Cisco Reader Comment Card

General Information

1 Years of networking experience ______ Years of experience with Cisco products ______

2 I have these network types: □ LAN □ Backbone □ WAN
□ Other: ____________________

3 I have these Cisco products: □ Switches □ Routers
□ Other: Specify model(s) ____________________

4 I perform these types of tasks: □ H/W Install and/or Maintenance □ S/W Config
□ Network Management □ Other: ____________________

5 I use these types of documentation: □ H/W Install □ H/W Config □ S/W Config
□ Command Reference □ Quick Reference □ Release Notes □ Online Help
□ Other: ____________________

6 I access this information through: ____% Cisco.com ____% CD-ROM
____% Printed docs ____% Other: ____________________

7 Which method do you prefer? ____________________

8 I use the following three product features the most:

Document Information

Document Title: *ATM and Layer 3 Switch Router Troubleshooting Guide*

Part Number: 78-11614-01

On a scale of 1–5 (5 being the best) please let us know how we rate in the following areas:

_____ The document was written at my technical level of understanding.

_____ The information was accurate.

_____ The document was complete.

_____ The information I wanted was easy to find.

_____ The information was well organized.

_____ The information I found was useful to my job.

Please comment on our lowest score(s):

Mailing Information

Company Name Date

Contact Name Job Title

Mailing Address

City State/Province ZIP/Postal Code

Country Phone ( ) Extension

Fax ( ) E-mail

Can we contact you further concerning our documentation? □ Yes □ No

You can also send us your comments by e-mail to **bug-doc@cisco.com**, or fax your comments to us at **(408) 527-8089**.

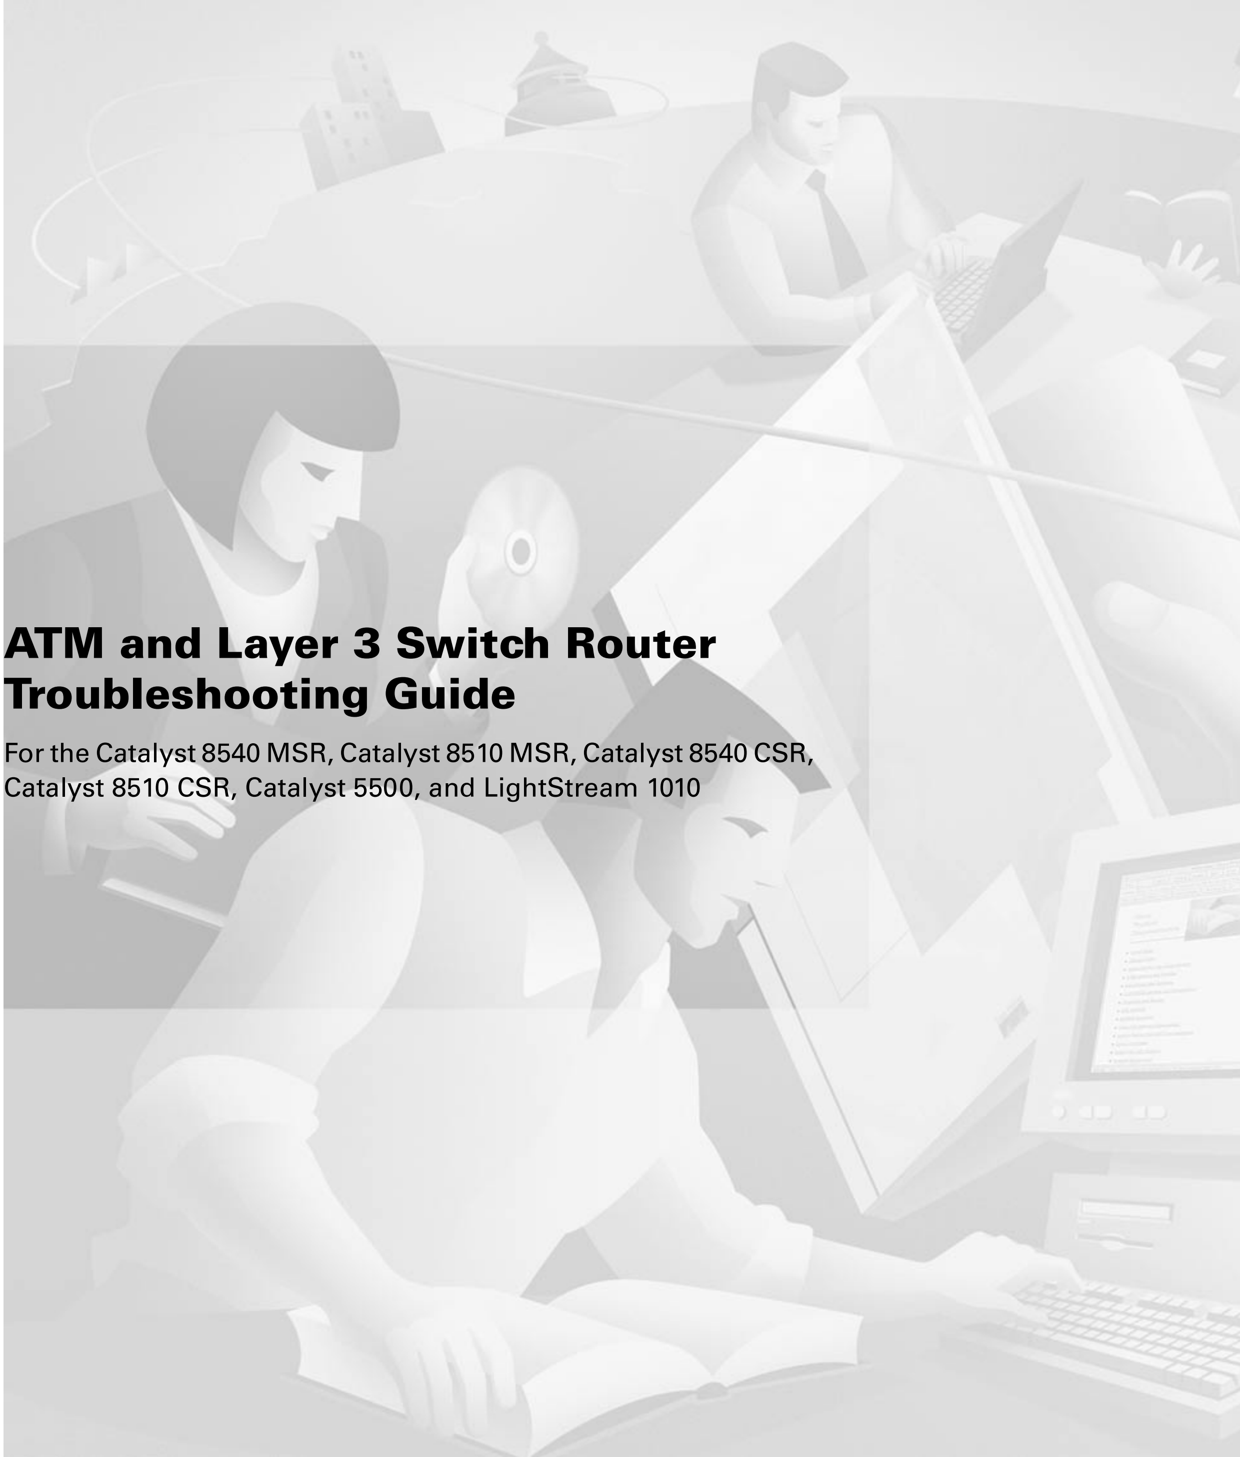

ATM and Layer 3 Switch Router Troubleshooting Guide

For the Catalyst 8540 MSR, Catalyst 8510 MSR, Catalyst 8540 CSR, Catalyst 8510 CSR, Catalyst 5500, and LightStream 1010

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Customer Order Number: DOC-7811614=
Text Part Number: 78-11614-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

ATM and Layer 3 Switch Router Troubleshooting Guide

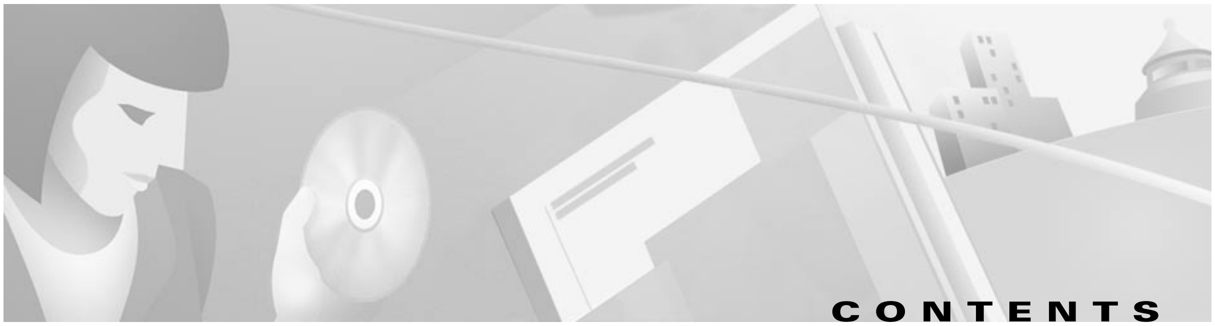

CONTENTS

CHAPTER 4

PART 1

ATM-to-ATM Connection Troubleshooting

CHAPTER 5

PART 2 Layer 3-to-Layer 3 Connection Troubleshooting

CHAPTER 11

PART 4 Appendixes

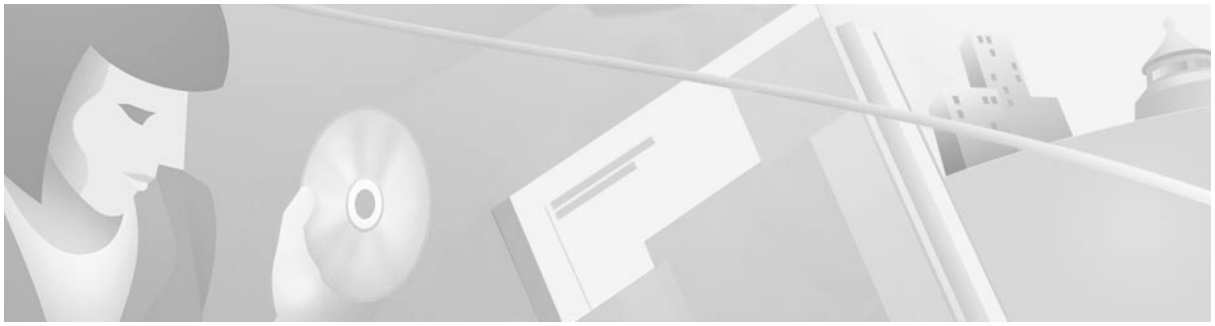

Preface

This preface describes the purpose, audience, organization, and conventions for the *ATM and Layer 3 Switch Router Troubleshooting Guide* and provides information on how to obtain related documentation.

Purpose

Failures in internetworks are characterized by certain symptoms. These symptoms might be general (clients that are unable to access specific servers) or more specific (routes that are not in the routing table). Each symptom can be traced to one or more problems or causes by using specific troubleshooting tools and techniques. Once you know what the problem is, you can take steps to fix it.

The goal of this guide is to help you isolate and resolve the most common connectivity and performance problems with your Cisco switch router. This guide describes how to define symptoms, identify problems, and implement solutions in Cisco switch router environments. This guide does not describe troubleshooting router connections and configurations. For router troubleshooting refer to the *Internetwork Troubleshooting Guide.*

This preface describes who should read the *ATM and Layer 3 Switch Router Troubleshooting Guide*, how it is organized, and its document conventions.

Audience

This publication is a stand-alone document for experienced network administrators responsible for configuring and maintaining the ATM and Layer 3 switch router.

Administrators should have hands-on experience in configuring, administering, and troubleshooting a network, should know how to configure routers, switches, and other internetwork devices, and should be familiar with the protocols and media that their hardware supports. Awareness of the basic topology of their network is also essential.

Organization

The major sections of this guide are as follows:

Chapter	Title	Description
Chapter 1	Troubleshooting Overview	Contains an overview of the switch router troubleshooting features and functions
Chapter 2	Troubleshooting Tools	Describes the wide variety of tools available to assist you in troubleshooting your internetwork
Chapter 3	Initial Troubleshooting	Describes the first steps you should take when you start troubleshooting your switch router
Chapter 4	Example Network	Describes the example network used to illustrate the hardware and configuration troubleshooting problems throughout this guide
Part 1, ATM-to-ATM Connection Troubleshooting		
Chapter 5	Troubleshooting Switch Router ATM Interface Connections	Presents troubleshooting information for connectivity and performance problems of physical interfaces of switch routers
Chapter 6	Troubleshooting Switch Router ATM Network Connections	Presents troubleshooting information for connectivity and performance problems in ATM switching network connections
Chapter 7	Troubleshooting LAN Emulation Switching Environments	Presents troubleshooting information for connectivity and performance problems in LAN emulation switching environments
Chapter 8	Troubleshooting Tag Switching Connections	Presents troubleshooting information for connectivity and performance problems in tag switching environments
Chapter 9	Troubleshooting CES Connections and Network Clocking	Presents troubleshooting information for connectivity problems in circuit emulation service (CES) environments and network clocking
Part 2, Layer 3-to-Layer 3 Connection Troubleshooting		
Chapter 10	Troubleshooting Ethernet, ATM Uplink, and POS Uplink Interfaces	Presents troubleshooting information about connectivity and performance problems in the Ethernet physical interfaces of a switch router.
Chapter 11	Troubleshooting Layer 3 Network Connections	Presents troubleshooting information about connectivity and performance problems in the Layer 3 network connections of the Layer 3 enabled ATM switch router
Chapter 12	Troubleshooting Layer 2 Interfaces	Presents troubleshooting information about connectivity and performance problems in the Layer 2 network connections of an ATM switch router

Chapter	Title	Description
Part 3, Layer 3-to-ATM Connection Troubleshooting		
Chapter 13	Troubleshooting ATM Router Module Connections	Presents troubleshooting information about connectivity and performance problems in the ATM router module
Part 4, Appendixes		
Appendix A	Debugging a Switch Router	Describes helpful **debug** commands to use when troubleshooting your switch router
Appendix B	Troubleshooting TACACS+ and Recovering Passwords	Presents troubleshooting information relating to security implementations
Appendix C	ATM Cell Structures	Describes the various ATM cell types and their configuration that can be helpful when troubleshooting your switch router
Appendix D	Creating a Core Dump	Describes procedures used to obtain a full copy of the memory image (or core dump) to identify the cause of a crash
Appendix E	Technical Support	Describes the process used to contact and provide your technical support representative with the information about the symptoms and the problem

Related Documentation

Use the following books as supplements to this guide:

- *Internetwork Troubleshooting Guide*
- *Debug Command Reference*
- *Guide to ATM Technology*
- *ATM Switch Router Software Configuration Guide*
- *ATM Switch Router Command Reference*
- *Layer 3 Switching Software Feature and Configuration Guide*
- *Catalyst 8540 CSR Route Processor and Interface Module Installation Guide*
- *Site Preparation and Safety Guide*
- *Catalyst 8540 Chassis Installation Guide*
- *Hardware Installation Guide (Catalyst 8510 MSR and LightStream 1010)*
- *Processor Installation Guide (Catalyst 8510 MSR and LightStream 1010)*
- *ATM Port Adapter and Interface Module Installation Guide*
- *Configuration Fundamentals Command Reference*

Conventions

This document uses the following conventions:

Convention	Description
boldface font	Commands and keywords are in **boldface**.
italic font	Arguments for which you supply values are in *italics*.
[ ]	Elements in square brackets are optional.
{ x | y | z }	Alternative keywords are grouped in braces and separated by vertical bars.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
`screen` font	Terminal sessions and information the system displays are in `screen` font.
`boldface screen` font	Information you must enter is in **`boldface screen`** font.
italic screen font	Arguments for which you supply values are in *italic screen* font.
➔	This pointer highlights an important line of text in an example.
^	The symbol ^ represents the key labeled Control—for example, the key combination ^D in a screen display means hold down the Control key while you press the D key.
< >	Nonprinting characters, such as passwords are in angle brackets.

Notes use the following convention:

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the publication.

Cautions use the following convention:

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Obtaining Documentation

The following sections provide sources for obtaining documentation from Cisco Systems.

World Wide Web

You can access the most current Cisco documentation on the World Wide Web at the following sites:

- http://www.cisco.com
- http://www-china.cisco.com
- http://www-europe.cisco.com

Documentation CD-ROM

Cisco documentation and additional literature are available in a CD-ROM package, which ships with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or as an annual subscription.

Ordering Documentation

Cisco documentation is available in the following ways:

- Registered Cisco Direct Customers can order Cisco Product documentation from the Networking Products MarketPlace:

  http://www.cisco.com/cgi-bin/order/order_root.pl
- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:

  http://www.cisco.com/go/subscription
- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco corporate headquarters (California, USA) at 408 526-7208 or, in North America, by calling 800 553-NETS(6387).

Documentation Feedback

If you are reading Cisco product documentation on the World Wide Web, you can submit technical comments electronically. Click **Feedback** in the toolbar and select **Documentation**. After you complete the form, click **Submit** to send it to Cisco.

You can e-mail your comments to bug-doc@cisco.com.

To submit your comments by mail, use the response card behind the front cover of your document, or write to the following address:

Attn Document Resource Connection
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com as a starting point for all technical assistance. Customers and partners can obtain documentation, troubleshooting tips, and sample configurations from online tools. For Cisco.com registered users, additional troubleshooting tools are available from the TAC website.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information and resources at anytime, from anywhere in the world. This highly integrated Internet application is a powerful, easy-to-use tool for doing business with Cisco.

Cisco.com provides a broad range of features and services to help customers and partners streamline business processes and improve productivity. Through Cisco.com, you can find information about Cisco and our networking solutions, services, and programs. In addition, you can resolve technical issues with online technical support, download and test software packages, and order Cisco learning materials and merchandise. Valuable online skill assessment, training, and certification programs are also available.

Customers and partners can self-register on Cisco.com to obtain additional personalized information and services. Registered users can order products, check on the status of an order, access technical support, and view benefits specific to their relationships with Cisco.

To access Cisco.com, refer to the following website:

http://www.cisco.com

Technical Assistance Center

The Cisco TAC website is available to all customers who need technical assistance with a Cisco product or technology that is under warranty or covered by a maintenance contract.

Contacting TAC by Using the Cisco TAC Website

If you have a priority level 3 (P3) or priority level 4 (P4) problem, contact TAC by going to the TAC website:

http://www.cisco.com/tac

P3 and P4 level problems are defined as follows:

- P3—Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- P4—You need information or assistance on Cisco product capabilities, product installation, or basic product configuration.

In each of the above cases, use the Cisco TAC website to quickly find answers to your questions.

To register for Cisco.com, refer to the following website:

http://www.cisco.com/register/

If you cannot resolve your technical issue by using the TAC online resources, Cisco.com registered users can open a case online by using the TAC Case Open tool at the following website:

http://www.cisco.com/tac/caseopen

Contacting TAC by Telephone

If you have a priority level 1(P1) or priority level 2 (P2) problem, contact TAC by telephone and immediately open a case. To obtain a directory of toll-free numbers for your country, refer to the following website:

http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

P1 and P2 level problems are defined as follows:

- P1—Your production network is down, causing a critical impact to business operations if service is not restored quickly. No workaround is available.
- P2—Your production network is severely degraded, affecting significant aspects of your business operations. No workaround is available.

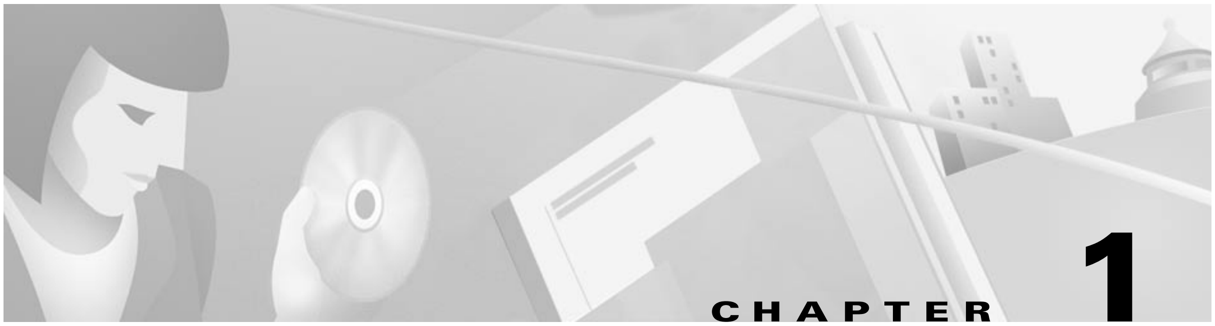

CHAPTER 1

Troubleshooting Overview

This chapter gives a brief overview of the various areas where you might need troubleshooting and contains the following sections:

- Using this Troubleshooting Guide, page 1-1
- General Model of Problem Solving, page 1-2
- Preparing for Network Failures, page 1-3
- Troubleshooting General Problems, page 1-3

Using this Troubleshooting Guide

This *ATM and Layer 3 Switch Router Troubleshooting Guide* describes troubleshooting procedures for the following ATM switch routers:

- Catalyst 8540 MSR
- Catalyst 8510 MSR
- Catalyst 8540 CSR
- Catalyst 8510 CSR
- LightStream 1010

The Catalyst 8500 series and LightStream 1010 hardware and software provide flexibility and performance in a single integrated ATM switch router. For example, your switch router could be configured to provide Layer 3 enabled ATM functionality, delivered through the ATM router module (ARM), which provides routing between ATM and Layer 3 interfaces on a single platform. Or, your switch could be configured to function as a simple ATM backbone switch with no Layer 3 connectivity at all.

To eliminate redundancy and allow you to quickly find your troubleshooting information, this guide is separated into the following parts:

- Chapters 1 through 4 describe how to use this guide, troubleshooting tools, and example networks. Theses chapters should be read first.
- Part 1 (Chapters 5 through 9) describes ATM-to-ATM connection troubleshooting.
- Part 2 (Chapters 10 through 12) describes Layer 3-to-Layer 3 connection troubleshooting.
- Part 3 (Chapter 13) describes Layer 3-to-ATM connection troubleshooting.
- Part 4 (Appendixes A through E) provides general information helpful for troubleshooting, and information about technical support.

Basic troubleshooting processes, such as troubleshooting Ethernet connections, not specific to the ATM switch router, are not described in this document. This information is found online in other troubleshooting guides such as the *Internetwork Troubleshooting Guide.*

General Model of Problem Solving

When troubleshooting a network environment, define the specific symptoms, identify all potential problems that could be causing the symptoms, and then systematically eliminate each potential problem (from most likely to least likely) until the symptoms disappear.

Figure 1-1 illustrates the general problem-solving model. This process is not a rigid outline for troubleshooting an internetwork. It is a foundation on which you can build a problem-solving process for your environment.

Figure 1-1 General Model of Problem Solving

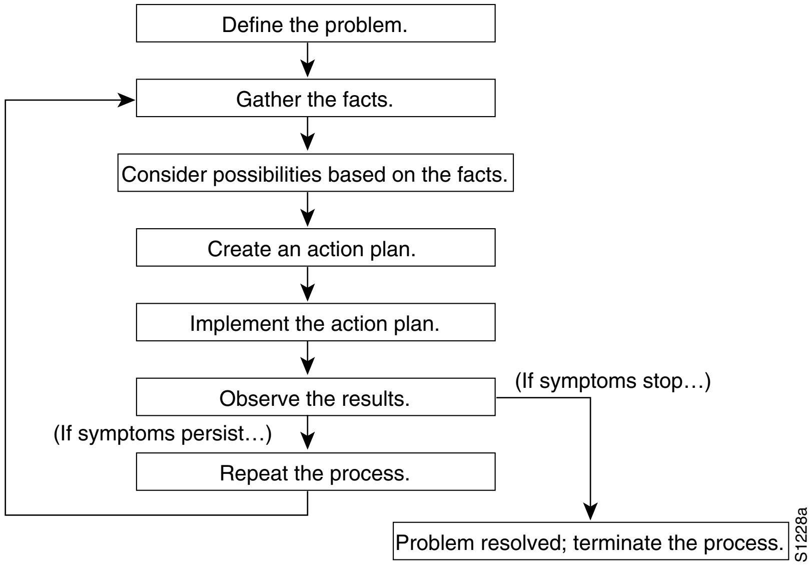

The following steps detail the problem-solving process outlined in Figure 1-1:

Step 1 Analyze the network problem and create a clear problem statement. Define symptoms and potential causes.

Step 2 Gather the facts you need to help isolate possible causes.

Step 3 Consider possible causes based on the facts you gathered.

Step 4 Create an action plan based on those causes. Begin with the most likely problem and devise a plan in which you manipulate only *one* variable.

Step 5 Implement the action plan, performing each step carefully while testing to see whether the symptom disappears.

Step 6 Analyze the results to determine whether the problem has been resolved. If it has, the process is complete.

Step 7 If the problem has not been resolved, create an action plan based on the next most probable cause on your list. Return to Step 4 and repeat the process until the problem has been solved.

Make sure that you undo anything you changed while implementing your action plan. Remember that you want to change only one variable at a time.

Note If you exhaust all the common causes and actions (either those outlined in this publication or others that you have identified in your environment), contact customer service. See Appendix E, "Technical Support," for additional information.

Preparing for Network Failures

It is always easier to recover from a network failure if you are prepared ahead of time. To determine if you are prepared for a network failure, answer the following questions:

1. Do you have an accurate physical and logical map of your internetwork that outlines the physical location of all of the devices on the network and how they are connected, as well as a logical map of network addresses, network numbers, and subnetworks?
2. Do you have a list of all network protocols implemented in your network for each of the protocols implemented and a list of the network numbers, subnetworks, zones, and areas that are associated with them?
3. Do you know which protocols are being routed and the correct, up-to-date configuration information for each protocol?
4. Do you know which protocols are being bridged? Are there any filters configured in any of these bridges, and do you have a copy of these configurations?
5. Do you know all the points of contact to external networks, including any connections to the Internet? For each external network connection, do you know what routing protocol is being used?
6. Has your organization documented normal network behavior and performance so that you can compare current problems with a baseline?

If you can answer *yes* to these questions, then you should be able to recover from a failure quickly and easily.

Troubleshooting General Problems

This section describes where to find troubleshooting, installation, and configuration information for non-ATM related switch router problems, and includes:

- Troubleshooting Hardware and Booting Problems
- Troubleshooting Ethernet Media Problems
- Troubleshooting Console, Auxiliary Line, and Modem Problems

Refer to the following publications for ATM-specific software configuration and command reference information:

- *ATM Switch Router Software Configuration Guide*
- *ATM Switch Router Command Reference*

Troubleshooting Hardware and Booting Problems

Refer to the following publications for more detailed information about booting problems or specific hardware, including descriptions of specific LEDs and configurations and additional troubleshooting information:

- *Hardware Installation Guide (Catalyst 8510 MSR and LightStream 1010)*
- *Processor Installation Guide (Catalyst 8510 MSR and LightStream 1010)*
- *ATM Port Adapter and Interface Module Installation Guide*
- *Catalyst 8540 CSR Route Processor and Interface Module Installation Guide*
- Refer to the *Internetwork Troubleshooting Guide* for general information describing booting problems.

Troubleshooting Ethernet Media Problems

The information referred to in this guide is by no means comprehensive. Instead, it offers solutions to the problems most commonly encountered when using Ethernet media.

Note To troubleshoot ATM physical interface connections, refer to Chapter 5, "Troubleshooting Switch Router ATM Interface Connections."

Refer to the following publications for more detailed information about booting problems or specific hardware, including descriptions of specific LEDs, configuration, and additional troubleshooting information:

- *Processor Installation Guide (Catalyst 8510 MSR and LightStream 1010)*
- *Hardware Installation Guide (Catalyst 8510 MSR and LightStream 1010)*

Refer to the *Internetwork Troubleshooting Guide* for general information describing Ethernet LAN media problems.

Troubleshooting Console, Auxiliary Line, and Modem Problems

Refer to the *Processor Installation Guide (Catalyst 8510 MSR and LightStream 1010)* for more detailed information about specific Ethernet, auxiliary and console port connections, including descriptions of specific LEDs and configurations, and additional troubleshooting information.

Refer to the *Internetwork Troubleshooting Guide* for general information describing Ethernet, auxiliary, and console port connection problems.

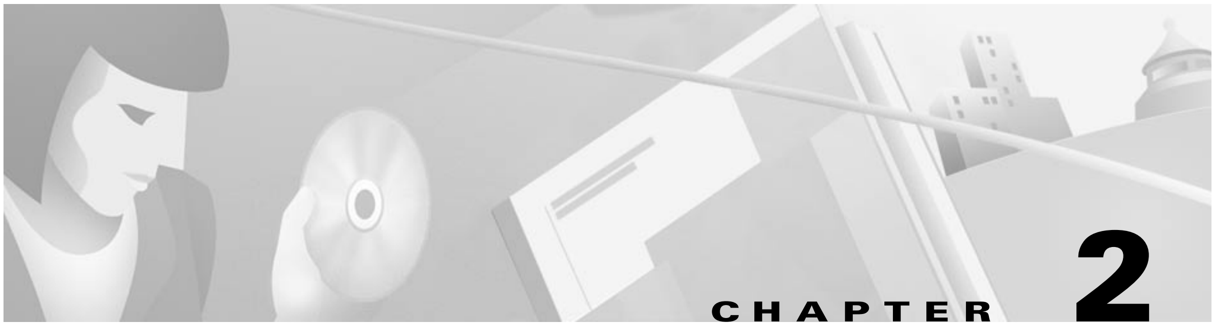

CHAPTER 2

Troubleshooting Tools

This chapter describes the tools available to assist you in troubleshooting your switch router, and contains the following sections:

- Using Diagnostic Commands, page 2-1
- Third-Party Troubleshooting Tools, page 2-4

Using Diagnostic Commands

You can use the **show**, **debug**, **ping**, and **traceroute** commands to monitor and troubleshoot your internetwork.

show Commands

You can use the **show** commands to perform many functions:

- Monitor switch router behavior during initial installation
- Monitor normal network operation
- Isolate problem interfaces, nodes, media, or applications
- Determine when a network is congested
- Determine the status of servers, clients, or other neighbors

The following are some of the most commonly used **show** commands:

Table 2-1 Useful Diagnostic Commands

Command	Purpose
show interfaces **show interfaces atm** **show interfaces atm-p** **show interfaces cbr** **show interfaces FastEthernet** **show interfaces GigEthernetWAN** **show interfaces GigabitEthernet** **show interfaces POS** **show interfaces Port-channel** **show interfaces Tunnel** **show interfaces stats**	Displays statistics for the network interfaces.
show controllers **show controllers atm** **show controllers ethernet** **show controllers FastEthernet** **show controllers GigEthernetWAN** **show controllers GigabitEthernet** **show controllers POS** **show controllers c8500**	Displays statistics for port adapter interface controllers.
show lane	Displays the LAN emulation configuration.
show running-config	Displays the switch router configuration currently running.
show startup-config	Displays the switch router configuration stored in nonvolatile RAM (NVRAM).
show flash	Displays the layout and contents of Flash memory.
show buffers	Displays statistics for the buffer pools on the switch router.
show memory	Shows statistics about the switch router memory, including free pool statistics.
show processes	Displays information about the active processes on the switch router.
show stacks	Displays information about the stack utilization of processes and interrupt routines, and the reason for the last system reboot.
show version	Displays the configuration of the system hardware, the software version, the names and sources of configuration files, and the boot images.

For further information about **show** commands, refer to the *ATM Switch Router Command Reference* for your specific software version.

Converted show epc Commands

In the 12.0(10)W5(18) system software release, some of the **show** commands commonly used for troubleshooting have been converted from **show epc** commands to **show controllers** commands. The **show controllers** commands are described in the next section. Table 1 provides the mapping of the command syntax conversion.

Table 1 show Command Conversion

Release 12.0(5)W5(13d) and Earlier show Command Syntax	Release 12.0(10)W5(18) show Command Syntax
show epc if-entry interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **all**	**show controllers** {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **if-entry all**
show epc if-entry interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **entry** {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port*	**show controllers** {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **if-entry entry** {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port*
show epc freecam interface *slot*/*subslot*/*port*	**show controllers** {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **freecam**
show epc ipmcast *groupaddr* **interface** {**fastethernet** | **gigabitethernet**} [**cam** {**0** | **1**}]	**show controllers** {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **ipmcast** *groupaddr* [**cam** {**0** | **1**}]
show epc ipmcast *groupaddr* **all interface** {**fastethernet** | **gigabitethernet**}	**show controllers** {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **ipmcast** *groupaddr* **all**
show epc ipmcast *groupaddr* **detail interface** {**fastethernet** | **gigabitethernet**}	**show controllers** {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **ipmcast** *groupaddr* **detail**
show epc counters	**show controllers c8500 counters**
show epc queuing	**show controllers c8500 queuing**

debug Commands

The **debug** privileged EXEC commands provide a wealth of information about the traffic seen (or *not* seen) on an interface, error messages generated by nodes on the network, protocol-specific diagnostic packets and cells, and other useful troubleshooting data.

Caution Be careful when using **debug** commands. Many of these commands are processor intensive and can cause serious network problems (such as degraded performance or loss of connectivity) if they are enabled on an already heavily loaded switch router. When you finish using a **debug** command, remember to disable it with its specific **no debug** command (or use the **no debug all** command to turn off all debugging).

Appendix A, "Debugging a Switch Router," provides an overview of **debug** commands, including how to use them when you are troubleshooting the switch router.

In many situations, third-party diagnostic tools can be more useful and less intrusive than using **debug** commands. For more information, see the "Third-Party Troubleshooting Tools" section on page 2-4.

ping Commands

To check host reachability and network connectivity, use the **ping** user EXEC or privileged EXEC command. This command can be used to confirm basic network connectivity on IP networks.

For IP, the **ping** command sends Internet Control Message Protocol (ICMP) echo messages. If a station receives an ICMP echo message, it sends an ICMP echo reply message back to the source.

Using the extended command mode of the privileged EXEC mode **ping** command, you can specify the supported IP header options, which allow the switch router to perform a more extensive range of test options. To enter **ping** extended command mode, enter the **ping** command at the command prompt followed by a return.

To see how the command works under normal conditions, use the **ping** command when the network is functioning properly. When you are troubleshooting, you can then see the difference between normal and abnormal operation.

For detailed information about using the **ping** and extended **ping** commands, refer to the *Cisco IOS Configuration Fundamentals Command Reference* publication.

traceroute Command

The **traceroute** user EXEC command discovers the routes packets follow when traveling to their destinations. With the **traceroute** privileged EXEC command, the supported IP header options are specified, and the switch router can perform a more extensive range of test options.

The **traceroute** command works by using the error message generated by switch routers when a datagram exceeds its time-to-live (TTL) value. First, probe datagrams are sent with a TTL value of one. This causes the first switch router to discard the probe datagrams and send back “time exceeded” error messages. The **traceroute** command then sends several probes, and displays the round-trip time for each. After every third probe, the TTL increases by one.

Each outgoing packet can result in one of two error messages. A “time exceeded” error message indicates that an intermediate switch router has seen and discarded the probe. A “port unreachable” error message indicates that the destination node has received the probe and discarded it because it could not deliver the packet to an application. If the timer goes off before a response comes in, the **traceroute** command displays an asterisk (*).

The **traceroute** command terminates when the destination responds, when the maximum TTL is exceeded, or when the user interrupts the **traceroute** command with the escape sequence.

To see how the command works under normal conditions, use the **traceroute** command when the network is functioning properly. When you are troubleshooting, you can then see the difference between normal and abnormal operation.

For detailed information about using the **traceroute** and extended **traceroute** commands, refer to the *ATM Switch Router Command Reference*.

Third-Party Troubleshooting Tools

In many situations, third-party diagnostic tools can be helpful. For example, attaching a network analyzer to a network is less intrusive and is more likely to yield useful information, without interrupting the operation of the switch router, than using the **debug** commands, which are processor intensive.

Some typical third-party tools used for troubleshooting internetworks are described in the following sections:

- Volt-Ohm Meters, Digital Multimeters, and Cable Testers
- TDRs and OTDRs
- Network Monitors
- Network Analyzers

Volt-Ohm Meters, Digital Multimeters, and Cable Testers

Volt-ohm meters and digital multimeters measure parameters such as AC and DC voltage, current, resistance, capacitance, and cable continuity. They check physical connectivity.

Using cable testers (scanners), you can also check physical connectivity. Cable testers are available for foil twisted-pair (FTP), unshielded twisted-pair (UTP), 10BaseT, and coaxial and twinax cables. A given cable tester can perform any of the following functions:

- Test and report on cable conditions, including near-end crosstalk (NEXT), attenuation, and noise
- Perform time domain reflectometer (TDR) functions, traffic monitoring, and wire map functions
- Display media access control (MAC)-layer information about LAN traffic, provide statistics such as network utilization and packet error rates, and perform limited protocol testing (for example, TCP/IP tests such as **ping**).

Similar testing equipment is available for fiber-optic cable. Due to the relatively high cost of fiber cable and its installation, test fiber-optic cable both before installation (on-the-reel testing) and after installation. Continuity testing of the fiber requires either a visible light source or a reflectometer. Light sources capable of providing light at the three predominant wavelengths, 850 nanometers (nm), 1300 nm, and 1550 nm, are used with power meters that can measure the same wavelengths and test attenuation and return loss in the fiber.

TDRs and OTDRs

TDRs quickly locate open circuits, short circuits, crimps, kinks, sharp bends, impedance mismatches, and other defects in metallic cables.

A TDR reflects a signal off the end of the cable. Opens, shorts, and other problems reflect back the signal at different amplitudes, depending on the problem. A TDR measures the time it takes for the signal to reflect and calculates the distance to a fault in the cable. TDRs can also measure the length of a cable, and some TDRs can calculate the rate of propagation based on a configured cable length.

Fiber-optic measurement is performed by an optical time domain reflectometer (OTDR). OTDRs can accurately measure the length of the fiber, locate cable breaks, measure the fiber attenuation, and measure splice or connector losses. An OTDR can take the signature of a particular installation, noting attenuation and splice losses. This baseline measurement can then be compared with future signatures when you suspect a problem in the system.

Network Monitors

Network monitors continuously track packets crossing a network, providing an accurate picture of network activity. Network monitors do not decode the contents of frames. They are useful for creating a baseline of normal performance.

Monitors collect information such as packet sizes, the number of packets, error packets, overall usage of a connection, the number of hosts and their MAC addresses, and details about communications between hosts and other devices. This data can be used to create profiles of LAN traffic and assist in locating traffic overloads, planning for network expansion, detecting intruders, and distributing traffic more efficiently.

Network Analyzers

To accurately troubleshoot your ATM network, you should have the following analyzers:

- Simple cell generators and analyzers to test high-speed ATM and Broadband Integrated Services Digital Network (BISDN) transmission and protocols.
- Signalling generators to test ATM equipment, service installation, and the interworking of broadband services. They help manage the performance of broadband networks, and guarantee end-to-end quality of service (QoS).
- Physical layer analyzers to provide physical, convergence, and ATM cell testing capabilities and transmission test functionality.

  Most physical layer analyzers can perform many of the following functions:

  - Traffic generation
  - Cell error and cell loss measurements
  - Cell delay measurements
  - Traffic capture and playback
- Network analyzers (or protocol analyzers) decode the various protocol layers in a recorded frame and present them as readable abbreviations or summaries, detailing which layer is involved (physical, data link, and so forth) and what function each bit or byte content serves.

  Most network analyzers can perform many of the following functions:

  - Filter traffic that meets certain criteria so that, for example, all traffic to and from a particular device is captured
  - Time-stamp captured data
  - Present protocol layers in an easily readable form
  - Generate frames and transmit them to the network
  - Incorporate an “expert” system in which the analyzer uses a set of rules, combined with information about the network configuration and operation, to diagnose, solve, or offer potential solutions to network problems

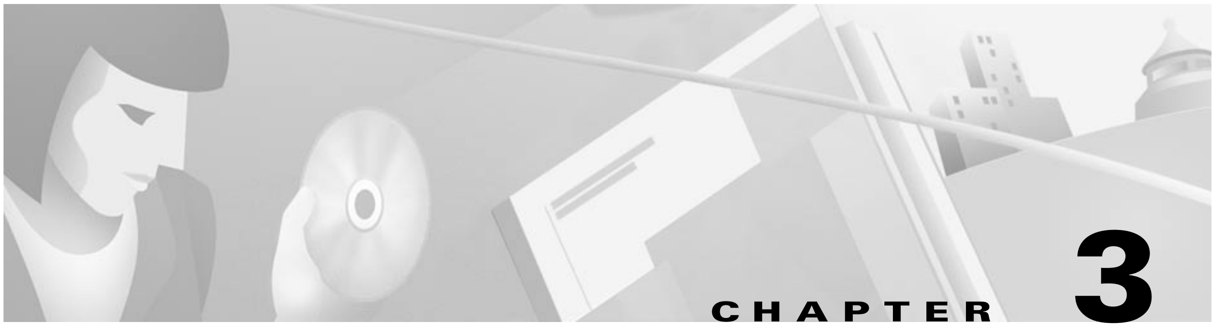

CHAPTER 3

Initial Troubleshooting

This chapter describes the first steps you should take when you start troubleshooting your switch router, and contains the following sections:

- Online Diagnostics, page 3-1
- Checking DDTs Database and Release Notes for Workarounds, page 3-5
- Troubleshooting Hardware and Software Version Problems, page 3-6
- Troubleshooting Route Processor Redundancy (Catalyst 8540 CSR and Catalyst 8540 MSR), page 3-14
- Troubleshooting Switch Processor Redundancy with HSRP (Catalyst 8540 CSR and Catalyst 8540 MSR), page 3-18

Online Diagnostics

This section describes the online diagnostics available for troubleshooting your switch router. Online diagnostics provide the following types of tests:

- Processor Loopback Test (Catalyst 8540 CSR)
- Accessibility tests between the route processor and the interface modules and the route processor and the switch processor.
- Online insertion and removal (OIR) diagnostic tests.
- Snake tests through the switch router to ensure connectivity between the ports.

The switch router displays an error message on the console when it detects a hardware failure or problem.

Note Online diagnostic tests only run on the primary route processor.

Processor Loopback Test (Catalyst 8540 CSR)

The processor loopback test detects failures in the route processor to switch processor interface. The test sends a packet from the route processor to each switch processor which then loops back to the route processor. This test can run on the switch router without any port adapters or interface modules installed.

Note The size of the packet and frequency of the test are configurable to minimize the impact on system performance.

Accessibility Test

The accessibility tests ensure connectivity, at a configurable interval, between all of the following:

- Interface modules
- Active switch processor
- Standby switch processor, if it is present

OIR Test

Online insertion and removal (OIR) tests check the functioning of the switch fabric and interfaces on a per-port basis. The switch router performs these tests when the system boots up and when you insert an interface module into a slot. The OIR test sends a packet to the interface loopback and expects to receive it within a certain time period. If the packet does not reach the port within the expected time period, or the received packet is corrupted, an error is registered and the port is changed to an administrative down state. Packets that are 1000 bytes in size are used in the test.

The OIR tests support all the enhanced Gigabit Ethernet interface modules, with the exception of the OC-3c and OC12c ATM uplink interface module. In addition, OIR tests are not supported on the Fast Ethernet or Gigabit Ethernet interface modules.

OIR is enabled by default on the Catalyst 8540 CSR. To disable it, enter the **no diag online oir** command. To enable OIR, refer to commands in the “Configuring Online Diagnostics (Catalyst 8540 CSR)” section on page 3-2.

Snake Test

The snake test detects and reports port-to-port connectivity failures. The snake test establishes connections across all the active ports in the switch router, originating and terminating at the primary route processor. The route processor establishes a connection by sending a packet to each port in turn, which then terminates at the route processor. If the packet does not reach the route processor within the expected time period, or the received packet is corrupted, further testing is performed to isolate and disable the port causing the problem. The frequency of the test is configurable to minimize the impact on system performance.

The snake test supports all the enhanced Gigabit Ethernet interface modules on the Catalyst 8540 MSR and Catalyst 8540 CSR.

Configuring Online Diagnostics (Catalyst 8540 CSR)

To configure online diagnostics, use the following global configuration commands:

Command	Purpose
diag online	Enables all of the online diagnostic tests.
diag online access	Enables only the accessibility diagnostic test.

Command	Purpose
diag online access freq [*seconds*]	Configures the frequency of the accessibility diagnostic tests. The default frequency is every 10 seconds.
diag online oir	Enables only the OIR test.
diag online oir pktsize [*bytes*]	Specifies the packet size for the OIR test. The default size is 1000 bytes.
diag online snake	Enables only the snake test.
diag online snake timer [*seconds*]	Specifies the time interval for the snake test. The default interval is 10 seconds.
no diag online [**access** | **oir** | **snake**]	Disables the online diagnostic tests.
clear counters diag online {**access** | **oir** | **snake** | **all**}	Clears the online diagnostic test counters.
debug diag online [**access** | **oir** | **snake**]	Enables debugging of online diagnostic tests.
no debug diag online [**access** | **oir** | **snake**]	Disables debugging of online diagnostic tests.

Examples

The following example shows how to enable all online diagnostic tests:

```
Router(config)# diag online
 ONLINE-DIAG: Enabling all Online Diagnostics tests
```

The following example shows how to change the frequency of the access test:

```
Router(config)# diag online access freq 20
ONLINE-DIAG: Online Access Test Frequency set to 20 sec
```

Displaying the Online Diagnostics Configuration and Results (Catalyst 8540 CSR)

To display the online diagnostics configuration and results, use the following EXEC command:

Command	Purpose
show diag online [**details** | **status**] [**access** | **oir** | **snake**]	Displays information about the online diagnostics tests and the test results.

Example

The following example shows how to display detailed access test information:

```
Switch# show diag online details access
======== Online Access Test Details ========

Current Test Status : Test is Enabled
Current Frequency of Access Test : 10 seconds

Slot Card-Type       Iteration    Success    Failure    Last Failure
---- ----------      ----------   -------    -------    ------------
 0/* K1 GIGETHERNET  114536       114536     0          ----
 1/* CMPM Card       114536       114536     0          ----
 1/0 XPIF POS OC12 P 114536       114536     0          ----
 2/* CMPM Card       114536       114536     0          ----
 2/0 XPIF POS OC12 P 114536       114536     0          ----
 2/1 XPIF GIGE PAM   114536       114536     0          ----
 3/* CMPM Card       114536       114536     0          ----
 3/0 XPIF ATM OC3 PM 114536       114536     0          ----
 3/1 XPIF GIGE PAM   114536       114536     0          ----
 5/* Switch Card     114536       114536     0          ----
 6/* Switch Card     27126        0          27126      1w6d
 7/* Switch Card     114536       114536     0          ----
10/* ETHERNET PAM    27119        27119      0          ----
11/* CMPM Card       114536       114536     0          ----
11/0 XPIF GIGE PAM   114536       114536     0          ----
11/1 XPIF GIGE PAM   114536       114536     0          ----
12/* CMPM Card       114537       114537     0          ----
12/0 XPIF ATM OC12 P 114537       114537     0          ----
12/1 XPIF GIGE PAM   114537       114537     0          ----
```

Example

The following example shows how to display OIR test status and details:

```
Switch# show diag online oir
========  OIR Test Status and Details ========
======== Online OIR Test Status ========
Current Test Status :Test is Enabled
-------- Bootup OIR status --------
Port    Card Type   Pkt Size  Result                    Test Time LOOP
_______ ___________ _________ ___________________  ______________ ____
00/0/00 ETHERNET PA      1000 OIR_SUCCESS                6 00:01:54  PIF
00/0/01 ETHERNET PA      1000 OIR_SUCCESS                 00:01:52  PIF
00/0/02 ETHERNET PA      1000 OIR_SUCCESS                 00:01:50  PIF
00/0/03 ETHERNET PA      1000 OIR_SUCCESS                 00:01:47  PIF
00/0/04 ETHERNET PA      1000 OIR_SUCCESS                 00:01:55  PIF
00/0/05 ETHERNET PA      1000 OIR_SUCCESS                 00:01:53  PIF
00/0/06 ETHERNET PA      1000 OIR_SUCCESS                 00:01:51  PIF
00/0/07 ETHERNET PA      1000 OIR_SUCCESS                 00:01:49  PIF
00/0/08 ETHERNET PA      1000 OIR_SUCCESS                 00:02:04  PIF
00/0/09 ETHERNET PA      1000 OIR_SUCCESS                 00:02:01  PIF
00/0/10 ETHERNET PA      1000 OIR_SUCCESS                 00:01:59  PIF
00/0/11 ETHERNET PA      1000 OIR_SUCCESS                 00:01:56  PIF
00/0/12 ETHERNET PA      1000 OIR_SUCCESS                 00:02:05  PIF
00/0/13 ETHERNET PA      1000 OIR_SUCCESS                 00:02:03  PIF
00/0/14 ETHERNET PA      1000 OIR_SUCCESS                 00:02:00  PIF
00/0/15 ETHERNET PA      1000 OIR_SUCCESS                 00:01:58  PIF

01/0/00 XPIF POS OC      1000 OIR_SUCCESS                 00:01:48  PIF
01/1/00 XPIF GIGE P      1000 OIR_SUCCESS                 00:01:57  PIF

02/0/00 GIGETHERNET      1000 OIR_SUCCESS                 00:02:07  PIF
02/0/01 GIGETHERNET      1000 OIR_SUCCESS                 00:02:10  PIF

03/0/00 XPIF GIGE P      1000 OIR_SUCCESS                 00:02:08  PIF
03/1/00 XPIF GIGE P      1000 OIR_SUCCESS                 00:02:11  PIF

-------- Latest OIR status --------
********No Other OIR tests not performed ********


======== Online OIR Test Details ========
Current Test Status :Test is Enabled
-------- Previous failure details ----------
******* No failures in OIR tests *******

-------- Complete details --------
Port    Tx Pkt    Rx Pkt    Success   Failure   Total Tests
_______ _________ _________ _________ _________ ___________
```

```
00/0/00         1         1          1         0           1
00/0/01         1         1          1         0           1
00/0/02         1         1          1         0           1
00/0/03         1         1          1         0           1
00/0/04         1         1          1         0           1
00/0/05         1         1          1         0           1
00/0/06         1         1          1         0           1
00/0/07         1         1          1         0           1
00/0/08         1         1          1         0           1
00/0/09         1         1          1         0           1
00/0/10         1         1          1         0           1
00/0/11         1         1          1         0           1
00/0/12         1         1          1         0           1
00/0/13         1         1          1         0           1
00/0/14         1         1          1         0           1
00/0/15         1         1          1         0           1

01/0/00         1         1          1         0           1
01/1/00         1         1          1         0           1

02/0/00         1         1          1         0           1
02/0/01         1         1          1         0           1

03/0/00         1         1          1         0           1
03/1/00         1         1          1         0           1
```

Checking DDTs Database and Release Notes for Workarounds

This section describes different methods you can use to check for IOS software bugs (defect tracking tool numbers [DDTs]) in your version of IOS software. Often, your problems with the switch router have been fixed or a workaround has been determined in a more recent version of software.

There are two ways to check for known bugs in the IOS software:

- Using Bug Navigator II, page 3-5
- Checking IOS Release Notes, page 3-6

Using Bug Navigator II

Bug Navigator II is a DDT search tool you can use to search the DDT database and ask either of two types of questions:

- Symptom Diagnostics (for example, “What defect is causing my current symptoms?”)
- Upgrade Planning (for example, “What software release is best for the features I am interested in?”)

To search the DDT database, you can access Bug Navigator II on the World Wide Web at http://www.cisco.com/support/bugtools/bugtool.shtml. Then perform the following steps:

Step 1 If you are not already logged in to Cisco.com, enter your user name and password at the login prompt.

Step 2 Read the “Bug Navigator II Help” instructions.

Step 3 From the “Cisco Hardware” list select your switch router under “Catalyst 8500 Series Switches.” The Bug Navigator search tool replaces “Bug Navigator II Help” (in the right frame of the page).

Step 4 Select the following from the drop down menus:

- Version
- Revision
- Severity

Note As an option, you can enter words or phrases (separated by commas) in the data entry field to limit your search.

Step 5 Click the “Search” button.

The entire window will be replaced with a “Bug Search Results” window with a list of DDTs containing your search criteria. Look at the Bug reports listed in the “titles” column. An existing bug entry that describes the problem you are having may have been fixed in a more recent version of the IOS software. Look in the “Fixed-in” column for a later version of the IOS software. All you might have to do to solve your problem is upgrade your software.

If a software upgrade is not listed as a way to solve your problem, double-click on the bug title and read the DDT details; a workaround might be listed there.

Checking IOS Release Notes

Release notes describe the features and caveats for Cisco IOS software releases. The release notes are listed by both product and IOS release number.

Note All information pertains to both the Catalyst 8540 CSR and Catalyst 8510 CSR platforms, unless differences between the platforms are noted in the text.

The “Caveats” section of the release note lists known caveats for the switch router by tracking the DDTS number and the release number, and indicates whether the caveat has been corrected.

The “Caveat Symptoms and Workarounds” section summarizes caveat symptoms and suggested workarounds for the switch routers. You can also search thorough this section online, using either a word string or the DDTS number.

Troubleshooting Hardware and Software Version Problems

A common error you may encounter is the incompatibility of hardware modules and the IOS software version needed to perform a particular function. Check the *Hardware and Software Compatibility Matrix* document, available online, to confirm that you are using IOS software that supports the various hardware components installed in your switch router.

Verifying Hardware and Software Versions

Display the hardware and software versions to ensure that they are the most recent. Very old hardware and software versions (two or three versions back) can have caveats that have been fixed in more recent versions. Use the following EXEC commands to display the version information:

Command	Purpose
show version	Displays the software version information.
show hardware [**detail**]	Displays detailed hardware information including revision level and version.
show functional-image slot *slot*	Displays functional image information.

To verify the hardware and software versions, use the following steps:

Step 1 Display the system software version:

```
Switch# show version
Cisco Internetwork Operating System Software
➔ IOS (tm) L3 Switch/Router Software (C8540CSR-IN-M), Version 12.0(10)W5(18c)
Copyright (c) 1986-2000 by cisco Systems, Inc.
Compiled Sat 19-Aug-00 00:11 by integ
Image text-base: 0x60010930, data-base: 0x608CA000

ROM: System Bootstrap, Version 12.0(4.6)W5(13) RELEASE SOFTWARE

8540CSR uptime is 2 minutes
System restarted by reload
System image file is "slot0:cat8540c-in-mz.120-10.W5.18c.bin"

cisco C8540CSR (R5000) processor with 262144K/256K bytes of memory.
R5000 processor, Implementation 35, Revision 2.1 (512KB Level 2 Cache)
Last reset from power-on
1 Ethernet/IEEE 802.3 interface(s)
16 FastEthernet/IEEE 802.3 interface(s)
13 Gigabit Ethernet/IEEE 802.3z interface(s)
3 ATM network interface(s)
2 Packet over SONET network interface(s)
505K bytes of non-volatile configuration memory.

20480K bytes of Flash PCMCIA card at slot 0 (Sector size 128K).
16384K bytes of Flash internal SIMM (Sector size 256K).
Configuration register is 0x0
```

Step 2 Use the **show hardware** command to display the hardware revision levels.

```
  Switch# show hardware

  C8540 named Router, Date: 13:59:13 UTC Mon Jan 3 2000

➔ Slot Ctrlr-Type    Part No.  Rev  Ser No  Mfg Date   RMA No. Hw Vrs  Tst EEP
  ---- ------------  ---------- -- -------- --------- -------- ------- --- ---
   0/* K1 GIGETHERN  73-3324-03 A0 0336441Y Oct 13 99 0          3.4
   1/* CMPM Card     73-3944-03 09 03445724 Nov 09 99            3.0
   1/0 XPIF POS OC1  73-4462-01 09 034558YP Nov 09 99            1.1
   2/* CMPM Card     73-3944-03 A0 04087BW8 Mar 22 00 0          3.0
   2/0 XPIF POS OC1  73-4462-01 A0 04046NRQ Mar 22 00 0          2.0
   2/1 XPIF GIGE PA  73-4167-05 A0 04097GRJ Mar 22 00 0          1.0
   3/* CMPM Card     73-3944-03 A0 04087BXK Mar 15 00 0          3.0
   3/0 XPIF ATM OC3  73-3889-03 A0 040879AA Mar 15 00 0          1.0
   3/1 XPIF GIGE PA  73-4167-05 A0 04097GQA Mar 15 00 0          1.0
   4/* Route Proc    73-3775-04 A0 03201VCZ Oct 04 99 0          5.7
   5/* Switch Card   73-3327-08 A0 032428ZR Jun 15 99 0          8.0
   7/* Switch Card   73-3327-08 A0 032428ZE Jun 15 99 0          8.0
  10/* ETHERNET PAM  73-3753-04 A0 03020FCA Sep 22 99 0          4.1
  11/* CMPM Card     73-3944-03 A0 04087BY5 Mar 13 00 0          3.0
  11/0 XPIF GIGE PA  73-4415-05 A0 04087AZE Mar 13 00 0          1.0
  11/1 XPIF GIGE PA  73-4415-05 A0 04087AZL Mar 13 00 0          1.0
  12/* CMPM Card     73-3944-03 A0 04087BWS Mar 14 00 0          3.0
  12/0 XPIF ATM OC1  73-3889-03 A0 040879AO Mar 14 00 0          1.0
  12/1 XPIF GIGE PA  73-4167-05 A0 04107N8R Mar 14 00 0          1.0

  DS1201 Backplane EEPROM:
  Model  Ver.  Serial  MAC-Address  MAC-Size  RMA  RMA-Number   MFG-Date
  ------ ---- -------- ------------ --------  ---  ----------  -----------
  C8540  2   12237014 00D0BA1D3200   1024      0        0      Jun 18 1999
  cubi version : 11

  Power Supply:
  Slot Part No.         Rev  Serial No.  RMA No.     Hw Vrs  Power Consumption
  ---- ---------------- ---- ----------- ----------- ------- -----------------
  0          34-0829-02 A000 APQ02450080 00-00-00-00   1.0            2746 cA
```

Step 3 Verify that the hardware version is listed in the Hw Vrs column.

Step 4 Use the **show hardware detail** command to display detailed information about the hardware, including the functional image versions.

```
Switch# show hardware detail

<Information deleted>

slot:  3/0  Controller-Type : XPIF ATM OC3 PM - 1 Port SM_IR
  Part Number: 73-3889-03                       Revision: A0
Serial Number: CAB040879AA                      Mfg Date: Mar 15 00
   RMA Number: 0                             H/W Version: 1.0
 FPGA Version: 1.14

 XPIF Version: 3001                             CAM size: 64 KB
Ucode Version: 1.0                              CAM Type: Private TCAM

Port Phy Setup
      Port  0: DONE

Optical Line Daughter Card Serial EEPROM:
  Part Number: 73-3975-02                       Revision: A0
Serial Number: CAB0407768M                      Mfg Date: 2000/04/10
   RMA Number: 0                                HW Rever: 1.0

TCAM Daughter Card Serial EEPROM:
  Part Number: 73-3970-02                       Revision: A0
Serial Number: CAB04087BXK                      Mfg Date: 2000/03/04
   RMA Number: 0                                HW Rever: 2.0

slot:  3/1  Controller-Type : XPIF GIGE PAM
  Part Number: 73-4167-05                       Revision: A0
Serial Number: CAB04097GQA                      Mfg Date: Mar 15 00
   RMA Number: 0                             H/W Version: 1.0
 FPGA Version: 20.72

 XPIF Version: 3001                             CAM size: 64 KB
Ucode Version: 1.0                              CAM Type: Private TCAM

Port Phy Setup
      Port  0: DONE                          GBIC Vendor: No vendor info.

slot:  4/*  Controller-Type : Route Proc
  Part Number: 73-3775-04                       Revision: A0
Serial Number: CAB03201VCZ                      Mfg Date: Oct 04 99
   RMA Number: 0                             H/W Version: 5.7
 FPGA Version: 4.8

slot:  5/*  Controller-Type : Switch Card
  Part Number: 73-3327-08                       Revision: A0
Serial Number: CAB032428ZR                      Mfg Date: Jun 15 99
   RMA Number: 0                             H/W Version: 8.0
 FPGA Version: 1.2

slot:  7/*  Controller-Type : Switch Card
  Part Number: 73-3327-08                       Revision: A0
Serial Number: CAB032428ZE                      Mfg Date: Jun 15 99
   RMA Number: 0                             H/W Version: 8.0
 FPGA Version: 1.2

<Information deleted>
```

Step 5 Use the **show functional-image** command to display detailed information about the functional images for the route processors, switch processors, and Fast Ethernet and Gigabit Ethernet interface modules for the switch router (in this example, the Catalyst 8540 CSR). The following example shows how to display the functional image for the route processor in slot 4:

```
  Switch# show functional-image slot 4

  Details for cpu Image on slot: 4

➔ Functional Version of the FPGA Image: 4.8
   #Jtag-Distribution-Format-B
➔  #HardwareRequired: 100(3.0-19,4.0-19,5.0-19)
   #FunctionalVersion: 4.8
   #Sections: 1
   #Section1Format: MOTOROLA_EXORMAX

    Copyright (c) 1996-00 by cisco Systems, Inc.
    All rights reserved.
    generated by:       holliday
    on:                 Mon Mar  6 13:59:17 PST 2000
    using:              /vob/cougar/bin/jtag_script Version 1.13
    config file:        cpu.jcf

    Chain description:
    Part type Bits Config file
    10k50     10   ../cidrFpga2/max/cidr_fpga.ttf
    xcs4062   3    ../cubiFpga2/xil/cubi.bit
    xcs4062   3    ../cubiFpga2/xil/cubi.bit
    generic   2
    XC4005    3    /vob/cougar/custom/common/jtcfg/xil/jtcfg_r.bit
    Number devices             = 5
    Number of instruction bits = 21

    FPGA config file information:
    Bitgen date/time  Sum   File
    100/03/02 19:14:49 7068  ../cidrFpga2/max/cidr_fpga.ttf
    1999/04/15 18:46:32 36965 ../cubiFpga2/xil/cubi.bit
    1999/04/15 18:46:32 36965 ../cubiFpga2/xil/cubi.bit
    98/06/11 16:56:44 49904 /vob/cougar/custom/common/jtcfg/xil/jtcfg_r.bit

   #End-Of-Header
```

Step 6 Verify the FunctionalVersion and #HardwareRequired fields to determine the FPGA version and the hardware version required for the FPGA. Compare this with the hardware version using the **show hardware** command output displayed in Step 2 and Step 4. If the FPGA version does not support the hardware version, download a new FPGA image, upgrade the hardware, or both.

Finding the Image

Use the *Hardware and Software Compatibility Matrix* (found on Cisco.com) to determine the correct IOS software image for your hardware configuration.

If you determine that you need to upgrade your IOS image, refer to the Cisco.com web page, and follow these steps to find and download the image you need:

Step 1 Login to the Cisco.com; the window will change and display additional features.

Step 2 Under Service and Support, select Software Center.

Step 3 The Technical Assistance Center page is displayed. Under Tools, select IOS Upgrade Planner.

The Cisco IOS Planner page provides greater flexibility to browse for your preferred software. You can view all major releases, all platforms, and all software features from a single interface. Choosing a platform, a maintenance release, or a software feature the planner automatically limits the other menu choices, based on your selections until you arrive at your preferred software.

Step 4 In the Select Platform column, select either of the following:

- "C8540C" (for Catalyst 8510 CSR and Catalyst 8540 CSR platforms)
- "C8540M" for Catalyst 8510 MSR and Catalyst 8540 MSR platforms)

Step 5 In the Select Release column, select the IOS software release you want to download.

Step 6 Read the requirements, and if your hardware configuration meets the requirements, click the agreement button.

Step 7 Select the file to download to your switch router.

To upgrade the IOS image on your platform, continue with the following section.

IOS Upgrade Procedures

If your IOS image is not the most recent, you can download the IOS image from Cisco.com to the switch router.

Refer to the *Configuration Fundamentals Configuration Guide* for details on the following frequently performed tasks:

- Formatting Flash memory on a new PCMCIA card or on any Flash memory device that has locked blocks or failed sectors.
- Managing files on file systems, including setting the default file system, listing files on a file system and deleting and recovering files.

Modifying, Downloading, and Maintaining System Images

The following tasks are performed frequently to maintain system image files:

- Copy images from Flash memory to a network server. You can store system images for backup, or other purposes, by copying them from a Flash memory device to a TFTP or rcp server.
- Copy images from a network server to Flash memory. You perform this procedure when upgrading your system image or functional image.
- Copy images between local Flash memory devices. You perform this procedure when moving a system image from one switch router to another or to synchronize the functional images on switch routers with redundant route processors.

Maintaining Functional Images

You can load functional images used by certain hardware controllers in the switch router. The following sections describe the function and maintenance of functional images:

- Understanding Functional Images
- FPGA Upgrade Procedures

Understanding Functional Images

Functional images provide the low-level operating functionality for various hardware controllers. On hardware controllers within system programmable devices, such as Field Programmable Gate Arrays (FPGAs) and Erasable Programmable Logic Devices (EPLDs), the hardware functional images can be reprogrammed independently of loading the system image, and without removing the devices from the controller.

All new hardware you purchase is shipped with the functional images loaded. Loading a different functional image is required only when upgrading or downgrading functional image versions.

FPGA Upgrade Procedures

If the functional image is not the most recent, you can download the functional image to the switch router from Cisco.com. Use the **reprogram** command to update the functional image to the processor or interface module. The following example shows how to reprogram the route processor in slot 4 with the functional image fi-c8540-rp.A.4-8.bin from the Flash PC card in slot 0:

```
Switch# reprogram slot0:fi-c8540-rp.A.4-8.bin 4
```

Note You can only enter the **reprogram** command from the console session prompt.

You can find the functional images and release notes for one of the following on Cisco.com, on the LAN Switching Products site:

- Catalyst 8540C Functional Software (for the Catalyst 8510 CSR and Catalyst 8540 CSR)
- Catalyst 8540M Functional Software (for the Catalyst 8510 MSR and Catalyst 8540 MSR).

Note After you have determined the hardware and software versions on the switch router, check the release notes and DDTS database for symptoms resembling those you are observing. Often, the problem has already been discovered and a workaround has been provided.

Troubleshooting Switch Route Processor Redundancy and Enhanced High System Availability (Catalyst 8540 CSR and Catalyst 8540 MSR)

The Catalyst 8540 MSR and Catalyst 8540 CSR support, redundant route processor operation with dual route processors. In addition, Enhanced High System Availability (EHSA) is provided in the switching fabric when three switch processors are installed in the chassis. The troubleshooting processes for these features are described in the following sections:

- Route Processor Redundant Operation (Catalyst 8540 CSR and Catalyst 8540 MSR), page 3-13
- Troubleshooting Route Processor Redundancy (Catalyst 8540 CSR and Catalyst 8540 MSR), page 3-14
- Troubleshooting Switch Processor Redundancy with HSRP (Catalyst 8540 CSR and Catalyst 8540 MSR), page 3-18

Route Processor Redundant Operation (Catalyst 8540 CSR and Catalyst 8540 MSR)

The Catalyst 8540 MSR and Catalyst 8540 CSR supports fault tolerance by allowing a secondary route processor to take over if the primary fails. This secondary, or redundant, route processor runs in standby mode. In standby mode, the secondary route processor is partially booted with the Cisco IOS software; however, no configuration is loaded.

At the time of a switchover, the secondary route processor takes over as primary and loads the configuration as follows:

- If the running configurations on the primary and secondary route processors match, the new primary uses the running configuration file
- If the running configurations on the primary and secondary route processors do not match, the new primary uses the last saved configuration file in its nonvolatile random-access memory (NVRAM) (not the NVRAM of the former primary)

The former primary then becomes the secondary route processor.

Note If the secondary route processor is unavailable, a major alarm is reported. Use the **show facility-alarm status** command to display the redundancy alarm status.

For detailed redundant route processor configuration information, refer to the “Initially Configuring the ATM Switch Router” chapter in the *ATM Switch Router Software Configuration Guide*.

Troubleshooting Route Processor Redundancy (Catalyst 8540 CSR and Catalyst 8540 MSR)

For redundant route processors to function correctly, your switch router's route processors must meet all of the following requirements:

- Route processors must have identical hardware configurations
- ROMMON must be version 12.0(4.6)W5(13) or later
- Both route processors must have identical releases of IOS software

A common error you may encounter is the incompatibility of hardware modules and the IOS software version needed to perform a particular function. Check the *Hardware and Software Compatibility Matrix* document, available on-line, to confirm that you are using IOS software that supports the various hardware components installed in your switch router.

Troubleshooting Hardware and Software Versions of Redundant Route Processors

To troubleshoot the route processor hardware and software versions for redundancy, use the following commands:

Command	Purpose
show version	Displays the configuration register value.
show hardware detail	Displays the hardware and software configurations of the primary and secondary route processors.
show redundancy	Displays the hardware and software configurations of the primary and secondary route processors.

To confirm that your switch router route processors meets the redundancy requirements, complete the following steps:

Step 1 Use the **show version** command to confirm the system hardware and software status of the primary route processor.

```
Switch# show version
Cisco Internetwork Operating System Software
IOS (tm) PNNI Software (cat8540m-WP-M), Version 12.1(2.3)W6(1.33), CISCO DEVELOP
MENT TEST VERSION
Copyright (c) 1986-2000 by cisco Systems, Inc.
Compiled Fri 20-Oct-00 23:39 by
Image text-base: 0x60010958, data-base: 0x60D30000

➔ ROM: System Bootstrap, Version 12.0(4.6)W5(13), RELEASE SOFTWARE

8540MSR uptime is 8 weeks, 3 days, 10 hours, 0 minutes
System returned to ROM by reload
System image file is "slot0:cat8540m-wp-mz.121-2.3.PE33"

cisco C8540MSR (R5000) processor with 262144K/256K bytes of memory.
R5000 CPU at 200Mhz, Implementation 35, Rev 2.1, 512KB L2 Cache
Last reset from power-on
3 Ethernet/IEEE 802.3 interface(s)
16 FastEthernet/IEEE 802.3 interface(s)
14 ATM network interface(s)
505K bytes of non-volatile configuration memory.

20480K bytes of Flash PCMCIA card at slot 0 (Sector size 128K).
8192K bytes of Flash internal SIMM (Sector size 256K).
Configuration register is 0x0

Switch#
```

Step 2 Verify the ROM field. It indicates the release of IOS software loaded and running on the primary route processor.

Step 3 Use the **show hardware detail** command to compare the hardware versions of the primary and secondary route processors.

```
    8540MSR# show hardware detail

    C8540 named 8540MSR, Date: 18:42:12 UTC Fri Jan 5 2001
    .
    (Information Deleted)
    .

➔   slot:  4/*  Controller-Type : Route Proc
      Part Number: 73-2644-05                       Revision: A0
    Serial Number: MIC03140NXK                      Mfg Date: Apr 04 99
       RMA Number: 0                             H/W Version: 5.7
     FPGA Version: 4.8

➔   slot:  4/0  Controller-Type : Netclk Module
      Part Number: 73-2868-03                       Revision: A0
    Serial Number: MIC03140NSU                      Mfg Date: Apr 04 99
       RMA Number: 0                             H/W Version: 3.1
     FPGA Version: 4.8
    .
    (Information Deleted)
    .
➔   slot:  8/*  Controller-Type : Route Proc
      Part Number: 73-2644-05                       Revision: A0
    Serial Number: MIC03140NXH                      Mfg Date: Apr 04 99
      RMA Number: 0                             H/W Version: 5.7
     FPGA Version: 4.8

➔   slot:  8/0  Controller-Type : Netclk Module
      Part Number: 73-2868-03                       Revision: A0
    Serial Number: MIC03140NVT                      Mfg Date: Apr 04 99
      RMA Number: 0                             H/W Version: 3.1
     FPGA Version: 4.8
    .
    (Information Deleted)
    .
```

Step 4 In the slots labeled Controller-Type : Route Proc, compare the Part Number, FPGA, and H/W Version fields. These numbers must all match, or redundancy will not function correctly on your switch router.

Step 5 Use the **show redundancy** command to check the configuration and status of the route processors.

```
  Switch# show redundancy

  This CPU is the PRIMARY
→ Primary
  -------
→ Slot:                          8
  Uptime:                        8 weeks, 4 days, 11 hours, 31 minutes
  Image:                         PNNI Software (cat8540m-WP-M), Version 12.1(2.3)W
  12.0(4a)W5(11a)  RELEASE SOFTWARE

  Time Since :
→    Last Running Config. Sync:   2 weeks, 4 days, 2 hours, 28 minutes
→    Last Startup Config. Sync:   2 weeks, 4 days, 2 hours, 34 minutes
→ Last Restart Reason:           Normal boot

→ Secondary
  ---------
→ State:                         UP
→ Slot:                          4
  Uptime:                        1 day, 1 hour, 59 minutes
  Image:                         PNNI Software (cat8540m-WP-M), Version 12.0(4a)W5(11a)
  RELEASE SOFTWARE
```

Step 6 Verify the Primary, Secondary, and Slot fields. They indicate in which slot the primary route processor is configured.

Step 7 Verify the Last Running Config. Sync and Last Startup Config. Sync fields. They indicate the last time the running configuration and startup configuration were synchronized between the route processors.

Troubleshoot Redundant Route Processor Functions

To troubleshoot the route processor functions for redundancy, use the following commands:

Command	Purpose
show atm vc interface atm 0	Displays the VC status between the primary and secondary route processors.
show atm status	Displays the status of the primary and secondary route processors.

Follow these steps to troubleshoot the route processor redundancy on the switch router:

Step 1 Use the **show atm vc interface atm 0** command to confirm that the permanent virtual circuit (PVC) between switch route processors (SRPs) is up.

```
  Switch# show atm vc interface atm 0
  Interface         VPI  VCI   Type   X-Interface      X-VPI X-VCI Encap  Status
  ATM0               0     35     PVC     ATM0/0/0           0     16    ILMI   DO
  .
  (Information Deleted)
  .
→ ATM0               0     245    PVC     ATM-SEC0           0     29    IPC    UP
  Switch#
```

Step 2 Verify the ATM-SEC0 field and confirm the PVC is up.

Step 3 Use the **show atm status** command to confirm the status of the PVC and other states.

```
  Switch# show atm status
  NUMBER OF INSTALLED CONNECTIONS: (P2P=Point to Point, P2MP=Point to MultiPoint,
  MP2P=Multipoint to Point)

  Type     PVCs SoftPVCs     SVCs     TVCs     PVPs SoftPVPs     SVPs      Total
  P2P       982        0        0        0        0        0        0        982
  P2MP       36        0        0        0        0        0        0         36
  MP2P        0        0        0        0        0        0        0          0
                                        TOTAL INSTALLED CONNECTIONS =       1018

  PER-INTERFACE STATUS SUMMARY AT 16:42:51 UTC Fri Dec 8 2000:
     Interface      IF         Admin  Auto-Cfg    ILMI Addr     SSCOP    Hello
       Name       Status      Status    Status    Reg State     State    State
  ------------- -------- ------------ -------- ------------ --------- --------
  .
  (Information Deleted)
  .
➔ ATM0                UP           up      n/a  UpAndNormal      Idle      n/a
➔ ATM-SEC0            UP           up     done          n/a      none      n/a
  .
  (Information Deleted)
  .
```

Step 4 Verify that the IF Status field is up.

Step 5 Verify that the ILMI Addr Reg State field is UpAndNormal.

If you determine that redundancy is configured incorrectly, refer to the "Initially Configuring the ATM Switch Router" chapter in the *ATM Switch Router Software Configuration Guide*.

Troubleshooting Switch Processor Redundancy with HSRP (Catalyst 8540 CSR and Catalyst 8540 MSR)

Enhanced High System Availability (EHSA) is provided in the switching fabric when three switch processors are installed in the chassis. These features and their configuration are described in the "Initially Configuring the ATM Switch Router" chapter in the *ATM Switch Router Software Configuration Guide*.

To troubleshoot the EHSA of the switch processor, use the following commands:

Command	Purpose
show preferred-switch-card-slots	Displays the configuration of the switch processors.
show switch fabric	Displays the switch fabric details of the switch router.

Follow these steps to troubleshoot the redundant switch processors on the switch router:

Step 1 Use the **show preferred-switch-card-slots** command to confirm the configuration of the switch processors.

```
  Switch# show preferred-switch-card-slots
➔ The currently preferred switch card slots are slot: 5 and slot: 7
➔ The currently active switch card slots are slot: 5 and slot: 7

  Switch#
```

Step 2 Use the **show switch fabric** command to confirm the status of the switch processors and their location.

```
  Switch# show switch fabric
  swc_presence_mask: 0x5
  Switch mode: NR_20G
➔ Number of Switch Cards present in the Chassis: 3

  SWC SLOT             SWC_TYPE         SWC_STATUS
  =================================================

➔      5               EVEN              ACTIVE
➔      6               STANDBY           STANDBY
➔      7               ODD               ACTIVE


  MMC Switch Fabric (idb=0x62146E7C)
  .
  (Information Deleted)
  .
```

Step 3 Check the field Number of Switch Cards present in the Chassis. If you only have two switch processors, your switch router does not have EHSA. With three switch processors installed, if either of the two active switch processors fail, the third processor takes over.

Step 4 Check the SWC_STATUS field. This fields identifies the active and standby switch processors.

If you determine that redundancy is configured incorrectly, refer to the "Initially Configuring the ATM Switch Router" chapter in the *ATM Switch Router Software Configuration Guide*.

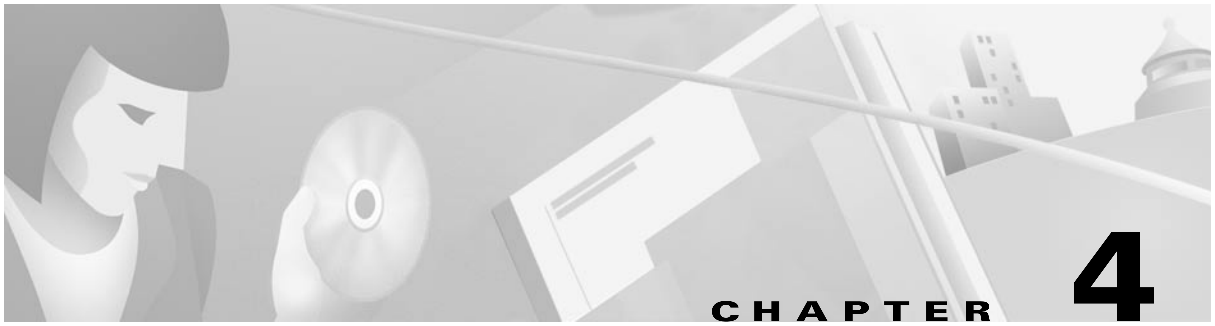

CHAPTER 4

Example Network

This chapter uses an example network to illustrate hardware configuration problems and configuration troubleshooting. The example network describes the most common connections you need to troubleshoot, and enables you to apply your own enterprise network configuration to the examples.

This chapter contains the following sections:

- Network Hierarchy, page 4-1
- Example Networks, page 4-2

Network Hierarchy

In a well-formed hierarchical network, there are three easily defined layers, traditionally referred to as the access, distribution, and core layers.

In an enterprise network, each layer provides different functions. Because these layers are not always recognized by their traditional names, the names have been modified to access or workgroup, distribution or policy, and core or backbone.

The access or workgroup layer connects users. Other functions of this layer are shared bandwidth, switched bandwidth, MAC-layer filtering, and micro segmentation. LAN switches, such as the Catalyst 5000 and Catalyst 6000 family of switches, exist most commonly in this layer of the network.

The distribution or policy layer performs the complex, CPU-intensive calculations such as filtering, access lists, inter-VLAN routing, Group Multicast Protocol (GMP), broadcast and multicast domain definition, and address or area aggregation. This layer might also contain the local servers. Routers, switch routers, and occasionally LAN switches reside in the distribution layer.

The core or backbone layer is the backbone of the network. It is high-speed and concerned with quick traffic switching. It does not get involved in expensive packet manipulation. In the following example network, ATM connections function together as the core backbone, and Fast Ethernet and Gigabit Ethernet connections function together as the redundant backbone core. The central servers might also be attached to the high-speed backbone in the core. Switch routers, high-speed routers, and occasionally LAN switches can be found in the core.

Example Networks

This section includes the following example networks:

- Example ATM Network, page 4-2
- Example Mixed ATM and Layer 3 Network, page 4-10

Example ATM Network

This section uses a fictitious network to describe actual problems in troubleshooting ATM switched networks.

Figure 4-1 provides a high level overview of the campus and remote networks.

Figure 4-1 Example Network Overview

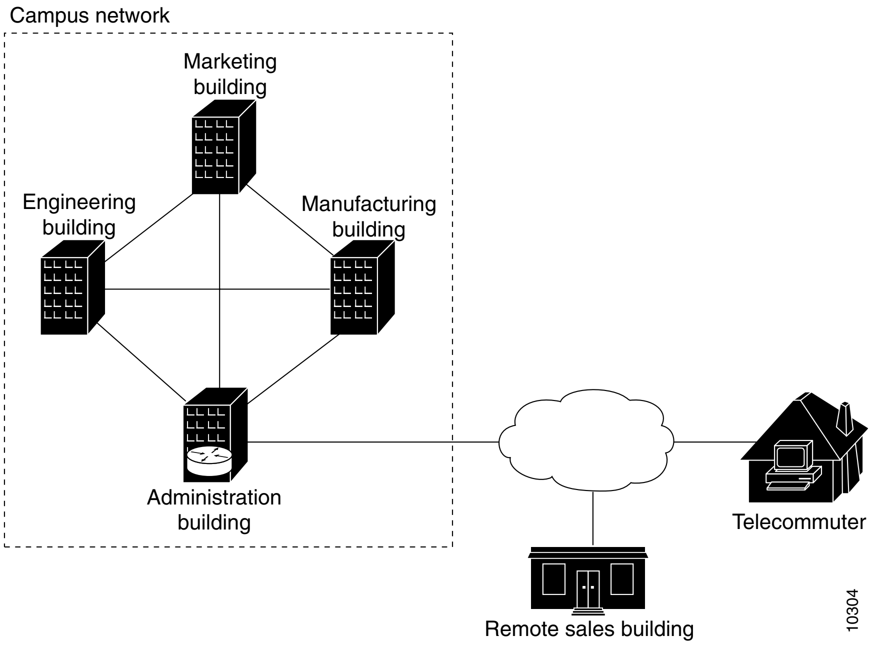

The example network in Figure 4-1 has the following components:

- A campus network of four 10-story buildings
- A remote sales building
- A telecommuter
- 4,000 employees on campus
- 4 buildings with 1,000 employees per building
- 5,000 total ports
- Microsoft NT servers and IP as the primary protocol
- Dynamic Host Configuration Protocol (DHCP) used to automatically allocate IP addresses to clients

- Approximately 100 users per Catalyst 5000 or Catalyst 5500 switch. This example network requires approximately 50 Catalyst 5000 or Catalyst 5500 switches:
  - One intermediate equipment closet per building that connects buildings with the ATM distribution switch routers
  - Fiber-optic connections between wiring closets and intermediate equipment closets
  - One-half of the users are on VLAN 2; the other half are on VLAN 3
- Network 10.0.0.0 255.255.255.0
- 254 hosts per subnet
- Spanning tree and root bridges enabled
- No single point of failure
- Workgroup servers that are connected using either ATM or Fast Ethernet in Layer 2
- Enterprise servers (e-mail, Web, and meeting scheduling) located in the administration building with the edge routers and firewall protection
- Switch routers that provide the following:
  - 155-Mbps unshielded twisted-pair (UTP) Optical Carrier 3 (OC-3) connections to servers and high-bandwidth users (computer-aided design [CAD], video, and voice) to the backbone
  - 2,488-Mbps single-mode fiber (OC-48) connections to the core between buildings in the intermediate wiring closets creating the backbone
  - T3 coaxial connections to the WAN
- Catalyst 5000 or Catalyst 600 family LAN switches provide the following:
  - Access and workgroup connection to individual users of the network
  - Workgroup server connections
  - Spanning-tree loop protection and network redundancy
- The remote site switch router has the following:
  - 500 employees
  - 750 total ports
- The telecommuter router has the following:
  - Dialup connections
  - ISDN
  - Frame Relay

Physical Connections

The example network contains the following physical connections:

- 155-Mbps UTP—Using permanent virtual path (PVP) and LAN emulation (LANE), connect distribution switch routers to Catalyst 5000 or Catalyst 6000 family LAN switches
- 622-Mbps multimode fiber and single-mode fiber —Using PVP, connect core switch routers with tag switching enabled

- T1 or E1—Using PVPs, connect to the WAN to reach remote sites such as WWW, FTP, Telnet, and e-mail
- T3 or E3—Using PVPs, connect to the WAN to reach remote sites such as WWW, FTP, Telnet, and e-mail
- T1 circuit emulation switch—Using PVP, connect to private branch exchange (PBX) or using switched virtual circuit (SVC), connect to coder/decoder (CODEC) for constant bit rate (CBR) video
- 25 Mbps—Connect to computer-aided design/computer-aided manufacturing (CAD/CAM) using a soft permanent virtual circuit (soft PVC) that provides the following QoS:
  - 10 Mbps: video
  - 5 Mbps: audio
  - 5 Mbps: unspecified bit rate (UBR) for data
- Frame Relay—Using PVC, connect to a telecommuter

Virtual Connections

The example network in Figure 4-2 has the following virtual connections:

- PVPs—Connections between buildings
- PVP tunnels—Connect to the remote site through the public network to avoid signalling
- SVCs—Connect to nodes that require longer data exchanges but infrequent connections (for example, e-mail server, CAD/CAM connections)
- PVC—Connect to nodes that need quick, short access without signalling delay (for example, Domain Name System [DNS] server connections)
- Soft PVC—Connect to the UNIX network interface cards (NICs) that do not support signalling (for example, SGI workstations)
- LAN emulation (LANE), which has the following connection types:
  - LAN emulation client (LEC)—Typical application from Catalyst 5000 and Catalyst 5500 to the switch router
  - LAN emulation configuration server/broadcast and unknown server (LECS/BUS)—Configure on a low-usage switch router, because the application is very route processor intensive
- Tag switching—Connect all core switches

Figure 3-2 shows the equipment overview of the example network, including the connection types of the network.

Figure 4-2 Equipment Overview of the Example ATM Network

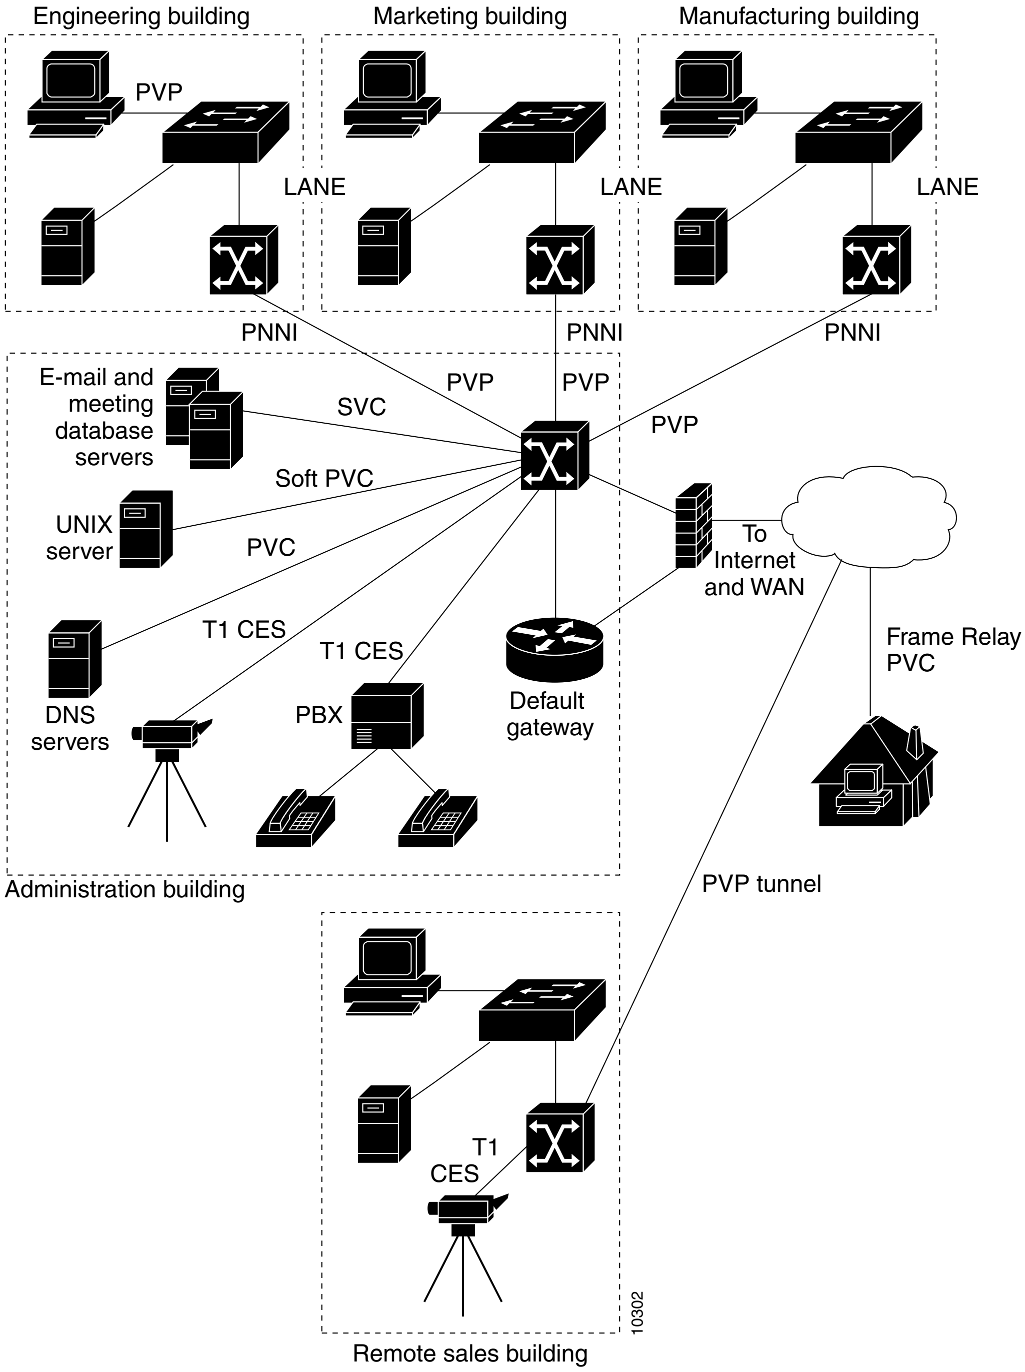

The engineering building in Figure 4-3 shows the following connections:

- 622-Mbps single-mode fiber connections between the ATM core switch router on Floor 1 and the campus backbone
- T1 CES access connection to CBR and QoS video CODEC for the video conference room
- 155-Mbps UTP SVC connections from the access switch router to the enterprise servers
- 155-Mbps UTP, multimode fiber, or single-mode fiber LANE SVC connection from distribution ATM switch routers in each wiring closet to Fast Ethernet access switch routers

Figure 4-3 Engineering Building ATM Connections

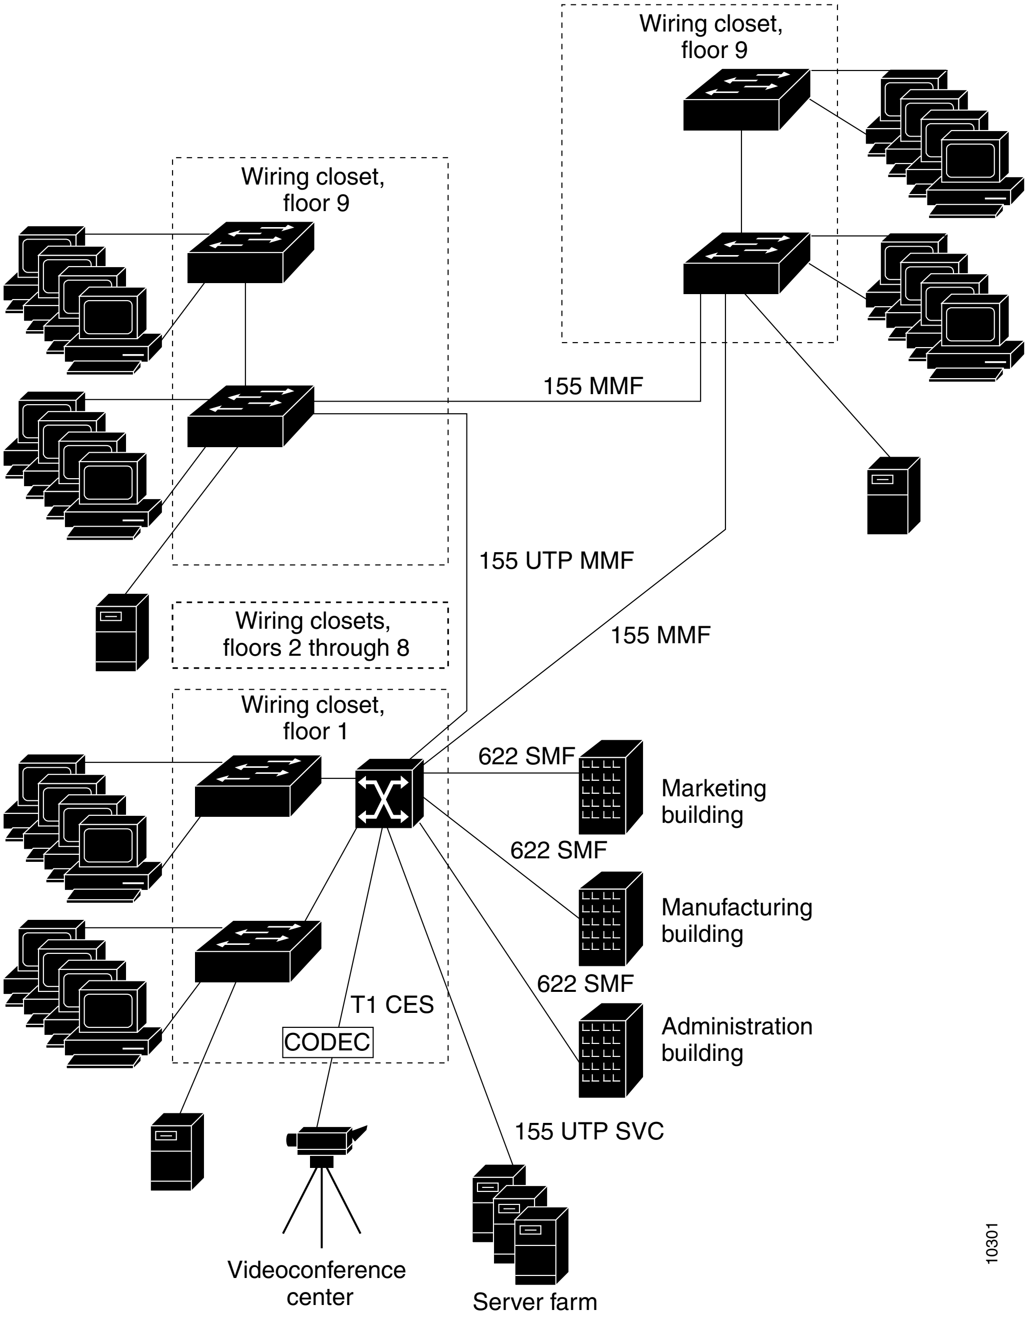

The typical Floor 1 wiring closet in Figure 4-4 shows the following connection examples:

- 622-Mbps single-mode fiber ATM core switch router connections to the backbone
- 25-Mbps port adapter providing 12 PVC access connections to CAD/CAM users with SGI workstations whose NICs do not support signalling
- T1 CES connection access connections to CBR and QoS video CODEC
- 155-Mbps UTP connection through LANE SVC to Fast Ethernet access switch router

Note Each Fast Ethernet distribution switch connection has a redundant link. (See Figure 4-4.)

