit.
- *RIP-1 Compatibility*: Causes RIPv2 to broadcast its messages instead of multicasting them so that RIPv1 hosts can receive them.
- *RIP-2*: RIPv2 messages are multicast to destination address 224.0.0.9.
- *None*: No updates are sent.

### RIP Limitations

- **Limitations:**

- **Maximum network diameter = 15**
- **Regular updates include entire routing table approximately every 30 seconds**
- **Poison reverse increases size of routing updates**
- **Count to infinity slows route-loop prevention**
- **Metrics only involve hop count**
- **Broadcasts between neighbors (RIPv1 only)**
- **Classful routing means no prefix length carried in route updates (RIPv1 only)**
- **No authentication mechanism (RIPv1 only)**
- **Poor convergence**

Copyright © 2006, Juniper Networks, Inc.

### RIP Limitations

- The designers believe that the basic protocol design is inappropriate for larger networks. Assuming a cost of 1, the protocol is limited to networks whose longest path involves 15 hops. If the system administrator chooses to use larger costs, the upper bound of 15 can become a problem easily.
- Updates occur every 30 seconds by default, and the entire routing table is sent in an update. In addition, a triggered update, resulting from a network change, occurs immediately and involves sending the entire routing table.
- Poison reverse aids in network convergence, but it also increases the size of update messages, which include valid and poisoned routes.
- The protocol depends upon counting to infinity to resolve certain unusual situations. Resolving a loop with counting to infinity involves time because a route's metric is increased by two in each update interval, and the loop is only broken when the count reaches 16.
- This protocol uses fixed metrics to compare alternative routes. This method is not appropriate, however, for situations where routes must be chosen based on real-time parameters, such as measured delay, reliability, or load.
- Broadcasting between neighbors forces processing of packets by each host, whether involved in the routing process or not.
- RIP cannot distinguish between subnets. RIPv1 cannot advertise destination prefix lengths; thus, all networks involved in an RIPv1 network must use the same mask.
- RIP provides no authentication mechanism, so a RIP router accepts all RIP-compliant updates.
- Convergence on the network can be slow leading to loops and suboptimal paths.

## **JUNOS Software RIP Support**

---

- **JUNOS software supports:**
  - **RIPv1**
  - **RIPv2**
  - **Peer groups**
    - **Neighbors defined in peer group**
  - **Default setting of no route export**
    - **Need export policy to readvertise RIP or to advertise other routes**
  - **Default preference of 100**
  - **Modification of metrics in or out**

Copyright © 2006, Juniper Networks, Inc.

### **JUNOS Software RIP Support**

The slide shows aspects of RIPv1/v2 supported by JUNOS software.

### Configuring RIP—Incoming

---

#### Minimum RIP configuration

```
protocols {
  rip {
    group group-name {
      neighbor interface-name;
    }
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

#### Minimum RIP Configuration

The minimum configuration, as shown on the slide, starts RIP on the interface specified. After you commit this configuration, this router will understand and evaluate any RIP routes advertised by another router on the same network segment as the interface specified.



### Configuring RIP Advertisements

- Determine which routes to advertise and create export policy

```
policy-options {  
  policy-statement statics-to-rip {  
    from protocol static;  
    then accept;  
  }  
}
```

- Apply export policy to RIP neighbors

```
protocols {  
  rip {  
    group rip-neighbors {  
      export statics-to-rip;  
      neighbor fe-0/0/0.0;  
      neighbor fe-0/0/1.0;  
    }  
  }  
}
```

Copyright © 2006, Juniper Networks, Inc.

#### Advertising and Creating Export Policy

By default, a router with the configuration from the previous page cannot create and advertise any RIP routes that it has not learned from another router. To create and advertise RIP routes to its neighbors, you must configure a router using policy to advertise the specific routes.

#### Exporting Policy

The example on the slide uses policy to advertise this router's static routes using RIP.

## Configuring Juniper Networks Routers

### Monitoring RIP (1 of 4)

- Show the state of your RIP interfaces using the `show rip neighbor` command

Neighbor	State	Source Address	Destination Address	Send Mode	Receive Mode	In Met
fe-0/0/1.0	Up	10.0.31.2	224.0.0.9	mcast	both	1
fe-0/0/0.0	Up	10.0.13.2	224.0.0.9	mcast	both	1

- Show all routes learned via RIP using the `show route protocol rip` command

inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
172.16.0.0/16      [RIP/100] 00:07:02, metric 2
                  > to 10.0.13.1 via fe-0/0/0.0
192.168.8.1/32    *[RIP/100] 00:07:02, metric 2
                  > to 10.0.13.1 via fe-0/0/0.0
224.0.0.9/32      *[RIP/100] 00:11:25, metric 1
...
```

Copyright © 2006, Juniper Networks, Inc.

#### State of RIP Interfaces

The `show rip neighbor` command lists interfaces currently running RIP. The output fields of this command are:

- Neighbor: Displays the name of RIP neighbor.
- State: Displays the state of the connection. The interface can be either up or down.
- Source Address: Displays the source address.
- Destination Address: Displays the destination of RIP updates, which can be either broadcast or multicast.
- Send Mode: Displays the send options, which can be broadcast, multicast, none, or version 1.
- Receive Mode: Displays the type of packets to accept, which can be both, none, version 1, or version 2.
- In Met: Displays the metric added to incoming routes when advertising into RIP routes that were learned from other protocols.

#### RIP Routes

To view the routes in the unicast routing table, issue the `show route protocol rip`. This command filters your routing table and shows only entries learned using RIP.

### Monitoring RIP (2 of 4)

- Display routes that a RIP interface sends using the `show route advertising-protocol rip neighbor` command
  - neighbor is the IP address of local RIP interface
  - Watch for poison reverse behavior when making policy changes!

```
user@host> show route advertising-protocol rip 10.0.21.1

inet.0: 12 destinations, 12 routes (11 active, 1 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

3.0.0.0/8          *[Static/5] 00:10:56
                   Reject
30.0.0.0/16       *[Static/5] 00:12:20
                   Reject
```

Copyright © 2006, Juniper Networks, Inc.

#### Advertised RIP Routes

Use the `show route advertising-protocol rip neighbor` command to view the routes that are advertised out a RIP interface as a result of your RIP export policy. The neighbor argument in this command takes the form of the IP address assigned to the local router's RIP interface.

Note that to help guard against routing loops, the RIP protocol requires that a router continue to advertise a newly unreachable prefix with an infinite metric for a period of time after the route's status changes.

This poison reverse behavior can make it *seem* as though export policy changes are not taking effect because you might see the continued advertisement of prefixes that the current export policy should be rejecting when using the `show route advertising-protocol rip neighbor` command. When you adjust RIP export policy to reject routes previously being accepted, you should expect to see ongoing advertisement of the rejected prefixes for three RIP update cycles (approximately 90 seconds).

### Monitoring RIP (3 of 4)

- Show routes that a RIP interface receives using the `show route receive-protocol rip neighbor` command

– neighbor is the IP address of remote RIP neighbor

```
user@host> show route receive-protocol rip 10.100.3.2
```

```
inet.0: 17 destinations, 18 routes (17 active, 0
holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
172.20.4.0/24      *[RIP/100] 00:01:01, metric 2
                  > to 10.100.3.2 via fe-0/0/0.0
192.168.28.0/24   *[RIP/100] 00:01:01, metric 2
                  > to 10.100.3.2 via fe-0/0/0.0
```

Copyright © 2006, Juniper Networks, Inc.

#### Received RIP Routes

Issue a `show route receive-protocol rip neighbor` command to view the routes being received on a RIP interface from the neighbor address specified. Note that the neighbor argument, in this case, is the IP address of the remote RIP neighbor.

Also note that the routes are displayed *before* your RIP import policy has a chance to manipulate their attributes, but *after* the RIP protocol has discarded nonbest routes. To confirm the operation of your RIP import policy, display the properties of the routes as they reside in the routing table with a `show route protocol rip` command.

### Monitoring RIP (4 of 4)

#### Use the `show rip statistics` command to display various statistics

```
user@host> show rip statistics

RIP info: port 520; update interval 30s; holddown 180s; timeout 120s.
  rts learned  rts held down  rqsts dropped  resps dropped
           1             1             0             0

fe-0/0/0.0: 1 routes learned; 3 routes advertised
Counter      Total      Last 5 min  Last minute
-----
Updates Sent          28          11          2
Triggered Updates Sent  1           0           0
Responses Sent        0           0           0
Bad Messages          0           0           0
RIPv1 Updates Received  0           0           0
RIPv1 Bad Route Entries  0           0           0
RIPv1 Updates Ignored   0           0           0
RIPv2 Updates Received  14          11          3
RIPv2 Bad Route Entries  0           0           0
RIPv2 Updates Ignored   0           0           0
Authentication Failures  0           0           0
RIP Requests Received   0           0           0
RIP Requests Ignored    0           0           0
```

Copyright © 2006, Juniper Networks, Inc.

#### RIP Statistic Gathering

The `show rip statistics` command displays general RIP protocol statistics. The output fields of this command are:

- `RIP info`: Displays the information about RIP on the specified interface.
- `port`: Displays the UDP port number used for RIP.
- `update interval`: Displays the number of seconds since last update.
- `holddown`: Displays the hold-down interval, in seconds.
- `timeout`: Displays the timeout interval, in seconds.
- `bad msgs`: Displays the number of bad messages received.
- `rts learned`: Displays the number of routes learned through RIP.
- `rts held down`: Displays the number of routes held down by RIP.
- `rqst dropped`: Displays the number of request messages dropped by RIP.
- `resp dropped`: Displays the number of response messages dropped by RIP.
- `Counter`: Displays the list of counter types.
- `Total`: Displays the total number of packets for the selected counter.

See the Command Reference Guide for a description of all the various counters.

### Tracing RIP

- A typical RIP tracing configuration:

```
[edit protocols rip]
user@host# show
traceoptions {
  file rip-trace;
  flag error detail;
  flag update detail;
}
```

- Monitor the resulting *rip-trace* log file using `monitor start log-file-name` or the `show log log-file-name` commands

Copyright © 2006, Juniper Networks, Inc.

### RIP Tracing

To perform debugging functions on the RIP routing process, use the JUNOS software `traceoptions` function. The trace output (debug information) is directed to the named log file, which is stored in the `/var/log` directory on the router's hard drive. You can view the log file using the `monitor start` or `show log` operational mode commands.

In addition to specifying the trace file, you also must tell the router what information you want to trace. You accomplish this by specifying one or more `flag` keywords.

While you can only direct tracing to a single file, you can trace many options by using the `flag` keyword multiple times. In addition, you can add granularity by using the `detail`, `receive`, and `send` flag modifiers.

Available tracing flags for RIP include:

<code>all</code>	Trace everything
<code>auth</code>	Trace RIP authentication
<code>error</code>	Trace RIP errors
<code>expiration</code>	Trace RIP route expiration processing
<code>general</code>	Trace general events
<code>holddown</code>	Trace RIP hold-down processing
<code>normal</code>	Trace normal events
<code>packets</code>	Trace all RIP packets
<code>policy</code>	Trace policy processing
<code>request</code>	Trace RIP information packets
<code>route</code>	Trace routing information
<code>state</code>	Trace state transitions
<code>task</code>	Trace routing protocol task processing
<code>timer</code>	Trace routing protocol timer processing
<code>trigger</code>	Trace RIP triggered updates
<code>update</code>	Trace RIP update packets

## Review Questions

---

1. By default, what RIP version does JUNOS software use?
2. Do you need export policy to have RIP export RIP routes learned from other neighbors?
3. Which version of RIP supports VLSM?
4. What command would you use to verify that your router is sending and receiving RIP updates?

Copyright © 2006, Juniper Networks, Inc.

### This Module Discussed:

- The operation of RIP;
- RIP configuration in JUNOS software; and
- Using operational-mode commands to monitor and troubleshoot RIP operation.

## Configuring Juniper Networks Routers

### Lab 4: RIP Configuration

---

#### Lab Objective:

**Configure a Juniper Networks M-series router to run RIP, and monitor its operation**

Copyright © 2006, Juniper Networks, Inc.





**Introduction to Juniper Networks Routers**

***Module 5: Routing Policy***

Copyright © 2006, Juniper Networks, Inc.

CJNR-M-7.a.7.6.1

### Module Objectives

---

- **After successfully completing this module, you will be able to:**
  - **State the purpose of routing policy**
  - **Explain the difference between import and export policies**
  - **Describe the default policy for OSPF, IS-IS, and BGP**
  - **Compare route filter match types**
  - **Write multiterm policies**
  - **Correctly apply policy to BGP**
  - **Use the CLI to monitor policy operation**
  - **Describe advanced policy capabilities**

Copyright © 2006, Juniper Networks, Inc.

#### **This Module Discusses:**

- Policy overview;
- Import versus export policies;
- Default policies for common protocols;
- Route filters and match types;
- The creation of multiterm policies;
- Using operational mode commands to monitor policy operation; and
- Advanced policy capabilities.

## **Routing Policy**

---

- **Where we are going...**
  - Overview
  - When to use policy
  - Import vs. export policy
  - Routing policy flow
  - Generic policy syntax
  - Match conditions
  - Match actions
  - Default policies
  - Policy examples
  - Applying policy
  - Route filters
  - Advanced policy overview

Copyright © 2006, Juniper Networks, Inc.

### **Routing Policy**

The slide shows the topics discussed on the following pages.

### Policy Overview

---

- **Controls routing information transferred into and out of the routing table**
  - Can ignore or change incoming routing information
  - Can suppress or change outgoing routing information
- **Policies are made up of match/action pairs**
  - Match conditions can be protocol specific

Copyright © 2006, Juniper Networks, Inc.

#### Concept of Routing Policy

The concept of routing policy has been around for many years and is not specific to Juniper Networks platforms. Policy is a very powerful tool that lets you manipulate routes that you receive and/or send. In other words, you can manipulate the default decision process of the router by changing route attributes or ignoring and suppressing routes. As we look at policy in more detail, note that policy evaluation is centered on the routing table. Subsequent slides address this fact.

#### Match/Action Pairs

JUNOS software policies are sets of match and action pairs. The match section is a listing of criteria; the action section defines what to do if the match criteria are satisfied.

### When to Apply Policy

---

- **Apply policy when:**
  - You do not want to import all learned routes into the routing table
  - You do not want to advertise all learned routes to neighboring routers
  - You want one protocol to receive routes from another protocol
  - You want to modify information associated with a route

Copyright © 2006, Juniper Networks, Inc.

### Applying Policy

Generically speaking, you use JUNOS software policies when you want to alter the default behavior of the router. More specifically, you might want to filter routing information from a neighbor, filter routes to a neighbor, or redistribute routes between routing protocols.


The filtering of routing information is one major use of the policy language. Based on criteria such as protocols or individual routes, you have the ability to allow or deny information to neighboring routers.

If a situation exists in your networking environment where information from a particular protocol (such as static routes) must be sent to another protocol (such as BGP), you need a policy. Due to the match/action pairing within a policy, you can select the criteria of *all static routes* and the action to perform *send out via BGP* with relative ease.

Lastly, you can alter and modify attribute information within the routes by using a policy. You can change things such as metric values and JUNOS software route preference.

### Import and Export Policies

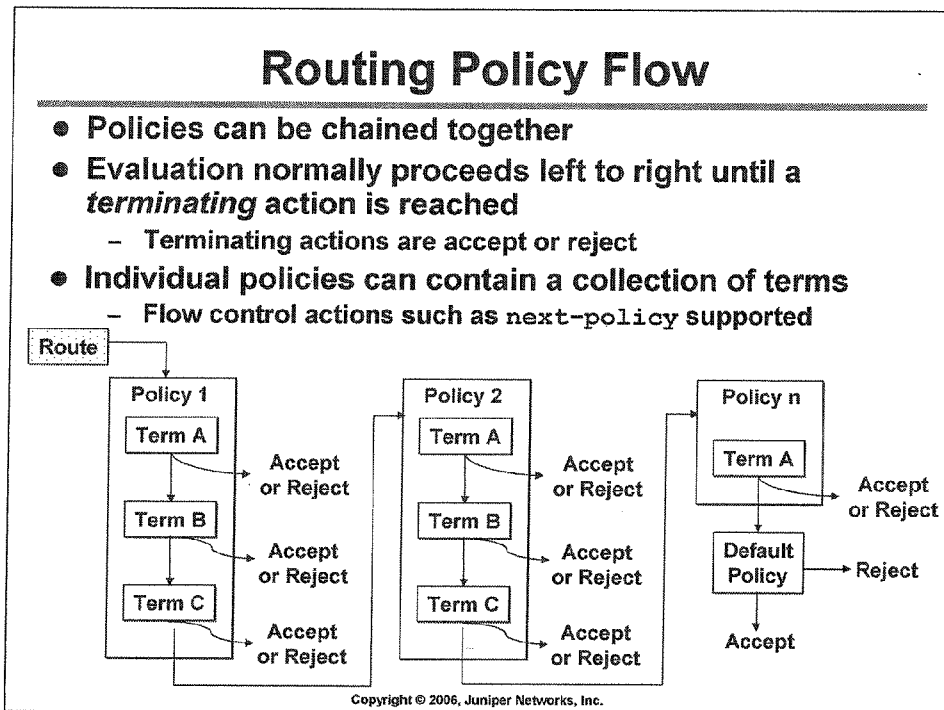
- Perform policy filtering with respect to the JUNOS software routing table
  - JUNOS software applies import policy prior to inclusion in the routing table
  - JUNOS software applies export policy *only* to active routes in the routing table



Copyright © 2006, Juniper Networks, Inc.

### Policy Filtering

All policy processing on Juniper Networks M-series and T-series platforms occurs with respect to the routing table. JUNOS software applies policies as the routing table adds and removes routing information. The keywords *import* and *export* imply the direction of data flow with respect to the routing table.



### Policy Chaining

Policies can be cascaded to form a chain of policy processing. This is often done to solve a complex set of route manipulation tasks in a modular manner.

### Evaluation Process

JUNOS software evaluates policies from left to right based on the order in which they are applied to a routing protocol. JUNOS software checks each policy's match criteria and performs the associated action when a match occurs. If the first policy does not match or if the match is associated with a nonterminating action, the route is evaluated against the next policy in the chain. This pattern repeats itself for all policies in the chain. JUNOS software ultimately applies the default policy for a given protocol when no terminating actions occur while evaluating the user-defined policy chain.

Policy processing stops once a route meets a terminating action, unless you are grouping policies with Boolean operators. Grouping policies for logical operations, such as AND or OR, is a subject that is beyond the scope of this class.

### Individual Policies

Individual policies can be comprised of multiple entries called *terms*. Terms are individual match/action pairs and can be named numerically or symbolically.

JUNOS software lists terms sequentially from top to bottom and evaluates them in that manner. Each term is checked for its match criteria. When a match occurs, JUNOS software performs the associated action. If no match exists in the first term, JUNOS software checks the second term. If no match exists in the second term, JUNOS software checks the third term. This pattern repeats itself for all terms. If no match exists in the last term, JUNOS software checks the next applied policy.

When a match is found within a term, JUNOS software takes the corresponding action. When that action is taken, the processing of the terms and the applied policies stops.

### Generic Policy Syntax

#### Basic policy syntax:

```
policy-options {  
  policy-statement policy-name {  
    term term-name {  
      from {  
        match-conditions;  
      }  
      then {  
        action;  
      }  
    }  
  }  
}
```

A policy  
can have  
multiple  
terms

Copyright © 2006, Juniper Networks, Inc.

#### Basic Policy Syntax

Up to this point, we have been referring to a policy term as a match/action pair. JUNOS software uses the keywords shown below for the match and the action:

- Match = `from`
- Action = `then`



### Match Conditions

- Policies typically contain some form of match criterion
- Possibilities include:
  - Neighbor address
  - Protocol (source of information)
    - BGP, direct, DVMRP, IS-IS, local, MPLS, OSPF, PIM, RIP, static, aggregate
  - Routing protocol information
    - OSPF area ID
    - IS-IS level number
    - BGP attributes
  - Regular expression-based matches for AS path and communities

Copyright © 2006, Juniper Networks, Inc.

#### Match Criterion

Match conditions in a policy are the criteria to be met for a policy action to take place. A policy term can have a single match condition, multiple conditions, or no conditions. The absence of any match conditions means that all possible routes match.

#### Match Conditions

As shown on the slide, match conditions can include the IP address of a neighbor in the AS, routing protocol-specific information (for example, the OSPF area), or the keyword `protocol`. JUNOS software also supports regex matches on AS path and communities. The discussion of regex matches is beyond the scope of this course.

In JUNOS software, the keyword `protocol` has special meaning. It is the *source of routing information*. This keyword is also used within the context of the routing table `inet.0` as to which protocol placed the IP prefix in the table. Within a policy, because we are implicitly referring to the routing table, the keyword `protocol` means the same thing: which protocol in the routing table would you like to reference? Possible protocols to reference in a policy include:

<code>aggregate</code>	<code>atmvpn</code>	<code>bgp</code>	<code>direct</code>	<code>dvmrp</code>
<code>esis</code>	<code>isis</code>	<code>l2circuit</code>	<code>l2vpn</code>	<code>ldp</code>
<code>local</code>	<code>msdp</code>	<code>ospf</code>	<code>pim</code>	<code>rip</code>
<code>ripng</code>	<code>rsvp</code>	<code>rtarget</code>	<code>static</code>	

### Match Actions

---

- The action associated with a given term/policy is performed for matching routes:
  - Terminating actions
    - Accept route
    - Reject (or suppress) route
  - Flow control actions
    - Skip to next policy
    - Skip to next term
  - Modify attributes actions
    - Metric
    - Preference
    - Color
    - Next-hop address

Copyright © 2006, Juniper Networks, Inc.

### Match Actions

When the match criteria are met, any possible action can take place. Within a policy term, you can specify one action, multiple actions, or no action. If you configure no action, policy processing continues with the next term in the policy.

- *Terminate*: The two terminating actions in a policy are *accept* and *reject*. Both of these actions stop the policy processing. If a term has multiple actions within it, the terminating action is completed last.
- *Flow control*: Flow control actions allow you to match some routes from the routing table and move them either to the next term in the policy or the next policy in a string of applied policies. The *next-policy* action is useful when a policy has a *global reject* at the end of it, which normally rejects all routes. This way, a small subset can be matched and moved to further policies for processing while the majority of the routes get rejected.
- *Modify attributes*: Lastly, you can modify routing protocol attributes within a policy action. Often these actions are specified along with a terminating action, such as *accept*. As noted above, JUNOS software performs the *accept* action last so that the modification can take place.

### Default Policies

- **Every protocol has a default policy**
  - The default policy is applied implicitly at the end of the policy chain; can be overridden with `default-action` statement
- **IS-IS and OSPF**
  - **Import: Accept all routes learned from that protocol**
    - Import policy is theoretically invalid for link-state protocols
  - **Export: Accept interface routes, reject all others**
    - Note: With OSPF, interface routes are not subject to policy
- **RIP**
  - **Import all learned RIP routes, export nothing**
    - RIP requires export policy to announce RIP (or other) routes
- **BGP**
  - **Import all routes learned from BGP neighbors**
  - **Export all active routes learned from BGP neighbors to all BGP neighbors**
    - EBGP-learned routes are exported to all BGP peers
    - IBGP-learned routes are exported to all EBGP peers (assumes logical IBGP full mesh)

Copyright © 2006, Juniper Networks, Inc.

#### Default Policy

The default policy always applied to a string of policies sounds very mysterious, but in reality it is not. In fact, every routing protocol that runs on a Juniper Networks M-series or T-series platform always applies the default policy for that protocol. Simply put, the default policy is the default operation of the protocol.

Starting with JUNOS software Release 5.5, you can override the default action intrinsic to a particular protocol by including a `default-action` [accept | reject] within a policy statement. The `default-action` statement is a nonterminating action modifier, which means that subsequent policy statements can continue to evaluate matching routes.

#### IS-IS and OSPF

The default export policy for IS-IS is to accept direct routes for all interfaces on which IS-IS is running and to reject all other routes. OSPF behaves similarly, except you cannot filter interface routes at all (the OSPF protocol uses subnet interconnectivity to describe the topology, so filtering it would break the network). The concept of import policy is invalid for link-state protocols because they use a flooding algorithm to ensure that all routers have identical information. IS-IS does not offer import policy. OSPF does offer import policy, but using it is difficult. The use of import policies in OSPF is beyond the scope of this course, and we strongly discourage you from using it.

#### RIP

The default RIP import policy is to accept all routes learned through RIP. The default export policy advertises no routes, not even those learned through RIP.

#### BGP

The default BGP import policy has all received BGP routes imported into the routing table. For export, all active BGP routes are sent to all peers, with the exception of not sending routes learned through IBGP to other IBGP speakers. This behavior is in accordance with BGP protocol requirements.

### A Policy Example

---

- Write a policy statement at the [edit policy-options] hierarchy:

```
[edit policy-options]
user@host# show policy-statement advertise-ospf
term pick-ospf {
    from protocol ospf;
    then accept;
}
```

- Apply the policy to one or more routing protocol in the import, export, or both directions:

```
[edit protocols bgp]
user@host# set export advertise-ospf
```

Copyright © 2006, Juniper Networks, Inc.

### Policy Configuration Example

As the slide shows, policies work best when you apply them to either a routing protocol or the forwarding table.

As far as the JUNOS software commit function is concerned, it is okay to define a policy under [edit policy-options] and not reference it elsewhere in the configuration. The opposite is *not* true, however; if you reference a policy under a protocol, you *must* configure a policy with the same name under the [edit policy-options] hierarchy.

### Another Policy Example

- Specifying multiple conditions in a `from` statement means that *all* criteria must match before the action is taken

```
[edit]
user@host# show policy-options
policy-statement isis-level2 {
  term find-level2-routes {
    from {
      protocol isis;
      level 2;
    }
    then accept;
  }
}
```

Logical AND Function

- To accomplish a logical OR, use separate terms

Copyright © 2006, Juniper Networks, Inc.

### Multiple Conditions

Within the `from` section of a policy term, you can reference multiple match criteria. In this case, all the match criteria must be satisfied before the action is taken. This is a logical AND function.

It is possible to write and commit a policy term that can never be matched. For example, you could configure a `from protocol static` and `from interface fe-0/0/0` match condition in the same policy term. While this configuration will commit, the term can never be matched because it is impossible for a static route to be associated with an incoming interface. Hence, nothing would match this policy term. This type of match criteria would be sensible for a direct routes, because direct routes are associated with an interface.

### Logical OR Function

You can accomplish a logical OR function by writing individual terms with different `from` criteria but identical `then` actions:

```
[edit policy-options policy-statement igp]
lab@SanJose# show
term isis {
  from protocol isis;
  then accept;
}
term ospf {
  from protocol ospf;
  then accept;
}
```

### Applying Policy

- You must apply policies before they can take effect
- Link-state protocols (IS-IS and OSPF) have only export filtering points
- BGP and RIP support both import and export policies

```
[edit protocols]
user@host# show
bgp {
    import bgp-import;
    export bgp-export;
}
ospf {
    export ospf-export;
}
```

Copyright © 2006, Juniper Networks, Inc.

#### Applying Policies

You can apply policies as import or export policies, as described earlier. To apply a policy, you must specify an `export` or `import` statement that references the desired policy.

#### Link-State Protocols

When the goal of a policy is to affect IS-IS or OSPF routes, you must apply the policy to that protocol. Link-state protocols support export policies only and do not allow the filtering of LSA/LSPs, as this behavior could violate the standards for the link-state protocol in question. You accomplish LSA/LSP filtering with configuration for a specific protocol, such as a multilevel IS-IS configuration that filters Level 2 LSPs from Level 1 areas.

Link-state routing protocols support export policy applications at the global level only; thus, the policy will apply to all neighbors/adjacencies.

As mentioned previously, OSPF does allow a form of import policy, but it does, in fact, conflict with the design of the protocol, and you should avoid using it.

#### BGP and RIP

The BGP and RIP protocols support both import and export policy applications. You can also apply these policies at the global, group, or neighbor levels in the case of BGP. For RIP, you can apply import policy at the global, group, or neighbor levels. You can only apply export policy at the group level, however.

## Apply Routing Policy to BGP

- BGP has three filtering points per direction:
  - Global
  - Groups of neighbors
  - Individual neighbors
- Only the *most* specific policies are applied to a particular peer
  - Neighbor policy overrides group and global policies
  - Group policy overrides global policy

Copyright © 2006, Juniper Networks, Inc.

### Filtering Points

BGP policies are *not* hierarchical. You can apply policy for BGP at the global, group, and neighbor levels.

### BGP Policy Evaluation

If you configure single or multiple policies at the neighbor level, any group-level policies will never be evaluated. In this case, you must reference the same policy in both places if you want both levels to evaluate it.

## Configuring Juniper Networks Routers

### BGP Policy Application Example

```
[edit protocols]
user@host# show
bgp {
  export local-customers;
  group meganet-inc {
    type external;
    import [ martian-filter long-prefix-filter as-47-filter ];
    peer-as 47;
    neighbor 1.2.2.4;
    neighbor 1.2.2.5;
  }
  group problem-child {
    type external;
    import [ as-47-filter long-prefix-filter martian-filter ];
    export kill-private-addresses;
    peer-as 54;
    neighbor 1.2.2.6;
    neighbor 1.2.2.7;
    neighbor 1.2.2.8 {
      import [ reject-unwanted as-666-routes ];
    }
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

### BGP Policy Application

The following list explains the details of the slide:

- Peer 1.2.2.4 will use the import policies *martian-filter*, *long-prefix-filter*, and *as-47-filter*. It will use the export policy *local-customers*.
- Peer 1.2.2.6 will use the import policies *as-47-filter*, *long-prefix-filter*, and *martian-filter*. It will use the export policy *kill-private-addresses*.
- Peer 1.2.2.8 will use the import policies *reject-unwanted* and *as-666-routes*. It will use the export policy *kill-private-addresses*.

Generally speaking, you can evaluate BGP policies (as well as other attributes) in the following way. Look at the neighbor level for configuration. If there is *nothing* at the neighbor level, go to the group level for configuration. If there is *nothing* at the group level, go to the global level for configuration. Put another way, JUNOS software always implements the most specific aspect of a configuration, and in the case of BGP, the most specific configuration item is applied at the exclusion of less specific items.



### Route Filters

- Use route filters to match an individual route (or groups of routes)
  - You can specify multiple route filters within a single term
  - General syntax in the form of:  
`route-filter prefix/prefix-length match-type actions;`
- Route filter evaluation has special rules according to the match type
  - Match types specify different sets of routes:
    - exact
    - orlonger
    - longer
    - upto
    - through
    - prefix-length-range
  - Policy test function is useful for route-filter debugging

Copyright © 2006, Juniper Networks, Inc.

#### Route Filters

Use route filters as a match condition in a policy where the goal is to match criteria for a particular route or group of routes. You can have multiple route filters as a match condition. A group of route filters in a single match condition of a term in a policy are evaluated as a longest match, where only one of the route filters can result in a match.

#### Route Filter Evaluation

The following slides describe each of the different match types listed on the slide. You can use the CLI operational-mode `test` command to test a policy for matches against the active routes in the main routing table. The use of the `test` command is beyond the scope of this class.

### Route Filter Match Types (1 of 2)

- **exact**

- Match the specified prefix and mask exactly
- No other routes will be included

```
from route-filter 192.168/16 exact;
```

- **orlonger**

- Match the specified prefix and mask exactly
- Also match any routes that start with the same prefix and have longer masks

```
from route-filter 192.168/16 orlonger;
```

- **longer**

- Do not match the specified prefix and mask exactly
- Match only the routes that start with the same prefix and have longer masks

```
from route-filter 192.168/16 longer;
```

Copyright © 2006, Juniper Networks, Inc.

#### **exact**

The match type `exact` means that only routes that match the given prefix exactly will pass the filter statement. For example, on the slide, only the prefix `192.168/16`, and no other prefixes, will pass the filter statement.

#### **orlonger**

The match type `orlonger` means that routes greater than or equal to the given prefix will pass the filter statement, so the exact route `192.168/16` on the slide will match the statement. In addition, all routes that start with `192.168` and have bit-mask lengths between `/17` and `/32` will also pass. For example, the following prefixes match the statement: `192.168/16`, `192.168.65/24`, `192.168.24.89/32`, `192.168.128/18`, and `192.168.0/17`. The following prefixes do not match the statement: `10.0/16`, `192.167.0/17`, and `200.123.45/24`.

#### **longer**

The match type `longer` means that only routes greater than the given prefix will pass the filter statement, so from the example on the slide, all routes that start with `192.168` and have bit-mask lengths between `/17` and `/32` will pass the filter statement. The following prefixes match the statement: `192.168.65/24`, `192.168.24.89/32`, `192.168.128/18`, and `192.168.0/17`. The following prefixes do not match the statement: `10.0/16`, `192.167.0/17`, `200.123.45/24`.

### Route Filter Match Types (2 of 2)

- **upto**
  - Match the specified prefix and mask exactly
  - Also match any routes that start with the same prefix and have a mask no longer than the second value specified

```
from route-filter 192.168/16 upto /24;
```
- **through**
  - Match the first specified prefix and mask exactly
  - Match the second specified prefix and mask exactly
  - Match all prefixes *directly* between the two prefixes

```
from route-filter 192.168/16 through 192.168.16/20;
```
- **prefix-length-range**
  - Match only routes that start with the same prefix and have a mask between the two values specified (inclusive match)

```
from route-filter 192.168/16 prefix-length-range /20-/24;
```

Copyright © 2006, Juniper Networks, Inc.

#### **upto**

The match type `upto` means that routes greater than or equal to the first specified prefix, but less than or equal to the second specified prefix, will pass the filter statement. Thus, using the example on the slide, the exact route 192.168/16 will match the statement. All routes that start with 192.168 and have bit-mask lengths between /17 and /24 will also pass. The following prefixes match the statement: 192.168/16, 192.168.65/24, 192.168.128/18, and 192.168.0/17. The following prefixes do *not* match the statement: 192.168.24.89/32, 10.0/16, 192.167.0/17, and 200.123.45/24.


#### **through**

The match type `through` represents a string of exact matches. In other words, instead of specifying multiple exact route-filter statements, you can use the `through` statement. However, the `through route-filter` statement requires that the exact matches follow a specific pattern described by the prefixes given. In the example on the slide, the first specified route in the statement is matched exactly (192.168/16), and the second specified route in the statement is matched exactly (192.168.16/20). Then prefixes along the path in a radix-like tree, called the *J-tree*, also match the policy statement. The following prefixes also match the statement: 192.168.0/17, 192.168.0/18, 192.168.0/19. All other prefixes do *not* match the statement.

#### **prefix-length-range**

The match type `prefix-length-range` works very much like the match type `upto`. The difference is that you are specifying both a starting bit-mask length and an ending length. So, all routes that start with 192.168 and have bit masks between /20 and /24 will pass the filter statement. In the example on the slide, the following prefixes match the statement: 192.168.0/20, 192.168.64/24, 192.168.128/21. The following prefixes do *not* match the statement: 192.168.24.89/32, 10.0/16, 192.167.0/17, 200.123.45/24, and 192.168.128/18.

## Configuring Juniper Networks Routers



### Prefix Match

The slide shows a graphical summary of each of the route-filter match types.

### Route Filter Actions

```
term term-name {  
  from {  
    route-filter dest-prefix match-type actions;  
    route-filter dest-prefix match-type actions;  
  }  
  then actions;  
}
```

] Longest-Match  
Lookup

- **Only one route filter in a given term can be considered a match**
  - Longest-match lookup is performed on the prefix being evaluated
- **If an action is specified to a route filter, it takes effect immediately**
  - The global **then** portion of the term is ignored
    - If specific actions are not defined, the **then** portion of the term is executed for matching prefixes

Copyright © 2008, Juniper Networks, Inc.

#### Multiple Route Filters

Although having multiple match criteria in a single term is a logical AND function, the presence of multiple route filters is different. In this case, JUNOS software evaluates the prefixes on a route-by-route basis and performs a longest-match lookup. This lookup is similar to the one done by the forwarding table. Once that longest match is found, JUNOS software only evaluates that one route filter to determine a match for the policy term.

If there are multiple route filters and other match criteria in a single policy term, the one longest-match route filter must still be applied to the logical AND process with the other match criteria for the action to be taken.

#### Specified Route Filter Actions

You can specify policy actions at the route-filter level in addition to the policy term level. When there is an action at the route-filter level, that action is taken above all else, and the policy-term action is ignored. If there is not a route-filter action specified, JUNOS software applies the policy term.

### Test Your Knowledge (1 of 2)

Which action is taken when this policy evaluates 10.0.67.43/32?

```
[edit policy-options policy-statement pop-quiz]
user@host# show
from {
    route-filter 10.0.0.0/16 orlonger accept;
    route-filter 10.0.67.0/24 orlonger;
    route-filter 10.0.0.0/8 orlonger reject;
}
then {
    metric 10;
    accept;
}
```

Copyright © 2006, Juniper Networks, Inc.

#### Policy Evaluation

While all three route-filter prefixes match the candidate route, the middle route filter of 10.0.67/24 is the longest match. Because this route filter does *not* have an action specified at its configuration level, JUNOS software applies the policy. This route has its metric changed to 10 and is accepted.

### Test Your Knowledge (2 of 2)

Which action is taken when this policy evaluates 10.0.55.2/32?

```
[edit policy-options policy-statement pop-quiz]
user@host# show
from {
    route-filter 10.0.0.0/16 orlonger accept;
    route-filter 10.0.67.0/24 orlonger;
    route-filter 10.0.0.0/8 orlonger reject;
}
then {
    metric 10;
    accept;
}
```

Copyright © 2008, Juniper Networks, Inc.

### Policy Evaluation

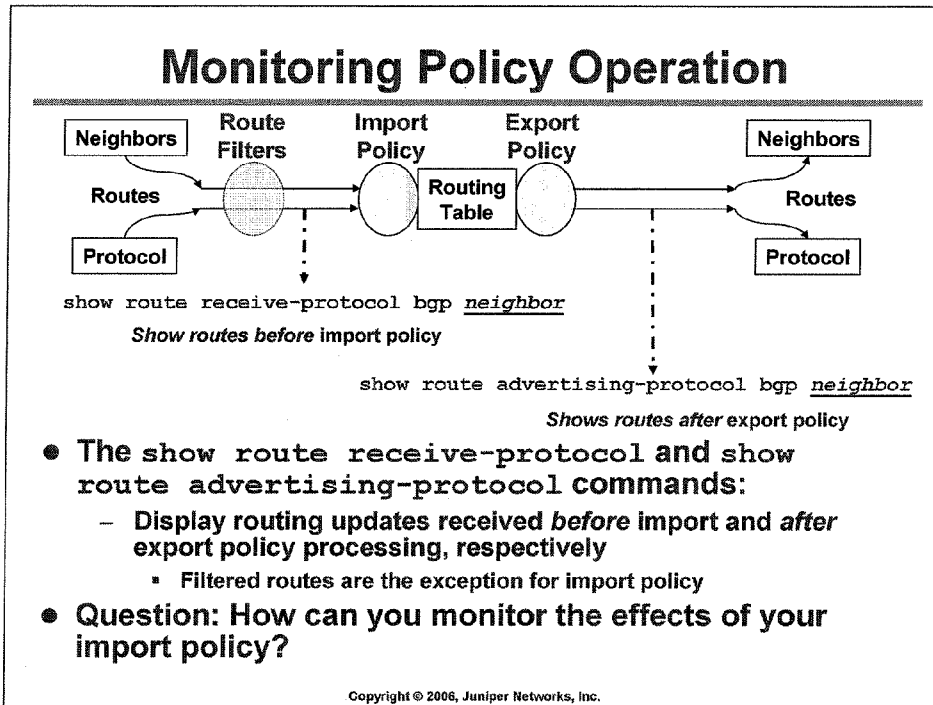
On this slide, only two of the route-filter prefixes match the candidate route: 10.0/16 and 10/8. The longest match for the candidate route is the 10.0/16 route filter. Because this route filter does have an action specified at its configuration level, JUNOS software takes that action and then ignores the policy. This route will be accepted with no other modifications made to its attributes.

For extra credit, what will happen to the 10.0.55.2/32 prefix if the 10.0/16 route filter statement has its match type changed to *exact*, as shown:

