

# Troubleshooting JUNOS Platforms

---

9.a

**Student Guide**



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

Course Number: EDU-JUN-TJP

Juniper Networks, the Juniper Networks logo, JUNOS, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOSe is a trademark of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

*Troubleshooting JUNOS Platforms Student Guide, Revision 9.a*

Copyright © 2009, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Revision History:

Revision 9.a—July 2009

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 9.5R1.8. 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 Juniper Networks software are described in the software license provided with the software, or to the extent applicable, in an agreement executed between you and Juniper Networks, or Juniper Networks agent. By using Juniper Networks software, you indicate that you understand and agree to be bound by its license terms and conditions. Generally speaking, the software license restricts the manner in which you are permitted to use the Juniper Networks software, may contain prohibitions against certain uses, and may state conditions under which the license is automatically terminated. You should consult the software license for further details.

# Contents

---

<b>Chapter 1:</b>	<b>Course Introduction</b> .....	<b>1-1</b>
<b>Chapter 2:</b>	<b>Overview of JUNOS Platforms</b> .....	<b>2-1</b>
	JUNOS Platforms Overview .....	2-3
	Installation and Handling Guidelines .....	2-18
	Platform Architecture and Components .....	2-28
	Interface Overview .....	2-61
<b>Chapter 3:</b>	<b>Troubleshooting Tool Kit for JUNOS Platforms</b> .....	<b>3-1</b>
	Caveats and Warnings .....	3-3
	Troubleshooting Methodology .....	3-6
	Troubleshooting Tools: The JUNOS Software CLI .....	3-14
	Troubleshooting Tools: The Craft Interface Panel .....	3-27
	Troubleshooting Tools: System Logs and Protocol Tracing .....	3-30
	Troubleshooting Tools: Interactive UNIX Shell .....	3-44
	Troubleshooting Tools: Core Files for Diagnostic Analysis .....	3-53
	Troubleshooting Tools: The JTAC Knowledge Base .....	3-62
	Best-Practices Case Study .....	3-67
	Lab 1: JUNOS Troubleshooting Tools .....	3-79
<b>Chapter 4:</b>	<b>JUNOS Platforms Hardware Troubleshooting</b> .....	<b>4-1</b>
	Hardware Troubleshooting Overview .....	4-3
	Power On, Power Off, and Boot Media .....	4-5
	Using the CLI to Troubleshoot .....	4-14
	Case Study .....	4-41
	Lab 2: Chassis Hardware Troubleshooting .....	4-54
<b>Chapter 5:</b>	<b>Interface Troubleshooting</b> .....	<b>5-1</b>
	Interface Configuration Overview .....	5-3
	General Interface Troubleshooting .....	5-9
	Media-Specific Interface Troubleshooting .....	5-30
	Case Study .....	5-79
	Lab 3: Interface Troubleshooting .....	5-97
<b>Chapter 6:</b>	<b>JTAC Processes, Guidelines, and Support Resources</b> .....	<b>6-1</b>
	Opening a Support Case .....	6-3
	Support Services .....	6-9
	How to Use FTP to Send Files to JTAC .....	6-16
<b>Appendix A:</b>	<b>JUNOS Platform Details</b> .....	<b>A-1</b>
	Components and Characteristics of Multiservice Routers .....	A-3
	Components and Characteristics of Ethernet Services Routers and Switches .....	A-32
	Primary Components and Characteristics of Security Services Gateways .....	A-56
	End-of-Life Products .....	A-66

**Appendix B: Packet Flow Details** .....**B-1**  
    RTOS Packet Flow ..... B-3  
    ABC Chipset Packet Flow ..... B-5  
    LMNR Chipset Packet Flow ..... B-12

**Appendix C: Acronym List** .....**C-1**

**Appendix D: Answer Key** .....**D-1**

Not for Reproduction



## Course Overview

---

This course provides students with the foundational knowledge required to troubleshoot Juniper Networks platforms running JUNOS Software. This two-day course provides a brief overview of the device families that run JUNOS Software and discusses the key architectural components of the devices. Additional key topics include discussions of the JUNOS Software troubleshooting toolkit, basic hardware and interface troubleshooting using the command-line interface (CLI), and Juniper Networks Technical Assistance Center (JTAC) processes, guidelines, and support resources.

Through demonstrations and hands-on labs, you will gain experience in troubleshooting and monitoring the JUNOS Software and basic device operations.

### Objectives

After successfully completing this course, you should be able to:

- Describe current JUNOS platforms offerings.
- Describe general installation procedures.
- Explain the architecture of JUNOS platforms.
- Describe the function of JUNOS platform components.
- Describe layered troubleshooting methodology.
- Use various troubleshooting tools.
- Explain JTAC recommendations for current troubleshooting best practices.
- Troubleshoot JUNOS platforms using visual indicators.
- Troubleshoot JUNOS platforms using the CLI.
- Troubleshoot JUNOS platform interfaces.
- Describe recommended JTAC troubleshooting processes and guidelines.

### Intended Audience

This course benefits individuals responsible for maintaining and monitoring devices running JUNOS Software.

### Course Level

The Troubleshooting JUNOS Platforms course is a two-day introductory course.

### Prerequisites

Students should have taken the Introduction to JUNOS Software (IJS) course, and should have basic networking knowledge, and an understanding of the OSI model and the TCP/IP protocol suite.

## Course Agenda

---

### Day 1

Chapter 1: Course Introduction

Chapter 2: Overview of JUNOS Platforms

Chapter 3: Troubleshooting Tool Kit for JUNOS Platforms

Lab 1: JUNOS Troubleshooting Tools

### Day 2

Chapter 4: JUNOS Platforms Hardware Troubleshooting

Lab 2: Chassis Hardware Troubleshooting

Chapter 5: Interface Troubleshooting

Lab 3: Interface Troubleshooting

Chapter 6: JTAC Processes, Guidelines, and Support Resources

Appendix A: JUNOS Platform Details

Appendix B: Packet Flow Details

## Document Conventions

---

### CLI and GUI Text

Frequently throughout this course, we refer to text that appears in a command-line interface (CLI) or a graphical user interface (GUI). To make the language of these documents easier to read, we distinguish GUI and CLI text from chapter text according to the following table.

Style	Description	Usage Example
Franklin Gothic	Normal text.	Most of what you read in the Lab Guide and Student Guide.
Courier New	Console text: <ul style="list-style-type: none"><li>• Screen captures</li><li>• Noncommand-related syntax</li></ul> GUI text elements: <ul style="list-style-type: none"><li>• Menu names</li><li>• Text field entry</li></ul>	<code>commit complete</code> <code>Exiting configuration mode</code> Select <code>File &gt; Open</code> , and then click <code>Configuration.conf</code> in the <code>Filename</code> text box.

### Input Text Versus Output Text

You will also frequently see cases where you must enter input text yourself. Often this will be shown in the context of where you must enter it. We use bold style to distinguish text that is input versus text that is simply displayed.

Style	Description	Usage Example
Normal CLI Normal GUI	No distinguishing variant.	<code>Physical interface:fxp0, Enabled</code> View configuration history by clicking <code>Configuration &gt; History</code> .
<b>CLI Input</b> <b>GUI Input</b>	Text that you must enter.	<code>lab@San_Jose&gt; <b>show route</b></code> Select <code>File &gt; Save</code> , and enter <b><code>config.ini</code></b> in the <code>Filename</code> field.

## Defined and Undefined Syntax Variables

Finally, this course distinguishes between regular text and syntax variables, and it also distinguishes between syntax variables where the value is already assigned (defined variables) and syntax variables where you must assign the value (undefined variables). Note that these styles can be combined with the input style as well.

Style	Description	Usage Example
<i>CLI Variable</i> <i>GUI variable</i>	Text where variable value is already assigned.	<code>policy my-peers</code> Click on <i>my-peers</i> in the dialog.
<u><i>CLI Undefined</i></u> <u><i>GUI Undefined</i></u>	Text where the variable's value is the user's discretion and text where the variable's value as shown in the lab guide might differ from the value the user must input.	Type <code>set policy <u><i>policy-name</i></u></code> . <code>ping 10.0.1.1</code> Select File > Save, and enter <u><i>filename</i></u> in the Filename field.

## 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/education/>.

### About This Publication

The *Troubleshooting JUNOS Platforms Student Guide* was developed and tested using software Release 9.5R1.8. Previous and later versions of software might 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:

- Go to <http://www.juniper.net/techpubs/>.
- 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).

Not for Reproduction



# **Troubleshooting JUNOS Platforms**

## **Chapter 1: Course Introduction**

Not for Reproduction

## Chapter Objectives

- After successfully completing this chapter, 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

### This Chapter Discusses:

- Objectives and course content information;
- Additional Juniper Networks, Inc. courses; and
- Juniper Networks Technical Certification Program (JNTCP).



## Introductions

- Before we get started...
  - 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?



© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 1-3

### Introductions

This slide asks several questions for you to answer during class introductions.

## Course Contents

### ■ Contents:

- Chapter 1: Course Introduction
- Chapter 2: Overview of JUNOS Platforms
- Chapter 3: Troubleshooting Tool Kit for JUNOS Platforms
- Chapter 4: JUNOS Platforms Hardware Troubleshooting
- Chapter 5: Interface Troubleshooting
- Chapter 6: JTAC Processes, Guidelines, and Support Resources
- Appendix A: JUNOS Platform Details
- Appendix B: Packet Flow Details

### Course Contents

The slide lists the topics we discuss in this course.

## Prerequisites

- The prerequisites for this course are the following:
  - Introduction to JUNOS Software (IJS) course
  - Basic networking knowledge
  - Understanding of the OSI model and TCP/IP

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services


www.juniper.net | 1-5


### Prerequisites

The slide lists the prerequisites for this course.

## Course Administration

- The basics:
  - Sign-in sheet
  - Schedule
    - Class times
    - Breaks
    - Lunch
  - Break and restroom facilities
  - Fire and safety procedures
  - Communications
    - Telephones and wireless devices
    - Internet access



© 2009 Juniper Networks, Inc. All rights reserved.  Juniper Education Services [www.juniper.net](http://www.juniper.net) | 1-6

### General Course Administration

This slide documents general aspects of classroom administration.

## Education Materials

- Available materials:
  - In class:
    - Lecture material
    - Lab guide
    - Lab equipment
  - Online:
    - eLearning courses



© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 1-7

### Training and Study Materials

This slide describes Education Services materials that are available for reference both in the classroom and online.

## Additional Resources

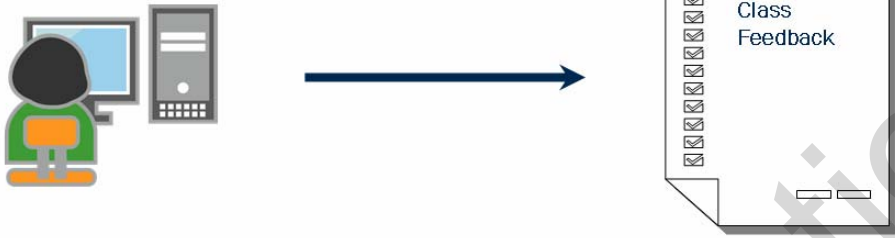
- For those who want more:
  - Juniper Networks Technical Assistance Center (JTAC)
    - <http://www.juniper.net/support/requesting-support.html>
  - Juniper Networks books
    - <http://www.juniper.net/training/jnbooks/>
  - Hardware and software technical documentation
    - Online: <http://www.juniper.net/techpubs/>
    - Image files for offline viewing:  
<http://www.juniper.net/techpubs/resources/cdrom.html>
  - Certification resources
    - <http://www.juniper.net/training/certification/resources.html>



### Additional Resources


This slide describes additional resources available to assist you in the installation, configuration, and operation of Juniper Networks products.

## Satisfaction Feedback



■ To receive your certificate, you must complete the survey

- Either you will receive a survey to complete at the end of class, or we will e-mail it to you within two weeks
- Completed surveys help us serve you better!

© 2009 Juniper Networks, Inc. All rights reserved.  Juniper Education Services [www.juniper.net](http://www.juniper.net) | 1-9

### 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 online 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.

## Juniper Networks Education Services Curriculum

- Consists of courseware for both enterprise and service provider environments
  - Complete list of courses
    - [http://www.juniper.net/us/en/training/technical\\_education/](http://www.juniper.net/us/en/training/technical_education/)

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 1-10

### Juniper Networks Education Services Curriculum

Juniper Networks Education Services can help ensure that you have the knowledge and skills to deploy and maintain cost-effective, high-performance networks for both enterprise and service provider environments. We have expert training staff with deep technical and industry knowledge, providing you with instructor-led hands-on courses as well as convenient, self-paced eLearning courses.

You can access the latest Education Services offerings covering a wide range of platforms at [http://www.juniper.net/us/en/training/technical\\_education/](http://www.juniper.net/us/en/training/technical_education/).



## Technical Certification Programs

- Demonstrate competence with Juniper Networks technology
  - Multiple tracks
  - Multiple certification levels
  - Written proficiency exams
  - Hands-on configuration and troubleshooting exams
  - For more information and details on how to prepare for the exams
    - <http://www.juniper.net/us/en/training/certification/>



### JNTCP

The Juniper Networks Technical Certification Program (JNTCP) consists of platform-specific, multitiered tracks that enable participants to demonstrate, through a combination of written proficiency exams and hands-on configuration and troubleshooting exams, competence with Juniper Networks technology. Successful candidates demonstrate thorough understanding of Internet and security technologies and Juniper Networks platform configuration and troubleshooting skills. You can learn more information about the JNTCP at <http://www.juniper.net/training/certification/>.

## Certification Levels

- Up to four levels per track:
  - Associate
    - Multiple choice exam
  - Specialist
    - Multiple choice exam
  - Professional
    - One-day, lab-based exam
  - Expert
    - One-day, lab-based exam

**JNCIA**  
Internet Associate

**JNCIS**  
Internet Specialist

**JNCIP**  
Internet Professional

**JNCIE**  
Internet Expert

© 2009 Juniper Networks, Inc. All rights reserved. [www.juniper.net](http://www.juniper.net) | 1-12

### Certification Levels

Each JNTCP track has one to four certification levels. Associate-level and Specialist-level exams are computer-based exams composed of multiple choice questions. These computer-based exams are administered at Prometric testing centers worldwide and have no prerequisite certification requirements.

Professional-level and Expert-level exams are composed of hands-on lab exercises that are administered at select Juniper Networks testing centers. Professional-level and Expert-level exams require that you first obtain the next lower certification in the track. Please visit the JNTCP Web site at <http://www.juniper.net/training/certification/> for detailed exam information, exam pricing, and exam registration.

## Certification Preparation

- How to prepare:
  - Training and study resources
    - JNTCP Web site  
<http://www.juniper.net/training/certification/>
    - Education Services training classes  
[http://www.juniper.net/training/technical\\_education/](http://www.juniper.net/training/technical_education/)
    - Juniper networks documentation and white papers  
<http://www.juniper.net/techpubs/>
  - Practical exams: lots of hands-on practice
    - On-the-job experience
    - Education Services training classes
    - Equipment access


© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

[www.juniper.net](http://www.juniper.net) | 1-13

### Prepping and Studying

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



### 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.



# **Troubleshooting JUNOS Platforms**

## **Chapter 2: Overview of JUNOS Platforms**

Not for Reproduction

## Chapter Objectives

- After successfully completing this chapter, you will be able to:
  - Describe current JUNOS platforms
  - Describe general installation procedures
  - Explain the architecture of JUNOS platforms
  - Describe the function of the RE, FPCs, PICs, SCBs, and Control Boards
  - Describe the operation of the Craft Interface
  - Describe interface naming conventions and the purpose of logical units

### This Chapter Discusses:

- The Juniper Networks platforms running JUNOS Software;
- Installation procedures for JUNOS platforms;
- General platform architecture;
- The function of major router components;
- Operation of the Craft Interface; and
- Interface naming conventions and the role of logical units.

## Agenda: Overview of JUNOS Platforms

- Overview of JUNOS Platforms
  - Installation and Handling Guidelines
  - Platform Architecture and Components
  - Interface Overview

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-3

### JUNOS Platforms Overview

The slide highlights the topics we cover in this chapter. We discuss the highlighted topic first.

## Categories of JUNOS Platforms

- Five product categories:
  - Multiservice routers
  - Ethernet services routers
  - Ethernet switches
  - Security services gateway
  - Multi-access gateway
- Industry-leading performance and reliability
- Legendary JUNOS Software

Categories of JUNOS Platforms

© 2009 Juniper Networks, Inc. All rights reserved. www.juniper.net | 2-4

### JUNOS Platform Categories

The categories of Juniper Networks JUNOS platforms are as follows:

- Multiservice routers (T Series, M Series, and J Series);
- Ethernet services routers (MX Series);
- Ethernet switches (EX Series);
- Security services gateways (SRX Series); and
- Multi-access gateway (BX7000).

Juniper Networks T Series core routers provide the highest possible forwarding performance density on the Internet today. They offer a wide selection of high-speed and extremely high-speed interface options suited for service provider cores, while maintaining feature richness and proven reliability.

Juniper Networks M Series Multiservice Edge Routers uniquely combine best-in-class IP and MPLS capabilities with unmatched reliability, stability, security, and service richness. These multiservice edge routers provide industry-leading port density across a wide range of medium-speed to high-speed interface options and price points.

Juniper Networks J Series Services routers offer predictable high performance and a variety of flexible interfaces that deliver secure, reliable network connectivity that is cost effective for remote, branch, and regional offices, and for small businesses.

*Continued on next page.*



## JUNOS Platforms Categories (contd.)

Juniper Networks MX Series Ethernet Services Routers provide Ethernet switching capabilities without sacrificing carrier-class routing features customers expect. MX Series devices have separate control and forwarding functions. Furthermore, the router's Ethernet switching separates Layer 2 and Layer 3 forwarding with the intelligence to bridge when possible and route when needed.

Juniper Networks EX Series Ethernet Switches offer flexible, powerful, and modular platforms that deliver performance, scalability, and high availability. You can deploy these products as a network access layer, as campus aggregation devices (within high-density data centers), or as core switches.

Juniper Networks SRX Series Services Gateways are the next generation security services gateways based on the Dynamic Services Architecture. The SRX Series gateways enable secure deployment of a wide range of business and residential applications and services ranging from small to large enterprises, at service provider premises and within data centers. The gateways offer native support for firewalls, virtual private networks (VPNs), switching and carrier-class Ethernet routing, and intrusion detection and prevention (IDP).

The Juniper Networks BX7000 Multi-Access Gateway operates within the environmental constraints of the cell site and features such common uplink types as copper, Ethernet, and DSL, to support both legacy and next-generation mobile technologies.

## Performance and Reliability

All Juniper Networks JUNOS platforms deliver deterministic forwarding performance with services enabled. All products employ the essential concept of separate forwarding and control planes. By performing all complex, computation-intensive tasks on an appropriately sized control plane, these products ensure that tasks go unhindered by any degree of forwarding requirements. Likewise, this separation guarantees that the custom-designed forwarding plane always operates at peak capacity, regardless of the complexity of control computations required at any given moment.

## JUNOS Software

A significant aspect of the JUNOS product line is that all JUNOS platforms run JUNOS Software with support for all features. Even the small enterprise-class J Series routers run the same JUNOS Software—repackaged to include the real-time operating system (RTOS) software and related interface drivers instead of the high-end M Series and T Series hardware drivers.

## Multiservice Routing Platforms

- **T Series Core Routers:**
  - T320, T640, T1600, and TX Matrix
- **M Series Multiservice Edge Routers:**
  - M7i, M10i, M40e, M120, and M320
- **J Series Services Routers:**
  - J2320, J2350, J4350, and J6350

### T Series Core Routers

T Series Core Routers provide Gigabit Ethernet, SONET/SDH, and other high-speed interfaces for large networks and network applications, such as those that ISPs support. Application-specific integrated circuits (ASICs) are a definitive part of the router design. ASICs enable the router to achieve data forwarding rates that match current fiber-optic capacity. The T Series includes the T320, T640, T1600, and TX Matrix platforms.

### M Series Multiservice Edge Routers

M Series Multiservice Edge Routers include the M7i, M10i, M40e, M120, and M320. Because the same scalable and production-hardened JUNOS Software runs on all M Series platforms, a consistent set of capabilities is available at all network layers, ranging from access layer to core. The platforms provide SONET/SDH, ATM, Ethernet, and channelized interfaces. ASICs, which are an integral part of device design, enable the routers to forward data at the wire rate.

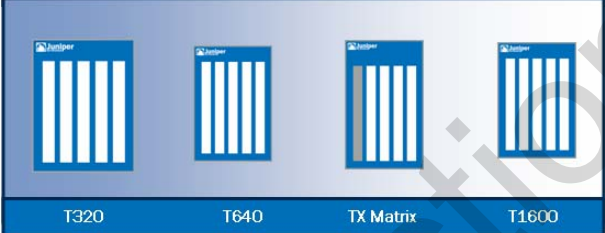
### J Series Services Routers

J Series Services Routers extend enterprise applications and deliver reliable connectivity to remote offices with a powerful blend of high-performance network protection and advanced services. J Series routers leverage the modular JUNOS Software, integrated WAN acceleration, and integrated voice gateway technology from Avaya. The router portfolio includes the J2320, J2350, J4350, and J6350.

## T Series Core Routers

- **Up to 1600 Gbps total throughput per chassis**
  - Half rack chassis
  - 3.2 Tbps in one 7-foot rack
  - 100 Gbps per slot
- **Wide range of interfaces:**
  - DS-3 to OC-768
- **Component level redundancy**
  - Routing Engines, Control Boards, Switch Interface Boards, and Power Entry Modules
  - Continuous operation for core

1/3 Rack platform Designed for small core	1 <sup>st</sup> 640 G platform (8) Ports in 1/2 rack	Optimal scale 3 <sup>rd</sup> dimension	1600 G platform 3.2 Tbps per rack
--	---	--	--------------------------------------



© 2009 Juniper Networks, Inc. All rights reserved. www.juniper.net | 2-7

### Total Chassis Throughput

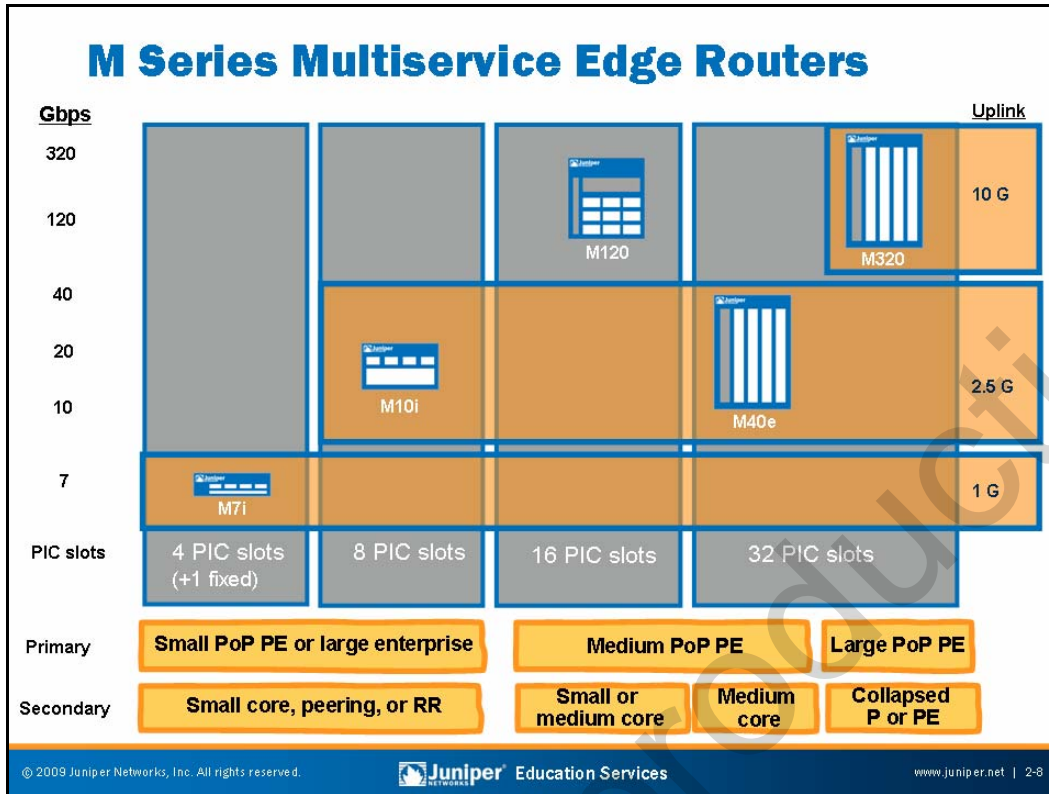
Juniper Networks T Series routers range from the T320, designed for a small network core, to the T1600, capable of offering 3.2 Tbps. The T1600 is a multichassis-capable core router that you can upgrade in-service from existing T640 routers. The scale and density of the T1600 allows service providers to increase capacity without adding equipment to the network, saving on valuable floor space, rack space, and power consumption.

### Range of Interfaces

T Series routers provide edge interfaces as well as the core functions required for consolidated point of presence (POP) solutions. All T Series routers, from the T320 to the T1600 and TX Matrix, support a wide range of interfaces, ranging from DS-3 to OC-768, including Asynchronous Transfer Mode (ATM), SONET, Ethernet, and both fixed and tunable dense wavelength-division multiplexing (DWDM) interfaces.

### Components Level Redundancy

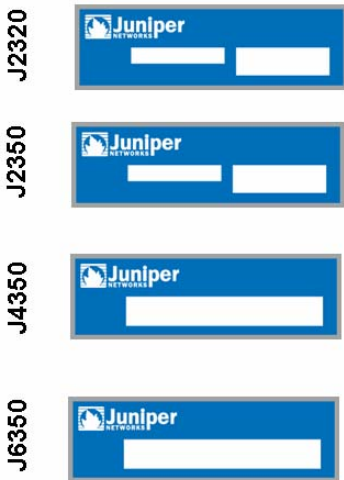
No single point of failure exists in the entire product line of T Series routers—including Routing Engines (REs), Control Boards (CBs), Switch Interface Boards (SIBs), and Power Entry Modules (PEMs). The redundancy in these components provides high availability and continuous operation in core routing, where loss of a single routing node can remove service for a wide geographical area.




### M Series Multiservice Edge Routers

The M Series portfolio ranges from 7 Gbps platforms to 320 Gbps platforms. You can deploy M Series routers in various roles within your network. The M Series portfolio uniquely combines IP/MPLS capabilities with service richness, stability, reliability, and security. The M Series routers allow service providers to consolidate multiple networks on a single IP/MPLS infrastructure. You can deploy the M Series platforms as a multiservice edge router, a small or a medium core router, a route reflector, or a peering device. It can also be deployed in multicast, mobile, or data center applications. The M10i and higher-end M Series routers offer RE and Control Board redundancy.

## J Series Services Routers



- Performance with services
  - Unmatched performance with services enabled
- Extensive connectivity
  - Four on-board Gigabit Ethernet ports
  - Expandable WAN and LAN interfaces using modules

© 2009 Juniper Networks, Inc. All rights reserved.  Education Services [www.juniper.net](http://www.juniper.net) | 2-9

### J Series Performance and Services

J Series routers maintain performance even when you enable advanced services such as Network Address Translation (NAT), access control lists, and stateful firewalls. Modular JUNOS Software provides key routing and security features such as stateful firewalls, IPsec, MPLS, and IPv6, defending against infrastructure attacks and fully protecting the processing resources.

### J Series Extensive Connectivity

J Series routers provide a large selection of connectivity options including T1 and E1, Serial, Fast Ethernet, Gigabit Ethernet, DS3, E3, ISDN, ADSL2+, and G.SHDSL. The product options include hardware encryption acceleration, power supplies, DRAM, compact flash, and feature licenses. All J2320, J2350, J4350, and J6350 routers ship with four fixed 10/100/1000 Ethernet ports. You can add more modular LAN and WAN interfaces.

## MX Series Ethernet Services Routers

- Dependable hardware
  - MX960, MX480, and MX240
- Designed for Ethernet routing with switching
  - Ethernet services edge and Layer 2 and Layer 3 Ethernet aggregation
  - DPCs with varying levels of QoS, Layer 2 switching, and Layer 3 routing services

Model	PIC slots	Capacity (Gbps)
MX240	4	120
MX480	8	240
MX960	14	480

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services [www.juniper.net](http://www.juniper.net) | 2-10

### MX Series Ethernet Services Routers Hardware

MX Series routers (MX960, MX480, and MX240) provide Ethernet switching capabilities without sacrificing the carrier-class routing features customers expect from Juniper Networks. MX Series routers surpass the requirements of carrier-grade Ethernet switches as defined by the Metro Ethernet Forum, leveraging the MPLS capabilities that have made Juniper Networks routers the platforms of choice for service providers seeking maximum performance, availability, and service agility. By extending the carrier-class routing functionality of JUNOS Software to include LAN switching functionality to facilitate migration and growth, Juniper Networks brings its traditional advantages to Ethernet aggregation. These advantages include high-performance routing capabilities such as nonstop routing (NSR), MPLS, fast reroute, and unified in-service software upgrade.


### Ethernet Routing with Switching Design

Running the same version of JUNOS Software shared by all Juniper Networks core and edge platforms, MX Series routers not only separate control and forwarding functions, but the Ethernet switching separates Layer 2 and Layer 3 forwarding with the intelligence to bridge when possible and route when needed. The platforms support Dense Port Concentrator (DPC) interface cards, offering enhanced queuing capabilities, QoS, L2 switching, and L3 routing services.



## EX Series Ethernet Switches

- **High-performance chassis platforms**
  - EX8216: Sixteen 200 Gbps line cards
  - EX8208: Eight 200 Gbps line cards
  - EX4200: 128 Gbps virtual backplane
  - EX3200: Fixed port system plus flexible 4-port and 10-port uplink modules
  - EX2500: 24 dual-mode 10 Gbps ports
  - EX2200: 24-port or 48-port systems—10/100/1000BASE-T
- **Proven Juniper Networks Technology**
  - Switch fabrics and control plane
  - JUNOS Software



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 2-11

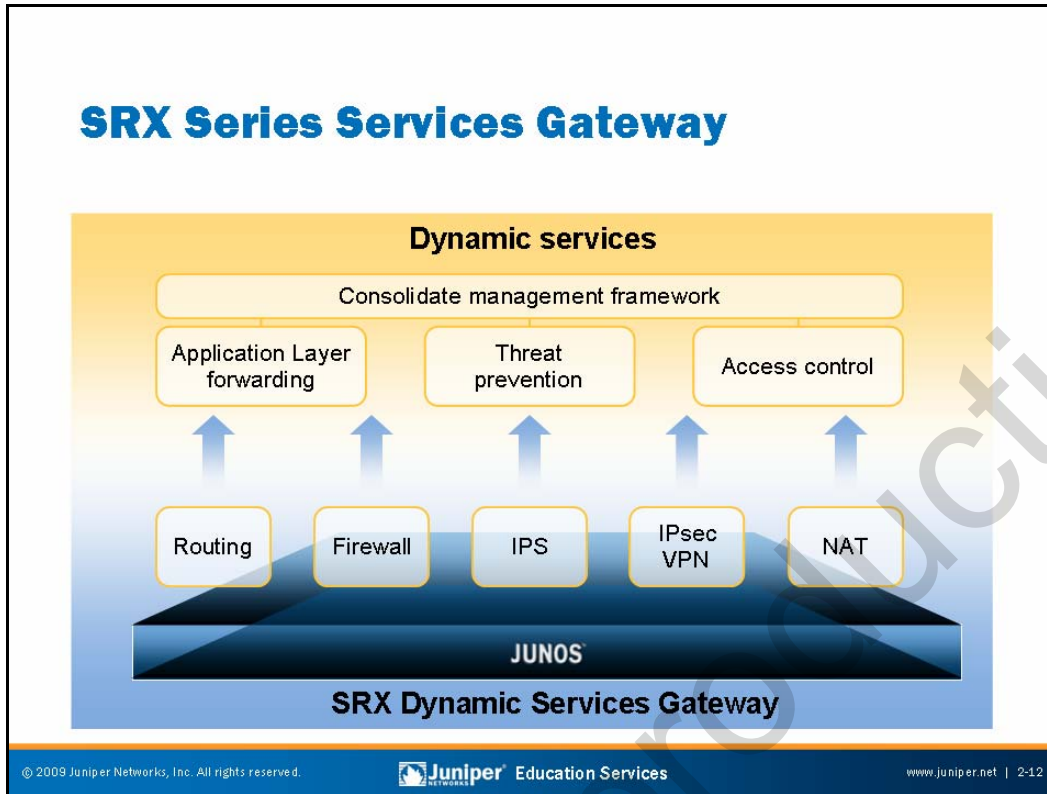
### EX Series Ethernet Switches

The EX Series Ethernet Switches offer flexible, powerful, and modular platforms that deliver the performance, scalability, and high availability required for today's high-density data center, campus aggregation, and core switching environments. The platforms range from the EX2200 to the EX8216 systems, which could include fixed ports and modular line cards of up to 10 Gbps.

### Proven JUNOS Software Technology

The EX8200 line Switch Routing Engines (SREs) process all Layer 2 and Layer 3 protocols and manage individual chassis components, while the switch fabric module provides the central crossbar matrix through which all data traffic passes. The SRE and switch fabric modules work together to fulfill all RE and switch fabric functions.

The EX Series Ethernet switches run the same JUNOS Software that Juniper Networks routers run. By utilizing a common operating system, Juniper Networks delivers a consistent implementation and operation of control plane features across all products. To maintain that consistency, JUNOS Software adheres to a highly disciplined development process that uses a single source code, follows a single quarterly release train, and employs a highly available modular architecture that prevents isolated failures from bringing an entire system down.



### SRX Series Services Gateway


Juniper Networks SRX Series Services Gateways are the next-generation solution for securing the ever increasing network infrastructure and application requirements for both enterprise and service providers. Designed from the ground up to provide flexible processing scalability, input-output scalability, and services integration, the SRX Series devices meet the network and security requirements of data center consolidation, managed services deployments, and aggregation of security solutions. Running JUNOS Software, which incorporates the Juniper Networks routing heritage and service provider reliability, the SRX Series also offers the high feature and service integration necessary to secure modern network infrastructure and applications.



## SRX Series Platforms (1 of 3)

- SRX5600 and SRX5800 share the same cards or blades

- Distributed multiblade chassis-based design, using Juniper Networks switch fabric technology
- Blades include SCB, RE, SPC, and DPC
- 2 SCB slots + 6 DPC or SPC slots in SRX5600
- 2 SCB slots + 12 DPC or SPC slots in SRX5800



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 2-13


### SRX5600 and SRX5800

SRX5600 and SRX5800 are next-generation services gateways based on new architecture that provides unprecedented scalability and service integration. The devices are ideally suited for large enterprise and service provider networks. Based on the Dynamic Services Architecture, the SRX5600 uses the same services processing cards (SPCs) and input/output cards (IOCs) as the SRX5800. The SRX5800 can support a more than 120 Gbps firewall and 30 Gbps IDP traffic. The SRX5600 can support up to a 60 Gbps firewall and 15 Gbps IDP traffic.

## SRX Series Platforms (2 of 3)

- SRX3400 and SRX3600 share the same cards or modules

- Modularized chassis-based design, using Juniper Networks switch fabric technology
- Cards include midplane, SCB, RE, SPC, NPC, and IOC
- SPC, NPC, and IOC are CFM form factor cards, and interchangeable in any CFM slots
- 7 CFM slots in SRX3400
- 12 CFM slots in SRX3600
- SRX5600 and SRX5800 cards are *not* interchangeable with SRX3400 and SRX3600 cards




### SRX3400 and SRX3600 Cards and Modules

The SRX3400 and SRX3600 offer native support for firewalls, VPNs, switching, carrier-class Ethernet routing, and IDP. Each device is flexible and scalable, with multiple, interchangeable physical interface cards and security cards, including the midplane, RE, SPCs, network processing cards (NPC), Switch Control Board (SCB), IOC, Common Form-Factor Modules (CFMs), and DPCs. The SRX3600 can support an approximately 10 Gbps firewall and 5 Gbps IDP traffic. The SRX3400 can support a 10 Gbps firewall and 2 Gbps IDP traffic. SRX3000 cards and SRX5000 cards are *not* interchangeable.

## SRX Series Platforms (3 of 3)

- SRX650, SRX240, and SRX210 support:

- Gigabit LAN or WAN connectivity to support high data rate over Ethernet
- Power over Ethernet support—power delivery to IP phones, cameras, and wireless devices
  - Simplified connectivity
  - No need for external power
- Maximum firewall performance
  - SRX650: 7 Gbps
  - SRX240: 1.5 Gbps
  - SRX210: 750 Mbps



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 2-15

### SRX650, SRX240, and SRX210

The SRX650, SRX240, and SRX210 provide full firewall and Unified Threat Management (UTM), encompassing antivirus, IDP, Web filtering, and antispam to secure branch offices with a wide range of interfaces for WAN connectivity.

The SRX210 Power over Ethernet (PoE) feature simplifies IP phone, camera, and wireless support by delivering power to those devices without any need for external power.

The maximum firewall performance for the SRX650 is 7 Gbps, while it is 1.5 Gbps for the SRX240, and 750 Mbps for the SRX210.

## BX Series Multi-Access Gateway

### ■ BX7000 details:

- 1.5U, compact unit—ideal for cell site deployment
- Expansion slots are available for flexibility
- Passive cooling for increased reliability and power conservation
- DC option is available



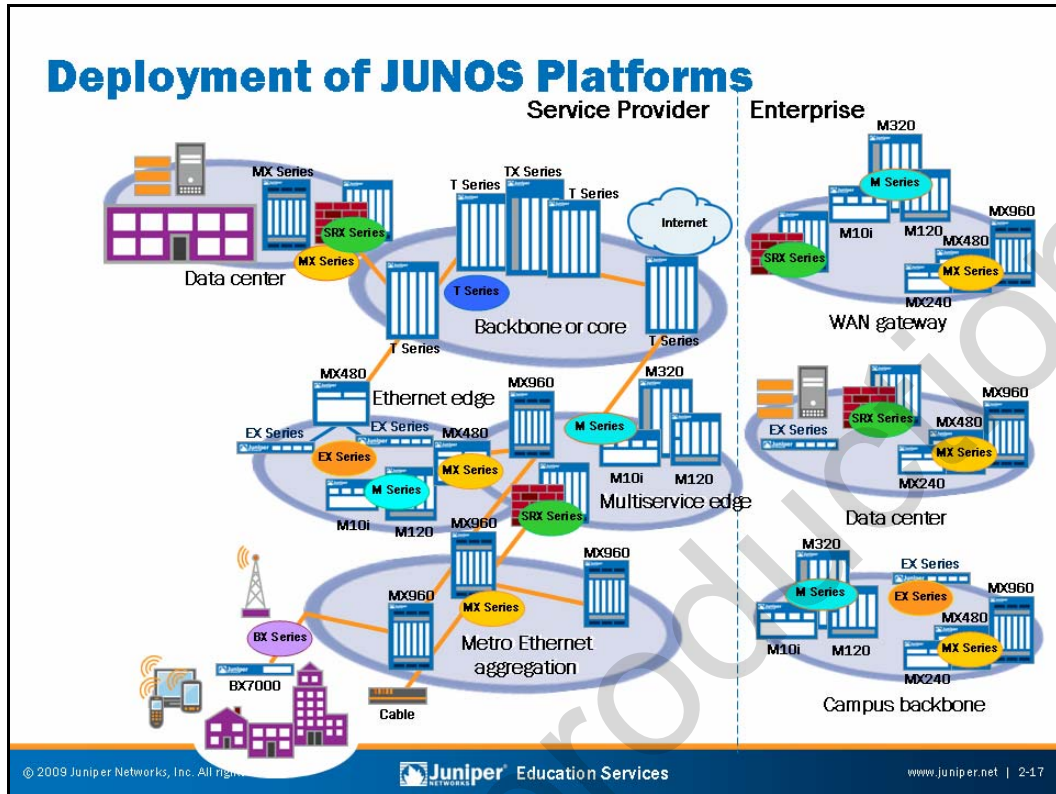
© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-16

### BX Series Multi-Access Gateway

The BX7000 Multi-Access Gateway effectively addresses the backhaul evolution challenge for mobile operators with the most operationally efficient approach. The BX7000 is a compact platform that enables ease of its deployment in a cell site cabinet with limited rack space. Additionally, it is available in an environmentally hardened form factor for deployment scenarios where cabinets could have additional exposure to natural elements.



### Deployment of JUNOS Platforms

The slide summarizes our review of all JUNOS platforms.

## Agenda: Overview of JUNOS Platforms

- Overview of JUNOS Platforms
- Installation and Handling Guidelines
- Platform Architecture and Components
- Interface Overview

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-18

### Installation and Handling Guidelines

The slide highlights the topic we discuss next.



## Chassis Installation Guidelines

- Follow site-preparation guidelines
  - Space, environment, power, grounding, and so forth
- Lifting requires three or more people for some platforms
  - We recommend a mechanical lift
  - Selected M Series, T Series, MX Series, and SRX Series maximum product weights:
    - T640: 565 pounds (256.28 kg)
    - T320: 369.9 pounds (167.78 kg)
    - M320: 439 pounds (199.6 kg)
    - MX960: 334 pounds (151.7 kg)
    - SRX5800: 334 pounds (151.7 kg)
- Remove heavier components before lifting
  - Power supplies, FPCs, fan trays, and system boards
- Lift into rack
  - Replace components

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-19

### Follow Site Preparation Guidelines

To ensure a smooth installation, each 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. We also provide an installation checklist form to help ensure that you are ready to press your new platform into service the day that it arrives.

For hardware manuals, go to: <http://www.juniper.net/techpubs/hardware/index.html>.

### Heavy Lifting Requires More People

A Juniper Networks platform with a maximum configuration can weigh anywhere from 61 lbs (27.7 kg) for an M7i router, all the way up to 565 lbs (256.28 kg) for a T640 Core Router.

### Remove the Heavier Components Before Lifting

Removing the power supplies, FPCs, fan trays, REs, and CBs 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 weighs 565 pounds (256 kg).

*Continued on next page.*

### Carefully Lift Device into Rack

Use extreme care when lifting heavy platforms! We recommend a mechanical lift. Once properly secured in the rack, you should replace the components that you removed from the chassis to reduce weight.

Not for Reproduction



## Powering On and Powering Off

- Power on:
  - Connect all cables
  - Turn on one power supply
  - Turn on second power supply
    - It can take up to 60 seconds for accurate power supply and PEM status indications
- Power off:
  - Shut down JUNOS Software
    - CLI `request system halt` command
  - Turn off power supplies and remove power

### Powering On JUNOS Platforms

You can equip each device can be equipped with two redundant, load-sharing power supplies of the same type; in most cases these power supplies can be either AC or DC. 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 that plugs into a grounded 100–240 VAC power receptacle. For sites with a DC power source, power normally travels 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 you intend to install the router. A pair of cables (–48 V and RTN) connects each DC supply to the power distribution panel. The rear enclosure has grounding studs. 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 and PEM vary by platform, you should note that after a power supply powers 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 Off JUNOS Platforms

You must always perform a graceful shutdown of JUNOS Software before removing power. We discuss this procedure in detail on the next slide.

## Reboots and Shutdowns

- 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:

```
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
```

### Graceful Shutdowns and Rebooting

JUNOS Software runs on a multitasking, multi-user operating system based on FreeBSD. As with any UNIX system, you should always gracefully shut down the system before removing power (as opposed to performing an abrupt hard power off, in which you simply remove the power from the system). Failing to shut down in a graceful manner can result in file system corruption that prolongs the next boot, and in the worst of cases, it might actually prevent a successful boot. Note that on some platforms the chassis has a button that gracefully shuts down the RE when depressed for 3–5 seconds. Once the system properly halts, you can safely remove power.

The **request system halt** command gracefully stops the software and prepares the device to shut down. Note that you must either cycle power or press the Enter key on the terminal attached to the device's console port to effect the reboot of a device that has executed a shutdown command.