Figure 4-4 Typical Floor 1 ATM Wiring Closet

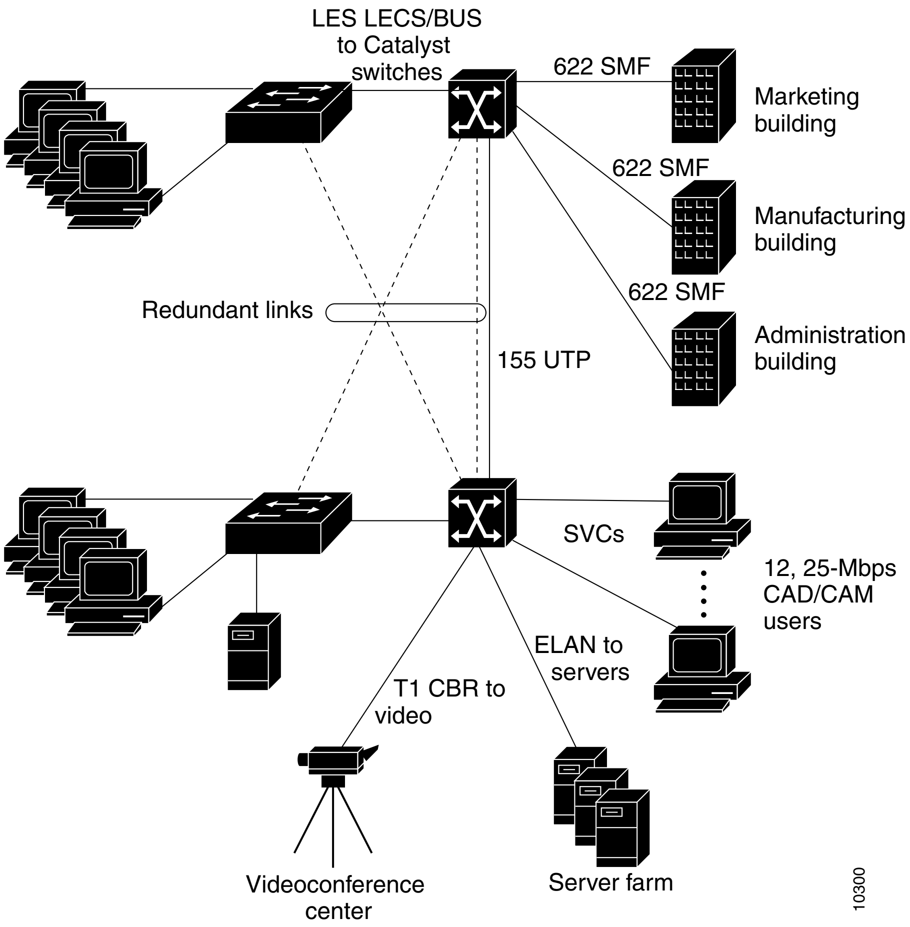

The typical core switch router configuration in Figure 4-5 shows the following connections:

- 622-Mbps single-mode fiber core connection through PVC for Private Network-Network Interface (PNNI) redundancy to other buildings
- 155-Mbps single-mode and multimode fiber distribution connection through PVC to ATM distribution switch routers within the building

Figure 4-5 Typical Core Switch Router ATM Configuration

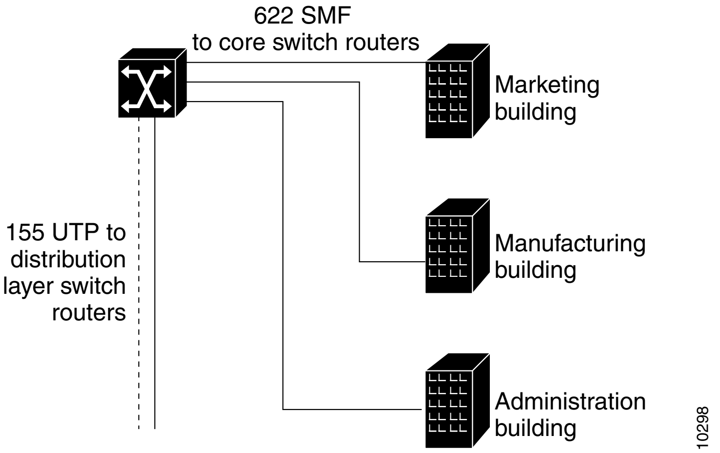

The typical distribution switch router configuration in Figure 4-6 shows the following connections:

- 155-Mbps UTP distribution connection through PVC PNNI between core switch routers
- 155-Mbps UTP distribution connection through LANE SVCs to the Cisco Systems Catalyst 5000 switches running LECS/BUS
- 155-Mbps UTP access connection through ELAN SVCs to individual servers
- 155-Mbps UTP or multimode fiber access connection through SVC with a CBR connection to CODEC for videoconferencing

Figure 4-6 Typical Distribution Switch Router ATM Configuration to Floor 1

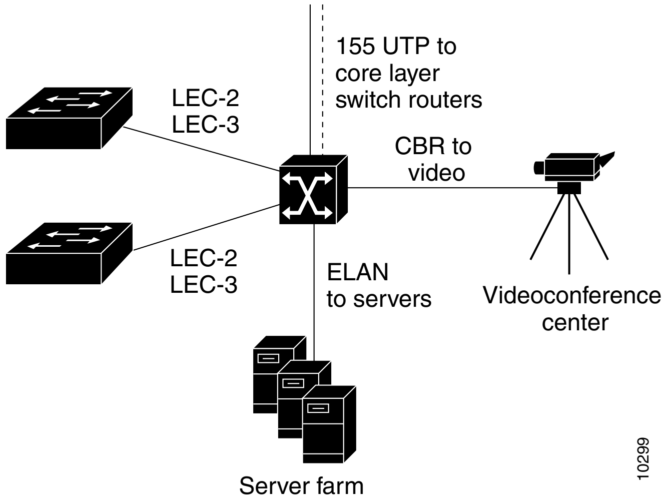

The administration building configuration in Figure 4-7 shows the following connections:

- 155-Mbps UTP connections using LANE SVC connections to e-mail servers, for example, that allow "bursty" traffic requiring signalling and less frequent use
- 155-Mbps UTP connections using PVC connections to DNS servers, for example, that allow short duration connections without signalling
- T3 connection to WAN with access filtering to Hypertext Transfer Protocol (HTTP) and other users
- 155-Mbps UTP connection to edge router or default gateway with ATM Interface Processor (AIP) installed and tag switching enabled
- T1 CES connection to PBX
- Video CBR using LANE SVC connections and T1 CES port adapters providing multicast connections to selected users
- Soft PVC from source video connection to a destination at a remote site
- PVP tunnel to the remote sales building
- Frame Relay PVC to the telecommuter

Figure 4-7 Administration Building ATM Connections

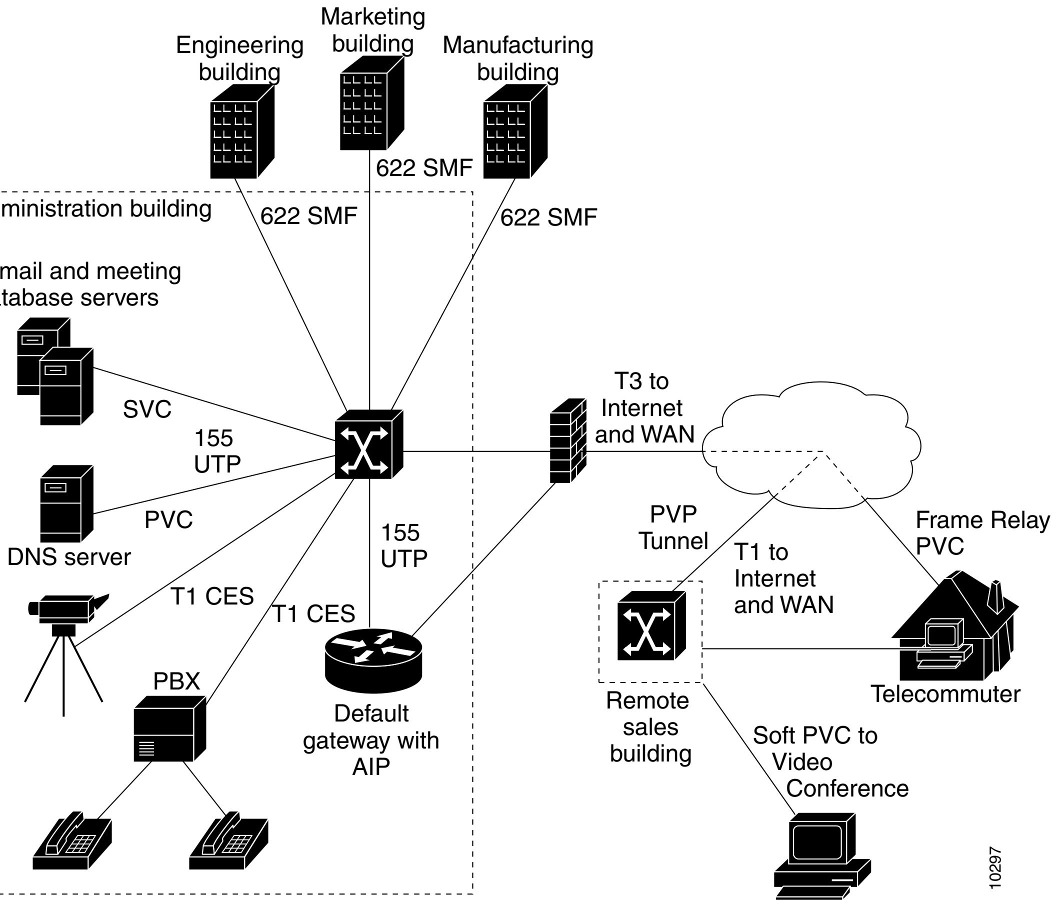

Example Mixed ATM and Layer 3 Network

This section uses the fictitious network described in the "Example ATM Network" section on page 4-2 to illustrate actual problems in troubleshooting a mixed ATM and Layer 3 switched network.

While the example network overview is the same as shown in Figure 4-1, there are additional redundant Layer 3 Gigabit Ethernet connections between buildings, LAN switches, and some high-usage servers. These redundant Gigabit Ethernet and Gigabit EtherChannel provide the high-capacity trunks needed to connect these gigabit switches if the primary ATM connections should fail.

Physical Connections

The example network contains the following physical connections:

- Gigabit Ethernet—Connect distribution Layer 3 switch routers to Catalyst 5000 or Catalyst 6000 family switches
- T1 or E1—Using PVPs, connect to the WAN to reach remote sites such as WWW, FTP, Telnet, and e-mail
- T3 or E3—Using PVPs, connect to the WAN to reach remote sites such as WWW, FTP, Telnet, and e-mail
- T1 circuit emulation switch—Using PVPs, connect to private branch exchange (PBX) or, using switched virtual circuit (SVC), connect to coder/decoder (CODEC) for constant bit rate (CBR) video
- Frame Relay—Using PVC, connect to a telecommuter

Virtual Connections

The example network in Figure 4-8 has the following virtual connections:

- PVP tunnels—Connect to the remote site through the public network to avoid signalling
- T1 CES access connection to CBR and QoS video CODEC for the video conference room

Figure 4-8 shows the equipment overview of the example network, including the connection types of the network.

Figure 4-8 Equipment Overview of the Example Mixed Layer 3 and ATM Network

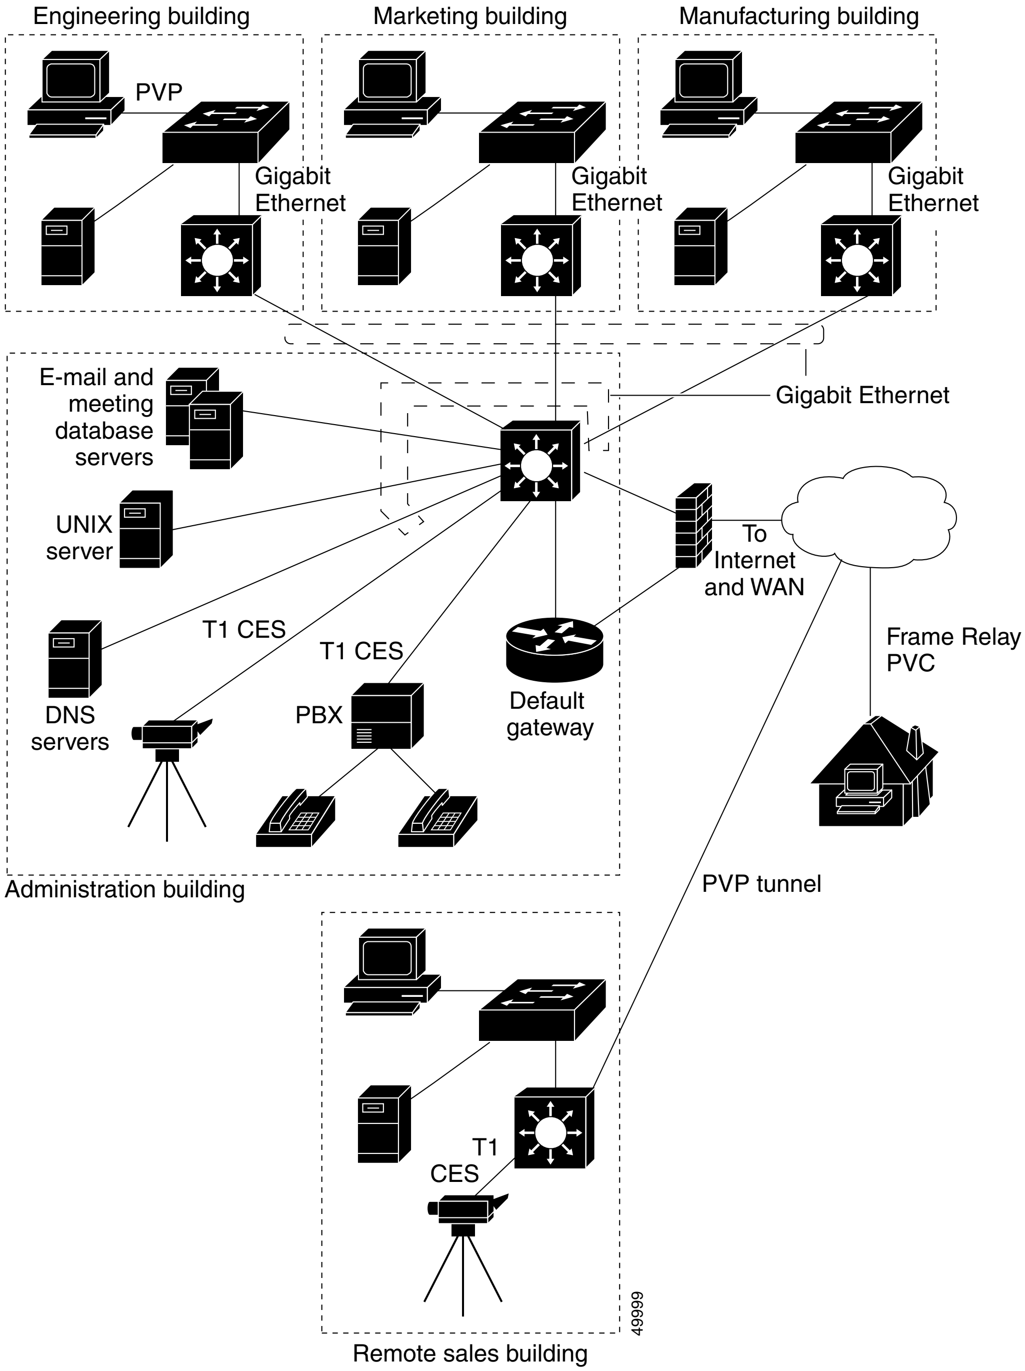

The engineering building in Figure 4-9 shows the following connections:

- Gigabit Ethernet single-mode fiber connections between the ATM core switch router on Floor 1 and the campus backbone
- T1 CES access connection to CBR and QoS video CODEC for the video conference room
- Gigabit Ethernet UTP SVC connections from the access switch router to the enterprise servers
- Gigabit Ethernet UTP, multimode fiber, or single-mode fiber connection from distribution ATM switch routers in each wiring closet to Fast Ethernet access switch routers

Figure 4-9 Engineering Building Layer 3 and ATM Connections

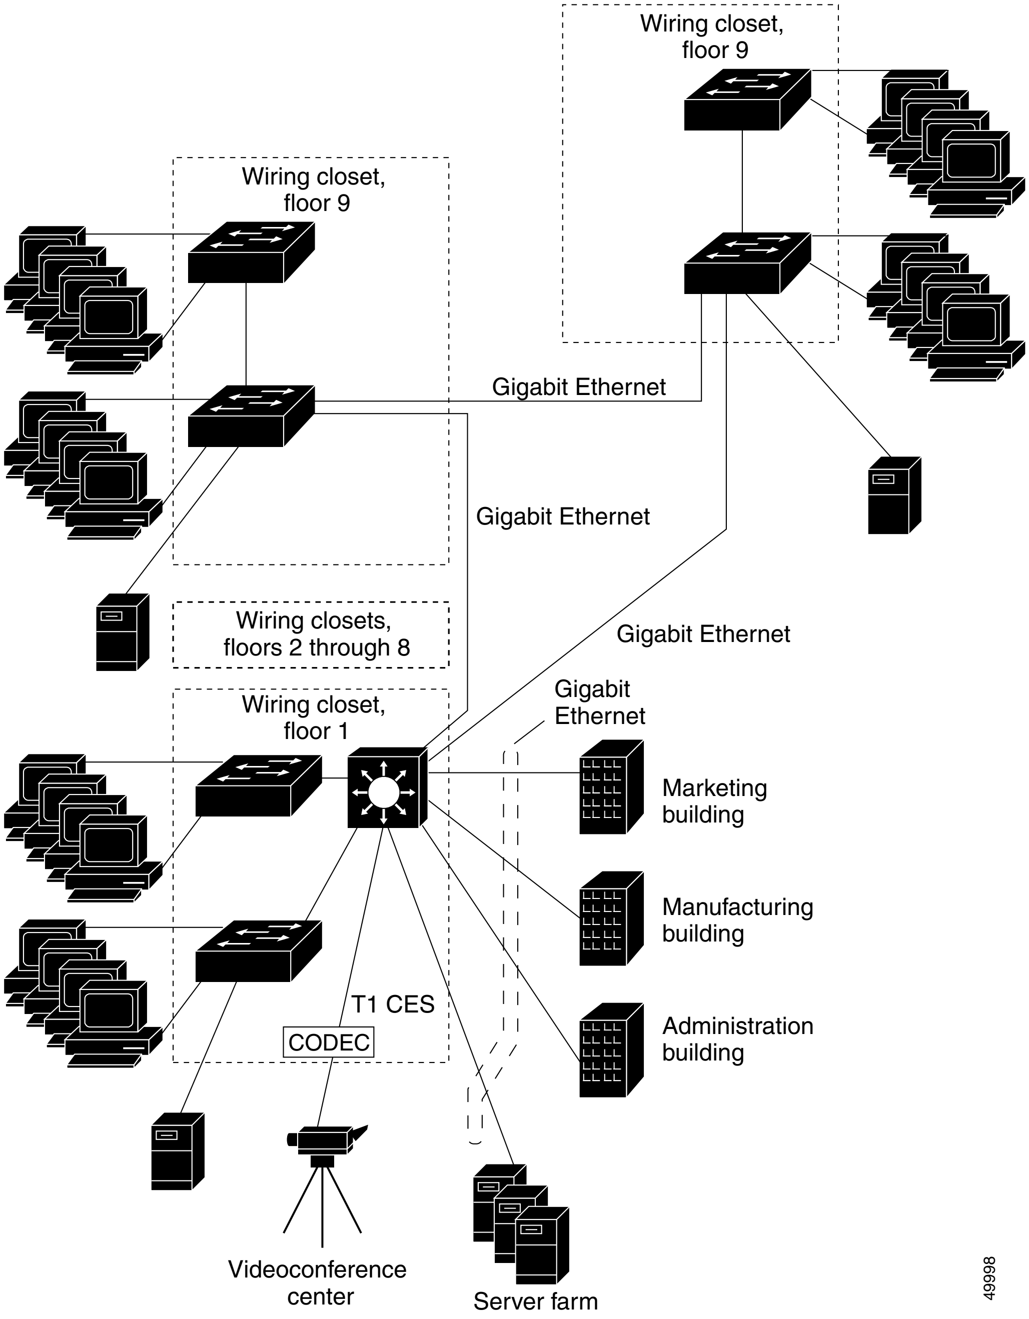

The typical Floor 1 wiring closet in Figure 4-10 shows the following connection examples:

- Gigabit Ethernet single-mode fiber ATM core switch router connections to the backbone
- T1 CES connection access connections to CBR and QoS video CODEC
- Gigabit Ethernet UTP connection to Fast Ethernet access switch

Note Each Gigabit Ethernet distribution switch connection has a redundant link. See Figure 4-10.

Figure 4-10 Typical Floor 1 Layer 3 and ATM Wiring Closet

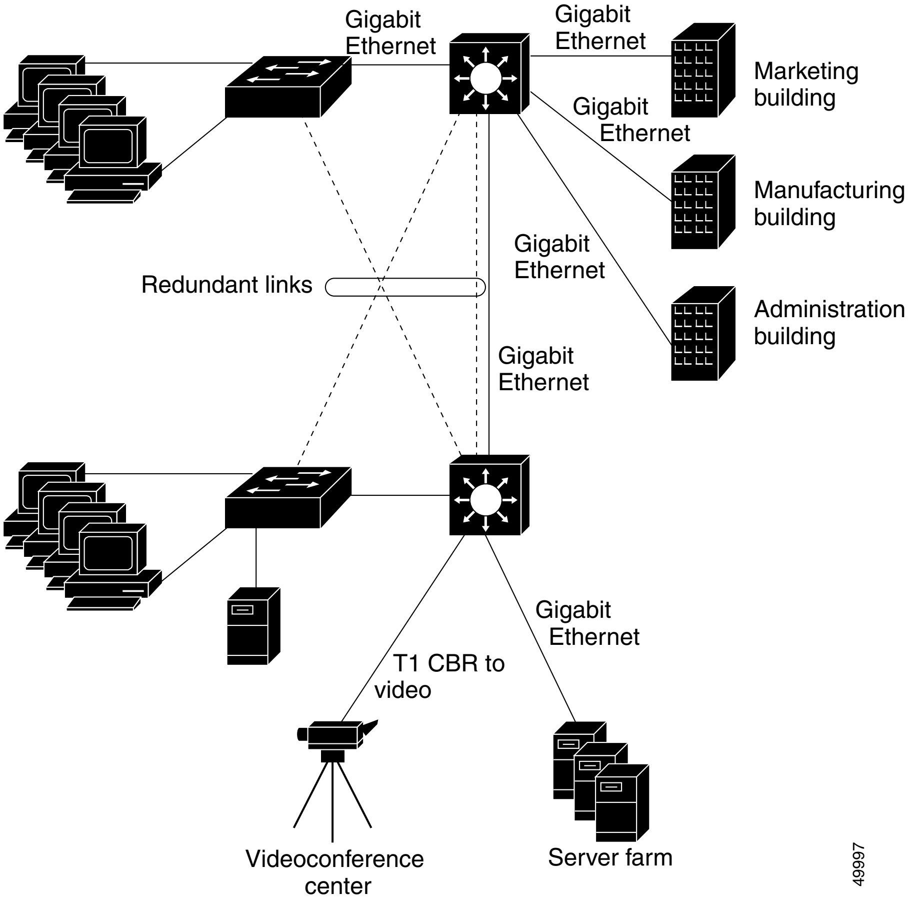

The typical core switch router configuration in Figure 4-11 shows the following connections:

- Gigabit Ethernet single-mode fiber core connection for redundancy to other buildings
- Gigabit Ethernet single-mode and multimode fiber distribution connection to ATM distribution switch routers within the building

Figure 4-11 Typical Layer 3 and ATM Core Switch Router Configuration

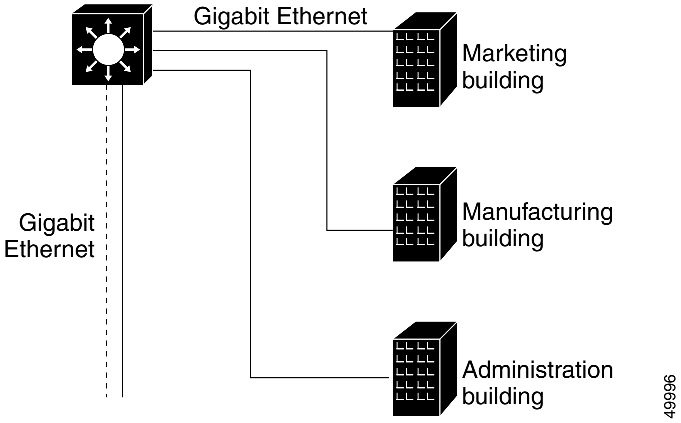

The typical distribution switch router configuration in Figure 4-12 shows the following connections:

- Gigabit Ethernet UTP distribution connection between core switch routers
- Gigabit Ethernet UTP distribution connection to the Cisco Systems Catalyst 5000 switches
- Gigabit Ethernet UTP access connection to individual servers
- 155-Mbps UTP or multimode fiber access connection through SVC with a CBR connection to CODEC for videoconferencing

Figure 4-12 Typical Layer 3 and ATM Distribution Switch Router Configuration to Floor 1

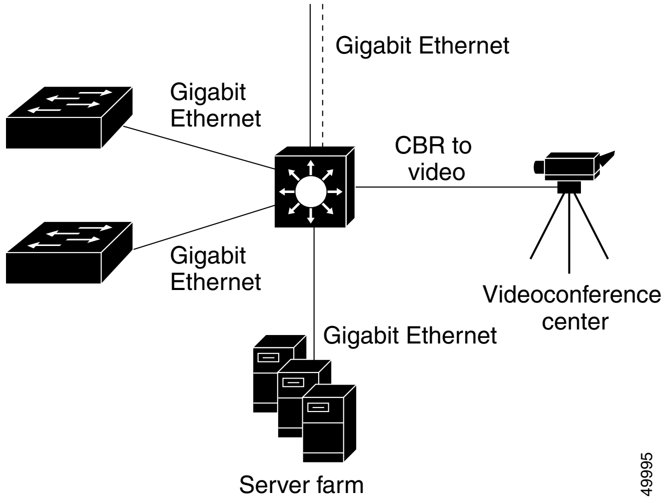

The administration building configuration in Figure 4-13 shows the following connections:

- Gigabit Ethernet UTP connections to e-mail servers
- Gigabit Ethernet UTP connections to DNS servers
- T3 connection to WAN with access filtering to Hypertext Transfer Protocol (HTTP) and other users
- Gigabit Ethernet UTP connection to edge router or default gateway
- T1 CES connection to PBX
- Video CBR using LANE SVC connections and T1 CES port adapters providing multicast connections to selected users
- Soft PVC from source video connection to a destination at a remote site
- PVP tunnel to the remote sales building
- Frame Relay PVC to the telecommuter

Figure 4-13 Administration Layer 3 and ATM Building Connections

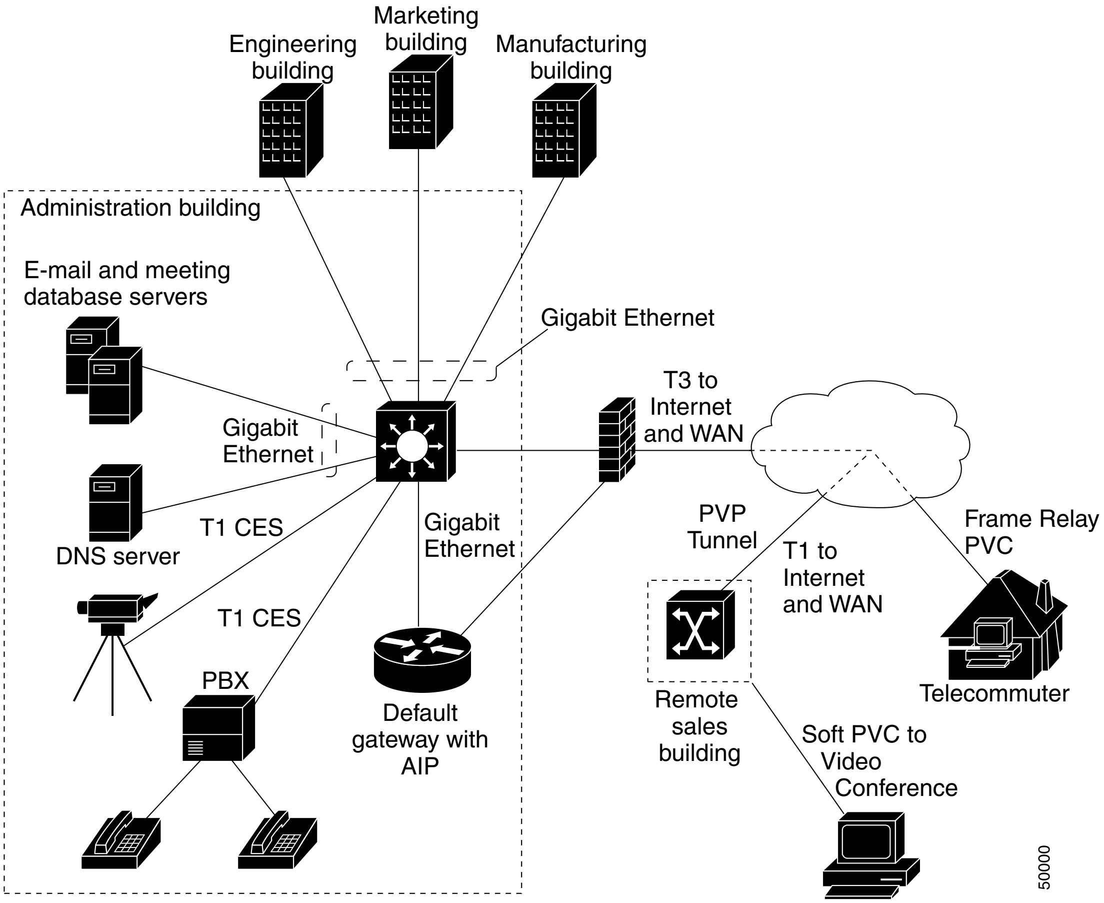

PART 1

ATM-to-ATM Connection Troubleshooting

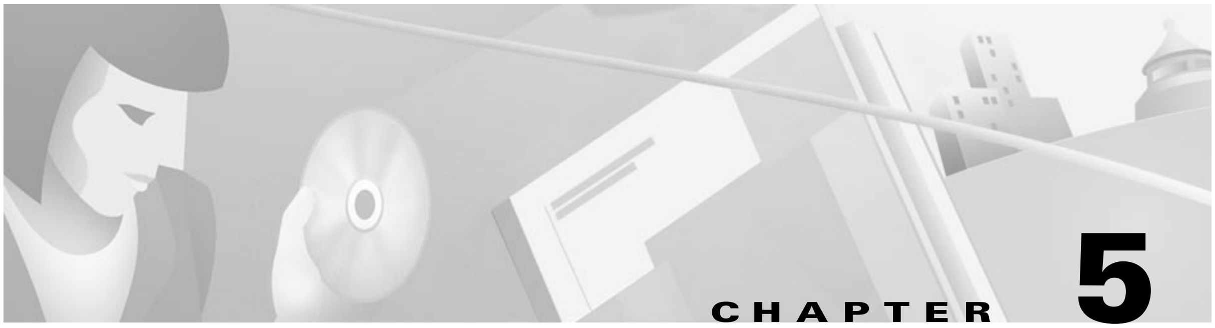

CHAPTER 5

Troubleshooting Switch Router ATM Interface Connections

This chapter provides troubleshooting information about connectivity and performance problems in the physical interfaces of a switch router.

The chapter includes the following sections:

- Performing Basic Interface Checks, page 5-1
- Determining Network Connectivity, page 5-5
- Performing OAM Loopback Tests, page 5-6
- Troubleshooting 155-Mbps and 622-Mbps Interfaces, page 5-14
- Troubleshooting OC-3c, OC-12c, and OC-48c Interfaces, page 5-18
- Troubleshooting T1 and E1 Interfaces, page 5-23
- Troubleshooting DS3 and E3 Interfaces, page 5-27
- Troubleshooting CES T1 and CES E1 Interfaces, page 5-32
- Troubleshooting 25-Mbps Interfaces, page 5-35
- Troubleshooting CDS3 Frame Relay Interfaces, page 5-38

Note For detailed cabling and hardware information for each port adapter, refer to the *ATM Port Adapter and Interface Module Installation Guide.* The default configurations for the various port adapters are described in the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide.*

Performing Basic Interface Checks

This section outlines the steps for performing basic interface checks and verifies that an ATM switch router interface is enabled and functioning correctly.

Always check the following when an interface fails:

- The port free of dust and debris?
- Are the cables inserted properly?
- Do the transmit and receive cable pairs match?
- Are the cables the proper type?

Note Just because the connector fits, does not mean the wires in the cable are necessarily cross-connected correctly, and the cable is not necessarily the correct type.

- Are the cables reliable? If not, try a different cable.
- Are the interfaces on both sides of the cables enabled and in no-shutdown mode?
- Are the interfaces configured properly (for example, framing mode, line coding, scrambling mode)?
- Are the interfaces on both ends of the cable the same type of interface?

Use the following command to check the ATM physical interface configuration:

Command	Purpose
show interfaces atm *card***/***subcard***/***port*	Shows the status of the physical interface.

Follow these steps to troubleshoot the physical interface:

Step 1 Use the **show interfaces atm** *card***/***subcard***/***port* command to display status and error information about an interface.

```
  Switch# show interfaces atm 1/0/0
➔ ATM1/0/0 is up, line protocol is up
    Hardware is oc3suni
    MTU 4470 bytes, sub MTU 4470, BW 156250 Kbit, DLY 0 usec, rely 255/255, load 1/255
➔   Encapsulation ATM, loopback not set, keepalive not supported
➔   Last input 00:00:00, output 00:00:00, output hang never
    Last clearing of "show interface" counters never
    Queueing strategy: fifo
    Output queue 0/40, 0 drops; input queue 0/75, 0 drops
    5 minute input rate 0 bits/sec, 0 packets/sec
    5 minute output rate 0 bits/sec, 0 packets/sec
       527152 packets input, 27939056 bytes, 0 no buffer
       Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
➔      0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
       527246 packets output, 27944038 bytes, 0 underruns
➔      0 output errors, 0 collisions, 0 interface resets
       0 output buffer failures, 0 output buffers swapped out
  Switch#
```

Step 2 Check the ATM or constant bit rate (CBR) field to see whether the interface is up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the line protocol field to see whether the status is up.

If down, check for the following:

- The line protocol software processes might have determined that the line is unusable. Try swapping the cable.
- The local or remote interface might be misconfigured. Check the interface configuration.
- The clocking might be misconfigured or the source interface might have failed. Refer to the "Initially Configuring the ATM Switch Router" chapter in the *ATM Switch Router Software Configuration Guide* for configuration information.
- The hardware might have failed. Try swapping the port adapter.

Step 4 Check the Encapsulation field. Confirm that the encapsulation method matches the interface type.

Step 5 Check the Last input and Last output fields. They show the number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface.

Step 6 Check the output hang field. It shows the number of hours, minutes, and seconds since the last reset caused by a lengthy transmission.

Step 7 Check the cyclic redundancy check (CRC) field. The presence of many CRC errors, but not many collisions, indicates excessive noise. If the number is too high, check the cables for damage. If you are using unshielded twisted-pair (UTP) cables, make sure you are using category 5 cable and not another type, such as category 3.

Note Errors and the input and output difference should not exceed 0.5 to 2.0 percent of traffic on the interface.

Step 8 Check the collisions field. It shows the total number of collisions compared to the total number of output packets and should be approximately 0.1 percent or less. If the number is too high, perform the following tasks:

- Use a protocol analyzer to check for late collisions. Collisions do not occur in a properly designed network. Collisions usually occur when cables are too long.
- Check the diameter of the network and make sure it is within specification.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide* for configuration information.

If you still have not determined the problem, continue with the next phase of basic interface troubleshooting.

Checking Cell Rates

This procedure determines if the cell rate for an interface is correctly configured.

Use the following command to check the cell rate on a physical interface:

Command	Purpose
show atm interface atm *card/subcard/port*	Confirms the ATM interface configuration.

Use the following steps to check the cell rate of the interface:

Step 1 Use the **show atm interface atm** command to display information about an interface.

```
  Switch# show atm interface atm 1/0/0
  Interface:      ATM1/0/0        Port-type:      oc3suni
→  Status:      UP              Admin Status:   up
  Auto-config:    enabled         AutoCfgState:   completed
  IF-Side:        Network         IF-type:        NNI
  Uni-type:       not applicable  Uni-version:    not applicable
  Max-VPI-bits:   8               Max-VCI-bits:   14
  Max-VP:         255             Max-VC:         16383
  ConfMaxSvpcVpi: 255             CurrMaxSvpcVpi: 255
  ConfMaxSvccVpi: 255             CurrMaxSvccVpi: 255
  ConfMinSvccVci: 33              CurrMinSvccVci: 33
  Svc Upc Intent: pass            Signalling:     Enabled
  ATM Address for Soft VC: 47.0091.8100.0000.0040.0b0a.2a81.4000.0c80.8000.00
  Configured virtual links:
    PVCLs SoftVCLs   SVCLs   TVCLs   PVPLs SoftVPLs   SVPLs Total-Cfgd Inst-Conns
        4        0       0       0       0        0       0          4          4
  Logical ports(VP-tunnels):     0
→ Input cells:    528135          Output cells:   528235
  5 minute input rate:             0 bits/sec,       0 cells/sec
  5 minute output rate:            0 bits/sec,       0 cells/sec
→ Input AAL5 pkts: 344844, Output AAL5 pkts: 344878, AAL5 crc errors: 0
  Switch#
```

Step 2 Check the IF (interface) Status and Admin (administration) Status fields to see whether they are up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the Input cells and Output cells fields. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, the interface is experiencing congestion and dropping cells.

Step 4 Check the AAL5 CRC errors field. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, check for the following:

- Many CRC errors, but not many collisions. This indicates excessive noise.
- Cable damage. If you are using UTP cables, make sure you are using category 5 cable and not another type, such as category 3.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide* for configuration information. If you still have not determined the problem, continue with the next phase of basic interface troubleshooting.

Determining Network Connectivity

To check ATM connection reachability and network connectivity, use the **ping atm interface atm** command in either privileged or user mode. You can use either an IP address or an ATM address prefix as a ping destination. You can also ping a neighbor switch router by selecting the segment loopback option. In privilege extended command mode, you can select other parameters such as repeat count and timeout values.

Command	Purpose
ping atm interface atm *card/subcard/port vpi vci* {**atm-prefix** *prefix* | **end-loopback** | **ip-address** *ip-address* | **seg-loopback**}	Checks the interface connection.

Follow these steps to ping a specific ATM prefix in both normal and extended mode:

Step 1 Use the **ping atm interface atm** command, in normal mode, to confirm connectivity through a specific interface to an ATM address prefix.

```
Switch# ping atm interface atm 1/0/0 0 5 atm-prefix 47.009181000000000000000001
Type escape sequence to abort.
Sending 5, 53-byte OAM Echoes to 47.0091.8100.0000.0000.0000.0001..., timeout is
 5 seconds:
!!!!!
➔ Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 2 Check the Success rate field. It should be 100 percent. If not, check the interface configuration.

Step 3 Use the **ping atm interface atm** command, in extended mode, to confirm connectivity through a specific interface to an ATM address prefix and modify the default repeat or timeout.

```
Switch# ping
Protocol [ip]: atm
Interface [card/subcard/port]: 1/0/0
VPI [0]: 0
VCI [0]: 5
Send OAM-Segment-Loopback ? [no]:
Target IP address:
Target NSAP Prefix: 47.009181000000000000000001
Repeat count [5]:
Timeout in seconds [5]:10
Type escape sequence to abort.
Sending 5, 53-byte OAM Echoes to 47.0091.8100.0000.0000.0000.0001..., timeout is
 10 seconds:
!!!!!
➔ Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 4 Check the Success rate field. It should be 100 percent. If not, check the interface configuration.

Note If you skip both destination IP address and the ATM prefix fields, the extended ping considers its neighbor switch as its destination and uses a segment-Loopback operation, administration, and maintenance (OAM) cell. In an IP address or ATM prefix case, the **ping** command always uses an end-to-end OAM loopback cell.

If the success rate is less than 100 percent, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide* and confirm the interface configuration.

Performing OAM Loopback Tests

OAM performs fault management and performance management functions at ATM management-plane (M-plane) layer.

Note Current OAM implementation supports only the fault management function, which includes connectivity verification and alarm surveillance.

The ATM switch router fully supports the following ATM OAM cell flows:

- F4 flows—OAM information flows between network elements used within virtual paths to report an unavailable path or a virtual path (VP) that cannot be guaranteed.
- F5 flows—OAM information flows between network elements used within virtual connections to report degraded virtual channel (VC) performance such as late arriving cells, lost cells, and cell insertion problems.

You can configure both F4 and F5 flows as either end-to-end or segment-loopback, and they can be used with alarm indication signal (AIS) and remote defect indication (RDI) functions.

Note Cells can be sent either on demand or periodically to verify link and connection integrity.

In addition to the standard OAM functions, the ATM switch router can also send OAM pings. See the "Determining Network Connectivity" section on page 5-5. Using OAM cells containing the ATM node addresses or IP addresses of intermediate switch routers, you can determine the integrity of a chosen connection at any intermediate point along that connection. With this information, you can debug and troubleshoot the network connection.

OAM Operation

OAM software implements ATM Layer F4 and F5 OAM fault management functions. OAM performs standard loopback (end-to-end or segment) and fault detection and notification (AIS and RDI) for each connection. It also maintains a group of timers for the OAM functions. When there is an OAM state change such as loopback failure, OAM software notifies the connection management software. The network operator can enable or disable OAM operation for the following ATM switch router components:

- The entire ATM switch router
- A specific ATM interface
- Each ATM connection

OAM AIS, RDI, and loopback operations are enabled or disabled for the entire switch router using the **atm oam** command in global configuration mode. Use the **atm oam** command in interface mode to configure OAM on a specific connection. For more information about configuring OAM operations, refer to the "Configuring Operation, Administration, and Maintenance" chapter in the *ATM Switch Router Software Configuration Guide*.

Note These OAM configuration commands are not stored in the nonvolatile random-access memory (NVRAM).

If OAM operation is disabled, outgoing OAM cells are not generated, and all incoming OAM cells are discarded.

To support various OAM operations, the ATM switch router hardware provides OAM cell routing functions on a per-connection basis for each direction and for different OAM cell spans (segment and end-to-end). The hardware OAM cell routing determines the destination of an OAM cell received from the link or the network and then determines whether OAM cells are processed by the ATM switch router software.

The hardware can perform the following functions on OAM cells:

- Intercept—Intercepted to the route processor queue and processed by the ATM switch router software
- Relay—Relayed along user cells by hardware without any software processing
- Discard—Discarded by hardware

An ATM connection consists of a group of network points, which are the edges of each ATM switch router or end system.

Each point can be one of the following:

- Connection endpoint—The end of a connection where the user ATM cells are terminated
- Segment endpoint—The end of a connection segment
- Connecting point—The middle point of a connection segment

Figure 5-1 shows the various loopback operations available.

Figure 5-1 OAM Loopback Operations

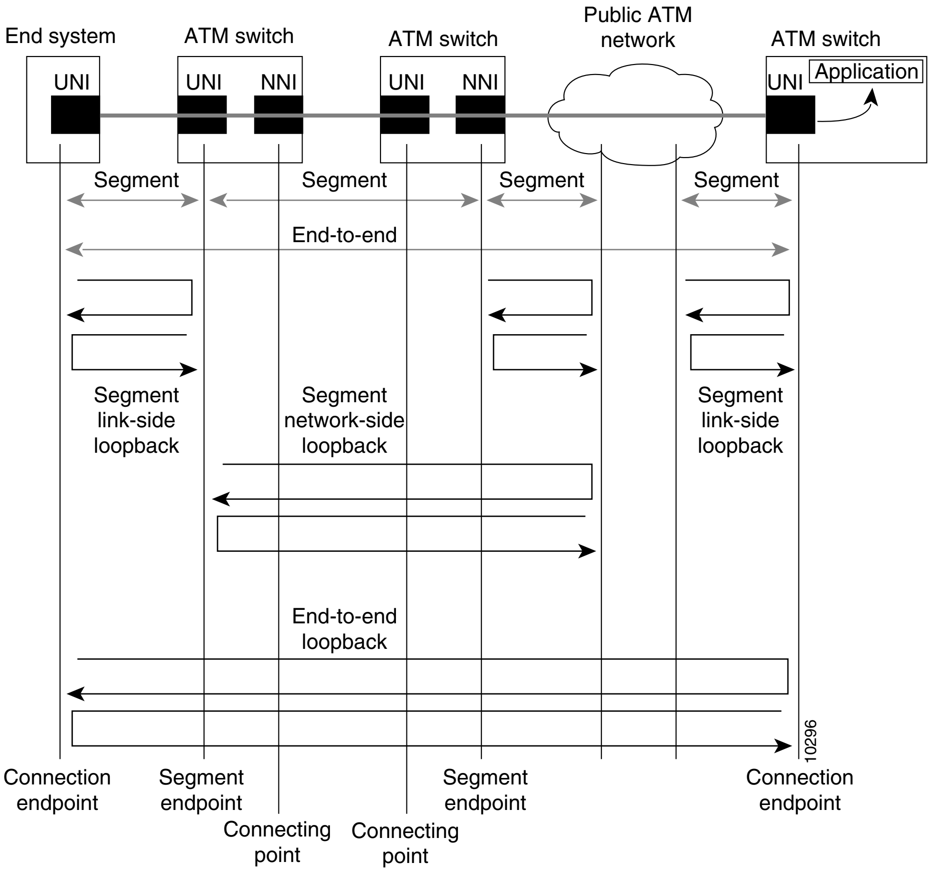

OAM Loopback Testing

You can use the loopback test to pinpoint faults by looping a signal at various points in the network. Use the loopback test before and after the initiation of service. Figure 5-2 shows how ATM OAM cell loopbacks are performed first across the interface and then across different segments of the connection.

Figure 5-2 Loopback Testing Process

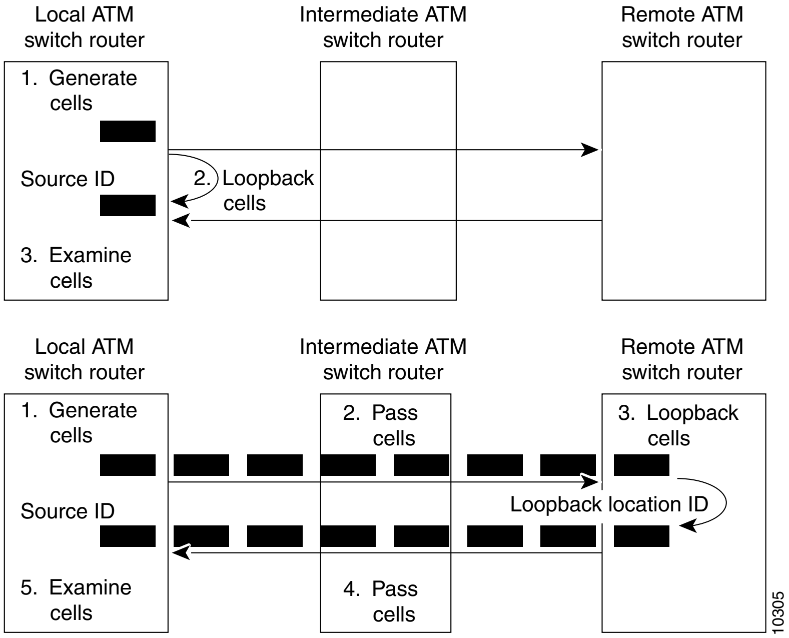

An ATM switch router generates the OAM cells and forwards them to another network element, which is responsible for returning them to the generating network elements.

Each loopback cell contains the ID of the generating network element and the ID of the network element that is looping the cells back to the originator. Any intermediate site must pass the cells on to the loopback site (the farthest point to which the cells progress) and the generating site (the point to which the cells return).

See Appendix C, “ATM Cell Structures,” for a format description of the OAM loopback cell.

The ATM switch router provides the following three types of loopback tests:

- Diagnostic
- Line
- PIF (physical interface)

Note If the loopback test is successful, data is reaching the I/O module properly. However, a successful test does not verify whether the I/O module correctly encodes the data sent onto the line.

Configuring Loopback Examples

The following examples show how to perform loopback tests on the interfaces shown in Figure 5-3.

If users connected to the Fast Ethernet Catalyst 5000 switch in the manufacturing building are not able to connect to the other users outside their building (including the DNS server in the administration building), you should try a loopback test. Use the procedures described in this section to test the ATM switch router connections starting at the middle section and proceeding outward.

Figure 5-3 Loopback Test Configuration Example

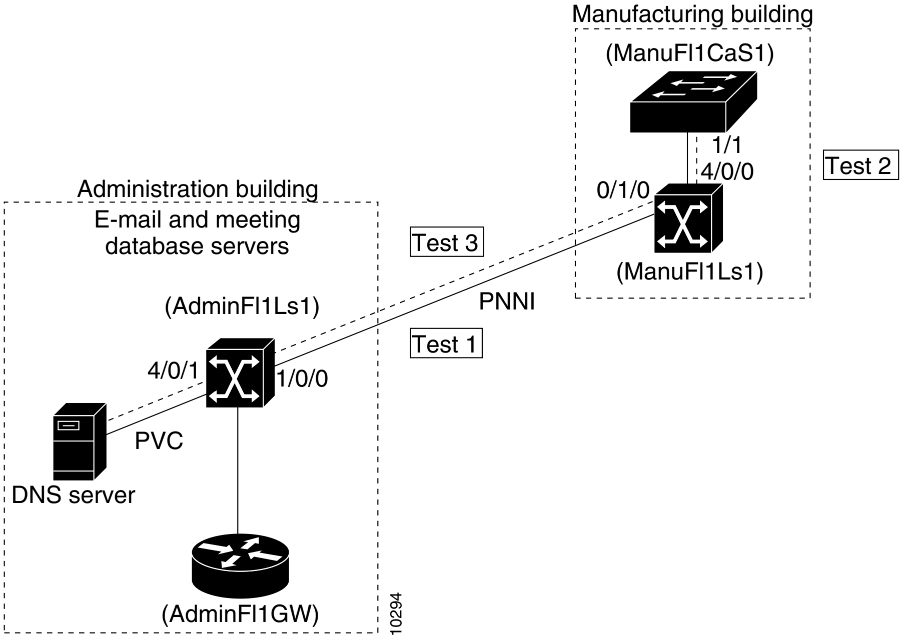

Use the **atm ping** command to confirm the ATM connection between the administration and manufacturing buildings.

Command	Purpose
ping atm interface atm *card/subcard/port vpi* [*vci*] [**atm-prefix** *prefix*] | [**end-loopback**] | [**ip-address** *ip-address*] | [**seg-loopback**]}	Checks the interface connection.

Perform the interface loopback tests in the following order:

Test 1—Segment network-side loopback between ATM switch router AdminFl1Ls1, interface 1/0/0, and ATM switch router ManuFl1Ls1, interface 4/0/1

Test 2—Segment link-side loopback between the DNS server and ATM switch router AdminFl1Ls1, interface 4/0/0

Test 3—End-to-end loopback between the DNS server and the Catalyst 5000 Fast Ethernet switch, ManuFl1CaS1, interface 1/1

Test 1—Segment Network-Side Loopback Process

Follow these steps to ping the ATM virtual channel 2, 130 between the administration and manufacturing buildings, with a segment loopback signal in normal mode:

Step 1 Use the **ping atm interface atm** *card/subcard/port* command to confirm the VP connectivity.

```
AdminFl1Ls1# ping atm interface atm 1/0/0 2 seg-loopback

Type escape sequence to abort.
Sending Seg-Loopback 5, 53-byte OAM Echoes to a neighbor, timeout is 5 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
AdminFl1Ls1#
```

Step 2 Use the same command to ping the ATM virtual channel 2, 130 between the administration and manufacturing buildings with a segment loopback signal in normal mode:

```
AdminFl1Ls1# ping atm interface atm 1/0/0 2 130 seg-loopback

Type escape sequence to abort.
Sending Seg-Loopback 5, 53-byte OAM Echoes to a neighbor, timeout is 5 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
AdminFl1Ls1#
```

Step 3 Check the Success rate field. If the success rate is less that 100 percent, you have a problem on the 622-Mbps connection between the administration and manufacturing buildings.

Step 4 Check the cables and the interface configuration, using the procedures in the "Performing Basic Interface Checks" section on page 5-1**.**

If the success rate is 100 percent, then this segment of the connection is not the problem. Proceed with the next phase of the interface loopback test.

Test 2—Segment Link-Side Loopback Process

Log in to the ATM switch router in the manufacturing building and use the **ping atm interface atm** command again to confirm the ATM connection between the ATM switch router and the Catalyst 5000 switches in the manufacturing building.

Use the following steps to ping the ATM virtual path 2 between the ATM switch router and the Catalyst 5000 switches in the manufacturing building, with a segment loopback signal in normal mode:

Step 1 Use the **ping atm interface atm** *card/subcard/port* command to confirm the VP connectivity.

```
ManuFl1Ls1# ping atm interface atm 4/0/0 2 seg-loopback

Type escape sequence to abort.
Sending Seg -Loopback 5, 53-byte OAM Echoes to a neighbor, timeout is 5 seconds:
!!!!!
➔ Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
ManuFl1Ls1#
```

Step 2 Check the Success rate field. If the success rate is less that 100 percent, there is a problem on the OC-3 155-Mbps connection between the ATM switch router and the Catalyst 5000 switch in the manufacturing building.

Step 3 Check the cables and the interface configuration, using the procedures in the "Performing Basic Interface Checks" section on page 5-1.

If the success rate is 100 percent, then this segment of the connection is not the problem. Proceed with the next phase of the interface loopback test.

Test 3—End-to-End Loopback Process

Check the end-to-end connection between the DNS server and the Catalyst 5000 switch in the manufacturing building.

Following is an example of the steps to ping the entire ATM virtual path between the administration and manufacturing buildings, with an end-to-end loopback signal in normal mode:

Step 1 Use the **ping atm interface atm** *card/subcard/port* command to confirm the VP connectivity.

```
AdminFl1Ls1# ping atm interface atm 4/0/0 2 end-loopback

Type escape sequence to abort.
Sending end-Loopback 5, 53-byte OAM Echoes to a neighbor, timeout is 5 seconds:
!!!!!
➔ Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Step 2 Check the Success rate field. If the success rate is less that 100 percent, you have a problem on the OC-3 155-Mbps connection between the ATM switch router and the Catalyst 5000 switch in the manufacturing building.

Step 3 Check the cables and the interface configuration using the procedures in the "Performing Basic Interface Checks" section on page 5-1.

If the success rate is 100 percent, then this segment of the connection is not the problem. Continue with the next phase of the interface test.

Using the debug Commands to Troubleshoot an Interface

The debug privileged EXEC commands can provide a wealth of information about the traffic being seen (or *not* seen) on an interface.

Caution Exercise care when using **debug** commands. Many of these commands are processor intensive and can cause serious network problems (such as degraded performance or loss of connectivity) if they are enabled on an already heavily loaded switch router. When you finish using a **debug** command, remember to disable it with its specific **no debug** command (or use the **no debug all** command to turn off all debugging).

For detailed information about using the **debug** commands, see Appendix A, "Debugging a Switch Router."

To isolate problems and troubleshoot the physical connections of the ATM switch router, use the following **debug** commands in privileged EXEC mode. Use the **no** form of these commands to disable debugging.

Command	Purpose
debug ports {**aal5** [**interface atm** *card*/*subcard*/*port*] | **dcu** | **ds3e3** | **netclock** | **oc12** | **oc3** | **t1e1**}	Starts debugging at the driver level for a specific port.
debug atm oam-all	Starts debugging, using generic OAM cells.
debug atm oam-pkt	Starts debugging, using OAM packets.
debug atm errors	Starts debugging to display all ATM errors.
no debug all	Disables all debugging.

Refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide* to confirm the interface configuration.

Troubleshooting 155-Mbps and 622-Mbps Interfaces

This section describes specific processes and commands used to troubleshoot the 155-Mbps and 622-Mbps port adapters.

Port Adapter LEDs

The port adapter faceplate LEDs provide status information for individual 155-Mbps and 622-Mbps single-mode and multimode fiber-optic and UTP interface connections of the port adapter. The LEDs are described in Table 5-1.

Note Use the **show controllers** command to display the LED status.

Table 5-1 155-Mbps and 622-Mbps Port Adapter LED Descriptions

LED	Status	Description
RX (Receive)	Off Flashing green Red	LOS[1] or port adapter is shut down. Cells are being received. LED blinks every 5 seconds and pulse rate increases with data rate. Alarm (LOF[2], LCD[3], AIS[4]).
TX (Transmit)	Off Flashing green Flashing yellow Steady yellow	No transmit line activity indication. Cells are being transmitted. LED pulse rate increases with data rate. Loopback. FERF[5] alarm.

[1] LOS = loss of signal

[2] LOF = loss of frame

[3] LCD = loss of cell delineation

[4] AIS = alarm indication signal

[5] FERF = far-end receive failure

Note Single-mode fiber-optic interface connectors are blue, and multimode connectors are black.

Displaying Interface Port Configuration

To display the interface configuration, use the following commands:

Command	Purpose
show interfaces atm *card/subcard/port*	Shows the status of the physical interface.
show atm interface atm *card/subcard/port*	Shows the interface configuration.
show controllers atm *card/subcard/port*	Shows the interface memory management and error counters.

Follow these steps to troubleshoot a 155-Mbps or 622-Mbps physical interface:

Step 1 Use the **show interfaces atm** *card/subcard/port* command to check the configuration.

```
   Switch# show interfaces atm 1/0/0
➔  ATM1/0/0 is up, line protocol is up
     Hardware is oc3suni
     MTU 4470 bytes, sub MTU 4470, BW 156250 Kbit, DLY 0 usec, rely 255/255, load 1/255
➔    Encapsulation ATM, loopback not set, keepalive not supported
➔    Last input 00:00:00, output 00:00:00, output hang never
     Last clearing of "show interface" counters never
     Queueing strategy: fifo
     Output queue 0/40, 0 drops; input queue 0/75, 0 drops
     5 minute input rate 2000 bits/sec, 6 packets/sec
     5 minute output rate 3000 bits/sec, 9 packets/sec
        4703704 packets input, 249296312 bytes, 0 no buffer
        Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
➔       54 input errors, 55 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
        5737496 packets output, 304087288 bytes, 0 underruns
➔       0 output errors, 0 collisions, 0 interface resets
        0 output buffer failures, 0 output buffers swapped out
   Switch#
```

Step 2 Check the ATM field to see whether the interface is up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the line protocol field to see whether the status is up.

If down, check for the following:

- The line protocol software processes might have determined that the line is unusable. Try swapping the cable.
- The local or remote interface might be misconfigured. Check the interface configuration.
- Clocking might be misconfigured or the source interface might have failed. Refer to the "Initially Configuring the ATM Switch Router" chapter in the *ATM Switch Router Software Configuration Guide*.
- Hardware might have failed. Try swapping the port adapter.

Step 4 Check the Encapsulation field. Confirm that the encapsulation method matches the interface type.

Step 5 Check the Last input and Last output fields. They show the number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface.

Step 6 Check the output hang field. It shows the number of hours, minutes, and seconds since the last reset caused by a lengthy transmission.

Step 7 Check the CRC field. The presence of many CRC errors, but not many collisions, indicates excessive noise. If the number of errors is too high, check the cables for damage. If you are using UTP cable, make sure you are using category 5 cables and not another type, such as category 3.

Note Errors and the input and output difference should not exceed 0.5 to 2.0 percent of traffic on the interface.

Step 8 Check the collisions field. It shows the total number of collisions compared to the total number of output packets, and it should be approximately 0.1 percent or less. If the number is too high, perform the following tasks:

- Use a protocol analyzer to check for late collisions. Collisions do not occur in a properly designed network. Collisions usually occur when cables are too long.
- Check the diameter of the network and make sure it is within specification.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide.*

Follow these steps to show the configuration of a 155-Mbps or 622-Mbps interface:

Step 1 Use the **show atm interface atm** *card/subcard/port* command to check the configuration.

```
  Switch# show atm interface atm 1/0/0
  Interface:      ATM1/0/0        Port-type:      oc3suni
➔ IF Status:      UP              Admin Status:   up
  Auto-config:    enabled         AutoCfgState:   completed
  IF-Side:        Network         IF-type:        UNI
  Uni-type:       Private         Uni-version:    V3.1
  Max-VPI-bits:   2               Max-VCI-bits:   10
  Max-VP:         255             Max-VC:         16383
  ConfMaxSvpcVpi: 255             CurrMaxSvpcVpi: 3
  ConfMaxSvccVpi: 255             CurrMaxSvccVpi: 3
  ConfMinSvccVci: 33              CurrMinSvccVci: 33
  Svc Upc Intent: pass            Signalling:     Enabled
  ATM Address for Soft VC: 47.0091.8100.0000.0000.0000.0001.4000.0c80.8000.00
  Configured virtual links:
    PVCLs SoftVCLs   SVCLs   TVCLs   PVPLs SoftVPLs   SVPLs Total-Cfgd Inst-Conns
        2        0      12       0       0        0       0         14         16
  Logical ports(VP-tunnels):     0
➔ Input cells:    4703972         Output cells:   5737883
  5 minute input rate:          2000 bits/sec,       4 cells/sec
  5 minute output rate:         4000 bits/sec,       9 cells/sec
➔ Input AAL5 pkts: 169899, Output AAL5 pkts: 644764, AAL5 crc errors: 0
  Switch#
```

Step 2 Check the IF Status and Admin Status fields to see whether they are up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the Input cells and Output cells fields. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, the interface is experiencing congestion and dropping cells.

Step 4 Check the AAL5 crc errors field. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, check for the following:

- Many CRC errors, but not many collisions. This indicates excessive noise.
- Cable damage. If you are using UTP cables, make sure you are using category 5 cables and not another type, such as category 3.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to display the memory management and error counters of a 155-Mbps or 622-Mbps ATM interface:

Step 1 Use the **show controllers atm** *card/subcard/port* command to check memory management and error counters.

```
  Switch# show controllers atm 1/0/0
  IF Name: ATM1/0/0    Chip Base Address: A8A08000
  Port type: OC3    Port rate: 155 Mbps    Port medium: MM Fiber
➔ Port status:Good Signal    Loopback:None    Flags:8308
➔ TX Led: Traffic Pattern    RX Led: Traffic Pattern  TX clock source: network-derived
➔ Framing mode:  sts-3c
➔ Cell payload scrambling on
➔ Sts-stream scrambling on
  OC3 counters:
    Key: txcell - # cells transmitted
         rxcell - # cells received
         b1     - # section BIP-8 errors
         b2     - # line BIP-8 errors
         b3     - # path BIP-8 errors
         ocd    - # out-of-cell delineation errors - not implemented
         g1     - # path FEBE errors
         z2     - # line FEBE errors
         chcs   - # correctable HEC errors
         uhcs   - # uncorrectable HEC errors

  <Information Deleted>

  phy_tx_cnt:4789577, phy_rx_cnt:4704918
  Switch#
```

Step 2 Check the Port status field. It should read "Good Signal."

Step 3 Check the Loopback field. It should read "None."

Step 4 Check the TX Led field. It should read "Traffic Pattern." If it does not, see Table 5-1 for LED descriptions.

Step 5 Check the RX Led field. It should read "Traffic Pattern." If it does not, see Table 5-1 for LED descriptions.

Step 6 Check the Framing mode field. It should match the framing mode configuration of the destination port.

Step 7 Check the Cell payload scrambling field. It should match the cell payload scrambling mode configuration of the destination port.

Step 8 Check the Sts-stream scrambling field. It should match the STS stream scrambling mode configuration of the destination port.

If you determine that the physical interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *ATM Switch Router Software Configuration Guide*.

If the interface is still not operating correctly, proceed with the troubleshooting process in Chapter 6, “Troubleshooting Switch Router ATM Network Connections.”

Troubleshooting OC-3c, OC-12c, and OC-48c Interfaces

This section describes specific processes and commands used to troubleshoot the OC-3c, OC-12c, and OC-48c interface modules.

Interface Module LEDs

The interface module faceplate LEDs provide status information for individual single-mode and multimode fiber-optic interface connections of the interface module. The LEDs are described in Table 5-2.

Note Use the **show controllers** command to display the LED status.

Table 5-2 OC-3c, OC-12c, and OC-48c Interface Module LED Descriptions

LED	Status	Description
LINK	Off Green	Carrier detect signal not received. Carrier detect signal received.
RX (Receive)	Off Flashing green Red	LOS or interface module is shut down. Cells are being received. LED blinks every 5 seconds and pulse rate increases with data rate. Alarm (LOF[1], OCD[2], AIS[3], LOP[4], RDI[5], LCD[6], UNEQ[7], PLM[8]).
TX (Transmit)	Off Flashing green Flashing yellow Steady yellow	No transmit line activity indication. Cells are being transmitted. LED pulse rate increases with data rate. Loopback. RDI.

[1] LOF = loss of frame

[2] OCD = out of cell delineation

[3] AIS = alarm indication signal

[4] LOP = loss of pointer

[5] RDI = remote defect indicator

[6] LCD = loss of cell delineation (OC-48c)

[7] UNEQ = unequipped code (OC-48c)

[8] PLM = payload label mismatch (OC-48c)

Note Single-mode fiber-optic interface connectors are blue, and multimode connectors are black.

Displaying Interface Port Configuration

To display the interface configuration, use the following commands:

Command	Purpose
show interfaces atm *card/subcard/port*	Shows the status of the physical interface.
show atm interface atm *card/subcard/port*	Shows the interface configuration.
show controllers atm *card/subcard/port*	Shows the interface memory management and error counters.

Follow these steps to troubleshoot an OC-3c, OC-12c, or OC-48c physical interface:

Step 1 Use the **show interfaces atm** *card/subcard/port* command to check the configuration.

```
➔ Switch# show interfaces atm 11/0/0
  ATM11/0/0 is down, line protocol is down
    Hardware is oc48c
    MTU 4470 bytes, sub MTU 4470, BW 2488320 Kbit, DLY 0 usec, rely 0/255, load 15
➔   Encapsulation ATM, loopback not set, keepalive not supported
➔   Last input never, output never, output hang never
    Last clearing of "show interface" counters never
    Queueing strategy: fifo
    Output queue 0/40, 0 drops; input queue 0/75, 0 drops
    5 minute input rate 0 bits/sec, 0 packets/sec
    5 minute output rate 0 bits/sec, 0 packets/sec
       0 packets input, 0 bytes, 0 no buffer
       Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
➔      0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
       0 packets output, 0 bytes, 0 underruns
➔      0 output errors, 0 collisions, 0 interface resets
       0 output buffer failures, 0 output buffers swapped out
  Switch#
```

Step 2 Check the ATM field to see whether the interface is up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the line protocol field to see whether the status is up.

If the interface is down, check for the following:

- The line protocol software processes might have determined that the line is unusable. Try swapping the cable.
- The local or remote interface might be misconfigured. Check the interface configuration.
- Clocking might be misconfigured or the source interface might have failed. Refer to the "Initially Configuring the ATM Switch Router" chapter in the *ATM Switch Router Software Configuration Guide*.
- Hardware might have failed. Try swapping the interface module.

Step 4 Check the Encapsulation field. Confirm that the encapsulation method matches the interface type.

Step 5 Check the Last input and Last output fields. They show the number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface.

Step 6 Check the output hang field. It shows the number of hours, minutes, and seconds since the last reset caused by a lengthy transmission.

Step 7 Check the CRC field. The presence of many CRC errors, but not many collisions, indicates excessive noise. If the number of errors is too high, check the cables for damage. If you are using UTP cable, make sure you are using category 5 cables and not another type, such as category 3.

Note Errors and the input and output difference should not exceed 0.5 to 2.0 percent of traffic on the interface.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide.*

Follow these steps to show the configuration of an OC-3c, OC-12c, or OC-48c interface:

Step 1 Use the **show atm interface atm** *card/subcard/port* command to check the configuration.

```
  Switch# show atm interface atm 1/0/0
  Interface:      ATM1/0/0        Port-type:      oc3suni
→ IF Status:      UP              Admin Status:   up
  Auto-config:    enabled         AutoCfgState:   completed
  IF-Side:        Network         IF-type:        UNI
  Uni-type:       Private         Uni-version:    V3.1
  Max-VPI-bits:   2               Max-VCI-bits:   10
  Max-VP:         255             Max-VC:         16383
  ConfMaxSvpcVpi: 255             CurrMaxSvpcVpi: 3
  ConfMaxSvccVpi: 255             CurrMaxSvccVpi: 3
  ConfMinSvccVci: 33              CurrMinSvccVci: 33
  Svc Upc Intent: pass            Signalling:     Enabled
  ATM Address for Soft VC: 47.0091.8100.0000.0000.0000.0001.4000.0c80.8000.00
  Configured virtual links:
    PVCLs SoftVCLs   SVCLs   TVCLs   PVPLs SoftVPLs   SVPLs Total-Cfgd Inst-Conns
        2        0      12       0       0        0       0         14         16
  Logical ports(VP-tunnels):     0
→ Input cells:    4703972         Output cells:   5737883
  5 minute input rate:          2000 bits/sec,       4 cells/sec
  5 minute output rate:         4000 bits/sec,       9 cells/sec
→ Input AAL5 pkts: 169899, Output AAL5 pkts: 644764, AAL5 crc errors: 0
  Switch#
```

Step 2 Check the IF Status and Admin Status fields to see whether they are up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the Input cells and Output cells fields. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, the interface is experiencing congestion and dropping cells.

Step 4 Check the AAL5 crc errors field. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, check for the following:

- Many CRC errors, but not many collisions. This indicates excessive noise.
- Cable damage. If you are using UTP cables, make sure you are using category 5 cables and not another type, such as category 3.

If you determine that the physical interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to display the memory management and error counters of an OC-3c, OC-12c, or OC-48c ATM interface:

Step 1 Use the **show controllers atm** *card/subcard/port* command to check memory management and error counters.

```
  Switch# show controllers atm 1/0/0
  IF Name: ATM1/0/0    Chip Base Address: A8A08000
  Port type: OC3    Port rate: 155 Mbps    Port medium: MM Fiber
➔ Port status:Good Signal    Loopback:None    Flags:8308
➔ TX Led: Traffic Pattern    RX Led: Traffic Pattern  TX clock source: network-derived
➔ Framing mode:  sts-3c
➔ Cell payload scrambling on
➔ Sts-stream scrambling on
  OC3 counters:
    Key: txcell - # cells transmitted
         rxcell - # cells received
         b1     - # section BIP-8 errors
         b2     - # line BIP-8 errors
         b3     - # path BIP-8 errors
         ocd    - # out-of-cell delineation errors - not implemented
         g1     - # path FEBE errors
         z2     - # line FEBE errors
         chcs   - # correctable HEC errors
         uhcs   - # uncorrectable HEC errors

  <Information Deleted>

  phy_tx_cnt:4789577, phy_rx_cnt:4704918
  Switch#
```

Step 2 Check the Port status field. It should read “Good Signal.”

Step 3 Check the Loopback field. It should read “None.”

Step 4 Check the TX Led field. It should read “Traffic Pattern.” If it does not, see Table 5-1 for LED descriptions.

Step 5 Check the RX Led field. It should read “Traffic Pattern.” If it does not, see Table 5-1 for LED descriptions.

Step 6 Check the Framing mode field. It should match the framing mode configuration of the destination port.

Step 7 Check the Cell payload scrambling field. It should match the cell payload scrambling mode configuration of the destination port.

Step 8 Check the Sts-stream scrambling field. It should match the STS stream scrambling mode configuration of the destination port.

If you determine that the physical interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *ATM Switch Router Software Configuration Guide*.

If the interface is still not operating correctly, proceed with the troubleshooting process in Chapter 6, “Troubleshooting Switch Router ATM Network Connections.”

Troubleshooting T1 and E1 Interfaces

This section describes specific processes and commands used to troubleshoot the T1 and E1 port adapters.

Port Adapter LEDs

The port adapter faceplate LEDs provide status information for individual T1 and E1 coaxial and UTP interface connections of the port adapter. The LEDs are described in Table 5-3.

Note Use the **show controllers** command to display the LED status.

Table 5-3 T1 and E1 Port Adapter LED Descriptions

LED	Status	Description
RX (Receive)	Off Flashing green Red	LOS[1] or port adapter is shut down. Cells are being received. LED blinks every 5 seconds and pulse rate increases with data rate. Alarm (LOF[2], LCD[3], AIS[4]).
TX (Transmit)	Off Flashing green Flashing yellow Steady yellow	No transmit line activity indication. Cells are being transmitted. LED pulse rate increases with data rate. Loopback. FERF[5] alarm.

[1] LOS = loss of signal

[2] LOF = loss of frame

[3] LCD = loss of cell delineation

[4] AIS = alarm indication signal

[5] FERF = far-end receive failure

Displaying Interface Port Configuration

To display the T1 and E1 interface configuration, use the following commands:

Command	Purpose
show interfaces atm *card/subcard/port*	Shows the status of the physical interface.
show atm interface atm *card/subcard/port*	Shows the interface configuration.
show controllers atm *card/subcard/port*	Shows the interface memory management and error counters.

Follow these steps to troubleshoot the T1 or E1 physical interface:

Step 1 Use the **show interfaces atm** *card/subcard/port* command to check the configuration.

```
  Switch# show interfaces atm 0/1/0
➔ ATM0/1/0 is down, line protocol is down
    Hardware is t1suni
    MTU 4470 bytes, sub MTU 0, BW 1500 Kbit, DLY 0 usec, rely 0/255, load 1/255
➔   Encapsulation ATM, loopback not set, keepalive not supported
➔   Last input never, output never, output hang never
    Last clearing of "show interface" counters never
    Queueing strategy: fifo
    Output queue 0/40, 0 drops; input queue 0/75, 0 drops
    5 minute input rate 0 bits/sec, 0 packets/sec
    5 minute output rate 0 bits/sec, 0 packets/sec
       0 packets input, 0 bytes, 0 no buffer
       Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
➔      0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
       0 packets output, 0 bytes, 0 underruns
➔      0 output errors, 0 collisions, 0 interface resets
       0 output buffer failures, 0 output buffers swapped out
  Switch#
```

Step 2 Check the ATM field to see whether the interface is up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the line protocol field to see that the status is up.

If the status is down, check for the following:

- The line protocol software processes might have determined that the line is unusable. Try swapping the cable.
- The local or remote interface might be misconfigured. Check the interface configuration.
- Clocking might be misconfigured or the source interface might have failed. Refer to the “Initially Configuring the ATM Switch Router” chapter in the *ATM Switch Router Software Configuration Guide*.
- Hardware might have failed. Try swapping the port adapter.

Step 4 Check the Encapsulation field. Confirm that the encapsulation method matches the interface type.

Step 5 Check the Last input and Last output fields. They show the number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface.

Step 6 Check the output hang field. It shows the number of hours, minutes, and seconds since the last reset caused by a lengthy transmission.

Step 7 Check the CRC field. The presence of many CRC errors, but not many collisions, indicates excessive noise. If the number is too high, check the cables for damage. If you are using UTP cables, make sure you are using category 5 cables and not another type, such as category 3.

Note Errors and the input and output difference should not exceed 0.5 to 2.0 percent of traffic on the interface.

Step 8 Check the collisions field. It shows the total number of collisions compared to the total number of output packets, and it should be approximately 0.1 percent or less. If the number is too high, perform the following tasks:

- Use a protocol analyzer to check for late collisions. Collisions do not occur in a properly designed network. Collisions usually occur when cables are too long.
- Check the diameter of the network and make sure it is within specification.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to troubleshoot the configuration of a T1 or E1 interface:

Step 1 Use the **show atm interface atm** *card/subcard/port* command to check the configuration.

```
  Switch# show atm interface atm 0/1/0
  Interface:      ATM0/1/0        Port-type:      t1suni
➔ IF Status:      DOWN            Admin Status:   down
  Auto-config:    enabled         AutoCfgState:   waiting for response from peer
  IF-Side:        Network         IF-type:        UNI
  Uni-type:       Private         Uni-version:    V3.0
  Max-VPI-bits:   8               Max-VCI-bits:   14
  Max-VP:         255             Max-VC:         16383
  ConfMaxSvpcVpi: 255             CurrMaxSvpcVpi: 255
  ConfMaxSvccVpi: 255             CurrMaxSvccVpi: 255
  ConfMinSvccVci: 33              CurrMinSvccVci: 33
  Svc Upc Intent: pass            Signalling:     Enabled
  ATM Address for Soft VC: 47.0091.8100.0000.0000.0000.0001.4000.0c80.1000.00
  Configured virtual links:
    PVCLs SoftVCLs   SVCLs   TVCLs   PVPLs SoftVPLs   SVPLs Total-Cfgd Inst-Conns
        2        0       0       0       0        0       0          2          0
  Logical ports(VP-tunnels):     0
➔ Input cells:    0               Output cells:   0
  5 minute input rate:             0 bits/sec,       0 cells/sec
  5 minute output rate:            0 bits/sec,       0 cells/sec
➔ Input AAL5 pkts: 0, Output AAL5 pkts: 0, AAL5 crc errors: 0
  Switch#
```

Step 2 Check the IF Status and Admin Status fields to see whether they are up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the Input cells and Output cells fields. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, the interface is experiencing congestion and dropping cells.

Step 4 Check the AAL5 crc error field. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, check for the following:

- Many CRC errors, but not many collisions. This indicates excessive noise.
- Cable damage. If you are using UTP cables, make sure you are using category 5 cables and not another type, such as category 3.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to display the memory management and error counters of a T1 or E1 ATM interface:

Step 1 Use the **show controllers atm** *card/subcard/port* command to check memory management and error counters.

```
Switch# show controllers atm 0/1/0
IF Name: ATM0/1/0, SUNI PDH Chip Base Address: A8908000
IF Name: ATM0/1/0, framer Base Address: A8909000
Port type: T1    Port rate: 1.5 Mbps    Port medium: UTP
➔ Port status:Good signal Loopback:None    Flags:8000
 showdow clk reg value AA
➔ TX Led: Traffic Pattern    RX Led: Traffic Pattern   CD Led: off
TX clock source:  network-derived
T1 Framing Mode:  ESF PLCP format
FERF on AIS is on
FERF on LCD is on (n/a in PLCP mode)
FERF on RED is on
FERF on OOF is on
FERF on LOS is on
LBO: between 0-110
Counters:
  Key: txcell   - # cells transmitted
       rxcell   - # cells received
       lcv      - # line code violations
       ferr     - # framing bit error event counter
       bee     - # bit error event, CRC-6 in ESF, Framing bit error in SF
       b1       - # PLCP BIP errors
       fe       - # PLCP framing pattern octet errors
       plcp_febe- # PLCP FEBE errors
       hcs      - # uncorrectable HEC errors
       uicell   - # unassigned/idle cells dropped
<Information Deleted>
Dump of internal registers for mask
 9 9 9 9 1 1 0 0
Switch#
```

Step 2 Check the Port status field. It should read "Good Signal."

Step 3 Check the Loopback field. It should read "None."

Step 4 Check the TX Led field to see that it reads "Traffic Pattern." If it does not, see Table 5-3 for LED descriptions.

Step 5 Check the RX Led field to see that it reads “Traffic Pattern.” If it does not, see Table 5-3 for LED descriptions.

Step 6 Check the CD field to see that it reads “Traffic Pattern.” If it does not, see Table 5-3 for LED descriptions.

If you determine that the physical interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *ATM Switch Router Software Configuration Guide*.

If the interface is still not operating correctly, proceed with the troubleshooting process in Chapter 6, “Troubleshooting Switch Router ATM Network Connections.”

Troubleshooting DS3 and E3 Interfaces

This section describes specific processes and commands used to troubleshoot the DS3 and E3 port adapters.

Port Adapter LEDs

The port adapter faceplate LEDs provide status information for individual DS3 and E3 coaxial interface connections of the port adapter. The LEDs are described in Table 5-4.

Note Use the **show controllers** command to display the LED status.

Table 5-4 DS3 and E3 Port Adapter LED Description

LED	Status	Description
RX (Receive)	Off Flashing green Red	LOS[1] or port adapter is shut down. Cells are being received. LED blinks every 5 seconds and pulse rate increases with data rate. Alarm (LOF[2], LCD[3], AIS[4]).
TX (Transmit)	Off Flashing green Flashing yellow Steady yellow	No transmit line activity indication. Cells are being transmitted. LED pulse rate increases with data rate. Loopback. FERF alarm.[5]

[1] LOS = loss of signal

[2] LOF = loss of frame

[3] LCD = loss of cell delineation

[4] AIS = alarm indication signal

[5] FERF = far-end receive failure

Displaying Interface Port Configuration

Use the following commands to display the DS3 or E3 interface configuration:

Command	Purpose
show interfaces atm *card/subcard/port*	Shows the status of the physical interface.
show atm interface atm *card/subcard/port*	Shows the interface configuration.
show controllers atm *card/subcard/port*	Shows the interface memory management and error counters.

Follow these steps to troubleshoot the DS3 or E3 physical interface:

Step 1 Use the **show interfaces atm** *card/subcard/port* command to check the configuration.

```
  Switch# show interfaces atm 0/1/0
→ ATM0/1/0 is down, line protocol is down
    Hardware is ds3suni_Quad
    MTU 4470 bytes, sub MTU 4470, BW 45000 Kbit, DLY 0 usec, rely 0/255, load 1/255
→   Encapsulation ATM, loopback not set, keepalive not supported
→   Last input never, output never, output hang never
    Last clearing of "show interface" counters never
    Queueing strategy: fifo
    Output queue 0/40, 0 drops; input queue 0/75, 0 drops
    5 minute input rate 0 bits/sec, 0 packets/sec
    5 minute output rate 0 bits/sec, 0 packets/sec
       0 packets input, 0 bytes, 0 no buffer
       Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
→      0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
       0 packets output, 0 bytes, 0 underruns
→      0 output errors, 0 collisions, 0 interface resets
  0 output buffer failures, 0 output buffers swapped out
  Switch#
```

Step 2 Check the ATM field to see whether the interface is up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Replace faulty hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the line protocol field. The status should be up.

If the interface is down, check for the following:

- The line protocol software processes might have determined that the line is unusable. Try swapping the cable.
- The local or remote interface might be misconfigured. Check the interface configuration.
- Clocking might be misconfigured or the source interface might have failed. Refer to the "Initially Configuring the ATM Switch Router" chapter in the *ATM Switch Router Software Configuration Guide*.
- Hardware might have failed. Try swapping the port adapter.

Step 4 Check the Encapsulation field. Confirm that the encapsulation method matches the interface type.

Step 5 Check the Last input or the Last output fields. They show the number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface.

Step 6 Check the output hang field. It shows the number of hours, minutes, and seconds since the last reset caused by a lengthy transmission.

Step 7 Check the CRC field. The presence of many CRC errors but not many collisions is an indication of excessive noise. If the number is too high, check the cables to determine if any are damaged. If you are using UTP cables, make sure you are using category 5 cables and not another type, such as category 3.

Note Errors and the input and output difference should not exceed 0.5 to 2.0 percent of traffic on the interface.

Step 8 Check the Collisions field. This value indicates the total number of collisions compared to the total number of output packets and should be approximately 0.1 percent or less. If the number is too high, perform the following tasks:

- Use a protocol analyzer to check for late collisions. Collisions do not occur in a properly designed network. Collisions usually occur when cables are too long.
- Check the diameter of the network and make sure it is within specification.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to show how to troubleshoot the configuration of a DS3 or E3 interface:

Step 1 Use the **show atm interface atm** *card/subcard/port* command to check the configuration.

```
  Switch# show atm interface atm 0/1/0
  Interface:      ATM0/1/0        Port-type:      ds3suni_Quad
➔ IF Status:      UP              Admin Status:   up
  Auto-config:    enabled         AutoCfgState:   waiting for response from peer
  IF-Side:        Network         IF-type:        UNI
  Uni-type:       Private         Uni-version:    V3.0
  Max-VPI-bits:   8               Max-VCI-bits:   14
  Max-VP:         255             Max-VC:         16383
  ConfMaxSvpcVpi: 255             CurrMaxSvpcVpi: 255
  ConfMaxSvccVpi: 255             CurrMaxSvccVpi: 255
  ConfMinSvccVci: 33              CurrMinSvccVci: 33
  Svc Upc Intent: pass            Signalling:     Enabled
  ATM Address for Soft VC: 47.0091.8100.0000.0040.0b0a.2a81.4000.0c80.1000.00
  Configured virtual links:
    PVCLs SoftVCLs   SVCLs   TVCLs   PVPLs SoftVPLs   SVPLs Total-Cfgd Inst-Conns
        2        0       0       0       0        0       0          2          0
  Logical ports(VP-tunnels):     0
➔ Input cells:    0               Output cells:   0
  5 minute input rate:             0 bits/sec,       0 cells/sec
  5 minute output rate:            0 bits/sec,       0 cells/sec
➔ Input AAL5 pkts: 0, Output AAL5 pkts: 0, AAL5 crc errors: 0
  Switch#
```

Step 2 Check the IF Status and Admin Status fields to see whether they are up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the Input cells and Output cells fields. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, the interface is experiencing congestion and dropping cells.

Step 4 Check the AAL5 crc error field. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, check for the following:

- Many CRC errors, but not many collisions. This indicates excessive noise.
- Cable damage. If you are using UTP cables, make sure you are using category 5 cables and not another type, such as category 3.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to display the memory management and error counters of a DS3 or E3 ATM interface:

Step 1 Use the **show controllers atm** *card/subcard/port* command to check memory management and error counters.

```
Switch# show controllers atm 0/1/0
IF Name: ATM0/1/0, Chip Base Address: A8908000
Port type: DS3    Port rate: 45 Mbps    Port medium: Coax
➔ Port status:Good Signal Loopback:None    Flags:8000
➔ TX Led: Traffic Pattern    RX Led: Traffic Pattern  TX clock source:  network-de
rived
DS3 Framing Mode:  cbit adm
FERF on AIS is on
FERF on LCD is on (n/a in PLCP mode)
FERF on RED is on
FERF on OOF is on
FERF on LOS is on
LBO: <= 225'
PDH counters:
  Key: txcell   - # cells transmitted
       rxcell   - # cells received
       lcv      - # line code violations
       ferr     - DS3: # F-bit/M-bit errors; E3: # framing errors
       exzs_ier - T3: # excessive zeros; E3 G.832: # iec errors
       perr     - DS3: # P-bit errors; E3 G.832: # BIP-8 errors
       cperr    - DS3: # path parity errors
       febe     - DS3 or E3 G.832: # FEBE errors
       b1       - # PLCP BIP errors
       fe       - # PLCP framing pattern octet errors
       plcp_febe- # PLCP FEBE errors
       hcs      - # uncorrectable HEC errors
       uicell   - # unassigned/idle cells dropped

<Information Deleted>

Netclock Reg1 Shadow:55, Netclock Reg2 Shadow:1,
Interrupt Status:DF, ASP ClkSel:C7FF
Switch#
```

Step 2 Check the Port status field. It should read “Good signal.”

Step 3 Check the Loopback field. It should read “None.”

Step 4 Check the TX Led field. It should read “Traffic Pattern.” If it does not, see Table 5-4 for LED descriptions.

Step 5 Check the RX Led field. It should read “Traffic Pattern.” If it does not, see Table 5-4 for LED descriptions.

If you determine that the physical interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *ATM Switch Router Software Configuration Guide*.

If the interface is still not operating correctly, proceed with the troubleshooting process in Chapter 6, “Troubleshooting Switch Router ATM Network Connections.”

Troubleshooting CES T1 and CES E1 Interfaces

This section describes specific processes and commands used to troubleshoot T1 and E1 circuit emulation service (CES) port adapters.

Port Adapter LEDs

The port adapter faceplate LEDs provide status information for individual CES T1 and CES E1 UTP and coaxial interface connections of the port adapter. The LEDs are described in Table 5-5.

Table 5-5 CES T1 and CES E1 Port Adapter LED Descriptions

LED	Status	Description
RX (Receive)	Off Flashing green Red	LOS[1] or port adapter is shut down. Cells are being received. LED blinks every five seconds and pulse rate increases with data rate. Alarm (LOF[2], LCD[3], AIS[4]).
TX (Transmit)	Off Flashing green Flashing yellow Steady yellow	No transmit line activity indication. Cells are being transmitted. LED pulse rate increases with data rate. Loopback. FERF alarm.[5]

[1] LOS = loss of signal

[2] LOF = loss of frame

[3] LCD = loss of cell delineation

[4] AIS = alarm indication signal

[5] FERF = far-end receive failure

Note Single-mode fiber-optic interface connectors are blue, and multimode connectors are black.

Displaying Interface Port Configuration

To display the CES T1 and CES E1 interface configuration, use the following commands:

Command	Purpose
show interfaces cbr *card/subcard/port*	Shows the status of the physical interface.
show ces interface atm *card/subcard/port*	Shows the interface configuration.

Follow these steps to troubleshoot the CES physical interface:

Step 1 Use the **show interfaces cbr** *card/subcard/port* command to check the configuration.

```
  Switch# show interfaces cbr 3/0/0
➔ CBR3/0/0 is up, line protocol is up
    Hardware is DCU
    MTU 53 bytes, BW 1544 Kbit, DLY 0 usec, rely 0/255, load 1/255
➔   Encapsulation ATMCES-T1, loopback not set
➔  Last input never, output never, output hang never
    Last clearing of "show interface" counters never
    Queueing strategy: fifo
    Output queue 0/40, 0 drops; input queue 0/75, 0 drops
    5 minute input rate 0 bits/sec, 0 packets/sec
    5 minute output rate 0 bits/sec, 0 packets/sec
       0 packets input, 0 bytes, 0 no buffer
       Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
➔      0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
       0 packets output, 0 bytes, 0 underruns
➔      0 output errors, 0 collisions, 0 interface resets
       0 output buffer failures, 0 output buffers swapped out
  Switch#
```

Step 2 Check the CBR field to see whether the interface is up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the line protocol field to see whether the status is up.

If the interface is down, check for the following:

- The line protocol software processes might have determined that the line is unusable. Try swapping the cable.
- The local or remote interface might be misconfigured. Check the interface configuration.
- Clocking might be misconfigured or the source interface might have failed. Refer to the “Initially Configuring the ATM Switch Router” chapter in the *ATM Switch Router Software Configuration Guide*.
- Hardware might have failed. Try swapping the port adapter.

Step 4 Check the Encapsulation field. Confirm that the encapsulation method matches the interface type.

Step 5 Check the Last input and Last output fields. They show the number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface.

Step 6 Check the output hang field. It shows the number of hours, minutes, and seconds since the last reset caused by a lengthy transmission.

Step 7 Check the CRC field. The presence of many CRC errors but not many collisions indicates excessive noise. If the number is too high, check the cables for damage. If you are using UTP cables, make sure you are using category 5 cables and not another type, such as category 3. Also check the clock mode, framing, and line coding configuration for each end of the connection.

Note Errors and the input and output difference should not exceed 0.5 to 2.0 percent of traffic on the interface.

Step 8 Check the collisions field. It shows the total number of collisions compared to the total number of output packets and should be approximately 0.1 percent or less. If the number is too high, perform the following tasks:

- Use a protocol analyzer to check for late collisions. Collisions do not occur in a properly designed network. Collisions usually occur when cables are too long.
- Check the diameter of the network and make sure it is within specification.

If you determine that the physical interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to troubleshoot the configuration of a CES interface:

Step 1 Use the **show ces interface cbr** *card/subcard/port* command to check the configuration.

```
  Switch# show ces interface cbr 3/0/0
  Interface:      CBR3/0/0        Port-type:T1-DCU
→ IF Status:      UP              Admin Status: UP
  Channels in use on this port:
  LineType: ESF          LineCoding: B8ZS  LoopConfig: NoLoop
  SignalMode: NoSignalling   XmtClockSrc: network-derived
→ DataFormat: UnStructured   AAL1 Clocking Mode: Synchronous  LineLength: 0_110
  LineState:  XmtAIS LossOfSignal
  Errors in the Current Interval:
    PCVs        0 LCVs        0   ESs         0   SESs        0   SEFSs        0
    UASs        0 CSSs        0   LESs        0   BESs        0   DMs          0
  Errors in the last 24Hrs:
    PCVs        0 LCVs        0   ESs         0   SESs        0   SEFSs        0
    UASs        0 CSSs        0   LESs        0   BESs        0   DMs          0
→ Input  Counters: 0 cells, 0 bytes
→ Output Counters: 0 cells, 0 bytes
  Switch#
```

Step 2 Check the IF Status and Admin Status fields to see whether they are up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 If the DataFormat field indicates that the circuit is unstructured, check the AAL1 Clocking Mode field to ensure that it matches the AAL1 clocking mode of the destination interface.

Step 4 Check the LineLength field to see if the value is correct. Measure the distance between the ATM switch router and the customer provided equipment (CPE) or regenerating device. The maximum supported distance for CES T1 interfaces is 650 feet, or 198 meters. The maximum supported distance for CES E1 interfaces and 820 feet, or 248.5 meters. The default value is 0 to 110 feet.

Note For detailed cabling and hardware information, refer to the "CES T1 and E1 Port Adapters" chapter in the *ATM Port Adapter and Interface Module Installation Guide.*

Step 5 Check the Input Counters and Output Counters fields. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, the interface is experiencing congestion and dropping cells.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

If the interface is still not operating correctly, continue with the troubleshooting process in Chapter 6, "Troubleshooting Switch Router ATM Network Connections."

Troubleshooting 25-Mbps Interfaces

This section describes specific processes and commands used to troubleshoot the 25-Mbps port adapter.

Port Adapter LEDs

The port adapter faceplate LEDs provide status information for individual 25-Mbps UTP interface connections of the port adapter. The LEDs are described in Table 5-6.

Note Use the **show controllers** command to display the LED status.

Table 5-6 25-Mbps UTP Port Adapter LED Descriptions

LED	Status	Description
TX (Transmit)	Off	No receive line activity indication.
	Flashing green	Cells are being received. LED blinks every 5 seconds and pulse rate increases with data rate.
	Flashing yellow	Loopback.
	Steady yellow	FERF alarm.[1]
	Red	Alarm indication (LOF[2], LCD[3]).

[1] FERF = far-end receive failure

[2] LOF = loss of frame

[3] LCD = loss of cell delineation

Displaying Interface Port Configuration

To display the 25-Mbps interface configuration, use the following commands:

Command	Purpose
show interfaces atm *card/subcard/port*	Shows the status of the physical interface.
show atm interface atm *card/subcard/port*	Shows the interface configuration.
show controllers atm *card/subcard/port*	Shows the interface memory management and error counters.

Follow these steps to troubleshoot the 25-Mbps physical interface:

Step 1 Use the **show interfaces atm** *card/subcard/port* command to check the configuration.

```
  Switch# show interfaces atm 1/0/0
→ ATM1/0/0 is UP, line protocol is UP
    Hardware is ATM25
    MTU 4470 bytes, sub MTU 4470, BW 25600 Kbit, DLY 0 usec, rely 0/255, load 1/255
→   Encapsulation ATM, loopback not set, keepalive not supported
→   Last input never, output never, output hang never
    Last clearing of "show interface" counters never
    Queueing strategy: fifo
    Output queue 0/40, 0 drops; input queue 0/75, 0 drops
    5 minute input rate 0 bits/sec, 0 packets/sec
    5 minute output rate 0 bits/sec, 0 packets/sec
       0 packets input, 0 bytes, 0 no buffer
       Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
→      0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
       0 packets output, 0 bytes, 0 underruns
→      0 output errors, 0 collisions, 0 interface resets
       0 output buffer failures, 0 output buffers swapped out
  Switch#
```

Step 2 Check the ATM field to see whether the interface is up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the line protocol field to see that the status is up.

If the interface is down, check for the following:

- The line protocol software processes might have determined that the line is unusable. Try swapping the cable.
- The local or remote interface might be misconfigured. Check the interface configuration.
- Clocking might be misconfigured or the source interface might have failed. Refer to the "Initially Configuring the ATM Switch Router" chapter in the *ATM Switch Router Software Configuration Guide*.
- Hardware might have failed. Try swapping the port adapter.

Step 4 Check the Encapsulation field. Confirm the encapsulation method matches the interface type.

Step 5 Check the Last input or the Last output fields. They show the number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface.

Step 6 Check the output hang field. It shows the number of hours, minutes, and seconds since the last reset caused by a lengthy transmission.

Step 7 Check the CRC field. The presence of many CRC errors, but not many collisions, indicates excessive noise. If the number is too high, check the cables for damage. If you are using UTP cables, make sure you are using category 5 cables and not another type, such as category 3.

Step 8 Errors and the input and output difference should not exceed 0.5 to 2.0 percent of traffic on the interface.

Step 9 Check the collisions field. It shows the total number of collisions compared to the total number of output packets and should be approximately 0.1 percent or less. If the number is too high perform the following tasks:

- Use a protocol analyzer to check for late collisions. Collisions do not occur in a properly designed network. Collisions usually occur when cables are too long.
- Check the diameter of the network and make sure it is within specification.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to show how to troubleshoot the configuration of a 25-Mbps interface:

Step 1 Use the **show atm interface atm** *card/subcard/port* command to check the configuration.

```
  Switch# show atm interface atm 1/0/0
  Interface:      ATM1/0/0        Port-type:    ATM25
→ IF Status:      UP              Admin Status: up
  Auto-config:    enabled         AutoCfgState: waiting for response from peer
  IF-Side:        Network         IF-type:      UNI
  Uni-type:       Private         Uni-version:  V3.0
  Max-VPI-bits:   2               Max-VCI-bits: 14
  Max-VP:         4               Max-VC:       16383
  Svc Upc Intent: pass            Signalling:   Enabled
  ATM Address for Soft VC: 47.0091.8100.0000.00e0.4fac.b401.4000.0c84.8000.00
  Configured virtual links:
    PVCLs SoftVCLs   SVCLs   PVPLs SoftVPLs   SVPLs  Total-Cfgd  Installed-Conns
        2        0       0       0        0       0           2                0
  Logical ports(VP-tunnels):     0
→ Input cells:    0               Output cells: 0
  5 minute input rate:             0 bits/sec,       0 cells/sec
  5 minute output rate:            0 bits/sec,       0 cells/sec
→ Input AAL5 pkts: 0, Output AAL5 pkts: 0, AAL5 crc errors: 0
  Switch#
  Switch#
```

Step 2 Check the IF Status and Admin Status fields to see whether they are up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the Input cells and Output cells fields. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, the interface is experiencing congestion and dropping cells.

Step 4 Check the AAL5 crc errors field. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, check for the following:

- Many CRC errors, but not many collision, indicates excessive noise.
- Cables damage. If you are using UTP cables, make sure you are using category 5 cables and not another type, such as category 3.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

If the interface is still not operating correctly proceed with the troubleshooting process in Chapter 6, "Troubleshooting Switch Router ATM Network Connections."

Troubleshooting CDS3 Frame Relay Interfaces

This section describes specific processes and commands used to troubleshoot the channelized DS3 Frame Relay port adapter (CDS3).

Port Adapter LEDs

The port adapter faceplate LEDs provide status information for individual channelized DS3 Frame Relay port adapter (CDS3) coaxial interface connections of the port adapter. The LEDs are described in Table 5-7.

Table 5-7 CDS3 Frame Relay Adapter LEDs

LED	State	Description
CD (carrier detect)	Off Green	Carrier detect signal not received Carrier detect signal received
RX (receive)	Off Flashing green Red	LOS[1] or shutdown Cells being received Alarm (LOF[2], LCD[3], AIS[4])
TX (transmit)	Off Flashing green Steady yellow	No transmit line activity Cells being transmitted Alarm FERF[5]

[1] LOS = loss of signal

[2] LOF = loss of frame

[3] LCD = loss of cell delineation

[4] AIS = alarm indication signal

[5] FERF = far-end receive failure

Displaying Interface Port Configuration

To display the channelized DS3 Frame Relay port adapter (CDS3) interface configuration, use the following commands:

Command	Purpose
show interfaces atm *card/subcard/port*	Shows the status of the physical interface.
show atm interface atm *card/subcard/port*	Shows the interface configuration.
show controllers e1 *card/subcard/port*	Shows the interface memory management and error counters.

Follow these steps to troubleshoot the DS3 or E3 Frame Relay physical interface:

Step 1 Use the **show interfaces atm** *card/subcard/port* command to check the configuration.

```
  Switch# show interfaces atm 4/0/0
➔ ATM-P4/0/0 is up, line protocol is up
    Hardware is ATM-PSEUDO
    MTU 4470 bytes, sub MTU 4470, BW 6000 Kbit, DLY 0 usec, rely 0/255, load 1/255
➔   Encapsulation ATM, loopback not set, keepalive not supported
    Encapsulation(s):
    2000 maximum active VCs, 0 current VCCs
    VC idle disconnect time: 300 seconds
➔   Last input never, output never, output hang never
    Last clearing of "show interface" counters never
    Queueing strategy: fifo
    Output queue 0/40, 0 drops; input queue 0/75, 0 drops
    5 minute input rate 0 bits/sec, 0 packets/sec
    5 minute output rate 0 bits/sec, 0 packets/sec
       0 packets input, 0 bytes, 0 no buffer
       Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
➔      0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
       0 packets output, 0 bytes, 0 underruns
➔      0 output errors, 0 collisions, 0 interface resets
       0 output buffer failures, 0 output buffers swapped out
  Switch#
```

Step 2 Check the ATM-P field to see whether the interface is up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Replace faulty hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the line protocol field to see if the status is up.

If the interface is down, check for the following:

- The line protocol software processes might have determined that the line is unusable. Try swapping the cable.
- The local or remote interface might be misconfigured. Check the interface configuration.
- Clocking might be misconfigured or the source interface might have failed. Refer to the "Initially Configuring the ATM Switch Router" chapter in the *ATM Switch Router Software Configuration Guide*.
- Hardware might have failed. Try swapping the port adapter.

Step 4 Check the Encapsulation field. Confirm that the encapsulation method matches the interface type.

Step 5 Check the Last input or the Last output fields. They show the number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface.

Step 6 Check the output hang field. It shows the number of hours, minutes, and seconds since the last reset caused by a lengthy transmission.

Step 7 Check the CRC field. The presence of many CRC errors but not many collisions is an indication of excessive noise. If the number is too high, check the cables to determine if any are damaged. If you are using UTP cables, make sure you are using category 5 cables and not another type, such as category 3.

Note Errors and the input and output difference should not exceed 0.5 to 2.0 percent of traffic on the interface.

Step 8 Check the Collisions field. This value indicates the total number of collisions compared to the total number of output packets and should be approximately 0.1 percent or less. If the number is too high, perform the following tasks:

- Use a protocol analyzer to check for late collisions. Collisions do not occur in a properly designed network. Collisions usually occur when cables are too long.
- Check the diameter of the network and make sure it is within specification.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to show how to troubleshoot the configuration of a DS3 or E3 Frame Relay interface:

Step 1 Use the **show atm interface atm** *card/subcard/port* command to check the configuration.