```
[edit policy-options policy-statement pop-quiz]
user@host# show
from {
    route-filter 10.0.0.0/16 exact accept;
    route-filter 10.0.67.0/24 orlonger;
    route-filter 10.0.0.0/8 orlonger reject;
}
then {
    metric 10;
    accept;
}
```

In this case, the longest match is still the 10.0/16 router filter statement. However, because the match type is now *exact*, the prefix does not match the 10.0/16 statement. Because each prefix is evaluated only against the longest match within a particular term, this prefix matches nothing in the term. Therefore, it will continue to be evaluated by any remaining policies (possibly the default policy).

## Configuring Juniper Networks Routers



### Monitoring Effects of Policy

The commands on the slide show routing updates received before import policy processing and the routing updates sent after export policy processing.

Use the `show route receive-protocol protocol neighbor` command to show the specified protocol-type route advertisements that a particular neighbor is advertising to your router *before* import policy is applied. Use the `show route advertising-protocol protocol neighbor` command to show the protocol-type route advertisements that you are advertising to a particular neighbor *after* export policy is applied.

The use of route filters marks an exception to the behavior documented above. JUNOS software evaluates route filters before the output of a `show route receive-protocol` command is generated. This means that you must specify the hidden `switch` to the `show route receive-protocol` command to display received routes filtered by your import policy.

### Answer

After import policy processing, use the `show route protocol protocol` command to monitor the effects of your import policy. This command shows all routes from the protocol type specified that are installed in the routing table.



## Review Questions

---

1. What is the purpose of routing policy?
2. The terms import and export are based on the perspective of which entity within the router?
3. How does the default policy for OSPF differ from that of BGP?
4. What types of match conditions are supported in policy?
5. What types of match actions can you use in policy?
6. Explain the difference between applying policy at the global, group, and peer levels of BGP.
7. What command would you use to monitor the effects of your import policy?

Copyright © 2006, Juniper Networks, Inc.

### This Module Discussed:

- Policy overview;
- Import versus export policies;
- Default policies for common protocols;
- Route filters and match types;
- The creation of multiterm policies;
- Using operational mode commands to monitor policy operation; and
- Advanced policy capabilities.

## **Lab 4: Routing Policy**

---

### **Lab Objective:**

**Configure routing policy on your router using JUNOS software. You will complete this lab by configuring a policy to the RIP configuration left in place from the last lab.**

Copyright © 2006, Juniper Networks, Inc.



## **Configuring Juniper Networks Routers**

### ***Module 6: OSPF Operation, Configuration, and Troubleshooting***

Copyright © 2006, Juniper Networks, Inc.

CJNR-M-7.a.7.8.1

## Configuring Juniper Networks Routers

### Module Objectives

---

- After successfully completing this module, you will be able to:
  - Explain general OSPF operation
  - Describe the role of a DR
  - List and describe three types of OSPF areas
  - Configure OSPF
  - Monitor and troubleshoot OSPF
  - Display and interpret the LSDB

Copyright © 2006, Juniper Networks, Inc.

#### **This Module Discusses:**

- General OSPF operation;
- The role of the designated router (DR);
- OSPF area types and rationale;
- Configuring OSPF in JUNOS software;
- Using operational mode commands to monitor and troubleshoot OSPF; and
- Displaying and interpreting the link-state database (LSDB).

## **OSPF Agenda**

---

- **Where we are going...**
  - **OSPF general review**
  - **Protocol scalability**
  - **Adjacency formation and designated router election**
  - **OSPF configuration**
  - **OSPF monitoring and troubleshooting**

Copyright © 2006, Juniper Networks, Inc.

### **OSPF Agenda**

The slide shows the topics examined on the following pages.

### OSPF Protocol Review

---

- OSPF is a link-state IGP that routes packets within a single AS
- OSPF reliably floods LSAs to distribute link-state information once an adjacency is formed
- Each router uses these LSAs to create a complete database for the network
- OSPF uses the SPF algorithm within the database to calculate the best route to every node in the network
- OSPF is defined in:
  - RFC 2328, *OSPF Version 2*
  - RFC 1587, *The OSPF NSSA Option*

Copyright © 2006, Juniper Networks, Inc.

#### Link-State Protocol

OSPF is a link-state routing protocol designed for use within an autonomous system (AS). It is considered an interior gateway protocol (IGP). Link-state protocols allow for faster reconvergence, support larger internetworks, and are less susceptible to bad routing information than distance vector protocols.

#### LSA Flooding

Routers running OSPF send out information about their network links and the state of those links to other routers in the AS. This information is transmitted reliably to all other routers in the AS via link-state advertisements (LSAs). The other routers receive this information and store it locally on each router. This total set of information now contains all possible links in the network.

#### Link-State Database

In addition to flooding LSAs and discovering neighbors, a third major task of the link-state routing protocol is establishing the link-state database. The link-state, or topological, database stores the LSAs as a series of records. The important information for the shortest path determination process is the advertising router's ID, its attached networks and neighboring routers, and the cost associated with those networks or neighbors.

#### Shortest Path First Algorithm

OSPF uses the shortest path first (SPF) algorithm or Dijkstra algorithm to calculate all at once the shortest paths to all destinations. It does this calculation by calculating a tree of shortest paths incrementally and picking the best candidate from that tree.

#### Standards

RFC 2328 defines OSPF version 2; RFC 1587 defines the OSPF NSSA option.

### The Link-State Database

---

- Every router in an OSPF network maintains a copy
  - LSDB on all routers must match within an area
  - Made up of received link-state advertisements
  - Contains all known information
- SPF algorithm uses the LSDB as input data to calculate network paths

Copyright © 2006, Juniper Networks, Inc.

#### Routers Maintain Identical LSDBs

Each router in an OSPF area stores received link-state updates in a database of information known as the link-state database (LSDB). This database contains knowledge of all known network links in the network. According to the OSPF RFC, each router in the area must have an identical LSDB to ensure accurate routing knowledge. The router uses the LSDB as input to the SPF algorithm to determine the shortest route (that is, least cost) to each other node in the network.

#### SPF Algorithm

The OSPF router runs the SPF algorithm against this common data store. The results of this computation are then handed to the router's routing table for the actual forwarding of data packets.

### OSPF Packet Types

---

- **Five OSPF packet types**
  - Hello packets (Type 1)
  - Database description packets (Type 2)
  - Link-state request packets (Type 3)
  - Link-state update packets (Type 4)
  - Link-state acknowledgment packets (Type 5)
  - **Link-state advertisements are flooded reliably using:**
    - Link-state requests
    - Link-state updates
    - Link-state acknowledgments

Copyright © 2006, Juniper Networks, Inc.

#### **Packet Types Used in OSPF**

The slide lists the OSPF packet types.



### OSPF Hello Packets

- Hello packets

- Routers periodically send multicast hello packets out all OSPF interfaces to establish and maintain neighbor relationships
- Sent to 224.0.0.5—all OSPF routers address
- Hello packets consist of the OSPF header plus the following fields:
  - Network mask\*
  - Hello interval\*
  - Dead interval\*
  - Options\*
  - Router priority
  - Designated router
  - Backup designated router
  - Neighbor

\* Fields that must match to form an adjacency over a broadcast medium; a matching network mask is not required for point-to-point links

Copyright © 2006, Juniper Networks, Inc.

### OSPF Hello Packets

All OSPF routers send hello packets on all their links on a regular cycle, which is 10 seconds by default. The hello packet is multicast to the all OSPF routers multicast address of 224.0.0.5. In addition to the OSPF header, OSPF includes additional information specific to that link.

Some of the link-specific information must match between neighbors before they can form an adjacency. On the slide, an asterisk (\*) marks the fields that must match to form an adjacency. The following list contains the details of these fields:

- *Network mask*: This field is evaluated only on broadcast media links (that is, Ethernet). All routers on the segment must agree on the subnet mask of the link.
- *Hello interval*: The two routers must agree on how often to send hello packets, which this field determines.
- *Dead interval*: Also referred to as the keepalive, this timer states how long to wait before removing an adjacency to a neighbor. The two routers must agree on this timer.
- *Options*: This 8-bit field represents such things as the ability to be a stub area. These options are critical to correct OSPF operation, so they, too, must match between neighbors.

### Database Description Packets

---

- **Exchanged during adjacency formation**
  - Determine which router is in charge of the database exchange
  - Describe the contents of the link-state database
- **Consist of:**
  - OSPF header
  - Sequence number
  - LSA headers

Copyright © 2006, Juniper Networks, Inc.

### Database Description Packet Exchange

OSPF uses database description (DD) packets only during the adjacency formation process between two OSPF routers. The DD packets serve two main purposes: determining who is in charge of the database synchronization, and actually transferring the LSA headers between the two systems.

For database synchronization, the two OSPF routers must decide on who is in charge. The router with the highest router ID (RID) always is the router in charge. This router is known as the *master*; the other router is the *slave*. Simply put, it is the job of the master to set and maintain the sequence numbers used during the database transfer. After the transfer is complete, the knowledge of the master/slave relationship is forgotten.

OSPF also uses DD packets actually to transmit the LSA headers between the systems. This transmission serves as the minimum synchronization between the routers.

### Database Description Packet Fields

The details of the DD packet fields are:

- *OSPF header*: This header is 24 bytes.
- *Sequence number*: This field ensures that the full sequence of DD packets is received in the database synchronization process. The sequence number is set by the master to some unique value in the first DD packet, and the sequence is incremented in subsequent packets.
- *LSA header*: This header lists some or all of the headers of the LSAs in the originator's link-state database. The header contains enough information to uniquely identify the LSA and the particular instance of the LSA.

### Link-State Request Packets

---

- Sent when router detects database is stale
  - Request precise version of database
- Consist of:
  - OSPF header
  - Link-state type
  - Link-state ID
  - Advertising router

Copyright © 2006, Juniper Networks, Inc.

#### Link-State Request Packets

After receiving a number of DD packets, a router might find that its neighbor is advertising an LSA header that is not currently in its own database. In that situation, the receiving router sends out a link-state request packet to the sending router. In simplest terms, the link-state request packet is the request for information; it contains the LSA header for the missing link.

#### Link-State Request Packet Fields

The following list contains the details of the link-state request packet fields:

- *OSPF header*: This header is 24 bytes.
- *Link-state type*: This field contains the LSA type number, which identifies the type of LSA (for example, a router LSA or network LSA).
- *Link-state ID*: This field is type-dependent on the LSA header.
- *Advertising router*: This field contains the router ID of the router that originated the LSA.

### Link-State Update Packets

- **Carry one or more LSAs**
  - **Multicast on physical networks that support multicast or broadcast mode**
    - 224.0.0.5 (all OSPF routers address) or 224.0.0.6 (all DRs address)
  - **All received link-state updates are acknowledged**
    - Retransmissions are sent unicast
- **Consist of:**
  - **OSPF header**
  - **Number of advertisements**
  - **Link-state advertisements**

Copyright © 2006, Juniper Networks, Inc.

### Link-State Update Packets

A link-state update packet is the basic information block in OSPF. It can contain multiple LSAs. Link-state update packets are reliably transmitted via multicast to either the all OSPF routers address (224.0.0.5) or the all DRs address (224.0.0.6), depending on the link type.

OSPF sends link-state update packets in two different ways: in response to a link-state request packet during the adjacency database synchronization, and after an adjacency is formed, if information about that link changes.

### Link-State Update Packet Fields

The following list provides the details of the link-state update packet fields:

- *OSPF header*: This header is 24 bytes.
- *Number of advertisements*: This field specifies the number of LSAs included in this packet.
- *Link-state advertisements*: This field contains the full LSAs as described in OSPF LSA formats. Each update can carry multiple LSAs, up to the maximum packet size allowed on the link.

### Link-State Acknowledgment Packets

---

- Sent in response to link-state update packets
  - Acknowledge successful receipt of the update packets
  - Can include responses to multiple update packets
- Consist of:
  - OSPF header
  - LSA header

Copyright © 2006, Juniper Networks, Inc.

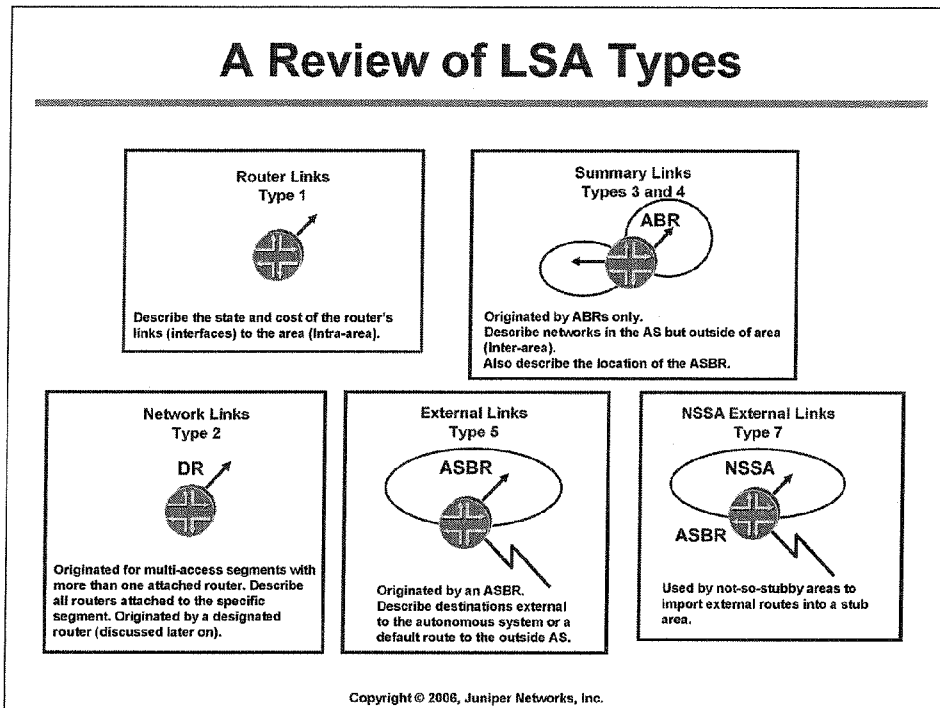
#### Link-State Acknowledgement Packets

OSPF sends link-state acknowledgment packets after the receipt of a link-state update packet. The link-state acknowledgement packet is sent back in unicast fashion to the originating system and ensures it that the link-state update packet was received. This acknowledgement is the basis for the reliable flooding in OSPF. An individual link-state acknowledgement packet can contain an acknowledgment for a single link-state update packet or for multiple link-state update packets.

#### Link-State Acknowledgement Packet Fields

The link-state acknowledgement packet consists of nothing more than an OSPF packet header and a list of LSA headers.

## Configuring Juniper Networks Routers



### A Review of LSA Packet Types

The OSPF specification details 11 different LSA types (listed below). These LSAs each represent a portion of the OSPF network and are flooded into the AS based on the function of the router. Modern OSPF networks use the LSA types marked with an asterisk (\*).

- Type 1: Router LSA\*
- Type 2: Network LSA\*
- Type 3: Summary LSA\*
- Type 4: ASBR summary LSA\* (also referred to as ASBR reachability LSA)
- Type 5: AS external LSA\*
- Type 6: Multicast OSPF LSA
- Type 7: NSSA external LSA\*
- Type 8: External attributes LSA
- Type 9: Opaque LSA (link scope)
- Type 10: Opaque LSA (area scope—used for traffic engineering)\*
- Type 11: Opaque LSA (AS scope)

## OSPF Scalability

---

- Where we are going...
  - Protocol scalability
  - OSPF areas
  - OSPF router terminology
  - OSPF area relationships
  - OSPF area types

Copyright © 2006, Juniper Networks, Inc.

### OSPF Scalability

The slide shows the topics examined on the following pages.

### Protocol Scalability

---

- **Every router must flood LSAs and maintain an identical LSDB**
  - As more routers are added to the network, these requirements place a load on the protocol
  - Resources are spent constantly flooding new information and rerunning the SPF calculation
- **Solution is to shrink the size of the LSDB**
  - Network needs hierarchy
  - Creation of OSPF areas

Copyright © 2006, Juniper Networks, Inc.

#### OSPF Scalability

With a link-state protocol, flooding link-state update packets and running the SPF algorithm consumes router resources. As the size of the network grows and more routers are added to the AS, this use of resources becomes a bigger and bigger issue. This issue stems from the RFC requirement that all OSPF routers share an identical LSDB. Eventually, the routers spend so much time flooding and running the SPF algorithm that they cannot route data packets. Clearly, this scenario is not optimal.

#### Shrinking the Link-State Database

The solution to this problem (and link-state protocol scalability in general) is to reduce the size of the LSDB through the segmentation of the network. OSPF accomplishes this segmentation by using areas. The LSDB is now reduced to the size of an individual OSPF area—and not the entire AS.



### OSPF Areas

---

- **Areas**
  - Single AS can be divided into smaller groups called *areas*
  - Reduces the LSDB because LSA flooding is now constrained to the area
  - Routers maintain a separate LSDB on a per-area basis
  - Each LSDB within an area still must be identical on all routers
- **Special OSPF area called the *backbone* area**
  - Backbone area (0.0.0.0) distributes routing information between areas
  - All other OSPF areas must connect to the backbone area
  - All user traffic from one area to another must traverse the backbone

Copyright © 2006, Juniper Networks, Inc.

#### OSPF Areas

Using areas achieves the OSPF hierarchy. As mentioned previously, areas reduce the size of the LSDB on an individual router. Now, each router maintains a separate LSDB for each area to which it is connected.

#### Backbone Area

To ensure correct routing knowledge and connectivity, OSPF maintains a special area called the backbone area. It is designated as area 0. All other OSPF areas must connect themselves to the backbone for connectivity. All data traffic between OSPF areas must transit the backbone.

### OSPF Router Terminology

---

- **Internal router** has all OSPF links in the same area
  - Within area 0, also called a *backbone router*
- **Backbone router**
  - Any router with a link to area 0
- **Area border routers (ABRs)**
  - Routers that belong to more than one area are called *area border routers*
  - Connect OSPF areas to the backbone area 0
- **Autonomous system boundary routers (ASBRs)**
  - Routers that inject routing information from outside the OSPF domain are called *AS boundary routers*

Copyright © 2006, Juniper Networks, Inc.

#### **Internal Routers**

An OSPF router with all its links within an area is known as an internal router. If that router is located within the backbone area (0.0.0.0), it is also known as a backbone router.

#### **Backbone Routers**

Any OSPF router with a link to area 0 (the backbone) is considered to be a backbone router. This router may also be an internal or area border router depending on whether it has links to other, nonbackbone areas.

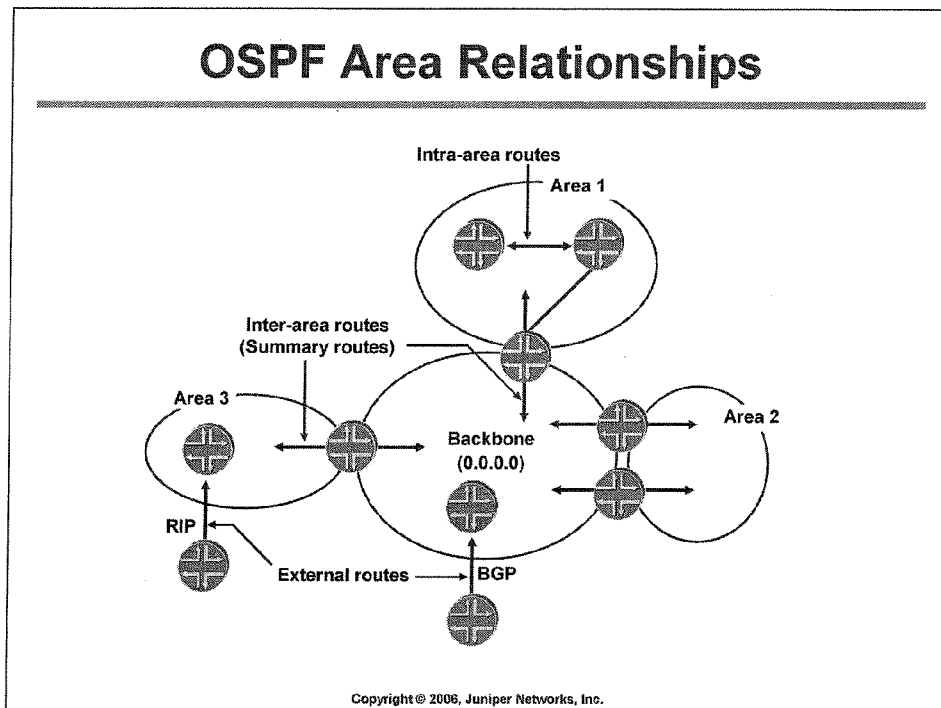
#### **Area Border Router**

An OSPF router with links in two areas is called an ABR. The ABR is responsible for connecting OSPF areas to the backbone. It transmits network information between the backbone and the other areas.

#### **Autonomous System Boundary Router**

An OSPF router that injects routing information from outside the OSPF AS is known as an area system boundary router (ASBR). Typically, an ASBR is located in the backbone, but the OSPF specification allows an ASBR in other areas as well.

## Configuring Juniper Networks Routers



### OSPF Area Relationships

- On the slide, all areas are connected directly to the backbone. In the rare situations where a new area is introduced that cannot have direct physical access to the backbone, you must configure a virtual link.

OSPF classifies different types of routing information as follows:

- Routes that are generated from within an area, where the destination belongs to the area are called *intra-area*, or *internal*, routes.
- Routes that originate from other areas are called *inter-area* or *summary* routes.
- Routes that originate from other routing protocols, or different OSPF processes, and that are injected into OSPF through redistribution, are called *external* routes.

### OSPF Area Types

- **Stub areas**
  - Do not carry external routes
  - Do not allow the configuration of virtual links across them
  - Cannot contain ASBRs
- **Totally stubby areas**
  - Stub areas that only receive the default route from the backbone
- **Not-so-stubby areas**
  - Allow external routes to be advertised from the area but not received from another area
- **Transit areas**
  - Used to pass traffic from one adjacent area to the backbone

Copyright © 2006, Juniper Networks, Inc.

#### Stub Areas

Stub areas are areas through which, or into which, AS external advertisements are not flooded. You might want to create stub areas when much of the topological database consists of AS external advertisements. Doing so reduces the size of the topological databases and, therefore, the amount of memory required on the internal routers in the stub area.

When you configure an area border router for stub area operation, a default route is normally advertised in the place of the external routes that are blocked from stub areas. The default route provides routers in the stub area with reachability to external destinations. In JUNOS software ABRs require explicit configuration for default route generation.

Note that you cannot create a virtual link through a stub area, and a stub area cannot contain an AS boundary router.

#### Totally Stubby Area

A totally stubby area is a stub area that only receives the default route from the backbone.

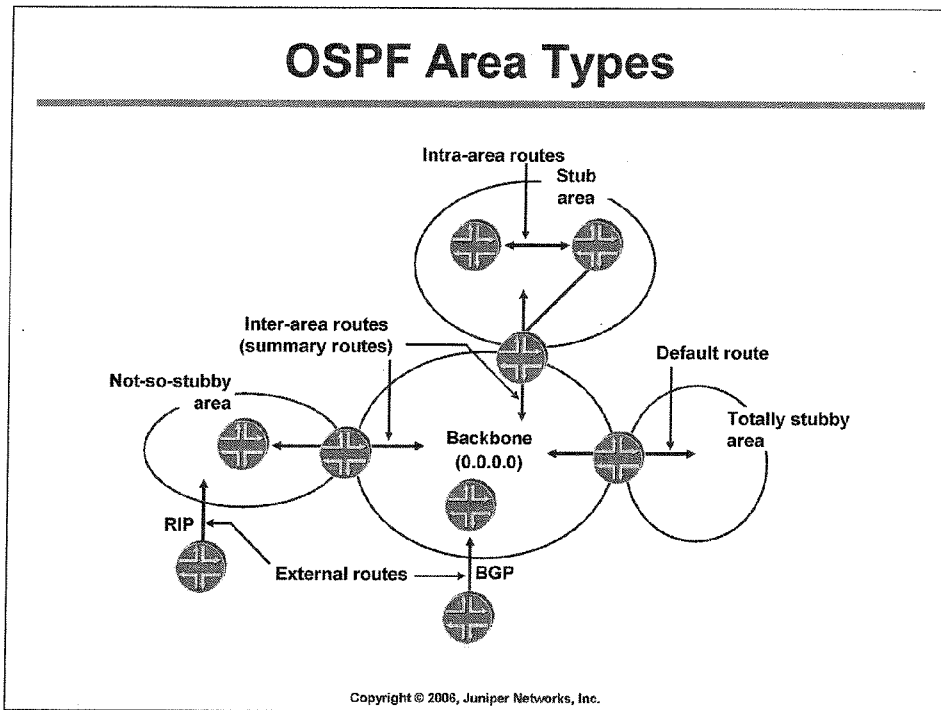
#### Not-So-Stubby Area

An OSPF stub area has no external routes in it, so you cannot redistribute routes from another protocol into a stub area. A not-so-stubby-area (NSSA) allows external routes to be flooded within the area. These routes are then leaked into other areas. However, external routes from other areas still do not enter the NSSA.

#### Transit Areas

Transit areas pass traffic from one adjacent area to the backbone, or to another area if the backbone is more than two hops away from an area.

## Configuring Juniper Networks Routers



### Illustration of Area Types

The slide illustrates the different OSPF area types.

### Adjacency Formation and DR Election

---

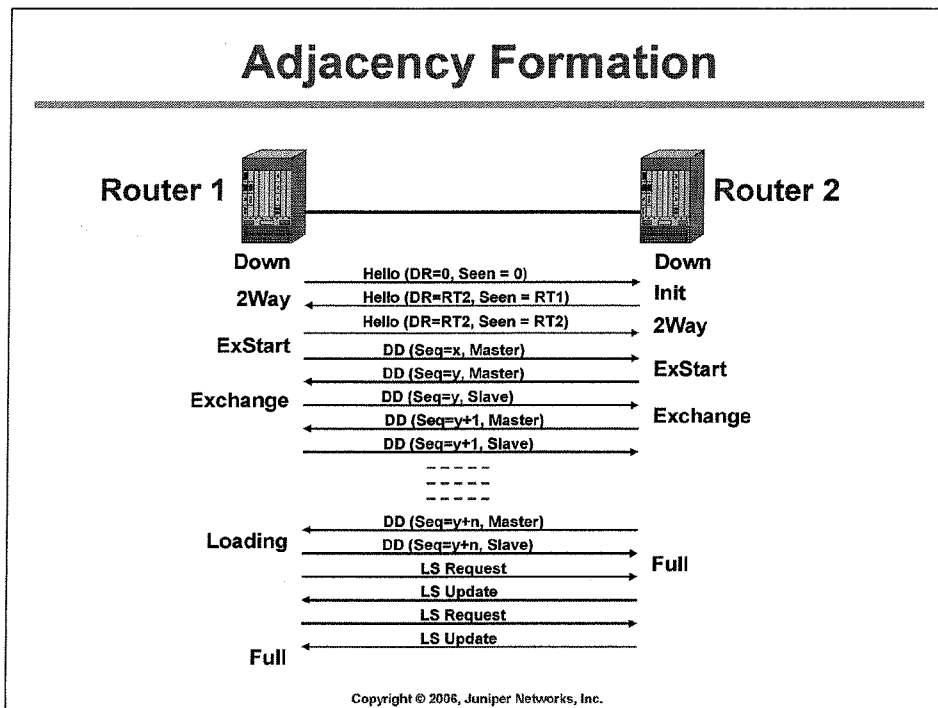
- Where we are going...
  - Adjacency formation
  - Adjacency optimization: the DR
  - Designated router election
  - Backup designated router
  - OSPF neighbors vs. adjacency

Copyright © 2006, Juniper Networks, Inc.

### Adjacency Formation and DR Election

The slide lists the topics examined on the following pages.

## Configuring Juniper Networks Routers



### Forming an Adjacency


Before any link-state information can be flooded by an OSPF router, that network link must be eligible for flooding. The key to eligibility for OSPF is the adjacency, which is the relationship formed between two OSPF-speaking routers. This adjacency ensures that both routers know about each other and can agree on certain parameters about the link. This agreement assures that data traffic can be transmitted reliably across that link.

Seven possible adjacency states in OSPF exist. These states describe the process of an OSPF adjacency formation:

1. *Down*: This initial state indicates that OSPF is waiting for a start event.
2. *Init*: This state indicates that a hello packet has been sent, but bidirectional communication has not been established.
3. *2Way*: This state indicates that a hello packet has been received with the router's RID listed in the neighbor section. Bidirectional communication has been achieved.
4. *ExStart*: This state indicates that the routers negotiate between themselves to determine which router is in charge of the database synchronization process.
5. *Exchange*: This state indicates that the routers exchange LSA headers describing their own databases. If a router does not know about a received LSA header, it can transmit a link-state request for the complete information.
6. *Loading*: This state indicates a router has finished transmitting its database to its peer but is still receiving database information.
7. *Full*: This state indicates that the databases are now fully synchronized. The network link now can be advertised to the OSPF network.

### Adjacency Optimization: The DR (1 of 2)

- By default, OSPF attempts to form an adjacency with all neighbors discovered on all interfaces
  - On a broadcast media like Ethernet, this desire is suboptimal because it would require a full mesh of adjacencies



Copyright © 2006, Juniper Networks, Inc.


### The Designated Router

OSPF routers want to form an adjacency with all routers with which they exchange hello packets. On a broadcast medium such as Ethernet, this desire can pose quite a problem. As more routers are added to the link, more adjacencies must be formed. This full-mesh requirement places extra load on the routers with little extra benefit because they all are advertising the same link information.



## Adjacency Optimization: The DR (2 of 2)

- OSPF elects a single router to represent the segment
  - Significantly reduces OSPF traffic on segment
  - Minimizes OSPF processes



Copyright © 2006, Juniper Networks, Inc.

### Using the DR

To avoid the problem described on the previous page, OSPF has a single router represent the broadcast link to the rest of the network. This router is called the designated router (DR). It is the DR's job to form an adjacency to all other routers on the link and to advertise the link-state information to the AS.

### Designated Router Election

---

- Every OSPF router has a DR election priority
  - Range is 0–255
  - JUNOS software default is 128
  - Priority is the first tie breaker in a DR election
  - If two routers share the highest priority, the router with the highest router ID (RID) is elected
  - The election of a DR is a nondeterministic event
    - An existing DR will not be replaced
    - First router on the segment within 40 seconds wins

Copyright © 2006, Juniper Networks, Inc.

#### Designated Router Priority

OSPF bases the election of the DR on two election criteria: priority and RID. OSPF DR priorities can range from 0 through 255, with the JUNOS software default being 128. A router with a higher priority has a better chance of becoming the DR, as this is the first tiebreaker in a DR election. A router with a DR priority of 0 is not eligible for election and never becomes the DR. In the event of a priority tie, the second tiebreaker for DR elections is the RID of the routers—the higher the value of the RID, the better the chance of becoming the DR for the segment.

The election of the DR on an OSPF link is a nondeterministic event. To avoid network instabilities, the current DR always operates until it leaves the network. Thus, if a router with a DR priority of 250 comes online and sees that there is already a DR present on the segment, it does not assume the DR role.

The election of the very first DR on a segment occurs within 40 seconds of the first router transmitting a hello packet. This wait time is honored every time an election is held.

### Backup Designated Router

- If a DR fails, a new DR must be elected
  - The default 40-second wait time means the segment is unreachable within that period
  - To ease the transition, a backup designated router (BDR) is elected on every segment
  - The BDR watches and waits for a DR failure
  - Election rules for BDR are the same as for the DR

Copyright © 2006, Juniper Networks, Inc.

#### The Backup Designated Router

One of the reasons that an OSPF DR maintains its role *forever* is to avoid network instabilities. However, if the DR does leave the network, the 40-second wait time for a new election should take place. During this 40-second wait time, OSPF can form no adjacencies and cannot advertise the link to the AS. Therefore, no data traffic can use the link. This is clearly in violation of the stability goal!

To avoid this problem, each broadcast network segment also elects a backup designated router (BDR). The BDR is elected after the DR is chosen. It is the BDR's responsibility to monitor the DR and ensure that it is working properly. If the BDR notices that the DR has left the network, it automatically assumes the role of the DR. A new BDR is then elected on the segment.

The election rules for a BDR are the same as for the DR. Again, the BDR election is nondeterministic.

## Configuring Juniper Networks Routers

### OSPF Neighbors versus Adjacencies

```
user@host> show ospf neighbor extensive
Address      Intf      State      ID          Pri  Dead
172.16.30.254 ge-0/0/0.0 Full       10.250.240.8 128  30
area 0.0.0.5, opt 0x42, DR 172.16.30.254, BDR 172.16.30.253
Up 00:10:50, adjacent 00:10:50

172.16.30.253 ge-0/0/0.0 Full       10.250.240.35 128  30
area 0.0.0.5, opt 0x42, DR 172.16.30.254, BDR 172.16.30.253
Up 00:10:50, adjacent 00:10:52

172.16.30.252 ge-0/0/0.0 2Way      10.250.240.32 64   38
area 0.0.0.5, opt 0x42, DR 172.16.30.254, BDR 172.16.30.253
Up 00:08:10
```

Copyright © 2006, Juniper Networks, Inc.

### OSPF Neighbor Relationship

As soon as an OSPF router sees a hello packet on an interface, it starts to retain knowledge of that neighbor. You can display this information with the operational CLI command `show ospf neighbor`.

On the slide, this router has three neighbors on the `ge-0/0/0.0` interface. Two of the three routers are the DR and the BDR; full adjacencies exist with them. Each of the hello packets received from all three routers lists their addresses.

The router that is in a 2-way state is a neighbor on the link, but it is not the DR/BDR. This router reaches the 2-way state because the DR and BDR can see its hello packets, and this router's own RID is located in the received hello. For broadcast media, it is acceptable to have some neighbors in the 2-way state.

The address column is the interface IP address of the neighboring router. The ID column is the RID of the neighboring router.

## Configuring OSPF

---

- Where we are going...
  - JUNOS software OSPF support
  - Configuration examples
  - Router ID algorithm

Copyright © 2006, Juniper Networks, Inc.

### Configuring OSPF

The slide lists the topics discussed on the following pages.

### JUNOS Software OSPF Support

- OSPF Version 2, including:
  - Virtual links
  - Stub areas, not-so-stubby areas (NSSA), and totally stubby areas
  - Authentication
  - Summarization
  - Traffic engineering (LSA Type 10 support)
  - Graceful restart
  - External prefix limits
  - Interacts with Bidirectional Forwarding Detection for rapid convergence

Copyright © 2006, Juniper Networks, Inc.

### JUNOS Software OSPF Support

The slide details JUNOS software support for OSPF.

Graceful restart basically allows for forwarding while routing (*rpd*) restarts. The Juniper Networks implementation is based on *Hitless OSPF Restart* by John Moy (*draft-ietf-ospf-hitless-restart-<version>.txt*).

You can place a limit on the number of external prefixes advertised by OSPF. When this limit is reached, the router appears *overloaded* to its peers. The feature is useful when you want to protect your IGP from possible meltdown in the event that a policy mistake results in the redistribution of an entire BGP table into OSPF. You configure this option with a `set prefix-export-limit number-of-prefixes` statement at the `[edit protocol ospf]` hierarchy. Once the router enters the overload state, you can clear the condition with a `clear ospf database` command. Note that the router will once again become overloaded if the number of external prefixes still exceeds the configured limit.

Bidirectional Forwarding Detection (BFD) is a protocol designed to provide subsecond failure detection between routers connected with various media types. BFD is not a routing protocol in itself; BFD integrates with the OSPF protocol to expedite reconvergence when the BFD daemon has detected a failure. You enable BFD for OSPF with a `set area area-number interface interface-name bfd-liveness-detection minimum-interval value` statement issued at the `[edit protocols ospf]` hierarchy. The BFD interval is in milliseconds. Setting the value too low might result in loss of adjacencies, but the protocol is designed to automatically increase timers to avoid erroneous adjacencies failures. *Draft-katzward-bfd-<version>.txt* currently defines BFD.

## Configuring Juniper Networks Routers

### Configuring OSPF—Single Area

---

```
[edit]
user@host# set protocols ospf area 0 interface ge-0/0/0

[edit]
user@host# show protocols ospf
ospf {
  area 0.0.0.0 {
    interface ge-0/0/0.0;
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

#### Basic OSPF Configuration

- The screen capture on the slide shows the minimum configuration needed for a single OSPF area network on a Juniper Networks M-series or T-series platform.

## Configuring Juniper Networks Routers

### Configuring OSPF—Multiple Areas

```
[edit]
user@host# show protocols ospf
ospf {
  area 0.0.0.0 {
    interface ge-0/0/0.0;
  }
}
[edit]
user@host# set protocols ospf area 1 interface at-0/1/1.100
[edit]
user@host# show protocols ospf
ospf {
  area 0.0.0.0 {
    interface ge-0/0/0.0;
  }
  area 0.0.0.1 {
    interface at-0/1/1.100;
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

#### Multiple Area OSPF Configuration

The screen capture on the slide shows a configuration for an ABR. Notice that there are two interfaces specified in two different areas. Because of this specification, this router maintains two link-state databases—one for area 0.0.0.0, and one for area 0.0.0.1.



### Router ID Selection

- You can set RID manually under `routing-options`, although explicit configuration is not necessary
  - JUNOS software uses the first non-127/8 address it finds as the RID
  - `lo0` is the first interface activated, so a non-127/8 configured here serves as the RID
    - JUNOS software automatically advertises a stub route to the RID, so you do not need to run OSPF (passive or active) on the interface that sources the RID
  - If JUNOS software does not find a suitable address on `lo0`, it examines the next interface activated (normally `fxp0`)
    - `fxp0` does not support transit traffic, so sourcing the RID from `fxp0` can result in the inability to ping the RID

Copyright © 2008, Juniper Networks, Inc.

#### Determining the Router ID

Every OSPF router has a single RID; it is a 32-bit number in dotted quad notation. OSPF uses this RID for several purposes, including DR election and LSA originator identification.

You can explicitly configure the RID under the `routing-options` portion of the hierarchy. However, if you do not specify it, JUNOS software will use one of the IP addresses configured on the router at the time the OSPF process starts. If the loopback interface is configured with an IP address, JUNOS software uses that address first. If a loopback interface is not configured, JUNOS software uses the next suitable address, typically `fxp0`. Remember, `fxp0` does not support transit traffic, so sourcing the RID from `fxp0` can result in the inability to ping the RID.

### OSPF Monitoring and Troubleshooting

---

- Where we are going...
  - Various show commands exist to provide detailed information on the operation of OSPF
    - `show ospf interface`
    - `show ospf neighbor`
    - `show ospf log`
    - `show ospf statistics`
    - `show ospf io-statistics`
    - `show ospf overview`
    - `show ospf route`
    - `show ospf database`
  - OSPF tracing

Copyright © 2006, Juniper Networks, Inc.

### OSPF Monitoring and Troubleshooting

The following pages provide the details of various `show ospf...` commands.

### Displaying OSPF Interface Parameters

Use `show ospf interface` to display the OSPF parameters associated with an interface

```
user@host> show ospf interface ?
Possible completions:
<[Enter]>          Execute this command
brief              Show brief status
detail             Show detailed status
extensive          Show extensive status
. . .
```

```
user@host> show ospf interface brief
Intf      State   Area   DR ID      BDR ID      Nbrs
ge-1/2/0.0 DRother 0.0.0.0 10.250.240.8 10.250.240.35 3
ge-2/0/0.0 DR       0.0.0.1 10.250.240.17 10.250.240.11 2
ge-2/1/0.0 DR       0.0.0.1 10.250.240.17 10.250.240.9 1
ge-4/1/0.0 DR       0.0.0.1 10.250.240.17 10.250.240.10 1
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying OSPF Interface Parameters

The output fields of the `show ospf interface` command are:

- **Intf:** Displays the name of the interface running OSPF.
- **State:** Displays the state of the interface. It can be BDR, Down, DR, DRother, Loop, PtToPt, or Waiting.
- **Area:** Displays the number of the area in which the interface is located.
- **DR ID:** Displays the address of the area's DR.
- **BDR ID:** Displays the BDR for a particular subnet.
- **Nbrs:** Displays the number of neighbors on this interface.
- **Type (detail and extensive output only):** Displays the type of interface. It can be LAN, NBMA, P2MP, P2P, or Virtual.
- **address (detail and extensive output only):** Displays the IP address of the neighbor.
- **mask (detail and extensive output only):** Displays the mask of the interface.
- **MTU (detail and extensive output only):** Displays the interface's MTU.
- **cost (detail and extensive output only):** Displays the interface's cost (metric).
- **DR addr (detail and extensive output only):** Displays the address of the DR.
- **BDR addr:** Displays the address of the BDR.
- **adj count (detail and extensive output only):** Displays the number of adjacent neighbors.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Displaying OSPF Interface Parameters (contd.)

- `Flood list` (extensive output only): Displays the list of LSAs pending flood on this interface.
- `Ack list` (extensive output only): Displays the list of pending acknowledgments on this interface.
- `Descriptor list` (extensive output only): Displays the list of packet descriptors.
- `Dead` (detail and extensive output only): Displays the configured value for the dead timer.
- `Hello` (detail and extensive output only): Displays the configured value for the hello timer.
- `ReXmit` (detail and extensive output only): Displays the configured value for the retransmit timer.
- `OSPF area type` (detail and extensive output only): Displays the type of OSPF area, which can be Stub, Not Stub, or NSSA.

### Displaying OSPF Adjacency Information

Use `show ospf neighbor` to display OSPF parameters adjacency information

```
user@host> show ospf neighbor ?
Possible completions:
<[Enter]>          Execute this command
brief              Show brief status
detail            Show detailed status
extensive         Show extensive status

user@host> show ospf neighbor
Address          Intf           State  ID              Pri  Dead
192.168.254.225  ge-1/2/0.0    2Way   10.250.240.32   128  36
192.168.254.230  ge-1/2/0.0    Full   10.250.240.8    128  38
192.168.254.229  ge-1/2/0.0    Full   10.250.240.35   128  33
10.1.1.129       ge-2/0/0.0    Full   10.250.240.12   128  23
10.1.1.131       ge-2/0/0.0    Full   10.250.240.11   128  24
10.1.2.1         ge-2/1/0.0    Full   10.250.240.9    128  32
10.1.2.81        ge-2/1/0.0    Full   10.250.240.10   128  33
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying OSPF Adjacency Information

The output fields of the `show ospf neighbor` command are:

- Address: Displays the address of the neighbor.
- Intf: Displays the interface through which the neighbor is reachable.
- State: Displays the state of the neighbor, which can be Attempt, Down, Exchange, ExStart, Full, Init, Loading, or 2Way.
- ID: Displays the RID of the neighbor.
- Pri: Displays the priority of the neighbor to become the DR.
- Dead: Displays the number of seconds until the neighbor becomes unreachable.
- area (detail and extensive output only): Displays the area in which the neighbor is located.
- opt (detail and extensive output only): Displays the option bits from the neighbor.
- DR (detail and extensive output only): Displays the address of the DR.
- BDR (detail and extensive output only): Displays the address of the BDR.
- Up (detail and extensive output only): Displays the length of time since the neighbor came up.
- adjacent (detail and extensive output only): Displays the length of time since the adjacency with the neighbor was established.

### Troubleshooting Adjacency Problems

---

- **Problem:**
  - **No neighbor**
    - Physical and data link layer connectivity
    - Mismatched IP subnet/mask, area number, area type, authentication, hello/dead interval, or network type
  - **Stuck in 2-way state**
    - Normal for DR-other neighbors
  - **Stuck in exchange start**
    - Mismatched IP MTU

Copyright © 2006, Juniper Networks, Inc.

### Troubleshooting

When attempting to find the reason why an OSPF adjacency does not come up to full state, look for some of the items listed in the slide.

## Clearing OSPF Adjacencies

---

Use the `clear ospf neighbor` to clear OSPF adjacencies

```
user@host> clear ospf neighbor 192.168.254.225
```

Copyright © 2006, Juniper Networks, Inc.

### Clearing OSPF Adjacencies

This command clears the neighbor adjacency. The adjacency should be reformed immediately. Use it as necessary, but remember that it causes a momentary disruption.

### Displaying OSPF SPF-Related Information

Use `show ospf log` to display OSPF SPF-related information

```
user@host> show ospf log

      Last instance of each event type
When      Type      Elapsed
00:03:15  SPF       0.000089
00:03:15  Stub     0.000097
00:03:00  Interarea 0.000126
00:03:00  External  0.000200
00:03:00  NSSA     0.000001
00:03:00  Cleanup  0.000460

      Maximum length of each event type
When      Type      Elapsed
00:03:15  SPF       0.000089
00:03:15  Stub     0.000097
00:03:00  Interarea 0.000126
00:03:15  External  0.000235
00:04:05  NSSA     0.000001
00:03:15  Cleanup  0.000787

      Last 100 events
When      Type      Elapsed
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying OSPF SPF-Related Information

The `show ospf log` command displays the entries in the OSPF log of SPF calculations. Note that the SPF algorithm performs multiple calculations on different portions of the link-state database. Use this command to verify that you have a stable OSPF routing domain. Multiple recalculations in a short period of time indicate potential instability. The output fields of this command are:

- **When:** Displays the time when the SPF calculation was made.
- **Type:** Displays the type of calculation, which can be Cleanup, External, Interarea, NSSA, Redist, SPF, Stub, Total, or Virtuallink.
- **Elapsed:** Displays, in seconds, the how much time has passed since the calculation was made.



## Configuring Juniper Networks Routers

### Displaying OSPF Statistics

#### Use show ospf statistics

```
user@host> show ospf statistics
```

Packet type	Total		Last 5 seconds		
	Sent	Received	Sent	Received	
Hello	5	4	0	0	
DbD	2	3	0	0	
LSReq	1	1	0	0	
LSUpdate	3	6	0	0	
LSAck	5	3	0	0	
DBDs retransmitted	:		0, last 5 seconds	:	0
LSAs flooded	:		3, last 5 seconds	:	0
LSAs flooded high-prio	:		0, last 5 seconds	:	0
LSAs retransmitted	:		0, last 5 seconds	:	0
LSAs transmitted to nbr	:		5, last 5 seconds	:	0
LSAs requested	:		51, last 5 seconds	:	0
LSAs acknowledged	:		60, last 5 seconds	:	0
Flood queue depth	:	0			
Total rexmit entries	:	0			
db summaries	:	0			
lsreq entries	:	0			

Copyright © 2006, Juniper Networks, Inc.

#### Displaying OSPF Statistics

The first part of this command's output displays details of the number and type of OSPF packets sent and received by the router. The output fields of this command are:

- Packet type: Displays the type of OSPF packet.
- Total Sent/Received: Displays the total number of packets sent and received.
- Last 5 seconds Sent/Received: Displays the total number of packets sent and received in the last five seconds.
- LSAs retransmitted: Displays the total number of link-state advertisements transmitted and number retransmitted in the last five seconds.
- Receive errors: Displays the number and type of receive errors.

The second part of the output shows more specific details about the packets' contents.

## Configuring Juniper Networks Routers

### Displaying OSPF Route Information

Use the `show ospf route` command to display routes learned from, and advertised to, OSPF

```
user@host> show ospf route ?
Possible completions:
<[Enter]>          Execute this command
abr                Display OSPF routes to area border routers
asbr              Display OSPF routes to AS border routers
detail            Display detailed output
extern            Display external OSPF routes
instance          Name of OSPF instance
inter             Display interarea OSPF routes
intra             Display intraarea OSPF routes
logical-router    Name of logical router, or 'all'
```

```
user@host> show ospf route detail
Prefix          Route/Path/NextHop Type  Metric  Next hop i/f  NH addr/Label
192.168.0.1/32  Intra Router IP 1      so-0/1/2.0
  area 0.0.0.0, options 0x0x0, origin 192.168.0.1
192.168.28.1/32 Intra Router IP 1      so-0/1/1.0
  area 0.0.0.0, options 0x0x0, origin 192.168.28.1
10.0.0.0/24     Intra Network IP 2      so-0/1/2.0
  area 0.0.0.0, options 0x0x0, origin 192.168.0.1
```

Copyright © 2008, Juniper Networks, Inc.

### Displaying OSPF Route Information

This command displays only those routes in the unicast routing table, `inet.0`, installed by OSPF. The use of additional keywords allows you to display only OSPF routes learned by specific LSA types. The output fields of the `show ospf route` command are:

- Prefix: Displays the destination of the route.
- Route/Path Type: Displays how the route was learned:
  - ABR: Route to area border router;
  - ASBR: Route to AS border router;
  - Ext: External router;
  - Inter: Inter-area route;
  - Intra: Intra-area route; or
  - Network: Network router.
- Metric: Displays the route's metric value.
- Next hop i/f: Displays the interface through which the route's next hop is reachable.
- Next hop addr: Displays the address of the next hop.
- area: (detail output only) Displays the area ID of the route.
- options: (detail output only) Displays the option bits from the LSA.
- origin: (detail output only) Displays the router from which the route was learned.

### Displaying Routes Learned from OSPF

- Use the `show route protocol ospf` command to display routes learned from OSPF

- Command does not list OSPF interface routes

```
user@host> show route protocol ospf
```

```
inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, * = Both
```

```
10.0.0.0/24      *[OSPF/10] 00:09:28, metric 2  
                 > via so-0/1/2.0  
10.0.1.0/24      *[OSPF/10] 00:09:28, metric 2  
                 > via so-0/1/2.0  
10.0.16.0/24     *[OSPF/10] 00:09:28, metric 2  
                 > via so-0/1/2.0
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying Routes Learned from OSPF

This command displays only those routes in the unicast routing table, `inet.0`, installed by OSPF. The use of additional keywords allows you to display only OSPF routes learned by specific LSA types.

## Configuring Juniper Networks Routers

### Displaying Entries in the OSPF LSDB

Use the `show ospf database` command to display entries in the OSPF LSDB

```
user@host> show ospf database ?
Possible completions:
<[Enter]>           Execute this command
advertising-router  Router ID of advertising router
area                OSPF area ID
asbrsummary         Show summary AS boundary router link-state database
brief              Display brief output (default)
detail             Display detailed output
extensive          Display extensive output
extern             Show external link-state database
instance           Name of OSPF instance
link-local         Show link local link-state database
logical-router     Name of logical router, or 'all'
lsa-id            Link-state advertisement ID
netsummary         Show summary network link-state database
network           Show network link-state database
nssa              Show not-so-stubby area link-state database
router            Show router link-state database
summary           Display summary output
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying Entries in the OSPF LSDB

The `show OSPF database` command displays the OSPF database. The display is organized by LSA types. ABRs have a separate database for each area. The options for this command are:

- `brief` (optional): Displays a brief listing of all entries in the OSPF link-state database. This is the default setting.
- `detail` (optional): Displays detailed information about the entries in the OSPF link-state database.
- `extensive` (optional): Displays extremely detailed information about the entries in the OSPF link-state database.
- **LSA filters** (optional): Displays one or more of the following LSA filters. If you specify more than one filter, only LSAs that match all the filters are displayed. For example, the command `show ospf database detail router lsa-id 10.0.0.1` displays all router LSAs in all areas that have an LSA identifier of 10.0.0.1.
  - `advertising-router address`: Displays the LSAs advertised by a particular router.
  - `area area-id`: Displays the LSAs in a particular area.
  - `lsa-id lsa-id` (optional): Displays the LSA with the specified LSA identifier.
  - `lsa-type`: Displays specific types of LSAs. You can specify `asbrsummary`, `extern`, `netsummary`, `network`, `nssa`, or `router`.
- `summary` (optional): Displays summary information about the OSPF link-state database.

### Displaying OSPF LSDB Brief Entries

```
user@host> show ospf database brief
OSPF link state database, area 0.0.0.0
Type      ID                Adv Rtr           Seq              Age      Cksum  Len
Router    *10.250.240.8     10.250.240.8     0x800001fc      2388    0x3684 36
Router    10.250.240.17    10.250.240.17    0x80000217      1835    0x444c 36
Router    10.250.240.32    10.250.240.32    0x80000232      1876    0x0158 36
Router    10.250.240.35    10.250.240.35    0x80000291      1100    0x4aa5 36
Network   *192.168.254.230 10.250.240.8     0x800001cc      117     0xab67 40
Summary   10.1.2.0          10.250.240.17    0x80000216      1535    0x1729 28
Summary   *10.1.3.34        10.250.240.8     0x8000013a      2217    0x842f 28

OSPF link state database, area 1.0.0.0
Type      ID                Adv Rtr           Seq              Age      Cksum  Len
Router    10.250.240.9     10.250.240.9     0x80000267      116     0x1bb3 36
[additional information]
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying OSPF LSDB Brief Entries

This command displays a brief entry for each entry in the database. The asterisk (\*) indicates that this router originated this entry. The sequence numbers are used to determine if an LSA is new; they start with a value of 0x80000001. According to the OSPF specification, an LSA's originator should refresh all valid LSAs every 30 minutes, or 1800 seconds, to prevent LSAs from aging out. The JUNOS software implementation increases the refresh rate to every 50 minutes to help minimize the consumption of network bandwidth and CPU cycles. *Stale* LSAs are LSAs with an age greater than the implementation's refresh rate. Any LSA older than 3600 are purged automatically. The output fields of the `show ospf database brief` command are:

- **area:** Displays the area number. Area 0.0.0.0 is the backbone area.
- **Type:** Displays the type of link advertisement, which can be ASBRSum, Extern, Network, NSSA, Router, or Summary.
- **ID:** Displays the link identifier included in the advertisement. An asterisk (\*) preceding the identifier marks database entries that originated from the local router.
- **Adv Rtr:** Displays the address of the router that sent the advertisement.
- **Seq:** Displays the link sequence number of the advertisement.
- **Age:** Displays the time elapsed since the LSA was originated, in seconds.
- **Cksum:** Displays the checksum value of the LSA.
- **Len:** Displays the length of the advertisement, in bytes.

### Displaying OSPF LSDB Extensive Entries

```
user@host> show ospf database extensive

OSPF link state database, area 0.0.0.0
Type      ID          Adv Rtr      Seq          Age  Opt  Cksum  Len
Router   192.168.0.1 192.168.0.1 0x80000003   740 0x2  0x20d7 108
  bits 0x0, link count 7
  id 172.25.8.0, data 255.255.255.0, type Stub (3)
  TOS count 0, TOS 0 metric 1
  id 10.0.16.0, data 255.255.255.0, type Stub (3)
  TOS count 0, TOS 0 metric 1
  id 192.168.0.1, data 255.255.255.255, type Stub (3)
  TOS count 0, TOS 0 metric 0
  id 10.0.1.0, data 255.255.255.0, type Stub (3)
  TOS count 0, TOS 0 metric 1
  id 10.0.0.0, data 255.255.255.0, type Stub (3)
  TOS count 0, TOS 0 metric 1
  id 192.168.20.1, data 10.0.18.2, type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 10.0.18.0, data 255.255.255.0, type Stub (3)
  TOS count 0, TOS 0 metric 1
Aging timer 00:47:39
Installed 00:12:18 ago, expires in 00:47:40, sent 00:12:17 ago
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying OSPF LSDB Extensive Entries

This slide shows the same command used on the previous page except that we appended it with the extensive modifier. It shows a much more detailed view of the OSPF link-state database. The following list shows additional output fields from this command:

- **bits**: Displays the flags describing the router that generated the LSP.
- **link count**: Displays the number of links in the advertisement.
- **Each link contains the following output fields:**
  - **id**: Displays the ID of a router or subnet on the link.
  - **data**: For stub networks, displays the subnet mask. Otherwise, it displays the IP address of the router that generated the LSP.
  - **type**: Displays the type of link. It can be PointToPoint, Transit, Stub, or Virtual.
  - **TOS count**: Displays the number of type-of-service (ToS) entries in the advertisement.
  - **TOS 0 metric**: Displays the metric for ToS 0.
- **Each ToS entry contains the following output fields:**
  - **TOS**: Displays the ToS value.
  - **metric**: Displays the metric for the ToS.
  - **Aging timer (extensive output only)**: Displays how long until the LSA expires (displayed as hrs:min:sec).

*Continued on next page.*

## Configuring Juniper Networks Routers

### Displaying OSPF LSDB Extensive Entries (contd.)

- **Installed:** (extensive output only) Displays how long ago the route was installed.
- **expires** (extensive output only): Displays how long until the route expires (displayed in hrs:min:sec).
- **Ours:** (extensive output only) Indicates that this is a local advertisement.

Table of type content in LSDB:

<u>Type</u>	<u>Link ID</u>	<u>Link Data</u>
1 point-to-point	Neighbor's router	IP of originating router's interface
2 link to transit network	Interface address of DR	IP of originating router's interface
3 link to stub network	IP network number	Network's subnet mask or host address (/32)
4 virtual link	Neighbor's router ID	MIB-II Index of originating router's interface

### Clearing the OSPF Database

- **Clears OSPF database**

- The **purge switch** forces all LSAs to maximum age so they must be refreshed

```
lab@host> clear ospf database ?
Possible completions:
  <[Enter]>          Execute this command
  instance          Name of OSPF instance
  logical-router    Name of logical router, or 'all'
  purge            Purge database entries instead of deleting them
  |               Pipe through a command
lab@host> clear ospf database purge
```

Copyright © 2006, Juniper Networks, Inc.

### Clearing the OSPF Database

The `clear ospf database` command clears the OSPF database for a particular router.

The `clear ospf database` command supports an optional `purge` switch. By including the `purge` switch you force the local router to set *all* LSAs in its database to `max-age`. These LSAs are then re-flooded according to the OSPF specification, which states that a router must flood any LSA that it has aged to `max-age`, regardless of whether that LSA was generated by the local router. All routers receive the newly flooded LSAs, which are now set to `max-age`, because of the reliable flooding mechanisms used by the OSPF protocol; the router that originated a given LSA is compelled to refresh that LSA when it receives an updated copy that indicates the LSA has reached `max-age`.

Albeit somewhat disruptive, this procedure tends to eliminate stale/bogus database entries without having to wait for the normal aging out process, which can take as long as 3,600 seconds (one hour). Note that the `purge` function is not supported in the `clear isis database` command.

Some readers have argued that purging another router's LSA is a violation of the OSPF standard. In fact, this behavior is well within the specification, which calls out the fact that because each OSPF routers runs on its own notion of local time, one router will always set a given LSA to `max-age` before all others, because one router will have a local oscillator that runs faster than all others. Using the `purge` switch simply exaggerates this behavior by causing the local router to *act* as if its internal clock is running really, really, fast.



### Tracing OSPF

● A typical OSPF tracing configuration:

```
[edit protocols ospf]
user@host# show
traceoptions {
  file ospf-trace;
  flag error detail;
  flag hello detail;
  flag lsa-update detail;
}
```

- Monitor the resulting *ospf-trace* log file using `monitor start log-file-name` or the `show log log-file-name` commands

Copyright © 2006, Juniper Networks, Inc.

### OSPF Tracing

To perform debugging functions on the OSPF routing process, use the JUNOS software `traceoptions` function. The trace output (debug information) is directed to the named log file, which is stored in the `/var/log` directory on the router's hard drive. You can view the log file using the `monitor start` or `show log` operational mode commands. In addition to specifying the trace file, you also must tell the router what information you want to trace. You can accomplish this specifying one or more `flag` keywords.

While you can only direct tracing to a single file, you can trace many options by using the `flag` keyword multiple times. In addition, you can add granularity by using the `detail`, `receive`, and `send` flag modifiers.

Available tracing flags for OSPF include:

<code>all</code>	Trace everything
<code>database-description</code>	Trace database description packets
<code>error</code>	Trace errored packets
<code>event</code>	Trace OSPF state machine events
<code>flooding</code>	Trace LSA flooding
<code>general</code>	Trace general events
<code>hello</code>	Trace hello packets
<code>lsa-ack</code>	Trace LSA acknowledgement packets
<code>lsa-request</code>	Trace LSA request packets
<code>lsa-update</code>	Trace LSA update packets
<code>normal</code>	Trace normal events
<code>on-demand</code>	Trace demand circuit extensions
<code>packet-dump</code>	Dump the contents of selected packet types
<code>packets</code>	Trace all OSPF packets
<code>policy</code>	Trace policy processing
<code>route</code>	Trace routing information
<code>spf</code>	Trace SPF calculations
<code>state</code>	Trace state transitions
<code>synchronization</code>	Trace NSR synchronization events

### Review Questions

---

1. What is the purpose of the OSPF DR?
2. What is the rationale behind a multi-area OSPF design?
3. How would you configure a Juniper Networks M-series or T-series platform to run OSPF area 0 on all of its interfaces?
4. List and describe three commands you can use to monitor and troubleshoot the OSPF protocol.