*Continued on next page.*

## Graceful Shutdowns and Rebooting (contd.)

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!
```

The **request system reboot** command causes the device to reboot. Reboot requests record 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 from which to boot. Media options include compact-flash and disk. The device always attempts to boot from removable media when such media is present and you perform a power cycle (cold boot).

## General Handling: Mechanical

- FPCs are heavy and you can damage them with improper handling
  - Fully loaded FPC3 weighs up to 32 lbs (11.3 kg)
    - Be prepared for the weight of the FPC when removing it from the midplane
    - Always handle appropriately



Examples of proper FPC handling




© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-24

## FPC Handling


Some Flexible PIC Concentrators (FPCs) are heavy and can receive damage with improper handling. The following figures illustrate some common forms of FPC abuse that can damage your expensive hardware!




## General Handling: Electrostatic Damage

- Always observe ESD precautions
  - Use a grounded wrist or ankle strap connected to a chassis ESD lug

ESD lugs are normally located on the front and back of each chassis



- Always place FRUs in ESD-approved packaging



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 2-25

### General Handling: Electrostatic Damage

The field-replaceable units (FRUs) associated with JUNOS platforms contain electrostatic discharge (ESD) sensitive parts; some components can receive damage 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.

Follow these 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 Mega Ohms.
- When handling any component that you remove from the chassis, make sure the equipment end of your ESD strap attaches 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

- Distinguish FRUs by whether you can remove and install them without causing system disruption
  - You can swap hot-insertable and hot-removable FRUs without disrupting global routing functions
    - You must follow procedures for online and offline of hot-swappable FRUs to ensure that global routing does not experience disruption
  - Hot-pluggable FRUs interrupt global routing functions when you remove or install the component
  - Some FRU types require that you remove power 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

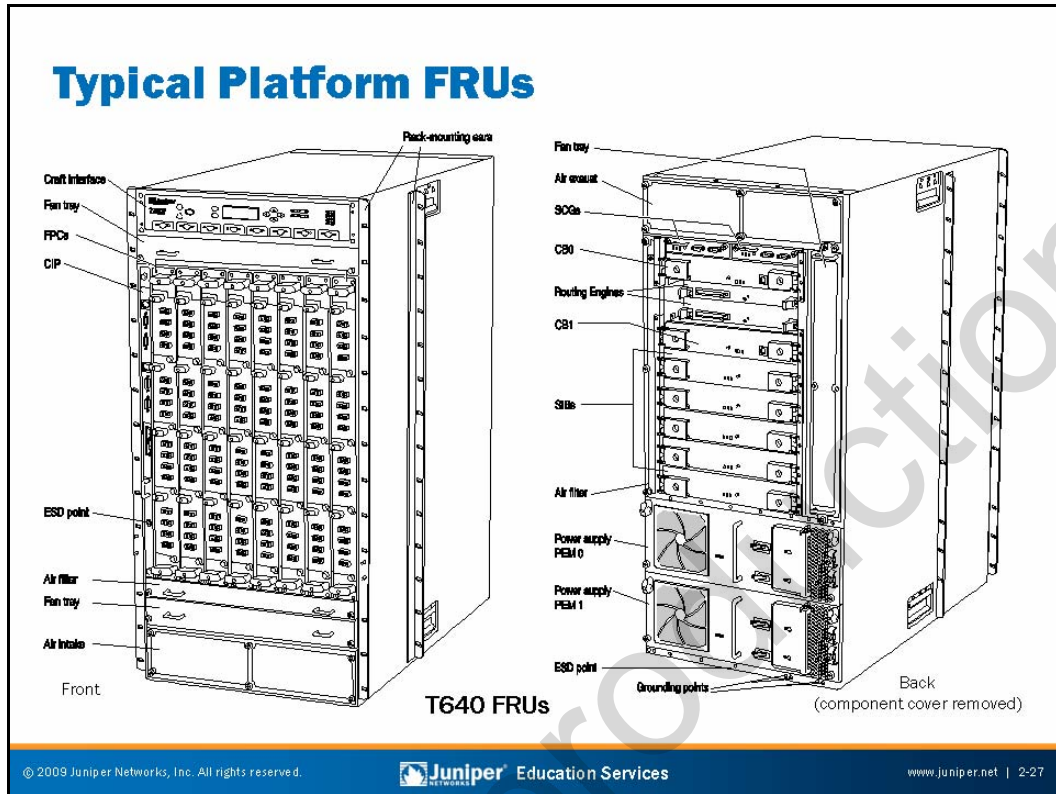
**T640 routing node FRU classification**

## FRU Types

FRUs are platform components that you can replace at the customer site. Replacing FRUs requires minimal routing node downtime. Generally speaking, the three types of FRUs are the following:

- *Hot-insertable and hot-removable FRUs:* You can remove and replace these components without powering off the routing node or disrupting routing functions. We often refer to this type of FRU as being *hot-swappable*.
- *Hot-pluggable FRUs:* You can remove and replace these components without powering off the device, but you interrupt the routing functions of the system when you remove the component.
- *FRUs that require power off:* In rare cases, an FRU requires that you remove power from the chassis before removing or inserting it. An example of this type of FRU is the M7i and M10i platform Compact Forwarding Engine Board (CFEB) and RE. In the case of these platforms, the lack of redundant Packet Forwarding Engine (PFE) and RE capabilities made the engineering of hot-swappability for these components a nonissue.





### Typical Platform FRUs

All JUNOS platforms consist of a chassis and one or more FRUs. The slide shows the primary FRUs associated with the T640.

## Agenda: Overview of JUNOS Platforms

- Overview of JUNOS Platforms
- Installation and Handling Guidelines
- Platform Architecture and Components
- Interface Overview

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-28


### Platform Architecture and Components


The slide highlights the topic we discuss next.



## Fundamentals of Juniper Networks Advantage

- JUNOS Software innovation:
  - Modular control plane software
  - Memory protected processes

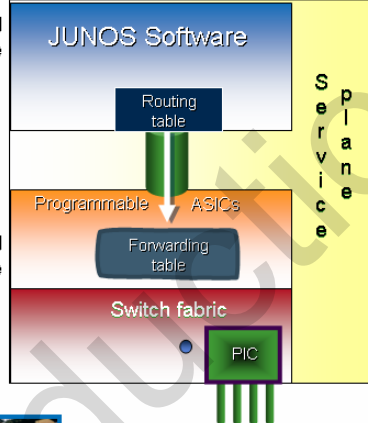






- Breakthrough architecture
  - Performance, intelligence, and flexibility

**Routing Engine**

**Packet Forwarding Engine**



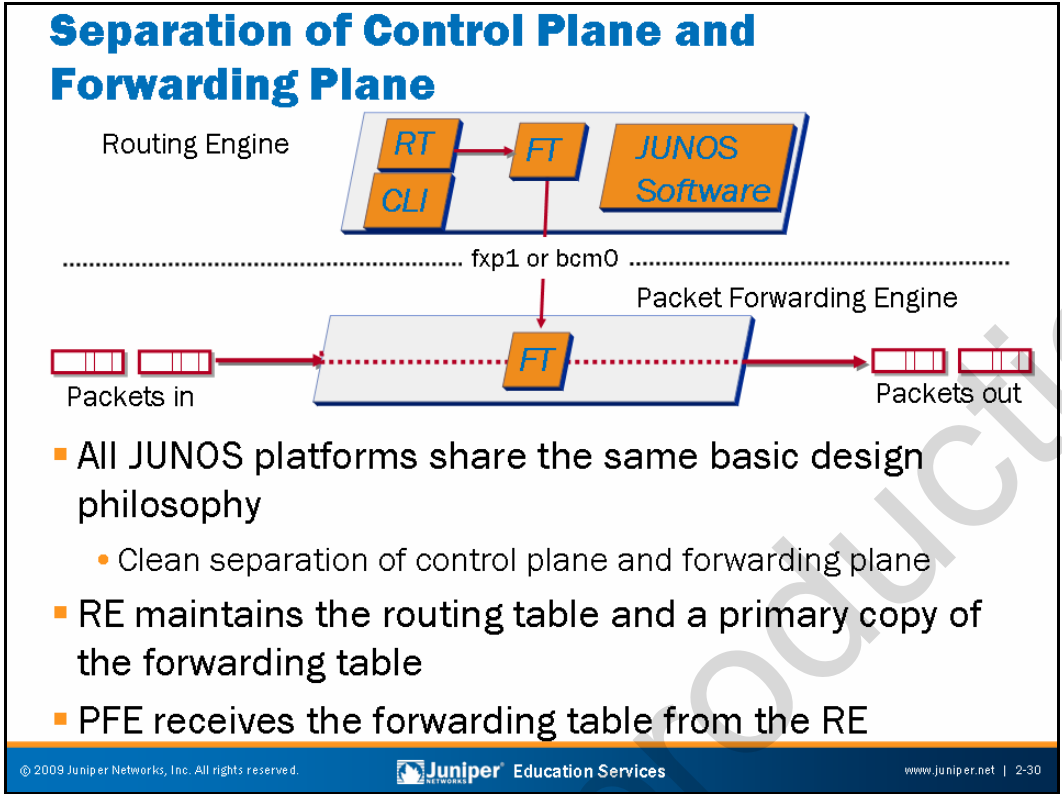


© 2009 Juniper Networks, Inc. All rights reserved.
 Juniper Education Services
www.juniper.net | 2-29

### JUNOS Software Innovation

The modular design of the architecture of JUNOS platforms results in a clean separation of the control plane and the forwarding plane, facilitating rich packet processing through ASIC construction. The control plane is where routers talk to each other to determine topology and to set up services. The Routing Engine, which boasts dedicated hardware, represents the control plane. Running on top of the Routing Engine is the industry's most proven and scalable operating system, JUNOS Software. The operating system might be the single most critical component of a high-performance device. JUNOS Software is a modular fault-protected operating system that can handle thousands of VPNs and BGP sessions and other types of routing and signaling control information. Once the Routing Engine determines routing topology and services, it pushes that information down to the forwarding plane, or PFE. The forwarding plane consists of programmable ASICs, which deliver very rich packet processing on a per customer basis, and it enables unparalleled forwarding performance.

Smaller JUNOS platforms, for example J Series Services Routers, continue the practice of separation of forwarding and control planes, implementing this separation entirely in software. The routers use an RTOS to simultaneously implement JUNOS Software, packet forwarding, and advanced services—each in separate real-time threads.



### Architectural Philosophy

Architecturally, all JUNOS platforms share a common design that separates the router's control plane and forwarding plane. JUNOS platforms consist of the following two major components:

- *Routing Engine:* 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. It connects to the PFE through an internal 100 Mbps connection identified as fxp1 on most platforms. Some larger JUNOS platforms (such as the M320) use a Gigabit Ethernet link, referred to as bcm0, between the RE and a Fast Ethernet switch, which in turn has dedicated 100 Mbps links to each FPC.
- *Packet Forwarding Engine:* The PFE is responsible for forwarding transit packets through the router using an ASIC-based switching path. 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.

*Continued on next page.*

## Routing and Forwarding Table Interaction

The JUNOS Software routing protocol process implements the various routing protocols that run on the router. The routing protocol process starts all configured routing protocols and handles all routing messages. The routing protocol process (rpd) maintains one or more routing tables that consolidate 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.

## RE and PFE Synchronization

The PFE receives the forwarding table from the RE. In the majority of cases, the PFE's forwarding table and the RE's forwarding table synchronize over the 100 Mbps fxp1 Ethernet link, which interconnects the two entities. This synchronization ensures that a change in topology produces identical forwarding tables in the RE and PFE. In the case of the M320 platform, a 100 Mbps Ethernet switch provides a dedicated link to each FPC. These 100 Mbps links then present to the M320 RE as a single Gigabit Ethernet uplink named bcm0. Forwarding table updates are a high priority for the JUNOS Software kernel and it performs them incrementally.

## T Series Architecture

- T Series offers:
  - Separate routing and forwarding planes
  - Crossbar switch fabric
  - Distributed Packet Forwarding Engines on each FPC

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | 2-32

### T Series Architecture

The T Series architecture cleanly separates control operations from packet forwarding operations. This design eliminates processing and traffic bottlenecks, permitting the router to achieve high performance. The router's host subsystem—running JUNOS Software to handle routing protocols, traffic engineering, policy, monitoring, and configuration management—performs the control operation. The router's PFEs—consisting of ASICs—perform the forwarding operations.

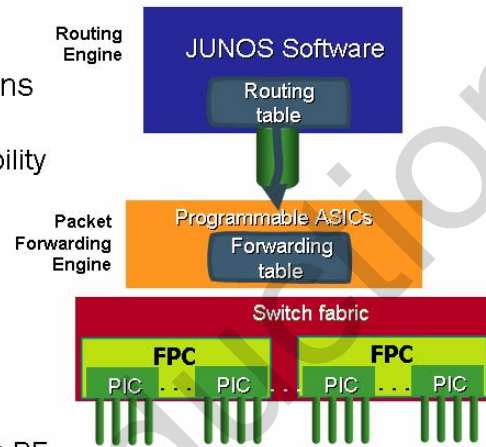
The use of distributed memory and crossbar switching in the forwarding plane contributes to a functional redundancy in the T series. Part of the distributed packet forwarding architecture includes positioning of a PFE on each FPC, which adds to the resiliency of the T Series devices.

Note that M120 and M320 routers deploy the T Series distributed memory architecture.

## Logical Platform View of M Series

### ■ The M Series:

- Clean separation of routing and packet forwarding functions
  - Consistent performance
  - Stability, reliability, and availability
- Routing Engine
  - Intel-based processor runs JUNOS Software
  - Maintains routing table and constructs forwarding table
- Packet Forwarding Engine
  - ASIC-based design
  - Receives forwarding table from RE
  - Conducts incremental table updates without forwarding interruption
- Flexible PIC Concentrator



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 2-33

### Logical Platform View of M Series

We designed the M Series architecture with scale and stability in mind, including the modular and fault-protected design of JUNOS Software. M Series architecture offers a clean separation of routing and packet forwarding functions, resulting in stable, predictable, and reliable systems. The Intel-based RE is responsible for running JUNOS Software, for performing routing updates, and for system management. The PFE is the logical entity responsible for packet forwarding. It physically consists of programmable ASICs—each of which is dedicated to a specific task. The PFE contains a copy of the forwarding table that matches the forwarding table copy in the RE. Similar to the T Series, the M Series device architecture includes FPCs, which could house Physical Interface Cards (PICs). The M Series devices (with the exception of the M120 and M320) use a shared memory architecture.

## J Series RTOS Technology

- Designed to meet the needs of small enterprises
  - Extremely competitive pricing
  - Nearly all features from M Series routers
  - Single CPU with real-time threads and interface processor
  - Port adapters versus FPCs
  - Many ease-of-use features

**RE** JUNOS Software

UNIX socket

**PFE** Host Shared memory

INQ JEXEC Result processing OUTQ

rt threads

J23X0  
J4350  
J6350

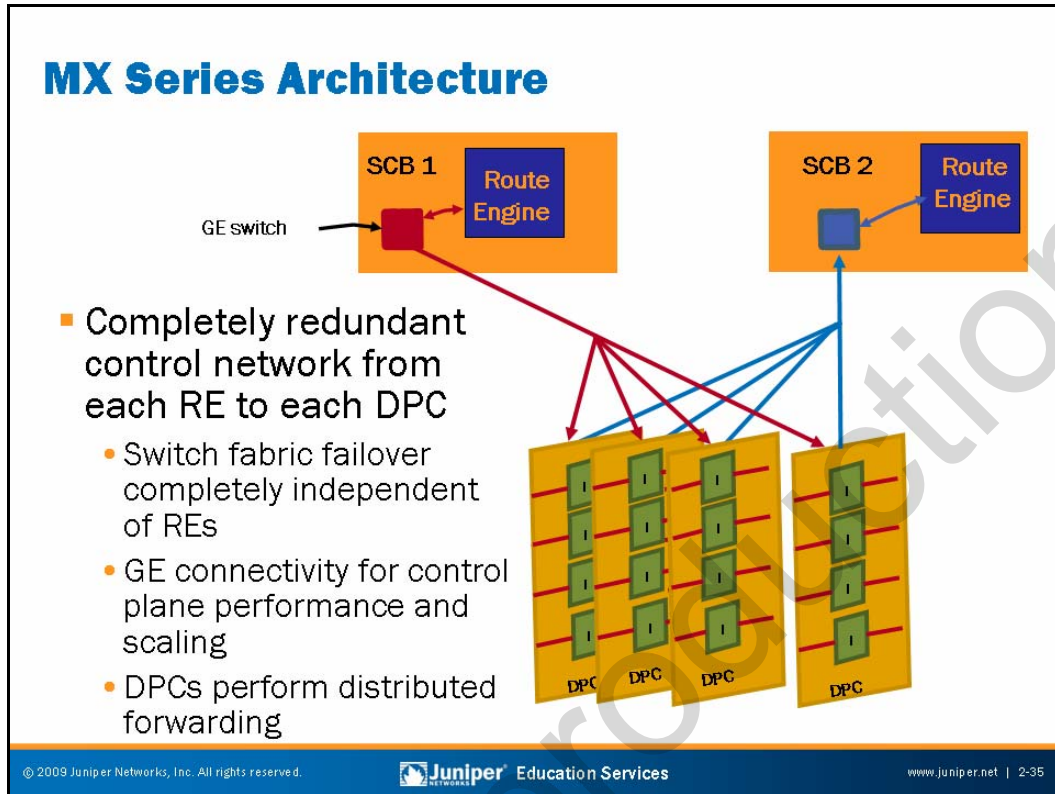
The J Series Services Router family is specifically designed for smaller enterprise environments including remote, branch, and regional offices.

© 2009 Juniper Networks, Inc. Juniper Education Services www.juniper.net | 2-34

### J Series Routers

The J Series brings deterministic forwarding performance and carrier-class stability to small enterprise and remote offices. These routers continue the practice of separation of the forwarding plane and control plane, but they feature the economical innovation of implementing this separation entirely in software. J Series routers use an RTOS 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 Physical Interface Modules (PIMs) instead of FPCs. A PIM is like an FPC with built-in PICs of a common media type. The PIMs have numerous low-speed WAN interfaces 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 also proportionally add the necessary processing power to maintain the router's performance capabilities with the additional connectivity.

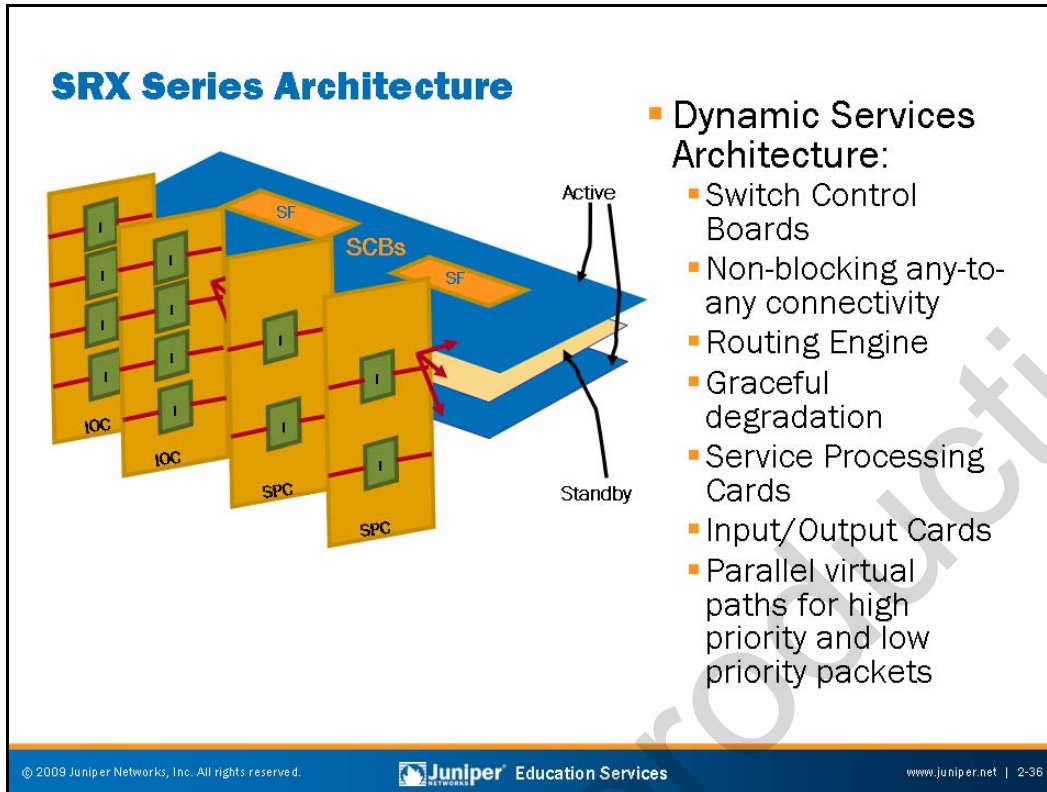


### MX Series Architecture

Similar to core routers, MX Series devices are packet based and deploy distributed memory architecture. The architecture includes RE redundancy, the control plane, the SCB, and the DPCs. The control plane of the platform's chassis consists of Gigabit Ethernet links between the SCBs or REs and each DPC. Each SCB has a connection to every DPC in the chassis, which in turn means that each RE has a redundant connection to each DPC. The switch fabric failover and the RE failover are independent from each other.

The host subsystem consists of an RE functioning together with an SCB. The router can have one or two host subsystems. If it has two host subsystems installed, one functions as the master and the other functions as the backup. If the master host subsystem (or either of its components) fails, the backup can take over as the master. To operate, each host subsystem requires an RE installed directly into an SCB. If you configured the REs for graceful switchover, the backup RE automatically synchronizes its configuration and state with that of the master RE. Any update to the master RE state replicates on the backup RE. If the backup RE assumes mastership, packet forwarding continues through the router without interruption.





### SRX Series Architecture

SRX Series devices deploy Dynamic Services Architecture that includes the management and control necessary to incorporate individual blades into a powerful collective solution. Rather than housing disparate cards, the Dynamic Services Architecture adds each blade into a growing pool of resources. The SRX Series device can utilize these resources as necessary for optimal processing of traffic.

At the heart of the Dynamic Services Architecture is the switch fabric and SCB. The SCB transforms the chassis from a simple blade enclosure into a highly effective mesh network. The purpose of the SCB is to allow all blades in the chassis to send traffic at extremely high bandwidth.

The RE tightly couples with the functionality of the SCB and we can consider it the central nervous system of the architecture. The RE is the control plane of the chassis and provides overall management and communications to and from system administrators, as well as calculating route tables for routing network traffic. JUNOS Software, which includes key chassis functionality, also runs on the RE. In the case of networking and security, functionalities such as advanced routing, switching, flow-based security, zone-based management, and screens are available in the software.

*Continued on next page.*



**SRX Series Architecture (contd.)**

If the RE is the central nervous system of the chassis, the SPC is the brain. SPCs are blades that provide the capacity to perform the heavy lifting of processing network packets. The chassis must have at least one SPC to operate.

The chassis slots are card-agnostic, allowing you to configure the architecture for their specific needs up to the limits of the chassis itself. Based on the agnostic design, the IOCs can scale independently.

Not for Reproduction

## Routing Engine Overview

- JUNOS Software resides in flash memory
  - Backup copy available on hard disk
- Provides forwarding table to the PFE
  - Not directly involved with packet forwarding
  - Runs various routing protocols
- Implements CLI
- Manages PFE

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-38

### JUNOS Software

The primary copy of JUNOS Software resides on the flash memory of the JUNOS device. A backup copy is available on the hard disk when you issue a **request system snapshot** command.

### Routing Engine Intelligence

The RE handles all the routing protocol processes as well as other software processes that control the interfaces on the device. It also handles a few of the chassis components, system management, and user access to the device. These routing and software processes run on top of a kernel that interacts with the PFE. JUNOS Software directs all routing protocol packets from the network to the RE.

### Command-Line Interface

The RE provides the command-line interface (CLI). The CLI runs on top of the kernel; the management process (mgd) controls it.

*Continued on next page.*

## Packet Forwarding Engine Management

The RE controls the PFE by providing an accurate and up-to-date forwarding table and by downloading microcode and managing software processes that live in the PFE's microcode. The RE receives hardware and environmental status messages from the PFE and acts upon them as appropriate.

Not for Reproduction

## 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
- J Series uses real-time threads
  - Like ASICs, real-time threads are unimpeded simultaneous functions in other real-time threads

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-40

### Custom ASICs

ASICs enable the router to achieve data forwarding rates that match current fiber-optic capacity. The router achieves these rates by distributing packet processing tasks across highly integrated ASICs. As a result, ASICs-based JUNOS platforms do not require a general purpose processor for packet forwarding. The custom ASICs provide enhanced services and features, such as multicast, class of service (CoS) and queuing, and firewall filtering in hardware so that you can enable services on production devices without concern of significant performance hits. Each ASIC provides a piece of the forwarding puzzle, allowing a single ASIC to perform its specific task optimally.

### Real-Time Threads

Real-time operating systems give guaranteed processor cycles to each real-time thread. Careful designers can guarantee that the processor never becomes oversubscribed. As a result, these threads have the same benefits as ASICs, albeit at a lower total performance capability.

## PFE Components: ABC Chipsets

- The system midplane
- System control names vary by platform
- 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 packets and control packets
- Flexible PIC Concentrators
- Physical Interface Cards

### The Midplane

The system midplane is the component of the PFE that distributes power and electrical signals to each card in the system. Typically the midplane is passive in JUNOS platforms.

### System Control Boards

On M Series platforms, the System Control Boards provide the route lookup component of the PFE using the Internet Processor II ASIC. Each System Control Board on M Series routers provides the same function, despite each having different names. On the M7i and M10i, the FPC and Control Board components combine onto a single board named the Compact Forwarding Engine Board (CFEB). On the M7i and M10i, the Fixed Interface Card (FIC) or High-Availability Chassis Manager (HCM) perform chassis control functions, such as PIC online and offline and chassis monitoring, respectively. The M Series System Control Board—FEB, CFEB, or Switching and Forwarding Module (SFM)—also houses the buffer management ASICs on all models.

*Continued on next page.*

## General System Board Functions

The system board functions of JUNOS platforms deploying ABC chipsets include forwarding table updates and route lookup, system control, PFE clock generation, exception packet and control packet handling, and environmental monitoring:

- *Route lookups and forwarding table maintenance:* The Internet Processor ASIC performs route lookups using a forwarding table stored in the chip's SSRAM. The System Board updates its copy of the forwarding table when instructed by the JUNOS Software kernel.
- *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 RE, which sets the appropriate alarm. For example, if a temperature sensor exceeds the first internally defined threshold, the RE issues a `high temp` alarm. The System Board handles the power on and power off of PFE components with diagnostic errors reported to the RE over the 100 Mbps fpx1 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.
- *Transfer of exception and control packets:* The Internet Processor ASIC passes exception packets to a microprocessor on the System Board, which processes almost all of them. JUNOS Software sends the remaining packets to the RE for further processing. If the System Board detect errors originating in the PFE, the software logs them and makes them available to the CLI.

## Flexible PIC Concentrator

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.

## Physical Interface Cards

Juniper Networks M Series routers provide a complete range of fiber-optic and electrical transmission interfaces to the network through a variety of PICs. These space-efficient modules offer exceptional flexibility and high port density.

## PFE Components: LMNR Chipsets (1 of 2)

- 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
- SIBs perform switching between PFEs
  - Three SIBs comprise a T320 switch fabric—two active and one spare
  - Five SIBs comprise the T640 switch fabric—four active and one spare
- The system midplane

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-43

### Physical Interface Cards

As with the M Series routers, PICs provide T Series PFEs with a large range of fiber-optic and electrical transmission interfaces to the network.

### The T Series PFE

T Series routers implement either one or two complete PFE complexes on each FPC. On the T640 a single PFE is present on FPC2 while two PFEs are present on FPC3. We designed the latter FPC type specifically for native T Series PICs. Packets that ingress and egress on the same PFE complex (for example, on PICs 0 and 1 or PICs 2 and 3 of a given FPC) do not leave that PFE. The SIBs switch packets between PFEs across the T Series switch fabric as needed.

*Continued on next page.*

## The T Series Switch Fabric

T Series routers make use of a shared memory switch fabric for communications between and across FPCs and PFEs. In addition, inter-FPC communications require transit of the T Series crossbar switch fabric, which the system's SIBs represent. The T320 can support up to three SIBs, while the T640 supports five. In the case of the T640, four SIBs provide the necessary speedup for a nonblocking architecture. The fifth SIB comes into use only 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 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 has only one high-speed line to SIB 0 (two high-speed lines interconnect the FPCs 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 the switch fabric.



## PFE Components: LMNR Chipsets (2 of 2)

- **M320 FPCs contain one or two PFE complexes**
  - PFEs interface to other PFEs through the switch fabric
    - Nonblocking crossbar switch matrix with high-speed lines to each FPC
    - Switch fabric redundancy
  - SIBs perform switching between PFEs
    - Up to four SIBs comprise the M320 switch fabric
    - Adding SIBs enables line-rate performance for more or higher-capacity PICs
- **M120 FPCs translate packets from PICs to FEBs**
  - FPCs can map to any FEB
  - Control Boards carry SIB functionality
  - Data path is PIC > FPC > FEB > CB > FEB > FPC > PIC

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-45

### The M320 FPC and Switch Fabric

As with the T Series platforms, the M320 FPCs contain from one to two complete PFE complexes. As was also the case with T Series, packets that ingress and egress on the same PFE complex (for example, on PICs 0 and 1 or PICs 2 and 3 of a given FPC) do not leave that PFE. The SIBs switch packets between PFEs across the crossbar switch fabric as needed.

The M320 also makes use of a shared memory switch fabric for communications between and across FPCs and PFEs. In addition, inter-FPC communications require transit of the T Series-based crossbar switch fabric, which the system's four SIBs form. 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 a detailed performance breakdown). The M320 platform can operate with a SIB in standby mode for fault tolerance, but it does not have the space 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 performance declines until you replace the affected SIB.

### The M120 FPCs and Switch Fabric

Each M120 FPC contains a translator, a crossbar connection to the FEBs, power subsystem, and the physical PIC connectors. The translation component converts midplane signals to signals required by the types of supported PICs. The FEBs then, if necessary, switch the packet through the Control Board (CB) switch fabric to get to a different FEB.

## 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



© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-46

### Internet Processor II

The Internet Processor ASIC, which first shipped with the M40 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 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. 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.

All T Series routers make use of the latest Internet Processor ASIC technologies. In fact, a T640 might contain as many as 16 Internet Processor II chips because on a T Series device, each FPC can contain from one to two complete PFE complexes, and each PFE receives service from its own Internet Processor II ASIC.

On ABC chipset platforms, the C chip—actually the Cf chip (f is for *filtering*)—is the Internet Processor II chip. On LMNR chipset platforms, the R chip is the Internet Processor II chip.

## I-Chip

- I-chip is the Juniper Networks next-generation Internet processor
  - Inherits all the features of the M320 chipset
  - Packet Forwarding Engine on a chip
- New product features and capabilities:
  - Scaling enhancements
    - Larger number of logical interfaces
    - Route lookup and next hop
    - Interface accounting
  - Enhanced QoS capabilities
    - Delay-sensitive traffic
    - Multiplay networks
  - Improved multicast performance
  - Additional microcode capacity



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

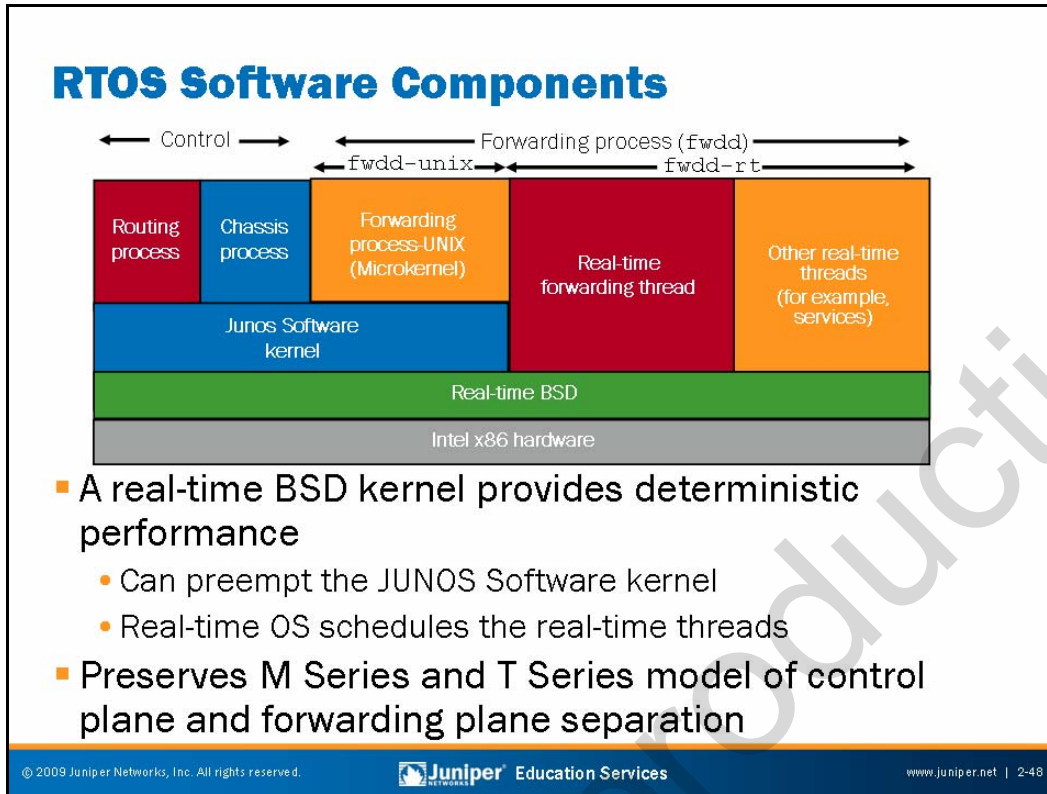
www.juniper.net | 2-47

### I-Chip Overview

Inheriting all the features of the LMNR chips, the I-chip is the Juniper Networks next-generation Internet processor delivering PFE on a chip. Each I-chip can send 20 Gbps of bandwidth to the fabric and can receive 17 Gbps of bandwidth from the fabric. The I-chip's flexibility comes through programmability, a rich instruction set, and silicon development. Various Juniper Networks platforms deploy the I-chip, including the M120, the MX Series, and the SRX Series.

### New Capabilities

The I-chip provides industry-leading scalability, allowing significant headroom in multiple dimensions, including VLANs, logical interfaces, routes, counters, number and size of firewall filters, and policing and shaping technologies. It provides complete control of any traffic management attributes of a packet, allowing manipulation of QoS attributes in a very sophisticated fashion. Layer 2 and Layer 3 classification can mix on the same physical port. This flexibility provides intelligent end to end QoS, because network elements in other segments might use different markings to determine classification. The data structures used in the I-chip allow it to scale multicast traffic at port speed without compromising performance.



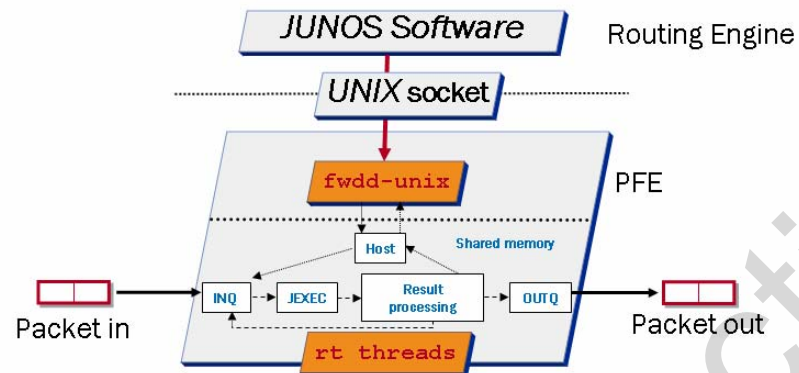
### Deterministic Performance

The J Series Routing Engine and software PFE both implement on the primary x86 architecture microprocessor. A real-time operating system kernel mediates access to the underlying hardware. The real-time kernel ensures that operating system services delivery occurs in a constant, load-independent, amount of time. This process ensures that the forwarding and services real-time threads deliver predictable packet forwarding performance.

### Control and Forwarding Separation

Separate real-time processes maintain logical separation between the control plane and forwarding plane. Control plane processes continue to run on the traditional JUNOS Software kernel that is a client of the real-time kernel. Forwarding and services threads run directly on the real-time kernel.

## The Net Result: A Virtualized PFE



- The `fwdd-unix` process emulates the microkernel found in M Series or T Series PFEs
  - Allows use of existing application programming interfaces between the RE and the forwarding process
  - RE communicates to `fwdd-unix` through a UNIX socket

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 2-49

### J Series Virtual Packet Forwarding Engine

The J Series software PFE maintains many of the benefits of the microkernel and ASIC-based PFE found on M Series and T Series platforms at a fraction of the cost. A UNIX socket provides the internal link between the RE and PFE and allows reuse of the JUNOS Software control plane from the M Series and T Series platforms on the J Series platform.



## 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\***
  - T640: 80+ Gbps
  - T320: 40+ Gbps
  - M7i, M10i, and M40e: 6.4 Gbps per FPC

\* The numbers quoted are two times the unidirectional (simplex) capacity of each FPC.

© 2009 Juniper Networks, Inc. All rights reserved.
 Juniper Education Services
www.juniper.net | 2-50

### 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; it does not require any specific order. 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 OC192 interface on an M120 router goes into a Type 3 FPC, which has only one slot—such a high-speed interface consumes all available FPC bandwidth. Most FPCs support four PIC connectors; current exceptions are the T320 and the M320, which support two PIC connectors per Type 3 FPC, the M120 as mentioned, and the M40e Type 2 FPCs, which also support only a single PIC.

When you install an FPC into a running system, the FPC requests its operating software from the Routing Engine, runs its diagnostics, and enables its PICs on the FPC slot. FPCs are hot-swappable on all platforms except the M7i and the M10i because these routers have FPCs that combine with the system board components to create a CFEB. The CFEB on the M7i and the M10i is hot-insertable but not hot-removable.

Note that when you remove or install an FPC on an M Series router, the system must re-partition 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 removal or insertion of FPCs does not affect packet forwarding on other FPCs.

*Continued on next page.*

### General FPC Characteristics and Features (contd.)

A portion of the memory associated with each FPC pools 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 at least 100 milliseconds of delay buffer exists (for each transmit and receive, yielding a total of 200 milliseconds of delay buffering). Currently, the amount of memory present on a given FPC ranges from 256 MB on the M7i FPC to 1.2 GB on the T640 FPC3. In the latter case, this memory yields approximately 600 MB per PFE complex.

### Industry-Leading Throughput

ABC chipset-based routers have an aggregate slot throughput of 6.4 Gbps. LMNR chipset-based platforms increase aggregate slot throughput to a respectable 20 Gbps for the M120, 40 Gbps for the M320 and the T320, and 80 Gbps for the T640.

## MX Series High Density DPC Architecture

- Dense Port Concentrators:
  - Line rate connectivity to the switch fabric
  - 4 PFEs per DPC
  - Three DPC types:
    - Switching and routing
    - Switching and limited scaling for routing
    - Enhanced queuing

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | 2-52

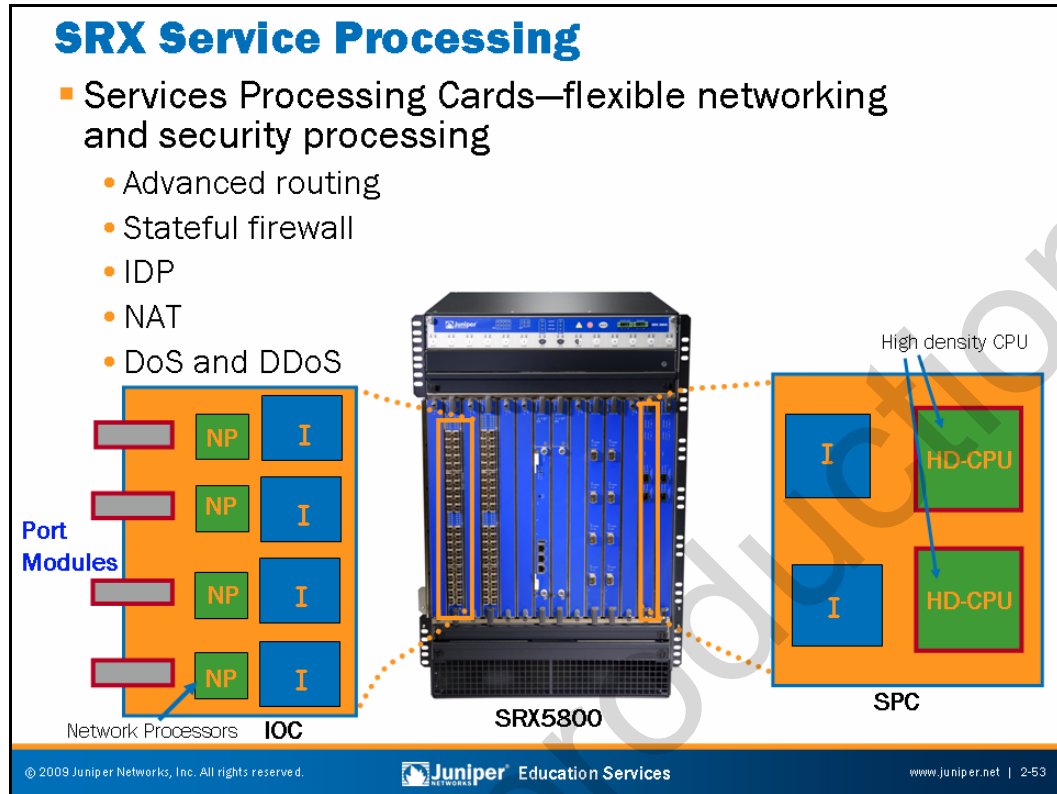
### MX Series Dense Port Concentrators

All DPCs come in either a small form-factor pluggable transceiver (SFP) or a 10-gigabit small form-factor pluggable transceiver (XFP), which uses the most optimal port density with cost efficiency. Each DPC has connections to the switch fabric providing line-rate connectivity for each port on the card.

DPCs provide multiple physical interfaces and PFEs on a single board that installs in a slot in the MX Series router. A DPC receives incoming packets from the network and sends outgoing packets to the network. Each DPC contains four PFEs. The PFEs on a DPC have purpose-built ASICs that perform packet processing and forwarding. Each PFE consists of one I-chip for Layer 3 processing and one Layer 2 network processor. Multiple PFEs contribute to the system's full packet forwarding redundancy and resiliency.

The three types of DPCs are switching and routing (DPCE-R), switching and limited scaling for Layer 3 (DPCE-X), and enhanced queuing (DPCE-Q). The DPCs support a wide range of Layer 2 and Layer 3 Ethernet functionality, including 802.1Q VLAN, link aggregation, circuit cross-connect, Virtual Router Redundancy Protocol (VRRP), Layer 2 to Layer 3 mapping, and port monitoring. Additionally, the DPCs support filtering, sampling, load balancing, rate limiting, class of service, and other key features necessary for deployment of dependable, high-performance Ethernet services.






### Service Processing Cards

SPCs are blades that provide the capacity to perform the heavy lifting of processing network packets. The chassis must have at least one SPC to operate. You realize the true elegance of this design when your system has more than one SPC installed. Rather than the chassis having two or more “brains”, as in traditional network architecture, the addition of a new SPC essentially results in a larger system that can perform many more tasks at a given time. To ensure the highest level of reliability, the SPCs and REs are separate—both physically and logically. This separation of the control and data planes ensures that a fault on any of the SPCs does not result in catastrophic failure of the entire chassis. You can see the importance of this concept in a security situation such as a denial of service (DoS) attack. When the attack launches, your efforts to contact the system do not simply become part of network traffic. Because the control plane remains separate from traffic flow, you can immediately respond to network-threatening situations to divert the attack, while all the SPCs continue to process network traffic.

## 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, security, and so on)
    - Services PICs normally have no physical ports
- Media-specific ASICs
- Status indicators
- Hot-swappable on most platforms



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 2-54

### 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 intermix within the same FPC. The number of ports on a given PIC varies with the PIC and platform type. For example, 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 over 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, yielding a service offering ranging from 1.5 Mbps through 12 Mbps (T1) or 2 Mbps through 16 Mbps (E1).

*Continued on next page.*

## Media-Specific ASIC

Each PIC has 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 field-programmable gate arrays (FPGAs)—that are specifically suited to the particulars of each medium.

## 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.

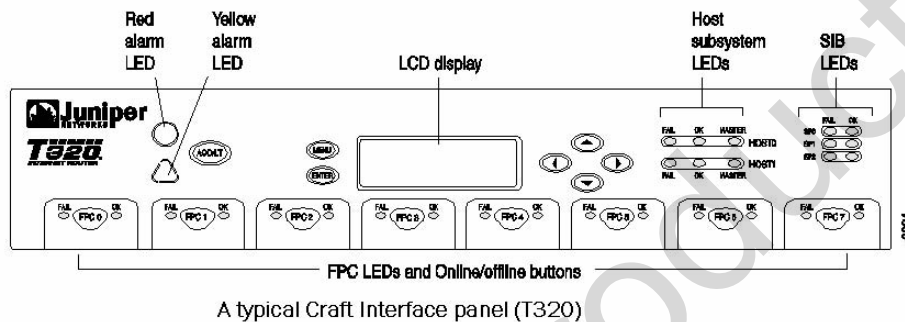
## Hot-Pluggable in Most Platforms

You can replace or install PICs without removing the associated FPC on most platforms.

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 Craft Interface—LCD Display Example

- Craft Interface overview:
  - LCD display (certain platforms only)
  - FPC online and offline buttons
  - PIC online and offline buttons
  - Status LEDs



A typical Craft Interface panel (T320)

© 2009 Juniper Networks, Inc. All rights reserved.

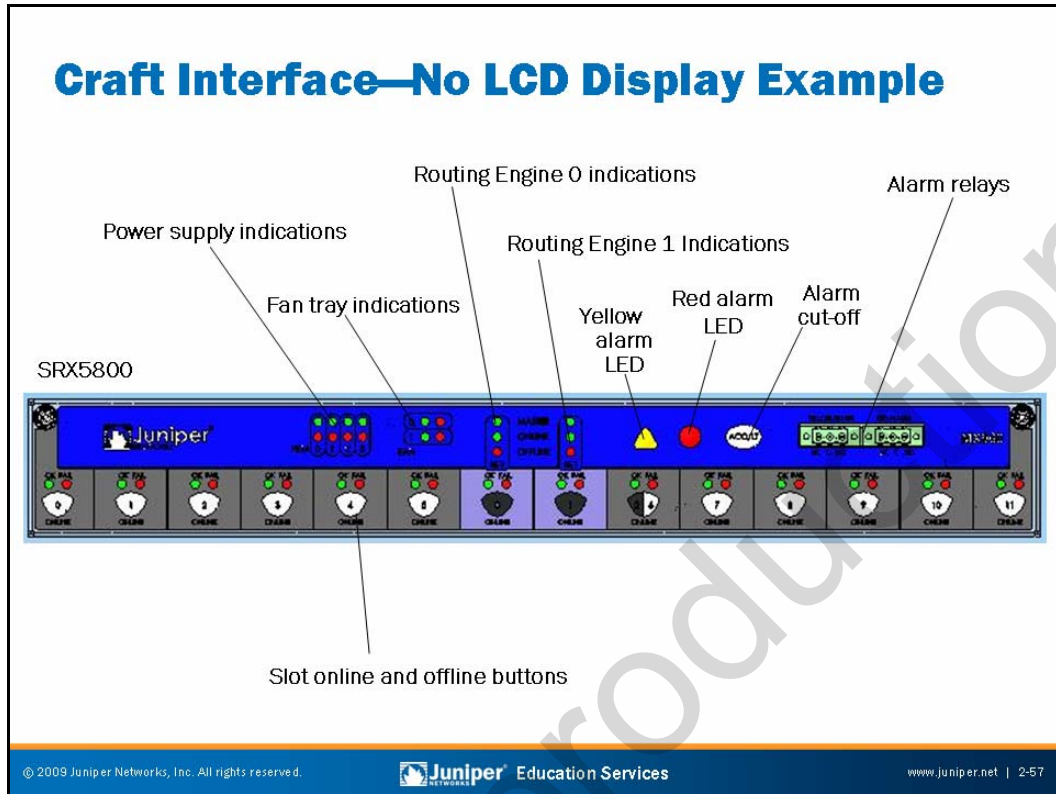
Juniper Education Services

www.juniper.net | 2-56

### Craft Interface

The Craft Interface is the collection of mechanisms on some JUNOS 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 and offline buttons. On supported platforms the Craft Interface includes an LCD screen that provides status reporting for the entire system.

The M7i's FIC and the M10i's HCM card provide PIC offline and online functionality.



### Craft Interface—No LCD Display Example

The slide illustrates a typical view of the SRX5800 Craft Interface. Although the Craft Interface does not have the LCD display, system component LEDs provide enough information to identify their status.

## Status LEDs and FRU Offline and Online

▪ **Status LEDs:**

- OK
  - Blinking = starting
  - Solid = running
- FAIL
  - Solid = taken offline because of failure

▪ **Online and offline buttons:**

- Press and hold for three seconds to take FPC (or PIC) offline

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | 2-58

### System Status LEDs

The system status LEDs include the following:

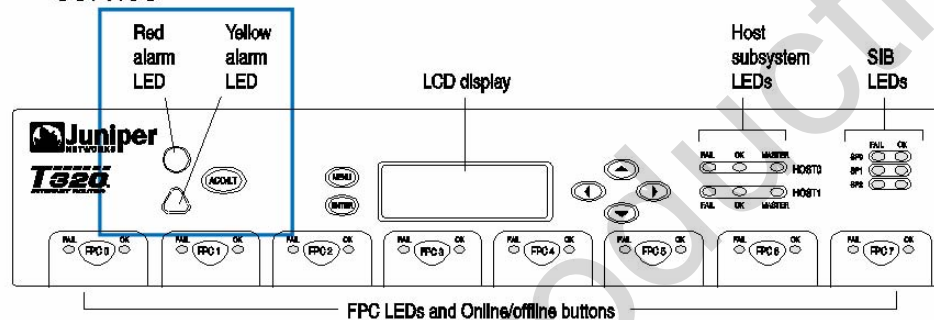
- *FPC LEDs:* Two LEDs exist—one green for OK and one red for fail. These lights indicate the status of each FPC. Each LED pair on the Craft Interface aligns with the corresponding FPC module slot.
- *Routing Engine or Host LEDs:* A red fail LED and a green OK LED on the Craft Interface indicate the status of the Routing Engines or 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 M7i and M10i, the online and offline buttons help prepare a PIC for removal from the system.

## Alarm Indications

- Red alarm:
  - Major failure that affects service and safety
- Yellow alarm:
  - Minor failure that needs attention but does not affect service



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 2-59

### Red Alarm

The red alarm LED indicates a system failure likely to cause an interruption in service. Examples of red alarms include the following:

- 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 include the following:


- Maintenance alert;
- FPC with recoverable errors; and
- Cooling system problems.

You can configure the mapping of various events to alarm actions of ignore, yellow, or red. Environmental and safety-related alarms are not mappable.



## The LCD Display

- LCD display is available only on M40e, M320, T320, and T640 routers
  - Displays general system status when no alarms are present
  - Displays alarm information when alarms are present
    - Identifies the total number and types of active alarms
  - Currently, the navigation buttons are only for obtaining 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 or *idle mode*, displays the current system status until *alarm mode* preempts it. The following list contains the basic status information in the LCD display:

- 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 configure 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.



## Agenda: Overview of JUNOS Platforms

- Overview of JUNOS Platforms
- Installation and Handling Guidelines
- Platform Architecture and Components
- Interface Overview

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-61

### Interface Overview

The slide highlights the topic we discuss next.

## Permanent Interfaces

- Device has several permanent interfaces:
  - Out-of-band management interface
    - Most platforms use fpx0
    - EX Series use me0
    - Smaller SRX platforms use ge-0/0/0
    - Requires configuration
  - Internal RE to PFE connection is named fpx1 or bcm0
  - Internal RE to RE connection is fpx2 or em0
    - Internal interfaces do not require any configuration; do not attempt to modify these interfaces!

### Permanent Interfaces

Each JUNOS platform has several permanent interfaces. One, the management Ethernet interface, provides an out-of-band method for connecting to the device. 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 out-of-band management interface requires configuration to operate. The names of out-of-band management interfaces vary depending on the platform. For example, most JUNOS platforms use fpx0 for out-of-band management, while the EX Series uses me0.

A second permanent interface provides internal Ethernet-based connectivity between the RE and the PFE. This interface is named fpx1 on most platforms. In the case of the M320, the interface operates at 1 Gbps and is named 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 cannot egress on a PFE port, and vice versa. This design isolates the out-of-band network from transient traffic and preserves bandwidth on the internal fpx1 or bcm0 control interface.


Platforms with redundant REs use a third interface (fpx2 or em0) to interconnect the REs for heart-beat and internal communications exchanges.

Internal interfaces like fpx1 and fpx2 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 the 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:

`et1-5/2/3:27` = A channelized DS3 interface in FPC slot 5, PIC position 2, port 3, and time-slot 27



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 2-63

### 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 device's transient interfaces. We refer to them as transient because you can hot-swap FPCs and PICs on most platforms at any time. From the point of view of the PFE, you can place any FPC into any slot, and you can generally place any combination of PICs in any location on an FPC. These characteristics makes PFE interfaces transient.

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

*Continued on next page.*

## 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 you install the FPC, the location of the PIC, and for some PICs, the port to which you connect. When dealing with a channelized PIC you must also reference the correct channel and time slot value using a `:xx` form of syntax.

## Interface Naming Example

The slide shows an example of a channelized interface name that makes use of the colon (:) delimiter to identify a time slot within the channel bundle.

Not for Reproduction

## Logical Units

- **Logical units are like subinterfaces in other equipment**
  - JUNOS Software requires a logical unit when configuring logical interface parameters
- **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 or HDLC encapsulations support only one logical unit**
  - Must configure unit number as zero for these encapsulations
- **A single logical unit supports multiple protocol addresses**
  - Typing in additional addresses does not override previous address
    - Watch for multiple addresses when correcting addressing mistakes!

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-65

### 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 referred to as 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 identifies a logical partition of the physical interface.

Although not required, we generally consider it best practice to keep the unit number and circuit identifier identical. 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.

*Continued on next page.*

### Addressing Issues

A JUNOS 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.

Not for Reproduction

## 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 associate 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

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 2-67

## 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 in the following list:

- 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;

*Continued on next page.*

### IPServices PICS (contd.)

- *sp*: Adaptive services interfaces; and
- *vt*: Virtual loopback tunnel interface.

Actual coverage of the services provided by these PICS is beyond the scope of this class.

Internally generated and nonconfigurable interfaces include the following:


- *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 FreeBSD continues to support it.



## Typical FPC and PIC Placement

Typical FPC and PIC numbering (T640)



- We identify transient interfaces according to FPC, PIC, and port convention
- FPC and PIC numbering varies by platform
- FPC slot and PIC port numbers have labels!

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 2-69

### Identifying Transient Interfaces

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, identify transient interfaces. Channelized interfaces identify a particular subchannel with the addition of a suffix in the form of :sub-channel-number. We identify a logical unit (also referred to as a subinterface) 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 graphic shows typical FPC and PIC numbering in the T640, which makes use of vertically aligned FPCs.

### FPC and PIC Slot Labeling

Each platform has labels that clearly identify the FPC slot number and PIC number. Furthermore, each PIC has a label to identify the number associated with that PIC's physical ports.

## Summary

- In this chapter, we:
  - Described current JUNOS platforms
  - Described general installation procedures
  - Explained the architecture of JUNOS platforms
  - Described the function of the RE, FPCs, PICs, SCBs, and Control Boards
  - Described the operation of the Craft Interface
  - Described interface naming conventions and the purpose of logical units

### This Chapter Discussed:

- Juniper Networks platforms running JUNOS Software and their typical applications;
- General installation procedures;
- General platform architecture;
- The functions of major components of the platforms;
- Operation of the Craft Interface; and
- Interface naming conventions and the role of logical units.

## Review Questions

1. How do you safely power off an M Series router?
2. What are the primary responsibilities of the RE and the PFE?
3. How can the Craft Interface assist in troubleshooting?
4. Describe what each field means in the interface name at-0/1/1.100.

## Review Questions

- 1.
- 2.
- 3.
- 4.

Not for Reproduction



# **Troubleshooting JUNOS Platforms**

## **Chapter 3: Troubleshooting Tool Kit for JUNOS Platforms**

Not for Reproduction

## Chapter Objectives

- After successfully completing this chapter, you will be able to:
  - Explain why some of the information in this chapter can be disruptive to a production network
  - Describe the layered troubleshooting methodology
  - Use various troubleshooting tools
  - Explain JTAC recommendations for current best practices that promote troubleshooting

### This Chapter Discusses:

- Warnings and caveats regarding potentially disruptive commands and techniques;
- Layered troubleshooting methodology;
- Various troubleshooting tools supported by JUNOS Software; and
- Juniper Networks Technical Assistance Center (JTAC) recommended configuration settings for ease of troubleshooting.

## Agenda: Troubleshooting Tool Kit for JUNOS Platforms

- Caveats and Warnings
- Troubleshooting Methodology
- Troubleshooting Tools
  - The JUNOS Software CLI
  - The Craft Interface Panel
  - System Logs and Protocol Tracing
  - Interactive UNIX Shell
  - Core Files for Diagnostic Analysis
  - The JTAC Knowledge Base
- Best-Practices Case Study

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-3

### Caveats and Warnings

The slide lists the topics we cover in this chapter. We discuss the highlighted topic first.

## Caveats and Warnings

- Some of the information presented in this chapter could disrupt a production network if used in the wrong way
  - Working in a UNIX shell is potentially dangerous and Juniper Networks does not officially support doing so
    - Only perform escape to a shell under the guidance of JTAC
  - Connecting to PFE components and executing arbitrary commands can be disruptive
  - Hidden commands are hidden for a reason
    - Some commands have the potential to disrupt your system, which is why they are hidden
    - Use hidden commands only under the guidance of JTAC

### Potential Disruptions

Troubleshooting a communications network involves the use of numerous tools. As with all tools, a potential exists for disruption or harm should you not use a given tool in the manner for which the designers created it.

In this chapter we discuss several potentially dangerous commands and troubleshooting actions that only qualified technicians who understand the potential ramifications associated with these commands and techniques should undertake. We advise you to pay special attention to the warnings and cautions we provide in this chapter.

The goal of this chapter is to prepare you to work with a JTAC engineer, who, in an effort to resolve a problem, might request that you perform one or more of the actions we document.



## When in Doubt: RTFM

- The materials in this course are somewhat general so as to appeal to a wide audience
  - You should always plan on *Reading the Fine Manual* that relates to your hardware platform and software release
    - <http://www.juniper.net/techpubs/index.html>
- Make use of the network operations guides
  - <http://www.juniper.net/techpubs/software/nog/index.html>
  - Focused on operation and troubleshooting
    - Interfaces, chassis, and baseline operations

### Generalized Content

In an effort to appeal to the wide range of customers that deploy, operate, and troubleshoot JUNOS platforms, the materials in this course are somewhat generalized. We always recommend that you consult the specific documentation for your particular hardware platform and software release before taking any specific actions. You should always defer to the specifics documented in a particular manual in the event of a conflict between the information presented in this course and that found in your manuals.

### Use the Network Operations Guides

The Juniper Networks Technical Publications group has prepared a series of operations guides to assist you with day-to-day operation and troubleshooting of JUNOS platforms. These guides provide operational information helpful for the most basic tasks associated with running a network using Juniper Networks products. The guides do not directly relate to any particular release of JUNOS Software and make excellent reference companions to this course. The material in this course augments and expands upon the information contained in these operator guides.

## Agenda: Troubleshooting Tool Kit for JUNOS Platforms

- Caveats and Warnings
- Troubleshooting Methodology
- Troubleshooting Tools
  - The JUNOS Software CLI
  - The Craft Interface Panel
  - System Logs and Protocol Tracing
  - Interactive UNIX Shell
  - Core Files for Diagnostic Analysis
  - The JTAC Knowledge Base
- Best-Practices Case Study

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-6

### Troubleshooting Methodology

The slide highlights the topic we discuss next.

## General Troubleshooting Tips

- Start with a visual inspection
  - Check power, grounds, connections, and configurations
- You must know what constitutes *normal* for your system
- Verify the problem before attempting repair action
- A divide-and-conquer approach is ideal when multiple faults can lead to a common symptom
  - Reduce the system to the minimum components necessary for testing
- Failure hypotheses should be testable—be definitive about what you are or are not testing with a given test
  - Each test should reduce the number of possible causes for the problem regardless of pass or fail status
- Do not be blinded by subjectivity—keep an open mind

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-7

### Begin with a Visual Inspection

The slide provides a few general troubleshooting tips. For example, it is generally a good idea to begin hardware or platform troubleshooting with a visual inspection. This approach uses the *keep it simple* philosophy of life. If you happen to notice a black smear that is indicative of smoke or fire damage near a component, you have most likely brought yourself closer to the source of a problem with little effort.

### Know What Constitutes Normal Status

It might seem pretty basic, but how can you spot signs of anomalous behavior if you are not confident of what behavior you expect in the first place? Put another way, how can you know if 30% CPU utilization on a system's Control Board is a sign of a problem, or an indication of normality, if the first time you display the component's CPU usage is during a troubleshooting operation?

### Always Confirm the Symptom

Many problems are transient by nature, and in some cases, testing causes more disruption than the problem itself. If a transient condition has already cleared, conducting disruptive testing benefits you very little. It is better to plan on long-term monitoring with testing occurring when the problem next manifests.

*Continued on next page.*

## The Art of War: Divide and Conquer

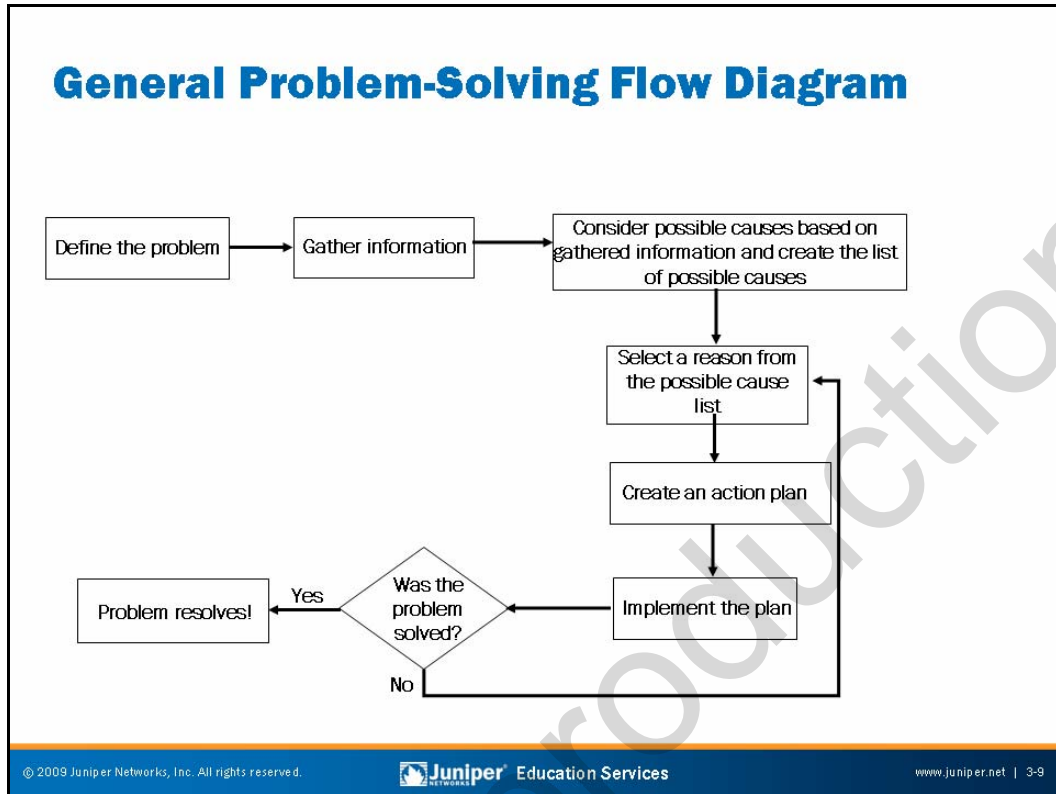
Over 2,500 years ago, Sun Tzu wrote a book named *The Art of War*, in which he told us to divide and conquer the enemy. This general approach works well when troubleshooting a problem that is generic enough to have numerous possible causes. In many cases you get closer to the real cause of a problem when you can effectively eliminate things that are not causing the problem. For example, if you do not need a joy-stick card to boot a PC, and the PC does not boot, then perhaps you should start by removing such unnecessary components for a successful boot.

## Each Hypothesis Should Be Testable

It does little good to dream up possible causes for a problem if you cannot definitively test whether the hypothesis is valid. You should try to formulate possible causes that, when tested, tend to eliminate possible causes for the problem, regardless of the actual outcome of the test. For example, conducting a local loopback on an interface eliminates the transmission line as a possible cause when the test fails. At the same time, this test eliminates the interface as a possible cause should the test succeed.

## Open Your Mind

Operators often overlook a potential source for a problem because of their subjective experiences. While leveraging your memory and past actions against a current problem is a good thing, you should never close your mind to new possibilities.



### General Problem-Solving Flow Diagram

Before embarking on your troubleshooting effort, be sure to have a plan in place to identify potential problems, isolate the likely causes of those problems, and then systematically eliminate each potential cause.

This page presents a general problem-solving flow diagram that you might want to follow during your troubleshooting. Although the presented diagram is not a rigid cookbook for troubleshooting, you can use it as a foundation from which you can build more detailed problem-solving plans.

## A Layered Troubleshooting Approach

- Designers model modern communications networks around layered architectures
  - Each layer depends on the services of the underlying layers
- Matching a symptom to the root-cause layer is a critical step in rapid diagnosis and restoration
  - Numerous failure scenarios might result in a common symptom like no route to the remote host
  - Allows escalation and hand-off to the appropriate group
- Identify the specific fault
  - Problem resolution is essentially fault confirmation and root-cause layer determination

### Modern Communications Networks Are Layered

Modern communications networks are complex. In 1977, the International Standards Organization developed a standard way of viewing these functions in the form of the Open Systems Interconnection (OSI) model. While the specifics of the OSI model are now more or less irrelevant given that TCP/IP is generally favored, the concept of a layered communications architecture is still quite valid.

Understanding the role that each layer plays and how each layer depends upon the services of the layers that lie below it, can greatly simplify the task of locating the elusive possible cause of problems. Put simply, it is a waste of time to troubleshoot a failed Layer 3 connectivity when the Link Layer protocol (Layer 2) running over that circuit is in a down state because the underlying Physical Layer is experiencing a loss of light alarm.

### Matching Symptoms to the Root-Cause Layer Is Job Number 1

A chain is only as strong as the weakest link, and so, too, is a layered communications system. The net result is that many common symptoms, for example, no route, can tie to failures that can occur at numerous layers. In these cases, you must question whether the route is missing because of a Physical Layer fault, a malfunction of the Data Link Layer, a failed Layer 3 adjacency, or other network layer problem—or if it is an upper-layer problem like a policy that is rejecting the route in question.

*Continued on next page.*

## Matching Symptoms to the Root-Cause Layer Is Job Number 1 (contd.)

By conducting tests that accurately isolate a symptom to the root-cause layer, you ensure that the problem escalates (as appropriate) to the correct group, and you avoid wasting time testing layers that are not at fault.

### Identify the Specific Fault

Once you correctly identify the root-cause layer, the next step is to isolate the problem at that layer so you can take the appropriate corrective actions. For example, knowing that the issue relates to mismatched T1 and DS1 framing (Physical Layer) allows you to correct the problem by configuring both devices for compatible framing to actually resolve the fault.

## Layered Troubleshooting Case Study

- Symptom: No HTTP connectivity between subscriber sites
  - Identify the layers that might account for this symptom and indicate their scope on the diagram
  - Identify specific faults that could lead to the symptom at each identified layer

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | 3-12

### No HTTP Connectivity: A Rather Generic Symptom

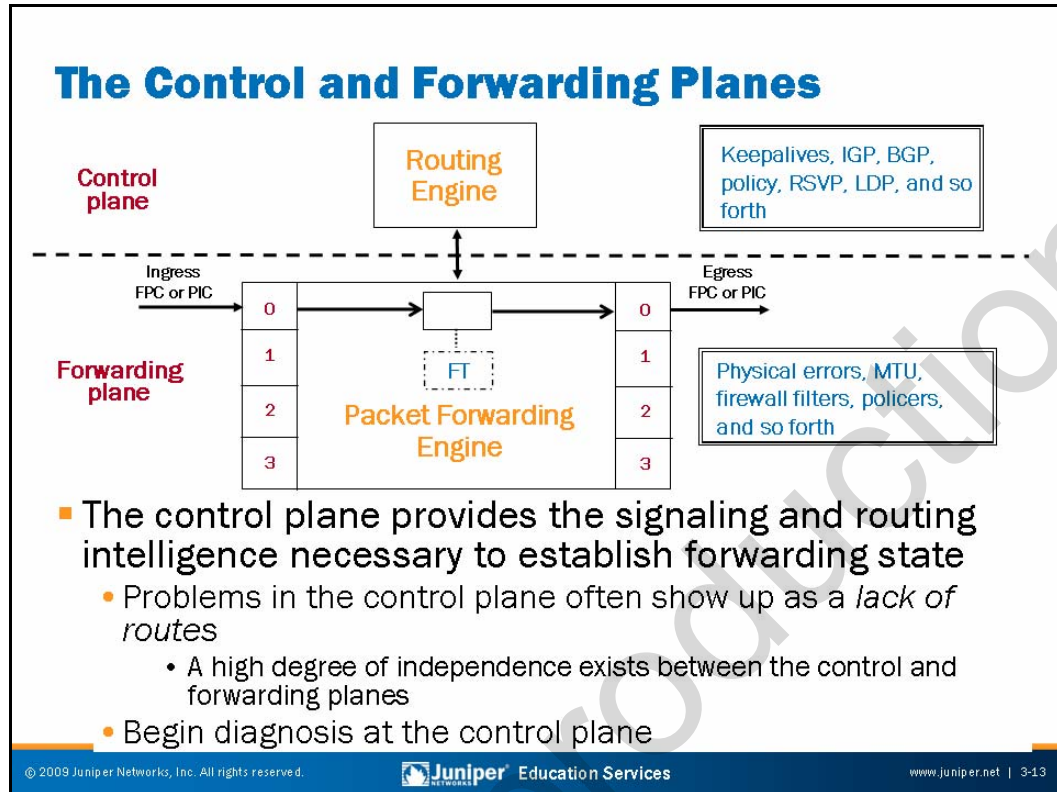
The slide helps illustrate the layered approach to troubleshooting by providing a typical communications topology and a rather generic symptom.

As far as what layers can account for this problem, the best answer is *all of them*. Specifically, a fault at the Physical, Data Link, Network, Transport, or Application Layers might exist.

Examples of possible faults and their scope include the following:

- *Physical Layer:* Broken wires or glass, power levels, framing, transmission line, or router or interface hardware could all be possible faults. This layer operates on a link-by-link basis.
- *Data Link Layer:* Mismatched framing, lack of keepalives, or invalid connection identifiers (data-link connection identifiers [DLCI] or virtual channel identifiers [VCI]) could all be possible faults. This layer operates on a link-by-link basis.
- *Network Layer:* Incompatible addressing, subnet masks, filters, or interior gateway protocol (IGP) parameters that prevent adjacency formation could all be possible faults. This layer operates end to end involving both routers and end systems (hosts).
- *Transport Layer:* Invalid ports, maximum transmission unit (MTU), lack of related service (Hypertext Transfer Protocol process not running), or authentication could all be possible faults. This layer operates end to end and involves only end systems.





### Understanding Control and Forwarding Plane Separation

When troubleshooting JUNOS platforms, you must understand the separation of the control and forwarding planes, regardless if the separation occurs in hardware or software. Generally speaking, problems with a routed network come down to either a control plane issue or a forwarding plane issue. It is extremely rare to find a fault in both planes simultaneously because of the completely different role that each plane plays.

The control plane primarily deals with the installation of routes in the forwarding table. This function relies on routing protocols, configuration, authentication of routing peers, and so forth. The most common symptom of a control plane problem is the lack of one or more routes.

Once the software installs a route into the forwarding table, the forwarding plane of the platform simply uses that route as a next hop for matching traffic using a switching path. Problems in the forwarding plane tend to take the form of bad hardware (for hardware-based platforms), policers, or firewall filters that prevent or impair communications despite valid routes existing in the control plane. (We can argue that the last two items—policers and filters—are really control plane problems that manifest themselves in the forwarding plane.)

While application-specific integrated circuits (ASICs) and higher-end platform packet forwarding engines are complex, they tend to work. Thus, the majority of problems you encounter when troubleshooting high-end platforms relate to the control plane of the device, which is why the slide suggests that you begin fault analysis by examining the control plane first.

## Agenda: Troubleshooting Tool Kit for JUNOS Platforms

- Caveats and Warnings
- Troubleshooting Methodology
- Troubleshooting Tools
  - The JUNOS Software CLI
    - The Craft Interface Panel
    - System Logs and Protocol Tracing
    - Interactive UNIX Shell
    - Core Files for Diagnostic Analysis
    - The JTAC Knowledge Base
- Best-Practices Case Study

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-14

### Troubleshooting Tools: The JUNOS Software CLI

The slide highlights the topics we discuss next.

## The JUNOS Software CLI

- The JUNOS Software CLI:
  - A variety of operational mode commands report on hardware, software, and protocol status
    - Piped output accommodates searching, counting, and so forth on command output
  - Process restart and hardware restart
    - You can restart software processes
    - You can take FPCs and PICs online and offline
    - You can perform full commits
  - Hardware redundancy control
    - Switch mastership on redundant platforms
    - Log in to a backup RE using an internal communications path
  - Network utilities
    - Ping and traceroute utilities with a rich set of options
    - Telnet, SSH, FTP, and SCP
    - Monitor traffic (tcpdump)

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-15

### The JUNOS Software CLI

The JUNOS Software command-line interface (CLI) is the primary mechanism for troubleshooting and operational analysis. Using the CLI, it is easy to determine hardware, software, protocol, and general operational status. The following are some key CLI features:

- Support for piped output to functions like count or match for all commands and in all modes (configuration or operational mode);
- The ability to restart software processes and take hardware online or offline;
- The ability to control redundant hardware; and
- Various network utilities like ping and traceroute, and the ability to monitor local traffic in a manner similar to tcpdump.

## Key Operational Mode Commands

- Key operational mode commands include:
  - **show chassis**
    - alarms, cluster, environment, firmware, forwarding, fpc, hardware...
  - **show system**
    - alarms, statistics, storage, connections, users...
  - **show interfaces**
    - terse, detail, filters, policers...
  - **show route**
    - protocol, aspath-regex, community, hidden, resolution, advertising-protocol, receive-protocol, detail...
  - **monitor interface**
  - **monitor traffic**
- All commands and all CLI modes support piped output:
  - **count, match, find, except, request**, and so on

### Key Operational Mode Commands

Depending on the type of problem with which you are dealing, numerous JUNOS Software CLI commands might exist that can assist you in problem determination. The slide calls out the main classes of operational mode commands that prove particularly useful in most troubleshooting situations:

- The various **show chassis** commands are well suited to assisting you in performing operational and fault analysis of hardware-related issues;
- The family of **show system** commands are useful in detecting configuration and operational status of system protocols and users;
- The **show interfaces** commands are useful when your focus is on physical or link-level operational analysis, and when you suspect interface hardware-related faults.
- The **show route** commands are invaluable when testing the control plane to determine what routes are present, from where the router learned of them, and where they direct matching traffic;
- The **monitor interface** command provides detailed, real-time snapshots of the traffic patterns, error counts, and alarm status for the monitored interface; and
- The **monitor traffic** command makes tcpdump protocol analysis capabilities for local traffic available to the user.

*Continued on next page.*

### Enhanced Features Are Always Available

As noted previously, the CLI supports piped output to value-added features that make potentially arduous tasks, such as counting or comparing, really easy. Having these features supported in all modes, and for all commands, is a real plus!

Not for Reproduction

## Restarting a Software Process

- You can restart most software processes from the CLI
  - Restarting other processes requires escape to a shell

```

user@host> restart ?
Possible completions:
 802.1x-protocol-daemon  Port based Network Access Control
audit-process           Audit process
autoinstallation        Autoinstallation process
chassis-control         Chassis control process
class-of-service        Class-of-service process
database-replication    Database Replication process
dhcp                    Dynamic Host Configuration Protocol process
dhcp-service            Dynamic Host Configuration Protocol process
dialer-services         Dial-Out On Demand process
. . .
redundancy-interface-process  Redundancy interface management process
remote-operations        Remote operations process
routing                  Routing protocol process
. . .
user@host> restart routing
Routing protocols process started, pid 5042
    
```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-18

### Restarting Software Processes

You can restart most JUNOS Software processes from the CLI. This capability leverages the modular nature of JUNOS Software and avoids the need for a system reboot when a particular process encounters a problem.

Processes that are not listed in the CLI output, such as the `init` process (which is the meta-process that controls the starting of all other processes), require that you escape to a shell to restart them. It is also necessary to escape to a shell to pass the process a signal such as a `kill -1 (SIGHUP)`. The `kill -1` signal forces that process to reread its configuration file but does not terminate the process.

When restarting a process, the default behavior is a soft kill, or graceful shutdown, in which the process receives a signal that it should terminate but is given time to clean up its state first. In contrast, a hard kill is equivalent to issuing a `kill -9 pid`, in that it terminates the process immediately.

The `init` process restarts any process that has failed, so after killing a process, a new instance of that process starts. However, if a process fails repeatedly in rapid succession, the `init` process disables it to prevent thrashing. Once `init` disables a process, you must reboot, or force `init` to reread its configuration before it allows that process to restart. Issuing a `commit` with the hidden `full` switch passes the `init` process a `SIGHUP` that causes it to restart all configured processes, regardless of previous thrashing behavior. However, if the process still thrashes, `init` disables it.

## Bouncing an `rpd` Component

- The routing protocol process currently handles all routing protocols
  - Bouncing `rpd` with a `restart routing` disrupts *all* `rpd` processes
  - Use `deactivate` to bounce a specific `rpd` component; the example bounces BGP while leaving OSPF untouched:

```
[edit]
user@host# show protocols
bgp {
  group test {
    vpn-apply-export;
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-0/3/0.0;
    interface at-0/1/0.100;
    interface so-0/2/0.0;
  }
}
. . .

[edit]
user@host# deactivate protocols bgp

[edit]
user@host# commit
commit complete

[edit]
user@host# rollback 1
load complete

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

© 2009 Juniper Networks, Inc. All rights reserved.



www.juniper.net | 3-19

### Bouncing a Component of `rpd`

Currently, the routing protocol process (`rpd`) is responsible for handling all routing protocol functions. If you detect a problem in the OSPF protocol, for example, then a **restart routing** command might resolve the issue. The problem is that restarting routing affects all routing protocols, which include BGP, IS-IS, RIP, and so forth.

When the goal is to minimize overall disruption (which it always is), you might consider the technique shown on the slide, which involves deactivating a particular protocol, rather than restarting all routing functionality. The downside to this approach is that configuration privileges are necessary.

The example on the slide shows the operation bouncing BGP by deactivating the `bgp` stanza and issuing a **commit**. During the process, the OSPF protocol remains untouched and continues to operate as before. After the **commit** and a **rollback 1**, the user issued another **commit** that restored the `bgp` stanza to its previous (active) state. The BGP protocol now initializes, just as if you had restarted the `rpd` process. Rather than using the **rollback** function, you can also issue an **activate protocol bgp** command from the `[edit]` hierarchy, followed by a **commit** to achieve the same results.



## Full Commits

- Juniper Networks optimized the `commit` function
  - Goal is to avoid disruption to processes not affected by a configuration change
- The hidden `full` switch affects all processes
  - Forces reread of configuration, reactivating the entire configuration
  - An excellent way to restart a process that is disabled because of thrashing

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

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

Hidden switch

### Performing a Full Commit

JUNOS Software optimizes the process of committing a candidate configuration so it does not disrupt processes when their portion of the configuration has not changed. While a great idea to be sure, the situation is rare in which a particular process fails to wake up with a commit, and as a result, the modified configuration does not go into effect.

By including the hidden `full` switch, when issuing a `commit`, you force all processes to reread their configuration, which ensures the honoring of changes. A `commit full` also signals the `init` process with a `kill -1 (SIGHUP)` that forces it to reread its configuration.

### Shaking It Up

Because a full commit places a processing strain on a router with a complex configuration, you should only perform a full commit when conditions warrant.



## Hardware Restart

- You can restart FPCs and PICs or bring them offline or online using the CLI:

```

user@host> request chassis ?
Possible completions:
  cfeb          Change Compact Forwarding Engine Board status
  pic          Change Physical Interface Card status
  routing-engine Change Routing Engine status

user@host> request chassis cfeb ?
Possible completions:
  master      Set CFEB mastership
  offline     Take CFEB offline
  online      Bring CFEB online
  restart     Restart CFEB

user@host> request chassis cfeb offline
CFEB Offlined

user@host> show chassis alarms
2 alarms currently active
Alarm time      Class  Description
2009-01-08 17:40:40 UTC Major  CFEB not online, the box is not
forwarding
2009-01-08 16:47:22 UTC Minor  Host 0 Boot from alternate media

```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-21

### Hardware Restart

The slide shows how you can use the JUNOS Software CLI to take a Compact Forwarding Engine Board (CFEB) (in some models), Flexible PIC Concentrator (FPC), or PIC offline and online. In some cases, you can clear problems by *bouncing* a piece of hardware, which means taking the device offline and then bringing it back online again.

The commands shown on the slide have the same effect as if you depressed the CFEB offline button on the physical router to bring it offline.

## Hardware Redundancy

- Defaults and configuration determine hardware mastership
  - Control RE and system Control Board mastership from the CLI:
 

```
user@host> request chassis ?
Possible completions:
  cfeb          Change Compact Forwarding Engine Board status
  pic           Change Physical Interface Card status
  routing-engine Change Routing Engine status
user@host> request chassis routing-engine master ?
Possible completions:
  acquire      Attempt to become the master Routing Engine
  release      Request that the other Routing Engine become master
  switch       Toggle mastership between Routing Engines
```
- Log in to *other* RE using internal communications path:
 

```
user@host-re0> request routing-engine login rel
e
--- JUNOS 9.3R2.8 built 2008-12-17 23:25:33 UTC
user@host-rel>
```

  - Most designs require a common configuration
    - Use `commit synchronize` to mirror changes to backup RE

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-22

### Hardware Mastership

On platforms that support hardware redundancy, the determination of a component's status as either master or standby is a function of software defaults and explicit configuration. The slide shows that you can use the CLI to determine mastership status and to effect a change in status of a redundant component.

Although many hardware faults and some software faults trigger a mastership change automatically (when so configured), instances exist in which a marginal failure does not result in the affected components relinquishing their mastership role. In cases such as this one, or when you must perform routine maintenance on a redundant component, you might want to force a change in mastership by using the CLI. Note that depending upon what is being switched—Routing Engine (RE) versus system Control Board—and the specific configuration (such as graceful restart enabled), switching mastership status might result in a disruption to packet forwarding.

### Login to Other RE

On systems equipped with redundant REs, you can establish a login to the other RE using an internal communications path. In most cases, you should ensure that the software replicates configuration changes made on the active RE to the configuration file used by the backup RE. When you issue a `commit synchronize` command, the software uses the same internal path used for RE-to-RE logins to synchronize the configuration file to the backup REs.

## Network Utilities and Applications (1 of 3)

```

user@host> ping ?
Possible completions:
<host>
  atm                Hostname or IP address of remote host
                    Ping remote Asynchronous Transfer Mode node
  bypass-routing    Bypass routing table, use specified interface
  cls               Ping ISO node
  count            Number of ping requests to send (1..2000000000 packets)
  detail          Display incoming interface of received packet
  do-not-fragment  Don't fragment echo request packets (IPv4)
  inet            Force ping to IPv4 destination
  inet6          Force ping to IPv6 destination
  interface      Source interface (multicast, all-ones, unrouted packets)
  interval       Delay between ping requests (seconds)
+ loose-source   Intermediate loose source route entry (IPv4)
  mpls          Ping label-switched path
  no-resolve     Don't attempt to print addresses symbolically
  pattern        Hexadecimal fill pattern
  rapid         Send requests rapidly (default count of 5)
  record-route   Record and report packet's path (IPv4)
  routing-instance Routing instance for ping attempt
  size         Size of request packets (0..65468 bytes)
  source        Source address of echo request
  strict       Use strict source route option (IPv4)
+ strict-source Intermediate strict source route entry (IPv4)
  tos         IP type-of-service value (0..255)
  ttl        IP time-to-live value (IPv6 hop-limit value) (hops)
  verbose    Display detailed output
  vpls      Ping VPLS MAC address
  wait     Number of seconds to wait after sending last packet
  
```

Shaded options are particularly useful for fault isolation

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 3-23

### Network Utilities and Applications: Part 1

As you might expect, JUNOS Software supports standard network utilities like ping and traceroute. As shown on the slide (for the case of ping) these utilities support a rich set of optional switches that can prove especially useful when troubleshooting. The following are some of the key switches:

- **atm:** Generates special Asynchronous Transfer Mode (ATM) pings that use Operation, Administration, and Maintenance (OAM) cells.
- **count:** Limits the number of ping attempts.
- **do-not-fragment:** Useful in diagnosing MTU-related problems by preventing the fragmentation of large packets.
- **pattern:** By altering the payload of ping packets, you can detect error conditions that are triggered by data patterns.
- **record-route:** Allows you to trace the set of egress interfaces the packet encounters. Note that this process differs from traceroute, which displays the set of *ingress* interfaces.
- **routing-instance:** Use this switch to provide routing instance and virtual private network (VPN) context for a ping (or similar) command. By default, a command is issued in the context of the main routing instance unless you use this switch.

*Continued on next page.*

### Network Utilities and Applications: Part 1 (contd.)

- **size:** By altering the size of packets, you can detect MTU-related and capacity-related problems.
- **source:** This switch lets you control the source address placed in the resulting packet. This capability can help diagnose routing problems because you can make the packet appear to come from any address owned by the device (spoofing is not permitted).
- **tos:** This switch lets you alter the type-of-service (ToS) bits in the packet when testing a class-of-service (CoS) issue.

Not for Reproduction

## Network Utilities and Applications (2 of 3)

### ■ Telnet, SSH, and FTP support

- Ability to specify nonstandard ports and source address for Telnet:

```
user@host> telnet ?
Possible completions:
<host>           Hostname or address or remote host
8bit             Use 8-bit data path
bypass-routing   Bypass routing table, use specified interface
inet            Force telnet to IPv4 destination
inet6           Force telnet to IPv6 destination
interface        Name of interface for outgoing traffic
no-resolve       Don't attempt to print addresses symbolically
port            Port number or service name on remote host
routing-instance Name of routing instance for telnet session
source          Source address to use in telnet connection
```

- Use the `no-resolve` switch when Telnet sessions take a long time to establish because of name lookup failures
- You must enable the related service under the `[edit system services]` hierarchy to support incoming connections

## Network Utilities and Applications: Part 2

JUNOS Software offers support for Telnet, SSH or SCP, and FTP. As with the ping and traceroute utilities, these applications support switches that are useful in troubleshooting. The following are some of the key switches:

- **no-resolve:** This switch disables the normal reverse lookup performed on the host address specified in a `telnet` command. Use this switch when sessions take a long time to open because of the inability to perform the reverse lookup.
- **port:** The port switch allows you to specify a destination port other than the default port normally associated with that service.
- **routing-instance:** This switch supports VPN and routing instance context for applications like Telnet and FTP. A classic use would be to establish a Telnet connection from a provider edge (PE) router to an attached customer edge (CE) device, which, being part of a VPN, would reside in a specific routing table and instance.
- **source:** As with ping, altering the source address used in a connection request might uncover problems with routing that prevent connection establishment when sourcing traffic from the egress interface (the default).

## Network Utilities and Applications (3 of 3)

- The `monitor traffic` command provides CLI access to the `tcpdump` utility
  - Only displays traffic originating or terminating on the local RE
    - The best way to perform analysis of Layer 2 protocols with JUNOS Software
    - Protocol filtering currently requires writing and reading from a file (hidden `write-file` and `read-file` options)

```

user@host> monitor traffic interface se-1/0/0 detail
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on se-1/0/0, capture size 1514 bytes
. . . . .
02:18:43.121184 In IP (tos 0xc0, ttl 1, id 21998, offset 0, flags [none],
  proto: OSPF (89), length: 68) 172.18.36.2 > 224.0.0.5: OSPFv2, Hello, length 48
  Router-ID 192.168.36.1, Backbone Area, Authentication Type: none (0)
  Options [External]
  Hello Timer 10s, Dead Timer 40s, Mask 255.255.255.252, Priority 128
  Neighbor List:
    192.168.24.1
. . . . .
02:18:46.280403 Out LCP, Echo-Request (0x09), id 177, length 10
  encoded length 8 (=Option(s) length 4)
  Magic-Num 0x92da0b79
    
```

OSPF hello

ICMP ping traffic

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-26

### Network Utilities and Applications: Part 3

The `monitor traffic` command provides CLI-based access to the `tcpdump` utility. This command monitors only traffic originating or terminating on local the RE. This capability is the best way to monitor and diagnose problems at Layer 2 with JUNOS Software because `tracing`, which is similar to `debug` on equipment from other vendors, does not function for Layer 2 protocols. We cover tracing on subsequent pages that deal with system logging.

Note that protocol filtering functions (for example, matching on only UDP traffic sent from a specific port) are currently not supported for real-time monitoring because in real-time mode, the Layer 2 headers are stripped at ingress, which prevents filtering on protocol types. As a workaround, you can write the monitored traffic to a file and then read the file with a `tcpdump`-capable application like `Ethereal`. We provide an example of how to achieve protocol filtering with the JUNOS Software `monitor traffic` command in a subsequent case study.



## Agenda: Troubleshooting Tool Kit for JUNOS Platforms

- Caveats and Warnings
- Troubleshooting Methodology
- Troubleshooting Tools
  - The JUNOS Software CLI
  - The Craft Interface Panel
  - System Logs and Protocol Tracing
  - Interactive UNIX Shell
  - Core Files for Diagnostic Analysis
  - The JTAC Knowledge Base
- Best-Practices Case Study

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services


www.juniper.net | 3-27

### Troubleshooting Tools: The Craft Interface Panel

The slide highlights the topic we discuss next.

## The Craft Interface Display

- LED with LCD screen indicates system and hardware status
  - Can view remotely with a **show chassis craft-interface** command



FPC status

```
user@host> show chassis craft-interface
```

```
Red alarm: LED on, relay on
Yellow alarm: LED off, relay off
Routing Engine OK LED: On
Routing Engine fail LED: Off

FPCs 0
-----
Green *
Red .

LCD screen:
```

Red alarm active

```
Host
1 Alarm active
R: Supply A FAIL
```

### The Craft Interface

The craft interface panel for systems that support the LCD status screen is an excellent troubleshooting and operational analysis tool because it provides component and system alarm status in a manner that is easy to interpret. When working remotely you can issue a **show chassis craft-interface** command to obtain an ASCII representation of the LEDs and messages that the craft interface displays.