```
  Switch# show atm interface atm 4/0/0

  Interface:      ATM-P4/0/0      Port-type:      ATM-PSEUDO
➔ IF Status:      UP              Admin Status:   up
  Auto-config:    enabled         AutoCfgState:   waiting for response from peer
  IF-Side:        Network         IF-type:        UNI
  Uni-type:       Private         Uni-version:    V3.0
  Max-VPI-bits:   8               Max-VCI-bits:   14
  Max-VP:         255             Max-VC:         16383
  ConfMaxSvpcVpi: 255             CurrMaxSvpcVpi: 255
  ConfMaxSvccVpi: 255             CurrMaxSvccVpi: 255
  ConfMinSvccVci: 35              CurrMinSvccVci: 35
  Svc Upc Intent: pass            Signalling:     Enabled
  ATM Address for Soft VC: 47.0091.8100.0000.00e0.4fac.b401.4000.0c82.0000.00
  Configured virtual links:
    PVCLs SoftVCLs   SVCLs   TVCLs   PVPLs SoftVPLs   SVPLs Total-Cfgd Inst-Conns
        7        0       0       0       0        0       0          7          7
  Logical ports(VP-tunnels):     0
➔ Input cells:    0               Output cells:   0
  5 minute input rate:             0 bits/sec,       0 cells/sec
  5 minute output rate:            0 bits/sec,       0 cells/sec
➔ Input AAL5 pkts: 1, Output AAL5 pkts: 0, AAL5 crc errors: 0

  Switch#
```

Step 2 Check the IF Status and Admin Status fields to see whether they are up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the Input cells and Output cells fields. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, the interface is experiencing congestion and dropping cells.

Step 4 Check the AAL5 crc error field. If the errors and the input and output difference exceed 0.5 to 2.0 percent of traffic on the interface, check for the following:

- Many CRC errors, but not many collisions. This indicates excessive noise.
- Cable damage. If you are using UTP cables, make sure you are using category 5 cables and not another type, such as category 3.

If you determine that the physical interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to display the memory management and error counters of a DS3 or E3 ATM interface:

Step 1 Use the **show controllers atm** *card/subcard/port* command to check memory management and error counters.

```
Switch# show controllers e1 4/0/0
E1 4/0/0 is down.
  PAM state is Up
  FPGA Version:  fi-c8510-4e1fr.A.3.2
  Firmware Version: fi-c8510-4e1fr.A.2.3
  Transmitter is sending LOF Indication (RAI).
➔ Receiver has loss of signal.
  Framing is crc4, Line Code is HDB3, Clock Source is line.
  Data in current interval (143 seconds elapsed):
     0 Line Code Violations, 0 Path Code Violations
     0 Slip Secs, 0 Fr Loss Secs, 0 Line Err Secs, 0 Degraded Mins
     0 Errored Secs, 0 Bursty Err Secs, 0 Severely Err Secs
     143 Unavail Secs
  Data in Interval 1:
     0 Line Code Violations, 0 Path Code Violations
     0 Slip Secs, 0 Fr Loss Secs, 0 Line Err Secs, 0 Degraded Mins
     0 Errored Secs, 0 Bursty Err Secs, 0 Severely Err Secs
     900 Unavail Secs
.
(Information Deleted)
.
     Total Data (last 95 15 minute intervals):
     2 Line Code Violations,0 Path Code Violations,
     0 Slip Secs, 0 Fr Loss Secs, 2 Line Err Secs, 0 Degraded Mins,
     0 Errored Secs, 0 Bursty Err Secs, 0 Severely Err Secs
     1833290 Unavail Secs, 0 Stuffed Secs
Switch#
```

Step 2 Check the Receiver field. If the receiver has a loss of signal check the following:

- Confirm the cable between the interface port and the Service Provider equipment or terminal equipment is connected correctly. Confirm the cable is hooked up to the correct ports. Correct the cable connections if necessary.
- Check the cable integrity by looking for breaks or other physical abnormalities in the cable. Ensure the pinouts are set correctly. Replace the cable if necessary.

If you determine that the physical interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *ATM Switch Router Software Configuration Guide*.

If the interface is still not operating correctly, proceed with the troubleshooting process in Chapter 6, “Troubleshooting Switch Router ATM Network Connections.”

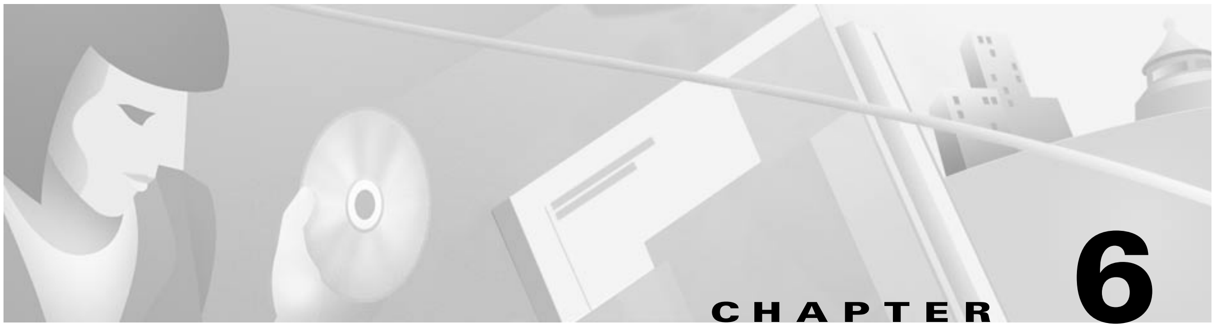

CHAPTER 6

Troubleshooting Switch Router ATM Network Connections

This chapter describes troubleshooting information about connectivity and performance problems in ATM switching network connections, and contains the following sections:

Checking Network Connections

Before you begin, make sure that all physical port connections are working correctly. See Chapter 5, “Troubleshooting Switch Router ATM Interface Connections.” Confirm the following:

- Proper cable insertion. Be sure that transmit and receive cable pairs match.
- Proper cable types. Connector fit does not ensure that the cables are the proper types or are cross-connected correctly.
- Reliable cables.
- No-shutdown mode on all interfaces on both ends of the cable.
- Proper interface configuration (for example, framing and line coding).
- Proper interface types on both ends of the cable.

Troubleshooting PVP and PVC Connections

This section describes how to troubleshoot permanent virtual paths (PVPs) and permanent virtual channels (PVCs). PVP and PVC connections are used primarily between buildings as the backbone connection and between frequently accessed hosts, such as the Domain Name System (DNS) server.

In the example network in Figure 6-1, the primary PVC configured as the backbone connection between the switch router on Floor 1 in the administration building and the switch router on Floor 1 in the manufacturing building has the following virtual path identifier (VPI) and virtual channel identifier (VCI) numbers:

- AdminFl1Ls1, ATM interface 3/1/0, VPI 50, and VCI 100
- ManuFl1Ls1, ATM interface 0/1/0, VPI 75, and VCI 150

Figure 6-1 PVC VPI and VCI Test in the Example Network

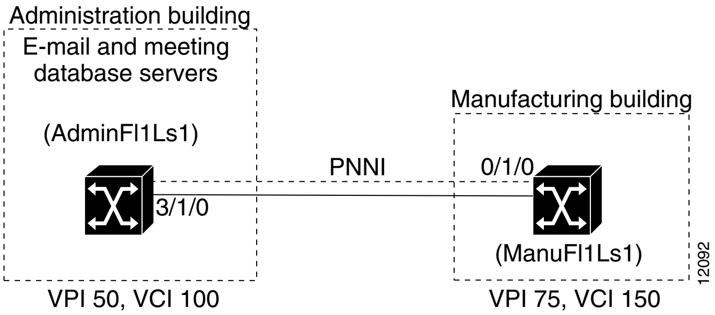

This section contains the following procedures:

- Checking the PVC Interface Status
- Checking the VPI and VCI Numbers
- Checking the VPI and VCI Ranges
- Checking the UBR Resources
- Checking the VBR and CBR Resources
- Debugging the PVC Connection Management

For detailed configuration information, refer to the "Configuring Virtual Connections" chapter in the *ATM Switch Router Software Configuration Guide*. For detailed information about configuring PVCs and traffic shaping on the Catalyst 5000 and 6000 ATM modules, refer to the *ATM Configuration Guide and Command Reference: Catalyst 5000 and 6000 ATM Modules.*

Checking the PVC Interface Status

Use the following command to confirm that the configured PVC interface status is up:

Command	Purpose
show atm status	Confirms the interface status.

Follow these steps to check the interface status:

Step 1 Use the **show atm status** command to check the status of the interface PVP.

```
Switch# show atm status

NUMBER OF INSTALLED CONNECTIONS: (P2P=Point to Point, P2MP=Point to MultiPoint)
Type     PVCs SoftPVCs     SVCs     TVCs     PVPs SoftPVPs     SVPs      Total
P2P        26        0        1        0        0        0        0         27
P2MP        0        0        0        0        0        0        0          0
                                      TOTAL INSTALLED CONNECTIONS =         27
PER-INTERFACE STATUS SUMMARY AT 16:02:57 UTC Mon May 11 1998:
   Interface      IF         Admin  Auto-Cfg    ILMI Addr     SSCOP    Hello
     Name       Status      Status    Status    Reg State     State    State
------------- -------- ------------ -------- ------------ --------- --------
ATM10/0/0           UP           up     done  UpAndNormal    Active  LoopErr
ATM10/0/1           UP           up     done  UpAndNormal    Active      n/a
ATM10/0/2         DOWN         down  waiting          n/a      Idle      n/a
ATM10/0/3           UP           up     done  UpAndNormal    Active  LoopErr
➔ ATM10/0/3.80        UP           up     done  UpAndNormal    Active  LoopErr
ATM10/1/0         DOWN         down  waiting          n/a      Idle      n/a
ATM10/1/1           UP           up     done  UpAndNormal    Active      n/a
ATM10/1/2           UP           up     done  UpAndNormal    Active  LoopErr
ATM10/1/3           UP           up     done  UpAndNormal    Active  LoopErr
➔ ATM10/1/3.80        UP           up     done  UpAndNormal    Active  LoopErr
ATM13/0/0           UP           up      n/a  UpAndNormal      Idle      n/a
Switch#
```

Step 2 Check the IF (Interface) Status field to confirm that the interface is up. If it is not, see Chapter 5, "Troubleshooting Switch Router ATM Interface Connections."

Step 3 Check the Admin (Administration) Status field to confirm that the interface is up. If it is not, see Chapter 5, "Troubleshooting Switch Router ATM Interface Connections."

Checking the VPI and VCI Numbers

Use the following command to confirm the configured PVC interface VPI and VCI numbers:

Command	Purpose
show atm vc interface atm *card/subcard/port vpi vci*	Confirms the interface status.

Follow these steps to check the VPI and VCI numbers configured for the PVC connection:

Step 1 Use the **show atm vc interface atm** command to confirm the numbers at both ends of the connection between the administration building and the manufacturing building:

```
  AdminFl1Ls1# show atm vc interface atm 3/1/0 50 100
  Interface: ATM3/1/0, Type: oc12suni
➔ VPI = 50 VCI = 100
  Status: UP
  Time-since-last-status-change: 5w1d
  Connection-type: PVC
  Cast-type: point-to-point
  Packet-discard-option: disabled
  Usage-Parameter-Control (UPC): pass
  Wrr weight: 32
  Number of OAM-configured connections: 0
  OAM-configuration: disabled
  OAM-states:  Not-applicable
➔ Cross-connect-interface: ATM0/1/0, Type: oc12suni
➔ Cross-connect-VPI = 75
➔ Cross-connect-VCI = 150
  Cross-connect-UPC: pass
  Cross-connect OAM-configuration: disabled
  Cross-connect OAM-state:  Not-applicable
  Threshold Group: 5, Cells queued: 0
  Rx cells: 0, Tx cells: 0
  Tx Clp0:0,  Tx Clp1: 0
  Rx Clp0:0,  Rx Clp1: 0
  Rx Upc Violations:0, Rx cell drops:0
  Rx Clp0 q full drops:0, Rx Clp1 qthresh drops:0
  Rx connection-traffic-table-index: 1
  Rx service-category: UBR (Unspecified Bit Rate)
  Rx pcr-clp01: 7113539
  Rx scr-clp01: none
  Rx mcr-clp01: none
  Rx      cdvt: 1024 (from default for interface)
  Rx       mbs: none
  Tx connection-traffic-table-index: 1
  Tx service-category: UBR (Unspecified Bit Rate)
  Tx pcr-clp01: 7113539
  Tx scr-clp01: none
  Tx mcr-clp01: none
  Tx      cdvt: none
  Tx       mbs: none
  AdminFl1Ls1#
```

Step 2 Check the VPI and VCI fields. They show the VPI and VCI of the PVC connection at the administration building.

Step 3 Check the Cross-connect-interface and Cross-connect-VPI and Cross-connect-VCI fields. They indicate the VPI and VCI of the PVC connection at the manufacturing building.

Checking the VPI and VCI Ranges

Use the following commands to check the VPI and VCI ranges of the PVC connection:

Command	Purpose
show atm ilmi-status atm *card/subcard/port*	Confirms the range configuration of the PVC and its VPI and VCI numbers.

Follow these steps to check the VPI and VCI ranges of the PVC connection at the administration building:

Step 1 Use the **show atm ilmi-status atm** command to confirm the ranges of the connection at the administration building.

```
  AdminFl1Ls1# show atm ilmi-status atm 3/1/0
  Interface : ATM3/1/0 Interface Type : Private NNI
  ILMI VCC : (50, 100) ILMI Keepalive : Disabled
  ILMI State:       UpAndNormal
  Peer IP Addr:     172.20.41.93    Peer IF Name:     ATM0/1/1
  Peer MaxVPIbits:  8               Peer MaxVCIbits:  14
➔ Peer MaxVPCs:     255             Peer MaxVCCs:     16383
  Peer MaxSvccVpi:  255             Peer MinSvccVci:  33
  Peer MaxSvpcVpi:  255
  Configured Prefix(s) :
  47.0091.8100.0000.0040.0b0a.2a81
  AdminFl1Ls1#
```

Step 2 Check the Peer MaxVPCs and Peer MaxVCCs fields. They indicate the VPI and VCI ranges of the PVC connection at the manufacturing building.

Step 3 Use the **show atm ilmi-status atm** command to confirm VPI and VCI ranges of the PVC connection at the manufacturing building.

```
ManuFl1Ls1# show atm ilmi-status atm 0/1/0
Interface : ATM0/1/0 Interface Type : Private NNI
ILMI VCC : (75, 150) ILMI Keepalive : Disabled
ILMI State:       UpAndNormal
Peer IP Addr:     172.20.41.93    Peer IF Name:     ATM0/1/0
Peer MaxVPIbits:  8               Peer MaxVCIbits:  14
➔ Peer MaxVPCs:     255             Peer MaxVCCs:     16383
Peer MaxSvccVpi:  255             Peer MinSvccVci:  33
Peer MaxSvpcVpi:  255
Configured Prefix(s) :
47.0091.8100.0000.0040.0b0a.2a81
ManuFl1Ls1#
```

Step 4 Check the Peer MaxVPCs and Peer MaxVCCs fields. They indicate the VPI and VCI ranges of the PVC connection at the administration building.

Step 5 If either the VPI or VCI of the PVC are configured incorrectly, refer to Chapter 6, “Configuring Virtual Connections,” of the *ATM Switch Router Software Configuration Guide*.

Checking the UBR Resources

Use the following commands to confirm unspecified bit rate (UBR) for the PVP and PVC best-effort connection limit configuration:

Command	Purpose
show atm interface resource atm *card/subcard/port*	For UBR connections, confirms connection admission control (CAC) best-effort limit configuration.
show atm resource	For VBR and CBR connections, confirms that the resources requested are available.

Follow these steps to confirm UBR for the PVC and PVP best-effort connection limit configuration on the interface.

Step 1 Use the **show atm interface resource atm** *card/subcard/port* command to confirm the maximum number best-effort connection limit configuration number.

```
Switch# show atm interface resource atm 10/0/0
Resource Management configuration:
    Output queues:
        Max sizes(explicit cfg): none cbr, none vbr-rt, none vbr-nrt, none abr-ubr
        Max sizes(installed): 256 cbr, 256 vbr-rt, 4096 vbr-nrt, 12032 abr-ubr
        Efci threshold: 25% cbr, 25% vbr-rt, 25% vbr-nrt, 25% abr, 25% ubr
        Discard threshold: 87% cbr, 87% vbr-rt, 87% vbr-nrt, 87% abr, 87% ubr
        Abr-relative-rate threshold: 25% abr
    Pacing: disabled   0 Kbps rate configured, 0 Kbps rate installed
    Service Categories supported: cbr,vbr-rt,vbr-nrt,abr,ubr
    Link Distance: 0 kilometers
    Controlled Link sharing:
        Max aggregate guaranteed services: none RX,  none TX
        Max bandwidth: none cbr RX, none cbr TX, none vbr RX, none vbr TX,
                       none abr RX, none abr TX, none ubr RX, none ubr TX
        Min bandwidth: none cbr RX, none cbr TX, none vbr RX, none vbr TX,
                       none abr RX, none abr TX, none ubr RX, none ubr TX
➔   Best effort connection limit: 10 max connections
    Max traffic parameters by service (rate in Kbps, tolerance in cell-times):
        Peak-cell-rate RX: none cbr, none vbr, none abr, none ubr
        Peak-cell-rate TX: none cbr, none vbr, none abr, none ubr
        Sustained-cell-rate: none vbr RX, none vbr TX
        Minimum-cell-rate RX: none abr, none ubr
        Minimum-cell-rate TX: none abr, none ubr
        CDVT RX: none cbr, none vbr, none abr, none ubr
        CDVT TX: none cbr, none vbr, none abr, none ubr
        MBS: none vbr RX, none vbr TX
Resource Management state:
    Cell-counts: 0 cbr, 0 vbr-rt, 0 vbr-nrt, 0 abr-ubr
    Available bit rates (in Kbps):
        147743 cbr RX, 147743 cbr TX, 147743 vbr RX, 147743 vbr TX,
        0 abr RX, 0 abr TX, 0 ubr RX, 0 ubr TX
    Allocated bit rates:
        0 cbr RX, 0 cbr TX, 0 vbr RX, 0 vbr TX,
 0 abr RX, 0 abr TX, 0 ubr RX, 0 ubr TX
➔   Best effort connections: 1 pvcs,  0 svcs
Switch#
```

Step 2 Check the Best effort connection limit field max (maximum) connections number. If the number is too low, increase it using the **atm cac best-effort-limit** interface command.

Step 3 Check the Best effort connection field to determine the number of established connections. If no connections are available, the connection fails.

To modify the best-effort connection limit, refer to the "Configuring Resource Management" chapter in the *ATM Switch Router Software Configuration Guide*.

Checking the VBR and CBR Resources

Use the following commands to confirm the VBR and CBR resources of the configured PVP:

Command	Purpose
show atm interface resource atm *card*/*subcard*/*port*	For UBR connections, confirms CAC best-effort-limit configuration.
show atm resource	For VBR and CBR connections, confirms that the resources requested are available.

The Catalyst 5000 and 6000 ATM modules do not support the **show atm interface resource atm** command. To check the status of a virtual connection on a Catalyst 5000 or 6000 ATM module, use the **show atm vc** command. For detailed information about configuring traffic shaping on the Catalyst 5000 and 6000 ATM modules, refer to the *ATM Configuration Guide and Command Reference: Catalyst 5000 and 6000 ATM Modules*.

The following example shows the status for all configured VCs on a Catalyst 5000 or 6000 ATM module:

```
ATM-5500# show atm vc
                                     AAL /         Peak   Avg.  Burst
Interface     VCD   VPI   VCI Type  Encapsulation  Kbps   Kbps  Cells Status
0.1             1     0     1  PVC  AAL5-SNAP     155000 100000     0 ACTIVE
0               2     0    16  PVC  AAL5-ILMI          0      0     0 ACTIVE
ATM-5500#
```

The following example shows how to display the status for a specific VCD on a Catalyst 5000 or 6000 ATM module:

```
ATM-5500# show atm vc 1
ATM0.1: VCD: 1, VPI: 0, VCI: 1, etype:0x0, AAL5 - LLC/SNAP, Flags: 0x830
PeakRate: 155000, Average Rate: 100000, Burst Cells: 0, VCmode: 0x0
OAM frequency: 60 second(s), InARP DISABLED
InPkts: 20972, OutPkts: 6924, InBytes: 6778670, OutBytes: 6210607
InPRoc: 20972, OutPRoc: 0, Broadcasts: 0
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
OAM F5 cells sent: 6924, OAM cells received: 0
ATM-5500#
```

Debugging the PVC Connection Management

Use the following commands to debug the PVC connection management:

Command	Purpose
debug atm conn errors	Enables connection management error debugging.
debug atm conn events	Enables connection management event debugging.
no debug all	Disables all debugging.

Troubleshooting Soft PVC Connections

This section describes how to troubleshoot a soft PVC configuration. Soft PVCs are used primarily to connect hosts that do not support signalling and cannot use SVCs.

In the example network in Figure 6-2, the connection between the switch router on Floor 1 in the administration building and the e-mail server has the following VPI and VCI numbers and ATM address:

- AdminFl1Ls1, ATM interface 4/0/0, VPI 150, and VCI 250
- E-mail server, ATM interface VPI 100, VCI 200, with an ATM address 11.1111.1111.00.1111.1111.1111.1111.1111.1111.00

Figure 6-2 Soft PVC Test in the Example Network

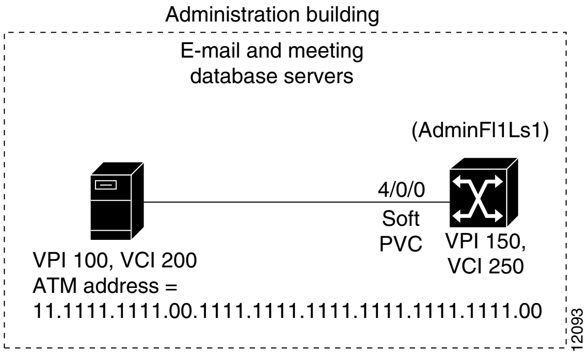

This section contains the following procedures:

- Checking the Interface Status
- Checking the VPI Number, VCI Number, and ATM Address
- Checking the Connection Management
- Debugging the Connection Management

For detailed information, refer to the “Configuring Virtual Connections” chapter in the *ATM Switch Router Software Configuration Guide*.

Checking the Interface Status

Use the following command to check soft PVC connection interface status:

Command	Purpose
show atm status	Confirms the interface status is up.

Follow these steps to confirm the soft PVC interface is up:

Step 1 Use the **show atm status** command to check the interface status.

```
Switch# show atm status
NUMBER OF INSTALLED CONNECTIONS: (P2P=Point to Point, P2MP=Point to MultiPoint)
Type     PVCs SoftPVCs     SVCs     TVCs     PVPs SoftPVPs     SVPs      Total
P2P        17        0        0        0        0        0        0         17
P2MP        0        0        0        0        0        0        0          0
                                      TOTAL INSTALLED CONNECTIONS =         17
PER-INTERFACE STATUS SUMMARY AT 13:41:00 UTC Tue May 12 1998:
   Interface      IF         Admin  Auto-Cfg    ILMI Addr     SSCOP    Hello
     Name       Status      Status    Status    Reg State     State    State
------------- -------- ------------ -------- ------------ --------- --------
ATM-P0/0/3          UP           up  waiting          n/a      none      n/a
ATM0/1/0          DOWN         down  waiting          n/a      Idle      n/a
ATM0/1/1          DOWN         down  waiting          n/a      Idle      n/a
ATM0/1/2          DOWN         down  waiting          n/a      Idle      n/a
ATM0/1/3          DOWN         down  waiting          n/a      Idle      n/a
ATM1/0/0            UP           up      n/a  UpAndNormal    Active      n/a
ATM1/0/0.80         UP           up  waiting  WaitDevType      Idle      n/a
ATM1/0/1          DOWN         down  waiting          n/a      Idle      n/a
ATM1/0/2          DOWN         down  waiting          n/a      Idle      n/a
ATM1/0/3            UP           up     done  UpAndNormal    Active      n/a
ATM1/1/0            UP           up      n/a  UpAndNormal    Active      n/a
ATM1/1/1          DOWN         down  waiting          n/a      Idle      n/a
ATM1/1/2          DOWN         down  waiting          n/a      Idle      n/a
ATM1/1/3            UP           up     done  UpAndNormal    Active      n/a
➔ ATM1/1/3.80       UP           up  waiting  WaitDevType      Idle      n/a
ATM2/0/0            UP           up      n/a  UpAndNormal      Idle      n/a
ATM4/1/0          DOWN         down  waiting          n/a      Idle      n/a
ATM4/1/1          DOWN         down  waiting          n/a      Idle      n/a
ATM4/1/2          DOWN         down  waiting          n/a      Idle      n/a
ATM4/1/3          DOWN         down  waiting          n/a      Idle      n/a
Switch#
```

Step 2 Confirm that the IF Status field corresponding to the soft PVC interface is up. If it is down, see Chapter 5, "Troubleshooting Switch Router ATM Interface Connections."

Step 3 Confirm that the Admin Status field is up. If it is down, see Chapter 5, "Troubleshooting Switch Router ATM Interface Connections."

Step 4 If both fields are up, continue with the following troubleshooting sections.

Checking the VPI Number, VCI Number, and ATM Address

Use the following command to confirm the VPI, VCI, and ATM address of the configured soft PVC:

Command	Purpose
show atm vc interface atm *card/subcard/port*	Confirms the configuration of VPI, VCI, and ATM address numbers of a soft PVC.

Follow these steps to confirm the VPI, VCI, and ATM address of the configured soft PVC:

Step 1 Use the **show atm vc interface atm** command to confirm the numbers at both ends of the connection between the administration building switch router and the e-mail server that does not support signalling:

```
  AdminFl1Ls1# show atm vc interface atm 4/0/0 150 250
  Interface: ATM4/0/0, Type: oc3suni
→ VPI = 150 VCI = 250
  Status: NOT CONNECTED
  Time-since-last-status-change: 00:00:45
  Connection-type: SoftVC
  Cast-type: point-to-point
  Soft vc location: Source
→ Remote ATM address: 11.1111.1111.00.1111.1111.1111.1111.1111.1111.00
→ Remote VPI: 100
→ Remote VCI: 200
→ Soft vc call state: Active
→ Number of soft vc re-try attempts: 4
  Slow-retry-interval: 60 seconds
  Next retry in: 29 seconds
  Aggregate admin weight: 0
  Packet-discard-option: disabled
  Usage-Parameter-Control (UPC): pass
  Wrr weight: 32
  Number of OAM-configured connections: 0
  OAM-configuration: disabled
  OAM-states:  Not-applicable
  Threshold Group: 5, Cells queued: 0
  Rx cells: 0, Tx cells: 0
  Tx Clp0:0,  Tx Clp1: 0
  Rx Clp0:0,  Rx Clp1: 0
  Rx Upc Violations:0, Rx cell drops:0
  Rx Clp0 q full drops:0, Rx Clp1 qthresh drops:0
  Rx connection-traffic-table-index: 1
  Rx service-category: UBR (Unspecified Bit Rate)
  Rx pcr-clp01: 7113539
  Rx scr-clp01: none
  Rx mcr-clp01: none
  Rx tolerance: 1024 (from default for interface)
  Tx connection-traffic-table-index: 1
  Tx service-category: UBR (Unspecified Bit Rate)
  Tx pcr-clp01: 7113539
  Tx scr-clp01: none
  Tx mcr-clp01: none
  Tx tolerance: none
  AdminFl1Ls1#
```

Step 2 Check the Remote ATM address. This address should match the ATM address at the other end of the soft PVC connection.

Step 3 Check the VPI and VCI fields. They indicate the VPI and VCI configuration of this interface.

Step 4 Check the Remote VPI and Remote VCI fields. They indicate the VPI and VCI configuration of the interface in the e-mail server.

If you determine that the VPI and VCI configurations are incorrect, refer to the “Configuring Virtual Connections,” of the *ATM Switch Router Software Configuration Guide*.

Step 5 Check the Soft vc call state field. This field should be Active.

Step 6 Check the Number of soft vc re-try attempts field. The number should be 0.

Checking the Connection Management

Use the following command to check soft PVC connection management:

Command	Purpose
show atm interface atm *card/subcard/port*	Confirms the interface status and configuration.

Debugging the Connection Management

Use the following commands to debug the PVC connection management:

Command	Purpose
debug atm conn errors	Enables connection management error debugging.
debug atm conn events	Enables connection management event debugging.
no debug all	Disables all debugging.

Troubleshooting SVC Connections on a PNNI Routing Network

This section describes how to troubleshoot switched virtual channel (SVC) connections, using the **show** command and **debug** command. These commands can be used to troubleshoot problems with SVC setup between end systems. The SVCs are automatically configured on the switch router when the cables are connected and the switch router is powered on.

In the example network in Figure 6-3, EndSys1 originates the signalling messages, which attempt to establish an SVC connection to EndSys2. In this example, Endsys1 connects directly to the switch router, named RemDvLs1, over the User-Network Interface (UNI) connection at ATM interface 3/1/1. Endsys2 is connected directly to the switch router, EngFl1Ls1, over the UNI connection at ATM interface 0/0/0. Both switch routers connect to other switch routers using network-to-network interface (NNI) connections.

Figure 6-3 SVC `Connection Example

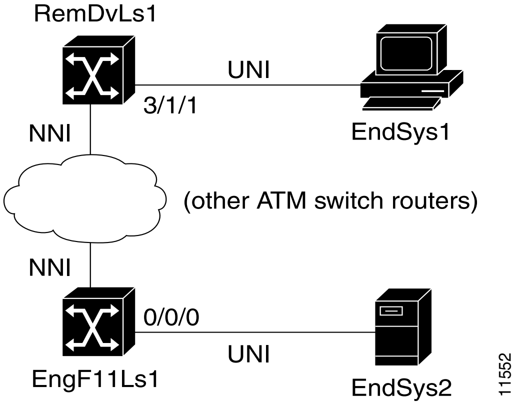

This section contains the following procedures:

- Checking the SVC Status at the End UNI Interface
- Checking UNI Interfaces
- Debugging SVC Signalling
- Alternate SVC Diagnostics
- Debugging PNNI SVC Routing
- Checking ATM Routes
- Checking PNNI Topology
- Checking SVC Downstream

Checking the SVC Status at the End UNI Interface

Use the following commands to check SVC interface status:

Command	Purpose
show atm vc signalling interface atm *card/subcard/port* **detail**	Confirms the SVC connection to the intended destination ATM NSAP address.
show atm vc interface atm *card/subcard/port vpi vci*	Confirms the destination UNI connection is up, and confirms the correct traffic characteristics are being used.

Follow these steps to confirm whether there is a new SVC connection from the originating side of the UNI interface to the intended remote or destination ATM network service access point (NSAP) address:

Step 1 Use the **show atm vc signalling interface atm** *card/subcard/port* **detail** command on the originating side.

```
RemDvLs1# show atm vc signalling interface atm 3/1/1 detail
➔ interface = ATM3/1/1, call remotely initiated, call reference = 19
vcnum = 0, vpi = 0, vci = 18, state = Active(EngFl1Ls1), point-to-point call

<Information Deleted>

timer currently inactive, timer value = 00:00:00
  Remote Atm Nsap address: 47.0091810000000060705BD900.123412344321.11
  local , Req Connect Ack -> Active(EngFl1Ls1),

<Information Deleted>

RemDvLs1#
```

Step 2 If the connection is up, confirm the correct traffic characteristics by using the VPI and VCI listed in the previous command display for the SVC to the target ATM NSAP address.

```
RemDvLs1# show atm vc interface atm 3/1/1 0 18
Interface: ATM3/1/1, Type: oc3suni
VPI = 0  VCI = 18
Status: UP

<Information Deleted>

Rx connection-traffic-table-index: 2147483647
Rx service-category: UBR (Unspecified Bit Rate)
Rx pcr-clp01: 7113539

<Information Deleted>

Tx connection-traffic-table-index: 2147483647
Tx service-category: UBR (Unspecified Bit Rate)
Tx pcr-clp01: 7113539

<Information Deleted>

RemDvLs1#
```

Step 3 If the connection is not UP, or not shown, continue with the following section, “Checking UNI Interfaces.”

Checking UNI Interfaces

Use the following commands to check the UNI configuration on the *originating* and *terminating* interfaces of the end systems:

Command	Purpose
show atm interface atm *card*/*subcard*/*port*	Confirms the interface status, UNI type, and UNI version.
show atm interface atm *card*/*subcard*/*port* **status**	Confirms the interface ILMI[1] and active signalling SSCOP[2] status.
show running-config	Confirms the interface configuration is valid.
show atm ilmi-status atm *card*/*subcard*/*port*	Confirms the end systems ATM addresses are registered for the UNI interface.

[1] ILMI = Interim Local Management Interface

[2] SSCOP = Service Specific Connection Oriented Protocol

Follow these steps to confirm that the *originating* end of the SVC connection (RemDvLs1 ATM 3/1/1 in this example) has the correct interface status, type, and UNI version compatible with the end system:

Step 1 Use the **show atm interface** command on the RemDvLs1 ATM 3/1/1 in this example.

```
  RemDvLs1# show atm interface atm 3/1/1
  Interface:      ATM3/1/1        Port-type:      oc3suni
➔ IF Status:      UP              Admin Status:   up
  Auto-config:    enabled         AutoCfgState:   completed
➔ IF-Side:        Network         IF-type:        UNI
➔ Uni-type:       Private         Uni-version:    V3.1
  Max-VPI-bits:   2               Max-VCI-bits:   10
  Max-VP:         255             Max-VC:         16383
  ConfMaxSvpcVpi: 255             CurrMaxSvpcVpi: 3
  ConfMaxSvccVpi: 255             CurrMaxSvccVpi: 3
  ConfMinSvccVci: 33              CurrMinSvccVci: 33
  Svc Upc Intent: pass            signalling:     Enabled
  ATM Address for Soft VC: 47.0091.8100.0000.00e0.4fac.b401.4000.0c85.0000.00
  Configured virtual links:
    PVCLs SoftVCLs   SVCLs   TVCLs   PVPLs SoftVPLs   SVPLs Total-Cfgd Inst-Conns   2
  0       0       0        0        0       0          2          2
  Logical ports(VP-tunnels):     0
  Input cells:    113971          Output cells:   98053
  5 minute input rate:          2000 bits/sec,       4 cells/sec
  5 minute output rate:         2000 bits/sec,       4 cells/sec
  Input AAL5 pkts: 64732, Output AAL5 pkts: 80752, AAL5 crc errors: 0
  EngFl1Ls1#
```

Step 2 Check to see whether the IF Status is UP. If it is not, see Chapter 5, "Troubleshooting Switch Router ATM Interface Connections."

Step 3 Check to see whether the IF-type is UNI. If it is not, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

Step 4 Check to see whether the UNI-version is compatible at both end systems. If it is not, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

Step 5 Next use the **show atm interface** command to confirm the EngFl1Ls1 ATM0/0/0 in this example:

```
  EngFl1Ls1# show atm interface atm 0/0/0
  Interface:      ATM0/0/0        Port-type:      oc3suni
→ IF Status:      UP              Admin Status:   up
  Auto-config:    enabled         AutoCfgState:   completed
→ IF-Side:        Network         IF-type:        UNI
→ Uni-type:       Private         Uni-version:    V3.1
  Max-VPI-bits:   2               Max-VCI-bits:   10
  Max-VP:         255             Max-VC:         16383
  ConfMaxSvpcVpi: 255             CurrMaxSvpcVpi: 3
  ConfMaxSvccVpi: 255             CurrMaxSvccVpi: 3
  ConfMinSvccVci: 33              CurrMinSvccVci: 33
  Svc Upc Intent: pass            signalling:     Enabled
  ATM Address for Soft VC: 47.0091.8100.0000.00e0.4fac.b401.4000.0c85.0000.00
  Configured virtual links:
    PVCLs SoftVCLs   SVCLs   TVCLs   PVPLs SoftVPLs   SVPLs Total-Cfgd Inst-Conns   2
  0       0       0       0        0       0          2          2
  Logical ports(VP-tunnels):     0
  Input cells:    113971          Output cells:   98053
  5 minute input rate:          2000 bits/sec,       4 cells/sec
  5 minute output rate:         2000 bits/sec,       4 cells/sec
  Input AAL5 pkts: 64732, Output pkts: 80752, AAL5 crc errors: 0
  EngFl1Ls1#
```

Step 6 Check to see whether the IF Status is UP. If it is not, see Chapter 5, “Troubleshooting Switch Router ATM Interface Connections.”

Step 7 Check to see whether the IF-type is UNI. If it is not, refer to the “Configuring Interfaces” chapter in the *ATM Switch Router Software Configuration Guide*.

Step 8 Check to see whether the Uni-version is compatible at both end systems. If it is not, refer to the “Configuring Interfaces” chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to confirm that the SVC connections have the correct ILMI and active signalling SSCOP status:

Step 1 Use the **show atm interface atm** command to confirm the *originating* end of the SVC connection:

```
RemDvLs1# show atm interface atm 3/1/1 status
   Interface      IF         Admin  Auto-Cfg    ILMI Addr      SSCOP    Hello
      Name      Status      Status    Status    Reg State      State    State
------------- -------- ------------ -------- ------------ --------- --------
ATM3/1/1            UP           up     done  UpAndNormal    Active  n/a
RemDvLs1#
```

Step 2 Use the **show atm interface atm** command to confirm the *terminating* end of the SVC connection:

```
EngFl1Ls1# show atm interface atm 0/0/0 status
   Interface      IF         Admin  Auto-Cfg    ILMI Addr      SSCOP    Hello
     Name       Status      Status    Status    Reg State      State    State
------------- -------- ------------ -------- ------------ --------- --------
ATM0/0/0            UP           up     done  UpAndNormal    Active   n/a
EngFl1Ls1#
```

Step 3 Check to see whether the IF Status is UP. If it is not, see Chapter 5, “Troubleshooting Switch Router ATM Interface Connections.”

Step 4 Confirm the ILMI Addr Reg State is Up And Normal.

Step 5 Confirm the SSCOP State is Active.

If either of these steps indicate a problem, use the **show running-config** command to check both the terminating and originating ends of the SVC connection for a valid interface configuration. Otherwise, continue with the following checks.

Follow these steps to check the addresses registered for the UNI interfaces:

Step 1 If the interfaces support ILMI, use the **show atm ilmi-status** command on the *originating* end of the SVC to verify that the expected end-system ATM addresses are registered for the UNI interfaces.

Step 2 If the interfaces support ILMI, use the **show atm ilmi-status** command on the *terminating* end of the SVC to verify that the expected end-system ATM addresses are registered for the UNI interfaces.

Step 3 Confirm the expected end-system ATM addresses are registered for the UNI interfaces.

For interfaces that do not support ILMI, use the **show running-config** command to verify that a static route has been configured with the correct end-system ATM address.

If static route has not been configured, refer to the "Initially Configuring the switch router" chapter in the *ATM Switch Router Software Configuration Guide*. Otherwise, continue with the next phase of SVC troubleshooting if you still have not determined the problem with the SVC configuration.

Debugging SVC Signalling

Use the following debug commands to check SVC signalling:

Command	Purpose
debug atm sig-all atm *card/subcard/port*	Confirms the SVC connection to the intended destination ATM NSAP address.
no debug all	Turns off debugging.

Follow these steps to turn on signalling debugging and then retry the setup of the SVC from EndSys1.

Step 1 Use the **debug atm sig-all atm** *card/subcard/port* command to enable signalling debugging for the originating end switch router UNI interface (on RemDvLs1 ATM 3/1/1).

Step 2 Retry to set up the SVC from EndSys1.

If no debug printouts occur on the switch router (RemDvLs1 in this example), then the problem is upstream on either the originating UNI interface, on the originating switch router itself, or in EndSys1.

Note Confirm that terminal monitor has been enabled on the switch router by entering the **terminal monitor** EXEC command.

Step 3 If debug printouts do occur, turn off further printouts by using the **no debug all** command.

Step 4 Scroll up to the beginning of the debug printouts to confirm the following:

- Check for a valid Called Party Address and Calling Party Address. If these are not valid or are not displayed, recheck the EndSys1 configuration.
- Check for the message `ROUTING INTERFACE: err_code = PNNI_SUCCESS`. If you do not see this message, continue to the "Debugging PNNI SVC Routing" section on page 6-20.
- Check to see whether there is an `Input Event: Rcvd Release` printout indicating a received release and look at the cause = *reason* and *location*. This indicates that the problem is downstream of the originating UNI, so proceed to the "Debugging PNNI SVC Routing" section on page 6-20 and then proceed to the "Checking SVC Downstream" section on page 6-23.

Alternate SVC Diagnostics

This section describes an alternate method you can use to troubleshoot SVC signalling, using the **atm signalling diagnostics** command.

Use the following commands starting at the privileged EXEC prompt to check SVC signalling:

	Command	Purpose
Step 1	Switch# **configure terminal** Switch (config)#	Enters configuration mode from the terminal.
Step 2	Switch (config)# **atm signalling diagnostics enable** Switch (config-atmsig-diag)#	Enables ATM signalling diagnostics.
Step 3	Switch (config-atmsig-diag)# **atm signalling diagnostics** *filter-index-number*	Starts ATM signalling diagnostics, using an index number from 1 to 50, and changes the prompt to ATM signalling diagnostics configuration mode.
Step 4	Switch (config-atmsig-diag)# **incoming-port atm** *card/subcard/port*	Configures the incoming port to filter.
Step 5	Switch (config-atmsig-diag)# **called-nsap-addres***s* *NSAP-address*	Sets the full called side NSAP address to filter.
Step 6	Switch (config-atmsig-diag)# **status active**	Activates the filter.
Step 7	Switch (config-atmsig-diag)# **end** Switch#	Exits signalling diagnostic configuration mode.
Step 8	Switch# **show atm signalling diagnostics filter** *filter-index-number*	Displays the configuration of the ATM signalling diagnostics filter.
Step 9	Switch# **show atm signalling diagnostic records** *filter-index-number*	Displays any captured records for this signalling diagnostics filter.

	Command	Purpose
Step 10	Switch# **configure terminal** Switch (config)#	At the privileged EXEC prompt, enters configuration mode from the terminal.
Step 11	Switch (config)# **no atm signalling diagnostics enable**	Disables ATM signalling diagnostics.

Follow these steps to check SVC signalling:

Step 1 Use the **atm signalling diagnostics enable** command to enable ATM signalling diagnostics.

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# atm signalling diagnostics enable
```

Step 2 Use the **atm signalling diagnostics** *filter-index-number* command to configure an ATM signalling diagnostics filter number.

```
Switch(config)# atm signalling diagnostics 1
```

Step 3 In ATM signalling diagnostics mode, use the **incoming-port atm** command to configure an ATM port for filtering.

```
Switch(cfg-atmsig-diag)# incoming-port atm 0/0/0
```

Step 4 Use the **called-nsap-address** command with the 20-octet called NSAP address to configure an ATM NSAP address for filtering.

```
Switch(cfg-atmsig-diag)# called-nsap-address
47.0091.8100.0000.00e0.4fac.b401.4000.0c80.8020.00
```

Step 5 Use the **status-active** command to start capturing records for this filter.

```
Switch(cfg-atmsig-diag)# status active
```

Step 6 Exit ATM signalling diagnostic mode, and use the **show atm signalling diagnostic filter** command to confirm that the filter is properly configured and active.

```
Switch(cfg-atmsig-diag)# end
Switch# show atm signalling diagnostics filter 1
F I L T E R   I N D E X    1
------------------------------
Scope: all, Cast Type: all
Connection Kind:   all
Service Category:  all
Clear Cause: 0, Initial TimerValue: 600
Max Records: 20,   NumMatches: 0,   Timer expiry: 557
Incoming Port: ATM0/0/0, Outgoing Port: 0
Calling Nsap Address:NULL
Calling Address Mask:NULL
Called Nsap Address :47.00918100000000E04FACB401.40000C808020.00
Called Address Mask :NULL
Status : active
Switch#
```

Step 7 Retry to set up the SVC from the end system.

Step 8 Use the **show atm signalling diagnostic record** command to examine the first filter record (labelled as: `DISPLAYINDEX 1`).

```
Switch# show atm signalling diagnostic records 1

<Display Omitted>

Switch#
```

Note No signalling diagnostic records are captured if the signalling setup is successful, or if the connection is immediately released by the End System.

If no captured records appear for an unsuccessful setup, the problem is at the originating UNI, or end system.

Step 9 Check the Calling-Address field. If the address is wrong, check the end system configuration.

If no list of DTLs are shown, see the following section, "Debugging PNNI SVC Routing."

If there is a Crankback type listed, see the "Checking SVC Downstream" section on page 6-23.

Step 10 In privileged EXEC mode, use the **no atm signalling diagnostic enable** command to disable ATM signalling diagnostics.

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# no atm signalling diagnostics enable
```

Debugging PNNI SVC Routing

Use the following commands to debug Private Network-Network Interface (PNNI) SVC routing:

Command	Purpose
debug atm pnni route-all atm	Confirms the SVC connection PNNI routing.
no debug all	Turns off debugging.

Follow these steps to enable PNNI routing debugging for the originating end switch router UNI interface.

Step 1 Use the **debug atm pnni route-all atm** command to enable PNNI routing debugging for RemDvLs1 ATM 3/1/1.

Step 2 Retry to set up the SVC from EndSys1.

Step 3 Turn off further debug printouts with the **no debug all** command.

```
  PNNI: Rcvd UBR Route Req to addr 47.0091810000000060705BD900.123412344321.11
  PNNI: Looking For Nodes That Advertise This Prefix
  PNNI: Best Match Is 47.0091810000000060705BD900.000000000000.00/104
→ PNNI: Found 1 POAs
        priority: 2  (12 0            ) pnni-remote-internal
  PNNI: Compute On-Demand Route Based On Admin Weight
  PNNI: Found A 1 Hop Route To Destination
→ PNNI: SOURCE ROUTE
        DTL 1> 2 Nodes
        DTL 1> 56:160:47.00918110000000613E7B2F01.00613E7B2F99.00 ATM0/1/1
        DTL 1> 56:160:47.009181100000006122222222.006122222222.00 ATM0/3/1
        DTL 2> 2 Nodes
        DTL 2> 24:40:47.009181100000000000000000.0060705BAD01.00 4276000
               24:160:47.0091810000000060705BD900.0060705BD900.00 0
→ PNNI: Found 1 Ports To Next DTL Node 12 ATM0/1/1
→ PNNI: Send Source Route Reply To Requestor: Code PNNI_SUCCESS
```

Step 4 Check printouts for correct service class, correct target address, and for at least 1 POA (Point of Attachment) at the target node. If no best match or POAs were found, proceed to the "Checking ATM Routes" section on page 6-21.

Step 5 Check to see whether at least one Ports to Next DTL Node *n* was found. If no ports were found, check for proper UNI/NNI interface configuration and status on the interfaces to the next indicated node *n*.

Note Use the **show atm pnni identifiers** command to determine the node that node *n* represents.

Step 6 If the initial Source Route Reply code is PNNI_SUCCESS and there are further tries with Crankback Set, the problem is downstream of this switch router. Note the original SOURCE ROUTE, shown as a list of DTLs (which are lists of node IDs and ports), as well as any calculated port list to the next node. Continue with the "Checking SVC Downstream" section on page 6-23.

If the Source Route Reply code is other than PNNI_SUCCESS, the actual code gives information about the nature of the problem when routing constraints are not met.

Checking ATM Routes

Use the following command to list the routes and destination prefixes:

Command	Purpose
show atm route	Displays the destination prefixes the originating switch router has learned.

Follow these steps to list the routes learned by the *originating* end of the switch router on the UNI interface:

Step 1 Use the **show atm route** command to display a list of routes learned by the originating end switch router UNI interface on RemDvLs1 ATM 3/1/1 as shown in Figure 5-3.

Step 2 Confirm that a prefix matching the intended target address is shown with a ST (State) UP. If there is more than one prefix that exactly matches the corresponding prefix of the target address, PNNI will choose the longest matching prefix.

If the longest matching prefix ST is DN (Down) for a node other than node 1, it indicates that there is no connectivity to that node. Continue to the following section "Checking PNNI Topology."

Note If the State is DN for a desired prefix on node 1 (this node), then check for proper status for the terminating UNI interface on this node. The ILMI Auto-Cfg (auto configuration) status must be shown as done, or auto configuration must be turned off for the prefix state to be UP.

Step 3 Confirm that the Node *n* shown for the longest matching prefix is the terminating switch router (EngFl1Ls1 for this example). If PNNI Hierarchy is being used, the node can instead be a logical group node (LGN) ancestor of the terminating switch router.

Note Use the **show atm pnni identifiers** command to determine which node *n* represents.

If the wrong node is listed with a matching prefix, check for proper ATM address configuration for the destination switch router (EngFl1Ls1 in this example), as well as for its UNI interface and for any hierarchy ancestor LGN.

Step 4 If there is no matching prefix appearing in the list of prefixes reachable from the originating end switch router (RemDvLs1 in this example), use the **show atm route** command on the terminating node (EngFl1Ls1 in this example).

If the prefix appears correctly on the terminating node, continue to the following section, "Checking PNNI Topology."

Checking PNNI Topology

The **show atm pnni topology** command and **show atm pnni election peers** command display the actual topology of connected switch router nodes that the originating node (RemFl1Ls1 in this example) has learned. Confirm that an unbroken path of nodes and links with the status up can be found between the originating and terminating switch routers (or for hierarchy, to a terminating end ancestor LGN).

Use the following commands to examine the node PNNI topology and switch router connectivity:

Command	Purpose
show atm pnni topology	Displays the actual topology of the connected nodes.
show atm pnni election peers	Displays the connectivity to a specific node within a peer group.

Follow these steps to display the actual topology of the connected nodes that the originating switch router has learned:

Step 1 Use the **show atm pnni topology** command to display the actual topology of the connected switch router nodes.

Step 2 If the terminating node is not shown or if necessary links are down or missing for an unbroken path, it indicates that the originating switch router (RemDvLs1 in this example) cannot find a path to the terminating node. Either a physical problem exists at the indicated network failure location, or else PNNI is unable to update its database to reflect the actual network condition.

Step 3 Use the **show atm pnni election peers** command to confirm whether this node has connectivity to any particular node within the same peer group.

Note Use the **show atm pnni identifiers** command to determine which nodes are represented by the node numbers that are internally assigned.

If a peer node is missing or is shown as NO for the Connected column, then PNNI considers that there is no path to that node.

Step 4 Check for physical problems by executing the **show atm pnni interface** command on the indicated failing nodes. If no physical problems are shown for the indicated failing nodes, proceed to the "Troubleshooting the PNNI Database" section on page 6-26.

If an unbroken path does exist based on the topology display, but debugging the PNNI routing showed that the destination was not initially PNNI_SUCCESS, it might mean that there are routing restrictions based on QoS, CAC, scope, or other path constraints that could not be met.

Checking SVC Downstream

This section is separated into two subsections:

- Flat Network
- Hierarchical Network

Proceed to the sub-section that best describes your PNNI network configuration.

Flat Network

Use the following commands to check ATM signalling events on the terminating switch router:

Command	Purpose
debug atm sig-events atm *card/subcard/port*	Confirms the SVC connection from the destination end of the SVC.
no debug all	Turns off debugging.

Note This process also applies to troubleshooting an SVC connection downstream in a terminating end peer group in a PNNI hierarchy.

Follow these steps to enable ATM signalling events debugging for the terminating end switch router UNI interface (on EngFl1Ls1 ATM 0/0/0):

Step 1 Use the **debug atm sig-events atm** *card/subcard/port* command to display signalling events at the terminating end of the switch router on the UNI interface.

Step 2 Alternately, you can set up a signalling diagnostic filter by using the appropriate called and calling end NSAP address, and examine the diagnostic record you receive.

Step 3 Retry to set up the SVC from EndSys1.

Step 4 If no debug printouts occur on the terminating switch router (EngFl1Ls1 in this example), then the signalling messages are not reaching the terminating node. Check for valid signalling status on the NNI links interconnecting the switch router nodes, using the **show atm status** command and **show atm interface** command.

Note Confirm that the terminal monitor has been enabled on the switch router by entering the **terminal monitor** EXEC command.

If debug printouts are shown on the terminating switch router (EngFl1Ls1 in this example), the problem has been isolated to either the terminating switch router, UNI, or the end system.

Step 5 Turn off further debug printouts with the **no debug all** command and scroll up to the beginning of the printouts to check the validity of party addresses and the occurrence of repeat events.

Step 6 Check for a valid Called Party Address and Calling Party Address (or a valid target address in the ROUTING INTERFACE information). If these are not valid, the printout might be for some other SVC setup.

If `ROUTING INTERFACE: err_code` (error codes) shows an err_code other than PNNI success, see the "Debugging PNNI SVC Routing" section on page 6-20 for the terminating switch router node (EngFl1Ls1 in this example).

Step 7 Confirm that there is an `Input Event: Rcvd Release` printout indicating a receive release and look at the cause = *reason* and *location*. This indicates that the problem is downstream on the terminating end system.

Hierarchical Network

Use the following commands to troubleshoot an SVC connection if the network supports PNNI hierarchy and the terminating node is in another peer group:

Command	Purpose
debug atm sig-events atm *card/subcard/port*	Determines the exit border node for the local peer group.
no debug all	Turns off debugging.
show atm pnni identifiers	Determines the internal node number and name corresponding to the exit border node ID.

Note To troubleshoot an SVC connection downstream at the terminating end peer group, see the previous section, "Flat Network."

Follow these steps to enable debugging ATM signalling events for the terminating end switch router on the UNI interface (on EngFl1Ls1 ATM 0/0/0):

Step 1 Start debugging signalling events with the **debug atm sig-events atm** *card/subcard/port* command to display signalling events on the terminating end switch router on the UNI interface.

Step 2 Retry to set up the SVC from EndSys1.

Step 3 If no debug printouts occur on the terminating switch router (EngFl1Ls1 in this example), then the signalling messages are not reaching the terminating node. Check for a valid signalling status on the NNI links interconnecting the nodes, using the **show atm status** command and **show atm interface** command.

Note Confirm that the terminal monitor is enabled on the switch router by entering the **terminal monitor** EXEC command.

If debug printouts are shown on the terminating switch router (EngFl1Ls1 in this example) the problem has been isolated to either the terminating switch router, UNI, or the end system.

Step 4 Turn off further debug printouts using the **no debug all** command.

```
EngFl1Ls1# debug atm sig-events atm 0/0/0

<Information Deleted>

→ PNNI: SOURCE ROUTE
        DTL 1> 2 Nodes
        DTL 1> 56:160:47.00918110000000613E7B2F01.00613E7B2F99.00 ATM0/1/1
        DTL 1> 56:160:47.009181100000006122222222.006122222222.00 ATM0/3/1
        DTL 2> 2 Nodes
        DTL 2> 24:40:47.009181100000000000000000.0060705BAD01.00 4276000
               24:160:47.009181000000060705BD900.0060705BD900.00 0

<Information Deleted>

EngFl1Ls1# no debug all
```

Step 5 Examine the initial SOURCE ROUTE. The last node ID listed for the lowest level DTL (shown as DTL 1>) is the exit border node for the local peer group. Make a note of the exit border node ID and port.

Follow these steps to determine the internal node number and name corresponding to the exit border node ID for the terminating end switch router on the UNI interface (EngFl1Ls1 ATM 0/0/0 in this example):

Step 1 Use the **show atm pnni identifiers** command to determine the internal node number and name corresponding to the exit border node ID.

The lowest level neighbor node on the other end of the exit border port is the entry border node for the next peer group.

Note The **show atm pnni topology node** *exit-border-node-number* command shows the neighbor node name of the entry border node if the interface is up.

Step 2 After determining the next entry border node, repeat the troubleshooting steps in the following sections on that node:

- Debugging SVC Signalling, page 6-17
- Debugging PNNI SVC Routing, page 6-20
- Checking ATM Routes, page 6-21
- Checking PNNI Topology, page 6-22
- Checking SVC Downstream, page 6-23

Step 3 Repeat these steps on that node and continue until either the terminating peer group is reached or the problem is isolated.

Troubleshooting the PNNI Database

This section outlines how to troubleshoot the PTSE (PNNI topology state element) database. When the PNNI topology or prefixes do not accurately reflect the state of other nodes in the network, you have problems with the PTSE database. All knowledge about other PNNI nodes is contained in the PTSE databases, which exist independently for each PNNI node in the network.

This section contains the following:

- Checking PNNI Neighbor Database Synchronization
- Checking the Flat Network or the Database Within the Same Peer Group
- Checking the PNNI Hierarchical Network Database for Different Peer Groups

Checking PNNI Neighbor Database Synchronization

When a node first initializes, it exchanges PTSEs with its immediate neighbor peer nodes. The progress of the database synchronization is tracked by the neighboring peer states.

Use the following commands to check the neighbor nodes and their corresponding states:

Command	Purpose
show atm pnni neighbor	Confirms the neighbor nodes and their corresponding PNNI states.
debug atm pnni adj-event	Confirms individual PNNI events being exchanged.
debug atm pnni adj-packet	Confirms individual PNNI packets being exchanged.
no debug all	Turns off all debugging.

Follow these steps to troubleshoot PNNI neighbor database synchronization problems:

Step 1 Use the **show atm pnni neighbor** command to confirm the neighbor nodes and their corresponding PNNI states.

Step 2 Check whether a neighboring peer node can reach the full state. If the neighboring peer node does not reach the full state, the following subset of neighboring peer states might indicate problems if they remain unchanged for an extended period:

- NPdown—There are no active links (for example, hello state 2way_in) to the neighboring peer. See the "Checking the PNNI Lowest Level Interface" section on page 6-32 to debug a known interface to a neighbor node unable to reach the 2way_in state.
- Negotiating—During this transient state, the neighbors agree upon which is the master (for example, the higher node ID) and the DS sequence number.
- Exchanging—During this state, the node describes its database by sending database summary packets containing PTSE headers only. When both adjacent nodes have the complete list of PTSE headers from the neighbor node, they transition to another state.
- Loading—During this state, the nodes are requesting PTSEs from the neighbor, but at least one has not been received.

If the neighbor machine remains in the Negotiating, Exchanging, or Loading state, turn on debugging by using the **debug atm pnni adj-event** command and **debug atm pnni adj-packet** command to see the individual events and the packets being exchanged.

Enter the **no debug all** command to turn off debug messages.

Checking the Flat Network or the Database Within the Same Peer Group

Use the following command to check the nodes in the peer group:

Command	Purpose
show atm pnni database [*internal-node-number*] [**detail**]	Confirms all nodes in the peer group with the PTSEs that each node originates.

Follow these steps to list all nodes in the peer group along with the PTSEs that each node originates:

Step 1 Use the **show atm pnni database** command to list all nodes in the peer group.

```
Switch# show atm pnni database 1
Node 1 ID 96:160:47.00918100000000E04FACB401.00E04FACB401.00 (name: Switch)
  PTSE ID  Length  Type  Seq no.   Checksum  Lifetime   Description
  1        92      97    117       37853     3143       Nodal info
  2        52      224   3331      18077     3016       Int. Reachable Address
Switch#
```

Step 2 Use the **show atm pnni database** command (again), with the **detail** command option.

```
Switch# show atm pnni database 1 detail
Node 1 ID 96:160:47.00918100000000E04FACB401.00E04FACB401.00 (name: Switch)
  PTSE ID  Length  Type  Seq no.   Checksum  Lifetime   Description
  1        92      97    117       37853     3135       Nodal info
    Time to refresh 1441, time to originate 0
    Type 97 (Nodal info), Length 48
    ATM address 47.00918100000000E04FACB401.00E04FACB401.00
    priority 0, leader bit NOT SET
    preferred PGL 0:0:00.000000000000000000000000.000000000000.00
  2        52      224   3331      18077     3008       Int. Reachable Address
    Time to refresh 1478, time to originate 0
    Type 224 (Int. Reachable Address), Length 32, Port 0, vp capable
    Scope (level) 0, Address info length (ail) 16, Address info count 1
    Pfx: 47.0091.8100.0000.00e0.4fac.b401..., length 104
Switch#
```

These commands should display similar information when the command is used on any other node in the same peer group.

The only differences are the internal node numbers (Node *n*), which are independently assigned by each node so that node 1 represents the node itself and other numbers are assigned as new nodes are discovered. The PTSE information might also differ for the valid case where some nodes have received more recent information than other nodes. A redisplay of the information on the node, which originally displayed older information for some PTSEs, normally shows more recent information, but might also have even newer information for other PTSEs.

In the output from the **show atm pnni database** command in Step 1 and Step 2, check the following:

Step 1 Whether all nodes in the peer group are shown. If no overlapping sets of partial nodes are shown for two different nodes in a peer group, it might indicate a peer group partition. Examine the interface status, using the **show atm pnni interface** card/subcard/port and **show atm pnni neighbor** commands for links and nodes that should connect to the nearest missing node to further isolate the problem.

Step 2 Whether the same PTSEs and similar sequence numbers appear on displays for different switch router nodes. If they do not, redisplay for the node with the older seq no (sequence number) to see if it gets updated. If there are differences, use the **debug atm pnni flood-packet** command on the originating and other nodes to see when PTSEs are being sent and received, along with any error conditions detected.

Step 3 Whether topology or other types of information for a node are incorrect, when displayed on another node. If they are not, use the **detail** option for the **show atm pnni database** command to display the complete PTSE contents both on the originating node and on any other node in the peer group. Determine whether the PTSE originates incorrectly or if there is a problem in synchronizing and flooding the PTSE to the other node.

Checking the PNNI Hierarchical Network Database for Different Peer Groups

A logical group node (LGN) originates PTSEs, which summarize the information from the entire child peer group it represents. The PTSEs that an LGN receives from its peer LGNs are flooded down to its child peer group leader (PGL), which then floods the PTSEs to its peers.

Use the following commands to check the PNNI hierarchy network database configuration:

Command	Purpose
show atm pnni database [*internal-node-number*] [**detail**]	Confirms that the PTSEs originated by all lowest level nodes in its peer group, its higher level ancestor LGNs, and all peers of the ancestor LGNs.
show atm pnni election local-node *node-index* **peers**	Confirms the PNNI PGL election process configuration.
show atm pnni database local-node [*internal-node-number*]	Confirms the contents of the PNNI topology database of the specified node.
debug atm pnni flood-packet local-node *node-index*	Debugs PNNI flood related packets for the local node.

Follow these steps to troubleshoot hierarchy database problems:

Step 1 Use the **show atm pnni database** *internal-node-number* **detail** command on the lowest level node to confirm that the PTSEs were originated by all lowest level nodes in the peer group, its higher level ancestor LGNs, and the PTSEs from all peers of the ancestor LGNs.

Note Use the **show atm pnni hierarchy network** command to determine the higher level ancestors for a node.

If there are problems with nodes or PTSEs within the same peer group, see the troubleshooting information in the "Checking the Flat Network or the Database Within the Same Peer Group" section on page 6-27 earlier in this chapter.

If there are problems with PTSEs from higher level LGNs, confirm the following for the output display:

Step 2 In addition to its peer nodes, check that the display shows all ancestor nodes. If some ancestor nodes are missing, see the next section, "Troubleshooting PNNI Peer Group Leaders."

Step 3 If all ancestor nodes are present, but other peer LGNs are missing at one of the higher levels, check which switch router is acting as the ancestor LGN for the affected level, using the **show atm pnni hierarchy network detail** command.

Step 4 Use the **show atm pnni database local-node** *node-index* command on the ancestor LGN switch router after determining the locally assigned node number for the affected LGN node. This command shows the subset of PTSEs that the higher level LGN has in its database.

Step 5 If the peer LGNs are missing from its database, use the **show atm pnni election local-node** *node-index* **peers** command to check connectivity to the missing LGNs.

Step 6 If there is no connectivity shown for some LGNs, see the "Troubleshooting PNNI Hierarchical Networks" section on page 6-44 to isolate problems with the child peer group leader for the missing uplink. Also, see the "Troubleshooting PNNI SVCC-RCC and Higher Level Links" section on page 6-39.

Step 7 If PTSEs originated by a higher level LGN show up incorrectly when displayed for a lowest level LGN, use the **show atm pnni database local-node** *node-index* command to display the higher level PTSEs for the ancestor LGN of the affected lowest level node and for the originating LGN node.

Step 8 If there are differences, use the **debug atm pnni flood-packet local-node** *node-index* command on the originating LGN and on any other affected LGN and child node.

This command shows when PTSEs are being sent and received, along with any error conditions detected.

Step 9 Check to see whether topology or other types of information for a higher level LGN are incorrect when displayed on a lowest level node in another peer group. Use the **detail** option for the **show atm pnni database local-node** *node-index* command.

This command shows the complete PTSE contents. Determine if the PTSE originates incorrectly or a problem exists transporting the PTSE to other LGNs or to the lowest level node.

Step 10 If the PTSE contents for the LGN originator do not accurately represent its child peer group information, see either the "Troubleshooting PNNI Hierarchical Networks" section on page 6-44 or the "Debugging Summary Addresses" section on page 6-51, depending on the type of affected PTSE.

Troubleshooting PNNI Peer Group Leaders

This section describes how to troubleshoot the PNNI peer group leader (PGL). In a PNNI network supported hierarchy, one node within the peer group is elected as the PGL. It summarizes and aggregates information from the entire peer group and passes that information to its parent LGN node, which advertises the information in PTSEs to its peer LGNs at the higher hierarchy level.

Use the following commands to check the PGL configuration:

Command	Purpose
show atm pnni hierarchy network [**detail**]	Confirms configured PNNI hierarchy and its status in detail.
show atm pnni election [**local-node** *node-index*]	Confirms PGL election process for the local node.
show atm pnni hierarchy local-configured	Confirms configured PNNI hierarchy for the local node.
show atm pnni election peers	Confirms PGL election priority and preferred PGL as advertised by all peers in the peer group.

Follow these steps to troubleshoot the PNNI PGL:

Step 1 Use the **show atm pnni hierarchy network detail** command to display the PGL and ancestor LGN for all higher hierarchy levels.

Step 2 If no active parent LGNs are shown, use the **show atm pnni election local-node** *node-index* command on the node (or nodes) that is configured to allow operation as the PGL. If the problem occurs for elections on a higher level, use the local-node option to specify the node index number of the higher level node.

```
Switch# show atm pnni election local-node 1
PGL Status.............: Not PGL
Preferred PGL..........: NULL
Preferred PGL Priority.: n/a
Active PGL.............: NULL
Active PGL Priority....: n/a
Active PGL For.........: n/a
Current FSM State......: PGLE Operating: Not PGL
Last FSM State.........: PGLE Calculating
Last FSM Event.........: Preferred PGL Is Not Self
Configured Priority....: 0
Advertised Priority....: 0
Conf. Parent Node Index: NONE
PGL Init Interval......: 15 secs
Search Peer Interval...: 75 secs
Re-election Interval...: 15 secs
Override Delay.........: 30 secs
Switch#
```

Step 3 Confirm that the election leadership-priority is configured to a nonzero value and that the expected primary PGL has the highest priority.

Step 4 Confirm that the PGL has a parent node configured that is enabled and running. Use the **show atm pnni hierarchy local-configured** command to view the locally configured parent nodes.

```
Switch# show atm pnni hierarchy local-configured
  Locally configured parent nodes:
  Node          Parent
  Index  Level  Index   Local-node Status     Node Name
  ~~~~~  ~~~~~  ~~~~~~  ~~~~~~~~~~~~~~~~~~~~  ~~~~~~~~~~~~~~~~~~~~~~
  1      96     N/A     Enabled/ Running      Switch
Switch#
```

Step 5 Use the **show atm pnni election peers** command to see which other peer nodes are known by a local node. Only those nodes listed as connected are eligible to be the preferred PGL for a local node.

```
Switch# show atm pnni election peers
  Node No.   Priority   Connected   Preferred PGL
  ~~~~~~~~   ~~~~~~~~   ~~~~~~~~~   ~~~~~~~~~~~~~
  1          0          Yes         NONE
Switch#
```

Step 6 If the expected leader still does not become PGL, check the current FSM state by using the **show atm pnni election** command (preferably on the switch router that acts as the PGL). The following subset of election states might indicate possible user correctable conditions if they remain unchanged for an extended period:

- PGLE Starting—Waiting for the first interface hello state machine to be started on a link. Be sure that at least one NNI is connected to another switch router (or LGN) and that the hello state machine on at least one interface is in a state other than down.

- PGLE Awaiting—Waiting for the first interface to reach the hello 2way_in Hierarchical state. It automatically transitions to the calculating state after waiting for the search peer interval as displayed by the **show atm pnni election** command.
- PGLE Awaiting Full, PGLE Initial Delay—Waiting for the first neighbor state machine to reach the full state and for an initial delay to allow peers to exchange election information. If the election gets stuck in the awaiting full state, proceed to the "Checking PNNI Neighbor Database Synchronization" section on page 6-26 to debug neighbor state machine problems.
- PGLE Awaiting Unanimity—This node prefers itself as PGL and is waiting for other nodes to reach unanimity. It automatically transitions to another state after waiting for the override delay as displayed by the **show atm pnni election** command.
- PGLE Hung Election: Not PGL—After waiting for the override delay, less than two-thirds of the other nodes are advertising it as their preferred PGL. This might result from a change in the topology or other network parameters. In that case, it should recover by itself. It can also indicate a defective node or link. Use the **show atm pnni election peers** command to check for the current connectivity to other nodes within the same peer group.
- PGLE Awaiting Reelection—The node has lost connectivity to the current PGL. It automatically transitions to another state after waiting for the re-election interval as displayed with the **show atm pnni election** command. Use the **show atm pnni election peers** command to check for connectivity to the original PGL and to the other nodes within the same peer group.
- PGLE Operating: Not PGL—The node has lost the election for PGL. To force this node to be the PGL, reconfigure the election priority to a value higher than the current PGL as listed with the **show atm pnni election** command, or else lower the election priority of the current PGL.

For other PGL election problems not isolated by these steps, use the **debug atm pnni election** command to turn on debugging messages that show the election events and state changes leading up to the election outcome as well as some additional election error conditions.

Turn off debugging messages with the **no debug all** command.

Troubleshooting the PNNI Lowest Level Interface

This section describes how to troubleshoot the lowest level PNNI interface connection problems.

This section contains the following procedures:

- Checking the PNNI Lowest Level Interface
- Checking the PNNI and Signalling Control Channels
- Checking PNNI PVC Status on Lowest Level Interfaces
- Checking PNNI Interface Metric Configuration for Lowest Level Interfaces
- Debugging PNNI Hello State at the Lowest Level

Checking the PNNI Lowest Level Interface

Use the following commands to check the lowest level PNNI interface status:

Command	Purpose
show atm interface atm *card/subcard/port* [**status**]	Confirms PNNI interface and administration status plus the hello state.
show atm status	Displays status information for all of the interfaces.
show atm routing-mode	Checks the switch router routing mode.
no atm routing-mode static	If needed, configures the switch router to allow PNNI operation.
show atm interface *card/subcard/port*	Confirms that the interface is configured with: • Auto configuration enabled (or as NNI) • IF-type is NNI • signalling: Enabled

Follow these steps to troubleshoot the lowest level PNNI interface status:

Step 1 Use the **show atm interface atm** *card/subcard/port* **status** command to confirm PNNI interface and administration status plus the hello state.

Note You can use the **show atm status** command to show status information for all of the interfaces.

If the IF status and admin status are not up, make sure that the interface is not configured as shutdown. If they still do not change to the UP state, see Chapter 5, "Troubleshooting Switch Router ATM Interface Connections."

If the PNNI hello state is n/a or not shown for an NNI interface between two switch routers, check the routing mode by using the **show atm routing-mode** command. If it is static mode, use the **no atm routing-mode static** command to allow PNNI operation. If this does not work, confirm that the installed software version allows PNNI operation.

Note For UNI interfaces, the PNNI Hello protocol is not used. The Hello state is not applicable for UNI interfaces.

If the hello state reads "LoopErr," it means that the line side is connected to another port on the same switch router, or to another switch router that has an identical node ID.

Step 2 Check the output of the **show atm interface atm** *card/subcard/port* command to confirm that the interface is configured with the following:

- Auto-config is enabled (or as NNI)
- IF-type is NNI
- signalling is enabled

If the interface is port-type: vp tunnel, confirm that the VP tunnel is configured correctly at both ends. see the "Troubleshooting Virtual Path Tunnel Connections" section on page 6-52.

Also check whether the listed port-adapter port-type supports ATM VCs on its line side. If it does not, then this interface will not be usable either as an NNI or as a UNI signalling interface.

Checking the PNNI and Signalling Control Channels

Use the following command to check the status of the PNNI Routing Control Channel (RCC) and signalling control channels:

Command	Purpose
show atm interface atm *card/subcard/port*	Confirms the PNNI signalling control channels status.

Follow these steps to check that the PNNI RCC and signalling control channel VCs are up:

Step 1 Use the **show atm vc interface atm** *card/subcard/port* command to confirm the status of the signalling control channels:

```
Switch# show atm vc interface ATM 0/0/1
Interface    VPI   VCI   Type    X-Interface  X-VPI X-VCI  Encap Status
ATM0/0/1     0     5      PVC     ATM2/0/0     0     43    QSAAL  UP
ATM0/0/1     0     16     PVC     ATM2/0/0     0     35    ILMI   UP
ATM0/0/1     0     18     PVC     ATM2/0/0     0     107   PNNI   UP

Switch#
```

Step 2 Check the command display for the following:

- Whether VCs with Encap (Encapsulation) types of PNNI and Q.2931 Signalling ATM Adaptation Layer (QSAAL) are not shown, check the interface configuration to confirm that signalling is enabled.
- If the interface has the manual-well-known-vc mode enabled, either disable it, or if that is the preferred mode, then manually configure PVCs with encapsulation types PNNI and QSAAL.

Note NNI interfaces require both QSAAL and PNNI PVCs, but UNI interfaces only require the QSAAL PVC along with the Interim Local Management Interface (ILMI) PVC.

- If VCs with PNNI and QSAAL are shown, but the status is not UP on an interface with an IF status that is UP, confirm that the interface has the manual-well-known-VC mode disabled and that the interface is type NNI.

Note If the neighbor node has multiple hierarchy levels and if one of its higher levels matches the level and peer group ID of the lowest level local node, then it is normal for a PNNI SVCC-RCC to be set up to communicate to the same level LGN, in addition to the PNNI PVC that communicates to the lowest level PNNI node of the neighbor.

Checking PNNI PVC Status on Lowest Level Interfaces

Use the following commands to check the PNNI PVC status:

Command	Purpose
show atm pnni interface atm *card/subcard/port* **detail**	Confirms the PNNI PVC status.
show atm pnni interface	Confirms the status of all PNNI interfaces.
show atm pnni local-node	Confirms that the lowest level peer group IDs match.
show atm pnni hierarchy network detail	Confirms that a common higher level peer group ID exists.

Follow these steps to troubleshoot the PNNI PVC status:

Step 1 Use the **show atm pnni interface atm** *card/subcard/port* **detail** command to confirm PNNI PVC status.

Note You can use the **show atm pnni interface** command to show PNNI information for all of the interfaces.

Step 2 Check for the following hello states. They can indicate possible user-correctable conditions if they remain unchanged for an extended period:

- DOWN—Lower level protocols have indicated that the link is not usable. See the previous sections for debugging low level interface problems.
- ATTEMPT—No hello messages have been (recently) received from the neighbor, even though the PNNI RCC PVC is up. Confirm that the remote end of the line (or VP Tunnel) is connected to the correct port on the intended remote switch router. Also check the status of the interface at the remote end of the line. For further analysis, see the "Debugging PNNI Hello State at the Lowest Level" section on page 6-38.
- 1-WAY INSIDE—Hellos have been recently received from a neighbor in the same peer group, but the neighbor has not yet acknowledged the information sent from this end. Confirm that the listed remote node and remote port ID are correct. See the "Debugging PNNI Hello State at the Lowest Level" section on page 6-38.
- 1-WAY OUTSIDE or 2-WAY OUTSIDE—Hellos have been recently received from a neighbor in another peer group, but no common higher level peer group has been found.

Step 3 If the neighbor was expected to be in the same peer group, confirm that the remote node has the expected peer group ID. Use the **show atm pnni local-node** command on this node and on the neighbor node to confirm that the lowest level peer group IDs match.

Note If the neighbor node has multiple hierarchy levels and if one of its higher levels matches the level and peer group ID of the lowest level local node, then it is normal for the Hello to the lowest level neighbor to reach the COMMON OUTSIDE state and for a PNNI SVCC-RCC to also be set up to communicate to the LGN that is at the same level as this node.

Step 4 If the neighbor was supposed to be in another peer group, but the COMMON OUTSIDE state has not been reached, use the **show atm pnni hierarchy network detail** command on this node and on the neighbor node to confirm that a common higher level peer group ID exists.

Step 5 It might take a minute or two for the higher level LGNs to come up for some hierarchy configurations that have multiple higher levels or do not have interfaces fully up yet at the higher levels. If a common higher level cannot be found after several minutes, see the "Debugging PNNI Hello State at the Lowest Level" section on page 6-38. Confirm that the peer group IDs appearing in the nodal hierarchy lists were sent in the individual hello messages on the outside link.

Step 6 If the peer group IDs do not have the expected values, use the **show atm pnni local-node** command on the switch routers where the higher level LGNs are running to confirm that peer group IDs have the expected values. If not, verify that the peer group IDs have not been configured to nondefault values.

Step 7 Also verify that if the active ATM address has been changed on one of the switch routers, that the lowest level node has been disabled and reenabled to reassign the node ID and peer group IDs based on the active ATM address (unless nondefault values are preferred).

Step 8 If common higher levels are not running, see the "Troubleshooting PNNI Peer Group Leaders" section on page 6-30.

Checking PNNI Interface Metric Configuration for Lowest Level Interfaces

Use the following commands to check the PNNI interface metric configuration:

Command	Purpose
show atm pnni interface atm *card/subcard/port*	Confirms PNNI interface metric configuration.
show running-config	Confirms administrative weight (AW) value, which shows the significant change boundaries.
show controllers atm *card/subcard/port*	Confirms minimum cell rate (MCR) value, port type, and port rate.
show atm pnni resource-info *card/subcard/port*	Confirms significant change boundaries.

Note Some resource metrics are valid only for a subset of the service classes.

Follow these steps to troubleshoot PNNI interface metric configuration and resource availability information for the lowest level interfaces:

Step 1 Use the **show atm pnni interface atm** *card/subcard/port* **detail** command to confirm the PNNI interface metric configuration.

```
Switch1# show atm pnni interface atm 0/0/0 detail
PNNI Interface(s) for local-node 1 (level=96):
Port ATM0/0/0 RCC is up  , Hello state 2way_in    with node Switch Error: Port
Looped back
  Next hello occurs in 0 seconds, Dead timer fires in 68 seconds
  CBR    : AW 5040 MCR 155519 ACR 147743 CTD 154 CDV 138 CLR0 10 CLR01 10
  VBR-RT : AW 5040 MCR 155519 ACR 155519 CTD 707 CDV 691 CLR0 8 CLR01 8
  VBR-NRT: AW 5040 MCR 155519 ACR 155519 CLR0 8 CLR01 8
  ABR    : AW 5040 MCR 155519 ACR 0
  UBR    : AW 5040 MCR 155519
  Aggregation Token: configured 0 , derived 0, remote 0
Switch#
```

Step 2 Check the administrative weight (AW) configuration. If the AW value is not what you expect, use the **show running-config** command to check the administrative-weight mode (for the ATM router PNNI configuration on the switch router).

Also, check whether the AW has been configured to a nondefault value for the specific interface.

Step 3 Check the minimum cell rate (MCR) configuration. If the MCR value is not what you expect, check the port type and port rate, using the **show controllers atm** *card/subcard/port* command (for physical interfaces only).

Step 4 Use the **show running-config** command to check the ATM pacing configuration. For VP tunnels, check the configuration of the corresponding PVP connection.

Step 5 Check the available cell rate (ACR), cell transfer delay (CTD), and cell delay variation (CDV) configuration. Use the **show atm pnni resource-info** *card/subcard/port* command to see the significant change boundaries.

Note Changes that are within the significant change boundaries do not trigger updates to the hello metrics or horizontal link PTSEs.

Step 6 Check the allocated bit rates (which affect ACR) by using the **show atm interface resource atm** *card/subcard/port* command.

Step 7 Check the CLR0 and CLR01 (CLR for CLP=0 and for CLP=0+1) configuration. Use the **show controllers atm** *card/subcard/port* command to see detailed error information for a specific interface.

Debugging PNNI Hello State at the Lowest Level

Use the **debug atm pnni hello-packet atm** *card/subcard/port* command at both the local end and (if possible) the remote end of the interface to see the actual hello messages being transmitted with some additional error condition messages.

Command	Purpose
debug atm pnni hello-packet atm *card/subcard/port*	Confirms the actual hello messages being transmitted.
no debug all	Turns off all debugging.

Follow these steps for further PNNI hello debugging at the lowest interface level:

Step 1 Use the **debug atm pnni hello-packet atm** command at the *local* end of the interface to see the actual hello messages being transmitted with some additional error condition messages.

```
Switch1# debug atm pnni hello-packet atm 0/0/1
<display omitted>
```

Step 2 Use the **debug atm pnni hello-packet atm** *card/subcard/port* command at the *neighbor* node of the interface (if possible) to see the actual hello messages being transmitted.

Step 3 After the display prints out two screens full of information, turn off further printouts by using the **no debug all** command.

Step 4 Scroll back up the screen display and confirm the following:

If no printouts are shown, be sure debugging is on. Confirm that this is an NNI interface and recheck the interface debugging steps in the "Checking the PNNI Lowest Level Interface" section on page 6-32, the "Checking the PNNI and Signalling Control Channels" section on page 6-34, and the "Checking PNNI PVC Status on Lowest Level Interfaces" section on page 6-35.

Step 5 Confirm that transmit messages are shown and have the expected local peer group ID and port ID. The transmit message contains the word "Tx."

Hello messages to peer group neighbors should look like this:

```
PNNI:56.1 Hello at ATM0/0/1: Tx, state 2way_in    with node Switch2
NodeId: 56:160:47.00918100000000613E7B2F01.00613E7B2F99.00 Address:
47.00918100000000613E7B2F01.00613E7B2F99.01 PgId: 56:47.0091.8100.0000.0000.0000.0000
Remote: port: ATM0/0/1 (80001000),
NodeId: 56:160:47.0091810000\0000400B0A3081.00400B0A3081.00

Local port: ATM0/0/1 (80001000)all
```

Hello messages on outside links to another peer group should have the same information as the previous example, but should include ULIA sequence number, hierarchy list, and aggregation token value.

Step 6 Confirm that receive messages are shown from the neighbor.

The receive message contains the word "Rx."

Hello messages received from peer group neighbors should look like the following:

```
PNNI:56.1 Hello at ATM0/0/1: Rx, state 2way_in    with node Switch1
NodeId: 56:160:47.00918100000000400B0A3081.00400B0A3081.00
Address: 47.00918100000000400B0A3081.00400B0A3081.01
PgId: 56:47.0091.8100.0000.0000.0000.0000
Remote: port: ATM0/0/1 (80001000), NodeId: 56:160:47.0091810000
Local port: ATM0/0/1 (80001000)
```

If no receive messages are shown on the local node, but the remote neighbor shows that it is transmitting them, there is a problem with transporting the message across the PNNI PVC.

When receive messages are shown, but do not match the transmit messages of the remote neighbor, it indicates that the line (or VP Tunnel) is connected to some remote port, but it is the wrong port.

Hello messages received on outside links from another peer group should have the same information as in the previous example, but in addition they should show a ULIA sequence number, a hierarchy list and sequence number, and an aggregation token value.

The hierarchy list can be examined to confirm whether a common peer group ID exists at some level.

Step 7 Look for other PNNI hello debugging error messages that might give further indication of internal or configuration problems.

Troubleshooting PNNI SVCC-RCC and Higher Level Links

This section describes how to troubleshoot PNNI routing control channel (RCC) between LGNs.

For a network that supports PNNI hierarchy, the PNNI RCC between LGNs (or between an LGN and a lowest level node), is a special type of SVC connection (referred to as an SVCC-RCC). After the SVCC-RCC is set up between the higher level LGN peers, PNNI hello messages are sent across it.

Each hello message contains information about all of the aggregated links between the local and remote LGN. Therefore, the following three types of states are kept independently, and all are important for higher level links:

- The SVCC-RCC setup state—Tracks the progress of requests to signalling to set up the SVCC-RCC.
- The RCC Hello state—An overall hello state for the RCC link, based on hello messages sent between the local and remote LGNs.
- Horizontal link states—Kept independently for each of the aggregation tokens that exist between a pair of LGNs. There is a unique (hexadecimal) port ID assigned for each of the aggregation tokens between a pair of LGNs, even though they are included in one common hello message.

For detailed configuration information, refer to the "Configuring ATM Routing and PNNI" chapter in the *ATM Switch Router Software Configuration Guide*.

This section contains the following procedures:

- Checking the PNNI Aggregated Horizontal Link Interface Status
- Checking SVCC-RCC Status

- Checking SVCC-RCC Hello State
- Debugging SVCC-RCC and Higher Level Link Problems

Checking the PNNI Aggregated Horizontal Link Interface Status

Use the following commands to check the status of all PNNI aggregated horizontal links and induced uplinks:

Command	Purpose
show atm pnni interface local-node *node-index*	Confirms the status of all PNNI aggregated horizontal links and induced uplinks.
show atm pnni neighbor	Verifies that the neighbor peer LGN has reached the full state for its database synchronization.

Follow these steps to troubleshoot a higher level LGN and the status of all PNNI aggregated horizontal links and induced uplinks at that level:

Step 1 Use the **show atm pnni interface local-node** *node-index* command to check all PNNI aggregated horizontal links and induced uplinks at the LGN level:

```
Switch1# show atm pnni interface local-node 2

   PNNI Interface(s) for local-node 2 (level=40):
    Local Port    Type  RCC HrzLn St Deriv Agg  Remote Port   Rem Node(No./Name)
    ~~~~~~~~~~~~~ ~~~~~ ~~~ ~~~~~~~~ ~~~~~~~~~~ ~~~~~~~~~~~~~ ~~~~~~~~~~~~~~~~~
    2C49000       HrzLn UP  2way     0          2230000       10 Switch2.2.40
    2C49003       HrzLn UP  2way     3          2230003       10 Switch2.2.40
    2276000       UpLnk -   n/a      0          FFFFFFFF      11 Switch4
   Switch1#
```

Note You can leave off the **local-node** option to show information for interfaces at all node levels present on the switch router.

If all of the expected interfaces between a pair of LGNs are missing, or if the RCC is not UP, proceed with the following SVCC-RCC checks.

If the RCC is listed as UP, but ports are missing for some expected aggregation tokens, proceed to the "Troubleshooting PNNI Hierarchical Networks" section on page 6-44.

For example, if the RCC is listed as UP, but the HrzLn State (Horizontal Link State) is other than 2way, the following subset of aggregated horizontal link states might indicate conditions that you can correct if they remain unchanged for an extended period:

- DOWN—No uplink PTSEs have been received that include an uplink to the neighboring peer LGN with the same aggregation value as listed in a hello message. See the "Troubleshooting PNNI Hierarchical Networks" section on page 6-44.
- ATTEMPT—Although at least one uplink PTSE has been received, including an uplink to the neighbor peer LGN with the listed aggregation token, no valid confirming information has recently been received piggybacked onto the peer hello message.

- 1 WAY—Hellos have recently been received confirming the aggregation token from a neighbor peer LGN, but the neighbor has not acknowledged the information that the local side transmitted. See the “Debugging SVCC-RCC and Higher Level Link Problems” section on page 6-43.

Step 2 Verify that the neighbor peer LGN has reached the full state for its database synchronization by using the **show atm pnni neighbor** command. If it has not, see the “Checking PNNI Neighbor Database Synchronization” section on page 6-26. If the neighbor has reached the full state but the horizontal link remains in the attempt state, see the “Checking SVCC-RCC Status” section on page 6-41.

Checking SVCC-RCC Status

This section describes troubleshooting the status of SVCC-RCCs from a local LGN node to all of its LGN peers.

Use the following commands to confirm the status of SVCC-RCCs from a local LGN node:

Command	Purpose
show atm pnni svcc-rcc *local-node* **detail**	Confirms the status of SVCC-RCCs from a local LGN node to all of its LGN peers.
show version	Confirms the route processor card has sufficient memory to support the software version.

Follow these steps to troubleshoot the status of SVCC-RCCs from a local LGN node to all of its LGN peers:

Step 1 Use the **show atm pnni svcc-rcc** *local-node* **detail** command to confirm the status of SVCC-RCCs from a local LGN node.

Step 2 Confirm that the RCC state is UP (with SVCC setup state shown as SVCC_UP for the detailed display). Make a note of whether this is the calling side or the called side as shown on the same display line where the SVCC setup state appears.

Step 3 If the RCC state is up, but the hello state is not 2way_in, proceed to the next section, “Checking SVCC-RCC Hello State.” Otherwise, if the RCC state is not up, continue with these checks.

Note The LGN with the higher node ID is the calling side originator of the signalling messages that set up the SVCC-RCC.

Step 4 Check whether the SVCC has the intended remote (LGN) node and rem-node (remote node) name. If it does not, verify the LGN ancestor information for the child PGL that was the intended remote LGN node.

Step 5 If for an extended period the SVCC setup state (listed for the **detail** option) is not SVCC_UP, the following subset of SVCC setup states might indicate correctable conditions:

- WAIT_INITIAL—An approximate four second delay after an uplink PTSE to a new upnode LGN with a common peer group ID has been received before the SVCC-RCC setup is attempted. The state should change automatically to another one of the setup states, and no user actions are required.

- WAIT_CONNECT, WAIT_DELAY_RETRY, or WAIT_IMMED_RETRY—An SVC setup request has been sent to signalling, but the SVC connection has either been released or the connect confirmation has not been received.

  Proceed to the "Troubleshooting SVC Connections on a PNNI Routing Network" section on page 6-12, keeping in mind that the originating interface is the route processor port for this node (ATM 0 for a standard switch router). However, if those steps show no signalling debugging messages, proceed to the "Debugging SVCC-RCC and Higher Level Link Problems" section on page 6-43.
- WAIT_SETUP—A called side node that has received an uplink PTSE to a new upnode LGN with a common PG ID, but it is still waiting for the SVC signalling setup message to arrive from the remote SVC originator LGN node. Further debugging should take place on the remote side switch router where the remote LGN resides. Repeat these debugging steps for the remote side switch router.
- WAIT_REL_SYS_BUSY—An internal software condition resulting from a shortage of route processor processor memory. If this condition persists or recurs, be sure that the route processor card has sufficient memory to support the software version listed by the **show version** command.

Checking SVCC-RCC Hello State

Use the following command to check the SVCC-RCC hello state:

Command	Purpose
show atm pnni svcc-rcc [**local-node** *node-index*]	Confirms the SVCC-RCC hello state.

Follow these steps to check the SVCC-RCC hello state:

Step 1 Use the **show atm pnni svcc-rcc local-node** *node-index* command without the **detail** option to determine the SVCC-RCC hello state.

Step 2 Use the **show atm pnni svcc-rcc local-node** *node-index* command to check the following:

If the SVCC-RCC state is up, but the hello state is other than 2way_in, the following subset of RCC Hello States might indicate possible user correctable conditions if they remain unchanged for an extended period:

- DOWN—The SVCC-RCC has not indicated that it is up. See the previous sections for debugging SVCC-RCC setup problems.
- ATTEMPT—No hello messages have been (recently) received from the LGN peer, even though the SVCC-RCC is up. Check for correct status at the remote LGN. For further analysis, see the "Debugging SVCC-RCC and Higher Level Link Problems" section on page 6-43.
- 1-WAY INSIDE—Hello messages have been recently received from an LGN peer, but the LGN has not yet acknowledged the information sent from this end. Confirm that the remote node listed is correct. See the following section, "Debugging SVCC-RCC and Higher Level Link Problems."

Debugging SVCC-RCC and Higher Level Link Problems

If the previous steps cannot isolate the cause of a problem with higher level link status, this section describes the **debug** command and **show** command that recognize the following:

- SVCC-RCC setup state transitions
- RCC hello state transitions
- Aggregated horizontal link state transitions
- Full hello message contents

Use the following commands to debug and check the SVCC-RCC setup:

Command	Purpose
debug atm pnni svcc-rcc remote-node *internal-node-number*	Confirms the SVCC-RCC setup, the RCC hello, and aggregated horizontal link state transitions, plus full hello message contents.
no debug all	Turns off all further bugging.
show atm pnni topology	Confirms if PNNI shows a route to the remote LGN.

Follow these steps to debug and check the SVCC-RCC setup:

Step 1 Use the **debug atm pnni svcc-rcc remote-node** *internal-node-number* command on the switch router that has the higher numbered LGN node ID, because that is the SVC originator (for example, the calling side).

Caution Because this debugging mode controls extensive information, it is best to filter the output by specifying either a local node or preferably the target remote node number, if it is known.

The **debug atm pnni svcc-rcc remote-node** *internal-node-number* command is normally more helpful when used on the switch router that has the higher numbered LGN node ID, because that is the SVC originator (for example, the calling side). However, it is useful to display the debugging hello messages at both ends for debugging RCC hello problems.

Step 2 Wait approximately one minute to allow any SVCC setup retries to be listed and turn off debugging, using the **no debug all** command.

Step 3 Scroll to the top of the screen and confirm the following:

If the SVCC-RCC has not yet reached the SVCC_UP state, confirm that a queued ATM_SETUP line is displayed and make note of the service category. Normally the service category is VBR-NRT, except for cases where the SVC must traverse a VP tunnel or some link that does not support VBR-NRT.

When the service category changes with each attempt, release messages are being received. This indicates that no path exists for each attempted service category. Check the topology, using the **show atm pnni topology** command to see whether PNNI shows a route to the remote LGN.

Note Of course, the horizontal link between these LGNs is not listed as up in the topology, but the special case of SVCC-RCC setup does not require an UP status.

Step 4 Check whether the debugging line following the queued ATM_SETUP phrase shows the ATM address of the intended remote LGN node. If it does not, these messages might belong to the SVCC-RCC for another remote LGN.

If there are no setup attempts for the case where an expected SVCC-RCC has not yet reached the SVCC_UP state, proceed to the, "Troubleshooting PNNI Hierarchical Networks" section.

If there are setup attempts, but release messages are received for each attempt, note the cause code which might explain the problem. Proceed to the "Troubleshooting SVC Connections on a PNNI Routing Network" section on page 6-12, keeping in mind that the origination interface is the route processor port.

Step 5 Note any other debugging error messages that might be printed in SVCC-RCC debugging mode.

For RCC hello FSM problems, the full hello messages are listed, along with the horizontal link extension entries for all aggregation tokens. By listing the hello messages at both ends of the SVCC-RCC, it is possible to locate where missing or mismatched information takes place.

Troubleshooting PNNI Hierarchical Networks

This section describes how to troubleshoot PNNI uplink and aggregation problems for hierarchical networks by the LGN.

Links that connect border nodes between two different peer groups are referred to as outside links. When the hello state finds a common higher level ancestor LGN on an outside link, it transitions to the common outside state. At this time each border node advertises an uplink PTSE to its peer nodes. The uplink PTSE contains the resource information for both directions of the outside link along with the node ID and peer group ID of the upnode and the aggregation token for the link.

The PGL uses the uplink PTSE information to aggregate the resource information from all outside links with the same aggregation token that connects to another peer group. The PGL notifies its parent LGN whenever there are changes to an uplink status. The parent LGN creates either an induced horizontal link or an induced uplink for each aggregation token to an upnode at the same or higher level.

For detailed configuration information, refer to the "Configuring ATM Routing and PNNI" in the *ATM Switch Router Software Configuration Guide*.

This section contains the following procedures:

- Checking Uplinks for Peer Group
- Checking Missing Upnode or Aggregation Token Pairs

Checking Uplinks for Peer Group

To see the table that summarizes all of the uplinks for a peer group, enter the **show atm pnni aggregation link** command on the switch router acting as the PGL. The display also shows the port identity of the induced horizontal or uplink for the parent LGN.

Use the following command to check the uplinks for the peer group on the PGL:

Command	Purpose
show atm pnni aggregation link	Confirms the PGL summary uplink to the peer group.

Follow these steps to check the PGL uplink summaries:

Step 1 Use the **show atm pnni aggregation link** command.

```
Switch# show atm pnni aggregation link

     PNNI link aggregation for local-node 2 (level=40, name=Switch5.2.40)
       Upnode Number: 10  Upnode Level: 40  Upnode Name: Switch7.2.40
        AggToken  InducPort BorderPort    Border Node(No./Name)
       ~~~~~~~~~~ ~~~~~~~~~ ~~~~~~~~~~~~~ ~~~~~~~~~~~~~~~~~~~~~
       0          2C49000   ATM0/1/0      9 Switch6
                            ATM0/1/2      1 Switch5
       3          2C49003   ATM0/0/0      11 Switch3

       Upnode Number: 11  Upnode Level: 24  Upnode Name: Switch8.3.24
        AggToken  InducPort BorderPort    Border Node(No./Name)
       ~~~~~~~~~~ ~~~~~~~~~ ~~~~~~~~~~~~~ ~~~~~~~~~~~~~~~~~~~~~
       0          2276000   ATM0/1/1      9 Switch6
```

Step 2 If an expected Upnode and AggToken (Aggregation Token) pair is missing from the PGL uplink summaries table, proceed to the "Checking Missing Upnode or Aggregation Token Pairs" section to determine whether the PGL knows about an uplink PTSE originated by one or more known border nodes.

Checking Missing Upnode or Aggregation Token Pairs

Use the following commands to check the uplink PTSEs and derived aggregation token configuration:

Command	Purpose
show atm pnni database [*internal-node-number*] [**detail**]	Confirms uplink PTSEs for the border nodes.
show atm pnni interface atm *card/subcard/port* [**detail**]	Confirms the hello state and derived aggregation token value on the node.

Follow these steps to troubleshoot missing upnode and aggregation token pairs on border nodes:

Step 1 Use the **show atm pnni database** command to check the PGL for an uplink PTSE originated by border nodes.