Copyright © 2006, Juniper Networks, Inc.

#### **This Module Discussed:**

- General OSPF operation;
- The role of the DR;
- OSPF area types and rationale;
- Configuring OSPF in JUNOS software;
- Using operational mode commands to monitor and troubleshoot OSPF; and
- Displaying and interpreting of the LSDB.

## Configuring Juniper Networks Routers

### Lab 5: OSPF

---

#### Lab Objective:

**Configure OSPF and use the CLI to monitor and troubleshoot its operation**

Copyright © 2006, Juniper Networks, Inc.





## **Configuring Juniper Networks Routers**

### ***Module 7: IS-IS Operation, Configuration, and Troubleshooting***

Copyright © 2006, Juniper Networks, Inc.

CJNR-M-7.a.7.6.1

### Module Objectives

---

- **After successfully completing this module, you will be able to:**
  - Describe IS-IS operation
  - Explain the reason for a designated intermediate system
  - Describe IS-IS areas and levels
  - Configure IS-IS
  - Monitor and troubleshoot IS-IS using the CLI
  - Display and interpret the LSDB

Copyright © 2006, Juniper Networks, Inc.

#### **This Module Discusses:**

- IS-IS operation;
- The IS-IS DIS;
- IS-IS areas and levels;
- Configuring IS-IS in JUNOS software;
- Using operational-mode commands to monitor and troubleshoot IS-IS; and
- Displaying and interpreting the IS-IS link-state database (LSDB).

## **IS-IS**

---

- **Where we are going...**
  - **What is IS-IS?**
  - **Integrated IS-IS**
  - **IS-IS concepts**
  - **IS-IS and OSPF comparison**
  - **IS-IS protocol data units**
  - **Neighbors and adjacencies**
  - **Designated intermediate system (DIS)**
  - **IS-IS metrics**
  - **Configuring IS-IS**
  - **Monitoring and troubleshooting IS-IS**

Copyright © 2006, Juniper Networks, Inc.

### **IS-IS**

The slide shows the topics examined on the following pages.

### What Is IS-IS?

- **An interior gateway protocol (IGP) based on the shortest path first (SPF) algorithm**
  - Uses link-state information to make routing decisions
- **Developed for routing International Organization for Standardization (ISO) Connectionless Network Protocol (CLNP) packets**
  - IP was added later
  - Defined in ISO/IEC 10589, RFCs 1142, 1195, and 2763

Copyright © 2006, Juniper Networks, Inc.

#### The IS-IS Protocol

The Intermediate Systems-to-Intermediate Systems (IS-IS) protocol is an IGP that uses link-state information to make routing decisions. It also uses an SPF algorithm, similar to OSPF.

#### The ISO

The International Organization for Standardization (ISO) developed IS-IS to be the routing protocol for the ISO's Connectionless Network Protocol (CLNP) and is described in ISO 10589. The protocol was developed by Digital Equipment Corporation for its DECnet Phase V. The ISO was working on IS-IS at the same time as the Internet Advisory Board (IAB) was working on OSPF. ISO proposed that IS-IS be adopted as the routing protocol for TCP/IP in place of OSPF. This proposal was driven by the opinion that TCP/IP was an interim protocol suite that would eventually be replaced by the OSI suite.



### Integrated IS-IS

- Implementation of IS-IS for routing IP in addition to CLNP
  - Also called *Dual IS-IS*
- All routers run a single routing algorithm
- IS-IS routers exchange link-state PDUs (LSPs)
  - Similar to OSPF LSAs/link-state update packets
  - IS-IS PDUs are the Layer 2 protocol used to transmit the information—NOT IP
  - IP reachability information is included in the updates
- M-series and T-series platforms do not support CLNP/CLNS routing
  - J-series platforms do

Copyright © 2006, Juniper Networks, Inc.

#### Dual IS-IS

To support the predicted transition from TCP/IP to OSI, an extension to IS-IS, known as *Integrated IS-IS*, was proposed. The purpose of integrated IS-IS, also known as dual IS-IS, was to provide a single routing protocol with the capabilities of routing both Connectionless Network Service (CLNS) and IP. The protocol was designed to operate in a pure CLNS, pure IP, or dual CLNS/IP environment.

#### Single Algorithm

Like all integrated routing protocols, Integrated IS-IS requires that all routers run a single routing algorithm. LSPs sent by routers running Integrated IS-IS include all destinations running either IP or CLNP network layer protocols. Routers running Integrated IS-IS still must support protocols such as ARP, ICMP for IP, and the ES-IS protocol for CLNP.

#### LSPs

Standard IS-IS packets must be modified to support multiple network layer protocols. IS-IS packet formats were designed to support the addition of new fields without a loss of compatibility with nonintegrated versions of IS-IS.

IS-IS adds type/length/value (TLV) objects (discussed further on following pages) to support integrated routing. These TLVs tell ISs which network layer protocols are supported by other ISs and whether end stations running other protocols can be reached. They also include any other required network layer, protocol-specific information.

#### JUNOS Support

M-series and T-series routers only support IP routing with IS-IS. They do not support CLNP or CLNS routing. J-series routers do support full IS-IS CLNP routing, including ES-IS (End System-to-Intermediate System routing).]

### IS-IS Concepts

- **IS-IS network is a single autonomous system (AS)**
  - End systems: network entities (that is, hosts) that send and receive packets
  - Intermediate systems: network entities (that is, routers) that send and receive packets and relay (that is, forward) packets
  - Protocol data units (PDUs): term for IS-IS packets
- **A single AS can be divided into smaller groups called areas, which are organized hierarchically**
  - Level 1 intermediate systems route within an area or toward a Level 2 system
    - Attached bit
  - Level 2 intermediate systems route between areas and toward other ASs

Copyright © 2006, Juniper Networks, Inc.

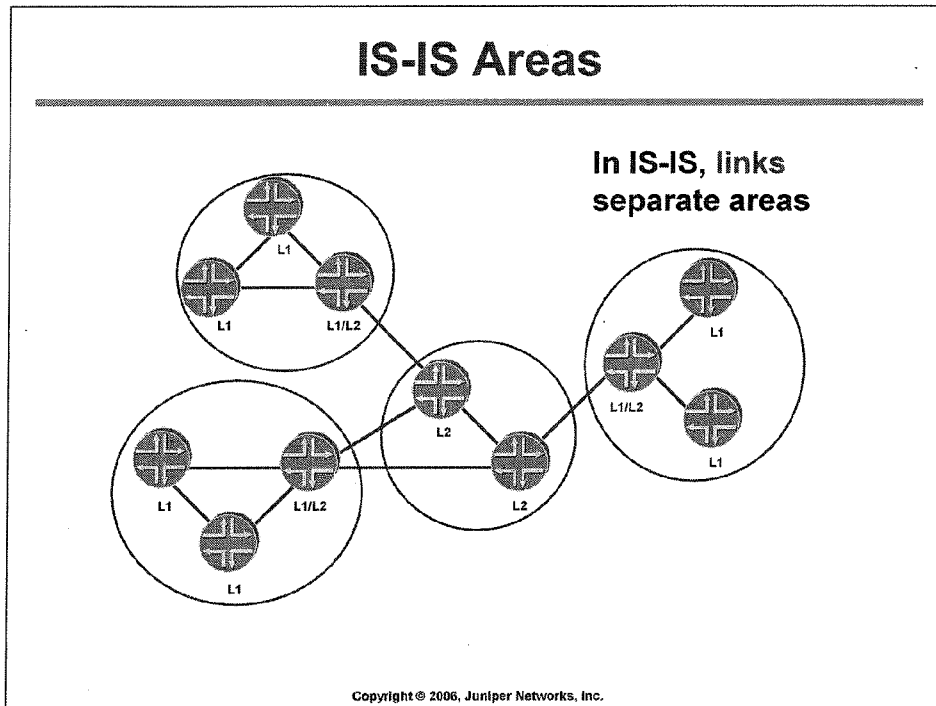
### Operation of IS-IS

An IS-IS network is a single AS, also called a *routing domain*, that consists of *end systems* (ESs) and *intermediate systems* (ISs). End systems are network entities that send and receive packets. Intermediate systems—which is the OSI term for a router—send and receive packets and relay, or forward, packets. IS-IS packets are called *protocol data units* (PDUs).

### IS-IS Areas

In IS-IS, a single AS can be divided into smaller groups called *areas*. Routing between areas is organized hierarchically, allowing a domain to be divided administratively into smaller areas. IS-IS accomplishes this organization by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area, and Level 2 intermediate systems route between areas and toward other ASs. A Level1/Level 2 system routes within an area on one interface and between areas on another.

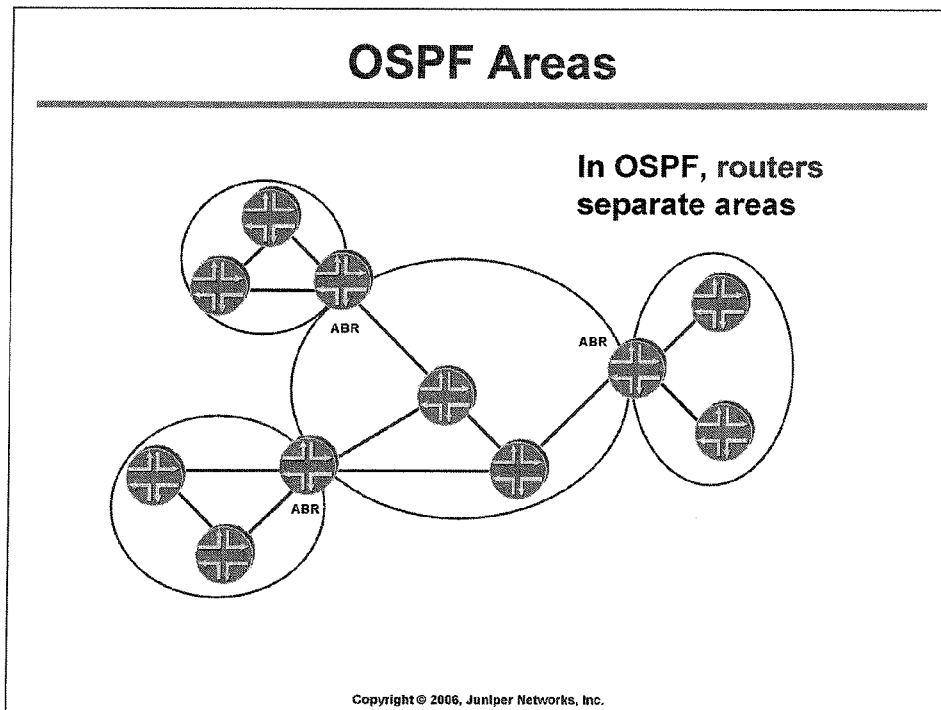
A Level 1/Level 2 system sets the attached bit in the Level 1 PDUs that it generates into a Level 1 area to indicate that it is a Level 2-attached backbone router and that it can be used to reach prefixes outside the Level 1 area. Level 1 routers create a default route for inter-area prefixes, which points to the closest (in terms of metrics) Level1/Level2-attached router.



### IS-IS Areas

Both IS-IS and OSPF are link-state routing protocols with many similarities. There are, of course, differences as well. In IS-IS areas, a Level 1/Level 2 router fulfills the same purpose as an area border router in OSPF. Likewise, the collection of Level 2 routers in IS-IS is the backbone, while area 0 is the backbone in OSPF. However, in IS-IS, all routers are completely within an area, and the area borders are on the links, not on the routers. The routers that connect areas are Level 2 routers, and routers that have no direct connectivity to another area are Level 1 routers. An intermediate system can be a Level 1 router, a Level 2 router, or both (an L1/L2 router).

## Configuring Juniper Networks Routers



### OSPF Areas

This slide shows the same network as the previous slide, only configured with OSPF instead of IS-IS. Compare the two slides to identify the similarities and differences. OSPF area borders are marked by routers. Some interfaces are in one area, and other interfaces are in another area. When an OSPF router has interfaces in more than one area, it is an ABR.

## IS-IS and OSPF

- **Both IS-IS and OSPF:**
  - Maintain link-state databases and construct a shortest path tree
    - Dijkstra algorithm
  - Use hello packets to form and maintain adjacencies
  - Use a two-level hierarchy
  - Provide for address summarization between areas
  - Elect a designated router
  - Have authentication capabilities
- **An excellent comparison of the two protocols is at <http://www.nanog.org/mtg-0006/katz.html>**

Copyright © 2006, Juniper Networks, Inc.

### IS-IS and OSPF Common Features

✳ The slide lists the commonalities between IS-IS and OSPF.

### ✳ A Comparison

An excellent comparison of the OSPF and IS-IS protocols is currently located at <http://www.nanog.org/mtg-0006/katz.html>. The presentation was written by Dave Katz, who is a principal developer of IGP code at Juniper Networks.

## Configuring Juniper Networks Routers

### Network Entity Titles

---

- IS-IS uses ISO NSAP addressing
  - Even in an IP-only environment, all routers must have an ISO address
  - ISO address is a network address or network entity title (NET)

<Area ID>	<System ID>	<n-selector>
49.0001.	00a0.c96b.c490	.00
49.0001.	1921.6816.9018	.00
  - First byte of area ID is the authority and format indicator (AFI)

1-13 Bytes	6 Bytes	1
Area ID	System ID	SEL
Level 2 Routing	Level 1 Routing	

Copyright © 2006, Juniper Networks, Inc.

### IS-IS Addressing

IS-IS uses ISO network addresses. Each address identifies a point of connection to the network, such as a router interface, and is called a *network service access point* (NSAP).

An end system can have multiple NSAP addresses, in which case the addresses differ only by the last byte, called the *n-selector*. Each NSAP represents an available service at that node. In addition to having multiple services, a single node can belong to multiple areas.

Each network entity also has a special network address called a *network entity title* (NET). Structurally, a NET is identical to an NSAP address but has an *n-selector* of 00. Most end systems and intermediate systems have one NET. Intermediate systems that participate in multiple areas can have multiple NETs.

The following ISO addresses illustrate the IS-IS address format:

- 49.0001.00a0.c96b.c490.00
- 49.0001.1921.6816.9018.00

The first portion of the address is the area number, which is a variable number from 1 through 13 bytes. The first byte of the area number (49) is the authority and format indicator (AFI). The next bytes are the assigned domain—or area—identifier, which can be from 0 through 12 bytes. In the examples above, the area identifier is 0001.

The next 6 bytes form the system identifier, or *sysid*. The *sysid* can be any 6 bytes that are unique throughout the entire domain. The system identifier commonly is the MAC address (as in the first example, 00a0.c96b.c490), or the IP address expressed in binary-coded decimal (BCD) (as in the second example, 1921.6816.9018, which corresponds to IP address 192.168.169.18). The last byte (00) is the *n-selector*.

*Continued on next page.*

## Configuring Juniper Networks Routers

### IS-IS Addressing (contd.)

To provide help with IS-IS debugging, JUNOS software supports dynamic mapping of ISO sysids to the host name. You can configure each system with a host name, which allows the sysid-to-host name mapping to be carried in a dynamic host name TLV in IS-IS LSP packets. This dynamic mapping permits any IS in the routing domain to learn about the ISO sysid of a particular IS.

You can assign multiple ISO NETs to the router's loopback interface. Such an action might be wanted when migrating two previously independent IS-IS domains into a single routing domain.

### IS-IS PDUs

---

- Where we are going...
  - Hello PDUs
  - Link-state PDUs
  - Sequence number PDUs
    - Complete
    - Partial
  - The building blocks for PDUs
    - Type/length/values (TLVs)

Copyright © 2006, Juniper Networks, Inc.

### IS-IS PDUs

IS-IS uses the following PDUs to exchange protocol information:

- *IS-IS hello (IIH) PDUs*: IS-IS broadcasts these PDUs to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems.
- *Link-state PDUs (LSPs)*: These PDUs contain information about the state of adjacencies to neighboring IS-IS systems. LSPs are flooded periodically throughout an area.
- *Complete sequence number PDUs (CSNPs)*: CSNPs contain a complete description of all LSPs in the IS-IS database. IS-IS periodically sends CSNPs on all links, and the receiving systems use the information in the CSNP to update and synchronize their LSP databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each LSP.
- *Partial sequence number PDUs (PSNPs)*: A receiver multicasts these PDUs when it detects that it is missing an LSP or when its LSP database is out of date. The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing LSP be transmitted. That router, in turn, forwards the missing LSP to the requesting router.
- *Type/length/values (TLVs)*: IS-IS PDUs use TLV encoding as the basic structure for all routing information. TLV encoding requires that the length of any field be defined explicitly when the field is used in a PDU.



### IS-IS Hello PDUs

- Hello mechanism for neighbor discovery used to build and maintain adjacencies
  - Similar to the hello mechanism in OSPF
- Separate hellos for:
  - Broadcast networks
    - Well-known multicast address used for Level 1 and Level 2 hellos
  - Point-to-point networks
- IS-IS transmits hello PDUs at regular intervals
- Hello PDUs:
  - Identify the device
  - Describe its capabilities
  - Describe the parameters of the interface

Copyright © 2006, Juniper Networks, Inc.

#### Purpose of IS-IS Hello PDUs

The purpose of the IS-IS hello PDU is to allow IS-IS routers to discover IS-IS neighbors on a link. Once the neighbors have been discovered and are adjacent, the hello PDU acts as a keepalive to maintain the adjacency and to inform the neighbors of any changes in the adjacency parameters.

#### IS-IS Hello PDU Types

Two kinds of IS-IS hellos PDUs exist: LAN hello PDUs and point-to-point hello PDUs. The LAN hello PDUs can be divided further into Level 1 and Level 2 hello PDUs. The format of the two types of LAN hello PDUs is identical. Note that on broadcast networks IS-IS Level 1 and Level 2 hellos are coded with multicast address 01-80-C2-00-00-14 or 01-80-C2-00-00-15 respectively.

#### Hello Transmission

Routers send hello packets at a fixed interval on all interfaces to establish and maintain neighbor relationships. The hello interval field advertises this interval in the hello packet. By default, a designated intermediate system (DIS) router sends hello packets every 3 seconds, and a non-DIS router sends hello packets every 9 seconds.

*Continued on next page.*

## Configuring Juniper Networks Routers

### PDU Fields

The following list provides the details of the PDU fields:

- *Circuit type*: Defines the router as an Level 1, Level 2 or Level 1/Level 2 router.
- *Source ID*: Identifies the system ID of the router that originated the hello PDU.
- *Holding time*: Specifies the period a neighbor should wait to receive the next hello PDU before declaring the originating router dead.
- *PDU length*: Specifies the length of the entire PDU in octets.
- *Priority*: Provides a value between 0 and 127 used for DIS election.
- *LAN ID*: Identifies the system ID or the DIS plus one more octet (the pseudo-node ID) to differentiate this LAN ID from another LAN ID that might have the same DR.

### IS-IS Link-State PDUs

---

- **Link-state PDUs (LSPs)**
  - **Used to build the link-state database**
    - Similar to LSAs in OSPF
  - **Separate LSPs for:**
    - Level 1 systems
    - Level 2 systems
  - **Sent as a result of a network change, during adjacency formation, and in response to a sequence number PDU (described on next page)**
  - **Link-state PDUs:**
    - Identify an IS's adjacencies
    - Describe the state of its adjacencies
    - Describe its reachable address prefixes (routes)

Copyright © 2006, Juniper Networks, Inc.

#### IS-IS Link-State PDUs

IS-IS sends LSPs at regular intervals and when an IS discovers that:

- Its link to a neighbor is down;
- It has a new neighbor; and
- The cost of a link to an existing neighbor has changed.

Once LSPs are distributed appropriately, IS-IS runs the Dijkstra algorithm to compute optimal paths to each ES. This algorithm iterates on the length of a path, examining the LSPs of all ISs working outward from the host IS. At the end of the computation, IS-IS forms a connectivity tree yielding the shortest paths, including all intermediate hops, to each IS.

### IS-IS Sequence Number PDUs

- **Partial sequence number PDUs (PSNPs)**
  - Used to:
    - Maintain the link-state database synchronization
    - Acknowledge LSPs from a neighbor on a point-to-point network
    - Request a copy of a missing LSP on a broadcast network
  - Separate type for Level 1 and 2 systems
  - Contain specific header information for a particular LSP being acknowledged or requested
- **Complete sequence number PDUs (CSNPs)**
  - Used to maintain the link-state database synchronization
    - Sent periodically by all ISs on point-to-point networks
    - Only sent by DIS on broadcast networks
  - Separate type for Level 1 and 2 systems
  - Contain header information for all LSPs in IS's link-state database

Copyright © 2006, Juniper Networks, Inc.

#### Partial Sequence Number PDUs

A receiver multicasts partial sequence number PDUs (PSNPs) when it detects that it is missing an LSP or when its LSP database is out of date. The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing LSP be transmitted. That router, in turn, forwards the missing LSP to the requesting router.

#### Complete Sequence Number PDUs

Complete sequence number PDUs (CSNPs) contain a complete description of all LSPs in the IS-IS database. IS-IS sends CSNPs periodically on all links. The receiving systems use the information in the CSNP to update and synchronize their LSP databases. The DR router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each LSP.

## Type/Length/Values

---

- **IS-IS information objects**

- Each piece of IS-IS information is defined as an object with three attributes:
  - Object type: a predefined code for the type of information contained in the object
  - Object length: the length of the information (allows for variable-length objects)
  - Object value: the actual information defined by the type attribute
- TLVs are the building blocks of IS-IS PDUs, which are used for information exchange
  - Some TLVs are used in multiple PDUs
  - Some TLVs are PDU-specific
- Similar to OSPF packet-specific and LSA-specific fields

Copyright © 2006, Juniper Networks, Inc.

### IS-IS Information Objects

IS-IS PDUs use TLV encoding as the basic structure for all routing information. TLV encoding requires that the length of any field be defined explicitly when a PDU uses that field. IS-IS ignores all unknown TLVs, making the protocol easily extensible.

## Configuring Juniper Networks Routers

### Some Common TLVs

TLV Code	Defined in	Used for
1	ISO 10589	Area Addresses
2	ISO 10589	IS Neighbor Metrics
6	ISO 10589	Neighbor LAN ID
8	ISO 10589	Padding
9	ISO 10589	LSP Entries
10	ISO 10589	Authentication
128	RFC 1195	IP Prefix, Mask, and Metrics
129	RFC 1195	Protocols Supported
130	RFC 1195	IP External Information
132	RFC 1195	IP Interface Address
137	RFC 2763	Dynamic Hostname Mapping

Copyright © 2006, Juniper Networks, Inc.

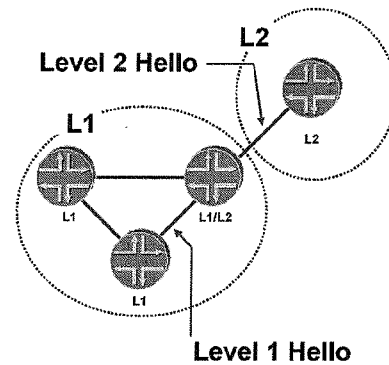
#### Common TLVs

The accompanying table shows the TLVs commonly used when using IS-IS as an IP routing protocol. Other TLVs are defined for MPLS traffic engineering that serve a purpose similar to OSPF Type 10 opaque LSAs.

## Configuring Juniper Networks Routers

### Neighbors and Adjacencies

- IS-IS adjacency rules:
  - Level 1 routers never form an adjacency with a Level 2 router
    - The reverse is also true
  - For Level 1 adjacencies, Area IDs must be the same
  - For Level 2 adjacencies, Area IDs can be different



Copyright © 2006, Juniper Networks, Inc.

### Neighbors and Adjacencies

- ▣ The slide lists the rules Level 1 and Level 2 routers must follow when forming an adjacency.

### Designated Intermediate System

- **IS-IS elects a designated intermediate system (DIS) on broadcast/multi-access networks**
  - In case of a priority tie, the system with the highest SNPA/MAC address wins the DIS election
  - Separate DIS is elected for L1 and L2 (could be the same router)
- **The IS-IS network is considered a router—called a pseudo-node**
  - Each router advertises a single link to the pseudo-node, including the DIS
  - Each router forms an adjacency with each of its neighbors on a broadcast, multi-access network
- **DIS characteristics:**
  - DIS acts as representative of the pseudo-node and advertises the pseudo-node to all attached routers
  - No backup DIS in IS-IS

Copyright © 2006, Juniper Networks, Inc.

### Designated Intermediate System Election

The IS-IS DIS election process is achieved by assigning a Level 1 priority and a Level 2 priority on every IS-IS router interface, with a range of 0 through 127. JUNOS software uses a default priority of 64 for both levels. The router advertises its priority in the hello PDUs sent from each interface. The L1 priority is advertised in L1 hello PDUs, and the L2 priority is advertised in L2 hello PDUs. If the priority is 0, the router is ineligible to become the DR. Interfaces to nonbroadcast networks automatically have a priority of 0. The router with the higher priority value becomes the DR. In the event of a tie, the router with the numerically highest Sub-network Point of Attachment (SNPA), which is a fancy name for a MAC address, wins the election.

### Pseudo-Node

IS-IS elects a DR on broadcast and multi-access networks for the same reason as OSPF. Rather than having each router connected to the LAN advertise an adjacency with every other router connected to the LAN, the network itself is considered a router—a pseudo-node. Each router, including the DR, advertises a single link to the pseudo-node. The DR also advertises, as the representative of the pseudo-node, a link to all of the attached routers.

### DIS Characteristics

Unlike OSPF, however, an IS-IS router attached to a broadcast, multi-access network establishes adjacencies with all of its neighbors on the network, not just the DIS. Each router multicasts its LSPs to all of its neighbors, and the DIS uses a system of PDUs—called sequence number PDUs—to ensure that the flooding of LSPs is reliable. Also unlike OSPF, there is no election of a backup designated router in IS-IS. If the IS-IS DIS fails, a new DIS is elected. Another characteristic is that if a new router with a higher priority than the existing DIS becomes active, or if the new router has an equal priority and a higher subnetwork point of attachment (MAC address), it becomes the new DIS. When the DIS changes, a new set of LSPs must be flooded.



### Troubleshooting IS-IS Adjacencies

---

- IS-IS adjacencies do not require or rely on IP configuration
  - Unlike OSPF
- If no adjacency, check for:
  - Physical/data link layer connectivity
  - Mismatched areas (if L1 router) and levels
  - Failure to support minimum MTU of 1492
  - Lack of, or malformed, ISO-NET
    - No NET configured
    - Failure to include 100 as an IS-IS interface

Copyright © 2006, Juniper Networks, Inc.

#### IP Configuration not Necessary

- When establishing adjacencies in OSPF, all routers on a link must agree upon the IP subnet to which they belong, except on point-to-point links, which may be unnumbered or use /32 addressing. This agreement is not necessary with IS-IS. IS-IS does not rely on IP; it simply carries IP information within its TLVs. Thus, it is possible for an IS-IS adjacency to come up, even if the routers on a link do not agree on the IP subnet to which they belong. As a result, it is possible to have an IS-IS adjacency established and not be able to ping the neighboring routers.

#### Troubleshooting No Adjacency

The slide provides a checklist to use when troubleshooting IS-IS adjacency problems.

### IS-IS Metrics

- **Metrics:**

- Any single link can have a maximum value of 63
- Arbitrary and assigned by network administrator
- IS-IS calculates path links by summing link values
- IS-IS interprets a single, required, default metric, with a maximum path value of 1023 as cost

- **Wide metrics:**

- Wide metrics are always sent, but...
  - Their value is limited to 63 unless you...
  - ...Disable standard metrics with level `level wide-metrics-only`
- Maximum value for any one link is  $2^{24}$  (~16 million)
- Maximum path value is  $2^{32}$
- Maximum network diameter is 256

Copyright © 2006, Juniper Networks, Inc.

### IS-IS Metrics

IS-IS uses a single, required, default metric with a maximum path value of 1023. The metric is arbitrary and typically is assigned by a network administrator.

Any single link can have a maximum value of 63; path metrics are calculated by summing link values. Maximum metric values were set originally at these levels to provide the granularity to support various link types while at the same time ensuring that the shortest-path algorithm used for route computation would be reasonably efficient.

IS-IS also defines three optional metrics, or *costs*:

- The *delay* cost metric reflects the amount of delay on the link;
- The *expense* cost metric reflects the communications cost associated with using the link; and
- The *error* cost metric reflects the error rate of the link.

IS-IS maintains a mapping of these four metrics to the quality-of-service (QoS) option in the CLNP packet header. IS-IS uses these mappings to compute routes through the internetwork.

### Wide Metrics

IS-IS also has wide metrics. In most cases, a range of 1–63 is insufficient to properly distinguish between high- and low-speed links. You can set the metric up to  $2^{24}$  (approximately 16 million). Wide metrics allow for the network diameter to be up to 256. This diameter results in a maximum total path value of  $2^{32}$ , or around 4.2 billion. JUNOS software sends both wide and standard metrics by default. However, to ensure a consistent topology, the software limits the wide metric value to 63 if it is also sending standard metrics. To benefit from the increased metric values, you must disable the sending of standard metrics on a per-level basis, with the `wide-metrics-only` knob.

### JUNOS Software IS-IS Support

---

- IS-IS support with:
  - Multiple areas and multiple levels
  - Authentication
  - Route summarization and route leaking
  - Traffic engineering TLVs and wide metrics
  - Mesh groups
  - Overload
  - Hello PDU interval, LSP lifetime, etc.
  - Graceful restart
  - External prefix limits

Copyright © 2006, Juniper Networks, Inc.

### JUNOS Software IS-IS Support

JUNOS software supports:

- Level 2 routers using the area ID field in the NET to support multiple area routing with Level 1, Level 2, or Level 1/Level 2 routing domains;
- Authentication, including simple or HMAC-MD5 formats;
- Policy-based route summarization for efficient inter-level routing and route leaking, which allows Level 2 prefix injection into Level 1 areas;
- Traffic engineering TLVs and support for wide metrics, which expand the metric range above 63;
- Mesh groups to reduce LSP flooding;
- The configuration of the router to appear overloaded, so that it cannot be used for transit traffic;
- The change of IS-IS timer intervals;
- Graceful restart based on *Restart Signaling for IS-IS* by Mike Shand (*draft-ietf-isis-restart-00.txt*);
- External prefix limits, which allows you to limit the number of external prefixes that are advertised by IS-IS; and
- Bidirectional Forwarding Detection (BFD) support. BFD is a new protocol designed to provide sub-second failure detection between routers connected with various media types. BFD is not a routing protocol in itself; BFD integrates with the IS-IS protocol to expedite reconvergence when the BFD daemon detects a failure. You enable BFD for IS-IS with a `set interface interface-name bfd-liveness-detection minimum-interval value` statement issued at the `[edit protocols isis]` hierarchy.

## Configuring Juniper Networks Routers

### Configuring IS-IS (1 of 2)

---

```
[edit protocols]
user@host# set isis interface ge-0/0/0.0 level 1 disable

[edit protocols]
user@host# set isis interface at-0/1/1.100 level 2 disable

[edit protocols]
user@host# show
isis {
  interface ge-0/0/0.0 {
    level 1 disable;
  }
  interface at-0/1/1.100 {
    level 2 disable;
  }
  interface lo0.0 {
    level 2 disable;
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

### Configuring IS-IS

By default, all interfaces specified as IS-IS interfaces are Level 1 and Level 2 interfaces. You might need to disable a particular level on a given interface. Notice that configuring IS-IS on a Juniper Networks M-series or T-series platform requires minimal configuration.

### Configuring IS-IS (2 of 2)

**You must include the ISO family on all interfaces on which you want to run IS-IS and a NET on one of the router's interfaces (usually lo0)**

```
[edit]
user@host# show interfaces
ge-0/1/0 {
  unit 0 {
    family iso;
    family inet {
      address 10.0.24.1/24;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.2.1/32;
    }
    family iso {
      address 49.0001.0192.0168.0201.00;
    }
  }
}
```

Copyright © 2006, Juniper Networks, Inc.

#### Configuring IS-IS

For IS-IS to run on the router, you must enable IS-IS on the router, configure a NET on one of the router's interfaces (preferably the loopback interface, lo0), and configure the ISO family on all interfaces on which you want IS-IS to run.

JUNOS software supports the assignment of multiple ISO NETs to the router's loopback interface. Such a configuration might prove helpful when migrating two previously independent IS-IS domains into a single routing domain.

### Monitoring IS-IS Operation

---

- Where we are going...
  - Various `show` commands exist to provide detailed information on the operation of IS-IS
    - `show isis interface`
    - `show isis adjacency`
    - `show isis spf log`
    - `show isis statistics`
    - `show isis route`
    - `show isis database`
  - IS-IS tracing

Copyright © 2006, Juniper Networks, Inc.

### Monitoring IS-IS Operation

The slide lists the topics discussed on the following pages.

### Displaying IS-IS Interface Status

Use the `show isis interface` command to display the IS-IS parameters associated with an interface

```
user@host> show isis interface ?
Possible completions:
  <[Enter]>           Execute this command
  <interface-name>   Interface name
  brief              Show brief status
  detail             Show detailed status

user@host> show isis interface
IS-IS interface database:
Interface      L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0          0  0x1 Disabled             Passive          0/0
so-0/1/0.0     2  0x1 Disabled             Point to Point   10/10
so-0/1/1.0     2  0x1 Disabled             Point to Point   10/10
so-0/1/2.0     2  0x1 Disabled             Point to Point   10/10
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying Interface Status

The `show isis interface` command displays the status of an interface. Use this command to ensure that IS-IS is configured correctly on the router. The output fields of this command are:

- `interface-name` (detail output only): Displays the name of the interface.
- `Index` (detail output only): Displays the interface index assigned by the JUNOS kernel.
- `State` (detail output only): Displays the internal implementation information.
- `Circuit ID` (detail output only): Displays the circuit identifier.
- `Circuit type` (detail output only): Displays the circuit type, which can be 1—Level 1 only, 2—Level 2 only, or 3—Level 1 and Level 2.
- `LSP interval` (detail output only): Displays the interface's LSP interval.
- `Sysid` (detail output only): Displays the system identifier.
- `Interface` (brief output only): Displays the interface through which the adjacency is made.
- `Level 1 DR/Level 2 DR` (brief output only): Displays the Level 1 or Level 2 designated intermediate system.
- `L1/L2 Metric`: Displays the interface's metric for Level 1 and Level 2. If there is no information, the metric is 0.
- `Adjacencies` (detail output only): Displays the number of adjacencies established on this interface.
- `Priority` (detail output only): Displays the priority value for this interface.
- `Metric` (detail output only): Displays the metric value for this interface.
- `Hello (s)` (detail output only): Displays the interface's hello interval.
- `Hold (s)` (detail output only): Displays the interface's hold time.

### Displaying IS-IS Adjacency Status

Use the `show isis adjacency` command to display IS-IS parameters adjacency status

```
user@host> show isis adjacency ?
Possible completions:
  <[Enter]>      Execute this command
  brief          Show brief status
  detail         Show detailed status
  system-id     Display entries for specified system

user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State          Hold (secs) SNPA
so-0/1/2.0     Denver          2 Up             24
so-0/1/3.0     SanFran         2 Up             28
so-0/2/2.0     Toronto         2 Up             23
so-0/2/3.0     Amsterdam      2 Up             26
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying IS-IS Adjacency Status

The `show isis adjacency` command displays the status of IS-IS adjacencies. The output fields of this command are:

- **Interface:** Displays the interface through which the neighbor is reachable.
- **System:** (brief output only): Displays the system identifier (sysid), printed as a name if possible.
- **L:** Displays the level, which can be 1—Level 1 only; 2—Level 2 only; or 3—Level 1 and Level 2. An exclamation point (!) preceding the level number indicates that the adjacency is missing an IP address.
- **State:** Displays the state of the adjacency. It can be Up, Down, New, One-way, Initializing, or Rejected.
- **Hold (secs)** (brief/standard output only): Displays the remaining hold time of the adjacency. Note that the `show isis adjacency` command returns brief output by default.
- **SNPA** (brief output only): Displays the subnetwork point of attachment (MAC address of the next hop).



### Displaying Detailed Adjacency Information

#### Display detailed IS-IS adjacency information with the detail switch

```
user@host> show isis adjacency detail
IS-IS adjacency database:

Denver
  Interface: so-0/1/2.0, Level: 2, State: Up, Expires in 25 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:15:27 ago
  Circuit type: 2, Speaks: IP
  IP addresses: 10.0.18.2
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying Detailed Adjacency Information

The `show isis adjacency detail` command displays detailed IS-IS adjacency information. The output fields of this command are:

- `Expires in` (detail output only): Displays how long until the adjacency expires, in seconds.
- `Priority` (detail output only): Displays the priority to become the designated intermediate system.
- `Up/Down transitions` (detail output only): Displays the count of adjacency status changes from up to down or from down to up.
- `Last transition` (detail output only): Displays the time of the last up/down transition.
- `Circuit type` (detail output only): Displays the bit mask of levels on this interface, which can be 1—Level 1 router, 2—Level 2 router, or 1/2—both Level 1 and Level 2 routers.
- `Speaks` (detail output only): Displays the protocols supported by this neighbor.
- `IP addresses` (detail output only): Displays the IP address of this neighbor.

### Clearing and Restarting IS-IS Adjacencies

#### Clear IS-IS adjacencies with the `clear isis adjacency` command

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System      L State      Hold (secs) SNPA
so-0/1/2.0     Denver     2 Up         24
so-0/1/3.0     SanFran    2 Up         28
so-0/2/2.0     Toronto    2 Up         23
so-0/2/3.0     Amsterdam  2 Up         26
```

```
user@host> clear isis adjacency Toronto
```

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System      L State      Hold (secs) SNPA
so-0/1/2.0     Denver     2 Up         22
so-0/1/3.0     SanFran    2 Up         26
so-0/2/3.0     Amsterdam  2 Up         24
```

Copyright © 2006, Juniper Networks, Inc.

#### Clearing and Restarting IS-IS Adjacencies

The `clear isis adjacency` command clears and restarts the adjacency process with a particular IS-IS neighbor.

### Displaying SPF Operation

Display information about SPF operation with the `show spf log` command

```
user@host> show isis spf log
IS-IS level 1 SPF log:
Start time           Elapsed (secs)  Count Reason
Tue Apr 17 16:13:26  0.000039       1 Reconfig
Tue Apr 17 16:13:32  0.000062       1 Updated LSP host.00-00
Tue Apr 17 16:28:26  0.000064       1 Periodic SPF

IS-IS level 2 SPF log:
Start time           Elapsed (secs)  Count Reason
Tue Apr 17 16:13:26  0.000025       1 Reconfig
Tue Apr 17 16:13:32  0.000051       1 Updated LSP host.00-00
Tue Apr 17 16:13:57  0.000087       2 New adjacency Denver on so-
0/1/2.0
Tue Apr 17 16:14:03  0.000241       6 Updated LSP Atlanta.00-00
Tue Apr 17 16:26:38  0.000298       1 Periodic SPF
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying SPF Operation

The `show isis spf log` command displays the elapsed time to perform SPF calculations and the reasons why they were triggered. The output fields of this command are:

- Node: Displays the sysid of a node.
- Metric: Displays the metric to the node.
- Interface: Displays the interface of the next hop.
- Via: Displays the sysid of the next hop.
- SNPA: Displays the subnetwork point of attachment (MAC address of the next hop).
- Start time (log output only): Displays the time that the SPF computation started.
- Elapsed time (log output only): Displays the length of time required to complete the SPF computation in seconds.
- Count (log output only): Displays the number of times the SPF was triggered.
- Reason (log output only): Displays the reason that the SPF computation was completed.

### Displaying IS-IS Statistics

Display various IS-IS statistics information with the `show isis statistics` command

```
user@host> show isis statistics
IS-IS statistics for host:
PDU type      Received  Processed    Drops      Sent      Reremit
LSP           17        17           0           2         0
IIH           190       190          0          571       0
CSNP          99        99           0          294       0
PSNP          3         3            0           12        0
Unknown       0         0            0           0         0
Totals        309       309          0          879       0

Total packets received: 309 Sent: 879

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0
SPF runs: 17
Fragments rebuilt: 5
LSP regenerations: 2
Purges initiated: 0
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying IS-IS Statistics

The `show isis statistics` command displays statistics about IS-IS traffic. The output fields of this command are:

- PDU type: Displays the protocol data unit type.
- Received: Displays the number of PDUs received since IS-IS started or since the statistics were zeroed.
- Processed: Displays the number of PDUs received less the number dropped.
- Drops: Displays the number of PDUs dropped.
- Sent: Displays the number of PDUs transmitted since IS-IS started or since the statistics were zeroed.
- Reremit: Displays the number of PDUs retransmitted since IS-IS started or since the statistics were zeroed.
- Total packets received/sent: Displays the total number of PDUs received and transmitted since IS-IS started or since the statistics were zeroed.
- SNP queue length: Displays the number of CSNPs and PSNPs sitting on the SNP queue waiting for processing. This value is almost always 0.
- LSP queue length: Displays the number of LSPs sitting on the LSP queue waiting for processing. This value is almost always 0.
- SPF runs: Displays the number of SPF calculations performed. If this number is incrementing rapidly, it indicates that the network is unstable.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Displaying IS-IS Statistics (contd.)

- `Fragments rebuilt`: Displays the number of LSP fragments that the local system has computed.
- `LSP regenerations`: Displays the number of LSPs that have been regenerated. An LSP is regenerated when it is nearing the end of its lifetime and it has not changed.
- `Purges initiated`: Displays the number of purges that the system initiated. A purge is initiated if the software decides that an LSP must be removed from the network.

### Displaying IS-IS Routes

```
user@host> show isis route
IPv4/IPv6 Routes
-----
IS-IS routing table                               Current version: L1: 3 L2: 5
Prefix      L Version Metric Type Interface  Via
10.0.0.0/24 2      5    20 int  so-0/1/2.0  Denver
10.0.1.0/24 2      5    20 int  so-0/1/2.0  Denver
10.0.2.0/24 2      5    30 int  so-0/1/2.0  Denver
10.0.8.0/24 2      5    30 int  so-0/1/2.0  Denver
. . .
user@host> show route protocol isis

inet.0: 64 destinations, 64 routes (64 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/24      *[IS-IS/18] 00:00:59, metric 30
> to 10.0.21.1 via fe-0/0/2.0
10.0.1.0/24      *[IS-IS/18] 00:00:59, metric 30
> to 10.0.21.1 via fe-0/0/2.0
10.0.2.0/24      *[IS-IS/18] 00:00:59, metric 40
to 10.0.21.1 via fe-0/0/2.0
> to 10.0.22.2 via so-0/1/1.0
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying IS-IS Routes

The `show isis route` command displays the routes in the IS-IS routing table. The output fields of this command are:

- **Current version:** Displays the number of the current version of the IS-IS routing table.
- **L1:** Displays the version of Level 1 SPF that was run.
- **L2:** Displays the version of Level 2 SPF that was run.
- **Prefix:** Displays the destination of the route.
- **L:** Displays the level, which can be 1—Level 1 only; 2—Level 2 only; and 3—Level 1 and Level 2.
- **Version:** Displays the version (or run) of SPF that generated the route.
- **Metric:** Displays the metric value associated with the route.
- **Type:** Displays the metric type. It can be `int` (internal) or `ext` (external).
- **Interface:** Displays the interface to the next hop.
- **Via:** Displays the system identifier of the next hop, displayed as a name if possible.

### Displaying IS-IS Database Entries

#### Display IS-IS database entries with the `show isis database` command

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
host.00-00            0x3     0x48ba    1132 L1 L2
1 LSPs

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
Denver.00-00          0x6     0xb978    1154 L1 L2
host.00-00            0x4     0x140c    1154 L1 L2
wash-dc.00-00        0x1d7   0x7f37    683 L1 L2
Atlanta.00-00        0x1d3   0xe603    1024 L1 L2
Atlanta.02-00        0x9     0xea81    1175 L1 L2
Houston.00-00         0x1c8   0x8b1c    677 L1 L2
Dallas.00-00         0x1ca   0x1571    1103 L1 L2
NewYork.00-00        0x1dd   0x178a    779 L1 L2
8 LSPs
```

Copyright © 2006, Juniper Networks, Inc.

#### Displaying IS-IS Database Entries

The `show isis database` command displays a brief view of all the LSPs in the IS-IS link-state database. The output fields of this command are:

- LSP ID: Displays the link-state PDU (LSP) identifier.
- Sequence: Displays the sequence number of the LSP.
- Checksum: Displays the checksum value of the LSP.
- Lifetime (secs): Displays the remaining lifetime of the LSP, in seconds.
- IP prefix: (detail and extensive output only) Displays the prefix advertised by this LSP.
- IS neighbor: (detail output only): Displays an IS-IS neighbor of the advertising system.
- Metric (detail and extensive output only): Displays the metric of the prefix or neighbor.