## Displaying a Message on the Craft Interface

- Can display a user-defined message on the craft interface panel's LCD screen
  - Useful for identifying the correct system when relying on *remote-hands*
  - User message alternates with normal display for five minutes
    - Use the `permanent` switch to prevent alternation of normal craft display
    - Maximum of four lines with a 20-character limit per line

```

user@host> set chassis display message "M320 unit for RE swap"
message sent
lab@host> show chassis craft-interface
Red alarm:          LED off, relay off
Yellow alarm:       LED off, relay off
. . .
LCD screen:        +-----+
                    | "M320 unit for RE |
                    | swap"           |
                    +-----+
  
```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-29

### Displaying Messages on the LCD Screen

Displaying messages on the craft interface panel's LCD screen can be helpful when you want to identify a system or communicate in some way with a person that is local to that machine. By default, the custom user message alternates with the normal LCD message display (system status messages that alternate every few seconds). Use the **permanent** switch with the **set chassis display** operational mode command to force only the display of the custom message.

Note that the custom message times out after five minutes, and the display returns to the default system status message rotation. This command is applicable only to platforms that have an LCD screen.

## Agenda: Troubleshooting Tool Kit for JUNOS Platforms

- Caveats and Warnings
- Troubleshooting Methodology
- Troubleshooting Tools
  - The JUNOS Software CLI
  - The Craft Interface Panel
  - System Logs and Protocol Tracing
    - Interactive UNIX Shell
    - Core Files for Diagnostic Analysis
    - The JTAC Knowledge Base
- Best-Practices Case Study

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-30

### Troubleshooting Tools: System Logs and Protocol Tracing

The slide highlights the topic we discuss next.

## Syslog and Tracing

- **Syslog:**
  - Standard UNIX syslog configuration syntax
    - Primary syslog file is /var/log/messages
    - Most processes also write to individual log files
  - Supports numerous facilities and severity levels
    - The facility defines the class of log message while the severity level determines the level of logging detail
  - Local and remote syslog support
    - We recommend remote logging (and archiving) for troubleshooting
- **Tracing decodes protocol packets and certain router events:**
  - Some other vendors refer to tracing as *debug*
  - Tracing operations include:
    - Global routing behavior
    - Router interfaces
    - Protocol-specific information

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-31

### Syslog

Syslog operations use a UNIX syslog-style mechanism to record system-wide, high-level operations, such as interfaces going up or down or users logging in to or out of the router. You configure these operations by using the **syslog** statement at the [edit system] hierarchy level and the **options** statement at the [edit routing-options] hierarchy level.

The results of tracing and logging operations go in files that the router stores in the /var/log directory. You use the **show log file-name** command to display the contents of these files.

### Tracing Operations

Tracing operations allow you to monitor the operation of routing protocols by decoding the sent and received routing protocol packets. In many ways, tracing is synonymous with the debug function on equipment made by other vendors. Note that because of the design of some hardware-based Juniper Networks platforms, you can enable reasonably detailed tracing in a production network without negative impact on overall performance or packet forwarding.

## Syslog Configuration Example

```

[edit system syslog]
user@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;
}
    
```

*Annotations on the slide:*

- Explicit priority support:** points to `explicit-priority;` in the `file errors` block.
- Comments:** points to the multi-line comment `/* send all error messages to file "errors" with explicit priority */`.
- The log file name:** points to `file security`.
- The level at which to begin logging:** points to `authorization info;` in the `file security` block.
- The syslog facility:** points to `interactive-commands info;` in the `file security` block.
- Archive and permission settings for the messages file:** points to `archive size 10m files 20 no-world-readable;` in the `file messages` block.

© 2009 Juniper Networks, Inc. All rights reserved. www.juniper.net | 3-32

### Syslog Options Example

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

- **archive:** Configures 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 that designers developed to parse the logs generated by equipment from other vendors.

*Continued on next page.*

**Syslog Options Example (contd.)**

The following table illustrates the mapping of numeric codes to message severity:

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 disk 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

The following are examples of a syslog message, both with and without an explicit priority, respectively:

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

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

## Process and Miscellaneous Log Files

- **Key process and miscellaneous log files include:**
  - `apspd`: Automatic protection switching process
  - `bfd`: Bidirectional failure detection process
  - `chassisd`: Chassis management process
  - `commits`: Records the username and date for each commit
  - `cosd`: Class of service process
  - `dcd`: Device control process
  - `eccd`: Error checking and correction process
  - `mastership`: Records changes in hardware mastership
  - `sampled`: Sampling process (cflowd)
  - `snmpd`: SNMP process
  - `vrp`: Virtual Router Redundancy Protocol process
- **Entries also write to the messages file**
  - Parsing a process log file helps reduce clutter when you are dealing with a particular process

© 2009 Juniper

www.juniper.net | 3-34

### Process and Miscellaneous Log Files

The primary system log file is the messages file. However, some of the processes that run under JUNOS Software maintain their own log files named after their respective process. No requirement exists to configure the router to keep these logs. Note that in many cases, the software also writes the entries found in these logs to the main messages file. Key process log files include the following:

- `apspd`: The automatic protection switching process handles events relegated to SONET Automatic Protection Switching (APS). View this log when you are dealing with an APS issue.
- `bfd`: The bidirectional failure detection process functions to provide rapid detection of failures in the forwarding plane to expedite routing protocol convergence.
- `chassisd`: The `chassisd` process is responsible for monitoring and managing the hardware present in the physical router chassis, including ASICs, power supplies, fans, and temperature sensors, as well as managing hot-swap events.
- `commits`: This log file records the commit activities on the router in the form of date and time, user, and mode.
- `cosd`: The class of service process monitors class-of-service events in the chassis.

*Continued on next page.*

## Process and Miscellaneous Log Files (contd.)

- `dcd`: The device control process communicates with the Packet Forwarding Engine (PFE) to track the status and condition of the router's interfaces. The `dcd` configures interfaces on the basis of information in the configuration file and the hardware present in the device. You can configure physical interfaces before the hardware is present; likewise, a router can contain unconfigured FPCs and PICs. Check the `dcd` log for interface-related entries when troubleshooting interface problems.
- `eccd`: The error correction control process deals with memory errors. If you suspect bad or failing memory, check this log.
- `mastership`: The `mastership` log records events related to hardware redundancy.
- `mgdd`: The management process controls the CLI process. No log file associated with this process exists.
- `sampled`: The sampling process handles tasks related to packet sampling. Check this log when troubleshooting or monitoring a sampling configuration.
- `snmpd`: The SNMP process handles tasks related to SNMP. Check this log when troubleshooting or monitoring SNMP. Note that wherever possible, the SNMP `ifIndex` values are persistent across reboots or in the event of hardware additions and deletions that result from PIC or FPC insertion and removal. This persistence is the default behavior and is achieved by storing SNMP indexes in the `/var/db/dcd.snmp_ix` file.
- `vrrpd`: The virtual router redundancy protocol (VRRP) process handles the activities related to this protocol. Check this log when troubleshooting or monitoring VRRP.

## Entries Also Written to the Main Syslog File

The entries written to individual process log files also write into the main syslog file (`messages`). Generally speaking, you begin by analyzing the `messages` file for signs of trouble. Once you identify trouble relating to a particular process, you can parse or monitor the files of that process to reduce the amount of information through which you must wade.

## Interpreting Syslog Messages

- Standard log entries consist of the following fields:
  - Timestamp, platform name, software process name or PID, a message code, and the message text:
 

```
Apr 29 09:43:08 host chassisd[2320]: CHASSISD_FRU_EVENT: scb_recv_slot_detach: FPC 1 detach
```
  - Using **explicit-priority** alters the message format to include a numeric priority value:
 

```
Apr 29 09:41:27 *DAEMON-5-CHASSISD_FRU_EVENT: host chassisd[2320]: scb_recv_slot_detach: FPC 1 detach
```
  - Consult the *System Log Messages Reference* documentation for details on log entries
    - Use **help syslog ?** for help in decoding message codes:

```
user@host> help syslog CHASSISD_IFDEV_DETACH_FPC
Name:          CHASSISD_IFDEV_DETACH_FPC
Message:      ifdev_detach(<fpc-slot>)
Help:         chassisd detached all PIC interface devices on FPC
Description:  The chassis process (chassisd) detached the interface devices
              for all Physical Interface Cards (PICs) installed in the
              indicated Flexible PIC Concentrator (FPC).
Type:         Event: This message reports an event, not an error
Severity:     notice
```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-36

### Interpreting System Log Entries

When using the standard syslog format, each log entry written to the messages file consists of the following fields:

- *timestamp*: Time of logging the message.
- *name*: The configured system name.
- *Process name* or *PID*: The name of the process (or the Process ID [PID] when a name is not available) that generated the log entry.
- *message-code*: A code that identifies the general nature and purpose of the message. In the example shown, the message code is `CHASSISD_IFDEV_DETACH_FPC: ifdev_detach(1)`.
- *message-text*: Additional information related to the message code.

When you add the **explicit-priority** statement, the syslog message format alters to include a numeric priority value. In this case the value 0 is for the most significant and urgent messages (emergency), while 7 denotes debug level messages.

Consult the *System Log Messages Reference* documentation for a full description of the various message codes and their meanings—better yet, use the CLI's **help** function to obtain this information. The example shows the operator obtaining help on the meaning of the `CHASSISD_IFDEV_DETACH_FPC: ifdev_detach(1)` message code. Based on the output, it becomes relatively clear that the message code relates to the `chassisd` processing disconnecting the interfaces associated with a given FPC, and that this process is considered an event rather than an error.



## Tracing Overview

- Tracing is the JUNOS Software equivalent of *debug*
  - You can enable tracing on a production network
  - Requires configuration
  - Can trace multiple options (flags) to a single file
- Generic tracing configuration syntax:

```
[edit protocols protocol-name]
user@host# show
  traceoptions {
    file filename [size size] [files number]
                                [world-readable | no-world-readable];
    flag flag [flag-modifier] [disable];
  }
```

The protocol or function being traced

Where to write the trace results

Flags identify what aspects of the protocol the software traces and at what level of detail

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 3-37

### Hear Tracing and Think Debug

Tracing is the JUNOS Software term for what other vendors sometimes call *debug*. In most cases when you enable tracing (through configuration), you create a trace file that stores decoded protocol information. You analyze these files using standard CLI log file syntax like **show log logfile-name**. Because of the design of Juniper Networks routing platforms, you can enable detailed tracing in a production network without significantly impacting performance. Even so, you should always remember to turn tracing off once you complete your testing to avoid unnecessary resource consumption.

### Generic Tracing Configuration

The slide shows a generic tracing stanza that, if applied to the [edit protocols] portion of the configuration hierarchy, would result in tracing of the specified routing protocol's events. Specified routing protocol tracing operations track the flagged routing operations and record them in the specified log file.

*Continued on next page.*

## Generic Tracing Configuration (contd.)

The following are configuration options for tracing:

- **file *filename***: Specifies the name of the file in which to store information.
- **size *size***: Specifies the maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named `trace-file` reaches this maximum size, it receives compression and is renamed `trace-file.0.gz`. When the trace file again reaches its maximum size, `trace-file.0.gz` is renamed `trace-file.1.gz`, and `trace-file` is compressed and renamed `trace-file.0.gz`. This renaming scheme continues until it reaches the maximum number of allowable trace files. The software then overwrites the oldest trace file. If you do not specify a maximum number of trace files with the `files` option, the default number of files to keep is ten. If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option. You can use `yk`, `ym`, or `yg` to specify kilobytes, megabytes, or gigabytes, respectively. The default range is 128 KB.
- **flag *flag***: Specifies a tracing operation to perform. You can specify multiple flags.
- **files *number***: Specifies the maximum number of trace files. When a trace file named `trace-file` reaches its maximum size, JUNOS Software renames it `trace-file.0`, then `trace-file.1`, and so forth, until it reaches the maximum number of trace files. The software then overwrites the oldest trace file. The default is ten files.

Including the `traceoptions` statement at the `[edit interfaces interface-name]` hierarchy level allows you to trace the operations of individual router interfaces. You can also trace the operations of the interface process, which is the device control process.

When tracing a specific interface, the software does not support the specification of a trace file. The JUNOS Software kernel does the logging in this case, so the software places the tracing information in the system's `messages` file. In contrast, global interface tracing supports an archive file; by default it uses `/var/log/dcd` for global interface tracing.

## Protocol Tracing

- Include the `traceoptions` statement at the `[edit protocols protocol-name]` hierarchy
  - Useful when troubleshooting configuration and interoperability problems
- A typical OSPF tracing configuration along with sample output:

```
[edit protocols ospf]
user@host# show
traceoptions {
  file ospf-trace;
  flag hello detail;
  flag lsa-request detail;
  flag lsa-update detail;
}
user@host> show log ospf-trace
. . .
Oct 9 22:41:45.233671 OSPF built router LSA, area 0.0.0.1
Oct 9 22:41:45.233715 ospf_set_lsdb_state: Router LSA 192.168.24.1 adv-rtr
192.168.24.1 state GEN_PENDING->QUIET
Oct 9 22:41:45.233732 OSPF built router LSA, area 0.0.0.11
Oct 9 22:41:45.233865 OSPF sent Hello 10.222.100.1 -> 224.0.0.5 (ge-0/0/3.100,
IFL 70)
Oct 9 22:41:45.233885 Version 2, length 44, ID 192.168.24.1, area 0.0.0.1
. . .
```

© 2009 Juniper Networks, Inc. All rights reserved.



www.juniper.net | 3-39

### Protocol Tracing

You trace the operations of a specific protocol by including the `traceoptions` statement at the `[edit protocols protocol-name]` hierarchy. In most cases you should be selective in what you trace because selecting the `all` keyword can overwhelm you with endless minutia. The sample OSPF stanza on the slide reflects a typical tracing configuration that provides details about important events like hello message or OSPF link-state advertisement (LSA) details. In most cases you should use the `detail` switch with a given protocol flag for the added information often needed in troubleshooting scenarios.

### Sample Output

The slide shows a sampling of the results obtained with the tracing configuration. As with any log file, you simply enter a `show file trace-file-name` command to view the decoded protocol entries. The sample trace output reflects the receipt of an OSPF hello message from 10.222.100.1 and goes on to show some of the hello protocol parameters.

## Analyzing Log and Trace Files

- The software stores log and trace files 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!

```
user@host> show log ospf-trace | match "packet ignored"
Oct 9 22:41:46.318641 OSPF packet ignored: area stubness mismatch from 10.222.100.2
Oct 9 22:41:48.209141 OSPF packet ignored: area stubness mismatch from 10.222.200.2
Oct 9 22:41:55.389251 OSPF packet ignored: area stubness mismatch from 10.222.100.2
Oct 9 22:41:57.768850 OSPF packet ignored: area stubness mismatch from 10.222.200.2
```

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

```
user@host> show log ospf-trace | match "packet ignored" | match virtual
Oct 18 07:43:10.651694 OSPF packet ignored: can't find virtual link for 10.222.200.2
Oct 18 07:43:19.235156 OSPF packet ignored: can't find virtual link for 10.222.200.2
```

- Use quotes and the pipe character to invoke 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

### 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 it has 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
```

*Continued on next page.*

### Viewing Logs and Traces (contd.)

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 cascades so that only lines containing the words `packet ignored` and `virtual` are displayed; this cascading 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 it is that you are seeking. The slide provides two examples of a logical OR search. The basis for the former is 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 requires enabling syslog priority with the **explicit-priority** keyword.

## Miscellaneous Log File Commands

- Monitor a log or trace in real time with the CLI's `monitor` command:

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

- Shows updates to monitored files until canceled, with matching piped output!
- Use Esc-q to enable or disable real-time output to screen
- Issue a `monitor stop` to cease all monitoring

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

```
[edit protocols ospf traceoptions]
user@host# delete flag open
```

- Log and trace file manipulation:

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

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

- Use the `file delete` command to delete log and trace files:

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

### 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. You can identify the messages from each log by *filename*, where *filename* is the name of the file that displays entries. This line displays 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 displaying only entries that match your search criteria. To make use of the functionality, use a command in the following form:

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

### Stopping Tracing Through Configuration

If you do not delete or disable all trace flags, tracing continues in the background, and the output continues to write to the specified file. The file remains on the Routing Engine hard disk until it is either 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.

*Continued on next page.*

## Log and Trace File Manipulation

To truncate files used for logging, use the `clear log filename` command. To delete a file, use the `file delete` command. You can also use wildcards with `delete`, `compare`, `copy`, `list`, and `rename` operations.

Not for Reproduction

## Agenda: Troubleshooting Tool Kit for JUNOS Platforms

- Caveats and Warnings
- Troubleshooting Methodology
- Troubleshooting Tools
  - The JUNOS Software CLI
  - The Craft Interface Panel
  - System Logs and Protocol Tracing
  - Interactive UNIX Shell
    - Core Files for Diagnostic Analysis
    - The JTAC Knowledge Base
- Best-Practices Case Study

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-44

### Troubleshooting Tools: Interactive UNIX Shell

The slide highlights the topic we discuss next.



## The Interactive Shell

- **Interactive UNIX shell support:**
  - CLI users can escape to an interactive shell when permitted by their login class
  - Juniper Networks does not support the shell, and it is potentially dangerous
    - Use only under JTAC guidance
- **Some things to do when in a shell:**
  - Access standard UNIX utilities such as `ls`, `tar`, `gzip`, `vi`, `tcpdump`, and so forth
  - Display and modify kernel variables using `sysctl`
  - Establish connections (`vtty/ctty`) to PFE components to display NVRAM and other diagnostic data

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-45

### Interactive Shell Support

Based on a FreeBSD operating system, the JUNOS Software CLI supports an escape to a UNIX shell. While the possibilities can seem endless, we stress that designers highly customized JUNOS Software, and did not design it to act as a Web server or as some type of UNIX device. You can do serious damage to the Juniper Networks platform if you do not observe great care and caution when operating in a shell. Access to an interactive shell is controllable through login class permissions. Once at a shell, you can `su` to root, if you know the root password, or if you have not set it.

Juniper Networks does not officially support use of the shell because the CLI offers all that you should need in normal circumstances. For advanced troubleshooting activities, or for advanced functionality like automated shell scripts (for which Juniper Networks support is not expected or sought), the shell can be a real boon.

Users who wish to add production scripting functionality to their networks should consider operational scripts, commit scripts, and event scripts. The coverage of these scripts is outside of the scope of this course.

*Continued on next page.*

## Some Shell Support Details

From a troubleshooting and operational analysis perspective, a few good reasons for escaping a shell exist. These reasons include the following:

- Access to standard utilities and programs like `tar`, `gzip`, `top`, `ps`, `kill`, `vi`, and so forth, offer experienced UNIX users the tools they need to perform advanced troubleshooting tasks like compressing a core file or manually editing a configuration file when the CLI is not available;
- Use `sysctl` to access and modify (under the guidance of JTAC) various kernel parameters like TCP window sizes, the number of available protocol sockets, and so forth; and
- Establish a connection to the embedded hosts (controllers) within the PFE complex to access diagnostic and log data held in NVRAM.

## Shell Case Study 1: tcpdump

- Use the CLI to monitor traffic that writes to a file
  - Escape to a shell and use tcpdump to read the file with protocol filtering

- You can use the CLI with the hidden `read-file` switch:

```

user@host> monitor traffic interface se-1/0/0 write-file dump-file
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on se-1/0/0, capture size 96 bytes
^C
23 packets received by filter
0 packets dropped by kernel
user@host> start shell
% tcpdump -r dump-file ip proto 89
02:36:14.083877 In IP 172.18.36.2 > 224.0.0.5: OSPFv2, Hello, length 48
02:36:18.734087 Out IP 172.18.36.1 > 224.0.0.5: OSPFv2, Hello, length 48
. . .
% exit
exit
user@host> monitor traffic read-file dump-file matching "ip proto 89"
02:36:14.083877 In IP 172.18.36.2 > 224.0.0.5: OSPFv2, Hello, length 48
02:36:18.734087 Out IP 172.18.36.1 > 224.0.0.5: OSPFv2, Hello, length 48
. . .
    
```

A hidden CLI switch writes output to a named file; use with care so your hard disk does not fill!

tcpdump invoked with the -r switch to read the named file

### Interactive Shell Case Study: tcpdump

The slide shows how writing monitored traffic to a file allows the use of protocol filtering and standard protocol analysis tools like tcpdump. Note that in real-time mode, protocol filtering does not function because the Layer 2 headers are stripped in hardware before the monitoring of traffic occurs. To work around this problem, JUNOS Software writes pseudo-Layer 2 headers when writing monitored traffic to a file. The presence of these headers accommodates protocol filtering actions.

On the slide, we start by issuing a `monitor traffic interface se-1/0/0` command using the hidden `write-file` switch and a target file name of `dump-file`. Note that the `write-file` switch is hidden because failing to stop the traffic monitoring could result in the `/var` file system becoming full. While this condition should not crash the router, it impacts the router's ability to conduct on-going logging and tracing activities.

After the traffic monitoring ceases, we escape to a UNIX shell and invoke `tcpdump` with the `-r` switch to tell it to read the contents of the named file.

*Continued on next page.*

### Interactive Shell Case Study: tcpdump (contd.)

The following example shows how similar results are possible within the CLI using the hidden **read-file** switch and a protocol filter expression:

```
user@host> monitor traffic read-file dump-file matching ?
Possible completions:
  <matching>          Expression for headers of receive packets to match
user@host> monitor traffic read-file dump-file matching "ip proto 89"
19:00:18.203725 Out IP 10.0.13.2 > 224.0.0.5: OSPFv2, Hello, length: 48
19:00:22.938474 Out IP 10.0.13.1 > 224.0.0.5: OSPFv2, Hello, length: 48
. . . .
```

Not for Reproduction

## Shell Case Study 2: Who Is Listening?

- A `show system connections` command (or port scan) indicates that a process is listening on port 6154:

```
user@host> show system connections | match 6154
tcp4      0      0 128.0.1.16.6154      *.*      LISTEN
tcp4      0      0 *.6154              *.*      LISTEN
```

- Q: How can you determine what process is listening at that port?

- A: Escape to a shell and use the `netstat` and `fstat` utilities!
  - The `-A` switch tells `netstat` to display protocol control block information that then feeds into the `fstat` program
  - The `-a` switch displays all sockets; the default displays only nonlistening (connected) sockets

```
% netstat -Aa | grep 6154
c3ad31e4 tcp4      0      0 128.0.1.16.6154      *.*      LISTEN
c3ad3790 tcp4      0      0 *.6154              *.*      LISTEN
% fstat | grep c3ad31e4
root      fwdd      4309    37* internet stream tcp c3ad31e4
```

The PCB of the process in question

The listening process is revealed!

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 3-49

### What Is Listening on TCP Port 6154?

In some cases, you must determine exactly what process has opened a TCP or UDP port to listen for connections. Such an event might stem from a curious operator who questions the output of a `show system connections` command, as shown on the slide, or as a result of a security audit involving port scanning.

### How to Determine What Process Opened a Port

By escaping to a shell and using the standard BSD `netstat` and `fstat` commands, you can determine what process is listening on a given port using the steps outlined on the slide. You begin by issuing a `netstat -Aa` command that displays all listening and connected sockets (the `-a` switch), along with the related protocol control block (PCB) information (the `-A` switch). In this example, the `grep` utility saves some parsing work by matching only lines containing the value 6154. The result of this command is the PCB information needed for the subsequent `fstat` command. Once again, `grep` matches only the lines of interest.

The output of the `fstat` command makes it clear that the culprit is the `fwdd` process, which is the packet forwarding engine forwarding process.

## Connecting to PFE Components

- You can establish connections from the RE to PFE components
  - Use to display NVRAM crash data, perform module diagnostics, upgrade FPGA, and so forth
  - Uses internal connection and proprietary TNP protocol
    - Connections are Ethernet (`vtty`) or Console (`ctty`)
    - Use `start shell pfe [network | direct] ?`
- Older versions of software require a root shell and knowledge of `tnp` address assignment
  - Use the CLI command `show tnp addresses` to display the `tnp` addresses for your system

© 2009 Juniper Networks | www.juniper.net | 3-50

### Connecting to PFE Components

You can use the internal connectivity between the RE and PFE, along with the Trivial Network Protocol (TNP), to establish connections to embedded hosts (controllers) within the PFE complex. The term *embedded host* refers to a PFE component with its own microprocessor and microkernel. Examples include system Control Boards and FPCs.

In most cases, the only reason to connect to a PFE component is to access diagnostic information in the form of log entries or core files retained in the affected component's NVRAM. On most platforms you use virtual terminal (`vtty`) connectivity over an Ethernet communications channel. The use of `vtty` requires that you specify the correct `tnp` address. Some platforms also support console (asynchronous) access using a serial type of connection known as a `ctty`.

By parsing entries in the syslog, you can determine what PFE component has reported a crash, and therefore to which embedded host you must connect to obtain crash and log data for submission to JTAC.

*Continued on next page.*

## Connecting to PFE Components (contd.)

To connect to a PFE component, issue use the **start shell pfe** command with the **network** switch for Ethernet access or the **direct** switch for console access. Use the context sensitive help feature to display tnp address or name assignment:

```
user@host> start shell pfe network ?
Possible completions:
  fpc0          Connect to Flexible PIC Concentrator 0
  fpc1          Connect to Flexible PIC Concentrator 1
  fpc2          Connect to Flexible PIC Concentrator 2
  fpc3          Connect to Flexible PIC Concentrator 3
  fpc4          Connect to Flexible PIC Concentrator 4
  fpc5          Connect to Flexible PIC Concentrator 5
  fpc6          Connect to Flexible PIC Concentrator 6
  fpc7          Connect to Flexible PIC Concentrator 7
user@host> start shell pfe network fpc1
```

## Older Versions of Software

If you are running older versions of JUNOS Software, whether connecting by `vtty` or `ctty`, you might need to be at a root shell prompt to forge a connection from the RE to a PFE component. When using a `vtty` connection, you should first issue the **show tnp addresses** CLI command so that you know which address to specify. You can also use the **tnpdump** command, which is an alias to the **show tnp addresses** command at the shell prompt.



## Shell Case Study 3: Display NVRAM

- Goal: Escape to a root shell and establish a `vtty` connection to the M10i router's CFEB to display NVRAM contents

```

user@host> show tnp addresses
Name      TNPaddr      MAC address  IF  MTU  E  H  R
master    0x1 02:00:00:00:00:04 fxp1 1500 3 0 3
cfeb      0x2 02:00:00:00:00:02 fxp1 1500 5 0 3
re0       0x4 02:00:00:00:00:04 fxp1 1500 3 0 3
. . .
user@host> start shell
% su
Password:
root@shost% vty 2
CSBR platform (266Mhz PPC 603e processor, 256MB memory, 512KB flash)
CSBR1(my-junos-device vty)# show nvram
System NVRAM :
32751 available bytes, 56 used, 32695 free
Contents:
CSBR: Reset reason (0xc): RE Initiated Reset, Power On
. . .
CSBR1(my-junos-device vty)# quit
root@host%

```

CLI command displays tnp address of each PFE host module

Shell prompt and su to root

vtty connection requested

Truncation of remainder of output

### Shell Case Study: Display NVRAM

The example on the slide begins by issuing the `show tnp addresses` command to obtain the list of `tnp` endpoints for the platform in question (an M10i router). In this example the goal is to connect to the M10i router's CFEB, which is currently using `tnp` address 2.

Armed with the knowledge of the CFEB's `tnp` address, we escape to a shell and issue an `su` to the root so as to execute a `vtty 2` command. The slide shows that the connection is successful by virtue of receiving the login banner from the CFEB. Once connected, we issue a `show nvram` command to obtain diagnostic information for use by JTAC. When done, we break out of the `vtty` connection and return to the root shell with a `quit` (or `exit`) command.



## Agenda: Troubleshooting Tool Kit for JUNOS Platforms

- Caveats and Warnings
- Troubleshooting Methodology
- Troubleshooting Tools
  - The JUNOS Software CLI
  - The Craft Interface Panel
  - System Logs and Protocol Tracing
  - Interactive UNIX Shell
  - Core files for Diagnostic Analysis
    - The JTAC Knowledge Base
- Best-Practices Case Study

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-53

### Troubleshooting Tools: Core Files for Diagnostic Analysis

The slide highlights the topic we discuss next.

## Complexity

- Modern computing environments are complex and therefore, have complex bugs
  - Transient software failures are extremely hard to reproduce and, therefore, difficult to fix
    - Hardware errors can also trigger software failures
  - Well-written code dumps a core file for diagnostic analysis when a fatal fault (panic) occurs
    - The stack trace identifies the name of the offending process, memory pointers, and register data at the time of the fault
  - In JUNOS Software numerous entities can dump a core at panic or upon command
    - The kernel, software processes, and embedded hosts in the PFE

### Complexity of Modern Computers and Operating Systems

The complexity of modern computers and operating systems leads to equally complex bugs! It is very difficult to diagnose transient software failures (for example, a random crash or reboot), because so many potential causes for these types of faults exist. In most cases, a crash is the result of a programming error or the failure to anticipate a particular set of events and the software interaction that ensues. However, a crash can also stem from hardware-related causes. In the latter case, a memory error might corrupt a memory pointer or result in an illegal instruction.

Because transient software failures are so difficult to diagnose, well-written code incorporates the ability to *dump* the program's environment in the form of memory pointers, instructions, and register data to a file in the event of a panic or other serious malfunction. A software engineer using a debugger and a version of the executable containing debugging symbols can analyze the resulting core file. The result of this analysis is generally a very good idea of the sequence of events that led to the crash, and armed with this information, you can take corrective actions. For example, you can perform a software patch or hardware Return Materials Authorization (RMA).

While it might sound bad, it is actually quite beneficial that JUNOS Software has the ability to dump various types of core files for diagnostic use. In most cases, core files generate automatically as a result of a failure, but you can also generate cores on demand. JUNOS Software can generate core files relating to the JUNOS Software kernel itself, to the processes that run above that kernel, or to the embedded host modules within the PFE.

## Core Files and Memory Dumps

- Core files are critical to diagnosing software faults
  - In some cases, faulty hardware can also lead to a core dump
- Technical support engineers deal with three types of core files:
  - JUNOS Software kernel (also known as RE cores)
  - JUNOS Software processes
  - PFE-embedded host cores (FPCs and system boards)
- Three distinct locations where JUNOS Software writes core files:
  - JUNOS Software kernel (RE) cores write to `/var/crash`
  - JUNOS Software process cores write to `/var/tmp`
  - PFE stack trace writes to NVRAM on the affected component or system board
    - Core also copies to `/var/crash` when you enable `chassis dump-on-panic`

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-55

### Core Files Are Critical

Today's internetworking software is exceedingly complex. As a result, equally complex bugs that result from unforeseen circumstances can result in a fatal error within a software process. Most of these software faults relate to illegal memory operations caused by the process attempting to read or write data from a memory area outside the boundaries allocated for that process. In some cases, faulty hardware, such as failing memory, can cause stack or register corruption, which leads to a fatal error in a software process. You can use core and log file analysis to determine if hardware errors have led to software problems.

In a monolithic operating system, such a fault results in a crash of the entire operating system. In contrast, the protected memory environment of JUNOS Software ensures that faulty processes do not affect other aspects of the operating system.

Even so, it can be very difficult to diagnose the exact set of events that lead up to a process crash without a core file for forensic analysis. A core file represents the set of memory locations and stack data that was in place at the time of the fault. A software engineer then runs a copy of the binary image that left the core file (with debug symbols included) against the actual core file using a debugger to enable problem diagnosis.

*Continued on next page.*

## Three Types of Core Files

Juniper Networks support engineers typically deal with three types of core files. These files are the following:

- *JUNOS Software kernel (RE) cores:* A kernel core file is left by the JUNOS Software kernel when it encounters a panic condition. The software also saves a copy of the virtual memory state (which can be quite large).
- *JUNOS Software process cores:* Each process, such as the chassis management or automatic protection switching processes (*chassisd* or *apsd*), is capable of leaving a core when a panic occurs.
- *PFE cores:* Various components in the PFE contain their own microprocessors that run a microkernel. Examples include the CFEB on M7i and M10i platforms, FPCs, the Forwarding Engine Boards (FEB) on the M120, and others. Each of the PFE's embedded hosts is capable of dumping a core file when a crash (panic) occurs.

## Core File Locations

Depending upon the JUNOS Software version, you might need to explicitly configure core file storage. When enabled, the process that generates the core determines the actual location of a core file.

Core files created by a kernel panic are stored in the `/var/crash` location when you enable the **system dump-on-panic** option (hidden) at the `[edit system]` hierarchy. The software enables this option by default.

Core files generated by a process are stored in the `/var/tmp` directory. This behavior is the default in all JUNOS Software releases.

When a PFE component dumps a core, the resulting stack trace writes into that component's NVRAM. If you enable **chassis dump-on-panic** (hidden) at the `[edit chassis]` hierarchy, a copy of the core is also stored in the `/var/crash` directory on the RE. We recommend this option, and it is the default.

You can use the CLI command **show system core-dumps** to quickly determine if any core files are stored on the RE.

## Forcing Cores

- Forcing a running process to write a core can help diagnose certain problems
- Two methods exist:
  - Use the hidden `request system core-dump` command to force a core dump
    - Use with caution! By default, the software creates a copy of the running process; this copy can result in excessive memory paging if the memory footprint of the process is large

```
user@host> request system core-dump routing
Generating core dump for routing process using running method
user@host> show system core-dumps
/var/crash/*core*: No such file or directory
-rw-rw---- 1 root wheel 275275 Jan 13 05:11 /var/tmp/rpd.core.0.gz
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
total 1
```

Hidden command

- Use `gcore -s path/binary-name pid` at a root shell
  - This approach suspends the process during core writing
  - Uses less memory, but process suspension can lead to other problems
- You can use different procedures to force kernel and PFE cores

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 3-57

## Forcing Process Cores

In certain rare situations, a software engineer might want to obtain a core file from a process that appears to be running normally. Note that forcing software processes to write cores might impact system, performance, and operation. Only perform these steps under the guidance of JTAC.

## Two Methods

In most cases, you obtain a running core file by using the hidden `request system core-dump process-name` CLI command. By default, this process forks off a copy of the running process (a running core), which has the upside of leaving the original process free to do its process duties. The downside is that if the process in question is large (for example, `rpd`) it might tax system memory, because the system must support two instances of that process. A system that is low on memory begins paging to the swap file, and this procedure can slow things down to the extent that keepalives are lost or `rpd` scheduler slips begin to occur. For the routing process (`rpd`), you can specify whether a fatal (the process is stopped and then restarted) or running core should be generated. For most processes, a running core is the only option. Note that either type of core can be disruptive, and that a running core does not generate a `.tar` archive with context.

*Continued on next page.*

## Two Methods (contd.)

You can also instruct a process to generate a core file from a root shell using the `gcore` utility. The main advantage to this approach is that you can instruct `gcore` to suspend the process in question during the core dump. Because the software does not create a copy of the process, less taxation occurs on the system's memory. However, because the process suspends during what can be a somewhat lengthy period (10 seconds or so for a busy system with a large process), other problems, like `rpd` scheduler slips, might occur.

The slide shows an example of the recommended `gcore` syntax. The `-s` argument tells `gcore` to suspend the process during the dump. You must also specify the full path and binary name of the processes, as well as the PID of the currently running processes.

You can use the `which` command to obtain the path of a process, and the output of a `ps ax` command to obtain the PID associated with that process. You should change into the `/var/tmp` directory before running `gcore` because it writes the core file to the current working directory by default. Note that using `gcore` from a root shell never produces a `.tar` archive with context information.

The following output shows an operator using `gcore` to obtain an `rpd` core:

```
root@host% cd /var/tmp
root@host% ls *core*
ls: No match.
root@host% which rpd
/usr/sbin/rpd
root@host% ps ax | grep rpd
 2275 ?? S    0:09.08 /usr/sbin/rpd -N
 2280 ?? I    0:00.40 /usr/sbin/vrrpd -N
root@host% gcore -s /usr/sbin/rpd 2275
root@host% ls *core*
core.2275
```

The procedures outlined on these pages are for the generation of core files from processes only. The forcing of a JUNOS Software kernel core is beyond the scope of this class because it requires that you enter complex `sysctl` syntax at a root shell.

You can issue a `write coredump` command when connected to an embedded host to force a PFE component to write a core file, as shown in the case of an M10i router's CFEB 0 (with `chassis dump-on-panic` enabled):

*Continued on next page.*

**Two Methods (contd.)**

```
root@host% vty cfeb0
```

```
CSBR platform (266Mhz PPC 603e processor, 256MB memory, 512KB flash)
```

```
CSBR0(host vty)# write core
```

```
[Jan 23 18:32:03.002 LOG: Info] Dumping core-CSBR0 to 1
```

```
[Jan 23 18:32:08.003 LOG: Err] Coredump write - saw ack 18038, expected 18039
```

```
CSBR0(host vty)# [Jan 23 18:32:58.005 LOG: Info] Coredump finished!
```

```
CSBR0(host vty)# exit
```

```
root@host% ls -l /var/crash
```

```
total 507780
```

```
-rw-r--r-- 1 root wheel 259885052 Jan 23 18:32 core-CSBR0.core.0
```

```
-rw-rw-r-- 1 root wheel          5 Sep  9 2004 minfree
```

```
root@host%
```

## Submit the Core to JTAC

- Use FTP for core files larger than 10 MB
  - Attach to case with case manager when less than 10 MB
  - Recommended procedures for FTP:
    1. Log in to the case manager at <https://www.juniper.net/cm/> to open a case
    2. Escape to a root shell and change to the directory containing the core file
    3. Rename the file according to case-number.core.sequence-number
    4. Use `chmod 444 file-name` to ensure all users have read permissions
    5. Compress the core file if needed (especially important for kernel memory images that can be quite large)
    6. Log in to Juniper Networks anonymous FTP site at <ftp.juniper.net>, and change into the `\pub\incoming` directory
    7. Create a directory using your case number as a name; change into this directory
    8. Enable binary mode transfer and enable hash mark printing for progress indication
    9. Transfer the core file using a `put` or `mput` command

## Transferring Core Files to Juniper Networks

You should always submit core files to JTAC for fault analysis. The following are the recommended procedures for transferring core files to JTAC:

1. Log in to the case manager at <https://www.juniper.net/cm/> to open a support case and obtain a case number.
2. Escape to a root shell and change to the directory containing the core file.
3. Rename (or copy) the file using a name in the form of case number-core-sequence number.
4. Although not strictly necessary, we recommend that you `chmod` the core file with 444 to ensure that all users (root, owner, and other) have read permissions for the file.
5. In some cases, the core file is already in a compressed state, as indicated by a `.tgz` or `.gz` file extension. If not compressed, you should compress the file to reduce transfer and storage requirements. This compression is especially important when dealing with the `vmcore.0` file associated with a kernel crash, because this memory image file can be quite large.

*Continued on next page.*



**Transferring Core Files to Juniper Networks (contd.)**

6. Log in to the Juniper Networks anonymous FTP site at `ftp://ftp.juniper.net` and change into the `/pub/incoming` directory.
7. Create a directory named with your assigned case number and change into this directory.
8. Ensure that you set your FTP client for a binary transfer. In many cases the client defaults to the correct transfer type. Issue a **type** command to confirm the current transfer setting and use the **image** or **binary** command to enable binary transfer mode as needed. Enabling hash-mark printing provides transfer progress indication.
9. Upload the file using a **put** or **mput** command.

## Agenda: Troubleshooting Tool Kit for JUNOS Platforms

- Caveats and Warnings
- Troubleshooting Methodology
- Troubleshooting Tools
  - The JUNOS Software CLI
  - The Craft Interface Panel
  - System Logs and Protocol Tracing
  - Interactive UNIX Shell
  - Core Files for Diagnostic Analysis
  - The JTAC Knowledge Base
- Best-Practices Case Study

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

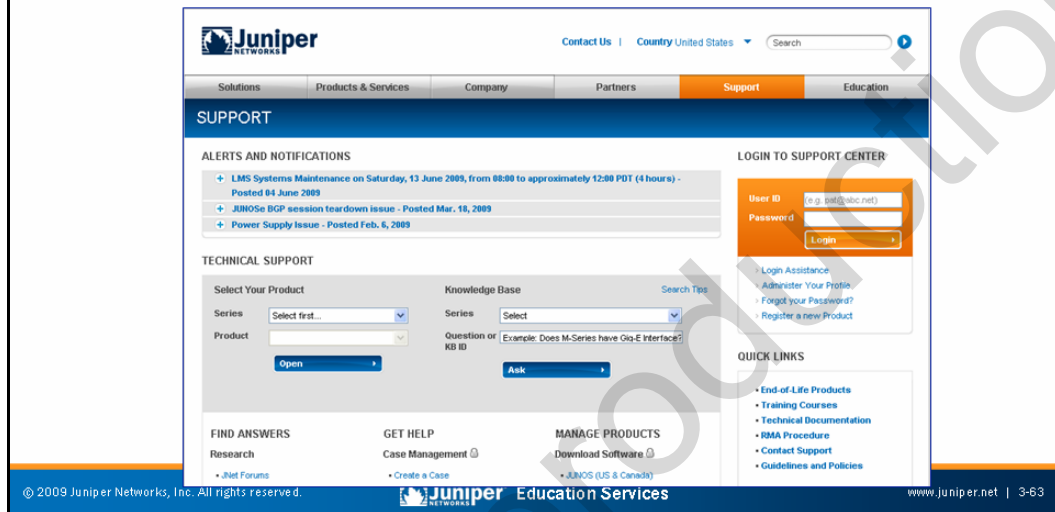
www.juniper.net | 3-62

### Troubleshooting Tools: The JTAC Knowledge Base

The slide highlights the topic we discuss next.

## JTAC Knowledge Base and Problem Report Search Tool

- Customers can access the JTAC Knowledge Base and the Problem Report search tool
  - <http://www.juniper.net/customers/support/>



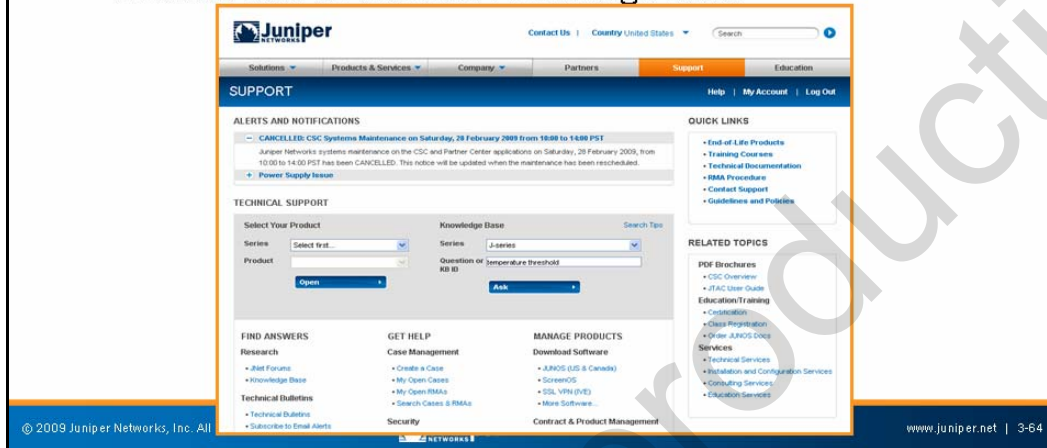
### The JTAC Knowledge Base and Problem Report Search Tools

Customers with support contracts can access the JTAC Knowledge Base and Problem Report search tool to assist themselves in problem diagnosis. The graphic on the slide shows the current Customer Support Center (CSC) welcome page that greets the user.

The Knowledge Base contains various entries on technology, troubleshooting, and recommended procedures. The Problem Report database, on the other hand, contains a listing of known bugs along with their status and any known workarounds.

## Knowledge Base Case Study (1 of 3)

- You worry about the potential for heat damage and thermal shutdown of your J Series router
  - Question: At what point do temperature alarms activate, and when does a thermal shutdown initiate?
  - Answer: Search the JTAC Knowledge Base!



### JTAC Knowledge Base Case Study: Part 1

The next set of pages illustrates how the JTAC Knowledge Base and the Problem Report search tool can help you find your own answers. The stage is set with the rather common question of “just how hot is too hot for a typical J Series platform?”

The slide shows a user just about to search the Knowledge Base for the keywords *temperature threshold*.

## Knowledge Base Case Study (2 of 3)

- When the search completes, scan the results for promising matches
  - The first entry looks promising

The screenshot shows the Juniper Knowledge Base interface. At the top, there's a navigation bar with 'Solutions', 'Products & Services', 'Company', 'Partners', and 'Support'. Below that, the 'KNOWLEDGE BASE' header is visible. A search bar contains the text 'temperature threshold'. The search results are displayed in a list format. The first result is highlighted in blue and has a red arrow pointing to it from the text 'The first entry looks promising' in the adjacent list. The highlighted result is:
   
**[1971] - Routing Engine temperature thresholds**
  
 What is the **temperature threshold** at which an alarm is generated?
   
 By Product: Router Products → J-series
   
 Source: Knowledge Base

### JTAC Knowledge Base Case Study: Part 2

The slide shows a portion of the results returned from the Knowledge Base search. The highlights suggest that ID number 1971 is quite promising in that it seems to deal with temperature threshold values on J Series platforms.

## Knowledge Base Case Study (3 of 3)

- Enjoy your newfound wisdom!

The screenshot shows the Juniper Knowledge Base interface. The article title is "Routing Engine temperature thresholds". The synopsis asks: "What is the temperature threshold at which an alarm is generated? What is the temperature threshold at which the router will shut down?". The solution section states: "This note gives details about the different temperature thresholds for each type of Routing Engine (RE):".

- Chassis will generate an alarm (logtrap) when the detected temperature reaches the Yellow Alarm level. The temperature of Yellow Alarm activation varies depending on the type of RE in use.
- If the temperature exceeds the applicable *Shutdown Temperature* for four minutes or longer, DCD will **non-gracefully** shut down the router by disabling the power supplies.
- The following figures are in **degree Celsius**; this reading is **measured by the RE CPU built in sensor**. The rest of the RE board (and chassis) will be at a lower level.

Metadata for KB1971:

- Version: 4.0
- Published: 07 Oct 2008
- Updated: 07 Oct 2008
- Categories: JUNOS, J-series, M-series, T-series, Alarms
- Former Article ID: (blank)
- Owner: Vincent Regis
- Reputation Lvl/Pts: (57)
- Available To: Customer, Partner, PTAC, Employee

At the bottom, there is a table of RE models and their corresponding temperature thresholds:

RE-1.0	RE-2.0	RE-3.0	RE-4.0	RE-5.0	RE-5.0+	RE-1000	RE-2000	J2300	J4300	J6300	J4350	J6350
--------	--------	--------	--------	--------	---------	---------	---------	-------	-------	-------	-------	-------

### JTAC Knowledge Base Case Study: Part 3

The contents of ID number 1971 are helpful, and with your newfound wisdom, you are well on your way to *J Series Guru* status.

Note that you can provide feedback to the maintainers of the CSC to indicate whether particular entries were helpful to you.

## Agenda: Troubleshooting Tool Kit for JUNOS Platforms

- Caveats and Warnings
- Troubleshooting Methodology
- Troubleshooting Tools
  - The JUNOS Software CLI
  - The Craft Interface Panel
  - System Logs and Protocol Tracing
  - Interactive UNIX Shell
  - Core Files for Diagnostic Analysis
  - The JTAC Knowledge Base
- **Best-Practices Case Study**

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services


www.juniper.net | 3-67

### Best-Practices Case Study

The slide highlights the topic we discuss next.

## An Out-of-Band Management Network

- An out-of-band management network is critical in times of network outage
  - Built in out-of-band support with the fxp0 interface
    - Juniper Networks does not support transit routing over fxp0
    - Define a `backup-gateway` to support out-of-band routing when `rpcd` is not running, and mark the default route used for out-of-band as `no-readvertise`
    - Enable remote access services (Telnet, SSH, or FTP) only as needed!
  - Console access recommended for maintenance activities



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 3-68

### Deploying an Out-of-Band Management Network

Relying on in-band methods to manage your network might seem like a good idea up until the point that a circuit or hardware outage prevents you from accessing your network, and as a result, prolongs corrective actions. We highly recommend deploying an out-of-band management network because it provides you with a *back door* into your network during times of outage or disruption.

All JUNOS platforms come with a built-in out-of-band interface in the form of fxp0. Note that fxp0 is an out-of-band interface because transit traffic cannot be routed over this interface. Put another way, if a packet arrives on fxp0 it can never egress on a PFE interface, and vice versa. Because of this behavior, we do not recommend running a routing protocol over the fxp0 interface in most cases. Instead, we recommend a static route flagged with `no-readvertise`. This flag ensures that the static route used for out-of-band connectivity does not advertise over any routing protocol.

We also recommend the use of a `backup-gateway`, especially when your hardware supports redundant REs. You use the `backup gateway` entry whenever `rpcd` is not running, such as in the case of a backup RE or a system that has had `rpcd` shutdown because of thrashing.

*Continued on next page.*



### Deploying an Out-of-Band Management Network (contd.)

Your out-of-band connectivity should provide both Ethernet (fpx0-based) and console access to your routers. You normally gain console access through some type of terminal server. We recommend console access whenever you perform serious maintenance activities, like upgrading or downgrading the system software, because if something goes wrong, or the system somehow returns to a factory default, you might no longer have Ethernet-based access to the system. Having console access is the only way that you can reload software from removable media or recover a lost root password.

Not for Reproduction

## Recommended Syslog Settings

- Where possible, configure your syslog to do the following:
  - Write entries to both a local file and to a remote host
    - Remote archiving proves invaluable when the local hard disk fails
    - Configure remote syslog service to retain log entries for at least one month
  - Use archive settings for your messages file to maintain at least 20 copies with a minimum 1 MB file size
    - Default is 10 copies of files; default size is platform specific
    - 128 Kb maximum size on J Series routers
    - 10 MB on TX Matrix
    - 1 MB on all other platforms
    - Especially important if remote syslog is not in effect
  - Log interactive CLI commands and configuration changes
    - Achieved with the `interactive-commands` and `change-log` facilities using the `info severity` level
    - Provides an audit trail of who did what, and when

© 2009 Juniper Netw

www.juniper.net | 3-70

### Recommended System Log Settings

Wherever possible, you should place the following system logging recommendation into effect:

- *Use a remote syslog host:* This recommendation helps in archiving syslog messages, and ensures that these valuable messages are available even in the event of a catastrophic failure of a router.
- *Archive logs:* You should configure syslog archive settings that ensure retaining entries for at least two weeks. This suggestion is especially important when remote system logging is not in place. We recommend configuring 20 copies of the messages file with each copy being at least 1 MB in size, except on J Series routers, which have limited storage space.
- *Log CLI commands and configuration changes:* We have all seen the joke about what to do if you break something while no one is watching—just walk away. While this advice is perhaps sound, it is futile when the system configuration logs interactive CLI commands. When combined with unique user logins, the logging of all commands issued on the machine provides an excellent audit trail of who did what, and when.

## Clock Synchronization

- We recommend synchronizing router clocks with NTP
  - Correlated timestamps in log files assist fault analysis
    - Also useful in forensic analysis of security incidents
- JUNOS Software cannot provide primary time reference
  - You need an external device for synchronization
    - A simple UNIX device using an undisciplined local clock suffices
  - Support for client, server, or symmetric modes, with or without authentication
  - Use the **show ntp associations** command to confirm synchronization status

Uses boot-server to set initial synchronization at boot

```
[edit system]
user@host# show
ntp {
  boot-server 10.0.1.201;
  server 10.0.1.201;
}
```

The configured list of possible synchronization sources (one server shown)

A simple NTP client-mode configuration

### Synchronize Router Clock

We recommend using the Network Time Protocol (NTP) to synchronize all routers to a common, and preferably accurate, time source. By synchronizing all routers, you ensure that time stamps on log messages are both accurate and meaningful, which is especially important when conducting security-related forensics where you must correlate events that might have occurred on numerous machines.

### JUNOS Software Needs a Reference

The basis for the NTP protocol is a series of timing hierarchies, with a Stratum 1 (atomic) timing source at the very top. While accuracy is desirable, you do not need to synchronize to Stratum 1 reference to benefit from having synchronized views as to the time of day. JUNOS Software cannot provide its own timing source because it does not support the definition of a local, undisciplined clock source (for example, the local crystal oscillator). If needed, you can always obtain a commodity UNIX device of some type with a configuration that provides a timing reference based on its local clock. Again, remember that any synchronization, even if based on an inaccurate local clock, is better than none.

JUNOS Software supports client, sever, and symmetric modes of NTP operation, and can also support broadcast and authentication. We recommend the use of authentication to ensure that an attacker cannot compromise your synchronization.

*Continued on next page.*

### JUNOS Software Needs a Reference (contd.)

The slide provides a typical NTP-related configuration stanza. While complete coverage of NTP is beyond the scope of this course, note that two machines can synchronize only when their current clocks are relatively close. A `boot-server` can set a router's clock at boot time to ensure that it is close enough to later synchronize to the configured time server. You can also issue a `set date ntp address` command as a substitute for a `boot-server`. Use the `show ntp associations` command to display synchronization status.

Not for Reproduction

## Enable Dump on Panic

- The software might not enable core dumps for the JUNOS kernel and embedded PFE hosts by default
  - JTAC recommends that you enable kernel and chassis dumps for added diagnostic capabilities
    - `system dump-on-panic` is the default
    - `chassis dump-on-panic` is the default
  - Generally requires a reboot to activate a `chassis dump-on-panic` change
    - Can connect with `vtty` or `ctty` to PFE and enable with a `set coredump enable` command to avoid a reboot
- Configuration requires hidden commands

```
[edit]
user@host# show system
host-name host;
dump-on-panic;
. . .
[edit]
user@host# show chassis
dump-on-panic;
```

 Juniper Education Services

www.juniper.net | 3-73

### Enable Core Dumps

Based on feedback from JTAC, `system dump-on-panic` is now enabled by default. Depending upon the JUNOS Software version, you might need to use hidden configuration commands to enable core dumps. Note that a change in chassis dump status normally requires a reboot of the PFE before placing the settings into effect. You can place the change into effect immediately by issuing a `set coredump enable` command on each PFE component that contains an embedded host. The `chassis dump-on-panic` statement enables core dumps on all PFE components (at reboot).

### Configuration Requires Hidden Commands

The slide shows the hidden configuration statements that you need to enable system and chassis core dumps.

## Confirming Dump Settings

- Use an interactive shell to confirm kernel dump setting
 

```

user@host> start shell
% sysctl -a | grep coredump
kern.sugid_coredump: 1
kern.coredump: 1
debug.elf_legacy_coredump: 0
%
      
```

Kernel dump confirmed
- Connect to the PFE's system board to confirm PFE dump settings
  - Example shown uses `cty` to connect to an M10i router's FEB
 

```

% su
Password:
root@host% cty cfeb0

CSBR0(host uart0)# show coredump
coredump is enabled
server address : 1
file name      : core-CSBR0
compression   : 0
          
```

Chassis dump confirmed

© 2009 Juniper Networks, Inc. All rights reserved. Education Services www.juniper.net | 3-74

### Confirming Dump Settings

Because core dumps are so critical when dealing with transient software failures, it is worthwhile to confirm the current settings for both system and chassis dumps. Waiting until after a crash to find out that you needed a reboot to enable a PFE core file is no way to begin the day.

The first code example shows the operator using the shell to obtain kernel parameters via the `sysctl` command. The output pipes to `grep` with the match criteria of `coredump`. In this example it is clear that we enabled kernel core dumps.

### Confirming PFE Dump Settings

You can establish a `vtty` or `cty` connection from the RE to an embedded host on the PFE to confirm the dump status of a given PFE component by issuing a `show coredump` command. In this example, we see a confirmation that we enabled core dumps for an M10i router's CFEB. Note that this setting tells the PFE component that it should place a copy of its core file onto the RE's `/var/tmp` directory, where you can easily access the core file.

## Best-Practices Case Study (1 of 2)

- Describe how each stanza assists in problem determination:

```
[edit]
user@host# show
system {
  host-name host;
  backup-router 10.250.0.254 destination 0.0.0.0/0;
  dump-on-panic;
  root-authentication {
    encrypted-password "$1$Nbh7fzlh$FFjUiCv4EjWwcfBBest6//."; # SECRET-DATA
  }
  services {
    ssh;
  }
  syslog {
    user * {
      any emergency;
    }
    host 10.0.1.1 {
      any notice;
      change-log info;
      interactive-commands info;
    }
    file messages {
      any notice;
      authorization info;
      archive size 1m files 20;
    }
  }
}
```

### Best-Practices Case Study: Part 1

This sequence of slides illustrates a basic configuration that reflects the current best-practices recommendations made by JTAC. Can you provide insight as to how the various configuration stanzas can help you when performing fault analysis?

## Best-Practices Case Study (2 of 2)

```

. . .
file cli-commands {
    interactive-commands any;
    archive size 1m files 10;
}
file config-changes {
    change-log info;
    archive size 1m files 10;
}
file errors {
    any error;
    explicit-priority;
}
}
archival {
    configuration {
        transfer-on-commit;
        archive-sites {
            "ftp://user;password@...
            archive-host/configs";
        }
    }
}
. . .

. . .
ntp {
    boot-server 10.1.10.2;
    server 10.1.10.2;
}
}
interfaces {
    fxp0 {
        description "OoB management
        interface";
        unit 0 {
            family inet {
                address
                10.250.0.136/16;
            }
        }
    }
}
routing-options {
    static {
        route 0.0.0.0/0 {
            next-hop 10.250.0.254;
            no-readvertise;
        }
    }
}
}
. . .

```

### Best-Practices Case Study: Part 2

The slide continues the case study started on the previous page.



## Summary

- In this chapter, we:
  - Explained why some of the information in this chapter can be disruptive to a production network
  - Described the layered troubleshooting methodology
  - Used various troubleshooting tools, including:
    - The JUNOS Software CLI
    - The Craft Interface
    - System logs and tracing
    - Interactive shell
    - Core files
    - JTAC Knowledge Base and Problem Report search
  - Explained JTAC recommendations for current best practices that promote troubleshooting

### This Chapter Discussed:

- Warnings and caveats regarding potentially disruptive commands and techniques;
- A layered troubleshooting methodology;
- Various troubleshooting tools supported by JUNOS Software; and
- JTAC recommended configuration settings for ease of troubleshooting.

## Review Questions

1. Describe the concept of a layered troubleshooting approach.
2. Describe three ways in which the CLI can help you perform fault analysis.
3. List two reasons to escape to an interactive shell.
4. How can you reset the OSPF process without also affecting other routing protocols?

## Review Questions

- 1.
- 2.
- 3.
- 4.

## Lab 1: JUNOS Troubleshooting Tools

- Gain experience using JUNOS Software troubleshooting tools.

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 3-79

### Lab 1: JUNOS Troubleshooting Tools

The slide shows the objective for this lab.

Not for Reproduction



# **Troubleshooting JUNOS Platforms**

## **Chapter 4: JUNOS Platforms Hardware Troubleshooting**

Not for Reproduction

## Chapter Objectives

- After successfully completing this chapter, you will be able to:
  - Provide an overview of hardware troubleshooting tools
  - Describe power on and power off procedures and boot media options
  - Troubleshoot JUNOS platforms using visual indicators
  - Troubleshoot JUNOS platforms using the CLI
  - Parse log files for symptoms of hardware problems

### This Chapter Discusses:

- An overview of hardware troubleshooting tools;
- Power on, power off, and boot media options;
- Troubleshooting based on visual indicators;
- Troubleshooting based on the JUNOS Software command-line interface (CLI); and
- Parsing log files for indications of hardware problems.

## Agenda: JUNOS Platforms Hardware Troubleshooting

- Hardware Troubleshooting Overview
- Power On, Power Off, and Boot Media
- Using the CLI to Troubleshoot
- Case Study

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-3

### Hardware Troubleshooting Overview

The slide lists the topics we cover in this chapter. We discuss the highlighted topic first.

## Hardware Troubleshooting Overview

- The Craft Interface and visual indicators
  - Red LEDs indicate failure
  - LCD panel displays all major and minor alarms
    - Issue a **show chassis craft-interface** command to view the display remotely on all platforms
  - Many individual components have their own status indicators
- The JUNOS Software CLI and J-Web tools
  - Interactive failure analysis using `show` commands
  - Monitor log files using `monitor` command
  - J-Web displays diagnostics about the platform
- System logs (syslog)
  - Log files contain a wealth of invaluable information
    - CLI **show log log-file-name** command
    - Remember to use pipe for added functionality
  - Chapter 3 discusses syslog details

### Troubleshoot Using the Craft Interface

You can use the Craft Interface to troubleshoot chassis problems. Some Juniper Networks platforms use LEDs on the Craft Interface to indicate the status of various chassis components. The M40e, M320, T320, T640, and TX Matrix platforms use the LCD to display general system status and a listing of any alarms that are currently active.

### Troubleshoot Using the CLI and J-Web

The primary means of controlling and troubleshooting JUNOS Software, routing protocols, network connectivity, and the hardware is to execute various operational mode commands from the CLI. The CLI provides commands that let you display the status of the various hardware components and monitor the log files. Also, you could use the J-Web to monitor, diagnose, and analyze hardware problems.

### Troubleshoot Using Syslog Messages

The various system logs maintained by JUNOS Software and the various processes that run on top of the JUNOS kernel contain a wealth of information regarding the operational status of a given system. The information stored in system logs is normally more detailed than that displayed on the Craft Interface. Do not forget to leverage the CLI pipe function to simplify the task of parsing through large log files for symptoms of abnormal operation or hardware failure.



## Agenda: JUNOS Platforms Hardware Troubleshooting

- Hardware Troubleshooting Overview
- Power On, Power Off, and Boot Media
- Using the CLI to Troubleshoot
- Case Study

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-5

### Power On, Power Off, and Boot Media

The slide highlights the topic we discuss next.

## Powering On and Powering Off

- Powering on:
  - Connect all cables
  - Turn on one power supply
  - Turn on the second power supply
    - It can take up to 60 seconds for accurate power supply and PEM status indications
- Powering off:
  - Shutdown JUNOS Software
    - CLI `request system halt` command
  - Turn off power supplies

### Powering On JUNOS Platforms

Each device can be equipped with two redundant, load-sharing power supplies of the same type, either AC or DC. 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 plugs 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 you intend to install the device. 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.

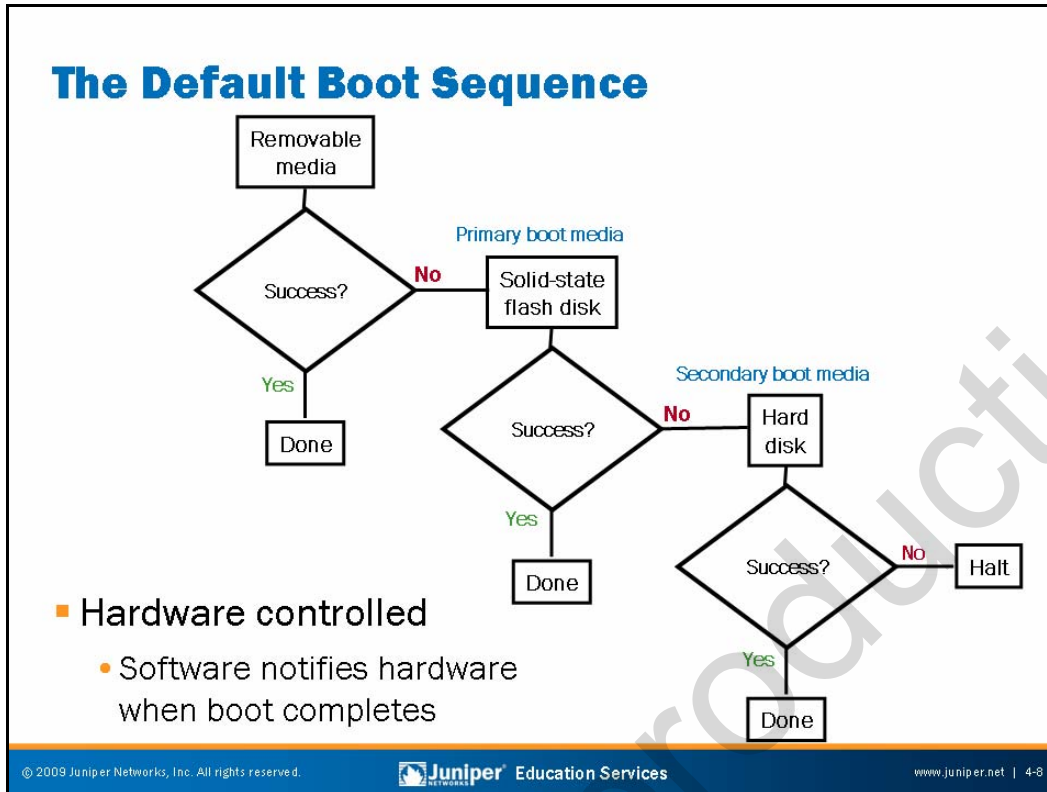
Although the specifics of each power supply and Power Entry Module (PEM) vary by platform, note that after a power supply is powered on, it can take up to 60 seconds for status indicators—such as LEDs on the power supply and `show chassis` commands—to indicate that the power supply is functioning normally. You should ignore error indicators that appear during the first 60 seconds.

*Continued on next page.*

## Powering Off JUNOS Platforms

If you want to power off a device, we recommend that you do a *graceful shutdown* as opposed to simply removing power from the system. A graceful shutdown initiates when you issue a **request system halt** command. Once the system properly halts, you can safely turn off the power supplies. Note that file system corruption can occur if you remove power before executing a graceful shutdown because of the multitasking nature of JUNOS Software.

Not for Reproduction



### Hardware Controls the Boot Sequence

At power on, as the device begins the boot process, it first attempts to start the image of software from the removable media if it is installed in the Routing Engine. If this attempt fails or if no media is present, the device next tries to boot from the image of software on the flash disk and then finally from the hard disk.

This sequence is controlled by hardware that waits for a special signal from the JUNOS Software kernel, indicating a successful boot. If the hardware does not receive the signal after a few minutes, it forces the system to boot from the next available device in the boot chain.

## Boot Devices and Media

- Three boot media options:
  - Removable media
    - Used for install and upgrade, normally left empty
  - Flash disk
    - Solid-state nonrotating media
    - Primary source for booting software
  - Hard disk
    - Traditional rotating media
    - Secondary source for booting software
- CLI option to identify boot source at next reboot:

<pre>user@host&gt; request system reboot media ? Possible completions: compact-flash disk</pre>	<pre>user@host&gt; request system reboot media ? Possible completions: internal usb</pre>
<pre>Standard boot off flash device Boot off hard disk</pre>	<pre>Boot from internal NAND flash Boot off USB device</pre>
M10i	SRX210

### Three Forms of Boot and Storage Media

JUNOS platforms generally support three forms of boot and storage media:

- *Removable media:* Depending on the platform model, your device might have a PC card slot (which reads flash disks), a compact-flash slot, or a universal serial bus (USB) port. A copy of JUNOS Software on removable media ships with most JUNOS platforms.
- *Flash disk (nonrotating disk):* Some JUNOS platforms ship with JUNOS Software preinstalled on the flash disk. For example, the flash disk is the primary boot device for M Series and T Series platforms.
- *Hard disk (rotating disk):* Some JUNOS platforms have a backup copy of JUNOS Software preinstalled on the hard disk. This disk also stores system log files and diagnostic dump files.

A JUNOS platform typically boots either from the flash disk or from the hard disk. (While it is possible to boot the device from the removable media disk, you typically do not do so.) We refer to these disks as the boot media. The disk from which the device boots is named the *primary boot medium*, and other disks are *secondary boot media*. Depending on the platform, the primary boot medium is generally the flash disk, and the secondary boot medium is generally the hard disk.

*Continued on next page.*

### CLI Option to Control Boot Media at the Next Reboot

The slide illustrates the choices that you could make when selecting from which media to boot for two types of platforms—the M10i and the SRX210. When issuing a **request system reboot** command, you can flag which boot medium should be used next by including the **media** switch. Note that this flag is temporary and affects only the next reboot.

Not for Reproduction

## Selecting Boot Medium at Boot Time

- Select the boot device at the boot loader prompt
  - Your computer must be attached through the console as the device boots:

```
Award Modular BIOS v4.51PG, An Energy Star Ally
Copyright (C) 1984-98, Award Software, Inc.