```
Switch# show atm pnni database 1
Node 1 ID 96:160:47.00918100000000E04FACB401.00E04FACB401.00 (name: Switch)
  PTSE ID  Length  Type  Seq no.   Checksum  Lifetime   Description
  1        92      97    117       37853     3143       Nodal info
  2        52      224   3331      18077     3016       Int. Reachable Address
Switch#
```

Step 2 Check to see whether an expected upnode and aggregation token pair is missing from the table listed on the PGL. Also check to see whether the PGL receives an uplink PTSE originated by one or more of the known border nodes.

Step 3 Use the *internal-node-number* option and the **detail** option to examine the contents of the uplink PTSE for a border node.

```
Switch# show atm pnni database 1
Node 1 ID 96:160:47.00918100000000E04FACB401.00E04FACB401.00 (name: Switch)
  PTSE ID  Length  Type  Seq no.   Checksum  Lifetime   Description
  1        92      97    117       37853     3143       Nodal info
  2        52      224   3331      18077     3016       Int. Reachable Address
Switch# show atm pnni database 1 detail
Node 1 ID 96:160:47.00918100000000E04FACB401.00E04FACB401.00 (name: Switch)
  PTSE ID  Length  Type  Seq no.   Checksum  Lifetime   Description
  1        92      97    117       37853     3135       Nodal info
    Time to refresh 1441, time to originate 0
    Type 97 (Nodal info), Length 48
    ATM address 47.00918100000000E04FACB401.00E04FACB401.00
    priority 0, leader bit NOT SET
    preferred PGL 0:0:00.000000000000000000000000.000000000000.00
  2        52      224   3331      18077     3008       Int. Reachable Address
    Time to refresh 1478, time to originate 0
    Type 224 (Int. Reachable Address), Length 32, Port 0, vp capable
    Scope (level) 0, Address info length (ail) 16, Address info count 1
    Pfx: 47.0091.8100.0000.00e0.4fac.b401..., length 104
Switch#
```

Step 4 If an expected uplink PTSE is missing, enter the same command on the border node switch router.

Step 5 If the uplink PTSE is present in the border node database but not in the PGL database, see the "Troubleshooting the PNNI Database" section on page 6-26 for further debugging.

Step 6 If the uplink PTSE is missing from the border node database, use the **show atm pnni interface atm** command to verify the hello state for the interface to the other peer group.

```
Switch# show atm pnni interface atm 10/0/0
PNNI Interface(s) for local-node 1 (level=96):
Port ATM10/0/0 RCC is up  , Hello state 2way_in    with node Switch Error: Port
Looped back
  Next hello occurs in 2 seconds, Dead timer fires in 67 seconds
Switch#
```

Step 7 If a lower level outside link interface is not in the common outside state, proceed to the "Checking the PNNI Lowest Level Interface" section on page 6-32.

Step 8 If the missing interface is a higher level induced uplink, perform the same checks at the next lower hierarchy level on the switch router acting as the LGN (and child PGL) node.

Step 9 If the derived aggregation token does not have the expected value, use the **detail** option to show additional interface information.

```
Switch# show atm pnni interface atm 10/0/0 detail
PNNI Interface(s) for local-node 1 (level=96):
Port ATM10/0/0 RCC is up  , Hello state 2way_in    with node Switch Error: Port
Looped back
  Next hello occurs in 6 seconds, Dead timer fires in 69 seconds
  CBR    : AW 5040 MCR 155519 ACR 147743 CTD 154 CDV 138 CLR0 10 CLR01 10
  VBR-RT : AW 5040 MCR 155519 ACR 155519 CTD 707 CDV 691 CLR0 8 CLR01 8
  VBR-NRT: AW 5040 MCR 155519 ACR 155519 CLR0 8 CLR01 8
  ABR    : AW 5040 MCR 155519 ACR 0
  UBR    : AW 5040 MCR 155519
➔ Aggregation Token: configured 0 , derived 0, remote 0
Switch#
```

Step 10 Check for correct Aggregation Token local and remote configuration.

Note Mismatched aggregation configuration results in a derived aggregation token value of zero.

Step 11 Verify the expected upnode node ID and common peer group ID.

Checking the Induced Port on the LGNs

If the induced port value is missing or does not appear to be functional for an aggregate token and upnode combination, use the following command to check the higher level interfaces for the parent LGN local node.

Use the following command to check the induced port on the LGN:

Command	Purpose
show atm pnni interface local-node *node-index*	Confirms the higher level interfaces for the parent LGN local node configuration.

Follow these steps to troubleshoot the port on the LGN:

Step 1 Use the **show atm pnni interface local-node** *node-index* command to check the induced port on the LGN.

Step 2 If the interface port does not appear up, see the “Troubleshooting PNNI SVCC-RCC and Higher Level Links” section on page 6-39.

Checking Link Aggregation

Use the following commands to check link aggregation:

Command	Purpose
show atm pnni aggregation link border-detail	Confirms per service class aggregation mode (best-link or aggressive).
show atm pnni database *internal-node-number* **detail**	Confirms uplink PTSEs.

Follow these steps to troubleshoot the aggregated metrics along with the border node interface metrics for each aggregation token:

Step 1 Use the **show atm pnni aggregation link border-detail** command to confirm per-service class aggregation mode.

Step 2 Check the per-service class aggregation mode (best-link or aggressive). The aggregation mode can be changed to control the resulted aggregated metrics.

- If the aggregated metrics are inaccurate because of the contribution of one or more border node interfaces that are significantly different from all others, new aggregation tokens can be configured to treat those aggregated links separately at the higher level.
- If the contribution from a border node interface is not as expected, check the border node uplink PTSE, using the **show atm pnni database** *n* **detail** command where *n* is the node number of the border node.

Step 3 Use the same **show atm pnni database** *n* **detail** command entered on the border node switch router to verify the same uplink PTSE information.

- If the two **show atm pnni database** *n* **detail** commands display significantly different information, see the "Troubleshooting the PNNI Database" section on page 6-26. Otherwise, see "Checking PNNI Interface Metric Configuration for Lowest Level Interfaces" section on page 6-36.
- If the border node is a higher level LGN with an induced uplink, check the link aggregation at the next lower hierarchy level on the switch router acting as the border node.

Troubleshooting PNNI Addresses and Address Summarization

This section describes how to troubleshoot PNNI address and address summarization problems. Summary addresses can be used to decrease the amount of information advertised by a PNNI node, and thereby contribute to scaling in large networks.

This section contains the following procedures:

- Checking PNNI Address Prefix Configurations
- Debugging Summary Addresses

Checking PNNI Address Prefix Configurations

A single default summary address is configured for each logical group node in the PNNI hierarchy. The length of that summary for any LGN equals the level of the child peer group, and its value is equal to the first level bits of the child peer group identifier. This address prefix is advertised into the peer group LGN.

Use the following commands to check the PNNI address prefix configuration:

Command	Purpose
show atm route	Confirms the list of all prefixes known by the node.
show atm addresses	Confirms that the correct prefix is present for the active ATM Address.
show atm pnni database *internal-node-number ptse-id* **detail**	Confirms the actual prefixes being advertised.
show atm pnni hierarchy network detail	Confirms which switch router is acting as the ancestor LGN.
show atm route	Confirms that the child PGL is up and that the scope is appropriate to allow advertising at the higher level.
show atm pnni scope	Confirms the configuration of the UNI scope map.

Follow these steps to troubleshoot PNNI address prefix configuration:

Step 1 Use the **show atm route** command to confirm the list of all prefixes known by the node.

```
Switch# show atm route
Codes: P - installing Protocol (S - Static, P - PNNI, R - Routing control),
       T - Type (I - Internal prefix, E - Exterior prefix, SE -
                 Summary Exterior prefix, SI - Summary Internal prefix,
                 ZE - Suppress Summary Exterior, ZI - Suppress Summary Internal)
P  T Node/Port        St Lev Prefix
~ ~~ ~~~~~~~~~~~~~~~~ ~~ ~~~ ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
P SI 1   0            UP 0   47.0091.8100.0000.00e0.4fac.b401/104
R  I 1   ATM10/1/0    UP 0   47.0091.8100.0000.00e0.4fac.b401.0080.1c93.8060/152
R  I 1   ATM10/0/1    UP 0   47.0091.8100.0000.00e0.4fac.b401.00e0.4fac.b030/152
R  I 1   ATM10/0/1    UP 0   47.0091.8100.0000.00e0.4fac.b401.00e0.4fac.b031/152
R  I 1   ATM10/0/1    UP 0   47.0091.8100.0000.00e0.4fac.b401.00e0.4fac.b032/152
R  I 1   ATM13/0/0    UP 0   47.0091.8100.0000.00e0.4fac.b401.00e0.4fac.b401/152
R  I 1   ATM13/0/0    UP 0   47.0091.8100.0000.00e0.4fac.b401.00e0.4fac.b402/152
R  I 1   ATM13/0/0    UP 0   47.0091.8100.0000.00e0.4fac.b401.00e0.4fac.b403/152
R  I 1   ATM13/0/0    UP 0   47.0091.8100.0000.00e0.4fac.b401.00e0.4fac.b404/152
R  I 1   ATM13/0/0    UP 0   47.0091.8100.0000.00e0.4fac.b401.00e0.4fac.b405/152
R  I 1   ATM13/0/0    UP 0   47.0091.8100.0000.00e0.4fac.b401.4000.0c/128
Switch#
```

Step 2 Verify that any expected address, prefix, or summary address is in the list of prefixes.

If the interface addresses do not have the expected prefix, verify that the correct prefix is present for the active ATM address, using the **show atm address** command.

Step 3 To see the actual prefixes being advertised by a local node, use the **show atm pnni database** *internal-node-number* command to get the PTSE ID number for the internal reachable address PTSE.

Step 4 Use the **show atm pnni database** *internal-node-number ptse-id* **detail** command to see the full contents of the PTSE.

```
Switch# show atm pnni database 1 1 detail
  1        92      97    551       37417     2116       Nodal info
    Time to refresh 646, time to originate 0
    Type 97 (Nodal info), Length 48
    ATM address 47.00918100000000E04FACB401.00E04FACB401.01
    priority 0, leader bit NOT SET, restricted transit bit NOT SET
    complex node bit NOT SET, restricted branching bit NOT SET
    non-transit for PGL election bit NOT SET
    preferred PGL 0:0:00.000000000000000000000000.000000000000.00
    Type 640 (System Capabilities Info), Length 24
    System Type: ls1010, Major Version: 11, Minor Version: 3
    System Name: Switch
Switch#
```

Step 5 To see if a prefix is being advertised at higher levels, determine which switch router is acting as the ancestor LGN by using the **show atm pnni hierarchy network detail** command.

Step 6 Use the **show atm pnni database** *internal-node-number ptse-id* **detail** command on the switch router acting as the ancestor LGN.

Step 7 If the expected prefix is not being advertised at the higher level, display the same information for the child PGL. If it is not present at the child PGL level, but was present at the originating node, see the "Troubleshooting the PNNI Database" section on page 6-26.

Step 8 If the prefix is present at the child PGL, but is missing for the parent LGN, verify that it is listed as up by using the **show atm route** command.

> **Note** It is normal for prefixes to be missing at the higher level if there is a matching summary or a suppressed summary present at its level.

Step 9 Verify that the scope (level) is appropriate to advertise at the desired higher levels.

Step 10 If the scope (level) does not have the expected value for a local prefix, check the configuration of the UNI scope map, using the **show atm pnni scope** command. If it is not the desired map, the mode can be changed to manual, and the desired scope translation levels can be configured.

It is normal for the prefixes to be missing if there is a shorter matching summary prefix configured at its level. The summary prefix will be advertised instead of any longer prefixes that match.

However, if the summary prefix is configured for suppress, none of the prefixes that match it will be advertised.

Debugging Summary Addresses

Use the following commands to show summary addresses:

Command	Purpose
show atm pnni summary [**local-node** *node-index*]	Confirms summary information advertised by PNNI node.
show running-config	Confirms that no auto summary is not configured and summary address has not been manually configured.
show atm addresses	Confirms active and no active switch router addresses.
show atm pnni local-node	Confirms node and peer group IDs of higher local nodes.

Follow these steps to troubleshoot summary addresses and suppressed summary addresses for all of the local nodes on an switch router:

Step 1 Use the **show atm pnni summary** command to display the PNNI summary address configuration.

Step 2 Check to see whether any expected summary addresses appear in the list for the expected local node for the correct Int (Internal) or Ext (External) type with the expected suppressed or non suppressed attribute.

Step 3 Verify that all longer prefixes and addresses matching any summary addresses are reachable at the local node or at a node in a child peer group. Otherwise, some addresses might be unreachable.

Step 4 Verify that the scope (level) is appropriate to advertise at all desired higher levels.

Step 5 If the default switch router address summary is missing, use the **show running-config** command to make sure that no auto summary is not configured for the affected local node.

Step 6 If an automatically generated ATM summary address is not the expected address, use the **show atm address** command to show the configured active and nonactive switch router addresses.

Step 7 Use the **show atm pnni local-node** command to check the node IDs and peer group IDs of higher level local nodes. If they are not based on the prefix of the ATM address, verify that no other peer group IDs have been manually configured. Also, verify that the lowest level node on the switch router has been disabled and reenabled since the last time the active switch router ATM address was reconfigured.

Step 8 If an unexpected summary address appears in the list, use the **show running-config** command to make sure that the summary address has not been manually configured.

If a summary prefix has been configured, but it is not possible to route to the summarized addresses from another peer group, check for an overlapping summary address within the other peer group. If the overlapping summary is for an automatically generated prefix, it could mean the ATM node addresses need to be modified to give unique prefixes for the ancestors of the two peer groups.

Troubleshooting Virtual Path Tunnel Connections

This section describes how to troubleshoot virtual path (VP) tunnels. VP tunnels are used primarily between private ATM networks across public ATM networks, such as telecom carriers, that do not yet support ATM signalling. Signalling traffic is mapped into the VP tunnel and the switch routers that allocate virtual channel connections (VCCs) on that VP instead of the default VP=0. With these connections, signalling can travel transparently through the public network.

In the example network in Figure 6-4, the PVC tunnel connection configured between the switch router on Floor 1 of the administration building and the switch router on Floor 1 of the remote sales building has the following interface and subinterface numbers:

- AdminFl1ls1, ATM interface 1/0/0, PVP 99
- RsalFl1Ls1, ATM interface 4/0/0, PVP 99

Figure 6-4 VP Tunnel Test in the Example Network

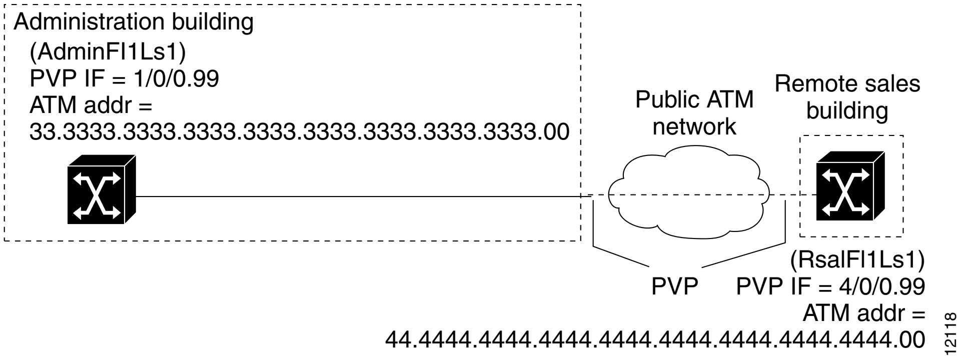

This section contains the following procedures:

- Checking VP Tunnel Configuration
- Checking Virtual Path PVP Configuration
- Debugging VP Tunnel Connection Management

For detailed configuration information, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

Checking VP Tunnel Configuration

If the permanent virtual path (PVP) subinterface numbers do not match on both ends of the VP tunnel, the connection is not established.

To show the ATM virtual interface configuration, use the following command:

Command	Purpose
show atm interface [**atm** *card/subcard/port*[**.***vpt#*]]	Shows the ATM interface configuration.

Follow these steps to troubleshoot VP tunnel connections:

Step 1 Use the **show atm interface atm** *card/subcard/port* command to display the configuration of switch router AdminFl1Ls1, located in the headquarters building at subinterface 1/0/0.99.

```
  AdminFl1Ls1# show atm interface atm 1/0/0.99

  Interface:      ATM1/0/0.99     Port-type:    vp tunnel
➔ IF Status:      UP              Admin Status: up
  Auto-config:    enabled         AutoCfgState: waiting for response from peer
  IF-Side:        Network         IF-type:      UNI
  Uni-type:       Private         Uni-version:  V3.0
  Max-VPI-bits:   0               Max-VCI-bits: 14
  Max-VP:         0               Max-VC:       16383
  Signalling:     Enabled
➔ ATM Address for Soft VC: 44.4444.4444.4444.4444.4444.4444.4444.4444.00
  Configured virtual links:
    PVCLs SoftVCLs   SVCLs  Total-Cfgd  Installed-Conns   4      0     0        4
  4

  AdminFl1Ls1#
```

Step 2 Check the IF Status field to confirm the interface is up. If it is not, see Chapter 5, “Troubleshooting Switch Router ATM Interface Connections.”

Step 3 Check the Admin Status field to confirm that the interface is up. If it is not, see Chapter 5, “Troubleshooting Switch Router ATM Interface Connections.”

Step 4 Check the interface and ATM Address for Soft VC fields. These values indicate that the VP tunnel is configured correctly.

Step 5 Use the **show atm interface atm** *card/subcard/port* command to display the configuration of the ATM switch router RsalFl1Ls1, located in the remote sales building at subinterface 4/0/0.99:

```
  RsalFl1Ls1# show atm interface atm 4/0/0.99

  Interface:      ATM4/0/0.99     Port-type:    vp tunnel
➔ IF Status:      UP              Admin Status: up
  Auto-config:    enabled         AutoCfgState: waiting for response from peer
  IF-Side:        Network         IF-type:      UNI
  Uni-type:       Private         Uni-version:  V3.0
  Max-VPI-bits:   0               Max-VCI-bits: 14
  Max-VP:         0               Max-VC:       16383
  Signalling:     Enabled
➔ ATM Address for Soft VC: 33.3333.3333.3333.3333.3333.3333.3333.3333.00
  Configured virtual links:
    PVCLs SoftVCLs   SVCLs  Total-Cfgd  Installed-Conns   4      0     0        4
  4

  RsalFl1Ls1#
```

Step 6 Check the IF Status field to confirm the interface is up. If it is not, see Chapter 5, “Troubleshooting Switch Router ATM Interface Connections.”

Step 7 Check the Admin Status field to confirm the interface is up. If it is not, see Chapter 5, “Troubleshooting Switch Router ATM Interface Connections.”

Step 8 Check the Interface and ATM address for Soft VC fields. These values indicate that the VP tunnel is configured correctly.

If you determine that the PVP is not configured correctly, refer to the “Configuring Virtual Connections” chapter in the *ATM Switch Router Software Configuration Guide*.

Continue with the next phase of VP tunnel troubleshooting if you still have not determined the problem.

Checking Virtual Path PVP Configuration

To confirm the ATM virtual connection (VC) interface configuration, use the following command:

Command	Purpose
show atm vc interface atm *card/subcard/port*	Shows the ATM VC interface configuration.

The following example shows how to confirm the configuration of ATM subinterface 1/0/0.99 on the switch router AdminFl1Ls1 located in the administration building:

```
  AdminFl1Ls1# show atm vc interface atm 1/0/0
  Interface    VPI   VCI   Type    X-Interface  X-VPI X-VCI  Encap Status
  ATM1/0/0     0     5      PVC     ATM2/0/0     0     41    QSAAL  UP
  ATM1/0/0     0     16     PVC     ATM2/0/0     0     33    ILMI   UP
➔ ATM1/0/0     99    40     PVC     ATM4/0/0.99  99    50    UP
  AdminFl1Ls1#
```

The interface ATM 1/0/0 field indicates that the cross-connect is configured correctly.

The following example shows how to confirm the configuration of ATM subinterface 1/0/0.99 on the switch router RsalFl1Ls1 located in the remote sales building:

```
  RsalFl1Ls1# show atm vc interface atm 4/0/0
  Interface    VPI   VCI   Type    X-Interface  X-VPI X-VCI  Encap Status
  ATM4/0/0     0     5      PVC     ATM2/0/0     0     41    QSAAL  UP
  ATM4/0/0     0     16     PVC     ATM2/0/0     0     33    ILMI   UP
➔ ATM4/0/0     99    40     PVC     ATM1/0/0.99  99    50           UP
  RsalFl1Ls1#
```

The interface ATM 4/0/0 field indicates that the cross-connect is configured correctly.

If you determine that the PVP is not configured correctly, refer to the “Configuring Virtual Connections” chapter in the *ATM Switch Router Software Configuration Guide* for configuration information.

Debugging VP Tunnel Connection Management

Use the following commands to debug the VP tunnel connection management:

Command	Purpose
debug atm conn errors	Enables connection management error debugging.
debug atm conn events	Enables connection management event debugging.
no debug all	Disables all debugging.

Troubleshooting Dropped Connections

This section describes how to troubleshoot the PVC traffic being dropped. In the example network in Figure 6-5, the connection between the DNS and e-mail servers and the switch router on Floor 1 of the administration building and the Catalyst 5000 switch on Floor 1 of the manufacturing building is dropping cells at some node in the connection.

This connection includes the following interfaces:

- AdminFl1Ls1, ATM interface 1/0/0
- ManuFl1Ls1, ATM interface 0/1/0
- ManuFl1Ls1, ATM interface 4/0/0
- ManuFl1CaS1, ATM LANE interface 1/1
- AdminFl1Ls1, ATM interface 4/0/1

Figure 6-5 PVC VPI and VCI Test in the Example Network

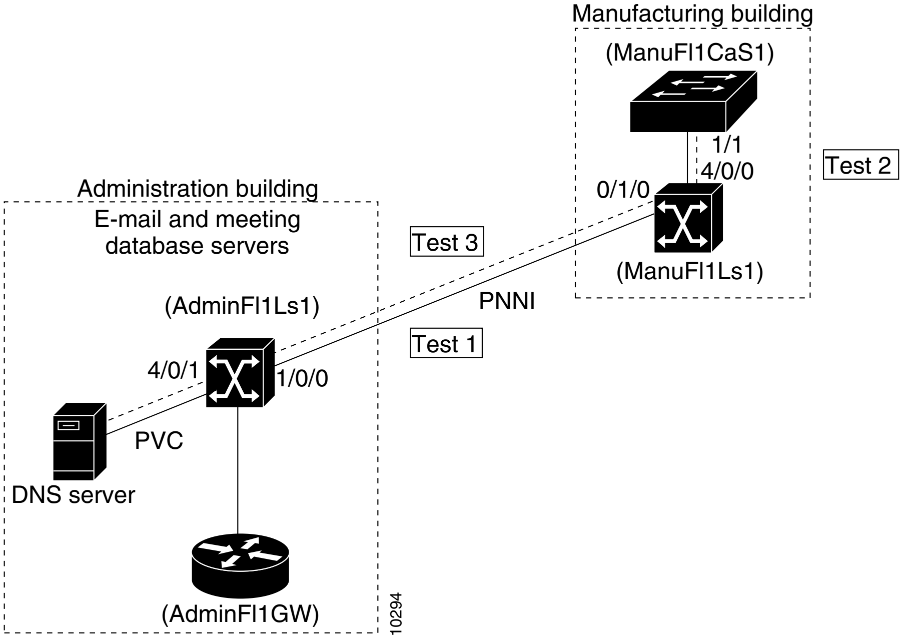

This section contains the following procedures:

- Determining Cell Drop Location
- Checking Line and Circuit Oversubscription
- Checking Traffic Priority
- Checking Network Circuit Timing

For detailed configuration information, refer to the “Configuring Resource Management” chapter in the *ATM Switch Router Software Configuration Guide*.

Determining Cell Drop Location

Use the following command to determine where the cells are being dropped.

Command	Purpose
show atm vc traffic interface atm *card/subcard/port*	Checks the VCs for interface where cells are being dropped.

Note The recommended procedure is to start at the center of the circuit and work outward until you find an switch router with mismatched receive and transmit cell counts.

Follow these steps to troubleshoot a VC to determine where the cells are being dropped along the length of the circuit:

Step 1 Use the **show atm vc traffic interface atm** command to look for mismatching numbers on both ends of the cable starting with the backbone interface connection at the switch router in the administration building and ending with the backbone interface connection at the manufacturing building.

```
AdminFl1Ls1# show atm vc traffic interface atm 1/0/0
Interface    VPI     VCI      Type      rx-cell-cnts     tx-cell-cnts
ATM1/1/0     0       5         PVC          672286           672286
ATM1/1/0     0       16        PVC              45               45
→ ATM1/1/0     0       18        PVC          730020           730155
ATM1/1/0     12      67        PVC               0                0
AdminFl1Ls1#
```

Step 2 Use the **show atm vc traffic interface atm** command to look for mismatching numbers on the interface connection at the switch router in the manufacturing building.

```
ManuFl1Ls1# show atm vc traffic interface atm 0/1/0
Interface    VPI     VCI      Type      rx-cell-cnts     tx-cell-cnts
ATM0/1/0     0       5         PVC          672286           672286
ATM0/1/0     0       16        PVC              45               45
→ ATM0/1/0     0       18        PVC          730020           730155
ATM0/1/0     12      67        PVC               0                0
ManuFl1Ls1#
```

Step 3 Use the **show atm vc traffic interface atm** command to look for mismatching numbers on the interface connection between the ATM switch router and the Catalyst 5000 Fast Ethernet switch in the manufacturing building.

```
  ManuFl1Ls1# show atm vc traffic interface atm 4/0/0
  Interface    VPI     VCI      Type      rx-cell-cnts     tx-cell-cnts
  ATM4/0/0     0       5         PVC           672286           672286
  ATM4/0/0     0       16        PVC               45               45
➔ ATM4/0/0     0       18        PVC           730020              155
  ATM4/0/0     12      67        PVC                0                0
  ManuFl1Ls1#
```

Notice that the number of received and transmitted cell counts are vastly different, which indicates that this is the interface where the cells are being dropped.

Continue with the next phase of troubleshooting to determine why the cells are being dropped.

Checking Line and Circuit Oversubscription

Use the following commands to check for oversubscription of the line and circuit under test:

Command	Purpose
show interfaces atm *card/subcard/port*	Checks to see if the line is oversubscribed.
show atm interface atm *card/subcard/port*	Checks to see if the circuit is oversubscribed.

Follow these steps to troubleshoot a VC by checking for oversubscription of the line and circuit under test:

Step 1 Use the **show interface atm** command to check ATM interface 1/0/0 for oversubscription.

Step 2 Use the **show atm interface atm** command to check ATM interface 4/0/0 for oversubscription.

If the line or circuit is oversubscribed, causing cells to be dropped, add more interfaces or circuits between the switch routers.

If you determine that the line or circuit is oversubscribed, refer to the “Configuring Virtual Connections” chapter in the *ATM Switch Router Software Configuration Guide*.

Continue with the next phase of troubleshooting if you still have not determined why the cells are being dropped.

Checking Traffic Priority

If a circuit is configured with multiple traffic types and some have a higher priority or QoS, cells with a lower priority are going to be dropped on a congested circuit.

For detailed configuration information, refer to the "Configuring Resource Management" chapter in the *ATM Switch Router Software Configuration Guide*.

Use the following commands to determine the configuration cell traffic priority and policing:

Command	Purpose
show atm interface resource atm *card*/*subcard*/*port*	Confirms the configuration of resource management looking for traffic priority conflicts.
show atm vc interface atm *card*/*subcard*/*port vpi vci*	Confirms the configuration of the VC looking for policing conflicts.

Follow these steps to determine the configuration cell traffic priority and policing:

Step 1 Use the **show atm interface resource atm** command to confirm traffic priority.

Step 2 Use the **show atm vc interface atm** command to confirm traffic policing.

If you determine that traffic priority or policing is causing cells to be dropped, refer to the "Configuring Resource Management" chapter in the *ATM Switch Router Software Configuration Guide*.

Continue with the next phase of troubleshooting dropped cells if you still have not determined the cause of the problem.

Checking Network Circuit Timing

If the network timing is misconfigured, the network clock can become unsynchronized and the switch router can start dropping cells.

For detailed configuration information, refer to the "Initially Configuring the ATM Switch Router" chapter in the *ATM Switch Router Software Configuration Guide*.

Use the following commands to determine the clocking configuration of the interface:

Command	Purpose
show network-clocks	Shows the network clocking configuration.
show running-config	Shows the interface clock source configuration.
show controllers [**atm** *card/subcard/port*]	Shows the interface controller status.

Follow these steps to determine the clocking configuration of the interface:

Step 1 Use the **show network** command to display the clock source configuration.

```
ManuFl1Ls1# show network-clocks
➔ Priority 1 clock source: ATM0/0/0
Priority 2 clock source: ATM0/0/3
Priority 3 clock source: ATM1/0/0
Priority 4 clock source: No clock
Priority 5 clock source: System clock

Current clock source: ATM0/0/0, priority: 1

ManuFl1Ls1#
```

Step 2 Make note of the interface configured as Priority 1 clock source.

Step 3 Use the **show running-config** command to display the clock source configuration of ATM interface 4/0/0.

```
ManuFl1Ls1# show running-config
Building configuration...

Current configuration:
!
version 11.2
no service pad
service udp-small-servers
service tcp-small-servers
!
hostname Switch
!
boot bootldr bootflash:/tftpboot/ls1010-wp-mz.112-1.4.WA3.0.15
!
network-clock-select 2 ATM3/1/0

<Information Deleted>

!
interface ATM4/0/0
 no keepalive
 atm manual-well-known-vc
 atm access-group tod1 in
 atm pvc 0 35 rx-cttr 3 tx-cttr 3  interface  ATM2/0/0 0 any-vci  encap qsaal
 atm route-optimization soft-vc interval 360 time-of-day 18:0 5:0
➔ clock-source network-derived
!

<Information Deleted>

ManuFl1Ls1#
```

The clock source field indicates the clocking configuration of ATM interface 4/0/0.

Step 4 Use the **show controllers atm** *card/subcard/port* command to display the interface controller status of ATM interface 4/0/0.

```
ManuFl1Ls1# show controllers atm 4/0/0
IF Name: ATM4/0/0    Chip Base Address: A8808000
Port type: 155UTP    Port rate: 155 Mbps    Port medium: UTP
Port status:SECTION LOS    Loopback:None    Flags:8300
➔ TX Led: Traffic Pattern    RX Led: Traffic Pattern  TX clock source:  network-derived
Framing mode:  sts-3c
Cell payload scrambling on
Sts-stream scrambling on
OC3 counters:
  Key: txcell - # cells transmitted    rxcell - # cells received    b1     - # section
BIP-8 errors    b2     - # line BIP-8 errors    b3     - # path BIP-8 errors    ocd    -
# out-of-cell delineation errors - not implemented    g1     - # path FEBE errors    z2
- # line FEBE errors    chcs   - # correctable HEC errors    uhcs   - # uncorrectable HEC
errors

<Information Deleted>

phy_tx_cnt:0, phy_rx_cnt:0
ManuFl1Ls1#
```

Step 5 Check the TX clock source field. This field indicates that the clocking configuration of the interface is either internal or network derived.

If you determine that the clock configuration is causing cells to be dropped, refer to the “Initially Configuring the ATM Switch” chapter in the *ATM Switch Router Software Configuration Guide*.

For more information on troubleshooting network clocking, refer to the “Troubleshooting Network Clocking” section on page 9-7.

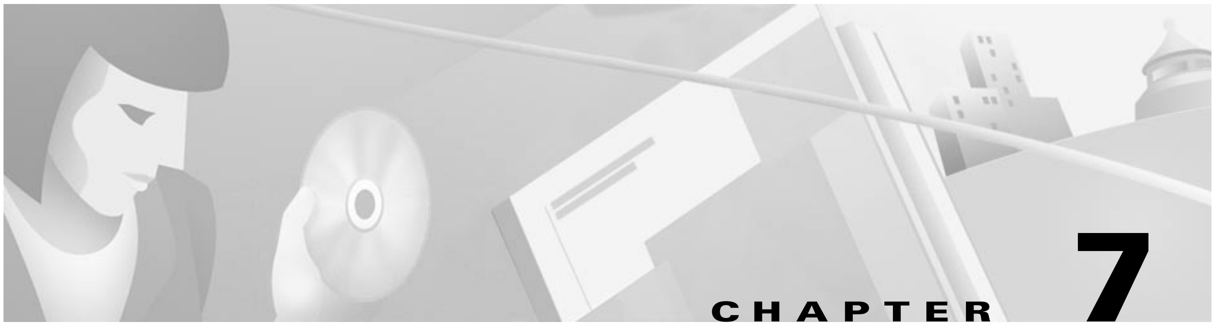

CHAPTER 7

Troubleshooting LAN Emulation Switching Environments

This chapter provides troubleshooting information for connectivity and performance problems in LAN emulation (LANE) switching environments. The ATM Forum defined the LANE specification so that legacy LAN users can take advantage of the benefits of ATM without requiring modifications to end-station hardware or software. For an overview of LANE on a switch router, refer to the “Configuring LAN Emulation” chapter in the *ATM Switch Router Software Configuration Guide*. For an overview of LANE on the Catalyst 5000 and 6000 ATM modules, refer to the “Configuring ATM LAN Emulation” chapter in the *ATM Software Configuration Guide and Command Reference*: *Catalyst 5000 and 6000 ATM Modules*.

Before you begin, make sure that all physical port connections are working correctly. See Chapter 5, “Troubleshooting Switch Router ATM Interface Connections.”

This chapter contains the following sections:

- Example of LANE Troubleshooting, page 7-1
- Initial Troubleshooting of LANE, page 7-3
- Checking the Configuration Server Database, page 7-8
- Debugging the LANE Connection, page 7-9

Example of LANE Troubleshooting

This section describes LANE troubleshooting, using the example network in Figure 6-1 and in Chapter 4, “Example Network.”

Figure 7-1 LANE Example Network

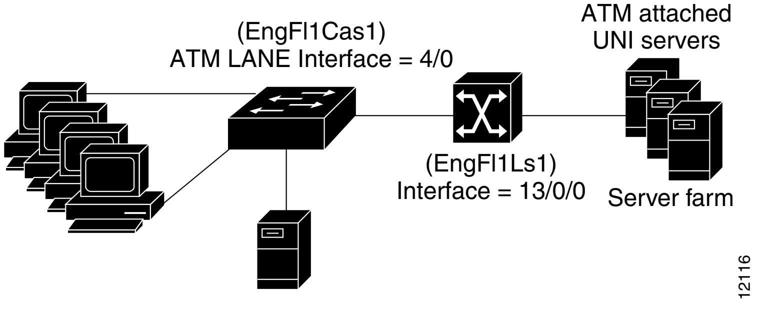

A single emulated LAN consists of a 155 multimode fiber physical connection between the switch router (EngFl1Ls1) and the Catalyst 5000 Fast Ethernet switch (EngFl1Cas1) in the engineering building. The switch router is configured as the LAN emulation configuration server (LECS), the LAN emulation server (LES), and the LAN emulation broadcast-and-unknown server (LANE BUS).

The LANE example network for the engineering building in Figure 7-1 is configured as follows:

- Switch router (EngFl1Ls1) LECS, LES, BUS configuration:
  - ATM 155 multimode fiber interface 13/0/0 connected to Catalyst 5000 ATM LANE module 4/0
  - Interface ATM address—47.009181000000000000000001.00400B0A2E82.01
  - Interface IP address—172.20.52.25 255.255.255.0
  - LANE client—47.00918100000000E04FACB401.00E04FACB402.00
  - LANE server—47.00918100000000E04FACB401.00E04FACB403.00
  - LANE BUS—47.00918100000000E04FACB401.00E04FACB404.00
  - LANE configuration server ATM address—47.00918100000000E04FACB401.00E04FACB405.00
  - Interface IP address—172.20.52.20 255.255.255.0
  - ELAN name—eng_elan
  - LANE database name—eng_database
- Catalyst 5000 switch (EngFl1Cas1) LANE client configuration:
  - ATM LANE module multimode interface 4/0 connected to switch router, ATM 155 multimode interface 13/0/0
  - Interface IP address—172.20.52.21 255.255.255.0
  - Interface ATM address—47.00918100000000E0F75D0401.00E0F75D041F.00
  - ELAN name—eng_elan
  - LANE database name—eng_database

Note The ATM addresses are examples only. The addresses on your switch might be different.

Note Emulated LAN entities coexist on one or more Cisco routers or switch routers. On Cisco routers or switch routers, each LANE server and broadcast-and-unknown server is always a single entity. Other LANE components include switch routers—those that support the Integrated Local Management Interface (ILMI) and signalling. Multiple emulated LANs can coexist on a single ATM network.

Initial Troubleshooting of LANE

This section describes how to use the Internet Control Message Protocol (ICMP) **ping** command to test for connectivity between the switch router and either a router or an Ethernet switch.

To test for Ethernet connectivity, perform the following task:

Command	Purpose
ping ip *ip-address*	Tests the configuration, using the **ping** command. The **ping** command sends an echo request to the host, which is specified in the command line.

Follow these steps to troubleshoot LANE connections:

Step 1 Use the **ping** command to confirm the connection between the switch router (EngFl1Ls1) and the Catalyst 5000 (EngFl1Cas1) in the engineering building:

```
EngFl1Ls1# ping 172.20.52.21
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.20.52.100, timeout is 2 seconds:
!!!!!
➔ Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
EngFl1Ls1# ping 172.20.52.20
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.20.52.20, timeout is 2 seconds:
!!!!!
➔ Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
EngFl1Ls1#
```

Step 2 Check the Success rate field. It should indicate 100 percent. This indicates that the ICMP packet was sent and returned.

If the success rate is less than 100 percent, continue with the following test to determine the problem with the LANE configuration. Refer to the "Configuring LAN Emulation" chapter in the *ATM Switch Router Software Configuration Guide*.

Checking Basic LANE Configuration

Use the following commands to test the LANE connection status and performance:

Command	Purpose
show lane default-atm-addresses	Shows the LANE default ATM address.
show lane	Shows the LANE connection status.
show lane client	Shows the LANE client connection status.

Follow these steps to confirm the LANE LECS, LES, and BUS configuration:

Step 1 Use the **show lane default-atm-addresses** command at the LANE component designated as the LECS, LES, and BUS to determine the addresses.

```
EngFl1Ls1# show lane default-atm-addresses
interface ATM13/0/0:
LANE Client:        47.00918100000000E04FACB401.00E04FACB402.**
LANE Server:        47.00918100000000E04FACB401.00E04FACB403.**
LANE Bus:           47.00918100000000E04FACB401.00E04FACB404.**
LANE Config Server: 47.00918100000000E04FACB401.00E04FACB405.00
note: ** is the subinterface number byte in hex
EngFl1Lsl#
```

Step 2 Use the **show lane** command to confirm the LANE configuration of the switch router.

Note This example show how to confirm the LANE configuration of a switch router. Use the same command to confirm the LANE configuration of a Catalyst 5000 or 6000 ATM module.

```
  EngFl1Lsl# show lane
  LE Config Server ATM13/0/0 config table: eng_dbase
➔ Admin: up  State: operational
  LECS Mastership State: active master
  list of global LECS addresses (7 seconds to update):
  47.00918100000000E04FACB401.00E04FACB405.00   <-------- me
➔ ATM Address of this LECS: 47.00918100000000E04FACB401.00E04FACB405.00 (auto)
               the above address has NOT yet been registered with ILMI
    actual user specified form: ...
  cumulative total number of unrecognized packets received so far: 0
  cumulative total number of config requests received so far: 0
  cumulative total number of config failures so far: 0

➔ LE Server ATM13/0/0.1  ELAN name: eng_elan  Admin: up  State: operational
  type: ethernet         Max Frame Size: 1516
➔ ATM address: 47.00918100000000E04FACB401.00E04FACB403.01
  LECS used: 47.00918100000000E04FACB401.00E04FACB405.00 connected, vcd 84

➔ LE BUS ATM13/0/0.1  ELAN name: eng_elan  Admin: up  State: operational
  type: ethernet         Max Frame Size: 1516
➔ ATM address: 47.00918100000000E04FACB401.00E04FACB404.01

  LE Client ATM13/0/0.1  ELAN name: eng_elan  Admin: up  State: initialState
  Client ID: unassigned        Next join attempt in 36 seconds
  Join Attempt: 18
  Last Fail Reason: Timeout on join request
  HW Address: 00e0.4fac.b402   Type: ethernet             Max Frame Size: 1516

  ATM Address: 47.00918100000000E04FACB401.00E04FACB402.01

   VCD  rxFrames  txFrames  Type       ATM Address
     0         0         0  configure  47.00918100000000E04FACB401.00E04FACB405.00
     0         0         0  direct     47.00918100000000E04FACB401.00E04FACB405.00
     0         0         0  distribute 00.000000000000000000000000.000000000000.00
     0         0         0  send       00.000000000000000000000000.000000000000.00
     0         0         0  forward    00.000000000000000000000000.000000000000.00
  EngFl1Lsl#
```

Step 3 Check the Admin (Administration) and State fields. The values should be up and operational, respectively.

Step 4 Check the ATM address of this LECS field. This ATM address should match the ATM address in the LANE configuration server displayed in Step 1.

Step 5 Within the LE BUS section of the display, check the ELAN name field. It should match the name configured. To determine the ELAN name, use the **show lane database** command and check the default ELAN field.

Step 6 Check the ATM Address field. This ATM address should match the address displayed in the LANE Server field shown in Step 1. The ATM address is appended with the corresponding subinterface number in hexadecimal.

If any of these fields do not match the actual LANE configuration, refer to the "Configuring LAN Emulation" chapter in the *ATM Switch Router Software Configuration Guide* and correct the configuration.

Checking LANE Client Configuration

Use the following command to troubleshoot the LANE client connectivity:

Command	Purpose
show lane client	Shows the connection status of the LANE client.

Follow these steps to confirm the configuration and status of the LANE client:

Step 1 Use the **show lane client** command to confirm the LANE client.

Note This example shows how to confirm the LANE client configuration at ATM subinterface 13/0/0.1 on a switch router. Use the same command to confirm LANE client configurations on the Catalyst 5000 or 6000 ATM modules.

```
  EngFl1Ls1# show lane client
➔ LE Client ATM13/0/0.1  ELAN name: eng_elan  Admin: up  State: initialState
  Client ID: unassigned       Next join attempt in 83 seconds
➔ Join Attempt: 11
➔ Last Fail Reason: Timeout on join request
  HW Address: 00e0.4fac.b402   Type: ethernet            Max Frame Size: 1516

  ATM Address: 47.00918100000000E04FACB401.00E04FACB402.01

   VCD  rxFrames  txFrames  Type       ATM Address
     0         0         0  configure  47.00918100000000E04FACB401.00E04FACB405.00
    87         1         2  direct     47.00918100000000E04FACB401.00E04FACB403.01
    90         1         0  distribute 47.00918100000000E04FACB401.00E04FACB403.01
    91         0         1  send       47.00918100000000E04FACB401.00E04FACB404.01
    94         0         0  forward    47.00918100000000E04FACB401.00E04FACB404.01
```

Note This same information appears in the previously described **show lane** command.

Step 2 Check the ELAN name field. This should match the ELAN name configured.

Step 3 Check the Admin and State fields. They should read up and operational, respectively.

If the Admin field shows “down”, the interface or subinterface is administratively shut down. To reenable the interface, use the **no shutdown** command on the interface or subinterface.

Step 4 Check the Join Attempt field. A high number of join attempts might mean that the LECS is unreachable. To determine the status of the LECS, use the **show lane config** command.

Step 5 Check the Last Fail Reason field for any of the following messages:

- Link went down

  The problem is on the physical layer. Check the cable and the module quality. Physically loopback the interface and check the status using the **show interfaces** command.

- Local config changed

  The switch detected a configuration change. Force the LANE client to join the ELAN with the **shutdown/no shutdown** command sequence on the interface or subinterface.

- Fail to set up config VC

  The LANE client cannot establish a VCC to the LECS. This failure might be caused by ILMI, which must be enabled to provide the ATM address prefix. If ILMI is not functioning, no ATM address prefix is distributed and no SVCs can be established. Check the LANE component addresses using the **show lane default-atm-addresses** command. The following example shows correct output:

```
EngFl1Ls1#show lane default-atm-addresses
interface ATM13/0/0:
LANE Client:        47.00918100000000E0F75D0401.00E0F75D0402.**
LANE Server:        47.00918100000000E0F75D0401.00E0F75D0403.**
LANE Bus:           47.00918100000000E0F75D0401.00E0F75D0404.**
LANE Config Server: 47.00918100000000E0F75D0401.00E0F75D0405.00
note: ** is the subinterface number byte in hex
```

  If ILMI is not suppling the ATM address prefix, the output appears as follows:

```
EngFl1Ls1#show lane default-atm-addresses
interface ATM13/0/0:
LANE Client:        ...00E0F75D0402.**
LANE Server:        ...00E0F75D0403.**
LANE Bus:           ...00E0F75D0404.**
LANE Config Server: ...00E0F75D0405.00
note: ** is the subinterface number byte in hex
```

  Use the **show atm ilmi-status** command to verify that the ILMI PVC is well defined. For more information on configuring ILMI, refer to the “Configuring ILMI” chapter in the *ATM Switch Router Software Configuration Guide*.

- Config VC being released

  The LECS address is incorrect or unreachable. If the LANE client could not reach the LECS, the **show lane client** command output looks like the following:

```
  EngFl1Ls1#show lane client
  LE Client ATM13/0/0.1  ELAN name: eng_elan  Admin: up  State: initialState
  Client ID: unassigned        Next join attempt in 93 seconds
  Join Attempt: 3909
  Known LE Servers: 0
➔ Last Fail Reason: Config VC being released
  HW Address: 0030.80ce.3a02   Type: ethernet             Max Frame Size: 1516
  ATM Address: 47.00918100000000E04FACB401.003080CE3A02.01
   VCD  rxFrames  txFrames  Type       ATM Address
➔    0         0         0  configure  47.007900000000000000000000.00A03E000001.00
```

Try the following actions to connect the LANE client with the LECS:

  - If the remote ATM switch is not a Cisco device, verify that the LECS address is advertised by ILMI. If it is not, use the well-known address of the LECS.
  - If the LECS ATM address is not explicitly configured on the LANE client switch, configure it with the **atm lecs-address-default** command.
  - If the LECS ATM address has been explicitly configured on the LANE client switch, compare the LECS ATM address on the LANE client switch, using the **show lane client** command, with the LECS ATM address on the switch hosting the LECS, using the **show lane config** command. Also check that the LECS is up and operational.

- Receiving negative config response

  The LECS refuses to connect to the ELAN. This is usually due to a configuration mistake, such as incorrectly specifying the ELAN type or name. Try the following actions to resolve the problem:

  - Check the LANE client ELAN type and name in the configuration, using the **show lane client** command. Remember that the ELAN names are case sensitive.
  - If the ELAN membership is restricted, use the **show lane database** command to verify that the LANE client ATM address is specified in the LANE database.
  - Use the **show lane server** command to verify that the LES connected to the LECS. The LES should have the same LECS ATM address information as a client.

- Control Direct VC being released

  The LANE client could not connect to the LES. The LES is either unreachable or misconfigured. If the LES address is hardcoded in the configuration, check the ATM address of the LES on the device where it is located, using the **show lane server** command. Compare that address with the ATM address configured for the LES on the LANE client switch, using the **show lane database** command.

- Receiving negative join response

  The LES refuses to connect. Try the following actions to resolve the problem:

  - If the ELAN is restricted, check the LANE database configuration, using the **show lane database** command, to ensure that it includes the ATM address of the LANE client.
  - If the LANE client and LES are configured on the same subinterface and the ATM address of the LES is explicitly configured with the **lane server-atm-address** command, the LANE client might be trying to contact a backup LES. These connection requests will be refused. Verify this with the **show lane client** and **show lane server** commands. To correct this situation, configure the LES on a different subinterface.

Step 6 Check the ATM Address column in the virtual channel circuit (VCC) configuration table. These ATM addresses should be propagated as follows:

- Configure, should be the same as the LECS ATM address
- Direct and distribute, should be the same as the LES ATM address
- Send and forward, should be the same as the LANE BUS ATM address

Step 7 Confirm that the rxFrames and txFrames columns contain values greater than 0.

Checking the Configuration Server Database

The LECS assigns individual clients to particular emulated LANs by directing them to the LANE server that corresponds to the emulated LAN. The LECS maintains a database of LANE client ATM or Media Access Control (MAC) addresses and their ELANs.

Note A LECS can serve multiple ELANs.

Use the following command in privileged EXEC mode to display the configuration of the LANE client database binding:

Command	Purpose
show lane database	Displays the LANE client database binding.

Follow these steps to confirm the configuration of the LANE database:

Step 1 Use the **show lane database** command to display the default binding of the LANE database of the switch router in the engineering building.

```
  EngFl1Ls1# show lane database
➔ LANE Config Server database table 'eng_dbase' bound to interface/s: ATM13/0/0
  default elan: eng_elan
  elan 'eng_elan': un-restricted
➔   server 47.00918100000000E04FACB401.00E04FACB403.01 (prio 0)
  EngFl1Ls1#
```

Step 2 Check the LANE Config Server database table field. It indicates the binding of the LANE client to the LANE database.

Step 3 Check the server field. The ATM address displayed should correspond to the ATM address shown in the **show lane default-atm-addresses** command in Step 1 of the previous section, “Checking Basic LANE Configuration.”

If the LANE client configuration server database is set up incorrectly, refer to the “Configuring LAN Emulation” chapter in the *ATM Switch Router Software Configuration Guide*.

Debugging the LANE Connection

This section outlines the **debug** commands used to troubleshoot the LANE setup and signalling.

Use the following **debug** commands to check the LANE configuration and setup processes:

Command	Purpose
debug lane client {**all** | **le-arp** | **packet** | **signalling** | **state** | **topology**}	Debugs all LANE client setup processes.
debug lane client signalling interface atm *card/subcard/port*	Debugs LANE client signalling processes for a specific ATM interface.
debug lane config {**all** | **events** | **packets**}	Debugs all LANE setup processes.
debug lane finder	Debugs all LANE LECS setup processes.
debug lane server interface atm *card/subcard/port*	Debugs LANE server processes for a specific ATM interface.
debug lane signalling interface atm *card/subcard/port*	Debugs LANE signalling processes for a specific ATM interface.
no debug all	Turns off all debugging.

Note Other helpful **debug** commands include the **debug atm ilmi** command and all the variations of the **debug atm sig** command.

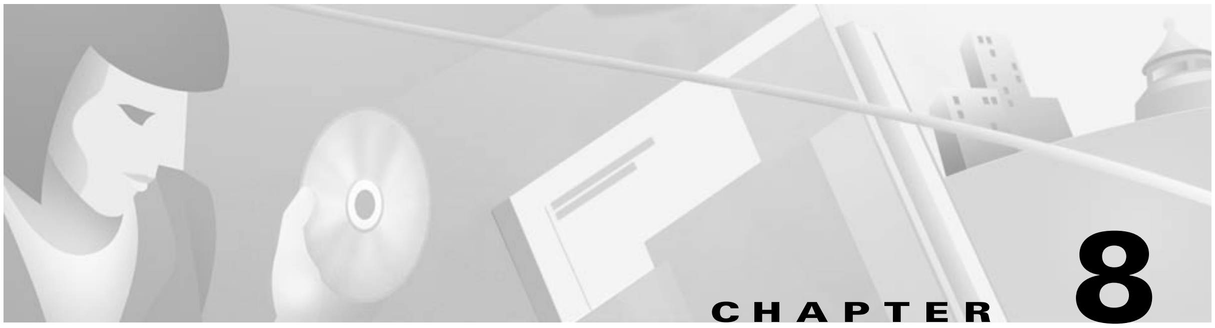

CHAPTER 8

Troubleshooting Tag Switching Connections

This chapter provides troubleshooting information for connectivity and performance problems in tag switching environments. For more information on tag switching, refer to the “Configuring Tag Switching” chapter in the *ATM Switch Router Software Configuration Guide*.

Before you begin, make sure that all physical port connections are working correctly. See Chapter 5, “Troubleshooting Switch Router ATM Interface Connections.”

This chapter contains the following sections:

Tag Switching Overview

Tag switching is a high-performance packet-forwarding technology that assigns tags to multiprotocol frames for transport across packet-based or cell-based networks.

In conventional Layer 3 forwarding, as a packet traverses the network, each router extracts forwarding information from the Layer 3 header. Header analysis is repeated at each router (hop) through which the packet passes.

In a tag switching network, the Layer 3 header is analyzed just once. It is then mapped into a short, fixed-length tag. At each hop, the forwarding decision is made by looking at the value of the tag only; there is no need to reanalyze the Layer 3 header. Because the tag is a fixed-length, unstructured value, looking it up is fast and simple.

A tag switching network consists of tag edge routers and tag switch routers, as shown in Figure 8-1. Tag edge routers are located at the edge of a tag switching network. They use standard routing protocols—such as Open Shortest Path First (OSPF)—to create routing tables that identify routes through the network. Based on the routing tables, tag edge routers use the Tag Distribution Protocol (TDP) to apply and distribute tags to other tag edge routers or tag switch routers. Tag switch routers are located at the core of a tag switching network. They receive TDP information from the tag edge routers and build their own forwarding database. Tag switch routers then switch the packets based on the tags only (without looking at the Layer 3 header).

How Tag Switching Works

When a tag edge router at the entry point of a tag switching network receives a packet for forwarding the following occurs:

1. The router analyzes the network layer header and performs any applicable network layer services such as security, accounting, or quality of service (QoS) classification.
2. The router chooses a route for the packet based on the information in its routing table, applies a tag, and forwards the packet to the next-hop tag switch router.
3. The tag switch router receives the tagged packet and switches the packet from switch router to switch router based on the tag only. The switch routers do not reanalyze the network layer header; they look only at the short, fixed-length tag.
4. The packet reaches the tag edge router at the exit point of the tag switched network, where the tag is removed and the packet is delivered.

Troubleshooting Tag Switching Example

In the example network in Figure 8-1, the primary campus network backbone is made up of two ATM switch routers connected to two Cisco routers:

- AdminFl1Rt1—Tag switching router located in the administration building
- AdminFl1Ls1—Tag switching switch router located in the administration building
- EngFl1Ls1—Tag switching switch router located in the engineering building
- EngFl1Rt1—Tag switching router located in the engineering building

Figure 8-1 Tag Switching Example Network

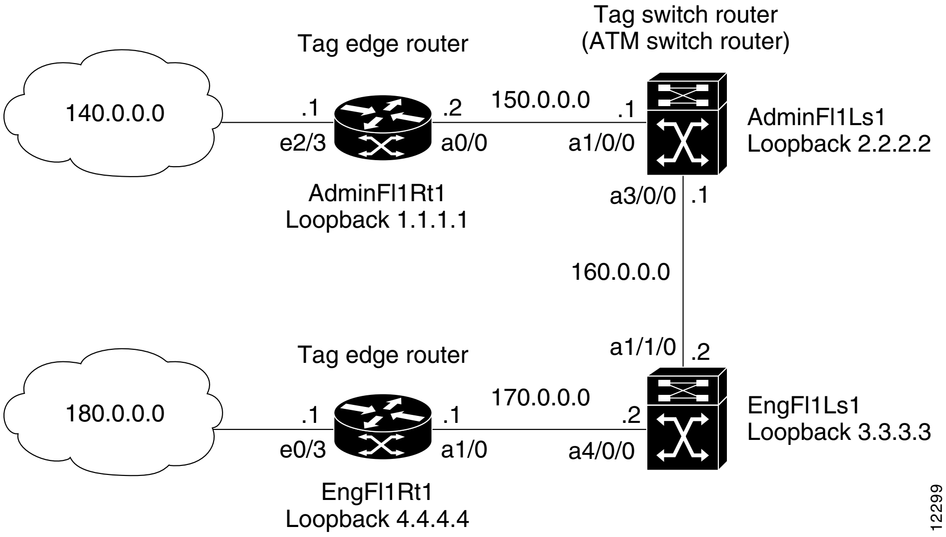

This network example is used to describe the troubleshooting examples in the rest of this chapter.

For detailed configuration information about tag switching, refer to the "Configuring Tag Switching" chapter in the *ATM Switch Router Software Configuration Guide*.

Initial Troubleshooting of Tag Switching

This section describes initial troubleshooting steps that you should perform when beginning to troubleshoot a tag switching connection.

At the switch router, use the following commands to check the tag switching configuration:

Command	Purpose
show tag-switching tdp discovery	Confirms the TDP identifier for the tag switching switch router or router that might be malfunctioning.
ping *tdp_id_of_neighbor*	Confirms that each tag switching switch router or router can connect to the TDP identifier of its neighbor.
show running-config	Confirms that tag switching is enabled on the switch router.
show tag-switching interfaces	Confirms the tag switching configuration on the ATM interface.
show tag-switching interfaces detail	Confirms the tag switching VPI[1] range on an interface.
show interfaces loopback 0	Confirms the loopback interface 0 configuration.
show ip ospf	Confirms the OSPF configuration.

[1] VPI = virtual path identifier

Follow these steps to confirm the TDP identifier for the routers or tag switching switch routers that might be malfunctioning:

Step 1 Use the **show tag-switching tdp discovery** command to determine the tag discovery protocol identifier of the tag switching switch router.

```
  AdminFl1Ls1# show tag-switching tdp discovery
➔ Local TDP Identifier:
      172.20.40.161:0
  TDP Discovery Sources:
➔ Interfaces:
         ATM1/0/0: xmit/recv
              TDP Id: 150.0.0.0:1
          ATM3/0/0.10: xmit/recv
              TDP Id: 160.0.0.0:1
  AdminFl1Ls1#
```

Step 2 Check the Local TDP Identifier field. This field indicates the TDP identifier for the local tag switching switch router or router for this session.

Step 3 Check the Interfaces field. This field displays the interfaces engaging in TDP discovery activity:

- xmit indicates that the interface is transmitting TDP discovery hello packets.
- recv indicates that the interface is receiving TDP discovery hello packets.

If either xmit or recv do not appear, refer to the “Configuring Tag Switching” chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to ping each tag switching switch router or router. This process confirms that each can connect to the TDP identifier of the neighbor:

Step 1 Use the **ping** command to confirm the connection to the TDP of the neighbor.

```
  AdminFl1Ls1# ping
  Protocol [ip]:
  Target IP address: 180.0.0.0
  Repeat count [5]:
  Datagram size [100]:
  Timeout in seconds [2]:
  Extended commands [n]: y
  Source address or interface: 140.0.0.0
  Type of service [0]:
  Set DF bit in IP header? [no]:
  Validate reply data? [no]:
  Data pattern [0xABCD]:
  Loose, Strict, Record, Timestamp, Verbose[none]:
  Sweep range of sizes [n]:
  Type escape sequence to abort.
  Sending 5, 100-byte ICMP Echoes to 180.0.0.0, timeout is 2 seconds:
  !!!!!
➔ Success rate is 100 percent (5/5), round-trip min/avg/max = 184/398/1188 ms
  AdminFl1Ls1#
```

Step 2 Check the Success rate field. This field should read “100 percent”. If it does not, continue with the following troubleshooting steps.

Follow these steps to confirm that tag switching is configured on the switch router and its interfaces:

Step 1 Use the **show running-config** command to confirm that tag switching is enabled on the ATM switch router.

```
AdminFl1Ls1# show running-config
Building configuration...
Current configuration:
!
version 11.3
no service pad
!

<Information deleted>

!
interface ATM0/1/1
 ip unnumbered Loopback0
➔ tag-switching ip
!
interface ATM1/0/0
 ip address 150.0.0.0 255.255.255.224
➔ tag-switching ip
!

<Information deleted>

!
end
AdminFl1Ls1#
```

Step 2 Check the tag switching switch router interface to confirm that tag switching is enabled on the connections.

For detailed interface configuration information about tag switching, refer to the "Configuring Tag Switching" chapter in the *ATM Switch Router Software Configuration Guide*.

The neighbor information branch can have information about all TDP neighbors or can be limited to the neighbor with a specific IP address, or TDP identifier, or to TDP neighbors known to be accessible over a specific interface.

Follow these steps to display the status of TDP sessions:

Step 1 Use the **show tag-switching tdp neighbor** command to display the status of TDP sessions.

```
  AdminFl1Ls1# show tag-switching tdp neighbor
➔ Peer TDP Ident: 1.0.12.12:2; Local TDP Ident 1.0.11.11:2
➔         TCP connection: 1.0.12.12.11008 - 1.0.11.11.711
➔         State: Oper; PIEs sent/rcvd: 2199/2198; Downstream on demand
➔         Up time: 02:31:58
          TDP discovery sources:
            ATM0/0/1
  Peer TDP Ident: 1.0.12.12:8; Local TDP Ident 1.0.11.11:7
          TCP connection: 1.0.12.12.11015 - 1.0.11.11.711
          State: Oper; PIEs sent/rcvd: 2119/2130; Downstream on demand
          Up time: 02:31:39
          TDP discovery sources:
            ATM0/1/0.19
  Peer TDP Ident: 1.0.12.12:7; Local TDP Ident 1.0.11.11:6
          TCP connection: 1.0.12.12.11016 - 1.0.11.11.711
          State: Oper; PIEs sent/rcvd: 2120/2119; Downstream on demand
          Up time: 02:31:38
          TDP discovery sources:
            ATM0/1/0.18
```

Step 2 Check the Peer TDP Ident field. This field indicates the TDP identifier of the neighbor (peer device) for this session.

Step 3 Check the Local TDP Ident field. This field indicates the TDP identifier for the local tag switching switch router or router for this session.

Step 4 Check the TCP connection field. This field indicates the TCP connection used to support the TDP session. The format for displaying the TCP connection is *peer IP address.peer port local IP address*.

Step 5 Check the PIEs sent/rcvd (Protocol Information Element sent or received) field. This field indicates the number of TDP PIEs sent to and received from the session peer device. The count includes the transmission and receipt of periodic keepalive PIEs, which are required for maintenance of the TDP session.

Step 6 Check the Up time field. This field indicates the length of time the TDP session has existed.

Follow these steps to confirm the tag switching interface configuration on the switch router:

Step 1 Use the **show tag-switching interfaces** command to confirm the configuration and connection of the tag switching interfaces.

```
AdminFl1Ls1# show tag-switching interfaces
Interface              IP    Tunnel   Operational
ATM1/0/0               Yes   No       Yes
ATM3/0/0               Yes   No       Yes
AdminFl1Ls1#
```

Step 2 Check the IP field. This field indicates whether the interface is configured to tag IP packets.

Step 3 Check the Operational field. This field shows whether the packets are being tagged.

Step 4 Use the **show tag-switching interfaces detail** command to confirm the tag switching VPI range on an interface.

```
   AdminFl1Ls1# show tag-switching interfaces detail
   Interface ATM1/0/0:
           IP tagging enabled
           TSP Tunnel tagging not enabled
           Tagging not operational
           MTU = 4470
           ATM tagging: Tag VPI = 1, Control VC = 0/32
   Interface ATM3/0/0:
           IP tagging enabled
           TSP Tunnel tagging not enabled
           Tagging not operational
           MTU = 4470
➔         ATM tagging: Tag VPI range = 5 - 6, Control VC = 6/32
   <Additional text omitted.>
```

Step 5 Check the IP tagging enabled field. This field indicates whether tag switching is enabled on this interface.

Step 6 Check the ATM tagging field. This field indicates the VPI range of the interface.

For detailed interface configuration information about tag switching, refer to the "Configuring Tag Switching" chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to confirm the loopback interface 0 configuration on the switch router:

Step 1 Use the **show interfaces loopback 0** command to confirm the loopback interface 0 configuration on the switch router.

```
   AdminFl1Ls1# show interfaces loopback 0
➔  Loopback0 is up, line protocol is up
     Hardware is Loopback
➔  Internet address is 2.2.2.2/24
     MTU 1500 bytes, BW 8000000 Kbit, DLY 5000 usec, rely 255/255, load 1/255
     Encapsulation LOOPBACK, loopback not set, keepalive set (10 sec)
     Last input 00:00:03, output never, output hang never
     Last clearing of "show interface" counters never
     Queueing strategy: fifo
     Output queue 0/0, 0 drops; input queue 0/75, 0 drops
     5 minute input rate 0 bits/sec, 0 packets/sec
     5 minute output rate 0 bits/sec, 0 packets/sec
        0 packets input, 0 bytes, 0 no buffer
        Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
        0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
        73 packets output, 0 bytes, 0 underruns
        0 output errors, 0 collisions, 0 interface resets
        0 output buffer failures, 0 output buffers swapped out
   AdminFl1Ls1#
```

Step 2 Check the Loopback 0 status field. It should be up.

Step 3 Check the line protocol field. It should be up.

Step 4 Check the Internet address field. It should display the IP address of the loopback interface on this switch router.

For detailed information, refer to the “Configuring Tag Switching” chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to confirm the OSPF configuration on the switch router:

Step 1 Use the **show ip ospf** command to confirm the OSPF configuration of the switch router.

```
AdminFl1Ls1# show ip ospf
➔ Routing Process “ospf 10000” with ID 150.0.0.0
 Supports only single TOS(TOS0) routes
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Number of DCbitless external LSA 0
 Number of DoNotAge external LSA 0
 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
    Area BACKBONE(0) (Inactive)
        Number of interfaces in this area is 3
        Area has no authentication
        SPF algorithm executed 2 times
        Area ranges are
        Link State Update Interval is 00:30:00 and due in 00:28:44
        Link State Age Interval is 00:20:00 and due in 00:18:44
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0

AdminFl1Ls1#
```

Step 2 Check the Routing Process field. The ospf field and ID fields should match the numbers configured. If they do not, refer to the “Configuring Tag Switching” chapter in the *ATM Switch Router Software Configuration Guide*.

Troubleshooting TDP Neighbors

This section describes how to troubleshoot TDP control channel VPI and virtual channel identifier (VCI).

Although not necessary for most configurations, you can change the default VPI and VCI of the TDP control channel if you want to use a nondefault value.

Note The default TDP control channel is on VPI 0 and VCI 32. TDP control channels exchange TDP hellos and PIEs to establish two-way TDP sessions. Tag virtual channels (TVCs) are created by the exchange of PIEs through TDP control channels.

Use the following command to check the tag switching TDP neighbor connections:

Command	Purpose
show tag-switching tdp neighbor	Confirms the tag switching TDP neighbor connection.

Follow these steps to check the tag switching TDP neighbor connections:

Step 1 Use the **show tag-switching tdp neighbor** command to confirm the tag switching TDP neighbor connections.

Step 2 Check the peer TDP identifier field. This field indicates the TDP identifier of the neighbor (peer device) for this session.

Step 3 Check the local TDP identifier field. This field indicates the TDP identifier for the local tag switching switch router or router for this session.

Step 4 Check the TCP connection field. This field indicates the TCP connection used to support the TDP session. The format for displaying the TCP connection is *peer IP address.peer port local IP address*.

Step 5 Check the PIEs sent/rcvd (sent or received) field. This field indicates the number of TDP PIEs sent to and received from the session peer device. The count includes the transmission and receipt of periodic keepalive PIEs, which are required for maintenance of the TDP session.

Step 6 Check the Up time field. This field indicates the length of time the TDP session has existed.

Follow these steps to confirm the VPI and VCI configuration of the tag switching interface on the switch router interface:

Step 1 Use the **show tag-switching interfaces atm** *card/subcard/port* **detail** command to confirm the configuration and connection of the tag switching interface VPI and VCI.

```
  AdminFl1Ls1# show tag-switching interfaces atm 0/0/1 detail
  Interface ATM0/0/1:
➔ IP tagging enabled
          TSP Tunnel tagging not enabled
➔ Tagging operational
          MTU = 8940
➔ ATM tagging: Tag VPI range = 2 - 5, Control VC = 6/32

  AdminFl1Ls1#
```

Step 2 Check the IP tagging field. This field shows whether the interface is configured to tag IP packets.

Step 3 Check the Tagging operational field. This field shows whether the packets are being tagged.

Step 4 Check the ATM tagging field. This field indicates the VPI range of the interface.

For detailed information, refer to the “Configuring Tag Switching” chapter in the *ATM Switch Router Software Configuration Guide*.

Troubleshooting Tag Switching on VP Tunnels

This section describes how to troubleshoot a tag switching connection configured on a VP tunnel.

For detailed information, refer to the "Configuring Tag Switching" chapter in the *ATM Switch Router Software Configuration Guide*.

To confirm VP tunnel configuration of tag switching, perform the following tasks in EXEC mode:

Command	Purpose
show atm vp	Confirms the VP tunnel configuration on an interface.
show tag-switching tsp-tunnels [*ip-address* | **all** | **head** | **middle** | **tail** | **remote**] [*interface-num*] [**brief**]	Confirms the TSP[1] tunnel status and configuration.

[1] TSP = tag switching path

Follow these steps to confirm the VP tunnel configuration of tag switching:

Step 1 Use the **show atm vp** command to confirm VP tunnel configuration.

```
EngFl1Ls1# show atm vp
Interface    VPI    Type  X-Interface     X-VPI     Status
ATM4/0/0     51      PVP     ATM1/1/0     101       UP
ATM1/1/0     101     PVP     ATM3/0/0     51        UP
EngFl1Ls1#
```

Step 2 Check the Status field. The PVP status should be UP. If it is not, check the VP tunnel configuration. Refer to "Configuring Tag Switching" chapter in the *ATM Switch Router Software Configuration Guide*.

Follow these steps to confirm the tag switching VP tunnel configuration:

Step 1 Use the **show tag-switching interfaces** command to confirm VP tunnel configuration on each router or switch router in the network. The following example starts at the head end:

```
EngFl1Rt1# show tag-switching tsp-tunnels
Signalling Summary:
            TSP Tunnels Process:            running
            RSVP Process:                   running
            Forwarding:                     enabled
TUNNEL ID               DESTINATION      STATUS           CONNECTION
10.106.0.6 0            10.2.0.12        up               up

EngFl1Rt1#
```

Step 2 Use the **show tag-switching tsp-tunnels** command to confirm VP tunnel configuration at the middle switch routers or routers:

```
AdminFl1Ls1# show tag-switching tsp-tunnels
Signalling Summary:
            TSP Tunnels Process:            running
            RSVP Process:                   running
            Forwarding:                     enabled
TUNNEL ID               DESTINATION      STATUS           CONNECTION
10.106.0.6 0            10.2.0.12        up               up

AdminFl1Ls1#
```

Step 3 Use the **show tag-switching tsp-tunnels** command to confirm VP tunnel configuration at the tail end switch router or router:

```
   AdminFl1Rt1# show tag-switching tsp-tunnels
   Signalling Summary:
➔              TSP Tunnels Process:            running
➔              RSVP Process:                   running
➔              Forwarding:                     enabled
   TUNNEL ID               DESTINATION      STATUS           CONNECTION
   10.106.0.6 0            10.2.0.12        up               up

   AdminFl1Rt1#
```

Step 4 Check whether the TSP Tunnels Process is running. If it is not, enter the **tag-switching tsp-tunnels** command to enable the process globally on the switch router or router.

Step 5 Check whether the RSVP Process is running. If it is not, enter the **tag-switching tsp-tunnels** command on the interfaces used by the tunnel to enable the process on the interface.

Step 6 If this is a router connection, check whether Forwarding is enabled on the router. If it is not, enter the **ip cef distributed switch** command or **ip cef switch** command to enable IP Cisco Express Forwarding (CEF) globally on the router.

Step 7 Use the **show tag-switching interfaces** command to check the VP tunnel interface configuration at each switch router or router in the tunnel. The following example starts at the head end:

```
EngFl1Rt1# show tag-switching interfaces
          Interface              IP     Tunnel   Operational
          ATM4/0/0               Yes    No       Yes
          ATM1/1/0               Yes    No       Yes
EngFl1Rt1#
```

Step 8 Use the **show tag-switching interfaces** command to check the VP tunnel interface configuration at the middle switch router or routers:

```
AdminFl1Ls1# show tag-switching interfaces
          Interface             IP    Tunnel   Operational
          ATM3/0/0              Yes   Yes      Yes
          ATM1/0/0              Yes   Yes      Yes
AdminFl1Ls1#
```

Step 9 Use the **show tag-switching interfaces** command to confirm VP tunnel configuration at the tail end switch router or router:

```
AdminFl1Rt1# show tag-switching interfaces
          Interface             IP    Tunnel   Operational
          ATM0/0                Yes   Yes      Yes
          Ethernet 2/3          Yes   Yes      Yes
AdminFl1Rt1#
```

Step 10 Check whether the interfaces used by the tunnel have "Yes" in the Tunnel column. If they do not, use the **tag-switching tsp-tunnels** command on the interfaces used by the tunnel to enable TSP tunnels, and refer to the "Configuring Tag Switching" chapter in the *ATM Switch Router Software Configuration Guide*.

Step 11 Verify that the interfaces used by the tunnel are operational. The interfaces should have "Yes" in the Operational column.

If not, check the interface configuration and refer to the "Configuring Tag Switching" chapter in the *ATM Switch Router Software Configuration Guide*.

Troubleshooting Tag Switching Using debug Commands

This section describes **debug** commands that you can use to troubleshoot tag switching connections on a switch router.

Use the following commands to debug tag switching connections on a switch router:

Command	Purpose
debug tag-switching adjacency	Debugs tag switching adjacency database events.
debug tag-switching atm-tdp {**api** | **routes** | **states**}	Debugs tag switching ATM Tag Distribution Protocol (TDP) events.
debug tag-switching packets {**atm** *card/subcard/port* | **atm-p** *card/subcard/port* | **cbr** *card/subcard/port* | **ethernet** *card/subcard/port* | **loopback 0** | **null**}	Debugs tag switching packets.
debug tag-switching tdp {**advertisements** | **bindings** | **directed-neighbors** | **pies** [**received** | **sent**] | **session** [**io** | **state**] | **transport** [**connections** | **events** | **timers**]}	Debugs TDP switching events.
debug tag-switching tfib {**cef** | **enc** | **state** | **struct** | **tsp**}	Debugs tag switching TFIB[1].

Command	Purpose
debug tag-switching traffic-eng {**events** | **interfaces** | **metrics** | **routing-table**}	Debugs tag switching traffic engineering.
debug tag-switching tsp-tunnels {**events** | **signalling** | **tagging**}	Debugs tag switching TSP tunnels.
no debug all	Turns off all debugging.

[1] TFIB = Tag Forwarding Information Base

For detailed interface configuration information, refer to the "Configuring Tag Switching" chapter in the *ATM Switch Router Software Configuration Guide*.

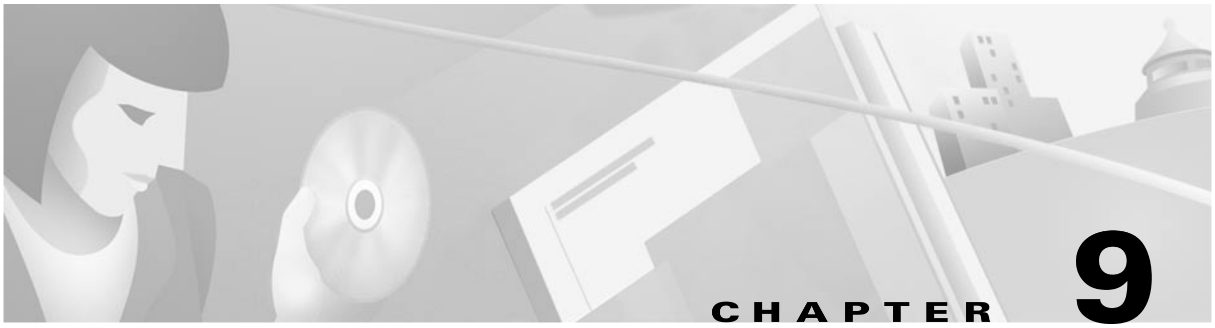

CHAPTER 9

Troubleshooting CES Connections and Network Clocking

This chapter provides troubleshooting information for connectivity problems in circuit emulation service (CES) environments and network clocking. For more information on CES, refer to the "Circuit Emulation Services and Voice over ATM" chapter in the *Guide to ATM Technology*, and "Configuring Circuit Emulation Services" chapter in the *ATM Switch Router Software Configuration Guide*.

Before you begin, make sure that all physical port connections are working correctly. For information on troubleshooting interfaces, refer to Chapter 5, "Troubleshooting Switch Router ATM Interface Connections."

This chapter contains the following sections:

- CES Overview, page 9-1
- CES Example Network, page 9-4
- Initial Troubleshooting of CES, page 9-4
- Troubleshooting CES Using debug Commands, page 9-7
- Troubleshooting Network Clocking, page 9-7

CES Overview

CES allows circuits to be transparently extended across an ATM network. CES is typically used to transport voice or video between the ATM switch router and non-ATM telephony devices, such as PBXs and TDMs, or video devices. Voice and video, unlike data traffic, are very sensitive to delay and delay variance. CES uses constant bit rate (CBR) virtual circuits (VCs), which guarantees acceptable delay and delay variation and thus satisfies the requirements of voice and video traffic.

Performing Basic Checks

This procedure outlines the steps for performing basic interface checks of the CES circuit configuration. Always check the following when a CES circuit fails to function:

- For hard PVCs, do the VPI and VCI numbers match those assigned by the service provider?
- For shaped VP tunnels, is the configured transmission rate within the range contracted with the service provider?

At the ATM switch router, use the following commands to check the CES configuration:

Command	Purpose
show ces circuit interface cbr *card*/*subcard*/*port circuit-id*	Confirms the configuration on the CES interface.
show atm vp interface atm *card*/*subcard*/*port vpi*	Confirms the configuration of the shaped VP tunnel.
show dcu leds	Confirms the status of the CES port LEDs.

Follow these steps to troubleshoot the CES configuration:

Step 1 Use the **show ces circuit** command to display the VPI and VCI configuration:

```
Switch# show ces circuit interface cbr 3/1/0 1
Circuit: Name CBR3/1/0:1, Circuit-state ADMIN_UP / oper-state UP Interface CBR3
Port Clocking network-derived, aal1 Clocking Method CESIWF_AAL1_CLOCK_SYNC
Channel in use on this port: 1
Channels used by this circuit: 1
Cell-Rate: 172, Bit-Rate 64000
cas OFF, cell_header 0x100 (vci = 16)
Configured CDV 2000 usecs, Measured CDV unavailable
De-jitter: UnderFlow unavailable, OverFlow unavaliable
ErrTolerance 8, idleCircuitdetect OFF, onHookIdleCode 0x0
state: VcLoc, maxQueueDepth       81, startDequeueDepth       64
Partial Fill:       47, Structured Data Transfer 1
HardPVC
src: CBR3/1/0 vpi 0, vci 16
➔ Dst: ATM1/1/1 vpi 1, vci 101
```

Step 2 Check the Dst field to confirm that the VPI and VCI values match those assigned by the service provider. If not, reconfigure the CBR interface using the **ces pvc** command.

Step 3 Use the **show atm vp** command to display the connection traffic table index configuration for the shaped VP tunnel:

```
  Switch# show atm vp interface atm 1/1/1 1

  Interface: ATM1/1/1, Type: oc3suni
  VPI = 1
  Status: SHAPED TUNNEL
  Time-since-last-status-change: 13:59:23
  Connection-type: PVP
  Cast-type: point-to-point
  Usage-Parameter-Control (UPC): pass
  Wrr weight: 2
  Number of OAM-configured connections: 0
  OAM-configuration: disabled
  OAM-states:  Not-applicable
  Threshold Group: 1, Cells queued: 0
  Rx cells: 0, Tx cells: 0
  Tx Clp0:0,  Tx Clp1: 0
  Rx Clp0:0,  Rx Clp1: 0
  Rx Upc Violations:0, Rx cell drops:0
  Rx Clp0 q full drops:0, Rx Clp1 qthresh drops:0
→ Rx connection-traffic-table-index: 10
→ Rx service-category: CBR (Constant Bit Rate)
→ Rx pcr-clp01: 4000
  Rx scr-clp01: none
  Rx mcr-clp01: none
  Rx      cdvt: 1024 (from default for interface)
  Rx       mbs: none
→ Tx connection-traffic-table-index: 10
→ Tx service-category: CBR (Constant Bit Rate)
→ Tx pcr-clp01: 4000
  Tx scr-clp01: none
  Tx mcr-clp01: none
  Tx      cdvt: none
  Tx       mbs: none
```

Step 4 Check the Rx service-category and Tx service-category fields for the CBR service.

Step 5 Check the Rx pcr-clp01 and Tx pcr-clp01 fields to ensure that the peak cell rate (PCR) is within the range contracted with the service provider.

Step 6 Use the **show dcu leds** command to display the status of the CES port LEDs:

```
NewLs1010# show dcu leds
CBR3/1/0 [20]:  idle
CBR3/1/1 [21]:  idle
CBR3/1/2 [22]:  Red (loss of signal and loss of cells)
CBR3/1/3 [23]:  Red (loss of cells)
```

Step 7 If the port LED status is RED, do the following:

- Check the cable for damage.
- Check the length of the cable. It should not be more than 1000 feet or 304.8 meters long.
- Check the interface configuration.

For detailed interface configuration information about CES, refer to the "Configuring Circuit Emulation Services" chapter in the *ATM Switch Router Software Configuration Guide*.

CES Example Network

In the example network in Figure 9-1, the ATM switch routers in the administration building and the remote sales building are each connected to a PBX and the WAN:

- AdminFl1Ls1—ATM switch router with CES interface located in the administration building
- SalesFl1Ls1—ATM switch router with CES interface located in the remote sales building

Figure 9-1 CES Example Network

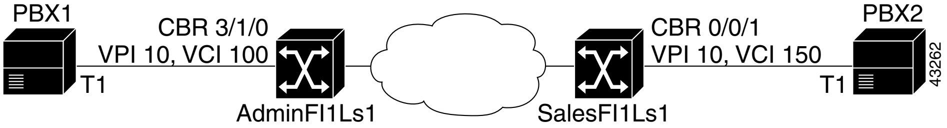

This network example is used to describe the CES troubleshooting examples in the rest of this chapter.

For detailed configuration information about CES, refer to the "Configuring Circuit Emulation Services" chapter in the *ATM Switch Router Software Configuration Guide*.

Initial Troubleshooting of CES

This section describes initial troubleshooting steps that you should perform when beginning to troubleshoot a CES connection.

At the ATM switch router, use the following commands to check the CES configuration:

Command	Purpose
show ces status	Confirms the status of the CES circuits.
show ces circuit	Confirms the configuration of the CES PVCs.
show ces circuit interface cbr *card***/***subcard***/***port circuit-id*	Confirms the configuration on the CES interface.
show interfaces cbr *card***/***subcard***/***port*	Confirms the status of the CES interface.

Checking the CES Circuit Status

Use the following command to confirm that the configured CES circuit is up:

Step 1 Use the **show ces status** command to check the status of the CES circuit.

```
AdminFl1Ls1# show ces status
   Interface      IF      Admin         Port  Channels in
     Name       Status   Status         Type      use
------------- -------- --------- ----------- -----------
     CBR3/1/0       UP        UP          T1 1-3,7,20-22,24
     CBR3/1/1     DOWN        UP          T1 1-24
     CBR3/1/2     DOWN        UP          T1 24
     CBR3/1/3       UP        UP          T1 10-13
AdminFl1Ls1#
```

Step 2 Check the IF and Admin Status fields to confirm that they are up.

If the interface is down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

For detailed interface configuration information about CES, refer to the "Configuring Circuit Emulation Services" chapter in the *ATM Switch Router Software Configuration Guide*.

Checking CES Circuit Configuration

Follow these steps to confirm that CES circuits are configured correctly:

Step 1 Use the **show ces interface** command to confirm the configuration of the circuit:

```
  Switch# show ces circuit interface cbr 3/1/0 0
➔ Circuit: Name CBR-PVC-A, Circuit-state ADMIN_UP / oper-state UP Interface CBR3/
➔ Port Clocking network-derived, aal1 Clocking Method CESIWF_AAL1_CLOCK_SYNC
  Channel in use on this port: 1-31
  Channels used by this circuit: 1-31
  Cell-Rate: 5447, Bit-Rate 2048000
  cas OFF, cell_header 0x100 (vci = 16)
  Configured CDV 2000 usecs, Measured CDV 1769 usecs
➔ De-jitter: UnderFlow 42, OverFlow 0
  ErrTolerance 8, idleCircuitdetect OFF, onHookIdleCode 0x0
  state: VcAlarm, maxQueueDepth      823, startDequeueDepth      435
  Partial Fill:       47, Structured Data Transfer 0
  Passive SoftVC
➔ Src: atm addr 47.0091.8100.0000.00e0.f75d.0401.4000.0c81.9030.10 vpi 0, vci 16
  Switch#
```

Step 2 Check the oper-state field to confirm that it is up.

If down, verify that the aal1 Clocking Method field for each end of the circuit shows the same configuration.

Step 3 Check the Underflow and Overflow fields to confirm that the network clocking is synchronized.

Buffer overflows and underflows indicate a slight clocking difference between the devices. Buffer overflows occur when the transmitting device is faster than the receiving device; such a condition results in frame drops. Buffer underflows occur when the transmitting device is slower than the receiving device; such a condition results in frame resends. Check with your service provider to reduce the cell delay variation (CVD).

Step 4 Check the Src and Dst fields for the correct addresses.

Use the **show atm status**, **show ces status**, **show atm address**, and **show ces address** commands to confirm the source and destination address configuration.

For detailed interface configuration information about CES, refer to the "Configuring Circuit Emulation Services" chapter in the *ATM Switch Router Software Configuration Guide*. For detailed physical interface troubleshooting information, see the "Troubleshooting CES T1 and CES E1 Interfaces" section on page 5-32.

Follow these steps to confirm the CES interface configuration on the ATM switch router:

Step 1 Use the **show ces circuit** command to confirm the connection of the CES interfaces.

```
Switch# show ces circuit
Interface  Circuit  Circuit-Type     X-interface   X-vpi   X-vci Status
 CBR3/1/0    0       Passive SoftVC  ATM-P3/1/3      0      3088  UP
 CBR3/1/1    0       Passive SoftVC  ATM-P3/1/3      0      2064  UP
 CBR3/1/2    0       Active SoftVC   ATM-P3/1/3      0      1040  UP
 CBR3/1/3    0       Active SoftVC   ATM-P3/1/3      0        16  UP
Switch#
```

Step 2 Check the Circuit-Type field. It should contain the correct type for the circuit.

Step 3 Check the X-interface field. It should contain the correct destination interface for the circuit.

Step 4 Check the Status field. It should read "UP."

If the rest of the fields for the interface are correct but the status is "DOWN," then check the X-interface status using the **show interfaces** command. If the interface is administratively down, use the **no shutdown** command to reenable it.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

For detailed interface configuration information about CES, refer to the "Configuring Circuit Emulation Services" chapter in the *ATM Switch Router Software Configuration Guide*.

Troubleshooting CES Using debug Commands

This section describes **debug** commands that you can use to troubleshoot CES circuits on an ATM switch router.

Use the following commands to debug CES connections on an ATM switch router:

Command	Purpose
debug ces-iwf connection	Debugs CES circuit connection events.
debug ces-iwf createloc cbr *card/subcard/port* [*vpi*] {**on** | **off**}	Enables cell loss for debugging purposes.
debug ces-iwf dcu	Debugs CBR-DCU internal events.
debug ces-iwf internal	Debugs CES internal events.
no debug all	Turns off all debugging.

For detailed interface configuration information, refer to the “Configuring Circuit Emulation Services” chapter in the *ATM Switch Router Software Configuration Guide*.

Troubleshooting Network Clocking

This section describes how to troubleshoot problems with the network clocking configuration. For more information on network clocking, refer to the “Network Clock Synchronization” chapter in the *Guide to ATM Technology* and the “Initially Configuring the ATM Switch Router” chapter in the *ATM Switch Router Software Configuration Guide*.

Overview of Network Clocking

Network clocking facilities generate or derive a clock signal and distribute it throughout a network to ensure synchronized network operation. This is important in delay-sensitive data types, such as voice and video, because these types of data must be received and transmitted at the same rate at every step, or hop, in a connection. If network synchronization is lost, data might be lost due to buffer overflow or underflow; cyclic redundancy check (CRC) errors might also occur.

Table 9-1 provides a summary of network clocking features.

Table 9-1 Network Clocking Feature Summary

Platform	Up/Down Detection	Loss of Synchronization Detection	Phase Adjustment Cutover	Stratum 3 Clock	BITS[1] Port	Clock Source Preference
Catalyst 8540 MSR with network clock module	Yes	Yes	Yes	Yes	Yes	Best
Catalyst 8510 MSR	Yes	Yes	Yes	No	No	Medium
LightStream 1010 with FC-PFQ	Yes	Yes	Yes	No	No	Medium
Catalyst 8540 MSR without network clock module	Yes	No	No	No	No	Poor
LightStream 1010 with FC-PCQ	Yes	No	No	No	No	Poor

[1] BITS = Building Integrated Timing Supply

Network Clock Module LEDs

The network clock module faceplate LEDs provide status information for the BITS ports and the alarm port. The LEDs are described in Table 9-2.

Table 9-2 Network Clock Module LED Descriptions

LED	Status	Description
POWER	Green Off	The switch is powered on and the processor is functioning. The switch is powered off.
STATUS	Green Red Orange	The clock module is the primary network clock source. The processor has crashed. The clock module is operating in standby mode.
MAJOR ALARM	Red Off	A major alarm condition has occurred. No major alarm reported.
MINOR ALARM	Red Off	A minor alarm condition has occurred. No minor alarm reported.
CRITICAL ALARM	Red Off	Not supported. No critical alarm reported.

The following are major alarm conditions:

- A switchover from the primary clock source to the default clock source occurred.
- A switchover from the secondary clock source to the default clock source occurred.
- A loss of all references to the network clock source occurred while the network clock source was set to free running.

- A route processor failure caused the network clock module to fail.
- The network clock module is in holdover mode.

The following are minor alarm conditions:

- A switchover from the primary clock source to the secondary clock source occurred.
- A loss of a single reference to the network clock source occurred while the network clock source was set to free running.

Checking the Network Clock Source Configuration

Use the following commands to troubleshoot the network clock source configuration:

Command	Purpose
show network-clocks	Confirms the network clock configuration.
debug ports netclock	Debugs network clock events.
no debug all	Turns off all debugging.

Follow these steps to troubleshoot the network clocking configuration on an ATM switch that does not have the network clock module installed:

Step 1 The network clock type configured on the switch must match the current clock source. Use the **show network-clocks** command to display the network clocking configuration.

```
  Switch# show network-clocks
➔ clock configuration is NON-Revertive
➔ Priority 1 clock source: ATM0/0/0(down)
  Priority 2 clock source: No clock
  Priority 3 clock source: No clock
  Priority 4 clock source: No clock
  Priority 5 clock source: System clock
➔ Current clock source:System clock, priority:5
```

Step 2 Check the clock configuration field to confirm that the clock switchover mode is configured correctly. If it is not, use the **network-clock-select** command to correctly configure the clock switchover mode.

Step 3 Check the clock source fields to confirm that the desired clock sources are configured. Use the **network-clock-select** command to configure the clock sources. See Table 9-1 for a list of network clock source features.

Step 4 Check the status of the clock sources. If the clock source is listed as “down,” check the interface status using the procedure described in Chapter 5, “Troubleshooting Switch Router ATM Interface Connections.”

Note Once the clock problem is solved, reconfigure the network clock using the **network-clock-select** command to make sure that the primary clock becomes the present clock source. The **no shutdown** command does not affect the network clock source status.

Follow these steps to troubleshoot the network clocking configuration on a switch that uses the network clock module:

Step 1 Use the **show network-clocks** command to display the network clocking configuration.

```
   Switch# show network-clocks
   Network clocking information:
   ---------------------------------------
➔  Source switchover mode:    non-revertive
   Netclkd state:             Active
   Source selection method:   provisioned
   NCLKM hardware status:     installed & usable
   NCLKM status:              software enabled
➔  Primary   clock source:    ATM11/1/2 Unlockable
   Secondary clock source:    ATM12/0/0 one shot triggered
➔  Present   clock source:    ATM12/0/0 (2) Locked
```

Step 2 Check the Source switchover mode field to confirm that the clock switchover mode is configured correctly. If it is not, use the **network-clock-select** command to correctly configure the network clocking mode.

Step 3 Check the clock source fields to confirm that the desired clock sources are configured. Use the **network-clock-select** command to configure the clock sources. See Table 9-1 for a list of network clock source features.

Step 4 Check the status of the primary and secondary clock sources. The status follows the clock source name.

- If the status is "Unlockable," check the network clock configuration on all the switches in the network for possible clock loops. A clock loop occurs when two switches derive clocking information from the same interface.
- If the status is "one shot triggered," the clock source switchover mode is non-revertive and the primary clock source has gone down once. When the primary clock source comes back up, it does not become the present clock source. Correct this situation by manually setting the primary clock source as the present clock source, or the source switchover mode to nonrevertive, by using the **network-clock-select** command.
- If the status is blank, compare the Primary clock source field with the Present clock source field. If the present clock source does not match the primary clock source, check the interface status using the procedures described in Chapter 5, "Troubleshooting Switch Router ATM Interface Connections."

Note Once the clock problem is resolved, reconfigure the network clock using the **network-clock-select** command to make sure that the primary clock becomes the present clock source. The **no shutdown** command does not affect the network clock source status.

Follow these steps to troubleshoot the network clocking configuration on a switch that uses the Building Integrated Timing Supply (BITS) port on the network clock module.

Step 1 Use the **debug ports netclock** and **show network-clocks** commands to display the configuration and status of the BITS ports.