## Configuring Juniper Networks Routers

### Displaying Detailed IS-IS Database Information

```
lab@Tokyo> show isis database extensive
IS-IS level 1 link-state database:

Tokyo.00-00 Sequence: 0x1, Checksum: 0x1c90, Lifetime: 864 secs
  IP prefix: 192.168.20.0/24      Metric:      0 External Up
  IP prefix: 192.168.21.0/24      Metric:      0 External Up
  IP prefix: 192.168.22.0/24      Metric:      0 External Up
  IP prefix: 200.0.0.0/24         Metric:      0 External Up

Header: LSP ID: Tokyo.00-00, Length: 140 bytes
  Allocated length: 1492 bytes, Router ID: 192.168.20.1
  Remaining lifetime: 864 secs, Level: 1, Interface: 0
  Estimated free bytes: 1352, Actual free bytes: 1352
  Aging timer expires in: 864 secs
  Protocols: IP, IPv6

Packet: LSP ID: Tokyo.00-00, Length: 140 bytes, Lifetime : 1200 secs
  Checksum: 0x1c90, Sequence: 0x1, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 18, Packet version: 1, Max area: 0

TLVs:
  Area address: 49.0001 (3)
  Speaks: IP
  Speaks: IPv6
  IP router id: 192.168.20.1
  IP address: 192.168.20.1
  Hostname: Tokyo
  IP external prefix: 192.168.20.0/24, Internal, Metric: default 0, Up
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying Detailed IS-IS Database Information

The `show isis database extensive` command provides detailed output for the contents of the IS-IS link-state database. The output fields of this command are:

- **LSP ID:** Displays the link-state PDU (LSP) identifier.
- **Sequence:** Displays the sequence number of the LSP.
- **Checksum:** Displays the checksum value of the LSP.
- **Lifetime (in seconds):** Displays the remaining lifetime of the LSP, in seconds.
- **IP prefix (detail and extensive output only):** Displays the prefix advertised by this LSP.
- **IS neighbor (detail output only):** Displays an IS-IS neighbor of the advertising system.
- **Metric (detail and extensive output only):** Displays the metric of the prefix or neighbor.

The command output then displays detailed packet content information.



### Tracing IS-IS

- A typical IS-IS tracing configuration:

```
[edit protocols isis]
user@host# show
traceoptions {
  file isis-trace;
  flag error detail;
  flag hello detail;
  flag lsp detail;
}
```

- Monitor the resulting *isis-trace* log file using `monitor start log-file-name` or the `show log log-file-name` commands

Copyright © 2006, Juniper Networks, Inc.

#### IS-IS Tracing

To perform debugging functions on the IS-IS routing process, use the JUNOS software `traceoptions` feature. The trace output (debug information) is directed to the named log file, which is stored in the `/var/log` directory on the router's hard drive. You can view the log file using the `monitor start` or `show log` operational mode commands. In addition to specifying the trace file, you also must tell the router what information you want to trace. You accomplish this by specifying one or more `flag` keywords.

While you can only direct tracing to a single file, you can trace many options by using the `flag` keyword multiple times. In addition, you can add granularity by using the `detail`, `receive`, and `send` flag modifiers.

Available tracing flags for IS-IS include:

<code>all</code>	Trace everything
<code>csn</code>	Trace complete sequence number (CSN) packets
<code>error</code>	Trace errored packets
<code>general</code>	Trace general events
<code>graceful-restart</code>	Trace graceful restart events
<code>hello</code>	Trace hello packets
<code>lsp</code>	Trace link-state packets
<code>lsp-generation</code>	Trace LSP generation
<code>normal</code>	Trace normal events
<code>packets</code>	Trace IS-IS packets
<code>policy</code>	Trace policy processing
<code>psn</code>	Trace partial sequence number (PSN) packets
<code>route</code>	Trace routing information
<code>spf</code>	Trace SPF events
<code>state</code>	Trace state transitions
<code>task</code>	Trace routing protocol task processing
<code>timer</code>	Trace routing protocol timer processing

## Configuring Juniper Networks Routers

### Review Questions

---

1. What is the purpose of the IS-IS DIS?
2. What is the rationale behind a multi-level IS-IS design?
3. How would you configure a Juniper Networks M-series or T-series platform to run only Level 2 IS-IS on an interface?
4. List and describe three CLI commands you can use to monitor and troubleshoot the IS-IS protocol.

Copyright © 2006, Juniper Networks, Inc.

#### This Module Discussed:

- IS-IS operation;
- The IS-IS DIS;
- IS-IS areas and levels;
- Configuring IS-IS in JUNOS software;
- Using operational-mode commands to monitor and troubleshoot IS-IS; and
- Displaying and interpreting the IS-IS LSDB.

## Configuring Juniper Networks Routers

### Lab 6: IS-IS Configuration Lab

---

#### Lab Objective:

Configure IS-IS and use the CLI to monitor and troubleshoot its operation

Copyright © 2006, Juniper Networks, Inc.





## Configuring Juniper Networks Routers

### *Module 8: BGP Operation, Configuration, and Troubleshooting*

Copyright © 2006, Juniper Networks, Inc.

CJNR-M-7.a.7.6.1

### Module Objectives

---

- **After successfully completing this module, you will be able to:**
  - Compare the operation of EGPs to IGPs
  - Describe the use of common BGP attributes, such as local preference, MED, and AS path
  - Describe the JUNOS software BGP route selection algorithm
  - Explain why IBGP peering normally is done between loopback addresses
  - Describe the default BGP route advertisement rules
  - Configure BGP
  - Monitor and troubleshoot BGP operation using the CLI

Copyright © 2006, Juniper Networks, Inc.

#### **This Module Discusses:**

- EGP to IGP comparison;
- The use of common BGP attributes;
- The JUNOS software BGP route selection procedure;
- The rationale behind loopback peering for IBGP;
- Default BGP route advertisement rules;
- Configuring BGP in JUNOS software; and
- Using operational-mode commands to monitor and troubleshoot BGP.

# Border Gateway Protocol

---

- **Where we are going...**
  - **What is BGP?**
  - **BGP fundamentals and connections**
  - **BGP attributes and active route selection**
  - **BGP peering**
    - **Internal BGP (IBGP)**
    - **External BGP (EBGP)**
  - **BGP route advertisement rules**
  - **JUNOS support for BGP**
  - **Basic BGP configuration**
  - **Monitoring BGP operation**

Copyright © 2006, Juniper Networks, Inc.

## Border Gateway Protocol

The following pages examine the topics shown on the slide.

### What Is BGP?

---

- **BGP**

- Is an inter-domain routing protocol that communicates prefix reachability
- Is a *path vector* protocol
- Views the Internet as a collection of autonomous systems
- Supports CIDR
- Exchanges routing information between *peers*
- Is defined in RFC 1771

Copyright © 2006, Juniper Networks, Inc.

### What Is BGP?

The Border Gateway Protocol (BGP) is an inter-AS routing protocol and is sometimes called a *path-vector routing protocol* because it uses an autonomous system path, used as a vector, to prevent inter-domain routing loops. The term *path vector*, in relation to BGP, means that BGP routing information includes a series of AS numbers, indicating the path that a route takes through the network.

BGP exchanges routing information among ASs or domains. An AS is a set of routers that operate under the same administration. BGP routing information includes the complete route to each destination. BGP uses the routing information to maintain a database of network layer reachability information (NLRI), which it exchanges with other BGP systems. BGP uses the NLRI to construct a graph of AS connectivity, thus allowing BGP to remove routing loops and enforce policy decisions at the AS level.

BGP is classless routing protocol, which supports prefix routing, regardless of the class definitions of IPv4 addresses. BGP routers exchange routing information between peers. The peers must be connected directly for inter-AS BGP routing (unless certain configuration changes are done). The peers are TCP peers, which is addressed later in this module.

RFC 1771 defines BGP.



### BGP Fundamentals

---

- Each BGP update contains one path advertisement and attributes
  - Many prefixes can share the same path
- Routes consist of a destination prefix along with an AS path and other BGP-specific attributes
- BGP compares the AS path and other attributes to choose the best path

Copyright © 2006, Juniper Networks, Inc.

#### BGP Updates

Routes in BGP consist of destination networks and attributes associated with those routes. Each BGP update contains one path advertisement. However, many destinations can share the same path.

#### BGP Routes

Once a connection is open and active, BGP sends routes. BGP routes consist of destination prefixes, each associated with BGP attributes. (For IGPs, *metrics* is the term used to describe their attributes.) Some of the complexities of BGP are the variety of these metrics (or *attributes*), the order of their execution, and various rules that can be applied to the attributes.

#### Selecting the Active BGP Route

Both you and BGP itself can associate one or more attributes with a route advertisement. Attributes carry descriptive information about the route and are used in choosing the best path to a destination.

BGP attributes describe:

- The next hop for a packet sent to a particular destination;
- Various numeric-type attributes;
- The path through ASs that a routing announcement has traversed to arrive at the destination where it is now; and
- The method of generation for the prefix, or which protocol originated the route.

BGP re-advertises a prefix as *unfeasible* when that route should be removed from a peers RIB-In. A single BGP update message can contain both feasible and unfeasible route, when these routes share a common set of attributes.

### BGP Connections

- **BGP updates are incremental**
  - No regular refreshes
  - Except at session establishment, when volume of routing can be high
- **BGP runs over TCP connections**
  - TCP port 179
  - TCP services
    - Fragmentation
    - Acknowledgments
    - Checksums
    - Sequencing
    - Flow control
  - No automatic neighbor discovery

Copyright © 2006, Juniper Networks, Inc.

### BGP Updates

BGP updates are incremental; BGP does not require that routing table information be refreshed. The receiver considers a route to be valid until the route originator withdraws the route or the connection is terminated.

### BGP Uses TCP for Transmission

BGP uses TCP as its transport protocol (port 179). TCP provides a full-duplex, connection-oriented, reliable, byte-stream service to BGP. BGP considers a connection between two peers to be idle until a TCP connection is established between them. With the TCP connection established, the endpoints are assured of a reliable connection. The following are TCP services:

- *Fragmentation*: TCP takes the BGP data and fragments it, if necessary, into the best size *chunks*, or segments, to send.
- *Acknowledgments*: When it sends a segment, TCP maintains a timer while waiting for a receipt acknowledgment from the other end of the connection. If an acknowledgment is not received, TCP retransmits the segment, although unlike UDP, it does not retransmit the entire message. Upon successful receipt of the data, the receiving end acknowledges the data. The acknowledgment actually is delayed up to 200 milliseconds to determine if there is any data to send along with the acknowledgment.
- *Checksums*: TCP also maintains a checksum on its header and data. The sender includes the checksum in the TCP datagram, and upon receipt, the receiver calculates a checksum. If the checksums differ, TCP discards the datagram, sends no acknowledgement, and thus, the sender initiates a retransmission.

*Continued on next page.*

## Configuring Juniper Networks Routers

### BGP Uses TCP for Transmission (contd.)

- *Sequencing:* TCP also includes a sequence number in the header to ensure that data is received in the proper sequence. TCP requires this sequence number because TCP segments are sent in IP datagrams, which can arrive out of sequence because IP is connectionless. TCP also discards any duplicate IP datagrams that it encounters.
- *Flow control:* TCP provides flow control. The receiving end communicates to the sender how much buffer space it has available and allows the sender to send only that much data.

### BGP Neighbor States

---

- **TCP connectivity**
  - Idle
  - Connect
  - Active
- **BGP connectivity**
  - OpenSent
  - OpenConfirm
  - Established

Copyright © 2006, Juniper Networks, Inc.

#### TCP Connectivity

To exchange traffic using BGP, we first must establish a TCP connection. The following list shows the possible states for these sessions:

- *Idle*: This is the initial neighbor state. All BGP connections are refused. BGP waits for a start event, which is usually caused by configuring a BGP session or resetting an existing session. The router also might try to initiate a start event.
- *Connect*: In this state, BGP waits for a TCP connection to be completed.
- *Active*: In this state, a TCP timeout has occurred and the TCP session is not established yet. The ConnectRetry timer is restarted, and the system continues to listen for completion of TCP session. This state is bypassed when a TCP session is established before the ConnectRetry expires.

#### BGP Connectivity

After establishing a TCP session, we now can attempt to create communication at the BGP level. The following list shows the states for BGP connectivity:

- *OpenSent*: The TCP connection is established and an open message has been sent to the remote peer. BGP waits to receive an open message from the peer.
- *OpenConfirm*: An open message has been received from the remote peer, and BGP now waits for either a keepalive message or a notification message. If a keepalive is received, the state machine transitions to established. If a notification message is received, the state machine transitions to idle.
- *Established*: The BGP peers are fully adjacent and can exchange keepalive, update, and notification messages.

### **BGP Attributes: Local Preference**

---

- **Determines the preferred path out of the AS**
- **All BGP traffic in an AS flows toward the peer with the highest local preference value**
- **Values are used only within an individual AS**
  - **Nothing is sent across EBGP links**

Copyright © 2006, Juniper Networks, Inc.

#### **The Local-Preference Attribute**

Local preference typically is used when an AS wants to direct all outbound traffic to a certain peer. Before sending the traffic to the internal peers, the designated peer sets the local preference value on all routes received. Then all the peers use those routes in their RIB-local tables.

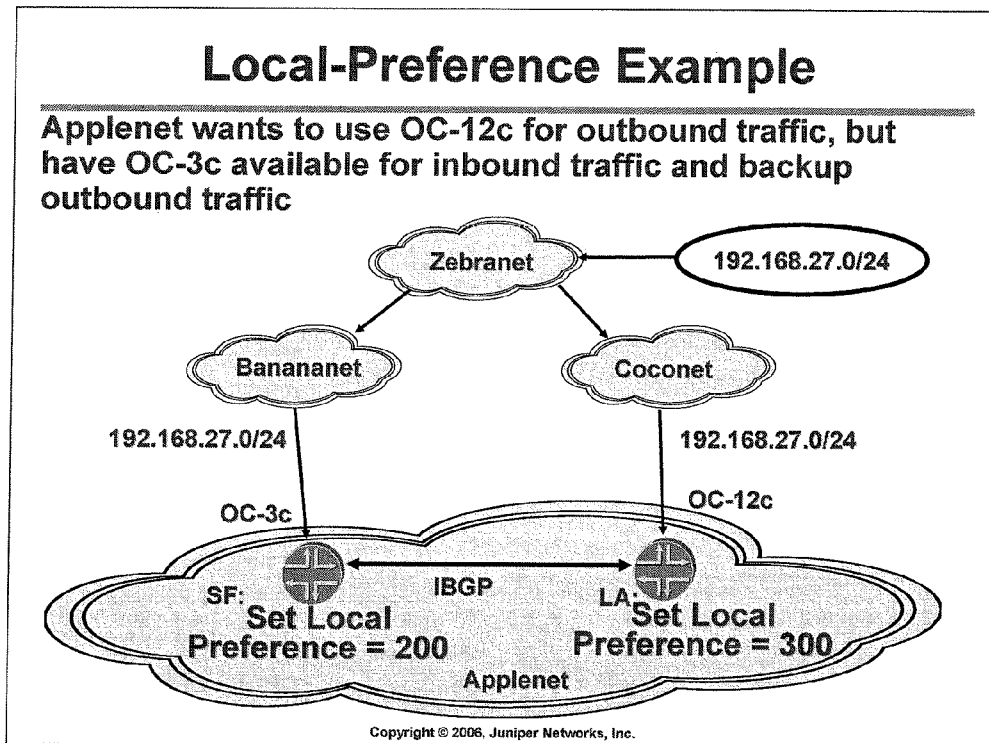
#### **Best Value of Local-Preference Attribute**

The local-preference attribute is a numeric value. Higher values indicate a better metric.

#### **Scope of Local-Preference Attribute**

BGP uses the local preference attribute only within an AS. Local-preference values are not transmitted across EBGP links. In JUNOS software, you can set the local-preference value with BGP configuration or through routing policy.

## Configuring Juniper Networks Routers



### One Local-Preference Example

The slide shows an example of how the local-preference attribute is used within an AS. The network administrators in Applenet have decided that the routers in Applenet should use the Los Angeles router to reach the 192.168.27.0/24 network in Zebranet. This decision is based on the greater bandwidth capacity available on that link: an OC-12c runs at 622 Mbps while an OC-3c runs at 155 Mbps.

To use the Los Angeles path, the Applenet network administrators set the local-preference attribute on the route to 192.168.27.0/24 advertised to the San Francisco router by Banananet to 200, and they set the local-preference attribute on the route to 192.168.27.0/24 advertised to the Los Angeles router by Coconet to 300 through the use of import policy on the EBGP peering sessions. BGP then sends these routes with the modified local-preference value to all IBGP speakers in Applenet. The default local-preference value is 100, which is assumed when a route does not carry the local-preference attribute.

In contrast to almost every other metric associated with routing protocols, the highest local preference is better. So for this route, the exit point of the AS is through Los Angeles. IBGP shares these values with all other routers in Applenet.

The link on the San Francisco router still can be used for inbound traffic flows from Zebranet and for outbound failover traffic if the OC-12c is not useable because of a router or link failure.

### BGP Attributes: AS Path

---

- Provides a path back to the source of the route, preventing routing loops
  - Routes with the router's own AS number in the path must be looped; these routes are dropped immediately
- Each router on the edge of the AS adds its AS number to the front of the path, for example:
  - 34 67 195 6743 701
- AS-path attribute is always present and is transmitted across all BGP links

Copyright © 2006, Juniper Networks, Inc.

#### Purpose of AS-Path Attribute

The AS-path attribute describes the path of autonomous systems that the route has been through since it was sourced into BGP. In addition, individual AS numbers can be added to the AS path. When a BGP router receives routes in an update message, the first action is to examine the current AS path to see if the local AS number is in the path. If it is, this indicates that the route has been through the AS already; accepting the route would cause a loop. Therefore, BGP drops the route.

#### AS-Path Attribute Creation

By default, the AS-path value is changed as a route transitions between autonomous systems. As the route leaves an AS, the BGP router adds the local AS number to the front of the path before sending it to a peer. Any BGP router is allowed, using a policy, to add information to the AS-path attribute. A router is *never* allowed to remove information from the path. Typical convention for the router is to add its local AS number to the path multiple times. These multiple entries discourage routers in other autonomous systems from using that router because it now has a longer AS-path attribute.

#### AS-Path Attribute Presence in BGP Updates

The AS-path attribute is mandatory; thus, it always must be present for all BGP routes.

### BGP Attributes: Origin

---

- Added by the router that added the route to BGP
- Describes where the first router received the information
  - I = IGP (0)
  - E = EGP (1)
  - ? = Incomplete (2)
- Attribute is always present and is transmitted across all BGP links

Copyright © 2006, Juniper Networks, Inc.

#### The Origin Attribute

The BGP router that injects a route into the BGP process is responsible for placing the origin attribute into the route.

#### Values of Origin Attribute

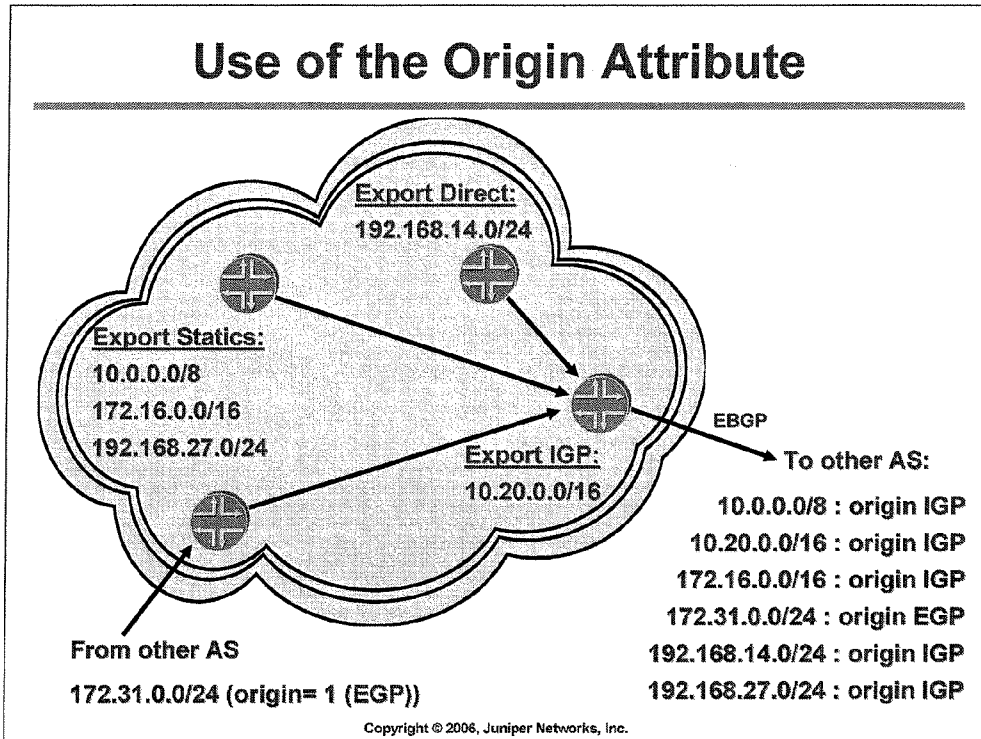
The origin attribute describes where the original router received the route. The possible choices are:

- *IGP*: BGP assigns an IGP route a value of 0. Examples of IGP routes include OSPF, IS-IS, static, and aggregate.
- *EGP*: BGP assigns EGP routes a value of 1. EGP routes are from the original EGP protocol, which was the predecessor to BGP.
- *Unknown*: BGP assigns unknown routes a value of 2. Unknown routes are those that did not come from an IGP or from EGP.

#### Origin Attribute Presence in the BGP Updates

The origin attribute is a mandatory attribute; it is transmitted across all BGP links. A BGP router can modify this attribute through the use of policy.





### Origin and JUNOS Software

The slide shows the default BGP behavior within JUNOS software with regard to the origin code.

Within the AS shown, the static routes of 10.0.0.0/8, 172.16.0.0/16, and 192.168.27.0/24 exist. An export policy placed these routes into BGP. There is a direct route of 192.168.14.0/24 that was exported into BGP. The route to 172.31.0.0/24 was learned from another AS altogether, and this route contains an origin attribute coded to 1, which indicates an EGP origin. Finally, there is an IGP-learned route of 10.20.0.0/16 in the network. The router does not know whether this route is an OSPF route or an IS-IS route, but the route had the appropriate export policy and, therefore, was placed into BGP.

To the Juniper Networks M-series or T-series router, it does not matter that these routes are advertised to another AS through EBGP; the BGP origin code is not altered as the routes are advertised to an EBGP peer. On the basis of the origin attribute alone, the 172.31.0/24 prefix appears less attractive to the remote AS.

### BGP Attributes: MED

---

- By default, used only when multiple links exist between the same two ASs
- Used to help influence the preferred path back into an AS
  - Lower values are better
- Presence on a route not necessary
  - When present, it is transmitted on all BGP links

Copyright © 2006, Juniper Networks, Inc.

#### Use of MED Attribute

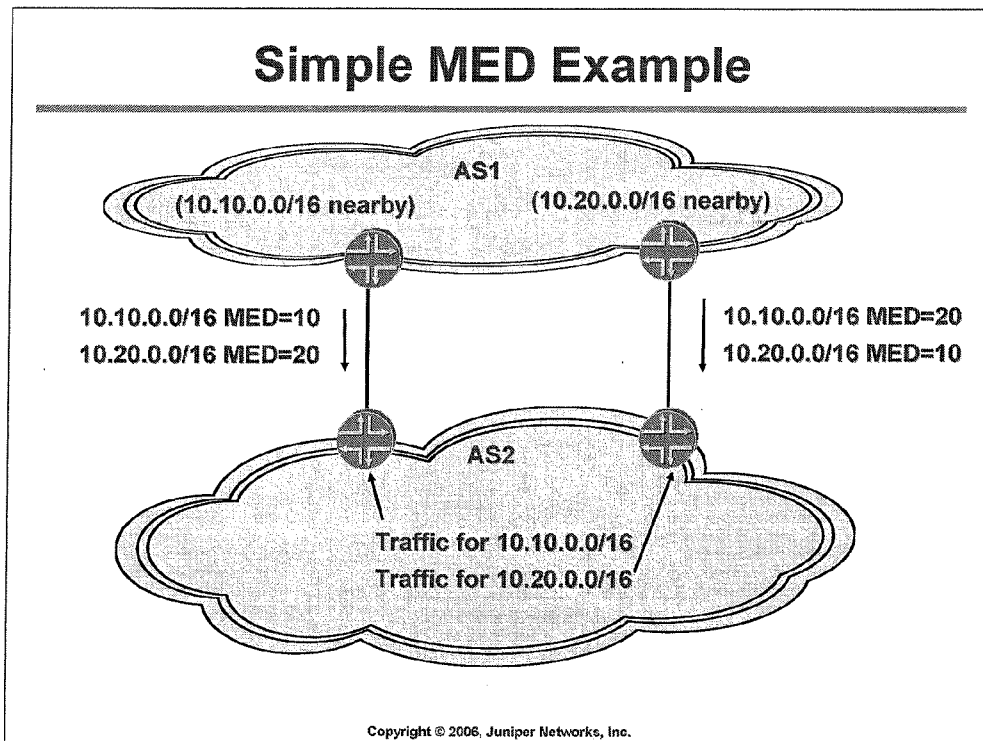
By default, BGP uses the multiple exit discriminator (MED) value only when the BGP router's AS has two or more connections to the same upstream AS. The existence of multiple connections can be determined by examining the AS-path attribute to find that multiple routes from the same AS being advertised by multiple, different peers exist.

#### MED Attribute Value

An AS uses the MED value in an attempt to influence data traffic headed back toward the AS. The local AS sets the MED value differently on separate peers headed toward the same remote AS. The remote AS picks routes based on the lowest MED value it finds. The remote AS then uses that local peer to route traffic back to the local AS.

#### MED Attribute Is Optional

BGP routes do not require the MED attribute. If it is missing, BGP assumes the MED value to be 0.



### How to Use MED

The slide shows a very basic example of the use of the BGP MED attribute to influence inter-AS traffic flows.

AS1 has assigned its IP address spaces so that it can summarize its network into two major segments. Furthermore, AS1 is divided relatively cleanly into networks that are near the left-most router (10.10.0.0/16 networks are nearby) and networks that are near the right-most router (10.20.0.0/16 networks are nearby). This split could be between Eastern and Western operations or many other alternatives.

AS1 has two EBGP sessions to AS2 and advertises both the 10.10.0.0/16 and the 10.20.0.0/16 networks to AS2 on each EBGP session, as shown on the slide.

Naturally, AS1 would like AS2 to return traffic to the closest point in the AS1 network so that *timely* packet delivery and low latency can be achieved. Ordinarily, AS1 would have no real way to convey this desire to AS2, and AS2 simply would send traffic to AS1 over whichever router AS2 decided to use, based on its own routing policies.

However, the MED attribute offers a way for AS1 to influence the routing policy of AS2 for traffic sent to AS1. To accomplish this *closest point* goal, AS1 alters the MED values on the routes that it advertises to AS2 with a BGP export routing policy.

AS1 advertises both networks across both links for redundancy. All things being equal, the routers in AS2 still see multiple network paths to the routers in AS1 as the AS1 routes are passed along throughout AS2. However, the AS2 routers will use the MED values of 10 and 20 (10 being preferred) to choose the BGP path to install in their local routing tables.

Thus, within AS2, traffic to 10.10.0.0/16 networks flows to the left-most router and out to AS1, while traffic to 10.20.0.0/16 networks flows to the right-most router and out to AS1. AS1 has influenced AS2 and, at the same, has time achieved a primitive type of load balancing.

### BGP Attributes: Community

---

- Generic mechanism for tagging routes
- Communities can be:
  - Used by policy to perform an action on a particular set of routes tagged with a community
  - Added to the community list (`community add command`)
  - Deleted from current community list (`community delete command`)
  - Set to the community list (`community set command`)

Copyright © 2006, Juniper Networks, Inc.

#### The Community Attribute

A BGP community is a group of destinations that share a common property. Communities are used to tag certain routes that can be identified easily later for a variety of purposes. BGP includes community attribute information as a path attribute in BGP update messages.

#### Use of the Community Attribute

You can define routing policy matches based on the BGP communities. You can associate multiple communities with the same BGP route.

### BGP Route Selection Part 1

---

1. Can the BGP next hop be resolved?
2. Prefer the highest local preference value
3. Prefer the shortest AS-path length
4. Prefer the lowest origin value
5. Prefer the lowest MED value
6. Prefer routes learned using EBGP over routes learned using IBGP

Copyright © 2006, Juniper Networks, Inc.

### Choosing the Active Route Part 1

When a BGP router has the same route in multiple RIB-in tables, it must decide which one to use and place into the RIB-local. The following describes the BGP route selection algorithm—the decision process taken by the local router to decide which route is best. We describe in the following pages the various attributes used in the selection process:

1. The router first must verify that it has a current route in the `inet.0` routing table to the IP address in the BGP next-hop attribute field.
2. BGP then compares routes for the *highest* local preference (the only choice based on a higher, rather than lower, value).
3. BGP evaluates the AS-path attribute next, where a shorter path is preferred. This attribute is often a common tiebreaker for routes.
4. BGP evaluates the origin code. The lowest origin code is preferred.
5. If any of the remaining routes are advertised from the same neighboring AS, BGP checks the MED attributes for a lowest value. The absence of a MED value is interpreted as a MED of 0.
6. If multiple routes remain, BGP prefers any routes learned via an EBGP peer over routes learned via an IBGP peer. If all remaining routes were learned through EBGP, skip to Step 9.

*Continued on next page.*

### BGP Route Selection Part 2

#### 7. Prefer routes with the lowest IGP metric

- Examine routing tables `inet.0` and `inet.3` for the BGP next hop, and install the physical next hop(s) for the route with the better preference
- For preference ties, install the physical next hop(s) found in `inet.3`
- For preference ties within the same routing table, install the physical next hop(s) where the greater number of equal-cost paths exist

#### 8. Prefer paths with the shortest cluster length

#### 9. Prefer routes from the peer with the lowest RID

- For external routes from different ASes, do not alter active route based on lowest RID (prevents MED oscillation)

#### 10. Prefer routes from the peer with the lowest peer ID

Copyright © 2006, Juniper Networks, Inc.

### Choosing the Active Route Part 2

7. If the remaining routes were learned through IBGP, use the path with the lowest IGP cost to the IBGP peer. For each IBGP peer, install a physical next hop(s) based on the following three rules:
  - a) BGP examines both the `inet.0` and the `inet.3` routing tables for the BGP next-hop value. The physical next hop(s) of the instance with the lowest JUNOS software preference is used. Often, this means that BGP uses the `inet.3` version of the next hop, via an MPLS LSP.
  - b) Should the preference values in the `inet.0` and the `inet.3` routing tables tie, the physical next hop(s) of the instance in `inet.3` is used.
  - c) When a preference tie exists and the instances are in the same routing table, BGP examines the number of equal-cost paths of each instance. The physical next hop(s) of the instance with more paths is installed. This might occur when you use the traffic engineering `bgp-igp` knob for MPLS.
8. BGP then examines the cluster list attribute for the shortest length. The cluster list is similar in function to an AS path.
9. At this point BGP normally prefers the route advertised from the peer with the lowest router ID (usually the loopback IP address). However, to prevent MED-related route oscillation, when comparing external routes from two neighboring ASs, the currently active route continues to be preferred when the routes are equal up to the router-ID comparison step.
10. The final tiebreaker is the peer ID value, which is the IP address of the BGP peering session.

### BGP Peering

- BGP sessions are established between peers
  - BGP speakers
- Two types of peering sessions
  - EBGP (external) peers with different ASs
  - IBGP (internal) peers within the same AS
    - IGP connects BGP speakers within the AS
    - IGP advertises internal routes

Copyright © 2006, Juniper Networks, Inc.

#### BGP Peers

BGP is a protocol in which routing information exchanges occur between exactly two nodes, called *peers*. These peers can be connected either directly or remotely.

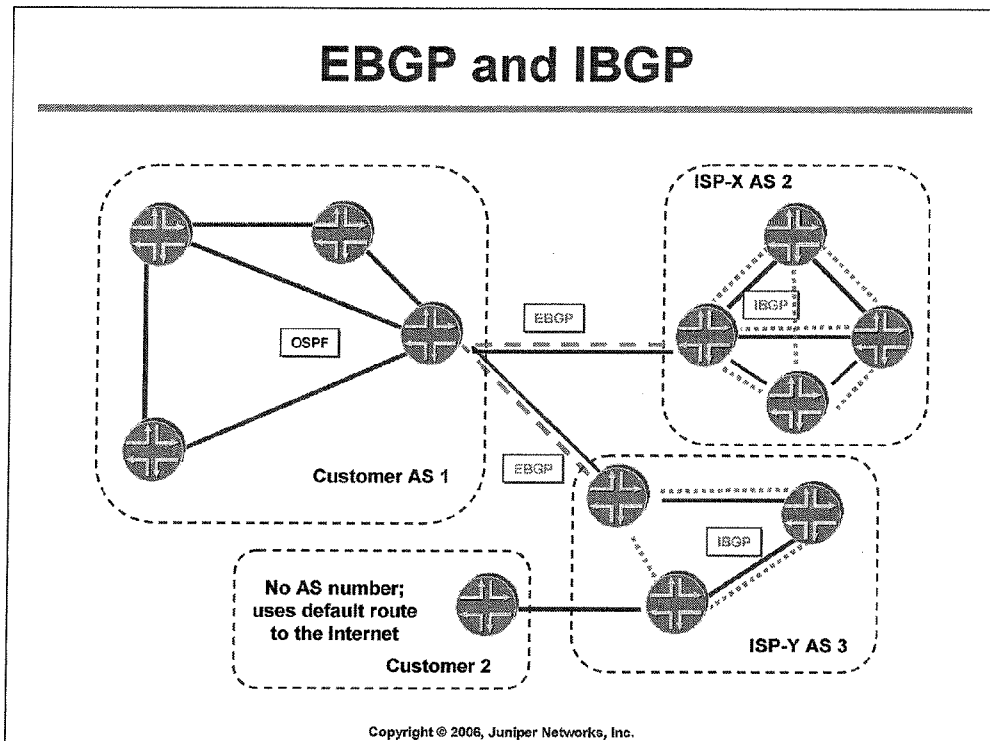
#### EBGP versus IBGP

BGP supports two different types of exchanges of routing information. Exchanges between ASs are called *external BGP* or *EBGP sessions* and handle inter-autonomous system routing. Exchanges within an AS are called *internal BGP* or *IBGP sessions*, and handle intra-autonomous system routing.

An EBGP peer connection is between a device in one AS and another device in a different AS. The connection between the two ASs consists of a physical connection and a BGP connection. The physical connection is a shared data link layer subnetwork between the two ASs. On this shared subnetwork, each AS has at least one border gateway belonging to that AS. The BGP connection exists between BGP speakers in each of the ASs. This session can communicate destinations that can be reached through the advertising AS. The EBGP connection typically is established between immediately connected devices located in two different ASs because the time-to-live (TTL) value of the EBGP packets is equal to 1, by default.

An IBGP connection is typically established between loopback interfaces of the routers not immediately connected (of course, everything depends on the AS's topology). BGP uses the loopback interfaces for stability reasons—these interfaces are always *alive*, unless the router itself *dies*. Because the IBGP connection typically exists between remotely connected routers, there is a requirement for an IGP within the AS. BGP's TCP session is established using regular routing tables.

## Configuring Juniper Networks Routers



### EBGP and IBGP

A BGP system shares network reachability information with adjacent BGP systems referred to as *neighbors* or *peers*.


BGP systems are arranged into groups. In an IBGP group, all peers in the group, called *internal peers*, are in the same AS. Internal peers can be anywhere in the AS and need not be directly connected to each other. Internal groups use routes from an internal gateway protocol (IGP) to resolve forwarding addresses. They also propagate external routes among all other internal routes running IBGP, and they compute the next hop by taking the BGP next-hop attribute received with the route and resolving it using information from one of the interior gateway protocols. A full-mesh configuration must be maintained between all IBGP peers in an AS.

In an EBGP peer group, the peers are called *external peers*, are in different ASs, and normally share a subnet. In an external group, the next hop is computed with respect to the interface shared between the external peers and the local router.



### IBGP Loopback Interfaces

- **IBGP peering is often done using loopback interfaces**
  - More stable
  - Not tied to a single physical path
- **The AS needs an IGP so that IBGP speakers can reach each others' loopback addresses**



#### IBGP Use of Loopback Interfaces

IBGP peers often use loopback interfaces. The advantage of using loopback interfaces is that they eliminate a dependency that would otherwise occur when you use the IP address of a physical interface to configure BGP. The slide shows a network in which using the loopback interface is advantageous.

On the slide, routers A and B run IBGP within AS 1. If router A were to specify the IP address of an Ethernet interface on router B in the remote neighbor with the router configuration command, and if the specified interface were to become unavailable, router A would not be able to establish a TCP connection with router B. Instead, router A specifies the IP address of the loopback interface that router B defines. When the loopback interface is used, BGP does not have to rely on the availability of a particular interface for making TCP connections. Router A specifies the IP address of the loopback interface (192.168.255.2) of router B.

Note that BGP rarely uses loopback interfaces between EBGP peers because EBGP peers usually are connected directly. EBGP peers therefore depend on a particular physical interface for connectivity (however, exceptions include parallel paths).

#### IGP Requirement

The AS needs IGP or static routes so that IBGP speakers can establish TCP sessions to each others' loopback interfaces.

### BGP Route Advertisement Rules

---

- **Advertise only the *active* BGP routes to peers**
  - BGP next-hop attribute must be reachable
- **Never forward IBGP routes to IBGP peers**
  - Prevents loops
- **Withdraw routes if active BGP routes become unreachable**

Copyright © 2006, Juniper Networks, Inc.

#### **BGP Advertises Only Active Routes**

Prior to route advertisement, BGP always checks if the routes are active. Should the route be known only to BGP (not IGP, because preference for IGP is lower than for BGP), BGP checks if the route's BGP next-hop attribute is reachable. If it is, only then the route becomes active.

#### **IBGP-to-BGP Route Updates**

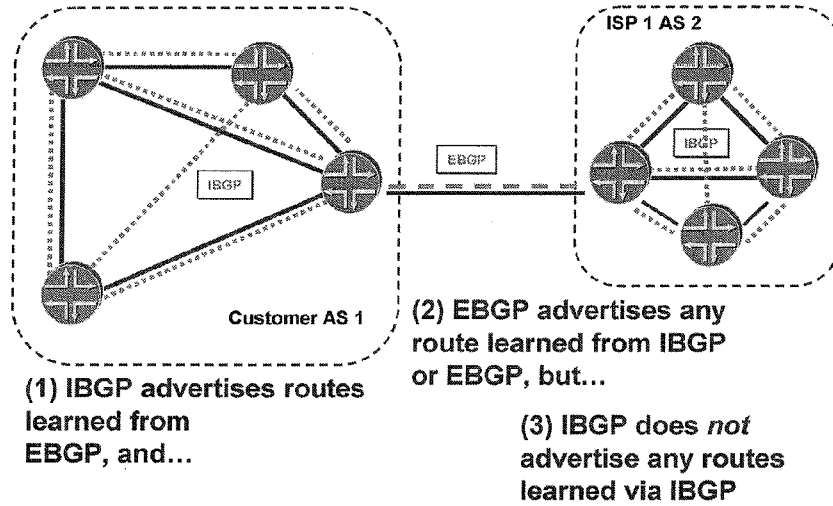
All IBGP connections (unless you use route reflections or confederations) must be fully meshed. Every pair of routers within an AS must have an IBGP session.

The full mesh of IBGP connections is totally independent of the physical connectivity. In other words, a physical link does not have to connect directly two IBGP peers. An IBGP session can be in place between two IBGP peers even if the session is established over multiple links to establish connectivity. This requirement is necessary for loop prevention. Remember, EBGP loop prevention is achieved by keeping track of the AS-path attribute for every advertised route. The AS-path attribute for IBGP is meaningless because it does not change within the AS. Consequently, this strict rule for fully meshed IBGP peering is a must.

#### **BGP Routes Withdrawal**

BGP withdraws all the routes that become unreachable.

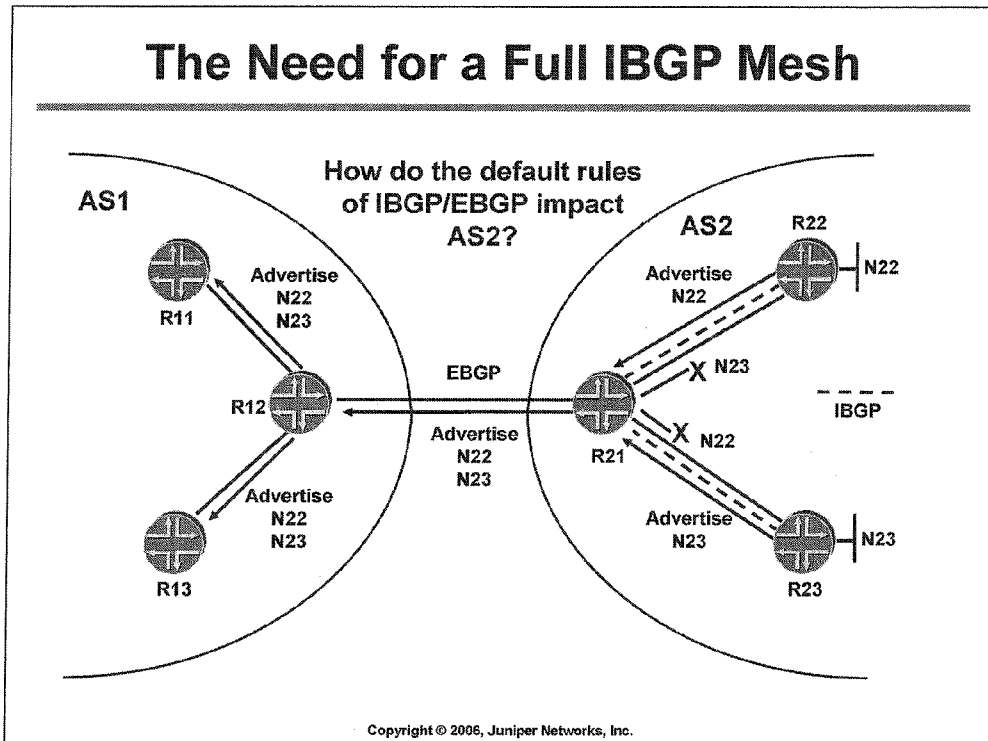
### Default BGP Advertisement Rules



Copyright © 2006, Juniper Networks, Inc.

### Default BGP Advertisement Rules

JUNOS software implements the default BGP advertisement rules as follows: EBGP always advertises any learned route, and IBGP advertises only routes learned from EBGP. The purpose of the advertisement rules is to prevent routing loops on a BGP network.



## The Need for a Full IBGP Mesh

Because IBGP only advertises routes learned from EBGP, on the slide, router 22 does not have a route to network 23. Also, router 23 does not have a route to network 22 without a direct IBGP link between router 22 and router 23. Routers 11, 12, and 13, on the other hand, have paths to both networks 22 and 23 because router 12 learns the routes from EBGP.

## JUNOS Software and BGP

---

- **Where we are going...**
  - JUNOS software support for BGP
  - Configure BGP
  - Monitor BGP

Copyright © 2006, Juniper Networks, Inc.

### JUNOS Software and BGP

- ☛ The following pages describe JUNOS software support for BGP, BGP configuration, and operational analysis of BGP.

### JUNOS Software BGP Routing Table

---

- BGP routes are placed in the JUNOS software main routing table (`inet.0`)
- Routing table stores:
  - Routing information learned from update messages
  - Routing information that passes sanity check (for instance, AS loop detection)
  - Local routing information selected by applying local policies to routes received in update messages

Copyright © 2006, Juniper Networks, Inc.

#### JUNOS Software BGP Routing Table

JUNOS software places BGP routes that pass the import policy in the local routing table: `inet.0`. Recall that the `inet.0` routing table contains entries from all unicast routing protocols.

#### BGP Routes Stored in `inet.0`

JUNOS software stores BGP routes in the `inet.0` table when they are learned from the update messages, checked for their validity, and conform with the import policy.

### A Basic BGP Configuration

```
routing-options {
    autonomous-system 64;
}
...
protocols {
    bgp {
        group external-peer1 {
            type external;
            peer-as 1;
            neighbor 10.0.3.6;
        }
        group internal-peers {
            type internal;
            local-address 192.168.24.1;
            neighbor 192.168.16.1;
            neighbor 192.168.6.1;
        }
    }
}
```

Copyright © 2006, Juniper Networks, Inc.

#### Basic BGP Configuration

- A BGP system must know which routers are its peers. You define the peer relationships explicitly by configuring the neighbor routers that are the peers of the local BGP system. After you establish peer relationships, the BGP peers exchange update messages to advertise network reachability information.

You arrange BGP routers into groups of peers. Different peer groups must have different group types, or AS numbers. On the slide, two groups are defined; *group external-peer1* and *group internal-peers*. The *type* field of the groups is different: *external* and *internal*, respectively.

The `protocols BGP` statement enables BGP on the router. The group, *external-peer1*, defines an EBGp group that recognizes neighbor 10.0.3.6 in AS 1 as a peer. The group *internal-peers* defines an IBGP group with the IP address on this local machine of 192.168.24.1, and explicitly defines two neighbors, 192.168.16.1 and 192.168.6.1.

With JUNOS software, you must configure at least one peer in each group. JUNOS software gives you flexibility to configure multiple peers belonging to different ASs. Should you choose to do this configuration, you must configure each peering neighbor with its own peer AS.

### Monitoring BGP Operation

---

Several commands display a wide variety of BGP information, either from the protocol itself or from BGP routes

```
user@host> show bgp ?
Possible completions:
  group                Show the BGP group database
  neighbor             Show the BGP neighbor database
  summary              Show an overview of the BGP
  information
```

Copyright © 2006, Juniper Networks, Inc.

### Monitoring BGP Operation

JUNOS software has a wide variety of BGP monitoring commands. This slide displays the following commands:

- **show bgp group**: Shows the BGP group database.
- **show bgp neighbor**: Shows the BGP neighbor database.
- **show bgp summary**: Displays the overall BGP information, including the state of BGP peer session establishment.

The next few slides elaborate on the details for each command.



### Displaying BGP Group Information

#### View information about a BGP group

```
user@host> show bgp group
Group Type: Internal   AS: 1           Local AS: 1
Name: ibgp             Index: 0         Flags: <Export Eval>
Export: [ ibgp-export ]
Holdtime: 0
Total peers: 1         Established: 1
192.168.16.1+3744
Trace options: detail open detail update detail keepalive
Trace file: /var/log/bgp size 131072 files 10
inet.0: 12/12/0

Group Type: External   Local AS: 1
Name: as2              Index: 1         Flags: <Export Eval>
Export: [ ebgp-export ]
Holdtime: 0
Total peers: 1         Established: 1
10.0.22.2+2460
Trace options: detail open detail update detail keepalive
Trace file: /var/log/bgp size 131072 files 10
inet.0: 3/10/0

Groups: 2 Peers: 2 External: 1 Internal: 1 Down peers: 0 Flaps: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 22 15 0 0 0 0 0
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying BGP Group Information

Use the `show bgp group` command to view BGP group information. Some of the output fields of this command are:

- **Group Type:** Displays the type of BGP group. It can be either Internal or External.
- **AS:** Displays the number of the remote AS. For IBGP, this number should be the same as the local AS number.
- **Local AS:** Displays the number of the local AS.
- **Export:** Displays the export policies configured for the BGP group with the export statement.
- **Total peers:** Displays the total number of peers in the group.
- **Established:** Displays the number of peers within the group that are in the established state.
- **ip addresses:** Displays the list of peers that are members of the group. The address is followed by the peer's port number.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Displaying BGP Group Information (contd.)

- Options: Displays configured BGP options. These options can be one or more of the following.
  - Local address: Displays the address configured with the `local-address` statement.
  - NLRI: Displays the configured MBGP state for the BGP group. It can be either multicast or unicast, or both if you have configured `nlri any`.
  - Hold time: Displays the hold time configured with the `hold-time` statement. The default hold time is 90 seconds.
  - Preference: Displays the preference value configured with the `preference` statement. The default preference value is 170.

### Displaying BGP Neighbors

```
user@host> show bgp neighbor
Peer: 192.168.16.1+3744 AS 1   Local: 192.168.20.1+179 AS 1
Type: Internal   State: Established   Flags: <ImportEval Sync>
Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Export: [ ibgp-export ]
Options: <Preference LocalAddress Refresh>
Local Address: 192.168.20.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 192.168.16.1   Local ID: 192.168.20.1   Active Holdtime: 90
Keepalive Interval: 30   Peer index: 0
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10001
RIB State: BGP restart is complete
Send state: in sync
Active prefixes:          12
Received prefixes:        12
Suppressed due to damping: 0
Advertised prefixes:      7
Last traffic (seconds): Received 29   Sent 29   Checked 29
Input messages: Total 16   Updates 9   Refreshes 0   Octets 640
Output messages: Total 16   Updates 8   Refreshes 0   Octets 652
Output Queue[0]: 0
Trace options: detail open detail update detail keepalive
Trace file: /var/log/bgp size 131072 files 10
Copyright © 2006, Juniper Networks, Inc.
```

### Displaying BGP Neighbors

Some of the output fields of the `show bgp neighbor` command are:

- Peer: Displays the address of each BGP peer. Each peer has one line of output.
- Type: Displays the type of peer (Internal or External).
- State: Displays the BGP state for this neighbor.
- Flags: Displays the internal peer-specific flags for this neighbor.
- Last State: Displays the BGP state of this neighbor prior to the current state.
- Last Event: Displays the last BGP state transition event.
- Last Error: Displays the last notification message sent to the neighbor.
- Options: Displays the configuration options in effect for this neighbor.
- Holdtime: Displays the configured hold time for this neighbor.
- Preference: Displays the configuration preference for routes learned from the neighbor.
- Peer ID: Displays the neighbor's router ID.
- Local ID: Displays the local system's router ID.
- Active Holdtime: Displays the hold-time value that was negotiated during the BGP open.
- Table inet.0 Bit: Displays the Internal bit used for the peer group.
- Send state: Displays whether all peers in the group have received all their updates (in sync or out of sync).

*Continued on next page.*

## Configuring Juniper Networks Routers

### Displaying BGP Neighbors (contd.)

- **Active Prefixes:** Displays the number of prefixes accepted as active from this neighbor.
- **Last traffic (seconds):** Displays how recently a BGP message was sent or received between the local system and this neighbor.
- **Output Queue:** Displays the number of BGP update messages pending for transmission to the neighbor.
- **Deleted routes:** Displays the prefixes queued for withdrawal through pending update messages.
- **Queued AS Path:** Displays an AS path queued for transmission in an update message.

### Displaying BGP Neighbor Information

Use the `show bgp summary` command to view basic information about all BGP neighbors

```
Groups: 2 Peers: 3 Down peers: 1
Table          Tot Paths  Act Paths  Suppressed  History Damp State  Pending
inet.0         0          0          0           0       0       0       0
inet.2         0          0          0           0       0       0       0
Peer           AS         InPkt     OutPkt     OutQ     Flaps Last Up/DwnState
192.168.16.1   65412     39        40         0        0      18:41 0/0/0
10.0.3.6       1         16        17         0        0      7:15 0/0/0
10.0.29.1      1         0         0          0        0      18:53 Active
```

Copyright © 2008, Juniper Networks, Inc.

### Displaying BGP Neighbor Information

The output fields of the `show bgp summary` command are:

- **Groups:** Displays the number of BGP groups.
- **Peers:** Displays the number of BGP peers.
- **Down peers:** Displays the number of unestablished BGP peers.
- **Peer:** Displays the address of each BGP peer. Each peer has one line of output.
- **AS:** Displays the peer's AS number.
- **InPkt:** Displays the number of packets received from the peer.
- **OutPkt:** Displays the number of packets sent to the peer.
- **OutQ:** Displays the count of the number of BGP packets queued to be transmitted to a particular neighbor. It usually is 0 because the queue is emptied quickly.
- **Last Up/Down:** Displays the last time since the neighbor transitioned to or from the established state.
- **State:** Displays either the BGP state or, if the neighbor is connected, the number of paths received from the neighbor, the number of these paths that have been accepted as active and are being used for forwarding, and the number of routes being damped.