BIOS Version 1.2
11/02/2004-i440GX-SMC67X-2A69TU00C-00

Will try to boot from :
PCMCIA ATA Flash Card
Compact Flash
Primary IDE Hard Disk
Ethernet

Trying to Boot from PCMCIA ATA Flash Card

Trying to Boot from Compact Flash
. . . .
FreeBSD/i386 bootstrap loader, Revision 1.1
(builder@elliath.juniper.net, Wed Dec 17 23:00:37 UTC 2008)
Loading /boot/defaults/loader.conf
/kernel text=0xa2f02c data=0x6d248+0x81120 syms=[0x4+0x9f750+0x4+0xaae32]

Hit [Enter] to boot immediately, or space bar for command prompt.
Type '?' for a list of commands, 'help' for more detailed help.
OK ?
Available commands:
  reboot      Reboot the system
  . . . .
  nextboot   set next boot device
  . . . .
```

Reboot from next device in the boot order

Select boot device used at the next reboot

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 4-11

### Choose the Boot Device at Boot Time


At times you might want to manually choose to boot from alternative media even though the flash file system is bootable. You can tell the system to boot from a given medium to override its desire to always boot from flash disk.

The preferred way to force a boot from the hard disk is to issue a **request system reboot** CLI command with the **media** switch and **disk** argument. You can also select a boot device during the actual boot process when you connect through the console. This technique is useful if you are unable to access the CLI, perhaps because of a software fault in the flash file system (a fault that leaves the flash medium bootable but unusable from a JUNOS Software CLI perspective).


To select an alternative medium at the time of boot, watch for the boot loader screen shown on the slide. When prompted, enter a space to access the boot loader command prompt. From the prompt you can get some help with the ? key. As indicated in the initial help screen, typing **reboot** causes the device to reboot from the next device on the boot list. This device should be the hard disk because the flash disk is normally the first boot device. When needed, you can manually select the device to use when the **reboot** command executes by first using the **nextboot** command with the desired medium type specified as an argument.

## Visible Activity at Startup—Typical

- Craft Interface LCD display:
  - Idle mode: Cycles through various status displays
  - Alarm mode: Displays alarms in order of severity
- Craft Interface LEDs:
  - LEDs for FPCs, DPCs, PICs, RE, CBs, and others
    - Blinking green indicates test is in progress
    - Solid green indicates success; solid red indicates failure
- Craft Interface: Front panel alarm LEDs



The T640 Craft Interface panel



The SRX5000 Series front panel

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 4-12

### Craft Interface Liquid Crystal Display

As supported JUNOS platforms boot, the current status of the boot process displays on the Craft Interface LCD.

### Craft Interface LEDs

A series of diagnostic tests are performed on the Flexible PIC Concentrators (FPC) during the boot process. Blinking LEDs indicate tests in progress. They become solid after conclusion of the testing period. Depending upon the platform, the Craft Interface might also support LEDs for host module (Routing Engine and Control Board combination) or Switch Interface Board (SIB) status.

### Alarm LEDs Illuminate as Needed


Should any red or yellow alarms arise, the corresponding alarm LED illuminates on the Craft Interface. To see the specifics relating to a given alarm, you can look at the LCD on the Craft Interface (when present) or use the command `show chassis alarms`.



## Power Supply and PEM Indicators

- Power supplies and PEMs have their own status indicators
  - Some platforms require at least 60 seconds for proper status indications
  - Each power supply also has LEDs on the Craft Interface or front panel

T640 power supply status indicators



SRX5000 Series power supply status indicators



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 4-13

### Power Supply and Power Entry Module LEDs

Depending upon the platform and power supply model, one or more status LEDs on each power supply or Power Entry Module (PEM) might exist that you can use to determine if a power supply is functioning normally. Note that for some platforms you must wait at least 60 seconds after applying power to a power supply before you can expect to see meaningful status indications. The self-test button present on some power supplies should never be used on a production system; this button is for engineering and Juniper Networks Technical Assistance Center (JTAC) use only.

## Agenda: JUNOS Platforms Hardware Troubleshooting

- Hardware Troubleshooting Overview
- Power On, Power Off, and Boot Media
- Using the CLI to Troubleshoot
- Case Study

### Using the CLI to Troubleshoot

The slide highlights the topic we discuss next.

## Displaying Chassis Inventory

```
(master)
user@t640> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
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
SCG 0         REV 09   710-003423   HF9311
SCG 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
```

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 4-15

### Displaying Chassis Inventory

The output of the **show chassis hardware** command displays the hardware components installed in the platform. This command is useful when troubleshooting or upgrading your device. The (edited) sample shown on the slide is from a T640 router with redundant Routing Engines (RE).

The following are the **show chassis hardware** command output fields:

- **Item:** Shows information for the chassis component about the backplane, the power supplies, the maxicab (the connection between the Routing Engine and the backplane), the System Control Board (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.

*Continued on next page.*

## Displaying Chassis Inventory (contd.)

JUNOS Software has a **show chassis hardware clei-models** command. The output from this command provides information in a format suitable for conducting inventory. The **clei-models** option means Common Language Equipment Identifier Code barcode and model number for orderable field-replaceable units (FRUs). The following sample output is from an M320 router:

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item                Version  Part number  CLEI code          FRU model number
Midplane            REV 07    710-009120   CHAS-BP-M320-S
FPM Display         REV 05    710-009351   CRAFT-M320-S
CIP                 REV 05    710-005926   CIP-M320-S
PEM 0               Rev 05    740-009149   PWR-M-AC-S
PEM 1               Rev 05    740-009149   PWR-M-AC-S
PEM 2               Rev 05    740-009149   PWR-M-AC-S
PEM 3               Rev 05    740-009149   PWR-M-AC-S
Routing Engine 0    REV 07    740-014082   RE-A-2000-4096-S
Routing Engine 1    REV 07    740-014082   RE-A-2000-4096-S
. . .
```

## Displaying Alarm Conditions

```
user@host> show chassis alarms
1 alarm is currently active

Alarm time                Class  Description
2009-02-09 21:30:07 UTC   Major  Power Supply B not providing
power
```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-17

### Listing Alarm Conditions

The **show chassis alarms** command lists all of the alarm conditions that currently exist in the device. 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 device) nor does it extinguish the alarm LEDs. In addition, new alarms that occur after silencing an external device reactivate the external device.

The following are the **show chassis alarms** output fields:

- **Alarm time:** Displays the date and time of the alarm;
- **Class:** Displays the severity class for this alarm (it can be minor or major); and
- **Description:** Displays 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 normal speed
      . . .
```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-18

### Displaying Environmental Information

The **show chassis environment** command displays environmental information about the device chassis, including the temperature, and information about the fans, power supplies, and Routing Engine. The truncated example is from a T640 platform.

The following are the output fields:

- **Power:** Displays information about each power supply. The status can be **OK**, **Testing** (during initial power-on), **Failed**, or **Absent**. For the M120, M320, and T Series platforms, information displays about the PEM.
- **Temp:** Displays the temperature of air flowing through the chassis.
- **Fans:** Displays information about the fans. The status can be **OK**, **Testing** (during initial power-on), **Failed**, or **Absent**. **Measurement** indicates whether the fans are spinning at normal or high speed.
- **Other:** Depending upon the platform, various other fields might be present. For example, for T Series platforms, the display includes information on the SONET Clock Generator (SCG), Control Board (CB), SIBs, the Switch Processor Mezzanine Board (SPMB) and the Connector Interface Panel (CIP). The sample on the slide does not show these fields.

## Displaying CPU Temperature

```

user@M10i> show chassis routing-engine
Routing Engine status:
Routing Engine status:
Slot 0:
  Current state           Present
Slot 1:
  Current state           Master
  Election priority       Backup (default)
  Temperature              38 degrees C / 100 degrees F
  CPU temperature         39 degrees C / 102 degrees F
  DRAM                    1536 MB
  Memory utilization      21 percent
  CPU utilization:
    User                   0 percent
    Background              0 percent
    Kernel                  3 percent
    Interrupt               0 percent
    Idle                    97 percent
  Model                   RE-850
  Serial ID               1000591462
  Start time              2009-01-29 18:25:59 UTC
  Uptime                  4 days, 3 hours, 6 minutes, 2 seconds
  Last reboot reason      Router rebooted after a normal shutdown.
  Load averages:         1 minute  5 minute  15 minute
                          0.14      0.18      0.09
...

```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-19

### CPU Temperature

In addition to the ambient temperature surrounding the system components, you can see the actual CPU temperature of the Routing Engine.

## Displaying Craft Interface

```

(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 *  *  *  *  *
    
```

© 2009 Juniper Networks, Inc. All rights reserved.  www.juniper.net | 4-20

### Displaying Craft Interface

The **show chassis craft-interface** command shows all current messages. The capture shown is from a T640. Output fields include the following:

- **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 the status of the Front Panel System LEDs. A dot (.) indicates the LED is not lit. An asterisk (\*) indicates that the LED is lit.
- **Front Panel Alarm Indicators:** Displays the status of the Front Panel Alarm Indicators. A dot indicates the relay is off. An asterisk indicates that the relay is active.

*Continued on next page.*



### Displaying Craft Interface (contd.)

- **Front Panel FPC LEDs:** Displays the status of the Front Panel FPC LEDs. A dot indicates the LED is not lit. An asterisk indicates that the LED is lit.
- **MCS, SFM, SCG, CB, and SIB LEDs:** Displays the status of the Miscellaneous Control Subsystem (MCS), SCG, CB, Switching and Forwarding Module (SFM), and SIB LEDs as supported by a given platform. A dot indicates that the LED is not lit. An asterisk indicates that the LED is lit. When neither a dot nor an asterisk displays, no board is present in that slot.

## System Control Board Status

- `show chassis (feb|scb|cfeb|sfm slot|ssb slot)`

- Displays information about controller boards (FEB, CFEB, SCB, SFM, or SSB)

```

user@m10i> show chassis cfeb
CFEB status:
Slot 0 information:
State
Intake temperature      37 degrees C / 98 degrees F
Exhaust temperature     43 degrees C / 109 degrees F
CPU utilization          2 percent
Interrupt utilization    0 percent
Heap utilization         4 percent
Buffer utilization       25 percent
Total CPU DRAM           256 MB
Internet Processor II   Version 2, Foundry IBM, Part number 164
Start time:              2009-02-02 21:28:51 UTC
Uptime:                  9 minutes, 9 seconds
Slot 1 information:
State                    Backup
    
```

High CPU utilization levels normally indicate a high volume of exception traffic

### Displaying System Control Board Status

The `show chassis (feb | scb | | cfeb | sfm slot | ssb slot)` command displays information about the system controller boards—either Forwarding Engine Board (FEB), Compact Forwarding Engine Board (CFEB), SCB, SFM, or System and Switch Board (SSB). The following are the output fields:

- `Intake temperature`: Displays the temperature of the air passing by the controller in both Celsius and Fahrenheit;
- `Exhaust temperature`: Displays the temperature of the air flowing past the exhausts;
- `CPU utilization`: Displays the total percentage of the CPU used by the controller's processor;
- `Interrupt utilization`: Out of the total CPU percentage in use by the controller's processor, displays the percentage in use for interrupts;
- `Heap utilization`: Displays the percentage of heap space in use by the controller's processor;
- `Buffer utilization`: Displays the percentage of buffer space in use by the controller's processor;
- `DRAM`: Displays the total DRAM available to the controller's processor;

*Continued on next page.*

### Displaying System Control Board Status (contd.)

- `Start time`: Displays the time when the controller started running; and
- `Uptime`: Displays how long the controller has been running.

Because the controller board 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 for processing exception traffic. This traffic tends to take the form of Packet Forwarding Engine (PFE)-generated ICMP error messages or traffic that requires direction to the host RE for processing. Examples of exception traffic include Internet Control Message Protocol (ICMP) echo exchanged packets with expired time to live (TTL), which is one of the IP options, or traffic that is sampled or counted (or both) as part of a firewall filter.

## Displaying FPC Status

```
{master}
user@t640> show chassis fpc
```

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Empty		0	0	0	0	0
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						

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-24

### Displaying FPC Status

The **show chassis fpc** command displays the status of the installed FPCs. The sample is from a T640 platform. The following are the output fields:

- **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 Celsius.
- **CPU Utilization (%):** Displays the total percentage of CPU in use by the FPC's processor.
- **Interrupt CPU Utilization (%):** Of the total CPU percentage used by the FPC's processor, displays the percentage in use for interrupts.
- **Memory DRAM:** Displays the total DRAM (in megabytes) available to the FPC's processor.
- **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 (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@m10i> show chassis fpc detail 0
Slot 0 information:
  State                               Online
  Logical slot                         0
  Total CPU DRAM                       256 MB
  Total SRAM                           4 MB
  Total SDRAM                          64 MB
  Total notification SDRAM            12 MB
  I/O Manager ASIC information         Version 3.1, Foundry IBM,
  Part number 0
  Start time                           2009-02-02 21:28:53 UTC
  Uptime                               11 minutes, 51 seconds

```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-25

### Displaying the Status of a Specific FPC

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 sample output fields are as follows:

- State: Displays the state of the FPC slot:
  - 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 that the FPC is online and running.
  - Present: Indicates that the chassis process detects the FPC but it is either not supported by the current version of JUNOS Software or it is in the wrong slot. The output also states either `Hardware Not Supported` or `Hardware Not In Right Slot`. FPC is coming up but it is not yet online.

*Continued on next page.*

### Displaying the Status of a Specific FPC (contd.)

- `Probed`: Indicates that the probe is complete and awaits PFE restart.
- `Probe-wait`: Indicates that the slot is waiting for probing.
- `Logical slot`: Displays the slot number.
- `Total CPU DRAM`: Displays the amount of DRAM available to the FPC's CPU.
- `Total SRAM`: Displays the amount of SRAM in use by the FPC's CPU.
- `Total SDRAM`: Displays the total amount of memory used for storing packets and notifications.

Please note that depending on the platform, the output field details might vary.

## Displaying PIC Status

- Display information for all PICs or for PICs in a particular slot:

```

user@M10i> show chassis fpc pic-status
Slot 0   Online      E-FPC
  PIC 0   Online      4x F/E, 100 BASE-TX
  PIC 1   Online      2x OC-3 SONET, MM
Slot 1   Online      E-FPC
  PIC 0   Online      1x G/E, 1000 BASE-SX
  PIC 3   Online      Adaptive Services

user@M10i> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
Type                2x OC-3 SONET, MM
ASIC type            D chip
State                Online
PIC version          1.4
Uptime              14 minutes, 19 seconds

```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-27

### Displaying PIC Status

The **show chassis fpc pic-status** command displays information for all PICs. The following are the output fields:

- State: Indicates the state of the FPC slot, which can be one of the following:
  - Dead: Indicates that the FPC is held in reset because of errors.
  - Diag: Indicates that the slot is being ignored while the FPC runs 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.
  - Present: Indicates that the chassis process detects the FPC but it is either not supported by the current version of JUNOS Software or it is in the wrong slot. The output also states either `Hardware Not Supported` or `Hardware Not In Right Slot`. FPC is coming up but it is not yet online.

*Continued on next page.*

### Displaying PIC Status (contd.)

- `Probed`: Indicates that the probe is complete and that the slot awaits PFE restart.
- `Probe-wait`: Indicates that the slot is waiting for probing.
- `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.

Not for Reproduction



## Displaying Routing Engine Status

```

user@m10i> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Present
Slot 1:
  Current state           Master
  Election priority       Backup (default)
  Temperature              38 degrees C / 100 degrees F
  CPU temperature          39 degrees C / 102 degrees F
  DRAM                     1536 MB
  Memory utilization       21 percent
  CPU utilization:
    User                   0 percent
    Background              0 percent
    Kernel                  3 percent
    Interrupt               0 percent
    Idle                    97 percent
  Model                   RE-850
  Serial ID                1000591462
  Start time               2009-01-29 18:25:59 UTC
  Uptime                   4 days, 3 hours, 18 minutes, 57 seconds
  Last reboot reason       Router rebooted after a normal shutdown.
  Load averages:          1 minute   5 minute   15 minute
                          0.00       0.02      0.03

```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-29

### Displaying Routing Engine Status

The **show chassis routing-engine** command displays information about the Routing Engine. The following are the output fields:

- **Slot:** Indicates the slot number for the RE on systems that support RE redundancy;
- **Current state:** Indicates the current state of the RE on systems that support RE redundancy;
- **Election priority:** Specifies election priority for the RE on systems that support RE redundancy;
- **Temperature:** Displays the temperature of the air flowing past the RE;
- **DRAM:** Displays the total DRAM available to the RE's processor;
- **CPU utilization:** Displays information about the RE's CPU utilization, which include the following:
  - **User:** Displays the percentage of CPU time in use by user processes;
  - **Background:** Displays the percentage of CPU time in use by background processes;
  - **Kernel:** Displays the percentage of CPU time in use by kernel processes;

*Continued on next page.*

## Displaying Routing Engine Status (contd.)

- `Interrupt`: Displays the percentage of CPU time in use by interrupt processes; and
- `Idle`: Displays the percentage of CPU time that is idle.
- `Model`: Displays the RE model;
- `Serial ID`: Provides the identification number of the RE in this slot for systems that support RE redundancy;
- `Start time`: Displays the time at which the Routing Engine started running;
- `Uptime`: Displays how long the Routing Engine has been running;
- `Last reboot reason`: Provides a reason for the last reboot; and
- `Load averages`: Displays the Routing Engine load averages for the last 1, 5, and 15 minutes.

For systems with redundant REs installed, you can specify the RE slot number or see information about all REs installed in the system, as in the following example taken from an M320 platform:

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority      Master (default)
  Temperature             44 degrees C / 111 degrees F
  CPU temperature        51 degrees C / 123 degrees F
  DRAM                   3584 MB
  Memory utilization     11 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel                3 percent
    Interrupt            0 percent
    Idle                 97 percent
  Model                  RE-A-2000
  Serial ID              1000702757
  Start time             2009-02-06 07:56:24 PST
  Uptime                 11 days, 1 hour, 20 minutes, 36 seconds
  Last reboot reason     Router rebooted after a normal shutdown.
  Load averages:        1 minute   5 minute  15 minute
                       0.00         0.05     0.05

```

*Continued on next page.*

## Displaying Routing Engine Status (contd.)

Routing Engine status:

Slot 1:

Current state	Backup
Election priority	Backup (default)
Temperature	47 degrees C / 116 degrees F
CPU temperature	59 degrees C / 138 degrees F
DRAM	3584 MB
Memory utilization	13 percent
CPU utilization:	
User	47 percent
Background	0 percent
Kernel	24 percent
Interrupt	1 percent
Idle	29 percent
Model	RE-A-2000
Serial ID	1000699981
Start time	2008-12-30 16:07:53 PST
Uptime	48 days, 17 hours, 9 minutes, 2 seconds
Last reboot reason	0x1:power cycle/failure

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: 1.5
```

## Displaying PFE Errors

- The `show pfe statistics error` command displays PFE ASIC-related errors
  - Use `clear pfe statistics` (hidden command) to reset:

```

user@host> show pfe statistics error
PFE error statistics:
-----
C chip      A1 chip      A2 chip
-----
0           0           N/A scan fail
0           0           N/A A1<->C CRC error
0           N/A         0 A2<->C CRC error
N/A        0           0 A<->B CRC error
B chip slots:
-----
0           1           2           3
-----
0           0           0           0 scan fail
0           0           0           0 A1->B CRC error
0           0           0           0 A2->B CRC error
0           0           0           0 correctable ECC error
0           0           0           0 uncorrectable ECC error
0           0           0           0 multiple ECC errors
0           0           0           0 B->HS link error
0           0           0           0 A1->Bm error
0           0           0           0 A2->Bo error
0           0           0           0 write buffer overflow
0           0           0           0 Bo FIFO sync error
0           0           0           0 Bo FIFO size error
0           0           0           0 Bo stream stuck error
0           0           0           0 Bo SRAM parity error
    
```

Shared fabric and route lookup errors

PIC and FPC I/O errors

### Displaying PFE Errors

The `show pfe statistics error` command displays information about errors that might occur within the PFE’s application-specific integrated circuits (ASIC) or internal communications paths. To clear PFE statistics, use the CLI’s hidden `clear pfe statistics` command.

The slide provides a sample display for error statistics taken from an M10i platform.

The key reason for issuing the `show pfe statistics error` command is to determine if a system has a chronic error condition versus a transient burst of errors, as might be caused by incorrect FPC removal. Put another way, PFE errors are primarily of concern when you observe the error counts to be incrementing when the system is in an otherwise stable state (for example, you are not removing or inserting any FRUs).

## Bouncing PICs and FPCs

- Can take PICs and FPCs online and offline from the CLI
  - The resulting *soft-boot* might prevent the need for more drastic measures such as a system reboot

```

user@host> request chassis ?
Possible completions:
  cluster          Chassis cluster related requests
  fpc              Change Flexible PIC Concentrator status
user@host> request chassis fpc ?
Possible completions:
  offline         Turn FPC offline
  online          Bring FPC online
  restart         Restart FPC
  slot           FPC slot number (0..1)

user@host> request chassis fpc slot 0 restart
Restart initiated, use "show chassis fpc" to verify

```

© 2009 Juniper Networks, Inc. All rights reserved.



www.juniper.net | 4-33

### Restarting Hardware Components

You can restart or take an FPC online and offline from the CLI with a **request chassis fpc slot-number [restart | online | offline]** command. Similarly, a PIC can be bounced with a **request chassis pic fpc-slot slot-number pic-slot slot-number [online | offline]** command. The example illustrates the basic CLI syntax for restarting the FPC in slot 0. In some cases, bouncing an FPC or a problematic PIC might alleviate the need for more drastic actions like a system reboot. The following capture shows the result of restarting an FPC as it progresses through the *offline*, *present*, and *online* status through the output of a series of **show chassis fpc** commands:

*Continued on next page.*

### Restarting Hardware Components (contd.)

user@host> **show chassis fpc**

Slot	State	Temp (C)	CPU Utilization (%) Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Utilization (%) Buffer
0	Offline	---Restarted by cli command---					
1	Online	37	1	0	1024	2	49
2	Online	38	1	0	1024	2	49
3	Online	40	1	0	1024	2	49
4	Online	39	1	0	1024	2	49
5	Empty						
6	Empty						
7	Empty						

user@host> **show chassis fpc**

Slot	State	Temp (C)	CPU Utilization (%) Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Utilization (%) Buffer
0	Present	39					
1	Online	37	1	0	1024	2	49
2	Online	38	1	0	1024	2	49
3	Online	40	1	0	1024	2	49
4	Online	39	1	0	1024	2	49
5	Empty						
6	Empty						
7	Empty						

user@host> **show chassis fpc**

Slot	State	Temp (C)	CPU Utilization (%) Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Utilization (%) Buffer
0	Online	37	0	0	0	0	0
1	Online	37	1	0	1024	2	49
2	Online	38	1	0	1024	2	49
3	Online	40	1	0	1024	2	49
4	Online	39	1	0	1024	2	49
5	Empty						
6	Empty						
7	Empty						

## Viewing Boot Messages

```

user@host> show system boot-messages
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 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 9.5R1.8 #0: 2009-04-13 19:11:52 UTC
    builder@ianath.juniper.net:/volume/build/junos/9.5/release/9.5R1.8/obj-1386/bsd/sys/compile/JSR
Timecounter "i8254" frequency 1193182 Hz quality 0
CPU: Intel(R) Celeron(R) CPU 2.53GHz (2533.44-MHz 686-class CPU)
    Origin = "GenuineIntel" Id = 0xf49 Stepping = 9

Features=0xbfebfbff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,APIC,SEP,MTRR,PGE,MCA,CMOV,PAT,PSE36,CLFLUSH,DTS,ACPI,MM
X,FXSR,SSE,SSE2,SS,HTT,TM,PBE>
    Features2=0x651d<SSE3,RSVD2,MON,DS_CPL,TM2,CNTX-ID,CX16,<b14>>
    AMD Features=0x20100000<NX,LM>
real memory = 1073741824 (1024 MB)
avail memory = 541659136 (516 MB)
Initializing JSR platform properties ..
cpu0 on motherboard
pcib0: <Intel E7221 Memory Controller Hub (MCH) for DDR2> pcibus 0 on motherboard
pcid0: <PCI bus> on pcib0
pcib1: <PCI-PCI bridge> irq 10 at device 1.0 on pcid0
pcil1: <PCI bus> on pcib1
pcib2: <PCI-PCI bridge> at device 0.0 on pcil1
. . .
da0 at umass-sim0 bus 0 target 0 lun 0
da0: <SanDisk U3 Cruzer Micro 3.21> Removable Direct Access SCSI-2 device
da0: 1.000MB/s transfers
da0: 979MB (2006673 512 byte sectors: 64H 32S/T 979C)
Trying to create bootdev, rootpartition ad0s1a
Trying to mount root from ufs:/dev/ad0s1a

```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-35

### Viewing Boot Messages

The slide shows an example of an operator displaying the contents of the boot log by issuing a **show system boot-messages** command. JUNOS Software writes this file during the system boot, and the file contains the various boot-up messages generated during the last power cycle and boot or reboot.

In some cases, JUNOS Software reports hardware errors and device malfunctions at boot time. The truncated capture on the slide does not show any abnormal events.

## Displaying System Storage

- Displays amount of storage available on flash and rotating disks:

```
user@M320-re0> show system storage
```

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/ad0s1a	885M	143M	671M	18%	/
devfs	1.0K	1.0K	0B	100%	/dev
devfs	1.0K	1.0K	0B	100%	/dev/
/dev/md0	30M	30M	0B	100%	/packages/mnt/jbase
/dev/md1	157M	157M	0B	100%	/packages/mnt/jkernel-9.5R1.8
/dev/md2	56M	56M	0B	100%	/packages/mnt/jpfe-M320-9.5R1.8
/dev/md3	4.4M	4.4M	0B	100%	/packages/mnt/jdocs-9.5R1.8
/dev/md4	44M	44M	0B	100%	/packages/mnt/jroute-9.5R1.8
/dev/md5	13M	13M	0B	100%	/packages/mnt/jcrypto-9.5R1.8
/dev/md6	25M	25M	0B	100%	/packages/mnt/jpfe-common-9.5R1.8
/dev/md7	2.0G	8.0K	1.8G	0%	/tmp
/dev/md8	2.0G	748K	1.8G	0%	/mfs
/dev/ad0s1e	98M	50K	90M	0%	/config
procs	4.0K	4.0K	0B	100%	/proc
/dev/ad2s1f	34G	7.5G	24G	24%	/var

/ = Partition on flash (885 MB)

/config = Partition on flash (98 MB)

/var = Partition on hard disk (34 GB)

Memory virtual disks

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 4-36

## 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 various file systems used by the device. Values display in 512 byte blocks. This command is equivalent to the UNIX **df** command.

The highlights on the slide indicate the device names used by the flash and rotating storage devices. In this case, the flash medium is device **ad0s1a**, while the hard disk is device **ad2s1f**. You can also see that JUNOS Software makes use of FreeBSD's virtual file system support to mount images of jbundle components on memory virtual disks. These RAM disk devices always indicate being 100% full because of their read-only nature.

*Continued on next page.*



## Displaying System Storage (contd.)

Note that the output of the **show system storage** command might list the same flash and hard disk devices with different names due to changes in the underlying FreeBSD distribution on which the JUNOS Software version is based. For the curious, the device name `ad0s1a` has the following meaning:

- `ad` = IDE hard disk (the flash device emulates an IDE disk).
- `0` = The unit number for that device type—for example, the first IDE disk is unit 0.
- `s1` = Slice 1 for PC BIOS partition 1.
- `a` = The root (`/`) partition. A `b` partition type is for swap space, while a `c` partition type is used in dedicated mode (native BSD slice mode). Other partition types are for general use, such as the `e` designation for the `/config` partition.

## Displaying System Uptime and Users

- Can use system uptime to determine platform stability
  - Also indicates the user who committed the current configuration:

```
user@host> show system uptime
Current time: 2009-02-05 20:22:41 UTC
System booted: 2009-01-29 18:25:59 UTC (1w0d 01:56 ago)
Protocols started: 2009-02-02 21:28:10 UTC (2d 22:54 ago)
Last configured: 2009-01-23 18:07:20 UTC (1w6d 02:15 ago) by lab
8:22PM up 7 days, 1:57, 1 user, load averages: 0.06, 0.03, 0.02
```

- Knowing who is currently logged in and what they are doing might affect maintenance plans:

```
user@host> show system users
2:27PM UTC up 20 days, 22:05, 2 users, load averages: 0.01, 0.01, 0.00
USER      TTY      FROM          LOGIN@  IDLE WHAT
root      d0       -             2:23PM  1  -csh (csh)
lab       p0       10.0.24.1    2:27PM  -  -cli (cli)
```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-38

### Displaying Uptime

The **show system uptime** command displays the current time and information about how long the device, its software, and routing protocols have been running. The following are the output fields:

- **Current time:** Displays the current system time in UTC;
- **System booted:** Displays the date and time when the device last booted and how long it has been running;
- **Protocols started:** Displays the date and time when the routing protocols last started and how long they have been running;
- **Last configured:** Displays the date and time a configuration last activated (either by booting the device or issuing the **commit** command in configuration mode);
- **Time:** Displays the current time, in the local time zone;
- **Up:** Displays how long the device has been operational;
- **user:** Displays the number of users logged into the device; and
- **load averages:** Displays the load averages for the last 1 minute, 5 minutes, and 15 minutes.

*Continued on next page.*

## Display Users

The **show system users** command displays the currently logged in users, and displays what they are doing. The example on the slide shows that the root user is logged in at a c-shell while the lab user is in the CLI. Knowing that folks are actively logged into the device might factor in to a decision to perform disruptive actions like software upgrades. Note that the terminal name for a console port begins with d, while virtual terminal (vty) ports, such as result from Telnet or SSH connections, use p-style tty connections. This information is helpful when the goal is to disconnect a user with the **request system logout user** command because you must specify both the user name and the related terminal port, unless you use the **all** keyword.

## Parsing System Logs

- The CLI pipe function makes parsing log files easy
  - Search the messages and chassisd logs for entries like fail, kernel, core, error, and so forth
    - Use quotes and the pipe function to search for multiple items:
 

```
show log messages
| match "fpc | sfm | kernel | tnp"
```
  - Can you describe the nature of the hardware fault from these log entries?

```
user@Bangkok-re1> show log messages | match fail
Jan  8 16:33:16 Bangkok chassisd[2850]: snmp_ipc_try_connect: connect to master (unix
sock) failed: Connection refused, retry in 1
. . .
Feb  2 21:28:12 Bangkok-re1 chassisd[4446]: CHASSISD_BLOWERS_SPEED_FULL: Fans and
impellers being set to full speed [fan/blower missing/failed]

user@Bangkok-re1> show log messages | match fail | match Feb
Feb  2 21:28:10 Bangkok-re1 pfed: SNMP_NS_LOG_WARNING: Warning: Failed to connect to the
agentx master agent (/var/run/snmpd_agentx): Unknown host (/var/run/snmpd_agentx)
(Connection refused)
Feb  2 21:28:11 Bangkok-re1 cosd[5041]: SNMP_NS_LOG_WARNING: Warning: Failed to connect
to the agentx master agent (/var/run/snmpd_agentx): Unknown host (/var/run/snmpd_agentx)
(Connection refused)
. . .
```

### Parsing System Logs

The slide shows examples of how you can use the CLI to rapidly locate signs of trouble within a given log file. The syslog samples shown on the slide come from the messages file. The samples illustrate how you can use the CLI match function, which allows you to easily and effectively parse the system log files.

## Agenda: JUNOS Platforms Hardware Troubleshooting

- Hardware Troubleshooting Overview
- Power On, Power Off, and Boot Media
- Using the CLI to Troubleshoot
- Case Study

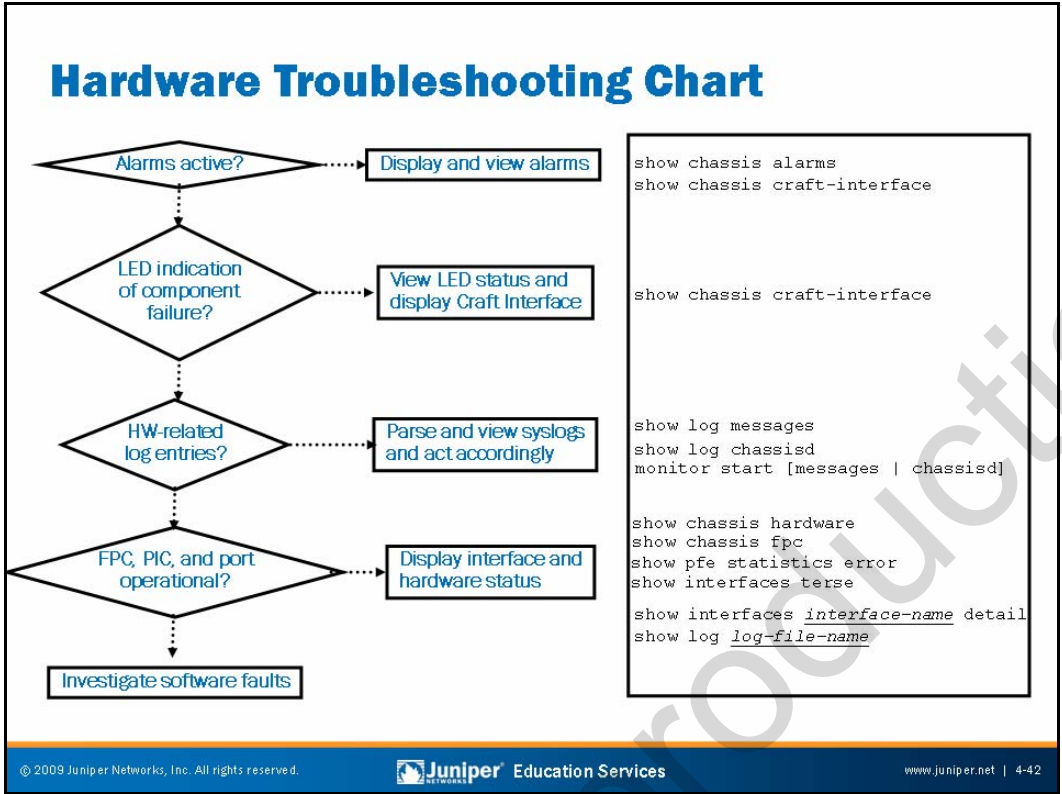
© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-41

### Case Study

The slide highlights the topic we discuss next.



### Hardware Troubleshooting Flow Chart

Troubleshooting is an artfully applied science. The intent of this statement is to highlight that even though many aspects of fault isolation have a basis in straightforward facts and physics, a certain degree of artistic license that determines how a particular technician decides to approach a problem always remains. Put another way, two individuals working with the same sets of tools and a common symptom might approach the act of fault analysis in completely different ways. For example, one person might always start with visual inspection while another opts to begin with interface loopback tests. In the end, it is hard to say that one approach is better than another, assuming that both individuals arrive at a similar conclusion, in a similar amount of time, with similar levels of minimal disruption.

The determination that troubleshooting is a mix of science and artwork is important because most agree that art is a subjective concept that must come from within the artist; although artistic skills and concepts can be taught, the receipt of such edification does not imply the student will actually create master artworks.

*Continued on next page.*

### Hardware Troubleshooting Flow Chart (contd.)

The artistic aspect of troubleshooting and the myriad ways in which a modern communications device might malfunction combine to make a definitive set of troubleshooting steps and procedures an unobtainable goal. The purpose of the hardware troubleshooting flow-chart shown on the slide is simply to provide a set of high-level steps designed to get you started with hardware fault analysis. Note that reasonable people might disagree on the exact ordering of the steps or on the particulars of the CLI commands that you could use to help isolate a hardware failure (for example, some might prefer the **extensive** switch to the **show interfaces** command, while the sample chart calls out the **terse** and **detail** switches).

## Hardware Case Study A (1 of 4)

- Case study background:

- The platform is an M320
- You see high-speed link errors in the system log

```

. . . .
host chassisd[4503]: CHASSISD_GBUS_RESET_EVENT: SIB#3 - Assert Board Reset
host chassisd[4503]: CHASSISD_FCHIP_HSR_ERROR: Fchip high-speed receiver (HSR) error:
SR:name=SIB3_F0,05, index=53, cookie=5 is still not cell aligned
host chassisd[4503]: CHASSISD_FM_ERROR: fm_ack_updt_nf_to_f: High-speed receiver (HSR)
link failed (SIB#3, Packet Forwarding Engine 1 on FPC 2)
. . . .

```

- SIB 3 goes offline and becomes unusable for forwarding traffic
- For a short time all FPCs that are still trying to use SIB 3 to send traffic through the switch fabric generate destination errors:

```

. . . .
host chassisd[4503]: CHASSISD_FM_ERROR: update_fe_planes: Packet Forwarding Engine got
destination errors on SIB (SIB#3, Packet Forwarding Engine 0 on FPC 7)
host fpc7 ichip_f_check_dest_errors: Fabric request time out for plane 0 dest 5
. . . .

```

### Hardware Case Study A: Part 1

The slide sets the stage for a sample hardware troubleshooting case study. We begin with the general description of the problem, which in this case indicates that you see high speed link (HSL) errors as reported in the system log file. The HSL interconnects the FPC, the PFE, the SIB, and the midplane of the M320 router. In addition to HSL error messages, SIB 3 goes offline and becomes unusable for forwarding the traffic. In fact, for a short time, all FPCs that are still trying to use SIB 3 to send traffic through the switch fabric generate destination errors as illustrated on the slide.



## Hardware Case Study A (2 of 4)

- What is wrong?
  - Faulty hardware components can cause HSL errors either on the FPC, the midplane, or the SIB
  - Based on this information, what possible causes can you eliminate?
  - What CLI commands might you issue to narrow down a possible cause?

### Hardware Case Study A: Part 2

Faulty hardware components can cause HSL errors either on the FPC, the midplane, or the SIB. Therefore, following the methodology of troubleshooting, you can eliminate possible causes of the problem, narrowing down your choices.

## Hardware Case Study A (3 of 4)

- Sample course of action:

1. Identify which SIB and FPC interconnect with the faulty HSL (CLI method shown):

```

user@host> show chassis fabric topology 3
backup@CR&PARTV1> show chassis fabric topology 3
fchip (mode)
in-links          state      out-links          state
-----
Sib #3 :
FPC0_T->SIB3_F0,00  UP        SIB3_F0,00->FPC3_B  UP
FPC0_B->SIB3_F0,01  RESET     SIB3_F0,01->FPC3_T  UP
FPC1_T->SIB3_F0,02  RESET     SIB3_F0,02->FPC2_B  UP
FPC1_B->SIB3_F0,03  RESET     SIB3_F0,03->FPC2_T  UP
FPC2_T->SIB3_F0,04  UP        SIB3_F0,04->FPC1_B  JB
    
```

The faulty HSL interconnects FPC2 with SIB3

2. Replace the SIB and bring it online:

```

user@host> request chassis sib slot 3 offline
Offline initiated, use "show chassis sibs" to verify
user@host> request chassis sib slot 3 online
Online initiated, use "show chassis sibs" to verify
    
```

3. If the new SIB comes online without HSL errors, then the cause of the problem was the faulty SIB → problem solved!

### Hardware Case Study A: Part 3

To find the defective hardware, you decide to identify the SIB and the FPC that interconnect with the faulty HSL. The **show chassis fabric topology** command illustrates that the faulty HSL interconnects FPC2 with SIB3. Next you decide to replace the SIB and bring it online. If the new SIB comes online without HSL errors, then you found the cause of the problem—the previous SIB in slot 3 was faulty. If the error persists, follow step 4 on the following slide.

## Hardware Case Study A (4 of 4)

- Sample course of action (contd.):
  4. If the error still persists on the same HSL, then replace the FPC and bring it online:

```
user@host> request chassis fpc slot 2 offline
Offline initiated, use "show chassis fpc" to verify
user@host> request chassis fpc slot 2 online
Online initiated, use "show chassis fpc" to verify
```

5. If the new FPC comes online without HSL errors, then the cause of the problem was the faulty FPC → problem solved!
6. If the error still persists on the same HSL, then replace the midplane → problem solved!

### Hardware Case Study A: Part 4

Now you must check the FPC. Specifically, you replace the FPC and bring it online. Once you bring the new FPC online, you check for HSL errors. If the new FPC comes online without HSL errors, then the cause of the problem was the faulty FPC and you solved the problem. Otherwise, you must replace the midplane. Finally, the problem should be solved.

## Hardware Case Study B (1 of 4)

### ■ Case study background:

- You replaced PIC 2 in FPC 4 in an M320 router
- You want to obtain the new PIC's serial number using the `show chassis hardware` command
- You issue the `show chassis hardware` command, but you do not see the new PIC's serial number
- What CLI command can help you obtain the new PIC's serial number?

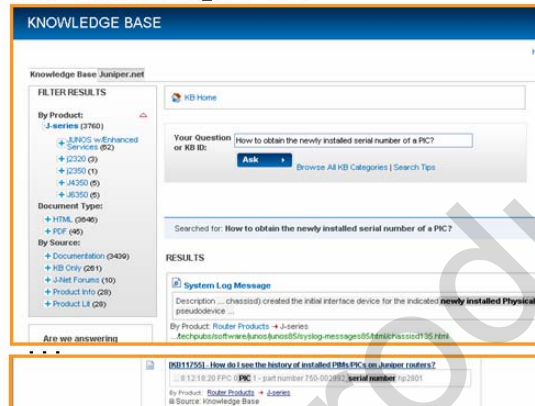
### Hardware Case Study B: Part 1

The slide sets the stage for another sample hardware troubleshooting case study. You replaced the PIC in slot 2 within FPC 4 in an M320 router. Now you want to obtain the new PIC's serial number. You are aware of the CLI **show chassis hardware** command that lists serial numbers for all the hardware components of the router. However, when you issue the command, you do not see the new PIC's serial number. How else can you obtain the new PIC's serial number?

## Hardware Case Study B (2 of 4)

- Course of action:

1. Search the technical documentation to determine if a command exists that can list the PIC serial number
2. Search the Knowledge Base



### Hardware Case Study B: Part 2

First try a search for the necessary information in the technical documentation. Remember the use of the Ctrl+click to select multiple products to search—this technique is useful if you become overwhelmed with results from a wide-open search but are still not exactly sure in which category your result will appear.

If you do not find the information, remember to search the JTAC Knowledge Base. The slide illustrates the result of a Knowledge Base search.

## Hardware Case Study B (3 of 4)

- Course of action (contd.):

- Read the details of KB11755:

**KNOWLEDGE BASE**

[KB Home](#)
[Back to Answers](#)
[Browse Knowledge Base Categories](#)
[Printer Friendly](#)

[Rate this Page](#)
[Recommend Change](#)
[Subscribe](#)
[Manage Document](#)

**How do I see the history of installed PIMs/PICs on Juniper routers?**

**SYNOPSIS:**  
The inventory log file keeps a record of serial numbers for parts that are or were in a chassis. This could be useful if you removed a part from the chassis but are unable to get the serial number off of the part. The inventory log will show you for example what FPC slot and PIC slot the part was in.

**PROBLEM:**  
One issue this could also help with includes if the part is in the chassis but the serial number is no longer showing in the command: `show chassis hardware`.

**SOLUTION:**  
Enter command `'show log inventory'` and match around the most recent time to get the serial number of the part.

**Knowledge Base ID:** KB11755

**Version:** 2.0

**Published:** 07 Oct 2008

**Updated:** 07 Oct 2008

**Categories:**


- [JUNOS](#)
- [J-series](#)
- [JUNOS\\_FIPS](#)
- [JUNOS-ES](#)
- [Logging/Syslog](#)

**Former Article ID:**

**Owner:** [Andree Ducey](#)

**Reputation Lvl/Pts:**  (0)

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 4-50

### Hardware Case Study B: Part 3

Browsing through the Knowledge Base, you locate useful information in KB11755. Once you read the details of that Knowledge Base entry, you realize that it reflects your situation and provides you with instructions on how to find the newly installed PIC's serial number.

## Hardware Case Study B (4 of 4)

- Course of action (contd.):
  4. Obtain the new PIC's serial number using the **show log inventory** CLI command
  5. Match on the month you installed the new PIC

```

user@host-re0> show log inventory | match "Jun 3"
Jun  3 07:57:07 CHASSISD release 9.5R1.8 built by builder on 2009-04-13
19:11:52 UTC
. . .
Jun  3 07:57:07 Midplane - part number 710-009120, serial number RB8559
Jun  3 07:59:07 FPC 4 PIC 0 - part number 750-013168, serial number WD5325
Jun  3 07:59:08 FPC 4 PIC 2 - part number 750-005634, serial number DG1725
Jun  3 07:59:08 FPC 4 PIC 3 - part number 750-003034, serial number RG4040
. . .

```

The new PIC's serial number



### Hardware Case Study B: Part 4

Following the directions from the Knowledge Base, you use the **show log inventory** command. To locate the information faster, you narrow down the displayed information to the installation date of the new PIC—June 3. The output of the command provides the information you need.

## Summary

- In this chapter, we:
  - Provided an overview of hardware troubleshooting tools
  - Described power on and power off procedures and boot media options
  - Troubleshoot JUNOS devices using visual indicators
  - Troubleshoot JUNOS devices using the CLI
  - Parsed log files for symptoms of hardware problems

### This Chapter Discussed:

- An overview of hardware troubleshooting tools;
- Power on, power off, and boot media options;
- Troubleshooting based on visual indicators;
- Troubleshooting based on the JUNOS Software CLI; and
- Parsing log files for indications of hardware problems.



## Review Questions


1. What is the correct procedure for powering off a JUNOS platform?
2. List three ways that you can use the JUNOS Software CLI to assist in hardware troubleshooting.
3. Describe two ways of determining if any chassis alarms are present.
4. What CLI command searches the `messages` file for all lines matching `fail` and `error`?

## Review Questions

- 1.
- 2.
- 3.
- 4.

## **Lab 2: Chassis Hardware Troubleshooting**

- Troubleshoot chassis hardware problems.

© 2009 Juniper Networks, Inc. All rights reserved.  Juniper Education Services www.juniper.net | 4-54

### **Lab 2: Chassis Hardware Troubleshooting**

The slide lists the objective of this lab.



# **Troubleshooting JUNOS Platforms**

## **Chapter 5: Interface Troubleshooting**

Not for Reproduction

## Chapter Objectives

- After successfully completing this chapter, you will be able to:
  - Describe physical and logical interface properties that require configuration
  - Deactivate and disable interfaces
  - Configure loopbacks and BERT testing
  - Use operational mode commands to monitor and troubleshoot a variety of interface types

### This Chapter Discusses:

- Physical and logical interface properties;
- Deactivating and disabling interfaces;
- Configuring loopbacks and the bit error rate test (BERT); and
- Using operational mode commands to monitor and troubleshoot a variety of interfaces and media types.

## Agenda: Interface Troubleshooting

- Interface Configuration Overview
- General Interface Troubleshooting
- Media-Specific Interface Troubleshooting
- Case Study

### Interface Configuration Overview

The slide lists the topics we cover in this chapter. We discuss the highlighted topic first.

## Interface Properties

- **Physical properties:**
  - Clocking
  - Scrambling
  - Frame check sequence
  - Maximum transmission unit
  - Data Link Layer protocol and keepalives
  - Diagnostic capabilities
    - Local, remote, and facility loopback
    - BERT
- **Logical properties:**
  - Protocol family (Internet, ISO, and MPLS)
  - Addresses (IP address and ISO NET address)
  - Virtual circuits (VCI, VPI, and DLCI)
  - Other characteristics

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-4

### Physical Properties

The following list provides details of the interface 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 and keepalives*: You can change the Data Link Layer protocol for the particular media type—for example, Point-to-Point Protocol (PPP) to Cisco High-Level Data Link Control (or Cisco HDLC)—and you can turn keepalives on or off; and
- *Diagnostic characteristics*: You can enable local or remote loopbacks or set up during a BERT test (see Chapter 6).

*Continued on next page.*

## Logical Properties

The following list provides details of the interface logical properties:

- *Protocol family*: Refers to the protocol family you want to use—`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 data-link connection identifier (DLCI), Virtual Path Identifier (VPI), Virtual Channel Identifier (VCI), or Virtual Local Area Network (VLAN) tag; and
- *Other characteristics*: Some other configurable options include Inverse ARP, traps, and accounting profiles.

## Deactivating Versus Disabling (1 of 2)

### ■ Deactivating

- Add the `inactive:` tag to a statement, effectively commenting out the statement or identifier from the configuration

```
[edit]
user@host# run show interfaces so-2/0/0 terse
Interface      Admin Link Proto  Local      Remote
so-2/0/0       up    up
so-2/0/0.0     up    up   inet    10.2.1.21/30
```

```
[edit]
user@host# deactivate interfaces so-2/0/0
[edit]
user@host# show interfaces
...
inactive: so-2/0/0 ( ...
```

- Upon the `commit`, you see the following output:

```
[edit]
user@host# run show interfaces se-2/0/0 terse
Interface      Admin Link Proto  Local      Remote
so-2/0/0       up    up
```

### Deactivating an Interface

In a configuration, you can deactivate statements and identifiers so that they do not take effect when you issue the `commit` command. Any deactivated statements and identifiers have the `inactive:` tag. They remain in the configuration but are not active when you issue a `commit` command.

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 Versus Disabling (2 of 2)

### ■ Disabling

- Disable an interface or a logical unit:

```
[edit]
user@host# run show interfaces so-2/0/0 terse
Interface      Admin Link Proto  Local      Remote
so-2/0/0       up    up
so-2/0/0.0     up    up   inet    10.2.1.21/30
```

```
[edit]
pe@cartman-re0# set interfaces so-2/0/0 disable
```

```
[edit]
user@host# show interfaces so-2/0/0
disable;
unit 0 {
  family inet {
    address 10.2.1.21/30;
  }
}
```

- Upon the **commit**, you see the following:

```
[edit]
user@host# run show interfaces so-2/0/0 terse
Interface      Admin Link Proto  Local      Remote
so-2/0/0       down   up
so-2/0/0.0     up    down inet    10.2.1.21/30
```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-7

### Disable Versus Deactivate

In some portions of the configuration hierarchy, you can include a **disable** statement to disable functionality. One example is disabling an interface by including the **disable** statement at the [edit interface *interface-name*] hierarchy level. When you deactivate a statement, JUNOS Software completely ignores that specific object or property and does not apply it when you issue a **commit** command. When you disable a functionality, it is activated when you issue a **commit** command but the software treats it as being down or administratively disabled.

## Interface Configuration Examples

### ATM interface with multiple units

```
[edit interfaces]
user@host# show at-0/2/1
description "SY to HK and DE";
atm-options {
  vpi 0 {
    maximum-vcs 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;
  }
}
```

### Gigabit Ethernet with inet and mpls support

```
[edit interfaces]
user@host# show ge-0/0/2
unit 0 {
  family inet {
    address 10.0.13.1/24;
  }
  family mpls;
}
```

### SONET interface running Frame Relay with keepalives (LMI) disabled

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

## Interface Configuration Examples

The 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 created a given configuration stanza.

Note that each configuration example makes use of at least one logical unit, and that you specify a protocol family and related logical properties at the unit level. The commands used to configure the Asynchronous Transfer Mode (ATM) interface shown on the slide are shown in the following output:

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

## Agenda: Interface Troubleshooting

- Interface Configuration Overview
- General Interface Troubleshooting
- Media-Specific Interface Troubleshooting
- Case Study

### General Interface Troubleshooting

The slide highlights the topic we discuss next.

## Interface Troubleshooting Overview

- **Understanding the demarcation:**
  - North America typically includes the CSU/DSU (CPE perspective) because it is owned by the customer
  - Other parts of the world, Europe, for example, typically excludes the CSU/DSU (CPE perspective) because equipment is owned by the telco
- **Topology determines troubleshooting approach—essentially three topology types to consider when troubleshooting:**
  - LAN or broadcast multiaccess (Fast Ethernet and Gigabit Ethernet)
  - Point-to-point (SONET/SDH, T3 and E3, T1 and E1, PPP, or Cisco HDLC)
  - Point-to-multipoint (SONET/SDH, T3 and E3, T1 and E1, Frame Relay or ATM-VC)
- **Tools available for each type vary**

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-10

### Understanding the Demarcation

Understanding the demarcation is important when troubleshooting a given problem. The model in North America is based on the customer providing, and thereby being responsible for, the CSU/DSU function. The telco in this environment does not have any means of verifying the *local-loop* or *tail* without getting the subscriber to set a loop back to the provider.

In Europe, the telco supplies the CSU/DSU device and is responsible for the verification and testing of the local-loop in addition to whatever segments might exist between the customer premises equipment (CPE).

### Topology Determines Approach

Three topology types exist for you to consider when troubleshooting:

- LAN or broadcast multiaccess (Fast Ethernet or Gigabit Ethernet);
- Point-to-point (SONET/SDH, T3 and E3, T1 and E1, PPP, or Cisco HDLC); and
- Point-to-multipoint (SONET/SDH, T3 and E3, T1 and E1, Frame Relay or Asynchronous Transfer Mode-virtual circuit (ATM-VC).

### Tools Available

The following pages discuss the troubleshooting tools available in JUNOS Software.

## Displaying Terse Interface Status

- The `show interfaces terse` command provides a quick view of interface status:

```

user@host> show interfaces so* terse
Interface      Admin Link Proto Local                               Remote
so-1/1/0       down  up
so-1/1/0.0     up    down inet 1.1.1.1/30
               up    down iso
so-1/1/1       up    down
so-1/1/1.0     up    down inet 2.2.2.2/30
               up    down iso
so-1/1/2       up    up
so-1/1/2.0     up    up  inet 3.3.3.3/30
. . .
    
```

Administratively disabled

Data Link Layer down

Data Link Layer up

Admin	Link	Meaning
down	down	Admin disabled
up	down	Router interface problem Interface misconfigured (encapsulation) Keepalive sequencing not incrementing CSU/DSU failure Carrier problem (noisy line or timing mismatches)

### Displaying Interface Status at a Glance

Use the `show interfaces terse` command to display a terse listing of all interfaces installed in the device along with their administrative and Data Link Layer status. The table on the slide explains the meaning of the Admin and Link status indications.

When an interface is administratively disabled, the physical interface has an Admin status of down and a Link status of up, and the logical interface has an Admin status of up and a Link status of down. The physical interface has a Link status of up because the physical link is healthy (no alarms). The logical interface has a Link status of down because the Data Link Layer cannot establish end to end.

When an interface is not administratively disabled and the Data Link Layer between the local device and the remote device is not functioning, the physical interface has an Admin status of up and a Link status of up while the logical interface has an Admin status of up and a Link status of down. The physical interface has a Link status of up because the physical link is healthy (no alarms). The logical interface has a Link status of down because the Data Link Layer cannot establish end to end.

If the Data Link Layer between the local device and the remote device is up and running, both the physical and logical interfaces have an Admin status of up and a Link status of up, as shown in the case of the so-0/1/2 interface on the slide.

## Standard Interface Display

```

lab@host> show interfaces so-0/1/2
Physical interface: so-0/1/2, Enabled, Physical link is Up
Interface index: 134, SNMP ifIndex: 28
Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
Loopback: None, FCS: 16, Payload scrambler: Enabled
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Link flags   : Keepalives
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 90939 (00:00:07 ago), Output: 90879 (00:00:04 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
CoS queues : 4 supported, 4 maximum usable queues
Last flapped : 2009-02-17 03:13:49 UTC (1w3d 11:12 ago)
Input rate : 0 bps (0 pps)
Output rate : 280 bps (0 pps)
SONET alarms : None
SONET defects : None
Logical interface so-0/1/2.0 (Index 67) (SNMP ifIndex 75)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
Protocol inet, MTU: 4470
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.0.31/24, Local: 10.0.31.2, Broadcast: 10.0.31.255
    
```

Physical device indexes

Device configuration and operational flags

Traffic load and alarm status

Logical device indexes

Logical device settings

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 5-12

### Standard Interface Status

Use the **show interfaces** command without the **terse** or **detailed** switches to display standard information about the named interface (or all interfaces when you do not identify a specific interface). The slide provides sample output for an OC-3C SONET interface. The callouts on the slide help illustrate how interfaces partition into physical devices and logical units in JUNOS Software.

Each physical and logical interface is referenced by two index numbers within JUNOS Software. An interface index is assigned to each interface at boot time depending upon the order in which that interface activates. The SNMP *ifIndex* is used to identify and reference that interface when performing SNMP MIB walks. Note that the indexes assigned to the physical interface device (*ifd*) differ from the index used to identify the logical device (*ifl*). Wherever possible, the SNMP *ifIndex* values are persistent across reboots or in the event of hardware additions and deletions that result from PIC or Flexible PIC Concentrator (FPC) insertion and removal.

The output of a **show interfaces** command also includes a section on the device-level configuration and its operational flags.

*Continued on next page.*



## Standard Interface Status (contd.)

The output of a **show interfaces** command displays the device-level configuration and provides additional information about the device operation through various flags. These flags include the following:

- **Down:** Device was administratively disabled;
- **Hear-Own-Xmit:** Device hears its own transmissions;
- **Link-Layer-Down:** The link-layer protocol failed to successfully connect with the remote endpoint;
- **Loopback:** Device is in physical loopback;
- **Loop-Detected:** The Data Link Layer received frames that it sent and suspects a physical loopback;
- **No-Carrier:** Where the media supports carrier recognition, this flag indicates that no carrier is currently visible;
- **No-Multicast:** Device does not support multicast traffic;
- **Present:** Device is physically present and recognized;
- **Promiscuous:** Device is in promiscuous mode and sees frames addressed to all physical addresses on the medium;
- **Quench:** Device is satiated because it overran its output buffer;
- **Recv-All-Multicasts:** No multicast filtering (promiscuous); and
- **Running:** Device is active and enabled.

One or more flags help communicate the status of the interface. These flags include the following:

- **Admin-Test:** Interface is in test mode, which means that some sanity checking, such as loop detection, is disabled;
- **Disabled:** Interface is administratively disabled;
- **Hardware-Down:** Interface is nonfunctional or incorrectly connected;
- **Link-Layer-Down:** Interface keepalives indicate that the link is incomplete;
- **No-Multicast:** Interface does not support multicast traffic;
- **Point-To-Point:** Interface is point to point;
- **Promiscuous:** Interface is in promiscuous mode and sees frames addressed to all physical addresses;
- **Recv-All-Multicasts:** No multicast filtering (promiscuous);
- **SNMP-Traps:** SNMP traps are enabled; and
- **Up:** Interface is enabled and operational.

*Continued on next page.*

### Standard Interface Status (contd.)

Flags also indicate the operational status of the device link layer protocol. These flags include the following:

- **Give-Up:** Link protocol does not continue to attempt to connect after repeated failures;
- **Keepalives:** Link protocol keepalives are enabled;
- **Loose-LCP:** PPP does not use Link Control Protocol (LCP) to indicate whether the link protocol is up;
- **Loose-LMI:** Frame Relay will not use Local Management Interface (LMI) to indicate whether the link protocol is up;
- **Loose-NCP:** PPP does not use Network Control Protocol (NCP) to indicate whether the device is up; and
- **No-Keepalives:** Link protocol keepalives are disabled.

The output also summarizes the device-level traffic load, which displays in both bits and packets per second, as well as any alarms that might be active. The final portion of the command output displays the configuration and status of each logical unit defined on that device. In this example, a single unit is present with support for the `inet` protocol family.



## Displaying Input and Output Errors

- The `show interfaces extensive` command produces similar output for all interfaces types
  - Use `clear interfaces statistics` to reset counters

```
user@host> show interfaces so-0/1/1 extensive
Physical interface: so-0/1/1, Enabled, Physical link is Up
Interface index: 133, SNMP ifIndex: 25, Generation: 16
. . .
Statistics last cleared: Never
```

When counters were last cleared

```
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
Bucket drops: 0, Policed discards: 0, L3 incompletes: 0,
L2 channel errors: 0, L2 mismatch timeouts: 0, HS link CRC errors: 0,
HS link FIFO overflows: 0
```

Input errors

```
Output errors:
Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 0,
HS link FIFO underflows: 0, MTU errors: 0
```

Output errors

```
SONET alarms : None
SONET defects : None
SONET PHY:
Seconds Count State
PLL Lock 0 0 OK
PHY Light 139 2 OK
SONET section:
BIP-B1 0 0
SEF 143 157 OK
```

Media errors

Note: Policed discards count the receipt of unrecognized protocol types (for example, CDP or STP)

### Displaying Input and Output Errors for the Interface

Use the `show interfaces extensive` command to display input errors (extensive output only) on the interface. Use the `clear interfaces statistics interface-name` command to reset the counters for the specified interface; omit an interface name to clear all interface statistics. The following list explains some of the less obvious counters:

- `Errors`: Displays the sum of the incoming frame aborts and FCS errors.
- `Policed discards`: Displays the frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that JUNOS Software does not handle, such as Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or any protocol type that JUNOS Software does not understand. (On an Ethernet network, numerous possibilities exist.)
- `L3 incompletes`: This counter increments when the incoming packet fails Layer 3 (usually IPv4) checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded, and this counter would increment.
- `L2 channel errors`: This counter increments when the software cannot find a valid logical interface (such as e3-1/2/3.0) for an incoming frame.

*Continued on next page.*

### Displaying Input and Output Errors for the Interface (contd.)

- `L2 mismatch timeouts`: Displays the count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable.
- `SRAM errors`: This counter increments when a hardware error occurs in the static RAM on the PIC. The value in this field should always be 0. If it increments, the PIC is malfunctioning.

The **show interface extensive** command also displays the output errors on the interface. The following list explains the less obvious counters:

- `HS link CRC errors`: Displays the count of errors on the high-speed links between the application-specific integration circuits (ASICs) responsible for handling the router interfaces.
- `Carrier transitions`: Displays the number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system powers on and off, or a similar problem occurs. If it does increment quickly (perhaps every 10 seconds), then either the transmission line, the far-end system, or the PIC is broken.
- `Errors`: Displays the sum of the outgoing frame aborts and FCS errors.
- `Drops`: Displays the number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that the ASIC's random early detection (RED) mechanism drops.
- `Aged packets`: Displays the number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly due to malfunctioning hardware.

## Monitoring an Interface

```

user@router                               Seconds: 55                               Time: 19:12:00
                                           Delay: 0/0/66
Interface: ge-0/2/0, Enabled, Link is Down
Encapsulation: Ethernet, Speed: 1000mbps
Traffic statistics:
  Input bytes:          17707053 (0 bps)
  Output bytes:        10369709 (0 bps)
  Input packets:       292046 (0 pps)
  Output packets:     147886 (0 pps)
Error statistics:
  Input errors:        0
  Input drops:         0
  Input framing errors: 0
  Policed discards:   14355
  L3 incompletes:     261
  L2 channel errors:  0
  L2 mismatch timeouts: 156
  Carrier transitions: 2
  Output errors:      0
  Output drops:       0
  Aged packets:       0
Active alarms : LINK
Active defects: LINK
Input MAC/Filter statistics:
  Unicast packets      104547
  Broadcast packets    40494
  Multicast packets    67917
  . . .
  
```

Current Delta

[0]
[312]
[0]
[4]
[0]
[0]
[0]
[0]
[0]
[0]
[0]
[1]
[0]
[0]
[0]

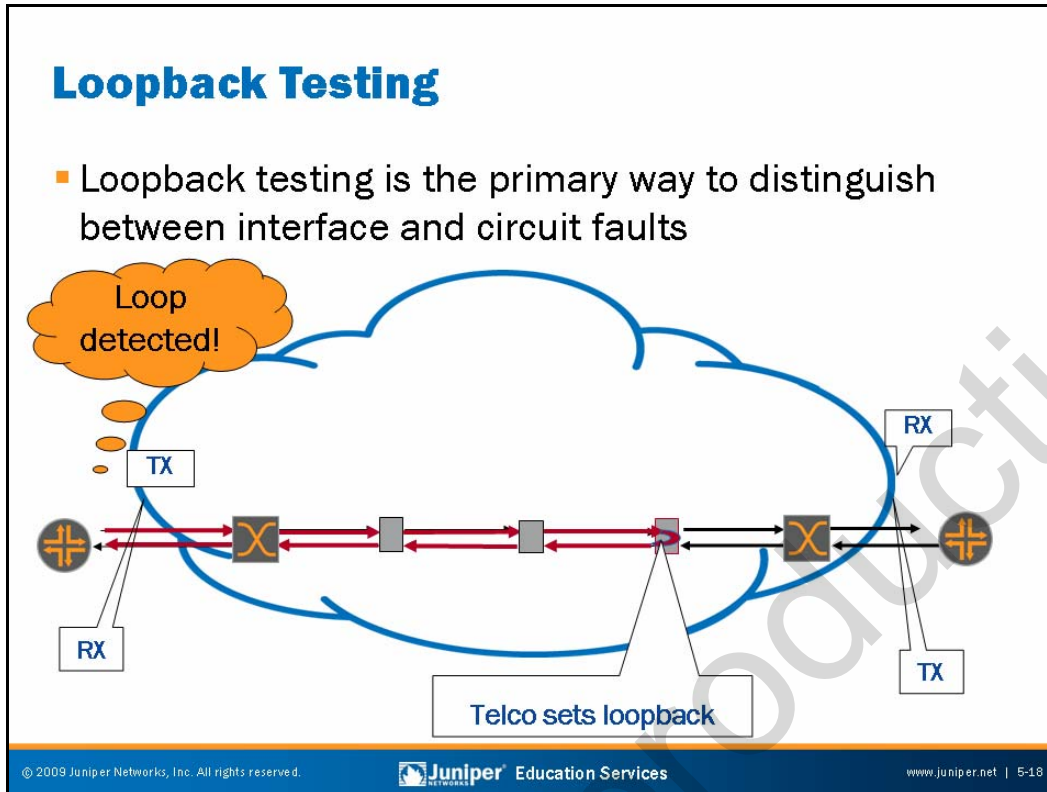
Real-time traffic and error counts

Types of packets

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | 5-17

### Monitoring an Interface

The slide depicts a typical output from the **monitor interface** command. You must set your terminal session to VT100 for the screen to display correctly. This command provides real-time packet and byte counters as well as displaying error and alarm conditions.



### Loopback Testing

The physical path of a leased line usually consists of a number of segments or spans interconnected by devices that repeat and regenerate the signal. When a fault occurs on the circuit that takes the form of either a break or signal corruption due to noise, it is possible to localize the problem by testing the line on a segment-by-segment basis or an end-to-end basis, as needed.

Each circuit is symmetric in that a transmit path from one device connects to the receive path on the remote side, and vice versa. Looping is the process of connecting the transmit path of a router or intermediate device to the receive path. If this device is one of the routers, the loop is either detected if the looped segment is operational, or not detected if a break occurs. The device achieves this detection by detecting its own Data Link Layer keepalive packets (for example, the magic number when the encapsulation is PPP).