```
  Switch# debug ports netclock
  Rhino network clocks debugging is on
  Switch# show network-clocks
  Network clocking information:
  ---------------------------------------
  Source switchover mode:   non-revertive
  Netclkd state:            Active
  Source selection method:  provisioned
  NCLKM hardware status:    installed & usable
  NCLKM status:             software enabled
➔ Primary   clock source:   BITS 0 in T1 mode
  Secondary clock source:   ATM11/1/2  one shot triggered
  Present   clock source:   ATM11/1/2 (2) Locked
➔ bits 0 state             :down (32827388)
  bits 0 admin state       :up
➔ bits 1 state             :down (32827388)
  bits 1 admin state       :up
  do_not_switch  flag      :0
  other_holdover flag      :0
  p_one_shot     flag      :0
  s_one_shot     flag      :0
  p_l_state                :Reset
  s_l_state                :Reset
  other_priority           :0
  other_type               :0
  ncdp in use              :0
  Hello        tx seq no   :0
  Hello        rx seq no   :0
  Clock update tx seq no   :0
  Clock update rx seq no   :0
  Hello timer              :running:1,
  Ref lock timer           :running:1, time_left:580
  <information deleted>
```

Step 2 Check the state field for the BITS port being used. If it is down, check the cable for damage.

Step 3 Check the BITS port LED. If it is not on, check the interface mode type. The mode type is either T1 or E1. T1 is the default mode. Use the **network-clock-select** command to change the BITS interface mode type.

For detailed interface configuration information on network clocking, refer to the "Initially Configuring the ATM Switch Router," and Chapter 18, "Configuring Circuit Emulation Services" chapter in the *ATM Switch Router Software Configuration Guide*.

Checking the CES Interface Clocking Configuration

The clocking configuration of a CES interface might affect the traffic flow on the circuit. If quality of the transmission has degraded, follow these steps to troubleshoot the CES interface clocking for unstructured services:

Step 1 Check the Underflow and Overflow fields. Use the **show ces interface** command to display the network clocking configuration:

```
Switch# show ces interface cbr 3/1/0
Interface:      CBR3/1/0        Port-type:E1-120ohms-DCU
IF Status:      UP              Admin Status: UP
Channels in use on this port: 1
LineType: E1_LT        LineCoding: HDB3  LoopConfig: Payload
SignalMode: NoSignalling   XmtClockSrc: network-derived
➔ DataFormat: Structured     AAL1 Clocking Mode: Synchronous  LineLength: 330_440
e1InternationalBits 0x3, e1NationalBits 0x1F, e1MultiFrameBits 0xB
LineState:  RcvAIS LoopbackState
Errors in the Current Interval:
  PCVs        0 LCVs        0   ESs         0   SESs        0   SEFSs        0
  UASs        0 CSSs        0   LESs        0   BESs        0   DMs          0
Errors in the last 24Hrs:
  PCVs        0 LCVs        0   ESs         0   SESs        0   SEFSs        0
  UASs        0 CSSs        0   LESs        0   BESs        0   DMs          0
Input  Counters: 0 cells, 0 bytes
Output Counters: 0 cells, 0 bytes
```

Step 2 Check the AAL1 Clocking Mode field. Use the **ces aal1 clock** command to modify the clocking mode to **adaptive**. Adaptive clocking does not require an external clock source. If the problem ceases when the clocking mode is changed from synchronous to SRTS, the reference clock is the problem.

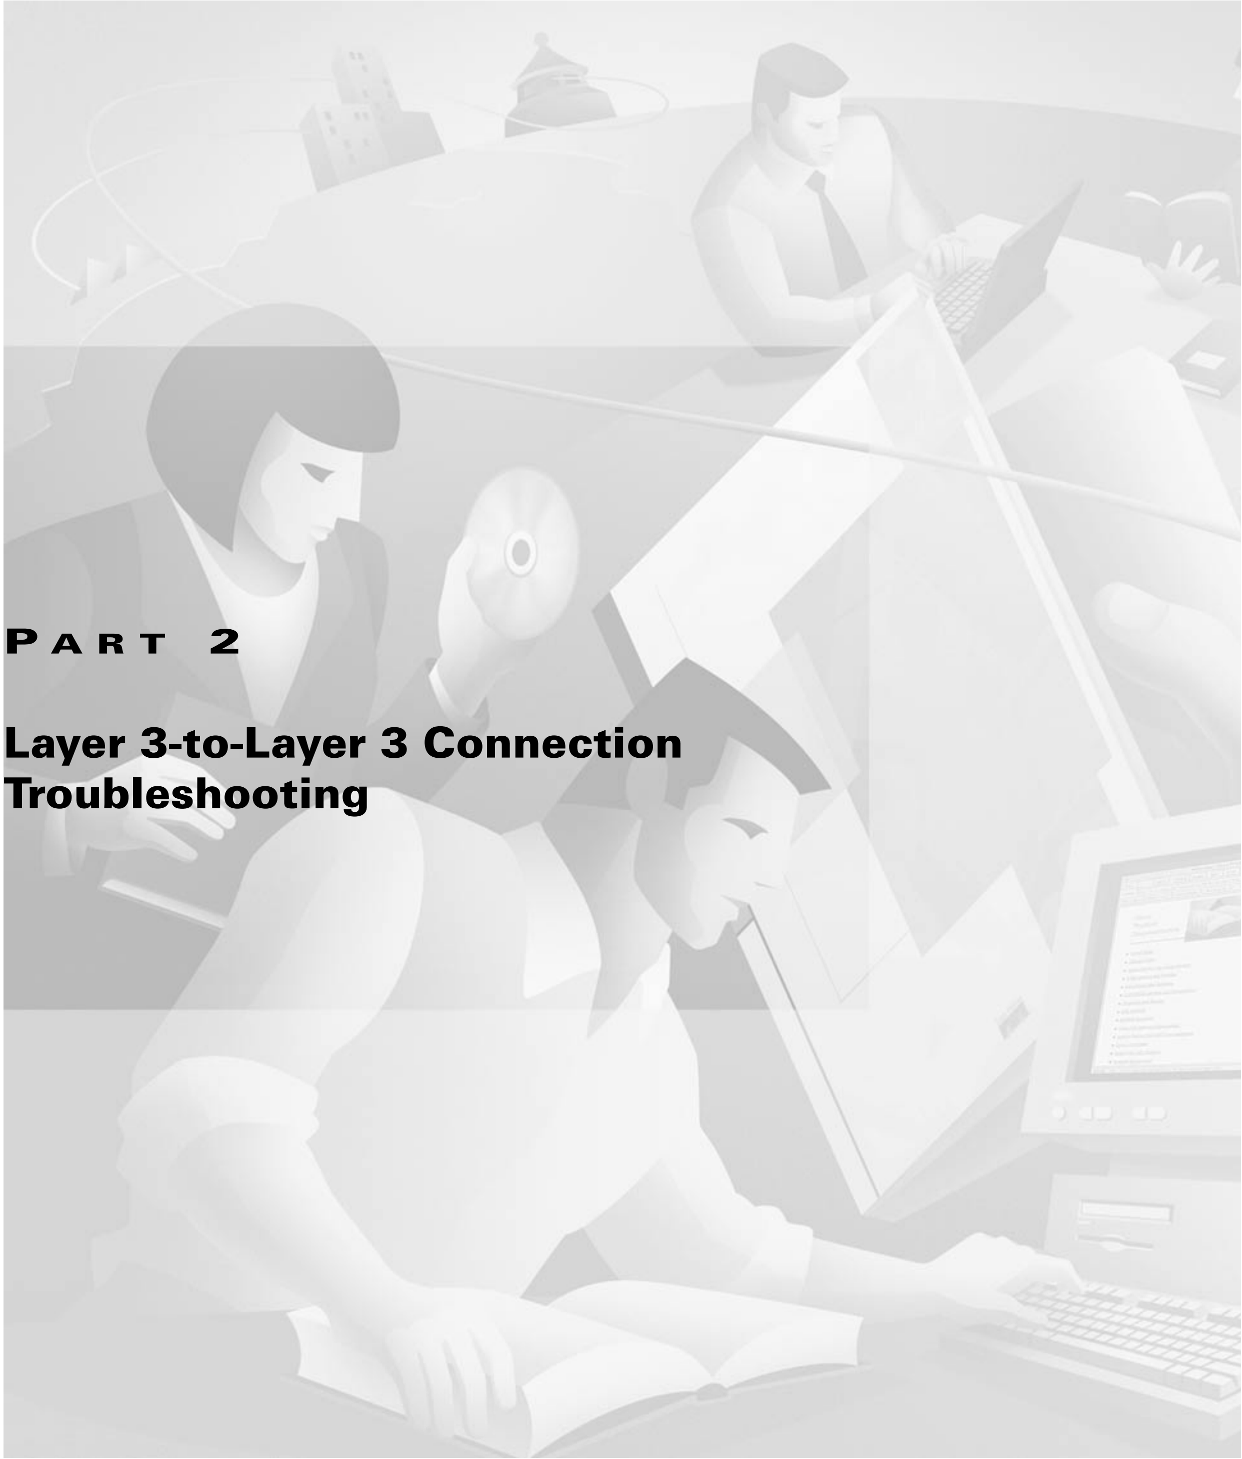

PART 2

Layer 3-to-Layer 3 Connection Troubleshooting

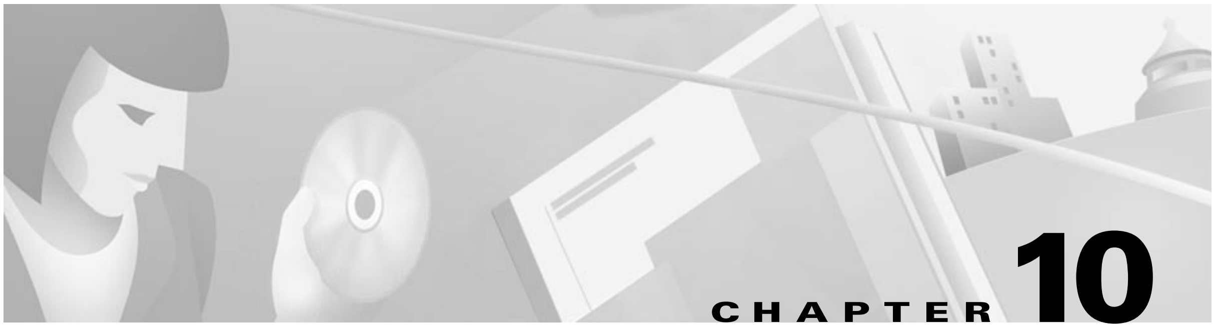

CHAPTER 10

Troubleshooting Ethernet, ATM Uplink, and POS Uplink Interfaces

This chapter provides troubleshooting information about connectivity and performance problems in the Ethernet, ATM uplink, and POS uplink physical interfaces of a Layer 3 enabled ATM switch router.

The chapter includes the following sections:

- Troubleshooting General Ethernet Interface Problems, page 10-1
- Troubleshooting 10/100 Ethernet Interface Modules, page 10-8
- Troubleshooting Gigabit Ethernet Interface Modules, page 10-15
- Troubleshooting ATM Uplink with Enhanced Gigabit Ethernet Interface Modules, page 10-34
- Troubleshooting Packet-over-SONET Uplink with Enhanced Gigabit Ethernet Interface Modules, page 10-41
- Troubleshooting ACL Daughter Card, page 10-46

Note For detailed cabling and hardware information for each port adapter, refer to the *Catalyst 8540 CSR Route Processor and Interface Module Installation Guide.*

Troubleshooting General Ethernet Interface Problems

You might see problems of cell transmission through the switch router, detected by a buildup of cells on an internal virtual channel (VC). These problems occur for the following reasons:

- Troubleshooting Switch Card Failures, page 10-2
- Troubleshooting Port Stuck Problems, page 10-4
- Connectivity Troubleshooting Commands, page 10-8

Troubleshooting Switch Card Failures

A switch processor can have a cell stuck problem in internal virtual channels (VCs), resulting from timing issues in the hardware and software on the Catalyst 8540 CSR. You might see more than one port affected on one or more interface modules. Online insertion and removal (OIR) of the interface module will temporarily fix the problem.

Follow these steps to troubleshoot cell stuck problems:

Step 1 Under the lightest possible traffic, issue the **show switch fabric** command on the switch router to clear the counters.

```
Switch# show switch fabric

<Information deleted>

MMC Switch Fabric (idb=0x61DD8F0C)

  Key: Rej. Cells  - # cells rejected due to lack of resources
                            or policing (16-bit)
       Inv. Cells    - # good cells that came in on a non-existent conn.
       Mem Buffs     - # cell buffers currently in use
       RX Cells      - # rx cells (16-bit)
       TX Cells      - # tx cells (16-bit)
       Rx HEC        - # cells Received with HEC errors
       Tx PERR       - # cells with memory parity errors

  MSC#     Rej. Cells    Inv. Cells    Mem. Buffs    Rx Cells     Tx Cells     Rx HEC         Tx PErr
 -----    -----------   ------------   -----------  -----------  ----------   ----------   ----------
MSC 0:             0         110018          0           0           0           0           0
MSC 1:             0         231044          0           0           0           0           0
MSC 2:             0         234283          0           0           0           0           0
MSC 3:             0         232492          0           0           0           0           0
MSC 4:             0         242004          0           0           0           0           0
MSC 5:             0         120995        345           0           0           0           0
MSC 6:             0         111466          0           0           0           0           0
MSC 7:             0         334398          0           0           0           0           0

Switch Fabric Statistics

➔    Rejected Cells: 0
➔    Invalid Cells: 1616700
     Memory Buffers: 345
     Rx Cells: 0
     Tx Cells: 0
     RHEC: 0
     TPE: 0

<Information deleted>
```

The **show switch fabric** command clears the counters after it displays. Entering the command again shows the current activity on the switch router.

Step 2 Issue the **show switch fabric** command again to show new activity.

```
Switch# show switch fabric

<Information deleted>

MMC Switch Fabric (idb=0x60CF1788)
  Key: Rej. Cells    - # cells rejected due to lack of resources
                            or policing (16-bit)
       Inv. Cells    - # good cells that came in on a non-existent conn.
       Mem Buffs     - # cell buffers currently in use
       RX Cells      - # rx cells (16-bit)
       TX Cells      - # tx cells (16-bit)
       Rx HEC        - # cells Received with HEC errors
       Tx PERR       - # cells with memory parity errors

  MSC#     Rej. Cells    Inv. Cells    Mem. Buffs    Rx Cells     Tx Cells     Rx HEC         Tx PErr
 -----    -----------   ------------   -----------  -----------  ----------    ----------   ----------
MSC 0:          2189               6       14177           0           0            0           0
MSC 1:             0              36        2070           0           0            0           0
MSC 2:             0               0           0           0           0            0           0
MSC 3:             0               0           0           0           0            0           0
MSC 4:             0               0           0           0           0            0           0
MSC 5:             0               0           0           0           0            0           0
MSC 6:             0               6        1351           0           0            0           0
MSC 7:             0              10        1280           0           0            0           0

Switch Fabric Statistics

➔    Rejected Cells: 2189
➔    Invalid Cells: 58
     Memory Buffers: 18878
     Rx Cells: 0
     Tx Cells: 0
     RHEC: 0
     TPE: 0
<Information deleted>
```

Look at the values in the Rejected Cells and Invalid Cells fields. Note that the Rejected Cells and Invalid Cells field counters are increasing. This means there might be a problem in the switch fabric.

Step 3 Verify that no ports are involved by issuing the **show epc queuing** and **show epc status** commands.

```
Switch# show epc queuing
INT         X-INT       VCI   QCNT    VCI   QCNT

Switch# show epc status
Status of GigabitEthernet0/0/0: OK
Status of GigabitEthernet0/0/1: OK
Status of GigabitEthernet1/0/0: OK
Status of GigabitEthernet1/0/1: OK
Status of GigabitEthernet2/0/0: OK
Status of GigabitEthernet2/0/1: OK
Status of GigabitEthernet3/0/0: OK
Status of GigabitEthernet3/0/1: OK
Status of GigabitEthernet9/0/0: OK
Status of GigabitEthernet9/0/1: OK
Status of GigabitEthernet10/0/0: OK
Status of GigabitEthernet10/0/1: OK
Status of GigabitEthernet11/0/0: OK
Status of GigabitEthernet11/0/1: OK
Status of GigabitEthernet12/0/0: OK
Status of GigabitEthernet12/0/1: OK
```

If the queues are empty and all of the ports show OK status, then the problem is not the ports, it is the switch processor.

You can resolve this problem by upgrading your system software image to release Cisco IOS Release version 12.0(4a)WX5(11) or later, by replacing the switch processors, or by doing both.

Troubleshooting Port Stuck Problems

If one or more Fast Ethernet or Gigabit Ethernet ports are not transmitting cells, then the failure might be a port stuck problem.

Follow these steps to troubleshoot a port stuck problem:

Step 1 Use the **show switch fabric** command to display the activity in the switch processors.

Note Be sure to use the **show switch fabric** command during the lightest possible traffic conditions because actual traffic might be using the memory buffers.

```
Switch# show switch fabric

<Information deleted>

MMC Switch Fabric (idb=0x60CF1788)
  Key: Rej. Cells    - # cells rejected due to lack of resources
                            or policing (16-bit)
       Inv. Cells    - # good cells that came in on a non-existent conn.
       Mem Buffs     - # cell buffers currently in use
       RX Cells      - # rx cells (16-bit)
       TX Cells      - # tx cells (16-bit)
       Rx HEC        - # cells Received with HEC errors
       Tx PERR       - # cells with memory parity errors

  MSC#     Rej. Cells    Inv. Cells    Mem. Buffs    Rx Cells     Tx Cells     Rx HEC         Tx PErr
 -----    -----------   ------------   -----------  -----------  ----------   ----------   ----------
MSC 0:        389023            7896       14177           0           0           0           0
MSC 1:             0           32709        2070           0           0           0           0
MSC 2:             0               0           0           0           0           0           0
MSC 3:             0               0           0           0           0           0           0
MSC 4:             0               0           0           0           0           0           0
MSC 5:             0               0           0           0           0           0           0
MSC 6:             0            6170        1351           0           0           0           0
MSC 7:             0            9624        1280           0           0           0           0

Switch Fabric Statistics

➔    Rejected Cells: 389023
➔    Invalid Cells: 56399
     Memory Buffers: 18878
     Rx Cells: 0
     Tx Cells: 0
     RHEC: 0
     TPE: 0
```

The **show switch fabric** command clears the counters after it displays. Entering the command again shows the current activity on the switch router.

Step 2 Enter the **show switch fabric** command again.

```
Switch# show switch fabric

<Information deleted>

MMC Switch Fabric (idb=0x60CF1788)
  Key: Rej. Cells    - # cells rejected due to lack of resources
                            or policing (16-bit)
       Inv. Cells    - # good cells that came in on a non-existent conn.
       Mem Buffs     - # cell buffers currently in use
       RX Cells      - # rx cells (16-bit)
       TX Cells      - # tx cells (16-bit)
       Rx HEC        - # cells Received with HEC errors
       Tx PERR       - # cells with memory parity errors

  MSC#     Rej. Cells    Inv. Cells    Mem. Buffs    Rx Cells     Tx Cells     Rx HEC         Tx PErr
 -----    -----------   ------------   -----------  -----------  ----------   ----------   ----------

MSC 0:          2189               6       14177          0            0           0           0
MSC 1:             0              36        2070          0            0           0           0
MSC 2:             0               0           0          0            0           0           0
MSC 3:             0               0           0          0            0           0           0
MSC 4:             0               0           0          0            0           0           0
MSC 5:             0               0           0          0            0           0           0
MSC 6:             0               6        1351          0            0           0           0
MSC 7:             0              10        1280          0            0           0           0

Switch Fabric Statistics

➔    Rejected Cells: 2189
➔    Invalid Cells: 58
     Memory Buffers: 18878
     Rx Cells: 0
     Tx Cells: 0
     RHEC: 0
     TPE: 0
```

Look at the values in the Rejected Cells and Invalid Cells fields. Note that the Rejected Cells and Invalid Cells field counters are increasing. This means there might be a problem in the switch fabric.

Step 3 Use the **show epc queuing** and **show epc status** command to display interface queues and status.

```
Switch# show epc queuing
INT         X-INT       VCI   QCNT    VCI   QCNT
Gi0/0/0     Gi1/0/0      67    640     62      0
Gi0/0/0     Gi1/0/0      71    546     66      0
Gi0/0/1     Gi1/0/0      67    135    147      0
Gi0/0/1     Gi1/0/0      69     18    149      0
Gi1/0/0     SRP          35      0    342   1791
Gi1/0/0     Gi0/0/0      62      0     67    640
Gi1/0/0     Gi0/0/0      66      0     71    546
Gi1/0/0     Gi0/0/1     147      0     67    135
Gi1/0/0     Gi0/0/1     149      0     69     18
Gi1/0/0     Gi1/0/1     152      0     67    639
Gi1/0/0     Gi12/0/0    577      0     67    640
Gi1/0/0     Gi12/0/0    578      0     68     16
Gi1/0/0     Gi12/0/0    579      0     69     38
Gi1/0/0     Gi12/0/0    580      0     70     16
Gi1/0/0     Gi12/0/1    662      0     67    640
Gi1/0/0     Gi12/0/1    666      0     71    640
Gi1/0/1     Gi1/0/0      67    639    152      0
Gi12/0/0    Gi1/0/0      67    640    577      0
Gi12/0/0    Gi1/0/0      68     16    578      0
Gi12/0/0    Gi1/0/0      69     38    579      0
Gi12/0/0    Gi1/0/0      70     16    580      0
Gi12/0/1    Gi1/0/0      67    640    662      0
Gi12/0/1    Gi1/0/0      71    640    666      0

Switch# show epc status
Status of GigabitEthernet0/0/0: OK
Status of GigabitEthernet0/0/1: OK
➔ Status of GigabitEthernet1/0/0: not OK
Status of GigabitEthernet1/0/1: OK
Status of GigabitEthernet2/0/0: OK
Status of GigabitEthernet2/0/1: OK
Status of GigabitEthernet12/0/0: OK
Status of GigabitEthernet12/0/1: OK
```

The **show epc queuing** command output shows that no activity is going across interface GigabitEthernet 1/0/0. This is verified in the **show epc status** command output, which indicates that interface GigabitEthernet 1/0/0 is “not OK.” You have confirmed that the problem is a stuck port.

Note You might see a few cells in the QCNT column in the **show epc queuing** command output. That is normal. Issue the command several times to verify that traffic is moving through the queues. If the QCNT column values are incrementing and incrementing for the VCIs belonging to a particular interface, the problem is probably a stuck port.

You can remedy the port stuck condition by removing and reinserting the interface module. A **shutdown/no shutdown** command sequence on the problem interface does not resolve the problem.

Configuring Automatic Port Stuck Failure Recovery

To recover from a port stuck failure, perform the following tasks:

- Detect port stuck failure.
- Isolate the cell stuck failure.
- If it is only a port stuck failure, isolate the port from the other functional ports.
- Depending on the configuration option for reset of the stuck port, one of the following actions will occur:
  - Default behavior

    If the switch router is not configured to reset the port upon detecting a port stuck failure, the port will be isolated, thus preserving the integrity of the switch router.
  - Nondefault behavior

    If the switch router is configured to reset the port upon detection of a stuck port failure, the switch router will isolate the port from the rest of the functioning ports, and reset the port. This might affect other ports on the interface module.

Note If you configure the switch router as described in the nondefault behavior section after a port stuck failure is detected, the switch router will not reset the Ethernet ports. The Ethernet interface must be configured to reset before the port stuck failure occurs. Also, the default behavior is to not reset the port if a port stuck failure is detected. If the Ethernet interface is not configured to reset when a port stuck failure is detected, schedule the switch router for downtime to remove and reinsert the module.

To configure the switch router to automatically recover from port stuck failures, use the following interface configuration commands:

Command	Purpose
Switch(config-if)# **epc port-reload**	Enables automatic resetting and reloading of the interface module microcode after detecting a port stuck failure.
Switch(config-if)# **epc portstuck-wait** *seconds*	Specifies the delay before signalling a port stuck failure (from the time the failure is detected). The default is 180 seconds. The range for seconds is 0 to 200. A value of 0 seconds causes a port stuck failure to not be detected.

Caution Because of the nature of the microcode architecture, do not configure low values for the wait time in the **epc portstuck-wait** command. The default value of 180 seconds has been carefully chosen, allowing for the hello intervals of protocols such as HSRP, EIGRP, and OSPF. Configuring a low value might lead to incorrectly detecting temporary port stuck failures as real port stuck failures, and can cause a temporary loss of connectivity. It is highly recommended to keep this value to at 60 seconds, at a minimum. Lower values are provided to allow for some specific network designs when you can absolutely rule out temporary port stuck failure scenarios, and also as a debugging aid. For most networks, 180 seconds works very well.

Connectivity Troubleshooting Commands

To troubleshoot a connectivity problem between a port and another port or end-station, use the following commands:

Command	Purpose
show interfaces {**fastethernet** | **gigabitethernet**} *card/subcard/port*	Displays interface configuration, status, and statistics.
show controllers {**fastethernet** | **gigabitethernet**} *card/subcard/port*	Displays controller status for the specified interface.
show epc if-entry interface {**fastethernet** | **gigabitethernet**} *card/subcard/port* **all**	Displays all interface entry information for the specific interface.
show epc ip-prefix interface {**fastethernet** | **gigabitethernet**} *card/subcard/port* **all-entries**	Displays all ip prefix entries for the specified interface.
show epc ip-address interface {**fastethernet** | **gigabitethernet**} *card/subcard/port* **all-entries**	Displays all adjacent IP addresses for the specified interface.
show epc patricia interface {**fastethernet** | **gigabitethernet**} *card/subcard/port* **ipucast detail**	Displays IP unicast patricia tree for the specified interface.
show epc patricia interface {**fastethernet** | **gigabitethernet**} *card/subcard/port* **mac detail**	Displays the MAC patricia tree for the specified interface.

Troubleshooting 10/100 Ethernet Interface Modules

This section describes specific processes and commands used to troubleshoot the 10/100BASE-T and BASE-FX Ethernet interface modules.

The Catalyst 8500 CSR supports two different interface modules. The 10/100BASE-T Ethernet interface module supports 100-Mbps Layer 2 or Layer 3 UTP connections. The 100BASE-FX Ethernet interface module supports 100-Mbps Layer 2 or Layer 3 multimode fiber connections.

This section includes the following:

- 10/100BASE-T Interface Modules, page 10-9
- 100BASE-FX Interface Modules, page 10-9
- Displaying 10/100BASE-T and 100BASE-FX Interface Module Configurations, page 10-10

10/100BASE-T Interface Modules

The 10/100BASE-T Ethernet interface module supports 16 10-Mbps or 100-Mbps Layer 2 or Layer 3 unshielded twisted-pair (UTP) ports. This module supports full-duplex or half-duplex connections and Fast EtherChannel operation. The 10/100BASE-T interface module is available with 16K or 64K of memory. Routing tables use this memory.

10/100BASE-T Interface Module LEDs

Table 10-1 describes the LEDs used to confirm and troubleshoot the operation of interface modules. The LEDs on interface modules indicate the status of the modules and their ports.

Table 10-1 10/100BASE-T Interface Module LED Descriptions

LED	State	Description
Lk	Green	Port is operational (a signal is detected).
	Off	No signal is detected.
Sp	Green	Port is operating at 100 Mbps.
	Off	Port is operating at 10 Mbps.

100BASE-FX Interface Modules

The 100BASE-FX Ethernet interface module supports 100-Mbps Layer 2 or Layer 3 multimode fiber connections. This module supports full-duplex connections and Fast EtherChannel operation. It provides 16 multimode fiber ports that have MT-RJ connectors. The 100BASE-FX interface module is available with 16K or 64K of memory. Routing tables use this memory.

100BASE-FX Interface Module LEDs

Table 10-2 describes the LEDs used to confirm and troubleshoot the operation of interface modules. The LEDs on interface modules indicate the status of the modules and their ports.

Table 10-2 100BASE-FX Interface Module LED Descriptions

LED	State	Description
Tx (Transmit)	Green	Port is transmitting a packet. Green for approximately 50 ms.
	Off	No signal is detected.
Rx (Receive)	Green	Port is receiving a packet. Green for approximately 50 ms.
	Off	No signal is detected.
Link	Green	Port is operational (a signal is detected).
	Off	No signal is detected.

Displaying 10/100BASE-T and 100BASE-FX Interface Module Configurations

To display the 10/100 Ethernet interface module configuration and status, use the following commands:

Command	Purpose
show interfaces FastEthernet *card/subcard/port*	Shows the status of the physical interface.
show controllers FastEthernet *card/subcard/port*	Shows the interface memory management and error counters.
show controllers c8500 counters	Shows the counters on the switch router's interfaces.

Follow these steps to troubleshoot a 10/100 Ethernet interface module:

Step 1 Use the **show interfaces FastEthernet** *card/subcard/port* command to check the configuration.

```
  Switch# show interfaces fastEthernet 3/0/0
→ FastEthernet3/0/0 is up, line protocol is up
    Hardware is epif_port, address is 0090.2156.d837 (bia 0090.2156.d837)
    Internet address is 172.20.52.36/27
    MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
       reliability 255/255, txload 1/255, rxload 1/255
    Encapsulation ARPA, loopback not set
    Keepalive set (10 sec)
→   Auto-duplex, Auto Speed, 100BaseTX
    ARP type: ARPA, ARP Timeout 04:00:00
→   Last input 00:00:29, output never, output hang never
    Last clearing of "show interface" counters never
    Queueing strategy: fifo
    Output queue 0/40, 0 drops; input queue 0/75, 0 drops
    5 minute input rate 0 bits/sec, 0 packets/sec
    5 minute output rate 0 bits/sec, 0 packets/sec
       21584 packets input, 7591871 bytes
       Received 3 broadcasts, 0 runts, 0 giants, 0 throttles
→      0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored 0 abort
       0 watchdog, 21563 multicast
       0 input packets with dribble condition detected
       26882 packets output, 7764915 bytes, 0 underruns(0/0/0)
→      0 output errors, 0 collisions, 0 interface resets
→      0 babbles, 0 late collision, 0 deferred
→      0 lost carrier, 0 no carrier
→      0 output buffer failures, 0 output buffers swapped out
  Switch#
```

Step 2 Check the FastEthernet field to see whether the interface is up. If it is down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the line protocol field to see whether the status is up.

If the interface is down, check for the following:

- The line protocol software processes might have determined that the line is unusable. Try swapping the cable.
- The local or remote interface might be misconfigured. Check the interface configuration.
- Hardware might have failed. Try swapping the interface module.

Step 4 Check the duplex mode field. It should match the speed of the interface and be configured as Auto-negotiation.

Step 5 Check the Last input and Last output fields. They show the number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface.

Step 6 Check the output hang field. It shows the number of hours, minutes, and seconds since the last reset caused by a lengthy transmission.

Step 7 Check the CRC field. The presence of many CRC errors, but not many collisions, indicates excessive noise. If the number of errors is too high, check the cables for damage. If you are using UTP cable, make sure you are using category 5 cables and not another type, such as category 3.

Note Errors and the input and output difference should not exceed 0.5 to 2.0 percent of traffic on the interface.

Step 8 Check the collisions fields. These numbers indicate packet collisions and these numbers should be very low. The total number of collisions, with respect to the total number of output packets, should be 0.1 percent or less.

Step 9 Check the late collisions fields. Late collisions should never occur in a properly designed Ethernet network. They usually occur when Ethernet cables are too long or when there are too many repeaters in the network.

Step 10 Check carrier fields. These numbers indicate a lost carrier detect signal and can be caused by a malfunctioning interface that is not supplying the transmit clock signal, or by a cable problem. If the system notices that the carrier detect line of an interface is up, but the line protocol is down, it periodically resets the interface in an effort to restart it. Interface resets can also occur when an interface is looped back or shut down.

Step 11 Check the buffer fields. These numbers indicate the number of received packets discarded because there was no buffer space. Broadcast storms on Ethernet networks, and bursts of noise on serial lines, are often responsible for no-input buffer events.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Follow these steps to troubleshoot the status of 10/100BASE-T and BASE-FX interfaces:

Step 1 Use the **show controllers FastEthernet** *card/subcard/port* command to check the configuration.

```
   Switch# show controllers fastEthernet 3/0/0

   IF Name: FastEthernet3/0/0
   Port Status UP
   Loopback Reg [3-0]|[7-4]: 0x8|0x8
   Duplex/Speed Reg [3-0]|[7-4]: 0xFFF7|0x0
   FPGA Rev : 3.8
   Internal Reset Trigger Count: 0

   Slicer registers
   SMDR 0x0060 (Tx En,Rx En)
   SSTR 0x1000
   EVER 0x1704 (C1)
   SSMR 0x4000 SIMR 0x0000 MBXW 0x0000 MBXR 0x0000
   SPER 0xF000 GMUX VER 0xF000 MARKER 0x0000

   MAC registers
   CMCR : 0x00000443 (Tx Enabled,Rx Enabled,Full)
   CMPR : 0x140A0E60

   MII registers:
   Control Register              (0x0): 0x1000 (Auto negotiation enabled)
   Status Register               (0x1): 0x782D (Auto negotiation complete)
   PHY Identification Register 1 (0x2): 0x7810

   PHY Identification Register 2 (0x3): 0x43
   Auto Neg. Advertisement Reg   (0x4): 0x1E1 (Speed 100 ,Duplex Full )
   Auto Neg. Partner Ability Reg (0x5): 0x81 (Speed 100 ,Duplex Half )
   Auto Neg. Expansion Register  (0x6): 0x0
   Mirror Register              (0x10): 0x630
   Interrupt Enable Register    (0x11): 0x0
   Interrupt Status Register    (0x12): 0x4000
   Configuration Register       (0x13): 0x0 (UTP, Tx Enabled)
➔  Chip Status Register         (0x14): 0x28C8 (Link Up,a-Half,a-100  )
   Link Status Register    [3-0]|[7-4]: 0x1|0x0

   Counters :

   MAC Receive Counters:
   Bytes                 =7592927
   pkt64                 =22
   pkt65to127            =0
   pkt128to255           =0
   pkt256to511           =21564
   pkt512to1023          =1
   pkt1024to1518         =0
   good_giants           =0
   error_giants          =0

   good_runts            =0
   error_runts           =0
   ucast_pkts            =18
   mcast_pkts            =21566
   bcast_pkts            =3
   align_errs            =0
   fcs_errs              =0
   overruns              =0

   MAC Transmit Counters:
   Bytes                 =7771055
   pkt64                 =1998
   pkt65to127            =3264
   pkt128to255           =39
   pkt256to511           =21597
```

```
pkt512to1023          =29
pkt1024to1522         =0
ucast_pkts            =1342
mcast_pkts            =21640
bcast_pkts            =3945
fcs_errs              =0
giants                =0
underruns             =0

one_collision         =0
mult_collisions       =0
excess_collisions     =0
Ingress Markers       =46522
Egress Markers        =27508

Slicer Receive Counters:
Cells                 =22528410
Frames                =40502
Header Sequence Errors=0
fcs_errs              =0
Length                =0

Slicer Transmit Counters:
Cells                 =13186868
Frames                =34787

Switch#
```

Step 2 Check the Chip Status Register field. It should match the link status, duplex mode, and speed shown in the previous **show interface** command. If it does not, see the "Troubleshooting Half- or Full-Duplex Negotiation" section on page 11-21.

Follow these steps to troubleshoot the counters of the Fast Ethernet interface module physical interface:

Step 1 Use the **show controllers c8500 counters** command to check the Fast Ethernet interface module counters.

```
  Switch# show controllers c8500 counters
➔ Interface   Input     Runts Giants  Input       CRC  Frame Output      Output
      State   Packets                 Errors                 Packets     Errors
  -----------------------------------------------------------------------------
  G0/1/0  U   0           0     0     0           0    0     136972      0
  G0/1/1  U   0           0     0     0           0    0     20          0
  P1/0/0  AD  0           19600630    2271017     3    0     0           0
  P2/0/0  AD  0           2     0     139         2    0     0           0
  G2/0/1  AD  1           0     0     0           0    0     1           0
  A3/0/0  AD  0           0     0     0           0    0     0           0
  G3/0/1  AD  1           0     0     0           0    0     1           0
  F9/0/0  U   14364       0     0     0           0    0     14367       0
  F9/0/1  AD  1           0     0     0           0    0     1           0
  F9/0/2  AD  1           0     0     0           0    0     1           0
  F9/0/3  AD  1           0     0     0           0    0     1           0
  F9/0/4  AD  1           0     0     0           0    0     1           0
  F9/0/5  AD  1           0     0     0
  .
  (Information Deleted)
  .
  A12/0/0 AD  0           0     0     0           0    0     0           0
  G12/0/1 AD  1           0     0     0           0    0     1           0
  -----------------------------------------------------------------------------
  AD - Admin Down, D - Down, F - Fail, U - Up

  Switch#
```

Step 2 Check the Interface State field. It should indicate the interfaces are up.

Step 3 Check the Input Packets and Output Packets fields. The **show controllers c8500 counters** command should be entered at least twice. The counters in the Input Packets and Output Packets fields should be incrementing. This information can also be displayed using the **show interfaces** command.

Note The **clear counters** command does not clear the **show controllers c8500 counters** command display.

If you determine that the interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Troubleshooting Gigabit Ethernet Interface Modules

This section describes specific processes and commands used to troubleshoot the Gigabit Ethernet interface modules.

The Catalyst 8500 CSR supports three different interface modules for Gigabit Ethernet transmission over fiber connections: the eight-port Gigabit Ethernet interface module, the two-port Gigabit Ethernet interface module, and the two-port enhanced Gigabit Ethernet interface module.

This section includes the following:

- Ethernet and Gigabit Ethernet Processor Differences, page 10-15
- Eight-Port Gigabit Ethernet Interface Modules, page 10-20
- Two-Port Gigabit Ethernet Interface Modules, page 10-21
- Troubleshooting Two-Port Enhanced Gigabit Ethernet Interface Modules, page 10-26
- Displaying Enhanced Gigabit Ethernet Interface Module Configurations, page 10-28

Ethernet and Gigabit Ethernet Processor Differences

The switch router uses two different Gigabit Ethernet interface-modules hardware types.

The two- and eight-port Gigabit Ethernet interface modules use the Ethernet processor interface (EPIF) that has an internal binary CAM search engine built into the processor.

The enhanced Gigabit Ethernet interface module uses the Gigabit processor interface (XPIF) that has a faster external search engine using a Cisco Systems proprietary FPGA and Ternary CAM (TCAM) for the Layer 3 routing and Layer 2 switching functionality.

The two-port and eight-port Gigabit Ethernet interface modules are full-width modules. The Ethernet interface processors support 1000-Mbps Layer 2 or Layer 3 fiber-optic connections. They provide Gigabit Ethernet ports that have Gigabit Interface Converter (GBIC) modular transceivers and SC-type fiber connectors. The two-port Gigabit Ethernet interface module is available with 16K or 64K of memory. The eight-port Gigabit Ethernet interface module is available with 16K of routing table memory.

The two-port enhanced Gigabit Ethernet interface module with Gigabit interface processors supports 1000-Mbps multimode and single-mode Layer 2 and Layer 3 fiber-optic connections. It consists of two one-port Gigabit Ethernet port adapters attached to a carrier module. The port adapters are not hot-swappable, but the complete interface module is hot-swappable. The port adapters have GBIC modular transceivers and SC-type fiber connectors. The interface module is full-duplex, supports Fast EtherChannel operation, and provides built-in ACL functionality. It is available with 16K, 64K, or 256K of routing table memory.

The troubleshooting procedures are slightly different for the Ethernet processor interface and Gigabit processor interface modules. You need to determine which type of Gigabit interface module you are troubleshooting.

To display the Gigabit Ethernet interface and enhanced Gigabit Ethernet interface modules installed in your switch router, use the following commands:

Command	Purpose
show hardware [**detail**]	Shows physical interfaces and their type.
show interfaces gigabitEthernet *card/subcard/port*	Shows the status of the physical interfaces and the type.

Follow these steps to determine which type of Gigabit Ethernet interface modules are installed in your switch router:

Step 1 Use the **show hardware** command to check the Gigabit interface type.

```
  Switch# show hardware

  C8540 named Switch, Date: 12:23:27 PST Sat Feb 26 2000

  Slot Ctrlr-Type    Part No.  Rev  Ser No  Mfg Date   RMA No. Hw Vrs  Tst EEP
  ---- ------------  ---------- -- -------- --------- -------- ------- --- ---
   0/* K1 GIGETHERN  73-3324-03 A0 0336441Y Oct 13 99 0          3.4
   1/* CMPM Card     73-3944-03 09 03445724 Nov 09 99            3.0
   1/0 XPIF POS OC1  73-4462-01 09 034558YP Nov 09 99            1.1
   2/* CMPM Card     73-3944-03 A0 04087BW8 Mar 22 00 0          3.0
   2/0 XPIF POS OC1  73-4462-01 A0 04046NRQ Mar 22 00 0          2.0
   2/1 XPIF GIGE PA  73-4167-05 A0 04097GRJ Mar 22 00 0          1.0
   3/* CMPM Card     73-3944-03 A0 04087BXK Mar 15 00 0          3.0
   3/0 XPIF ATM OC3  73-3889-03 A0 040879AA Mar 15 00 0          1.0
   3/1 XPIF GIGE PA  73-4167-05 A0 04097GQA Mar 15 00 0          1.0
   4/* Route Proc    73-3775-04 A0 03201VCZ Oct 04 99 0          5.7
   5/* Switch Card   73-3327-08 A0 032428ZR Jun 15 99 0          8.0
   7/* Switch Card   73-3327-08 A0 032428ZE Jun 15 99 0          8.0
   9/* ETHERNET PAM  73-3754-06 C0 04239U9B Jun 16 00 0          5.1
→ 10/* GIGETHERNET   73-3375-03 04 031215VT Apr 06 99            3.0
  11/* CMPM Card     73-3944-03 A0 04087BY5 Mar 13 00 0          3.0
→ 11/0 XPIF GIGE PA  73-4415-05 A0 04087AZE Mar 13 00 0          1.0
→ 11/1 XPIF GIGE PA  73-4415-05 A0 04087AZL Mar 13 00 0          1.0
  12/* CMPM Card     73-3944-03 A0 04087BWS Mar 14 00 0          3.0
  12/0 XPIF ATM OC1  73-3889-03 A0 040879AO Mar 14 00 0          1.0
  12/1 XPIF GIGE PA  73-4167-05 A0 04107N8R Mar 14 00 0          1.0
  .
  (Information Deleted)
  .
```

Step 2 Check the Ctrlr-Type field of the **show hardware** command.

The interface module installed in slot 11 has the following components:

- slot 11/* is the CMPM controller type or carrier module with the two enhanced Gigabit Ethernet interface modules installed.
- slot 11/0 is the XPIF GIGE PA or enhanced Gigabit Ethernet interface module installed in the left side of the carrier module.
- slot 11/1 is the XPIF GIGE PA or enhanced Gigabit Ethernet interface module installed in the right side of the carrier module.

Note The individual enhanced Gigabit Ethernet interface modules are not hot-swappable, but the entire carrier module with both interface modules installed is hot-swappable.

In the previous **show hardware** command example, the GIGETHERNET interface module installed in slot 10/* does *not* have "XPIF" preceding its controller type, and is an Ethernet processor (EPIF) type Gigabit Ethernet interface module.

Step 3 Use the **show hardware detail** command to check the Gigabit interface processor type in greater detail.

```
Switch# show hardware detail

C8540 named Switch, Date: 12:25:53 PST Sat Feb 26 2000

Slot Ctrlr-Type    Part No.  Rev  Ser No  Mfg Date   RMA No. Hw Vrs  Tst EEP
---- ------------  ---------- -- -------- --------- -------- ------- --- ---
 0/* K1 GIGETHERN  73-3324-03 A0 0336441Y Oct 13 99 0          3.4
 1/* CMPM Card     73-3944-03 09 03445724 Nov 09 99            3.0
 1/0 XPIF POS OC1  73-4462-01 09 034558YP Nov 09 99            1.1
 2/* CMPM Card     73-3944-03 A0 04087BW8 Mar 22 00 0          3.0
 2/0 XPIF POS OC1  73-4462-01 A0 04046NRQ Mar 22 00 0          2.0
 2/1 XPIF GIGE PA  73-4167-05 A0 04097GRJ Mar 22 00 0          1.0
 3/* CMPM Card     73-3944-03 A0 04087BXK Mar 15 00 0          3.0
 3/0 XPIF ATM OC3  73-3889-03 A0 040879AA Mar 15 00 0          1.0
 3/1 XPIF GIGE PA  73-4167-05 A0 04097GQA Mar 15 00 0          1.0
 9/* ETHERNET PAM  73-3754-06 C0 04239U9B Jun 16 00 0          5.1
10/* GIGETHERNET   73-3375-03 04 031215VT Apr 06 99            3.0
11/* CMPM Card     73-3944-03 A0 04087BY5 Mar 13 00 0          3.0
11/0 XPIF GIGE PA  73-4415-05 A0 04087AZE Mar 13 00 0          1.0
11/1 XPIF GIGE PA  73-4415-05 A0 04087AZL Mar 13 00 0          1.0
12/* CMPM Card     73-3944-03 A0 04087BWS Mar 14 00 0          3.0
12/0 XPIF ATM OC1  73-3889-03 A0 040879AO Mar 14 00 0          1.0
12/1 XPIF GIGE PA  73-4167-05 A0 04107N8R Mar 14 00 0          1.0

➔ slot: 10/*  Controller-Type : GIGETHERNET PAM
  Part Number: 73-3375-03                        Revision: 04
Serial Number: CAB031215VT                       Mfg Date: Apr 06 99
   RMA Number:                                H/W Version: 3.0
 FPGA Version: 2.3

 EPIF Version: 1704                              CAM size: 64 KB
Ucode Version: 0.0                               CAM Type: Dual

Port Phy Setup
      Port  0: DONE                           GBIC Vendor: No vendor info.
      Port  1: DONE                           GBIC Vendor: No vendor info.

➔ slot: 11/*  Controller-Type : CMPM Card
  Part Number: 73-3944-03                        Revision: A0
Serial Number: CAB04087BY5                       Mfg Date: Mar 13 00
   RMA Number: 0                              H/W Version: 3.0
 FPGA Version: 1.4

➔ slot: 11/0  Controller-Type : XPIF GIGE PAM
  Part Number: 73-4415-05                        Revision: A0
Serial Number: CAB04087AZE                       Mfg Date: Mar 13 00
   RMA Number: 0                              H/W Version: 1.0
 FPGA Version: 20.72

 XPIF Version: 3001                              CAM size: 16 KB
Ucode Version: 1.0                               CAM Type: Private TCAM
```

```
Port Phy Setup
      Port  0: DONE                          GBIC Vendor: No vendor info.

➔ slot: 11/1  Controller-Type : XPIF GIGE PAM
  Part Number: 73-4415-05                       Revision: A0
Serial Number: CAB04087AZL                      Mfg Date: Mar 13 00
   RMA Number: 0                             H/W Version: 1.0
 FPGA Version: 20.72

 XPIF Version: 3001                             CAM size: 16 KB
Ucode Version: 1.0                              CAM Type: Private TCAM

Port Phy Setup
      Port  0: DONE                          GBIC Vendor: No vendor info.

.
(Information Deleted)
.
```

Step 4 Check the Controller-Type field in the **show hardware detail** command.

The interface module installed in slot 11 has the following components:

- In slot 11/* the Controller-Type field is the CMPM controller type or carrier module with the two enhanced Gigabit Ethernet interface modules installed.
- In slot 11/0 the Controller-Type field is the XPIF GIGE PA or enhanced Gigabit Ethernet interface module installed in the left side of the carrier module.
- In slot 11/1 the Controller-Type field is the XPIF GIGE PA or enhanced Gigabit Ethernet interface module installed in the right side of the carrier module.

In the previous **show hardware detail** command example, the interface module installed in slot 10/* has GIGETHERNET PAM listed in the Controller-Type field and is an Ethernet processor type interface.

Step 5 Use the **show interfaces GigabitEthernet** *card/subcard/port* command as another way to check the Gigabit Ethernet interface processor type.

```
Switch# show interfaces gigabitEthernet 11/0/0
GigabitEthernet11/0/0 is administratively down, line protocol is down
➔ Hardware is xpif_port, address is 00d0.ba1d.3267 (bia 00d0.ba1d.3267)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
.
(Information Deleted)
.
```

Step 6 Check the Hardware field. In this example, the hardware is listed as xpif_port, indicating this interface module uses the Gigabit processor interface.

Step 7 Use the **show interfaces GigabitEthernet** *card/subcard/port* command on a different interface to check the Gigabit Ethernet interface processor type.

```
Switch# show interfaces gigabitEthernet 10/0/0
GigabitEthernet10/0/0 is up, line protocol is up
➔ Hardware is epif_gigether_port, address is 00d0.5845.1257 (bia 00d0.5845.1257)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, 1000Base-SX, Auto-negotiation
  output flow-control is unsupported, input flow-control is unsupported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output never, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     345 packets input, 119370 bytes
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored 0 abort
     0 watchdog, 349 multicast
     0 input packets with dribble condition detected
     688 packets output, 238736 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
Switch#
```

Step 8 In this **show interface GigabitEthernet** *card/subcard/port* command example, again check the Hardware field. The hardware is listed as epif_gigether_port, indicating this interface module uses the Ethernet processor interface type.

Troubleshooting the Gigabit interface module with the Ethernet interface processor is described in the following sections:

- Eight-Port Gigabit Ethernet Interface Modules, page 10-20
- Two-Port Gigabit Ethernet Interface Modules, page 10-21
- Displaying Gigabit Ethernet Interface Module Configurations, page 10-21.

Troubleshooting the enhanced Gigabit interface module with the Gigabit interface processor is described in the following troubleshooting sections:

- Troubleshooting Two-Port Enhanced Gigabit Ethernet Interface Modules, page 10-26
- Troubleshooting ATM Uplink with Enhanced Gigabit Ethernet Interface Modules, page 10-34
- Troubleshooting Packet-over-SONET Uplink with Enhanced Gigabit Ethernet Interface Modules, page 10-41

Eight-Port Gigabit Ethernet Interface Modules

The eight-port Gigabit Ethernet interface module supports 1000-Mbps Layer 2 or Layer 3 fiber-optic connections. It provides eight Gigabit Ethernet ports that have Gigabit Interface Converter (GBIC) modular transceivers and SC-type fiber connectors. The eight-port Gigabit Ethernet interface module is available with 16K of memory. Routing tables use this memory.

Figure 10-1 is a block diagram of the eight Gigabit Ethernet port interface module, and shows how the interface communicates with the route processor and switch fabric across the backplane.

Figure 10-1 Eight Gigabit Ethernet Port Interface Block Diagram

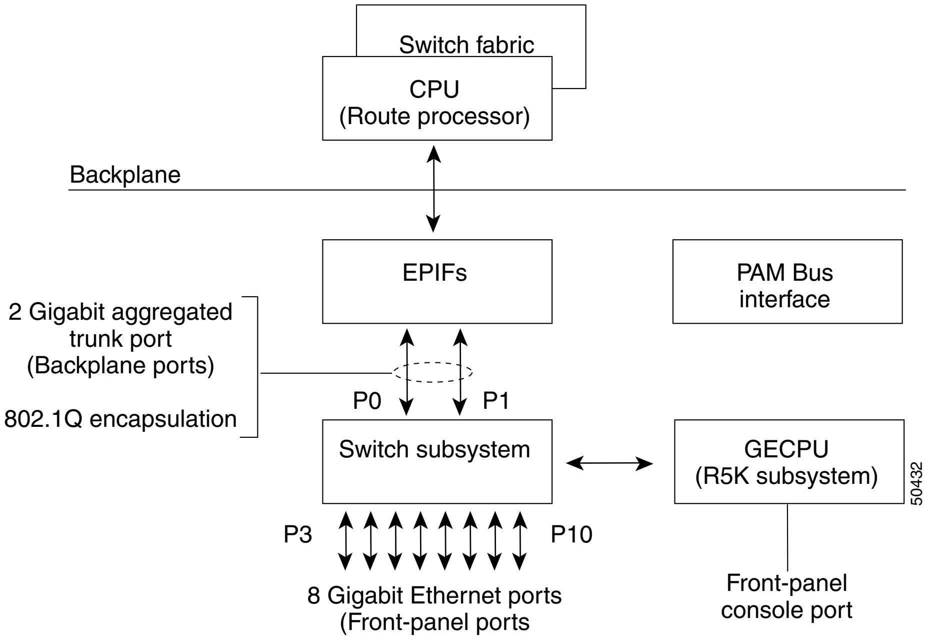

Eight-Port Gigabit Ethernet Interface Module LEDs

Table 10-3 describes the LEDs used to confirm and troubleshoot the operation of interface modules. The LEDs on interface modules indicate the status of the modules and their ports.

Table 10-3 Eight-Port Gigabit Ethernet Interface Module LED Descriptions

LED	State	Description
Status	Green	The system has passed internal self-tests and diagnostic tests.
	Red	The system has failed internal self-tests and diagnostic tests.
	Orange	The system is booting or a module is disabled.
(link)	Green	The Ethernet port is operational.
	Off	No signal is detected on the Ethernet port.

Two-Port Gigabit Ethernet Interface Modules

The two-port Gigabit Ethernet interface module supports 1000-Mbps Layer 2 or Layer 3 fiber-optic connections. It provides two Gigabit Ethernet ports that have GBIC modular transceivers and SC-type fiber connectors. The two-port Gigabit Ethernet interface module is available with 16K or 64K of memory. Routing tables use this memory.

Two-Port Gigabit Ethernet Interface Module LEDs

Table 10-4 describes the LEDs used to confirm and troubleshoot the operation of interface modules. The LEDs on interface modules indicate the status of the modules and their ports.

Table 10-4 Two-Port Gigabit Ethernet Interface Module LED Descriptions

LED	State	Description
Op-Det	On	An optical signal from another Gigabit Ethernet module is detected. It is steadily on when there is a Gigabit connection.
	Off	No Gigabit Ethernet optical signal is detected.
Tx (Transmit)	Green	A port is transmitting a packet. Green for approximately 50 ms.
	Off	No signal is detected.
Full-Duplex	On	A port is operating in full-duplex mode. This is always the case for an operational Gigabit Ethernet port.
Link	Green	A port is operational (a signal is detected).
	Off	No signal is detected.
Rx (Receive)	Green	A port is receiving a packet. Green for approximately 50 ms
	Off	No signal is detected.

Displaying Gigabit Ethernet Interface Module Configurations

To display the Gigabit Ethernet interface module using the Ethernet processor interface type configuration and status, use the following commands:

Command	Purpose
show interfaces gigabitEthernet *card/subcard/port*	Shows the status of the physical interface.
show controllers gigabitEthernet *card/subcard/port*	Shows the interface memory management and error counters.
show controllers c8500 counters	Shows the counters on the switch router's interfaces.

Follow these steps to troubleshoot a Gigabit Ethernet interface module physical interface:

Step 1 Use the **show interfaces GigabitEthernet** *card/subcard/port* command to check the configuration and status.

```
   Switch# show interfaces gigabitEthernet 10/0/0
→  GigabitEthernet10/0/0 is up, line protocol is up
     Hardware is epif_gigether_port, address is 00d0.5845.1257 (bia 00d0.5845.1257)
     MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
        reliability 255/255, txload 1/255, rxload 1/255
     Encapsulation ARPA, loopback not set
     Keepalive set (10 sec)
→    Full-duplex, 1000Mb/s, 1000Base-SX, Auto-negotiation
     output flow-control is unsupported, input flow-control is unsupported
     ARP type: ARPA, ARP Timeout 04:00:00
→    Last input 00:00:00, output never, output hang never
     Last clearing of "show interface" counters never
     Queueing strategy: fifo
     Output queue 0/40, 0 drops; input queue 0/75, 0 drops
     5 minute input rate 0 bits/sec, 0 packets/sec
     5 minute output rate 0 bits/sec, 0 packets/sec
        345 packets input, 119370 bytes
        Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
→       0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored 0 abort
        0 watchdog, 349 multicast
        0 input packets with dribble condition detected
        688 packets output, 238736 bytes, 0 underruns
→       0 output errors, 0 collisions, 0 interface resets
→       0 babbles, 0 late collision, 0 deferred
→       0 lost carrier, 0 no carrier
→       0 output buffer failures, 0 output buffers swapped out
   Switch#
```

Step 2 Check the GigabitEthernet field to see whether the interface is up.

If the interface is down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the line protocol field to see whether the status is up.

If the interface is down, check for the following:

- The line protocol software processes might have determined that the line is unusable. Try swapping the cable.
- The local or remote interface might be misconfigured. Check the interface configuration.
- Hardware might have failed. Try swapping the interface module.

Step 4 Check the duplex mode field. It should match the speed of the interface and be configured as Auto-negotiation.

Step 5 Check the Last input and Last output fields. They show the number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface.

Step 6 Check the output hang field. It shows the number of hours, minutes, and seconds since the last reset caused by a lengthy transmission.

Step 7 Check the CRC field. The presence of many CRC errors, but not many collisions, indicates excessive noise. If the number of errors is too high, check the cables for damage. If you are using UTP cable, make sure you are using category 5 cables and not another type, such as category 3.

Note Errors and the input and output difference should not exceed 0.5 to 2.0 percent of traffic on the interface.

Step 8 Check the collisions fields. These numbers indicate packet collisions, and these numbers should be very low. The total number of collisions with respect to the total number of output packets should be approximately 0.1 percent or less.

Step 9 Check the late collisions fields. Late collisions should never occur in a properly designed Ethernet network. They usually occur when Ethernet cables are too long or when there are too many repeaters in the network.

Step 10 Check the carrier fields. These numbers indicate a lost carrier detect signal, and can be caused by a malfunctioning interface that is not supplying the transmit clock signal, or by a cable problem. If the system notices that the carrier detect line of an interface is up, but the line protocol is down, it periodically resets the interface in an effort to restart it. Interface resets can also occur when an interface is looped back or shut down.

Step 11 Check the buffers fields. These numbers indicate the number of received packets discarded because there was no buffer space. Broadcast storms on Ethernet networks, and bursts of noise on serial lines, are often responsible for no-input buffer events.

If you determine that the physical interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Follow these steps to troubleshoot the status of a Gigabit Ethernet interface module:

Step 1 Use the **show controllers GigabitEthernet** *card/subcard/port* command to check the configuration.

Note The Catalyst 8540 CSR has no switch feature card. Consequently, you can not check the number of cells switched on an individual VC.

```
Switch# show controllers gigabitEthernet 10/0/0

IF Name: GigabitEthernet10/0/0
Port Status UP
FPGA Rev : 0.4
Gigabit Ether Status         : 0xFDE7(Optical Detect,Rx Sync,Link Up)
Mode Parallel Register       : 0x0
Port 0 Serial Mode Register  : 0x0
Port 1 Serial Mode Register  : 0x0
Link Interrupt Enable        : 0x1
Tx Disable                   : 0x0

Slicer registers
SMDR 0x0060 (Tx En,Rx En)
SSTR 0x1000
EVER 0x1704 (C1)
SSMR 0x4000 SIMR 0x0000 MBXW 0x0000 MBXR 0x0000
SPER 0xF000 GMUX VER 0x17B1 MARKER 0x17B1

MAC registers
CMCR : 0x00000423 (Tx Enabled,Rx Enabled,Half)
CMPR : 0x140A0E61

MII registers:

Control Register              (0x0): 0x1140
Status Register               (0x1): 0x16D
Auto Neg. Advt. Register      (0x4): 0x20
Auto Neg. Partner Ability Reg (0x5): 0x0
RX Configuration Register     (0xA): 0x17
TR_IPG_TIME Register         (0x10): 0x3
PAUSE_TIME Register          (0x11): 0x0
PAUSE_SA1 Register           (0x12): 0x0
PAUSE_SA2 Register           (0x13): 0x0
PAUSE_SA3 Register           (0x14): 0x0
Pause Watermark Register     (0x15): 0xC040
TX FIFO Watermark Register   (0x16): 0xFF02
PAUSE_STAT_SENT Register     (0x17): 0x0
PAUSE_STAT_RCVD Register     (0x18): 0x0
Memory Address Register      (0x19): 0x0
Memory Control Register      (0x1A): 0x1
Memory Data High Register    (0x1B): 0x0
Memory Data Low Register     (0x1C): 0x0
Sys Control Register         (0x1E): 0x70C
Sys Status Register          (0x1F): 0x80
Link Status Register    [3-0]|[7-4]: 0x1|0x0

Counters :


Channel 0:

MAC Receive Counters:
```

```
Bytes                 =71156278
pkt64                 =0
pkt65to127            =0
pkt128to255           =0
pkt256to511           =44
pkt512to1023          =0
pkt1024to1518         =0
good_giants           =0
error_giants          =0
good_runts            =8714
error_runts           =0
ucast_pkts            =8714
mcast_pkts            =44
bcast_pkts            =0
align_errs            =0
fcs_errs              =0
overruns              =0

MAC Transmit Counters:

Bytes                 =30189
pkt64                 =0
pkt65to127            =0
pkt128to255           =0
pkt256to511           =87
pkt512to1023          =0
pkt1024to1522         =0
ucast_pkts            =0
mcast_pkts            =87
bcast_pkts            =0
fcs_errs              =17429
giants                =0
underruns             =0
one_collision         =0
mult_collisions       =0
excess_collisions     =0
Ingress Markers       =8714
Egress Markers        =17429

Slicer Receive Counters:
Cells                 =8002496
Frames                =33557
Header Sequence Errors=0

fcs_errs              =0
Length                =0

Slicer Transmit Counters:
Cells                 =3951054
Frames                =12096

Switch#
```

Step 2 Check the Chip Status Register field. It should match the link status, duplex mode, and speed shown in the previous **show interface** command. If it does not, see the “Troubleshooting Half- or Full-Duplex Negotiation” section on page 11-21.

Follow these steps to troubleshoot the counters of the Gigabit Ethernet interface module physical interface:

Step 1 Use the **show controllers c8500 counters** command to check the Gigabit Ethernet interface module counters.

```
  Switch# show controllers c8500 counters
➔ Interface   Input     Runts Giants  Input       CRC  Frame Output      Output
      State   Packets                 Errors                 Packets     Errors
  ----------------------------------------------------------------------------
  F0/0/0  U   0           0     0     0           0     0     349         0
  F0/0/1  D   0           0     0     0           0     0     1           0
  .
  (Information Deleted)
  .
➔ G10/0/0 U   347         0     0     0           0     0     692         0
  G10/0/1 U   346         0     0     0           0     0     347         0
  .
  (Information Deleted)
  .
  ----------------------------------------------------------------------------
  AD - Admin Down, D - Down, F - Fail, U - Up

  Switch#
```

Step 2 Check the Interface State field. It should indicate the interfaces are up.

Step 3 Check the Input Packets and Output Packets fields. The **show controllers c8500 counters** command should be entered at least twice. The counters in the Input Packets and Output Packets fields should be incrementing. This information can also be displayed using the **show interfaces** command.

Note The **clear counters** command does not clear the **show controllers c8500 counters** command display.

If you determine that the interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Troubleshooting Two-Port Enhanced Gigabit Ethernet Interface Modules

The two-port enhanced Gigabit Ethernet interface module supports 1000-Mbps multimode and single-mode Layer 2 and Layer 3 fiber-optic connections. It consists of two one-port Gigabit Ethernet port adapters attached to a carrier module. The port adapters are not hot-swappable, but the complete interface module is hot-swappable. The port adapters have GBIC modular transceivers and SC-type fiber connectors. The interface module is full-duplex, supports Fast EtherChannel operation, and provides built-in ACL functionality. It is available with 16K, 64K, or 256K of routing table memory.

Note The port adapters within the two-port interface modules must have matching routing table memory. That is, if the ATM OC-12c port adapter has 64K of routing table memory, the Gigabit Ethernet port adapter must have 64K of routing table memory for the interface module to function properly.

Figure 10-2 is a block diagram of the enhanced Gigabit Ethernet interface module and shows how the interface communicates with the switch fabric across the backplane.

Figure 10-2 Enhanced Gigabit Ethernet Block Diagram

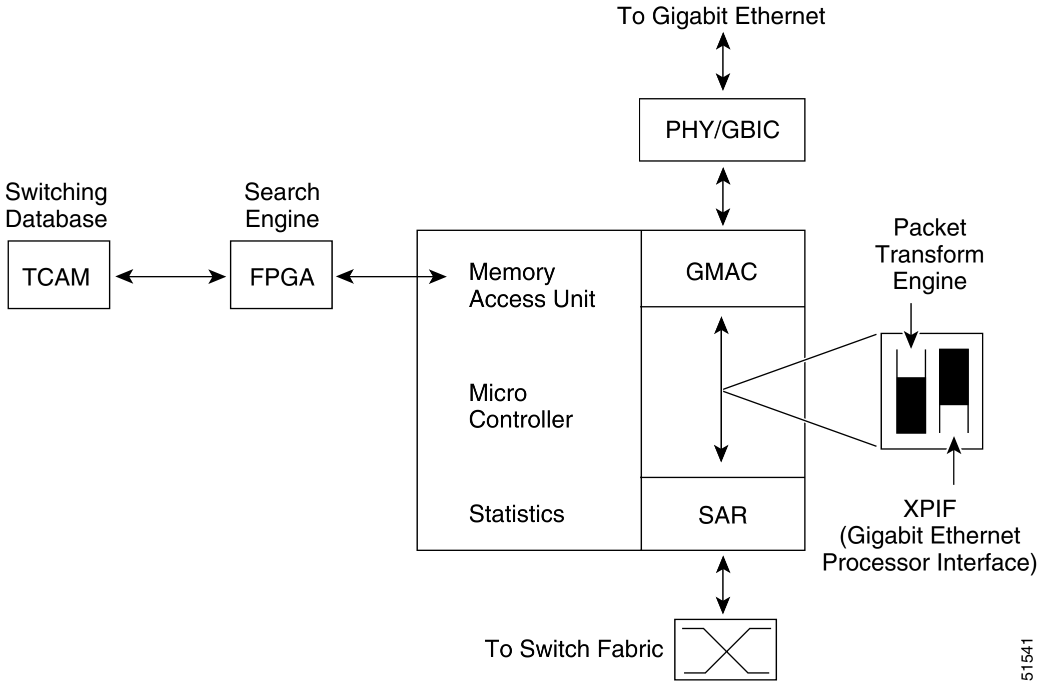

The enhanced Gigabit Ethernet interface module uses the Gigabit processor interface (XPIF) with a faster external search engine that has a Cisco Systems proprietary FPGA and Ternary CAM (TCAM) to provide the search engine for the Layer 3 routing and Layer 2 switching functionality.

The Gigabit Ethernet interface module with the Gigabit processor interface is used with all of the interface modules described in the troubleshooting sections:

- Troubleshooting Two-Port Enhanced Gigabit Ethernet Interface Modules, page 10-26
- Troubleshooting ATM Uplink with Enhanced Gigabit Ethernet Interface Modules, page 10-34
- Troubleshooting Packet-over-SONET Uplink with Enhanced Gigabit Ethernet Interface Modules, page 10-41

Two-Port Enhanced Gigabit Ethernet Interface Module LEDs

Table 10-5 describes the LEDs used to confirm and troubleshoot the operation of interface modules. The LEDs on interface modules indicate the status of the modules and their ports.

Table 10-5 Two-Port Enhanced Gigabit Ethernet Interface Module LED Descriptions

LED	State	Description
Link	Green	A port is operational (a signal is detected).
	Off	No signal is detected.

Table 10-5 Two-Port Enhanced Gigabit Ethernet Interface Module LED Descriptions (continued)

LED	State	Description
Full-Duplex	On	A port is operating in full-duplex mode. This is always the case for an operational Gigabit Ethernet port.
Rx (Receive)	Green	A port is receiving a packet. Green for approximately 50 ms.
	Off	No signal is detected.
Op-Det	Green	An optical signal from another Gigabit Ethernet module is detected. It is steadily on when there is a Gigabit connection.
	Off	No signal is detected.
Tx (Transmit)	Green	A port is transmitting a packet. Green for approximately 50 ms.
	Off	No signal is detected.
Rx Sync	Green	A port is synchronized with the port from which it is receiving data.

Displaying Enhanced Gigabit Ethernet Interface Module Configurations

To display the enhanced Gigabit Ethernet interface module configuration and status, use the following commands:

Command	Purpose
show interfaces gigabitEthernet *card/subcard/port*	Shows the status of the physical interface.
show controllers gigabitEthernet *card/subcard/port*	Shows the interface memory management and error counters.
show controllers c8500 counters	Shows the counters on the switch router interfaces.

Follow these steps to troubleshoot an enhanced Gigabit Ethernet interface module physical interface:

Step 1 Use the **show interfaces GigabitEthernet** *card/subcard/port* command to check the configuration and status.

```
  Switch# show interfaces gigabitEthernet 11/0/1
➔ GigabitEthernet11/0/1 is up, line protocol is up
    Hardware is xpif_port, address is 00d0.ba1d.3367 (bia 00d0.ba1d.3367)
    MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
       reliability 255/255, txload 1/255, rxload 1/255
    Encapsulation ARPA, loopback not set
    Keepalive set (10 sec)
➔   Full-duplex, 1000Mb/s, 1000Base-SX, Auto-negotiation
    output flow-control is unsupported, input flow-control is unsupported
    ARP type: ARPA, ARP Timeout 04:00:00
➔   Last input 00:00:12, output never, output hang never
    Last clearing of "show interface" counters never
    Queueing strategy: fifo
    Output queue 0/40, 0 drops; input queue 0/75, 0 drops
    5 minute input rate 0 bits/sec, 0 packets/sec
    5 minute output rate 0 bits/sec, 0 packets/sec
       21583 packets input, 7592700 bytes
       Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
➔      0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored 0 abort
       0 watchdog, 21582 multicast
       0 input packets with dribble condition detected
       41663 packets output, 14916014 bytes, 0 underruns(0/0/0)
➔      0 output errors, 0 collisions, 0 interface resets
➔      0 babbles, 0 late collision, 0 deferred
➔      0 lost carrier, 0 no carrier
➔      0 output buffer failures, 0 output buffers swapped out
  Switch#
```

Step 2 Check the GigabitEthernet field to see whether the interface is up.

If down, check for the following:

- Disconnected or faulty cabling—Check cables.
- Hardware failure—Swap hardware.

If administratively down, the interface has been administratively taken down. Use the **no shutdown** interface configuration command to reenable the interface.

Step 3 Check the line protocol field to see whether the status is up.

If the interface is down, check for the following:

- The line protocol software processes might have determined that the line is unusable. Try swapping the cable.
- The local or remote interface might be misconfigured. Check the interface configuration.
- Hardware might have failed. Try swapping the interface module.

Step 4 Check the duplex mode field. It should match the speed of the interface and be configured as Auto-negotiation.

Step 5 Check the Last input and Last output fields. They show the number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface.

Step 6 Check the output hang field. It shows the number of hours, minutes, and seconds since the last reset caused by a lengthy transmission.

Step 7 Check the CRC field. The presence of many CRC errors, but not many collisions, indicates excessive noise. If the number of errors is too high, check the cables for damage. If you are using UTP cable, make sure you are using category 5 cables and not another type, such as category 3.

Note Errors and the input and output difference should not exceed 0.5 to 2.0 percent of traffic on the interface.

Step 8 Check the collisions fields. These numbers indicate packet collisions, and these numbers should be very low. The total number of collisions with respect to the total number of output packets should be 0.1 percent or less.

Step 9 Check the late collisions fields. Late collisions should never occur in a properly designed Ethernet network. They usually occur when Ethernet cables are too long or when there are too many repeaters in the network.

Step 10 Check carrier fields. These numbers indicate a lost carrier detect signal and can be caused by a malfunctioning interface that is not supplying the transmit clock signal, or by a cable problem. If the system notices that the carrier detect line of an interface is up, but the line protocol is down, it periodically resets the interface in an effort to restart it. Interface resets can also occur when an interface is looped back or shut down.

Step 11 Check the buffers fields. These numbers indicate the number of received packets discarded because there was no buffer space. Broadcast storms on Ethernet networks and bursts of noise on serial lines are often responsible for no input buffer events.

If you determine that the physical interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Follow these steps to troubleshoot the status of a Gigabit Ethernet interface module:

Step 1 Use the **show controllers GigabitEthernet** *card/subcard/port* command to check the configuration.

Note The Catalyst 8540 CSR has no switch feature card and does not support the **show controllers GigabitEthernet** command and no individual VC statistics are available.

```
Switch# show controllers gigabitEthernet 9/0/0
IF Name: GigabitEthernet9/0/0
Port Status UP
FPGA Rev : 0.2
Gigabit Ether Status         : 0xF (Optical Detect,Rx Sync,Link UP)
Mode Parallel Register       : 0x0
Serial Mode Register         : 0x0
Link Interrupt Enable        : 0x1
Tx Disable                   : 0x0
Internal Reset Trigger Count : 0

Slicer registers
SMDR 0xFF78 SSTR 0x1202 SSMR 0x4002 EVER 0x3001
SIMR 0x0000 MBXW 0x0000 MBXR 0x0000 SPER 0xF000


F000  chan0 chan1 chan2 chan3 sstr 1202
      0006  0016  0006  0006
```

```
task0   11    11    11    11
task1  5EF   5EF   5EF   5EF
task2   11    11    11    11
task3  5EF   5EF   5EF   5EF

 GCR = 0x4       GICR = 0x2403


MII registers:

Direct Access:
Control Register               (0x0): 0x1140
Status Register                (0x1): 0x16D
Auto Neg. Advt. Register       (0x4): 0x1A0
Auto Neg. Partner Ability Reg (0x5): 0x4020
TR_IPG_TIME Register          (0x10): 0x7
PAUSE_TIME Register 1         (0x11): 0x100
PAUSE_TIME Register 2         (0x12): 0x18
PAUSE_SA1 Register            (0x13): 0x0
PAUSE_SA2 Register            (0x14): 0x0
PAUSE_SA3 Register            (0x15): 0x0
PAUSE_DA1 Register            (0x16): 0x180
PAUSE_DA2 Register            (0x17): 0xC200
PAUSE_DA3 Register            (0x18): 0x1
Pause Upper Watermark Reg.    (0x19): 0xC00
Pause Lower Watermark Reg.    (0x1A): 0x1000
TX FIFO Watermark Register    (0x1B): 0x40
Memory Address Register       (0x1C): 0xC004
Sync Status Address Register  (0x1D): 0x40
Sys Status Register           (0x1E): 0x3
Sys Control Register          (0x1F): 0x3FDA

Indirect Access:
Pause Frame Sent Counter(L)(0xF000): 0x0
Pause Frame Sent Counter(H)(0xF001): 0x0
Pause Frame Recv Counter(L)(0xF002): 0x0
Pause Frame Recv Counter(H)(0xF003): 0x0
Auto Neg. Control Register (0xF004): 0x7
Tx Phy Addr Register-GMAC0 (0xF005): 0x0
Rx Uinfo Registerter-GMAC0 (0xF006): 0x0
Tx Phy Addr Register-GMAC1 (0xC005): 0xFFFF
Link Status Register           [3-0]: 0x1


Xpif Counters:

MAC Receive Counters:
Bytes                  =63848
pkt64                  =0
pkt65to127             =0
pkt128to255            =0
pkt256to511            =184
pkt512to1023           =0
pkt1024to1518          =0
pkt1519to1530          =0
good_giants            =0
error_giants           =0
good_runts             =0
error_runts            =0
ucast_pkts             =0
mcast_pkts             =184
bcast_pkts             =0
sync_loss_errs         =0
overruns               =0
```

```
   fcs_errs              =0
   delimiter_seq_errs    =0
   gmac_dropcounts       =0
   symbol_errs           =0

   MAC Transmit Counters:
   Bytes                 =31620
   pkt64                 =0
   pkt65to127            =0
   pkt128to255           =0
   pkt256to511           =93
   pkt512to1023          =0
   pkt1024to1518         =0
   pkt1519to1530         =0
   good_giants           =0

   Slicer Receive Counters:
   Cells                 =87293
   Frames                =23312
   Header Sequence Errors=0
   fcs_errs              =0
   Length                =0

   Slicer Transmit Counters:
   Cells                 =0
   Frames                =0

➔  Status Registers:
   Rx_gmac_status        =0004015C
   Tx_gmac_status        =00000154
   Rx_slicer_status      =00000003
   Tx_slicer_status      =00000000
   IPC fail count        =0

   Switch#
```

Step 2 Check the Chip Status Register field. It should match the link status, duplex mode, and speed shown in the previous **show interface** command.

If not, see the “Troubleshooting Half- or Full-Duplex Negotiation” section on page 11-21.

Follow these steps to troubleshoot the counters of the Gigabit Ethernet interface module physical interface:

Step 1 Use the **show controllers c8500 counters** command to check the Gigabit Ethernet interface module counters.

```
Switch# show controllers c8500 counters
Interface   Input     Runts Giants  Input       CRC  Frame Output      Output
    State   Packets                 Errors                 Packets     Errors
-----------------------------------------------------------------------------
G0/1/0  U   0           0     0     0           0     0     136972      0
G0/1/1  U   0           0     0     0           0     0     20          0
P1/0/0  AD  0           19600630     2271017     3     0     0           0
P2/0/0  AD  0           2     0     139         2     0     0           0
G2/0/1  AD  1           0     0     0           0     0     1           0
A3/0/0  AD  0           0     0     0           0     0     0           0
G3/0/1  AD  1           0     0     0           0     0     1           0
F9/0/0  U   14364       0     0     0           0     0     14367       0
F9/0/1  AD  1           0     0     0           0     0     1           0
F9/0/2  AD  1           0     0     0           0     0     1           0
F9/0/3  AD  1           0     0     0           0     0     1           0
F9/0/4  AD  1           0     0     0           0     0     1           0
F9/0/5  AD  1           0     0     0
.
(Information Deleted)
.
A12/0/0 AD  0           0     0     0           0     0     0           0
G12/0/1 AD  1           0     0     0           0     0     1           0
-----------------------------------------------------------------------------
AD - Admin Down, D - Down, F - Fail, U - Up

Switch#
```

Step 2 Check the Interface State field. It should indicate the interfaces are up.

Step 3 Check the Input Packets and Output Packets fields. The **show controllers c8500 counters** command should be entered at least twice. The counters in the Input Packets and Output Packets fields should be incrementing. This information can also be displayed using the **show interfaces** command.

Note The **clear counters** command does not clear the **show controllers c8500 counters** command display.

If you determine that the interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Troubleshooting ATM Uplink with Enhanced Gigabit Ethernet Interface Modules

The ATM Uplink interface module has one ATM uplink port and one Enhanced Gigabit Ethernet port, and is designed specifically for large-enterprise and service-provider networks.

Up to eight of the ATM Uplink with Enhanced Gigabit Ethernet Modules may be installed in a Catalyst 8540 chassis, providing eight ATM uplink ports and eight ports of nonblocking, wire-speed Gigabit Ethernet capacity in the core of an Enterprise network.

An example application of the ATM uplink is traffic from a LAN switch being aggregated at the Catalyst 8540 CSR and then passed to the ATM network over the ATM uplink. The Layer 3 enabled ATM uplink supports RFC 1483 (Multiprotocol Encapsulation over ATM), which provides for the mapping of Layer 3 addresses to ATM virtual circuits, and traffic shaping. Refer to the *Guide to ATM Technology* for additional information on RFC 1483.

The two ATM uplink with enhanced Gigabit Ethernet interface modules are the OC-3c and the OC-12c. The ATM OC-3c or OC-12c uplink with enhanced Gigabit Ethernet interface modules consist of two port adapters that are attached to a carrier module. The port adapters are not hot-swappable, but the interface module as a whole is hot-swappable. The ATM OC-3c uplink port adapter or the OC-12c uplink port adapter resides on the left side of the interface module, and the one-port enhanced Ethernet Gigabit port adapter resides on the right side. This combination provides an Ethernet port for connection to, or within, a LAN and an ATM uplink port to a metropolitan-area network (MAN).

The ATM OC-3c uplink port adapter supports 155-Mbps multimode or single-mode intermediate-reach fiber connections. It supports Fast EtherChannel operation, uses SC-type connectors, and has built-in ACL functionality. The OC-3c has 64K of routing table memory.

The ATM OC-12c uplink port adapter supports 622-Mbps multimode or single-mode intermediate-reach fiber connections. It supports Fast EtherChannel, SC-type connectors, and has built-in ACL functionality. The OC-12c has 64K or 256K of routing table memory.

Note The port adapters within the ATM OC-12c or OC-3c uplink with enhanced Ethernet interface modules must have matching routing table memory. As an example, if the ATM OC-12c uplink port adapter has 64K of routing table memory, the enhanced Gigabit Ethernet port adapter must have 64K of routing table memory for the interface module to function properly.

ATM Uplink Interface Module LEDs

Table 10-6 describes the LEDs used to confirm and troubleshoot the operation of interface modules. The LEDs on interface modules indicate the status of the modules and their ports.

Table 10-6 ATM OC-3c and OC-12c Uplink With Enhanced Gigabit Ethernet Interface Module LED Descriptions

LED	State	Description
Tx (Transmit)	Green	Port is transmitting a packet. Green for approximately 50 ms.
	Off	No signal is detected.

Table 10-6 ATM OC-3c and OC-12c Uplink With Enhanced Gigabit Ethernet Interface Module LED Descriptions (continued)

LED	State	Description
Rx (Receive)	Green	Port is receiving a packet. Green for approximately 50 ms.
	Off	No signal is detected.
Alarm	Red	This alarm LED indicates one of the following conditions: LOS[1], LOF[2], LOP[3], AIS-L[4], AIS-P[5], RDI-L[6], RDI-P[7], UNEQ-P[8], PLM-P[9], or cell delineation error.
	Off	No error.
C/D (Carrier Detect)	Green	Carrier detect signal.
	Off	No carrier detect signal is detected.

[1] LOS = Loss of signal

[2] LOF = Loss of frame

[3] LOP = Loss of pointer

[4] AIS-L = Line alarm indication signal

[5] AIS-P = Path alarm indication signal

[6] RDI-L = Line remote defect indication

[7] RDI-P = Path remote defect indication

[8] UNEQ-P = Path unequipped

[9] PLM-P = Path payload label mismatch

Displaying ATM Uplink Interface Module Configurations

To display the interface configuration, use the following commands:

Command	Purpose
show running-config interfaces atm *card/subcard/port*{*.sub-interface*}	Shows the status of the physical interface.
show interfaces {**atm** | **gigabitEthernet**} *card/subcard/port*	Shows the status of the physical interface.
show controllers {**atm** | **gigabitEthernet**} *card/subcard/port*	Shows the interface memory management and error counters.

Follow these steps to troubleshoot an ATM uplink physical interface:

Step 1 Use the **show running-config interface atm** c*ard/subcard/port* command to check the interface status and configuration.

```
Switch# show running-config interface atm 12/0/0
Building configuration...

Current configuration:
!
interface ATM12/0/0
 no ip address
 no ip mroute-cache
 no atm ilmi-keepalive
 sonet ais-shut
end

Switch#
```

Step 2 Use the **show running-config interface atm** *card/subcard/port.sub-interface* command to check the subinterface status.

```
Switch# show running-config interface atm3/0/0.800

Current configuration:
!
interface ATM3/0/0.800 point-to-point
 ip address 10.6.85.253 255.255.255.252
 no ip directed-broadcast
 atm Pvt. 800 0 800 aal5snap
end
```

Step 3 Use the **show interface atm** *card/subcard/port* command to check the interface status.

```
   Switch# show interface atm 3/0/0
➔  ATM3/0/0 is up, line protocol is up
     Hardware is epif_port_garfield, address is 0090.2141.b037 (bia 0090.2141.b037)
➔    MTU 4470 bytes, sub MTU 4470, BW 622000 Kbit, DLY 10 usec, rely 110/255, load 1/255
     Encapsulation ATM, loopback not set, keepalive not supported
     Half-duplex, Unknown Speed
     ARP type: ARPA, ARP Timeout 00:15:00
     Encapsulation(s): AAL5 AAL3/4, PVC mode
➔    8191 maximum active VCs, 1024 VCs per VP, 1 current VCCs
     VC idle disconnect time: 300 seconds
➔    Last input 00:00:09, output never, output hang never
     Last clearing of "show interface" counters never
     Queueing strategy: fifo
     Output queue 0/40, 0 drops; input queue 0/75, 0 drops
     5 minute input rate 0 bits/sec, 0 packets/sec
     5 minute output rate 0 bits/sec, 0 packets/sec
        32 packets input, 2820 bytes, 0 no buffer
        Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
        0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
        10 packets output, 1120 bytes, 0 underruns
        0 output errors, 0 collisions, 0 interface resets
        0 output buffer failures, 0 output buffers swapped out
```

Step 4 Check the status of the interface. If the ATM interface is down and the line protocol is down, begin checking for active alarms and defects.

Step 5 Check the MTU field. This indicates the largest number of bytes of “payload” data a frame can carry, not counting the frame's header and trailer. For an ATM interface, the MTU should be 4470 bytes.

Step 6 Check the Last input and Last output fields. They show the number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface.

Step 7 Use the **show controllers atm** *card/subcard/port* command to check the interface memory status and configuration.

```
  Switch# show controllers atm 12/0/0
  slot:  7/0  Controller-Type : XPIF ATM OC12 PM - 1 Port MM

  F000  chan0 chan1 chan2 chan3 sstr 1202
        0006  0006  0006  0006
  task0   11    11    11    11
  task1  5E0   5E0   5E0   5E0
  task2   11    11    11    11
  task3  5E0   5E0   5E0   5E0
  SMDR 0xFF78 SSTR 0x1200 SSMR 0x4002 EVER 0x3001
  SIMR 0x0000 MBXW 0x0000 MBXR 0x0000 SPER 0xF000

→ TX SAR (Production 1.0.7) is Operational;
→ RX SAR (Production 1.0.7) is Operational;

→ SAR Counters:
       tx_paks         0, tx_abort_paks         0, tx_idle_cells   2975800744
       rx_paks         0, rx_drop_paks          0, rx_discard_cells         0

  Xpif Counters:

→ MAC Receive Counters:
  Bytes                 =0
  pkt64                 =0
  pkt65to127            =0
  pkt128to255           =0
  pkt256to511           =0
  pkt512to1023          =0
  pkt1024to1518         =0
  pkt1519to1530         =0
  good_giants           =0
  error_giants          =0
  good_runts            =0
  error_runts           =0
  ucast_pkts            =0
  mcast_pkts            =0
  bcast_pkts            =0
  sync_loss_errs        =0
  overruns              =0
  fcs_errs              =0
  delimiter_seq_errs    =0
  gmac_dropcounts       =0
  symbol_errs           =0

→ MAC Transmit Counters:
  Bytes                 =0
  pkt64                 =0
  pkt65to127            =0
  pkt128to255           =0
  pkt256to511           =0
  pkt512to1023          =0
  pkt1024to1518         =0
  pkt1519to1530         =0
  good_giants           =0

→ Slicer Receive Counters:
  Cells                 =21037265
  Frames                =5386756
  Header Sequence Errors=0
  fcs_errs              =0
  Length                =0
```

```
→ Slicer Transmit Counters:
  Cells                 =0
  Frames                =0

→ Status Registers:
  Rx_gmac_status        =00000000
  Tx_gmac_status        =00000000
  Rx_slicer_status      =00000003
  Tx_slicer_status      =00000000

→ Interface Configuration Mode:
       ATM clock line; STS-12c; Line is admin shutdown

→ Sonet overhead:
  k1/k2 = 0/6
  s1s0 = 00, c2 = 0xCF, s1 = 0x0
→ Contents of Section trace buffer:
  LLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLL
  Contents of Path trace buffer:

→ Active Defects: None
→ Alarm reporting enabled for: SF SLOS SLOF B1-TCA B2-TCA PLOP B3-TCA

→ Active ATM Payload Defect: LCD-P


→ OC12 counters:
    b1     - # section BIP-8 errors
    b2     - # line BIP-8 errors
    b3     - # path BIP-8 errors
    ocd    - # out-of-cell delineation errors - not implemented
    g1     - # path FEBE errors
    z2     - # line FEBE errors
    chcs   - # correctable HEC errors
    uhcs   - # uncorrectable HEC errors

  b1:0, b2:0, b3:0, ocd:0
  g1:0, z2:0, chcs:0, uhcs:0

→ OC12 errored secs:
  b1:0, b2:0, b3:0, ocd:0
  g1:0, z2:0, chcs:0, uhcs:0
  lineAIS:0, lineRDI:0, pathAIS:0, pathRDI:0

→ OC12 error-free secs:
  b1:0, b2:0, b3:0, ocd:0
  g1:0, z2:0, chcs:0, uhcs:0

  phy_tx_cnt:0, phy_rx_cnt:0

→ BER thresholds:  SF = 10e-0  SD = 10e-0
→ TCA thresholds:  B1 = 10e-6  B2 = 10e-6  B3 = 10e-6

  Switch#
```

Step 8 Check the Interface Configuration Mode field. This field indicates the clock configuration and the administrative status of the interface.

Step 9 Check Sonet Overhead fields. These fields indicate the following:

- k1/k2—used for Automatic Protection Switching (APS)
- s1s0—(2 bits) not used by SONET, may need to be configured for SDH

- c3— The value extracted from the SONET path signal label byte (C2)
- S1—(1 byte) Synchronization status byte

Step 10 Check the Content of Path trace field. The path trace buffer is used to communicate information regarding the remote hostname, interface name/number, and IP address. This is a Cisco-proprietary use of the J1 (path trace) byte.

Step 11 Check the Active defects field. It indicates the currently configured alarms with defects and is a primary troubleshooting indicator.

Step 12 Check the Alarm reporting enabled field—It is a list of alarms for which you enabled reporting by entering the **pos report interface** command.

Step 13 Check the Active Defects field—It is a list of all currently active defects.

Step 14 Check the OC12 Counters field. If this number is incrementing, this indicates a problem in the network.

Check for any BIP(B1)/BIP(B2)/BIP(B3) (Bit interleaved parity) error reported.

- For B1, the bit-interleaved parity error report is calculated by comparing the BIP-8 code with the BIP-8 code extracted from the B1 byte of the following frame. Differences indicate that section-level bit errors have occurred.
- For B2, the bit-interleaved parity error report is calculated by comparing the BIP-8/24 code with the BIP-8 code extracted from the B2 byte of the following frame. Differences indicate that line-level bit errors have occurred.
- For B3, the bit-interleaved parity error report is calculated by comparing the BIP-8 code with the BIP-8 code extracted from the B3 byte of the following frame. Differences indicate that path level-bit errors have occurred.

Check the FEBE (Far end block errors).

- Line far-end block errors (accumulated from the M0 or M1 byte) are reported when the downstream LTE detects BIP(B2) errors.
- Path far-end block errors (accumulated from the G1 byte) are reported when the downstream PTE detects BIP(B3) errors.

Step 15 Check the OC12 error secs field. This field shows the total seconds where there were one or more alarms since the switch was rebooted.

Check AIS (Alarm indication signal).

- The line alarm indication signal is sent by the section terminating equipment (STE) to alert the downstream line terminating equipment (LTE) that an LOS or LOF defect has been detected on the incoming SONET section.
- The path alarm indication signal is sent by the LTE to alert the downstream path terminating equipment (PTE) that it has detected a defect on its incoming line signal.

Check RDI (Remote defect indication).

- The line remote defect indication is reported by the downstream LTE when it detects LOF, LOS, or AIS.
- The path remote defect indication is reported by the downstream PTE when it detects a defect on the incoming signal.

Step 16 Check the OC12 error free secs field. It indicates the number of seconds since the last error.

Step 17 Check the BER thresholds field. It is a list of bit error rate (BER) thresholds that have been crossed.

Step 18 Check the TCA thresholds field. It is a list of threshold crossing alarms (TCA).

If you determine that the interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Troubleshooting Packet-over-SONET Uplink with Enhanced Gigabit Ethernet Interface Modules

The Packet-over-SONET (POS) uplink with enhanced Gigabit Ethernet interface module consists of two port adapters that are attached to a carrier module. The port adapters are not hot-swappable, but the interface module as a whole is hot-swappable. The Packet-over-SONET OC-12c uplink port adapter resides on the left side of the interface module, and the one-port enhanced Gigabit Ethernet port adapter resides on the right side. This combination provides an Ethernet port for connection to, or within, LANs, and a POS uplink port for connection to an ISP or MAN.

The Packet-over-SONET OC-12c uplink port adapter supports 622-Mbps single-mode intermediate and long-reach fiber connections. The Packet-over-SONET OC-12c uplink port adapter is a serial link, uses SC-type connectors, and has built-in ACL functionality. It is available with 64K or 256K of memory. Routing tables use this memory.

Note The port adapters within the Packet-over-SONET OC-12c interface module must have matching routing table memory. As an example, if the Packet-over-SONET OC-12c POS port adapter has 64K of routing table memory, the enhanced Gigabit Ethernet port adapter must have 64K of routing table memory for the interface module to function properly.

For detailed Cisco Packet over SONET/SDH (POS) technology information, see the following documents:

- *White Paper, Cisco's Packet over SONET/SDH (POS) Technology, Support; Mission Accomplished*
- *SONET Tech Tips*

Packet-over-SONET Uplink Interface Module LEDs

Table 10-7 describes the LEDs used to confirm and troubleshoot the operation of interface modules. The LEDs on interface modules indicate the status of the modules and their ports.

Table 10-7 Packet-over-SONET OC-12c Uplink With Enhanced Gigabit Ethernet Interface Module LED Descriptions

LED	State	Description
Rx (Receive)	Green; otherwise, it is off	Port is receiving a packet. Green for approximately 50 ms.
Tx (Transmit)	Green; otherwise, it is off	Port is transmitting a packet. Green for approximately 50 ms.
C/D (Carrier Detect)	Green	Carrier detect signal is received.
	Off	Carrier detect signal is not received.

Table 10-7 Packet-over-SONET OC-12c Uplink With Enhanced Gigabit Ethernet Interface Module LED Descriptions (continued)

LED	State	Description
Alarm	Red	This alarm LED indicates one of the following: LOS, LOF, LOP, AIS-L, AIS-P, RDI-L, RDI-P, UNEQ-P, or PLM-P.
	Off	No error.

Displaying POS Interface Module Configurations

To display the interface configuration, use the following commands:

Command	Purpose
show interfaces pos *card/subcard/port*	Shows the status of the physical interface.
show controllers pos *card/subcard/port*	Shows the interface memory management and error counters.

Follow these steps to troubleshoot the physical interface:

Step 1 Use the **show interfaces pos** *card/subcard/port* command to check the configuration.

```
Switch# show interfaces pos 3/0/0
POS3/0/0 is administratively down, line protocol is down
  Hardware is Packet Over SONET
  MTU 4470 bytes, BW 622000 Kbit, DLY 100 usec, rely 255/255, load 1/255
  Encapsulation HDLC, crc 32, loopback not set, keepalive not set
  Scramble enabled
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 386 bytes, 0 no buffer
     Received 0 broadcasts, 2 runts, 0 giants, 0 throttles
              0 parity
     3482907 input errors, 2 CRC, 0 frame, 3482903 overrun, 0 ignored, 0 abort
     0 packets output, 0 bytes, 0 underruns
     0 output errors, 0 applique, 3 interface resets
     0 output buffer failures, 0 output buffers swapped out
     0 carrier transitions
Switch#
```

Step 2 Check the POS field. If the link is down/down, start checking for active alarms and defects. Troubleshooting here is similar to serial interface troubleshooting.

Step 3 POS defects and alarms are similar to alarms occurring when troubleshooting and diagnosing T1/E1 and T3/E3 connections (for example, LOS, LOF, and AIS). For T1 connection troubleshooting procedures, refer to the *T1 Troubleshooting* at the following URL: http://www.cisco.com/warp/public/116/t1_flchrt_main.html

---

If you determine that the interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Follow these steps to troubleshoot the POS physical interface:

Step 1 Use the **show controllers pos** *card/subcard/port* command to continue checking the memory configuration.

```
  Switch# show controllers pos 3/0/0
  Interface POS3/0/0
  Hardware is Packet Over SONET, One-port OC12, Single Mode Intermediate Reach

  POS3/0/0
  SECTION
→    LOF = 1          LOS = 0                             BIP(B1) = 92
  LINE
→    AIS = 0          RDI = 0          FEBE = 342        BIP(B2) = 1179
  PATH
→    AIS = 0          RDI = 0          FEBE = 38         BIP(B3) = 52
→    LOP = 0
→    PLM-P = 1        UNEQ-P = 0

→ Active Alarms:  None
→ Active Defects: PLM-P
→ Alarm reporting enabled for: SF SLOS SLOF B1-TCA B2-TCA PLOP B3-TCA

→ Framing: SONET
  APS
    COAPS = 0          PSBF = 0
    State: PSBF_state = False
    Rx(K1/K2): 00/00  Tx(K1/K2): 00/00
    S1S0 = 0x00, C2 = 0x16
→ PATH TRACE BUFFER: UNSTABLE
    Remote hostname :
    Remote interface:
    Remote IP addr  :
    Remote Rx(K1/K2):   /    Tx(K1/K2):   /

→ BER thresholds:  SF = 10e-3  SD = 10e-6
→ TCA thresholds:  B1 = 10e-6  B2 = 10e-6  B3 = 10e-6

    Clock source:  Configured: internal  Current: internal

  Last valid pointer from H1-H2:  0x20A
  ----- XPIF PCS -----

  F000  chan0 chan1 chan2 chan3 sstr 1202
        0016  0006  0006  0006
  task0   11    11    11    11
  task1  4D8   4D8   4D8   4D8
  task2   11    11    11    11
  task3  4D8   4D8   4D8   4D8
```

```
----- XPIF SLICER Registers -----
SMDR 0xFF78 SSTR 0x1202 SSMR 0x4002 EVER 0x3001
SIMR 0x0000 MBXW 0x0000 MBXR 0x0000 SPER 0xF000

Xpif Counters:

MAC Receive Counters:
Bytes                 =0
pkt64                 =0
pkt65to127            =0
pkt128to255           =0
pkt256to511           =0
pkt512to1023          =0
pkt1024to1518         =0
pkt1519to1530         =0
good_giants           =0
error_giants          =0
good_runts            =0
error_runts           =0
ucast_pkts            =0
mcast_pkts            =0
bcast_pkts            =0
sync_loss_errs        =0
overruns              =0
fcs_errs              =0
delimiter_seq_errs    =0
gmac_dropcounts       =0
symbol_errs           =0

MAC Transmit Counters:
Bytes                 =0
pkt64                 =0
pkt65to127            =0
pkt128to255           =0
pkt256to511           =0
pkt512to1023          =0
pkt1024to1518         =0
pkt1519to1530         =0
good_giants           =0

Slicer Receive Counters:
Cells                 =89486
Frames                =23980
Header Sequence Errors=0
fcs_errs              =0
Length                =0

Slicer Transmit Counters:
Cells                 =0
Frames                =0

Status Registers:
Rx_gmac_status        =00000000
Tx_gmac_status        =00000000
Rx_slicer_status      =00000003
Tx_slicer_status      =00000000


Switch#
```

Note The numbers under the Section and Line are accumulators and tell you the number of times the condition has occurred, not if it is currently happening.

Step 2 Check LOF (loss of frame)—LOF is detected when a severe error framing (SEF) defect on the incoming SONET signal persists for 3 ms.

Step 3 Check LOS (loss of signal)—LOS is detected when an all-zeros pattern on the incoming SONET signal lasts 19 plus or minus 3 ms or longer. This defect might also be reported if the received signal level drops below the specified threshold.

Step 4 Check whether a Bit interleaved parity (BIP [B1]/BIP [B2]/BIP [B3]) error has been reported.

- For B1, the bit interleaved parity error report is calculated by comparing the BIP-8 code with the BIP-8 code extracted from the B1 byte of the following frame. Differences indicate that section level bit errors have occurred.
- For B2, the bit interleaved parity error report is calculated by comparing the BIP-8/24 code with the BIP-8 code extracted from the B2 byte of the following frame. Differences indicate that line level bit errors have occurred.
- For B3, the bit interleaved parity error report is calculated by comparing the BIP-8 code with the BIP-8 code extracted from the B3 byte of the following frame. Differences indicate that path level bit errors have occurred.

Step 5 Check the Alarm indication signal (AIS) field.

- Line alarm indication signal is sent by the section terminating equipment (STE) to alert the downstream line terminating equipment (LTE) that a LOS or LOF defect has been detected on the incoming SONET section.
- Path alarm indication signal is sent by the LTE to alert the downstream path terminating equipment (PTE) that it has detected a defect on its incoming line signal.

Step 6 Check the Remote defect indication (RDI) field.

- Line remote defect indication is reported by the downstream LTE when it detects LOF, LOS, or AIS.
- Path remote defect indication is reported by the downstream PTE when it detects a defect on the incoming signal.

Step 7 Check the Far end block errors (FEBE) field.

- Line far end block error (accumulated from the M0 or M1 byte) is reported when the downstream LTE detects BIP(B2) errors.
- Path far end block error (accumulated from the G1 byte) is reported when the downstream PTE detects BIP(B3) errors.

Step 8 Check the loss of pointer (LOP) Path field —LOP is reported as a result of an invalid pointer (H1, H2) or an excess number of new data flag (NDF) enabled indications.

Step 9 Check the NEWPTR Inexact count field for of the number of times the SONET framer has validated a new SONET pointer value (H1, H2).

Step 10 Check the PSE Inexact count field for of the number of times the SONET framer has detected a positive stuff event in the received pointer (H1, H2).

Step 11 Check the NSE Inexact count field for of the number of times the SONET framer has detected a negative stuff event in the received pointer (H1, H2).

Step 12 Check the Active Alarms field—It is a list of current Alarms as enforced by Sonet Alarm Hierarchy.

Step 13 Check the Active Defects field—It is a list of all currently active SONET defects.

Step 14 Check the Alarm reporting enabled field—It is a list of alarms that you enabled reporting for with the **pos report** interface command.

Step 15 Check the COAPS fields—These are an inexact count of the number of times a new APS value has been detected in the K1 and K2 bytes. These fields indicate the following:

- k1/k2—used for Automatic Protection Switching (APS)
- s1s0—(2 bits) not used by SONET, may need to be configured for SDH

Step 16 Check the PSBF field—It is an inexact count of the number of times a protection switching byte failure has been detected (no three consecutive SONET frames contain identical K1 bytes).

Step 17 Check the PSBF_state field—It lists protection switching byte failure state.

Step 18 Check the Rx(K1/K2)/Tx(K1/K2) field—It lists contents of the received and transmitted K1 and K2 bytes.

Step 19 Check the S1S0 field—It lists the two S bits received in the last H1 byte.

Step 20 Check the C2 field—It lists the value extracted from the SONET path signal label byte (C2).

Step 21 Check the PATH TRACE BUFFER field—It lists the SONET path trace buffer is used to communicate information regarding the remote hostname, interface name/number, and IP address. This is a Cisco-proprietary use of the J1 (path trace) byte.

Step 22 Check the BER thresholds field—It list of the bit-error rate (BER) thresholds you configured with the **pos threshold interface** command.

Step 23 Check the TCA thresholds field—It list of threshold crossing alarms (TCA) you configured with the **pos threshold interface** command.

If you determine that the interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Troubleshooting ACL Daughter Card

The access control list (ACL) daughter card implements data-plane access lists in hardware, providing high-speed performance. This extends the traffic control and security capabilities of the Catalyst 8540 beyond control-plane access lists, which are currently supported. See the “Comparing Data Plane and Control Plane Traffic” section on page 11-20.

The ACL daughter cards can be used with existing 10/100BASE-TX, 100BASE-FX, and Gigabit Ethernet interface modules on the switch router. The ACL daughter cards provide data-plane ACL functionality for both IP and IPX traffic.

The switch router supports control-plane access lists such as permit and deny IP and IPX routes and IPX Service Advertisement Protocol (SAP) filtering without the daughter card. The daughter card enables data-plane ACLs for IP and IPX traffic.

The ACL daughter card is a field-replaceable unit that can be mounted onto the following switch router interface modules:

- Two-port Gigabit Ethernet
- 10/100BASE-T Ethernet
- 100BASE-FX Ethernet

Note The eight-port Gigabit Ethernet interface module does not support the ACL daughter card. The enhanced Gigabit Ethernet interface modules have built-in ACL functionality.

The ACL daughter card allows you to create lists for network control and security that filter packet flow into or out of router interfaces.

Packet Flow through ACL Daughter Card

Following is a description of the packet flow through an ACL daughter card:

Step 1 Ethernet processor interface receives the packet.

Step 2 The appropriate information (for example, IP addresses, protocol, and port numbers) is extracted from the packet.

Step 3 The information described in Step 2 is passed to Access List Controller.

Step 4 The Access List Controller creates the ACL word for the ciscoCAM (132 bits).

Step 5 The ciscoCAM and associated RAM returns the access or deny bit and an index.

Step 6 The Ethernet processor interface accepts (forwards) or denies (drops) the packet.

Displaying ACL Daughter Card Configurations

To display the ACL daughter card interface module configuration, use the following commands:

Command	Purpose
show running-config interface {**fastethernet** | **gigabitethernet**} *card/subcard/port*	Displays the interface access list group configuration.
show access-lists {*list-name* | *list-name*}	Displays the access list configuration.
show epc acl lookup {**in** | **out** | **ipqos**} {**fastethernet** | **gigabitethernet**} *card/subcard/port* {*protocol*}*source-address destination-address*	Displays the ACL daughter card function.

Troubleshooting the ACL Daughter Cards

Follow these steps to troubleshoot the status of an ACL daughter card:

Step 1 Use the **show running-config interface** command to check the interface status and the access group enabled on the interface.