### BGP Route Advertisements Commands

- **show route receive-protocol bgp address**
  - Displays routes received by a peer *before* policy is applied

```
user@host> show route receive-protocol bgp 11.1.1.1
inet.0: 6 destinations, 6 routes (5 active, 0 holddown, 1 hidden)
Prefix                Nexthop      MED          Lclpref AS path
10.0.0.0/8            192.168.1.1 100          I
172.16.0.0/12         172.19.1.1  100          I
```

- **show route advertising-protocol bgp address**
  - Displays routes advertised to a specific peer

```
user@host> show route advertising-protocol bgp 10.1.1.2
inet.0: 10 destinations, 10 routes (8 active, 0 holddown, 2 hidden)
Prefix                Nexthop      MED          Lclpref AS path
10.0.0.0/8            Self          100          I
172.16.0.0/12         Self          100          I
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying Received BGP Routes

The `show route receive-protocol bgp peer-address` command displays the BGP routing information that was received from the specified BGP neighbor. The information reflects the routes *before* import policy processing, with the exception of route-filter processing, which causes any filtered routes to be hidden. Add the `hidden` switch to display routes that are hidden due to import policy route-filters.

Because this command displays received routes *before* the majority of your input policies have been evaluated (route-filters being the exception), you should expect to see the BGP attributes as they were originally advertised to the local router. By way of example, if you are removing community strings as part of your import policy, you should expect to see that the received routes still have community attributes attached. To judge the effects of your import policy display the routes as they exist in the routing table!

### Displaying Routes Advertised to a Peer

The `show route advertising-protocol bgp address` command displays the routing information as it has been prepared for advertisement to the specified BGP neighbor. The information reflects the routes that the routing table exported into the routing protocol, along with any attributes that have been attached. Note that the local AS-number is not shown in the output of this command.

## Display BGP Routes

```
user@host> show route protocol bgp ?
Possible completions:
<[Enter]>          Execute this command
<destination>    Destination prefix and prefix length information
advertising-protocol Information transmitted by a particular routing protocol
all              All entries including hidden entries
aspath-regex     Entries learned via a specific AS path
best            Show longest match
brief           Brief view
+ community      A community to match, possibly including wildcards
damping         Entries that have been subjected to route damping
detail          Detailed view
exact           Show exact match
extensive       Extensive view
hidden          Hidden entries
inactive        Inactive entries
label-switched-path Entries associated with a particular LSP tunnel
next-hop        Entries pointing to a particular next hop
output          Entries sending packets out a particular interface
range           Show entire prefix range
receive-protocol Information learned from a particular routing protocol
source-gateway  Entries learned from a particular router
table           Entries in a particular routing table
terse           Terse view
|              Pipe through a command
Copyright © 2006, Juniper Networks, Inc.
```

### Display BGP Routes

BGP routes received from other neighbors are stored in the main routing table, along with routes from other sources. You use the `show route protocol bgp` command to display only BGP routes. You can combine this command with the `extensive` or `detail` switches to display the various attributes associated with a route. You can also filter BGP routes based on AS path or community regular expressions.

Add the `hidden` switch to see prefixes that are hidden. A BGP route may be hidden because of import policy filtering, because it has an unreachable next hop, or because of damping. This slide shows only a partial list of options.

### Viewing Details for BGP Routes

Use the `detail` switch to display BGP route attributes

```
lab@router> show route protocol bgp 3/8 detail

inet.0: 134203 destinations, 134203 routes (134203 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

3.0.0.0/8 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Source: 172.24.0.31
    Nexthop: 172.24.230.33 via fxpl.0, selected
    State: <Active Int Ext>
    Local AS: 65222 Peer AS: 65222
    Age: 1w2d 14:20:56 Metric2: 0
    Task: BGP_65222.172.24.0.31+1417
    Announcement bits (2): 0-KRT 2-BGP_Sync_Any
    AS path: 10458 14203 2914 1239 80 I
    Communities: 2914:420
    BGP next hop: 207.17.136.192
    Localpref: 100
    Router ID: 172.24.0.31
```

BGP Attributes

Copyright © 2006, Juniper Networks, Inc.

### Viewing BGP Route Details

Use the `detail` (or `extensive`) switch to display the BGP attributes associated with a given prefix. These attributes include the AS path, origin, local preference, and MED, as well as any communities that are attached to the route. The slide shows a “live” BGP route taken from a Juniper Networks route server. This can be determined by examining the AS Path attribute. In this case the last AS number in the AS sequence (10458), is assigned to Juniper Networks Inc. In some cases the AS Path attribute will contain brackets, braces, or other characters, which are used to indicate AS path sets, sequences, confederations, etc. The meaning of these characters when they are present in an AS Path attribute are defined here:

- [ ]: Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the router.
- {}: Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.
- (): Parentheses enclose a confederation.
- ([ ]): Enclose a confederation set.

In some cases a BGP route may be hidden. This can occur because of AS path sanity checks detecting an AS loop, or because of an unreachable BGP next hop. Try adding the `hidden` switch when issuing a `show route` command to display any such hidden routes. Use the `show route resolution unresolved` command to determine the reason why a BGP prefix is hidden.



### Regular Expressions

---

- Regular expressions are a powerful *pattern matching* engine
- Combination of text and special operators that make up a *regular expression*
- Regular expressions allow for things to be found in context, not as isolated instances
- Used for AS-path and community matching in operational-mode *show* commands and routing policy
  - Detailed coverage of regex is beyond the scope of this class

Copyright © 2006, Juniper Networks, Inc.

#### Why Regex?

- ☞ A regular expression (often seen as regex or RE) is a powerful *pattern-matching* tool that you can apply in a routing policy to act on AS-path strings. This pattern matching engine can find specific strings of text or textual patterns.

#### Format of Regex

When used in a routing policy, regular expressions work not only on fixed strings like wildcard operators, such as an asterisk (\*), but regular expressions act on variable patterns of text, which is done through the combination of basic text patterns and special operators. This combination of basic text patterns and special operators are what make up the regular expression itself.

#### How Does It Help?

Regular expressions allow information in a string to be found within a specific context, not just in isolated instances. The use of regular expressions in conjunction with the BGP AS-path attribute can be used to match routes within a policy.

#### Regex with AS Paths and Communities

You can use regular expressions to match with AS paths and communities. These two concepts are examined on subsequent slides. Note that while regular expressions are defined as part of the Portable Operating System Interface (POSIX) standard, JUNOS software regular expressions, as used for AS-path matching, are not considered POSIX compliant. Community-based regular expressions adhere to the POSIX standard, however.

A detailed treatment of regex-based AS-path and community matching is beyond the scope of this course.

### Tracing BGP

- A typical BGP tracing configuration:

```
[edit protocols bgp]
user@host# show
traceoptions {
    file bgp-trace;
    flag open detail;
    flag update detail;
}
```

- Monitor the resulting *bgp-trace* log file using **monitor start log-file-name** or the **show log log-file-name** commands

Copyright © 2006, Juniper Networks, Inc.

### BGP Tracing

To perform debugging functions on the BGP routing process, use the JUNOS software `traceoptions` feature. The trace output (debug information) is directed to the named log file, which is located in the `/var/log` directory on the router's hard drive. You can view the log file using the `monitor start` or `show log` operational mode commands. In addition to specifying the trace file, you also must tell the router what information you want to trace. You accomplish this by specifying one or more `flag` keywords.

While you can only direct tracing to a single file, you can trace many options by using the `flag` keyword multiple times. In addition, you can add granularity by using the `detail`, `receive`, and `send` `flag` modifiers.

Available tracing flags for BGP include:

<code>all</code>	Trace everything
<code>aspath</code>	
<code>damping</code>	
<code>general</code>	Trace general events
<code>keepalive</code>	Trace BGP keepalive packets
<code>normal</code>	Trace normal events
<code>open</code>	Trace BGP open packets
<code>packets</code>	Trace all BGP protocol packets
<code>policy</code>	Trace policy processing
<code>refresh</code>	Trace BGP refresh packets
<code>route</code>	Trace routing information
<code>state</code>	Trace state transitions
<code>task</code>	Trace routing protocol task processing
<code>timer</code>	Trace routing protocol timer processing
<code>update</code>	Trace BGP update packets

### Review Questions

---

1. On what type of network would you implement BGP?
2. What are the default BGP route advertising rules?
3. How would a typical ISP design a network to support BGP? Draw a sample network.
4. How would you display the BGP attributes for a given route?
5. How would you determine the status of a BGP peer?
6. What is the purpose of the local address option in JUNOS software?

Copyright © 2008, Juniper Networks, Inc.

#### This Module Discussed:

- EGP to IGP comparison;
- The use of common BGP attributes;
- The JUNOS software BGP route selection procedure;
- The rationale behind loopback peering for IBGP;
- Default BGP route advertisement rules;
- Configuring BGP in JUNOS software; and
- Using operational-mode commands to monitor and troubleshoot BGP.

## **Lab 7: BGP**

---

### **Lab Objectives:**

**Configure BGP on a Juniper Networks M-series router and monitor the router's operation using CLI commands**

Copyright © 2006, Juniper Networks, Inc.



**Configuring Juniper Networks Routers**

***Appendix A: Overview of J-series,  
M-series and T-series Hardware***

Copyright © 2006, Juniper Networks, Inc.

CJNR-M-7.a.7.6.1

### Module Objectives

---

- **After successfully completing this module, you will be able to:**
  - Describe the current M-series and T-series platform offerings
  - Describe general installation procedures for M-series and T-series routers
  - Explain the architecture of Juniper Networks M-series and T-series platforms
  - Describe the function of the RE, FPCs, PICs, and System and Control Boards
  - Operate the Craft Interface
  - Describe packet flow through M-series and T-series platforms
  - Describe interface naming conventions and the purpose of logical units

Copyright © 2006, Juniper Networks, Inc.

#### **This Module Discusses:**

- The Juniper Networks, Inc. M-series, T-series, and J-series products and their typical applications;
- General platform architecture;
- The function of major router components;
- Operation of the Craft Interface;
- Packet flow through T-series and M-series platforms and;
- Interface naming conventions and the role of logical units.

## Agenda: Product Line Overview

---

### → M-series and T-series Overview

- Installation and Handling Guidelines
- Platform Architecture
- M-series Packet Flow and ASIC Functionality
- T-series Packet Flow and ASIC Functionality
- Interface Overview
- Additional Products

Copyright © 2006, Juniper Networks, Inc.

### Module Agenda

This slide outlines the module agenda and highlights the area that is covered next.

## Configuring Juniper Networks Routers

### M-series and T-series Product Line (1 of 2)

- Family of router platforms that deliver:
  - Industry-leading core and dedicated-access platforms
    - Solutions that scale in multiple dimensions with market-leading port density
  - Flexible and manageable traffic control
  - High reliability features

Copyright © 2006, Juniper Networks, Inc.

### The M-series and T-series Product Line: Part 1

The key application driving the new IP infrastructure is the Internet, which continues to grow despite recent market downturns and various *dot-bomb* explosions. The Internet has moved from a convenience to a mission-critical platform for conducting and succeeding in business. As reliance on the Internet grows, so do customer expectations for value-added services that require increased bandwidth. There is little doubt that these rising demands for new services and for the Internet will set the standard for the emerging IP infrastructure.

Juniper Networks delivers a family of router platforms that provide industry-leading performance with solutions that offer high availability, scalability in multiple dimensions, market-leading port density, and flexible control over traffic to achieve optimal bandwidth efficiency. This Juniper Networks M-series and T-series product line currently consists of the M5/M10, M7i/M10i, M20, M40, M40e, M160, and T320 routers, and the T640 routing node. All these platforms run a *common* JUNOS Internet software image, and all are based on purpose-built ASICs for packet forwarding, including the Internet Processor II ASIC, which has the unique ability to provide enhanced services on *all* interfaces without compromising performance.

Platform highlights include:

- The M40 Internet router provides more than 40-Gbps aggregate throughput and supports up to 32 PICs per chassis.
- The M20 Internet router also provides aggregate throughput of 20+ Gbps, supporting up to 16 PICs per chassis.
- The M10 Internet router supports up to eight PICs per chassis with an aggregate throughput of 10+ Gbps.

*Continued on next page.*

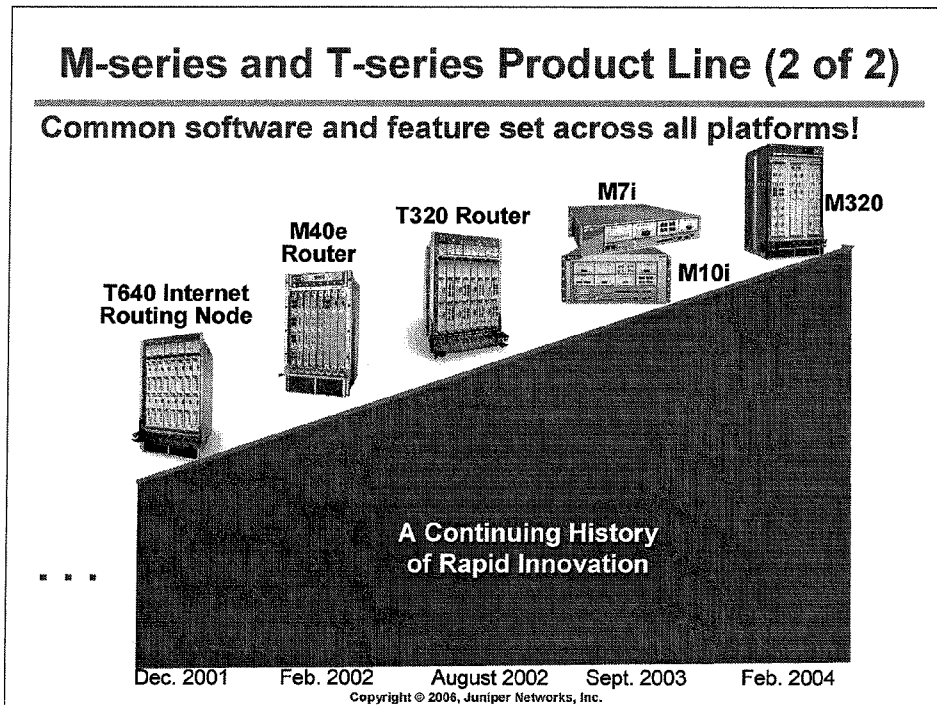


## Configuring Juniper Networks Routers

### The M-series and T-series Product Line: Part 1 (continued)

- The M5 Internet router supports up to four PICs per chassis with an aggregate throughput of 5+ Gbps.
- The T640 Internet routing node is Juniper Networks largest single-chassis router, offering a throughput of 320+ Gbps (640 Gbps aggregate) and support for up to 128 OC-48/STM-16, 32 OC-192/STM-64, or 128 Gigabit Ethernet ports for the router.

## Configuring Juniper Networks Routers



### The M-series and T-series Product Line: Part 2

The Juniper Networks product portfolio continues to grow with the release of ever-faster and higher-density routing platforms. A significant aspect of the M-series and T-series product line is that all platforms run JUNOS software with support for all features. This is significant when you consider that one router vendor currently has 5,800 images for the 2600/3600/3700 family of products alone! Even the small enterprise class J-series routers run the same JUNOS software. J-series software images are packaged separately to save storage space and to include support for software-based guaranteed performance forwarding and services.

Recent platform offerings include:

- The T320 router is based on T640 ASIC technology and features a 1/3 rack footprint with an impressive 160 Gbps of throughput (320 Gbps aggregate).
- The M160 Internet router offers an aggregate throughput of 160+ Gbps. It supports up to 32 OC-48c/STM-16 PICs or up to eight OC-192c/STM-64 PICs per chassis.
- The M40e Internet router delivers M40 throughput while meeting the needs of the edge aggregation market through its redundancy features and support of high-density PICs, such as the 48-port Fast Ethernet PIC.

*Continued on next page.*

## Configuring Juniper Networks Routers

### The M-series and T-series Product Line: Part 2 (contd.)

- The M7i and M10i platforms offer a migration path/alternative for customers using the previously released M5 and M10 platforms. Both platforms feature support for existing M5/M10 PICs. The M7i supports a Fixed Interface Card (FIC) with Ethernet connectivity for added value and flexibility as well as an integrated Tunnel Services PIC; an integral Adaptive Services PIC, which is in addition to the Tunnel PIC, is an optional CFEB upgrade. The M10i, on the other hand, offers RE and CFEB redundancy but does not support the FIC or integral services PICs. At only two rack units, up to 24 M7is can be mounted in a single 19-inch rack!
- The M320 boasts a highly reliable architecture with all common hardware fully redundant and hot swappable as a foundation for reliability features like in-service software upgrades and hitless switchover. The M320 platform includes unprecedented control plane scalability with its 1.6-GHz Routing Engine equipped with two Gigabytes of DRAM. The platform leverages the industry's most advanced programmable ASICs to support 385 million packets per second of packet processing and 320 Gigabits per second of aggregate throughput. The M320 is a half-rack chassis with a shallow profile to fit into smaller cabinets. Some of the key characteristics of the M320 platform include:
  - 1.6G GHz Routing Engine with Gigabit uplink to the PFE;
  - 385 Mpps forwarding;
  - 320 Gbps of throughput;
  - 32 Type 1 or Type 2 PICs, 16 Type 3 PICs
  - 16 x 10GE/OC-192/STM-64;
  - 64 x OC-48/STM-16;
  - 160 x 1 Gigabit Ethernet ; and
  - 32 x CHOC12/STM-4 (to DS0).
- See the end of this appendix for a discussion of the following products:
  - J-series small office routers;
  - M120 multiservice edge router; and
  - TX Matrix for T640 multichassis operation.

## Configuring Juniper Networks Routers

Feature		Platform								
		M5 Router	M7i Router	M10 Router	M10i Router	M20 Router	M40 Router	M40e Router	M160 Router	M320 Router
Chassis Throughput (Aggregate)	6.4 Gbps (40 Mpps)	8.4 Gbps (16 Mpps)	12.8 Gbps (40 Mpps)	12.8 Gbps (16 Mpps)	25.6 Gbps (40 Mpps)	51.2 Gbps (40 Mpps)	51.2 Gbps (40 Mpps)	204 Gbps (160 Mpps)	320 Gbps (386 Mpps)	
Slot Throughput (Aggregate)	6.4 Gbps	6.4 Gbps	6.4 Gbps	6.4 Gbps	6.4 Gbps	6.4 Gbps	6.4 Gbps	25.6 Gbps	40 Gbps (FPC3)	
Slots/PICs	1/4	1/6 (2 built-in PICs)	2/8	2/8	4/16	8/32	8/32	8/32	8/32 (FPC1/2) 8/16 (FPC3)	
Power	AC/DC	AC/DC	AC/DC	AC/DC	AC/DC	AC/DC	AC/DC	AC/DC	DC Only AC/DC	
Units per Rack	15 per rack	24 per rack	15 per rack	9 per rack	5 per rack	2 per rack	2 per rack	2 per rack	2 per rack	
RE/Control Redundancy	No	No	No	Yes	Yes	No	Yes	Yes	Yes	
Weight (Max)	81 Lbs/27.7 Kg	36.5 Lbs/16.6Kg	65 Lbs/29.5 Kg	66 Lbs/29.5 Kg	160 Lbs/68 Kg	280 Lbs/127 Kg	370.5 Lbs/168 Kg	370.5 Lbs/168 Kg	439Lbs/ 199.6 Kg	

\* Numbers quoted are two times the unidirectional (simplex) capacity for each FPC or chassis.

Copyright © 2006, Juniper Networks, Inc.

### Product Comparison: M-Series

This slide provides a matrix of key characteristics associated with the M-series product line. The model number for each router is based on the aggregate throughput capabilities of that router. For example, an M160 router can process 160 million packets per second. Also, all M-series routers support both AC/DC power (but not both simultaneously), except the M160, which requires DC power.

### Product Comparison: T-series

Feature	Platform	
	T640 Internet Routing Node	T320 Router
Chassis Throughput (Aggregate)	640+ Gbps (640 Mpps)	320+ Gbps (320 Mpps)
Slot Throughput (Aggregate)	FPC3 = 80+ Gbps FPC 2 and 3	FPC3 = 40+ Gbps FPC 1, 2, and 3
Slots/PICS	8/32	8/16
Power	DC only	DC only
Units Per Rack	2 per rack	3 per rack
RE/Control Redundancy	Yes	Yes
Weight (Max)	556Lbs/266.28Kg	359.9 Lbs/187.78Kg

\* Numbers quoted are two times the unidirectional (simplex) capacity for each FPC or chassis.

Copyright © 2006, Juniper Networks, Inc.

### Product Comparison: T-series

Juniper Networks is currently shipping two T-series platforms. The T640 Internet routing node is the flagship of the product line with its clustering capability and single-chassis forwarding rate of 640 million packets per second. The T320 platform is a smaller platform that supports 16 PICs (two per FPC).

To allow the reuse of existing M-series PICs, the T320 platform supports three FPC types; a T320 FPC1 supports original M-series PICs, while the FPC2 supports PICs native to the M160 platform. The FPC3 is the native FPC for the T-series and is designed to support high-speed T-series PICs. You can mix and match FPC types within a single chassis. The T640 platform supports only FPC types 2 and 3.

T-series platforms require DC power.

## Agenda: Product Line Overview

---

- M-series and T-series Overview
- Installation and Handling Guidelines
- Platform Architecture
- M-series Packet Flow and ASIC Functionality
- T-series Packet Flow and ASIC Functionality
- Interface Overview
- Additional Products

Copyright © 2006, Juniper Networks, Inc.

### Installation and Handling Guidelines

We discuss the following topics in subsequent pages:

- *General installation guidelines*: Things to keep in mind to ensure that all goes well when installing M-series and T-series platforms.
- *Power up and power down*: Steps associated with power up and safe power down of M-series and T-series platforms.
- *General hardware handling*: General hardware handling guidelines and an example of typical field-replaceable units (FRUs).

### Chassis Installation Guidelines

- **Follow site-preparation guidelines**
  - Space, environment, power, grounding, etc.
- **Lifting requires three or more people**
  - A mechanical lift is recommended
  - **Selected M-series and T-series maximum product weights**
    - T640 Internet routing node weighs 565 pounds (256.28 kg)
    - T320 weighs 369.9 pounds (167.78 kg)
    - M320 router weighs 439 pounds (199.6 kg)
    - M160 router weighs 370.5 pounds (168 kg)
    - M20 router weighs 150 pounds (68 kg)
    - M10i router weighs 79 pounds (36 Kg)
- **Remove heavier components before lifting**
  - Power supplies, FPCs, fan trays, system boards
- **Lift into rack**
  - Do not lift M40 router by Routing Engine handles
  - Replace components

Copyright © 2006, Juniper Networks, Inc.

#### Follow Site Preparation Guidelines

To ensure a smooth installation, each M-series and T-series platform has an associated hardware guide that provides installation instructions. The hardware guide provides critical details such as maximum component and system power draw, system weights, and the clearances needed for serviceability and proper cooling. An installation checklist form is also provided to help ensure that you will be ready to press your new M-series or T-series platform into service the day that it arrives.

For M-series and T-series hardware manuals, point your browser to:  
<http://www.juniper.net/techpubs/hardware/index.html>.

#### It's Not Heavy, It's My Router!

A Juniper Networks router with a maximum configuration can weigh anywhere from 61 lbs (27.7 kg) for an M5 Internet router all the way up to 565 lbs (256.28 kg) for a T640 Internet routing node.

#### Remove the Heavier Components before Lifting

Removing the power supplies, flexible PIC concentrators, fan trays, routing engine(s), and control board(s) can make a given platform considerably lighter. For example, a T640 chassis and midplane alone weighs only 205 pounds (93 kg). When fully loaded, the same chassis tips the scales at 565 pounds (256 kg).

#### Carefully Lift Router into Rack

Use extreme care when lifting M-series and T-series platforms! We recommend a mechanical lift. **DO NOT LIFT AN M40 ROUTER BY THE ROUTING ENGINE HANDLES!** Although they appear to be, at first glance, for the purpose of lifting the router, their purpose is only for the insertion and removal of the Routing Engine.

Once properly secured in the rack, you should replace in the chassis the components that were removed to reduce the weight of the chassis.

### Power up and Power down

- **Power up**
  - Connect all cables
  - Turn on one power supply
  - Turn on second power supply
    - It can take up to 60 seconds for accurate power supply/PEM status indications
- **Power down**
  - Shut down JUNOS Internet software
    - CLI `request system halt` command
  - Turn off power supplies/remove power

Copyright © 2006, Juniper Networks, Inc.

### Powering the M-series and T-series Platforms

Each router can be equipped with two redundant, load-sharing power supplies of the same type; in most cases these can be either AC or DC for M-series platforms but DC only for T-series platforms. Be sure to connect each power source properly. For example, each power supply requires a dedicated power source. For sites with an AC power source, each power supply has one power cord, which is plugged into a grounded 100–240 VAC power receptacle. For sites with a DC power source, power is normally carried around the site through a main conduit to frame-mounted DC power distribution panels, one of which might be located at the top of the rack where the router is to be installed. A pair of cables (–48 V and RTN) connects each DC supply to the power distribution panel. Grounding studs are provided at the rear enclosure. After connecting all cables, turn one power supply on first and then the second supply to avoid a large power spike.

Although the specifics of each power supply/Power Entry Module (PEM) vary by platform, you should note that after a power supply is powered on, it can take up to 60 seconds for status indicators—such as LEDs on the power supply and `show chassis` commands—to indicate that the power supply is functioning normally. You should ignore error indicators that appear during the first 60 seconds.

### Powering down the M-series and T-series Platforms

Because JUNOS software is a multitasking OS, we highly recommend that you perform a *graceful* power down when planning to remove power (as opposed to an abrupt *hard* power down, in which the power is simply removed from the system). To gracefully shutdown JUNOS software, issue a `request system halt` command. Once the system is properly halted, you can safely remove power. Note that on some platforms there is a button on the chassis that gracefully shuts down the RE when depressed for 3–5 seconds.



### Reboots and Shutdowns

- Always gracefully shutdown JUNOS software before removing power!

- Rebooting the system:

```
user@host> request system reboot ?
Possible completions:
<[Enter]>      Execute this command
at             Time at which to perform the operation
in            Number of minutes to delay before operation
media         Boot media for next boot
message       Message to display to all users
|             Pipe through a command
```

- Shutting down the system:

```
user@host> request system halt ?
Possible completions:
<[Enter]>      Execute this command
at             Time at which to perform the operation
both-routing-engines  Halt both Routing Engines
in            Number of minutes to delay before operation
media         Boot media for next boot
message       Message to display to all users
|             Pipe through a command
```

Copyright © 2006, Juniper Networks, Inc.

### Graceful Shutdowns and Rebooting

JUNOS software runs on a multitasking multi-user operating system based on Free-BSD. As with any UNIX system, you should always gracefully shut the system down before removing power. Failing to shutdown in a graceful manner can result in file system corruption that, in the best of cases, simply prolongs the next boot, and in the worst of cases, might actually prevent a successful boot.

The `request system reboot` command causes the router to reboot. Reboot requests are recorded to the system log files, which you can view with the `show log` command. The reboot command takes a variety of arguments that you can use to schedule the reboot, generate a system message, or specify the media that should be used to boot from. Media options include `compact-flash` and `disk`. The router always attempts to boot from the removable media when the media is installed and a power cycle (cold boot) is performed.

The `request system halt` command gracefully stops the router software and prepares the router to be shut down. Note that you must either cycle power or hit the Enter key on the terminal attached to the router's console port to effect the reboot of a router that has executed a shutdown command.

Both the reboot and shutdown command require user confirmation of the action:



```
user@host> request system reboot
Reboot the system ? [yes,no] (no) yes
```

```
*** FINAL System shutdown message from root@host ***
System going down IMMEDIATEL
```

```
shutdown: [pid 15049]
Shutdown NOW!
```

## General Handling: Mechanical

- FPCs are heavy and can be damaged with improper handling
  - Fully loaded FPC3 weighs up to 32 lbs/11.3 kg
    - Be prepared for the weight of the FPC when removing from the midplane
    - Always handle appropriately




Examples of Proper FPC Handling

Copyright © 2006, Juniper Networks, Inc.

### FPC Handling

T-series FPCs are heavy and can be damaged with improper handling. The figures below illustrate some common forms of FPC abuse that can damage your expensive hardware! The handling suggestions shown here hold true for M-series FPCs also.




Common FPC Handling Mistakes


### General Handling: Electrostatic

- Always observe ESD precautions
  - Use a grounded wrist/ankle strap connected to a chassis ESD lug

ESD lugs are normally located on the front and back of each chassis



- Always place field-replaceable units in ESD-approved packaging



Copyright © 2006, Juniper Networks, Inc.

### General Handling: Electro Static Damage

The FRUs associated with both M-series and T-series platforms contain ESD-sensitive parts; some components can be impaired by voltages as low as 30 V. You can easily generate potentially damaging static voltages whenever you handle plastic or foam packing material or if you move components across plastic or carpets. To avoid unnecessary downtime and added maintenance costs, you must always observe ESD handling precautions when handling FRUs.

General ESD guidelines:

- Always use an ESD wrist strap or ankle strap, and make sure that it is in direct contact with your skin. For safety, periodically check the resistance value of the ESD strap. The measurement should be in the range of 1 to 10 Mohms.
- When handling any component that is removed from the chassis, make sure the equipment end of your ESD strap is attached to one of the electrostatic discharge points on the chassis.
- Avoid contact between the component and your clothing. ESD voltages emitted from clothing can still damage components.
- When removing or installing a component, always place it component-side up on an antistatic surface, in an antistatic card rack, or in an electrostatic bag. If you are returning a component, place it in an electrostatic bag before packing it.

### FRU Types

---

- **FRUs are distinguished by whether they can be removed/installed without causing system disruption**
  - **Hot-insertable and hot-removable FRUs can be swapped without disrupting global routing functions**
    - Procedures for online/offline of hot-swappable FRUs must be followed to ensure that global routing is not disrupted
  - **Hot-pluggable FRUs interrupt global routing functions when the component is removed or installed**
  - **Some FRU types require that power be removed from the chassis**

Hot-Swappable	Hot-Pluggable
Air filters Flexible PIC Concentrators Front and rear fan trays Physical Interface Cards Power supplies SONET Clock Generators Switch Interface Boards	Control Boards Connector Interface Panel Routing Engine

**T640 Routing Node FRU Classification**


Copyright © 2006, Juniper Networks, Inc.

### FRU Types

Field-replaceable units (FRUs) are platform components that can be replaced at the customer site. Replacing FRUs requires minimal routing node downtime. Generally speaking, there are three types of FRUs:

- *Hot-insertable and hot-removable FRUs:* You can remove and replace these components without powering down the routing node or disrupting routing functions. This type of FRU is often referred to as being *hot-swappable*.
- *Hot-pluggable FRUs:* You can remove and replace these components without powering down the routing node, but the routing functions of the system are interrupted when the component is removed.
- *FRUs that require power down:* In rare cases an FRU requires that power be removed from the chassis before removal or insertion. An example of this type of FRU is the M5/M10 platforms' CFEB and RE. In the case of these platforms, the lack of redundant PFE and RE capabilities made the engineering of hot-swappability for these components a nonissue.

## Configuring Juniper Networks Routers



### Typical Platform FRUs

All Juniper Networks M-series and T-series platforms are comprised of a chassis and one or more FRUs. The related slide calls out the primary FRUs associated with the T640 routing node.

### **Agenda: Product Line Overview**

---

- M-series and T-series Overview
- Installation and Handling Guidelines
- **Platform Architecture and Hardware Overview**
- M-series Packet Flow and ASIC Functionality
- T-series Packet Flow and ASIC Functionality
- Interface Overview
- Additional Products

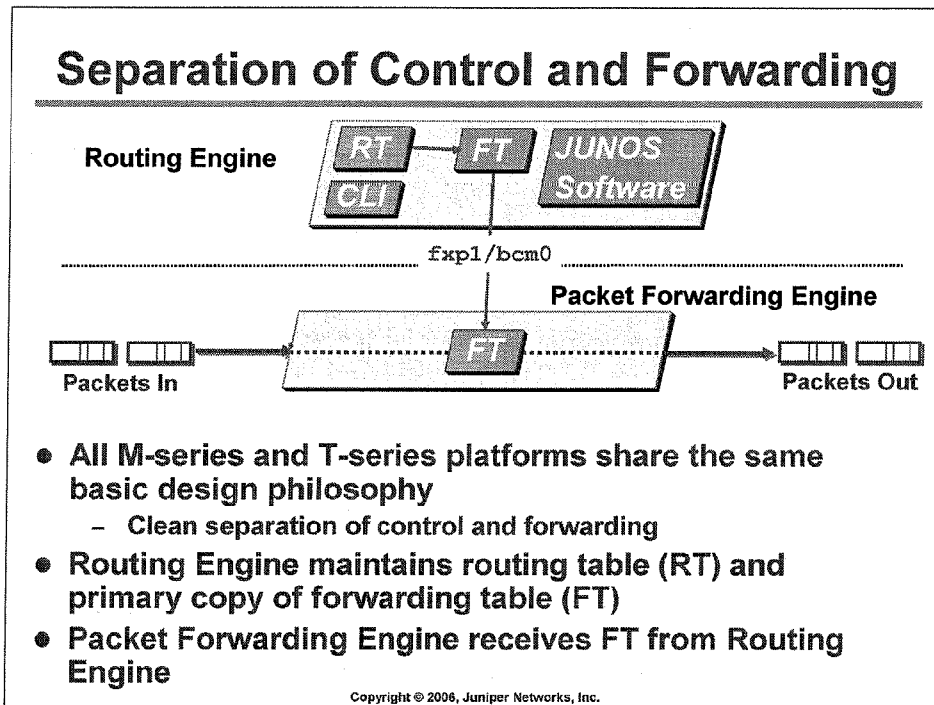
Copyright © 2006, Juniper Networks, Inc.

#### **Platform Architecture and Hardware Overview**

We discuss the following topics in subsequent pages:

- *General platform architecture:* The architecture of Juniper Networks M-series and T-series platforms is designed to separate the equally complex problems of control and packet forwarding.
- *Hardware overview:* The unique roles of the Routing Engine (RE) and the Packet Forwarding Engine (PFE) are fully explored, as are the components that constitute the Packet Forwarding Engine on M-series and T-series platforms.
- *Craft Interface:* The function and operation of the Craft Interface, including the LCD panel when so equipped, is described.

## Configuring Juniper Networks Routers



### Architectural Philosophy

Architecturally, all Juniper Networks M-series and T-series platforms share a common design that separates the router's control and forwarding planes. To this end, all M-series and T-series platforms consist of two major components:

- *The Routing Engine (RE):* The RE is the brains of the platform; it is responsible for performing routing updates and system management. The RE runs various protocol and management software processes that live inside a protected memory environment. The RE is a general-purpose computer platform based on an Intel microprocessor. The RE is connected to the PFE through an internal 100-Mbps connection identified as `fxp1` on most M-series and T-series platforms. On the M320 platform, a Gigabit Ethernet link referred to as `bcm0` ties the RE to a Fast Ethernet switch, which in turn has dedicated 100-Mbps links to each FPC.
- *The Packet Forwarding Engine (PFE):* The PFE is responsible for forwarding transit packets through the router using an ASIC-based switching path. The PFE is a high-performance switch capable of forwarding up to 160 Mpps, in the case of the M160 platform, for all packet sizes. By adding a cross-bar switching fabric between multiple PFEs, T-series platforms can achieve up to 640 Mpps of forwarding capacity in a single chassis!

Because this architecture separates control operations—such as routing updates and system management—from packet forwarding, the router can deliver superior performance and highly reliable Internet operation.

The simple fact that you can enable enhanced services without significantly impacting forwarding rates or system stability is a testament to the validity of M-series and T-series architecture.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Routing and Forwarding Table Interaction

The JUNOS software routing protocol process implements the various routing protocols that can be run on the router. The routing protocol process starts all configured routing protocols and handles all routing messages. The routing daemon (`rpd`) maintains one or more routing tables, which consolidates the routing information learned from all routing protocols into common tables. From this routing information, the routing protocol process determines the active routes to network destinations and installs these routes into the Routing Engine's forwarding table (FT).

### Routing Engine and Packet Forwarding Engine Synchronization

The Packet Forwarding Engine (PFE) receives the forwarding table from the Routing Engine. In the majority of cases, the Packet Forwarding Engine's FT and the Routing Engine's FT are kept synchronized over the 100-Mbps `fxp1` Ethernet link, which interconnects the two entities. This synchronization ensures that a change in topology produces identical FTs in the Routing Engine and Packet Forwarding Engine. In the case of the M320 platform, a 100-Mbps Ethernet switch provides a dedicated link to each FPC. These 100-Mbps links are then presented to the M320 RE as a single Gigabit Ethernet uplink called `bcm0`.

FT updates are a high priority for the JUNOS software kernel and are performed incrementally. The router's FT is large enough to hold over 800,000 entries. Thus, entries are never aged out of the FT to make room for new entries or because they have not been recently used. This behavior ensures that packets are switched in hardware not software, as can happen on the platforms of other vendors when FT entries are aged out of line cards, forcing process-based switching of packets destined to those addresses.



### Routing Engine Overview

---

- JUNOS software resides in flash memory
  - Backup copy available on hard drive
- Provides forwarding table to the Packet Forwarding Engine
  - Not directly involved with packet forwarding
  - Runs various routing protocols
- Implements CLI
- Manages Packet Forwarding Engine

Copyright © 2006, Juniper Networks, Inc.

#### JUNOS Software

The primary copy of JUNOS software resides on the flash memory of the router. A backup copy is available on the hard drive when you issue a `request system snapshot` command.

#### Routing Engine Intelligence

The Routing Engine handles all the routing protocol processes as well as other software processes that control the router's interfaces, a few of the chassis components, system management, and user access to the router. These routing and software processes run on top of a kernel that interacts with the Packet Forwarding Engine. All routing protocol packets from the network are directed to the Routing Engine.

#### Command-Line Interface

The Routing Engine provides the command-line interface (CLI). The CLI runs on top of the kernel; it is controlled by the management daemon (`mgd`). We examine the CLI and its features in more detail in a subsequent module.

#### Packet Forwarding Engine Management

The Routing Engine controls the Packet Forwarding Engine by providing an accurate and up-to-date forwarding table and by downloading microcode and managing software daemons that live in the Packet Forwarding Engine's microcode. The RE receives hardware and environmental status messages from the PFE and acts upon them as appropriate.

## Configuring Juniper Networks Routers

### Current Routing Engine Characteristics

		RE Model			
		RE-333	RE-400	RE-600	RE-1600
Feature	Processor/clock	Pentium III/333 MHz	Celeron/400 MHz	Pentium III/600 MHz	Pentium IV/1600 MHz
	Memory	768 MB	256, 512, 768 MB	512, 2 GB	2 GB
	Solid state flash storage	80 MB	256 MB (Optional)	128 MB/256 MB	256 MB
	Hard disk size	6.4+ GB	20 GB	30+ GB	30+ GB
	External media	PCMCIA flash card/LS-120*	PCMCIA flash card (Optional)	PCMCIA flash card/LS-120*	PCMCIA flash card
	Supported Platforms	Originally shipped on: M5/M10/20/40/40e, and M160	M7i/M10i Only	All M-series and T-series <i>except</i> M7i/M10i/M320	M320 T320/T640

\* The M40 router continues to use the original LS-120 drive for external storage regardless of RE model.

Copyright © 2006, Juniper Networks, Inc.

### Routing Engine Specifications

Juniper Networks has periodically released updated Routing Engines (REs) to enhance the performance of a given routing platform. For example, you can replace the original 233-MHz RE that shipped with the M40 (RE-M40) router by a higher performance RE-333 or RE-600 when you want increased memory and processor speed. The RE-333 and RE-600 are also known as RE2 and RE3 respectively when viewing the output of a `show chassis hardware` command.

At the time of this writing, the latest RE enhancement is reflected in the RE-1600, which features a 1.6-GHz clock rate and 2 GB of RAM (expandable to 4). The RE-1600 is required on the M320 platform and also supported on the T320 and T640 platforms as of Release 6.2. The RE-1600 uses a Broadcom chipset for Ethernet connectivity to the PFE. A Gigabit uplink is supported on the `bcm0` interface when installed in a M320 platform, while 100-Mbps operation is supported in T-series platforms.

The RE-600 is currently standard issue on all M-series and T-series platforms except the M7i, M10i, and M320 routers; RE-600 upgrades are available for systems that originally shipped with the RE-333 or RE-M40.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Routing Engine Specifications (contd.)

The M7i and M10i platforms make use of a specially designed RE model that is not based on a compact PCI platform. Juniper Networks chose the particulars of the RE-400 (RE5) to reduce platform costs (Celeron processor, optional compact flash and PCMCIA card support) while matching the processing requirements of the relatively small M7i and M10i routers. The RE-400 is only supported in the M7i and M10i platforms. Originally the RE-400 shipped with a 128-MB flash memory. New RE-400 units are now shipping with 256 MB of flash memory to better accommodate the new partition scheme used in JUNOS software Release 6.x. Juniper Networks will upgrade RE-400 units returned for maintenance with 256-MB of flash storage when needed.

With the exception of the M7i, M10i, M120, and M320 platforms, all remaining M-series and T-series platforms support the RE-600. This can have a significant impact on operational costs because it can simplify sparing logistics. Note that the M40 router continues to use the original LS-120-based external media, regardless of which RE version is installed.

The size of the JUNOS software can be large enough to inhibit an upgrade using the jbundle for systems with flash memory in the range of 76–128 MB. In these cases you must either increase the size of the flash memory or upgrade using either a jinstall package or removable media. Flash memory upgrades are available for the RE-600 model only.

Note that we recommend enabling configuration file compression to help reduce the amount of compact flash memory consumed by large router configurations. Current JUNOS Software enables configuration compression by default. If you are using older JUNOS software, you can enable configuration file compression with a `set system compress-configuration-files` statement issued at the `[edit]` hierarchy.

Self-Monitoring Analysis and Reporting Technology System (SMART) is automatically enabled on supported Routing Engines. SMART provides early warning of potential problems with the hard disk through entries that are written to the system's log file. You can disable SMART drive monitoring with a `set system processes disk-monitoring disable` command, which you issue at the `[edit]` configuration hierarchy when you do not want this feature.

### Packet Forwarding Engine Overview

---

- **Custom ASICs implement forwarding path**
  - No *process switching*
  - Value-added services and features implemented in hardware
    - Multicast
    - CoS/queuing
    - Firewall filtering
    - Accounting
- **Divide-and-conquer architecture**
  - Each ASIC provides a piece of the forwarding puzzle

Copyright © 2006, Juniper Networks, Inc.

#### Custom ASICs

ASICs enable the router to achieve data forwarding rates that match current fiber-optic capacity. Such high forwarding rates are achieved by distributing packet processing tasks across highly integrated ASICs. As a result, Juniper Networks M-series and T-series platforms do not require a general purpose processor for packet forwarding; this makes *process switching* (the software-based handling of packet forwarding) an alien concept for Juniper Networks routers. The custom ASICs provide enhanced services and features, such as multicast, CoS/queuing, and firewall filtering in hardware so that you can enable services on production routers without concern of significant performance hits.

#### Divide-and-Conquer Architecture

Each ASIC provides a piece of the forwarding puzzle, allowing a single ASIC to perform its specific task optimally. We examine the the role that each M-series and T-series ASIC plays in packet forwarding on following slides.

### PFE Components: M-series

- Physical Interface Cards (PICs)
- Flexible PIC Concentrators (FPCs)
- The system midplane
- For M5/M10, M7i/M10i, M20, and M40
  - System Control
    - M5/M10 and M7i—Forwarding Engine Board (FEB)
    - M7i /M10i routers—RE and Fixed Interface Card/High-Availability Chassis Manager (FIC/HCM)
    - M20 router—System Switching Board (SSB)
    - M40 router—System Control Board (SCB)
- For M40e and M160
  - Switching and Forwarding Module (SFM)
  - Miscellaneous Control Subsystem (MCS)
  - Packet Forwarding Engine Clock Generator

Copyright © 2006, Juniper Networks, Inc.

#### Physical Interface Cards

Juniper Networks M-series platforms provide a complete range of fiber-optic and electrical transmission interfaces to the network through a variety of Physical Interface Cards (PICs). These space-efficient modules offer exceptional flexibility and high port density.

#### Flexible PIC Concentrator

Flexible PIC Concentrators (FPCs) house the PICs and provide shared memory for the M-series switch fabric. These intelligent, high-performance interface concentrators allow you to mix and match PIC types within a given FPC.

#### The Midplane

The system midplane is the component of the Packet Forwarding Engine that distributes power and electrical signals to each card in the system. The midplane is passive in all M-series routers except the M40. The M40 router's midplane houses the Distributed Buffer Manager ASICs.

#### System Control Boards

On M-series platforms, the System Control Boards provide the route lookup component of the Packet Forwarding Engine using the Internet Processor II ASIC. Each System Control Board on M-series routers provides the same function, despite their having different names. On the M5 and M10 routers, the FPC and control board components are combined onto a single board called the Forwarding Engine Board (FEB).

*Continued on next page.*

## Configuring Juniper Networks Routers

### System Control Boards (contd.)

The M7i and M10i combine the same functionality into a smaller version of the FEB called a Compact FEB (CFEB); note that CFEBs are only interchangeable between the M7i and M10i when the optional Adaptive Services PIC (ASP) is not installed. On the M7i and M10i platforms chassis control functions, such as PIC online/offline and chassis monitoring is performed by the Fixed Interface Card (FIC) or High-Availability Chassis Manager (HCM), respectively.

The M-series System Control Board (FEB, CFEB, or SSB) also houses the buffer management ASICs on all models except the M40 router.

### The M40e and M160 Platforms

The M160 and the M40e platforms (the latter being a scaled down version of the M160) differ from the other M-series platforms in a number of ways. For example, the route lookup function associated with an M20 router's SSB is now performed by the Switch Fabric Module (SFM). Differences between the original M-series platforms and the M40e and M160 platforms include:

- *Route lookup:* On M40e and M160 platforms, the Internet Processor II route lookup ASIC is housed on the Switch Fabric Module (SFM). An M160 system supports a total of four SFMs. Each SFM is capable of performing 40 million packet lookup operations per second. The presence of four SFMs yields the 160-Mpps capacity of the M160 router. The failure of an SFM gracefully reduces total system throughput by approximately 25%. The M40e router supports a total of two SFMs, with only one SFM active at any given time. The failure of the active SFM results in the automatic switchover to the spare SFM when so equipped.
- *System control:* The Miscellaneous Control Subsystem (MCS) card works with the active Routing Engine to provide control and monitoring functions for the various components in the chassis.
- *PFE clock generation:* The active host module on an M40e or M160 platform requires a PFE Clock Generator (PCG) to provide a 125-MHz signal that clocks the various gates in the PFE complex. Redundant PCGs are supported.

### PFE Components: T-series

- **Physical Interface Cards (PICs)**
- **T-series FPCs contain one or two PFE complexes**
  - PFEs interface to other PFEs through the T-series switch fabric
    - Nonblocking crossbar switch matrix with high-speed lines to each FPC
    - Switch fabric redundancy
- **Switching between PFEs performed by Switch Interface Boards (SIBs)**
  - Three SIBs comprise a T320 switch fabric—two active, one spare
  - Five SIBs comprise the T640 switch fabric—four active, one spare
- **The system midplane**

Copyright © 2006, Juniper Networks, Inc.

#### Physical Interface Cards

As with the M-series routers, PICs provide the T-series PFE with a large range of fiber-optic and electrical transmission interfaces to the network.

#### The T-series PFE

T-series platforms implement either one or two complete PFE complexes on each FPC. On the T640 platform a single PFE is present on the FPC2 while two PFEs are present on the FPC3. The latter FPC type is designed specifically for native T-series PICs. Packets that ingress and egress on the same PFE complex (for example, on PICs 0 and 1, or 2, and 3 of a given FPC) do not leave that PFE. Packets are switched between PFEs across the T-series switch fabric as needed.

#### The T-series Switch Fabric

T-series platforms make use of a shared memory switch fabric for intra- and inter-FPC/PFE communications. In addition, inter-FPC communications requires transit of the T-series cross-bar switch fabric, which is instantiated by the system's Switch Interface Boards (SIBs). The T320 platform can support up to three SIBs while the T640 platform supports five. In the case of the T640 platform, four SIBs provide the necessary speedup for a nonblocking architecture. The fifth SIB is only used in the event of a SIB failure. The system's throughput gracefully degrades in the unlikely event of multiple SIB failures. In normal operation the T320 router makes use of SIBs 1 and 2 with SIB 0 functioning as a standby. In the event of a SIB failure, SIB 0 automatically becomes active. There might be a slight performance degradation when using SIB 0 because each FPC only has one high-speed line (HSL) to SIB 0 (two HSLs interconnect the FPCs to each of SIB 1 and SIB 2).

#### The Midplane

The midplane distributes power and electrical signals to the components and cards that make up the PFE and switch fabric.

### PFE Components: M320

- The M320 router is an M-series platform based on the T-series ASICs
  - Optimized for edge aggregation
- M320 FPCs contain one or two PFE complexes
  - PFEs interface to other PFEs through the a switch fabric
    - Nonblocking crossbar switch matrix with high-speed lines to each FPC
    - Switch fabric redundancy
- Switching between PFEs performed by Switch Interface Boards (SIBs)
  - Up to four SIBs comprise the M320 switch fabric
  - Adding SIBs enables line-rate performance for more or higher capacity PICs

Copyright © 2006, Juniper Networks, Inc.

#### M-series Platform with a T-series PFE

The M320 platform is somewhat unique in the M-series and T-series portfolio in that it marks the first instance of repurposing one platform group's technology for use in another platform group's application. Specifically, the M320 uses T-series ASICs, but it is optimized for dense edge aggregation and service creation, rather than core functionality.

#### The M320 FPC

As with the T-series platforms, the M320 FPCs contains from one to two complete PFE complexes. As was the case with T-series platforms, packets that ingress and egress on the same PFE complex (for example, on PICs 0 and 1, or 2, and 3 of a given FPC) do not leave that PFE. Packets are switched between PFEs across the cross-bar switch fabric as needed.

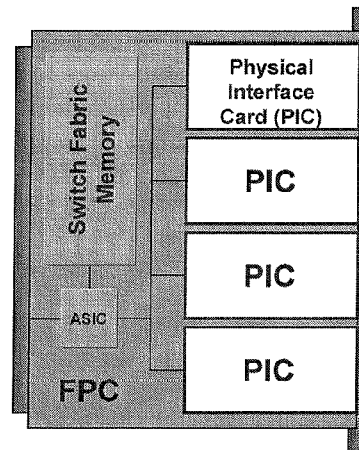
#### The M320 Switch Fabric

The M320 platform also makes use of a shared memory switch fabric for intra- and inter-FPC/PFE communications. In addition, inter-FPC communications requires transit of the T-series based cross-bar switch fabric, which is formed by the system's four Switch Interface Boards (SIBs). You can configure an M320 with one to four active SIBs. Adding and activating more SIBs enables you to maintain line-rate forwarding performance for larger numbers of higher-bandwidth PICs (see the *M320 Hardware Guide* for detailed performance breakdown). The M320 can operate with a SIB in standby mode for fault tolerance, but it does not have room for a fifth SIB. Thus, if you configure an M320 to have a bandwidth requirement for four SIBs, and you have a SIB failure, the router will have reduced performance until you replace the affected SIB.



### Physical Interface Cards

- PICs currently support from 0 to 48 physical ports
  - Some PICs support channelized and advanced CoS options
  - IP Service PICs (Tunnel, Multilink, Monitoring, etc.)
    - Services PIC normally have no physical ports
- Media-specific ASICs
- Status indicators
- Hot-swappable on all platforms except M20 and M40 routers
  - Slight disruption to other PICs in same FPC on M-series router



Copyright © 2006, Juniper Networks, Inc.

#### PIC Overview

PICs provide the physical connection to various network media types. PICs receive incoming packets from the network and transmit outgoing packets to the network. During this process, each PIC performs appropriate framing and signaling for its media type. Before transmitting outgoing data packets, the PIC adds media-specific framing to the packets received from the FPCs. You can install up to four PICs into slots on each FPC. PIC types can be intermixed within the same FPC. The number of ports on a given PIC varies with the PIC and platform type. As of this writing, M40e PICs are available with as many as 48 Fast Ethernet ports!

IP Services PICs enable a hardware assist for complex packet processing functions. Examples include the Tunnel Services and Multilink Services PICs. With the Tunnel Services PIC, routers can function as the ingress or egress point of an IP-IP unicast tunnel, a generic routing encapsulation (GRE) tunnel, or a Protocol Independent Multicast sparse mode (PIM-SM) tunnel. The Multilink PIC uses the Multilink Point-to-Point Protocol (MLPPP) and Multilink Frame Relay (MLFR, FRF 1.5) to group up to eight T1 or E1 links per bundle to yield a service offering ranging from 1.5 Mbps through 12 Mbps (T1) or 2 Mbps through 16 Mbps (E1).

#### Media-Specific ASIC

Each PIC is equipped with an ASIC that performs control functions tailored to the PIC's media type. For instance, an ATM PIC and a Fast Ethernet PIC each contain unique ASICs (or FPGA) that are specifically suited to the particulars of each medium.

*Continued on next page.*

## Configuring Juniper Networks Routers

### **PIC Status**

Each PIC supports one or more status LEDs that accommodate quick verification of the PIC, and in some cases, the port's operational status.

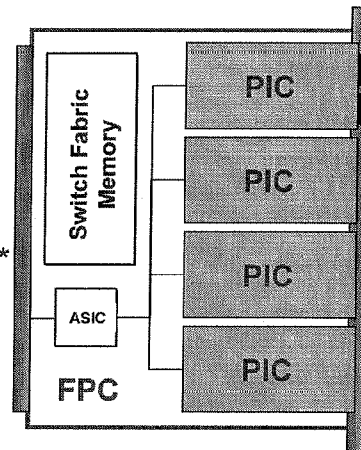
### **Hot-Pluggable in Most Platforms**

You can replace or install PICs without removing the associated FPC on all platforms except the M20 and M40, which require that the host FPC (which is hot-removable and hot-insertable) be removed to gain access to the PIC's mounting screws.

Note that you should always take care to take a PIC offline before removing it from its FPC to minimize system disruption. You should expect small amounts of packet loss on all PICs sharing the affected FPC when hot-swapping PICs on M-series platforms (excludes the M320, which is based on a T-series PFE). This momentary disruption is the result of the FPC undergoing a logical reset in reaction to the insertion and removal of a PIC. Failing to take a PIC offline before removing it from its FPC can result in damage to the system or a PFE reset.

### The Flexible PIC Concentrator

- **General FPC features**
  - Supports from 1 to 4 PICs
  - Hot-swappable on most platforms
  - PowerPC supervisory processor
    - Not used for packet forwarding
  - From 64 MB to 1.2 GB of memory
    - Pooled to create shared memory switch fabric on M-series platforms
- **High aggregate throughput rates\***
  - M5/M10, M7i/M10i, M20, M40, and M40e routers: 6.4 Gbps per FPC
  - M160 router: 25.6 Gbps per FPC2
  - T640 Internet Routing Node: 80+ Gbps with FPC3
  - T320 router: 40+ Gbps with FPC3



\* The numbers quoted are two times the unidirectional (simplex) capacity of each FPC.

Copyright © 2006, Juniper Networks, Inc.

### General FPC Characteristics and Features

FPCs install into the backplane from the front of the chassis. You can install an FPC into any FPC slot; a specific order is not required. If an FPC does not occupy a slot, you must install a blank FPC carrier to shield the empty slot so that cooling air can circulate properly through the card cage. FPCs can support from one to four PICs, depending upon specifics. For example, an OC-192 interface on an M160 platform has an FPC *built around it* because the high-speed interface consumes all four PIC positions. This yields the lowest possible density of one PIC/FPC. Most FPCs support four PIC connectors; the current exception is the T320 platform, which supports two PIC connectors per FPC.

When you install an FPC into a running system, the FPC requests its operating software from the Routing Engine, the FPC runs its diagnostics, and the PICs on the FPC slot are enabled. FPCs are hot-insertable and hot-removable on all platforms except the M5/M10 and the M7i/M10i platforms. This is because these routers have FPCs that are combined with the system board components to create a Forwarding Engine Board/Compact Forwarding Engine Board (FEB/CFEB). On the M5 and M10 platforms, the FEB is not hot-pluggable/insertable. You must remove power before inserting or removing the FEB. The CFEB on the M7i and M10i is hot-insertable, however.

*Note that when you remove or install an FPC on an M-series platform the system must repartition the shared memory pool; this process results in about 200 milliseconds of disruption to all packets associated with the affected PFE. T-series platforms contain from one to two complete PFEs on each FPC, and therefore packet forwarding on one FPC is not affected by the removal or insertion of other FPCs.*

*Continued on next page.*

## Configuring Juniper Networks Routers

### General FPC Characteristics and Features (contd.)

Some high-speed PICs, like the OC-48c/STM-16 for the M20/M40 and the M160 platform's OC-192c/STM-64 SONET/SDH PIC, are quad-wide and do not require an FPC because quad-wide PICs have FPC functionality built in.

A portion of the memory associated with each FPC is pooled together with the memory from other FPCs to create the M-series shared memory switch fabric. The actual amount of FPC memory varies by FPC type, but in all cases there is at least 100 milliseconds of delay buffer (for each transmit and receive yielding a total of 200 milliseconds of delay/bandwidth buffering). Currently, the amount of memory present on a given FPC ranges from 64 MB on the M5/M10 to 1.2 GB on the T640 FPC3. In the latter case, this yields approximately 600 MB per PFE complex.

Some routing platforms support multiple FPC types to allow customers to reuse PICs from earlier platforms. For example, the original FPC is the only FPC type supported on the M20, and M40 platforms. The M160's FPC1 is designed to support the reuse of M20 and M40 PICs in the higher-end M160 platform. The M160's FPC2, on the other hand, supports PICs designed specifically for the M160 platform's increased throughput like the OC48-c PIC. In a similar fashion, the T320 platform supports three types of FPCs (type 1, 2, and 3). The type 3 FPC offers support for native T-series PICs while the type 2 and type 1 FPCs offer support for M160 and M20/40 PICs respectively. The T640 platform supports only type 2 and type 3 FPCs at this time.

### Industry-Leading Throughput

Most M-series routers have a aggregate slot throughput of 6.4 Gbps; the M160 and M320 platforms support an aggregate capacity of 25.6 Gbps and 40 Gbps respectively.

T-series platforms increase aggregate throughput to a respectable 40 Gbps for the T320 platform and 80 Gbps for the T640 platform when using the native FPC3.

## M-series System Boards

- **General System Board functions:**
  - Forwarding table updates and route lookups
  - Management of ASICs and PFE hardware components
  - Environmental monitoring
  - Stratum 3 SONET clock generation
  - Handling exception/control packets
- **Names vary by platform**
  - M5 and M10—FEB
  - M20 and M40—SSB and SCB
  - M7i and M10i—CFEB
- **Enhanced System Boards feature the second generation Internet Processor II ASIC**

Copyright © 2006, Juniper Networks, Inc.

### General M-series System Board Functions

The M-series System Board houses the Internet Processor II ASIC and performs a variety of functions. These include:

- *Route lookups and forwarding table maintenance:* The Internet Processor ASIC performs route lookups using a forwarding table stored in the chip's synchronous SRAM (SSRAM). The System Board updates its copy of the forwarding table when instructed by the JUNOS software kernel. The M20's SSB contains the memory management ASICs as well as the Internet Processor II ASIC. The M40's SCB does not contain the memory management ASICs, as these are part of the M40's midplane.
- *Management of ASICs and PFE components:* The System Board monitors various system components for failures and alarm conditions. It collects statistics from all sensors in the system and relays them to the Routing Engine, which sets the appropriate alarm. For example, if a temperature sensor exceeds the first internally defined threshold, the Routing Engine issues a *high temp* alarm. The System Board handles the power up and power down of PFE components with diagnostic errors reported to the RE over the 100 Mbps *fxp1* interface.
- *Environmental monitoring:* The System Board monitors the various temperature sensors to control fan speed and over-temperature alarm generation.
- *SONET clock:* The System Board generates a Stratum 3 clock reference used to clock SONET interfaces.

*Continued on next page.*

## Configuring Juniper Networks Routers

### General M-series System Board Functions (contd.)

- *Transferring exception and control packets:* The Internet Processor ASIC passes exception packets to a microprocessor on the System Board, which processes almost all of them. Remaining packets are sent to the RE for further processing. Errors originating in the Packet Forwarding Engine detected by the System Board are sent to the RE where they are logged and made available to the CLI.

### The Names Vary

The System Board names vary by platform. The M40 platform has a System Control Board (SCB), while the M20 platform has a System Switching Board (SSB). SSB redundancy is supported on the M20 router. The M5/M10 and M7i/M10i platforms integrate System Board functionality into their Forwarding Engine Board/Compact Forwarding Engine Board (FEB/CFEB).

### Enhanced System Boards

Enhanced System Boards support the second generation Internet Processor II ASIC on M-series platforms (except the M5/M10 and M7i/M10i platforms). Enhanced system boards offer improved performance and scalability. For example, the size of the forwarding table is increased from approximately 420,000 entries to approximately 840,000 entries with an enhanced S-board. Some of the enhancements present in the second generation Internet Processor II ASIC are listed here:

- Doubles the amount of on-chip memory (now 16 MB).
- Increases memory to 128 MB in the CPU complex for the M40 router, and to 256 MB for the M20, M40e, and M160 routers.
- Increased CPU speed (now 256 MHz).

Enhanced System Boards first shipped with JUNOS software Release 5.5 circa September 2002.

### Control Boards: M-series and T-series

---

- **General Control Board functions:**
  - Component power up/down
  - Handling hardware faults
  - Controlling redundancy
  - Environmental monitoring
  - Distribution/generation of SONET clocking
- **M160/M40e control**
  - Control provided by Miscellaneous Control Subsystem (MCS); paired with a Routing Engine to form a Host Module
    - Host Module redundancy supported
- **M320, T640, and T320 control**
  - Control provided by Control Board (CB); the CB is paired with a Routing Engine to form a Host Subsystem
    - Host Subsystem redundancy supported

Copyright © 2006, Juniper Networks, Inc.

#### General Control Board Functions

Newer M-series routers and all T-series platforms make use of a Control Board to provide some of the functionality associated with the System Board found on previous M-series platforms. These functions include:

- *Management of ASICs and PFE components:* The Control Board handles the power up and power down of other PFE components. Diagnostic errors are reported to the RE over the 100-Mbps `fxp1` interface.
- *Environmental monitoring:* The Control Board monitors the various temperature sensors to control fan speed and over-temperature alarm generation.
- *SONET clock generation:* On M-series platforms the Miscellaneous Control Subsystem generates and distributes a Stratum 3 clock reference used to clock SONET interfaces. On T-series platforms the SONET Clock Generator (SCG) generates the Stratum 3 reference clock that is then distributed to the PFE components by the system's active Control Board.

#### M160/M40e Platforms


The M160 and M40e platforms use a Miscellaneous Control Subsystem (MCS) board to provide control functions. The MCS works in conjunction with a Routing Engine to form a Host Module. Host module redundancy is supported, as is redundant MCSs that are controlled by a common (nonredundant) RE.

#### M320 and T-series Platforms

The M320 and T-series platforms use a Control Board (CB) to provide control functions. The CB works in conjunction with an RE to form a Host Subsystem. Host Subsystem redundancy is supported in all cases.

### Internet Processor II ASIC

- **The Internet Processor II**
  - Provides industry-leading performance for longest-match packet lookup
  - Numerous packet processing features:
    - Filtering, sampling, logging, counting, and improved load balancing
  - Second generation Internet Processor II available on enhanced system boards
    - Standard on T-series FPCs



Copyright © 2006, Juniper Networks, Inc.

#### Internet Processor II

The Internet Processor ASIC, which first shipped with the M40 router in September 1998, heralded a breakthrough technology that facilitated longest-match traffic forwarding for virtually all packet sizes at or very near line rate. Performance tests in the lab, test networks, and on the Internet itself have all demonstrated 40 Mpps of 40-byte packets with 80,000 prefixes in the routing table!

Building on this tradition, the Internet Processor II ASIC continues to deliver best-of-class functionality for network core and edge applications. Simply put, as of this writing, T-series platforms are the highest performing systems on the market. While the Internet Processor II ASIC still delivers a 40-Mpps forwarding rate, the new ASIC adds rich packet processing features that include firewall filtering, sampling, logging, counting, and enhanced load balancing. The Internet Processor II ASIC maintains high performance in the presence of value-added feature sets and enhanced services.

For systems that did not originally ship with an Internet Processor II ASIC, a simple field upgrade of the system board is all that you need to enable Internet Processor II functionality on all existing interfaces. A second generation Internet Processor II ASIC is offered on enhanced system boards, which are now available for all M-series platforms except the M5 and M10. The enhanced S-board contains the second generation Internet Processor II processor, which sports a faster clock speed and increased memory.


All T-series platforms make use of the latest Internet Processor ASIC technologies. In fact, a T640 platform might contain as many as 16 Internet Processor II chips. This is because on a T-series platform each FPC can contain from one to two complete PFE complexes, and each such PFE is serviced by its own Internet Processor II (which is called the "R" chip internally).



## System Midplane Examples


- **M10/M10i system midplane:**

- FEB contains built-in FPCs, eight PIC slots




- **M40e, M160, T640, and T320 system midplane**

- Connector Interface Panel (CIP), eight FPC slots



- **M20 system midplane**

- System Switching Board slots, Craft Interface slot, four FPC slots



Copyright © 2006, Juniper Networks, Inc.

### M5/M10 and M7i/M10i System Midplanes

The M5/M10 and M7i/M10i systems are based on FPC and System Board functionality combined into the Forwarding Engine Board/Compact Forwarding Engine Board (FEB/CFEB). Each FPC hold four PICs. Thus, the M10/M10i routers support up to eight PICs, while the M5/M7i routers support up to four PICs. The M7i platform features a Fixed Interface Card (FIC). The M7i FIC supports two Fast Ethernet interfaces, or one Gigabit Ethernet interface, depending on the specific configuration, and also provides alarm LEDs and the PIC online/offline buttons. The M10i platform uses a Chassis Management Board (CMB) for PIC online/offline buttons and alarm indicators.

The M7i platform features an integral Tunnel Services PIC with optional support for an Adaptive Services PIC (ASP) (also internal). When so equipped, the ASP takes over the bandwidth and FPC/PIC position (1/2/0) associated with the integral Tunnel Services PIC to provide both conventional tunnel and adaptive services. While the M10i platform does not support a FIC or an integral Services PIC, the M10i platform features optional RE and CFEB redundancy.

### M40e, M160, T640 and T320 System Midplanes

The midplanes for these platforms support up to eight FPCs (0–7 counting from left to right). The midplane also contains the connector interface panel (CIP) slot.


### M20 System Midplane

The M20 system midplane can hold up to four FPC slots (0–3 counting from top to bottom). This midplane also contains the System Switching Board (SSB) slots (control board redundancy is supported) and the Craft Interface slot.

## Configuring Juniper Networks Routers

### The Craft Interface

- **Craft Interface overview**
  - LCD display (certain platforms only)
  - FPC online/offline buttons (M20, M40, M40e, M160, T640, and T320 platforms)
  - PIC online/offline buttons (M5/M10 and M7i/M10i routers)
    - For M7i/M10i PIC online/offline button are on the FIC/HCM respectively
  - Status LEDs



**A Typical Craft Interface Panel (T320)**

Copyright © 2006, Juniper Networks, Inc.

### Craft Interface

The Craft Interface is the collection of mechanisms on Juniper Networks M-series and T-series platforms that allows you to view system status messages and troubleshoot the router. The Craft Interface is located on the front of the chassis and typically consists of various system status LEDs and FPC (or PIC) online/offline buttons. On supported platforms the Craft Interface includes an LCD screen that provides status reporting for the entire system.

The M7i platform's Fixed Interface Card (FIC) and the M10i platform's High-Availability Chassis Manager (HCM) card provide PIC offline and online functionality.

## Configuring Juniper Networks Routers


### Status LEDs and FRU Offline/Online

#### ● Status LEDs

- OK
  - Blinking = starting
  - Solid = running
- FAIL
  - Solid = taken offline because of failure

#### ● Online/offline buttons

- Press and hold for three seconds to take FPC (or PIC) offline



#### System Status LEDs

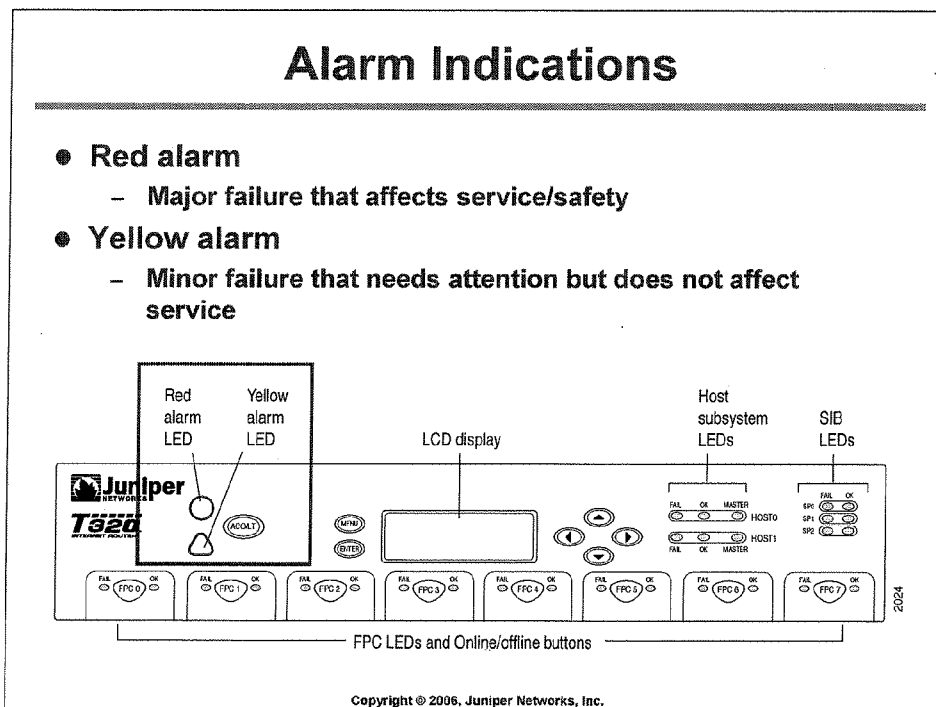
The system status LEDs include:

- **FPC LEDs:** Two LEDs exist—one green OK and one red fail. These lights indicate the status of each FPC. Each LED pair is located on the Craft Interface aligned with the corresponding FPC module slot.
- **Routing Engine/Host LEDs:** A red fail LED and a green OK LED on the Craft Interface indicate the status of the Routing Engines/Host modules.

#### FPC and PIC Offline Buttons

FPC (or PIC) offline buttons allow you to take an FPC (or PIC) offline gracefully. Press and hold the offline button near the FPC (or PIC) until the green OK LED extinguishes. For systems with fixed FPCs, like the M5, M10, M7i, and M10i, the online/offline buttons are used to prepare a PIC for removal from the system. Note again that for the M7i and M10i platforms, the PIC online/offline buttons are located on the FCM/HCM respectively.

## Configuring Juniper Networks Routers



### Red Alarm

The red alarm LED indicates a system failure likely to cause an interruption in service. Examples of red alarms are:

- Routing Engine failure;
- Cooling system failure; and
- Interface loss of light or framing.

### Yellow Alarm

The yellow alarm LED indicates a system warning not likely to interrupt service, but if left uncorrected, might eventually cause a service interruption. Examples of yellow alarms are:


- Maintenance alert;
- FPC with recoverable errors; and
- Cooling system problems.

You can configure the mapping of various events to an alarm action of ignore, yellow, or red. Environmental and safety-related alarms cannot be remapped, however.

## Configuring Juniper Networks Routers

### The LCD Display

- LCD display is available on M40, M40e, M160, M320, T640, and T320 platforms only
  - Displays general system status when no alarms are present
  - Displays alarm information when alarms are present
    - Identifies the total number and types of alarms that are active
  - Currently, the navigation buttons are only used to obtain the status of certain PICs



### LCD Display

The Craft Interface on selected platforms supports a four-line LCD screen with six navigation buttons. The LCD screen operates in one of two display modes. The default mode, *idle mode*, displays the current system status until it is preempted by *alarm mode*. The following list contains the basic status information displayed:

- Router's name, on the first line;
- Number of days, hours, minutes, and seconds that the system has been running, on the second line; and
- Status messages, on the fourth line, which are various system status messages that cycle at 3-second intervals.

You can alter the idle mode display by specifying a message of your choosing with the `set chassis display message` operational-mode command. The Craft Interface display cycles between the user-defined and standard display every 2 seconds unless you also give the `permanent` argument. The user-defined message only persists for 5 minutes, however. You can view the LCD display, along with an ASCII representation of the status LEDs, with a `show chassis craft-interface` operational-mode command.

*Continued on next page.*

## Configuring Juniper Networks Routers

### LCD Display (contd.)

The following example shows a custom user message being displayed:

```
lab@San_Jose-3> show chassis craft-interface
```

```
Red alarm:          LED off, relay off
Yellow alarm:       LED off, relay off
Routing Engine OK LED:  On
Routing Engine fail LED: Off
```

```
FPCs      0  1  2  3
-----
Green     *  *  .  .
Red       .  .  .  .
```

LCD screen:

```
+-----+
| "NOC contact Foo @ |
| 555-1212"          |
|                    |
|                    |
+-----+
```

*Alarm mode* displays alarm conditions whenever the red or yellow alarm LED is lit. When a red or yellow alarm occurs, alarm mode preempts idle mode and the LCD displays a message to alert you of serious alarm conditions. In alarm mode, the screen displays the following information:

- Router's name, on the first line;
- Number of alarms active on the router, on the second line; and
- Individual alarms, with the most severe condition shown first, on the third and fourth lines. Each line indicates whether the alarm is a red (R) or yellow (Y) alarm.

Idle and alarm present display samples are provided here:

```
godzilla
Up 49+11:21:31

System Chassis OK
```

**Normal (Idle) Mode**

```
godzilla
2
R: Fan failure
Y: Change air filter
```

**Alarms Active Mode**

Currently, the only use for the menu and navigation buttons associated with the LCD is to display the port status for certain high-density PICs, such as the 12-port and 48-port Fast Ethernet PICs supported on the M40e and M160 platforms. To display port status on such a PIC, follow these steps:

1. Locate the LCD and select MENU.
2. Choose `fe pic status` and press ENTER.
3. Scroll the arrow buttons to select the slot and PIC number. Then press Enter to see port status.
4. Read the port numbers vertically. You will see one of three symbols:
  - \* (asterisk—equivalent to green port LED): Port is active and receiving data.
  - - (minus sign—equivalent to flashing green port LED): Port might be active but is not receiving data.
  - Blank: Port is not active.

## Agenda: Product Line Overview

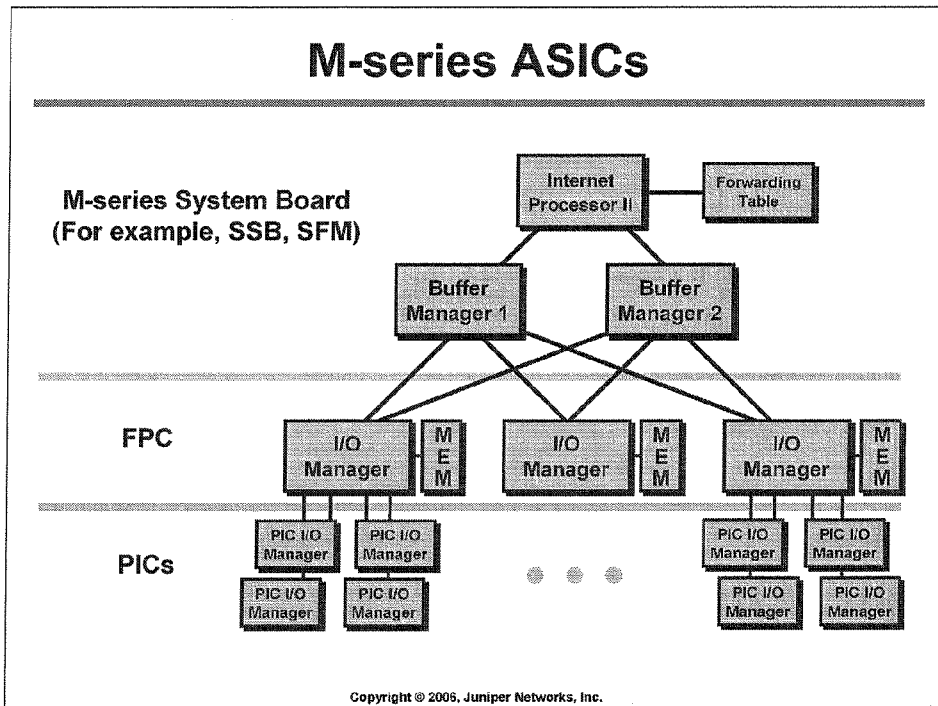
- M-series and T-series Overview
- Installation and Handling Guidelines
- Platform Architecture and Hardware Overview
- **M-series Packet Flow and ASIC Functionality**
- T-series Packet Flow and ASIC Functionality
- Interface Overview
- Additional Products

Copyright © 2006, Juniper Networks, Inc.

### **M-series Packet Flow and ASIC Functionality**

The following pages examine the ASICs and packet flow through the M-series platforms.

## Configuring Juniper Networks Routers

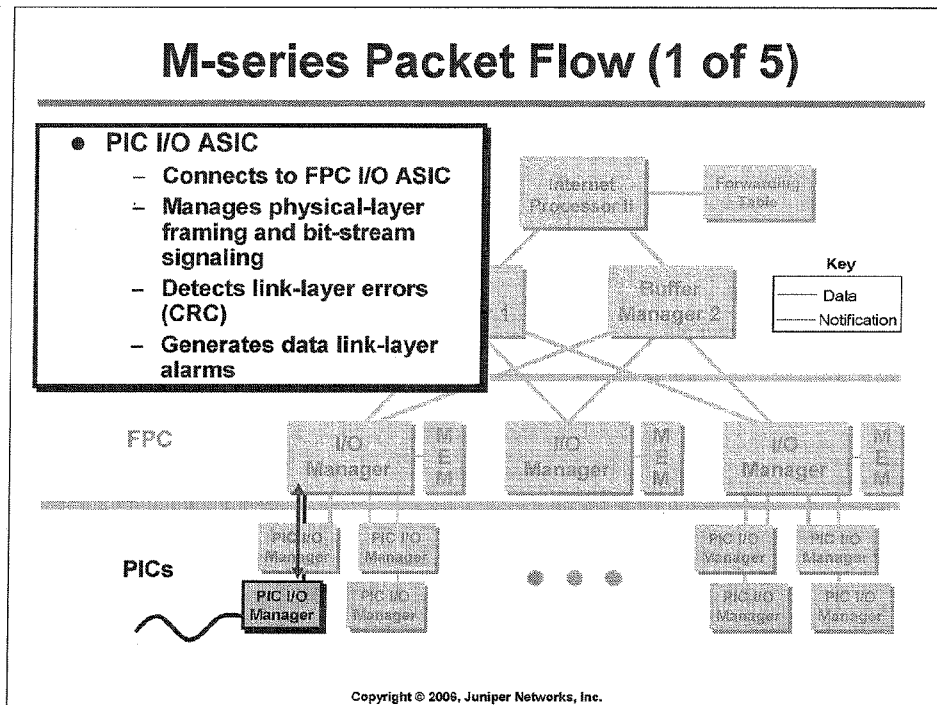


### M-series ASICs

This slide displays the ASICs that make up an M-series router's Packet Forwarding Engine (PFE). The function of each ASIC is detailed on subsequent pages. In M-series platforms the ASICs that comprise the PFE are located in the PICs, FPCs, and the System Board. On the M40 router the buffer management ASICs are mounted on the system midplane. The M5/M10 and M7i/M10i routers combine FPC and System Board functionality into the FEB/CFEB. The CFEB makes use of a combined I/O manager, Distributed Buffer Manager, and Internet Processor II ASIC to reduce cost and power consumption while also improving reliability. This ASIC is sometimes called the *ABC* ASIC in keeping with the internal ASIC designation of A, B, and C for the Distributed Buffer Manager, I/O Manager, and Internet Processor II ASICs respectively.



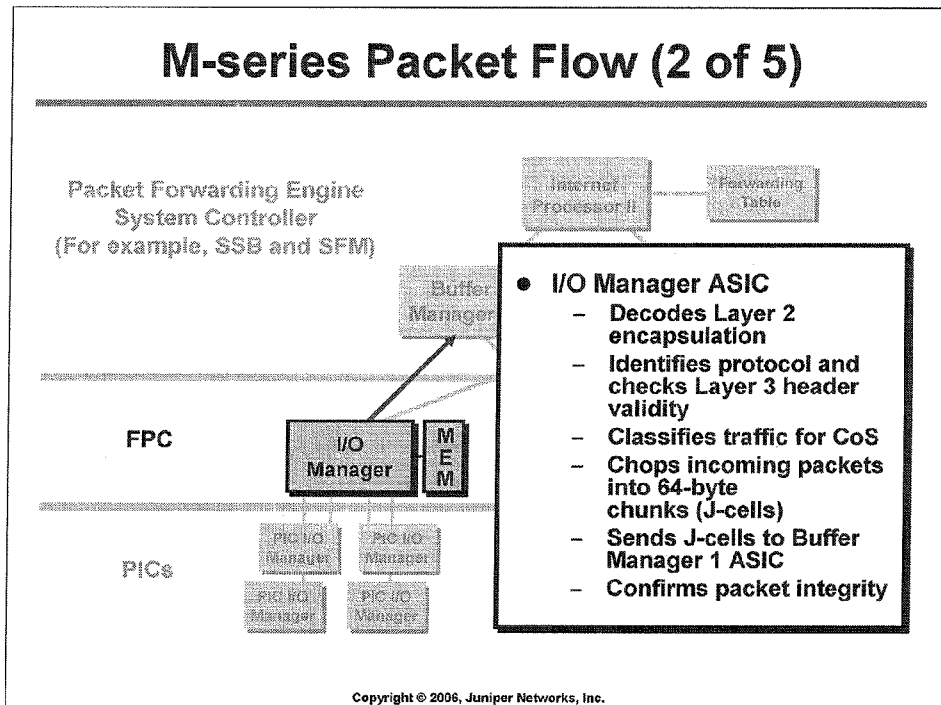
## Configuring Juniper Networks Routers



### M-series Packet Flow: Part 1

When a packet arrives on an input interface of the router, the PIC controller ASIC performs all the media-specific operations such as physical-layer framing and link-level FCS (CRC) verification. The PIC then passes a serial stream of bits to the I/O Manager ASIC on the FPC.

## Configuring Juniper Networks Routers



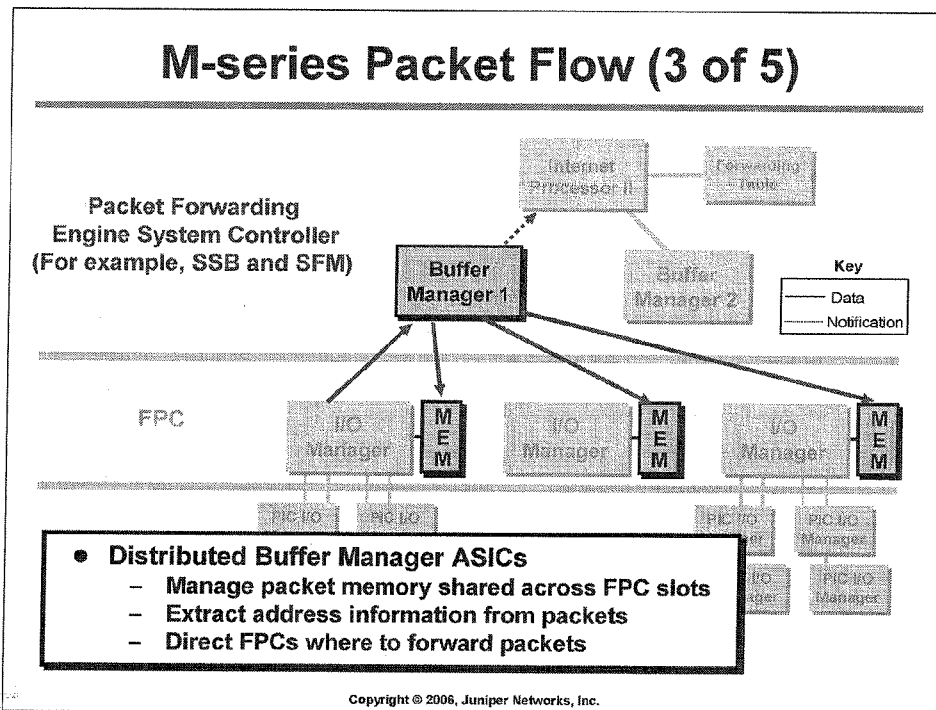
### M-series Packet Flow: Part 2

The I/O Manager ASIC parses the bit stream to locate the Layer 2 and Layer 3 encapsulation and chops the packet into 64-byte chunks called J-cells. These J-cells are then sent to the inbound Distributed Buffer Manager ASIC.

The I/O Manager ASIC also:

- Removes Layer 2 encapsulation to locate the beginning of the Layer 3 packet;
- Identifies incoming logical interface;
- Performs basic packet integrity checks;
- Counts packets and bytes for each logical circuit; and
- Performs BA-based traffic classification to associate traffic with a forwarding class for egress queuing and scheduling operations. Examples of BA classification include IP precedence, DiffServ code points, and MPLS EXP bit settings.

## Configuring Juniper Networks Routers

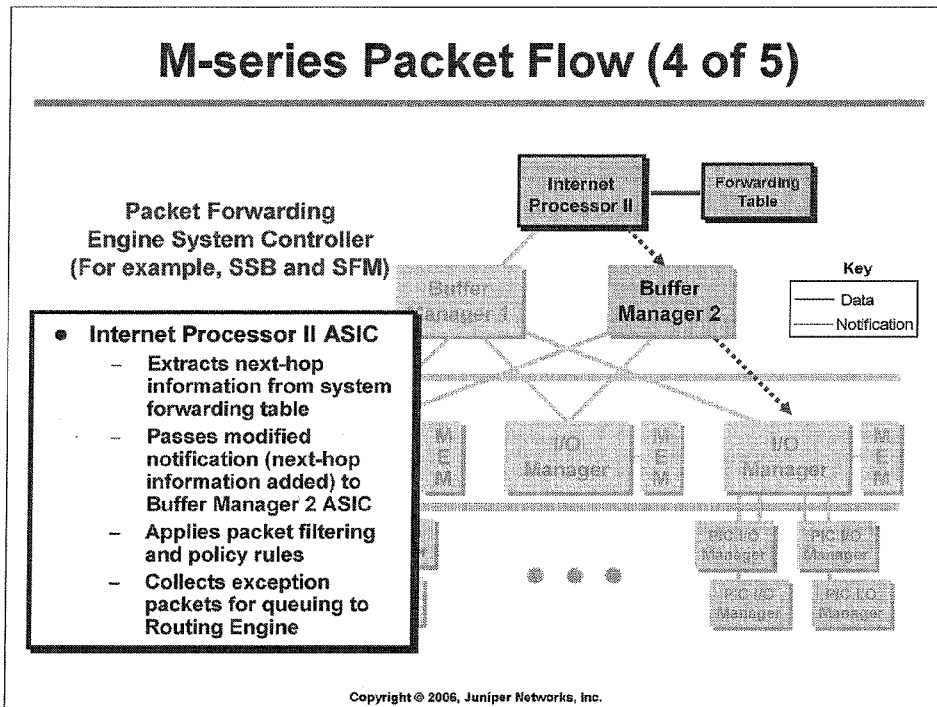


### M-series Packet Flow: Part 3

The Distributed Buffer Manager 1 ASIC receives J-cells from each FPC's I/O Manager ASIC and writes them into the shared memory bank. The shared memory bank is made up of memory contributed by each FPC installed in the router.

The Buffer Manager 1 ASIC also extracts the key information, which is normally the first 64-bytes of a Layer 3 packet, and passes this information to the Internet Processor II ASIC in the form of a notification cell. The Internet Processor II performs a longest-match route lookup against the forwarding table to identify the packet's outgoing interface and forwarding next hop.

## Configuring Juniper Networks Routers



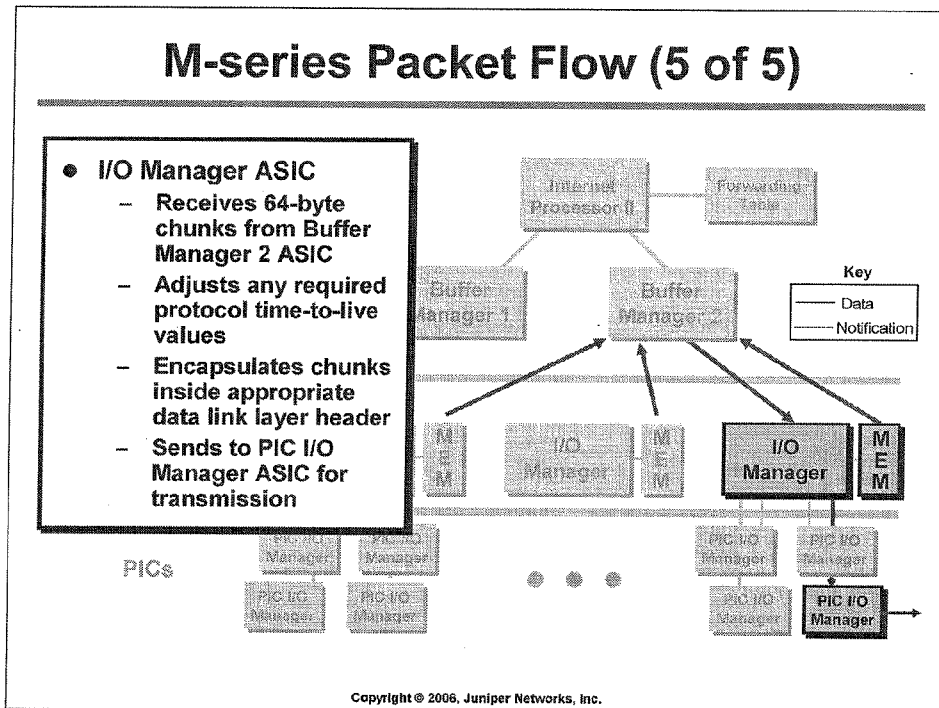
### M-series Packet Flow: Part 4

The Internet Processor II ASIC determines the ultimate destination for every packet arriving on a transit interface. The Internet Processor ASIC consults a copy of the forwarding table, which contains destination prefixes and their corresponding next hops. The forwarding table is constructed by the Routing Engine and maintained by the JUNOS software kernel.

After the Internet Processor II ASIC determines the packet's egress interface and forwarding next hop, it amends the notification cell with this information and passes the notification cell to the second Distributed Buffer Manager ASIC. The second Distributed Buffer Management ASIC then passes the notification cell to the I/O Manager ASIC on the egress FPC (as identified by the modified contents of the notification cell). The Distributed Buffer Manager 2 ASIC acts as an agent for the FPC's I/O Manager ASIC. Once the I/O Manager ASIC receives a notification cell indicating that a packet is waiting to be serviced, it issues read requests to the Buffer Manager 2 ASIC for the J-cells associated with this packet. As the I/O Manager receives the J-cells, it transmits them to the PIC Controller ASIC, which in turn transmits them out the appropriate port.

In the case of a multicast packet, multiple outgoing interfaces might exist, in which case the notification cell is directed to multiple FPCs or to the same FPC multiple times, once for each outgoing interface served by that FPC.

## Configuring Juniper Networks Routers



### M-series Packet Flow: Part 5

When the egress FPC is ready to service the packet, the I/O Manager ASIC issues read requests for the 64-byte J-cells that comprise the packet. In response, the Distributed Buffer Manager 2 ASIC retrieves the J-cells from shared memory and feeds them to the I/O Manager ASIC. The I/O Manager ASIC reassembles the packet, decrements the packet's TTL, adds the Layer 2 framing, and then sends the bit stream to the egress PIC.

The I/O Manager ASIC is responsible for CoS-related queuing, scheduling, and congestion avoidance operations at packet egress. Note that the packet itself is never queued on the FPC; rather, a pointer to the packet, in the form of a notification cell, is queued on the egress FPC. Each output port on a given PIC is associated with four forwarding classes (or queues). You configure schedulers to provide each forwarding class with some share of the port's bandwidth.

Note that traffic classification, which associates traffic with one of the defined forwarding classes, occurs at the ingress FPC. Once identified at ingress, the traffic is handled in accordance with the parameters configured for that traffic class by the I/O Manager on the egress FPC. The I/O Manager implements the random early detection (RED) algorithm during egress processing to avoid tail drops and the resulting risk of global synchronization of TCP retransmissions. A full coverage of JUNOS software CoS capabilities is beyond the scope of this class.

### **Agenda: Product Line Overview**

---

- M-series and T-series Overview
- Installation and Handling Guidelines
- Platform Architecture and Hardware Overview
- M-series Packet Flow and ASIC Functionality
- **T-series Packet Flow and ASIC Functionality**
- Interface Overview
- Additional Products

Copyright © 2006, Juniper Networks, Inc.

#### **T-series Packet Flow and ASIC Functionality**

The following pages examine the function of each T-series ASIC, the T-series switch fabric, and the flow of packets through the T-series PFE. Because the M320 platform is based on T-series ASIC technology, information regarding packet flow through a T-series PFE is also applicable to the M320 platform. The M320 router integrates the discrete functions associated with the various T-series ASICs that are described on subsequent pages into a single I or J chip, depending upon FPC type.

### T-series Packet Forwarding Engine

- Each T-series PFE consists of:
  - One or more media-specific PIC ASIC
    - Handles physical layer signaling, alarms, and CRC processing
  - Layer 2/Layer 3 Packet Processing ASIC
    - Provides link layer encapsulation and decapsulation
    - Manages division and reassembly of packets into J-cells
  - Queuing and Memory Interface ASICs
    - Manage data cell memory buffering
    - Manage notification queuing
  - Internet Processor II ASIC
    - Performs route lookups in forwarding table
  - Switch Interface ASICs
    - Extract route lookup keys
    - Manage cell flow across the switch fabric

Copyright © 2006, Juniper Networks, Inc.


### The T-series Packet Forwarding Engine

The term Packet Forwarding Engine (PFE) is used as a collective noun to describe the collection of components that work together to perform longest-match lookups and packet forwarding using a high-performance, silicon-based switching path. This slide lists the ASICs associated with the T-series PFE and provides a high-level description of the function performed by each ASIC. Subsequent pages delve into the role that each ASIC plays in packet forwarding in greater detail. Note that each T-series FPC provides one (FPC2) or two (FPC3) complete PFE complexes when the FPC is also equipped with one or more PICs:

- *Media-Specific ASIC:* Each PIC type is equipped with one or more ASICs specifically designed to handle the needs of a particular medium. For example, a SONET PIC is equipped with an ASIC that handles SONET framing and alarm generation.
- *Layer2/Layer3 Processing ASIC:* After the PIC performs the medium-specific functions, the bit stream is handed to the Layer2/Layer3 processing ASIC, which removes Layer 2 encapsulation, parses the Layer 3 header, and segments the bit stream into 64-byte chunks.
- *Queuing and Memory Interface ASIC:* The Queuing and Memory Interface ASIC is responsible for writing and reading the 64-byte chunks to the shared memory switch fabric present on each T-series PFE.
- *Internet Processor II ASIC:* The Internet Processor II ASIC performs longest-match route lookups using the information found in the notification cell (the first 64-byte chunk of a Layer 3 packet).
- *Switch Interface ASIC:* The Switch Interface ASICs handle the movement of data between T-series PFEs by facilitating the exchange of 64-byte chunks across the T-series cross-bar switch fabric.

### The T-series Switch Fabric

- Nonblocking topology with any-to-any connectivity
- No single point of failure, all SIBs fully redundant
  - Graceful degradation for multiple failures
    - T640 switch fabric consists of 5 Switch Interface Boards (SIBs) (5<sup>th</sup> is a spare)
    - T320 switch fabric consists of 3 Switch Interface Boards (SIBs) (3<sup>rd</sup> is a spare)
    - M320 fabric consists of four active SIBs
  - Packet order and CoS maintained across fabric



The T320 Switch Fabric  
Copyright © 2006, Juniper Networks, Inc.

#### The T-series Switch Fabric

T-series platforms use a nonblocking cross-bar switch fabric to switch traffic between the system's FPCs. The switch fabric is instantiated by the Switch Interface Board (SIB), which contains the F16 ASIC. SIBs interface to each FPC through high-speed lines (HSLs) that terminate on the SIB's F16 ASIC. The F16 ASIC provides a 16x16 matrix of high-speed input/output lines. Each HSL can support 10 Gbps of half-duplex traffic. By connecting each FPC to two of the F16's HSLs, 10 Gbps of full-duplex capacity (20-Gbps aggregate throughput) is achieved between that FPC and SIB. Each FPC is connected to multiple SIBs to provide the speedup needed for a nonblocking switch fabric and for redundancy reasons.

T-series FPCs interface to the switch fabric over the fabric side (f) of the Switch Interface ASIC; the WAN (w) side of the ASIC interfaces to the Layer2/Layer3 processing ASIC. The Switch Interface ASIC is also called the "N" chip. We use this terminology on the slide to save space.

*Continued on next page.*



## Configuring Juniper Networks Routers

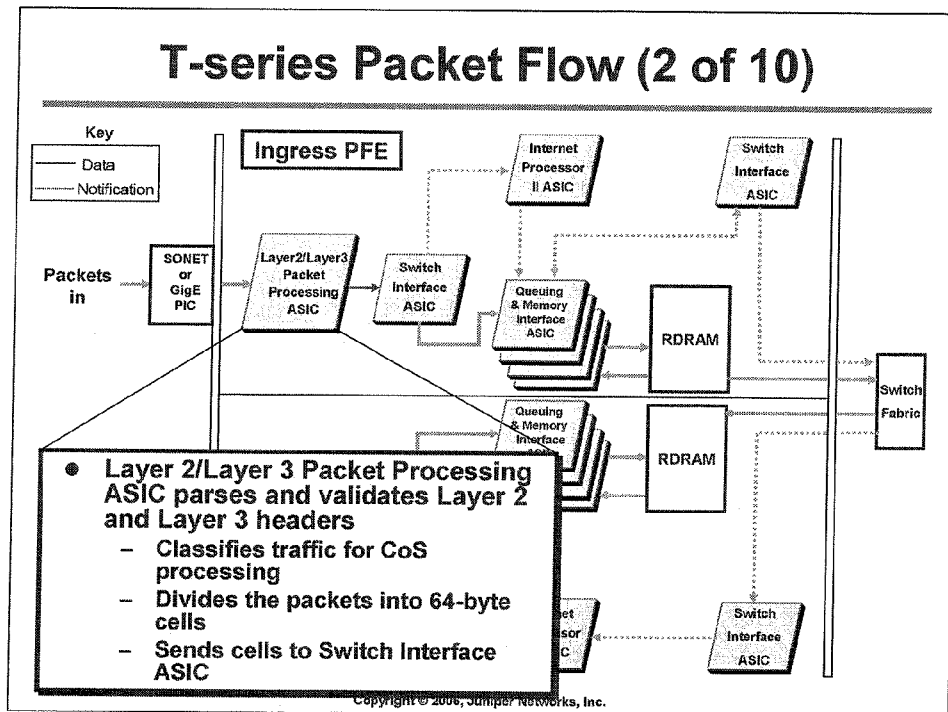
### Redundant Fabric

The graphic illustrates the specifics of a T320's switch fabric. Here, each FPC (or PFE) has four HSL connections to each of SIB 1 and SIB 2. This provides the T320 FPC with 40 Gbps of aggregate capacity. To accommodate SIB failures, each T320 FPC is also connected to a third SIB (SIB 0) using a single HSL. In normal operation, SIB 1 and 2 are active while SIB 0 functions in hot standby mode. SIB 0 automatically becomes active in the event of a SIB 1 or 2 failure. However, because each FPC is interconnected to SIB 0 through a single HSL, switch fabric speedup is reduced. The reduction in speedup results in a graceful degradation of the T320's switch fabric that might result in some packet loss. The T640 platform makes use of five SIBs in a similar configuration, with the exception that all FPCs are attached to all SIBs using two HSLs. The result is that a T640's switch fabric remains nonblocking despite the presence of a SIB failure. Multiple SIB failures results in graceful degradation of the T640 switch fabric capacity.

Note that the M320 platform is similar to the T320 in that the failure of one of its four SIBs results in graceful degradation of forwarding capacity.



## Configuring Juniper Networks Routers



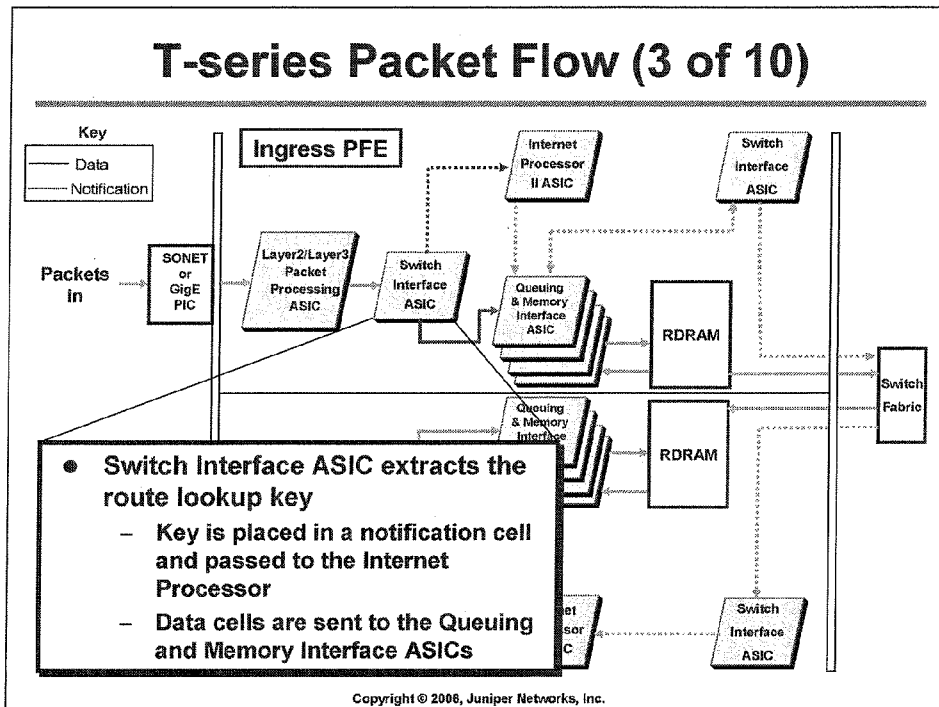
### T-series Packet Flow: Part 2

The Layer 2/Layer 3 Packet Processing ASIC performs Layer 2 and Layer 3 parsing. The Layer 2/Layer 3 ASIC also divides the packets into 64-byte chunks called *J-cells*. The J-cells are sent to the Switch Interface ASIC.

Errors detected during the Layer 2/Layer 3 parsing steps or when the Layer 2/Layer 3 processing ASIC receives an indication from the PIC that the received frame is corrupt results in error counter increments and an effective *no-op* flag for any J-cells relating to the corrupted frame still housed in shared memory.

The Layer 2/Layer 3 Processing ASIC also performs behavior aggregate (BA) traffic classification to associate traffic with a forwarding class for egress queuing and scheduling operations. Examples of BA classification include IP precedence and DiffServ code points.

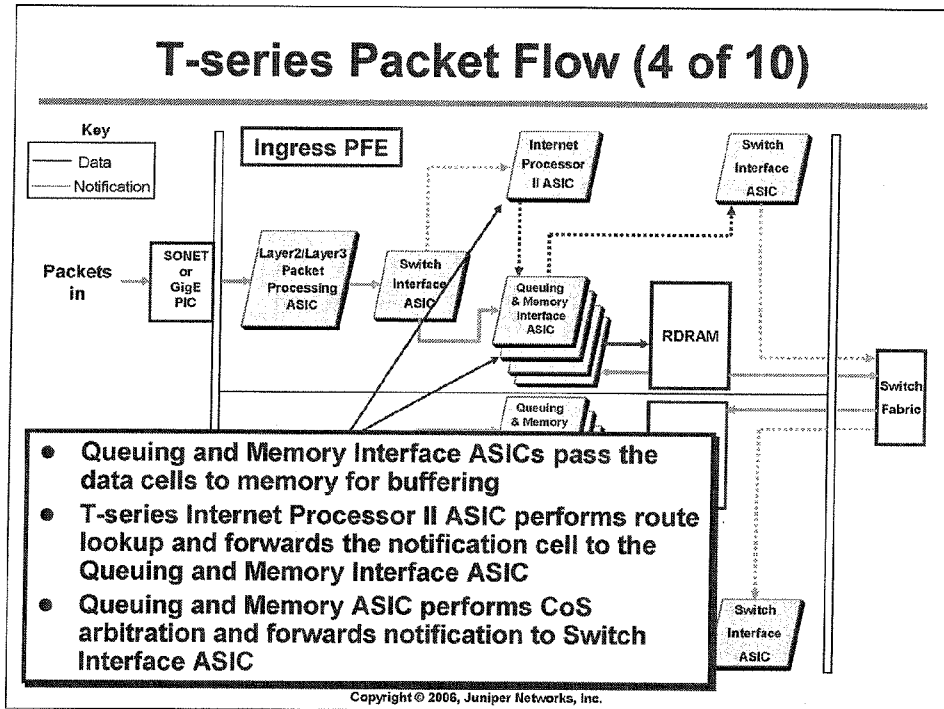
## Configuring Juniper Networks Routers



### T-series Packet Flow: Part 3

The Switch Interface ASIC extracts the route lookup key (comprised of the first 64 bytes of data in the Layer 3 packet), places it in a notification cell, and passes the notification to the T-series Internet Processor. The Switch Interface ASIC then passes the remaining data cells to the Queuing and Memory Interface ASICs. These ASICs manage the shared memory switch fabric associated with each T-series PFE. Note that the shared memory fabric facilitates the switching of packets within a specific PFE complex, such as occurs when the source and destination PICs share a PFE.

## Configuring Juniper Networks Routers



### T-series Packet Flow: Part 4

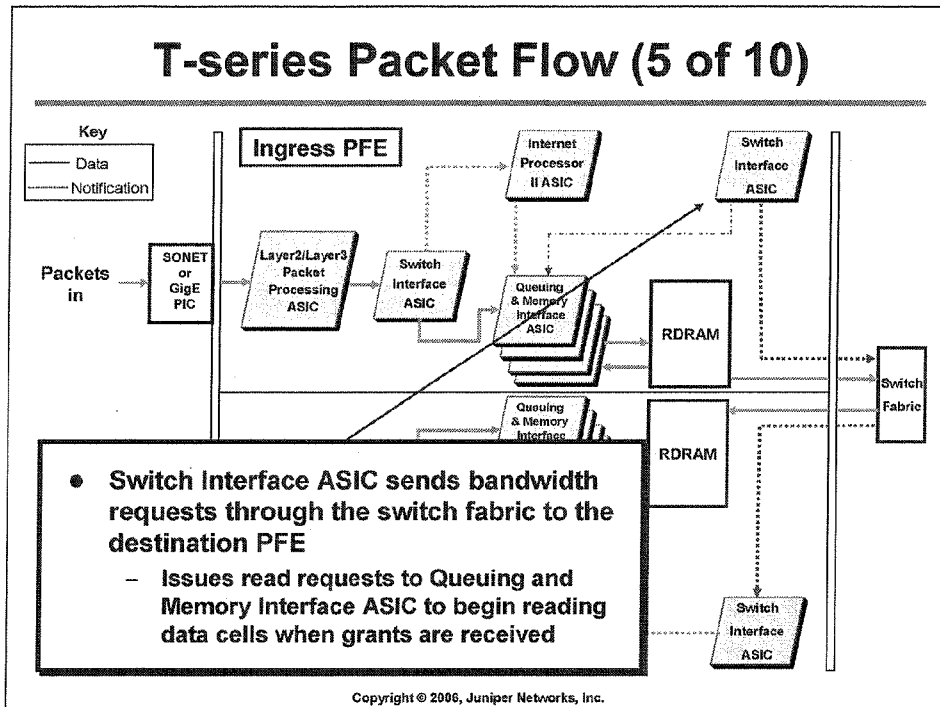
The Queuing and Memory Interface ASICs pass the received J-cells to the PFE's memory for buffering in the shared memory fabric within the PFE. Note that the cross-bar switch fabric is only used to exchange packets *between* PFE complexes.

While the J-cells are being written into shared memory, the T-series Internet Processor II ASIC performs a route lookup operation on the key data. The modified notification cell is then forwarded to the Queuing and Memory Interface ASIC.

In addition to queuing notifications, the Queuing and Memory Management ASIC also performs the following class-of-service (CoS) functions:

- Selection of notifications from the head of each queue for transmission to Switch Interface ASIC according to the priority level of each queue.
- Random early detection (RED): If a queue begins to fill up, it is desirable to randomly drop some packets from the queue before it is completely full. The drop probability is programmable, and this process is part of the TCP congestion control mechanism.

## Configuring Juniper Networks Routers

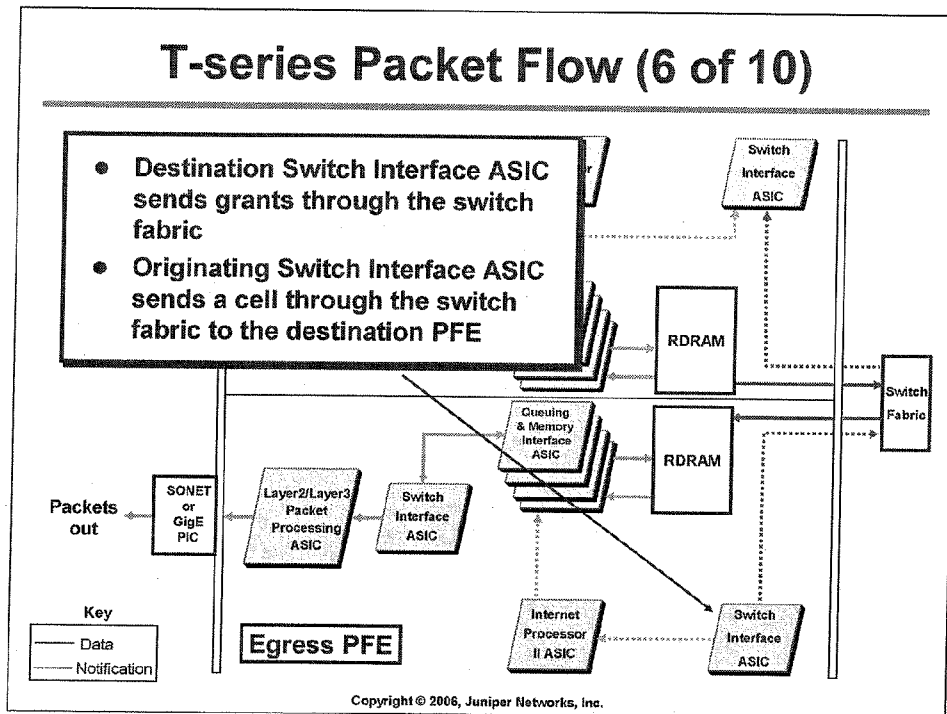


### T-series Packet Flow: Part 5

At this stage of the packet's processing, the Queuing and Memory Interface ASIC sends the notification cell to the Switch Interface ASIC that faces the switch fabric, unless the destination is a port on the same Packet Forwarding Engine. In this case, the notification is sent to the Switch Interface ASIC that faces the Layer 2/Layer 3 Processing ASIC. Packets exchanged between ports on a common PFE do not transit the switch fabric.

The Switch Interface ASIC sends bandwidth requests through the switch fabric to the destination PFE for those destinations that reside on another PFE. The Switch Interface ASIC also issues read requests to the Queuing and Memory Interface ASIC to begin reading data cells out of memory when the egress PFE (and the switch fabric) indicates it is ready to handle a given J-cell.

## Configuring Juniper Networks Routers

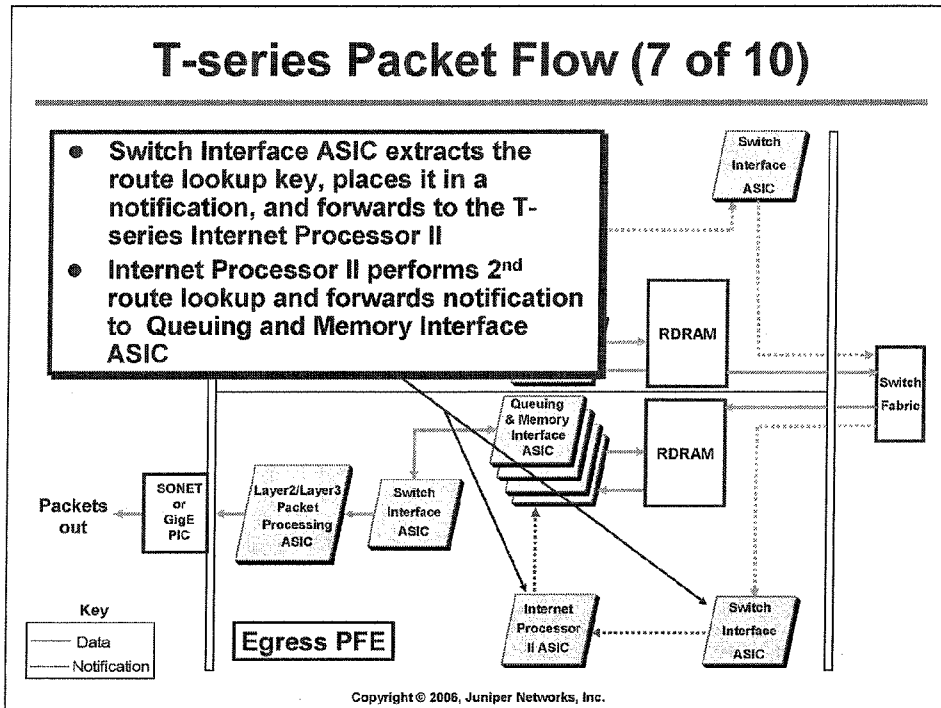


### T-series Packet Flow: Part 6

The destination Switch Interface ASIC returns bandwidth grants through the switch fabric to the originating Switch Interface ASIC in response to received bandwidth requests.

Upon receipt of each bandwidth grant, the originating Switch Interface ASIC sends a cell through the switch fabric to the destination PFE.

## Configuring Juniper Networks Routers

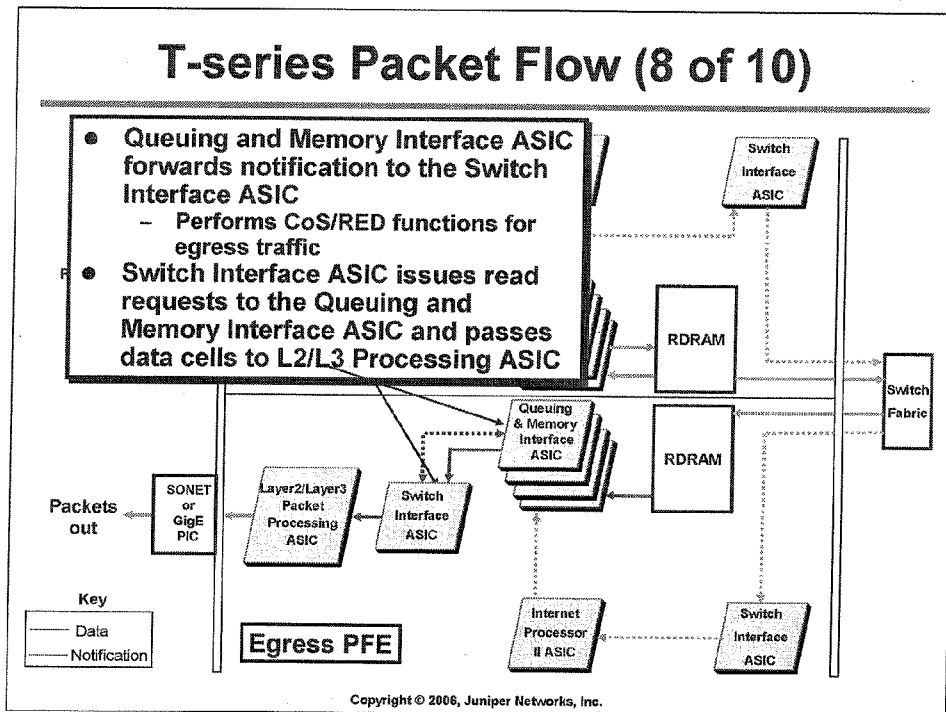


### T-series Packet Flow: Part 7

The destination Switch Interface ASIC receives cells from the switch fabric. Once again, the Switch Interface ASIC extracts the route lookup key and forwards the notification cell to that PFE's Internet Processor II ASIC for a second longest-match lookup operation. Note that this notification cell is modified to reflect the new memory locations for the related J-cells, owing to the fact that the memory locations for each chunk varies by PFE.

The T-series Internet Processor II ASIC in the destination PFE performs a second route lookup and forwards the notification to the Queuing and Memory Interface ASIC. The results of this route lookup include details regarding the egress PFE PIC, port, and required encapsulation.





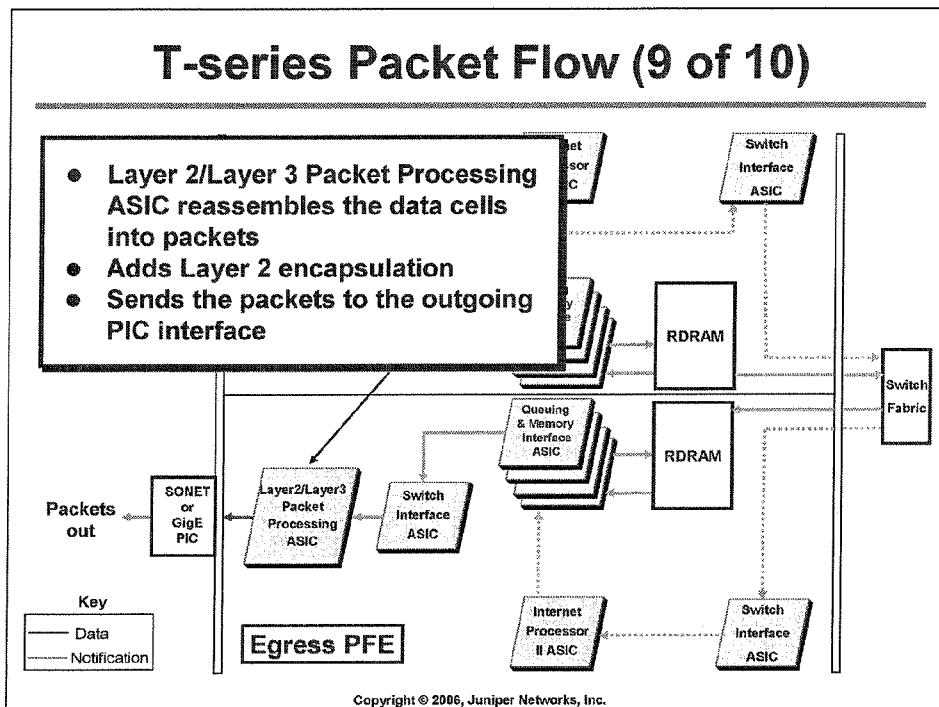
### T-series Packet Flow: Part 8

The Queuing and Memory Interface ASIC is responsible for performing class-of-service (CoS) functions for egress traffic. These CoS functions include the selection of notification cells from the head of each queue for submission to the Switch Interface ASIC, random early detection (RED) related discards, and policing and rate shaping.

The Queuing and Memory Management ASIC forwards a modified notification cell (the cell now includes next-hop information) to the Switch Interface ASIC when the CoS algorithms dictate that a given packet should now be serviced.

The Switch Interface ASIC sends read requests to the Queuing and Memory Interface ASIC to read the data cells out of memory and passes the cells to the Layer 2/Layer 3 Packet Processing ASIC.

## Configuring Juniper Networks Routers



### T-series Packet Flow: Part 9

The Layer 2/Layer 3 Packet Processing ASIC reassembles the data cells into packets. The Layer 2/Layer 3 processing ASIC then adds appropriate Layer 2 encapsulation and sends the resulting bit stream to the egress PIC.

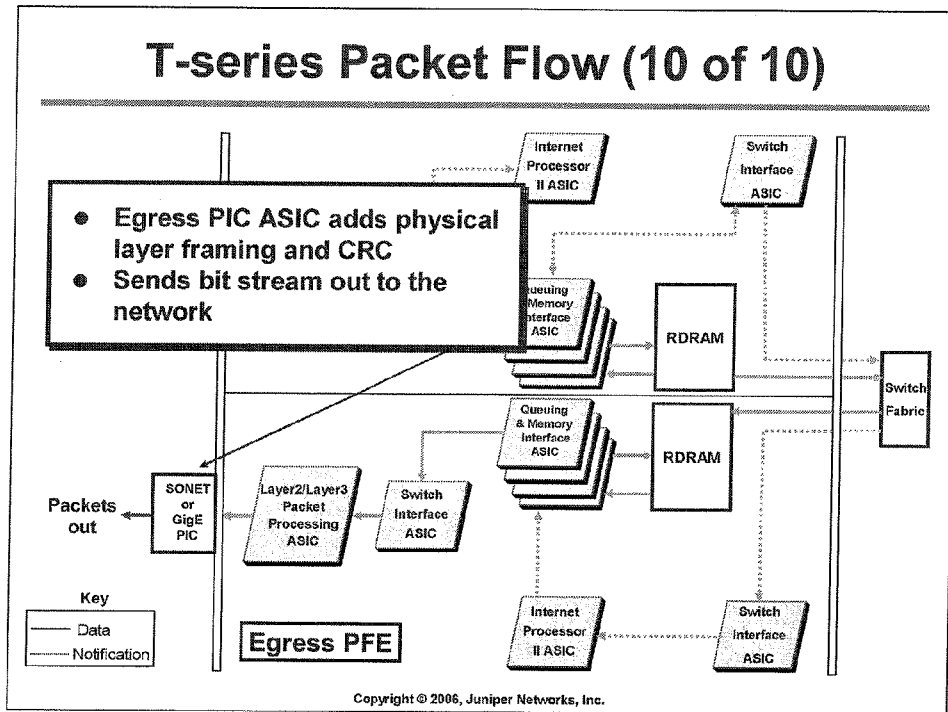
The Layer 2/Layer 3 Processing ASIC is responsible for CoS-related queuing, scheduling, and congestion avoidance operations at packet egress. Each output port on a given PIC is associated with four forwarding classes (or queues). You configure schedulers to provide each forwarding class with some share of the port's bandwidth.

Note that traffic classification, which associates traffic with one of the defined forwarding classes, occurs at the ingress FPC. Once so identified at ingress, the traffic is handled in accordance with the parameters configured for that traffic class by the Layer 2/Layer 3 Processing ASIC on the egress FPC. The random early detection (RED) algorithm is implemented by the Layer 2/Layer 3 Processing ASIC during egress processing to avoid tail drops and the resulting risk of global synchronization of TCP retransmissions.

The T-series Switch Interface ASICs handle switch fabric queuing and prioritization to extend CoS across the T-series switch fabric.

A full coverage of JUNOS software CoS capabilities is beyond the scope of this class.

## Configuring Juniper Networks Routers



### T-series Packet Flow: Part 10

The final steps in egress packet processing come into play when the egress PIC sends the packet out into the network with the appropriate physical layer signaling and medium-specific framing. The egress PIC also calculates and adds a CRC to the frame as needed for each particular medium.

### Exception Packets

- **Exception packets**
  - Local delivery
  - IP options
    - Source route, router alert, etc.
  - ICMP message generation
- **Generally processed by Packet Forwarding Engine control CPU**
  - Remaining traffic (local and control) sent to Routing Engine via internal link
    - Rate limiting
    - Hardware-based WRR ensures control traffic is not starved

Copyright © 2006, Juniper Networks, Inc.

### Exception Packets

Exception packets require some form of special handling. Examples of exception traffic include:

- Packets addressed to the chassis, such as routing protocol updates, Telnet sessions, pings, traceroutes, and replies to traffic sourced from the RE.
- IP packets with the IP options field. Options in the packet's IP header are rarely seen, but the Packet Forwarding Engine was purposely designed *not* to handle IP options. They must be sent to the Routing Engine for processing.
- Traffic that requires the generation of ICMP messages. ICMP messages are sent to the packet's source to report various error conditions and to respond to ping requests. Examples of ICMP errors include destination unreachable messages, which are sent when there is no entry in the forwarding table for the packet's destination address, or time-to-live (TTL) expired messages, which are sent when a packet's TTL is decremented to zero. In most cases, the PFE process handles the generation of ICMP messages.

### Packet Forwarding Engine CPU

The Internet Processor II ASIC passes exception packets to the microprocessor on the Packet Forwarding Engine Control Board, which in turn processes almost all of them. Certain exception packets are also sent to the Routing Engine for further processing. Exception traffic destined for the Routing Engine is sent over the 100 Mbps `fxp1` interface. Exception traffic is rate-limited by the PowerPC processor to protect the Routing Engine from denial-of-service attacks. During times of congestion, the router gives preference to the local and control traffic, with the latter being afforded a minimum of 5% of the `fxp1` interface's bandwidth through hardware-based weighted round-robin (WRR) queuing.

### **Agenda: Product Line Overview**

---

- M-series and T-series Overview
- Installation and Handling Guidelines
- Platform Architecture and Hardware Overview
- M-series Packet Flow and ASIC Functionality
- T-series Packet Flow and ASIC Functionality
- **Interface Overview**
- Additional Products

Copyright © 2006, Juniper Networks, Inc.

#### **Interface Overview**

The items covered in the following pages include:

- *Permanent and transient interfaces:* Juniper Networks M-series and T-series platforms carry the full range of network interfaces and also support built-in interfaces with dedicated functionality.
- *Interface names:* Interfaces are named by type and location in the chassis.
- *Interface media types:* Interfaces are named according to their media or service type.
- *Interface placement:* Transient interfaces are referenced according to their location in an FPC, which is in turn installed into a slot in a chassis.

### Permanent Interfaces

- Router has several permanent interfaces
  - Out-of-band management interface is called `fxp0`
    - Requires configuration
  - Internal Routing Engine-to-Packet Forwarding Engine connection is called `fxp1/bcm0`
  - Internal RE-to-RE connection is `fxp2` or `em0`
    - Internal interfaces do not require any configuration; do not attempt to modify these interfaces!