If a loop is set back towards a device and the device does not detect it, you can assume that the problem lies somewhere between the device and where the loop was set by the telco or provider. The next step is to set a loop somewhere closer to the device to localize the problem segment.

*Continued on next page.*

### Loopback Testing (contd.)

It is usually possible to loop the device interface locally by connecting the PIC's transmit and receive ports. You should take care to attenuate signal strength when dealing with intermediate-reach and long-reach fiber-optic interfaces.

You can use a similar approach to track down noise on a line by combining the looping process with a test that checks for bit rate errors, commonly known as BERT testing. Many of the interfaces on JUNOS platforms support BERT testing.

Not for Reproduction

## Supported Loopback Types

- Most PICs support internal loopbacks
  - Point-to-point type PICs also support remote loopbacks
  - A local loop does not also provide a remote loop
    - Can perform only one type of loopback at any given time

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | 5-20

### Supported Loopback Types

Most PICs supported on JUNOS platforms support local (internal) loopback tests. Where possible, it is best to perform local loopbacks using an external loopback plug because this setup also tests the PIC's transmit and receive circuitry. Point-to-point style interfaces (nonbroadcast types of technologies like SONET or T1/DS1), also support remote loopbacks. Note that configuring an interface for a remote loopback results in a line loop on the local device; it does not generate a remote loopback request to the remote device. Line loops can be remotely signalled for PICs with integral channel service unit (CSU) functionality (T1 or E1, and T3 or E3), but the generation of the remote loopback code requires telco interaction or test equipment. Again, configuring a remote loopback in JUNOS Software does not signal the remote end to perform a loopback; it creates a local line-loop condition.

For local loopback the PIC's transmit clocking should be set to internal, which is the default setting. A remote loopback (or line loop) allows telco testing on the local loop (also referred to as the *tail*) and also allows testing from the remote device.

## Configuring Loopbacks

- Local and remote loops require configuration on most PICs:

- External local loop and telco line loops do not require configuration

```
[edit interfaces so-0/1/1]
user@Tokyo# show
no-keepalives;
encapsulation frame-relay;
sonet-options {
  loopback local;
}
unit 100 {
  dlc1 100;
  family inet {
    address 10.0.22.1/24;
  }
}
```



- Display interface status to confirm that a configured loopback is in effect:

```
[edit interfaces so-0/1/1]
user@Tokyo# run show interfaces so-0/1/1 | match loop
Link-level type: Frame-Relay, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
Loopback: Local,
```

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 5-21

### Configuring Loopbacks

Interface loopbacks require configuration in JUNOS Software for most PICs and interface types. A small number of channelized DS3 and OC12 interfaces support the ability to initiate far-end alarm and control (FEAC)-based or T1 inband and FDL-based loopbacks using operational mode commands. Note that configuration is never necessary to effect an external local-loopback with a loopback plug, or when relying on the telco to provide a line loopback (which appears as a remote loopback to the attached router). The slide shows an example of a local-loopback configuration and the operational mode status display that confirms that the loopback is in place.

Note that when the telco provides a line loopback, no indication exists that a loopback is in place, unless the configured Layer 2 protocol has built-in loopback detection—for example, PPP. The routers used in this example are running Frame Relay with LMI-based keepalives disabled. As a result, a remote loopback goes undetected at the remote device, which is now talking to itself as indicated by the time-to-live (TTL) expiration messages shown (we cover the use of ping to test loopbacks on a forthcoming page):

*Continued on next page.*

### Configuring Loopbacks (contd.)


```
[edit interfaces so-0/1/1]
user@London# run ping 10.0.22.1 count 1
PING 10.0.22.1 (10.0.22.1): 56 data bytes
36 bytes from 10.0.22.2: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
 4 5 00 0054 601b 0 0000 01 01 198c 10.0.22.2 10.0.22.1
--- 10.0.22.1 ping statistics ---
1 packets transmitted, 0 packets received, 100% packet loss
```

```
[edit interfaces so-0/1/1]
user@London# run show interfaces so-0/1/1 | match loop
Link-level type: Frame-Relay, MTU: 4474, Clocking: Internal, SONET mode,
Speed: OC3, Loopback: None, FCS: 16,
```



## Layer 2 Protocols and Loopbacks

- Many Layer 2 protocols detect loopbacks resulting in link down status
  - Prevents diagnostic ping testing
  - Solution: Disable keepalives, change L2 encapsulation, or both
    - PPP always detects a loopback via IPCP, even when keepalives are disabled!



```
[edit interfaces so-0/1/1]
user@London# run show interfaces so-0/1/1
Physical interface: so-0/1/1, Enabled, Physical link is Up
. . .
Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3, Loopback:
None, FCS: 16,
Payload scrambler: Enabled
Device flags : Present Running Loop-Detected
. . .
Logical interface so-0/1/1.0 (Index 70) (SNMP ifIndex 26)
. . .
Protocol inet, MTU: 4470
Flags: Protocol-Down
```

```
[edit interfaces so-0/1/1]
user@London# show
no-keepalives;
encapsulation ppp;
unit 0 {
    family inet {
        address 10.0.22.2/24;
    }
}
```

Loop detected, link down—  
despite no-keepalives!

### Layer 2 Protocols and Loopbacks

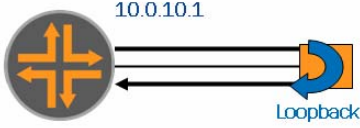
Many Layer 2 protocols make use of a keepalive mechanism that, among other things, can detect the presence of a loopback. Whether local or remote, the detection of a loop condition results in a link down declaration for that interface. When the interface is marked as down at the Data Link Layer, the related interface route is removed from the routing table, which prevents ping testing for the duration of the loopback. (We describe ping testing over a loopback on subsequent pages.)

In most cases you can work around this issue by configuring the interface with a `no-keepalives` statement, but, as shown on the slide, this workaround only works for the `frame-relay`, `atm`, and `cisco-hdlc` encapsulation types. Even with keepalives (LCP) disabled, PPP still detects the presence of a loopback when the Network Control Protocol (NCP) attempts to negotiate Layer 3 parameters. The only way around this conditions is to change the interface's encapsulation type for the duration of the loopback test.

Note that Ethernet-related technologies have no concept of a link-layer keepalive protocol, and they do not support the concept of a remote loopback. This information is only applicable to point-to-point interface types.

## Loopback Tests and Pings

Device from Other Vendor




10.0.10.1

Loopback

- Pings to local WAN address produce line traffic on equipment from some vendors

Device running JUNOS Software



10.0.10.1

Loopback

- Pings to local WAN address do not leave the chassis with Juniper Networks equipment

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | 5-24

### Equipment from Other Vendors

On equipment from some other vendors, you can test the operation of a WAN link by issuing pings to the router's local IP address. The top of the slide shows this mode of operation.

### Juniper Networks Equipment

JUNOS platforms do not exhibit this behavior. A ping sent to the device's local IP address does not exit the interface, and as such, cannot be used to ascertain the operational status of the line.

The next slide covers testing the line on JUNOS platforms.

## Testing a Looped Line with Ping

```
lab@router> ping 10.0.10.2
PING 10.0.10.2 (10.0.10.2): 56 data bytes
36 bytes from 10.0.10.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src
Det
 4 5 00 0054 211b 0 0000 01 01 708c 10.0.10.1
10.0.10.2
```

JUNOS Device



- **Loop line and ensure that interface stays up**
  - Needs AAL, Frame Relay, or HDLC encapsulation with keepalives disabled
- **Ping remote IP address**
  - If line is good, ping returns to router and is routed back out the interface with TTL decrease
- **When TTL expires, error message returns—this result is expected**
  - TTL expiration indicates that no packets were lost during the test—TTL setting determines number of loops (default TTL is 255)
  - Helpful to open another session and monitor the interface under test to display CRC errors

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 5-25

### Looping Line and Keeping Interface Up

In the example on the slide, we looped the line externally. This loop might be a hard loop, a telco loop, or a remote interface loop for the purposes of this example.

Because some Data Link Layer protocols detect the looped condition, and disable the interface as a result, you must use either ATM Adaptation Layer 5 (AAL5), Frame Relay, or Cisco HDLC encapsulation with keepalives turned off. PPP encapsulation generally does not work, because the looped condition prevents the NCP from completing its initialization, thereby preventing a declaration of up for the interface.

### Pinging Remote IP Address

With the loop in place and the interface up, we now issue a ping to the address associated with the remote end of the line. The address is 10.0.10.2 in this example.

### Error Returns When TTL Expires

If the line has a usable transmit and receive path, the packet returns to the local device as a result of the loop condition. Upon receiving this packet, the device once again sends the packet out the WAN interface a second time. The packet's TTL field decreases during this process. This operation continues until the packet's TTL reaches zero, or until a line error causes packet corruption and the resulting silent packet discard.

*Continued on next page.*

### Error Returns When TTL Expires (contd.)

Therefore, a good line should return Internet Control Message Protocol (ICMP) TTL expired messages for every packet sent, while a marginal line might return no TTL expired errors, or it might return TTL expired messages for a small subset of the packets sent. Packet size, TTL setting, and use of the **rapid** switch can affect your results as well.

Because the default TTL for locally generated pings is 255 on a JUNOS device, each TTL expiration message indicates 255 successful transmissions and receptions of the initial ping request, all at wire speed.

Setting the TTL to a lower value is useful when trying to determine marginality of a line. That is, a TTL of 1 requires only a single transmission and reception of the packet, which is similar to the type of test performed by other vendors when pinging a local WAN interface.

## BERT Testing

- BERT tests require a loop
  - Pattern received through a loop is verified against transmitted pattern
  - You can set a loop at various points in the transmission path

- BERT parameters:

```
[edit interfaces interface-name t3-options]
bert-algorithm algorithm;
bert-error-rate rate;
bert-period seconds;
```

- Starting and stopping the test:

- Interface must be administratively disabled before executing a BERT test

```
user@host> test interface t3-1/0/1 t3-bert-start
user@host> test interface t3-1/0/1 t3-bert-stop
```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-27

### BERT Tests Require a Loop

The pattern received through a loop is verified against the pattern sent. End-to-end testing is difficult to coordinate. By changing the position of the loop downstream from the device performing the test, you can locate the problem area easily. Common points for looping the line are the telco demarcation point (also named *demark*), the remote end, and the midpoint (with help from the carrier).

You can configure any of the following interfaces to execute a BERT test when the interface receives a request to run this test: E1, E3, T1, T3, the channelized DS-3, OC-3, OC-12, STM-1, the channelized DS-3 IQ, E1, and OC-12 IQ.

### BERT Parameters

You must configure the various parameters that influence a BERT test under the interface subject to testing. These options include the test duration (10 seconds is the default), the test pattern, and the error rate to include in the bit stream by including the `bert-period`, `bert-algorithm`, and `bert-error-rate` statements, respectively.

*Continued on next page.*

## Starting and Stopping the Test

Start and stop the BERT test with the **test interface *interface-name* bert-start** and **bert-stop** commands. Note that you cannot run a BERT test on an interface that is administratively enabled. To start a BERT test you must first disable the interface with a **set interfaces *interface-name* disable** statement. After the test completes, you can use a **rollback 1, commit** command sequence to re-enable the interface, or you can remove the disable statement with a **delete interface *interface-name* disable** statement.

Not for Reproduction

## Checking BERT Results

```

user@router> show interfaces t3-0/2/0 extensive | find BERT
BERT time period: 240 seconds, Elapsed: 240 seconds (completed)
Algorithm: All ones, Repetitive (22), Error rate: 10e-0
Bit count      :          0, Overflows: 0
Error bit count:          0, Overflows: 0
LOS status: OK, LOS count: 1, LOS seconds: 239
PFE configuration:
  Destination slot: 0, Stream number: 8, PLP byte: 1 (0x00)
  CoS transmit queue bandwidth:
    Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
  CoS weighted round-robin:
    Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
Logical interface t3-0/2/0.0 (Index 107) (SNMP ifIndex 29)
  Flags: Device-down Point-To-Point SNMP-Traps, Encapsulation: PPP
  . . .

```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-29

### Checking BERT Results

You can check the results of your BERT test using the **show interfaces extensive** command. The slide shows the formatting and fields associated with the results of a BERT test. Most of the fields are self-explanatory, but a few fields could use some additional explanation.

The `Error bit count` field displays the number of erroneous echo replies received from the remote end. The `LOS` field indicates pattern synchronization status. A working BERT test requires that the receiver be in sync with the transmitter. In this example, pattern synchronization was lost once during the test; furthermore, the loss of synchronization lasted for 239 seconds according to the `LOS seconds` field. The display also shows that no bits were received, and as a result, that no errors were detected. The lack of received bits is likely the result of the lack of test pattern synchronization.

Note that for a BERT test to be meaningful you must be able to inject and detect errors. Only by purposely injecting an error—and then witnessing that the injected error is detected—can you be sure that the test results are valid.

## Agenda: Interface Troubleshooting

- Interface Configuration Overview
- General Interface Troubleshooting
- Media-Specific Interface Troubleshooting
- Case Study

### Media-Specific Interface Troubleshooting

The slide highlights the topic we discuss next.



## LAN Topologies

- Port types:
  - fe-
  - ge-
  - xe-
  - fxp0
- Link mode (full duplex or half duplex)
- Tools:
  - ping
  - loopback (local)
  - show interfaces extensive
  - show interfaces media
  - show arp
  - monitor traffic
  - monitor interface
  - clear statistics
  - request ...
  - restart ...

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-31

### Media Types and Interface Naming

JUNOS platforms support several *flavors* of Ethernet. The following are the relevant media types:

- *Fast Ethernet*: fe-F/P/P;
- *Gigabit Ethernet*: ge-F/P/P;
- *10-Gigabit Ethernet*: xe-F/P/P; and
- *Management Ethernet*: fxp0 (10, 100, or 1000 Mbps).

### Link Mode

When troubleshooting LAN topologies, consider the link mode:

- Full duplex;
- Half duplex; or
- Link bonding (802.3ad)

Fast Ethernet interfaces can support half duplex or full duplex, but Gigabit Ethernet interfaces only function in full-duplex mode.

*Continued on next page.*

## JUNOS Software Tools

JUNOS Software provides the tools shown on the slide. The following pages examine these tools.

Not for Reproduction

## Ethernet Details (1 of 2)

### Local loopback support

- Monitor traffic while looped and look for the receipt of all transmitted ARP messages (broadcast) and LED status indications

```
[edit interfaces ge-0/0/0]
user@host# show
gigether-options {
  loopback;
}
```

- Better yet, use loopback plug and static ARP

- Static ARP must match the looped interface's own MAC address
- Must be in full-duplex mode

```
user@host> show interfaces ge-0/0/0 | match hardware
Current address: 00:17:cb:4e:ab:00, Hardware address: 00:17:cb:4e:ab:00
user@host> show configuration interfaces ge-0/0/0
unit 0 {
  family inet {
    address 200.0.0.1/24 {
      arp 200.0.0.20 mac 00:17:cb:4e:ab:00;
    }
  }
}
```

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 5-33

### Configuring Loopback Mode

To place an Ethernet interface in loopback mode, issue a **set gigether-options loopback** at the `[edit interfaces ge-interface-name]` hierarchy. You use a similar command for Fast Ethernet interfaces. When the interface loops, you can monitor traffic and expect to see all traffic that travels out—that is, an Address Resolution Protocol (ARP) request—coming right back in:

```
user@host> monitor traffic interface fe-0/0/0
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Listening on fe-0/0/0, capture size 96 bytes
21:14:04.424904 Out arp who-has 200.0.0.30 tell 200.0.0.1
21:14:04.425328 In arp who-has 200.0.0.30 tell 200.0.0.1
```

When operating in the default full-duplex mode, you can also attach an external loopback plug to effect an external local loopback. To see TTL expired messages (as expected for point-to-point interfaces), you must add a static ARP entry that matches the media access control (MAC) address for the looped interface for the target IP address. This addition is necessary so that the interface accepts returning traffic because a nonpromiscuous Ethernet interface only accepts broadcast and unicast traffic sent to its MAC address. When all is working you should see TTL errors:

```
user@host> ping 200.0.0.20 count 1
PING 200.0.0.20 (200.0.0.20): 56 data bytes
36 bytes from 200.0.0.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
4 5 00 0054 4e67 0 0000 01 01 db2c 200.0.0.1 200.0.0.20
```

## Ethernet Details (2 of 2)

- Ping a locally connected device
- Use the `show arp no-resolve` command:

```
user@host> show arp no-resolve
```

MAC Address	Address	Interface	Flags
00:d0:b7:3f:af:0f	10.0.1.200	ge-0/0/1	none

- Cable lengths and Physical Layer standards:
  - Cat 5 UTP copper: 100 meters
  - Multi-mode fiber: Check the port specifications
  - Single-mode fiber: Check the port specifications (IR and LR)
- Tips:
  - Check encapsulation types (802.3 LLC, 802.3 SNAP, and DIXv2)
  - Use the `show interfaces extensive` command
  - Use the `monitor interface` command

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 5-34

### Pinging a Locally Connected Device

A reply received from the device typically provides verification that the link and interface are operating correctly.

### Displaying ARP Table

The `show arp no-resolve` command displays the entries in the ARP table. Using the `no-resolve` option prevents the device from attempting to determine the host name that corresponds to the IP address.

### Verifying Cable Length

Ensure the cables used on the network do not exceed recommended lengths and meet all relevant specifications. If the cabling specifications are not met, the input errors on the interface will increase.

*Continued on next page.*

## Generic Tips

The following are a few generic tips:

- Ensure that encapsulation types are equivalent to other devices on link.
- Use the **show interfaces extensive** command to check the status the of interface.
- Use the **monitor interfaces** command to receive real-time statistics.
- Use the **monitor interface *interface-name*** command to display real-time statistics about a physical interface. The output updates every second. The output of this command also shows the amount that each field has changed since you started the command or since you cleared the counters by using the c key. This command also checks for and displays common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors. If the framing errors are increasing, this increase indicates that frames are being corrupted. If the input errors are increasing, check the cabling to the device and have the carrier verify the integrity of the line.

## Troubleshooting T3 and E3

- Check the port
  - Connect cables
  - Set clock internal
  - Test physical loopback between transmit and receive ports
- Check the subrate compatibility mode:

```
[edit interfaces t3-0/1/1 t3-options]
set compatibility-mode (adtran|digital-link|kentrox|larscom|verilink)
subrate value;
```

- Check miscellaneous settings at both ends:
  - Clocking (internal or external)
  - Frame checksum (FCS-32 and FCS-16)
  - HDLC payload scrambling (disabled by default)
  - T3 line buildout (copper only)
  - T3 C-bit parity mode (default is on)

### Checking the Port

To troubleshoot T3 and E3, set up a physical loopback between the transmit and receive ports. If the T3 interface is functioning properly, you should see a `Loop-Detected` flag in the device flags section of the `show interface` output. Also, the `monitor interface` command should show that the input packet count matches the output packet count. For Cisco HDLC encapsulation, the input keepalive packet count should also match the output keepalive packet count.

If you do not see the `Loop-Detected` flag, the PIC port might be bad. To isolate the problem, move the T3 link to another T3 interface on the router and verify whether the new port works. To move the configuration of the existing T3 interface to the new interface, use the `rename` command under the `[edit interfaces]` hierarchy in configuration mode:

```
[edit interfaces]
user@host# rename t3-1/0/0 to t3-1/0/1
```

*Continued on next page.*

## Checking the Compatibility Mode

The subrate of an E3 or T3 interface must match that of the remote CSU exactly. The purpose of a subrate is to fit more bandwidth into an E3 or T3 circuit. To specify the subrate, include the subrate option in the `compatibility-mode` statement. Follow these instructions for different types of CSUs:

- For Adtran CSUs, specify the subrate as a number from 1 through 588 that exactly matches the value configured on the CSU. A subrate value of 588 corresponds to 44.2 Mbps, or 100 percent of the HDLC-encapsulated payload. A subrate value of 1 corresponds to  $44.2 / 588$ , which is 75.17 Kbps, or 0.17 percent of the HDLC-encapsulated payload.
- For Digital Link CSUs, specify the subrate value as the data rate you configured on the CSU in the format `xKb` or `x.xMb`. For a list of specific rate values, use the command completion feature in the command-line interface (CLI). The range is 358 Kbps through 33.7 Mbps for E3 interfaces and 301 Kbps through 44.2 Mbps for T3 interfaces.
- For Kentrox CSUs, specify the subrate as a number from 1 through 69 that exactly matches the value configured on the CSU. A subrate value of 69 corresponds to 34.995097 Mbps, or 79.17 percent of the HDLC-encapsulated payload (44.2 Mbps). A subrate value of 1 corresponds to 999.958 Kbps, which is 2.26 percent of the HDLC-encapsulated payload.
- For T3 interfaces configured with Larscom CSUs, specify the subrate value as a number from 1 through 14 that matches the value configured on the CSU exactly. E3 interfaces do not support the subrate option with Larscom CSUs.
- For Verilink CSUs, specify the subrate as a number from 1 through 28 that exactly matches the value configured on the CSU. To calculate the maximum allowable peak rate, multiply the configured subrate by 1.578 Mbps. For example, a subrate value of 28 corresponds to  $28 \times 1.578$  Mbps, which is 44.2 Mbps—100 percent of the HDLC-encapsulated payload.

## Ensuring Compatible Settings

The following settings ensure compatibility:

- *E3 and T3 frame checksums:* By default, E3 and T3 interfaces use a 16-bit frame checksum. You can configure a 32-bit checksum that provides more reliable packet verification. However, some older equipment might not support 32-bit checksums.
- *HDLC payload scrambling:* E3 or T3 HDLC payload scrambling, which is disabled by default, provides better link stability. Both sides of a connection must either use or not use scrambling.
- *Line buildout:* To have the interface drive a line that is longer than 255 feet, include the `long-buildout` statement at the `[edit interfaces interface-name t3-options]` hierarchy level.

*Continued on next page.*

### Ensuring Compatible Settings (contd.)

- *T3 C-bit parity mode*: For T3 interfaces, the C-bit parity mode controls the type of framing present on the transmitted T3 signal. When C-bit parity mode is enabled, the C-bit positions are used for the far-end block error (FEBE), FEAC, the terminal data link, path parity, and mode indicator bits, as defined in American National Standards Institute (ANSI) T1.107a-1989. When you disable C-bit parity mode, the basic T3 framing mode is used. By default, C-bit parity mode is enabled. To disable it, include the **no-cbit-parity** statement at the [edit interfaces *interface-name* t3-options] hierarchy level.

Not for Reproduction



## Displaying T3 Active Alarms

```
user@router> show interfaces t3-1/0/0
```

```
Physical interface: t3-1/0/0, Enabled, Physical link is Down
```

Check alarms

```
Interface index: 9, SNMP ifIndex: 10
```

```
Link-level type: Cisco-HDLC, MTU: 4474, Clocking: Internal
```

Verify settings at both ends

```
Speed: T3, Loopback: None, CRC: 16, Mode: C/Bit parity
```

```
Device flags : Present Running Down
```

```
Interface flags: Hardware-Down Link-Layer-Down Point-To-Point SNMP-Traps
```

```
Link flags : Keepalives
```

```
Keepalive Input: 116 (00:02:32 ago), Output: 185 (00:00:02 ago)
```

```
Input rate : 0 bps (0 pps), Output rate: 0 bps (0 pps)
```

```
Active alarms : LOF, LOS
```

```
Active defects : LOF, LOS
```

```
Logical interface t3-1/0/0.0 (Index 12) (SNMP ifIndex 32)
```

```
Flags: Device-down Point-To-Point SNMP-Traps, Encapsulation: Cisco-HDLC
```

```
Protocol inet, MTU: 4470
```

```
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
```

```
Destination: 2.2.2.0/30, Local: 2.2.2.2
```

Alarms are active!

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 5-39

### Displaying Active Alarms

If the first line shows `Physical link is Up`, it means that the physical link is healthy and can pass packets. If the first line shows `Physical link is Down`, it means that the physical link is unhealthy and cannot pass packets. To display more extensive information about the T3 interface if the physical link is down, use the **`show interface t3-x/y/z extensive`** command. Look at the active alarms and active defects for the T3 interface, and troubleshoot the T3 media accordingly.

## Important DS-3 and E3 Alarms

- **Alarms and their meanings:**
  - *AIS*: Indicates a problem with the line upstream of the network equipment connected to the DS-3 interface
  - *LoF*: Indicates loss of DS-3 framing bits
  - *LoS*: Indicates that a signal could not be detected at the DS-3 interface
  - *IDLE*: Indicates that the line is not provisioned for service
  - *Yellow*: Indicates that the remote router reports that it is not receiving a signal on the DS-3 interface

### Important DS-3 Alarms

The following list provides descriptions of the important DS-3 alarms:

- *AIS*: An incoming alarm indication signal (AIS) indicates a problem with the line upstream of the network equipment connected to the DS-3 interface. Have the carrier check the equipment connected to the DS-3 interface and trace the problem.
- *LoF*: A loss of frame (LoF) alarm indicates loss of DS-3 framing bits. Verify the framing format configuration (C-bit parity mode) and check the DS-3 line. Bursts of Line Code Violations (LCVs) could indicate a timing problem.
- *LoS*: A loss of signal (LoS) alarm means that a signal could not be detected at the DS-3 interface. Check the cables connected to the DS-3 interface, and have the carrier verify the integrity of the line.
- *IDLE*: An idle alarm indicates that the line is not provisioned for service. Have the carrier ensure line-provisioning for service.
- *Yellow*: An incoming yellow alarm indicates that the router's DS-3 interface has a problem with the signal it is receiving from the remote equipment's DS-3 interface. Check the cable between the DS-3 interface and the directly connected network equipment.

*Continued on next page.*

## Important DS-3 Alarms (contd.)

The following are some DS-3 error events:

- *BPV*: A bipolar violation (BPV) error event, for bipolar 3 zero substitution (B3ZS—also referred to as high-density bipolar 3, or HDB3) coded signals, is the occurrence of a pulse of the same polarity as the previous pulse without being part of the zero substitution code. For B3ZS-coded signals, a bipolar violation error event might also include other error patterns, such as three (or four) or more consecutive zeros and incorrect polarity (see ANSI T1.231 Section 7.1.1.1.1).
- *EXZ*: An excessive zeros (EXZ) alarm indicates the occurrence of any zero string length equal to or greater than three for B3ZS, or greater than four for HDB3 (see ANSI T1.231 Section 7.1.1.1.2).
- *LCV*: The Line Code Violation (LCV) parameter (also known as CV-L) is a count of both BPVs and EXZs occurring over the accumulation period. An EXZ increments the LCV by one, regardless of the length of the zero string (see ANSI T1.231 Section 7.4.1.1).
- *PCV*: For all DS-3 applications, the P-bit Code Violation (PCV) error event is the same as a P-bit parity error event. In other words, the P-bit code on the DS-3 M-Frame does not match that code calculated locally (see ANSI T1.231 Section 7.1.1.2.1).
- *CCV*: For C-bit parity and SYNTRAN DS-3 applications, the C-bit Code Violation (CCV) is the count of coding violations reported through the C-bits. For C-bit parity, it is a count of C-bit parity errors occurring in the accumulation interval. For SYNTRAN, it is a count of CRC-9 errors occurring in the accumulation interval (see ANSI T1.231 Section 7.1.1.2.2).

For detailed definitions of the DS-3 performance parameters (LES, PES, PSES, CES, CSES, SEFS, UAS), see RFC 2496.

## T3 FEAC Response

- T3 interfaces can respond to remote loopback requests

- Must configure support for automatic response:

```
[edit interfaces interface-name t3-options]
user@host# set feac-loop-respond;
```

- Initiate and clear remote loopbacks with the operational mode **test** command

- Use **test interface *interface-name* feac-loop-initiate** to initiate a remote loopback
- Use **test interface *interface-name* feac-loop-terminate** to clear a remote loopback

### Responding to Loop Requests

For T3 interfaces, the T3 FEAC signal sends alarm or status information from the far-end terminal back to the near-end terminal and to initiate T3 loopbacks at the far-end terminal from the near-end terminal.

To allow the remote CSU to place the local router's integral CSU into loopback, you must configure the device to respond to the CSU's FEAC request by including the **feac-loop-respond** statement at the [edit interfaces *interface-name* t3-options] hierarchy level:

```
[edit interfaces interface-name t3-options]
user@host# set feac-loop-respond
```

By default, the device does not respond to FEAC requests.

### Remote Loopbacks

Issue the CLI operational mode command **test interface feac-loop-initiate *interface-name*** to activate a remote loopback. Use the **test interface feac-loop-terminate *interface-name*** command to clear a remote loopback.

## Troubleshooting T1 and E1

- Approach for troubleshooting T1 or E1 is similar to T3 and E3 troubleshooting
- Use ping testing for circuit error detection and verification
  - Functionality similar to BERT test with T3 or E3
- Loopback testing supported:
  - Local and remote
  - Check for loopback from telco with **show interface** command
- You must match configuration details at either end
- T1 and E1 configuration similarities:
  - Frame check sequence
  - Idle cycle flag
  - Timeslots (with channelized T1 or E1)
- Other than these settings, E1 and T1 configuration differ somewhat so we deal with them separately

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-43

### T1 and E1 Approach Is Similar to T3 and E3

The tests listed on this page and the following pages are very similar to those discussed in the T3 and E3 section.

### Ping Testing

You can use ping tests for data integrity testing. Alternatively, you also can perform BERT testing on T1 or E1 interfaces.

### Local and Remote Loopback Testing Supported

T1 or E1 loopback testing can be local or remote. Check for loopback using the **show interface** command.

### Matching Configuration at Both Ends

If the device on one side of a link differs from the device on the opposite side, the link has difficulty coming up. You should match all settings between endpoints.

*Continued on next page.*

## T1 and E1 Configuration Similarities

By default, E1 and T1 interfaces use a 16-bit frame checksum. You can configure a 32-bit checksum that provides more reliable packet verification. However, some older equipment might not support 32-bit checksums.

To configure a 32-bit checksum, include the **fcs 32** statement at the [edit interfaces *interface-name* e1-options] or [edit interfaces *interface-name* t1-options] hierarchy level:

```
[edit interfaces interface-name t1-options]
user@host# set fcs 32
```

Also by default, E1 and T1 interfaces transmit the value 0x7E in the idle cycles. To have the interface transmit the value 0xFF (all ones) instead, include the **idle-cycle-flag** statement at the [edit interfaces *interface-name* e1-options] or [edit interfaces *interface-name* t1-options] hierarchy level, specifying the **ones** option:

```
[edit interfaces interface-name t1-options]
user@host# set idle-cycle-flag ones
```

Channelized T1 and E1 applications also require the appropriate setting of time slots (or channels) that carry user data. A T1 interface can support up to 24 channels while an E1 interface supports up to 30 user channels. (Two of the 32 channels used in an E1 interface are for framing and alarm reporting.)

## T1 and E1 Specifics

While T3 and T1 troubleshooting have many similarities, a few pronounced differences exist. We describe the different methods for troubleshooting T1 and E1 interfaces on the following pages.

## T1 Configuration Specifics

- T1 buildout:
  - Distance of PIC port to CSU (default 0–133 feet)
- T1 byte encoding:
  - Default byte encoding of 8 bits per byte (nx64)
- T1 data inversion:
  - Default is data inversion disabled
- T1 framing:
  - Default T1 interfaces use extended super frame framing format
    - Alternative is super frame
- T1 line encoding:
  - Default T1 interfaces use B8ZS line encoding
    - Alternative is AMI (older)

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-45

### T1 Buildout

A T1 interface has five possible setting ranges for the T1 line buildout: 0–132, 133–265, 266–398, 399–531, or 532–655 feet. By default, the T1 interface uses the shortest setting (0–133). To have the interface drive a line at one of the longer distance ranges, include the **buildout** statement with the appropriate value at the [edit interfaces *interface-name* t1-options] hierarchy level:

```
[edit interfaces interface-name t1-options]
user@host# set buildout 532-655
```

### T1 Byte Encoding

By default, T1 interfaces use a byte encoding of 8 bits per byte (nx64). You can configure an alternative byte encoding of 7 bits per byte (nx56). To have the interface use 7 bits per byte encoding, include the **byte-encoding** statement at the [edit interfaces *interface-name* t1-options] hierarchy level, specifying the **nx56** option:

```
[edit interfaces interface-name t1-options]
user@host# set byte-encoding nx56
```

*Continued on next page.*

## T1 Data Inversion

By default, data inversion is disabled. To enable data inversion at the HDLC level, include the **invert-data** statement at the [edit interfaces interface-name t1-options] hierarchy level:

```
[edit interfaces interface-name t1-options]  
user@host# set invert-data
```

## T1 Framing

By default, T1 interfaces use extended superframe (ESF) framing format. You can configure superframe (SF) as an alternative. To have the interface use the SF framing format, include the **framing** statement at the [edit interfaces interface-name t1-options] hierarchy level, specifying the **sf** option:

```
[edit interfaces interface-name t1-options]  
user@host# set framing sf
```

## T1 Line Encoding

By default, T1 interfaces use B8ZS line encoding. You can configure alternate mark inversion (AMI) line encoding if necessary. You should use AMI coding in conjunction with the *nx56* byte encoding to prevent problems with ones-density:

```
[edit interfaces interface-name t1-options]  
user@host# set line-encoding ami
```



## E1 Configuration Specifics

- E1 framing:
  - By default, E1 interfaces use the G704 framing mode
    - Alternative is unframed

```
[edit interfaces e1-0/0/1 e1-options]
user@host# set framing (unframed | g704)
```

- So what is the difference?
  - Framed = 1984k
  - Unframed = 2048k
- Note: JUNOS Software timeslots begin with 1
  - TS1 = TS0 on platforms from some other vendors

### E1 Framing

E1 is a standard WAN digital communication format designed to operate over copper facilities at a rate of 2.048 Mbps. Widely used outside North America, it is a basic time-division multiplexing scheme used to carry digital circuits. The following standards apply to E1 interfaces:

- ITU-T Recommendation G.703, *Physical/electrical characteristics of hierarchical digital interfaces*, describes data rates and multiplexing schemes for the E series;
- ITU-T Recommendation G.751, *General Aspects of Digital Transmission Systems: Terminal Equipment*, describes framing methods; and
- ITU-T Recommendation G.775, *Loss of Signal (LoS) and Alarm Indication Signal (AIS) Defect Detection and Clearance Criteria*, describes alarm reporting methods.

To configure E1-specific physical interface properties, include the **e1-options** statement at the [edit interfaces *interface-name*] hierarchy level. By default, E1 interfaces use the G.704 framing mode; the alternative is unframed.

*Continued on next page.*

### What Is the Difference?

By framing an E1 signal, two of the 32 channels remain in reserve for framing and alarm reporting. These reserved channels result in a 128 Kbps reduction in payload capacity. When framed, the E1 operates at 1984 Kbps. When unframed, the E1 uses 2048 Kbps.

### Time Slots

JUNOS Software refers to the first time slot as 1. Some vendors refer to this time slot as time slot 0.

Not for Reproduction

## E1 Media and Alarms

```

user@host> show interfaces media e1-1/1/0
Physical interface: e1-1/1/0, Enabled, Physical link is Up
Interface index: 30, SNMP ifIndex: 130
Link-level type: PPP, MTU: 4474, Clocking: Internal
Speed: E1, Loopback: None, CRC: 16, Framing: G704
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps
Link flags : Keepalives
Input rate : 0 bps (0 pps), Output rate: 0 bps (0 pps)
Active alarms : None
Active defects : None
E1 errors:
    BPV: 2, EXZ: 1, LCV: 2, PCV: 7, CS: 0, FEBE: 561
  
```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-49

### E1 Media and Alarms

Active alarms and active defects can render an interface unable to pass packets. When a defect persists for a certain amount of time, it receives promotion to alarm status. Based on the router configuration, an alarm can trigger a red or yellow alarm on the Craft Interface. E1 alarms include the following:

- *LOS*: Loss of signal;
- *LOF*: Loss of frame;
- *AIS*: Alarm indication signal;
- *YLW*: Yellow alarm, indicating errors at the remote site receiver;

In addition to the alarms mentioned in the list, various errors exist that an E1 interface can report. E1 errors include the following:

- *BPV*: A BPV error event for an AMI-coded signal is the occurrence of a pulse of the same polarity as the previous pulse. (See ANSI T1.231 Section 6.1.1.1.1.) A BPV error event for a B8ZS-coded or HDB3-coded signal is the occurrence of a pulse of the same polarity as the previous pulse without being a part of the zero substitution code.
- *EXZ*: An EXZ error event for an AMI-coded signal is the occurrence of more than 15 contiguous zeros (see ANSI T1.231 Section 6.1.1.1.2). For a B8ZS-coded signal, the defect occurs when more than seven contiguous zeros are detected.

*Continued on next page.*

## E1 Media and Alarms (contd.)

- **LCV:** An LCV error event is the occurrence of either a BPV or an EXZ error event (also known as CV-L; see ANSI T1.231 Section 6.5.1).
- **PCV:** A PCV error event is a frame synchronization bit error in the D4 and E1-noCRC formats, or a CRC or frame synchronization bit error in the ESF and E1-CRC formats (also known as CV-P; see ANSI T1.231 Section 6.5.2.1).
- **CS:** A controlled slip error event is the replication or deletion of the payload bits of a DS-1 frame (see ANSI T1.231 Section 6.1.1.2.3.). A controlled slip might occur when a difference exists between the timing of a synchronous receiving terminal and the received signal. A controlled slip does not cause an out-of-frame defect.
- **OOF:** An out-of-frame defect is the occurrence of a particular density of framing error events (see ANSI T1.231 Section 6.1.2.2.1). For DS-1 links, JUNOS Software declares an out-of-frame defect when the receiver detects two or more framing errors within a 3 msec period for ESF signals and 0.75 msec for D4 signals, or two or more errors out of five or fewer consecutive framing bits. For E1 links, JUNOS Software declares an out-of-frame defect when three consecutive frame alignment signals arrive with an error (see ITU-T G.706 Section 4.1 [26]).

E1 framing alarms include the following:

- **TS16 Alarm Indication Signal Failure:** For E1 links, JUNOS Software declares this signal failure when time slot 16 is received as all 1s for all frames of two consecutive multiframes (see ITU-T G.732 Section 4.2.6). JUNOS Software never declares this condition for DS-1.
- **Loss of Multiframe Failure:** JUNOS Software declares this failure when two consecutive multiframe alignment signals (bits 4 through 7 of TS16 of frame 0) are received with an error. JUNOS Software clears this failure when the first correct multiframe alignment signal is received. The loss-of-multiframe failure can be declared only for E1 links operating with ITU-T G.732 [27] framing (sometimes referred to as *Channel Associated Signaling* mode).
- **Far End Loss of Multiframe Failure:** JUNOS Software declares this failure when bit 2 of TS16 of frame 0 is received set to 1 on two consecutive occasions. JUNOS Software clears this failure when bit 2 of TS16 of frame 0 is received set to 0. JUNOS Software can only declare the far-end loss of multiframe failure for E1 links operating in *Channel-Associated Signaling* mode (see ITU-T G.732).

## T1 and E1 Troubleshooting Procedures

- **Fault isolation**
  - If settings are the same on both ends and the problem persists, you must work with the telco
  - Set the CSU/DSU loopback to the T1 or E1 port
- **See RFC 2495 for more details on errors and performance-related errors**

### T1 and E1 Fault Isolation

When you confirm the settings on both ends but problems persist, you might want to involve the telco for line and loop testing. Before suspecting the transmission line, you should first perform local loopback testing at each end. You also should attempt remote loopback requests to the far-end internal CSU.

### RFC 2495

For additional information on E1 or T1 interface alarms and error conditions, you should consult RFC 2495, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*.

## Ping Test Patterns: Circuit Diagnosis

- Similar to BERT testing; data patterns help determine problems with transmission line
- Advisable to run tests specifying the maximum payload for problem determination
- Common patterns:

- All 1s pattern is maximum density on T1/DS1

```
ping 10.0.2.1 pattern ffff count 10000 rapid size 1500
```

- All zeros pattern is minimum density on T1/DS1 interface

```
ping 10.0.2.1 pattern 0000 count 10000 rapid size 1500
```

- Repeat with 0x5555, 0x8080, and 0x1111

### Ping Testing with Patterns

You can use ping testing to test a circuit, or, alternatively, to diagnose a problem with the transmission circuit for E1 and T1. By changing the payload pattern, you can detect problems with ones density and zero suppression code settings.

### Testing the MTU

You should set ping tests with patterns for payload with the size parameter to generate frames close to the interface's MTU.

*Continued on next page.*

## Common Patterns

Sending ping packets with the payload containing certain bit patterns might provide pointers as to what type of problem exists, depending on the ping failure rate associated with a particular pattern.

Patterns commonly used for this purpose include the following:

- *FFFF*: all ones;
- *0000*: all zeros; and
- *5555*: alternating ones and zeros.

JUNOS platforms also support BERT testing, as previously described on E1 and T1 interfaces. Because BERT testing is a far more definitive test, you should consider BERT testing when you suspect marginal performance or intermittent operation.

## Troubleshooting SONET/SDH

- SONET/SDH has section, line, and path OAM
  - Easier to determine where the problem lies than with T3 and E3 or T1 and E1
  - Numerous commands available for problem diagnosis
- You can localize the source of errors to a particular span
- JUNOS platforms have visibility over the entire path
  - Difficulty is learning what the various alarms indicate
  - Output of **show** commands facilitate pinpointing the problem on the transmission path
  - Note: J Series platforms do not offer SONET/SDH interfaces

```
user@host> show interfaces so-1/1/1 extensive
```

### SONET/SDH Has Embedded OAM

SONET/SDH transmission systems incorporate a multitude of Operation, Administration, and Maintenance (OAM) functionalities at the line, section, and path layers. Interpretation of the SONET/SDH errors is helpful in determining the source of the problem within the SONET/SDH network. The information within the various SONET/SDH counters is plentiful and, when properly interpreted, enables problem localization. The following pages cover the commands available for SONET/SDH troubleshooting in the CLI.

### Localizing Errors

Using the output from these commands, you can tell very easily where the problem lies on SONET/SDH links.

*Continued on next page.*



## Path Level Visibility


Because JUNOS platforms function as path-terminating equipment (PTE), they have end-to-end visibility. The difficulty is learning what the various alarms indicate. Using **show** commands, you can easily determine the nature of SONET alarms and error indications. Note that to change framing to SDH, you must configure framing under the [edit chassis] hierarchy level:

```
chassis {
  fpc 0 {
    pic 0 {
      framing sdh;
    }
  }
}
```

## Monitoring a SONET/SDH Interface

```

user@host> monitor interface so-1/1/1
enterprise                               Seconds: 168                               Time: 15:48:50
-----
Interface: so-1/1/1, Enabled, Link is Down
Encapsulation: Cisco-HDLC, Keepalives, Speed: OC3
Traffic statistics:
Input bytes:          375527568 (0 bps)          [0]
Output bytes:        6612857 (0 bps)           [475]
Input packets:       224001 (0 pps)            [0]
Output packets:      102090 (0 pps)           [20]
Encapsulation statistics:
Input keepalives:    0                        [0]
Output keepalives:  176                       [17]
Error statistics:
Input errors:        0                        [0]
Input drops:         0                        [0]
Input framing errors: 179                      [17]
Policed discards:   47                        [0]
L3 incompletes:     0                        [0]
L2 channel errors:  0                        [0]
L2 mismatch timeouts: 0                      [0]
Carrier transitions: 1                        [0]
Output errors:      0                        [0]
Output drops:       0                        [0]
F2      : 0x00 Z3      : 0x00 Z4      : 0x00
Interface warnings:
o Received keepalive count is zero
o Framing errors are increasing, check FCS configuration and link
    
```

© 2009 Juniper Networks, Inc. All rights reserved.  www.juniper.net | 5-56

### Monitoring SONET/SDH Interfaces

The **monitor interface** command can provide useful troubleshooting information for the SONET/SDH interface.

The statistics in the second column are the cumulative statistics from the last time the **clear interfaces statistics** command cleared them. The statistics in the third column are the statistics from the last execution of the **monitor interface** command.

If the framing errors increase, check the FCS and scrambling configuration. If the configuration is correct, check the cabling to the router, and have the carrier verify the integrity of the line.

If the input errors increase, check the cabling to the router, and have the carrier verify the integrity of the line.

## Displaying SONET/SDH Interface Status

```
user@host> show interfaces so-1/1/1
```

```
Physical interface: so-1/1/1, Enabled, Physical link is Down
Interface index: 17, SNMP ifIndex: 16
Description: router-02 pos 4/0
Link-level type: Cisco-HDLC, MTU: 4474, Clocking: Internal, SONET mode
Speed: OC3, Loopback: None, CRC: 32, Payload scrambler: Enabled
Device flags : Present Running Down
Interface flags: Hardware-Down Link-Layer-Down Point-To-Point SNMP Traps
Link flags : Keepalives
Keepalive Input: 621 (00:02:57 ago), Output: 889 (00:00:05 ago)
Input rate : 0 bps (0 pps), Output rate: 0 bps (0 pps)
Active alarms : LOL, LOS
Active defects : LOL, LOF, LOS, SEF, AIS-L, AIS-P, PLM-P
Logical interface so-1/1/1.0 (Index 18) (SNMP ifIndex 30)
Description: router-02 pos 4/0
Flags: Device-down Point-To-Point SNMP-Traps, Encapsulation: Cisco-HDLC
Protocol inet, MTU: 4470
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 10.10.10.48/30, Local: 10.10.10.50
Protocol iso, MTU: 4469
```

Check alarms

Verify settings at both ends

### Displaying SONET/SDH Interface Status

If the first line shows `Physical link is Up`, it means that the physical link is healthy and can pass packets. If the first line shows `Physical link is Down`, it means that the physical link is unhealthy and cannot pass packets. To display more extensive information about the SONET/SDH interface when the physical link is down, use the `show interface so-x/y/z extensive` command. Look at the active alarms and active defects for the SONET/SDH interface, and troubleshoot the SONET/SDH media accordingly.


## The SONET Path Trace

- SONET/SDH framing overhead includes support for a path-layer trace
  - JUNOS Software codes this field to indicate router and interface name by default
  - Useful for detecting problems with patching or loopbacks

```
user@San_Jose> show interfaces so-0/1/0 extensive | find trace
```

```
Received path trace: San_Jose so-0/1/0
53 61 6e 5f 4a 6f 73 65 20 73 6f 2d 30 2f 31 2f San_Jose so-0/1/
30 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0.....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0d 0a
Transmitted path trace: San_Jose so-0/1/0
53 61 6e 5f 4a 6f 73 65 20 73 6f 2d 30 2f 31 2f San_Jose so-0/1/
30 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0.....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

Transmitted and received path traces match!



Remote loopback is in effect

### SONET Path Trace

SONET and SDH framing overhead includes support for a path trace via the J1 byte in the path overhead. The path trace information is typically used to identify the device that is terminating the path layer. The path trace information is typically used to identify the device that is terminating the path layer. The slide shows the JUNOS Software default coding of the path trace field, which is coded to identify the router and SONET interface name. This information can prove invaluable when the goal is to confirm the correct patching of a transmission line, or when you suspect that a loopback might be in effect. In this example the transmitted and received path trace information confirms that *San\_Jose* is receiving its own transmitted path trace, which indicates that a loopback is in place somewhere in the SONET transmission path.

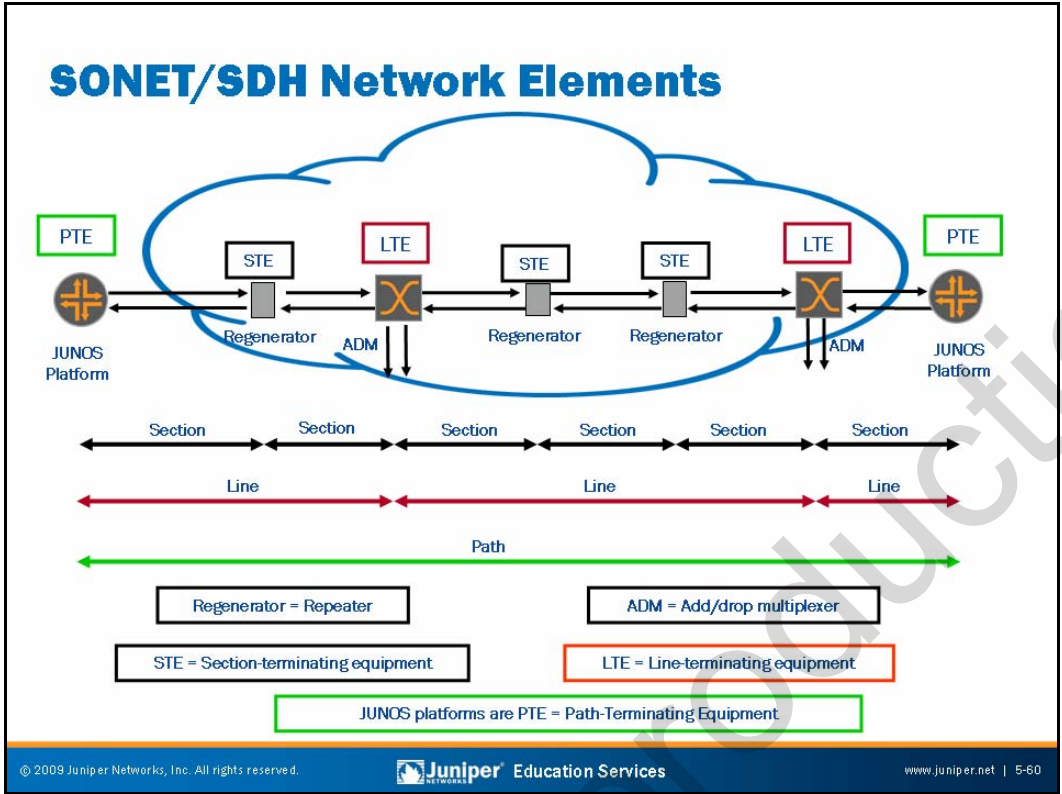
You can specify a custom path trace message with a **set sonet-options path-trace message** statement at the [edit interfaces **sonet-interface-name**] hierarchy. Note that custom path trace messages are not supported for ATM interfaces, which always use the default path trace coding.

*Continued on next page.*

### SONET Path Trace (contd.)

In SONET framing mode the path trace is 64 bytes, while in SDH mode the standards define a 16 byte trace. The difference in size can result in a truncated path trace when operating in SDH mode. The following is additional information from the ITU-T G.707 specification (G.707, ITU-T, March 1996) on the use of the J1 byte:

“This byte is used to transmit repetitively a path access point identifier so that a path receiving terminal can verify its continued connection to the intended transmitter. A 16-byte frame is defined for the transmission of an access point identifier. This 16-byte frame is identical to the 16-byte frame defined in 9.2.2.2 for the description of the byte J0. At international boundaries, or at the boundaries between the networks of different operators, the format defined in clause 3/G.831 shall be used unless otherwise mutually agreed by the operators providing the transport. Within a national network or within the domain of a single operator, this path access point identifier may use a 64-byte frame.”



### SONET/SDH Network Elements

The regenerators are considered section-terminating equipment (STE). STEs are responsible only for their particular section of the SONET/SDH span, as opposed to the entire SONET/SDH network path from one device designated as PTE to the other. They are responsible for simple regeneration of the SONET/SDH signal out to the next SONET/SDH equipment in the span. An STE only checks to ensure the incoming SONET/SDH frame arriving from its directly connected neighbor is good, while not having any knowledge of the rest of the span. STEs look only at the section overhead bytes of the SONET/SDH frame, even though they can rewrite the other overhead bytes if an alarm generates.

The add/drop multiplexers (ADMs) are considered line terminating equipment (LTE). LTEs have more knowledge of the SONET/SDH network than STEs, but they do not perform final processing of the SONET/SDH payload as PTEs do (although they can add and remove payloads). The SONET/SDH span from one LTE to another is referred to as a *line* span. LTEs mainly concern themselves with the line overhead bytes of the SONET/SDH frame.

JUNOS platforms are considered SONET/SDH PTEs. SONET/SDH PTEs are basically the endpoints of a typical SONET/SDH run. Because the SONET/SDH frame can traverse many regenerators and SONET/SDH multiplexers, the PTE is the final destination where the SONET/SDH frame terminates and the payload it carries receives processing. Hence, we consider the SONET/SDH span between two SONET/SDH PTEs a SONET/SDH path. PTEs pay particular attention to the path overhead bytes of the SONET/SDH frame.

*Continued on next page.*

### SONET/SDH Network Elements (contd.)

PTEs also play the LTE and STE role because they must look at the section overhead and line overhead bytes in addition to the path overhead bytes. LTEs also play the STE role because they must look at the section overhead bytes. STEs, on the other hand, must read only the section overhead bytes of the SONET/SDH frame, even though they do write into certain bytes of the line overhead at times.

All troubleshooting is from the perspective of the PTE (that is, the JUNOS platform). Although many situations exist where you cannot find the exact source of the problem unless you have access to the LTE or the STE, you can tell, at least from the PTE's perspective, that the problem is either somewhere upstream or local. With that said, the basic troubleshooting commands used to see any SONET/SDH line errors are **monitor interfaces so-0/0/1** and **show interfaces so-0/0/1 extensive**.

## Case Study 1: Local Fiber Break

Active alarms : LOL, PLL, LOS

Active defects : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P, PLM-P

SONET PHY:	Seconds	Count	State
PLL Lock	51	0	PLL Lock Error
PHY Light	51	0	Light Missing

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | 5-62

### Local Fiber Break

When a local fiber break occurs between Router B and add/drop multiplexer (ADM) B, the following alarms occur:

- **LOL:** A loss of light (LOL) alarm indicates a physical break in the connection to the router receive port. Check the connection between the router port and the first SONET/SDH network element (ADM A or ADM B in our example on the slide). The presence of an LOL alarm causes a cascading alarm effect, because no light means no signal, framing, and so forth.
- **LOS:** A loss of signal (LOS) alarm indicates a physical link problem with the connection to the router receive port. Check the connection between the router port and the first SONET/SDH network element (ADM in the example on the slide).
- **LOF:** A loss of frame (LOF) signal indicates detected errors in the A1 and A2 framing bytes. Check the connection between the router port and the first SONET/SDH network element.
- **AIS:** JUNOS Software sends an AIS downstream to signal an error condition. An AIS-L indicates that LOS or LOF is detected on an upstream STE. An AIS-P indicates that AIS-L, LOS, or LOF is detected on an upstream LTE. Work with the SONET/SDH network provider to locate the upstream SONET/SDH network element that detected the LOS or LOF.

*Continued on next page.*

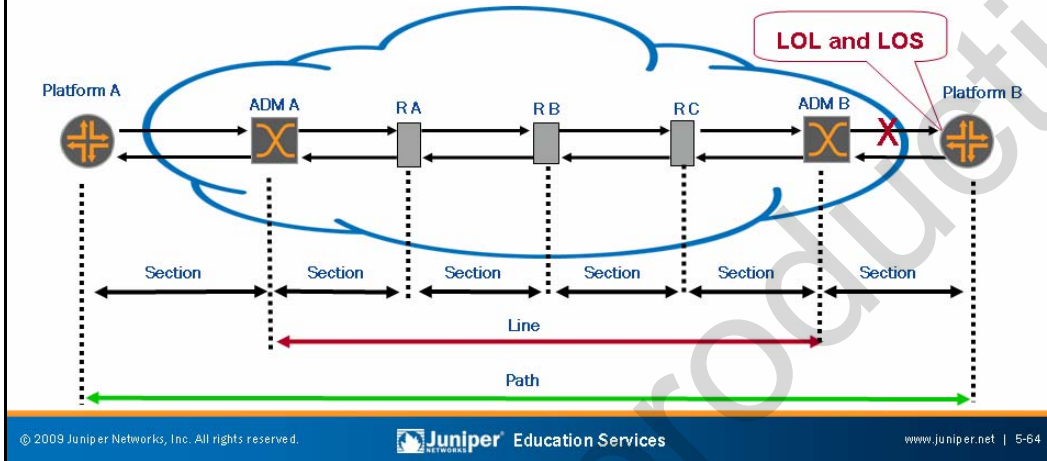


**Local Fiber Break (contd.)**

- *RDI*: JUNOS Software sends a remote defect indicator (RDI) upstream to signal an error condition. It sends an RDI-L upstream from LTE to LTE when the downstream LTE detects AIS-L, LOS, or LOF. It sends an RDI-P upstream from PTE to PTE when the downstream PTE detects AIS-P, AIS-L, LOS, or LOF. Work with the SONET/SDH network provider to locate the SONET/SDH network element that detected the LOS or LOF.
- *REI*: JUNOS Software sends a remote error indicator (REI) upstream to signal an error condition. It sends an REI-L to the upstream LTE when errors are detected in the B2 byte. It sends an REI-P to the upstream PTE when errors are detected in the B3 byte. Work with the SONET/SDH network provider to locate the source of the error condition. (Locate the section that causes the BIP-B1 errors.)

## Case Study 1: Why So Many Alarms?

Active alarms : LOL, PLL, LOS			
Active defects : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P, PIM-P			
SONET PHY:	Seconds	Count	State
PLL Lock	51	0	PLL Lock Error
PHY Light	51	0	Light Missing

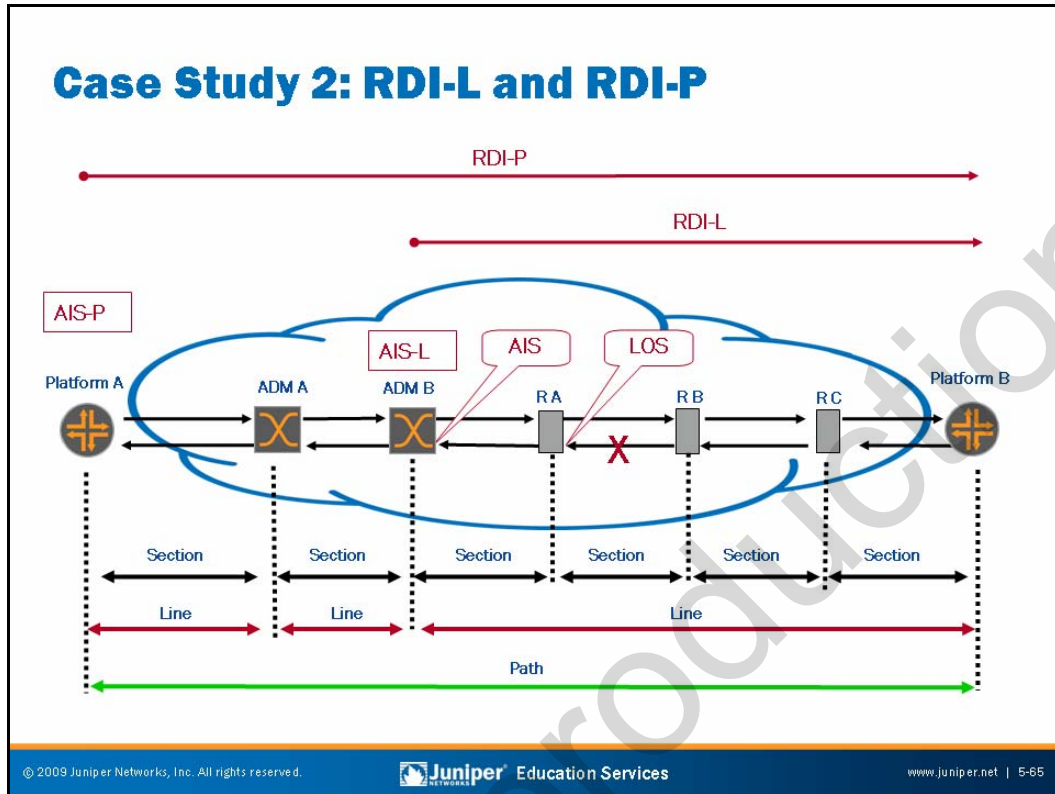


### Why So Many Alarms on the PTE?

The reasons why a broken cable should trigger an LOL and LOS alarm are obvious, but some might get confused when they also see AIS counters incrementing in the SONET/SDH line and SONET/SDH path. If no fiber plugs into the port, where do the AIS-L and AIS-P come from? In previous discussions, we mentioned that PTEs also have an LTE and STE component. Therefore, you must examine all SONET/SDH overhead.

When JUNOS Software raises the LOS alarm, the STE component sends SONET/SDH frames up to the LTE component. These frames have the K2 byte (in the section overhead) and the H1 and H2 bytes (in the line overhead) set to indicate AIS-L and AIS-P, respectively. When the LTE component sees the frame, it looks at the line overhead and finds the AIS-L signal in the K2 byte, and so the counter for AIS-L starts to increment. The frame is then handed to the PTE component, and the H1 and H2 bytes in the path overhead are examined and found to be AIS-P.

Thus, you must be able to prioritize the alarms present. After all, spending time on an AIS-P alarm when no light hits the interface does not lead to fast resolution of the problem.

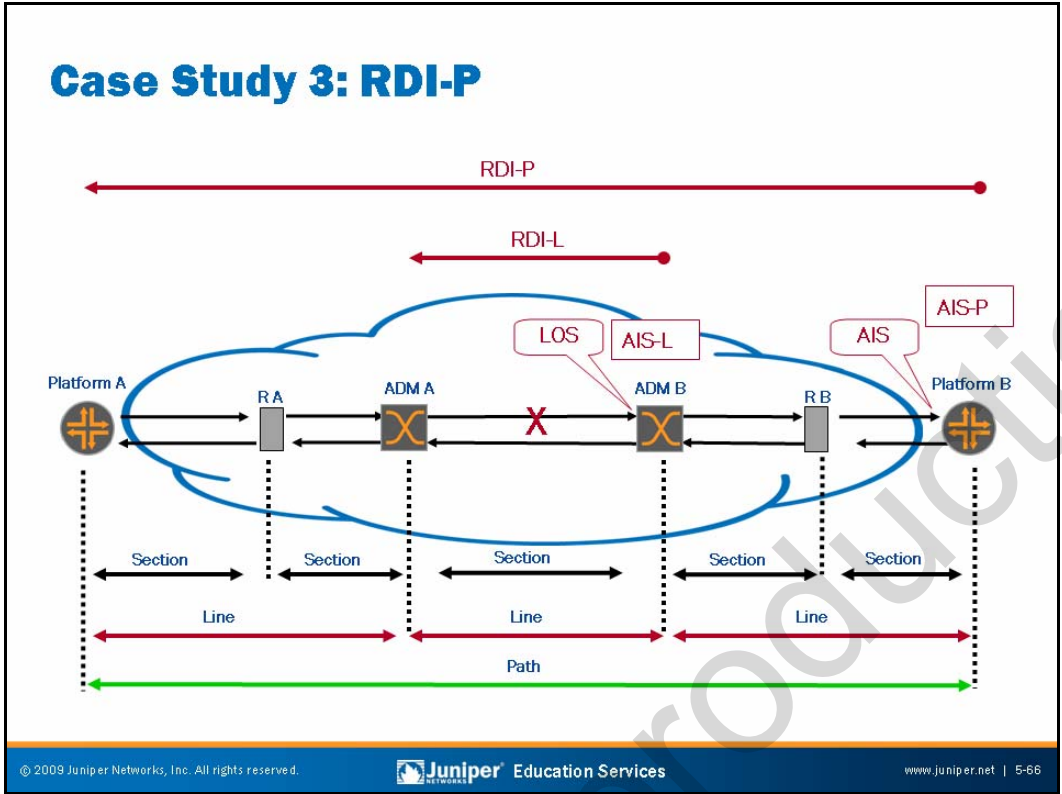


### Remote Defect Indicators (Line and Path)

Regenerator A (RA) notices that it receives no light, and after a very short time, it raises, at a minimum, an LOS alarm. RA rewrites clean section overhead bytes, but it writes all 1s in certain parts of the K2 byte residing in line overhead (bits 6–8), in the H1 and H2 bytes (pointer bytes used to find the path overhead and SONET/SDH payload envelope), and in the entire SONET/SDH payload envelope.

When ADM B receives the SONET/SDH frame, it sees clean section overhead, and because ADM B is an LTE, it also looks in the line overhead bytes and notices the 1s in the K2 byte (indicating an alarm indication signal, or AIS, in the SONET/SDH line span). If this alarm persists for a certain number of frames, it raises the AIS-L alarm and sends a remote defect indicator (RDI-L) back toward Platform B (in bits 5–7 of the SONET/SDH G1 byte). (Note that ADM A does not look at the path overhead because it is not a PTE.)

Finally, when Platform A receives the SONET/SDH frame, it looks at all overhead bytes because it is a PTE. It sees good section overhead, but when it looks in the line overhead, it notices all 1s in the H1 and H2 bytes (indicating AIS in the SONET/SDH path). Hence, it eventually raises an AIS-P alarm (when it receives enough frames in this state) and sends an RDI-P back toward JUNOS platform B. If the SONET/SDH path is clean all the way back to Platform B, it sees the RDI-P and raises this alarm in addition to the RDI-L it saw from ADM B.



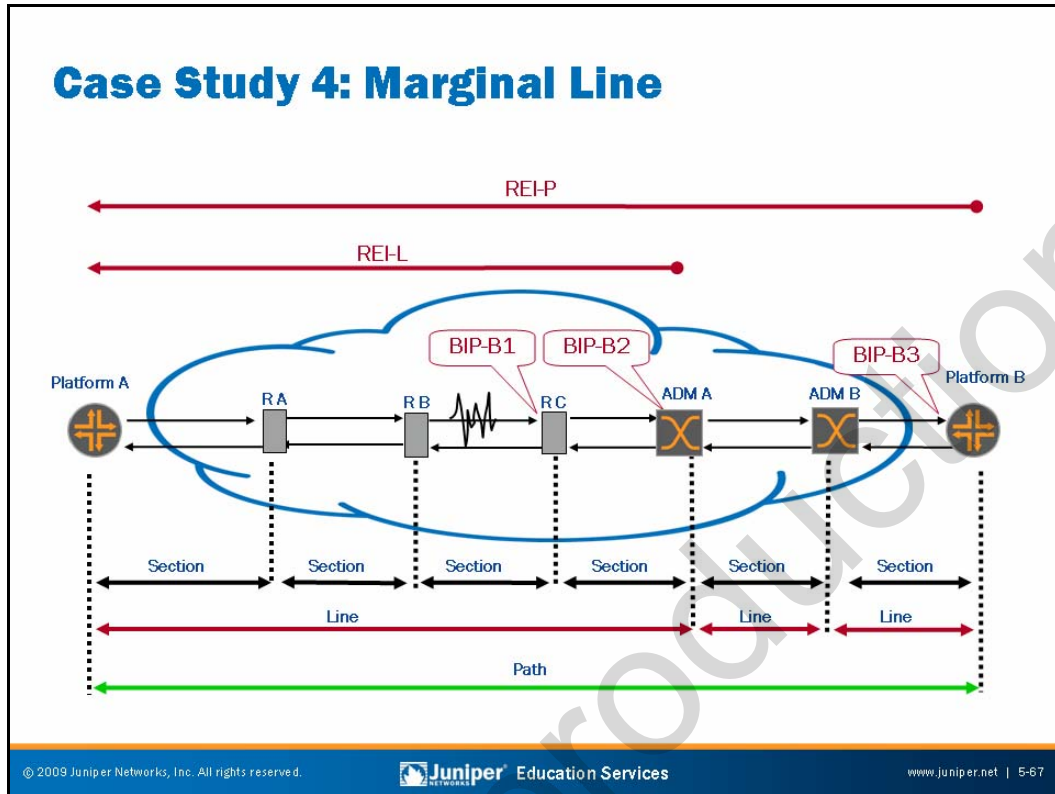
### Remote Defect Indicator (Path)

ADM B is an LTE that has the STE component built in. When the break shown on the slide occurs, the STE portion of ADM B notices the lack of light and eventually raises an LOS alarm by setting the appropriate bytes of the LOH with 1s. The SONET/SDH frame then traverses the internal components of ADM B.

ADM B's LTE portion notices the AIS-L and sends an RDI-L back to ADM A. Note that the RDI-L does not go back to Platform A; it terminates at ADM A. Basically, it is a SONET/SDH *line span error*, so it only remains within that line span. Moving forward, the section overhead and line overhead bytes (with the exception of the pointer bytes, at which PTEs look) are set back to normal, and the SONET/SDH frame travels out to Regenerator B (RB). Note that errors from section overhead and line overhead do not get passed on because section overhead and line overhead are rewritten clean as they leave ADM B.

RB sees good section overhead, so it sends the SONET/SDH frame out to Platform B. Note that it rewrites the section overhead in this process.

When Platform B receives the SONET/SDH frame, it notices clean section overhead and line overhead, except for all 1s in the pointer bytes and thus, raises the AIS-P alarm. Because AIS-P is raised, Platform B sends an RDI-P back towards Platform A. As it passes through ADM B and ADM A, both ADM A and ADM B rewrite clean section overhead and line overhead as usual, but the RDI-L is *not* considered clean as it leaves ADM A towards Platform A.



## Marginal Line

This example discusses the result of bit errors caused by a marginal line between Regenerator B and Regenerator C (RC). When RC receives the frame, it does a parity check on the B1 byte in the section overhead. It logs the occurrence of BIP-B1 errors before sending the SONET/SDH frame out to ADM A. However, as usual, it rewrites the section overhead; thus, the B1 byte should be clean.

Because B1 parity errors occurred, it is most likely that the rest of the SONET/SDH frame is also corrupt. When ADM A receives the SONET/SDH frame, it sees clean section overhead (thus, a good B1 byte). However, when it runs a parity check on the B2 byte in the line overhead, it likely sees BIP-B2 errors and raises an REI-L toward Platform A (REI-L is conveyed using the M0 and M1 byte in the line overhead).

When ADM A sends the SONET/SDH frame out to ADM B, it rewrites the section overhead and line overhead; thus, the B1 and B2 bytes are clean as the frame gets to ADM B. ADM B is an LTE, so it sees good B1 and B2 bytes (but does not look at the B3 byte in the path overhead) and forwards the frame to Platform B.

When the SONET/SDH frame reaches Platform B, the STE component runs a parity check on the B1 bytes and see it as good. It passes the frame internally up to its LTE component where the LTE component finds a good B2 byte when performing the parity check. However, when the PTE portion finally receives the SONET/SDH frame, it looks in the path overhead and calculates a bad B3 byte. When this calculation happens, it logs a BIP-B3 error and sends an REI-P back toward Platform A (certain bits of the G1 byte indicate REI-P).

## Point-to-Multipoint Topologies

- Port types:
  - e1
  - t1
  - e3
  - t3
  - so
  - at
- Encapsulation  
(`frame-relay`  
or `atm-pvc`)
- Tools:
  - **ping**
    - **atm-ping** for VC verification
  - **show interfaces extensive** (physical interface problems)
  - **show interfaces detail** (logical interface information)
  - ATM OAM cells
  - **monitor traffic**
  - **monitor interface** (physical interface only)
  - **show ilmi**
  - Loopback (physical interface only)

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-68

### Port Types

Point-to-multipoint and nonbroadcast multiaccess (NBMA) technologies exhibit unique characteristics and therefore require additional sophistication in your troubleshooting approach. You must differentiate whether the problem exists on the physical level or the logical level when dealing with point-to-multipoint and NBMA interfaces.

### Encapsulation

Encapsulation types for point-to-multipoint topologies are typically `frame-relay` and `atm-pvc`.

### JUNOS Software Tools

We discuss the JUNOS Software tools shown on the slide on the following pages.

## Frame Relay Configuration (1 of 2)

- The JUNOS device is typically DTE
  - Multilink support
- When connecting back to back:
  - Disable keepalives (LMI), or configure one device as a DCE
- Encapsulation set on physical interface, sometimes on the logical interface too (with circuit cross-connect)

```

user@host# show
encapsulation frame-relay;
unit 0 {
  point-to-point;
  dlci 512;
  family inet {
    address 192.168.1.1/30;
  }
}

```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-69

### Router Is Normally DTE

When you configure an interface with Frame Relay encapsulation, the router is assumed to be data terminal equipment (DTE). That is, JUNOS Software assumes the JUNOS platform to be at a terminal point on the network. To configure the platform to be data circuit-terminating equipment (DCE), include the **dce** statement at the [edit interfaces *interface-name*] hierarchy level:

```

[edit interfaces interface-name]
user@host# set dce

```

When you configure the device to be a DCE, keepalive generation is disabled by default, but the device responds to keepalive messages received from the DTE.

JUNOS Software supports Multilink Frame Relay (MLFR) as defined in FRF.16 for bonding T1 or E1 links. You cannot mix and match T1 and E1. JUNOS Software also supports Multilink Point-to-Point Protocol (MLPPP).

### Back-to-Back Connections

For back-to-back Frame Relay connections, either disable the sending of keepalives on both sides of the connection, or configure one side of the connection to function as a DCE (the default is a DTE line discipline).

*Continued on next page.*



## Encapsulation Configuration

As shown on the slide, except in the case of circuit cross-connect, you should specify the Frame Relay encapsulation configuration at the interface device level. Configure the specification of the connection type (point-to-point versus point-to-multipoint), and the DLCI at the logical device level. The slide shows a typical point-to-point configuration for a DTE device.

Not for Reproduction



## Frame Relay Configuration (2 of 2)

- JUNOS devices support LMI-type interoperability

- Default is ANSI T1-617 Annex D, both UNI and NNI modes are supported
- RFC 2427 (RFC 1490) encapsulation
  - Under Cisco IOS, configure `encapsulation frame-relay ietf`
- For ITU-T Q933 Annex A

```
[edit interfaces e3-0/1/1 lmi]
user@host# set lmi-type itu
```

- Enable support for Inverse ARP replies

- Described in RFC 2390, disabled by default

```
[edit interfaces t3-1/2/0 unit 201]
user@host# set inverse-arp
```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-71

### PVC Management Protocol

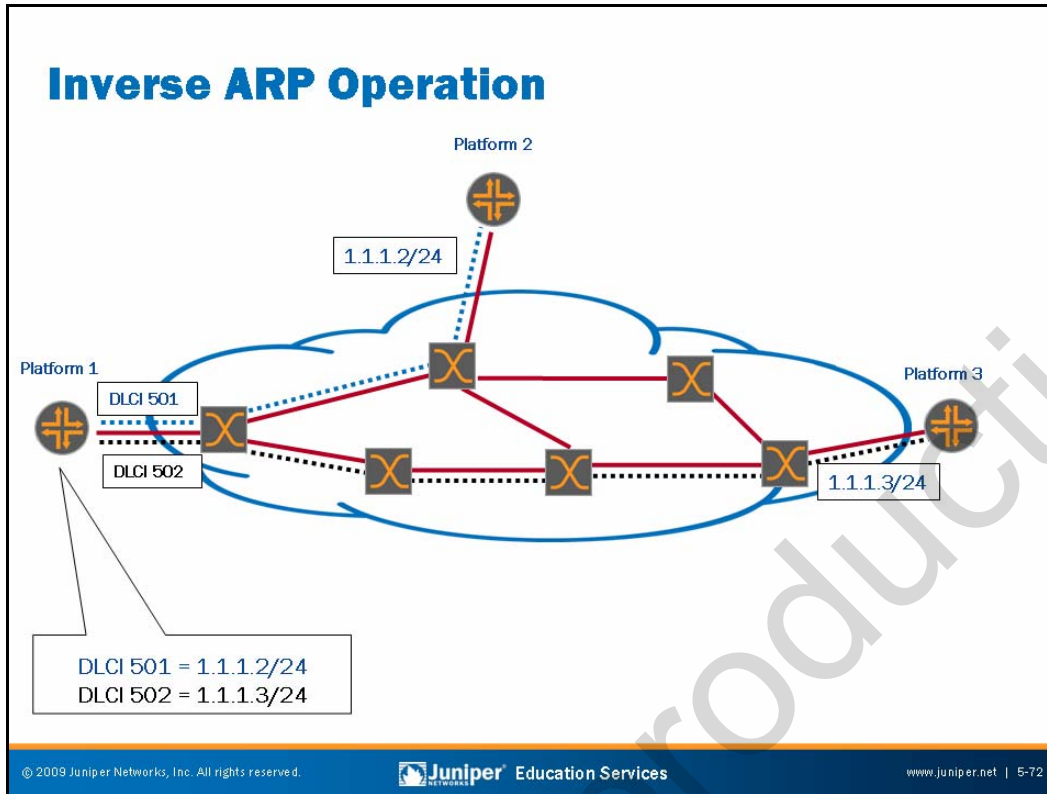
For LMI—a permanent virtual circuit (PVC) management protocol—interoperability, you might need to configure the LMI type to be either conformant to the ANSI T1-617 Annex D, or alternatively, to the ITU-T Q933 Annex A variant. The ANSI type is the default management type when encapsulation is set on the physical interface as Frame Relay. JUNOS platforms do not support the Frame Relay Forum's LMI specification, which is a default for equipment made by other vendors; thus, watch for interoperability issues. To alter the management protocol, issue the following commands:

```
[edit]
user@host# set interfaces so-0/1/1 encapsulation frame-relay
user@host# set interfaces so-0/1/1 lmi-type itu
```

### Enabling Inverse ARP

You can configure Inverse ARP, as defined in RFC 2390, on the logical interface with the following command. By default, Inverse ARP is disabled.