```
Switch# show running-config interface fastEthernet 11/0/0
Building configuration...

Current configuration:
!
interface FastEthernet11/0/0
 ip address 20.0.11.1 255.255.255.0
➔ ip access-group 110 in
 no ip directed-broadcast
end
```

This interface has access group 110 enabled.

Step 2 Use the **show access-lists** command to confirm the status and configuration of the access lists configured on the Layer 3 enabled ATM switch router.

```
Switch# show access-lists 110
➔ Extended IP access list 110
    permit ip host 20.0.11.1 host 20.0.11.2
    permit ip host 20.0.11.2 host 20.0.11.1
```

For detailed information about access list filters and their configuration, refer to the "Configuring IP Services" chapter in the *Cisco IOS IP and IP Routing Configuration Guide*.

Step 3 Use the **show epc acl lookup** command for a specific interface and IP addresses to display the response of the access list daughter card to a connection attempt from a source IP address to a destination.

In the following example, the packets are allowed to cross the switch router:

```
Switch# show epc acl lookup in fastEthernet 11/0/0 ip 20.0.11.1 20.0.11.2

Input IP ACL lookup on FastEthernet11/0/0:Label:1 Index:42
    DestIP:20.0.11.2 SrcIP:20.0.11.1 DestPort:0 SrcPort:0
    Proto:256 Precedence:0x0 TOS:0x0 TCPFLAGS:0x0
    ICMP type:0 code:0 IGMP type:0
Lookup Key:
 00000000 00000100 00001400 0B011400 0B020000 323A3337 45000000 2053756E 00000000
002A0001
TCAM Result:80420223 80400000
Lookup got hit at
[V:0x11C00000 M:0x12C00000][0 IP] permit ip host 20.0.11.1 host 20.0.11.2
➔ Packet will be permitted
```

In the following example, the packets are denied access to cross the switch router:

```
Switch# show epc acl lookup in fastEthernet 11/0/0 ip 20.0.11.1 20.0.11.3

Input IP ACL lookup on FastEthernet11/0/0:Label:1 Index:42
    DestIP:20.0.11.3 SrcIP:20.0.11.1 DestPort:0 SrcPort:0
    Proto:256 Precedence:0x0 TOS:0x0 TCPFLAGS:0x0
    ICMP type:0 code:0 IGMP type:0
Lookup Key:
 00000000 00000100 00001400 0B011400 0B030000 353A3439 45000000 2053756E 00000000 
002A0001
TCAM Result:80422441 00400010
Lookup got hit at
[V:0x11C00000 M:0x12C00000][0 IP] deny ip host 20.0.11.1 host 20.0.11.2

→ Packet will be denied
```

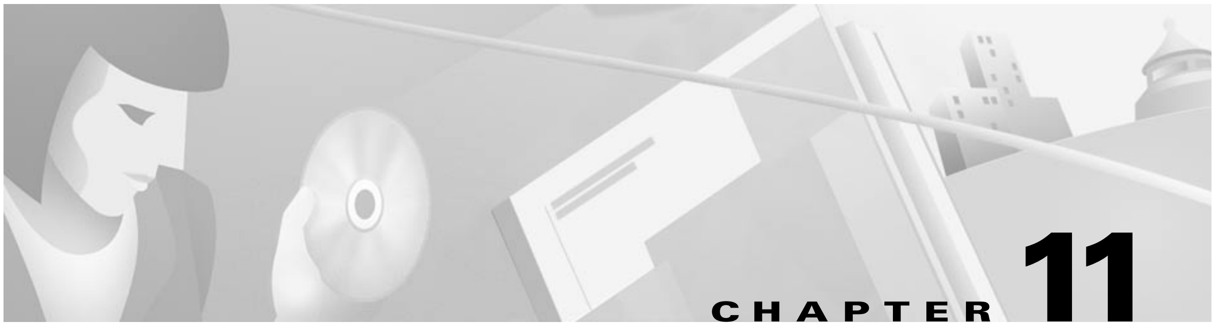

CHAPTER 11

Troubleshooting Layer 3 Network Connections

This chapter provides troubleshooting information about connectivity and performance problems in the Layer 3 network connections of the switch router.

The chapter includes the following sections:

Note For detailed cabling and hardware information for each port adapter, refer to the *Catalyst 8540 CSR Route Processor and Interface Module Installation Guide.*

Overview of Layer 3 Switching

This section provides an overview of Layer 3 switching using the switch router. It shows how a switch router fits into the network, the architecture of the switch router, and the course of a Layer 2 and Layer 3 packet through the switch router. Also included is a list of Layer 3 switching software features with brief descriptions of selected features.

Defining Layer 3 Switching

Layer 3 switching refers to a class of high-performance switch routers optimized for the campus LAN or intranet, providing both wirespeed Ethernet routing and switching services.

A Layer 3 switch router performs the following three major functions:

- Packet switching
- Route processing
- Intelligent network services

Compared to other routers, Layer 3 switch routers process more packets faster by using application-specific integrated circuit (ASIC) hardware instead of microprocessor-based engines. Layer 3 switch routers also improve network performance with two software functions—route processing and intelligent network services.

To simplify forwarding of the IP packets, route processing is usually executed during the initial call or session setup. At that point, the Layer 3 enabled ATM switch router determines the appropriate route, and forwards to the interfaces information describing the path to be used. In fact, data exchanged between the communicating source and destination end nodes may never need to flow to or through a conventional router.

Frame forwarding on subsequent packets in the same flow is performed using the Layer 3 switch functions at the line card. Once the route has been determined, all subsequent frames in the flow are simply switched or forwarded across the chosen path. This takes advantage of the high throughput and low latency characteristics of switching by enabling the traffic to bypass the route processor once a path calculation has been performed.

Understanding Packet Flow

Figure 11-1 shows and describes, in Steps 1 through 4, how the initial packet travels through the switch Layer 3 route processor to set up the network route.

Note When making Layer 3 switching decisions, the route processor does not reference the switch fabric, (that is, the PVC configuration). The interface map (where the switch maps an egress interface to a Broute VC) is programmed when the switch is booted up. At that time, the PVCs are automatically configured.

Figure 11-1 Phase 1 — Layer 3 Packet Flow

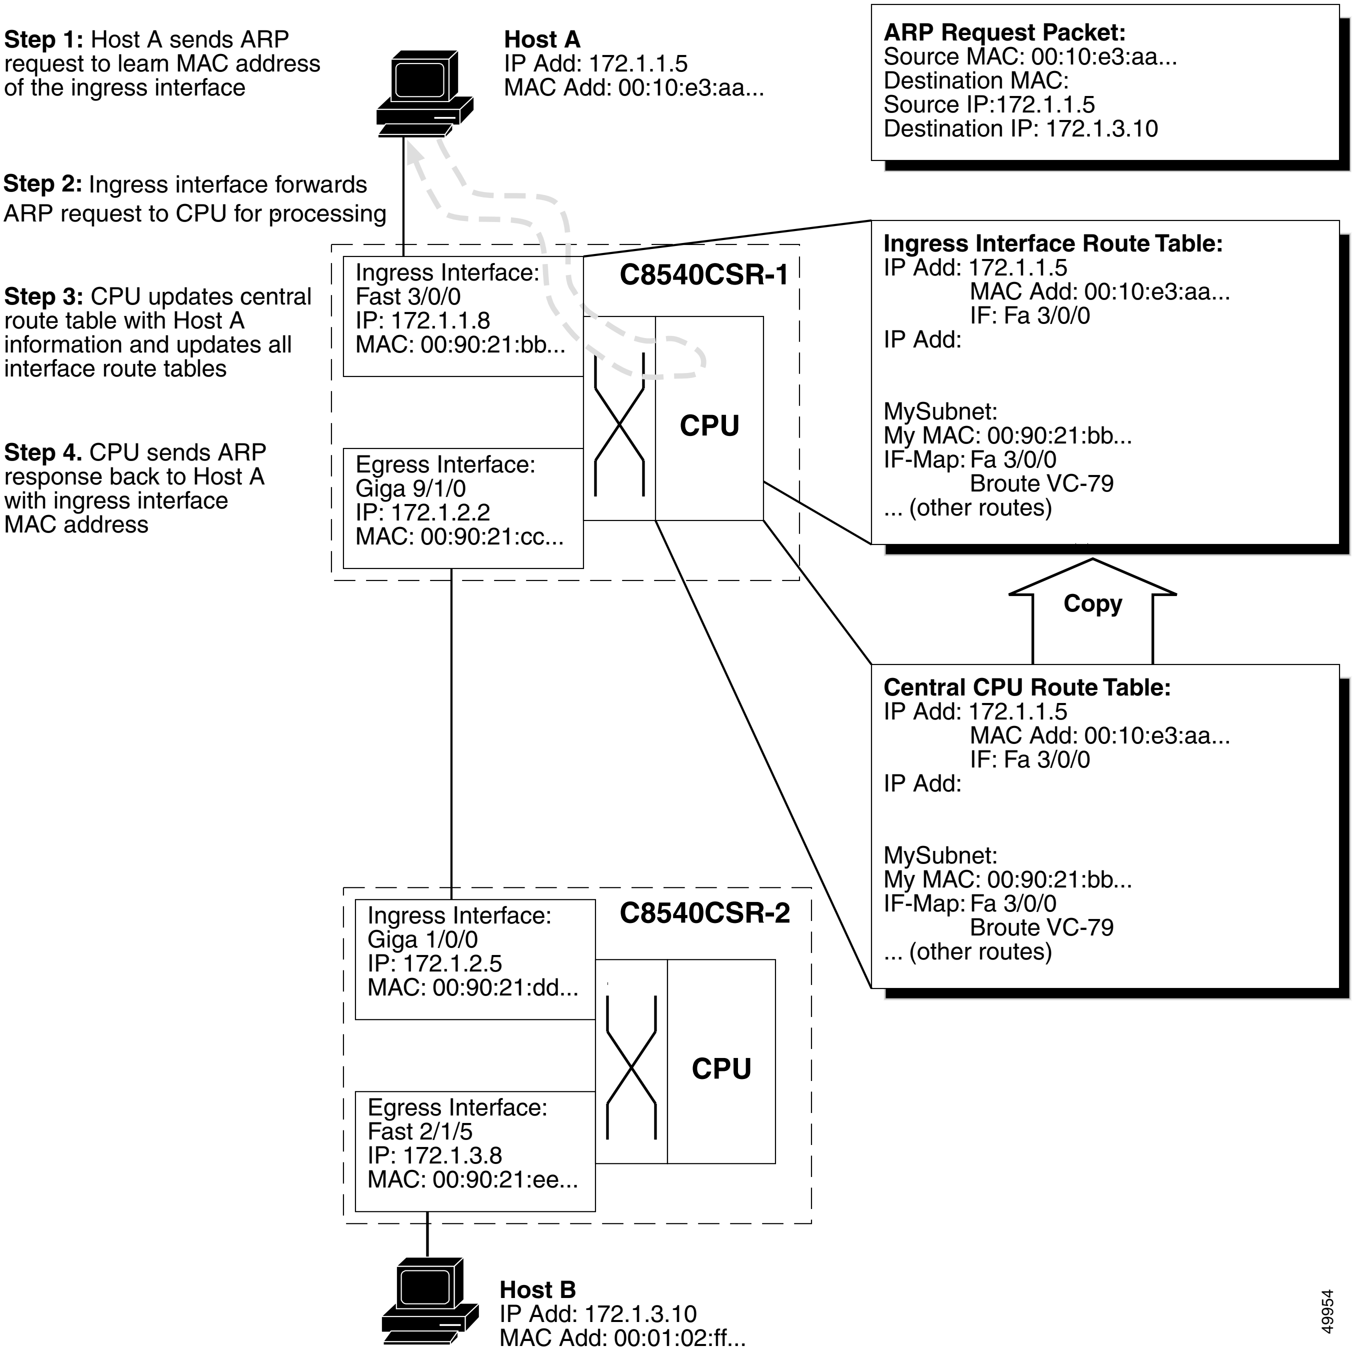

Figure 11-2 shows and describes, in Steps 5 through 7, how the route processor sends the ARP and propagates the updated routing tables to the interfaces.

Note In Figure 11-2, the ARP requests are described only for illustration purposes. In most cases, if you are running a dynamic protocol, the switches will have already sent and received ARP packets, and built the route tables.

Figure 11-2 Phase 2 — Layer 3 Packet Flow

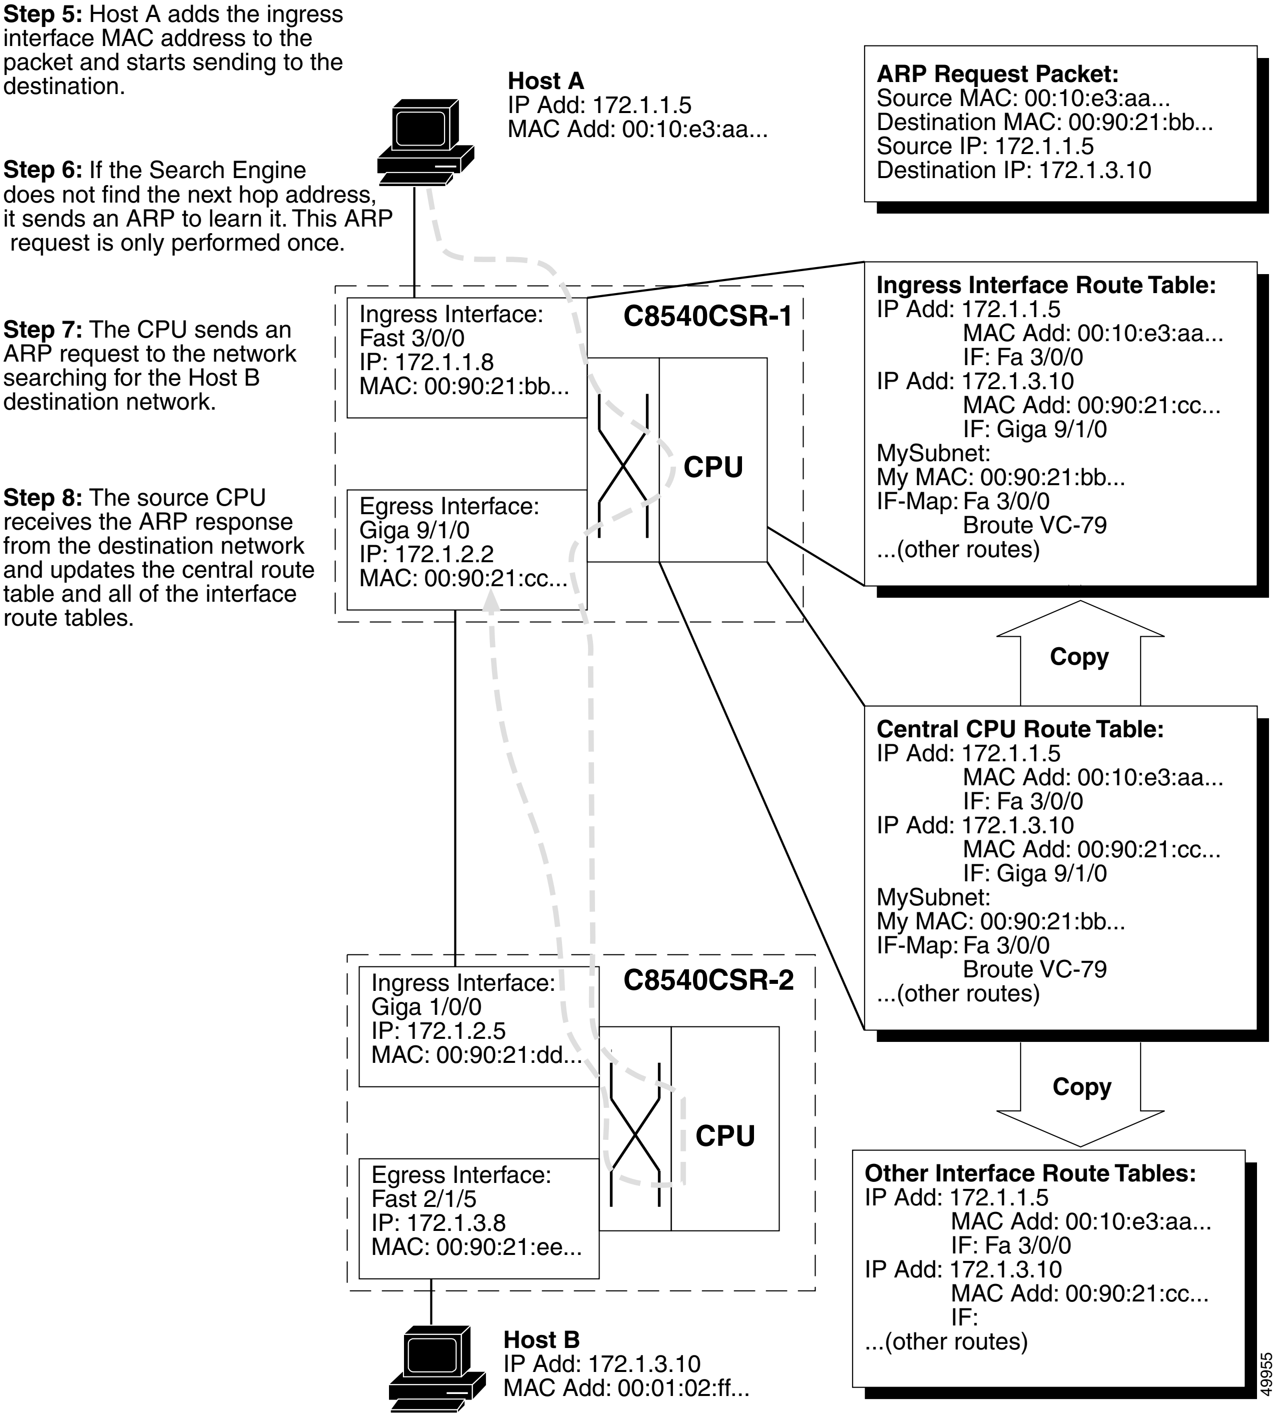

Figure 11-3 shows and describes, in Steps 8 and 9, how subsequent packets sent by Host A, to Host B, are switched without the help of the route processor.

Figure 11-3 Phase 3 — Layer 3 Packet Flow

Step 8: All subsequent packets received from Host A are Layer 3 switched using the ingress and egress route tables.

Step 9: When the packets reach the destination switch ingress interface the MAC address for Host B is rewritten to the destination MAC address and the packets are delivered.

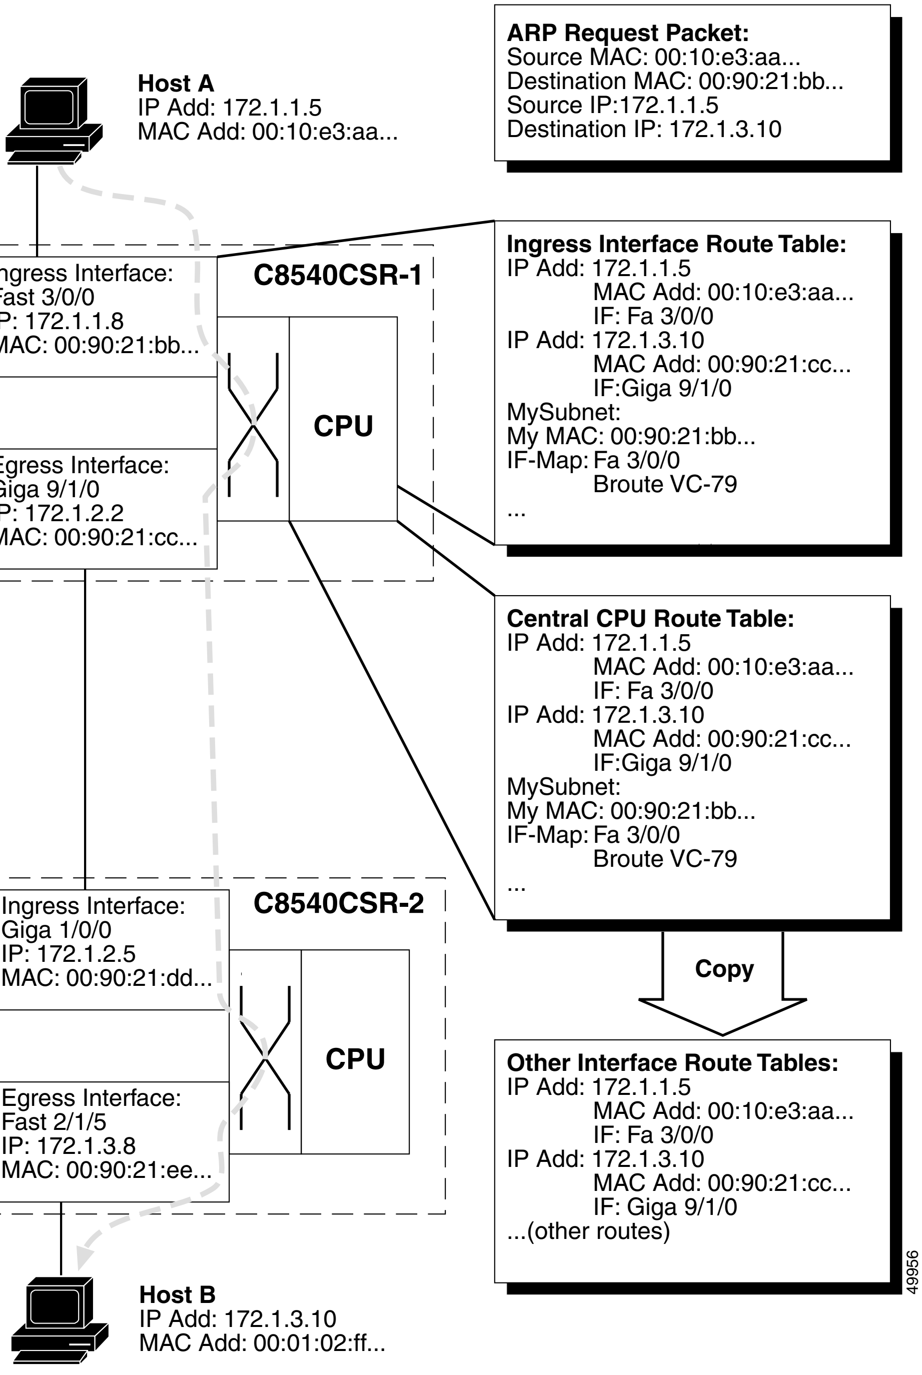

Layer 3 Forwarding

By using CEF, each of the line cards maintains a Forwarding Information Base (FIB) table downloaded from the switch processor. Any changes made to the route processor routing table, caused by additions or deletions of routes or route flaps, are updated in the central FIB, which in turn updates the line card FIBs. This means that, at all times, all line cards have a correct map of the network topology.

Packet switching in the Layer 3 enabled ATM switch router takes place as follows:

Step 1 A packet is received at the physical interface. The CEFA ASIC provides the MAC-layer functions, and the packet is stored in internal memory.

Step 2 As soon as the first 64 bytes of the frame are read, the microcode running on the microcontroller reads the source and destination IP addresses, or IPX network information. If the destination MAC address belongs to the switch router, the packet is routed. If not, it is bridged.

Step 3 The destination IP address information is used by the search engine to begin a lookup, in the CAM table, for the longest match entry.

Step 4 The destination network is matched within 64 clocks (or approximately 2.5 microseconds). The match is returned to the microcontroller, which in turn moves the frame from the internal memory to the Fabric Interface frame FIFO buffer. At the same time, the search engine returns relevant information such as quality of service (QoS) classifications, and MAC header rewrite information, to the Control FIFO buffer.

Step 5 Packet rewrite and QoS classifications take place at the input Ethernet processor interface or Cisco Express Forwarding ASIC (CEFA).

Step 6 The VPI and VCI are attached at the beginning of the packet. The VPI and VCI used corresponds to the particular QoS being requested. The packet then goes through the SAR (Segmentation and Reassembly), which segments the packet into 48-byte payloads. The previously retrieved VPI and VCI-value is written into the cell header to complete the 53-byte ATM Cell.

Step 7 As soon as the entire frame is received into the Frame FIFO buffer, the frame moves into the shared fabric and is stored with a pointer to the output port.

Step 8 If that output interface is currently busy transmitting a frame, the scheduler uses WRR to determine which packet should be sent next.

Step 9 The destination port is signaled, by the switching-fabric ASIC, to take the frame out of a known memory location. The destination port knows that it is receiving the correct frame because of the internal routing tag corresponding to a particular, internal, port-to-port circuit.

Step 10 The frame is sent out to the network.

Layer 2 Bridging

When a port or group of ports are running in bridging mode, the search engine initiates a lookup, in the CAM table, based on the Layer 2 MAC address. Because the Layer 3 enabled ATM switch router is a distributed switching system, each port (or in this case, CEFA) maintains a list of addresses and ports of exit that are of local significance. For example, if Address A is a destination learned on interface FastEthernet 0/0/1, the remaining interfaces on the switch do not have to have that address stored in their CAM tables unless they have a packet to send to Address A.

If the destination MAC address is a broadcast address (FFFF.FFFF.FFFF), the packet is tagged with the destination set as all ports in that bridge group, and it is sent out to the switching fabric. The fabric ASIC creates a pointer from that point in the memory to all ports in that bridge group. For example, if there were eight ports in a bridge group, all eight ports would receive that broadcast.

How MAC Addresses are Learned by the Switch

The following steps describe the MAC address learning process used by the switch router.

Step 1 When a port receives a packet with an unknown source and destination MAC address, it stores the source address as "locally learned" and forwards the packet, as an "unknown unicast," to all ports in the bridge group (similar to a broadcast).

Step 2 The receiving port also sends a LightStream InterProcess Communication (LSIPC) message to the route processor to allow it to update the bridging table on the route processor.

Note This bridging table in the route processor is only used to allow you check the learned MAC addresses using the **show bridge** command.

Step 3 All ports in the bridge group receive a copy of the "unknown unicast" and forward the packet.

Step 4 The receiving ports learn the new source address of the packet as a "remote entry."

Step 5 These receiving ports determine which interface sent the packet, based on the VPI and VCI header that points to a P2MP leaf, and the port already knows the corresponding P2MP root.

Step 6 Now all ports in the bridge port have learned the new source MAC address.

Step 7 The destination station for that frame responds.

Step 8 The port that receives the response learns the MAC address of the destination station (now the source address in the response). It has already learned the destination address, allowing it to forward the packet to the correct port.

Step 9 Only that egress port will then learn the new source address.

Step 10 The route processor is also notified of the new destination station source MAC address.

Step 11 Layer 2 switching then occurs between the two ports.

Note After 5 minutes of inactivity, MAC addresses are deleted from the CAM. The port sends another message to the route processor to remove the MAC from the bridging table.

After both the source and the destination MAC address have been learned, the following procedure occurs during Layer 2 frame switching:

Step 1 A packet is received at the physical interface. The CEFA ASIC provides the MAC-layer functions, and the packet is stored in internal memory.

Step 2 As soon as the first 64 bytes of the frame are read, the microcode running on the microcontroller reads the MAC source and destination addresses. If the destination MAC address is not that of the interface, Layer 2 switching is required. This information can now be used by the search engine.

Step 3 Because the packet has been received on a particular VLAN, the search engine begins a search for the MAC address and its corresponding port of exit.

Step 4 The destination MAC address is found. The microcontroller moves the frame from the internal memory to the switching fabric. At the same time, the search engine returns relevant information such as QoS classifications or ISL information to the switching fabric.

Step 5 The VPI and VCI are attached at the beginning of the packet. The VPI and VCI that are used correspond to the particular quality of service being requested, the appropriate port of exit. The packet then goes through the SAR (Segmentation and Reassembly), which segments the packet into 48-byte payloads. The previously retrieved VPI and VCI values are written into the cell header to complete the 53-byte ATM Cell.

Step 6 The frame moves into the shared fabric and is stored sequentially.

Step 7 The destination port is signaled by the switching-fabric ASIC to take the frame out of memory. The destination port knows that it is receiving the correct frame because of the internal routing tag.

Step 8 The frame is re-encapsulated via ISL, if necessary, and sent out to the network.

System Architecture

The best way to understand the architecture of the Layer 3 enabled ATM switch router is to divide the switch into the following three distinct, functional segments:

- Switch route processor
- Switch fabric
- Line cards

The switch route processor engine, show in Figure 11-4, is responsible for all address and route learning and distribution. Because the Layer 3 enabled ATM switch router is designed as a distributed switching system, the route processor (CPU) needs to ensure that all Layer 3 routes and Layer 2 MAC addresses are maintained and the line cards are updated as needed. The route processor is also responsible for handling all system management, including SNMP and remote monitoring (RMON) statistics.

Figure 11-4 High-Level Layer 3 Enabled ATM Switch Router Architecture

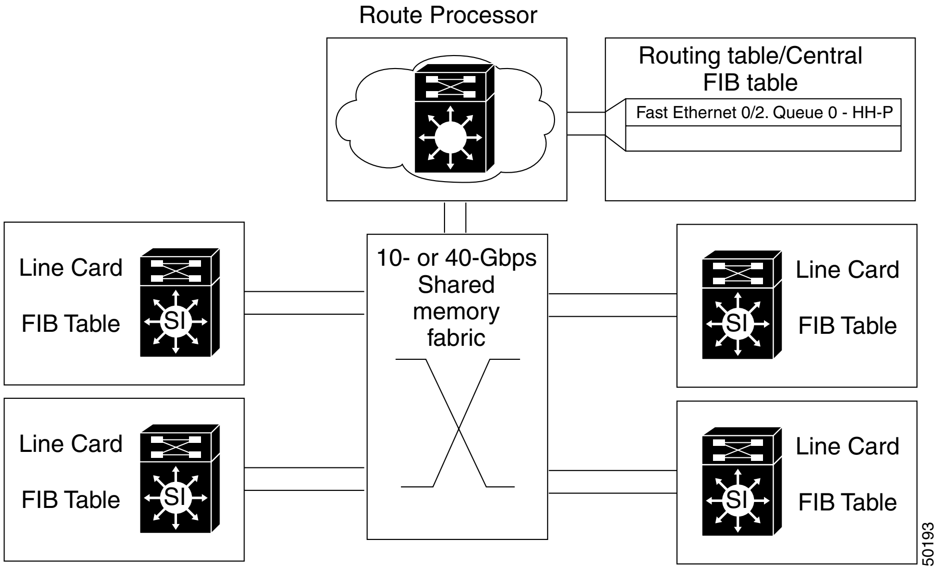
Route Processor
Routing table/Central FIB table
Fast Ethernet 0/2. Queue 0 - HH-P
Line Card
FIB Table
SI
10- or 40-Gbps Shared memory fabric
Line Card
FIB Table
Line Card
FIB Table
Line Card
FIB Table
50193

The switching fabric or shared memory fabric, show in Figure 11-4, differs for the two Catalyst 8500 CSR switches. The Catalyst 8540 includes 12-MB shared memory while the Catalyst 8510 includes 3-MB of shared-memory. This shared memory is dynamic, meaning that a packet stored in memory takes only as much memory as it needs. Access into and out of the shared memory is dynamically allocated by the direct memory access (DMA) ASIC. Because the switch fabric is nonblocking, it does not require per-port buffers; the fabric speed is faster than the combined speed of all the ports. Congestion, therefore, only occurs when an individual output port is congested.

The line cards, show in Figure 11-4, are designed to carry considerable intelligence for the switching system. Each line card contains ASICs designed to provide input and output into the fabric as well as to maintain a Layer 3 FIB or a Layer 2 MAC address table. These tables allow the Layer 3 enabled ATM switch router to make switching decisions very quickly prior to transmission across the switching fabric. The line cards, therefore, must work closely with the route processor to ensure that all address tables and routing information is current. The line cards are also responsible for presenting a uniform frame to the switching fabric for effective buffering, QoS policy enforcement, and packet switching.

Each of the three main components of the Catalyst 8540 CSR are described in detail in the following sections.

Route Processor

The system route processor is the first element of the Layer 3 enabled ATM switch router architecture and resides at the core of the switch. The route processor resides on the switch route processor (SRP) module, along with the shared memory fabric, described in the "Switching Fabric and Arbitration" section on page 11-13. The route processor for the Catalyst 8510 CSR is a 64-bit 100Mhz R4600 RISC processor. Its architecture is very similar to that of the Cisco 7500 Route Switch Processor (RSP). The route processor for the Catalyst 8540 CSR is a 200Mhz R5000 RISC processor, very similar to the RSP-4 engine. The Layer 3 enabled ATM switch router SRP runs the Cisco IOS Release 12.0 or later software.

Routing Protocols

The route processor is responsible for running all of the routing protocols shown in Table 11-1 on the Layer 3 enabled ATM switch router. Other protocols, such as AppleTalk, DECNet, and VINES are bridged in the switch.

Table 11-1 Supported Routing Protocols

IP Networks	IPX Networks	AppleTalk Networks
RIP	IPX RIP	RTMP
RIP-2	EIGRP	EIGRP
OSPF		AURP
IGRP		
EIGRP		
BGP		

Note The Catalyst 8540 CSR is designed to support multiprotocol routing.

Most importantly, the route processor is responsible for maintaining the routing table. By using Cisco Express Forwarding, the route processor creates a FIB, which contains a subset of the routing table. The FIB is based on a topology map of the network, allowing routing to take place via the network topology at high speed. The FIB is then downloaded to the line cards, allowing the line cards to make Layer 3 routing decisions without having to interrupt the route processor. This capability allows the Layer 3 enabled ATM switch router to forward all frames at wire speed for all ports. The FIB and Cisco Express Forwarding are also described in the "Line Card Architecture" section on page 11-16.

The route processor is also responsible for maintaining state information regarding multicast routing. The Layer 3 enabled ATM switch router supports PIM (sparse mode and dense mode) as well as Distance Vector Multicast Routing Protocol (DVMRP) interoperability. The route processor is responsible for responding to and forwarding joins and leaves as well as responding to pruning messages sent by PIM. Multicast forwarding takes place at the line card level.

Layer 2 VLAN and Switching

Although the switching decisions are made at the line cards, the route processor is still responsible for maintaining Layer 2 information. The route processor is responsible for bridge group configuration and spanning tree calculation.

Bridge groups are configured on the Layer 3 enabled ATM switch router in the same way they are in other Cisco routers. Instead of routing traffic to an outgoing interface, the traffic is bridged via its Layer 2 address. Integrated Routing and Bridging (IRB) is also supported in the Layer 3 enabled ATM switch router in order to support both bridging and routing at the same time.

Spanning tree information within the switch is maintained by the route processor. This includes calculation of the root bridge, optimum path determination to the root, and determining the forwarding and blocking links.

Cisco Express Forwarding

Cisco Express Forwarding (CEF) evolved to best accommodate the changing network dynamics and traffic characteristics resulting from increasing numbers of short-duration flows typically associated with Web-based applications and interactive multimedia sessions. Other Layer 3 switching paradigms use a route-cache model to maintain a fast lookup table for destination network prefixes (see Figure 11-5). The route-cache entries are traffic driven, in that the first packet to a new destination is routed via routing table information, and as part of that forwarding operation, a route-cache entry for that destination is added. This process allows subsequent packet flows to that same destination network to be switched based on a route-cache match. These entries are periodically aged out to keep the route cache current and can be immediately invalidated if the network topology changes.

Figure 11-5 Route-Cache and Distributed Routing Comparison

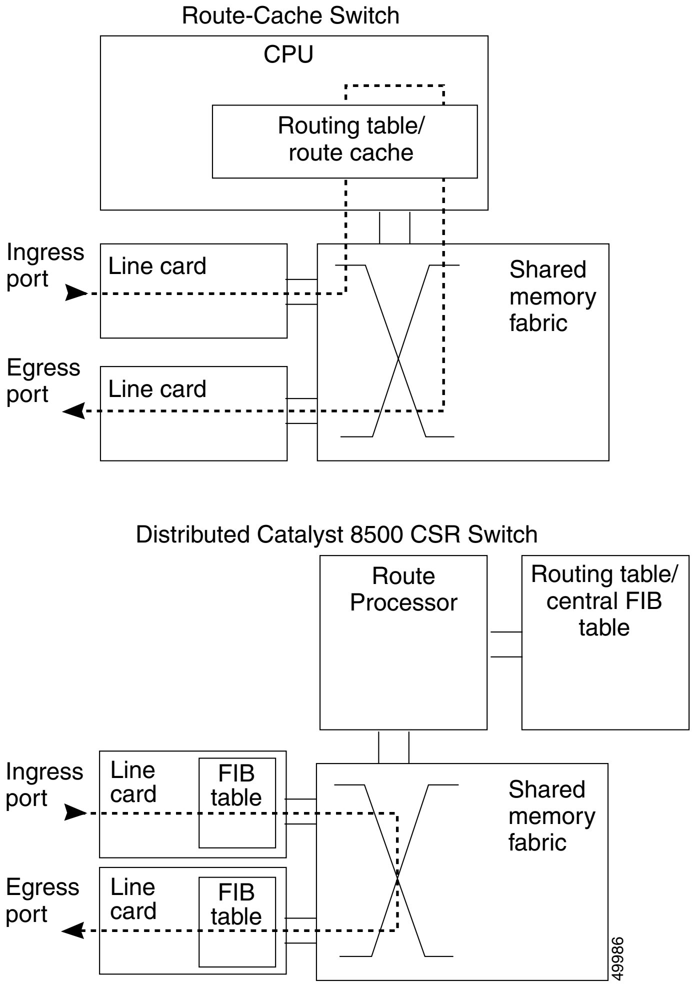

This “demand-caching” scheme used by other Layer 3 switches is optimized for networks where the majority of traffic flows are associated with a subset of destinations. Since the traffic profiles at the core of the Internet (and potentially within some large enterprise networks) no longer resemble this model, CEF was introduced. CEF eliminates the increasing cache maintenance problem resulting from growing numbers of topologically dispersed destinations and dynamic network changes.

CEF avoids the potential overhead of continuous cache churn by using a FIB on the line card for the destination switching decision. The FIB mirrors the entire contents of the IP and IPX routing table. This means that there is a one-to-one correspondence between FIB table entries and routing table prefixes; therefore, a route cache does not need to be maintained.

Note Although CEF has been specified for IP, it also applies to IPX as well.

CEF Operation

CEF provides features comparable to fast switching, including load sharing, recursive route resolution, and access lists. CEF uses two tables maintained in the SRP and downloaded to the line cards: the FIB and adjacency table. The FIB table is used for making forwarding decisions. The adjacency table maintains the adjacent nodes, and the link-layer information (such as packet rewrite information) necessary to reach that adjacent node. Every entry in the FIB table has a pointer to a corresponding entry in the adjacency table shown in Figure 11-6.

Figure 11-6 FIB and Adjacency Table

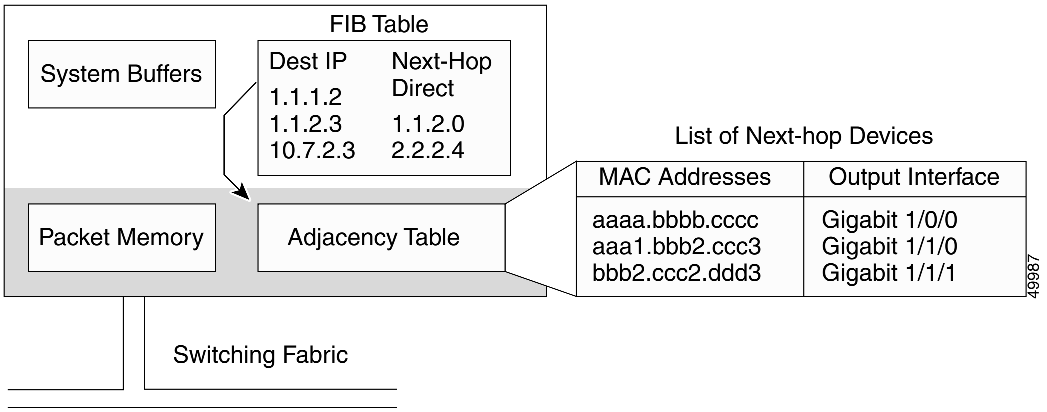

The FIB table is populated by callbacks (inputs) from the routing table. After a route is resolved, it points to a next hop, which should be an adjacency. This step is done at the SRP and then downloaded to the line cards, allowing the line cards to maintain a current topology of the network, which enables rapid switching decisions (within 10 ms) as well as fast convergence in the event of a routing topology change. The FIB is modified when a route is added, removed, or changed in the routing table. This information is immediately downloaded to the line cards.

The adjacency table is also populated by callbacks from the routing protocols, which include information such as next-hop information and (source, group [S,G]) interfaces for multicast groups. Adjacencies are added when a protocol detects that there is an adjacent node via the routing protocol. When a packet arrives at the ingress port, the CEF ASIC performs a FIB lookup based on the destination IP address. The matching FIB entry points to an adjacency entry, which in turn provides the valid link layer rewrite and outgoing interface. The packet is forwarded based on this information. Figure 11-6 shows the relation of the FIB to the adjacency table.

Switching Fabric and Arbitration

The Catalyst 8540 and Catalyst 8510 CSRs have different shared-memory architecture and system bandwidth. The Catalyst 8540 is based on a 12-MB shared-memory architecture with a total system bandwidth of 40 Gbps. The Catalyst 8510 is based on a 3-MB shared-memory architecture with a total system bandwidth of 10 Gbps. Both systems shared memory is completely nonblocking, meaning that all input ports have equal and full access into the shared memory for packet switching. The Layer 3 enabled ATM switch router also provides four queues per port, allowing the Frame Scheduler to make intelligent QoS decisions based on the priority of each queue.

In the Catalyst 8540, each line card has 5-Gbps access into the shared memory fabric as shown in Figure 11-7. This bandwidth is also divided into 2.5 Gbps transmit and 2.5 Gbps receive paths into the fabric. This allows for nonblocking switching capacity within the switching system by ensuring that each line card is given more bandwidth than all of the ports on the line card can generate. Each of the line cards in the Catalyst 8510 is allotted 2.5 Gbps of capacity into the fabric. The 2.5-Gbps bandwidth is divided into transmit and receive paths, each of 1.25 Gbps, to ensure that both reads and writes to the shared memory can be accomplished simultaneously.

Figure 11-7 Switching Bandwidth per Slot on Catalyst 8540 CSR

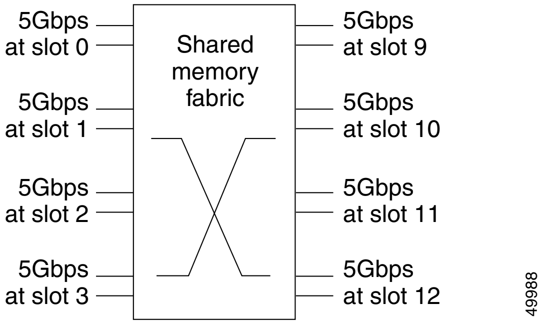

Because the Layer 3 enabled ATM switch router includes nonblocking memory, every port in the switch has full access to every other port. Each packet entering the switch fabric is tagged with an internal routing tag. This routing tag provides the switching fabric with the appropriate port of exit information, the QoS priority queue the packet is to be stored in, and the drop priority, shown in Figure 11-8.

Figure 11-8 Internal Routing Label Format

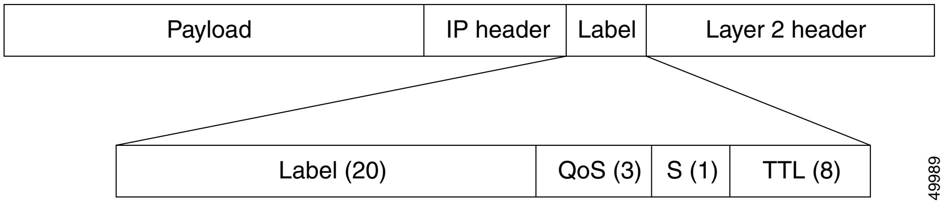

The 4 byte routing tag contains a 20-bit label value, a 3-bit QoS value, a 1-bit stack indicator, and an 8-bit TTL value.

The Fabric-Switching ASIC (FSA) then queues each packet into memory and creates a pointer, based on the internal routing tag, to the appropriate destination port. The Frame Scheduler is then responsible for scheduling the frame out of memory based on the queue where the packet is being stored.

Each port transmitting through the fabric is, by default, placed in the lowest-priority queue. This places all traffic at a "best-effort" QoS level. When you configure a policy, that traffic is transmitted in the queue corresponding to the specified IP precedence. That queue is granted more service, thereby reducing latency and the possibility that traffic on that queue will be dropped.

Note All management and control plane traffic, such as BDPU information, routing protocol updates, and management frames are placed in the highest-priority queue for transmission to the route processor.

The Frame Scheduler

The Frame Scheduler has two main responsibilities within the Layer 3 enabled ATM switch router: first, to schedule frames into the switching fabric based on the priority queue being requested, and second, to schedule frames out of the switching fabric based on the Weighted Round Robin (WRR) scheduling algorithm.

At the input to the switching fabric, the CEF ASIC posts a request to the Frame Scheduler for access to the fabric. The Frame Scheduler handles each request in a time-division multiplexing (TDM) fashion, meaning that each CEF ASIC will have the opportunity to clock an entire frame into the fabric when access has been granted. Because each CEF ASIC handles four ports, the Frame Scheduler allows the CEF ASIC to clock in a maximum of four packets into memory (see the "CEFA" section on page 11-17).

Each packet in memory has an internal routing tag added to the beginning which, as mentioned earlier, contains the port of exit, queuing priority, and drop priority. Based on the routing tag, the input Frame Scheduler places the packet in the correct queue (see Figure 11-9).

Figure 11-9 Input Scheduling and Queue Allocation

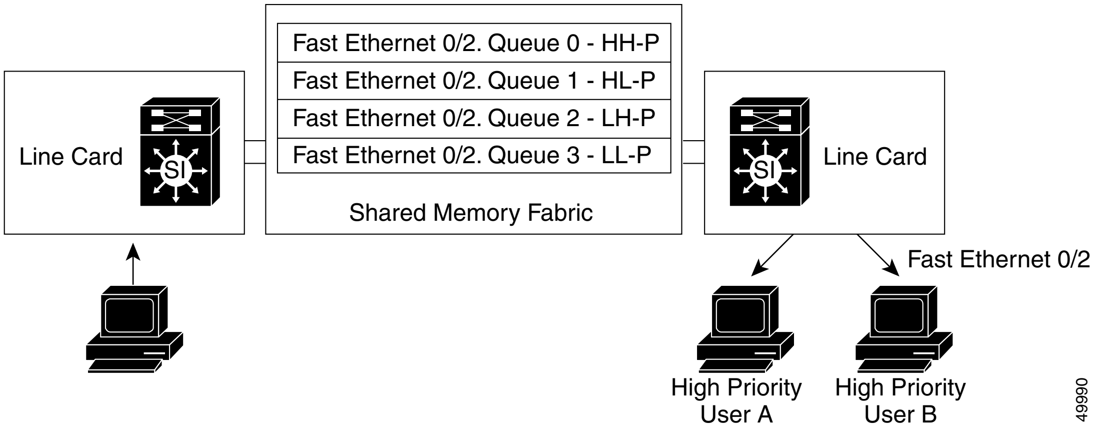
Fast Ethernet 0/2. Queue 0 - HH-P
Fast Ethernet 0/2. Queue 1 - HL-P
Fast Ethernet 0/2. Queue 2 - LH-P
Fast Ethernet 0/2. Queue 3 - LL-P
Shared Memory Fabric
Line Card
SI
Line Card
SI
Fast Ethernet 0/2
High Priority User A
High Priority User B
49990

The "HH," "HL," "LH," and "LL" designations refer to the IP precedence fields used by the Layer 3 enabled ATM switch router to determine the appropriate queue.

Note Although not shown, a fifth, critical high-priority routing tag is added to the beginning of all management and control plane packets for immediate delivery to the route processor.

On the output side, the Frame Scheduler is responsible for servicing each queue based on the WRR priority scheme. WRR allows the network manager to configure how much service each queue receives. In a situation where there is no congestion, WRR and the weights provided do not play a real part on how packets are switched out of the fabric, because there is plenty of bandwidth available. However, if a link is congested, WRR services each queue per port based on the priority set by the network manager. For example, look at the weights assigned by a network manager in Table 11-1.

Table 11-1 Sample WRR Priority Weights

Quality of Service Priority	Weight Given by Network Manager	Bandwidth Assignment Calculation	Bandwidth Assigned
QoS-0	8	=(8/(8+4+2+1)) x 100	53 Mbps
QoS-1	4	=(4/(8+4+2+1)) x 100	27 Mbps
QoS-2	2	=(2/(8+4+2+1)) x 100	13 Mbps
QoS-3	1	=(1/(8+4+2+1)) x 100	7 Mbps

Based on the priorities and weights provided, the Frame Scheduler services QoS-0 more often, granting queue 53 Mbps out of the 100 Mbps possible on the output link. The second queue, QoS-1, receives 27 Mbps of the bandwidth, and so forth. These commands are set globally on the switch router and function the same for all ports on the switch.

The switch router also allows you to override the global QoS settings by allowing port-to-port communications to have a different level of priority. You have the option of configuring bandwidth based on a source-destination, destination, or source basis and provide weights based on certain IP addresses having more bandwidth then others.

Note This feature is available with the hardware access list daughter card installed on an Ethernet interface module installed in the Catalyst 8510 CSR.

Figure 11-10 WRR Scheduling and Bandwidth Allocation

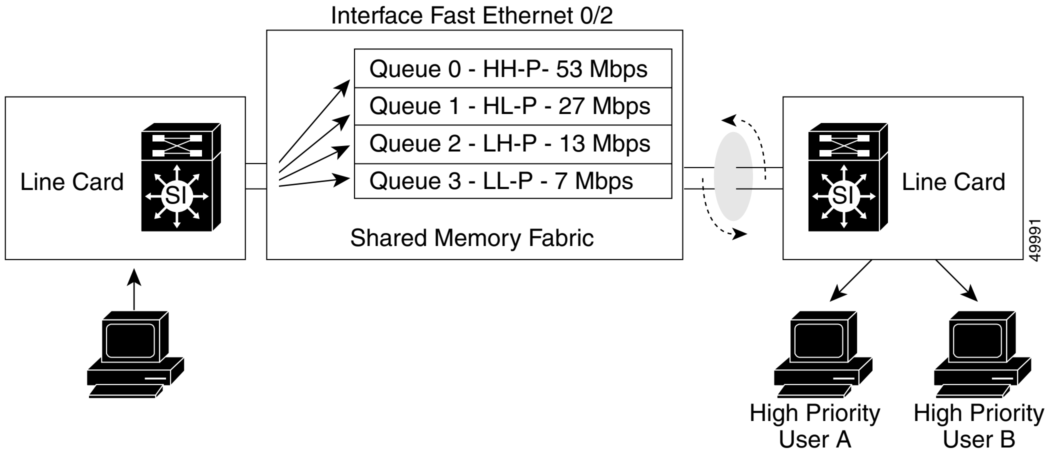

Line Card Architecture

The last major component of the Layer 3 enabled ATM switch router architecture is the line cards. Because the switch uses a distributed architecture, the line cards must be intelligent enough to make both Layer 3 and Layer 2 forwarding decisions at wire speed for all media types, as well as enforce QoS policies. Figure 11-11 details the architecture of the Layer 3 enabled ATM switch router line cards. In Figure 11-11, notice that the Catalyst 8540 uses four CEFAs per line card.

The Layer 3 enabled ATM switch router line cards are based on the Cisco Express Forwarding ASIC (CEFA). The CEF ASIC is based on the MMC Ethernet processor interface ASIC. It is called the CEF ASIC since the Cisco Express Forwarding mechanism is programmed into the ASICs. This ASIC is responsible for the Ethernet MAC layer functions, address or network lookup in the content-addressable memory (CAM) table, and forwarding of the packet with its correct rewrite information to the Fabric Interface. The Fabric Interface is also resident on the line card and is responsible for the packet rewrite, QoS classification, and signalling to the Frame Scheduler.

Figure 11-11 Catalyst 8540 CSR Line Card Architecture

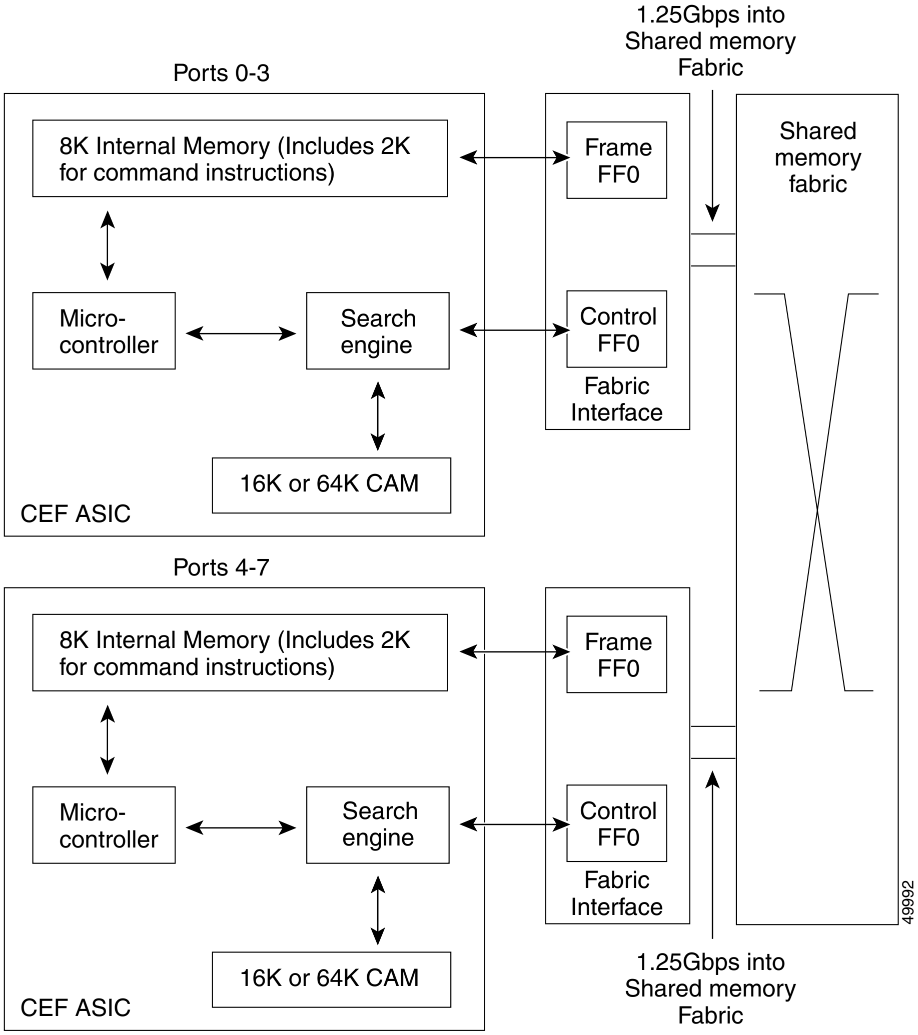

CEFA

The CEFA is at the heart of the line card architecture. This ASIC has several key components that will be discussed in detail. Each CEFA services four ports on the line card. In order to service eight ports, two CEFAs are used per line card. On the Catalyst 8540, four CEFAs are used in order to service 16 ports. Although not shown in Figure 11-11, the CEFA is responsible for all MAC layer functions. The MAC is 10/100 autosensing and autonegotiating, if so configured. The MAC can also be run in either full or half duplex mode.

Packets entering the switch port and having passed though MAC functions are stored in an internal block of SRAM. This memory is 8 kilobytes in size, with 2K reserved for command instructions. This memory is used to hold the packet while the appropriate lookups take place.

The CEFA microcontroller is a mini-route processor that is local to four ports on the Layer 3 enabled ATM switch router line module. The microcontroller is designed to handle the traffic on each of the ports in a fair manner. This means the CEFA must ensure that all packets have equal access into internal memory and that lookups via the search engine are done fairly by arbitrating service between the four ports. This is handled in a round-robin manner, meaning that the microcontroller cycles between each port, processing requests as needed.

The microprocessor also has the critically important task of forwarding system messages such as spanning tree BPDUs, routing advertisements, Cisco Discovery Protocol (CDP) packets, Address Resolution Protocol (ARP) frames, and other control-type messages back to the route processor. Those messages are forwarded by the CEFA to the route processor.

CEFA Search Engine

The search engine in the CEFA performs the address lookup or network output interface lookup. It performs its lookup in the CAM table, which can hold either 16,000 or an optional 64,000 entries. The search engine can make two types of switching decisions: Layer 2 based or Layer 3 based. With the hardware-based access list feature card, the search engine can also perform lookups based on Layer 4 information. The search engine is therefore responsible for maintaining the Layer 2 MAC address table and the Layer 3 FIB.

An incoming packet is placed into the internal memory. As soon as the first 64 bytes of the frame are read into memory, the microcode signals the search engine with the relevant source or destination MAC address, destination network, or Layer 4 port information. The search engine can then conduct a lookup in the CAM table for the corresponding entry. Using a binary tree lookup method, the search engine can hit a MAC address or perform a longest match on the destination network address very quickly. The corresponding rewrite information, which is stored in the CAM table, is then delivered to the control FIFO buffer of the Fabric Interface.

Fabric Interface

The final stage in packet switching within the Layer 3 enabled ATM switch router can now occur. The switching CEFA now knows the port-of-exit for the packet based either on its MAC address or on the Layer 3 IP or IPX network numbers. The packet must now be transferred across the switching fabric to the destination. The Fabric Interface is responsible for preparing the packet for its journey across the switching fabric.

The Fabric Interface consists of two main components: the frame FIFO buffer and the control FIFO buffer. Figure 11-11 shows the internal memory of the CEFA, its direct connection into the frame FIFO buffer, and the direct connection from the search engine into the control FIFO buffer. When the search engine completes the lookup, the packet moves from internal memory into the frame FIFO buffer. In parallel, the search engine returns to the control FIFO buffer all of the relevant rewrite and QoS information.

The Fabric Interface then rewrites the packet with the appropriate information and calculates the checksum. At the same time, the Fabric Interface adds to the beginning of the packet the internal routing tag containing port of exit, the QoS priority, and drop priority (see Figure 11-8). Once completed, the Frame Scheduler is signaled to place the frame into the fabric.

At the output port, the Fabric Interface forwards the packet to its output MAC. Since all rewrite and error checking has been done at the ingress port, no additional work needs to be performed on that frame.

Private, Shared, and Dual CAMs

Private CAM describes where each interface has its own CAM. The CAM space is used to store direct lookup tables, Layer 2 and Layer 3 forwarding tables that assist in the ASIC hardware forwarding. See Figure 11-12.

The various CAM types are described as follows:

- Private CAM
  - Each FastEthernet interface has its own CAM space
  - 1-to-1 ratio between hardware interface and CAM
- Shared CAM
  - One Ethernet processor interface (4 Ports) share CAM Space
  - 1-to-many ratio between hardware interface and CAM
- Dual CAM found on current Gigabit Ethernet module
  - One CAM per Ethernet processor interface (2 Ethernet processor interfaces per Gigabit Ethernet processor interface)
  - Many-to-1 ratio between hardware interface and CAM

Figure 11-12 Private CAM

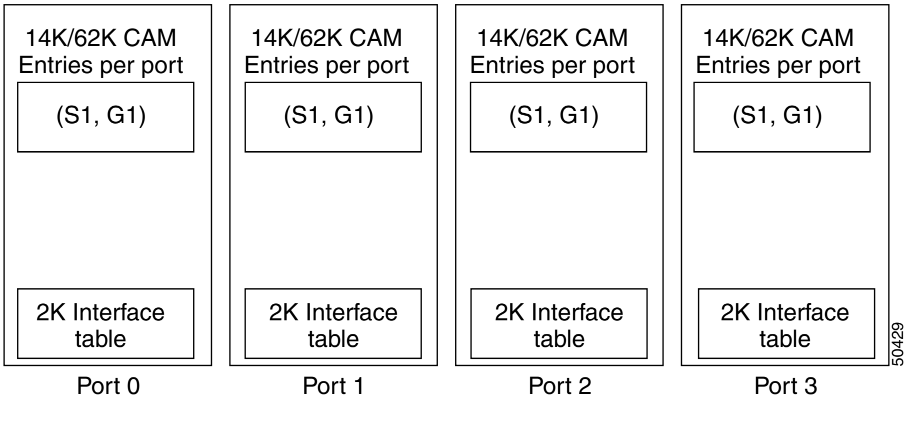

The shared CAM allows one single CAM space per Ethernet processor interface, and this CAM space is physically shared among all four ports within this interface. See Figure 11-13. Shared CAM space has implications in the way the direct lookup table and Layer 3 database are maintained in the CAM.

Note A shared CAM board and non-shared CAM board can co-exist in the same switch router.

Figure 11-13 Shared CAM

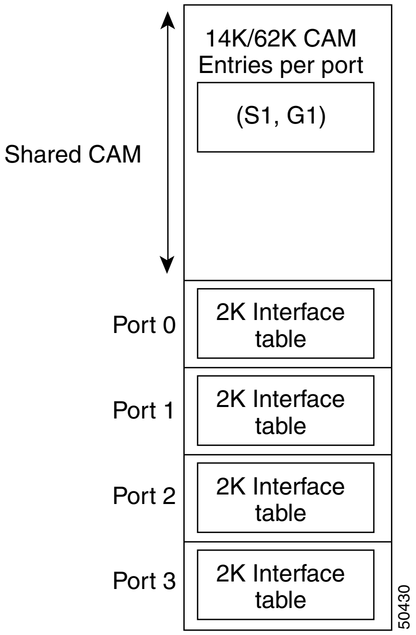

There are always five P2MP VCs in the switch router:

- One VC for all Gigabit processor interfaces, with two leaves for each Gigabit processor interface.
- Four P2MP VCs for each Ethernet processor interface, one corresponding to each channel.

With shared CAM, Gigabit processor interface P2MP remains the same. However, for Ethernet processor interfaces with shared CAM, only channel-0 leaf is created. Other channel leaves are not created. This allows a mix of private, shared, and dual CAM interfaces in the switch router.

To determine what type of CAM is installed on your interface use the **show hardware detail** command as shown in the following example:

```
Switch# show hardware detail

C8540 named Switch, Date: 10:41:12 UTC Thu Dec 7 2000

Slot Ctrlr-Type    Part No.  Rev  Ser No  Mfg Date   RMA No. Hw Vrs  Tst EEP
---- ------------  ---------- -- -------- --------- -------- ------- --- ---
 0/* Super Cam     73-2739-03 D0 03170TAL May 03 99 0          3.1
 0/0 8T1 IMA PAM   73-3367-02 B2 03100061 Mar 15 99 00-00-00   2.0     0   0
 0/1 8E1 IMA PAM   73-3378-02 B2 03120056 Mar 25 99 00-00-00   2.0     0   2
 2/* ARM PAM       73-4208-01 05 03150016 Apr 18 99            1.0
 3/* ETHERNET PAM  73-3754-06 B0 03282WBF Jul 13 99 0          5.1
 9/* OC48c PAM     73-3745-02 12 03190UXC Jun 28 99            2.1
10/* OCM Board     73-4165-01 04 03230ZZ2 Jun 28 99           10.1
10/0 QUAD 622 Gen  73-2851-05 A0 03160RVS Jun 16 99            5.0
11/* OC48c PAM     73-3745-02 12 03100015 Jun 28 99            2.1
12/* OCM Board     73-4165-01 04 03190UJV Jun 28 99           10.1
12/0 QUAD 622 Gen  73-2851-05 A0 03160S9J Jun 16 99 0          5.0

.
(Information Deleted)
```

```
  .
  slot:  2/*  Controller-Type : ARM PAM
    Part Number: 73-4208-01                       Revision: 05
  Serial Number: SCA03150016                      Mfg Date: Apr 18 99
     RMA Number:                               H/W Version: 1.0
   FPGA Version: 2.3

   EPIF Version: 1704                  CAM size: 64 KB EPIF Version: 1704
  CAM size: 64 KB
➔ Ucode Version: 0.0                              CAM Type: Dual

  Port Phy Setup
        Port  0: DONE                          GBIC Vendor: No vendor info.
        Port  1: DONE                          GBIC Vendor: No vendor info.

  slot:  3/*  Controller-Type : ETHERNET PAM
    Part Number: 73-3754-06                       Revision: B0
  Serial Number: CAB03282WBF                      Mfg Date: Jul 13 99
     RMA Number: 0                             H/W Version: 5.1
   FPGA Version: 3.2

  Chip 0 Reset Count: 0
  Chip 1 Reset Count: 0
  Chip 2 Reset Count: 0
  Chip 3 Reset Count: 0

   EPIF Version: 1704                             CAM size: 16 KB
➔ Ucode Version: 1.0                              CAM Type: Private

  --More--
```

In the previous example, the CAM Type field lists the CAM type for the ARM module in slot 2/* as Dual and the CAM type for the Ethernet module in slot 3/* as Private.

Comparing Data Plane and Control Plane Traffic

Data plane traffic is traffic between two endpoints (for example, a host on subnet A communicating with a host on subnet B). This data plane traffic will be typically switched by the Ethernet processor interface or Gigabit processor interface. Control Plane traffic is traffic which is handled by the route processor, typically Layer 2 and Layer 3 protocol updates.

The following is a list of traffic considered to be Control Plane traffic and handled by the route processor:

IP Packet Traffic on the Control Plane

IP packets are sent to the route processor in the following situations:

- Packets matching the switch router IP address
- No route found on the line card with “ICMP unreachable” is enabled
- Packets with TTL=0 after TTL decrement
- Packets with options set in IP header
- Packets in or out on the same interface and with ICMP redirect enabled
- ARP and Reverse ARP packets
- Certain multicast and broadcast packets (for example, OSFP/EIGRP route updates).

- RIP broadcasts
- HSRP hellos
- DHCP helper
- Invalid next hop

IPX Packet Traffic on the Control Plane

IPX packets are sent to the route processor in the following situations:

- Packets matching the switch router IPX address
- GNS packets
- Certain broadcast packets (for example, RIP/EIGRP/SAP route updates).
- Destination node broadcast
- Invalid next hop

Miscellaneous Packet Traffic on the Control Plane

The following packets are sent to the route processor on the control plane:

- SNMP Queries
- BPDUs
- Layer 2 Learning
- CAM Entry Overflows

Troubleshooting Half- or Full-Duplex Negotiation

Autonegotiation converges to using the minimum capability of the local interface and the peer interface. For example, if the local interface is capable of full-duplex transmission and the peer interface is only capable of half-duplex transmission, after the local interface performs autonegotiation the interface changes to operate in half-duplex mode.

If the peer interface does not have transmission mode autonegotiation capability, but the local interface has transmission mode autonegotiation capability and the local interface receives no response to its negotiation requests, the local interface changes to operate in half-duplex mode.

To Support half-duplex and full-duplex autonegotiation, your interface must confirm to the following:

- media type must be UTP
- Ethernet processor interface version must be C1 (Slicer Register EVER 0x1704)

Other interfaces (10/100Mbps Ethernet processor interface versions less than C1) have a default speed of 100Mbps, full duplex, and are not capable of autonegotiation.

To determine the installed interface version, use the **show controllers** {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* command and find the EVER field under the Slicer registers. The Ever field should be EVER 0x1704 (C1); or if it is not, your interface is not capable of autonegotiation.

```
   Switch# show controllers fastEthernet 3/0/0

   IF Name: FastEthernet3/0/0
   Port Status UP
   Loopback Reg [3-0]|[7-4]: 0x8|0x8
   Duplex/Speed Reg [3-0]|[7-4]: 0xFFF7|0x0
   FPGA Rev : 3.8
   Internal Reset Trigger Count: 0

→  Slicer registers
   SMDR 0x0060 (Tx En,Rx En)
   SSTR 0x1000
→  EVER 0x1704 (C1)
   SSMR 0x4000 SIMR 0x0000 MBXW 0x0000 MBXR 0x0000
   SPER 0xF000 GMUX VER 0xF000 MARKER 0x0000
   .
   (Information Deleted)
   .
```

Half- and Full-Duplex Troubleshooting Commands

To troubleshoot half- and full-duplex negotiation problem, use the following commands:

Command	Purpose
show interfaces {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port*	Displays interface configuration, status, and statistics.
show controllers {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port*	Displays controller status for the specified interface.

Follow these steps to troubleshoot the half- and full-duplex negotiation problem on an interface:

Step 1 Use the **show interfaces fastEthernet** *card*/*subcard*/*port* command to check the half-duplex and full-duplex autonegotiation configuration.

```
   Switch# show interfaces fastEthernet 3/0/0
   FastEthernet3/0/0 is up, line protocol is up
     Hardware is epif_port, address is 0090.2156.d837 (bia 0090.2156.d837)
     Internet address is 172.20.52.36/27
     MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
        reliability 255/255, txload 1/255, rxload 1/255
     Encapsulation ARPA, loopback not set
     Keepalive set (10 sec)
→    Auto-duplex, Auto Speed, 100BaseTX
     ARP type: ARPA, ARP Timeout 04:00:00
     Last input 00:00:01, output never, output hang never
     Last clearing of "show interface" counters never
     Queueing strategy: fifo
     Output queue 0/40, 0 drops; input queue 0/75, 0 drops
     5 minute input rate 0 bits/sec, 0 packets/sec
```

```
  5 minute output rate 1000 bits/sec, 2 packets/sec
     33684 packets input, 11817561 bytes
     Received 9 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored 0 abort
     0 watchdog, 33546 multicast
     0 input packets with dribble condition detected
     61232 packets output, 13584791 bytes, 0 underruns(0/0/0)
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
Switch#
```

Step 2 Check the Auto-duplex, Auto Speed, 100BaseTX fields. They should have the following default configuration:

- Auto-duplex—Auto duplex negotiation
- Auto-Speed—Auto speed negotiation
- 100BASE-TX—100-Mbps BASE-TX

If they do not, check the peer interface and determine whether it is capable of this configuration.

Step 3 Use the **show controllers fastEthernet** *card/subcard/port* command to check the half-duplex and full-duplex autonegotiation configuration.

Note The **show controllers** command for a specific interface has different information depending on the IOS software version running on your Layer 3 enabled ATM switch router.

Step 4 Use the **show controllers fastEthernet** *card/subcard/port* command to check the configuration. The following example uses the Cisco IOS Release 12.0(5)W5(13b) and later display:

```
  Switch# show controllers fastEthernet 3/0/0

  IF Name: FastEthernet3/0/0
  .
  (Information Deleted)
  .
  MII registers:
  Control Register              (0x0): 0x1000 (Auto negotiation enabled)
  Status Register               (0x1): 0x782D (Auto negotiation complete)
  PHY Identification Register 1 (0x2): 0x7810

  PHY Identification Register 2 (0x3): 0x43
➔ Auto Neg. Advertisement Reg   (0x4): 0x1E1 (Speed 100 ,Duplex Full )
➔ Auto Neg. Partner Ability Reg (0x5): 0x81 (Speed 100 ,Duplex Half )
  Auto Neg. Expansion Register  (0x6): 0x0
  Mirror Register              (0x10): 0x630
  Interrupt Enable Register    (0x11): 0x0
  Interrupt Status Register    (0x12): 0x4000
  Configuration Register       (0x13): 0x0 (UTP, Tx Enabled)
➔ Chip Status Register         (0x14): 0x28C9 (Link Up,a-Half,a-100  )
  Link Status Register    [3-0]|[7-4]: 0x1|0x0

  Counters :
  .
  (Information Deleted)
  .
```

Use the **show controllers fastEthernet** *card/subcard/port* command to check the configuration. The following example uses the Cisco IOS Release 12.0(5)W5(13) and earlier display:

```
  Switch# show controller fastEthernet 1/0/0

  IF Name: FastEthernet1/0/0
  .
  (Information Deleted)
  .
  MII registers:
  Control Register                (0x0): 0x1000
  Status Register                 (0x1): 0x782D
  PHY Identification Register 1   (0x2): 0x7810
  PHY Identification Register 2   (0x3): 0x43
➔ Auto Neg. Advertisement Reg     (0x4): 0x1E1
➔ Auto Neg. Partner Ability Reg   (0x5): 0x1E1
  Auto Neg. Expansion Register    (0x6): 0x1
  Mirror Register                (0x10): 0x30
  Interrupt Enable Register      (0x11): 0x0
  Interrupt Status Register      (0x12): 0x4000
  Configuration Register         (0x13): 0x0
  Chip Status Register           (0x14): 0x38C8
  Link Status Register     [3-0]|[7-4]: 0x1|0x0
  .
  (Information Deleted)
  .
```

Step 5 Check the Auto Neg. Advertisement Register (Reg 0x4). If it is set to 1, the following are the capabilities:

- Bit 8 - 100 BASE-TX Full Duplex
- Bit 7 - 100 BASE-TX
- Bit 6 - 10 BASE-T Full Duplex
- Bit 5 - 10 BASE-T

Step 6 Check the Auto Neg. Partner Ability Reg (Reg 0x5). If it is set to 1, the following are the status and capabilities:

- Bit 14 - Link Partner has received the Link code word from the local
- Bit 13 - Remote Fault
- Bit 8 - 100 BASE-TX Full Duplex
- Bit 7 - 100 BASE-TX
- Bit 6 - 10 BASE-T Full Duplex
- Bit 5 - 10 BASE-T

Step 7 Check the Chip Status Register field. It should match the link status, duplex mode, and speed shown in the **show interface** command in Step 2.

Troubleshooting IP Layer 3 Connections

The Layer 3 enabled ATM switch router uses Cisco Express Forwarding (CEF). Much of the internal troubleshooting determines whether the central CEF information in the route processor is consistent with the distributed information in the content addressable memory (CAM) on the interfaces.

Troubleshooting an IP Layer 3 connection is separated into the following processes:

- IP Layer 3 Connection Troubleshooting Commands, page 11-25
- Checking the IP Routing Table, page 11-27
- Checking the Interface Status, page 11-28
- Checking the IP CEF Adjacencies, page 11-30
- Checking the Interface CAM Table Entries, page 11-33

Figure 11-14 shows the example network used to troubleshoot an IP Layer 3 connection in the following examples.

Figure 11-14 IP Layer 3 Connection

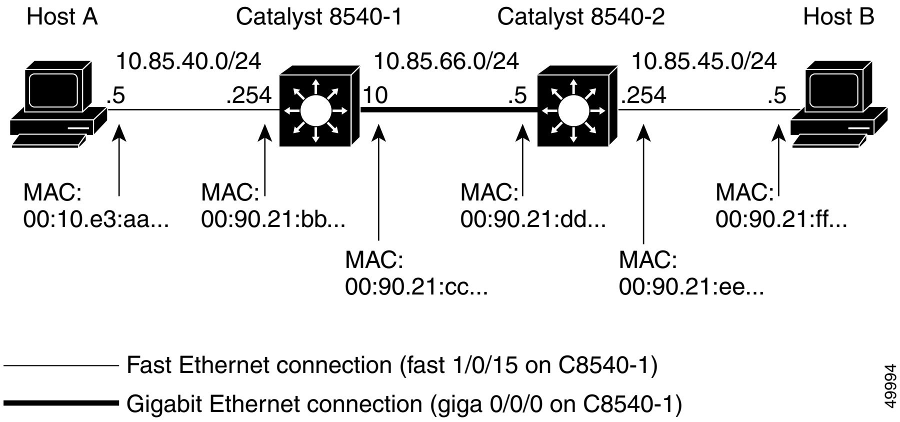

In Figure 11-14, Host A is the source end station trying to communicate with Host B, the destination end station.

IP Layer 3 Connection Troubleshooting Commands

To troubleshoot an IP Layer 3 connection problem, use the following commands:

Command	Purpose
show ip route	Displays routing table entries.
show controllers c8500 status	Displays the status of all Ethernet processor interfaces.
show controllers c8500 counters	Displays the counters of all Ethernet processor interfaces.
show ip cef	Displays Cisco Express Forwarding information.

Command	Purpose
show adjacency detail	Displays IP address table information for adjacent nodes.
show ip route summary	Displays summary information about the routing table entries.
show arp	Displays the ARP table.
show epc if-entry interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **all-entries**	Displays all interface entry information for the specific interface.
show epc ip-address interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port ip-address mask* (on the ingress interface)	Displays the IP addresses of adjacent interfaces.
show epc ip-address interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port ip-address mask* (on the egress interface)	Displays the IP addresses of adjacent interfaces.
show epc lsipc	Displays the LSIPC information.
show epc ifmapping	Displays interface mapping to CAM interface number.
show epc patricia interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **ipucast detail** (on the ingress interface)	Displays the patricia tree entries in the CAM.
show epc cam interface {**fastethernet** | **gigabitethernet**} [*CAM-start-address*] [*CAM-word-number*]	Displays the CAM table rewrite information.
show epc if-entry interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **entry** {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port*	Displays interface entry information for the specific interface.
show epc ip-prefix interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **all-entries** (on the egress interface)	Displays the IP network entries for the egress interface.
show epc ip-address interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port ip-address mask* (on the egress interface)	Displays the IP addresses of adjacent interfaces.
show epc patricia interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **ipucast detail** (on the ingress interface)	Displays the patricia tree entries in the CAM.
show epc patricia interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **ipucast detail** (on the egress interface)	Displays the patricia tree entries in the CAM.

Checking the IP Routing Table

Follow these steps to verify the IP routing tables in the IP Layer 3 connection shown in Figure 11-15.

Figure 11-15 Displaying Router Table Information

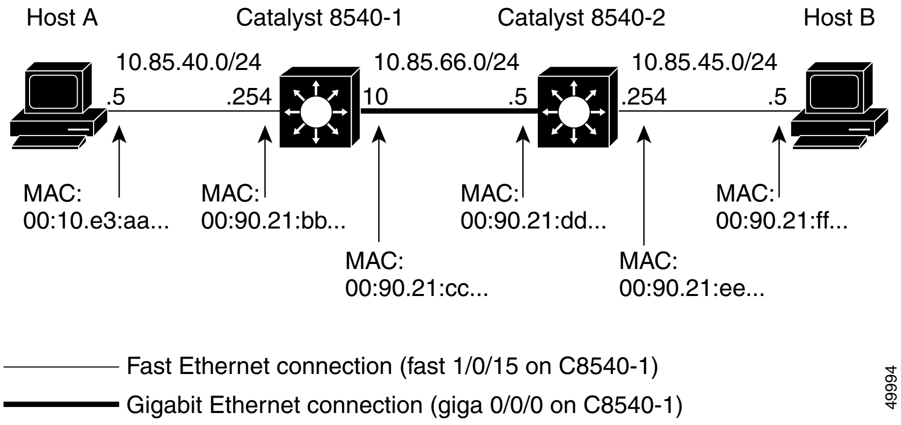

Step 1 From the Catalyst 8540-1, use the **show ip route** command to verify the status of the IP routing table for the example network shown in Figure 11-15.

```
C8540CSR-1# show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

C    10.85.40.0/24 is directly connected, Fast Ethernet 1/0/15
D    10.85.45.0/24 [90/30720] via 10.85.66.0, 01:22:23, Gigabit Ethernet 0/0/0
C    10.85.66.0/24 is directly connected, Gigabit Ethernet 0/0/0
8540CSR-1#
```

Note All the networks are directly connected except for 10.85.45.0, which was learned through EIGRP, via interface Gigabit Ethernet 0/0/0.

Step 2 From the Catalyst 8540-1, use the **show ip route** command to display the network connecting Host B to Catalyst 8540-2 with IP address 10.85.45.0.

```
C8540CSR-1# show ip route 10.85.40.0
➔ Last update from 10.85.66.5 on GigabitEthernet 0/0/0, 1d16h ago
    Routing Descriptor Blocks:
    * 10.85.66.5, from 10.85.45.9, 1d16h ago, via GigabitEthernet 0/0/0

C8540CSR-1#
```

The display confirms the route to network 10.85.45.0, which exists in the routing table and was learned via IP address 10.85.66.5 through Gigabit Ethernet interface 0/0/0.

Note If there are routes missing, continue with normal IP routing troubleshooting for the routing protocol you are using.

If you determine that the interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Checking the Interface Status

Follow these steps to verify the interface status in the IP Layer 3 connection shown in Figure 11-16.

Figure 11-16 Displaying the Interface Status Information

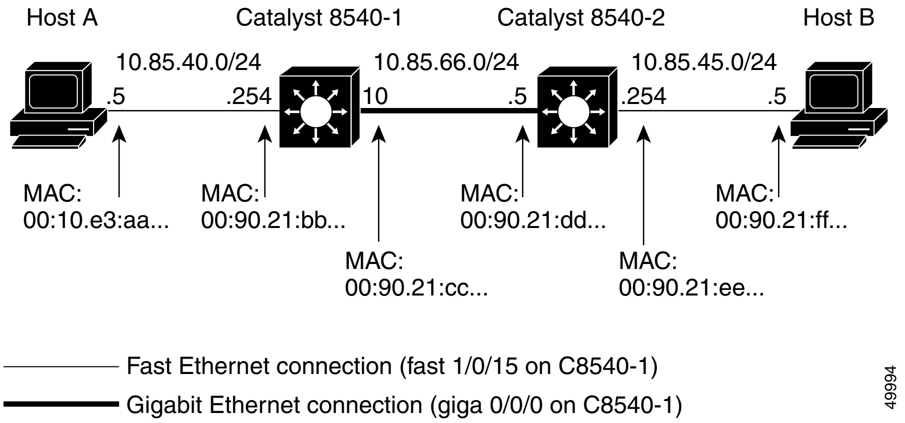

Step 1 Verify the status of the interfaces for the example network shown in Figure 11-16 using the **show controllers c8500 status** command.

Step 2 From the Catalyst 8540-1, use the **show controllers c8500 status** command to display the status of the interfaces used in this example.

```
  C8540CSR-1# show controllers c8500 status
➔ Status of GigabitEthernet0/0/0: OK
  Status of GigabitEthernet0/0/1: OK
  Status of FastEthernet1/0/0: OK
  Status of FastEthernet1/0/1: OK
  Status of FastEthernet1/0/2: OK
  Status of FastEthernet1/0/3: OK
  Status of FastEthernet1/0/4: OK
  Status of FastEthernet1/0/5: OK
  Status of FastEthernet1/0/6: OK
  Status of FastEthernet1/0/7: OK
  Status of FastEthernet1/0/8: OK
  Status of FastEthernet1/0/9: OK
  Status of FastEthernet1/0/10: OK
  Status of FastEthernet1/0/11: OK
  Status of FastEthernet1/0/12: OK
  Status of FastEthernet1/0/13: OK
  Status of FastEthernet1/0/14: OK
➔ Status of FastEthernet1/0/15: OK
  C8540CSR-1#
```

The OK in the **show controllers c8500 status** command display indicates the microcode was successfully downloaded to the Fast Ethernet processor interface and Gigabit processor interface.

Step 3 From the Catalyst 8540-1, use the **show controllers c8500 counters** command to display the status of the interfaces and the Input and Output Packet numbers.

```
  C8540CSR-1# show controllers c8500 counters
  Interface   Input     Runts Giants  Input       CRC  Frame Output      Output
      State   Packets                 Errors                 Packets     Errors
  -----------------------------------------------------------------------------
➔ G0/1/0  U   127286      0     0     0           0     0     137296      0
  G0/1/1  U   0           0     0     0           0     0     20          0
  F1/0/0  U   31849       0     0     0           0     0     31855       0
  F1/0/1  AD  1           0     0     0           0     0     1           0
  F1/0/2  AD  1           0     0     0           0     0     1           0
  F1/0/3  AD  1           0     0     0           0     0     1           0
  F1/0/4  AD  1           0     0     0           0     0     1           0
  F1/0/5  AD  1           0     0     0           0     0     1           0
  F1/0/6  AD  1           0     0     0           0     0     1           0
  F1/0/7  AD  1           0     0     0           0     0     1           0
  F1/0/8  AD  1           0     0     0           0     0     1           0
  F1/0/9  AD  1           0     0     0           0     0     1           0
  F1/0/10 AD  1           0     0     0           0     0     1           0
  F1/0/11 AD  1           0     0     0           0     0     1           0
  F1/0/12 AD  1           0     0     0           0     0     1           0
  F1/0/13 AD  1           0     0     0           0     0     1           0
  F1/0/14 AD  1           0     0     0           0     0     1           0
➔ F1/0/15 U   31968       0     0     0           0     0     54732       0
   --More--

  -----------------------------------------------------------------------------
  AD - Admin Down, D - Down, F - Fail, U - Up

  C8540CSR-1#
```

Step 4 Check the Interface State field. It should indicate the interfaces are up.

Step 5 Check the Input Packets and Output Packets fields. The **show controllers c8500 counters** command should be entered at least twice. The counters in the Input Packets and Output Packets fields should be incrementing. This information can also be displayed using the **show interfaces** command.

Note The **clear counters** command does not clear the **show controllers c8500 counters** command display.

If you determine that the interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Checking the IP CEF Adjacencies

Follow these steps to verify the IP CEF adjacencies in the IP Layer 3 connection shown in Figure 11-17.

Figure 11-17 Displaying the IP CEF Adjacency Information

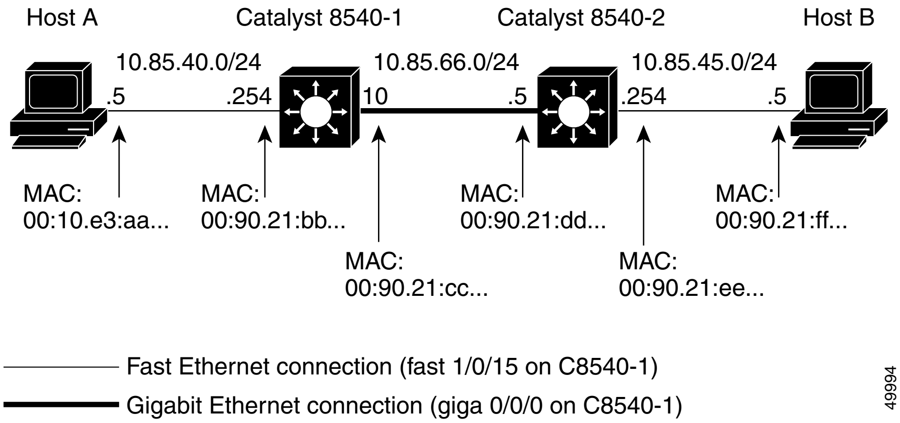

Step 1 Use the **show ip cef** command to verify that routes and attached devices appear in the table correctly and point to the correct next hop or outgoing interface.

```
  C8540CSR-1# show ip cef
  Prefix              Next Hop             Interface
  0.0.0.0/32          receive
  10.19.134.36/32     10.19.134.36         Ethernet0
  10.85.40.0/24       attached             FastEthernet1/0/15
➔ 10.85.40.0/32       receive
➔ 10.85.40.254/32     receive
➔ 10.85.40.5/32       10.85.40.5           FastEthernet1/0/15
➔ 10.85.40.255/32     receive
  .
  (Information Deleted)
  .
➔ 10.85.45.0/24       10.85.66.10          GigabitEthernet0/0/0
  10.85.66.0/24       attached             GigabitEthernet0/0/0
  10.85.66.0/32       receive
  10.85.66.10/32      receive
  10.85.66.255/32     receive
  224.0.0.0/4         drop
  224.0.0.0/24        receive
  255.255.255.255/32  receive
  C8540CSR-1#
```

The information in the **show ip cef** command display is built from the IP routing table and resides on the route processor.

The following is an explanation of the information in the Next Hop Column:

- Attached—This is a directly connected interface subnet. For example, 10.85.40.0/24 is the IP subnet assigned to interface Fast Ethernet1/0/15 with a 24-bit mask.
- Received—These entries are ARP entries for the directly connected interfaces. You will see three entries here for each directly connected interface. For example, prefix 10.85.40.254/32 is the IP address for interface Fast Ethernet 1/0/15. Prefix 10.85.40.0/32 using IP conventions means that this specific interface and prefix 10.85.40.255/32 is the broadcast address.
- xxxx.yyyy.zzzz.aaaa—These IP addresses belong to either the end station connected to the interface (ARP entries) or the next-hop router for a specific subnet. For example, prefix 10.85.40.5/32 is an end station. The prefix entry and next-hop entry are the same. Prefix entry 10.85.45.0/24 is a route learned via next-hop 10.85.66.5.

Step 2 From the Catalyst 8540-1, use the **show ip cef** command with the destination network IP address to display the CEF FIB table entry for the network connecting Host B to Catalyst 8540-2 with IP address 10.85.45.0.

```
  C8540CSR-1# show ip cef 10.85.45.0
  10.85.45.0/24, version 22, cached adjacency 10.85.66.5
    via 10.85.66.5, GigabitEthernet0/0/0, 0 dependencies
➔     next hop 10.85.66.5, GigabitEthernet0/0/0
➔     valid cached adjacency
```

The display confirms that the next hop IP address 10.85.66.5 is valid and has a valid cached adjacency.

Step 3 From the Catalyst 8540-1, use the **show adjacency** command to display the MAC address rewrite information for the connection from Catalyst 8540-1 to Catalyst 8540-2 with IP address 10.85.66.5.

```
  C8540CSR-1# show adjacency GigabitEthernet 0/0/0 detail
  Protocol   Interface            Address
➔ IP         GigabitEthernet0/0/0  10.85.66.5(9)
                                   0 packets, 0 bytes
➔                                  009021DDDDDD009021CCCCCC0800
                                   ARP        03:59:57
```

The display confirms the MAC address rewrite information as:

- **009021DDDDDD**—the Catalyst 8540-2 destination MAC address
- **009021CCCCCC**—the Catalyst 8540-1 source MAC address
- **0800**—protocol field, IP ARPA (IP Ethernet type [hex 0800])

Note The MAC addresses of the source and destination interfaces is displayed using the **show interface** command.

If the next hop interface does not display the correct MAC address rewrite information, use the **show arp** command to confirm the MAC addresses.

Step 4 From the Catalyst 8540-1, use the **show arp** command to display the ARP table.

```
  C8540CSR-1# show arp
  Protocol  Address          Age (min)  Hardware Addr   Type   Interface
  .
  (Information Deleted)
  .

➔ Internet  10.85.40.5            175   0010.e3aa.aaaa  ARPA
➔ Internet  10.85.40.254           -    0090.21bb.bbbb  ARPA   FastEthernet1/0/15

➔ Internet  10.85.66.10            -    0090.21cc.cccc  ARPA   GigabitEthernet0/0/0
➔ Internet  10.85.66.5            172   0090.21dd.dddd  ARPA   GigabitEthernet0/0/0

  C8540CSR-1#
```

The first entries in this ARP table are, from top to bottom:

- Host A end station
- Fast Ethernet interface connection to the end station
- Gigabit Ethernet interface out to the next hop
- next hop router interface

If you confirm the information is wrong using the **show adjacency** command to display the MAC address rewrite information, this might indicate a problem with CEF. You should confirm the interface CAM table entries using the following troubleshooting procedure.

If you determine that the interface is configured incorrectly, refer to the "Configuring Interfaces" chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Checking the Interface CAM Table Entries

Follow these steps to verify the interface CAM table entries in the IP Layer 3 connection shown in Figure 11-18.

Figure 11-18 Displaying the Interface CAM Table Information

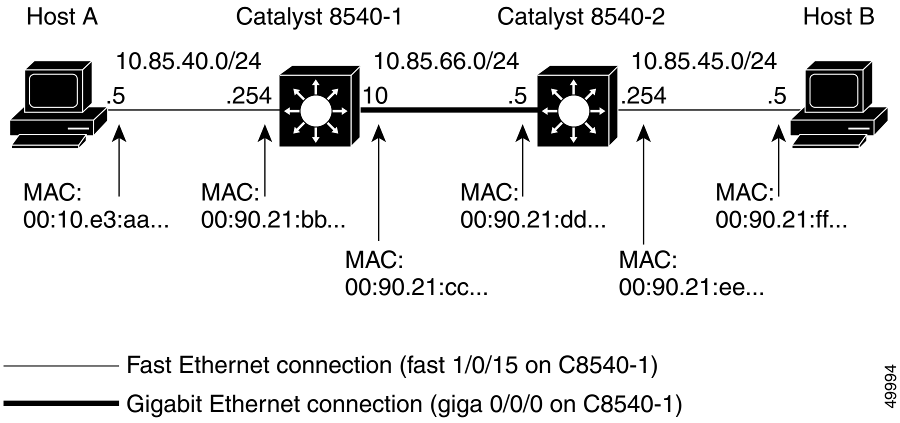

Step 1 From the Catalyst 8540-1, use the **show epc ip-prefix interface** command to display the status of the CAM table for the ingress interface in question.

```
   C8540CSR-1# show epc ip-prefix interface FastEthernet 1/0/15 all-entries
   Default Network Information:
→      Not configured
    Prefix/Masklen        Next Hop
→   0.0.0.0/32            not populated
→   10.0.0.7/32           20.0.0.1
    10.0.1.4/30           20.0.0.1
    10.0.1.12/30          20.0.0.1
    10.0.1.24/30          20.0.0.1
    10.0.1.124/30         20.0.0.1
    11.1.1.0/30           20.0.0.1
    11.1.2.0/30           20.0.0.1
    11.1.3.0/24           20.0.0.1
    11.1.9.0/24           20.0.0.1
    11.1.10.0/24          20.0.0.1
    11.1.40.0/24          20.0.0.1
    11.1.100.0/24         20.0.0.1
    11.1.120.0/24         20.0.0.1
    15.15.15.0/24         20.0.0.1
→   20.0.0.0/24           SRP
    20.0.0.0/32           SRP
    20.0.0.1/32           not populated
    20.0.0.2/32           SRP
    20.0.0.255/32         SRP
    20.0.1.0/24           SRP
    20.0.1.0/32           SRP
    20.0.1.2/32           SRP
    20.0.1.255/32         SRP
    172.17.110.0/24       20.0.0.1
    172.17.110.96/27      20.0.0.1
    224.0.0.0/4           not populated
    224.0.0.0/24          SRP
    255.255.255.255/32    not populated
→   Total IP Prefix Entries in CAM:25
    Missing IP Prefix Entries in CAM:0
    CEF entries not populated:4
   C8540CSR-1#
```

The Prefix/Masklen indicates the IP addresses and subnet masks of connections in the interface CAM table.

The Not configured field indicates no default route is known. If you added IP route 0.0.0.0 20.0.0.1 to that configuration, the display would change to include the following:

```
Default Network Information:
     Nexthop 1:
       IP addr:20.0.0.1   GigabitEthernet2/0/1 (58)
       Mac Addr:0090.2141.bd47
     Load Balancing:Off
```

Note Since there is only one route, the Load Balancing field is Off.

The next hop column contains the following descriptions:

- Not populated—Indicates this is an entry for either an end station or a next hop router interface. To display entries in the CAM for these connections, use the **show epc ip-address interface** command with the **all-entries** parameter.
- IP address—Indicates the next hop in the CAM is to this IP address from this source prefix.

- SRP—Indicates the prefix is a directly attached interface and these interfaces do not have a route in the table, therefore packets from this network are sent to the route processor for processing. See the earlier explanation in Step 7 for the entries per directly connected interface.

  All interfaces should have the same CAM entries, since the forwarding decision is made based on the information contained in the CAM table. This table is based on the network topology, not traffic flows. The **show epc ip-prefix** command used with any other interface on the switch should contain the same number of entries in the Total IP Prefix Entries in the CAM field (in this example, 25).
- CEF entries not populated—Indicates these network connections are missing the masklen /32. These connections appear if you use the **show epc ip-address** command. In this example, you should program all masklen as /30 and shorter prefixes.

Additionally, you can use other **show epc ip-prefix interface** command parameters to check the cam-summary as well as the **fail-entries** and **fail-summary**.

```
C8540CSR-1# show epc ip-prefix interface FastEthernet 1/0/15 ?
  A.B.C.D       IP prefix to display
  all-entries   All IP Prefix entries
  all-summary   IP Prefix summary
  fail-entries  missing IP prefix entries
  fail-summary  Summary of missing IP prefixes

C8540CSR-1#
```

Step 2 From the Catalyst 8540-1, use the **show epc ip-prefix interface** command to display the connection ingress interface to IP address 10.85.45.0.

```
  C8540CSR-1# show epc ip-prefix interface FastEthernet 1/0/15 10.85.45.0 255.255.255.0
  Prefix/Masklen        Gateway1        Gateway2
➔ 10.85.45.0/24         10.85.66.5
```

The Gateway IP address should match the next hop IP address in the **show epc cef** command output in Step 2 in the section “Checking the IP CEF Adjacencies” section on page 11-30.

To remove inconsistencies between the CEF table and the IP prefix table, use the **clear ip route** command to rebuild these tables. You can either clear a specific route or use an asterisk (*) to clear all routes.

Caution Use the **clear ip route** command carefully. It causes a temporary increase in switch router activity which can lead to traffic disruptions.

Step 3 From the Catalyst 8540-1, use the **show epc ip-address** command with the IP address of the egress interface to display the status of the MAC address rewrite.

```
  C8540CSR-1# show epc ip-address interface FastEthernet 1/0/15 10.85.66.5
➔ IPaddr: 10.85.66.5 MACaddr: 0090.21dd.dddd GigabitEthernet0/0/0 (4)
```

The information in this display should match the information shown using the **show adjacency** command to display the MAC address rewrite in Step 3 of the “Checking the IP CEF Adjacencies” section on page 11-30.

The display confirms the MAC address rewrite information as 0090.21dd.dddd, the Catalyst 8540-2 destination MAC address, and provides the interface index number, in this example “(4)”, to use in the command in Step 5.

Step 4 From the Catalyst 8540-1, use the **show epc lsipc** command to display the status of the interprocess information between the route processor and the Ethernet processor interfaces and Gigabit processor interfaces.

```
C8540CSR-1# show epc lsipc
LSIPC requested: Total: 214759866 Mlet: 214759866
Sent:Total: 214759866 Mlet: 214759866 No-resp: 214757881 Resp-required: 1985
Broadcast IPCs:Requested: 119 Sent: 119
  Queued: 119 Current qsize: 0 Max qsize-reached: 20

Received:Total: 246923174 Unsolicited: 214753326 Response: 1985
Recv Q size: 0

LSIPC Failures:
Toobig: 0 Memory Fail: 0 Packet fail: 0 Invalid VC: 0
Invalid resp: 0 Retries: 0 Timeouts: 0 Ack timeouts: 0
➔ Bcast: Failed: 0 Pkt failed: 0 enq failed: 0 discard: 0
Unicast: Enq failed: 0
C540CSR-1#
```

Check the Bcast fields for any failures. If messages are getting dropped, that could cause inconsistencies in the routing table transfers between the route processor and Ethernet processor interfaces and Gigabit processor interfaces. For example, the IPC communications could have failed.

Step 5 From the Catalyst 8540-1, use the **show epc ifmapping** command to display the status of the interface mapping of the egress interface.