Copyright © 2006, Juniper Networks, Inc.

### Permanent Interfaces

Each Juniper Networks M-series and T-series platform has several permanent interfaces. One—the management Ethernet interface—provides an out-of-band method for connecting to the router. You can connect to the management interface over a network using utilities such as SSH and Telnet, and SNMP can also use the management interface to gather statistics from the router. The `fxp0` management interface requires configuration to operate.


A second permanent interface provides internal Ethernet-based connectivity between the Routing Engine (RE) and the Packet Forwarding Engine (PFE). This interface is named `fxp1` on most platforms; In the case of the M320, the interface operates at 1 Gbps and is called `bcm0`. A proprietary protocol known as the Trivial Network Protocol (TNP) operates over these interfaces. Note that traffic arriving on the out-of-band interface is not permitted to egress on a PFE port, and vice versa. This design is intended to isolate the out-of-band network from transient traffic and to preserve bandwidth on the internal `fxp1` or `bcm0` control interface.

On platforms with redundant REs, a third interface (`fxp2` or `em0`) is used to interconnect the REs for heart-beat and internal communications exchanges. Internal interfaces like `fxp1` and `fxp2` do not require any configuration. You should not attempt to configure or modify these interfaces.

### Transient Interfaces

- PICs support transient interfaces
  - PICs plug into FPCs
  - FPC plugs into chassis
- Transient interfaces are named according to:
  - Interface media type
  - FPC slot number
  - PIC slot number within FPC
  - PIC port number
  - Channel number where applicable
- Example:

`ct1-5/2/3:27` = A channelized DS3 interface in FPC slot 5, PIC position 2, port 3, time-slot 27



Copyright © 2006, Juniper Networks, Inc.

### Transient Interfaces

Each FPC can support from two to four PICs, depending on the platform. PICs provide the actual physical interfaces to the network. These physical interfaces are the router's transient interfaces. We refer to them to as transient because you can hot-swap FPCs and PICs on most platforms at any time. From the point of view of the Packet Forwarding Engine, you can place any FPC into any slot, and you can generally place any combination of PICs in any location on an FPC. This characteristics makes PFE interfaces transient.

You can use a transient interface for networking or to provide hardware-assisted services within the router. Service examples include IPSec, stateful firewall, and GRE tunneling. You must configure each of the transient interfaces based on which slot the FPC is installed, in which location in the FPC the PIC is installed, and to which port you are connecting.

### Transient Interface Naming

JUNOS software uses a standard naming convention when naming interfaces. You must configure each of the standard interfaces based on the slot in which the FPC is installed, the location in which the PIC is installed, and for some PICs, the port to which you are connecting. When dealing with a channelized PIC you must also reference the correct channel/time-slot value using a `:n` form of syntax.

### Interface Naming Example

The slide shows an example of a channelized interface name that makes use of the `:` delimiter to identify a timeslot within the channel bundle.

### Logical Units

- **Logical units are like sub-interfaces in other equipment**
  - Except in JUNOS software, a logical unit is always required
    - Also used to support multipoint technologies like Frame Relay, ATM, or VLANs
- **Interface unit number is separate in meaning from the actual circuit identifier; can be any arbitrary value**
  - Suggested convention is to keep them the same
- **PPP/HDLC encapsulations support only one logical unit**
  - Must configure unit number as zero for these encapsulations
- **Multiple protocol addresses are supported on a single logical unit**
  - Typing in additional addresses does not override previous address
    - Watch for multiple addresses when correcting addressing mistakes!

Copyright © 2006, Juniper Networks, Inc.

### Logical Interfaces

Each physical interface descriptor can contain one or more logical interface descriptors. These descriptors allow you to map one or more logical (sometimes called virtual) interfaces to a single physical device. Creating multiple logical interfaces is useful for ATM and Frame Relay networks, where you can associate multiple virtual circuits or data link layer connections with a single physical interface.

### Circuit Identifier versus Unit Number

The unit number and the circuit identifier are different in meaning. The circuit identifier identifies the logical tunnel or circuit, while the unit is used to identify a logical partition of the physical interface.

Although not required, it is generally considered best practice to keep the unit number and circuit identifier the same. This practice can greatly aid in troubleshooting when you have many logical circuits.

### Point-to-Point Encapsulations

PPP and Cisco HDLC encapsulations support only a single logical interface, and its logical unit number must be zero. Frame Relay and ATM encapsulations support multiple logical interfaces, so you can configure one or more logical unit numbers.

### Addressing Issues

A Juniper Networks M-series or T-series platform can have more than one address on a single logical interface. Issuing a second `set` command does not overwrite the previous address but simply adds to that address. Use of the CLI's `rename` command is an excellent way to correct addressing mistakes.

Also note that JUNOS software will form IGP adjacencies over all logical interfaces when the IGP is configured on these interfaces; this behavior is worth noting because some vendors form an adjacency *only* over the primary address of an interface.



### Interface Media Types

- **Common media types:**
  - **at:** ATM-over-SONET/SDH ports
  - **e1:** E1 ports
  - **e3:** E3 ports
  - **fe:** Fast Ethernet ports
  - **so:** SONET/SDH ports
  - **t1:** T1 ports
  - **t3:** DS-3 ports
  - **ge:** Gigabit Ethernet ports
  - **ae:** Aggregated Ethernet ports
- **Various IP services and internally generated interface types**
  - **No ports are associated with IP services or internally generated interfaces**
    - Service PIC examples include Adaptive Services and encryption PIC
    - Internally generated interfaces include `tap`, `pime`, `pimd`, and `gre`

Copyright © 2006, Juniper Networks, Inc.

### Interface Media Types

The slide shows the list of interface media types.

### IP Services PICs

IP Services PICs provide value-added services such as tunneling or the management of multilink bundles. IP Services PICs do not have ports or media associated with them, but they do have two-letter interface type designations as shown below. Actual coverage of the services provided by these PICs is beyond the scope of this class.

- **es:** Encryption interface;
- **gr:** Generic route encapsulation tunnel interface;
- **ip:** IP-over-IP encapsulation tunnel interface;
- **ls:** Link services interface;
- **ml:** Multilink interface;
- **mo:** Passive monitoring interface;
- **mt:** Multicast tunnel interface;
- **sp:** Adaptive services interfaces; and
- **vt:** Virtual loopback tunnel interface.

Internally generated and nonconfigurable interfaces include:

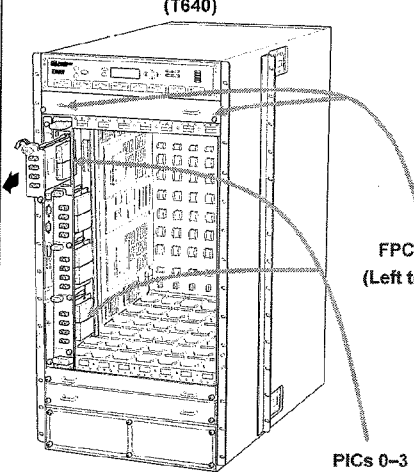
- `gre`;
- `mtun`;
- `ipip`;
- `tap`; and
- `pime/pimd`.

In most cases these interfaces require a corresponding services PIC to operate. The `tap` interface is no longer operational, but it continues to be supported in BSD.

## Configuring Juniper Networks Routers

### Typical FPC and PIC Placement

Typical FPC and PIC Numbering (T640)



- Transient interfaces identified according to FPC/PIC/port convention
- FPC and PIC numbering varies by platform
  - M40/M160 platforms support eight FPCs, numbered from left to right
    - PICs numbered from top to bottom (0-3)
  - M20 platform supports four FPCs numbered from top to bottom
    - PICs numbered from right to left (0-3)
- FPC slot and PIC port numbers are labeled!

Copyright © 2006, Juniper Networks, Inc.

### Identifying Transient Interfaces

Transient interfaces are identified by the interface's FPC slot number, the PIC slot number, and the PIC's physical port number in the form of *media-type-fpc-slot/pic-slot/port-number*. Channelized interfaces identify a particular subchannel with the addition of a suffix in the form of *:sub-channel-number*. A logical unit (aka, a subinterface) is identified with a suffix in the form of *.logical-interface-number*.

### FPC and PIC Slot Numbering Varies

The FPC and PIC slot numbering varies by platform due to some platforms using vertically aligned FPC slots while other platforms use a horizontal FPC arrangement. The slide details the differences in FPC and PIC slot numbering for the M160/M160 platforms versus the M20 platform. The graphic shows typical FPC and PIC numbering in the content of the T640 platform, which makes use of vertically aligned FPCs.

### The Upside?

The upside to this story is that each platform has labels that clearly identify the FPC slot number and PIC number. Further, each PIC has a label to identify the number associated with that PIC's physical ports.

## **Agenda: Product Line Overview**

---

- **M-series and T-series Overview**
- **Installation and Handling Guidelines**
- **Platform Architecture and Hardware Overview**
- **M-series Packet Flow and ASIC Functionality**
- **T-series Packet Flow and ASIC Functionality**
- **Interface Overview**
- **Additional Products**

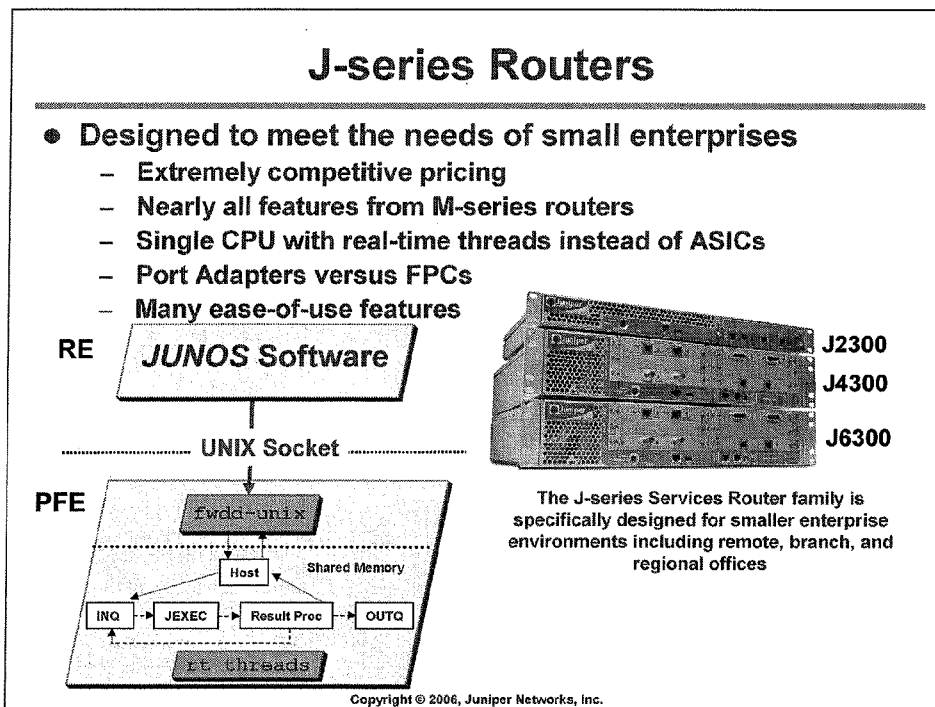
Copyright © 2006, Juniper Networks, Inc.

### **Interface Overview**

The items covered in the following pages include:

- J-Series router family for small offices and small businesses.
- M120 Multiservice Edge Router.
- TX Matrix for multi-chassis T-640 routing systems.
- Additional Routing Engine options.

## Configuring Juniper Networks Routers



The J-series Services Router family is specifically designed for smaller enterprise environments including remote, branch, and regional offices

### J-series Routers

Juniper Networks designed the J-series router family to leverage proven design principles and industry-leading JUNOS software to bring deterministic forwarding performance and carrier-class stability to the small enterprises and remote offices. With these routers, Juniper Networks continues the practice of separation of forwarding and control, but now with the economical innovation of implementing this separation entirely in software. J-series routers use a real-time operating system to simultaneously implement JUNOS software, packet forwarding, and advanced services, each in separate real-time threads. This practice enables these routers to perform each of these functions with guaranteed and deterministic performance, with no risk of one function degrading the performance of any other function.


J-series routers use *Port Interface Modules (PIMs)* instead of FPCs. A PIM is like an FPC with built-in PICs of a common media type. Numerous low-speed WAN interfaces exist, such as ISDN, xDSL and serial, as well as LAN and high-speed uplink interfaces, such as DS-3. All J-series routers include two embedded Fast Ethernet ports and an embedded software implementation of the Juniper Networks Adaptive Services PIC. Each port adapter contains Intel Interface Processors that offload some of the burden of packet processing from the CPU. As a result, when you add interfaces, you are also proportionally adding the necessary processing power to maintain the router's performance capabilities with the additional connectivity.

With the J-series platforms, Juniper Networks introduced the J-Web GUI. In addition to the standard JUNOS software CLI, you can now configure and operate the router entirely from a Web browser. J-Web also contains a number of Quick Configuration utilities, which speed the process of setting up basic or common functions. Juniper Networks has since expanded J-Web to run on M-series and T-series platforms and continues the practice of producing all software for all these products from a single train of source code.

### M120 Multiservice Edge Router

#### ● Features:

- 4 FPCs, each capable of:
  - 4 Type 1 PICs (4 Gbps)
  - 4 Type 2 PICs (16 Gbps)
  - 1 Type 3 PIC (10 Gbps)
- 2 cFPCs
  - Integrated PIC+FPC
  - OC192
  - 10 Gigabit Ethernet
- 1/4 rack
- Independent FEB/FPC Mapping
- 32,768 logical interfaces



Copyright © 2006, Juniper Networks, Inc.

#### Features

The M120 Multiservice Edge Router delivers support for 128 Gigabit Ethernet subscriber ports, with 10-GB Ethernet or OC 192 uplink capability in an affordable, compact form factor. The M120 hardware is designed to support logical interfaces numbering well beyond the current JUNOS software upper limit of 32,768.

The M120 supports four types of FPCs: up to four Type 1, Type 2, or Type 3 FPCs, and up to two compact FPCs (cFPCs). A cFPC is a combination of a PIC and an FPC. It contains the interface circuitry and the FPC as a single assembly in a space saving form factor that is about the same size as a PIC. Two cFPCs are available for the M120: a 10-Gigabit Ethernet cFPC and an OC192 cFPC. Both of these cFPCs provide receptacles for XFP optical transceivers. The 10-Gigabit Ethernet cFPC offers a WAN PHY mode that enables interconnection directly to SONET transport facilities, eliminating the need for a separate SONET interface device.

The M120 allows you to configure any Flexible PIC Concentrator (FPC) to use any Forwarding Engine Board (FEB). This configuration allows you to use a single FEB for a pair of Type 1 FPCs. It also allows you to provision for redundant FEBs. Type 2 and Type 3 FPCs and cFPCs require a dedicated FEB. The router supports up to 16 Type 1 or Type 2 PICs or four Type 3 (10-G) PICs.

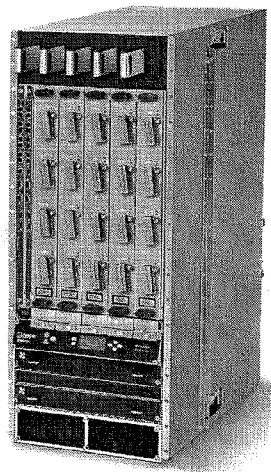
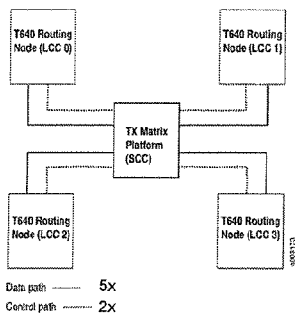
#### FPC specifications:

- FPC1: Rated at 4 Gbps full duplex, supports up to four PICs.
- FPC2: Rated at 16 Gbps full duplex, supports up to four PICs.
- FPC3: Rated at 20 Gbps full duplex, supports one PIC.
- cFPC: Rated at 20 Gbps full duplex, contains embedded PIC.

## Configuring Juniper Networks Routers

### TX Matrix

- **Multinode routing matrix**
  - Up to 4 T640s act as one router
  - 2.56 Tbps
  - 3 Gpps
  - Degradation-free failover
- **7 Interconnections per T640**



Copyright © 2006, Juniper Networks, Inc.

### Matrix Details

The TX Matrix platform is the centralized switch fabric of a T640 routing matrix, which is a terabit routing system interconnecting up to four T640 routing nodes to deliver up to 2.56 terabits per second (Tbps) of subscriber switching capacity. The routing matrix multichassis architecture provides scalable growth for aggregation and core services for voice, video, and data networks. A full T640 routing matrix can process 3 Gpps (billion packets per second) for the vast majority of probable packet size mixtures. It has degradation-free switch-fabric failover, made possible by both the T640s and the TX having a 25% switching capacity reserve in the form of an extra (fifth) switch plane. Five TX-SIPs connect to five T640 SIPs on each T640, but only four are required for maximum performance. A fully loaded T640 routing matrix consumes approximately 2.6 racks, weighs approximately 1.25 tons, and consumes slightly over 30 kW of power.

### Interconnection

The T640 routing nodes in a matrix are called Line Card Chassis (LCCs). The TX Matrix is called a Switch Card Chassis (SCC). The LCCs each connect to the SCC a total of seven times: A VCSEL (pronounced *VIX-e*) fiber-optic cable array connects each of the five switch fabrics, and a pair of UTP Ethernet cables redundantly connects the Control Boards. These connections form both a forwarding matrix and a control matrix. Both matrices are fully redundant.

## Routing Engines

---

- **Three Routing Engines:**
  - **RE-850:**
    - M7i/M10i upgrade
    - 850-MHz Celeron
    - 1.5-GB SDRAM
    - 20-GB hard drive
    - 256-MB compact flash drive
  - **RE-A-1000:**
    - Base RE for M120
    - 1.0-GHz Processor
    - 2-GB DRAM
  - **RE-A-2000:**
    - M120 upgrade

Copyright © 2009, Juniper Networks, Inc.

### Additional Routing Engines

Juniper Networks recently introduced three new routing engines. The RE-850 offers upgraded performance and capacity for the M7i/M10i platforms. The RE-A-1000 is the standard RE for the M120 platform, and the RE-A-2000 is a 2-GHz version of the same RE that is optionally available.

### Review Questions

---

1. How do you safely power down an M-series or T-series platform?
2. What are the primary responsibilities of the Routing Engine and the Packet Forwarding Engine?
3. List and describe at least four FRUs associated with typical M-series and T-series platforms
4. How can the Craft Interface assist in troubleshooting?
5. Describe packet flow through Juniper Networks M-series and T-series platforms
6. Describe the general procedure for hot-swapping FPCs and PICs
7. Describe what each field means in the interface name `at-0/1/1.100`

Copyright © 2006, Juniper Networks, Inc.

#### This Module Discussed:

- Juniper Networks, Inc. M-series, T-series, and J-series products and their typical applications;
- General platform architecture;
- The function of major router components;
- Operation of the Craft Interface;
- Packet flow through a T-series and M-series platform and;
- Interface naming conventions and the role of logical units.





**Configuring Juniper Networks Routers**

***Appendix B: Platform Troubleshooting***

Copyright © 2006, Juniper Networks, Inc.

CJNR-M-7.a.7.6.1

### Module Objectives

---

- After successfully completing this module, you will be able to:
  - Troubleshoot M-series and T-series hardware using the CLI
  - Troubleshoot M-series and T-series platforms using status indicators

Copyright © 2006, Juniper Networks, Inc.

#### **This Module Discusses:**

- M-series and T-series platform overview;
- Major router components;
- RE redundancy options;
- General maintenance issues; and
- Packet flow through an M-series router.

## Configuring Juniper Networks Routers

### Displaying Chassis Inventory

```
{master}
user@t640> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               54970          T640
Midplane      REV 05   710-002726   AX5084
FPM GBUS      REV 08   710-002901   HF5013
FPM Display   REV 04   710-002897   HF5248
CIP           REV 06   710-002895   HG0718
PEM 0         Rev 04   740-002595   ML12519       Power Entry Module
PEM 1         Rev 04   740-002595   ML12517       Power Entry Module
SCC 0         REV 09   710-003423   HF9311
SCC 1         REV 09   710-003423   HF9302
Routing Engine 0 REV 01   740-005022   210865700267  RE-3.0
Routing Engine 1 REV 01   740-005022   210865700264  RE-3.0
CB 0          REV 10   710-002728   HF9619
CB 1          REV 10   710-002728   HF9629
FPC 1         REV 05   710-007529   HL7538        FPC Type 3
CPU          REV 14   710-001726   HG2750
MMB 0         REV 02   710-005555   HL7476        MMB-288mbit
MMB 1         REV 02   710-005555   HL7126        MMB-288mbit
PPB 0         REV 04   710-002845   HJ7134        PPB Type 3
PPB 1         REV 04   710-002845   HJ7001        PPB Type 3
. . .
SPMB 0        REV 03   710-003229   HF5060
SPMB 1        REV 03   710-003229   HF5045
SIB 0         REV 01   750-005486   HG9926        SIB-I8-F16
SIB 1         REV 01   750-005486   HF7675        SIB-I8-F16
SIB 2         REV 01   750-005486   HF7734        SIB-I8-F16
SIB 3         REV 01   750-005486   HF7736        SIB-I8-F16
SIB 4         REV 01   750-005486   HG9299        SIB-I8-F16
Copyright © 2006, Juniper Networks, Inc.
```

### Displaying Chassis Inventory

The output of the `show chassis hardware` command displays the hardware components installed in the router. This command is useful when troubleshooting or upgrading your router. The (edited) sample shown is taken from a T640 platform with redundant REs.

The `show chassis hardware` command output fields are:

- **Item:** For the chassis component, information appears about the backplane, the power supplies, the maxicab (the connection between the Routing Engine and the backplane), the SCB, and each of the FPCs and their PICs.
- **Version:** Displays the revision level of the chassis component.
- **Part number:** Displays the part number of the chassis component.
- **Serial number:** Displays the serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis.
- **Description:** For the power supplies, it displays the type of supply; for the PICs, it displays the type of PIC.

Support for a `show chassis hardware frus` command started with JUNOS software Release 5.6. This command's output provides information in a format suitable for inventory and sparring purposes. Sample output is provided here:

```
user@host> show chassis hardware frus
Hardware inventory:
Item          Part number  Assem  Conn  Fiber  Driver  Description
Chassis                               0x0503  --   --   --      M20
Midplane      710-001517  0x0113  --   --   --
Power Supply A 740-001465  0x0404  --   --   --      AC
. . .
```

### Displaying Alarm Conditions

---

```
user@host> show chassis alarms
1 alarm is currently active

Alarm time                Class   Description
2003-05-09 21:30:07 UTC   Major  Power Supply B not providing power
```

Copyright © 2006, Juniper Networks, Inc.

#### Listing Alarm Conditions

The `show chassis alarms` command lists all of the alarm conditions that currently exist in the router. You can disable some alarms; however, you cannot disable safety-related and chassis component alarms.

Pressing the alarm cutoff (ACO) button located on the Craft Interface manually silences the alarm to an external device connected to the alarm relay, but it does not remove the alarm messages from the display (if present on the router) nor extinguish the alarm LEDs. In addition, new alarms that occur after silencing an external device reactivate the external device.

To clear all user-defined messages displayed on the M40 and M160 routers only, issue the `clear chassis craft-interface display` command after resolving causes of alarm.

The `show chassis alarms` output fields are:

- Alarm time: Displays the date and time the alarm was first recorded;
- Class: Displays the severity class for this alarm (it can be minor or major);  
and
- Description: Displays the information about the alarm.

### Displaying Environmental Information

```
{master}
user@t640> show chassis environment
Class Item                Status      Measurement
Temp  PEM 0                   OK          27 degrees C / 80 degrees F
      PEM 1                   OK          27 degrees C / 80 degrees F
      SCG 0                   OK          35 degrees C / 95 degrees F
      SCG 1                   OK          34 degrees C / 93 degrees F
      Routing Engine 0       OK          31 degrees C / 87 degrees F
      Routing Engine 1       OK          30 degrees C / 86 degrees F
      CB 0                    OK          34 degrees C / 93 degrees F
      CB 1                    OK          36 degrees C / 96 degrees F
      SIB 0                   OK          38 degrees C / 100 degrees F
      SIB 1                   OK          38 degrees C / 100 degrees F
      SIB 2                   OK          38 degrees C / 100 degrees F
      SIB 3                   OK          39 degrees C / 102 degrees F
      SIB 4                   OK          39 degrees C / 102 degrees F
      FPC 1 Top               Testing
      FPC 1 Bottom           Testing
      FPC 3 Top               OK          43 degrees C / 109 degrees F
      FPC 3 Bottom           OK          30 degrees C / 86 degrees F
      FPC 5 Top               OK          43 degrees C / 109 degrees F
      FPC 5 Bottom           OK          30 degrees C / 86 degrees F
      FPM GBUS                OK          28 degrees C / 82 degrees F
      FPM Display             OK          31 degrees C / 87 degrees F
Fans  Top Left Front fan     OK          Spinning at noxmal speed
      . . .
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying Environmental Information

The `show chassis environment` command displays environmental information about the router chassis, including the temperature and information about the fans, power supplies, and Routing Engine. The truncated example is taken from a T640 platform.

The output fields are:

- **Power:** Displays information about each power supply. Status can be OK, Testing (during initial power-on), Failed, or Absent. For the M160 and T-series platforms, information is displayed about the Power Entry Modules (PEM). Status can be OK, Testing (during initial power-on), Failed, or Absent.
- **Temp:** Displays the temperature of air flowing through the chassis.
- **Fans:** Displays information about the fans. Status can be OK, Testing (during initial power-on), Failed, or Absent. Measurement indicates if the fans are spinning at normal or high speed.
- **Other:** Depending upon the platform, various other fields might be present. For example, on the M160 router the `Misc` entry includes CIP status (Connector Interface Panel). In this case, OK indicates that the CIP is present. For T-series platforms, the display includes information on the SCG, CB, SIBs, and the Switch Processor Mezzanine Board (SPMB) and the CIP. These fields are not shown in the sample on the slide.