```
[edit]
user@host# set interfaces so-0/1/1 unit 25 inverse-arp
```



### Inverse ARP Operation

When a Frame Relay interface comes up, the Frame Relay switch announces the configured DLCIs to the router. Once these DLCIs become active, some vendors might attempt to map the remote Network Layer address (that is, the IP address) to the local DLCIs.

On the slide, Platform 1 learns about DLCIs 501 and 502 from the local Frame Relay switch. If the remote DLCIs are active, Platform 1 sends an Inverse ARP request over the active DLCIs to learn the IP addresses of the remote devices. Platform 1 then maps these Layer 2-to-Layer 3 addresses for use when routing packets between sites. Inverse ARP is not used on point-to-point interfaces.

With Inverse ARP you can resolve the IP addresses of directly connected Frame Relay peers. Thus, in the partial-mesh topology illustrated on the slide, reachability problems between Platform 2 and Platform 3 might exist. Note that these two stations are not directly connected through the Frame Relay cloud. You can resolve this problem in one of two ways: by configuring a full-mesh topology or by configuring a point-to-point operation on the logical interfaces. In the latter approach, each point-to-point link receives a unique IP number, whereas in the full-mesh scenario, all routers comprising the Frame Relay mesh share a common IP subnet. Defining Frame Relay connections as point-to-point eliminates the need for Inverse ARP and allows packets exchanged between Platform 2 and Platform 3 to route through Platform 1.

## Troubleshooting Frame Relay

- Perform local loopback (port testing)
- Verify logical settings
  - Verification of line status by checking LMI status using **show interfaces detail** command

ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds

LMI statistics:

```
Input : 67 (last seen 00:00:03 ago)
Output: 67 (last sent 00:00:03 ago)
```

Symmetric counts indicate  
no lost keepalives

- LMI has significance only between router and Central Office switch—ensure the LMI type is correct (ANSI or ITU-T)
- Check for DLCI mismatches or assignment to incorrect logical interface
- Verify both routers have Frame Relay encapsulation configured
- If problem does not show remotely, the telco must test its network

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-73

### Physical Interface Troubleshooting

To troubleshoot Frame Relay, verify whether the problem is on the physical port or the virtual circuit. You can troubleshoot using loops to determine where the problem lies. To verify the physical port, use the **show interfaces extensive** command and ensure that no SONET/SDH, E3 and T3, or E1 and T1 alarm is present. If more than one logical interface is configured, they will typically all be down if the port, cable, or CSU/DSU device is faulty.

### Logical Interface Troubleshooting

For problems on the logical interface level (that is, virtual circuit), check for the following:

- Mismatched LMI type;
- Incorrectly set encapsulation on one of the routers;
- Correctly configured DLCI values;
- DLCI-to-logical interface assignment; and
- If keepalives are sending and receiving.

You should use the **show interfaces brief** command to verify settings, LMI type, and packets sent and received. (LMI provides keepalive functionality in Frame Relay.)

## Troubleshooting ATM

- Determine if the problem is with the physical interface:
  - ATM runs over either SONET/SDH or T3 or E3, so troubleshooting these physical media is similar
  - All the VCs will be down (logical interfaces)
    - Use the **monitor interface at-0/1/1** command
    - Use the **show interfaces at-0/1/1 extensive** command
  - Loopback testing
- Determine if the problem is with the logical interface (virtual circuit):
  - Use the **show ilmi** command
  - Use the **show ilmi statistics** command
    - ILMI is not enabled by default—provides integrity testing of access link
  - Use the **show interfaces at-0/1/1 detail** command
    - Shows statistics for all logical interfaces
  - OAM F5 cells (loopback, RDI, and AIS)
  - Use the **ping atm** command

### Physical Interface Troubleshooting

The approach for troubleshooting ATM is similar to that for Frame Relay—verify whether the problem is on the physical port or the virtual circuit. To verify the physical port, use the **show interfaces extensive** command, and ensure that no SONET/SDH, E3 or T3, or alarm is present. If more than one logical interface is configured, they will typically all be down if the port, cable, or CSU/DSU device is faulty.

You also can use loopback testing to determine where the problem lies.

### Logical Interface Troubleshooting

For problems on the logical interface level (that is, the virtual circuit), check for the following:

- Correct operation of the Integrated Local Management Interface (ILMI) protocol, if enabled;
- Correct VPI and VCI;
- Operation of the Layer 2 ATM virtual circuit without requiring (or involving) IP functionality by using the ATM-based ping command; and
- Correct VPI and VCI to logical interface mappings.

*Continued on next page.*

## Logical Interface Troubleshooting (contd.)

You should use the **show interfaces brief** command to verify settings like VPI and VCI on the various logical interfaces. Another helpful command is **show ilmi all**.

The following is a typical point-to-point ATM interface configuration. This configuration includes support for the ILMI protocol and periodic OAM cell generation to provide keepalive functionality:

```
[edit interfaces at-0/2/0]
user@host# show
atm-options {
  vpi 0 {
    maximum-vcs 200;
  }
  ilmi;
}
unit 0 {
  vci 100;
  oam-period 10;
  oam-liveness {
    up-count 3;
    down-count 3;
  }
  family inet {
    address 10.0.16.1/24;
  }
}
```

## Logical Interface Problems

- Use ATM ping to determine whether the VC is broken
  - End-to-end or local segment option
- For logical interface statistics, use `show interfaces statistics` and `show interfaces detail` commands
  - Errors or drops on the VC statistics indicate a problem in the ATM cloud
- Use `show ilmi` command:

```
user@host> show ilmi all
Physical interface: at-6/2/1, VCI: 0.16 Peer IP address: 192.168.4.24,
Peer interface name: 1C4
Physical interface: at-6/3/0, VCI: 0.16 Peer IP address: 192.168.7.6,
Peer interface name: 2C3
Physical interface: at-6/4/0, VCI: 0.16 Peer IP address: 192.168.9.10,
Peer interface name: 1C2
```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-76

### ATM Pings

An ATM ping can be either end to end or local segment.

### Use the `show interfaces` Command

Errors or drops on the VC statistics indicate a problem in the ATM cloud or with local cabling.

### Verifying that ILMI Works

The ATM Forum specified the use of the ILMI protocol for the address registration of ATM edge devices and keepalive functions on the user-to-network interface (UNI).

You can configure ILMI to communicate with directly attached ATM switches to enable querying of the IP addresses and port numbers of the switches. To display ILMI statistics, use the command `show ilmi interface interface-name`. The router uses VC 0:16 to communicate with the ATM switch. This VPI and VCI pair is well known.

To enable ILMI communications between the router and its directly attached ATM switches, include the `ilmi` statement at the `[edit interfaces interface-name atm-options]` hierarchy level:

```
[edit interfaces interface-name atm-options]
user@host# set ilmi
```



## Displaying VC-Related Information

```

user@host> show interfaces brief at-0/3/0
Physical interface: at-0/3/0, Enabled, Physical link is Up
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode, Speed:
  OC3,
  Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  Logical interface at-0/3/0.104
    Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-SNAP
    inet 10.0.9.6/30
    VCI 0.104
      Flags: Active

Logical interface at-0/3/0.16384
  Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-VCMUX
  VCI 0.16
    Flags: Active, ILMI
    Total down time: 0 sec, Last down: Never
  
```

© 2009 Juniper Networks, Inc. All rights reserved.

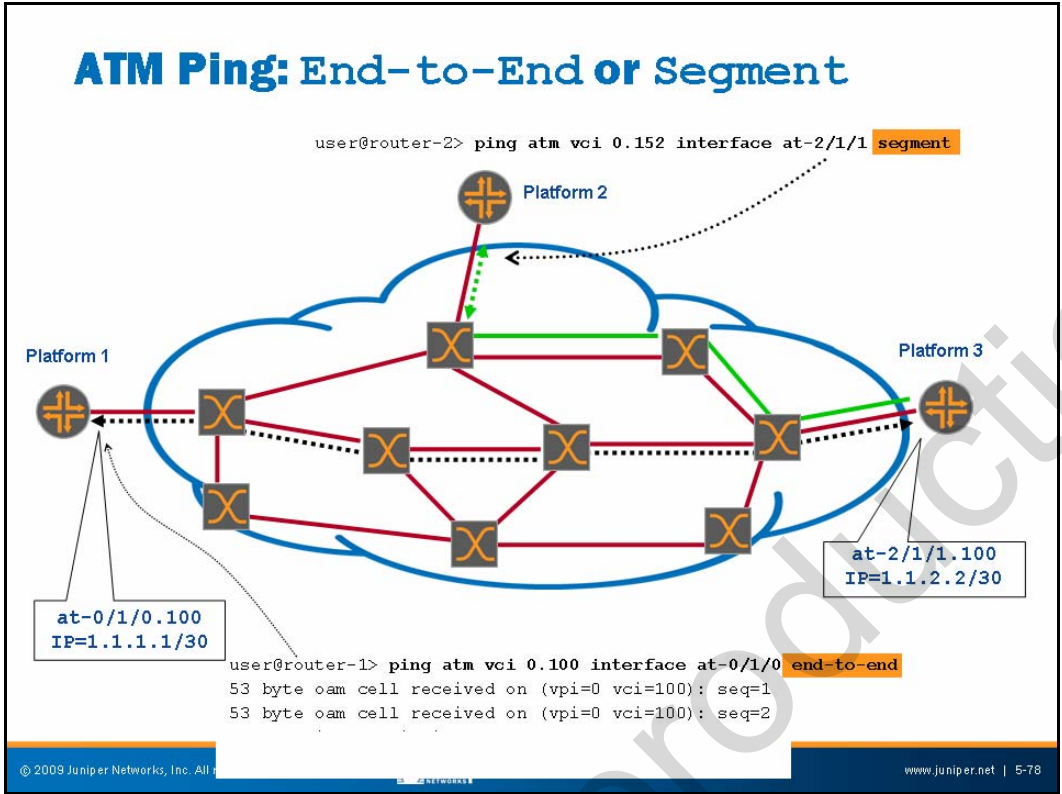
 Juniper Education Services

www.juniper.net | 5-77

### Command Enhancements

The JUNOS Software **show interfaces** command has several features that make dealing with multipoint interfaces easier:

- **show interfaces brief:** Includes VPI, VCI, and DLCI values for logical interfaces, including status and keepalive settings.
- **show interfaces detail:** Includes multipoint VPI and VCI to IP address mappings.



### ATM Ping: End-to-End or Segment

You can use the **atm** option with the **ping** command to verify that the virtual circuit is functional through the use of segment or end-to-end F5 (VC level) OAM cell flows. The key point is that these pings do not involve IP or ICMP, and, as such, are used to test the ATM layer itself. As shown on the slide, a segment-level ATM ping is returned by the device terminating that VC segment; typically this device is the ingress ATM switch port, as shown in the case of Platform 2. In contrast, the **end-to-end** switch causes the OAM cells to loop by the device, which terminates the VC, as shown in the case of Platform 1 and Platform 3.



## Agenda: Interface Troubleshooting

- Interface Configuration Overview
- General Interface Troubleshooting
- Media-Specific Interface Troubleshooting
- Case Study

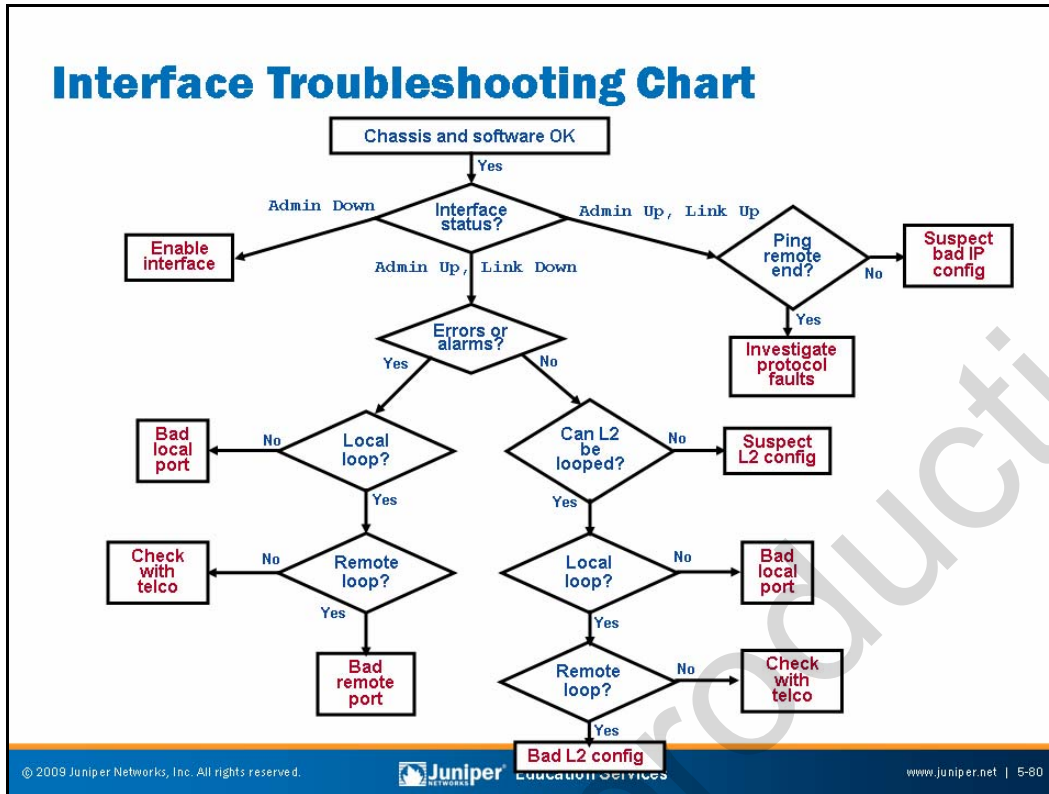
© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-79

### Case Study

The slide highlights the topic we discuss next.




### Interface Troubleshooting Flowchart

The purpose of the interface troubleshooting flow chart shown on the slide is simply to provide a set of high-level steps and decision points designed to get you started on the path of interface and transmission line troubleshooting. Note that reasonable people might disagree on the exact ordering of the steps or on the particulars of the CLI commands that could be used to help isolate an interface or circuit problem.

## Interface Case Study A (1 of 4)

- Case study background:
  - The SONET circuit between *Tokyo* and *London* is down at the link level
    - Configured for `cisco-hdlc` encapsulation with `no-keepalives`
  - No chassis hardware alarms or software malfunctions are evident
- What is wrong?
  - What CLI commands and fault analysis steps can help narrow down a possible cause?



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 5-81

### Interface Case Study A: Background

The slide sets the stage for a sample interface troubleshooting case study. We begin with a general description of the problem, which, in this case, indicates that the SONET link between the London and Tokyo devices appears to be down. In this example the link-level protocol is `cisco-hdlc` with `keepalives` disabled. You can assume that no chassis hardware problems or JUNOS Software faults exist.

### Feeling Lucky?

Based on this description you would be pretty lucky if you already knew the cause of the problem. After all, it could be the SONET transmission link or the PIC or port at either end, right? We suggest that you follow the general steps outlined on the sample interface troubleshooting flow chart to get things started. Put another way, it might be a good idea to start with the determination of interface status.

## Interface Case Study A (2 of 4)

- Sample course of action:
  1. Determine interface status:
 

```

user@Tokyo> show interfaces so-0/1/1 terse
Interface      Admin Link Proto Local
so-0/1/1      up   down  inet  10.0.22.1/24
so-0/1/1.0    up   down  inet  10.0.22.1/24
                    
```

Admin up, link level down
  2. Any errors or alarms?
 

```

user@Tokyo> show interfaces so-0/1/1
Physical interface: so-0/1/1, Enabled, Physical link is Down
Interface index: 133, SNMP ifIndex: 23
Link-level type: Cisco-HDLC, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
Loopback: None,
FCS: 16, Payload scrambler: Enabled
Device flags   : Present Running Down
Interface flags: Hardware-Down Point-To-Point SNMP-Traps
Link flags     : No-Keepalives
. . .
Input rate    : 0 bps (0 pps)
Output rate   : 0 bps (0 pps)
SONET alarms  : RDI-L
SONET defects : RDI-L, RDI-P
                    
```

Alarms present: RDI means a problem exists in the Tokyo to London direction

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | 5-82

### Interface Case Study A: Course of Action

The slide provides examples of troubleshooting steps based on the sample interface troubleshooting flow chart. In this case, you begin by displaying the interface status for the so-0/1/1 interface at the Tokyo station. The results indicate that the interface is up administratively but down at the link level.

You next display the status of the so-0/1/1 interface to determine if any media alarms or errors are present. As called out by the comments, SONET alarms and defects are present in the form of RDIs at the line and path levels. This media-specific alarm indicates that errors are occurring in the Tokyo-to-London direction. The fact that Tokyo is receiving RDI indications implies that it can receive information sent from London.

Note that at this time, the remote device (London) is displaying the following interface alarms:

```

user@London> show interfaces so-0/1/1 | match sonet
Link-level type: Cisco-HDLC, MTU: 4474, Clocking: Internal, SONET mode,
Speed: OC3, Loopback: None, FCS: 16,
SONET alarms   : LOL, LOS
SONET defects  : LOL, LOF, LOS, SEF, AIS-L, AIS-P
                    
```

The LOL indications are in keeping with the symptoms observed thus far. Because London is receiving no signal, it has lost SONET framing and generates an RDI signal back to Tokyo. Based on these results, you can confirm that a problem exists, but you are not yet able to eliminate any possible causes.

## Interface Case Study A (3 of 4)

- Sample course of action (contd.):

- Configure a local loopback:

```

user@Tokyo> configure
Entering configuration mode

[edit]
user@Tokyo# set interfaces so-0/1/1 sonet-options loopback local

[edit]
user@Tokyo# commit and-quit
commit complete
Exiting configuration mode
    
```

- Confirm local loopback results:

```

user@Tokyo> show interfaces so-0/1/1 terse
Interface      Admin Link Proto Local
so-0/1/1      up   up   inet  10.0.22.1/24
so-0/1/1.0    up   up

user@Tokyo> ping 10.0.22.2 count 1
PING 10.0.22.2 (10.0.22.2): 56 data bytes
36 bytes from 10.0.22.1: Time to live exceeded
Vr HL TOS Len ID Flg  off TTL Pro  cks  Src      Dst
 4  5  00 0054 6039  0 0000 01  01 196e 10.0.22.1 10.0.22.2

--- 10.0.22.2 ping statistics ---
1 packets transmitted, 0 packets received, 100% packet loss
    
```

Link is up, traffic is passing (TTL expired)

Local loop is possible because of L2 configuration

### Interface Case Study A: Configuration

The next step is to conduct a loopback test designed to eliminate either the local device or the transmission line and the remote device as possible sources for this problem. You can configure a local loopback at either end with similar results; because you are at the Tokyo station, you decide to configure a local loopback there first. Note that this loopback should succeed, assuming that the PIC and port are OK, because of the specifics of the link-level protocol currently in effect (*cisco-hdlc* with *no-keepalives*).

After configuration, the local loopback status of the *so-0/1/1* interface displays again. The output confirms that the interface is now up at the link level. With the interface up, you can test the loopback's ability to pass data with ping traffic destined to the remote end of the circuit (London's address), as shown on the slide. The TTL expired error message actually helps because it confirms that traffic is passing over the local loopback.

*Continued on next page.*

## Interface Case Study: Configuration (contd.)


Before moving on you should be sure to remove the local loopback at the Tokyo station. Use the following commands:

```
user@Tokyo> configure
Entering configuration mode
[edit]
user@Tokyo# rollback 1
load complete
[edit]
user@Tokyo# commit and-quit
commit complete
Exiting configuration mode
```

Note that because this loopback is internal and local, you have not fully tested the so-0/1/1 interface at Tokyo; it would be best to attach an external loopback plug to another device to conduct an external local loopback test. This technique has the added advantage of not requiring configuration to effect, and then again to remove, to the loopback.

## Interface Case Study A (4 of 4)

- What can you eliminate given the results obtained thus far?



- What test should you perform next?
- Assume that a local loop also passes at London
  - Where is the fault?

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 5-85

### What Do These Results Indicate?

The results of your testing thus far indicate that the problem is not in Tokyo's so-0/1/1 interface. Thus, you now must perform a similar test at the London station to isolate between its SONET interface and the transmission line.


### Assume Similar Results

Although not shown in the interest of brevity, you can assume that a local loopback test at the London station also passes. This test eliminates the London station and leaves the SONET transmission line as the most likely reason for the outage. Based on this information, you should decide to escalate the problem to the transmission or telco group.



## Interface Case Study B (1 of 3)

- **Case study background:**
  - The SONET circuit between the San Jose and Montreal devices is not passing traffic
    - Keepalives are disabled
    - Both ends indicate an Up status at the Data Link Layer
    - Local loops are successful at both ends, but the transmission line returns as no trouble found
  - No hardware alarms or software malfunctions are evident
  - You cannot view the configuration directly
- **What is wrong?**
  - What CLI commands and fault analysis steps can help narrow down a possible cause?



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 5-86

### Interface Case Study B: Background

The slide sets the stage for a second interface troubleshooting case study. We begin with a general description of the problem, which, in this case, indicates that the SONET link between the San Jose and Montreal devices is not passing traffic, despite both ends indicating an Up status at the Data Link Layer. The slide notes that keepalives are disabled at both ends. In this example you can assume that a local loop has been conducted at both ends with the expected results, which tends to identify the SONET transmission facility as the culprit. The problem is that the SONET transmission line returned as no trouble found (NTF), which leaves you scratching your head.

Note that in this example, you are not permitted to view the interface-related configuration at either end.

### Feeling Lucky?

Based on this description, you would be pretty lucky if you already knew the cause of the problem. While the symptoms indicate that the SONET transmission link is not the problem, the fact that both ends pass a local loopback tends to indicate that each router has a functional interface, yet for some reason the two ends do not want to communicate.

*Continued on next page.*



**Feeling Lucky? (contd.)**

This is tricky; we suggest that you follow the general steps outlined on the sample interface troubleshooting flow chart to get things started. Put another way, the lack of alarms and error indications, coupled with the knowledge that local loops passed and the transmission line was confirmed, pretty much points to Layer 2 configuration problems between the two stations. While displaying the configuration is a good place to start, the rules of engagement in effect prohibit you from viewing the configuration.

What will you do?

Not for Reproduction

## Interface Case Study B (2 of 3)

- Sample course of action:

- Look for I/O errors:

```

user@San_Jose> show interfaces so-0/1/0 extensive
Physical interface: so-0/1/0, Enabled, Physical link is Up
Interface index: 132, SNMP ifIndex: 65, Generation: 15
. . .
Input errors:
Errors: 1, Drops: 0, Framing errors: 3, Runts: 0, Giants: 1,
Bucket drops: 0, Policed discards: 86, L3 incompletes: 0,
. . .
user@Montreal> show interfaces so-0/1/0 extensive
Physical interface: so-0/1/0, Enabled, Physical link is Up
Interface index: 132, SNMP ifIndex: 48, Generation: 15
. . .
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
Bucket drops: 0, Policed discards: 0, L3 incompletes: 0,
L2 channel errors: 461, L2 mismatch timeouts: 0, HS link CRC errors: 0,
HS link FIFO overflows: 0
. . .
    
```

Policed discards occurring at San Jose

L2 channel errors occurring at Montreal

### Interface Case Study B: Course of Action—Part 1

Suspecting a Layer 2 configuration mismatch, you begin by displaying extensive interface information for the so-0/1/1 interfaces at the San Jose and Montreal stations while you attempt to ping the local router's address from the remote end of the circuit. The results indicate detection of input errors at both ends. Policed discards occur at the San Jose station. These errors indicates receipt of traffic for an unconfigured protocol. In the case of Montreal, Layer 2 channel errors are detected. This error indicates receipt of traffic for an unconfigured logical interface.

These symptoms definitely point to some sort of Layer 2 encapsulation mismatch between the two stations.

## Interface Case Study B (3 of 3)

- Sample course of action (contd.):

- Monitor traffic to reverse-engineer Layer 2 configuration

- Open a second session to each route so that ping traffic can generate while you monitor the results locally:

```

user@San_Jose> monitor traffic interface so-0/1/0 layer2-headers
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Listening on so-0/1/0, capture size 96 bytes

11:50:57.110708 Out unicast 88 0800: IP 10.0.1.2 > 10.0.1.1: icmp 64: echo
request
11:50:58.120706 Out unicast 88 0800: IP 10.0.1.2 > 10.0.1.1: icmp 64: echo
request
. . .
Packet length -----> EtherType for IP

user@Montreal> monitor traffic interface so-0/1/0 layer2-headers
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Listening on so-0/1/0, capture size 96 bytes

15:57:50.245935 Out 88 1841=DLCI(100) cc IP: IP 10.0.1.1 > 10.0.1.2: icmp
64: echo request
15:57:51.253562 Out 88 1841=DLCI(100) cc IP: IP 10.0.1.1 > 10.0.1.2: icmp
64: echo request
Frame Relay encapsulation -----> Direct NLPID for IP
    
```

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | 5-89

### Interface Case Study B: Course of Action—Part 2

You decide to use the JUNOS CLI monitor traffic utility to attempt reverse engineering of the Layer 2 encapsulation currently in effect at both ends of the circuit, because you cannot view the configuration directly. In this example, you first establish a second login to the tested device (using an out-of-band connection), so that you can monitor the egress traffic that results from local ping attempts at the remote end of the circuit.

The first such test that is conducted at the San Jose station, confirms that ping traffic leaves the so-0/1/0 interface with a rather basic form of encapsulation. Although the encapsulation type is not explicitly called out, the use of an EtherType code to identify IP traffic indicates that this interface is configured for Cisco HDLC encapsulation. A similar test conducted at the Montreal station clearly shows that its egress traffic makes use of Frame Relay encapsulation.

With these results you can conclusively determine that the Layer 2 encapsulation parameters are mismatched between the two stations. Note that this mismatch would have resulted in a link-down status if keepalives were not disabled.

*Continued on next page.*

## Interface Case Study B: Course of Action—Part 2 (contd.)


You can make a similar determination by simply displaying interface status with an eye towards the Layer 2 encapsulation settings. We did not show this approach in this case study so that we could demonstrate the utility of the **monitor traffic** command.

```
user@San_Jose> show interfaces so-0/1/0 | match Link
Physical interface: so-0/1/0, Enabled, Physical link is Up
  Link-level type: Cisco-HDLC, MTU: 4474, Clocking: Internal, SONET mode,
  Link flags      : No-Keepalives
user@Montreal> show interfaces so-0/1/0 | match link
Physical interface: so-0/1/0, Enabled, Physical link is Up
  Link-level type: Frame-Relay, MTU: 4474, Clocking: Internal, SONET mode,
  Link flags      : No-Keepalives DTE
```

## Interface Case Study C (1 of 4)

### ■ Case study background:

- The Ethernet circuit between the San Jose and Montreal routers fails to come up
  - Both ends indicate a `Down` status at the Physical Layer
- You cannot view the configuration directly
- What is wrong?
  - What CLI commands and fault analysis steps can help narrow down a possible cause?



### Interface Case Study C: Background

The slide sets the stage for a third interface troubleshooting case study. We begin with a general description of the problem, which, in this case, indicates that the Ethernet link between the San Jose and Montreal devices appears to be in `Down` state.

Note that in this example, you are not permitted to view the interface-related configuration at either end.

## Interface Case Study C (2 of 4)

### ■ Sample course of action:

1. Check if the San Jose Gigabit Ethernet link has successfully completed the autonegotiation with the Montreal router
2. Replace the PIC
3. Change the fiber

### Interface Case Study C: Course of Action

Several reasons might exist for an Ethernet links to be in the `Down` state. The reasons include:

- Misconfigured Ethernet parameters or incomplete auth-negotiations;
- Faulty PIC; or
- Faulty interface wiring.

In JUNOS Software, Ethernet link autonegotiation is enabled by default.

## Interface Case Study C (3 of 4)

- Check if the San Jose Gigabit Ethernet link has successfully completed the autonegotiation with the Montreal router

```

user@San_Jose> show interfaces ge-0/1/0 media
Physical interface: ge-0/1/0, Enabled, Physical link is Down
Interface index: 153, SNMP ifIndex: 135
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
Auto-negotiation: Enabled, Remote fault: Online
Device flags      : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
Link flags       : None
CoS queues       : 8 supported, 8 maximum usable queues
Schedulers      : 128
Current address:  00:19:e2:25:b2:7e, Hardware address: 00:19:e2:25:b2:7e
Last flapped    : 2009-03-04 17:15:05 PST (2d 16:36 ago)
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
Active alarms   : LINK
Active defects  : LINK
MAC statistics:
  Input bytes: 0, Input packets: 0, Output bytes: 0, Output packets: 0
Filter statistics:
  Filtered packets: 0, Padded packets: 0, Output packet errors: 0
Autonegotiation information:
  Negotiation status: Incomplete
    
```

Physical link is Down at San Jose

Autonegotiation is enabled by default in JUNOS Software

Autonegotiation is incomplete

### Interface Case Study C: Check the Ethernet Interface Details

Suspecting a Layer 2 configuration mismatch, you begin by displaying media interface information for the ge-0/1/0 interfaces at the San Jose and Montreal stations. The **show** command indicates that the interface's autonegotiation is enabled, which is the default behavior. Also, you see that the negotiation status is `Incomplete`.

## Interface Case Study C (4 of 4)

### ■ Solution:

- Ensure that both San Jose and Montreal devices can successfully autonegotiate the connection parameters, or explicitly configure both devices to operate in the desired mode with autonegotiation disabled

### Interface Case Study C: Solution

To fix the problem, you ensure that both local and remote devices can successfully autonegotiate the connection parameters. If that does not solve the problem, you explicitly configure San Jose and Montreal stations to operate in the desired mode with autonegotiation disabled. In our case the autonegotiation succeeded. Now the interface status is Up, as illustrated:

```
user@San_Jose> show interfaces ge-0/1/0 media
Physical interface: ge-0/1/0, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error:
None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
. . .
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
    Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote
fault: OK
  Local resolution:
    Flow control: Symmetric, Remote fault: Link OK
```



## Summary

- In this chapter, we:
  - Described physical and logical interface properties that require configuration
  - Deactivated and disabled interfaces
  - Configured loopbacks and BERT testing
  - Used operational mode commands to monitor and troubleshoot a variety of interface types

### This Chapter Discussed:

- Physical and logical interface properties;
- Deactivating and disabling interfaces;
- Configuring loopbacks and BERT tests; and
- Using operational mode commands to monitor and troubleshoot a variety of interfaces and media types.

## Review Questions

1. What is the difference between deactivating and disabling an interface?
2. How do you display the results of interface BERT testing?
3. What interface settings do you need to check when troubleshooting T3 or E3 interface problems?
4. What is the rationale behind using the `monitor interface interface-name` command?

## Review Questions

- 1.
- 2.
- 3.
- 4.

## Lab 3: Interface Troubleshooting

- Use the CLI to monitor and troubleshoot the operation of a variety of interface types.

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 5-97

### Lab 3: Interface Troubleshooting

The slide shows the objective for this lab.

Not for Reproduction



# **Troubleshooting JUNOS Platforms**

## **Chapter 6: JTAC Processes, Guidelines, and Support Resources**

Not for Reproduction

## Chapter Objectives

- After successfully completing this chapter, you will be able to:
  - Follow recommended procedures to open a support case
  - Access support resources
  - Use the Customer Support Center
  - Describe case management procedures
  - Research a problem using the KB or PR databases
  - Use the I2J tool
  - Download JUNOS Software and technical documentation
  - Use FTP to transfer files to JTAC

### This Chapter Discusses:

- Support entitlement and opening a support case;
- Support resources;
- The Customer Support Center (CSC);
- Case management;
- The Juniper Networks Technical Assistance Center (JTAC) Knowledge Base (KB) and Problem Report (PR) search tools;
- The IOS-to-JUNOS Conversion (I2J) tool;
- Downloading software and technical documentation; and
- Transferring files to JTAC using FTP.

## Agenda: JTAC Processes, Guidelines, and Support Resources

- Opening a Support Case
- Support Services
- How to Use FTP to Send Files to JTAC

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 6-3

### Opening a Support Case

The slide lists the topics we cover in this chapter. We discuss the highlighted topic first.

## Support Entitlement

- JTAC support is offered only to customers with a valid maintenance contract
  - A chassis serial number is necessary when opening a case so JTAC can determine support status
  - Details at: <http://www.juniper.net/support/requesting-support.html>

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 6-4

## Support Entitlement

Juniper Networks offers support only to customers with valid maintenance contracts. Therefore, you must provide a chassis serial number when opening a support case so that JTAC can verify your support status. Use a **show chassis hardware** command to obtain your chassis serial number:

```
user@host> show chassis hardware
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis              A9697
Midplane             REV 07   710-009120   RB8559         M320 Midplane
FPM GBUS             REV 05   710-005928   DE9089         M320 Board
FPM Display          REV 05   710-009351   DD4415         M320 FPM Display
CIP                  REV 05   710-005926   DC5728         M320 CIP
PEM 0                Rev 05   740-009149   SK52367        AC Power Entry Module
PEM 1                Rev 05   740-009149   SK52379        AC Power Entry Module
PEM 2                Rev 05   740-009149   SH52307        AC Power Entry Module
PEM 3                Rev 05   740-009149   SK52363        AC Power Entry Module
Routing Engine 0     REV 07   740-014082   1000702757    RE-A-2000
Routing Engine 1     REV 07   740-014082   1000699981    RE-A-2000
CB 0                 REV 12   710-009115   DG4583         M320 Control Board
CB 1                 REV 12   710-009115   DG4576         M320 Control Board
FPC 0                REV 04   710-013518   DG1608         M320 E2-FPC Type 3
  CPU                 REV 04   710-013562   DF5825         M320 FPC CPU
  PIC 0              REV 20   750-007141   DE4219         10x 1GE(LAN), 1000 BASE
. . .
```



## Opening a Case

- Can open cases several ways:
  - On the Web using Case Manager at: <https://www.juniper.net/cm/index.jsp>
  - By phone:
    - 1-888-314-JTAC (U.S., Canada, and Mexico customers)
    - Phone: 1-408-745-9500 (all other customers)
  - For Priority 1 cases, contact JTAC by phone, even when you opened the case with e-mail or Case Manager
- Documentation:
  - Illustrate all symptoms with relevant command output
  - The **request support information** command provides most of the information needed for most problems

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 6-5

### Opening a Case

Customers can open support cases using the Case Manager (from within the CSC), or over the telephone. Note that you should always follow up a Priority 1 case with a phone call to JTAC, even when you opened the case using the Case Manager or e-mail.

### Documenting the Problem

When you open a case, you should provide documentation of the problem or symptom by capturing the output of relevant CLI commands. You might consider also capturing the output of a **request support information** command; this command is actually a macro that executes a large number of operational mode commands. While the output is lengthy, adding this information might prevent the need to follow up with additional command output once your case is under analysis.

## Case Management

- Case management details are provided at:  
<https://www.juniper.net/support/guidelines.html>
  - Definition of case priorities, escalation, support levels, Return Materials Authorization handling, standard warranty, and inspection procedures
    - Case priority is mutually agreed upon between customer and JTAC:  
Priority 1 = Critical, Priority 2 = High, Priority 3 = Medium,  
Priority 4 = Low
  - Escalation management defines systematic escalation to ensure that the appropriate resources within Juniper Networks are used to resolve outstanding technical problems as efficiently as possible

### Case Management

Case management details and procedures are provided at <https://www.juniper.net/support/guidelines.html>. This document defines case priority levels, escalation management, warranty information, and so forth.

When you contact JTAC, a member of our support staff will work with you in assigning mutually agreeable priority levels to your problem that will be reflected in the support case opened on your behalf. The following are the priority levels:

- *Priority 1 (Critical):* Catastrophic impact to business operations. Examples of Priority 1 issues include network or system outages that cause customers to experience a total loss of service.
- *Priority 2 (High):* Significant impact to business operations. Examples of Priority 2 issues include network or system events that cause intermittent impact to end customers.
- *Priority 3 (Medium):* Limited impact to business operations. Examples of Priority 3 issues include network events that result in only limited impact to end customers.
- *Priority 4 (Low):* No impact to business operations. Examples of Priority 4 issues include information requests.

*Continued on next page.*

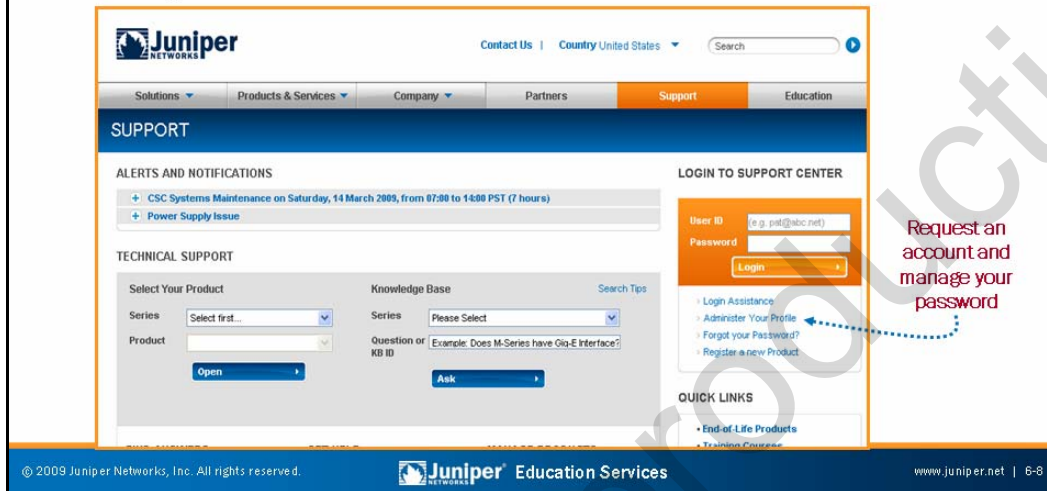
### Case Management (contd.)

Juniper Networks offers systematic escalation management to customers with current service agreements. This escalation management ensures that the appropriate resources within Juniper Networks are used to resolve outstanding technical problems as efficiently as possible.

Not for Reproduction

## Gaining Access to Support Services

- Request a CSC login at:  
<https://www.juniper.net/support/csc/>
  - Requires a valid maintenance contract



### Access Support Services

You must have a valid CSC login to access Juniper Networks support services over the Web. Point your browser to <https://www.juniper.net/support/csc/> to display the CSC login screen. Note that below the login, you can click links to request an account or to manage the passwords associated with your existing CSC account.

## Agenda: JTAC Processes, Guidelines, and Support Resources

- Opening a Support Case
- Support Services
- How to Use FTP to Send Files to JTAC

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 6-9

### Support Services

The slide highlights the topics we discuss next.

## The Customer Support Center

- The CSC provides easy access to support services and technical information:

The screenshot shows the Juniper Customer Support Center (CSC) website. The page is titled "SUPPORT" and features a navigation menu with "Solutions", "Products & Services", "Company", "Partners", "Support", and "Education". The "Support" menu is highlighted. The main content area is divided into several sections:

- ALERTS AND NOTIFICATIONS:** Displays a message about "CSC Systems Maintenance on Saturday, 14 March 2009, from 07:00 to 14:00 PST (7 hours)".
- TECHNICAL SUPPORT:** Includes a "Select Your Product" section with dropdown menus for "Series" and "Product", and a "Knowledge Base" section with a "Search Tips" field and a "Question or KB ID" field.
- QUICK LINKS:** Lists links such as "End-of-Life Products", "Training Courses", "Technical Documentation", "RMA Procedure", "Contact Support", and "Guidelines and Policies".
- RELATED TOPICS:** Lists "PDE Brochures", "CSC Overview", "JTAC User Guide", "Education/Training", "Certification", "Class Registration", "Order JUNOS Docs", and "Services".
- AWARDS AND RECOGNITION:** A section at the bottom right.

Annotations with arrows point to various features:

- "Search the Knowledge Base and Technical Bulletins" points to the Knowledge Base search field.
- "Download software and access technical information" points to the "Download Software" link in the "MANAGE PRODUCTS" section.
- "Subscribe to support services" points to the "Subscribe to Email Alerts" link in the "TECHNICAL BULLETINS" section.
- "Create and manage cases" points to the "Create a Case" link in the "GET HELP" section.

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | 6-10

### The Customer Support Center

The slide displays the top of the main CSC welcome page. From this page you can easily link to case management, technical research, software downloads, and so forth.

## Open a Case with Case Manager

- The Case Manager application provides a simple Web-based interface for creation and management of cases

The screenshot shows the 'Support' page in the Case Manager application. The page title is 'Support' and the breadcrumb trail is 'Home > Support > CEC > Case M'. The main heading is 'CASE MANAGEMENT' and the sub-heading is 'CASE CREATION - CASE TYPE SELECTION'. The page contains the following text and form elements:

For SBR and OAC (formerly FUNK) customers, you can look up your product serial number in the License Management System (see link below) using your license key for the search. Click here to [search for your license](#).

**Please select the case type you wish to create**

Choose Series / Serial number for your product (required for technical support cases, optional for Customer Care cases)

Series:

System Serial Number:

Click here to see your [products eligible for support](#).

Choose Technical Support for any support or RMA request.

Technical Support Case

Choose Customer Care for registration, license keys, login or other non-technical issues, also for J-ASSURE Proactive Services requests.

Customer Care Case

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | 6-11

### Using Case Manager

Many customers prefer to open and manage their cases with the point-and-click ease offered by the Case Manager tool. In addition to opening a case, the Case Manager provides a handy way of tracking the status of your open cases, Return Materials Authorization status, and so forth.

## Research a Problem

- Access to the JTAC Knowledge Base, PR search, technical bulletins, and other invaluable information is located at:
  - <https://www.juniper.net/customers/csc/research/index.jsp>

### Research Problems

When you access the CSC, you are privy to a wealth of technical support information in the form of the JTAC Knowledge Base, the PR (bugs) database, technical bulletins, and white papers that provide configuration examples and technology primers. You can access the research area of the CSC at <https://www.juniper.net/customers/csc/research/index.jsp>.



## The I2J Tool

- The I2J tool converts IOS configuration into a JUNOS Software configuration
  - Available at <https://i2j.juniper.net/release/index.jsp>
  - Not always perfect; you must inspect all output and heed all warnings!

The screenshot shows the I2J tool interface. On the left, there is a text area titled 'Paste a complete IOS config file' containing the following IOS configuration:

```
interface Ethernet0
ip address 200.2.2.2 255.255.255.0
ipx network E2
ipx encapsulation SAP
no mop enabled
no shutdown
```

Below the text area, there is an option to 'Upload an IOS config file' with a 'Browse...' button. Underneath, there are several checked options under 'Select option(s)':

- Output IOS lines that converted properly
- Output verbose IOS comments
- Consolidate firewall terms
- Consolidate policy terms
- Use my configuration for future I2J enhancements ([privacy information](#))

At the bottom of the tool interface, there are 'Reset Form' and 'Translate' buttons. On the right side of the screenshot, the resulting JUNOS configuration is shown:

```
interfaces {
  /* Created from IOS Interface: ethernet0 */
  fe-0/0/0 {
    unit 0 {
      family inet {
        address 200.2.2/24;
      }
    }
  }
}
```

A red arrow points from the 'Translate' button area to the resulting JUNOS configuration, which is labeled 'The result!'.

## The I2J Tool

The I2J tool is a utility that converts Cisco IOS configurations to the equivalent JUNOS Software configuration. The slide shows an example of the I2J tool in use to convert the Ethernet0 portion of a Cisco IOS configuration.

As noted, when accessing the tool, you must manually inspect the results of the I2J tool before using them. You must also carefully consider any warnings generated during the conversion process. The following capture shows the warnings generated during the sample Ethernet0 interface configuration conversion shown on the slide. The warnings regarding unsupported protocols, like AppleTalk and IPX, should not come as a great surprise:

**Lines that could not be converted are in red. Lines with warnings or comments are in blue.**

**Lines that have previously elicited an error or warning are in magenta and their error or warning is suppressed.**

```
FPC / PIC / Port numbers MUST ALWAYS be changed to match your Juniper Networks hardware.
1:interface Ethernet0 This interface was converted to a FAST ethernet which supports 100Mb ONLY
2: ip address 200.2.2.2 255.255.255.0
3: ipx network E2 Line not recognized by I2J
4: ipx encapsulation SAP Line not recognized by I2J
5: no mop enabled Line not recognized by I2J
6: no shutdown Line not recognized by I2J
[Interface Conversion Messages]
IOS interfaces without explicit clocking will have internal clocking in JUNOS by default
```

## Software Download

- Requires a valid CSC login
- Domestic versions have strong encryption capabilities and might require signed agreement
  - <https://www.juniper.net/customers/csc/software/index.jsp>

**DOWNLOAD SOFTWARE**

**NETSCREEN FIREWALL/SSG (SCREENOS)** [See all »](#)  
 ScreenOS runs on our popular Firewall/SEC VPN products such as the NS-200 and SSO-140 appliances.  
 • [IDP](#) • [NSMGlobal Pro](#)  
 • [NetScreen Remote VPN Client](#) • [ScreenOS](#)  
(secure & regular)

**CARRIER AAA**  
 Juniper's Carrier AAA software portfolio encompassing SBR Carrier, SBR SPE, SBR HA, SBR SIM, SBR MIM, SBR SLM, and IMS AAA.  
 • [Carrier AAA](#)

**SECURE ACCESS (SA) SSL VPN (IVE OS)**  
 SSL VPN (IVE OS) provides powerful remote-access solutions for the SA-Series products.  
 • [IVE OS](#) • [ESAP \(IVE\)](#)

**JUNOS**  
 JUNOS software for EX-series, J-series, M-series, MX-series SRX-series, and T-series routers delivers high levels of resiliency, stability and uptime.  
 • [JUNOS Canada and U.S.](#) • [JUNOS-FIPS](#)  
Encryption agreement Encryption agreement  
 • [JUNOS Worldwide](#)

**JUNOSE**  
 JUNOSe is the network operating system for the ERX products such as the new E320 Edge Router.  
 • [JUNOSE](#)  
Encryption agreement

**OTHER**  
 Application Acceleration, WAN Optimization and Circuit-to-Packet Solutions.  
 • [CTP-series](#) • [MXOS](#)  
(Formerly Accem) (including MX CMS)  
 • [DXOS](#)

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services [www.juniper.net](http://www.juniper.net) | 6-14

### Software Download

Customers with valid support contracts can download the latest version of JUNOS Software from <https://www.juniper.net/customers/csc/software/index.jsp>.

### Encryption = Weapon

The use of software with strong encryption might be subject to arms-related export restrictions. You might have to complete an encryption agreement form before being allowed to download the US domestic version of JUNOS Software.

## Documentation Download

- CSC login not necessary
  - <http://www.juniper.net/techpubs/>

The screenshot displays the 'TECHNICAL DOCUMENTATION' page on the Juniper website. The page is organized into several sections:

- Left Sidebar:** Contains navigation links such as 'Download Software', 'Research a Problem', 'Case Management', 'Contract & Product Management', 'Technical Documentation', 'Advanced Search', 'Enterprise MBs', 'Glossary', 'File Format Help', 'CD-ROM Documentation', 'Documentation Feedback Form', 'End-of-Life', 'End-of-Life Products', 'Contact Support', 'Guidelines and Policies', and 'Security Resources'.
- Main Content Area:**
  - TECHNICAL DOCUMENTATION:** A central heading.
  - ROUTING:** Lists platforms like BX-series, CTP-series, E-series, J-series (legacy services), JCS Platform, M-series, MX-series, and T-series.
  - SWITCHING:** Lists EX-series Platforms.
  - IDENTITY, POLICY, AND CONTROL:** Lists products like C-series Controllers, IMS AAA Server, Network and Security Manager, Odyssey Access Client, Security Threat Response Manager (STRM), Session and Resource Control Portfolio (SRC), Steel-Bebed RADIUS, and Unified Access Control (UAC).
  - SECURITY:** Lists products like Integrated Firewall/Peec VPN, Intrusion Prevention (IP) Software, J-series, NetScreen Remote Software, ScreenOS, Security Threat Response Manager (STRM), SRX-series, and SSL VPN (IVE OS).
  - MANAGEMENT SOFTWARE:** Lists Advanced Insight Solutions (AIS), JUNOScope, NMC-RX, and Service Deployment System (SDX).
  - APPLICATION AND ACCELERATION:** Lists DX Data Center Acceleration Platform, VXC WAN Application Acceleration Platforms, and VVX WAN Application.
- Right Sidebar:** Includes 'Related Topics' (File Format Help, Juniper Books, M. & T-series Study Guides) and 'What's New' (EX-series documentation, JUNOS 9.4 Documentation, JUNOS 9.0 Documentation).

At the bottom of the page, there is a footer with the copyright notice '© 2009 Juniper Networks, Inc. All rights reserved.', the Juniper logo, 'Education Services', and the URL 'www.juniper.net | 6-15'.

### Technical Documentation Download

You can download technical documentation for all JUNOS products from <http://www.juniper.net/techpubs/>. You do not need a CSC login.

## Agenda: JTAC Processes, Guidelines, and Support Resources

- Opening a Support Case
- Support Services
- How to Use FTP to Send Files to JTAC

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 6-16

### How to Use FTP to Send Files to JTAC

The slide highlights the topic we discuss next.

## Uploading Files with FTP

- Use FTP when transferring files larger than 10 MB
  - Attach smaller files directly to case using Case Manager
  - FTP procedures are specified at:  
[https://www.juniper.net/customers/csc/help/upload\\_files.jsp](https://www.juniper.net/customers/csc/help/upload_files.jsp)
    - Log in to Juniper Networks anonymous FTP site at `ftp.juniper.net` with e-mail address as your password
    - Change into the `/pub/incoming` directory
    - Create a directory using your case number as its name
    - Change into the directory you created; confirm with a `pwd` command
    - Ensure that the local directory contains the file (or files) you want to upload; use a `lpwd` command to determine the local directory
    - Enable binary mode transfer with an `image` or `binary` command
    - Enable hash-mark printing for status indication with the `hash` command
    - Transfer the file using a `put` or `mput` command
    - When all files transfer, break the FTP connection with the `quit` command

© 2009 Juniper Networks

Juniper Networks Education Services

www.juniper.net | 6-17

## Transferring Core Files to Juniper Networks

You should always submit core files to JTAC for fault analysis. The slide outlines the recommended procedures for transferring core files to JTAC:

1. First, send an e-mail to `support@juniper.net` to open a support case and obtain a case number.
2. Escape to a root shell and change to the directory containing the core file.
3. Rename (or copy) the file using a name in the form of `case_number-core-sequence_number`.
4. Although not strictly necessary, we recommend that you apply the `chmod` command to the core file with `444` to ensure that all users (root, owner, and other) have read permissions for the file.
5. In some cases, the core file will already be compressed—indicated by a `.tgz` or `.gz` file extension. If the file is not already compressed, you should compress the file to reduce transfer and storage requirements. This compression is especially important when dealing with the `vmcore.0` file associated with a kernel crash because this memory image file can be quite large.

*Continued on next page.*

### Transferring Core Files to Juniper Networks (contd.)

6. Log into the Juniper Networks anonymous FTP site at `ftp://ftp.juniper.net`, and change into the `/pub/incoming` directory.
7. Ensure that your FTP client is set for a binary transfer. In many cases the client defaults to the correct transfer type. Issue a **type** command to confirm the current transfer setting and use the **image** or **binary** command to enable binary transfer mode as needed.
8. Upload the compressed and renamed core or memory image file using a **put** or **mput** command.

Not for Reproduction

## Sample FTP Submission (1 of 2)

```

user@host> start shell
% su
Password:
root@host% cd /var/tmp
root@host% ls -la snmpd.core-tarball.0.tgz
-rw----- 1 root  field 137334 Mar 17 15:15 snmpd.core-tarball.0.tgz
root@host% chmod 444 snmpd.core-tarball.0.tgz
root@host% ls -la snmpd.core-tarball.0.tgz
-r--r--r-- 1 root  field 137334 Mar 17 15:15 snmpd.core-tarball.0.tgz
root@host% mv snmpd.core-tarball.0.tgz 2004-12345-snmpdcore-01.tgz


root@host% ftp ftp.juniper.net
Connected to ftp.juniper.net.
220 colo-ftp.juniper.net FTP server (Version 6.00LS) ready.
Name (ftp.juniper.net:lab): anonymous
331 Guest login ok, send your email address as password.
Password:
230 Guest login ok, access restrictions apply.
Remote system type is UNIX.
    