```
C8540CSR-1# show epc ifmapping 4
➔ GigabitEthernet0/0/0    (IF number: 4)
```

The IF number field, in this example “(4)”, indicates the interface index number is mapping correctly.

Step 6 From the Catalyst 8540-1, use the **show epc patricia interface** command with the **ipucast detail** parameters on the ingress interface to display the status of the Host Entry CAM location for the connection to Host B.

```
C8540CSR-1# show epc patricia interface FastEthernet 1/0/15 ipucast detail
.
(Information Deleted)
.
➔ 22#HOST Entry CAM location: 0x102D
   IP addr:10.85.66.5     Host  IF Number:4      Entry:Valid
➔   Mac Addr:0090.21dd.dddd
.
(Information Deleted)
.
C8540CSR-1#
```

The Mac Addr field in the command display indicates the correct MAC address for IP address 10.85.66.5, the next hop, is at the CAM entry location with hexadecimal address 0x102D.

Step 7 From the Catalyst 8540-1, use the **show epc cam interface** command with the CAM location hexadecimal address **0x102D** and the CAM word **2** parameters to display the status of the MAC rewrite for this interface.

```
C8540CSR-1# show epc cam interface FastEthernet 1/0/15 0x102D 2
➔ GigabitEthernet0/0/0 Addr:0x102D Word:2 Data[0]:0x009021DD Data[1]:0xDDDD0045
```

Figure 11-19 describes the CAM encoding information shown in the **show epc cam interface** command, using the CAM location hexadecimal address **0x102D** and the CAM word **2** parameters.

Figure 11-19 CAM Encoding Description

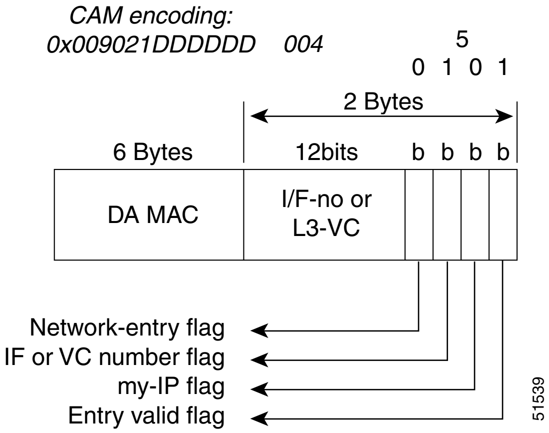

The Data fields in the display indicate the MAC address is written to the following:

- 0x**009021DD**—the first four bytes of the next-hop MAC address
- 0x**DDDD**0045—the last two bytes of the next-hop MAC address
- 0xDDDD**0045**—these last two bytes (in this example “0045”) indicate the following:
  - 004 (12 bits)—the interface or Layer 3 VC number
  - 5 (bit “0”)—Network-entry flag. A “1” indicates this is a host entry.
  - 5 (bit “1”)—ATM VC number flag. A “0” indicates the 12-bit field is an interface number.
  - 5 (bit “0”)—A “1” is a “My-IP flag” and indicates this is the IP address of this interface, and that the packet should be forwarded to the route processor.
  - 5 (bit “1”)—Entry valid flag. A “0” indicates this is an invalid entry.

Note The interface or VC number flag indicates the 12 bits are interpreted as either an interface or ATM VC number. If this were an ATM router module, you could configure the VC to transmit on the ATM side. The VC is then one of the following:
— a data direct VC for ATM LANE
— a PVC or SVC for 1483 or 1577, respectively

Step 8 If this process has not resolved the IP Layer 3 connection failure, repeat this same process for the reverse path from the destination host, and verify that all other interfaces have similar CAM table entries.

Caution Be aware that asymmetrical routing could lead to multicast delivery on an alternate, unintended path, if a forwarding algorithm based on Reverse Path Forwarding is used.

Step 9 From the Catalyst 8540-1, use the **show epc if-entry interface** command with the **entry** interface parameters to display the status of the Broute VC.

```
   C8540CSR-1# show epc if-entry interface FastEthernet 1/0/15 entry GigabitEthernet 0/0/0
   IF Entry for GigabitEthernet0/0/0 on FastEthernet1/0/15
➔         Mac(hex) - 00:90:21:CC:CC:CC
          isMyInteface : False isSubInterface : False
➔         Status Up Broute VC - 67 Bcast VC - 0
          Netmask: 24
          FEC disabled
          Trunking Disabled
          State : Not-Applicable/Listening/Blocking
          Bridge-Group disabled
➔         IP routing on bridging off
          IPX routing off bridging off
          Appletalk routing off
          In Encapsulation:
          ICMP Redirect enabled Unreachable enabled
          IP Multicast disabled: ttl-threshold: 0
```

Verify the following:

- MAC address shown is that of the entry interface.
- Broute VC field status is up.
- IP routing is on.

Workarounds

For inconsistencies between the adjacency table and the EPC IP address table, use the **clear arp** or **clear adjacencies** commands to rebuild the table. When you use one of these commands, the switch router sends an ARP request for all entries in the ARP cache. As replies come back, it will refresh the cache. If any entries time out, they will be cleared from the table. The switch router will then build the adjacency table using this information, and then populate the interface EPC IP address table.

If you find inconsistencies between the IP route table, CEF table, and the **epc ip-prefix** table, the **clear ip route** command will rebuild the entries in these tables. You can either clear a specific route using the **clear ip route** *ip-address* command or use the **clear ip route** * command to clear all routes. The routing protocol should relearn the routes and then rebuild the CEF table. The switch router then passes this information to the interfaces and into the ip-prefix tables.

Caution There will be a momentary spike in route processor activity and a corresponding traffic disruption. Use caution when performing the previously described workarounds in a production network.

If you determine that the interface is configured incorrectly, refer to the “Configuring Interfaces” chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Troubleshooting IPX Layer 3 Routing

Similar to troubleshooting IP Layer 3 routing connections, the key to troubleshooting IPX Layer 3 routing is to check on consistency between information contained in the route processor and what is in the CAM tables on the ports.

Troubleshooting an IPX Layer 3 connection is separated into the following processes:

- Checking the IPX Routing Table, page 11-40
- Checking the IPX CEF Adjacencies, page 11-41

Figure 11-20 shows the example network used to troubleshoot an IPX Layer 3 connection in the subsequent examples.

Figure 11-20 IPX Layer 3 Connection

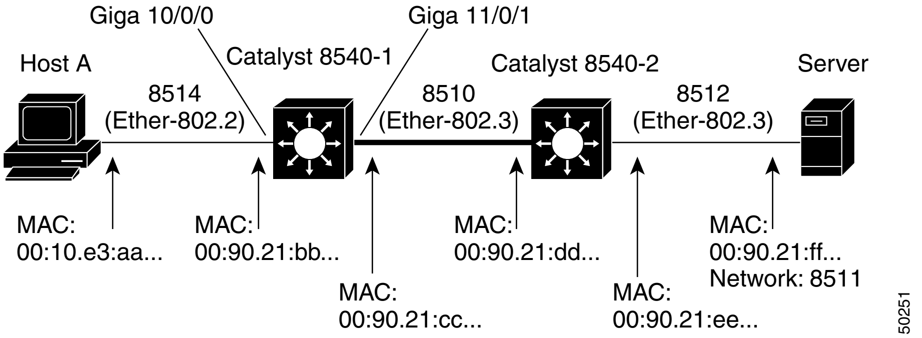

In Figure 11-20, Host A is the source end station trying to communicate with the Novell Server that is part of IPX network 8511, the destination end station.

IPX troubleshooting is similar to IP troubleshooting. The key is to check the consistency between the route processor table information and CAM tables on the ports.

IPX Layer 3 Connection Troubleshooting Commands

To troubleshoot IPX Layer 3 connection problems, use the following commands:

Command	Purpose
show ipx route	Displays the IPX routing table.
show ipx servers	Displays SAP server status information. **Note** Use this command only if you have a server or SAP reachability problem.
show epc ipx-prefix {*prefix-number*} {*netmask*} {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port*	Displays IPX prefix entries for the specified network, node, and interface.
show epc ipx-node {*network-number.node*} **cam** {*cam-address*}	Displays IPX node entry in interface CAM.
show epc ifmapping	Displays interface mapping to CAM interface number.

Command	Purpose
show epc patricia interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **ipx detail** (on the ingress interface)	Displays the IPX patricia tree for the ingress interface.
show epc patricia interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **ipx detail** (on the egress interface)	Displays the IPX patricia tree for the egress interface.
show epc if-entry interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **entry** {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port*	Displays interface entry information for the specific interface.

Checking the IPX Routing Table

Follow these steps to verify the IPX routing tables in the IP Layer 3 connection shown in Figure 11-21.

Figure 11-21 Displaying IPX Router Table Information

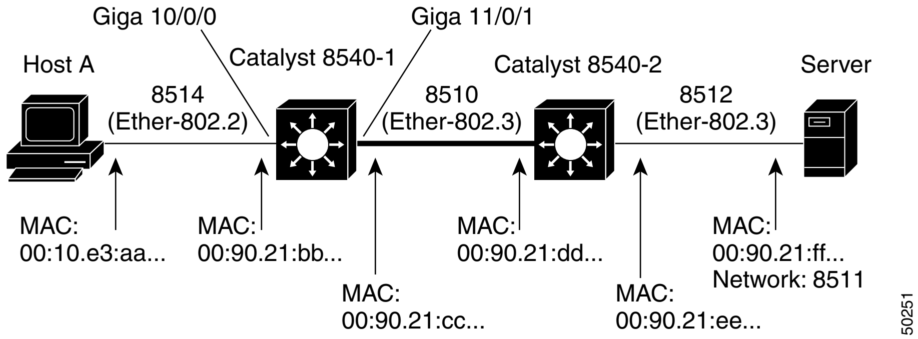

Step 1 From the Catalyst 8540-1, use the **show ipx route** command to verify the status of the IP routing table for the example network shown in Figure 11-21.

```
C8540CSR-1# show ipx route
Codes: C - Connected primary network,    c - Connected secondary network
       S - Static, F - Floating static, L - Local (internal), W - IPXWAN
       R - RIP, E - EIGRP, N - NLSP, X - External, A - Aggregate
       s - seconds, u - uses, U - Per-user static
5 Total IPX routes. Up to 1 parallel paths and 16 hops allowed.
No default route known.
➔ C       8510 (NOVELL-ETHER),   Gi11/0/1
➔ C       8540 (ISL vLAN),      Gi10/0/1.1
➔ C       8541 (SAP),           Gi10/0/0
  R       8511 [05/03] via     8510.0010.7bfa.5f1f,   12s, Gi11/0/1
  R       8512 [02/01] via     8510.0010.7bfa.5f1f,   12s, Gi11/0/1
```

Step 2 Check the Connected primary network (indicated with a “C”). The Novell network numbers 8511 and 8512 should appear as connected through interface Gigabit Ethernet 11/0/1.

Step 3 From the Catalyst 8540-1, use the **show ipx servers** command to verify the connection to the server in Novell network 8511, in the example network shown in Figure 11-21.

```
  C8540CSR-1# show ipx servers
  Codes: S - Static, P - Periodic, E - EIGRP, N - NLSP, H - Holddown, + = detail
  1 Total IPX Servers
  Table ordering is based on routing and server info
      Type   Name      Net    Address      Port    Route    Hops    Itf
➔ P   4      S_8510    8511.0000.0000.0001:0451    5/03     3       Gi11/0/1
```

Step 4 Confirm that the network (Net) number 8511 appears in the Periodic (indicated with a “P”) row of the connection list.

Note SAP entries reside in route processor Memory, not on the CAM tables.

Checking the IPX CEF Adjacencies

Follow these steps to check the IPX CEF adjacencies in the IPX Layer 3 connection shown in Figure 11-22.

Figure 11-22 Checking the IPX CEF Adjacency

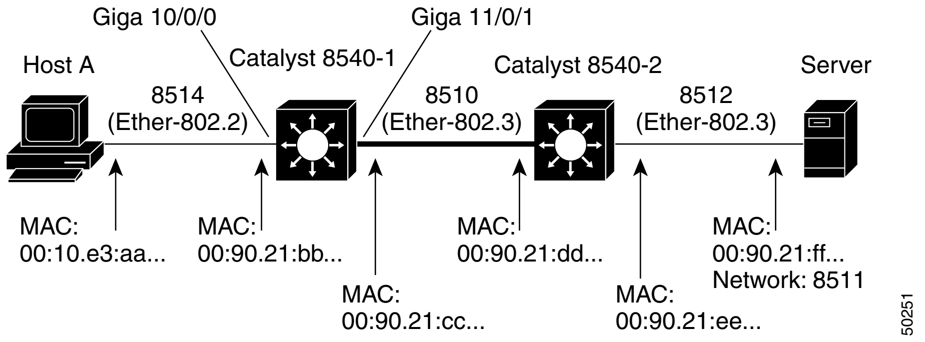

Step 1 Use the **show ipx ipx-prefix** command with the destination network number (8512), netmask, and interface parameters to display the status of the CAM table on the egress interface.

```
  C8540CSR-1# show epc ipx-prefix 8512 00 GigabitEthernet 11/0/1

  IPX Prefix Entries in CAM, Interface GigabitEthernet11/0/1
  -----------------------------------------------------------------
  Codes: C  - Connected network,       R - Remote network
         V  - valid entry,             N - Network entry
         L  - load balancing enabled,  D - default network
         E  - EIGRP enabled,           I - Internal network
         B  - BVI network,             M - My Mac Address
         VC - VCI
➔ GigabitEthernet11/0/1 net 8512 cptr 101D nhop1 101B nhop2 0 encap1 8 encap2 0 flags 9
```

Step 2 Confirm the following:

- net (network number) appears correctly.
- nhop (next hop) number is correct.

- 101B is the Entry CAM location confirmed (this is confirmed using the **show epc patricia interface** command with the **ipx** and **detail** parameters).
- encap (encapsulation) number is correct.

Step 3 This interface CAM information should match the information shown in Step 1 of the “Checking the IPX Routing Table” section on page 11-40.

Step 4 Use the **show ipx cef** command with the source network number (8541), netmask, and interface parameters to display the status of the CAM table on the egress interface.

```
  C8540CSR-1# show epc ipx-prefix 8541 00 GigabitEthernet 11/0/1

  IPX Prefix Entries in CAM, Interface GigabitEthernet11/0/1
  -----------------------------------------------------------------
  Codes: C  - Connected network,       R - Remote network
         V  - valid entry,             N - Network entry
         L  - load balancing enabled,  D - default network
         E  - EIGRP enabled,           I - Internal network
         B  - BVI network,             M - My Mac Address
         VC - VCI
➔ GigabitEthernet11/0/1 net 8541 cptr 1012 nhop1 1014 nhop2 1013 encap1 34 encap2 0 flags B
```

Step 5 Confirm the same parameters as in Step 3.

Step 6 From the Catalyst 8540-1, use the **show epc ipx-node** command with the IPX network and node addresses to display the status of the IPX network to node mapping.

```
  C8540CSR-1# show epc ipx-node 8510.0090.21cc.cccc
  Codes: V - valid entry, M - My-node, I - IF/VC flag
  Interface Network Node IF Number Flags

➔ GigabitEthernet11/0/1 network 8510, cptr 101B, node 0090.21cc.cccc flag 275
```

Step 7 From the Catalyst 8540-1, use the **show epc ifmapping** command to display the IF number mapped to the egress interface GigabitEthernet 11/0/1.

```
  C8540CSR-1# show epc ifmapping
  .
  (Information Deleted)
  .
➔ GigabitEthernet11/0/1    (IF number: 39)
```

The IF number field (in this example “39”) is used in Step 9.

Step 8 From the Catalyst 8540-1, use the **show epc patricia interface** command with the **ipx detail** parameters on the egress interface to display the status of the Host Entry CAM location for the connection to Host B.

```
  C8540CSR-1# show epc patricia interface gigabitEthernet 11/0/1 ipx detail
  0# CAM location: 0x0FF7  ROOT
➔ 2# Prefix Entry CAM location: 0x1018   Dirty
     Prefix 0x8510     CONNECTED NTP 0x101A  NTP 0x1019  Valid
    1. Node Entry CAM location: 0x101A  Dirty
       0090.21bb.bbbb interface 39 My-Node Valid
   2. Node Entry CAM location: 0x101B  Dirty
       0010.21aa.aaaa interface 39 Valid
  3# Prefix Entry CAM location: 0x101D   Dirty
➔    Prefix 0x8512     REMOTE NHOP1 0x101B  NOVELL_ETHER  Valid
  4# Prefix Entry CAM location: 0x101C   Dirty
➔   Prefix 0x8511     REMOTE NHOP1 0x101B  NOVELL_ETHER  Valid
  IPX Patricia Tree Summary:
     Number of IPX prefix entries: 5
     Number of Host Entries: 4
```

Look at entry 2#. The word “dirty” in the display is a normal entry type. The prefix (IPX network number) and node numbers are displayed. The entry marked “My-Node Valid” is for the directly connected interface on Catalyst 8540-1. The other node entry, marked Valid, is for host A on the network. Make a note of the hexadecimal address **0x101B** (converted to decimal 4123). You need that hexadecimal address, converted to decimal 4123, in Step 9.

Entries 3 and 4 are remote entries. NHOP1 means these are pointers to the adjacency entry for the next hop to get to the IPX networks Prefix 0x8512 and Prefix 0x8511. These are not the MAC addresses of the next hop. Valid means the entry is valid and usable.

Step 9 From the Catalyst 8540-1, use the **show epc cam interface** command with the CAM location hexadecimal address **0x101B** (converted to decimal **4123**) and the CAM word **2** parameters to display the status of the MAC rewrite for this interface.

```
  C8540CSR-1# show epc cam interface gigabitEthernet 11/0/1 4123 2
  .
  (Information Deleted)
  .
➔ GigabitEthernet11/0/1 Addr:0x101B Word:2 Data[0]:0x009021DD Data[1]:0xDDDD0275
```

The ingress interface fields in the display indicate the MAC address is written to the following:

- 0x**009021DD**—the first four bytes of the next-hop MAC address
- 0x**DDDD**0275—the last two bytes of the next-hop MAC address
- 0xdddd0**27**5—these last two bytes, in this example “0**27**5,” translates to hexadecimal 0x27 == 39 (decimal), which matches the interface number 39 that appears in the **show epc ifmapping** command mapped to the egress interface GigabitEthernet 11/0/1 in Step 7.

Step 10 From the Catalyst 8540-1, use the **show epc if-entry interface** command with the **entry** interface parameters to display the status of the Broute VC.

```
  C8540CSR-1# show epc if-entry interface gigabitEthernet 11/0/1 entry gigabitEthernet
  10/0/0
  IF Entry for GigabitEthernet10/0/0 on GigabitEthernet11/0/1
      Mac(hex) - 00:90:21:CC:CC:CC
      isMyInteface : False isSubInterface : False
➔     Status Up Broute VC - 412 Bcast VC - 0
      Netmask: 32
      FEC disabled
      Trunking Disabled
      State : Not-Applicable/Listening/Blocking
      Bridge-Group disabled
      IP routing off bridging off
➔     IPX routing on bridging off
      Appletalk routing off
➔     In Encapsulation: ET_SAP
      ICMP Redirect enabled Unreachable enabled
      IP Multicast disabled: ttl-threshold: 0

  C8540CSR-1#
```

Step 11 Check the following:

- Status field to ensure that the Broute VC is up.
- IPX routing field to ensure that it is on.
- Encapsulation field to ensure that it is set to ET_SAP.

If you have any problems with these fields, check the interface configuration. For information about configuring interfaces, refer to the *Layer 3 Software Feature and Configuration Guide*.

Troubleshooting Layer 3 IP Multicast Switching

IP multicast allows IP traffic to be sent from one source or multiple sources and delivered to multiple destinations. Instead of sending individual packets to each destination, which is highly taxing to the switch fabric, a single packet is sent to a multicast group, which is identified by a single IP destination group address. That IP destination group consists of a number of IP destinations that require that frame. From a router perspective, an input multicast feed from a given source must be sent out through (possibly) multiple output interfaces based on the information received by the multicast routing protocols such as PIM.

Layer 3 IP Multicast Overview

The Layer 3 enabled ATM switch router supports IP multicast at wire speed for all ports, allowing for high-speed switching of packets from input source ports to multiple destination ports. The Layer 3 enabled ATM switch router also supports IP multicast routing protocols such as PIM dense and sparse modes, as well as DVMRP interoperability.

Internet Group Management Protocol

The Internet Group Management Protocol (IGMP) provides a method for end stations to request multicast traffic as well as for switch router to determine who on a locally attached segment is requesting traffic. IGMP uses IP datagrams to allow IP multicast applications to join a multicast group. IGMP relies on Class D IP addresses for the creation of multicast groups and is defined in RFC 1112. Membership in a multicast group is dynamic, meaning that it changes over time as hosts join and leave the group. Multicast switch routers use IGMP host-query messages (sent to the group address 224.0.0.1 with a TTL of 1) to keep track of the hosts that belong to multicast groups. When switch router receives a packet addressed to a multicast group, it forwards the packet to those interfaces that have hosts belonging to that group. Switch routers periodically send host-query messages to refresh their multicast group membership knowledge.

The Catalyst 8500 supports both IGMP version 1, which most end stations currently support, and IGMP version 2, which, unlike version 1, provides support for clients informing the network that they are leaving a multicast group.

Protocol Independent Multicast

As networks increase in size, multicast routing becomes critically important in order to determine, in a large routed network, which segments require multicast traffic and which do not. PIM is a routing protocol for multicast that uses existing unicast routing protocols such as RIP or OSPF for path forwarding determination and network location. PIM can be operated in two modes. PIM dense mode and PIM sparse mode. The mode selected determines how the switch router populates its multicast routing table, and how the it forwards multicast packets it receives from its directly connected LANs.

Note Enabling PIM on an interface also enables IGMP operation on that interface.

Dense Mode

In dense mode, a switch router assumes that all other switch routers want to forward multicast packets for a group. Therefore, interfaces with PIM dense mode enabled receive the multicast feed as soon as a single user requests one. That segment will continue to receive the multicast until it times out. If a Catalyst 8500 receives a multicast packet and has no directly attached members or PIM neighbors present, a prune message is sent back to the source. Subsequent multicast packets are not flooded to this pruned branch. PIM builds source-based multicast distribution trees. PIM dense mode is most useful when:

- The senders are receivers are in close proximity to one another
- There are fewer senders than receivers
- Multicast traffic volume is high
- The stream of multicast traffic is constant

Sparse Mode

In sparse mode, a switch router assumes that other switch routers do not want to forward multicast packets for a group, unless there is an explicit request for the traffic. When hosts join a multicast group, the directly connected switch routers send PIM join messages to the rendezvous point (RP). The RP keeps track of multicast groups. Hosts that send multicast packets are registered with the RP by that host's first-hop switch router. The RP then send joins toward the source. At this point, packets are forwarded on a shared distribution tree. When the data stream begins to flow from sender to RP to receiver, the switch routers in the path optimize the path, automatically, to remove any unnecessary hops. Sparse mode assumes that no hosts want the multicast traffic unless they specifically ask for it.

Sparse mode PIM is optimized for environments where there are many multipoint data streams and each multicast stream goes to a relatively small number of LANs in the internetwork. PIM sparse mode is most useful when:

- There are few receivers in a group
- Senders and receivers are separated by WAN links
- The type of traffic is intermittent

There are two types of rendezvous points: statically configured and Auto-RP.

A statically configured PIM rendezvous point (RP) address is used by first-hop switch routers to send Register packets on behalf of source multicast hosts. The RP address is also used by switch routers on behalf of multicast hosts that want to become members of a group. These switch routers send Join and Prune messages toward the RP. A single RP can be configured for all multicast groups or a subset of the Class D address range as described by the access-list pointer.

Auto-RP automates the distribution of group-to-RP mappings in a PIM network. This feature has the following benefits:

- It is easy to use multiple RPs within a network to serve different group ranges.
- It allows load splitting among different RPs and arrangement of RPs according to the location of group participants.
- It avoids inconsistent, manual RP configurations that can cause connectivity problems.

Multiple RPs can be used to serve different group ranges or serve as hot backups of each other. To make Auto RP work, a Layer 3 enabled ATM switch router must be designated as an RP-mapping agent, which receives the RP-announcement messages from the RPs and arbitrates conflicts. The RP-mapping agent then sends the consistent group-to-RP mappings to all other switch routers. Thus, all switch routers automatically discover which RP to use for the groups they support.

One way to start is to place (preserve) the default route processor for all global groups at or near the border of your routing domain, while placing another route processor in a more centrally located switch router for all local groups using the administratively scoped addresses (239.x.x.x).

Note If you configure PIM in sparse mode or sparse-dense mode and do not configure Auto-RP, you must statically configure an RP.

Distance Vector Multicast Routing Protocol

DVMRP is the first-generation multicast routing protocol most known for its use in the Multicast Backbone (MBONE). DVMRP uses a flood-and-prune approach to multicast packet delivery. This means that DVMRP assumes that all other switch routers in a network want to forward multicast packets for a group. This creates huge scalability problems, as switch routers must now maintain state for multicast paths that may not require or want to handle multicast traffic. For that reason, the Cisco switch router does not support DVMRP, but does support DVMRP interoperability with PIM. This allows the Cisco switch router to interoperate with non-Cisco multicast switch routers that use DVMRP.

Cisco IOS software in the Catalyst 8500 supports dynamic discovery of DVMRP switch routers, and can interoperate with them over traditional media or over DVMRP-specific tunnels. When a DVMRP neighbor has been discovered, the switch router periodically transmits DVMRP report messages advertising the unicast sources reachable in the PIM domain.

When a Cisco switch router runs DVMRP over a tunnel, it advertises sources in DVMRP Report messages much as it does on real networks. In addition, the software caches DVMRP Report messages it receives and uses them in its Reverse Path Forwarding (RPF) calculation. This allows the software to forward multicast packets received over the tunnel.

Essential to multicast routing is the idea of spanning trees. Multicast routing procedures, for example PIM, construct these trees (with receivers as leafs), while multicast forwarding forwards multicast packets a long the trees.

To support a multicast forwarding function with tag switching, each tag switch associates a tag with a multicast tree as follows. When a tag switch creates a multicast forwarding entry (either for a shared or for a source-specific tree), and the list of outgoing interfaces for the entry, the switch also creates local tags (one per outgoing interface). The switch creates an entry in its TIB and populates (outgoing tag, outgoing interface, outgoing MAC header) with this information for each outgoing interface, placing a locally generated tag in the outgoing tag field. This creates a binding between a multicast tree and the tags. The switch then advertises over each outgoing interface associated with the entry the binding between the tag (associated with this interface) and the tree.

When a tag switch receives a binding between a multicast tree and a tag from another tag switch, if the other switch is the upstream neighbor (with respect to the multicast tree), the local switch places the tag carried in the binding into the incoming tag component of the TIB entry associated with the tree. When a set of tag switches are interconnected via a multiple-access subnetwork, the tag allocation procedure for multicast has to be coordinated among the switches. In all other cases tag allocation procedure for multicast could be the same as for tags used with destination-based routing.

Cisco Group Membership Protocol

Cisco Group Membership Protocol (CGMP) addresses the issue of efficiently forwarding IP multicast packets across Layer 2 switches. CGMP allows Layer 2 switches to leverage IGMP information recorded on the Catalyst 8500 to make intelligent Layer 2 forwarding decisions based on the destinations requesting the multicast traffic. The net result is that with CGMP, IP multicast traffic is delivered only to those Layer 2 switch ports that are interested in multicast traffic. All Layer 2 switch ports that have not requested the traffic do not receive it. When a Layer 3 enabled ATM switch router receives an IGMP join message, it records the source MAC address of the IGMP message, and turns around and issues a CGMP join message downstream, to a Layer 2 switch. The switch uses the CGMP message to dynamically build an entry in the switching table that maps the multicast traffic to the client switch port.

The Catalyst 8500 uses PIM, not CGMP, for multicast forwarding determination. However, the Catalyst 8510 does function as a CGMP server, meaning that on a per-interface basis, it informs the connected LAN switch of multicast groups that it needs to be aware of. The Catalyst 8500 responds to IGMP version 1 and 2 multicast join and leave (for IGMP v2) requests and forwards them on the multicast tree via PIM.

The Multicast Routing Table

The Cisco IOS software running on the switch router uses PIM and DVMRP interoperability to exchange IP multicast network information. Each routing protocol runs as a separate IOS process in the SRP. The multicast routing table is a centralized routing information database that is resident on the SRP. The packet forwarding engine consults the routing table to route the packets to appropriate destinations.

A multicast routing table is different than a unicast routing table. A multicast routing table maps an ordered pair consisting of a source IP address and a multicast group to an ordered pair consisting of an input interface and a set of output interfaces. Packets from the given source to the given multicast group that arrives over an input interface are appropriate output interfaces.

Packets that arrive on the wrong input interface are discarded.

The switch router maintains the central multicast routing table at the SRP. By using CEF and the associated distribution of the forwarding information base (FIB), the line cards can forward multicast traffic intelligently, based on the multicast topology of the network. This feature allows the input port to decide which output interfaces require the multicast traffic, and inform the switching fabric about which output ports to direct that packet to. Any change in the multicast routing table is instantly downloaded to the line cards, allowing the switch router to maintain a constant, up-to-date map of the network.

MSDP

In the PIM-SM model, multicast sources and receivers must register with their local RP. Actually, the switch router closest to the sources or receivers registers with the RP, but the key point to note is that the RP knows about all the sources and receivers for any particular group. RPs in other domains have no way of knowing about sources located in other domains. MSDP solves this problem.

MSDP allows RPs to share information about active sources. RPs know about the receivers in their local domain. When they hear about active sources through MSDP, they can pass on that information to their local receivers and multicast data can be forwarded between the domains directly. A useful feature of MSDP is that it allows each domain to maintain an independent RP that does not rely on other domains.

The RP in each domain establishes an MSDP peering session using a TCP connection with the RPs in other domains, or with border switch routers leading to the other domains. When the RP learns about a new multicast source within its own domain (through the normal PIM register mechanism), the RP encapsulates the first data packet in a Source Active (SA) message and sends the SA to all MSDP peers. The SA is forwarded by each receiving peer using a modified RPF check, until the SA reaches every MSDP switch router in the interconnected networks—theoretically the entire multicast internet. If the receiving MSDP peer is an RP, and the RP has a (*, G) entry for the group in the SA (there is an interested receiver), the RP will create (S, G) state for the source and join to the shortest path tree for the source. The encapsulated data will be decapsulated and forwarded down that shared tree of that RP.

When the packet is received by the last-hop switch router of a receiver, the last-hop switch router may also join the shortest path tree to the source. The MSDP speaker periodically sends source addresses that include all sources within that RP domain.

For detailed configuration information see the IOS document, *Configuring IP Multicast Routing.*

IP Multicast Troubleshooting Commands

To troubleshoot an IP multicast problem, use the following commands:

Command	Purpose
show interfaces {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port* (on the ingress interface)	Displays interface configuration, status, and statistics on the ingress interface.
show interfaces {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port* (on the egress interface)	Displays interface configuration, status, and statistics on the egress interface.
show ip mroute	Displays the IP multicast routing table.
show epc if-entry interface {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port* **all** (on the ingress interface)	Displays all interface entry information for the ingress interface.
show epc ipmcast *groupaddr* **all interface** {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port*	Displays the IP multicast routing table information stored on the ingress interface for a particular group IP address.
show epc ipmcast *groupaddr* **detail interface** {**fastethernet** | **gigabitethernet**}	Displays detailed IP multicast routing table information stored on the ingress interface for a particular group and source IP address.
show atm vc cast-type p2mp interface {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port*	Displays the ATM VC cast type point to multi-point configuration for a specific interface.

IP multicast troubleshooting is similar to IP troubleshooting. The key is to check the consistency between the route processor table information and CAM tables on the interfaces.

Follow these steps to troubleshoot IP multicast problems:

Step 1 Use the **show epc if-entry** command to display information about VC status:

```
    C8540CSR-1# show epc if-entry interface fastethernet 1/0/15 entry gigabitethernet 0/0/0
    IF Entry for GigabitEthernet0/0/0 on FastEthernet1/0/15
        Mac(hex) - 00:90:21:41:BC:07
        isMyInteface : False isSubInterface : False
➔       Status Up Broute VC - 67 Bcast VC - 0
        Netmask: 24
        FEC disabled
        Trunking Disabled
        State : Not-Applicable/Listening/Blocking
        Bridge-Group disabled
➔       IP routing on bridging off
        IPX routing off bridging off
        Appletalk routing off
        In Encapsulation:
        ICMP Redirect enabled Unreachable enabled
➔       IP Multicast enabled: ttl-threshold: 5
```

Step 2 Check the following:

- Status field to ensure that the Broute VC is up.
- IP routing field to ensure that it is on.
- IP Multicast field to ensure that it is enabled.

If you have any problems with these fields, check the interface configuration. For information about configuring interfaces, refer to the *Layer 3 Software Feature and Configuration Guide*.

Step 3 Display the IP multicast entries contained in the route processor using the **show ip mroute** command.

```
    C8540CSR-1# show ip mroute
    IP Multicast Routing Table
    Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned
           R - RP-bit set, F - Register flag, T - SPT-bit set, J - Join SPT
           X - Proxy Join Timer Running
           Outgoing Interface Flags: H - Hardware switched
    Timers: Uptime/Expires
    Interface state: Interface, Next-Hop or VCD, State/Mode

    (*, 224.2.236.92), 00:58:34/00:03:09, RP 10.6.11.10, flags: S
      Incoming interface: POS12/0/0, RPF nbr 10.6.11.10
      Outgoing interface list:
        FastEthernet3/0/13, Forward/Sparse, 00:57:56/00:03:09
        FastEthernet2/0/15, Forward/Sparse, 00:58:13/00:02:53

➔   (10.64.1.19, 224.2.236.92), 00:58:13/00:03:22, flags: T
➔     Incoming interface: POS12/0/0, RPF nbr 10.6.11.10
      Outgoing interface list:
        FastEthernet3/0/13, Forward/Sparse, 00:57:56/00:03:08
        FastEthernet2/0/15, Forward/Sparse, 00:58:13/00:02:53
```

Step 4 Use the address and interface information from the **show ip mroute** command output in Step 3 to display the CAM information with the **show epc ipmcast** command.

```
  C8540CSR-1# show epc ipmcast 224.2.236.92 10.64.1.19 detail interface pos 12/0/0
➔ MEMBER_ENTRY, root vc = 0/801, packet counter = 47
  (224.2.236.92, 10.64.1.19), CAM Loc 0x17102, 00 34 48 00 00 2F 32 11
➔ Send_to_cpu flag not set, SPT flag set

  p2mp vc:root   POS12/0/0,          VPI = 0, VCI = 801
          leaf   FastEthernet2/0/15, VPI = 0, VCI = 762
                 FastEthernet3/0/13, VPI = 0, VCI = 751
```

Step 5 Check the following:

- Multicast group 224.2.236.92 and source 10.64.1.19 has a CAM entry on interface POS 12/0/0.
- The Send_to_cpu flag is appropriately not set for a specified source (S, G) within a group indicating that the traffic is switched in the data plane by the interface. The Send_to_cpu flag is set for table entries for all sources within a group (*, G) to maintain the state for this entry on the control plane.

Step 6 Display the status of the VC for the incoming interface displayed in the **show ip mroute** command output in Step 3 using the **show atm vc cast-type p2mp** interface command.

```
  C8540CSR-1# show atm vc cast-type p2mp interface pos 12/0/0
  Interface         VPI  VCI   Type   X-Interface      X-VPI X-VCI Encap  Status
  .
  (Information deleted)
  .
➔ POS12/0/0         0    801   PVC    Fa2/0/15          0    762          UP
```

Step 7 Check the following:

- The VC identifier in the VPI and VCI columns matches the corresponding interface listed in the **show epc ipmcast** command output shown in Step 4.
- The VC identifier listed in the X-VPI and X-VCI columns matches the entry for the corresponding interface listed in the **show epc ipmcast** command output shown in Step 4.

If there are inconsistencies or non-zero invalid entries in the tables, you can use the **clear ip mroute *** command to rebuild the tables.

Caution Use the **clear ip mroute** command carefully. It causes a temporary increase in switch router activity, which can lead to traffic disruptions.

Troubleshooting IP and IPX Load Balancing

The Layer 3 enabled ATM switch router currently supports only two paths for IP and IPX. If there are more than two paths in the FIB table the switch router uses the first two.

Note To reduce the number of unnecessary IPC messages, use a maximum paths statement of two for both IP and IPX.

For IP: Load balancing is accomplished by using the Boolean function XOR on the least significant bit (LSB) of the source and destination IP addresses. If the bit is set, use the second path; if not, use the first path.

For IPX: Load balancing is accomplished by using the Boolean function XOR on the LSB of the IPX source network and destination IPX network. If the bit is set, use the second path; if not, use the first path.

By default IPX will only maintain one path in the IOS IPX routing table.

To get IPX to use more than one path use the global configuration command **ipx maximum-paths** *number*.

Note Even if you set the **ipx maximum-paths** command number to a number greater than two the interface module CAM still only maintains two paths.

Troubleshooting IP and IPX Load Balancing Commands

To display the interface configuration, use the following commands:

Command	Purpose
show epc ip-prefix interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **all-entries**	Displays all ip prefix entries for the specified interface.
show epc ipx-prefix prefix	Displays all IPX prefix entries for the specified interface.

Follow these steps to troubleshoot IP load balancing:

Step 1 Use the **show epc ip-prefix interface** command with the all-entries parameter to confirm the configuration of IP load balancing

```
C8540CSR-1# show epc ip-prefix interface FastEthernet 1/0/15 all-entries
Default Network Information:
     Nexthop 1:
       IP addr:20.0.0.1   GigabitEthernet2/0/1 (58)
       Mac Addr:0090.2141.bd47
➔    Load Balancing:Off
    Not configured
 Prefix/Masklen       Next Hop
 0.0.0.0/32           not populated
 10.0.0.7/32          20.0.0.1
 10.0.1.4/30          20.0.0.1
 10.0.1.12/30         20.0.0.1
.
(Information Deleted)
.
```

Step 2 Check the Not configured field. This indicates no default route is known. If you added IP route 0.0.0.0 20.0.0.1 to that configuration, the display would change to include the following:

```
Default Network Information:
     Nexthop 1:
       IP addr:20.0.0.1   GigabitEthernet2/0/1 (58)
       Mac Addr:0090.2141.bd47
     Load Balancing:Off
```

Note Since there is only one route in the example, the Load Balancing field is Off.

Follow these steps to troubleshoot IPX load balancing:

Step 1 Use the **show epc ipx-prefix interface** command with the source network number (8512), netmask, and interface parameters to display the status of the CAM table on the egress interface, to check the load balancing configuration.

```
C8540CSR-1# show epc ipx-prefix 8512 00 GigabitEthernet 11/0/1

IPX Prefix Entries in CAM, Interface GigabitEthernet11/0/1
-----------------------------------------------------------------
Codes: C  - Connected network,       R - Remote network
       V  - valid entry,             N - Network entry
       L  - load balancing enabled,  D - default network
       E  - EIGRP enabled,           I - Internal network
       B  - BVI network,             M - My Mac Address
       VC - VCI
➔ GigabitEthernet11/0/1 net 8512 cptr 101D nhop1 101B nhop2 0 encap1 8 encap2 0 flags 9
```

Step 2 Confirm that the load balancing enabled “L” code appears in the output.

Troubleshooting Route Processor Route Table and Utilization Problems

The following list describes common symptoms of high route processor utilization. If you notice any of these symptoms, follow the troubleshooting steps in this section to fix the problem.

- High percentages in the **show processes cpu** command output
- Input queue drops
- Slow performance
- Services on the switch router fail to respond, for instance:
  - Slow response in Telnet or unable to Telnet to the switch router
  - Slow response on the console
  - Slow or no response to ping
  - Switch router does not send routing updates

For additional information about troubleshooting route processor problems, see the IOS document *Troubleshooting High CPU Utilization on Cisco Routers*.

Troubleshooting Route Processor Route Table Problems Commands

To display the route processor route table statistics, use the following commands:

Command	Purpose
show processes cpu	Displays information about the active processes in the switch router and their corresponding route processor utilization statistics.
show {**ip** | **ipx**} **traffic**	Displays IP and IPX traffic statistics.
show ip spd	Displays selective packet discard configuration for the switch.
show epc spd	Displays selective packet discard configuration for the interfaces.

Troubleshooting Route Processor Route Table Problems

This section describes common symptoms and causes of, and solutions to, high route processor utilization on your switch router.

For additional information about high route processor utilization, refer to the Troubleshooting High CPU Utilization on Cisco Routers web page at the URL: http://www.cisco.com/warp/public/63/highcpu.html.

Follow these steps to troubleshoot route processor route table problems:

Step 1 Use the **show processes cpu** command to check the route processor route table and processes.

```
  Switch# show processes cpu
➔ CPU utilization for five seconds: 99%/24%; one minute: 25%; five minutes: 8%
  PID  Runtime(ms)  Invoked  uSecs    5Sec   1Min   5Min TTY Process
     1           8     2750      2   0.00%  0.00%  0.00%   0 Load Meter
     2       69168 14972355      4   0.00%  2.38%  0.88%   0 Exec
     3       13940     1771   7871   0.00%  0.10%  0.11%   0 Check heaps
     4         536      541    990   0.00%  0.00%  0.00%   0 Pool Manager
     5           0        2      0   0.00%  0.00%  0.00%   0 Timers
     6          36      301    119   0.00%  0.00%  0.00%   0 ARP Input
  .
  (Information Deleted)
  .
➔   63      196252     40503   4845   66.72% 25.93%  6.25%   0 IP-EIGRP Router
```

Step 2 Check the CPU utilization for five seconds field. In this example, it indicates the CPU has spiked to 99% with 24% at the interrupt level, where 99%/24% is equal to the following:

- 99%—Average total utilization during last five seconds
- 24%—Average utilization due to interrupts, during last five seconds
- 99 – 24 = 75—Percentage of traffic being process-switched

Note If the CPU utilization in the example indicated 99%/24%, that means the route processor is being consumed by interrupt-driven processes.

Step 3 Scan down the process list to identify which process is contributing to the 75% CPU process utilization. From this example of EIGRP convergence, you can see that the IP- EIGRP Router process is accounting for 66.72% of the 75% CPU process utilization. Use this same process to identify other processes.

Step 4 Use the **show ip traffic** command to check whether the packets sent to the route processor are being processed.

```
Switch# show ip traffic
IP statistics:
  Rcvd:  198650 total, 198639 local destination
         0 format errors, 0 checksum errors, 0 bad hop count
         0 unknown protocol, 0 not a gateway
         0 security failures, 0 bad options, 265609 with options
  Opts:  0 end, 0 nop, 265609 basic security, 0 loose source route
         0 timestamp, 0 extended security, 0 record route
         0 stream ID, 0 strict source route, 0 alert, 0 cipso
         0 other
  Frags: 0 reassembled, 0 timeouts, 0 couldn't reassemble
         0 fragmented, 0 couldn't fragment
  Bcast: 225 received, 134130 sent
  Mcast: 0 received, 166103 sent
  Sent:  291558 generated, 10 forwarded
  Drop:  44536 encapsulation failed, 0 unresolved, 0 no adjacency
         0 no route, 0 unicast RPF, 0 forced drop

ICMP statistics:
  Rcvd: 0 format errors, 0 checksum errors, 0 redirects, 10 unreachable
        110 echo, 14 echo reply, 0 mask requests, 0 mask replies, 0 quench
        0 parameter, 0 timestamp, 0 info request, 0 other
        0 irdp solicitations, 0 irdp advertisements
  Sent: 2 redirects, 8 unreachable, 25 echo, 110 echo reply
        0 mask requests, 0 mask replies, 0 quench, 0 timestamp
        0 info reply, 0 time exceeded, 0 parameter problem
        0 irdp solicitations, 0 irdp advertisements

IP-IGRP2 statistics:
  Rcvd: 158367 total
  Sent: 166052 total

UDP statistics:
  Rcvd: 19201 total, 0 checksum errors, 190 no port
  Sent: 108759 total, 0 forwarded broadcasts

TCP statistics:
  Rcvd: 20927 total, 0 checksum errors, 0 no port
  Sent: 16564 total

Probe statistics:
  Rcvd: 0 address requests, 0 address replies
        0 proxy name requests, 0 where-is requests, 0 other
  Sent: 0 address requests, 0 address replies (0 proxy)
        0 proxy name replies, 0 where-is replies

OSPF statistics:
  Rcvd: 0 total, 0 checksum errors
        0 hello, 0 database desc, 0 link state req
        0 link state updates, 0 link state acks
  Sent: 0 total

PIMv2 statistics: Sent/Received
  Total: 25/0, 0 checksum errors, 0 format errors
  Registers: 0/0, Register Stops: 0/0,  Hellos: 25/0
  Join/Prunes: 0/0, Asserts: 0/0, grafts: 0/0
  Bootstraps: 0/0, Candidate_RP_Advertisements: 0/0

IGMP statistics: Sent/Received
  Total: 27/0, Format errors: 0/0, Checksum errors: 0/0
  Host Queries: 13/0, Host Reports: 13/0, Host Leaves: 1/0
  DVMRP: 0/0, PIM: 0/0

IGRP statistics:
  Rcvd: 0 total, 0 checksum errors
  Sent: 0 total

ARP statistics:
  Rcvd: 6481 requests, 1388 replies, 0 reverse, 0 other
  Sent: 1465 requests, 29954 replies (42 proxy), 0 reverse
C8540CSR-1#
```

Step 5 Check, for example, TCP packets with options set, UDP broadcasts or packets with checksum errors, and ARP packets.

Step 6 For IPX routing, use the **show ipx traffic** command to check if the packets sent to the route processor are being processed.

```
  Switch# show ipx traffic
  System Traffic for 0.0000.0000.0001 System-Name: domino
➔ Rcvd: 0 total, 0 format errors, 0 checksum errors, 0 bad hop count,
  0 packets pitched, 0 local destination, 0 multicast
➔ Bcast: 0 received, 0 sent
  Sent: 0 generated, 0 forwarded
  0 encapsulation failed, 0 no route
  SAP: 0 Total SAP requests, 0 Total SAP replies, 0 servers
  0 SAP general requests, 0 ignored, 0 replies
➔ 0 SAP Get Nearest Server requests, 0 replies
  0 SAP Nearest Name requests, 0 replies
  0 SAP General Name requests, 0 replies
  0 SAP advertisements received, 0 sent, 0 Throttled
  0 SAP flash updates sent, 0 SAP format errors
```

Step 7 Check Bcast, Get Nearest Server (GNS), checksum errors, bad hop count.

Troubleshooting Route Processor Selective Packet Discard Problems

Selective Packet Discard (SPD) is used in the Layer 3 enabled ATM switch router when the following occurs:

- A route or interface flap will cause a burst of packets to be process-switched until the CAM or ternary CAM (TCAM) cache is repopulated.
- The switch router is switching a high volume of traffic, which may cause the input queue to overrun and eventually throttle the interface.
- Packets are blindly dropped, which may affect routing updates and keepalives, causing more invalidations.

SPD avoids dropping high precedence packets:

- Beyond a given threshold of queuing the input queue, packets with a low precedence are randomly dropped, and then always dropped if queuing persists.
- The Layer 3 enabled ATM switch router route processor does the selective discard.
- Packets with high precedence are queued in “headroom” and processed before low precedence packets. These high precedence packets are typically routing protocol packets.

Some of the information does not apply to CEF based forwarding which the Layer 3 enabled ATM switch router uses. However, you can use this information to see what is being dropped by the Fast Ethernet interfaces and the route processor.

Note SPD is enabled by default on the switch router.

Follow these steps to troubleshoot SPD:

Step 1 Use the **show ip spd** command to confirm the SPD settings.

```
Switch# show ip spd
Current mode: normal.
Queue min/max thresholds: 8/9, Headroom: 1024
IP normal queue: 0, priority queue: 0.
SPD special drop mode: none
Switch#
```

Step 2 Check the Queue min/max thresholds field. This determines when the lower-priority packets are discarded. Typically, lower-priority packets are discarded when the input queue size hits min-threshold. When the max-threshold is reached all lower-priority packets are dropped. For all the switch routers, the min/max queue thresholds are almost the same, if there are more than 75 packets in the input queue and all lower-priority packets will be discarded.

Step 3 Check the Headroom field. This indicates how many high-precedence packets will be enqueued over the normal input hold queue limit. This is to reserve room for incoming high precedence packets.

Since the switch router is a nonblocking switch, the lower-priority packets will actually be dropped by the route processor or the switch fabric, but the counters will be shown on the interfaces in Step 4.

Step 4 Use the **show epc spd** command to check the SPD on the interfaces.

```
Switch# show epc spd
INPUT-INT                 TOT-DROPS  PRIORITY-RCVD  PRIORITY-DROPS  NO-BUFS
FastEthernet3/0/0               0       7813353             0             0
FastEthernet3/0/1               0       7773376             0             0
FastEthernet3/0/2               0       7773593             0             0
FastEthernet3/0/3               0       7773568             0             0
FastEthernet3/0/4               0       7773593             0             0
FastEthernet3/0/5               0       7812594             0             0
FastEthernet3/0/6               0       7773593             0             0
FastEthernet3/0/7               0       7773569             0             0
FastEthernet3/0/8               0       7773592             0             0
FastEthernet3/0/9               0       7773592             0             0
FastEthernet3/0/10              0       7812705             0             0
FastEthernet3/0/11              0       7773567             0             0
FastEthernet3/0/12              0       7812538             0             0
FastEthernet3/0/13              0       7773642             0             0
FastEthernet3/0/14              0       7773591             0             0
FastEthernet3/0/15              0       7812666             0             0
ATM0                            0         45177             0             0
Ethernet0                       0             0             0             0
```

High precedence packets are Layer 2 and Layer3 control protocol traffic carried on Stream ID 35. Lower priority packets are carried on Stream ID 36, which would be used for traffic where there is no entry in the CAM table.

Troubleshooting SDM Problems

This section describes the switching database manager (SDM) features built into your switch router. This chapter includes the following topics:

- SDM Overview, page 11-57
- Troubleshooting SDM Problem Commands, page 11-59

The information in this section applies to the Catalyst 8540 CSR and Catalyst 8540 MSR with Layer 3 functionality.

For detailed SDM configuration information, refer to the "Configuring Switching Database Manager" chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

SDM Overview

SDM partitions TCAM space into multiple regions. Each region is protocol specific. SDM interacts with the individual protocol control layer to store Layer 3 switching information. SDM consists of the following types of regions:

- Exact-match region—The exact-match region consists of Layer 3 entries for multiple protocol regions such as IP adjacencies and IPX node.
- Longest-match region—Each longest-match region consists of multiple "buckets" or groups of Layer 3 address entries organized in decreasing order by mask length. All entries within a bucket share the same mask value and key size. The buckets can change their size dynamically by borrowing address entries from neighboring buckets. Although the size of the whole protocol region is fixed, you can reconfigure it. The reconfigured size of the protocol region is effective only at the next system reboot.

- First-match region—The first-match region consists of ACL entries; lookup stops at first match of the entry.

The enhanced Gigabit Ethernet interface module supports TCAM sizes of 32 KB, 64 KB, or 256 KB. Each entry in TCAM is 32 bits wide. Since SDM is responsible for managing TCAM space, SDM partitions the entire TCAM space for each protocol region based on user configuration. A change in the partition configuration takes effect only during the next system reboot.

Table 11-2 lists default partitioning for each protocol region in TCAM.

Table 11-2 TCAM Protocol Region Default Partitioning

Protocol Region	Lookup Type	Key Size	Default Size	No. of TCAM Entries
ipx-bvi-network	Exact-match	32 bits	32	32
ip-adjacency	Exact-match	32 bits	2048	2048
ipx-node	Exact-match	64 bits	2048	4096
ip-prefix	Longest-match	32 bits	8192	8192
ipx-network	Exact-match	32 bits	6144	6144
ip-mcast	Longest-match	64 bits	3072	6144
l2-switching	Exact-match	64 bits	1024	2048
udp-flooding	Exact-match	64 bits	256	512
access-list	First-match	128 bits	512	8192

The enhanced Gigabit Ethernet interface module is available with 32 KB, 64 KB, or 256 KB TCAM space. You can configure the various protocol regions in TCAM based on your requirements and on the size of TCAM on your Gigabit Ethernet interface module.

Figure 11-23 Dynamic CAM and TCAM Relationship

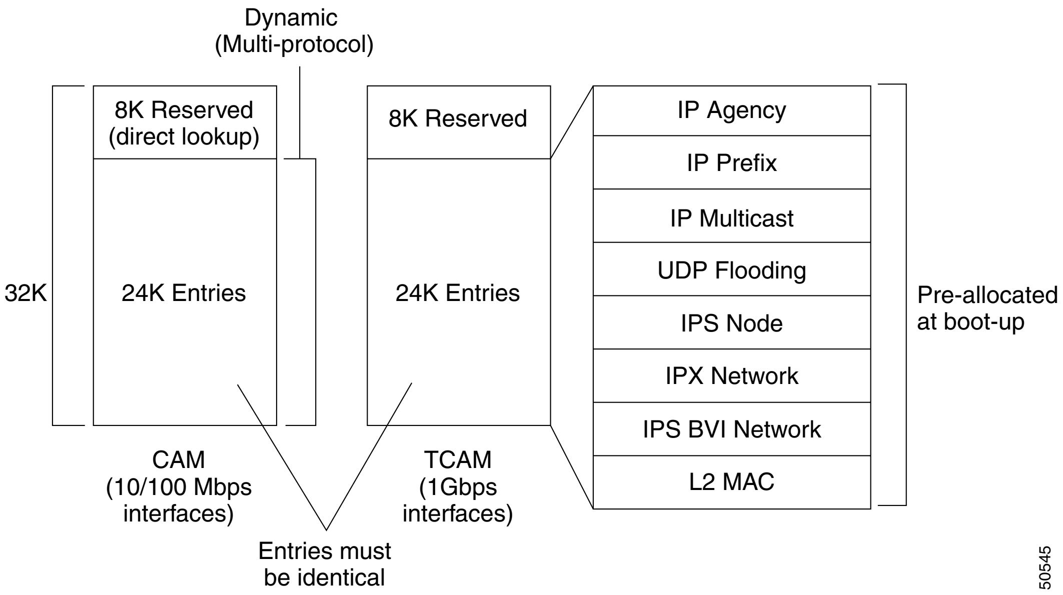

Note The enhanced Gigabit Ethernet interface module is available with 32 KB, 64 KB, or 256 KB TCAM space. The maximum SDM size is equal to the lowest TCAM size available among the interface modules present at the time of booting up the switch router. For example, if you have two interface modules with 64 KB and 256 KB TCAM sizes, then the maximum SDM size is 64 KB based on the lowest TCAM size available at bootup.

Troubleshooting SDM Problem Commands

To display and troubleshoot the SDM CAM configuration, use the following commands:

Command	Purpose
show sdm size	Displays the size of TCAM and the size of each protocol region.
show sdm internal {**all-region** | **ip-adjacency** | **ip-multicast** | **ip-prefix** | **ipx-network** | **ipx-node**}	Displays the SDM management information for each protocol region in TCAM
sdm size *region-name* {*num-entries* | **k-entries** *num-k-entries*}	Sets the name of the protocol region for which you want to configure the size.
sdm access-list *num-entries*	Sets the name of the protocol region for which you want to configure the size. You can enter the size as an absolute number of entries.

Configuring the Switching Database Manager

This section describes how to configure the SDM to change the size of the protocol-specific TCAM regions in the switch router.

To modify the default TCAM region sizes, use the following procedure:

Step 1 Based on your network protocol mix and the number of prefixes and stations in the network, determine the required size of the various protocol-specific TCAM regions.

Step 2 Modify the size of each region using the **sdm size** global configuration command.

Step 3 If desired, modify the SDM autolearn function using the [**no**] **sdm autolearn** global configuration command.

Step 4 Before reloading the system, verify that the desired sizing is reflected in the configuration (use the **show running-config** command).

Step 5 Reload the switch router to implement the new partitioning configuration.

The following process shows how to enlarge the size of the ip-prefix TCAM partition from 65,536 32-bit entries to 131,072 32-bit entries.

Note You must reload the system in order for the changes to take effect.

Follow these steps to check and configure the SDM size:

Step 1 Use the **show sdm size** command to see the configuration of the SDM CAM size.

```
Switch# show sdm size
Switching Database Region Sizes :
    IPX Direct          : 224     32-bit entries
    IPX Node            : 4096    64-bit entries
    IP Adjacency        : 4096    32-bit entries
➔   IP Prefix           : 65536   32-bit entries
    IP VRF Prefix       : 512     64-bit entries
    IP Multicast        : 32768   64-bit entries
    UDP Flooding        : 256     64-bit entries
    MAC Addr            : 1024    64-bit entries
    LFIB                : 1024    32-bit entries
    Label               : 8192    32-bit entries
    Access List         : 512    128-bit entries
Switch#
```

Step 2 Use the **sdm size ip-prefix k-entries** command to change the ip-prefix from 65,536 32-bit bytes to 131,072 32-bit bytes. Using the **k-entries** parameter with the **128** (Kbytes) * 1024 multiples, equals 131,072 32-bit entries.

```
Switch(config)# sdm size ip-prefix k-entries 128
```

Step 3 Use the **show running-config** command to confirm the new configuration.

```
Switch# show running-config
Building configuration...

Current configuration:
!
version 12.0
no service pad
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
no service password-encryption
!
hostname Switch
!
!
clock calendar-valid
sdm size ip-adjacency 4096
➔ sdm size ip-prefix 131072
sdm size ipx-network 16384
no sdm autolearn
ip subnet-zero
!
(Information Deleted)
!
```

Step 4 Use the **copy running-config startup-config** command to write the new configuration to the NVRAM.

```
Switch# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...

EHSA:Syncing monvars to secondary, : BOOT=
EHSA:Syncing monvars to secondary, : CONFIG_FILE=
EHSA:Syncing monvars to secondary, : BOOTLDR=[OK]
Switch#
```

Step 5 Use the **reload** command to restart the Layer 3 enabled ATM switch router and reallocate the memory partitions.

```
Switch# reload
     Proceed with reload? [confirm]

     Oct  9 18:54:55.294: %SYS-5-RELOAD: Reload requested

     ROMMON: Cold Reset frame @0x00000000
     ROMMON: Reading reset reason register
     ROMMON: Valid NVRAM config

     System Bootstrap, Version 12.0(7)W5(15) RELEASE SOFTWARE
     Copyright (c) 1998 by cisco Systems, Inc.
```

Step 6 Use the **show sdm size** command to see the new configuration of the SDM CAM size.

```
Switch# show sdm size
Switching Database Region Sizes :
    IPX Direct          : 224     32-bit entries
    IPX Node            : 4096    64-bit entries
    IP Adjacency        : 4096    32-bit entries
➔   IP Prefix           : 131072  32-bit entries
    IP VRF Prefix       : 512     64-bit entries
    IP Multicast        : 32768   64-bit entries
    UDP Flooding        : 256     64-bit entries
    MAC Addr            : 1024    64-bit entries
    LFIB                : 1024    32-bit entries
    Label               : 8192    32-bit entries
    Access List         : 512    128-bit entries
Switch#
```

If you determine that the SDM is configured incorrectly, refer to the "Configuring Switching Database Manager" chapter in the *Layer 3 Switching Software Feature and Configuration Guide*.

Troubleshooting Common Errors When Changing SDM Size

This section describes the following two common errors that might occur when you are trying to change the SDM size of the protocol-specific TCAM regions in the Layer 3 switch:

- The switch router generates a "Total protocol partitions exceed TCAM size!!" error when configuring the SDM.
- The switch router generates a "%LSS-1-SDM: Region reached limit. Cannot accept more entries" syslog message at startup, or during normal operation of the switch router.

Note You must reload the system in order for the changes to take effect.

Troubleshooting the “Total protocol partitions exceed TCAM size!!” Error

The switch router generates a “Total protocol partitions exceed TCAM size!!” error while you are configuring the SDM partition sizes for the following reasons:

- The command entered cannot be processed because the command you entered would cause the total size of the TCAM protocol partitions to exceed 32K.
- The command entered cannot be processed because the command you entered would cause the size of that specific TCAM protocol partition to exceed the maximum allowed size for that partition.

To solve the problem, specify a protocol partition size that does not exceed the total TCAM size, or specify the maximum size of the specified protocol partition.

In this example, the system generates an error when you attempt to specify more than 16,000 entries for the l2-switching region. The workaround is to ensure the specified size is less than or equal to the maximum region size, and that the sum of all of the protocol regions does not exceed 32K entries.

Follow these steps to eliminate the “Total protocol partitions exceed TCAM size!!” error while you are configuring the SDM partition sizes:

Step 1 While in EXEC configuration mode you use the **sdm size** command to modify the SDM partition sizes and receive a “Total protocol partitions exceed TCAM size!!” error.

```
Switch# configure terminal
     Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# sdm size l2-switching 17000
      Total protocol partitions exceed TCAM size!!
Switch(config)# sdm size l2-switching 16001
      Total protocol partitions exceed TCAM size!!
Switch(config)#
```

Step 2 Use the **show sdm size** command to display the size of the existing TCAM configuration.

```
Switch# show sdm size
Switching Database Region Sizes :
    IPX Direct          : 224     32-bit entries
    IPX Node            : 1024    64-bit entries
    IP Adjacency        : 2048    32-bit entries
    IP Prefix           : 8224    32-bit entries
    IPX Network         : 2016    32-bit entries
    IP VRF Prefix       : 512     64-bit entries
    IP Multicast        : 1024    64-bit entries
    UDP Flooding        : 256     64-bit entries
    MAC Addr            : 2048    64-bit entries
    LFIB                : 4096    32-bit entries
    Label               : 8192    32-bit entries
    Access List         : 0       128-bit entries
Switch#
```

Step 3 Use the **show sdm internal all-regions** command to display the size of the existing TCAM configuration for all regions.

```
Switch# show sdm internal all-regions
Address Map      :
Status             : Ready
TCAM Minimum Size  : 32768 entries
TCAM Required Size : 22272 entries
SRAM Sz            : 49152 entries
TCAM Start         : 32
Xinfo Start        : 45056
Xinfo Size         : 7424
Xinfo Used         : 3
Xinfo Free         : 7421
Name    : IPX Direct
Size    : 224
MinSize : 224
MaxSize : 224
FreeKey : 0x0
Start   : 0x20
End     : 0xFF
Entry   : 32-bit
Lookup  : Exact-Match
Events  :
Insert  : Success 2 Failure 0
Delete  : Success 0 Failure 0
Modify  : Success 0 Failure 0
IPCs    :
Insert  : Success 2 Failure 0
Delete  : Success 0 Failure 0
Modify  : Success 0 Failure 0
Move    : Success 0 Failure 0
Mask RW : Success 0 Failure 0

Name    : IPX Node
Size    : 1024
MinSize : 32
MaxSize : 16128
FreeKey : 0xF0000000
Start   : 0x100
End     : 0x8FE
Entry   : 64-bit
Lookup  : Exact-Match
Events  :
Insert  : Success 0 Failure 0
Delete  : Success 0 Failure 0
Modify  : Success 0 Failure 0
IPCs    :
Insert  : Success 0 Failure 0
Delete  : Success 0 Failure 0
Modify  : Success 0 Failure 0
Move    : Success 0 Failure 0
Mask RW : Success 0 Failure 0

Name    : IP Adjacency
Size    : 2048
MinSize : 32
MaxSize : 32768
FreeKey : 0xEEEEEEEE
Start   : 0x900
End     : 0x10FF
Entry   : 32-bit
Lookup  : Exact-Match
Events  :
Insert  : Success 36 Failure 0
Delete  : Success 6 Failure 0
Modify  : Success 2 Failure 0
IPCs    :
Insert  : Success 36 Failure 0
Delete  : Success 6 Failure 0
Modify  : Success 2 Failure 0
Move    : Success 0 Failure 0
Mask RW : Success 0 Failure 0

Name    : IP Prefix
Size    : 8224
MinSize : 32
```

```
MaxSize : 32768
FreeKey : 0xEEEEEEEEEEEEEEEE
Start   : 0x1100
End     : 0x30FF
Entry   : 32-bit
Lookup  : Longest-Match
Buckets : 33
Events  :
Insert  : Success 52 Failure 0
Delete  : Success 63 Failure 0
Modify  : Success 8 Failure 0
IPCs    :
Insert  : Success 52 Failure 0
Delete  : Success 63 Failure 0
Modify  : Success 8 Failure 0
Move    : Success 20 Failure 0
Mask RW : Success 8 Failure 0

Name    : IPX Network
Size    : 2016
MinSize : 32
MaxSize : 32768
FreeKey : 0x0
Start   : 0x3100
End     : 0x38DF
Entry   : 32-bit
Lookup  : Longest-Match
Buckets : 1
Events  :
Insert  : Success 0 Failure 0
Delete  : Success 0 Failure 0
Modify  : Success 0 Failure 0
IPCs    :
Insert  : Success 0 Failure 0
Delete  : Success 0 Failure 0
Modify  : Success 0 Failure 0
Move    : Success 0 Failure 0
Mask RW : Success 0 Failure 0

Name    : IP VRF Prefix
Size    : 512
MinSize : 32
MaxSize : 32768
FreeKey : 0xEEEEEEEE
Start   : 0x38E0
End     : 0x3C9E
Entry   : 64-bit
Lookup  : Longest-Match
Buckets : 33
Events  :
Insert  : Success 0 Failure 0
Delete  : Success 0 Failure 0
Modify  : Success 0 Failure 0
IPCs    :
Insert  : Success 0 Failure 0
Delete  : Success 0 Failure 0
Modify  : Success 0 Failure 0
Move    : Success 0 Failure 0
Mask RW : Success 0 Failure 0

Name    : IP Multicast
Size    : 1024
MinSize : 32
MaxSize : 16384
```

```
FreeKey : 0xF0000000F0000000
Start   : 0x3CA0
End     : 0x449E
Entry   : 64-bit
Lookup  : Longest-Match
Buckets : 34
Events  :
Insert  : Success 3 Failure 0
Delete  : Success 0 Failure 0
Modify  : Success 6 Failure 0
IPCs    :
Insert  : Success 3 Failure 0
Delete  : Success 0 Failure 0
Modify  : Success 6 Failure 0
Move    : Success 0 Failure 0
Mask RW : Success 2 Failure 0

Name    : UDP Flooding
Size    : 256
MinSize : 256
MaxSize : 256
FreeKey : 0xF0000000
Start   : 0x44A0
End     : 0x469E
Entry   : 64-bit
Lookup  : Exact-Match
Events  :
Insert  : Success 0 Failure 0
Delete  : Success 0 Failure 0
Modify  : Success 0 Failure 0
IPCs    :
Insert  : Success 0 Failure 0
Delete  : Success 0 Failure 0
Modify  : Success 0 Failure 0
Move    : Success 0 Failure 0
Mask RW : Success 0 Failure 0

Name    : MAC Addr
Size    : 2048
MinSize : 128
MaxSize : 16384
FreeKey : 0x0
Start   : 0x4700
End     : 0x56FE
Entry   : 64-bit
Lookup  : Reserved

Name    : Access List
Size    : 0
MinSize : 512
MaxSize : 16384
Entry   : 128-bit

Switch#
```

Step 4 Confirm that the value entered in Step 1 does not exceed the total existing TCAM size, and try again.

```
Switch(config)# sdm size l2-switching 16000
Switch(config)# ^Z
Switch#
```

Troubleshooting the "%LSS-1-SDM: Region reached limit. Cannot accept more entries" Syslog Message

The switch router generates the "%LSS-1-SDM: Region reached limit. Cannot accept more entries" syslog message at startup, or during normal system operation.

The following example shows that the system was unable to install one or more entries in the TCAM for the ip-adjacency and ip-prefix switching database regions. The following syslog messages indicate that the TCAM regions should be reconfigured to allow more entries for IP prefix and adjacency entries.

```
Oct 10 15:54:57.179: %LSS-1-SDM: IP Prefix     Region reached limit. Cannot accept more entries
Oct 10 16:12:45.275: %LSS-1-SDM: IP Adjacency  Region reached limit. Cannot accept more entries
```

The system generates the syslog message for a specific protocol region when the system fails to install one or more entries in the TCAM because the specified region is full.

To fix this problem you must increase the size of the specified protocol region, using the **sdm size** command, and reload the system.

Use the process described in the "Configuring the Switching Database Manager" section on page 11-59 to modify the ip-adjacency and ip-prefix switching database regions in the TCAM.

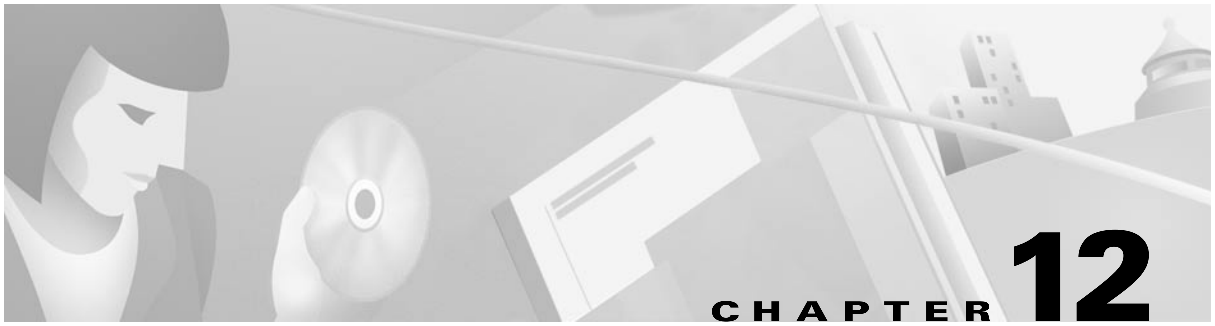

CHAPTER 12

Troubleshooting Layer 2 Interfaces

This chapter provides troubleshooting information about connectivity and performance problems in the Layer 2 network connections of an ATM switch router and includes the following sections:

- Layer 2 Switching and Bridging Overview, page 12-1
- Troubleshooting Layer 2 Switching, page 12-3
- Troubleshooting Integrated Routing and Bridging, page 12-13
- Troubleshooting Trunk Port Problems, page 12-16
- Troubleshooting Fast EtherChannel Problems, page 12-16

Note For detailed cabling and hardware information for each port adapter, refer to the *Catalyst 8540 CSR Route Processor and Interface Module Installation Guide.*

Layer 2 Switching and Bridging Overview

This section provides some overview information about Layer 2 switching and bridging.

Layer 2 Switching

The difference between Layer 2 and Layer 3 switching is the type of information inside the frame that is used to determine the correct output interface. With Layer 2 switching, frames are switched based on MAC address information. With Layer 3 switching, frames are switched based on network-layer information.

Layer 2 switching does not look inside a packet for network-layer information as does Layer 3 switching. Layer 2 switching is performed by looking at a destination MAC address within a frame. It looks at the frame destination address and sends it to the appropriate interface if the switch knows the destination address location. Layer 2 switching builds and maintains a switching table that keeps track of which MAC addresses belong to each port or interface.

If the Layer 2 switch does not know where to send the frame, it broadcasts the frame out to all its ports on the network, to learn the correct destination. When the frame reply is returned, the switch learns the location of the new address, and adds the information to the switching table.

The switch router performs Layer 2 switching using the following functions:

- Places path destination address, source address, and VLAN information are stored in CAM tables.
- Sends MAC Address update information to the route processor via Cisco IOS Interprocess Communications (IPC).
- Considers Ethernet processor interface learning as “normal”
- Updates the route processor from the Ethernet processor interface.
- Broadcasts use and Broute VC to communicate with ports on other modules in same bridge group.
- Sets spanning tree sets state of IOS-CPU.
- Transmitted by IOS-CPU, and received by a Layer 2 ASIC, bridge PDUs are tunneled in the IPC messages.
- Receives Bridge PDUs on trunk ports, and preserves port and tag information.
- During transmission, bridge PDUs are tagged by the ASICs on trunk ports

Bridging

Cisco IOS software supports transparent bridging for Ethernet. In addition, Cisco supports all the mandatory Management Information Base (MIB) variables specified for transparent bridging in RFC 1286.

Cisco IOS software bridging functionality combines the advantages of a spanning tree bridge and a full multiprotocol router. This combination provides the speed and protocol transparency of an adaptive spanning tree bridge, along with the functionality, reliability, and security of a router.

The switch router can be configured to serve as both an IP and IPX router and a MAC-level bridge, bridging any traffic that cannot otherwise be routed. For example, a router routing IP traffic can also bridge the Digital local-area transport (LAT) protocol or NetBIOS traffic.

To configure bridging, you must perform the following tasks:

- In global configuration mode:
  - Select Spanning Tree Protocol.
  - Assign a priority to the bridge (optional).
- In interface configuration mode:
  - Determine which interfaces belong to the same bridge group.

    These interfaces will be part of the same spanning tree. This allows the switch router to bridge all nonrouted traffic among the network interfaces comprising the bridge group. Interfaces not participating in a bridge group cannot forward bridged traffic.

    If the packet's destination address is known in the bridge table, it is forwarded on a single interface in the bridge group. If the packet's destination is unknown in the bridge table, it is flooded on all forwarding interfaces in the bridge group. The bridge places source addresses in the bridge table as it learns them during the process of bridging.

    A separate spanning tree process runs for each configured bridge group. Each bridge group participates in a separate spanning tree. A bridge group establishes a spanning tree based on the BPDUs it receives on only its member interfaces.

- Assign a port path cost on the outgoing interface (optional). When configuring POS uplink interfaces and ATM uplink interfaces in conjunction with parallel Gigabit Ethernet links, we recommend that you explicitly assign the port path cost to these interfaces because the default value might not yield the optimal spanning tree.

Troubleshooting Layer 2 Switching

This section describes Layer 2 Switching troubleshooting and includes the following sections:

- Layer 2 Switching Broadcasts, page 12-3
- Troubleshooting Layer 2 Switching, page 12-5
- Other Layer 2 Switching Problems, page 12-10
- Layer 2 Bridging Troubleshooting Commands, page 12-12

Layer 2 Switching Broadcasts

Figure 12-1 shows the broadcast process all interfaces use to update all the other CAM tables in the bridge group using the following processes:

- Each interface is the root of a P2MP VC.
- Leaves on all interfaces in the same bridge group are required to flood unknown unicast traffic received on that port.
- Broadcasts VC setup when the interface is added to a bridge group.
- Each interface is a leaf of point-to-multipoint (P2MP) VCs, with roots on all other interfaces of the bridge group.

Figure 12-1 Layer 2 Switching Broadcasts Used to distribute CAM Updates

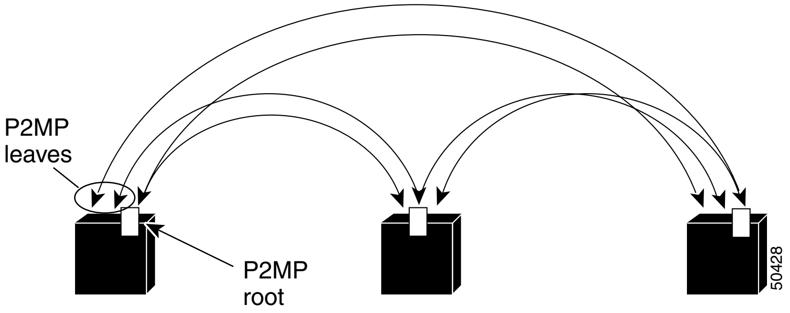

Following is the process the switch uses to learn the MAC addresses as show in Figure 12-2:

1. Interface module microcode learns a new MAC Address
2. Interface module microcode sends an IPC to the route processor to update the bridging table
3. Aged entries are removed from the CAM
4. Interface module microcode sends an IPC to the route processor to remove the aged entry from bridging tables

Layer 2 unknown unicast switching is accomplished by the following:

- Having all unknown unicast packets sent over the broadcast VC
- Depending on the IOS version installed, Layer 2 broadcasts are sent differently:
  - With Release 12.0(4a)W5(11a) and earlier, Layer 2 broadcasts are sent to the route processor
  - With Release 12.0(5)W5(13) and later, Layer 2 broadcasts are sent using the P2MP (BCAST) VC only

Note The route processor is not reachable using any of the bridged interfaces. Administrative tasks to the route processor are preformed using the Console or route processor Ethernet interface

Figure 12-2 Layer 2 Learning Process

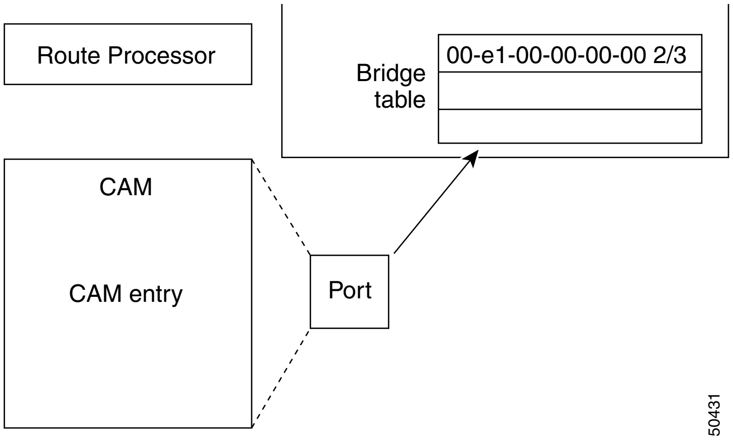

Bridging over Fast EtherChannel

MAC address learning in the switch router occurs differently depending on the IOS version installed.

For Cisco IOS Release 12.0(4a)W5(11a) and earlier, MAC address learning occurs as follows:

- The first member of Fast EtherChannel learns the source address via incoming traffic (Egress Learn).
- Other members learn the source address (Egress Learn) via IPCs from the Route Processor. These are static entries and will not age out in the CAM.
- Static entries are deleted explicitly via IPCs from route processor.

For Cisco IOS Release 12.0(5)W5(13) and later, MAC address learning occurs as follows:

- The first member of a Fast EtherChannel learns the source address via Traffic (Egress Learn).
- Other members learn the source address (Egress Learn) from the first member, via IPCs over the P2MP VC.
- Layer 2 entries age out in the CAM.

Troubleshooting Layer 2 Switching

To troubleshoot a Layer 2 switching problems, use the following commands:

Command	Purpose
show bridge group	Displays bridge group configuration and status information.
show bridge	Displays the status of all the bridge groups on the switch router.
show spanning-tree *number*	Displays the spanning tree topology for a bridge group.
show interfaces bvi *number*	Displays BVI interface configuration, status, and statistics. Use this command when the BVI is part of a bridge group.
show interfaces {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port* (on the ingress interface)	Displays interface configuration, status, and statistics on the ingress interface.
show interfaces {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port* (on the egress interface)	Displays interface configuration, status, and statistics on the egress interface.
show switch bridge-table *entry*	Displays bridge table entry summary.
show epc freecam interface {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port*	Displays information about free space in the content addressable memory.
show epc if-entry interface {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port* **all**	Displays all interface entry information for the specific interface.
show epc patricia interface {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port* **mac detail** (on the ingress interface)	Displays the MAC patricia tree for the ingress interface.
show epc patricia interface {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port* **mac detail** (on the egress interface)	Displays the MAC patricia tree for the egress interface.

As stated before, once the first interface learns a new destination address, source address, and VLAN, that information must be broadcasts to all other interfaces to allow them to update their CAM. That information is broadcast using the Broute VC to communicate with ports on other modules in the same bridge group. Figure 12-3 shows interface Fast Ethernet 0/0/0 using broadcast VC 0 to broadcast an update (to the CAM) to interfaces Fast Ethernet 0/0/1 and 0/0/2.

Figure 12-3 Broadcasting CAM Updates

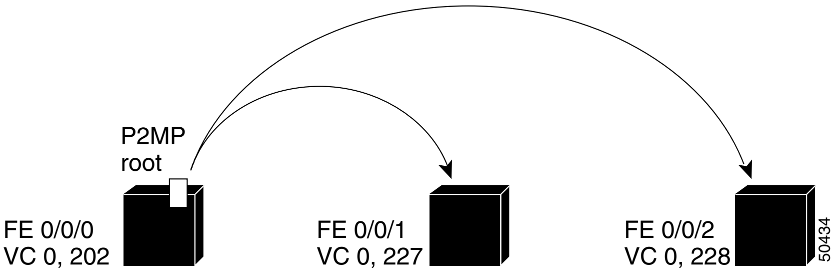

To confirm that Broute VC 0 is up, use the **show atm vc cast p2mp interface fastethernet** *slot*/*subslot*/*port* command as shown in the following example.

```
Switch# show atm vc cast p2mp interface fastethernet 0/0/0
Interface          VPI   VCI   Type    X-Interface       X-VPI X-VCI  Encap Status
FastEthernet0/0/0   0    202   PVC     FastEthernet0/0/1  0     227          UP
                                       FastEthernet0/0/2  0     228          UP
```

Using the cross-connect VCI information from the previous output, confirm that Broute VC 0 is configured on the other interfaces in the bridge group use the **show atm vc traffic interface fastethernet** *slot*/*subslot*/*port 0 X-VCI* command as shown in the following examples:

```
Switch# show atm vc traffic interface fastethernet 0/0/0 0 202
Interface    VPI     VCI      Type      rx-cell-cnts     tx-cell-cnts
FastEthernet 0       202       PVC                 0                0
Switch# show atm vc traffic interface fastethernet 0/0/1 0 227
Interface    VPI     VCI      Type      rx-cell-cnts     tx-cell-cnts
FastEthernet 0       227       PVC                 0                0
Switch# show atm vc traffic interface fastethernet 0/0/2 0 228
Interface    VPI     VCI      Type      rx-cell-cnts     tx-cell-cnts
FastEthernet 0       228       PVC                 0                0
```

Figure 12-4 is an example network of two Layer 2 switches, and is used in the following troubleshooting steps.

Figure 12-4 Layer 2 Troubleshooting Example Network

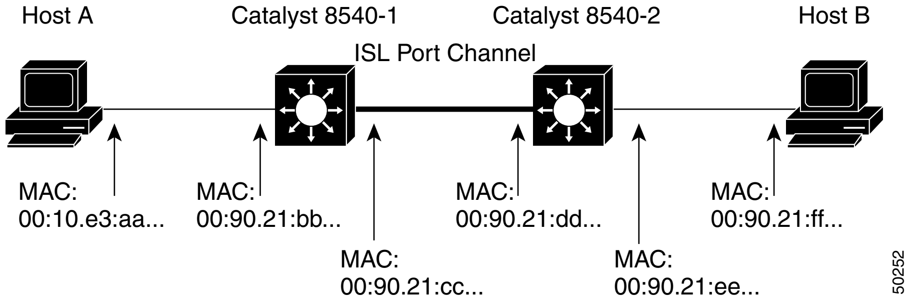

The following are the processes used to troubleshoot Layer 2 switching connections:

- Check the route processor "View" of the spanning tree.
- Verify information on interface modules.
- Verify VC status between ports.

To troubleshoot Layer 2 switching, perform the following steps:

Step 1 Use the **show spanning-tree** command to display bridge group information as shown in Figure 12-4.

```
  Switch# show spanning-tree 1

   Bridge group 1 is executing the IEEE compatible Spanning Tree protocol
  .
  (Information Deleted)
  .
➔ Port 43 (FastEthernet12/0/1) of Bridge group 1 is forwarding
  .
  (Information Deleted)
  .
➔ Port 58 (GigabitEthernet11/0/1.1 ISL) of Bridge group 1 is forwarding
```

Step 2 Find the entries in the **show spanning-tree** command output for the interfaces in question and check the Port fields to confirm the ports at Fast Ethernet12/0/1 and Gigabit Ethernet 11/0/1.1 ISL are forwarding.

Step 3 Use the **show bridge** command to display bridge group information.

```
  Switch# show bridge 1
  .
  (Information Deleted)
  .
  Bridge Group 1:
      Address       Action     Interface
➔ 0010.e3aa.aaaa   forward    Fa12/0/1
➔ 0090.21bb.bbbb   forward    Gi11/0/1.1
```

Find the entries for those interfaces in the **show bridge** command output. Note the Address fields to see that the MAC addresses being forwarded to the Fast Ethernet12/0/1 and Gigabit Ethernet 11/0/1.1 ISL connections are in the bridge table. These addresses are used in the following step.

Step 4 Use the **show epc mac interface FastEthernet** interface command with the MAC address parameter. Add the MAC address being forwarded to interface FastEthernet 12/0/1.

```
Switch# show epc mac interface FastEthernet 12/0/1 0010.e3aa.aaaa
MACaddr:0010.e3aa.aaaa    IF Number:43   MAC Local
```

Note the IF Number field; in this example, the number is “43.” You will need this number in the following command.

Step 5 Use the **show epc mac interface GigabitEthernet** interface command with the MAC address parameter. Add the MAC address being forwarded to interface GigabitEthernet 11/0/1.1 ISL.

```
Switch# show epc mac interface FastEthernet 12/0/1 0090.21bb.bbbb
MACaddr:0090.21bb.bbbb     IF Number:58
```

Note the IF Number field; in this example, the number is “58.” You will need this number in the following command.

Step 6 Use the **show epc patricia interface FastEthernet** interface command with the **mac detail** parameters.

```
Switch# show epc patricia interface FastEthernet 12/0/1 mac detail
.
(Information Deleted)
.

7# MAC addr:0090.21bb.bbbb   IF Number:58 Entry:Remote
      Learned 10450 times used
   CAM location: 101D
8# MAC addr:0010.e3aa.aaaa   IF Number:43 Entry:Local
      Learned 0 times used
   CAM location: 0FF8
 Total number of MAC entries: 8
```

Step 7 Verify that the information from this command is consistent with the command outputs in Step 5.

If there are inconsistencies or non-zero invalid entries in the tables, you can use the **clear bridge** command to rebuild the tables.

Step 8 Check the Entry field.

Interface Number 58 has learned of this entry via the flooding from IF Number 42 and it is marked as Remote. Interface number 43 has learned these entries as local entries, which is typical bridge behavior.

Note Entries marked as MyMac are for internal use. These are static entries and are for spanning tree BPDUs and CDP. The MAC address marked as HSRP is the actual BIA Mac address of the port. This entry is only present if there is a BVI defined for that bridge group.

Verify that the information from the **show epc patricia** command output is consistent with the command outputs in Step 4.

If there are inconsistencies or non-zero invalid entries in the tables, you can use the **clear bridge** command to rebuild the tables.

Caution Use the **clear bridge** command carefully. It causes a temporary increase in switch router activity which can lead to traffic disruptions.

Step 9 Use the **show epc patricia interface GigabitEthernet** interface command with the **vlan** number and **details** parameters.

```
Switch# show epc patricia interface GigabitEthernet 11/0/1 vlan 1 detail
15# MAC addr:0090.21bb.bbbb   IF Number:58 Entry:Local
      Learned 0 times used
   CAM location: 1034
16# MAC addr:0010.e3aa.aaaa   IF Number:43 Entry:Remote
      Learned 6029 times used
   CAM location: 101B
```

Step 10 Check the Entry field. The Entry field descriptions are the same as those in Step 8.

Verify that the information from the **show epc patricia** command output is consistent with the command outputs in Step 4.

If there are inconsistencies or non-zero invalid entries in the tables, you can use the **clear bridge** command to rebuild the tables.

Step 11 Use the **show epc if-entry interface FastEthernet** interface **entry GigabitEthernet** *interface.subinterface* command to display the CAM table entry.

```
Switch# show epc if-entry interface FastEthernet 12/0/1 entry GigabitEthernet 11/0/1.1
IF Entry for GigabitEthernet11/0/1.1 on FastEthernet12/0/1
    Mac(hex) - 00:90:21:dd:dd:dd
    isMyInteface : False isSubInterface : True
➔   Status Up Broute VC - 662 Bcast VC - 747
    Netmask: 32
    FEC disabled
    ISL, Vlan 1
➔   State : Forwarding
    Bridge-Group enabled
➔   IP routing off bridging on
➔   IPX routing off bridging on
➔   Appletalk routing off
    In Encapsulation:
    ICMP Redirect disabled Unreachable disabled
    IP Multicast disabled: ttl-threshold: 0
```

Step 12 Confirm the following:

- Brouter VC status is up.
- State is forwarding.
- All routing protocols, IP, IPX, and Appletalk are off.

Step 13 Use the same command again but with the interface entries reversed.

```
Switch# show epc if-entry interface GigabitEthernet 11/0/1.1 entry FastEthernet 12/0/1
IF Entry for FastEthernet12/0/1 on GigabitEthernet11/0/1
    Mac(hex) - 00:90:21:cc:cc:cc
    isMyInteface : False isSubInterface : False
➔   Status Up Broute VC - 622 Bcast VC - 747
    Netmask: 32
    FEC disabled
    Trunking Disabled
➔   State : Forwarding
    Bridge-Group enabled
➔   IP routing off bridging on
➔   IPX routing off bridging on
➔   Appletalk routing off
    In Encapsulation:
    ICMP Redirect disabled Unreachable disabled
    IP Multicast disabled: ttl-threshold: 0
```

Step 14 Confirm the following:

- Brouter VC status is up.
- State is forwarding.
- All routing protocols, IP, IPX, and Appletalk are off.

Troubleshooting Layer 2 connections differs from troubleshooting Layer 3 connections in the following ways:

- MAC entries are learned and not downloaded from route processor.
- Not all Ethernet processor interface CAM entries in the same bridge group contain all entries.
- You must verify the status of the Broute and the Bcast VC on both the ingress and egress ports.

Refer to the *Layer 3 Software Feature and Configuration Guide* if any changes are necessary to the configuration of the interface.

Other Layer 2 Switching Problems

This sections describes the following, less common, Layer 2 switching connection problems:

- Layer 2 Connection is Flooding, page 12-10
- Packets are Switched but are Not Appearing on the Wire, page 12-10
- Layer 2 CAM Display, page 12-11
- Check for Spanning Tree Loop, page 12-11

Layer 2 Connection is Flooding

If you determine the Layer 2 connection is flooding instead of switching, check the following:

- If the destination address of the traffic is known:
  - Check the Layer 2 CAM of the ingress interface for the specific destination address. If the entry exists, then note the destination interface number.
  - Use the **show epc ifmapping** command and note the name-string of the interface.
  - Display the interface table in the ingress interface for the destination interface (obtained from the Layer 2 CAM).
  - Get the BROUTE VC from the **show epc if-entry interface** command described earlier.
  - Use the **show atm vc traffic interface** *interface-namestring* **0** *BROUTE-VC* command.
  - Confirm that the “rx” and “tx” counters are increasing when traffic is switched (by entering the previous command, waiting, and entering the same command again).
- If the destination address is unknown:
  - Use the **show epc patricia interface** command to display the interface table in the ingress interface for all the interfaces in the bridge-group.
  - Get the BROUTE VCs from the **show epc if-entry interface** command described earlier, for all the interfaces in the bridge group.
  - Use the **show atm vc traffic interface** *interface-namestring* **0** *BROUTE-VC* command.
  - Confirm that the “rx” and “tx” counters are increasing when traffic is switched (by entering the previous command, waiting, and entering the same command again).

Packets are Switched but are Not Appearing on the Wire

If you are sure the packets are being switched but the connection does not appear on the wire, try the following:

- Use a Sniffer to see whether they are actually being sent on the wire.
- If the packets are not seen by Sniffer, then identify the output interface.

- Use the **show controller** {**FastEthernet | GigabitEthernet**} *card/subcard/port* command and check the following:
  - The MTx and SRx registers for doubts in transmission
  - The MRx and STx registers for doubts in reception

Layer 2 CAM Display

You might need to check for a specific MAC address on an interface. If so, use the **show epc patricia interface** {**FastEthernet | GigabitEthernet**} *card/subcard/port* **mac** command. The following is an example with a description of some useful MAC addresses:

```
  Switch# show epc patricia interface fastethernet 0/0/0 mac
  1# MAC addr:0000.0000.0000          VC:0     Entry:
  2# MAC addr:0900.2b01.0001  MyMAC   VC:4     Entry:
→ 3# MAC addr:0180.c200.0000  MyMAC   VC:4     Entry:
→ 4# MAC addr:0100.0ccc.cccc  MyMAC   VC:4     Entry:
  5# MAC addr:0010.073d.8207 HsrpMAC  VC:4     Entry:
→ 6# MAC addr:0008.e0bc.4190 MyMAC    VC:4     Entry:
   Total number of MAC entries: 6
```

In this example check the following:

- MAC address 3# and 4# are spanning tree BPDU addresses
- MAC address 6# is the interface fastEthernet 0/0/0 Mac-address (IRB only)

Check for Spanning Tree Loop

The spanning-tree algorithm in the IOS software is probably not the source of the spanning tree loop. The spanning-tree loop probably exists because of a problem in the end-to-end connectivity. Try one of the following tests to confirm you do not have a spanning-tree loop:

- Use the **show epc patricia interface** {**FastEthernet | GigabitEthernet**} *card/subcard/port* {**mac** | **vlan** *#*} command to confirm the spanning tree multicast address is present in the Layer 2 CAM. See the section "Layer 2 CAM Display" section on page 12-11.

  If the spanning tree MAC address is *not* in the output, then spanning tree is the problem.
- Use **show spanning-tree bridge-group** *bridge-group-number* command and check for the string "BPDU: sent 0, received 0" in the output. If both switch routers connected with back-to-back interfaces "received 0," then there is a physical layer connectivity problem between the switch routers.

Refer to the *Layer 3 Software Feature and Configuration Guide* if any changes are necessary to the configuration of the interface.

Layer 2 Bridging Troubleshooting Commands

To troubleshoot a Layer 2 bridging problem, use the following commands:

Command	Purpose
show bridge group	Displays bridge group configuration and status information.
show bridge	Displays the status of all the bridge groups on the switch router.
show spanning-tree *number*	Displays the spanning tree topology for a bridge group.
show interfaces bvi *number*	Displays BVI interface configuration, status, and statistics. Use this command when the BVI is part of a bridge group.
show interfaces {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* (on the ingress interface)	Displays interface configuration, status, and statistics on the ingress interface.
show interfaces {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* (on the egress interface)	Displays interface configuration, status, and statistics on the egress interface.
show switch bridge-table *entry*	Displays bridge table entry summary.
show epc freecam interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port*	Displays information about free space in the content addressable memory.
show epc if-entry interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **all**	Displays all interface entry information for the specific interface.
show epc patricia interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **mac detail** (on the ingress interface)	Displays the MAC patricia tree for the ingress interface.
show epc patricia interface {**fastethernet** | **gigabitethernet**} *slot*/*subslot*/*port* **mac detail** (on the egress interface)	Displays the MAC patricia tree for the egress interface.

If a BVI is involved, use the following commands:

Command	Purpose
show bridge group	Displays bridge group configuration and status information.
show interfaces irb	Displays integrated routing and bridging configuration and status for all interfaces.
show smf	Displays software MAC address information.
show interfaces bvi *number*	Displays BVI interface information.
show bridge *number* **group**	Displays the status of the member ports in the specified bridge group.

Command	Purpose
show bridge *number*	Displays the status of the bridge group.
show epc patricia interface {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port* **mac detail**	Displays the MAC patricia tree for the specified interface physical interface in the bridge group with a problem.

Troubleshooting Integrated Routing and Bridging

This section describes troubleshooting Integrated Routing and Bridging (IRB) on the Catalyst 8510 MSR and Catalyst 8540 MSR and Catalyst 8510 CSR and Catalyst 8540 CSR switches.

Your network may require you to bridge local traffic within several segments while having hosts on the bridged segments reach the hosts or routers on routed networks. For example, if you are migrating bridged topologies into routed topologies, you may want to start by connecting some of the bridged segments to the routed networks.

IP Switching with IRB Overview

Using the IRB feature, you can route a given protocol between routed interfaces and bridge groups within a single switch router. Specifically, local or unroutable traffic will be bridged among the bridged interfaces in the same bridge group, while routable traffic will be routed to other routed interfaces or bridge groups.

Because bridging is in the data-link layer (Layer 2) and routing is in the network layer (Layer 3), they have different protocol configuration models. With IP, for example, bridge group interfaces belong to the same network and have a collective IP network address. In contrast, each routed interface represents a distinct network and has its own IP network address. Integrated routing and bridging uses the concept of a Bridge-Group Virtual Interface (BVI) to enable these interfaces to exchange packets for a given protocol.

A BVI is a virtual interface within the campus switch router that acts like a normal *routed* interface. A BVI does not support bridging, but it actually represents the corresponding bridge group to routed interfaces within the switch router. The interface number is the link between the BVI and the bridge group.

Layer 3 switching software supports the routing of IP and IPX between routed interfaces and bridged interfaces in the same router, in both fast-switching and process-switching paths.

Note BVIs do not support IP multicast routing.

Before Configuring IRB

Consider the following before configuring IRB:

- The default route/bridge behavior in a bridge group (when IRB is enabled) is to bridge all packets. Make sure you explicitly configure routing on the BVI for protocols that you want routed.
- Packets of nonroutable protocols such as local-area transport (LAT) are always bridged. You cannot disable bridging for the nonroutable traffic.

- The protocol attributes should not be configured on the bridged interfaces when using IRB to bridge and route a given protocol. Bridging attributes cannot be configured on the BVI.
- A bridge links several network segments into one large, flat network. To bridge a packet coming from a routed interface among the bridged interfaces, the whole bridge group should be represented by one interface.
- The BVI has default data-link and network-layer encapsulations. These encapsulations are the same as on the Ethernet, except that you can configure the BVI with some encapsulations that are not supported on a normal Ethernet interface.

Troubleshooting IRB Connections

To troubleshoot the IRB configuration, use the following commands:

Command	Purpose
show interfaces bvi *interface-name*	Displays BVI information, such as the BVI MAC address and processing statistics.
show interfaces irb	Displays the following BVI information: • Protocols that this bridged interface can route to the other routed interface if this packet is routable • Protocols that this bridged interface bridges • Entries in the software MAC-address filter

Troubleshooting IRB is a combination of both Layer 2 and Layer 3 troubleshooting. Check the following:

- Use the processes in the “Troubleshooting Layer 2 Switching” section on page 12-5.
- Use the processes in Chapter 11, “Troubleshooting Layer 3 Network Connections.”

When using these processes be aware of the following differences:

- Each physical Interface of the Bridge-Group has two MAC Addresses.
  - BVI MAC Address (same as the first Member of the Bridge-Group)
  - Physical MAC Address
- Packets destined to any of these MAC addresses will be considered for routing.
- The CAM must be programmed so the ingress interface can forward frames to a specific egress interface.
- High route processor utilization is common when using IRB routing.
- Packets get switched by the route processor when the Layer3 (IP or IPX) adjacency gets invalidated in the Layer3 CAM. This is caused by MAC address aging in the Layer 2 CAM.

Follow these steps to troubleshoot the status of an IRB configuration:

Step 1 Use the **show interface bvi** *number* command to check the configuration and status of the BVI.

```
Switch# show interface bvi 1
➔ BVI1 is up, line protocol is up
  Hardware is BVI, address is 0000.0ccb.292c (bia 0000.0000.0000)
➔ Internet address is 172.20.52.123/29
  MTU 1500 bytes, BW 10000 Kbit, DLY 5000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  ARP type: ARPA, ARP Timeout 04:00:00
  Queueing strategy: fifo
  Output queue 0/0, 0 drops; input queue 0/75, 0 drops
Switch#
```

Step 2 Check the BVI field. It should indicate up.

Step 3 Check the Internet address field. It should include a valid address for routing.

Step 4 Use the **show interface irb** command to check the configuration and status of the IRB connections.