## Configuring Juniper Networks Routers

### Displaying Contents of LCD Display

```
{master}
user@t640> show chassis craft-interface
FPM Display contents:
+-----+
|daemon0
|Up: 21+04:07
|
|Temperature OK
+-----+

Front Panel System LEDs:
Routing Engine  0  1
-----
OK              *  *
Fail            .  .
Master          *  .

Front Panel Alarm Indicators:
-----
Red LED        .
Yellow LED     .
Major relay    .
Minor relay    .
. . .

. . .

Front Panel FPC LEDs:
FPC  0  1  2  3  4  5  6  7
-----
Red   .  .  .  .  .  .  .  .
Green .  .  .  *  .  *  .  .

CB LEDs:
CB  0  1
-----
Amber .  .
Green *  *
Blue  *  .

SCG LEDs:
SCG 0  1
-----
Amber .  .
Green *  *
Blue  *  .

SIB LEDs:
SIB 0  1  2  3  4
-----
Red   .  .  .  .  .
Green *  *  *  *  *
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying Contents of LCD Display

The `show chassis craft-interface` shows all current messages. The capture shown was taken from a T640 system. Output fields include:

- `FPM Display contents`: Displays contents of the Front Panel Module display.
- `router-name`: Shows the name of the router.
- `Up`: Shows how long the router has been operational in days, hours, minutes, and seconds.
- `message`: Displays information about the router traffic load, the power supply status, the fan status, and the temperature status. The display of this information changes every 2 seconds.
- `Front Panel System LEDs`: Displays status of the Front Panel System LEDs. A dot (.) indicates the LED is not lit. An asterisk (\*) indicates the LED is lit.
- `Front Panel Alarm Indicators`: Displays status of the Front Panel Alarm Indicators. A dot indicates the relay is off. An asterisk indicates the relay is active.
- `Front Panel FPC LEDs`: Displays status of the Front Panel FPC LEDs. A dot indicates the LED is not lit. An asterisk indicates the LED is lit.
- `MCS, SFM, SCG, CB, and SIB LEDs`: Displays status of the MCS, SCG, CB, SFM, and SIB LEDs as supported by a given platform. A dot indicates the LED is not lit. An asterisk indicates the LED is lit. When neither a dot nor an asterisk is displayed, no board is in that slot.

### System Controller Board Status

- `show chassis (feb|scb|cfcb|sfm slot|ssb slot)`

- Displays information about M-series system controller boards (FEB, SCB, SFM, or SSB)

High CPU utilization levels normally indicate a high volume of exception traffic

```
user@host> show chassis feb
FEB status:
  Temperature                26 degrees C / 78 degrees F
  CPU utilization             1 percent
  Interrupt utilization       0 percent
  Heap utilization           16 percent
  Buffer utilization          44 percent
  Total CPU DRAM             64 Mbytes
  Internet Processor II      Version 1, Foundry IBM, Part number 9
  Start time:                2001-05-09 17:31:52 UTC
  Uptime:                    27 days, 1 hour, 25 minutes, 9 seconds
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying System Controller Board Status

The `show chassis (feb | scb | cfcb | sfm slot | ssb slot)` command displays information about the system controller boards (either FEB, CFEB, SCB, SFM, or SSB). The output fields are:

- **Temperature:** Displays the temperature of the air passing by the controller, in degrees Centigrade.
- **CPU utilization:** Displays the total percentage of CPU being used by the controller's processor.
- **Interrupt utilization:** Of the total CPU being used by the controller's processor, displays the percentage being used for interrupts.
- **Heap utilization:** Displays the percentage of heap space being used by the controller's processor.
- **Buffer utilization:** Displays the percentage of buffer space being used by the controller's processor.
- **DRAM:** Displays the total DRAM available to the controller's processor.
- **Start time:** Displays the time when the controller started running.
- **Uptime:** Displays how long the controller has been running.

Because the controller board's CPU is not involved in actual packet forwarding, you should expect to see a very low level of CPU utilization in most cases. Aside from hardware and environmental monitoring, the controller board CPU is primarily used when processing exception traffic. This traffic tends to take the form of PFE-generated ICMP error messages or traffic that must be directed to the host RE for processing. Examples of exception traffic include ICMP echo exchanged, packets with expired TTL/IP options, or traffic that is being sampled and/or counted as part of firewall filter.

### Displaying FPC Status

#### Display status of all FPCs

```
{master}
user@t640> show chassis fpc
```

Slot	State	Temp	CPU Utilization (%)		Memory	Utilization (%)	
		(C)	Total	Interrupt	DRAM (MB)	Heap	Buffer
0	Empty						
1	Present	0	0	0	0	0	0
2	Empty						
3	Online	30	3	0	256	14	41
4	Empty						
5	Online	30	3	0	256	14	41
6	Empty						
7	Empty						

Copyright © 2006, Juniper Networks, Inc.

#### Displaying FPC Status

The `show chassis fpc` command displays the status of the installed FPCs. The sample is taken from T640 platform. The output fields are:

- Slot: Displays the FPC slot number.
- State: Displays the state of the FPC.
- Temp (C): Displays the temperature of the air passing by the FPC, in degrees Centigrade.
- CPU Utilization (%): Displays the total percentage of CPU used by the FPC's processor.
- Interrupt CPU Utilization (%): Of the total CPU used by the FPC's processor, displays the percentage being used for interrupts.
- Memory DRAM: Displays the total DRAM available to the FPC's processor, in megabytes.
- Heap Utilization (%): Displays the percentage of heap space (dynamic memory) used by the FPC's processor. If this number exceeds 80%, it might indicate a software problem (that is, a memory leak).
- Buffer Utilization (%): Displays the percentage of buffer space used by the FPC's processor for buffering internal messages.



### Displaying Status for a Specific FPC

#### Display detailed information for a specific FPC

```
user@host> show chassis fpc detail 0
Slot 0 information:
  State                               Online
  Logical slot                         0
  Temperature                          26 degrees C / 78 degrees F
  Total CPU DRAM                       8 Mbytes
  Total SRAM                           1 Mbytes
  Total SDRAM                          128 Mbytes
  Total notification SDRAM             24 Mbytes
  I/O Manager ASIC information         Version 2.0, Foundry IEM, Part number 0
  Start time:                         2001-05-29 21:25:02 UTC
  Uptime:                              8 days, 1 hour, 39 minutes, 19 seconds
```

Copyright © 2006, Juniper Networks, Inc.

#### Displaying a Specific FPC's Status

The `show chassis fpc detail` command shows detailed information about the FPCs installed in the system. Adding an FPC number, as in the example on the slide, limits the output to the specified FPC. The output fields are:

- State: Displays the state of the FPC slot, which include:
  - Dead: Indicates that the slot is held in reset because of errors;
  - Diag: Indicates that the slot is being ignored while the FPC is running diagnostics;
  - Dormant: Indicates that the slot is held in reset;
  - Empty: Indicates that no FPC is present;
  - Online: Indicates FPC is online and running;
  - Probed: Indicates that probe is complete and awaiting PFE restart; and
  - Probe-wait: Indicates that the slot is waiting to be probed.
- Logical slot: Displays the slot number.
- Temperature: Displays the temperature of the air passing by the FPC, in degrees Centigrade.
- Total CPU DRAM: Displays the amount of DRAM available to the FPC's CPU.
- Total SRAM: Displays the amount of SRAM used by the FPC's CPU
- Total SDRAM: Displays the total amount of memory used for storing packets and notifications.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Displaying a Specific FPC's Status (contd.)

- Total notification SDRAM: Displays the amount of memory used by the Packet Forwarding Engine for packet buffer and packet notification space.
- I/O Manager ASIC information: For the I/O Manager, identifies the version number, manufacturer, and part number.
- Start time: Displays the time when the Routing Engine noticed that the FPC was running.
- Uptime: Displays how long the Routing Engine has been connected to the FPC and how long the FPC has been up and running.

A sample of the detailed FPC output from a T-series platform is provided here:

```
{master}
user@t640> show chassis fpc detail 1
Slot 1 information:
State                               Present
Temperature                         0 degrees C / 32 degrees F
Total CPU DRAM                       0 MB
Total SRAM                           56 MB
Total SDRAM                          1280 MB
```

### Displaying PIC Status

Display information for all PICs or for PICs in a particular slot

```
user@host> show chassis fpc pic-status
Slot 0 Online
  PIC 0    4x F/E, 100 BASE-TX
  PIC 1    4x OC-3 SONET, MM
  PIC 3    1x Tunnel

user@host> show chassis pic fpc-slot 0 pic-slot 1
PIC fpc slot 0 pic slot 1 information:
Type                4x OC-3 SONET, MM
ASIC type           D chip
State               Online
PIC version         1.4
Uptime              16 days, 1 hour, 51 minutes, 3 seconds
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying PIC Status

The `show chassis fpc pic-status` command displays information for all PICs. The output fields are:

- **State:** Indicates the state of the FPC slot, which can be:
  - **Dead:** Indicates that the FPC is held in reset because of errors;
  - **Diag:** Indicates that the slot is being ignored while the FPC is running diagnostics;
  - **Dormant:** Indicates that the slot is held in reset;
  - **Empty:** Indicates that no FPC is present;
  - **Online:** Indicates that the FPC is online and running;
  - **Probed:** Indicates that the probe is complete and that the slot is awaiting PFE restart; and
  - **Probe-wait:** Indicates that the slot is waiting to be probed.
- **PIC type:** Displays the type of PIC at each PIC location and the number of ports on the PIC.

To display details about a specific PIC, use the `show chassis pic fpc-slot fpc slot number pic-slot pic slot number` command.

## Configuring Juniper Networks Routers

### Displaying Routing Engine Status

```
{master}
user@t640> show chassis routing-engine
Routing Engine status:
Slot 0:
Current state           Master
Election priority       Master
Temperature             31 degrees C / 87 degrees F
DRAM                   2048 MB
Memory utilization      8 percent
CPU utilization:
  User                  0 percent
  Background            0 percent
  Kernel                2 percent
  Interrupt             0 percent
  Idle                  97 percent
Model                  RE-3.0
Serial ID              210865700267
Start time             2003-08-14 17:29:13 UTC
Uptime                 21 days, 4 hours, 10 minutes, 38 seconds
Load averages:         1 minute  5 minute 15 minute
                       0.00      0.02   0.00
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying Routing Engine Status

The `show chassis routing-engine` command displays information about the Routing Engine. The output fields are:

- Slot: Indicates the slot number for the RE on systems that support RE redundancy.
- Temperature: Displays the temperature of the air flowing past the Routing Engine.
- DRAM: Displays the total DRAM available to the Routing Engine's processor.
- CPU utilization: Displays information about the Routing Engine's CPU utilization, which include:
  - User: Displays the amount of DRAM being used by user processes;
  - Background: Displays the amount of DRAM being used by background processes;
  - Kernel: Displays the amount of DRAM being used by kernel processes;
  - Interrupt: Displays the amount of DRAM being used by interrupt processes; and
  - Idle: Displays the amount of idle DRAM.
- Model: Display the RE model. Current RE models include RE-333 (RE2) RE-600 (RE3), and the RE-400 (RE5).
- Start time: Displays the time at which the Routing Engine started running.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Displaying Routing Engine Status (contd.)

- Start time: Displays the time at which the Routing Engine started running.
- Uptime: Displays how long the Routing Engine has been running.
- Load averages: Displays the Routing Engine load averages for the last 1, 5, and 15 minutes.

For system with redundant REs installed, you can specify the RE slot number or see information about all REs installed in the system, as in this example taken from a T640 platform:

```
{master}
user@t640> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state           Master
    Election priority       Master
    Temperature             31 degrees C / 87 degrees F
    DRAM                    2048 MB
    Memory utilization      8 percent
    CPU utilization:
      User                  0 percent
      Background           0 percent
      Kernel                2 percent
      Interrupt            0 percent
      Idle                  97 percent
    Model                   RE-3.0
    Serial ID               210865700267
    Start time              2003-08-14 17:29:13 UTC
    Uptime                  21 days, 4 hours, 10 minutes, 38 seconds
    Load averages:         1 minute  5 minute  15 minute
                           0.00      0.02     0.00

Routing Engine status:
  Slot 1:
    Current state           Backup
    Election priority       Backup
    Temperature             30 degrees C / 86 degrees F
    DRAM                    2048 MB
    Memory utilization      7 percent
    CPU utilization:
      User                  0 percent
      Background           0 percent
      Kernel                0 percent
      Interrupt            0 percent
      Idle                  100 percent
    Model                   RE-3.0
    Serial ID               210865700264
    Start time              2003-08-14 04:39:59 UTC
    Uptime                  21 days, 16 hours, 59 minutes, 34 seconds
```

Use the `show chassis routing-engine bios` command to display the revision level of the RE's BIOS:

```
user@host> show chassis routing-engine bios
Routing Engine BIOS Version: 0.9
```

### Displaying System Processes

Displays information and status of all system software processes

```
user@host> show system processes detail
PID  UID  PPID CPU PRI NI  RSS WCHAN  STARTED TT  STAT  TIME COMMAND
8695  0   3207  1  28  0  464 -      9:40AM ??  R    0:00.00 /bin/ps -a
1     0   0     0  0  10  0  520 wait   Thu12PM ??  ILs  0:00.45 /sbin/prei
2     0   0     0  0 -18  0  0  psleep Thu12PM ??  DL   0:00.36 (pagedaem
3     0   0     0  0  18  0  0  psleep Thu12PM ??  DL   0:00.00 (vmdaemon
4     0   0     0  0 -18  0  0  psleep Thu12PM ??  DL   0:00.92 (bufdaemo
5     0   0     0  0  18  0  0  syncer Thu12PM ??  DL   0:08.09 (syncer)
6     0   0     0  0  28  0  0  sleep  Thu12PM ??  DL   0:00.00 (netdaemo
7     0   0     0  0  2  0  0  pfeasel Thu12PM ??  IL   0:00.00 (if_pfe)
8     0   0     0  0  2  0  0  pfeacc  Thu12PM ??  IL   0:00.00 (if_pfe_l
9     0   0     0  0  2  0  0  cb-pol  Thu12PM ??  IL   0:00.00 (cb_poll)
10    0   0     0  0 -18  0  0  psleep  Thu12PM ??  DL   0:01.31 (vmuncach
11    0   0     0  0  2  0  0  scs_ho  Thu12PM ??  IL   0:00.00 (scs_hous
12    0   0     0  0  2  0  0  picacc  Thu12PM ??  IL   0:00.00 (if_pic_l
139   0   1     0  0  10  0  50600 mfsidl  Thu12PM ??  ILs  0:00.91 mfs -o noa
...
[additional process listings]
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying System Processes

The `show system processes extensive` command displays information and the status of all system software processes and software process information. The `show system processes` commands are equivalent to various UNIX `ps` and `top` commands. Options to the `show system processes` command include:

- `none`: Displays information about system processes. This option is equivalent to the UNIX `ps -ax` command.
- `brief (optional)`: Displays brief system process information, listing no processes themselves. This option is equivalent to the UNIX `top -bSd1 0` command.
- `detail (optional)`: Displays detailed system process information. This option is equivalent to the UNIX `ps -ax -uid, ppid, cpu, pri, nice, rss, wchan, start` command.
- `extensive (optional)`: Displays exhaustive system process information. This option is equivalent to the UNIX `top -bSd1 infinity` command.
- `summary (optional)`: Displays a summary of active system process information. This option is equivalent to the beginning of the UNIX `top` command.
- `wide (optional)`: Displays process information that might be wider than 80 columns. This option is equivalent to the UNIX `ps -ax -ww` command.

Please see the JUNOS software Operational Mode Command Reference for a description of the various fields reported by this command.

### View System Logs

```
user@host> show system boot-messages
Copyright (c) 1996-2001, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2001 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 6.0R1.3 #0: 2003-07-24 02:20:37 UTC
    builder@cecrops.juniper.net:/build/cecrops-d/6.0R1.3/obj-
1386/sys/compile/JUNIPER
Timecounter "i8254" frequency 1193182 Hz
```

```
. . .

sio0 at port 0x3f8-0x3ff irq 4 flags 0x90 on isa0
sio0: type 16550A, console
sio1 at port 0x3e8-0x3ef irq 5 on isa0
sio1: type 16550A
sio2 at port 0x2f8-0x2ff irq 3 on isa0
sio2: type 16550A
sio3: configured irq 7 not in bitmap of probed irqs 0
fxp0: Ethernet address 00:a0:a5:12:40:32
fxp1: Ethernet address 00:a0:a5:12:40:35
DEVFS: ready to run
ad0: 91MB <SanDisk SDCFB-96> [734/8/32] at ata0-master using PIO1
ad1: 19077MB <IBM-DJSA-220> [38760/16/63] at ata0-slave using UDMA33
Mounting root from ufs:/dev/ad0sla
```

Copyright © 2008, Juniper Networks, Inc.

### Viewing System Log Files

You can find a wealth of invaluable information in the various log files maintained by JUNOS software and the platform's syslog daemon. The slide shows an example of an operator displaying the contents of the boot log by issuing a `show system boot-messages` command. This file is written during system boot and contains the various boot-up messages generated during the last power cycle/boot or reboot.

General system login and tracing was covered in the main body of this volume. Here, the intent is to simply remind you that numerous log files are maintained by JUNOS software and can be viewed from the CLI using a `show log file-name` command. Most system log files end in the letter `d` to signify that some daemon is responsible for the contents of the file. A partial listing of the daemon log files that often prove useful when troubleshooting is provided here; note that the main system log file is called `messages`:

- `apsd`: Automatic Protection System (APS) daemon log file.
- `chassisd`: Chassis management daemon log file.
- `cosd`: Class of Service Daemon log file.
- `dcd`: Device Control Daemon log file. Manages interface devices.
- `ilmid`: Interim Local Management Interface daemon log file.
- `sampled`: Sampling process log file.
- `snmpd`: SNMP log file.
- `vrripd`: Virtual Router Redundancy Protocol daemon log file.

## Configuring Juniper Networks Routers

### Displaying System Statistics

#### Displays system-wide, protocol-related statistics

```
user@host> show system statistics
ip:
  99197 total packets received
  0 bad header checksums
  0 with size smaller than minimum
  0 with data size < data length
  0 with header length < data size
  0 with data length < header length
  0 with incorrect version number
  0 packets destined to dead next hop
  0 fragments received
  0 fragments dropped (dup or out of space)
  0 fragments dropped (queue overflow)
  0 fragments dropped after timeout
  0 fragments dropped due to over limit
  0 packets reassembled ok
  99197 packets for this host
  0 packets for unknown/unsupported protocol
  0 packets forwarded
  0 packets not forwardable
  0 redirects sent
  170220 packets sent from this host
  8375 packets sent with fabricated ip header
  0 output packets dropped due to no bufs
  0 output packets discarded due to no route
  0 output datagrams fragmented
  0 fragments created
  0 datagrams that can't be fragmented
  0 packets with bad options
  11165 packets with options handled without error
  0 strict source and record route options
  0 loose source and record route options
  0 record route options
  0 timestamp options
  0 timestamp and address options
  0 timestamp and prespecified address options
  0 option packets dropped due to rate limit
  11165 router alert options
  0 multicast packets dropped (no iflist)
icmp:
  0 drops due to rate limit
  . . . .
[additional protocol information]
tcp:
udp:
igmp:
arp:
ip6:
icmp6:
cni1:
esis:
tnp:
rdp:
tudp:
ttp:
mpls:
vpls:
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying System Statistics

The `show system statistics` command displays system-wide, protocol-related statistics. You can view statistics for a specific protocol by adding that protocol's keyword as an argument to the `show system statistics` command. You can obtain information about packet drops occurring within the PFE with a `show pfe statistics traffic` command:

```
user@host> show pfe statistics traffic
PFE Traffic statistics:
  344099 packets input (0 packets/sec)
  239208 packets output (0 packets/sec)
PFE Local Traffic statistics:
  344098 local packets input
  239208 local packets output
  0 software input high drops
  0 software input medium drops
  0 software input low drops
  0 software output drops
  0 hardware input drops
PFE Local Protocol statistics:
  0 hdlc keepalives
  0 atm oam
  0 fr lmi
  68411 ppp lcp/ncp
  77919 ospf hello
  0 rsvp hello
  0 isis iih
PFE Hardware Discard statistics:
  0 timeout
  0 truncated key
  . . . .
```



### Displaying System Storage

Displays amount of storage available on flash and rotating disks

```
user@host> show system storage
Filesystem      512-blocks      Used      Avail Capacity Mounted on
/dev/ad0s1a      158174          56438      89084      39% /
devfs            32              32          0          100% /dev/
/dev/vn0         18316           18316       0          100% /packages/mnt/jbase
devfs            32              32          0          100% /dev/
/dev/vn1         45448           45448       0          100% /packages/mnt/jkernel-6.0R1.3
/dev/vn2         22720           22720       0          100% /packages/mnt/jpfe-M40-6.0R1.3
/dev/vn3         3580            3580        0          100% /packages/mnt/jdocs-6.0R1.3
/dev/vn4         20728           20728       0          100% /packages/mnt/jroute-6.0R1.3
/dev/vn5         9252            9252        0          100% /packages/mnt/jcrypto-6.0R1.3
mfs:139         3048670         12          2804766    0% /tmp
/dev/ad0s1e     23742           14          21830     0% /config
procfs          8               8            0          100% /proc
/dev/ad1s1f     34635886        1746492     30118524   5% /var
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying System Storage

The `show system storage` command displays the amount of storage available on the flash and rotating disks. This command displays statistics about the amount of free disk space in the router's file systems. Values are displayed in 512-byte blocks. This command is equivalent to the UNIX `df` command.

### Displaying System Connections

#### Lists active IP sockets on the Routing Engine

```
user@host> show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address         (state)
ip4    0      0      *.*                    *.*                     *
ip4    0      0      *.*                    *.*                     *
ip4    0      0      *.*                    *.*                     *
ip4    0      0      *.*                    *.*                     *
ip4    0      0      *.*                    *.*                     *
ip4    2728  0      *.*                    *.*                     *
tcp4   0      0      192.168.0.1.2222       192.168.5.1.179       ESTABLISHED
tcp46  0      0      *.179                  *.*                     LISTEN
tcp4   0      0      *.179                  *.*                     LISTEN
tcp4   0      0      *.23                   *.*                     LISTEN
tcp4   0      0      *.22                   *.*                     LISTEN
tcp4   0      0      *.21                   *.*                     LISTEN
tcp4   0      0      192.168.0.1.23        10.0.0.2.2093        ESTABLISHED
tcp4   0      0      *.6153                 *.*                     LISTEN
tcp4   0      0      *.666                  *.*                     LISTEN
tcp4   0      0      *.31340                *.*                     LISTEN
tcp4   0      0      *.31341                *.*                     LISTEN
. . .
```

Copyright © 2006, Juniper Networks, Inc.

#### Displaying System Connections

The `show system connections` command lists the active IP sockets on the Routing Engine. Use this command to verify which services are active on a system and what connections are currently in progress. The output fields are:

- **Protocol:** Displays the protocol of the socket. It can be either TCP or UDP.
- **Recv-Q:** Displays the number of input packets received by the protocol and waiting to be processed by the application.
- **Send-Q:** Displays the number of output packets sent by the application and waiting to be processed by the protocol.
- **Local Address:** Displays the local address and port of the socket, separated by a dot (.). An asterisk (\*) indicates that the bound address is the wildcard address. Server sockets typically have the wildcard address and a well-known port bound to them.
- **Foreign Address:** Displays the foreign address and port of the socket, separated by a dot. An asterisk indicates that the address or port is a wildcard.
- **(state):** For TCP, displays the protocol state of the socket.

### Mapping Port Numbers to Services

- What service is associated with that port number?
  - Display `/etc/services`; pipe to match to simplify your search

```
lab@San_Jose-3> file show /etc/services
#
# Network services, Internet style
#
# Note that it is presently the policy of IANA to assign a single well-known
# port number for both TCP and UDP; hence, most entries here have two entries
# even if the protocol doesn't support UDP operations.
# Updated from RFC 1700, "Assigned Numbers" (October 1994). All ports
# are included.
#
# The latest IANA port assignments can be gotten from
# http://www.isi.edu/in-notes/iana/assignments/port-numbers
# The Well Known Ports are those from 0 through 1023.
# The Registered Ports are those from 1024 through 49151
# The Dynamic and/or Private Ports are those from 49152 through 65535
#
# Kerberos services are for Kerberos v4, and are unofficial. Sites running
# v5 should uncomment v5 entries and comment v4 entries.
#
# $FreeBSD: src/etc/services,v 1.62.2.3 2000/10/05 07:37:37 sheldonh Exp $
# From: @(#)services 5.8 (Berkeley) 5/9/91
#
# WELL KNOWN PORT NUMBERS
#
rtmp          1/ddp    #Routing Table Maintenance Protocol
tcpmux       1/tcp    #TCP Port Service Multiplexer
tcpmux       1/udp    #TCP Port Service Multiplexer
nbp          2/ddp    #Name Binding Protocol
compressnet  2/tcp    #Management Utility
...
```

Copyright © 2006, Juniper Networks, Inc.

#### Displaying Port Numbers

The file `show /etc/services` command displays most assigned protocol/port numbers.

On Unix systems, the file `/etc/services` lists most assigned protocol/port numbers. On a Juniper Networks M-series or T-series platform, this list is extensive—comparative to the IANA-assigned ports document.

If you are unsure of the port number seen in the `show system connections` output, you can check this in the `/etc/services` file.

As this information might have security implications for a given organization, the JUNOS software default is to have all services closed, unless specifically opened by configuration.

### Displaying System Uptime

**Displays system and uptime, and the user who committed the current configuration**

```
user@host> show system uptime
Current time: 2003-09-08 11:42:08 UTC
System booted: 2003-09-04 12:14:59 UTC (3d 23:27 ago)
Protocols started: 2003-09-04 12:15:56 UTC (3d 23:26 ago)
Last configured: 2003-09-08 11:08:46 UTC (00:33:22 ago) by lab
11:42AM UTC up 3 days, 23:27, 2 users, load averages: 0.03, 0.01, 0.00
```

Copyright © 2006, Juniper Networks, Inc.

### Displaying Uptime

The `show system uptime` command displays the current time and information about how long the router, router software, and routing protocols have been running. The output fields are:

- `Current time`: Displays the current system time in UTC.
- `System booted`: Displays the date and time when the router was last booted and how long it has been running.
- `Protocols started`: Displays the date and time when the routing protocols were last started and how long they have been running.
- `Last configured`: Displays the date and time when a configuration was last activated (either by booting the router or issuing the `commit` command in configuration mode).
- `Time`: Displays the current time, in the local time zone.
- `Uptime`: Displays how long the router has been operational.
- `users`: Displays the number of users logged into the router.
- `load averages`: Displays the load averages for the last 1 minute, 5 minutes, and 15 minutes.

### Bouncing PICs and FPCs

PICs and FPCs can be soft-booted and taken online/offline from the CLI

```

user@host> request chassis ?
Possible completions:
  fpc          Control operation of an FPC
  pic          Control the state of a PIC
  routing-engine Perform Routing Engine-specific operations
  ssb          Perform System Switch Board-specific operations

user@host> request chassis fpc ?
Possible completions:
  offline      Turn an FPC offline
  online       Turn an FPC online
  restart      Restart an FPC
  slot         FPC slot number (0..3)

user@host> request chassis fpc slot 0 restart
Restart initiated, use "show chassis fpc" to verify
    
```

Copyright © 2006, Juniper Networks, Inc.

### Restarting Hardware Components

You can restart or take online/offline FPC or PICs from the CLI with a `request chassis fpc slot-number [restart | online | offline]` or a `request chassis pic fpc-slot slot-number pic-slot slot-number [restart | online | offline]` command. The example illustrates the basic CLI syntax and an example of the operator restarting FPC 0. In some cases, bouncing an FPC or a problematic PIC can alleviate the need for more drastic actions, such as the dreaded reboot. The captures below show the progression of FPC status from dormant, to probed, to online in the output of a series of `show chassis fpc` commands:

```

user@host> show chassis fpc
Slot State      Temp  CPU Utilization (%)  Memory  Utilization (%)
              (C)  Total  Interrupt  DRAM (MB) Heap  Buffer
0 Dormant       32    0    0    0    0    0
1 Online        30    0    0    8    8    14
2 Empty
3 Empty

user@host> show chassis fpc
Slot State      Temp  CPU Utilization (%)  Memory  Utilization (%)
              (C)  Total  Interrupt  DRAM (MB) Heap  Buffer
0 Probed        32    0    0    0    0    0
1 Online        30    0    0    8    8    14
2 Empty
3 Empty

user@host> show chassis fpc
Slot State      Temp  CPU Utilization (%)  Memory  Utilization (%)
              (C)  Total  Interrupt  DRAM (MB) Heap  Buffer
0 Online        32    2    0    8    11   15
1 Online        30    0    0    8    8    14
2 Empty
3 Empty
    
```

### Restarting Processes

```
user@host> restart ?
Possible completions:
  adaptive-services  Adaptive services process
  audit-process      Audit process
  chassis-control    Chassis management process
  class-of-service   Class-of-service process
  disk-monitoring    Disk monitoring process
  dynamic-flow-capture Dynamic flow capture service
  ecc-error-logging  ECC parity errors logging process
  event-processing   Event processing process
  firewall           Firewall management process
  ggsn-services      GGSN services process
  gracefully         Gracefully restart the process
  immediately        Immediately restart (SIGKILL) the process
  interface-control  Interface process
  ipsec-key-management Key management process
  l2tp-service       Layer 2 Tunneling Protocol process
  lacp               Link Aggregation Control Protocol process
  mib-process        SNMP Management Information Base-II process
  . . .

user@host> restart routing
Routing protocol daemon started, pid 1539
```

Copyright © 2006, Juniper Networks, Inc.

### Restarting Selected Processes

The `restart process` command restarts the specified process. This is similar to issuing a UNIX `kill pid` command. Some of the options include:

- `adaptive-services`: Restarts the adaptive services process, which controls the stateful firewall and NAT functionality of an ASP PIC.
- `class-of-services`: Restarts the CoS daemon (`cosd`).
- `interface-control`: Restarts the interface's process, which controls the router's physical interface devices and logical interfaces.
- `mib-process`: Restarts the MIB II process, which provides the router's MIB II agent.
- `routing`: Restarts the routing protocol process, which controls the routing protocols that run on the router and maintains the routing tables.
- `snmp`: Restarts the SNMP process, which provides the router's SNMP master agent.
- `soft (optional)`: Rereads and reactivates the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. The `soft` option is the equivalent of a UNIX `SIGHUP` signal; omitting this option is the equivalent of a UNIX `SIGTERM` (kill) operation and is the behavior associated with the `gracefully` switch. Using the `immediately` switch restarts the process with the equivalent of a UNIX `SIGKILL` signal.

Note that having to restart a software process is a relatively rare event. Use caution with the `restart` command in a production network!

### Core Files and Memory Dumps

- Technical support engineers deal with two types of core files: **rpd** and **kernel**
  - Reside in `/var/tmp`
  - Transfer to `ftp://ftp.juniper.net/pub/incoming/core file`
    - For example: `case number-core file=1999-0101-001-rpd.core.0`
- Three distinct areas exist where JUNOS software can perform a crash
  - RE (`/var/crash`)
  - PFE (`nvrpm`)
  - Daemons (`/var/tmp`)

Copyright © 2006, Juniper Networks, Inc.

#### Two Types of Core Files

A core file is a snapshot of a program's state when the program is told to dump its state. The core file usually includes information about the program's stack, all variables, and register values at the instance the snapshot is taken. Usually a core file is created when a program encounters a fault while executing. Another way to create a core file is to send a signal to a program while it is running. A core file also is useful in finding out a problem, as long as the problem is not catastrophic.

Another type of core file is created by a kernel. This creation happens only when the kernel encounters a catastrophic fault; in JUNOS software, this type of core file usually is referred to as PANIC.

Core files usually reside in `/var/tmp` directory. Two ways exist to find out if there is any core file in that directory:

- Use the `file list /var/tmp` command from the CLI; and
- If you are in shell, issue the `ls -l/var/tmp` command.

The file name is `rpd.core.number`, where `number` ranges from 0 to 4.

#### Three Areas for a Crash

For core file created by kernel panic, the location will be under `/var/crash`. An indication that a kernel core file is created also exists when the router restarts—you can use the `dmesg shell` command to check for that. Because these core files are created when a program encountered a catastrophic failure, it is important to get these core files back to Juniper Networks for analysis. Usually, you can transfer these files to `ftp://ftp.juniper.net/pub/incoming/core file`. We recommend that the name of the core file is in the following format: `case number-core file`, for example, `1999-0101-001-rpd.core.0`.

*Continued on next page.*

## Configuring Juniper Networks Routers

### Three Areas for a Crash (contd.)

Because core files are highly compressible, we also recommend that you compress these core files with the `gzip` command before transferring the files. Assuming that the router can get to the Juniper Networks FTP server, `ftp.juniper.net`, use the `file copy ...` command. For example, if the core file is `/var/tmp/rpd.core.0`, and your case number is 1999-0101-001, the following is your CLI command: `file copy /var/tmp/rpd.core.0 ftp://ftp.juniper.net/pub/incoming/1999-0101-001-rpd.core.0`.



## Routing Engine Kernel and Cores

---

- Enabled by default
  - Files stored in `var/crash`
  - Two files generated
    - `Kernel.0` → copy of kernel
    - `Vmcore.0` → core of running system
  - Might need to gzip `vmcore` file

Copyright © 2006, Juniper Networks, Inc.

### Enabling Storage

Routing Engine kernel and core storage (that is, dumps) are enabled by default in JUNOS software Release 4.4 and later. This configuration enables `savecore` when the router is booted.

### Routing Engine Daemons and Core Files (1 of 2)

#### ● Core file storage

- Enabled by default
- Core file generated in `var/tmp`
- By default, five core files can be saved for each daemon
  - Modified by `set saved-core files filename` command
- Name → executable name.core.number
  - Rpd daemon → `rpdc.core.0`

```
user@router> file list detail /var/tmp
total 1292622
drwxrwxrwt  3 bin  field          512 Dec 31 06:48 ./
drwxr-xr-x 21 root wheel          512 Mar  5 1999 ../
-rw-rw----  1 root field 119713792 Nov 17 21:58 rpd.core.0
-rw-rw----  1 root field 120782848 Nov 17 22:12 rpd.core.1
```

Copyright © 2006, Juniper Networks, Inc.

#### Core File Storage

Daemon core file storage (that is, dumps) is enabled by default, so no additional CLI knob is needed. If a core file is generated, it is stored in the `var/tmp` directory. By default, you can save the five core files (for example, five for rpd, five for mgd). Once the file is generated, it follows the naming convention of executable name.core.core number, where number can range from 0 to 4. In this example, two core files were generated for the routing protocol daemon (rpd). The first file, `rpdc.core.0`, was created at 21:58, and the second core file, `rpdc.core.1`, was created at 22:12.

## Configuring Juniper Networks Routers

### Routing Engine Daemons and Core Files (2 of 2)

#### Transfer core files to ftp://ftp.juniper.net

```
user@router> start shell
% su
Password:
root@router% cd /var/tmp
root@router% ls -l
total 1292618
-rw-rw---- 1 root field 119713792 Nov 17 21:58 rpd.core.0
-rw-rw---- 1 root field 120782848 Nov 17 22:12 rpd.core.1

root@router% gzip rpd.core.0
root@router% ftp ftp.juniper.net
Connected to colo-ftp.juniper.net.
220 colo-ftp.juniper.net FTP server (Version 6.00LS) ready.
Name (ftp.juniper.net: root):
```

Copyright © 2006, Juniper Networks, Inc.

#### Transferring Core Files to Juniper Networks

To compress the file, which you can do before you transfer it, you must go to the shell level and issue the command (assuming it has the same file name) shown here:

```
user@host> start shell
% gzip /var/tmp/rpd.core.0
```

A file named `/var/tmp/rpd.core.0.gz` will be created.

Note the `.gz` extension denoting that it was compressed by the `gzip` program. After you compress the file you can transfer it using FTP, as shown on the slide. You can also use the `file copy` command with an FTP based URL to copy the file, if you want. An example of this approach, in the context of an anonymous FTP account with a specific upload directory in the form of `/pub/incoming`, is shown here:

```
user@host> file copy test ftp://anonymous:test@ftp.juniper.net/pub/incoming/test
Transferring ftp://anonymous:test@ftp.juniper.net/pu (2356 bytes): 100%
2356 bytes transferred in 0.0 seconds (5.48 MBps)
```

### PFE Microkernel Dumps

---

- Disabled by default, enable with:

```
lab@aurora# set chassis dump-on-panic
```

- Stored in `var/crash`
- SCB/SSB/SFM/FEB info is placed on NVRAM
- M160 FPC is placed in NVRAM of FPC, M20/M40 FPCs are placed in SCB/SSB NVRAM, due to no FPC NVRAM
- Name of core file: `core-pfecomponent.timestamp`
  - `Core-FPC4.100011808032`

Copyright © 2006, Juniper Networks, Inc.

### PFE Microkernel Dumps

Normally, the router does not try to dump a core. You can change this default behavior, however, by issuing the configuration command `set chassis dump-on-panic`. This command enables savecore when the router is booted. The files are stored in `/var/crash`.

## Configuring Juniper Networks Routers

### More on Cores (1 of 2)

#### JTAC might ask for NVRAM or other shell commands

```
user@router> start shell
% su
Password:
root@router% tnpdump
Name      TNPaddr  MAC address      IF      MTU      Expire  Hopcount
master    1        02:00:00:00:00:01  fxp1    1500     0       0
master    1        00:a0:a5:3d:0c:08  fxp2    1500     0       1
host0     4        02:00:00:00:00:01  fxp1    1500     0       0
host0     4        00:a0:a5:3d:0c:08  fxp2    1500     0       1
sfm0      8        02:00:00:00:00:08  fxp1    1500     3       0
sfm1      9        02:00:00:00:00:09  fxp1    1500     4       0
sfm2      10       02:00:00:00:00:0a  fxp1    1500     4       0
sfm3      11       02:00:00:00:00:0b  fxp1    1500     3       0
fpc0      16       02:00:00:00:00:10  fxp1    1500     4       0
fpc3      19       02:00:00:00:00:13  fxp1    1500     4       0
fpc5      21       02:00:00:00:00:15  fxp1    1500     4       0
bcast     255      ff:ff:ff:ff:ff:ff  fxp1    1500     0       0
bcast     255      ff:ff:ff:ff:ff:ff  fxp2    1500     0       1
```

Copyright © 2006, Juniper Networks, Inc.

#### NVRAM or Other Shell Commands

- To VTU or CTU to specific Packet Forwarding Engine components, you might have to perform a tnpdump to get the proper tnp address.

## Configuring Juniper Networks Routers

### More on Cores (2 of 2)

JTAC might ask for `show nvram` or other shell commands, such as `show syslog messages`

```
root@router% vty 8
or
root@router% vty sfm0
SFM platform (266Mhz PPC 603e processor, 64Mb memory, 512Kb flash)
SFM0(router vty)# show nvram
System NVRAM :
 4080 available bytes, 4080 used, 0 free
Contents:
mpc106 machine check caused by error on the PCI Bus
mpc106 error detect register 1: 0x08, 2: 0x00
mpc106 error ack count = 0
.....
SFM0(router vty)# show syslog messages
Oct 26 12:02:05 TL2 tnp_sfm_2 PFEMAN: sent Resync request to Master
Oct 26 12:02:07 TL2 tnp_sfm_3 CM(3): Slot 1: On-line
Oct 26 12:02:07 TL2 tnp_sfm_3 CM(3): Slot 2: On-line
```

Copyright © 2006, Juniper Networks, Inc.

#### Additional Shell Commands

JTAC might ask for certain Packet Forwarding Engine `show` commands. Two common commands are `show nvram` and `show syslog messages`. However, you should issue these commands only with JTAC assistance.

## Reinstalling JUNOS Software from Removable Media

---

- Reinstall JUNOS software if storage media fails or is corrupted
  - Future major software revisions might require a full installation
- Three steps:
  - Prepare to reinstall JUNOS software
  - Reinstall JUNOS software
  - Configure JUNOS software

Copyright © 2006, Juniper Networks, Inc.

### Reinstalling JUNOS Software

You perform a complete software reinstallation when the JUNOS system software on the router becomes damaged for some reason—for example, if the secondary storage on the system fails.

### Three Steps to Complete

The slide shows the three steps you must perform to reinstall the software.

### Reinstallation (1 of 2)

---

- **Preparation:**

- **Record basic information**
  - Router name
  - Management interface IP address and prefix length
  - Default router IP address
  - Domain name and DNS server IP address
- **Copy existing configuration file to a safe place on the network**
  - Located in `/config/juniper.conf`
  - Full installation erases both flash and rotating drives
- **Locate your Juniper Networks installation media**
  - LS-120 floppy or PCMCIA card contains entire JUNOS software distribution

Copyright © 2006, Juniper Networks, Inc.

### Preparing the Router for Reinstallation

Before you install the JUNOS software, you must perform the following steps:

1. If your router is currently running, take a quick inventory of your current network configuration. After you install the software, you must re-enter some of this information. Items to record include:
  - Name of the router;
  - IP address and prefix length information for the `fxp0` interface;
  - IP address of a default router; and
  - IP address of a DNS server.
2. If you want to continue using the existing configuration after you reinstall the software, copy the existing configuration, which is in the file `/config/juniper.conf`, from the router, either to another system or to a floppy disk.
3. Ensure that you have a Juniper Networks installation floppy or PCMCIA card. Contact customer support if you cannot locate the installation media that shipped with your router.



### Reinstallation (2 of 2)

- **Final steps:**

- **Insert installation media into Routing Engine**
  - M40 router—LS-120 floppy
  - All others—PCMCIA flash card
- **Reboot router**
  - Use the CLI from the serial console: `root@lab2> request system halt`
  - Power-cycle router
- **Follow prompts**
  - System reboots automatically after installation completes
- **Restore original configuration**
  - Configure management port and copy original configuration (FTP/SCP)
  - Use `load override terminal` to paste configuration using the console port

Copyright © 2006, Juniper Networks, Inc.

#### Final Steps of Reinstallation

To finish the installation process, perform the following steps:

1. Insert installation media into Routing Engine. The M40 router uses an LS-120 floppy disk; all other routers use PCMCIA flash cards.
2. Reboot the router using the CLI from the serial console:  
`root@lab2> request system halt`
3. Power-cycle the router and follow the prompts. The system reboots automatically after the installation completes.
4. Restore the original configuration. Configure the management port and copy the original configuration (FTP/SCP). Use `load override terminal` to paste the configuration using the console port.

We recommend that you remove the removable media after a reboot, which ensures that reformatting does not occur after an expected reboot.

### Updating Removable Media

---

- **Updating process:**

- FTP the image to the router's `/var/tmp` directory, and insert removable media into the router's drive
- Escape to the shell, and type `su` to become a super user, if not already in the super-user class
- Change to the `/var/tmp` directory, and issue the following commands (example is for PCMCIA media Release 4.x):

```
root@router% dd if=/dev/zero of=/dev/rwd4 count=20
root@router% tar xzOf 4.4R2.3-export.pcm110.tgz | dd
of=/dev/rwd4 bs=64k
```

- See the Juniper Networks Website for LS-120 syntax and Release 5.x

Copyright © 2006, Juniper Networks, Inc.

### Updating the Removable Media

The upgrade process for Release 5.x is different from previous processes. The procedure is located on the Juniper Networks Website, in the same location as the upgrade software.

### Review Questions

---

1. What is the difference between an FPC 1, 2, and 3?
2. Which Juniper Networks platforms support RE redundancy?
3. List the components that comprise the PFE.
4. What is the purpose of and relationship between FPCs and PICs?
5. What is the correct procedure for removing an FPC from a M40 router?
6. Which ASIC is responsible for parsing Layer 2 encapsulation?

Copyright © 2008, Juniper Networks, Inc.

#### This Module Discussed:

- M-series and T-series platform overview;
- Major router components;
- RE redundancy options;
- General maintenance issues; and
- Packet flow through an M-series router.





## **Configuring Juniper Networks Routers**

### ***Appendix C: Additional Features***

Copyright © 2006, Juniper Networks, Inc.

CJNR-M-7.a.7.6.1

### Module Objectives

---

- After successfully completing this module, you will be able to:
  - Describe additional JUNOS software features
  - Determine where to get more information on these features

Copyright © 2006, Juniper Networks, Inc.

#### This Module Discusses:

- Additional JUNOS software features relevant to the material in this course;  
and
- Where to go for more information on these features.

### Additional Features (1 of 4)

---

- **J-Web**

- GUI to configure and operate JUNOS software
- Quick Configuration utilities
- Menu-driven testing and troubleshooting

Copyright © 2006, Juniper Networks, Inc.

#### **J-Web**

J-Web is a graphic user interface (GUI) for configuring and operating JUNOS software-based routers. It was originally developed to assist enterprise customers in operating J-Series enterprise-class routers by enabling IT professionals responsible for a much wider array of network products than an IP backbone engineering group to bypass the need to learn to operate a new CLI. This feature offers a hierarchical display of configuration, quick configuration tools that coordinate multiple configuration statements required to support certain popular configuration tasks, and graphical menu-based operational network monitoring and testing. J-Web has proven very popular in practice, and as a result of extensive customer interest, it has been expanded to function on M-series and T-series routers. J-Web is offered as an optional install package on these routers because customers who do not want this capability are likely to be uncomfortable having an httpd binary installed on their routers that they do not need. J-Web includes weak (56-bit) https encryption by default so that it offers basic security on export versions of JUNOS software. If the j-crypto package is installed on the system, which it is for domestic versions of JUNOS software, it automatically overrides the weak encryption with strong encryption. For more information on J-Web, see the *JUNOS Internet Software Installation and Upgrade* and the *J-Web Interface User Guide*.

Note that unless otherwise specified, you can find all manuals and guides mentioned in this module at <http://www.juniper.net/techpubs>.

### Additional Features (2 of 4)

---

- CLI global replace
  - Search and replace on configuration text values
    - Values and identifiers only, not keywords
  - Uses regular expressions
  - Allows matched pattern substitution
  - Case sensitive only

Copyright © 2006, Juniper Networks, Inc.

#### Global Replace

Although it has always been fairly straightforward to save a JUNOS software configuration as a file and use local shell tools to manipulate its contents, JUNOS software now includes command-line ability to perform global textual search and replace operations. When you are in configuration edit mode, you can use the `replace pattern pattern1 with pattern2 [upto n]` command to perform regular expression-driven text matching and substitution. This command affects identifiers and values only, not keywords. For example, if you move a Fast Ethernet PIC from FPC 2 slot 1 to FPC 1 slot 0, you can replace all instances of `fe-2/1/` with `fe-1/0/`, or if you change the name of a policy, you can simultaneously change its name in every place where you used the policy. For more information, see "Using Global Replace in a Configuration" in the *JUNOS Internet Software CLI User Guide*, available at <http://www.juniper.net/techpubs>.



### Additional Features (3 of 4)

- **PIM register message filtering**
  - On DR send
  - On RP receive
    - RP sends register-stop for filtered messages
- **Advanced RIP timers**
  - Fine-tune RIP performance for specific situations
  - holddown
  - update-interval
  - route-timeout

Copyright © 2006, Juniper Networks, Inc.

#### Register Message Filtering

You can filter which register messages a multicast DR sends and which messages an RP receives at the `[edit pim rp]` hierarchy, using the `rp-register-policy` and `dr-register-policy` keywords. If an RP register policy rejects a particular (S,G), it sends a register-stop message to the DR, just as it would if it had received the message and processed it normally, to indicate to the DR that register messages are no longer necessary. You can read more about these features in the *JUNOS Internet Software Multicast Protocols Configuration Guide*, under "Configuring RP/DR Register Message Filtering."

*Continued on next page.*

## Configuring Juniper Networks Routers

### Advanced RIP Timers

In most networks using Juniper Networks routers, RIP, if used at all, is used in small localized situations to solve specific problems, rather than as a full IGP. This use of RIP is because it scales poorly in large environments. As an example, a service provider might allow a dual-homed customer to run RIP with the two attached provider edge routers because the customer does not have BGP knowledge but still wants to have dynamic routing and failover capability. This RIP environment will have only four routers and typically, a dozen or fewer prefixes shared between them. Ironically, in such a small environment, RIP is capable of performing exceptionally well. To realize this performance with RIP, you can tune several protocol timers, using the `holddown`, `route-timeout` and `update-interval` keywords at the global RIP, group, or interface level. Be careful and perform extensive testing before using these features in a production environment because decreasing these timers can cause instability. Conversely, if you are forced to use RIP in an inappropriately large environment, you can increase these timers to improve stability at the expense of slower failover performance. You can read more about these features under “Configuring RIP Timers” in the *JUNOS Internet Software Routing Protocols Configuration Guide*.

### Additional Features (4 of 4)

---

- **JUNOS software signed binaries**
  - Additional security to deter hackers
  - JUNOS Software only runs binaries with matching signature
- **PIC configuration check**
  - Automatic PIC/FPC distribution validity at upgrade time
  - Manually check logs for problems at PIC insertion time

Copyright © 2006, Juniper Networks, Inc.

#### **Signed Binaries**

JUNOS software is designed to make it extremely difficult for somebody such as a hacker to run foreign software on the router. JUNOS software distributions include a digitally signed software manifest and a digital signature for each included package. JUNOS software does not execute binaries that are not listed in the manifest or whose computed digital signatures do not match those in the manifest. See “JUNOS Software Media and Packages” in the *Installation and Upgrade* manual for more information.

#### **PIC Configuration Check**

With the advent of advanced capability PICs, particularly services PICs and PICs supporting large numbers of queues, some PICs have unusually large embedded software requirements. It is possible for a combination of different such PICs to require more memory for all of these software than some FPCs have available. This requirement causes a microcode (or uCode) overflow, and one or more PICs on the FPC will fail to become operational when installed. The specific requirements vary from one version of JUNOS software to another, as new features change the size of the microcode image for these PICs. You should be aware of two scenarios regarding this issue: first, when upgrading JUNOS software, it is possible that a currently valid PIC combination on an FPC will become invalid after the upgrade; and second, when inserting a PIC, it is possible that the PIC will fail to come online due to exhausted microcode buffer space.

*Continued on next page.*

## Configuring Juniper Networks Routers

### PIC Configuration Check (contd.)

When you attempt to upgrade JUNOS software, the software automatically detects PIC combinations that would become invalid after an upgrade, reports them to you, and refuses to perform the upgrade. When you insert a PIC, you can determine whether it is causing a microcode overflow by monitoring the syslog for the phrase uCode overflow, as follows:

```
Feb 6 17:57:40 CE1 feb BCHIP 0: uCode overflow - needs 129 inst
space to load b3_atm2_LSI_decode for stream 12
```

For more information, see "Verifying PIC Combinations" in the *Installation and Upgrade JUNOS Software* manual.