```

Only root can read and write by default

Read permissions for all

File renamed with case number

Anonymous FTP session established

© 2005 Juniper Networks, Inc. All rights reserved.  Juniper Education Services www.juniper.net | 6-19

### Transferring Core Files to Juniper Networks: Part 1

The slide illustrates the recommended procedures for transferring a core file to JTAC using FTP. The example begins with a user escaping to a root shell and then changing into the `/var/tmp` directory. We assume that the user already has knowledge that a process has left a core file in this directory.

In this example, the SNMP process has left a core file with context in the form of a compressed `.tgz` archive. Manual compression of the file (using `gzip file-name`) is not necessary in this case.

The user modified the original permissions on the core file with a `chmod 444` command to ensure that all users will have the read access required to perform stack trace analysis on the file.

The user renamed the file to reflect the assigned case number; in this example the sequence number is set to 01 to indicate the first core file submitted in conjunction with this case. Note that when renaming the file you should preserve any `.tgz` or `.gz` extensions so that the recipient knows whether the file is compressed, a tar archive, or both.

The slide ends with the successful establishment of an anonymous FTP session to the Juniper Networks FTP site.



## Sample FTP Submission (2 of 2)

```

. . .
Using binary mode to transfer files. <..... Binary transfer in effect by default in this case
ftp> cd pub/incoming <..... Change to pub/incoming directory
250 CWD command successful.
ftp> mkdir 2004-12345 <..... Create a directory with your case number
257 MKD command successful.
ftp> CD 2004-12345 <..... Change to new directory
250 CWD command successful.
ftp> mput 2004* <..... mput command allows use of wildcards
mput 2004-12345-snmpdcore-01.tgz? y <..... Confirm upload
200 PORT command successful.
150 Opening BINARY mode data connection for '2004-12345-snmpdcore-01.tgz'.
100% 134 KB 00:00 ETA
226 Transfer complete.
137334 bytes sent in 0.19 seconds (718.68 KB/s)
ftp> quit <..... Break FTP connection
221 Goodbye.

```

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | 6-20

### Transferring Core Files to Juniper Networks: Part 2

This slide continues the example started on the previous slide. Note that the user's FTP client defaults to a binary transfer mode, which is confirmed in the FTP server's output. Thus, the user does not have to switch into binary/image mode manually in this example.

After connecting to the Juniper Networks FTP site, the user changes into the `/pub/incoming` directory and uses the `mkdir` command to create a new directory that is named according to the assigned case number. The user then changes into this directory and makes use of the `mput` command to upload the file. The `mput` command supports wildcards (meta characters), which makes the exact specification of the file name unnecessary. The user confirms that the `2004-12345-snmpdcore-01.tgz` file should transfer, and the FTP client confirms successful upload. After the file transfers, the user breaks the FTP session by entering the `quit` command.



## Summary

- In this chapter, we:
  - Followed recommended procedures to open a support case
  - Accessed support resources
  - Used the Customer Support Center
  - Described case management procedures
  - Researched a problem using the KB or PR database
  - Used the I2J tool
  - Downloaded JUNOS Software and technical documentation
  - Used FTP to transfer files to JTAC

### This Chapter Discussed:

- Support entitlement and opening a support case;
- Support resources;
- The Customer Support Center;
- Case management;
- JTAC KB and PRs search tools;
- The I2J tool;
- Downloading software and technical documentation; and
- Transferring files to JTAC using FTP.

## Review Questions

1. What information do you need to open a support case?
2. What aspects of the CSC can help you research a problem?
3. Describe the purpose of the I2J tool.
4. When should you use FTP to transfer files to JTAC?

## Review Questions

- 1.
- 2.
- 3.
- 4.



# **Troubleshooting JUNOS Platforms**

## **Appendix A: JUNOS Platform Details**

Not for Reproduction

## Appendix Objectives

- After successfully completing this appendix, you will be able to:
  - List the primary components and characteristics of multiservice edge routers (M Series and T Series)
  - List the primary components and characteristics of Ethernet services routers and switches (MX Series and EX Series)
  - List the primary components and characteristics of security services gateways (SRX Series)
  - List end-of-life products

### This Appendix Discusses:

- Primary components and characteristics of multiservice routers (M Series and T Series);
- Primary components and characteristics of Ethernet services routers and switches (MX Series and EX Series);
- Primary components and characteristics of security services gateways (SRX Series); and
- End-of-life products.

## Agenda: JUNOS Platform Details

- Primary Components and Characteristics of Multiservice Edge Routers (M Series and T Series)
- Primary Components and Characteristics of Ethernet Services Routers and Switches (MX Series and EX Series)
- Primary Components and Characteristics of Security Services Gateways (SRX Series)
- End-of-Life Products

### Components and Characteristics of Multiservice Routers

The slide highlights the topics we cover in this appendix. We discuss the highlighted topic first.

## M7i and M10i Multiservice Edge Routers

- M7i and M10i routers:
  - Used for edge applications involving the aggregation of dedicated access circuits
  - Used for core applications in locations where space and power are at a premium
    - M7i router features integrated services and built-in Ethernet PICs
    - M10i router features RE and CFEB redundancy



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | A-4

### M7i and M10i Multiservice Edge Router Overview

The M7i and M10i Multiservice Edge Routers build on the proven performance and success of the M5 and M10 platforms by offering integrated IP services and Ethernet interface support, or Routing Engine (RE) and Packet Forwarding Engine (PFE) redundancy, in the case of the M7i and M10i, respectively.

The M7i and M10i forward packets at an aggregate throughput rate of 6.4 Gbps and 12.8 Gbps, respectively.

The M7i is the Juniper Networks high-performance customer premise equipment (CPE) offering, targeting campuses and large offices needing very secure, dependable, high-speed WAN connectivity. The M7i delivers unprecedented processing power for secure connectivity and the simultaneous support of multiple services in a single platform. Combined with modular flexibility and easy-to-manage JUNOS Software, the M7i is the choice for consolidating multiple platforms and bringing carrier-class IP routing and security capabilities to the enterprise.

The M7i features a Fixed Interface Card (FIC) for Ethernet support and an integral Smart IP Services PIC. The M7i's four PIC slots are open for additional interface support. You can order the M7i Compact Forwarding Engine Board (CFEB) with both an integral Smart IP Services PIC and an Adaptive Services PIC (AS PIC) to support additional service offerings such as Network Address Translation (NAT). Note that the integral AS PIC is not a field-replaceable unit (FRU); to add AS PIC functionality, order a new CFEB or install an AS PIC in one of the M7i's four PIC slots.


*Continued on next page.*

**M7i and M10i Router Overview (contd.)**

The M10i is targeted at dedicated access and core applications where space, power consumption, and high availability (HA) are wanted. The M10i does not include internal service PICs, but the platform does offer RE and CFEB redundancy.

The M7i is only two rack units (RUs) in height (2 inches or 8.9 cm), which allows as many as 24 M7i units in a single 19-inch equipment rack! At 4.9 RUs (8.7 inches or 22.1 cm), you can place as many as nine M10i units into a single rack.

Not for Reproduction



## M7i and M10i Hardware Components

The hardware components for the M7i and M10i are the following:

- *Sheet metal chassis.*
- *Two power supplies (AC or DC):* The maximum chassis power is 274 watts for the M7i and 427.2 watts for the M10i with a full complement of PICs (8) and full redundancy. Power supplies can be either AC or DC (but not both simultaneously); when two are installed, power is load-balanced.
- *Fan assemblies.*
- *Routing Engines:* The RE maintains the routing tables and controls the routing protocols, as well as the JUNOS Software processes that control the router's interfaces, 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 PFE. The M10i supports redundant REs.
- *Compact Forwarding Engine Boards:* The CFEB provides PFE functionality. The M7i and M10i PFE supports the same application-specific integrated circuit (ASIC) set used in other M Series routers, with the same support for enhanced features like filtering and accounting. The M10i supports redundant CFEBs.

*Continued on next page.*




### M7i and M10i Hardware Components (contd.)

- *Four (M7i) or eight (M10i) PICs:* In addition, the M7i supports an integral IP Tunneling Services PIC and an optional Adaptive Services PIC. Inclusion of the latter provides enhanced services, such as stateful firewall filters and NAT. When desired, you can equip the M10i with services PICs using one or more of its PIC slots.


Not for Reproduction

## M7i and M10i Craft Interface

- The M7i and M10i Craft Interface is on the RE at the front of the chassis
  - RE port (console port, AUX/modem port, and a single Fast Ethernet port)



- Alarm LEDs and PIC online and offline buttons on the M7i platform's FIC (shown) or the M10i platform's HCM



### M7i and M10i Craft Interface

The M7i and M10i Craft Interface provides connections for access to the local RE. You can achieve access using any of the following:

- *Ethernet management port:* Connects the RE to an out-of-band management network.
- *Console port:* Connects a system console to the RE with EIA/TIA-232 serial asynchronous cable. Use the system console to access the CLI to configure the attached router. This port is active by default.
- *Auxiliary port:* Connects a laptop or modem to the RE with EIA/TIA-232 serial asynchronous cable. You can also use this port to access the command-line interface (CLI). This port is disabled by default and you must activate it with a set **system ports auxiliary type terminal-type** command.

The M7i platform's Fixed Interface Card (FIC) or the M10i platform's High-Availability Chassis Manager (HCM) card supports PIC online and offline buttons and alarm LEDs. You should take PICs offline before removing them from the chassis using the respective PIC online and offline button located on the FIC or HCM. After inserting a PIC into the chassis, press and hold the online and offline button to activate power to that PIC. When installed, the Adaptive Services PIC shares PIC slot 1/2/0 with the internal Tunnel Services PIC. The two PICs contend for the approximate 1 Gbps of bandwidth associated with the M7i platform's 1/2/0 PIC slot.

## M7i and M10i Cooling

- M7i and M10i cool from side to side
  - M7i:
    - Single fan tray with redundant fans
  - M10i:
    - Redundant fan trays

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-9

### Side-to-Side Cooling

The M7i, M10i, and M20 incorporate side-to-side cooling using one to four fan trays. The M20's cooling system consists of the following subsystems:

- *Front cooling subsystem:* Three front fan trays that cool the Flexible PIC Concentrators (FPCs) and the System and Switch Board (SSB). These fan trays are on the left front side of the chassis.
- *Rear cooling subsystem:* One rear fan tray that cools the RE. This fan tray is immediately to the right of the RE.
- *Power supply integrated fan:* A built-in fan that cools each power supply.

In the case of the M20, the four fan trays work together to provide side-to-side cooling. The fan trays plug directly into the router midplane. Each front fan tray is a single FRU that contains three fans. The rear fan tray is a FRU that contains two fans. Both front and rear fan trays are hot-swappable.

The M7i implements a single fan tray along the side of the chassis for side-to-side cooling. The fan tray is a single unit containing four fans. It is hot-removable and hot-insertable, and it connects directly to the router midplane. The failure of any one fan does not effect the operation of the remaining fans, which can continue to run indefinitely in the face of a single fan failure.

*Continued on next page.*

### Side-to-Side Cooling (contd.)

The M10i supports redundant fan trays. Each fan tray is a single unit containing eight individually fault-tolerant fans. The left fan tray must be in place for proper cooling at all times. The right fan tray provides additional cooling and redundancy.

Juniper Networks M Series routers can operate with missing or failed fans. The system can measure the presence and rotational speed of all fan components. When chassis temperatures rise, fan speed increases, and a yellow alarm generates. If the temperature continues to rise, a red alarm generates; the system will shut itself down if the high-temperature condition persists.

Not for Reproduction

## M40e Multiservice Edge Router

- M40e router delivers:
  - Full hardware redundancy
  - The bandwidth required to grow networks from DS-0 to OC48c (STM16) speeds cost efficiently
  - OC48c (STM16) PICs (FPC2)



© 2009 Juniper Networks, Inc. All rights reserved.


 Juniper Education Services

www.juniper.net | A-11

### M40e Multiservice Edge Router Overview

Like the M40, the M40e is specifically for the specialized needs of high-growth Internet backbone providers, with its 40+ Gbps forwarding rate and enhanced redundancy.


The M40e is primarily for the edge aggregation market. The M40e provides RE and PFE redundancy and the ability to hot-swap PICs. The M40e can support true OC48C (STM16) PICs when equipped with FPC2. In contrast, an OC48C interface is a quad-wide PIC and FPC combination on the M40 and M20 platforms. Note that neither the M40 or M20 supports the hot-swapping of PICs.



## FPC1 and FPC2

The M40e supports two types of FPCs. The platform can operate with any combination of FPC1s and FPC2s installed. The best way to determine the FPC type is by the particulars of the PIC online and offline buttons:

- **FPC1:** Accommodates up to four PICs, including single-port OC12 (STM4) and Gigabit Ethernet interfaces. The PIC slots number top to bottom from 0 (zero) to 3.
- **FPC2:** Accommodates one higher-speed PIC, such as an OC48 (STM16) or a 4-port, single-wide Gigabit Ethernet interface. The PIC installs into the uppermost slot on the FPC—number 0 (zero)—and the other three slots have covers.



## M120 Multiservice Edge Router Major Components

The major components of the M120 from the front view are the following:

- **ESD point:** Consists of two electrostatic discharge (ESD) points (banana plug receptacles)—one front and one rear.
- **Craft Interface:** Allows you to view status and troubleshooting information at a glance and to perform many system control functions. It is hot-insertable and hot-removable. The Craft Interface is on the front of the router above the upper fan tray
- **FPCs:** The FPC slots on the router are two-PICs wide and two PICs tall, allowing four Type 1 or Type 2 PICs and one Type 3 PIC to share power circuits and a CPU and still fit in a quarter rack. Up to four FPCs install vertically in the front of the router. Additionally, the M120 has slots for two compact FPCs (cFPCs). These cFPCs contain an entire PFE and Type 3 PIC in a PIC-sized form factor.
- **Fan tray:** All chassis fans are redundant. In a failure situation, the Control Board is responsible for increasing the speed of the remaining fans to provide the necessary cooling.

*Continued on next page.*

## M120 Multiservice Edge Router Major Components (contd.)


The major components of the M120 from the rear view include the following:

- *Routing Engines:* The RE is an Intel-based Peripheral Component Interconnect (PCI) platform that runs JUNOS Software. Software processes that run on the RE maintain the routing tables, manage the routing protocols used on the router, control the router interfaces, control some chassis components, and provide the interface for system management and user access to the router. You can install one or two REs in the router. The REs install into the rear of the chassis, directly into the Control Board, in vertical slots labeled CBO and CB1.
- *Control Boards (CBs):* Each CB works with an installed RE to provide control and monitoring functions for the router. These functions include determining RE mastership, controlling power and reset for the other router components, connecting the FEBs, monitoring and controlling fan speed, and monitoring system status. You can install one or two CBs in the router.
- *Forwarding Engine Boards (FEBs):* The FEBs provide route lookup and forwarding functions from the PICs and cFPCs. The midplane architecture allows any FEB to carry traffic for any FPC. Depending on the bandwidth and streams of the FEB, it can carry the traffic of multiple FPCs. To provide redundancy, the FPC-FEB connections combine with 20 Gbps of bandwidth between each pair of cards. System software configuration determines which FPCs a given FEB supports and configures switches on the FEB or FPC accordingly. In addition, in the event of a FEB failure, a standby FEB can quickly take over packet forwarding.
- *Power Entry Modules (PEMs):* The M120 is configurable with either two AC power supplies or two DC power supplies. If one power supply fails or you remove it, the remaining power supply instantly assumes the entire electrical load. Each AC power supply has two AC appliance inlets that require a dedicated AC power feed. For 100-120 VAC, both inlets are in use. For 200-240 VAC, only one inlet is in use. Each AC power supply can draw up to 28 amps at 100 VAC. DC power supplies can draw up to 60 Amps at -48 VDC.



## M120 FPCs

- The M120 supports the following FPC types:
  - FPC1—Rated at 4 Gbps full duplex with 4 PICs
  - FPC2—Rated at 16 Gbps full duplex with 4 PICs
  - FPC3—Rated at 20 Gbps full duplex with 1 PIC
  - cFPC—Rated at 20 Gbps full duplex with 2 per chassis



10 Gigabit Ethernet cFPC

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | A-15

### M120 FPCs

M120 FPCs are classified identically to T Series FPCs, as types 1, 2, or 3, and have the same performance capabilities. In addition, the M120 supports the new cFPCs, which is an entire Type 3 FPC and PIC combination, leveraging the space savings of the I 2.0 chip over the LMNR chipset. As a result, it consumes approximately the same space as a traditional PIC. The M120 has two cFPC slots for uplinks. Sample interface types available as cFPCs are the OC192 and 10 Gigabit Ethernet.

## M120 Cooling

- M120 router uses front-to-rear cooling

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | A-16

### M120 Cooling

The diagram on the slide shows the airflow pathways in the M120. During normal operation, the fans in each fan tray function at less than full speed. The CB constantly monitors the temperatures detected by sensors on the midplane and router components, adjusting the speed of the fans as necessary. If the router temperature exceeds the acceptable maximum, the CB turns off the power supplies.

## M320 Multiservice Edge Router

- M320:
  - Industry's most scalable, secure, and reliable 10 Gbps-enabled multiservice edge platform
    - 385 Mbps to 320 Gbps aggregate throughput
    - New 1.6 GHz RE
    - Leverages T640 LMNR chipset
  - Optimized to consolidate the functions of multiple platforms into a single platform
  - Fully redundant hardware
  - Support for M Series and T Series PICs
  - High-density
    - 16 10GE/OC192c (STM64)
    - 160 Gigabit Ethernet



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | A-17


### M320 Multiservice Edge Router Overview

The M320 is a high performance, 10-Gbps-capable, distributed-architecture edge router. It offers up to 16 OC192c (STM64) PICs per chassis (32 per rack) or up to 64 OC48c (STM16) ports per chassis (128 per rack) with up to 320 Gbps throughput. The M320 is ideal for medium-sized backbone cores requiring predictable performance for feature-rich infrastructures, and it also supports provider edge services in 10 gigabit points of presence with the ability to support up to 32 Type 1 and Type 2 PICs and up to 16 Type 3 PICs for 10 Gbps uplinks. In addition, this platform is ideal where switching fabric and RE redundancy are necessary. All major components are field-replaceable, increasing system serviceability and reliability, and decreasing mean time to repair. The PICs are compatible with M40e and the T320 and T640 core routers.

The M320 leverages the same LMNR chipset as the T640 and T320.

## M40e and M320 Cooling

- M40e and M320 use front-to-rear cooling
  - Field-replaceable air filter and fan components
- Fan speed and temperature sensors monitored by JUNOS Software



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | A-18

### Front-to-Rear Cooling


The cooling systems on the M40e and M320 consist of two separate subsystems—front and rear. The front cooling subsystem consists of an upper impeller and a lower fan tray; it cools the FPCs, the PICs, and the midplane. The front impeller and fan tray are both hot-insertable and hot-removable.

The rear cooling subsystem consists of a pair of impellers that cool the Switching and Forwarding Modules (SFMs), RE, Miscellaneous Control Subsystem (MCS), the Packet Forwarding Engine Clock Generators (PCGs), and the power supplies. Each rear impeller is hot-insertable and hot-removable. The upper and lower impellers are not interchangeable.

Systems with front-to-back cooling generally support metal-screen air filters that might become contaminated with dirt or other debris. You should inspect and clean these filters periodically. You should never run the system with the air filters removed, because it might ingest foreign objects, causing damage to the system. In many cases, the air filters also provide electromagnetic interference (EMI) shielding.

### Temperature and Fan Speed Monitoring

Juniper Networks M Series routers are designed to operate with missing or failed fans. The system can measure the presence and rotational speed of all fan components. When chassis temperatures rise fan speed increases, and a yellow alarm generates. If it continues to rise, a red alarm generates; the system will shut itself down to prevent irreparable damage to the ASICs if the high-temperature condition persists.



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | A-19

## M320 Hardware Components

The hardware components for the M320 include the following:

- *Four AC or DC power supplies:* The maximum chassis power is 3150 watts (65 A at -48 VDC) when fully loaded. Each AC supply provides power to all components. In the DC power configuration, the DC power supplies in slots PEM0 and PEM2 are load-sharing and provide power to the FPCs in slots FPC3 through FPC7. The DC power supplies in slots PEM1 and PEM3 are load-sharing and provide power to the FPCs in slots FPC0 through FPC2, Switch Interface Boards (SIBs), CBs, and REs. All DC power supplies provide power to the fan trays. The DC power supplies are fully redundant. The DC power supplies in slots PEM0 and PEM1 can provide full power to the router. Likewise, the DC power supplies in slots PEM2 and PEM3 can also provide full power. The DC power supply in PEM2 serves as the backup to the DC power supply in slot PEM0, and the DC power supply in PEM3 serves as the backup to the DC power supply in slot PEM1.
- *Sheet metal chassis.*
- *Redundant cooling.*

*Continued on next page.*

### M320 Hardware Components (contd.)

- *One or two host subsystems:* The RE and CB work in pairs for host subsystem redundancy. The CB provides control of chassis power, environmental control systems, online and offline of system components, and Stratum 3 synchronization reference. Unlike most M Series and T Series platforms, the M320 uses a Gigabit Ethernet uplink (bcm0) to attach to the various components comprising the PFE.
- *Switch Interface Boards:* Four Switch Interface Boards comprise the switch fabric, and all are in use in normal operation. The failure of a SIB results in a a 25% reduction of switching capacity, which is handled gracefully.
- *Connector Interface Panel (CIP).*
- *Flexible PIC Concentrators:* Up to eight Type 1, Type 2, or Type 3 FPCs; you can mix FPC types can be mixed.

## M320 FPCs

### ■ Three types of FPCs

- FPC1: 1 PFE, 4 Gbps aggregate throughput
  - Supports four hot-swappable M40e-style or M160-style PICs
- FPC2: 2 PFE, 16 Gbps aggregate throughput full duplex
  - Supports four hot-swappable M160 or T640 PICs
- FPC3: 1 PFE, 20 Gbps aggregate throughput full duplex
  - Supports two native T Series PICs

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services


www.juniper.net | A-21

### Three Types of FPCs

The M320 associates with FPC Type 1, Type 2, and Type 3. The different FPC types accommodate the reuse of existing M Series PICs as well as the use of native T Series PICs. The M320 supports any combination of FPC1, FPC2, or FPC3 in a single chassis.

The M320 uses the same LMNR chipset as the T640 and T320 platform.





### Distinguishing M320 FPC Types

The M320 supports three types of FPCs. The platform can operate with any combination of FPC types installed. The best way to determine the FPC type is by the particulars of the PIC online and offline buttons:

- **FPC1:** Accommodates four M40e hot-swappable PICs, including single-port OC12 (STM4) and Gigabit Ethernet interfaces. The PIC slots number top to bottom from 0 (zero) to 3.
- **FPC2:** Accommodates four M40e high-speed, hot-swappable PICs such as a one port OC48 (STM16) or a 4-port Gigabit Ethernet interface. The PIC slots number top to bottom from 0 (zero) to 3.
- **FPC3:** Accommodates two high-speed T Series hot-swappable PICs, such as a one port 10 GE Ethernet interface or an OC192c (STM64) interface. The PIC slots number top to bottom from 0 (zero) to 1.

You can visually distinguish the three types of FPC by the different PICs installed in the FPC. PICs compatible with an FPC1 do not have an offline button on their faceplate. The offline buttons for FPC1 are on the FPC faceplate. PICs compatible with an FPC2 have an offline button on their faceplate. PICs compatible with an FPC3 have a plastic ejector handle at the top of their faceplate.



## The T640 Core Router

- Industry benchmark for scalability, density, and dependability
  - Designed with robust CoS and hardware-based MPLS for ATM, Frame Relay, and IP services
  - Unsurpassed 10 Gbps density (SONET/SDN and Gigabit Ethernet)
  - Rich, dependable services with uncompromising performance
  - Operational simplicity with a single JUNOS Software release across M Series and T Series platforms
  - Breakthrough asset longevity with T640 matrix technology



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | A-23

### The T640 Core Router Overview

Each half-rack T640 core router supports 32 10-Gbps ports or 128 OC48c (STM16) ports, and each slot handles 80 Gbps of aggregate throughput, fulfilling the need for current and future high-bandwidth services. Revolutionary matrix technology, which is an optical backplane extension, enables two-way connectivity: 640 Gbps of front-panel throughput and an oversized 1280 Gbps of rear-panel throughput for nonblocking, any-to-any connectivity to other T640 routers. These routers can interconnect to form a single logical routing entity. As such, service providers can craft architectures that preserve all front-panel ports for revenue generation. When connected using matrix technology, this future multichassis configuration eliminates intermediate layers and reduces complexity by requiring fewer hops. This scalable configuration increases equipment lifespan and further reduces capital expenditure costs.

The T640 is for the core of large service provider networks. It targets roles requiring a high density of 10 Gbps interfaces.


The T640 system is the first Juniper Networks product based on the T640 Series set of ASICs. The ASIC set allows for the building of a distributed system providing up to 320 Gbps of forwarding capacity in half of a 7 foot rack. Each slot in a T640 core router chassis can contain up to two T640 PFEs. Each PFE provides 20 Gbps of forwarding capacity. With eight FPC slots in the chassis, 320 Gbps of forwarding capacity per router are available.

*Continued on next page.*

### The T640 Core Router Overview (contd.)

In contrast to the centralized architecture associated with the M Series, the T640 systems have a distributed architecture. Each PFE is self-contained with its own hardware route lookup engine and its own delay bandwidth buffer. Two PFEs handle incoming unicast packets—one incoming to the router and one outgoing. Multicast packets travel to as many PFEs as necessary, based on the number of multicast destinations.

Not for Reproduction



## T640 Major Components

The major components of the T640 from the front view are the following:

- *ESD point:* Consists of two electrostatic discharge points (banana plug receptacles)—one on the front and one on the rear.
- *FPCs:* Up to eight FPCs install vertically in the front of the router. An FPC can hold up to four PICs. Each FPC contains from one to two complete PFEs. The PFEs receive incoming packets from the PICs installed on the FPC and forward them through the switch planes to the appropriate destination port. Each FPC contains data memory, which is managed by the Queuing and Memory Interface ASICs.
- *Craft Interface:* Allows you to view status and troubleshooting information at a glance and to perform many system control functions.
- *Connector Interface Panel:* Consists of Ethernet, console, and auxiliary connectors for the REs and alarm relay contacts. The front ESD point is near the bottom of the CIP.
- *Fan tray:* All chassis fans are redundant. In a failure situation, the Control Board is responsible for increasing the speed of the remaining fans to provide the necessary cooling.

*Continued on next page.*

## T640 Major Components (contd.)

The major components of the router chassis from the rear view are the following:

- *Switch Interface Board:* The T640 switching architecture is based on SIBs. Each of the system's five SIBs connects to each FPC in the router. Four SIBs are in use at any one time, with the fifth provided for redundancy.
- *Routing Engine:* The T640 contains up to two identical REs. One RE is required for line card chassis (LCC) operation and is a standard component of the base chassis. The second is necessary for redundant configurations. Each RE operates in a protected memory environment that limits the effect of a software process failing.
- *Control Board:* Provides control and monitoring functions for the chassis. A CB associates with each RE to provide control redundancy. The active CB always associates with the active RE. One CB is necessary for chassis operation and is a standard part of the LCC. The second is necessary for redundant system configurations where a second RE is present.
- *SONET Clock Generator:* The T640 provides SONET clocking through the SCGs. The SONET clocks generate on each SCG and distribute to all receiver modules along with a signal *mux* select, indicating which clock is active and, therefore, which clock is the master.
- *Power supplies:* The T640 uses a distributed power entry mechanism to allow for modularity, lower current, and easier hot-swapping. Each Power Entry Module (PEM) requires two 65 amp DC feeds. One PEM is necessary for chassis operation and runs the router indefinitely. The second PEM is necessary for redundancy. Both PEMs are a standard part of the chassis. They are hot-swappable and load-sharing. A fully loaded T640 can draw up to 6,500 watts (68 A at -48 VDC *per input*; note that each PEM has two inputs).

## The T320 Core Router

- The T320 is ideal for smaller sites
  - 1/3 rack chassis
  - Offers unprecedented density and double the power efficiency of full-rack alternatives
    - Three units per rack



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services


www.juniper.net | A-27

### The T320 Core Router

The T320 core router is the industry's most compact 10 Gbps routing platform, fitting three chassis to a rack. While fully leveraging the T Series architecture and leading-edge T Series ASICs, the T320 offers significant gains in form factor, power consumption, and bandwidth density. Designed for a breadth of applications, the T320 is ideal for multiservice transit, peering, metro Ethernet aggregation, data center aggregation, and carrier-of-carrier VPNs.

As an entry point into the T Series routing family, the T320 delivers unprecedented levels of flexibility, dependability, performance, density, and scalability. Network service providers can start with dense OC3 (STM1) and then linearly scale IP services to dense 10 Gbps rates in a single T320 chassis. If further capacity is necessary, service providers can cost-effectively migrate to the T640 because PICs are portable between platforms with identical feature support. Key characteristics of the T320 include the following:

- 320 Gbps throughput (640 Gbps aggregate throughput);
- 385 Mbps forwarding with features enabled;
- Space efficient with three chassis per rack;
- Cost-efficient scalability with M40e, T320, and T640 interface portability;
- Low power consumption of 60 A at -48 V and 2,880 watts; and
- Rich, dependable IP services on any port to meet advanced service-level agreements.



### T320 Major Components

The major components of the T320 from the front view are the following:

- *ESD point*: Consists of two ESD points (banana plug receptacles)—one in the front and one in the rear.
- *FPCs*: Up to eight FPCs install vertically in the front of the router. Each FPC can hold up to two PICs, and each FPC contains a complete PFE. The PFEs receive incoming packets from the PICs installed on the FPC and forward them through the switch planes to the appropriate destination port. Each FPC contains data memory, which is managed by the Queuing and Memory Interface ASICs.
- *Craft Interface*: Allows you to view status and troubleshooting information at a glance and to perform many system control functions.
- *Connector Interface Panel*: Consists of Ethernet, console, and auxiliary connectors for the REs and alarm relay contacts. The front ESD point is near the bottom of the CIP.
- *Fan tray*: All chassis fans are redundant. In a failure situation, the CB is responsible for increasing the speed of the remaining fans to provide the necessary cooling.

*Continued on next page.*

## T320 Major Components (contd.)

The major components of the router chassis from the rear view are the following:

- *Switch Interface Board:* The T320's switching architecture is based on SIBs. Each of the system's three SIBs connects to each FPC in the router. SIBs 1 and 2 are in use in normal operation with the third SIB (SIB 0) providing switch fabric redundancy. Throughput can degrade slightly when the T320 uses SIB 0 to compensate for a failure of either SIB 1 or SIB 2.
- *Routing Engine:* The T320 contains up to two identical REs. One RE is necessary for LCC operation and is a standard component of the base chassis. The second is necessary for redundant configurations. Each RE operates in a protected memory environment that limits the effect of a software process failing.
- *Control Board:* Provides control and monitoring functions for the chassis. A CB associates with each RE to provide control redundancy. The active CB always associates with the active RE. One CB is necessary for chassis operation and is a standard part of the LCC. The second is necessary for redundant system configurations where a second RE is present.
- *SONET Clock Generator:* The T320 provides SONET clocking through the SCGs. The SONET clocks generate on each SCG and distribute to all receiver modules along with a signal *mux* select, indicating which clock is active and, therefore, which clock is master.
- *Power supplies:* The T320 uses a distributed power entry mechanism to allow for modularity, lower current, and easier hot-swapping. Each Power Entry Module requires a single 65 amp DC feed. One PEM is necessary for chassis operation and runs the router indefinitely. The second PEM is necessary for redundancy. Both PEMs are a standard part of the chassis. They are hot-swappable and load sharing. A fully loaded T320 can draw up to 2,880 watts (60 A at -48 VDC).



## T Series FPCs

- **T Series FPC types:**
  - T320 FPCs support two PICs per FPC:
    - FPC1: One PFE, 3.2 Gbps aggregate throughput; supports M Series PICs
    - FPC2: One PFE, 10 Gbps aggregate throughput; supports FPC2-style PICs (M40e high-speed and formerly M160)
    - FPC3: One PFE, 40 Gbps aggregate throughput; supports native T Series PICs
  - T640 FPCs support four PICs per FPC:
    - FPC2: One PFE, 32 Gbps aggregate throughput; supports FPC2-style PICs
    - FPC3: Two PFEs, 80 Gbps aggregate throughput; supports native T Series PICs
- **FPC type distinguished by PIC offline button placement and presence of PIC ejector**

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-30

### T Series FPC Types

T Series platforms associate with FPC Type 1, Type 2, and Type 3. The different FPC types accommodate the reuse of existing M Series PICs in the newer T Series platforms.

The T320 supports any combination of FPC1, FPC2, or FPC3 in a single chassis. In all cases each T320 FPC supports a maximum of two PICs. The Type 1 FPC supports legacy M Series PICs and is rated at 3.2 Gbps of aggregate throughput. As noted on the slide, the FPC1 is supported only in the T320. The T320 FPC2 is rated at 10 Gbps of aggregate throughput. All T320 FPCs consist of a single PFE complex.

The T640 supports FPC2 and FPC3 only. FPC2 is rated at 32 Gbps of aggregate throughput and can support PICs native to the M40e. FPC3 is rated at 80 Gbps of aggregate throughput and can support native T Series PICs such as the OC192C (ST64) PIC. The router can operate with any combination of FPC2s and FPC3s installed. The T640 does not support FPC1. The T640 FPC2 contains a single PFE complex, while the FPC3 contains two complete PFEs.


### Differentiating T Series FPC Types

You can visually distinguish the three types of FPC by the different PICs installed in the FPC. PICs compatible with an FPC1 do not have an offline button on their faceplate. The offline buttons for FPC1 are on the FPC faceplate. PICs compatible with an FPC2 have an offline button on their faceplate. PICs compatible with an FPC3 have a plastic ejector handle at the top of their faceplate.



## T Series PICs and Port Numbers

- Sample port numbering for most T320 and T640 PICs (FPC2 and FPC3 compatible):



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | A-31

### T Series PIC Numbering

The slide details typical port numbering for common T Series PICs.

## Agenda: JUNOS Platform Details

- Primary Components and Characteristics of Multiservice Edge Routers (M Series and T Series)
- Primary Components and Characteristics of Ethernet Services Routers and Switches (MX Series and EX Series)
- Primary Components and Characteristics of Security Services Gateways (SRX Series)
- End-of-Life Products

### Components and Characteristics of Ethernet Services Routers and Switches

The slide highlights the topic we discuss next.

## MX960 Platform (1 of 3)

- **14 Slot Chassis**
  - 2 SCB and RE
    - Possible to install 1 additional SCB
  - 12 DPC slots (11 with SCB redundancy)
  - Up to 480 Gbps (full duplex) from 12 line
  - 17.5 x 8.7 x 23.8 in (16 U)
- **Dependable hardware**
  - Passive midplane
  - Redundant Routing Engines
  - Redundant switching fabric (2+1)
  - Distributed packet forwarding architecture
  - Redundant fan and power



© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-33

### Chassis Description

The MX960 Ethernet Services router is an Ethernet-optimized edge router that provides both switching and carrier-class Ethernet routing. The MX960 has a capacity of up to 480 gigabits per second (Gbps), full duplex. The MX960 enables a wide range of business and residential applications and services, including high-speed transport and VPN services, next-generation broadband multiplay services, and high-volume Internet data center internetworking. The MX960 is 16 rack units (RU) tall. You can stack three routers in a single floor-to-ceiling rack, for increased port density per unit of floor space. The router provides 14 slots that you can populate with up to 12 interface cards and two Switch Control Boards (SCBs) in nonredundant fabric configurations. Fully populated, the MX960 provides up to 480 Gigabit Ethernet or up to 48 10-Gigabit Ethernet ports. Two types of Dense Port Concentrator (DPC) interface cards are available, both of which consist of four PFEs and enable a throughput of 10 Gbps:

- A 40-port Gigabit Ethernet DPC with small form-factor pluggable transceiver (SFP) connectors (1000 Mbit copper and fiber only); and
- A 4-port 10-Gigabit Ethernet DPC with 10-gigabit small form-factor pluggable transceiver (XFP) connectors.

*Continued on next page.*

## Hardware Features

The midplane is in the center of the chassis and forms the rear of the DPC card cage. The DPCs and SCBs install into the midplane from the front of the chassis, and the power supplies install into the midplane from the rear of the chassis. The power supplies and cooling system components also connect to the midplane.

The MX960 chassis provides redundancy and resiliency. The hardware system is fully redundant, including power supplies, fan trays, REs, and SCBs.

Not for Reproduction

## MX960 Platform (2 of 3)

### ■ Power and cooling

- Front-to-back cooling with separate push-pull fan assemblies
- Holds up to 2 fan trays (1+1 redundancy)
- Holds up to 4 power supplies (2+2 DC, 3+1 AC)
- Rear-side power cabling

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-35

### Power and Cooling

In the AC power configuration, the router contains three or four AC power supplies, located vertically at the rear of the chassis in slots PEM0 through PEM3 (left to right). Each AC power supply provides power to all components in the router. When three power supplies are present, they share power almost equally within a fully populated system.

Four AC power supplies provide full power redundancy. If one power supply fails or you remove it, the remaining power supplies instantly assume the entire electrical load without interruption. Three power supplies provide the maximum configuration with full power for as long as the router is operational.

In the DC power configuration, the router contains either two or four DC power supplies located at the lower rear of the chassis in slots PEM0 through PEM3 (left to right). You can upgrade your DC power system from two to four power supplies. The DC power supplies in slots PEM0 and PEM2 provide power to the lower fan tray, DPC slots 6 through 11, and SCB slots 1 and 2. The DC power supplies in slots PEM1 and PEM3 provide power to the upper fan tray, DPC slots 0 through 5, and SCB slot 0.

Four power supplies provide full redundancy. If a DC power supply fails, its redundant power supply takes over without interruption.

The cooling system components work together to keep all router components within the acceptable temperature range. The router has two fan trays located in the front of the router that install horizontally above and below the DPC card cage. Each fan tray contains six fans. The fan trays are interchangeable and are hot-insertable and hot-removable.

## MX960 Platform (3 of 3)

- Three line card (DPC) types currently available:
  - DPCE-R
  - DPCE-X
  - DPCE -Q
- RE options:
  - 1.3 Ghz processor with 2 GB memory
  - 2.2 Ghz processor with 4 GB memory



### Types of Line Cards

Three types of DPCs are available for MX Series routers: DPCE-X, DPCE-R, and DPCE-Q.

The DPCE-X builds upon the Juniper Networks leadership position in providing high-performance networking hardware and software by extending the JUNOS Software to include traditional Layer 2 switching features such as Spanning Tree as well as more recent Ethernet standards developments including Operation, Administration, and Maintenance (OAM). The MX Series combined with the DPCE-X Layer 2 switching cards provide a wide range of MPLS and Ethernet functionality for cost-effective Layer 2 aggregation.

In addition to the Layer 2 switching and MPLS features offered on the DPCE-X cards, the DPCE-R cards provide a full suite of routing protocols and packet processing capabilities.


The DPCE-Q cards provide enhanced queuing capabilities with support of up to 64,000 individual queues, 4 level hierarchical weighted round-robin (WRR), 4 levels of per-VLAN queue priority, and changeable allocation of schedulers per port (8,000 scheduler nodes with 8 queues each or 16,000 nodes with 4 queues each).

*Continued on next page.*

### Types of Routing Engines

Two types of REs are available on the MX Series—an RE with a 1.3 Ghz processor and 2 GB of memory or an RE with a 2.2 Ghz processor and 4 GB of memory. Both REs contain a 30 GB hard disk and a flash disk that can be up to 1 GB. One port per RE exists for USB 2.0 media.

Not for Reproduction



### MX960 Component Placement

The slide shows the placement of the hardware components on the front of the device.



## MX480 Platform (1 of 2)

- **8 slot chassis**
  - 2 SCB and RE slots
  - 6 DPC slots
  - 17.5 x 14 x 23.5 in (8U)
  - Up to 240 Gbps (full duplex) from six line cards
- **Dependable hardware**
  - Passive midplane
  - Redundant Routing Engines
  - Redundant switching fabric (1+1)
  - Distributed packet forwarding architecture
  - Redundant power



© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-39

### Chassis Description

The MX480 Ethernet Services router is an Ethernet-optimized edge router that provides both switching and carrier-class Ethernet routing. The MX480 has a capacity of up to 240 Gbps, full duplex. The MX480 enables a wide range of business and residential applications and services, including high-speed transport and VPN services, next-generation broadband multiple services, and high-volume Internet data center internetworking.

The MX480 is eight RU tall. You can stack five routers in a single floor-to-ceiling rack, for increased port density per unit of floor space. The router provides eight slots that you can populate with up to six DPC cards and two SCBs in nonredundant fabric configurations.

Fully populated, the MX480 provides up to 240 Gigabit Ethernet or up to 24 10-Gigabit Ethernet ports. Six types of DPC cards are available, each of which consists of four PFEs and enables a throughput of 10 Gbps.

### Hardware Features

The midplane is in the center of the chassis and forms the rear of the DPC card cage. The DPCs and SCBs install into the midplane from the front of the chassis, and the power supplies install into the midplane from the rear of the chassis. The power supplies and cooling system components also connect to the midplane.

The MX chassis provides redundancy and resiliency. The hardware system is fully redundant, including power supplies, REs, and SCBs.

## MX480 Platform (2 of 2)

- Power and cooling
  - Side-to-side cooling
  - Holds a single fan tray
  - Holds up to 4 power supplies (2+2 DC, 3+1 AC)
  - Rear-side power cabling

### Power and Cooling

In the AC power configuration, the router contains three or four AC power supplies, located vertically at the rear of the chassis in slots PEM0 through PEM3 (left to right). Each AC power supply provides power to all components in the router. When three power supplies are present, they share equally within a fully populated system.

Four AC power supplies provide full power redundancy. If one power supply fails or you remove it, the remaining power supplies instantly assume the entire electrical load without interruption. Three power supplies provide the maximum configuration with full power for as long as the router is operational.

In the DC power configuration, the router contains either two or four DC power supplies located at the rear of the chassis in slots PEM0 through PEM3 (left to right). You can upgrade your DC power system from two to four power supplies. The DC power supplies in slots PEM0 and PEM2 provide power to the fan tray, DPC slots 0 and 1, and SCB slots 0 and 1. The DC power supplies in slots PEM1 and PEM3 provide power to the fan tray and DPC slots 2 through 5.

Four power supplies provide full redundancy. If a DC power supply fails, its redundant power supply takes over without interruption.

The cooling system components work together to keep all router components within the acceptable temperature range. The router has one fan tray and one air filter that installs vertically in the rear of the router. The fan tray contains six fans. The fan tray is hot-insertable and hot-removable.

## MX240 Platform (1 of 2)

- **4 slot chassis**
  - 3 DPC slots (2 with SCB redundancy)
  - 1 RE slots
  - 17.5 x 8.7 x 23.8 in (5U)
  - Up to 120 Gbps (full duplex) from four line cards
- **Dependable hardware**
  - Passive midplane
  - Redundant Routing Engines
  - Redundant switching fabric (1+1)
  - Distributed packet forwarding architecture
  - Redundant power



© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-41

### Chassis Description

The MX240 Ethernet Services router delivers increased port density over traditional carrier Ethernet platforms as well as performance of 200+ Gbps throughput, scalability, and reliability in a space-efficient package. The MX240 offers fully redundant hardware that includes redundant SCBs and REs to increase system availability.

The MX240 is a 4 slot chassis, with 3 DPC slots, and one SCB slot. If chassis fabric redundancy is required then one DPC slot can be used for an additional SCB and RE pair. The MX240 is 5U in size, which enables you to stack up to 8 devices in a standard floor-to-ceiling rack.

### Hardware Features

The midplane is in the center of the chassis and forms the rear of the DPC card cage. The DPCs and SCBs install into the midplane from the front of the chassis, and the power supplies install into the midplane from the rear of the chassis. The power supplies and cooling system components also connect to the midplane.

The MX chassis provides redundancy and resiliency. The hardware system is fully redundant, including power supplies, REs, and SCBs.

## MX240 Platform (2 of 2)

- Power and cooling
  - Side-to-side cooling
  - Holds a single fan tray
  - Holds up to 4 power supplies (2+2 DC, 3+1 AC)
  - Rear-side power cabling

### Power and Cooling

The MX240 supports either the low-line (110V) AC power configuration or the high-line (220V) AC power configuration.

In the low-line AC power configuration, the MX240 contains either two non-redundant AC power supplies—located horizontally at the rear of the chassis in slots PEM0 and PEM2 (left to right)—or four redundant AC power supplies, located in slots PEM0 through PEM3 (left to right). The low-line configuration requires two power supplies, and the third and fourth power supplies provide redundancy. Each AC power supply provides power to all components in the router. When two power supplies are present, they share power almost equally within a fully populated system. If one power supply in a redundant configuration fails or you remove it, the remaining power supplies assume the entire electrical load without interruption. Two power supplies provide the maximum configuration with full power for as long as the router is operational.

*Continued on next page.*


## Power and Cooling (contd.)

In the high-line AC power configuration, the MX240 contains one or two AC power supplies, located horizontally at the rear of the chassis in slots PEM0 and PEM2 (left to right). The high-line configuration requires one power supply, with the second power supply providing redundancy. Each AC power supply provides power to all components in the router. When two power supplies are present, they share power almost equally within a fully populated system. If one power supply fails or you remove it, the remaining power supply assumes the entire electrical load without interruption. One power supply can provide maximum configuration with full power for as long as the router is operational.

Each DC power supply has a single DC input (-48 VDC and return) that requires a dedicated 40 A (-48 VDC) circuit breaker for the maximum router hardware configuration. The DC power supply in PEM0 must receive power from dedicated power feeds derived from feed A, and the DC power supply in PEM2 must receive power from dedicated power feeds derived from feed B. This configuration provides the commonly deployed A and B feed redundancy for the system. Four power supplies provide full redundancy. If a DC power supply fails, its redundant power supply takes over without interruption.

The cooling system components work together to keep all router components within the acceptable temperature range. The router has one fan tray and one air filter that install vertically in the rear of the router. The fan tray contains six fans. The fan tray is hot-insertable and hot-removable.

## MX960 Switch Control Board



© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-44

### Switch Control Board

The slide displays a picture of the RE and SCB board as well as a graphic representation of the RE. The three RJ45 connectors on the RE contain the aux port, console port, and fxp0 port out-of-band management interface.

## MX960 SCB and RE Configuration

- **Fabric redundancy**
  - Redundant fabric configuration
    - 3 SCBs in SCB slot 0, SCB slot 1, and SCB slot 2
    - 11 DPCs with DPCs in DPC slots 0–5 and DPC slots 7–11
  - Nonredundant fabric configuration
    - 2 SCBs in SCB slot 0 and SCB slot 1
    - 12 DPCs in DPC slots 0–11
- **Control and Routing Engine redundancy**
  - Independent of fabric redundancy
  - Nonredundant Control and Routing Engine
    - Routing Engines in SCB slot 0 or SCB slot 1
  - Redundant Control and Routing Engine Redundancy
    - Routing Engines in SCB slot 0 and SCB slot 1

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-45

### Fabric Redundancy

You can install up to three SCBs in the router. Two SCBs are necessary to run the router and the third functions as the backup. The third installed SCB provides fabric redundancy, but no additional control or routing functions. If the master fails or you remove it, the backup restarts and becomes the master. In a nonredundant fabric configuration the loss of an SCB will not bring down the system but will cause a degradation in forwarding rate.

### RE Redundancy

You can install one or two REs in the router. The REs install into the front of the chassis in vertical slots directly into the SCBs labeled 0 and 1. If two REs are present, one functions as the master and the other acts as the backup. If the master RE fails or you remove it, and the backup configuration is appropriate, the backup takes over as the master.



## MX480 SCB and RE Configuration

- Fabric redundancy
  - Redundant fabric configuration
    - 2 SCBs in SCB slot 0 and SCB slot 1
    - 6 DPCs in DPC slots 0–5
  - Nonredundant fabric configuration
    - 1 SCB in SCB slot 0 or SCB slot 1
    - 6 DPCs in DPC slots 0–5
- Control and Routing Engine redundancy
  - Independent of fabric redundancy
  - Nonredundant Control and Routing Engine
    - Routing Engines in SCB slot 0 or SCB slot 1
  - Redundant Control and Routing Engine
    - Routing Engines in SCB slot 0 and SCB slot 1

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-46

### Switch Control Board

You can install one or two SCBs in the router. If two SCBs are present, one functions as the master SCB and the other as its backup. If the master fails or you remove it, the backup restarts and becomes the master.

### RE Redundancy

You can install one or two REs in the router. The REs install into the front of the chassis in vertical slots directly into the SCBs labeled 0 and 1. If two REs are present, one functions as the master and the other acts as the backup. If the master RE fails or you remove it, and the backup configuration is appropriate, the backup takes over as the master.



## MX240 SCB and RE Configuration

- Fabric redundancy
  - Redundant fabric configuration
    - 2 SCBs in SCB slot 0 and 1/0
    - 2 DPCs DPC slots 1/0, 1, or 2
  - Nonredundant fabric configuration
    - 1 SCB in SCB slot 0 or SCB slot 1
    - 3 DPCs in DPC slots 0–2
- Control and Routing Engine redundancy
  - Independent of fabric redundancy
  - Nonredundant Control and Routing Engine
    - Routing Engines in SCB slot 0 or SCB slot 1
  - Redundant Control and Routing Engine
    - Routing Engines in SCB slot 0 and SCB slot 1

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-47

### Switch and Control Board

You can install one or two SCBs in the router. If two SCBs are present, one functions as the master SCB and the other as its backup. If the master fails or you remove it, the backup restarts and becomes the master.

### RE Redundancy

You can install one or two REs in the router. The REs install into the front of the chassis in vertical slots directly into the SCBs labeled 0 and 1. If two REs are present, one functions as the master and the other acts as the backup. If the master RE fails or you remove it, and the backup configuration is appropriate, the backup takes over as the master.

## High Density DPC Architecture

- Dense Port Concentrator: SFPs or XFPs
- Line rate connectivity to the switch fabric
- 4 PFEs per DPC

The diagram illustrates the High Density DPC Architecture. It shows an MX480 router with a DPC card installed. The DPC card is highlighted in orange and contains four PFEs, each with an I-chip and an ESE chip. The MX480 router is shown with the DPC card inserted into a slot. The DPC card is labeled 'DPC' and contains four PFEs, each with an I-chip and an ESE chip. The MX480 router is labeled 'MX480'.

© 2009 Juniper Networks, Inc. All rights reserved. Juniper Education Services www.juniper.net | A-48

### DPC Media Types

All DPCs come as either SFPs or XFPs which use the most optimal port density with cost efficiency.

### Line Rate

Each DPC has connections to the switch fabric providing line-rate connectivity for each port on the card.

### Packet Forwarding Engines

DPCs provide multiple physical interfaces and PFEs on a single board that installs in a slot in the MX Series Ethernet Services Router. A DPC receives incoming packets from the network and sends outgoing packets to the network. Each DPC contains four PFEs. The PFEs on the DPC have purpose-built ASICs that perform packet processing and forwarding. Each PFE consists of one I-chip for Layer 3 processing and one Layer 2 network processor.

## MX Series Line Cards

- **Dense Port Concentrators:**
  - 4 port XFP
  - 40 port SFP
- **DPC-R line cards:**
  - Full JUNOS Software Layer 3 routing feature set
  - Layer 2 Ethernet switching features
  - Per-VLAN policing
  - Per-VLAN rewrite
  - Per-VLAN tri-color marking
  - Per-VLAN classification
  - Per-VLAN accounting
  - Per-VLAN filtering
  - Eight queues per port



DPCE-R-40GE-SFP  
 DPCE-R-4XGE-XFP  
 DPCE-R-40GE-SFP  
 DPCE-R-4XGE-XFP

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | A-49

### DPC Types

The DPCs are optimal for Ethernet density and can support up to 40 Gigabit Ethernet or 4 10-Gigabit Ethernet ports.

### DPC-R Features

The DPCs support a wide range of Layer 2 and Layer 3 Ethernet functionality, including 802.1Q virtual LAN (VLAN), link aggregation, circuit cross-connect, Virtual Router Redundancy Protocol (VRRP), Layer 2 to Layer 3 mapping, and port monitoring. Additionally, the DPCs support filtering, sampling, load balancing, rate limiting, class of service, and other key features necessary for the deployment of dependable, high-performance Ethernet services.

## EX8200 Platform (1 of 2)

### High-performance chassis platforms

- EX8208: Eight 200 Gbps line cards
- EX8216: Sixteen 200 Gbps line cards
- 100 Gigabit Ethernet ready
- Fully redundant routing engines with N+1 redundant switch fabrics
- Up to 256 wire-speed, non-blocking 10 Gigabit Ethernet ports in a rack



Module description	Max ports	Interface
48 port 10, 100, or 1000B-T	384 or 768	RJ45
48 port 100B-FX or 1000B-X	384 or 768	SFP
8 port 10 Gigabit Ethernet	64 or 128	SFP+

© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | A-50

### High Performance EX8200 Platform

EX8200 Series Ethernet switches come in two types, EX8208 and EX8216. The EX8200 Series switches are modular in nature and are designed for ultra high-density environments such as campus aggregation, data center, or high performance core switching environments. The EX8208 switch uses eight 200 Gbps line cards, while the EX8216 switch uses sixteen 200 Gbps line cards. The switches provide high availability and redundancy for all major hardware components, including REs, switch fabric, fan tray, and power supplies.

## EX8200 Platform (2 of 2)

- Fully redundant power and cooling
  - Redundant, load-sharing PSUs (AC and DC)
  - Hot-swappable fan tray with redundant fans
- Proven Juniper Networks Technology
  - Switch fabrics and control plane
  - PFE
  - JUNOS Software

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-51

### Redundant Power and Cooling

EX8200 switches provide high availability and redundancy for all major hardware components—including RE modules, switch fabric modules (SFMs), fan trays (with redundant fans), and load-sharing 2000 watt AC, 3000 watt AC, and 3000 watt DC power supplies.

### Proven Juniper Networks Technology

EX8200 switches deploy proven Juniper Networks technology, which includes REs, switch fabric, and PFEs. The RE modules are hot-insertable and hot-removable and install in the front of the chassis. The SFMs are also hot-insertable and hot-removable, and install in the rear of the chassis. The EX Series deploys the industry-proven JUNOS Software.

## EX4200 Ethernet Switch with Virtual Chassis Technology (1 of 2)

- Virtual chassis technology
  - 128 Gbps virtual backplane
  - Manage up to 10 switches as a single device
  - Extend over 10 Gigabit Ethernet or 1 Gigabit Ethernet uplinks
  - Master and backup route engines



Number of ports	Port type	PoE ports	Max power consumption (incl. PoE)
24	10, 100, and 1000B-T	8	129 (320) W
24	10, 100, and 1000B-T	24	160 (600) W
24	100B-FX and 1000B-X	N/A	108 (N/A) W
48	10, 100, and 1000B-T	8	181 (320) W
48	10, 100, and 1000B-T	48	224 (930) W

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-52

### EX4200 Virtual Chassis Technology

The Juniper Networks EX4200 Ethernet Switches with virtual chassis technology combine the high availability and carrier-class reliability of modular systems with the economics and flexibility of stackable platforms, delivering a high-performance, scalable solution for data center, and campus and branch office environments.

Up to 10 EX4200 switches can interconnect using virtual chassis technology to create a single logical device supporting up to 480 10, 100, or 1000BASE-T ports or 240 100BASE-FX or 1000BASE-X ports, plus an additional 40 Gigabit Ethernet or 20 10-Gigabit Ethernet uplink ports. Different models can mix in a virtual chassis configuration to provide a variety of port and density options.

## EX4200 Series Ethernet Switch with Virtual Chassis Technology (2 of 2)

- Flexible uplink modules
  - 4 port Gigabit Ethernet (SFP)
  - 2 port 10 Gigabit Ethernet (XFP)
- Fully redundant power and cooling
  - Dual, hot-swap AC and DC power supply unit
  - External RPS option
  - Fan FRU and multiple blowers
  - Full Class 3 PoE (15.4 W)
- Runs JUNOS Software with full OSPF and IP multicast in base license

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-53

### Flexible Uplink Modules

You can deploy a single 24-port or 48-port switch initially; as requirements grow, Juniper Networks virtual chassis technology allows up to 10 EX4200 switches to interconnect over a 128 Gbps backplane, which you can manage as a single device, delivering a scalable, pay-as-you-grow solution for expanding network environments. Flexible Gigabit Ethernet and 10-Gigabit Ethernet uplink options enable high-speed connectivity to aggregation-layer or core-layer switches which connect multiple floors or buildings.

### Redundant Power and Cooling

All EX4200 switches include high availability features such as redundant, hot-swappable internal power supplies and field-replaceable, multi-blower fan trays to ensure maximum uptime. In addition, the base EX4200 switch models offer Class 3 Power over Ethernet (PoE), delivering 15.4 watts on the first eight ports to support networked devices such as telephones, video cameras, and wireless LAN (WLAN) access points for low-density converged networks. Full PoE options delivering 15.4 watts on all 24 or 48 ports are also available, making them ideal for high-density IP telephony deployments.

### JUNOS Software

The EX4200 switches leverage the modular JUNOS Software, ensuring consistent implementation and operation.



## EX3200 Ethernet Switches

- Fixed, standalone configuration
- Flexible uplink modules
  - 4 port Gigabit Ethernet (SFP)
  - 2 port 10 Gigabit Ethernet (XFP)
- Modular power and cooling
  - Field-replaceable AC and DC power supply unit
  - External RPS option
  - Field-replaceable fan tray
  - Full Class 3 PoE (15.4 W)
- Runs JUNOS Software with full OSPF and IP multicast in base license



Number of ports	Port type	PoE ports	Max power consumption (incl. PoE)
24	10, 100, and 1000B-T	8	112 (320) W
24	10, 100, and 1000B-T	24	138 (600) W
48	10, 100, and 1000B-T	8	167 (320) W
48	10, 100, and 1000B-T	48	207 (930) W

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-54

### Fixed Configuration

The EX3200 fixed-configuration switches from Juniper Networks offer a high-performance standalone solution for access-layer deployments in branch and remote offices as well as campus networks.

### Uplink Modules

The EX3200 line also supports optional 4 port Gigabit Ethernet and 2 port 10-Gigabit Ethernet uplink modules with pluggable optics to provide high-speed connectivity to aggregation-layer switches or other upstream devices.

### Modular Power and Cooling

Featuring complete Layer 2 and Layer 3 switching capabilities, the EX3200 switches satisfy the wiring closet connectivity requirements of today's high-performance businesses. Four platform configurations are available offering 24 and 48 10, 100, and 1000BASE-T ports with either full or partial PoE. The base 24-port and 48-port EX3200 switches deliver a full 15.4 watts of Class 3 PoE on the first eight ports for supporting networked devices such as telephones, video cameras and WLAN access points in converged networks. The EX3200 switches with full PoE deliver 15.4 watts on all 24 or 48 ports to support high-density IP telephony and other converged network environments.

*Continued on next page.*



## JUNOS Software

The EX3200 switches run the same JUNOS Software used by Juniper Networks routers to power the world's largest and most complex networks.

By utilizing a common operating system, Juniper Networks delivers a consistent implementation and operation of control-plane features across all products—functions ranging from chassis management to Spanning Tree to OSPF. To maintain that consistency, the JUNOS Software adheres to a highly disciplined development process that utilizes a single source code, follows a single quarterly release train, and employs a highly available modular architecture that prevents isolated failures from bringing an entire system down.

Not for Reproduction

## Agenda: JUNOS Platform Details

- Primary Components and Characteristics of Multiservice Edge Routers (M Series and T Series)
- Primary Components and Characteristics of Ethernet Services Routers and Switches (MX Series and EX Series)
- Primary Components and Characteristics of Security Services Gateways (SRX Series)
- End-of-Life Products

### Primary Components and Characteristics of Security Services Gateways

The slide highlights the topic we discuss next.

## SRX Series Services Gateway

- **Feature integration**
  - Security: FW, IDP, DoS
  - Networking: high availability, NAT, QoS, routing, virtual routers
- **Scalability**
  - Linear scale with incremental cards
  - Superior OPEX
- **Based on JUNOS Software**



© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-57

### Feature Integration

The feature integration on SRX Series Services Gateways is enabled by JUNOS Software. The SRX line is equipped with a robust list of features that include firewall, Intrusion Detection and Prevention (IDP), denial of service (DoS) protection, Network Address Translation (NAT), and quality of service (QoS).

### Scalability

Based on the Dynamic Services Architecture, the SRX-series provides unrivaled scalability. Each services gateway can support almost linear scalability, with each additional services processing card (SPC) enabling a fully equipped SRX5800 to support more than 120 Gbps firewall throughput.

### JUNOS Software

The SRX Series runs on the proven JUNOS Software which is based on a modular design offering a single-source operating system for the network. Unlike any other product on the market, it provides one operating system, enhanced through one release, and developed based on a single modular architecture.

## Security Services Gateways Details (1 of 2)

- SRX5600 and SRX5800 share the same cards or blades
  - Distributed multiblade chassis-based design
  - Blades include SCB, RE, SPC, and DPC
  - 2 SCB slots + 6 DPC or SPC slots in SRX5600
  - 2 SCB slots + 12 DPC or SPC slots in SRX5800

### SRX5600 and SRX5800 Details

The SRX5600 uses the same SPCs and IOCs as the SRX5800.

As the brain behind the SRX Series services gateway, the SPCs process all available services on the gateway. Without the need for dedicated hardware for specific services or capabilities, no instances occur in which a piece of hardware is taxed to the limit while other hardware sits idle. All of the processing capabilities of the SPCs process all configured services on the gateway. The same SPCs are supported on both the SRX5600 and SRX5800. To provide the most flexible solution, SRX Series services gateways employ the same modular architecture for SPCs and IOCs. You can equip the SRX Series services gateway with one or several IOCs, with each IOC supporting 40 gigabit interfaces (4 x 10 Gigabit Ethernet or 40 x 1 Gigabit Ethernet).

## Security Services Gateways Details (2 of 2)

- SRX3400 and SRX3600 share the same cards or modules
  - Modularized chassis-based design
  - Cards include midplane, SFB, RE, SPC, NPC, and IOC
  - SPC, NPC, and IOC are form-factor modules and are interchangeable in any CFM slots
  - 7 CFM slots in SRX3400
  - 12 CFM slots in SRX3600
  - SRX5600 and SRX5800 cards are *not* interchangeable with SRX3400 and SRX3600 cards


### SRX3400 and SRX3600 Modules

The SRX3600 and SRX3400 support common form-factor modules (CFM). A single-wide module format and a double-wide module format are available. All single-wide CFMs can insert in any single-wide slot in the front and rear of the chassis. Similarly, all double-wide CFMs can insert in any double-wide slots. DPC, SPC, and Network Processing Card (NPC) modules are in single-wide CFM format. The SCB, the RE, and the CB are not in CFM format, and thus have assigned slots within the chassis. With the interchangeability among DPC, SPC, and NPC, you have more flexibility and scalability when deploying your networks based on the requirements in the field. For example, if you need more ports and a bigger oversubscription ratio, then you can load more slots with DPCs; on the other hand, if you need a smaller oversubscription ratio for better QoS behavior, or if you need more security services, then you can load more slots with SPCs. With this arrangement, only the onboard network interface ports are in use. SRX5600 and SRX5800 cards are not interchangeable with SRX3400 and SRX3600 cards.



### SRX5800 Major Components


The slide presents a pictorial view of the SRX5800 and its major components.



### SRX5600 Major Components

The slide presents a pictorial view of the SRX5600 and its major components.






### SRX Major Components


The slide presents a pictorial view of the SRX3600 and its major components.





### SRX3400 Front View

The slide presents a pictorial view of the SRX3400 and its major components from the front.



### SRX3400 Rear View

The slide presents a pictorial view of the SRX3400 and its major components from the rear.



### SRX210 Major Components

The slide presents a pictorial view of the SRX210 and its major components.

## Agenda: JUNOS Platform Details

- Primary Components and Characteristics of Multiservice Edge Routers (M Series and T Series)
- Primary Components and Characteristics of Ethernet Services Routers and Switches (MX Series and EX Series)
- Primary Components and Characteristics of Security Services Gateways (SRX Series)
- End-of-Life Products

### End-of-Life Products

The slide highlights the topic we discuss next.

## End-of-Life Products

- Details on the following routers:
  - M5 and M10
  - M20
  - M40
  - M160

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-67

### End-of-Life Products

The following slides are from earlier versions of this course. We include them for the benefit of customers with older products.

## M5 and M10 Routers

- M5 and M10 routers used for:
  - Edge applications involving the aggregation of dedicated access circuits
  - Core applications in locations where space and power are at a premium



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

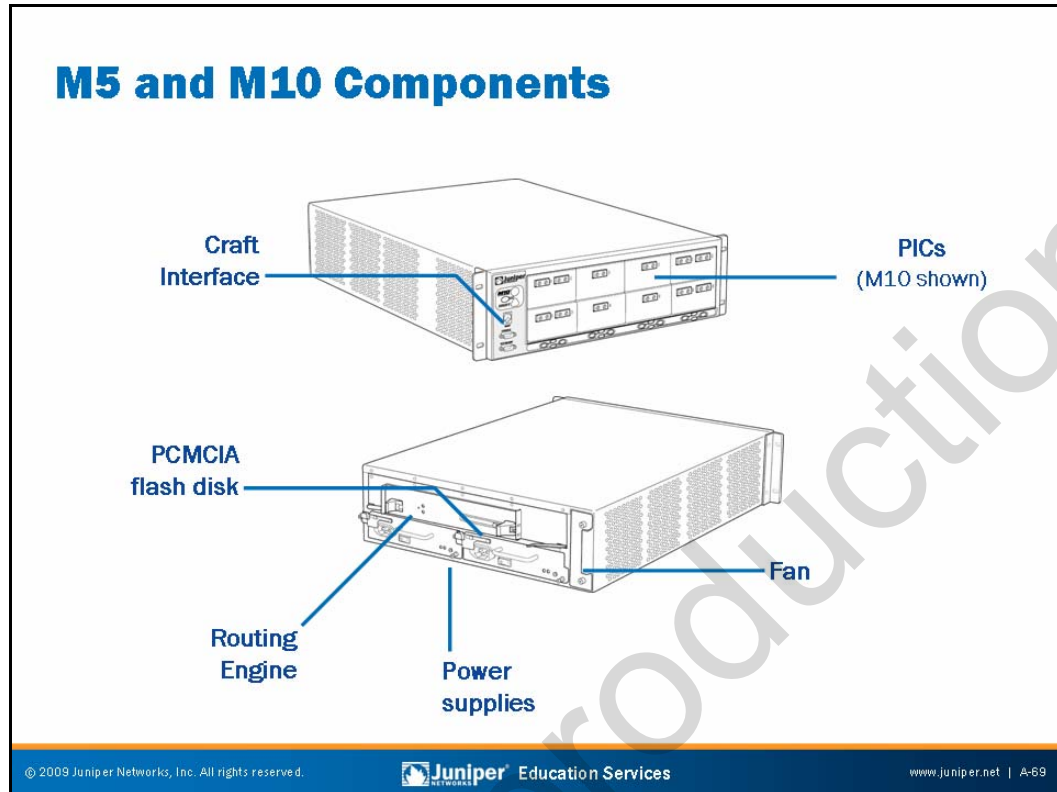
www.juniper.net | A-68

### M5 and M10 Overview

The M5 and M10 routers deliver high performance and highly-flexible interfaces in a space-efficient and power-efficient design. These routers use the same architecture, ASICs, and JUNOS Software as the already proven M Series routers. This Internet-tested core technology is now available at your network edge, along with value-added services, such as packet filtering and sampling.

The oversized forwarding performance of the Internet Processor II ASIC provides wire-speed forwarding with plenty of headroom. In fact, the M5 and M10 routers forward packets at an aggregate throughput rate of 6.4 Gbps and 12.8 Gbps, respectively.

The M5 and M10 are ideal for edge applications involving the aggregation of dedicated access circuits. They are also ideal for core applications in locations where space and power are at a premium, such as smaller metro points of presence. The ability to connect a wide range of high-performance interfaces from T1 and E1 through OC12c (STM4) ensures that you can scale the network easily and cost efficiently.




## M5 and M10 Hardware Components

The hardware components for the M5 and M10 are the following:

- Sheet metal chassis.
- Two power supplies (AC or DC): The maximum chassis power is 340 watts for the M5 and 434 watts for the M10. Power supplies can be either AC or DC (but not both simultaneously); when two are present, power is load-balanced.
- Fan assembly.
- *Routing Engine*: The RE maintains the routing tables and controls the routing protocols, as well as the JUNOS Software processes that control the router's interfaces, 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 PFE.
- *Forwarding Engine Board*: The PFE provides Layer 2 and Layer 3 packet switching, route lookups, and packet forwarding. The Internet Processor II ASIC forwards up to 40 Mbps for all packet sizes. The throughput is 5+ Gbps in an M5 and 10+ Gbps in an M10. The PFE supports the same ASIC-based features supported by all other M Series routers, including filtering and sampling for restricting access, increasing security, and analyzing network traffic.
- Four (M5) or eight (M10) PICs.

## The M5 and M10 Craft Interface

- The M5 and M10 platforms have similar Craft Interface panels
  - PIC online and offline buttons
  - Alarm LEDs
  - RE ports (console, AUX, and a single Fast Ethernet)



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | A-70

### M5 and M10 Craft Interface

The M5 and M10 Craft Interface supports PIC online and offline buttons, alarm LEDs, an alarm cutoff and lamp test button, and the ports used to connect to the routing engine. Ensure that you take PICs offline before removing them from the chassis using the PIC online and offline buttons located below each respective PIC. After inserting a PIC into the chassis, press and hold the online and offline button to activate power to that PIC.

You can achieve access to the local RE using any of the following:

- *Ethernet management port:* Connects the RE to an out-of-band management network.
- *Console port:* Connects a system console to the RE with EIA/TIA-232 serial asynchronous cable. Use the system console to access the CLI to configure the attached router. This port is active by default.
- *Auxiliary port:* Connects a laptop or modem to the RE with EIA/TIA-232 serial asynchronous cable. You can also use this port to access the CLI. This port is disabled by default and you must activate it with a **set system ports auxiliary type terminal-type** command.

Note that you must remove power to the chassis before removing or installing the FEB. You can swap the RE with power still applied, but doing so results in a system reboot.



## M20 Router

- Built for a variety of Internet applications:
  - Ideal for edge applications involving the aggregation of dedicated access circuits
    - Also deployed for high-speed access, public and private peering, hosting sites, and backbone core networks



© 2009 Juniper Networks, Inc. All rights reserved.

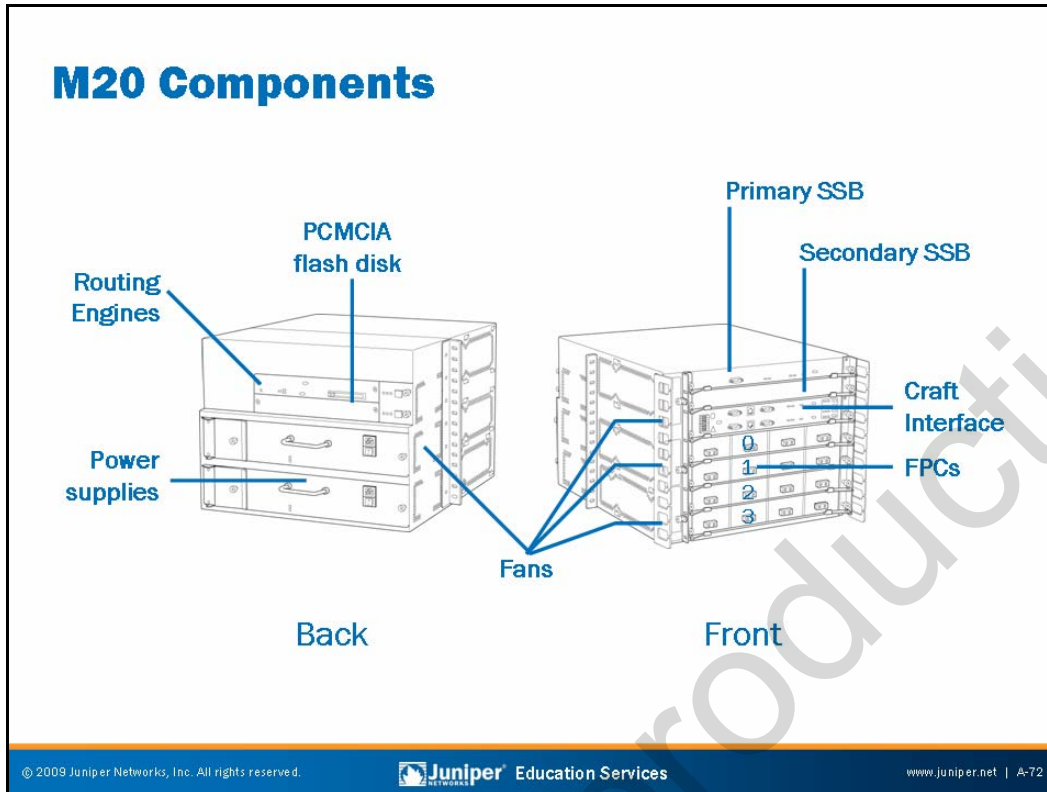
 Juniper Education Services

www.juniper.net | A-71

### M20 Overview

The M20 is a high-performance routing platform built for a variety of Internet applications, including high-speed access, public and private peering, hosting sites, and network applications. The primary application for the M20 is currently edge aggregation of dedicated access circuits.

The M20 leverages proven M Series ASIC technology to deliver wire-rate performance and rich packet processing, such as filtering, sampling, and rate limiting. It runs the same JUNOS Software and supports the same interfaces as other M Series routers, providing a seamless upgrade path that protects your investment. Its compact design (14 inches or 35.56 cm high) delivers market-leading performance and port density while consuming minimal rack space. You can install as many as five M20 units into a single 19-inch equipment rack.



## M20 Hardware Components

The hardware components for the M20 include the following:

- Sheet metal chassis.
- Two power supplies (AC or DC): The maximum chassis power draw is 1200 watts. Power supplies can be either AC or DC, but you cannot mix power supply types. When two power supplies are present, power is load-balanced between them.
- Four fan assemblies.
- One or two REs.
- One or two SSBs.
- Up to four FPCs: Each is populated with up to four PICs for various interface types, including OC48, OC12, OC3 SONET/SDH (including channelized OC12); OC12, OC3 ATM, DS3, Gigabit Ethernet, and Fast Ethernet.

Although the M20 supports RE and system board redundancy, hot PIC swapping is not available. To replace or install a PIC, you must first remove the target FPC from the system. Removing or inserting an FPC results in an approximately 100 millisecond period of packet loss. Be sure to gracefully remove the FPC by depressing the FPC offline button until the green OK light extinguishes. Upon insertion, the FPC automatically comes back online.

## M40 Router

- M40:
  - Original platform shipped by Juniper Networks
  - Replaced by platforms such as the M10i and M40e
    - Smaller form factors
    - Hardware redundancy



© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-73

### M40 Overview

The M40 was the first platform shipped by Juniper Networks. Although it is still supported, newer platforms, such as the M40e, are eclipsing the original M40 due to improvements in redundancy, form factor, and support for newer PIC types.

## M160 Router

- **M160:**
  - Ideal for large networks requiring predictable performance for feature-rich infrastructures
  - Built for large backbone cores, with features enabled for future migration to the backbone edge
  - Redundant hardware
  - OC192c (STM64) and 10 Gigabit Ethernet PICs (FPC2)



© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | A-74

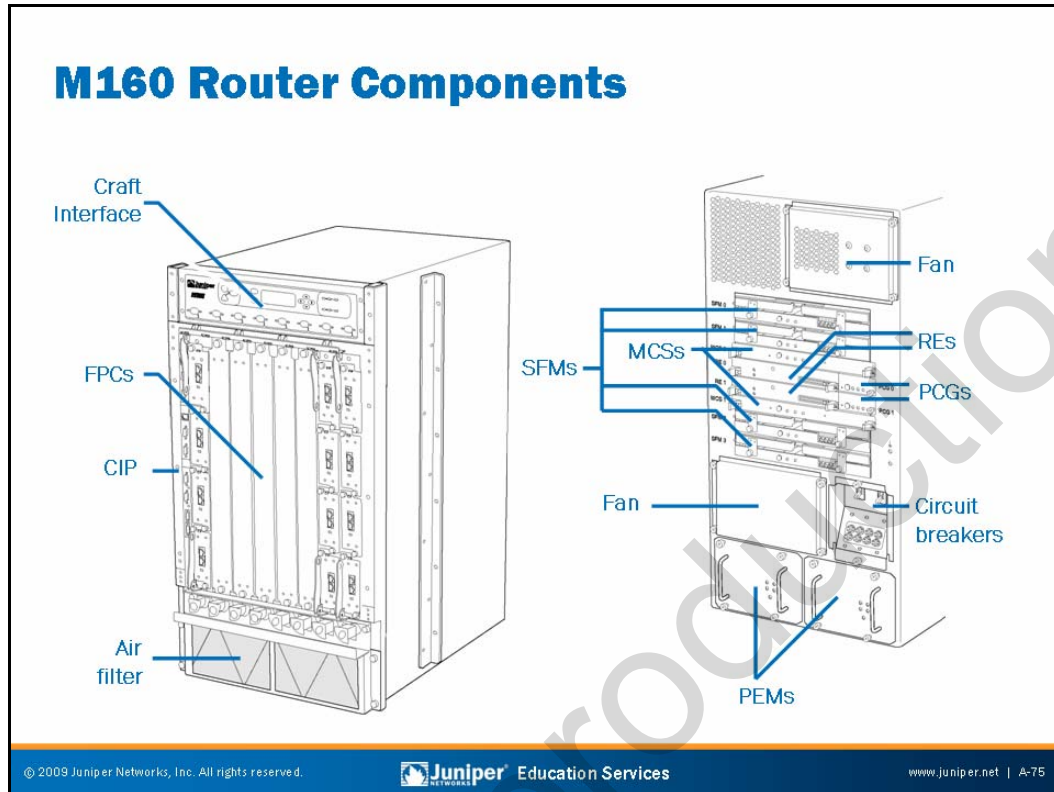
### M160 Overview

The M160 was the first full-performance OC192c (STM64) platform on the market. By providing aggregate forwarding rates of 160 Mbps and aggregate throughput exceeding 160 Gbps, this platform enables you to grow networks rapidly and reliably while taking advantage of increased optical bandwidth.

An ASIC-based packet forwarding path enables full wire-rate forwarding over 10 Gbps circuits. Key to this performance is the Internet Processor II ASIC and the custom made SONET processing ASIC used in the OC192c (STM64) interface.

The M160 offers a complete range of flexible and dense interfaces, allowing you to fit M160 routers into existing environments seamlessly. The platform supports up to 32 OC48c (STM16) PICs per chassis (64 per standard rack) or up to eight OC192c (STM64) PICs per chassis (16 per rack).

The M160 provides a scalable solution using the same JUNOS Software and the same ASICs, and supporting the same value-added services as all other M Series routers.



### M160 Hardware Components

The hardware components for the M160 include the following:

- Two DC-only Power Entry Modules: The maximum chassis power is 3150 watts (65 A at -48 VDC) when fully loaded. When both PEMs are present, power to the chassis is load-balanced between them.
- Sheet metal chassis.
- Redundant cooling.
- One or two host modules: RE and MCS work in pairs for host module redundancy. The MCS provides control of chassis power, environmental control systems, online and offline of system components and Stratum 3 synchronization reference.
- Four SFMs for redundant switch fabric.
- Connector Interface Panel.
- Redundant Packet Forwarding Engine Clock Generators provide redundant synchronization signals to the hardware components in the PFE.

*Continued on next page.*

### M160 Hardware Components (contd.)

- Up to eight FPCs: Populated with up to four PICs for various interface types, including OC48, OC12, O-3 SONET/SDH (including channelized OC12), OC12, OC3 ATM, DS3, Gigabit Ethernet, and Fast Ethernet. FPC1 supports legacy M20 and M40 PICs, while the native FPC2 is for high-speed PICs like the OC48c (STM16) and 48-port Fast Ethernet PICs.

Not for Reproduction

## Summary

- In this appendix, we:

- Listed the primary components and characteristics of multiservice edge routers (M Series and T Series)
- Listed the primary components and characteristics of Ethernet services routers and switches (MX Series and EX Series)
- Listed the primary components and characteristics of security services gateways (SRX Series)
- Listed end-of-life products

### This Appendix Discussed:

- Primary components and characteristics of multiservice edge routers (M Series and T Series);
- Primary components and characteristics of Ethernet services routers and switches (MX Series and EX Series);
- Primary components and characteristics of security services gateways (SRX Series); and
- End-of-life products.

Not for Reproduction





# **Troubleshooting JUNOS Platforms**

## **Appendix B: Packet Flow Details**

Not for Reproduction

## Appendix Objectives

- After successfully completing this appendix, you will be able to:
  - Describe the real-time operating system packet flow
  - Describe the ABC chipset packet flow
  - Describe the LMNR chipset packet flow

### This Appendix Discusses:

- Real-time operating system (RTOS) packet flow;
- ABC chipset packet flow; and
- LMNR chipset packet flow.

## Agenda: Packet Flow Details

- RTOS Packet Flow
- ABC Chipset Packet Flow
- LMNR Chipset Packet Flow

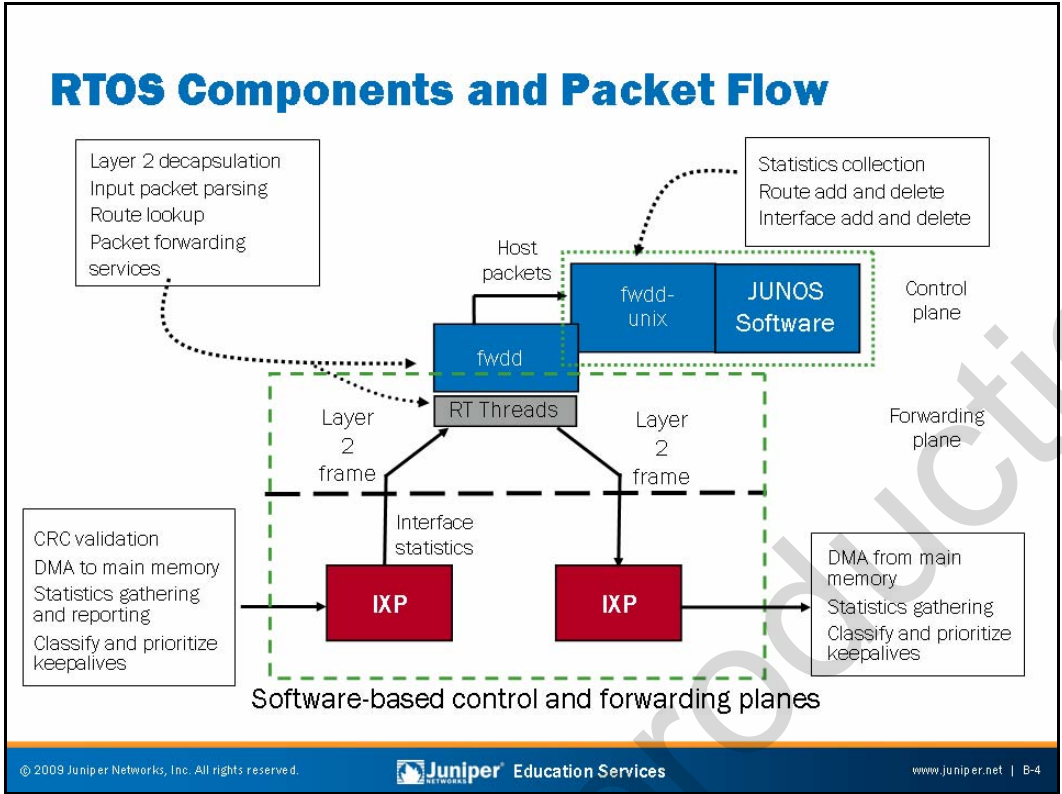
© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | B-3

### RTOS Packet Flow

The slide highlights the topics we cover in this appendix. We discuss the highlighted topic first.



## Packet Processing

The J Series Routing Engine (RE) and software Packet Forwarding Engine (PFE) are both implemented on the primary x86 architecture microprocessor. An RTOS kernel mediates access to the underlying hardware. The real-time kernel ensures that operating system services go out in a constant, load-independent, amount of time. This process ensures that the forwarding and services real-time threads deliver predictable packet forwarding performance.

While the software handles packet forwarding decisions with the virtual PFE, the Intel IXP network processors still provide performance scalability. These network processors handle Layer 2 functions such as cyclic redundancy check (CRC) validation, statistics gathering, classification, and keepalives.

An IXP network processor is included on each Physical Interface Module (PIM). As a result, overall router capability increases as you add PIMs.

## Agenda: Packet Flow Details

- RTOS Packet Flow
- ABC Chipset Packet Flow
- LMNR Chipset Packet Flow

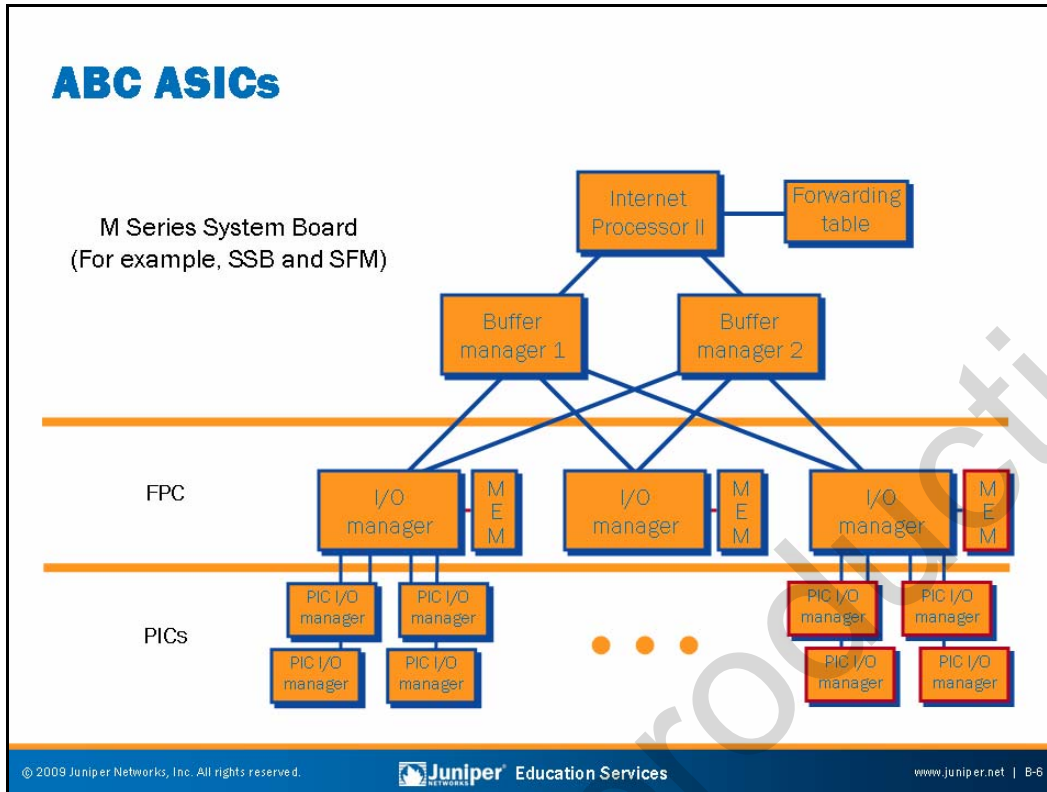
© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | B-5

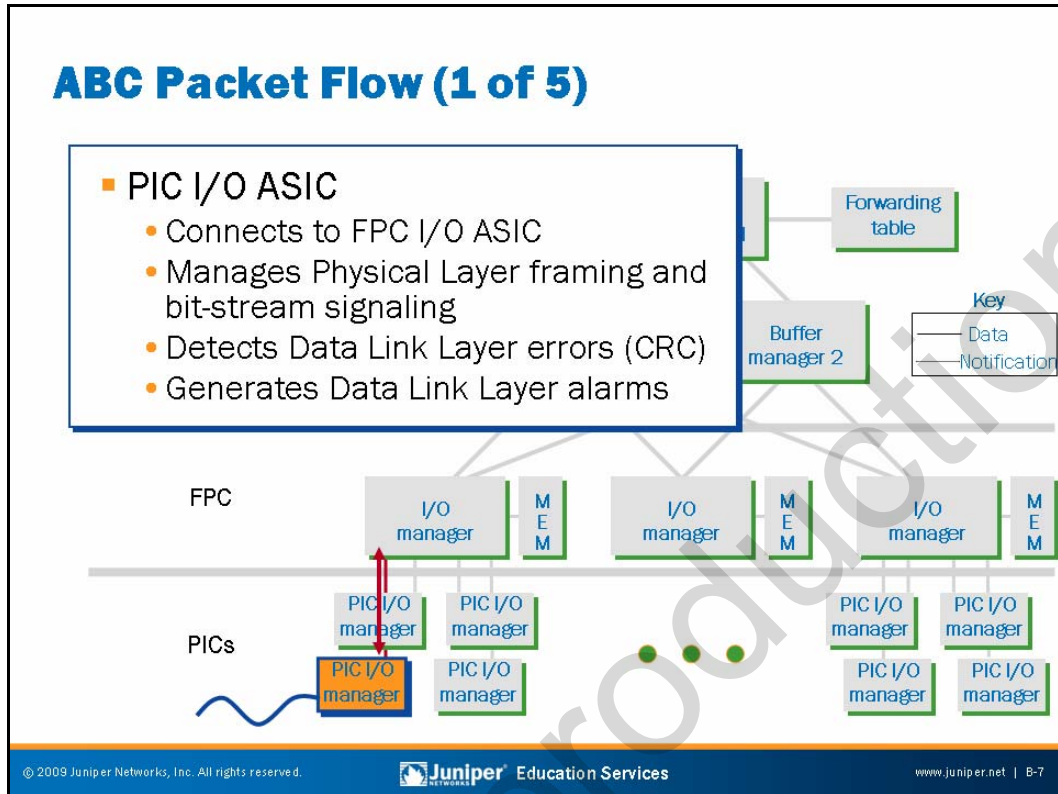
### ABC Chipset Packet Flow

The slide highlights the topic we discuss next.



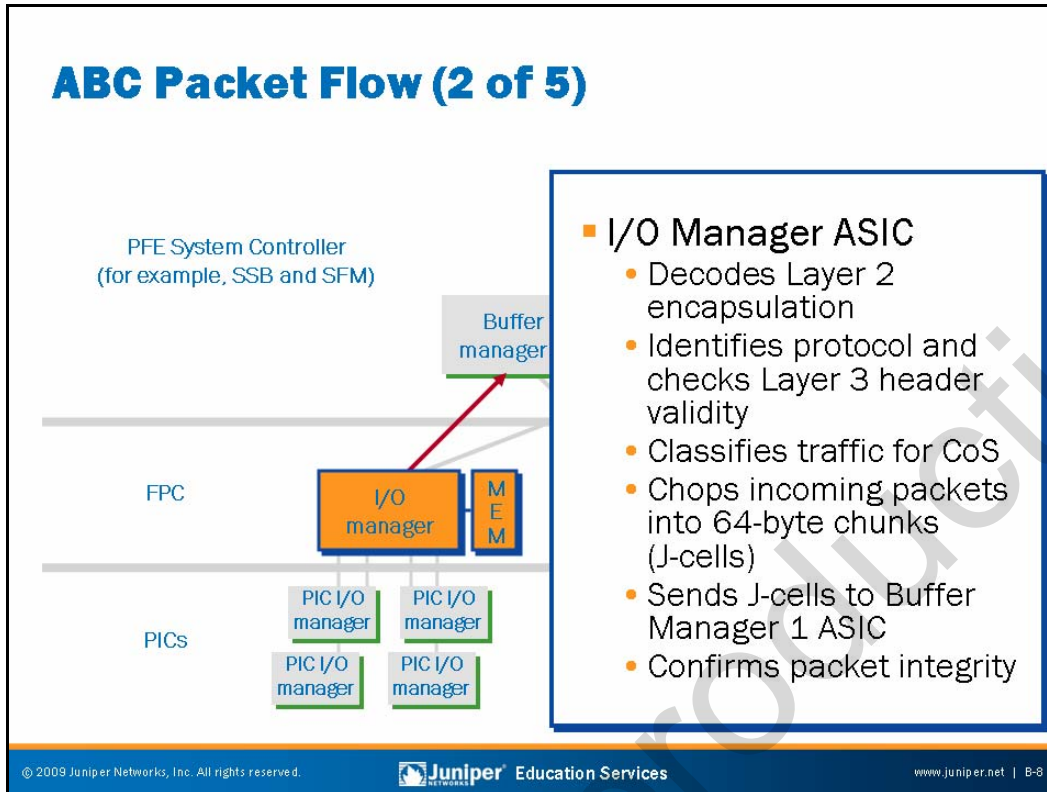
### ABC ASICs

The slide displays the application-specific integrated circuits (ASICs) that make up an ABC chipset router's PFE. We detail the function of each ASIC on subsequent slides. In ABC chipset platforms the ASICs that comprise the PFE are in the PICs, FPCs, and the System Board. The M7i and M10i routers combine Flexible PIC Concentrators (FPCs) and System Board functionality into the Compact Forwarding Engine Board (CFEB). The CFEB makes use of a combined I/O manager ASIC, Distributed Buffer Manager ASIC, and Internet Processor II ASIC to reduce cost and power consumption while also improving reliability. This ASIC set is sometimes referred to as 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.



### ABC Chipset 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 Data Link Layer FCS (CRC) verification. The PIC then passes a serial stream of bits to the I/O Manager ASIC on the FPC.



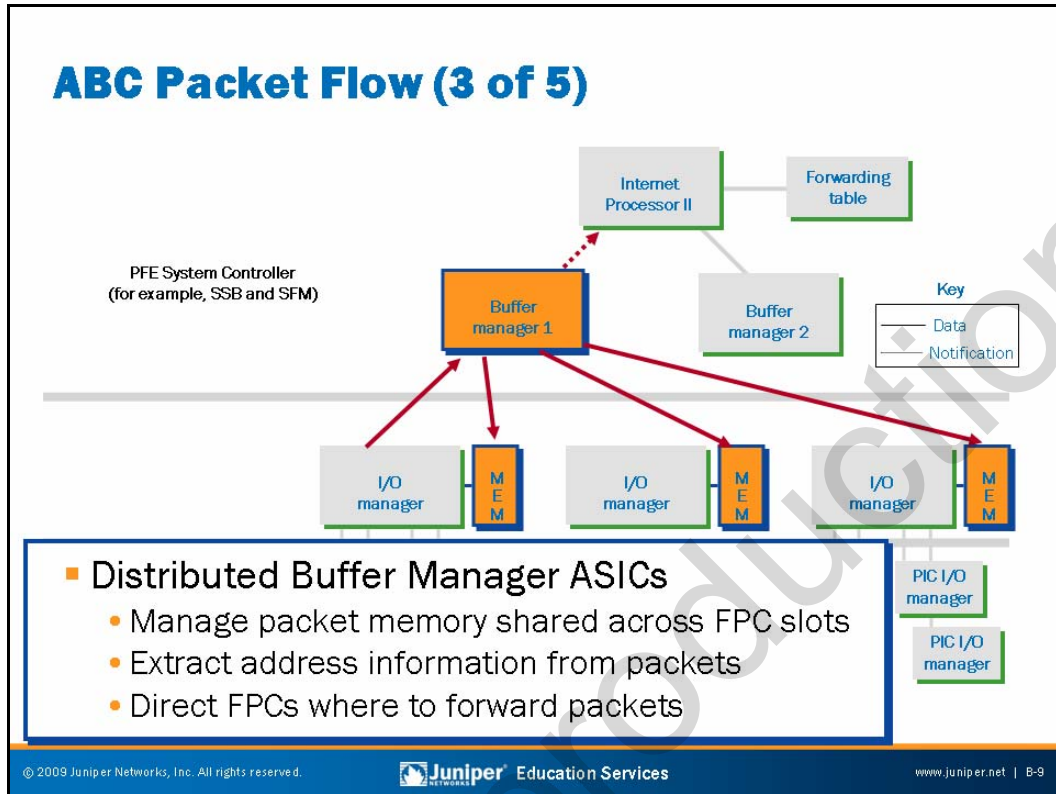
### ABC Chipset Packet Flow: Part 2

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 named *J-cells*. These J-cells then travel to the inbound Distributed Buffer Manager ASIC.

The I/O Manager ASIC also performs the following:

- 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.
- Performs behavior aggregate (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.

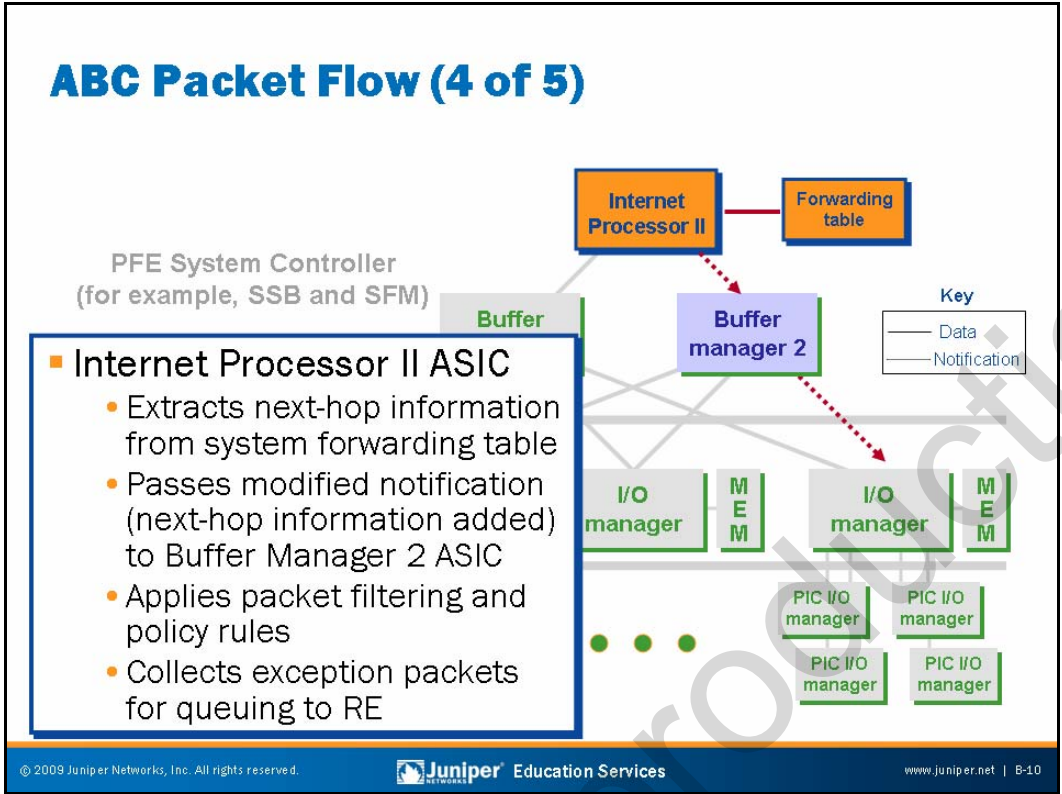




### ABC Chipset Packet Flow: Part 3

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.

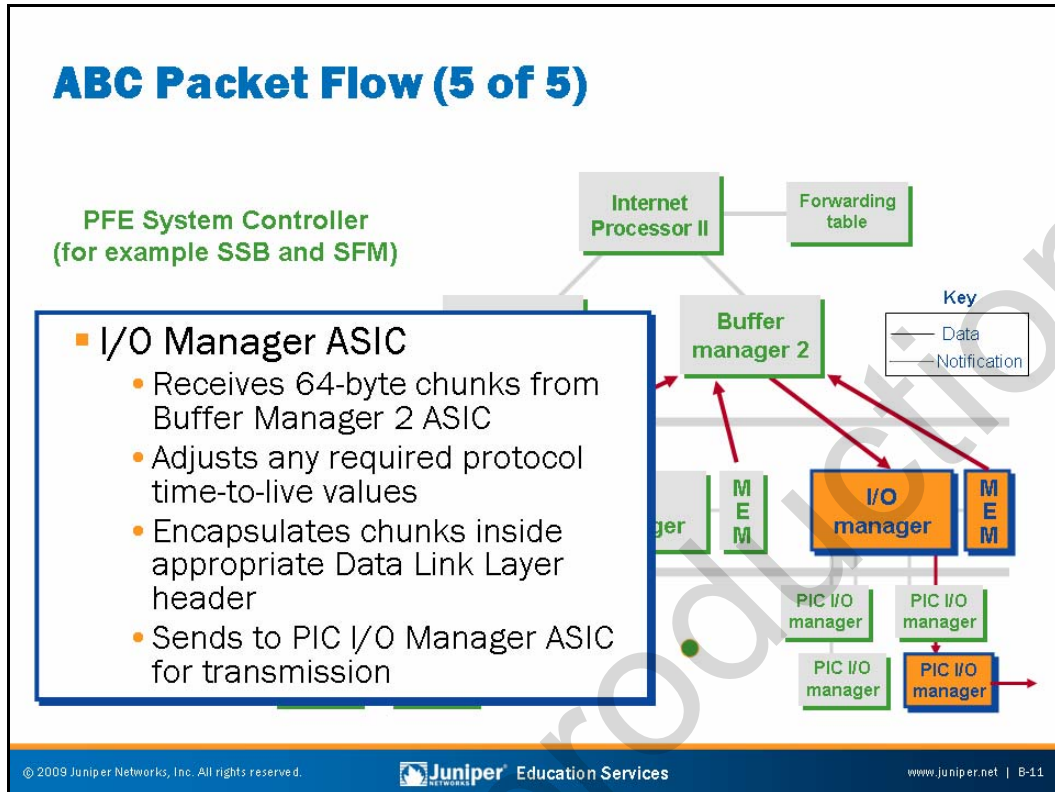


### ABC Chipset Packet Flow: Part 4

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 RE constructs the forwarding table and the JUNOS Software kernel maintains it.

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 for servicing, it issues read requests to the Buffer Manager 2 ASIC for the J-cells associated with the 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.



### ABC Chipset 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 class-of-service (CoS)-related queuing, scheduling, and congestion avoidance operations at packet egress. Note that the packet itself never queues on the FPC; rather, a pointer to the packet, in the form of a notification cell, queues on the egress FPC. Each output port on a given PIC associates 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: Packet Flow Details

- RTOS Packet Flow
- ABC Chipset Packet Flow
- LMNR Chipset Packet Flow

### LMNR Chipset Packet Flow

The slide highlights the topic we discuss next.

## LMNR PFE

- Each LMNR chipset PFE consists of the following:
  - One or more media-specific PIC ASICs
    - Handles Physical Layer signaling, alarms, and CRC processing
  - Layer 2 and Layer 3 Packet Processing ASIC
    - Provides Data 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

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | B-13

### LMNR Chipset PFE

The term Packet Forwarding Engine 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. The slide lists the ASICs associated with the LMNR chipset PFE and provides a high-level description of the function performed by each ASIC. Subsequent slides delve into the role that each ASIC plays in packet forwarding in greater detail. Note that each LMNR chipset 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.
- *Layer 2 and Layer 3 Processing ASIC:* After the PIC performs the medium-specific functions, the bit stream travels to the Layer 2 and Layer 3 processing ASIC, which removes Layer 2 encapsulation, parses the Layer 3 header, and segments the bit stream into 64-byte chunks.

*Continued on next page.*


### LMNR Chipset PFE (contd.)

- *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 LMNR chipset 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 LMNR chipset PFEs by facilitating the exchange of 64-byte chunks across the LMNR chipset cross-bar switch fabric.

Not for Reproduction

## LMNR Switch Fabric

- Nonblocking topology with any-to-any connectivity
- No single point of failure, all SIBs are fully redundant
  - Graceful degradation for multiple failures
    - T640 switch fabric consists of 5 SIBs (5th is a spare)
    - T320 switch fabric consists of 3 SIBs (3rd is a spare)
    - M320 fabric consists of four active SIBs
  - Packet order and CoS maintained across fabric



© 2009 Juniper Networks, Inc. All rights reserved.

Juniper Education Services

www.juniper.net | B-15

### The LMNR Chipset Switch Fabric

LMNR chipset platforms use a nonblocking cross-bar switch fabric to switch traffic between the system's FPCs. The Switch Interface Board (SIB) instantiates the switch fabric and 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 I/O 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) occurs between that FPC and SIB. Each FPC connects to multiple SIBs to provide the speed-up needed for a nonblocking switch fabric and for redundancy reasons.

LMNR chipset 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 Layer 2 and Layer 3 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.*

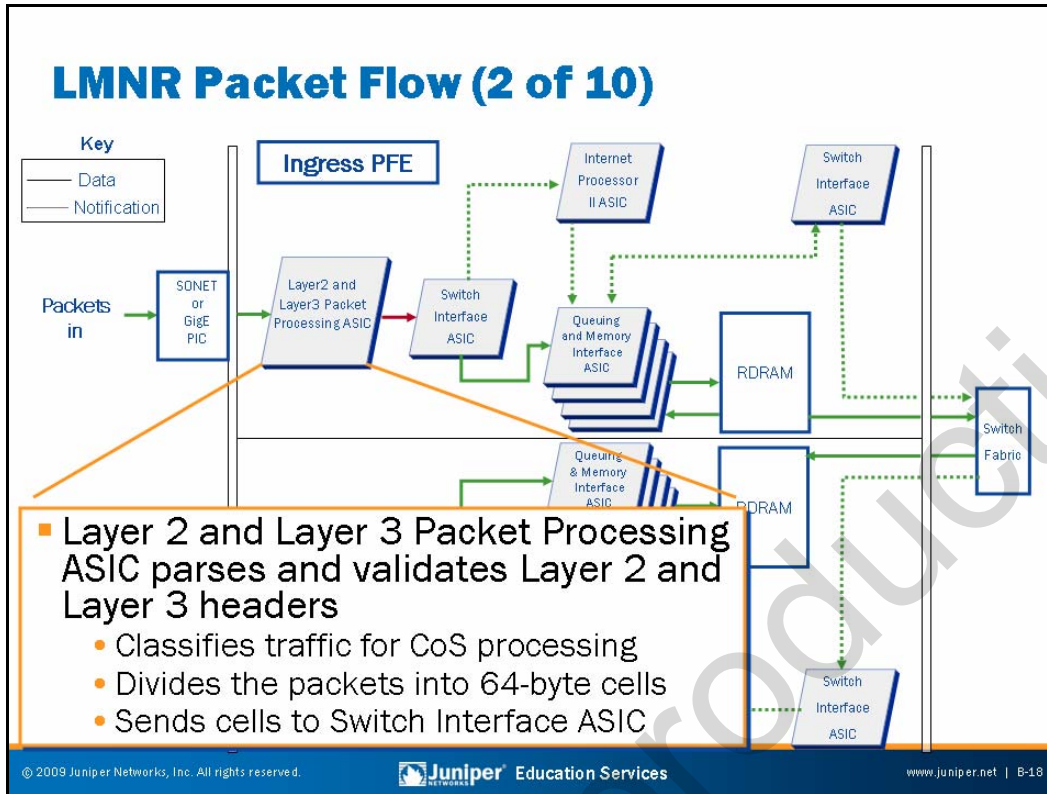
## Redundant Fabric

The slide illustrates the specifics of a T320 platform's switch fabric. In this case, each FPC (or PFE) has four HSL connections to both SIB 1 and SIB 2. This switch fabric provides the T320 FPC with 40 Gbps of aggregate capacity. To accommodate SIB failures, each T320 FPC also connects 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 SIB 1 or SIB 2 failure. However, because each FPC interconnects to SIB 0 through a single HSL, switch fabric speedup reduces. The reduction in speedup results in a graceful degradation of the T320 platform'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 attach to all SIBs using two HSLs. The result is that a T640 platform's switch fabric remains nonblocking despite the presence of a possible 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 platform in that the failure of one of its four SIBs results in graceful degradation of forwarding capacity.





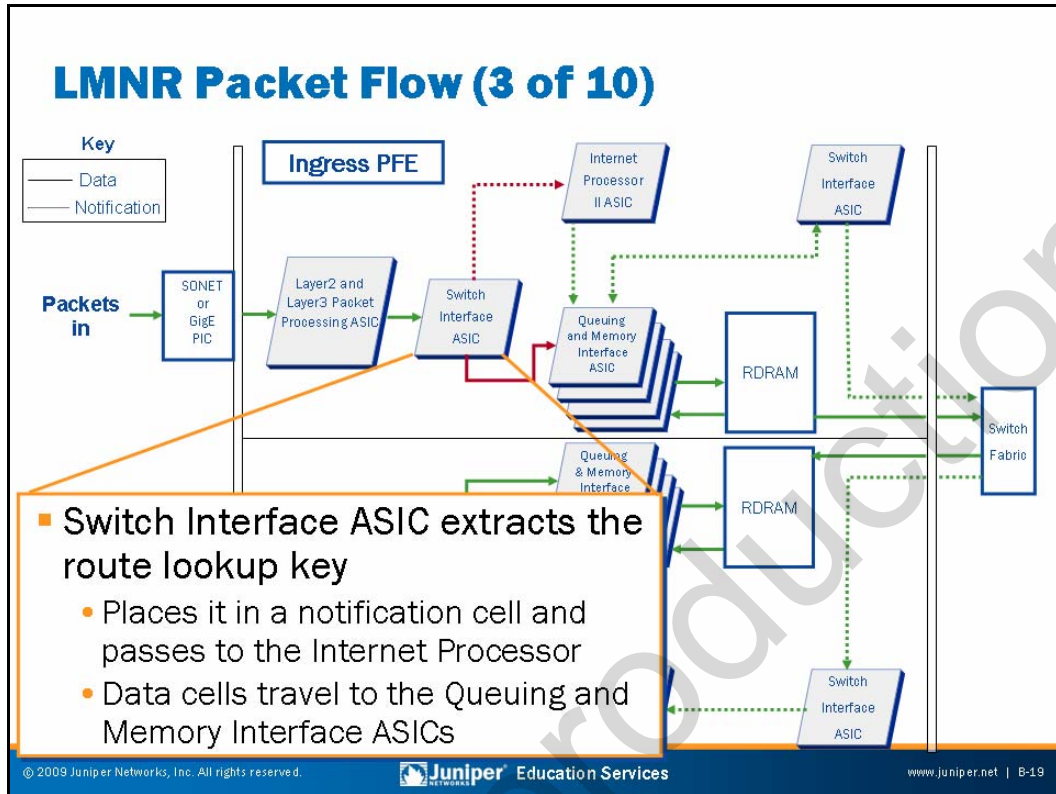


### LMNR Chipset Packet Flow: Part 2

The Layer 2 and Layer 3 Packet Processing ASIC performs Layer 2 and Layer 3 parsing. The Layer 2 and Layer 3 ASIC also divides the packets into 64-byte J-cells. The J-cells travel to the Switch Interface ASIC.

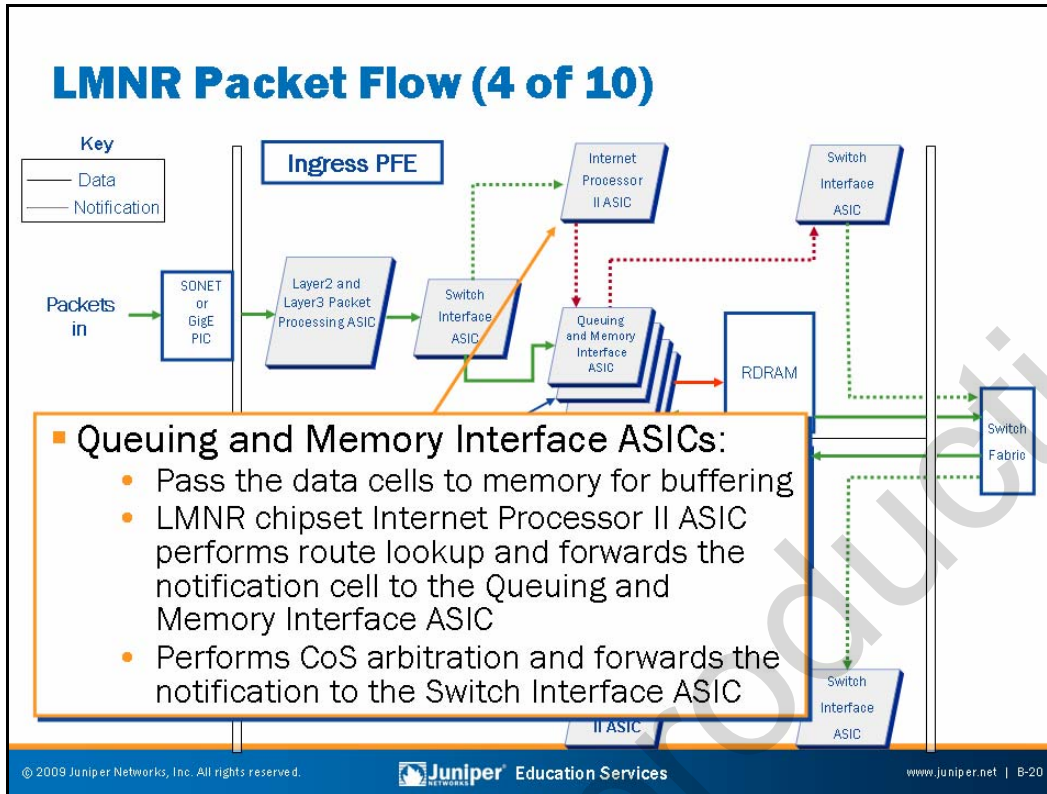
Errors detected during the Layer 2 and Layer 3 parsing steps or when the Layer 2 and Layer 3 processing ASIC receives an indication from the PIC that the received frame is corrupt results in error counter increments and an effective *no-op* flag for any J-cells relating to the corrupted frame still housed in shared memory.

The Layer 2 and Layer 3 Processing ASIC also performs 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.



### LMNR Chipset 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 LMNR chipset 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 LMNR chipset PFE. Note that the shared memory fabric facilitates the switching of packets within a specific PFE complex, such as the switching that occurs when the source and destination PICs share a PFE.



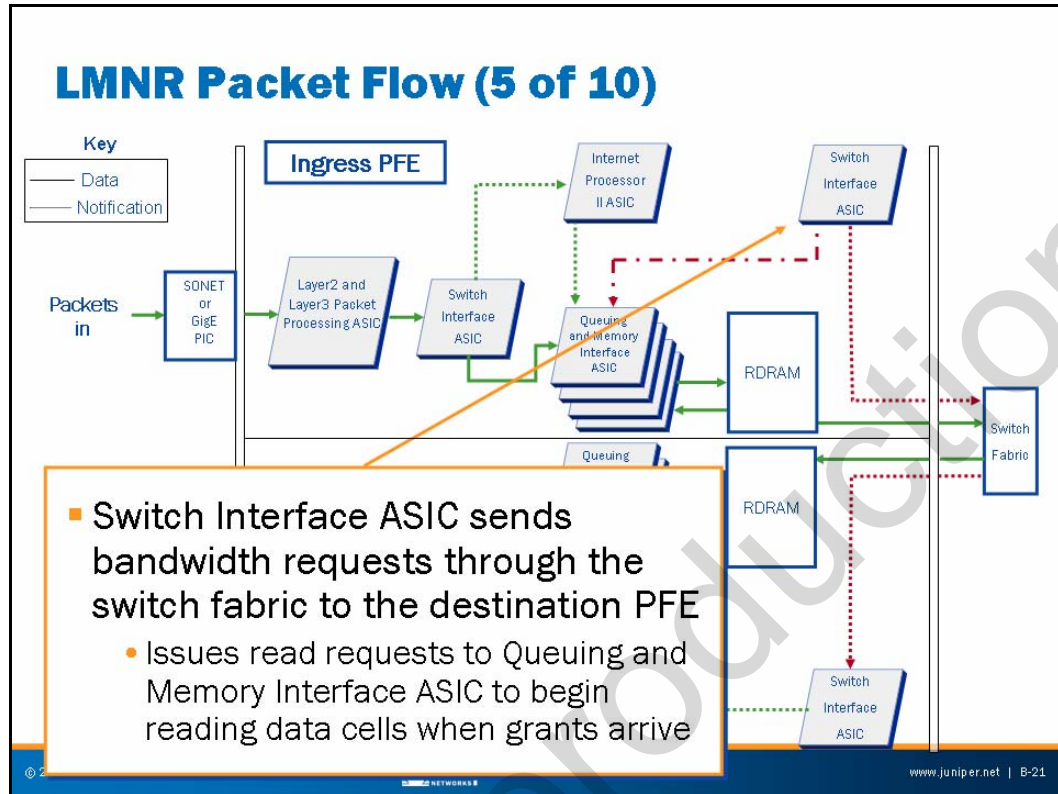
#### LMNR Chipset Packet Flow: Part 4

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 used only to exchange packets *between* PFE complexes.

While the J-cells write into shared memory, the LMNR chipset Internet Processor II ASIC performs a route lookup operation on the key data. The modified notification cell then forwards to the Queuing and Memory Interface ASIC.

In addition to queuing notifications, the Queuing and Memory Management ASIC also performs the following 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.
- 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.

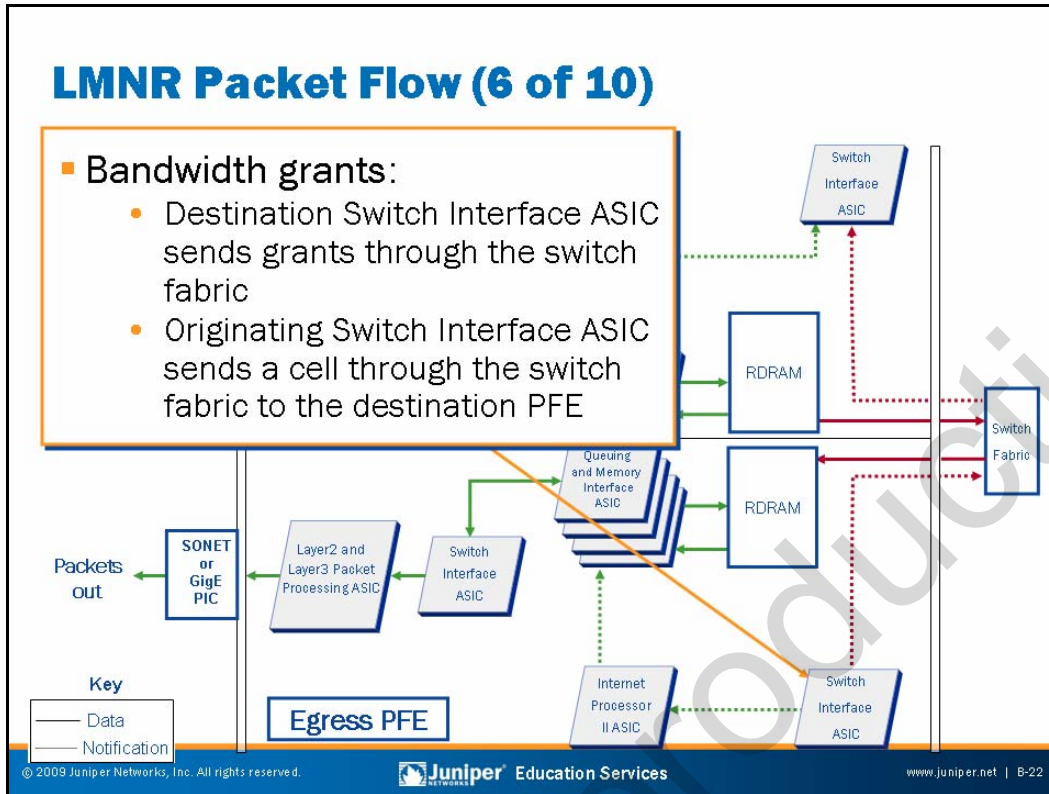


### LMNR Chipset Packet Flow: Part 5

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 PFE. In this case, the notification travels to the Switch Interface ASIC that faces the Layer 2 and 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.





### LMNR Chipset Packet Flow: Part 6

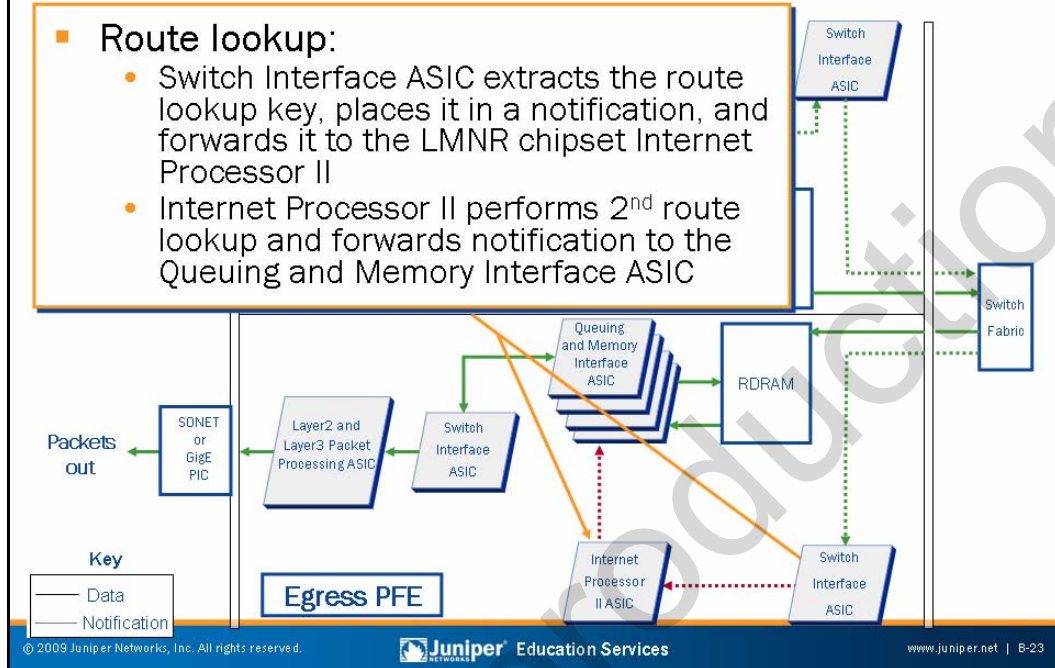
The destination Switch Interface ASIC returns bandwidth grants through the switch fabric to the originating Switch Interface ASIC in response to received bandwidth requests.

Upon receipt of each bandwidth grant, the originating Switch Interface ASIC sends a cell through the switch fabric to the destination PFE.

## LMNR Packet Flow (7 of 10)

- Route lookup:

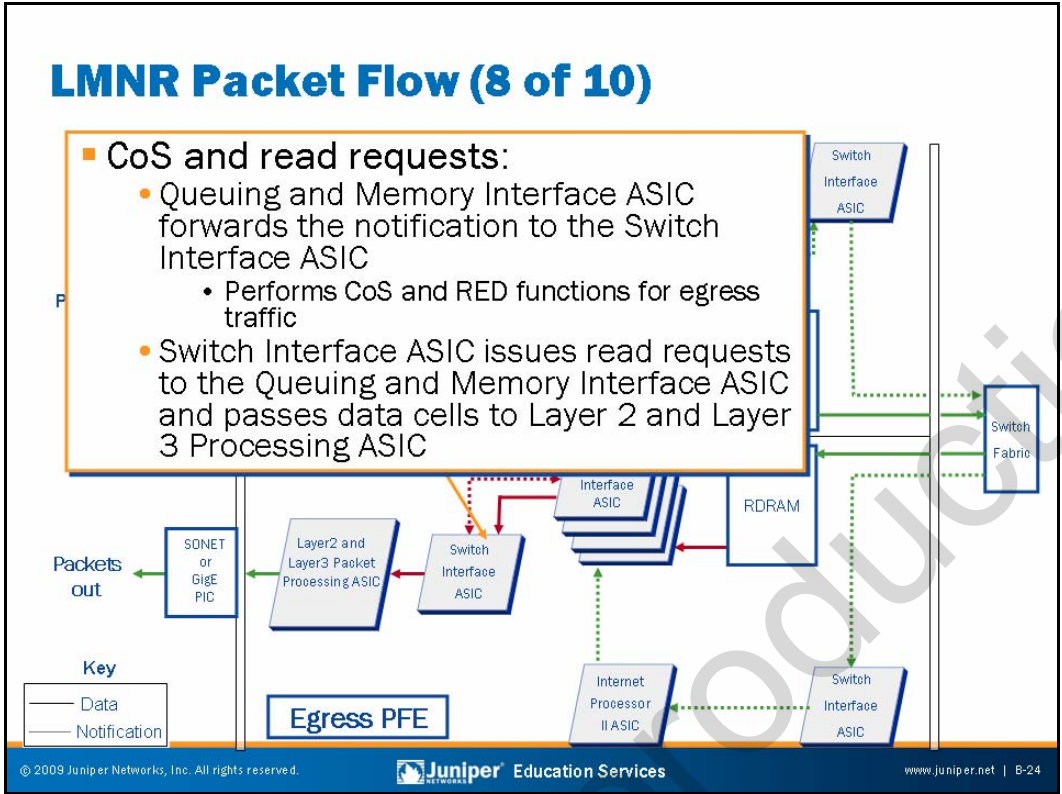
- Switch Interface ASIC extracts the route lookup key, places it in a notification, and forwards it to the LMNR chipset Internet Processor II
- Internet Processor II performs 2<sup>nd</sup> route lookup and forwards notification to the Queuing and Memory Interface ASIC



### LMNR Chipset 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, because the memory locations for each chunk varies by PFE.

The LMNR chipset 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.



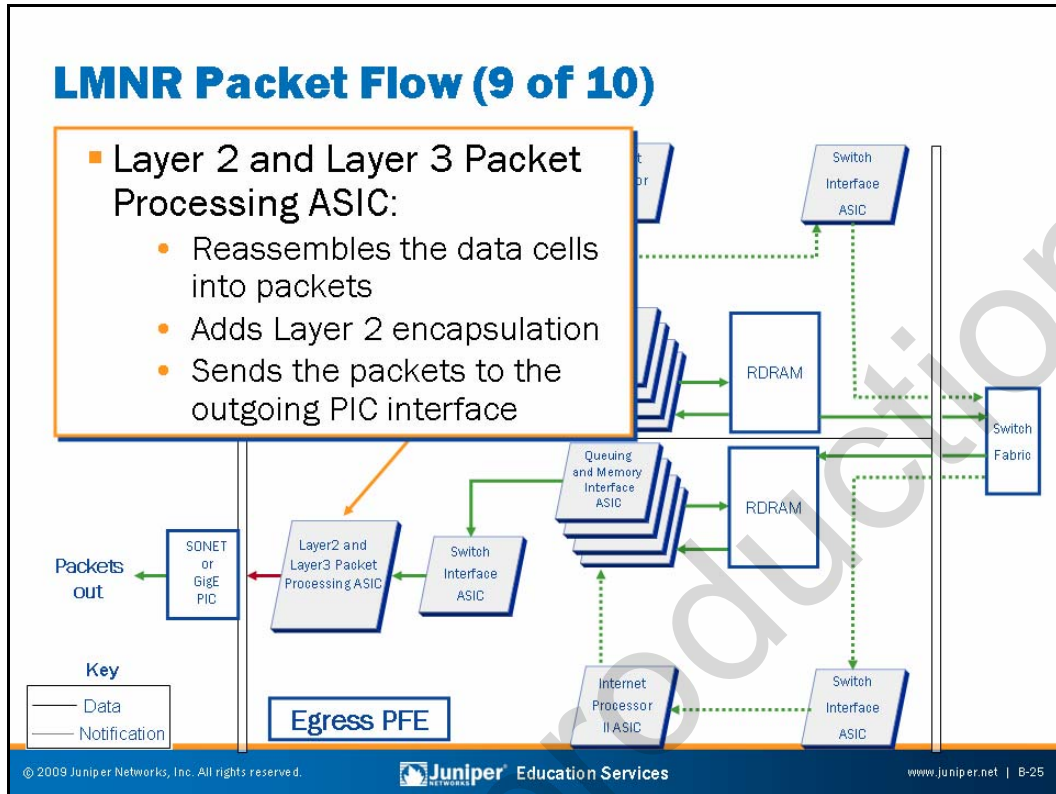
### LMNR Chipset Packet Flow: Part 8

The Queuing and Memory Interface ASIC is responsible for performing 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, 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 receive service.

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 and Layer 3 Packet Processing ASIC.





### LMNR Chipset Packet Flow: Part 9

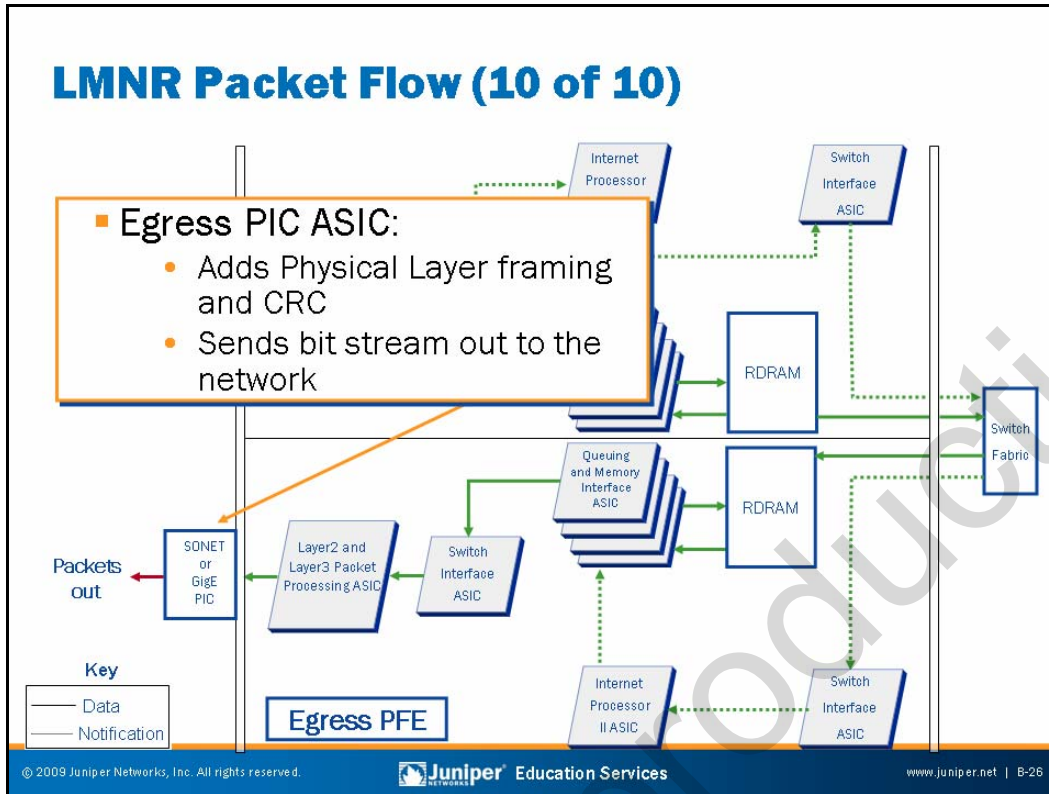
The Layer 2 and Layer 3 Packet Processing ASIC reassembles the data cells into packets. The Layer 2 and Layer 3 processing ASIC then adds appropriate Layer 2 encapsulation and sends the resulting bit stream to the egress PIC.

The Layer 2 and 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 associates 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 Layer 2 and Layer 3 Processing ASIC on the egress FPC. The Layer 2 and Layer 3 Processing ASIC implements the RED algorithm during egress processing to avoid tail drops and the resulting risk of global synchronization of TCP retransmissions.

The LMNR chipset Switch Interface ASICs handle switch fabric queuing and prioritization to extend CoS across the LMNR chipset switch fabric.

A full coverage of JUNOS Software CoS capabilities is beyond the scope of this class.



### LMNR Chipset Packet Flow: Part 10

The final steps in egress packet processing begin 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 cyclic redundancy check (CRC) to the frame as needed for each particular medium.

## Exception Packets

- Exception packets
  - Local delivery
  - IP options
    - Source route, router alert, and so forth
  - ICMP message generation
- Generally processed by PFE control CPU
  - Remaining traffic (local and control) sent to RE through an internal link
    - Rate limiting
    - Hardware-based WRR ensures control traffic does not starve

© 2009 Juniper Networks, Inc. All rights reserved.

 Juniper Education Services

www.juniper.net | B-27

### Exception Packets

Exception packets require some form of special handling. Examples of exception traffic include the following:

- 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 PFE was purposely designed not to handle IP options. They must travel to the RE for processing.
- Traffic that requires the generation of Internet Control Message Protocol (ICMP) messages. ICMP messages travel 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 no entry exists 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 decrements to zero. In most cases, the PFE process handles the generation of ICMP messages.

*Continued on next page.*

## Packet Forwarding Engine CPU

The Internet Processor II ASIC passes exception packets to the microprocessor on the PFE Control Board, which in turn processes almost all of them. Certain exception packets also travel to the RE for further processing. Exception traffic destined for the RE travels over the 100 Mbps fxp1 interface. Exception traffic is rate-limited by the PowerPC processor to protect the RE 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 percent of the fxp1 interface's bandwidth through hardware-based weighted round-robin (WRR) queueing.

Not for Reproduction

## Summary

- In this chapter, we:
  - Described the RTOS packet flow
  - Described the ABC chipset packet flow
  - Described the LMNR chipset packet flow

### This Appendix Discussed:

- RTOS packet flow;
- ABC chipset packet flow; and
- LMNR chipset packet flow.

Not for Reproduction

# Appendix C: Acronym List

---

AAL5	ATM Adaptation Layer 5
ADM	add/drop multiplexer
AIS	alarm indication signal
AMI	alternate mark inversion
ANSI	American National Standards Institute
APS	Automatic Protection Switching
ARP	Address Resolution Protocol
AS PIC	Adaptive Services PIC
ASIC	application-specific integrated circuit
ATM	Asynchronous Transfer Mode
BERT	bit error rate test
BPV	bipolar violation
CB	Control Board
CDP	Cisco Discovery Protocol
CE	customer edge
CFEB	Compact Forwarding Engine Board
CFM	cubic feet per minute
CIP	Connector Interface Panel
CLI	command-line interface
CoS	class of service
CoS	class-of-service
CPE	customer premises equipment
CRC	cyclic redundancy check
CSC	Customer Support Center
DCE	data circuit-terminating equipment
DLCI	data-link connection identifier
DoS	denial of service
DPC	Dense Port Concentrator
DTE	data terminal equipment
DWDM	dense wavelength-division multiplexing
EMI	electromagnetic interference
ESD	electrostatic discharge
ESF	extended superframe
FCS	frame check sequence
FEAC	far-end alarm and control
FEB	Forwarding Engine Board
FIC	Fixed Interface Card
FPC	Flexible PIC Concentrator
FPGA	field-programmable gate array
FRU	field-replaceable unit
GRE	generic routing encapsulation
HA	high availability
ICMP	Internet Control Message Protocol
IDP	initial domain part
IGP	interior gateway protocol
ILMI	Integrated Local Management Interface
IOC	input/output card
JNTCP	Juniper Networks Technical Certification Program
JTAC	Juniper Networks Technical Assistance Center
LCP	Link Control Protocol

LMI	Local Management Interface
LOF	loss of frame
LOS	loss of signal
LSA	link-state advertisement
MAC	media access control
MCS	Miscellaneous Control Subsystem
mgd	management process
MLFR	Multilink Frame Relay
MLPPP	Multilink Point-to-Point Protocol
MTU	maximum transmission unit
NAT	Network Address Translation
NBMA	nonbroadcast multiaccess
NCP	Network Control Protocol
NSR	nonstop active routing
NTP	Network Time Protocol
OAM	Operation, Administration, and Maintenance
PCD	protocol control block
PCG	Packet Forwarding Engine Clock Generator
PE	provider edge
PEM	Privacy-Enhanced Mail
PFE	Packet Forwarding Engine
PIM	Physical Interface Module
PoE	Power over Ethernet
POP	point of presence
PPP	Point-to-Point Protocol
PVC	permanent virtual circuit
RDI	remote defect indication
RE	Routing Engine
RED	random early detection
REI	remote error indication
RMA	Return Materials Authorization
rpd	routing protocol process
SCB	System Control Board
SCG	SONET Clock Generator
SFM	Switching and Forwarding Module
SFP	small form-factor pluggable transceiver
SIB	Switch Interface Board
SPC	services processing card
SPMB	Switch Processor Mezzanine Board
SSB	System and Switch Board
STP	Spanning Tree Protocol
TNP	Trivial Network Protocol
ToS	type of service
TTL	time to live
UNI	user-to-network interface
USB	universal serial bus
UTM	Unified Threat Management
VCI	virtual channel identifier
VLAN	virtual LAN
VPI	virtual path identifier
VPN	virtual private network
VRRP	Virtual Router Redundancy Protocol
WRR	weighted round-robin
XFP	10-gigabit small form-factor pluggable transceiver



# Appendix D: Answer Key

---

## Chapter 1: Course Introduction

This chapter does not contain any review questions.

## Chapter 2: Overview of JUNOS Platforms

1.

To safely power off an M Series router, you must perform a graceful shutdown of JUNOS Software before removing power. Use the **request system reboot** command, which causes the device to reboot.

2.

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 inside a protected memory environment. The PFE is responsible for forwarding transit packets through the router using an ASIC-based switching path.

3.

The Craft Interface is the collection of mechanisms on some JUNOS platforms that allow 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 and offline buttons. On supported platforms the Craft Interface includes an LCD screen that provides status reporting for the entire system.

4.

Interface name at-0/1/1.100 references a logical unit 100 of the ATM interface in FPC 0, PIC 1, and port 1.

## Chapter 3: Troubleshooting Tool Kit for JUNOS Platforms

1.

Because modern communications networks are complex, troubleshooting them can be challenging. To alleviate this challenge, we recommend a layered troubleshooting approach. Such an approach utilizes the OSI model to identify prospective problems, isolate the likely causes of those problems, and then systematically eliminate each potential cause. By conducting tests that accurately isolate a symptom to the root-cause layer, you avoid wasting time testing layers that are not at fault.

2.

The three ways in which the CLI can help to perform fault analysis are: (1) deployment of key commands; (2) process restart and bringing hardware online and offline; and (3) deployment of network and diagnostic utilities.

## Chapter 3: Troubleshooting Tool Kit for JUNOS Platforms (contd.)

3.

The two good reasons to escape to an interactive shell are: (1) to access standard utilities and programs, enabling experienced UNIX users to perform advanced troubleshooting tasks; and (2) the ability to establish a connection to embedded hosts within the PFE to perform complex diagnostics.

4.

You can reset the OSPF process without affecting other routing protocols by deactivating OSPF in the configuration mode and then committing the configuration file. This approach requires you to have configuration privileges.

## Chapter 4: JUNOS Platforms Hardware Troubleshooting

1.

The correct procedure for powering off a JUNOS platform is as follows: First shutdown the JUNOS Software using the CLI `request system halt` command. Next, turn off the power supplies.

2.

You can use the JUNOS Software CLI to display information about the chassis and the PFE using the `show chassis ...` and `show pfe ...` commands to display information about the system and software processes using the `show system ...` commands and to display system log files.

3.

The two ways of determining if chassis alarms are present are using the `show chassis alarms` command; and the `show chassis craft-interface` command.

4.

The following command searches the `messages` file for all lines matching `fail` and `error`:  
`show log messages | match fail | match error.`

## Chapter 5: Interface Troubleshooting

1.

When you deactivate an interface, JUNOS Software completely ignores that specific interface and does not apply it when you issue a `commit` command. When you disable an interface, it activates when you issue a `commit` command but is treated as being down or administratively disabled.

2.

You can check the results of your BERT test using the `show interfaces extensive` command.

3.

When troubleshooting T3 or E3 interface problems you must ensure that the two ends of the circuit have the following compatible parameters: clocking, frame checksum, HDLC payload scrambling (if using Cisco HDLC encapsulation), T3 line buildout, and T3 C-bit parity mode.

## Chapter 5: Interface Troubleshooting (contd.)

4.

You use the **monitor interface interface-name** command to display real-time statistics about a physical interface. The output updates every second. The output of this command also shows the amount that each field changes from the time you start the command or from the time you clear the counters by using the c key.

## Chapter 6: JTAC Processes, Guidelines, and Support Resources

1.

You need a chassis serial number when opening a case. The serial number will help JTAC staff to determine the support status.

2.

When you access the CSC, you have a wealth of technical support information in the form of the JTAC Knowledge Base, the PR database, technical bulletins, and white papers that provide configuration examples and technology primers.

3.

The I2J tool converts Cisco IOS configuration into a JUNOS Software configuration. It is a good idea to inspect the output, ensuring that you obtained the desired results.

4.

You should use FTP to transfer files to JTAC when the size of those files is larger than 10 MB.

Not for Reproduction