```
Switch# show interface irb

ATM0/0/0
.
(Information Deleted)
.

FastEthernet3/0/5

 Routed protocols on FastEthernet3/0/5:
  ip

 Bridged protocols on FastEthernet3/0/5:
  appletalk  clns       ip         ipx

 Software MAC address filter on FastEthernet3/0/5
  Hash Len    Address      Matches  Act      Type
  0x00:  0 ffff.ffff.ffff         0 RCV Physical broadcast
  0x1D:  0 0090.2156.d83c         0 RCV Interface MAC address
  0x20:  0 0000.0ccb.292c         0 RCV Bridge-group Virtual Interface
  0x2A:  0 0900.2b01.0001         0 RCV DEC spanning tree
  0xC0:  0 0100.0ccc.cccc      1185 RCV CDP
  0xC1:  0 0100.0ccc.cccd         0 RCV SSTP MAC address
  0xC2:  0 0180.c200.0000         0 RCV IEEE spanning tree
  0xC2:  1 0180.c200.0000         0 RCV IBM spanning tree
  0xC2:  2 0100.0ccd.cdce         0 RCV VLAN Bridge STP

FastEthernet3/0/6
.
(Information Deleted)
.
BVI1

 Routed protocols on BVI1:
  ip

Tunnel0
Switch#
```

If you determine that IRB is configured incorrectly, refer to the "Configuring Bridging" chapter in the *Layer 3 Switching Feature and Configuration Guide*.

Troubleshooting Trunk Port Problems

The switch router software provides several **show** commands that can be used for troubleshooting.

Troubleshooting Trunk Port Problems

To troubleshoot trunk port problems, use the following command:

Command	Purpose
show epc patricia interface {**fastethernet** | **gigabitethernet**} *slot/subslot/port* **vlan** *vlan* **detail**	Displays the patricia tree information for the specified VLAN (for each VLAN on trunk).

Troubleshooting Fast EtherChannel Problems

This section describes troubleshooting Fast EtherChannel problems and includes the following:

- Bridging Over Fast EtherChannel Overview, page 12-16
- Troubleshooting Bridging Over Fast EtherChannel, page 12-17

Bridging Over Fast EtherChannel Overview

Ether Channel is a trunking technology that groups together multiple full-duplex 802.3 Ethernet interfaces to provide fault-tolerant, high-speed links between switches, routers, and servers. EtherChannel is a logical aggregation of multiple Ethernet interfaces. EtherChannel forms a single higher bandwidth routing or bridging endpoint. EtherChannel is designed primarily for host-to-switch connectivity or Inter-Switch Link (ISL) switch-to-switch connectivity (for example, connectivity to a Catalyst 5500 switch).

In summary, EtherChannel provides the following benefits:

- Logical aggregation of bandwidth
- Load balancing
- Fault tolerance

The EtherChannel interface (consisting of up to four Ethernet interfaces) is treated as a single interface, which is called a *port channel*. You must configure EtherChannel on the EtherChannel interface rather than on the individual member Ethernet interfaces. You create the EtherChannel interface by using the **interface port-channel** interface configuration command. The switch router supports up to 64 port channels.

EtherChannel connections are fully compatible with Cisco IOS VLAN and routing technologies. The ISL VLAN trunking protocol can carry multiple VLANs across an EtherChannel, and routers attached to EtherChannel links can provide full multiprotocol routing with support for host standby using Host Standby Router Protocol (HSRP).

Your switch router supports Fast EtherChannel (FEC) and Gigabit EtherChannel (GEC).

Cisco Fast EtherChannel technology builds upon standards-based 802.3 full-duplex Fast Ethernet to provide a reliable high-speed solution for the campus network backbone. Fast EtherChannel provides bandwidth scalability within the campus by providing increments of 200 Mbps to 800 Mbps.

Cisco Gigabit EtherChannel technology provides bandwidth scalability within the campus by providing increments of 2 Gbps to 8 Gbps.

Note EtherChannel does not support IP/IPX filtering at Layer 3 with the ACL daughter card.

Note For more detailed information about EtherChannel, refer to the “Configuring LAN Interfaces” chapter in the *Cisco IOS Interface Configuration Guide*.

Troubleshooting Bridging Over Fast EtherChannel

To troubleshoot the EtherChannel status and configuration, use the following commands:

Command	Purpose
show interfaces port-channel *number*	Displays the status of the physical interface.
show epc fe-channel interface {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port* **channel port-channel** *number*	Displays all EPC interface information for the specific interface and port channel.
show epc if-entry interface {**fastethernet** | **gigabitethernet**} *slot***/***subslot***/***port* **entry port-channel** *number*	Displays interface entry information for the specific interface.

Follow these steps to troubleshoot the EtherChannel status and configuration:

Step 1 Use the **show interfaces port-channel** *number* command to confirm the EtherChannel status and configuration.

```
  Switch# show interfaces port-channel 1
➔ Port-channel1 is up, line protocol is up
➔   Hardware is FEChannel, address is 0010.073c.0513 (bia 0000.0000.0000)
    MTU 1500 bytes, BW 300000 Kbit, DLY 100 usec,
       reliability 255/255, txload 1/255, rxload 1/255
    Encapsulation ARPA, loopback not set
    Keepalive set (10 sec)
    Unknown duplex, Unknown Speed
    ARP type: ARPA, ARP Timeout 04:00:00
      No. of active members in this channel: 3
          Member 0 : FastEthernet1/0/4 , Full-duplex, 100Mb/s
          Member 1 : FastEthernet1/0/6 , Full-duplex, 100Mb/s
          Member 2 : FastEthernet1/0/7 , Full-duplex, 100Mb/s
    Last input 00:00:00, output never, output hang never
    Last clearing of "show interface" counters never
    Queueing strategy: fifo
    Output queue 0/120, 0 drops; input queue 0/225, 0 drops
    5 minute input rate 13000 bits/sec, 17 packets/sec
    5 minute output rate 2000 bits/sec, 1 packets/sec
       24335602 packets input, 2345055668 bytes
       Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
       0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored 0 abort
       0 watchdog, 0 multicast
       0 input packets with dribble condition detected
       1366573 packets output, 289782286 bytes, 0 underruns
       0 output errors, 0 collisions, 0 interface resets
       0 babbles, 0 late collision, 0 deferred
       0 lost carrier, 0 no carrier
       0 output buffer failures, 0 output buffers swapped out
  Switch#
```

Step 2 Check the Port-channel field. It should indicate up.

Step 3 Note the MAC address assigned to the port channel. It will be used in Step 4.

Step 4 Use the **show epc if-entry interface** command with the **entry** interface parameters to display the status of the Broute VCs.

```
Switch# show epc if-entry interface fastEthernet 1/0/4 entry port-channel 1
IF Entry for Port-channel1 on FastEthernet1/0/4
➔   Mac(hex) - 00:10:07:3C:05:13
    isMyInteface : True isSubInterface : False
➔   Status Up Broute VC - 97 Bcast VC - 0
    Netmask: 32
    FEC enabled ( Flow-based Load-balancing )
    Trunking Enabled
    State : Not-Applicable/Listening/Blocking
    Bridge-Group disabled
    IP routing off bridging off
    IPX routing off bridging off
    Appletalk routing off
    In Encapsulation:
    ICMP Redirect enabled Unreachable enabled
    IP Multicast disabled: ttl-threshold: 0
    ACL Indexs:
    Input ACL: 0 Output ACL: 0
    ACL Flags:
    Input IP: OFF Output IP: OFF
    Input IPX: OFF Output IPX: OFF
Switch#
```

Step 5 Confirm that the MAC address in this step matches the MAC address displayed in Step 1.

For inconsistencies between the adjacency table and the EPC IP address table, use the **clear arp** or **clear adjacencies** command to rebuild the tables. When you use these commands, the router will send an ARP request for all entries in the ARP cache. As replies come back, it will refresh the cache. If any entries time out, they will be cleared from the table. The router will then build the adjacency table using this information, and populate the interface EPC IP address table.

Step 6 Use the **show epc if-entry interface** command with the **entry** interface parameters to display the status of the VCs.

```
Switch# show epc fe-channel interface fastEthernet 1/0/4 channel port-channel 1
FEC Group (Port-channel1) Information on FastEthernet1/0/4
    Member 0 VC - 97
    Member 1 VC - 177
    Member 2 VC - 217
    Member 3 VC - 217
    Member Ship BitMap 0x7
Switch#
```

If you determine that the EtherChannel interface is configured incorrectly, refer to the "Configuring EtherChannel" chapter in the *Layer 3 Switching Feature and Configuration Guide*.

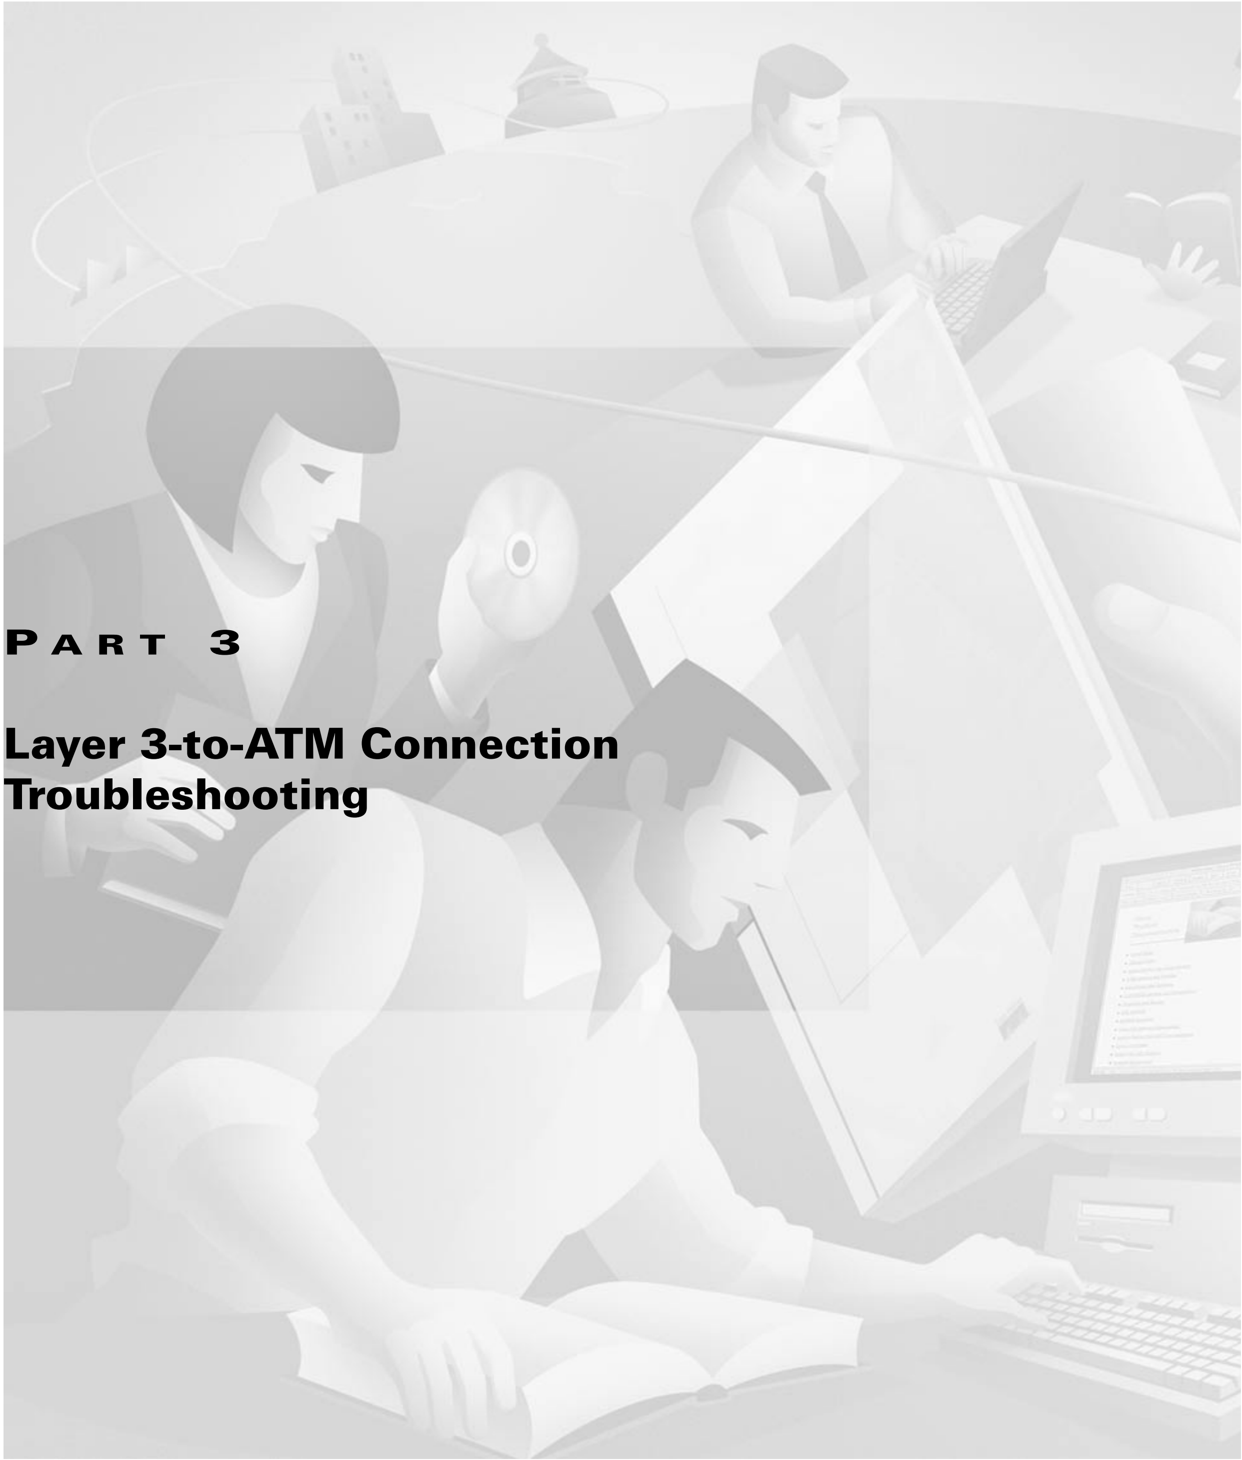

PART 3

Layer 3-to-ATM Connection Troubleshooting

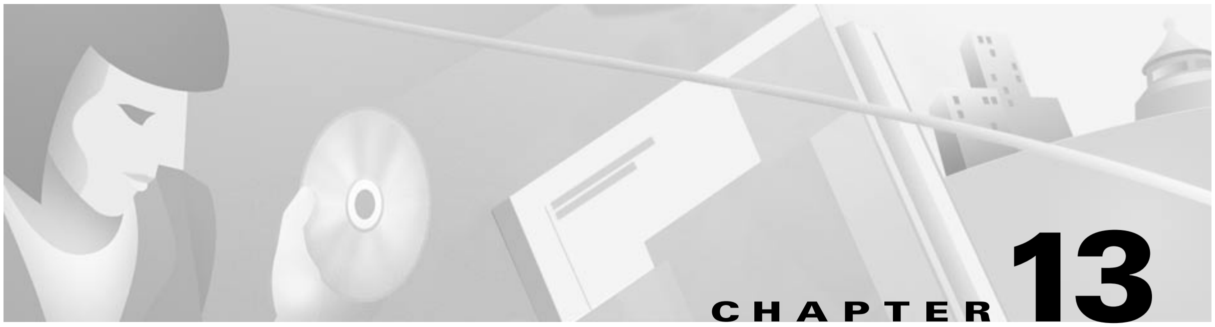

CHAPTER 13

Troubleshooting ATM Router Module Connections

This chapter provides troubleshooting information about connectivity and performance problems in the ATM router module (ARM) on the Catalyst 8540 CSR, Catalyst 8510 MSR, and LightStream 1010 ATM switch routers, and the enhanced ATM router module for the Catalyst 8540 CSR. The ATM router module allows you to integrate Layer 3 switching with ATM switching on the same switch router.

Note For complete descriptions of the commands mentioned in this chapter, refer to the *ATM Switch Router Command Reference*. For hardware installation and cabling instructions, refer to the *ATM and Layer 3 Port Adapter and Interface Module Installation Guide*.

Note The LightStream 1010 system software image does not include support for the ATM router module or Layer 3 features. You can download the Catalyst 8510 MSR image to a LightStream 1010 ATM switch router that has a multiservice ATM switch processor installed.

The chapter includes the following sections:

Note For detailed cabling and hardware information for each port adapter, refer to the *Catalyst 8540 CSR Route Processor and Interface Module Installation Guide.*

Overview of Integrated Layer 3-to-ATM Switching and Routing

This section describes the ATM router module that can be installed in Catalyst 8540 MSR and Catalyst 8540 CSR chassis that allows direct connections from the Layer 3 Ethernet network to the ATM backbone.

ATM Router Module Overview

The ATM router module allows you to integrate Layer 3 routing and ATM switching within a single chassis. When you install the ATM router module, you no longer need to choose either Layer 3 or ATM technology, as is frequently the case with enterprise, campus, and MAN applications.

The ATM router module can perform one or a combination of the functions described in Figure 13-1.

Figure 13-1 ATM Router Module Routing and Bridging Functions

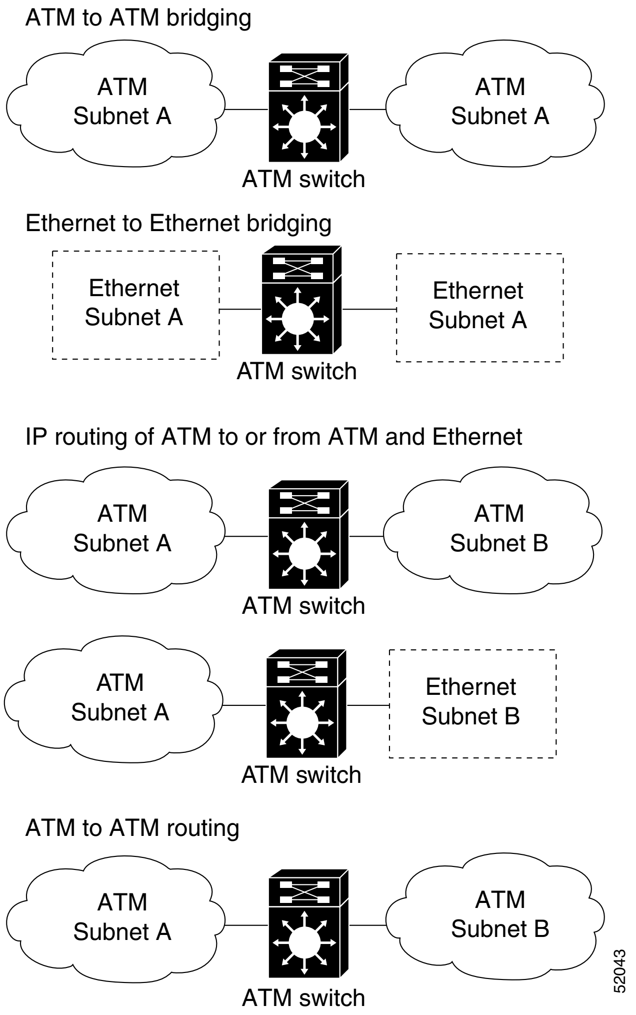

The ATM router module receives Address Resolution Protocol (ARP) messages and route broadcasts from connected ATM peers, and sends the appropriate control information to the route processor. On the ATM side, the ATM router module connects to the switching fabric as would any other interface module.

On the Catalyst 8540 CSR, the ATM router module supports LANE clients (LECs), but not the LANE servers (LES, LECS, and BUS). It separates the control and data path so that all LANE control messages are handled by the route processor, and data messages are switched on the ATM router module port, as shown in Figure 13-8. (See the “Comparing Data Plane and Control Plane Traffic” section on page 11-20 for a description of control and dataplane traffic.) The LEC is configured on the ATM router module interface, but control message traffic is sent to the route processor by the ATM router module. The ATM router module sends all ATM data traffic to the following VCs:

- In a LANE environment, the ATM router module sends all ATM data traffic to the Data Direct VCs.
- In an RFC 1483 environment, or multiprotocol encapsulation over ATM (MPOA), the ATM router module sends all ATM data traffic to the corresponding PVC.
- In an RFC 1577 environment, the ATM router module sends all ATM data traffic to the corresponding SVC.

Note The Catalyst 8540 CSR enhanced ATM router module does not support LANE clients.

The ATM router module has no external interfaces. All traffic is sent and received through internal interfaces to the switching fabric. The Catalyst 8540 CSR enhanced ATM router module has two internal ports. See the “Understanding Packet and Cell Flow” section on page 13-8 for a description of how the ATM router module interfaces connect to the other interfaces.

Hardware and Software Restrictions of the ATM Router Module

This section describes hardware and software restrictions for the ATM router module that could cause you connection or configuration problems.

Hardware Restrictions

The following hardware restrictions apply to the Catalyst 8540 CSR, Catalyst 8510 MSR, and LightStream 1010 ATM router modules, and the Catalyst 8540 CSR enhanced ATM router modules:

- You can install the ATM router module in any slot except a route processor slot and, in the case of the Catalyst 8540 CSR, a switch processor slot.
- The ATM router module is only supported on LightStream 1010 ATM switches that have a multiservice ATM switch route processor with FC-PFQ and the Catalyst 8510 MSR system software image.
- You can install up to two ATM router modules per chassis.
- When you hot swap an ATM router module, wait one minute after removing the module before inserting a new module.

Note The ATM router module is only supported on ATM switches that have a multiservice ATM switch processor installed.

ATM Router Module Software

This section describes software image requirements and restrictions that, if ignored could cause your ATM router module to malfunction.

Catalyst 8540 CSR Enhanced ATM Router Module Software Restrictions

The following software restrictions apply to the Catalyst 8540 CSR enhanced ATM router module:

- LANE is not supported.
- Use tag switching functionality with caution. Do not distribute routes learned through tag switching to Fast Ethernet (FE) or Gigabit Ethernet (GE), or vice versa. Otherwise, you might have unreachable route destinations.
- The ATM router module does not initialize if it replaces an ATM port adapter or interface module when hierarchical VP tunnels are globally enabled. Reboot the switch to initialize the ATM router module.
- ATM Director does not support any PVC commands.
- Up to 2048 external VCs can be configured on each ATM router module interface.
- Do not install an ATM router module in a slot pair with hierarchical VP tunnels configured. Slot pairs 0 and 1, 2 and 3, 9 and 10, and 11 and 12 use the same switching modules for scheduling. For example, do not install an ATM router module in slot 10 when hierarchical VP tunnels are configured on slot 9. For more information on hierarchical VP tunneling restrictions, refer to the "Configuring Virtual Connections" chapter in the *ATM Switch Router Software Configuration Guide*.

The Catalyst 8540 CSR enhanced ATM router modules do not support the following features:

- Point-to-point subinterfaces. Only point-to-multipoint subinterfaces are supported.
- Tag-edged router functionality
- Fast Simple Server Redundancy Protocol (SSRP)
- Bridging for multiplexing device encapsulation
- Protocol Independent Multicast (PIM) IP multipoint signalling
- PIM nonbroadcast multiaccess (NBMA)
- PIM over ATM multipoint signalling
- Translation from IP quality of service (QoS) to ATM QoS
- Resource Reservation Protocol (RSVP) to ATM SVC
- Access lists for ATM to ATM routing
- Half-bridge devices
- Layer 2 ACLs

Catalyst 8540 CSR ATM Router Module Software Restrictions

The following software restrictions apply to the Catalyst 8540 CSR ATM router module:

- Use tag switching functionality with caution. Do not distribute routes learned through tag switching to FE or GE, or vice versa. Otherwise, you might have unreachable route destinations.
- The ATM router module does not initialize if it replaces an ATM port adapter or interface module when hierarchical VP tunnels are globally enabled. Reboot the switch to initialize the ATM router module.
- ATM Director does not support any PVC commands.
- On an ATM router module interface, only LANE clients or RFC 1483 can be configured.

Note LANE clients or RFC 1483 can be configured on multiple ATM router module interface simultaneously.

- RFC 1483 on the ATM router module supports only AAL5 SNAP encapsulation.
- Up to 2048 external VCs can be configured on each ATM router module interface.
- You can have a maximum of 64 LECs per chassis.
- Do not install an ATM router module in a slot pair with hierarchical VP tunnels configured. Slot pairs 0 and 1, 2 and 3, 9 and 10, and 11 and 12 use the same switching modules for scheduling. For example, do not install an ATM router module in slot 10 when hierarchical VP tunnels are configured on slot 9. For more information on hierarchical VP tunneling restrictions, refer to the "Configuring Virtual Connections" chapter in the *ATM Switch Router Software Configuration Guide*.
- Token Ring LANE is not supported.

The Catalyst 8540 CSR ATM router modules do not support the following features:

- Point-to-point subinterfaces. Only point-to-multipoint subinterfaces are supported.
- Tag-edged router functionality
- Fast Simple Server Redundancy Protocol (SSRP)
- Bridging for multiplexing device encapsulation
- PIM IP multipoint signalling
- PIM NBMA
- PIM over ATM multipoint signalling
- Translation from IP QoS to ATM QoS
- RSVP to ATM SVC
- Access lists for ATM to ATM routing
- Half-bridge devices
- RFC 1483 MUX encapsulation
- ACL support for IP, and standard ACL support for IPX
- IP fragmentation support.
- IP 6-path load balancing support.

Catalyst 8510 MSR and LightStream 1010 ATM Router Module Software Restrictions

The following software restrictions apply to the Catalyst 8510 MSR ATM router module:

- Use tag switching functionality with caution. The switch router does not distribute routes learned through tag switching to FE or GE, or vice versa. If you use tag switching, you might have unreachable route destinations.

Note This is a temporary restriction. A soon-to-be-released software update will remove this limitation.

- The ATM router module does not initialize if it replaces an ATM port adapter or interface module when hierarchical VP tunnels are globally enabled. Reboot the switch to initialize the ATM router module.
- ATM Director does not support any PVC commands.
- RFC 1483 on the ATM router module supports only AAL5 SNAP encapsulation.
- Up to 2048 external VCs can be configured on each ATM router module interface.
- Do not install an ATM router module in a slot pair with hierarchical VP tunnels configured. Slot pair 0 and 1 and slot pair 3 and 4 use the same switching modules for scheduling. For example, do not install an ATM router module in slot 1 when hierarchical VP tunnels are configured on slot 0. For more information on hierarchical VP tunneling restrictions, refer to the "Configuring Virtual Connections" chapter in the *ATM Switch Router Software Configuration Guide*.
- RFC 1577 SVCs
- LANE clients are not supported.
- Only UBR PVCs are supported.

The Catalyst 8510 MSR and LightStream 1010 ATM router modules do not support the following features:

- Point-to-point subinterfaces. Only point-to-multipoint subinterfaces are supported.
- Tag-edged router functionality
- SSRP
- Bridging for multiplexing device encapsulation
- Protocol Independent Multicast (PIM) IP multipoint signalling
- PIM nonbroadcast multiaccess (NBMA)
- PIM over ATM multipoint signalling
- Translation from IP quality of service (QoS) to ATM QoS
- Resource Reservation Protocol (RSVP) to ATM SVC
- Access lists for ATM to ATM routing
- Half-bridge devices
- RFC 1483 MUX encapsulation
- ACL support for IP, and standard ACL support for IPX
- IP fragmentation support.
- IP 6-path load balancing support.

Note The ATM router module is only supported on ATM switches that have a multiservice ATM switch processor installed.

Understanding Packet and Cell Flow

This section describes packet flow through the ATM router module.

An ATM router module interface does not have any capabilities for ATM signalling. All ATM signalling is directed to the main route processor. The route processor is also responsible for setting up all ATM related VCs to enable the ATM router module to route any data traffic that it processes.

The ATM router module provides a packet-parsing, or look-up, engine that does not exist on the other ATM port adapter modules in the switch router. Data traffic coming from an ATM cloud targeted to hosts on the Ethernet side of the switch router are terminated on the ATM router module. The ATM router module processes the packets to identify the target port before the packets are sent to the Ethernet ports, ATM port, or route processor.

Note All LANE control frames are sent to route processor.

When an ATM router module encounters a spanning tree packet or an ARP request, it passes it to the route processor. Unlike an Ethernet module, the packet may have LANE or RFC 1483 encapsulation, and the packet must be transferred to the respective protocol layer once it reaches the route processor. For this to happen, one data VC per protocol is created when the ATM router module is initialized. These VCs are enabled as long as the ATM router module is present in the system.

The ATM router module port needs no external interfaces, such as cables, to come up. Each ATM router module interface has a unique MAC Address which is allocated by the route processor. You can configure subinterfaces on the ATM router module interfaces where the LECS or RFC 1483 clients are configured. The ATM interface allows limited ATM functionality; the subinterfaces on the ATM router module interface support full ATM functionality.

Note These subinterfaces are not created by default.

The ATM router module supports LANE clients (LECs), but not the LANE servers (LES, LECS, and BUS). It separates the control and data path so that all LANE control messages are handled by the route processor, and all data messages are switched on the ATM router module port, as shown in Figure 13-2. The LEC is configured on the ATM router module interface, but control message traffic is sent to the route processor by the ATM router module. The ATM router module then sends all ATM data traffic to the appropriate VCs.

Figure 13-2 ATM Router Module Traffic Flow (Catalyst 8540 CSR)

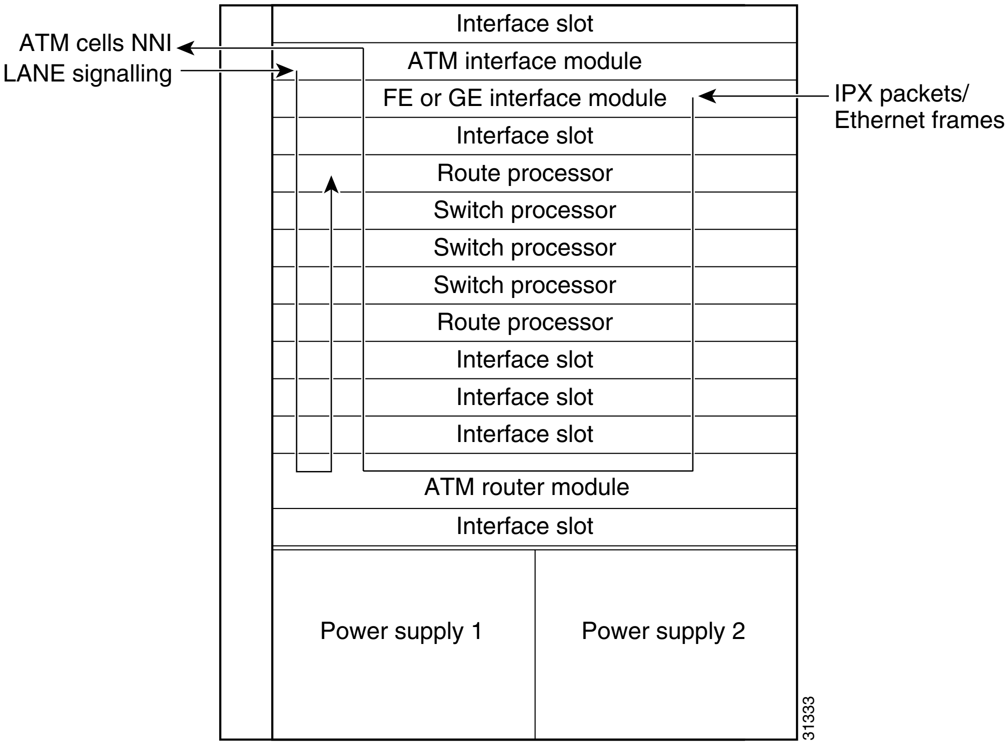

The design of ATM router module software is intended to separate the control and data paths so that all LANE control messages are handled by route processor, and all data is switched on the ATM router module port.

Note The LightStream 1010 ATM switch allows configuration of LECs only on the controller port subinterface (for example, the route processor atm2/0/0.subinterface). Thus, all VCs for signalling are terminated on the route processor.

Figure 13-3 shows the functional architecture of a switch router with an ATM router module installed. Traffic can enter the switch through any one of the ATM, Fast EtherChannel, or Gigabit Ethernet interfaces. Then the traffic is either:

- switched across the switch fabric to the route processor for initial route processing
- switched across the switch fabric to the ATM router module to be returned to the switch fabric for routing through any one of the remaining interfaces
- Layer 2 switched across the switch fabric to any one of the remaining interfaces

Figure 13-3 Packet Flow Through the ATM Router Module

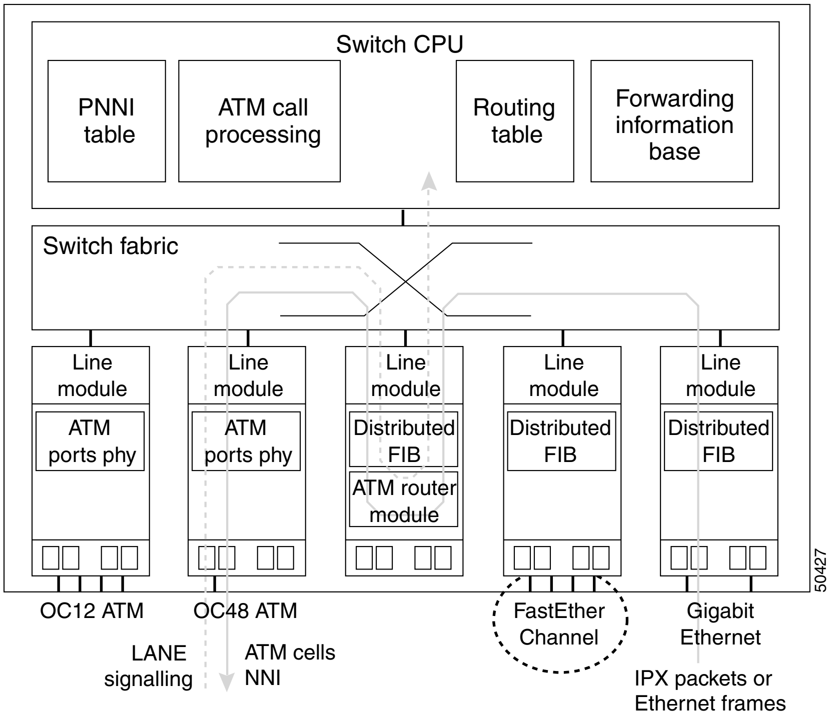

Logically, the ATM router module appears and functions like a router connected with both Gigabit Ethernet and ATM interfaces to the switch router on one side and the Ethernet and ATM networks connected to the other side. See Figure 13-4.

Figure 13-4 Logical View of the ATM Router Module in the Switch Router

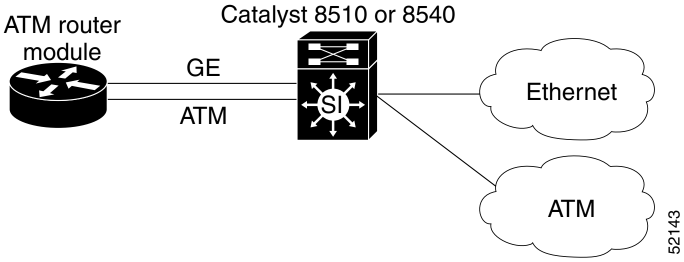

The ATM router modules for the switch routers have the following aggregate throughput:

- 2.5 Gb/s throughput for the Catalyst 8540 MSR ATM router module with two internal ATM interfaces
- 1.25 Gb/s throughput for the Catalyst 8510 MSR ATM router module with one internal ATM interface

Troubleshooting the ATM Router Module Hardware

The ATM router module for the Catalyst 8510 MSR is based on the single-port Gigabit Ethernet interface module. The ATM router module for the Catalyst 8540 CSR is based on the dual-port Gigabit Ethernet interface module. The ATM router module does not have the fiber transceivers in the faceplate. The Gigabit Ethernet processor interfaces are terminated on the board and only connect to the other interfaces on the ATM switch router through the backplane.

Note You can access the ATM router module interfaces using the standard CLI **show interface atm** *card/subcard/port* command and other interface commands.

On the faceplate, there is only one Status LED. If that LED appears green, the ATM router module is functioning properly, red means the ATM router module has failed its internal diagnostic self-tests.

Follow these steps to troubleshoot the ATM router module hardware:

Step 1 Use the **show hardware detail** command to confirm the ATM router module FPGA version and CAM configuration.

```
  Switch# show hardware detail

  Switch named Switch, Date: 18:23:14 UTC Tue Dec 5 2000

  Slot Ctrlr-Type    Part No.  Rev  Ser No  Mfg Date   RMA No. Hw Vrs  Tst EEP
  ---- ------------  ---------- -- -------- --------- -------- ------- --- ---
   0/* Super Cam     73-2739-03 D0 03170TAL May 03 99 0          3.1
   0/0 8T1 IMA PAM   73-3367-02 B2 03100061 Mar 15 99 00-00-00   2.0     0   0
   0/1 8E1 IMA PAM   73-3378-02 B2 03120056 Mar 25 99 00-00-00   2.0     0   2
➔  2/* ARM PAM       73-4208-01 05 03150016 Apr 18 99            1.0
   3/* ETHERNET PAM  73-3754-06 B0 03282WBF Jul 13 99 0          5.1
   9/* OC48c PAM     73-3745-02 12 03190UXC Jun 28 99            2.1
  10/* OCM Board     73-4165-01 04 03230ZZ2 Jun 28 99           10.1
  10/0 QUAD 622 Gen  73-2851-05 A0 03160RVS Jun 16 99            5.0
  11/* OC48c PAM     73-3745-02 12 03100015 Jun 28 99            2.1
  12/* OCM Board     73-4165-01 04 03190UJV Jun 28 99           10.1
  12/0 QUAD 622 Gen  73-2851-05 A0 03160S9J Jun 16 99 0          5.0

  .
  (Information Deleted)
  .

➔ slot:  2/*  Controller-Type : ARM PAM
    Part Number: 73-4208-01                       Revision: 05
  Serial Number: SCA03150016                      Mfg Date: Apr 18 99
     RMA Number:                               H/W Version: 1.0
➔  FPGA Version: 2.3

➔  EPIF Version: 1704                             CAM size: 64 KB
➔ Ucode Version: 0.0                              CAM Type: Dual

  Port Phy Setup
➔       Port  0: DONE                          GBIC Vendor: No vendor info.
➔       Port  1: DONE                          GBIC Vendor: No vendor info.

  slot:  3/*
  .
  (Information Deleted)
  .
```

Step 2 Check the Ctrlr-Type field. Find the slot where the ATM router module (shown as "ARM PAM") is installed.

Step 3 Check the FPGA Version field. It should match the version listed in the *Hardware and Software Compatibility Matrix.*

If it is not the correct version, update the FPGA image using the instructions in the "IOS Upgrade Procedures" section on page 3-11.

Step 4 Check the CAM size and type.

Note The GBIC Vendor field indicates no vendor information. These Gigabit interfaces, included with the ATM router module, are terminated on the board and only connect to the backplane.

If you determine that the interface is configured incorrectly, refer to the "Configuring ATM Router Module Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

Troubleshooting LANE Clients on ATM Router Module Connections

The troubleshooting process for LECs configured on the ATM router module is very similar to the troubleshooting process for ATM-to-ATM LANE connections described in Chapter 7, "Troubleshooting LAN Emulation Switching Environments," except for the following:

- All LANE VCs terminate on the ATM router module
- All ATM signalling is processed by the route processor and the ATM router module redirects LANE control traffic to route processor
- All Service Specific Connection Oriented Protocol (SSCOP) packets are forwarded directly to route processor

Figure 13-5 displays how the ATM router module installed in the Catalyst 8540 interacts with the other elements of the ATM network and allows connections to Ethernet networks.

Figure 13-5 ATM Router Module in the ATM Network

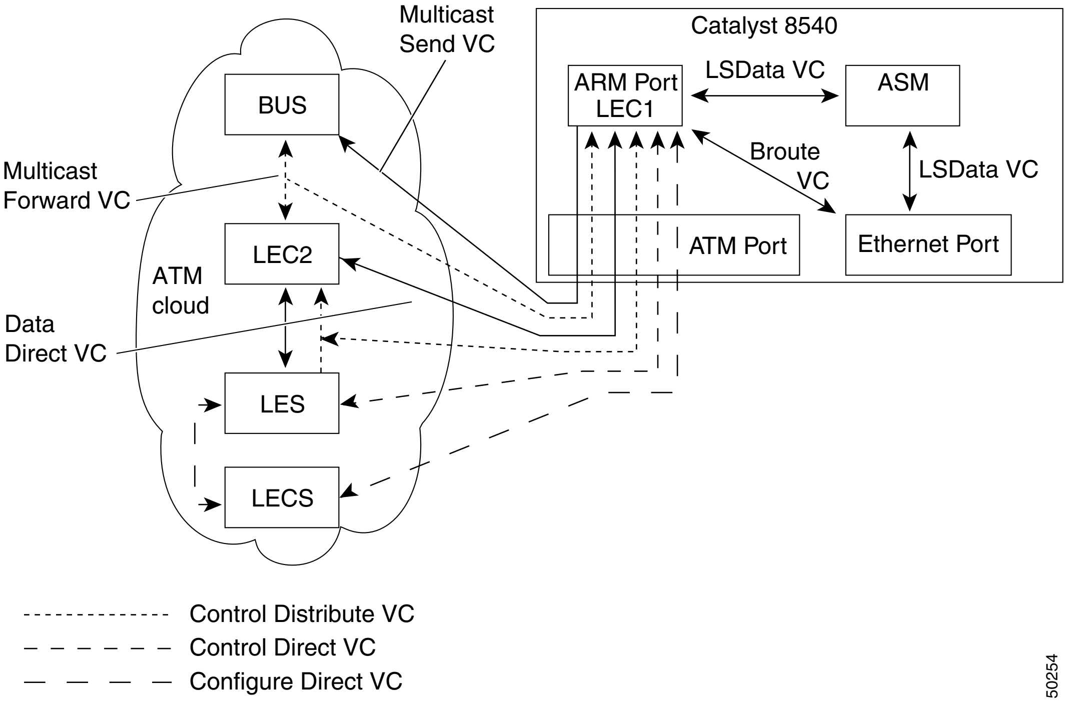

An ATM router module installed in the switch router allows the connection of Ethernet networks through the switch router to the ATM connections in the ATM cloud. The ATM router module has the following functions and limitations:

- Supports LEC configuration on the ATM router module interfaces and subinterfaces
- No LECS/LES/BUS configuration is allowed on the ATM router module
- Up to 32 LECs are allowed per ATM router module
- A maximum of two ATM router modules per chassis, allowing a maximum of 64 LECs
- Default ATM address prefix same as the route processor port
- All IP, IPX, and bridging commands are allowed on the ATM router module interfaces and subinterfaces
- **Shutdown** and **no shutdown** commands are allowed on the ATM router module interfaces and subinterfaces
- From the CLI, the ATM router module ATM interface is configured in the same manner as any other IOS switch router

Troubleshooting LECs Problems on the ATM Router Module Commands

To display the ATM router module and LECs configuration, use the following commands:

Command	Purpose
show lane client	Displays the LEC configuration and status
show atm vc interface *card/subcard/port.subinterface*	Displays the ATM layer connection information about the virtual connection.
show epc if-entry	Displays interface entry information for the specific interface.
show ip cef *ip-address*	Displays Cisco Express Forwarding information.
show epc ip-address interface {**fastethernet** | **gigabitethernet**} *slot/subslot/port ip-address*	Displays the IP addresses of adjacent interfaces.
show atm vc traffic interface atm *card/subcard/port VPI VCI*	Displays information about the ATM virtual connection.
ping *ip-address*	Confirms the IP connection and increments the transmit and receive cell counters.

To troubleshoot LECs configured on the ATM router module, refer to Chapter 7, "Troubleshooting LAN Emulation Switching Environments," and use normal LANE troubleshooting techniques.

Figure 13-6 is an example network of a switch router with an ATM router module configured with two LECs connecting an Ethernet network and an ATM network.

Figure 13-6 ATM Router Module LEC Example Network

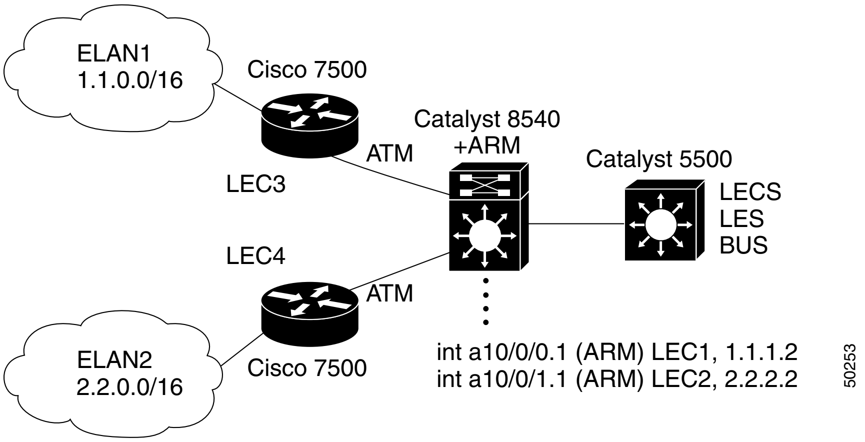

This example network is used in the following troubleshooting steps.

Follow these steps to troubleshoot the ATM router module LECs configured in the example:

Step 1 Use the following commands to configure the LECs on ATM interfaces 10.0.0.0 and 10.0.0.1.

```
Switch# config term
Switch(config)# interface atm10/0/0.1 multipoint
Switch(config-if)# lane client ethernet elan1
Switch(config-if)# ip address 1.1.1.2 255.255.0.0
Switch(config-if)# exit
Switch(config)# interface atm10/0/1.1 multipoint
Switch(config-if)# lane client ethernet elan2
Switch(config-if)# ip address 2.2.2.2 255.255.0.0
Switch(config-if)# end
Switch#
```

Step 2 Use the **show running-config** command to confirm the LEC configuration of the ATM router module interfaces.

```
Switch# show running-config
Building configuration...

Current configuration:
!
.
(Information Deleted)
.
!
interface ATM10/0/0
 no ip address
 logging event subif-link-status
!
→ interface ATM10/0/0.1 multipoint
 ip address 1.1.1.2 255.255.0.0
 lane client ethernet elan1
!
interface ATM10/0/1
 no ip address
 atm pvc 2 100 pd on  inarp 10
!
→ interface ATM10/0/1.1 multipoint
 ip address 2.2.2.2 255.255.0.0
 lane client ethernet elan2
!
.
(Information Deleted)
.
```

Step 3 Use the **show lane client** command to confirm the various LEC connections are up and the configuration is valid.

```
Switch# show lane client
LE Client ATM10/0/0.3  ELAN name: ELAN3  Admin: up  State: operational
→ Client ID: 4                LEC up for 1 hour 52 minutes 56 seconds
.
(Information Deleted)
.
```

Step 4 Check the LEC field. It should be up.

Step 5 Use the **show atm vc interface** command to confirm the connections are up and the configuration is valid.

```
  Switch# show atm vc interface atm 10/0/0.3
➔ Interface         VPI  VCI   Type   X-Interface      X-VPI X-VCI Encap  Status
  ATM10/0/0           0   35    PVC       ATM0                0      271  LSCNTL UP
  ATM10/0/0           0   36    PVC       ATM0                0      272  LSDATA UP
  .
  (Information Deleted)
  .
  ATM10/0/0           0   743    SVC      ATM9/0/0            0       53  LANE UP
  ATM10/0/0           0   744    SVC      ATM9/0/0            0       54  LANE UP
  ATM10/0/0           0   745    SVC      ATM9/0/0            0       55  LANE UP
  .
  (Information Deleted)
  .
  ATM10/0/0           0   322    PVC      Gi3/0/1             0       67  LSDATA UP
  ATM10/0/0           0   323    PVC      Gi3/0/1             0       68  LSDATA UP
  ATM10/0/0           0   325    PVC      Gi3/0/1             0       70  LSDATA UP
  ATM10/0/0           0   326    PVC      Gi3/0/1             0       71  LSDATA UP
```

Step 6 Check the Status field. It should appear up for all LECs ATM interfaces.

Step 7 Use the **show epc if-entry interface** command and test the CAM information between the egress Gigabit Ethernet interface from the entry ATM interface.

```
  Switch# show epc if-entry interface atm 10/0/0 entry gigabitEthernet 3/0/1
  IF Entry for GigabitEthernet3/0/1 on ATM10/0/0
      Mac(hex) - 00:90:21:41:88:38
      isMyInteface : False isSubInterface : False
➔     Status Up Broute VC - 322 Bcast VC - 0
      Netmask: 24
      FEC disabled
      Trunking Disabled
      State : Not-Applicable/Listening/Blocking
      Bridge-Group disabled
➔     IP routing on bridging off
      IPX routing off bridging off
      Appletalk routing off
      In Encapsulation:
      ICMP Redirect enabled Unreachable enabled
      IP Multicast disabled: ttl-threshold: 0
```

Check the following:

- Broute VC field status is up.
- Note the Broute VC number. In this example, the Broute VC is “322.”
- IP routing is on.

Step 8 Use the **show epc if-entry interface** command and test the CAM information in the opposite direction between the egress ATM interface from the entry Gigabit Ethernet interface.

```
   Switch# show epc if-entry interface gigabitEtherenet3/0/1 entry atm 10/0/0.3
   IF Entry for ATM10/0/0.3 on GigabitEthernet3/0/1
       Mac(hex) - 00:90:21:41:88:17
       isMyInteface : False isSubInterface : True
➔      Status Up Broute VC - 67 Bcast VC - 0
       Netmask: 25
       FEC disabled
       Trunking Disabled
       State : Not-Applicable/Listening/Blocking
       Bridge-Group disabled
➔     IP routing on bridging off
       IPX routing off bridging off
       Appletalk routing off
       In Encapsulation:
       ICMP Redirect enabled Unreachable enabled
       IP Multicast disabled: ttl-threshold: 0
      LECID - 0, Multicast Send VC - 0
```

Check the following:

- Broute VC field status is up.
- Note the Broute VC number. In this example, the Broute VC is “67.”
- IP routing is on.

Step 9 Use the **show ip cef** command to verify that routes and attached devices appear in the table correctly and point to the correct next hop or outgoing interface.

```
   Switch# show ip cef 128.250.0.1
   128.250.0.1/32, version 90, connected, cached adjacency 128.250.0.1
   0 packets, 0 bytes
     via 128.250.0.1, ATM10/0/0.1, 0 dependencies
➔      next hop 128.250.0.1, ATM10/0/0.1, valid cached adjacency
```

Step 10 Use the **show epc ip-address** command with the IP address of the egress interface to display the status of the MAC address rewrite and the VCI number.

```
   Switch# show epc ip-address interface atm 10/0/0 128.250.0.1
➔  IPaddr: 128.250.0.1 MACaddr: 0000.0c07.ac01  Routed to VC(940)
```

Check the Routed to VC field (in this example, the VC is “940”). The value is used in the next step.

Step 11 Use the **show atm vc traffic interface atm** command with the VPI and VCI parameters to see the receive and transmit cell counts.

```
   Switch# show atm vc traffic interface atm 10/0/0 0 940
   Interface         VPI  VCI   Type      rx-cell-cnts     tx-cell-cnts
➔  ATM10/0/0          0    940   SVC              18                25
```

Step 12 Use the **ping** command to confirm the connection and increment the receive and transmit cell counts.

```
Switch# ping 128.250.0.1
Sending 5, 100-byte ICMP Echos to 128.250.0.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Step 13 Again, use the **show atm vc traffic interface atm** command with the VPI and VCI parameters, to confirm that the receive and transmit cell counts are incrementing.

```
  Switch# show atm vc traffic interface atm 10/0/0 0 940
  Interface         VPI  VCI   Type      rx-cell-cnts     tx-cell-cnts
➔ ATM10/0/0           0    940   SVC              33               40
```

If you determine that the interface is configured incorrectly, refer to the “Configuring ATM Router Module Interfaces” chapter in the *ATM Switch Router Software Configuration Guide*.

ATM Router Module IP Switching Overview

This section describes IP switching using the ATM router module.

Figure 13-7 shows a network with a switch router that has an ATM router module installed and connected to an Ethernet subnet on one side and two ATM ELAN networks on the other.

Figure 13-7 ATM Router Module IP Switching Example Network

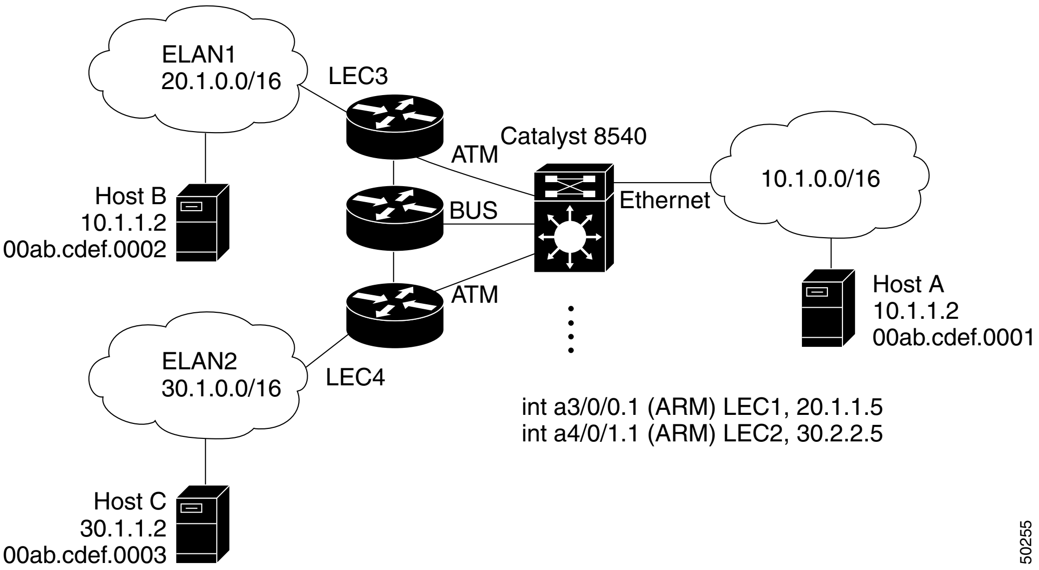

The following are the routing, CEF, and adjacency tables created for the network configuration on the Catalyst 8540 with an ATM router module, shown in Figure 13-7.

```
Routing Table:

10.1.0.0   is directly connected, FastEth 1/0/0
20.1.0.0   is directly connected, ATM3/0/0.1
30.1.0.0   is directly connected, ATM4/0/0.1

CEF Table

10.1.0.0    attached      FastEth 0/0/0
20.1.0.0    attached      ATM 3/0/0.1
30.1.0.0    attached      ATM 4/0/0.1

Adjacency Table:

FastEth   0/0/0    10.1.1.2   00ab.cdef.0001  Interface No.
ATM       3/0/0.1  20.1.1.2   00ab.cdef.0002  Data VC
ATM       4/0/0.1  30.1.1.2   00ab.cdef.0003  Data VC
```

Using this configuration, traffic entering the Catalyst 8540 through the Fast Ethernet interface 0/0/0 from Host A on network 10.1.0.0 propagates the CAM on the Ethernet interface with the following:

```
CAM Port FA 0/0/0

My-MAC=FE000
My-Subnet=10.1.0.0/16
Subnet    20.1.0.0/16
10.1.1.2=00ab.cdef.0001, S1.A,  FA0/0/0
20.1.1.5=LEC101,         S2.B,  ATM10/0/1

IF-MAP:
      ATM10/0/1   Broute VC    92
                  Bcast VC      0
```

Using Broute VC 92, the Ethernet packet is switched across the backplane to the ATM router module at ATM interface 1/0/1, where it propagates the CAM on the ATM interface with the following:

```
CAM Port atm 10/0/1

My-MAC=LEC101
My-Subnet=S2
Subnet S1
10.1.1.2=00ab.cdef.0001, S1.A, Fa 0/0/0
20.1.1.5=MAC_B, VPI0, VCI188
          (Data Direct VC)

IF-MAP:
      Fa 0/0/0   Broute VC    80
                 Bcast VC     0
```

Using VPI 10 and VCI 188, the ATM router module transfers the Ethernet packets across the backplane to ATM interface 9/0/3, for transmission out to the LANE cloud and subsequent delivery to the destination Host.

Troubleshooting IP switching with the ATM router module configured between the ATM and Ethernet interfaces is essentially the same as described in the “Troubleshooting IP Layer 3 Connections” section on page 11-25. However, you must confirm connections and adjacencies through the ATM router module.

IPX Switching Overview

This section describes IPX switching using the ATM router module.

Figure 13-8 shows a network with a switch router that has an ATM router module installed and connected to an Ethernet subnet on one side and two ATM ELAN networks on the other.

Figure 13-8 ATM Router Module IPX Switching Example Network

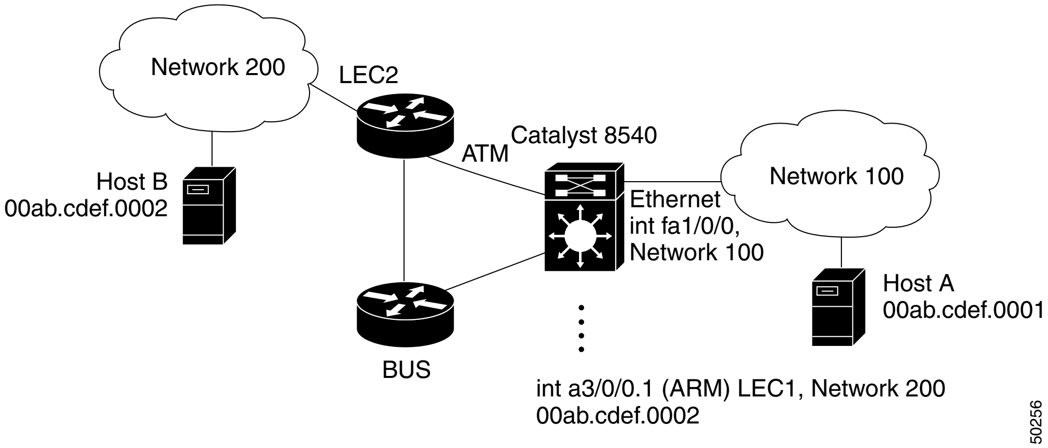

The following are the routing and node tables created for the network configuration on the Catalyst 8540 with an ATM router module shown in Figure 13-8.

```
Routing Table

Network 100, Primary network, FastEth 1/0/0
Network 200, Primary network, ATM 3/0/0.1

Node Table:

FastEth  1/0/0    100.00ab.cdef.0001  Interface no.
ATM      3/0/0.1  200.00ab.cdef.0002  Data VC
```

Troubleshooting IP switching with the ATM router module configured between the ATM and Ethernet interfaces is essentially the same as described in the "Troubleshooting IPX Layer 3 Routing" section on page 11-39. However, you must confirm connections and adjacencies through the ATM router module.

General ATM Router Module Troubleshooting

Use the following commands to troubleshoot general ATM router module connections:

Command	Purpose
ping *ip-address*	Tests the network node reachability.
show epc ifmapping *interface-map-number*	Displays interface mapping to CAM interface number.
show epc ip-address interface atm *card/subcard/port ip-address*	Displays all adjacent IP addresses for the specified interface.

Use the following commands to troubleshoot the ATM router module LECs and the configuration:

Step 1 Use the **ping** command to confirm the LEC connection to the end station.

```
Switch# ping 128.250.20.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 128.250.20.1, timeout is 2 seconds:

3d09h: LEC ATM10/0/0.3: received SETUP
3d09h: LEC ATM10/0/0.3:   callid         0x638CACB8
3d09h: LEC ATM10/0/0.3:   called party   39.036F100703060001010000000.009021418817.03
3d09h: LEC ATM10/0/0.3:   calling_party  39.036F100703060010073C1501.0010F66C841C.03
3d09h: LEC ATM10/0/0.3: sending CONNECT
3d09h: LEC ATM10/0/0.3:   callid         0x638CACB8
3d09h: LEC ATM10/0/0.3:   vcd            943
3d09h: LSSLEC:LANE-Cache_VC:if=45 vc=943 type=6
3d09h: LEC ATM10/0/0.3: received CONNECT_ACK.!!.!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms
3d09h: LSSLEC:ADJ:Host:128.250.20.1 MAC:0000.0c07.ac03 Valid:TRUE MyIP:FALSE if=8 vc=943 if_vc=45
3d09h: LSSLEC:ADJ:Host:128.250.20.1 MAC:0000.0c07.ac03 Valid:TRUE MyIP:FALSE if=8 vc=943 if_vc=45
```

Check the if_vc= field and make note of the interface VC number. In this case, the interface VC is “45.”

Step 2 Use the **show epc ifmapping** command with the interface number to confirm the interface VC is mapped correctly to the ATM interface.

```
  Switch# show epc ifmapping 45
➔ ATM10/0/0.3               (IF number: 45)
```

The IF number field (in this example, “45”) indicates the interface index number is mapping correctly.

Step 3 Use the **show epc ip-address interface** command with the IP address to confirm the VC number.

```
Switch# show epc ip-address interface atm 10/0/0 128.250.20.1
IPaddr: 128.250.20.1 MACaddr: 0000.0c07.ac03  Routed to VC(943)
```

The information in this display should match the information shown using the **show adjacency** command to display the MAC address rewrite.

If you determine that the interface is configured incorrectly, refer to the “Configuring ATM Router Module Interfaces” chapter in the *ATM Switch Router Software Configuration Guide*.

Troubleshooting RFC 1483 on ATM Router Module Connections

The primary use of RFC 1483, or multiprotocol encapsulation over ATM (MPOA), is to carry multiple Layer 3 and bridged frames over ATM. RFC 1483 traffic is routed through the ATM router module interface using static map lists. Static map lists provide an alternative to using the ATM Address Resolution Protocol (ARP) and ATM Inverse ARP (InARP) mechanisms.

Note Traffic shaping and policing are not supported on the ATM router module interfaces. Use VP tunnels as an alternative for traffic shaping on ATM connections.

Troubleshooting RFC 1483 Problems on ATM Router Module Commands

To display the ATM router module and RFC 1483 configuration, use the following commands:

Command	Purpose
show running-config	Shows the status of the configuration and physical interfaces.
show interfaces atm *card/subcard/port*	Shows the status of the physical interface.
show atm vc interface *card/subcard/port.subinterface*	Displays the ATM layer connection information about the virtual connection.
ping *ip-address*	Confirms the IP connection and increments the transmit and receive cell counters.

Figure 13-6 is an example network of a switch router with an ATM router module configured with two RFC 1483 interfaces connecting an Ethernet network and an ATM network.

Figure 13-9 ATM Router Module RFC 1483 Example Network

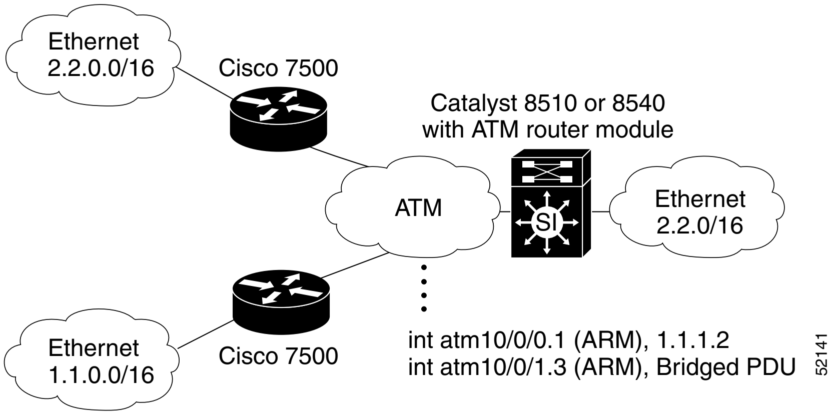

This example network is used in the following troubleshooting process.

To troubleshoot the ATM router module configured with RFC 1483 and with aggressive policing configured follow these processes:

- Verify VC status between ATM router module interfaces and the route processor.
- Display VC details between ATM router module interfaces and ATM interfaces
- Display VC details between ATM router module interfaces and Ethernet interfaces
- Verify policing or Usage Parameter Control (UPC) on switch and traffic shaping on routers
  - Policing UPC—Traffic is constantly monitored at the switch to ensure the contract is not violated. Non-conforming cells may be marked, dropped or passed
  - Traffic shaping—Typically done at edge devices to reduce burstiness. Decreases probability of cell loss at expense of occasional delay between ATM router module interfaces and Ethernet interfaces

Follow these steps to troubleshoot the ATM router module configured with RFC 1483 and very aggressive policing.

Step 1 Use the following commands to configure the aggressive policing, the ATM router module with RFC 1483, and the map list.

```
Switch(config)# atm connection-traffic-table-row index 110 ubr pcr 1
Switch(config)# interfaces atm10/0/1.3 multipoint
Switch(config-if)# ip address 2.2.0.2 255.255.255.0
Switch(config-if)# map-group RFC1483_2
Switch(config-if)# atm pvc 2 109 pd on interface ATM0/0/0 0 109 upc drop
Switch(config-if)# exit
Switch(config)# int atm10/0/1.3 multipoint
Switch(config-if)# atm pvc 2 101 pd on interface ATM0/0/1 0 101
Switch(config-if)# bridge-group 10
Switch(config-if)# exit
Switch(config)# map-list RFC1483_2
Switch(config-map-list)# ip 2.2.0.1 atm-vc 109 broadcast
Switch(config-map-list)# end
Switch(config)# bridge 10 protocol ieee
```

Step 2 Use the **show running-config** command to confirm your configuration.

```
   Switch# show running-config
   Building configuration...
   Current configuration:
   .
   (Information Deleted)
   .
   !
➔  atm connection-traffic-table-row index 110 ubr pcr 1
   .
   (Information Deleted)
   .
➔  interface ATM10/0/1.3 multipoint
    ip address 2.2.0.2 255.255.255.0
    map-group RFC1483_2
    atm pvc 2 109 pd on rx-cttr 110 tx-cttr 110 interface ATM0/0/0 0 109 upc drop
   !
   .
   (Information Deleted)
   .
➔  map-list RFC1483_2
    ip 2.2.0.1 atm-vc 109 broadcast
    ip 11.1.10.1 atm-vc 110 broadcast
```

Step 3 Use the **show interfaces atm** command to confirm the configuration of the ingress ATM interface connected to the Cisco 7500 router.

```
  Switch# show interfaces atm 0/0/0
➔ ATM0/0/0 is up, line protocol up
➔ Hardware is quad_oc12suni
    Internet address is 10.0.1.26/30
    MTU 4470 bytes, sub MTU 4470, BW 622080 Kbit, DLY 0 usec,
        rely 255/255, load 1/255
    Encapsulation ATM, loopback set, keepalive not supported
    Last input 00:00:00, output 00:00:00, output hang never
    Last clearing of "show interface" counters never
    Queueing strategy: fifo
    Output queue 0/40, 0 drops; input queue 0/75, 0 drops
    5 minute input rate 4000 bits/sec, 9 packets/sec
    5 minute output rate 4000 bits/sec, 9 packets/sec
       16300 packets input, 863900 bytes, 0 no buffer
       Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
➔      0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
       16808 packets output, 890824 bytes, 0 underruns
       0 output errors, 0 collisions, 1 interface resets
       0 output buffer failures, 0 output buffers swapped out
```

Step 4 Confirm the interface and protocol are both up.

Step 5 Use the **show interfaces atm** command to confirm the configuration of the ATM router module interface.

```
  Switch# show interfaces atm 10/0/1
➔ ATM10/0/1 is up, line protocol is up
    Hardware is ATM router module_port, address is 0090.2141.bc58 (bia 0090.2141.bc58)
    SVC idle disconnect time: 300 seconds
    MTU 1500 bytes, sub MTU 1500, BW 1000000 Kbit, DLY 10 usec,
       rely 255/255, load 1/255
    Encapsulation ATM, loopback not set, keepalive not supported
    Full-duplex, 1000Mb/s, 100BaseFX
    ARP type: ARPA, ARP Timeout 00:15:00
    Last input 00:02:13, output never, output hang never
    Last clearing of "show interface" counters never
    Queueing strategy: fifo
    Output queue 0/40, 0 drops; input queue 0/75, 0 drops
    5 minute input rate 0 bits/sec, 0 packets/sec
    5 minute output rate 0 bits/sec, 0 packets/sec
       23 packets input, 10352 bytes, 0 no buffer
       Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
➔      0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
       0 watchdog, 0 multicast
       23 packets output, 10352 bytes, 2 underruns
       2 output errors, 0 collisions, 0 interface resets
       0 output buffer failures, 0 output buffers swapped out
```

Step 6 Confirm the ATM router module interface and protocol are both up.

Step 7 Use the **show atm vc interfaces atm** command with VPI 0 and VCI 109 to confirm the configuration of the ingress ATM interface connected to the Cisco 7500 router.

```
  Switch# show atm vc interfaces atm 0/0/0 0 109

  Interface: ATM0/0/0, Type: quad_oc12suni
  VPI = 0  VCI = 109
➔ Status: UP
  Packet-discard-option: enabled
  Usage-Parameter-Control (UPC): drop
  Wrr weight: 2
➔ Cross-connect-interface: ATM10/0/1, Type: arm_port
  Cross-connect-VPI = 2
  Cross-connect-VCI = 109
  Threshold Group: 5, Cells queued: 0
  Rx cells: 0, Tx cells: 0
  Tx Clp0:0,  Tx Clp1: 0
  Rx Clp0:0,  Rx Clp1: 0
  Rx Upc Violations:0, Rx cell drops:0
  Rx pkts:0, Rx pkt drops:0
  Rx connection-traffic-table-index: 110
  Rx service-category: UBR (Unspecified Bit Rate)
```

Step 8 Confirm the Status is up.

Step 9 Confirm that the Cross-connect-interface is the ATM router module internal interface.

Step 10 Again use the **show atm vc interfaces atm** command with VPI 0 and VCI 109 on the ingress ATM interface connected to the Cisco 7500 router, to confirm that the cell and packet numbers are incrementing.

```
  Switch# show atm vc interfaces atm 0/0/0 0 109

  Interface: ATM0/0/0, Type: quad_oc12suni
  VPI = 0  VCI = 109
  Status: UP
  Packet-discard-option: enabled
  Usage-Parameter-Control (UPC): drop
  Wrr weight: 2
  Cross-connect-interface: ATM10/0/1, Type: arm_port
  Cross-connect-VPI = 2
  Cross-connect-VCI = 109
  Threshold Group: 5, Cells queued: 6
➔ Rx cells: 15, Tx cells: 0
  Tx Clp0:0,  Tx Clp1: 0
➔ Rx Clp0:15,  Rx Clp1: 0
➔ Rx Upc Violations:5, Rx cell drops:9
➔ Rx pkts:0, Rx pkt drops:5
  Rx connection-traffic-table-index: 110
  Rx service-category: UBR (Unspecified Bit Rate)
```

Step 11 Check the Rx cells fields. The numbers should have incremented from the previous display.

Step 12 From the downstream Cisco 7500 router use the **ping** command, with the IP address of the ATM router module, to send five ICMP messages.

```
  C7500# ping 2.2.0.2

  Type escape sequence to abort.
  Sending 5, 100-byte ICMP Echos to 2.2.0.2, timeout is 2 seconds:
  .....
➔ Success rate is 0 percent (0/3)
```

Step 13 Confirm that the Success rate is 0.

Step 14 From the downstream router, use the extended **ping ip** command, with the IP address of the ATM router module, to send five 64-byte ICMP messages.

```
  C7500#ping ip
  Target IP address: 2.2.0.2
  Repeat count [5]:
  Datagram size [100]: 64
  Timeout in seconds [2]:
  Extended commands [n]:
  Sweep range of sizes [n]:
  Type escape sequence to abort.
  Sending 5, 64-byte ICMP Echos to 2.2.0.2, timeout is 2 seconds:
  .!.!.
➔ Success rate is 40 percent (2/5), round-trip min/avg/max = 1/1/1 ms
  C7500#
```

Step 15 Check the Success rate field. Note that the success rate improved to 40 percent after you changed the ICMP datagram size to 64 bytes, from the default 100 bytes used in the previous **ping** command attempt.

Step 16 Again use the **show atm vc interfaces atm** command with VPI 0 and VCI 109 on the ingress ATM interface connected to the Cisco 7500 router.

```
  Switch# show atm vc interface atm 0/0/0 0 109

  Interface: ATM0/0/0, Type: quad_oc12suni
  VPI = 0  VCI = 109
  Status: UP
  Threshold Group: 5, Cells queued: 0
➔ Rx cells: 25, Tx cells: 4
  Tx Clp0:4,  Tx Clp1: 0
  Rx Clp0:25,  Rx Clp1: 0
➔ Rx Upc Violations:9, Rx cell drops:14
➔ Rx pkts:2, Rx pkt drops:8
  Rx connection-traffic-table-index: 110
```

Step 17 Check the Rx cells field. Notice that the number is incrementing.

Step 18 Check the values in Rx Upc Violations, Rx cell drops, and Rx pkt drops fields. These values are also incrementing proving that the aggressive policing, configured with the **atm connection-traffic-table-row index 110 ubr pcr 1** command setting the peak cell rate to 1, is working correctly.

If you determine that the ATM router module interface is configured incorrectly, refer to the "Configuring ATM Router Module Interfaces" chapter in the *ATM Switch Router Software Configuration Guide*.

Troubleshooting RFC 1577 on ATM Router Module Connections

Classical IP over ATM is an IETF protocol that uses high speed ATM connections to provide better connections between IP members. The classic logical IP subnet (LIS) is implemented over an ATM switching network using an ATMARP server to replace the broadcast ARP service. IP over ATM is a Layer 3 switching service, where IP and ARP datagrams are encapsulated in AAL5 using IETF RFC1483 LLC/SNAP encapsulation as the default. RFC1577 provides for "best effort" service only. However, Resource Reservation Protocol (RSVP) over ATM enhances classical IP to support RSVP signalling, allowing differentiated QoS over an ATM network.

Troubleshooting RFC 1577 Problems on the ATM Router Module Commands

To display and troubleshoot the ATM router module and RFC 1577 configuration, use the following commands:

Command	Purpose
show running-config	Shows the status of the configuration and physical interfaces.
show interfaces atm *card/subcard/port*	Shows the status of the physical interface.
show atm vc interface *card/subcard/port.subinterface*	Displays the ATM layer connection information about the virtual connection.
ping *ip-address*	Confirms the IP connection and increments the transmit and receive cell counters.
show atm map	Confirms the static connection to the ATM ARP server.

Figure 13-6 is an example network of a switch router with an ATM router module configured with two RFC 1577 connections between Ethernet networks and an ATM network.

Figure 13-10 ATM Router Module RFC 1577 Example Network

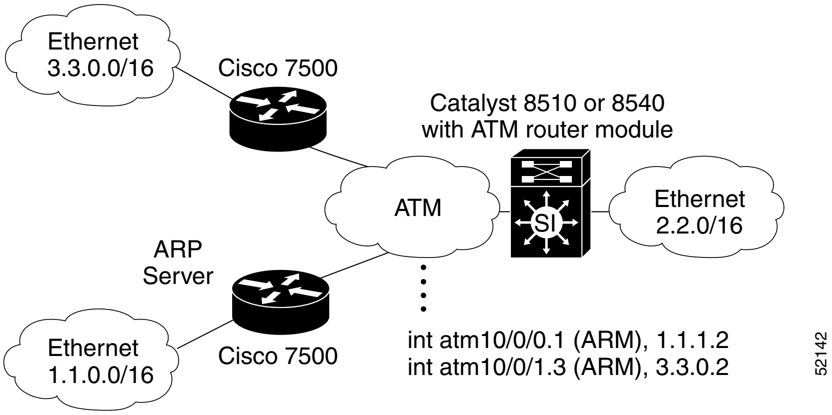

This example network is used in the following troubleshooting steps.

Follow these steps to troubleshoot the ATM router module, shown in Figure 13-10, configured with RFC 1577:

Step 1 Use the following commands to configure RFC 1577 for the example network shown in Figure 13-10.

```
Switch# config term
Switch(config)# interface atm 10/0/1.1 multipoint
Switch(config-if)# ip address 1.1.1.2 255.255.0.0
Switch(config-if)# atm arp-server nsap 47.009181000000009021418801.009021418801.00
Switch(config-if)# atm pvc 2 102 pd on inarp 5 interface ATM0/0/0 0 102
Switch(config-if)# exit
Switch(config)# interface atm 10/0/1.3 multipoint
Switch(config-if)# ip address 3.3.0.2 255.255.0.0
Switch(config-if)# atm arp-server nsap 47.009181000000009021418801.0050BD9B2160.0A
Switch(config-if)# atm pvc 2 103 pd on inarp 5 interface ATM0/0/1 0 103
Switch(config-if)# end
Switch#
```

Step 2 Use the **show running-config** command to confirm your configuration.

```
Switch# show running-config
Building configuration...

Current configuration:
!

!
interface ATM10/0/1.1 multipoint
 ip address 172.20.52.41 255.255.255.224
 atm arp-server nsap 47.00918100000000E04FACB401.00E04FACB401.00
 atm pvc 2 102 pd on inarp 5 interface ATM0/0/0 0 102
!
```

The following process describes troubleshooting basic connectivity problems with RFC 1577 networks. In the example shown in Figure 13-10, the Cisco 7500 router connected to Ethernet 1.1.0.0 is acting as the ARP server.

Step 1 Using the **show atm map** command, confirm that both the switch router and the Cisco 7500 router connected to Ethernet 3.3.0.0 have connections to the ARP server. If they are connected, they can ask the ARP server for an IP-to-ATM address resolution.

Step 2 To test the switch router configuration, use the **debug atm arp** command on the switch router, to see whether it is sending out an ARP request to the ARP server router.

Step 3 From the Cisco 7500 router connected to Ethernet 1.1.0.0 (and acting as the ARP server), confirm it is receiving the ARP request and responding to it with a positive acknowledgment by using the **debug atm arp** command.

Step 4 On the ARP server, use the **debug atm arp** command to confirm it is receiving the ARP requests and responding with a positive acknowledgment.

When the IP-to-ATM address is resolved, the Cisco 7500 router connected to Ethernet 3.3.0.0 should be able to make a call to the ATM address of the switch router ATM router module. If the Cisco 7500 router still can not connect to the switch router ATM router module, the problem is probably the call setup. Refer to the "Troubleshooting RFC 1483 on ATM Router Module Connections" section on page 13-21.

If you determine that RFC 1577 on the interface is configured incorrectly, refer to the "Configuring IP over ATM" chapter in the *ATM Switch Router Software Configuration Guide*.

Troubleshooting OAM on ATM Router Module Connections

OAM performs fault management and performance management functions at the ATM management (M)-plane layer.

Note Current OAM implementation supports only the fault management function, which includes connectivity verification and alarm surveillance.

The ATM switch router has full support for the following ATM OAM cell flows:

- F4 flows—OAM information flows between network elements (NEs) used within virtual paths, to report an unavailable path or a virtual path (VP) that cannot be guaranteed.
- F5 flows—OAM information flows between NEs used within virtual connections, to report degraded virtual channel (VC) performance such as late arriving cells, lost cells, and cell insertion problems.

Both F4 and F5 flows can be configured as either end-to-end loopback or segment-loopback and used with alarm indication signal (AIS) and remote defect indication (RDI) functions.

Troubleshooting OAM Problems on the ATM Router Module Commands

To display the ATM router module and OAM configuration, use the following commands:

Command	Purpose
show atm traffic	Displays the ATM layer traffic information for all of the ATM interfaces.
show atm vc interface atm *card/subcard/port.subinterface VPI VCI*	Displays the ATM layer connection information about the virtual connection.
show atm vc traffic interface atm *card/subcard/port VPI VCI*	Displays information about the ATM virtual connection.

To configure OAM on the ATM router module, refer to the "Configuring Operation, Administration, and Maintenance" chapter in the *ATM Switch Router Software Configuration Guide*.

Figure 13-6 is an example network of a switch router with an ATM router module having OAM configured on the ATM router module and the connecting ATM interfaces.

Figure 13-11 ATM Router Module OAM Example Network

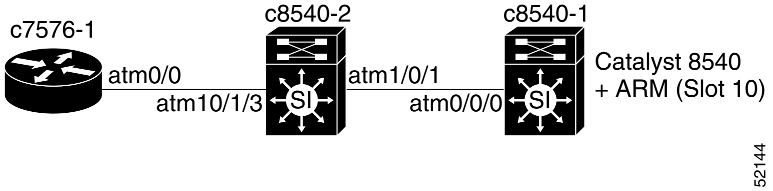

This example network is used in the following troubleshooting steps.

Follow these steps to configure and troubleshoot the ATM router module with OAM configured (see the example network shown in Figure 13-11):

Step 1 At switch router c8540-1, use the following commands to configure the ATM interface connected to switch router c8540-2, to perform OAM fault management.

```
  c8540-1(config)# interface atm 0/0/0.5 point-to-point
  c8540-1(config-subif)# atm pvc 2 105 int atm 0/0/0 0 105
  c8540-1(config-subif)# atm oam interfaces atm 0/0/0 0 105 end-loopback
➔ % OAM: Connection level end to end loopback is enabled

  c8540-1(config-subif)#
```

In this example, the system message “% OAM: Connection level end to end loopback is enabled” appears, indicating the subinterface is correctly enabled.

Step 2 At switch router c8540-1, use the **show running-config** command to confirm configuration of the ATM interface connected to switch router c8540-2 and the ATM route module interface.

```
  c8540-1# show running-config
  Building configuration...

  Current configuration:
  !
  .
  (Information Deleted)
  .
  interface ATM0/0/0
   no ip address
   no ip route-cache distributed
   no atm ilmi-keepalive
  !
➔ interface ATM0/0/0.5 point-to-point
   ip address 11.1.5.2 255.255.255.252
   pvc 0/105
    oam-pvc manage 5
    encapsulation aal5snap
   !
  .
  (Information Deleted)
  .
  interface ATM10/0/1
   no ip address
   no ip directed-broadcast
   logging event subif-link-status
   arp timeout 900
  !
➔ interface ATM10/0/1.5 point-to-point
   ip address 11.1.5.1 255.255.255.252
   no ip directed-broadcast
   atm pvc 2 105 pd on  interface  ATM0/0/0 0 105
   atm oam interface  ATM0/0/0 0 105 end-loopback
  !
```

Step 3 At router c7576-1, use the **show atm traffic** command to confirm the F5 OAM cells are being received.

```
  c7576-1# show atm traffic
  .
  (Information Deleted)
  .
➔ 348 OAM cells received
  F5 InEndloop: 348, […]
  586 OAM cells sent
  F5 OutEndloop: 586, […]
```

Step 4 At router c7576-1, use the **show atm traffic** command again, to confirm the F5 OAM cells are incrementing.

```
  c7576-1# show atm traffic
  .
  (Information Deleted)
  .
➔ 397 OAM cells received
  F5 InEndloop: 397, […]
  635 OAM cells sent
  F5 OutEndloop: 635, […]
```

Step 5 At router c8540-1, use the **show atm vc interfaces atm** command with the VPI and VCI parameters to confirm the ATM router module subinterface status and configuration.

```
  c8540-1# show atm vc interfaces atm 10/0/1.5 2 105

  Interface: ATM10/0/1, Type: arm_port
  VPI = 2  VCI = 105
➔ Status: UP
➔ Time-since-last-status-change: 00:34:41
  Connection-type: PVC
  .
  (Information Deleted)
  .
➔ Cross-connect OAM-configuration: End-to-end-loopback-on
➔ Cross-connect OAM-state:  OAM-Up
  OAM-Loopback-Tx-Interval: 5
```

Step 6 Check the Status field. It should be UP.

Step 7 Check the Time-since-last-status-change field. It should indicate the time since you enabled OAM on the subinterface.

Step 8 Check the Cross-connect OAM-configuration field. It should indicate End-to-end-loopback-on.

Step 9 Check the Cross-connect OAM-state field. It should indicate OAM-Up.

Step 10 To demonstrate an OAM failure, delete the PVC between switch router c8540-1 and switch router c8540-2 from the switch router c8540-2 end of the PVC.

```
c8540-2(config-if)# no atm pvc 0 105 int atm 1/0/1 0 105
```

Step 11 At router c7576-1, a system error message appears as in the following:

```
➔ c7576-1#00:43:55: %LINEPROTO-5-UPDOWN: Line protocol on Interface ATM0/0/0.5, changed
  state to down
```

The system error message "%LINEPROTO-5-UPDOWN: Line protocol on Interface ATM0/0/0.5, changed state to down" indicates the subinterface ATM 0/0/0.5 has changed status to down because the PCV was disabled on switch router c8540-2.

Step 12 At router c7576-1, use the **show atm traffic** command to check the OAM cells received and sent.

```
  c7576-1# show atm traffic
  .
  (Information Deleted)
  .

➔ 442 OAM cells received
➔ 708 OAM cells sent
```

Step 13 At router c7576-1, use the **show atm traffic** command a second time, and check the number of OAM cells received and sent.

```
  c7576-1# show atm traffic
  .
  (Information Deleted)
  .

➔ 442 OAM cells received
➔ 734 OAM cells sent
```

Step 14 Check the number before the OAM cells received field. Since the number of OAM cells received has not incremented since the previous display, this confirms the connection is down and the OAM cells are sent but not received.

Step 15 At router c8540-1, use the **show atm vc interfaces atm** command with the VPI and VCI parameters to confirm the ATM router module subinterface OAM status and configuration.

```
  c8540-1# show atm vc interfaces atm 10/0/1.5 2 105

  Interface: ATM10/0/1, Type: arm_port
  VPI = 2  VCI = 105
➔ Status: UP
➔ Time-since-last-status-change: 00:43:05
➔ Cross-connect OAM-configuration: End-to-end-loopback-on
➔ Cross-connect OAM-state:  OAM-Up End-to-end-loopback-failed
  OAM-Loopback-Tx-Interval: 5
```

Step 16 Check the Status field. It should be UP.

Step 17 Check the Time-since-last-status-change field. It should indicate the time since you enabled OAM on the subinterface.

Step 18 Check the Cross-connect OAM-configuration field. It should indicate End-to-end-loopback-on.

Step 19 Check the Cross-connect OAM-state field. It should indicate OAM-Up, but that the End-to-end-loopback-failed because the PCV was disabled on switch router c8540-2.

The effect of OAM failure on an interface or subinterface is as follows:

- On the ATM router module the interface or subinterface status and VC remain UP, causing the following:
  - the VC remains in switch fabric
  - OAM loopback cells are still sent, but not received
  - The routing or bridging layer is informed to remove routes or MAC addresses learned over the affected VC
- On the router, the subinterface status and VC status change to DOWN

If you determine that the OAM interface is configured incorrectly, refer to the “Configuring Operation, Administration, and Maintenance” chapter in the *ATM Switch Router Software Configuration Guide*.

PART 4

Appendixes

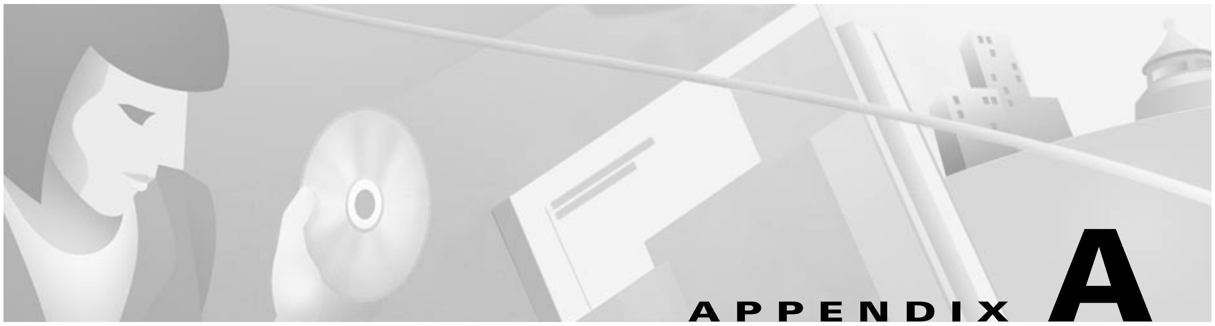

APPENDIX A

Debugging a Switch Router

This appendix provides an overview of the **debug** commands that might be helpful when troubleshooting your switch router.

This appendix consists of the following sections:

- Using the Debug Interface, page A-2
- Enabling Message Logging, page A-6

Use **debug** commands to isolate problems, not to monitor normal network operation. Because the high overhead of **debug** commands can disrupt switch router operation, use **debug** commands only when you are looking for specific types of traffic or problems and have narrowed your problems to a likely subset of causes.

Output formats vary with each **debug** command. Some generate a single line of output per packet, and others generate multiple lines of output per packet. Some generate large amounts of output, and others generate only occasional output. Some generate lines of text, and others generate information in field format.

Follow these steps to minimize the negative impact of using **debug** commands:

Step 1 Use the **no logging console** global configuration command on your switch router. This command disables all logging to the console terminal.

Step 2 Use Telnet to connect to a switch router port, and enter the **enable** EXEC command.

Step 3 Use the **terminal monitor** command to copy **debug** command output and system error messages to your current terminal display.

As a result, you can view **debug** command output remotely, without being connected through the console port.

Because the console port no longer has to generate character-by-character processor interrupts, following this procedure minimizes the load created by using **debug** commands.

If you intend to keep the output of the **debug** command, spool the output to a file. The procedure for setting up such a debug output file is described in the *Cisco IOS Debug Command Reference* publication. This publication provides complete details regarding the function and output of **debug** commands, and includes specific **debug** commands that are useful when troubleshooting specific problems.

Using the Debug Interface

This section explains how to diagnose and resolve internetworking problems by using **debug** commands, and covers the following topics:

- Entering debug Commands
- Using the debug ? Command
- Using the debug all Command: Warning
- Generating debug Command Output
- Redirecting Debugging and Error Message Output

Caution Because debugging output is assigned a high priority in the multiservice route processor, it can render the system unusable. For this reason, use **debug** commands only to troubleshoot specific problems or during troubleshooting sessions with customer support. Moreover, it is best to use **debug** commands during periods of lower network traffic and fewer users. Debugging during these periods decreases the likelihood that the increased overhead of **debug** command processing will affect system use.

Entering debug Commands

All **debug** commands are entered in privileged EXEC mode, and most **debug** commands take no arguments. For example, to enable the **debug atm packet** command, enter the following in privileged EXEC mode at the command line:

```
Switch# debug atm packet
ATM packets debugging is on
Displaying all ATM packets
```

To display the status of each debugging option, enter the **show debugging** command in privileged EXEC mode:

```
Switch# show debugging
Generic ATM:
  ATM packets debugging is on
Switch#
```

To turn off the **debug atm packet** command in privileged EXEC mode, enter the **no** form of the command at the command line:

```
Switch# no debug atm packet
ATM packets debugging is off
Switch#
```

Using the debug ? Command

To list and see a brief description of all the **debug** command options, at the command line enter the **debug ?** command in privileged EXEC mode, as shown in the following example:

```
Switch# debug ?
  aaa                 AAA Authentication, Authorization and Accounting
  all                 Enable all debugging
  arp                 IP ARP and HP Probe transactions
  async               Async interface information
  atm                 ATM interface packets
  broadcast           MAC broadcast packets
  callback            Callback activity
  cdp                 CDP information
  ces-iwf             CES-IWF Info
  chat                Chat scripts activity
  compress            COMPRESS traffic
  custom-queue        Custom output queueing
  dialer              Dial on Demand
  dnsix               Dnsix information
  domain              Domain Name System
  eigrp               EIGRP Protocol information
  ethernet-interface  Ethernet network interface events
  filesys             File system information
  general             Rhino General Debug
  ip                  IP information
  ipc                 Interprocess communications debugging
  lane                LAN Emulation
  list                Set interface or/and access list for the next debug command
  modem               Modem control/process activation
  ntp                 NTP information
  nvram               Debug NVRAM behavior
  packet              Log unknown packets
  ports               Rhino Ports Info
  ppp                 PPP (Point to Point Protocol) information
  priority            Priority output queueing
  probe               HP Probe Proxy Requests
  rif                 RIF cache transactions
  serial              Serial interface information
  snmp                SNMP information
  sscop               SSCOP
  standby             Hot standby protocol
  tacacs              TACACS authentication and authorization
  tag-switching       Tag Switching
  telnet              Incoming telnet connections
  tftp                TFTP packets
  token               Token Ring information
  tunnel              Generic Tunnel Interface
Switch#
```

Note Not all **debug** commands listed in the output of the **debug ?** command are described in this document. Commands included here assist you in diagnosing network problems.

Using the debug all Command: Warning

To enable all system diagnostics, you can enter the **debug all** command in privileged EXEC mode at the command line.

Caution Because debugging output takes priority over other network traffic, and because the **debug all** command generates more output than any other **debug** command, it can severely diminish switch router performance or render the switch router unusable. Use more specific **debug** commands.

The **no debug all** command turns off all diagnostic output. Using this command is a convenient way to ensure that you have not accidentally left any **debug** commands turned on.

Generating debug Command Output

Enabling a **debug** command can result in output similar to the following example for the **debug lane client all** command:

```
Switch# debug lane client all
Switch#
LEC ATM13/0/0.1: action A_RESEND_JOIN_REQ
LEC ATM13/0/0.1: sending LANE_JOIN_REQ on VCD 303
LEC ATM13/0/0.1:   Status               0
LEC ATM13/0/0.1:   LECID                0
LEC ATM13/0/0.1:   SRC MAC address      00e0.4fac.b402
LEC ATM13/0/0.1:   SRC ATM address      47.00918100000000E04FACB401.00E04FACB402
.01
LEC ATM13/0/0.1:   LAN Type     1
LEC ATM13/0/0.1:   Frame size   1
LEC ATM13/0/0.1:   LAN Name     eng_elan
LEC ATM13/0/0.1:   LAN Name size        8
LEC ATM13/0/0.1: state JOIN_CTL_DIST_CONN event LEC_TIMER_JOIN => JOIN_CTL_DIST_
CONN
%LANE-3-NOREGILMI: ATM13/0/0 LECS cannot register 47.00918100000000E04FACB401.00
E04FACB405.00 with ILMI
LEC ATM13/0/0.1: action A_RESEND_JOIN_REQ
LEC ATM13/0/0.1: action A_TEARDOWN_LEC
LEC ATM13/0/0.1: sending RELEASE
LEC ATM13/0/0.1:   callid               0x60EBB48C
LEC ATM13/0/0.1:   cause code   31
LEC ATM13/0/0.1: sending CANCEL
LEC ATM13/0/0.1:   ATM address  47.00918100000000E04FACB401.00E04FACB402.01
LEC ATM13/0/0.1: state JOIN_CTL_DIST_CONN event LEC_TIMER_JOIN => TERMINATING
LEC ATM13/0/0.1: received RELEASE_COMPLETE
LEC ATM13/0/0.1:   callid               0x60EBB48C
LEC ATM13/0/0.1:   cause code   31
LEC ATM13/0/0.1: action A_PROCESS_TERM_REL_COMP
LEC ATM13/0/0.1: state TERMINATING event LEC_SIG_RELEASE_COMP => IDLE
LEC ATM13/0/0.1: received CANCEL
LEC ATM13/0/0.1: state IDLE event LEC_SIG_CANCEL => IDLE
LEC ATM13/0/0.1: state IDLE event LEC_CTL_ILMI_SET_RSP_POS => IDLE
```

```
%LANE-3-NOREGILMI: ATM13/0/0 LECS cannot register 47.00918100000000E04FACB401.00
E04FACB405.00 with ILMI
%LANE-3-NOREGILMI: ATM13/0/0 LECS cannot register 47.00918100000000E04FACB401.00
E04FACB405.00 with ILMI
Switch#
%LANE-3-NOREGILMI: ATM13/0/0 LECS cannot register 47.00918100000000E04FACB401.00
E04FACB405.00 with ILMI
Switch# no debug lane client all
All possible debugging has been turned off
Switch#
```

The switch router continues to generate such output until you enter the corresponding **no debug** command (in this case, **no debug lane client all**).

If you enable a **debug** command and no output is displayed, consider the following possibilities:

- The switch router might not be properly configured to generate the type of traffic you want to monitor. Use the **show running-config** command to check the configuration of the switch router.
- Even a properly configured switch router might not generate the type of traffic you want to monitor when debugging is enabled. Depending on the protocol you are debugging, you can use commands such as the **ping atm interface atm** command to generate network traffic.

Redirecting Debugging and Error Message Output

By default, the network server sends the output from **debug** commands and system error messages to the console terminal. If you use this default, monitor debugging output with a virtual terminal connection, rather than the console port.

To redirect debugging output, use the variations of the **logging** command within configuration mode as described in the following sections.

Possible destinations include the console terminal, virtual terminals, internal buffer, and UNIX hosts running a syslog server. The syslog format is compatible with 4.3 Berkeley Standard Distribution (BSD) UNIX and its derivatives.

Note Be aware that the debugging destination you use affects system overhead. Logging messages to the console produces very high overhead, whereas logging messages to a virtual terminal produces less overhead. Logging messages to a syslog server produces even less, and logging them to an internal buffer produces the least overhead of any method.

To configure message logging, you need to be in configuration command mode. To enter this mode, use the **configure terminal** command at the EXEC prompt.

Enabling Message Logging

By logging messages you can redirect debug output to a file, memory, or a remote host connection.

This section contains the following:

- Setting the Message Logging Levels
- Limiting the Types of Logging Messages Sent to the Console
- Logging Messages to an Internal Buffer
- Limiting the Types of Logging Messages Sent to Another Monitor
- Logging Messages to a UNIX Syslog Server
- Limiting Messages to a Syslog Server
- Displaying the Logging Configuration

To configure message logging, use any of the following global configuration commands. Use the **no** form of the command to assign the default value:

Command	Purpose
logging on	Enables message logging.
logging *hostname* | *ip-address*	Enables logging to a sys server host.
logging buffered [*buffer-size*] [**0 - 7** | **alerts** | **critical** | **debugging** | **emergencies** | **errors** | **informational** | **notifications** | **warnings**]	Enables logging to an internal buffer instead of writing to the console, and configures the buffer size.
logging console [**0 - 7** | **alerts** | **critical** | **debugging** | **emergencies** | **errors** | **informational** | **notifications** | **warnings**]	Enables logging to the console, and configures the level displayed.
logging facility {**auth** | **cron** | **daemon** | **kern** | **local0** - **local7** | **lpr** | **mail** | **news** | **sys9** - **sys14** | **syslog** | **user** | **uucp**}	Enables logging to a specific facility.
logging monitor [**0 - 7** | **alerts** | **critical** | **debugging** | **emergencies** | **errors** | **informational** | **notifications** | **warnings**]	Enables logging to a monitor, and configures the level displayed.
logging source-interface {{**atm** | **atm-p**} *card/subcard/port* | **async 1** | **bvi** *number* | **dialer** | **ethernet** *number* | **lex** *number* | **loopback** *number* | **null 0** | **tunnel** *number* | **virtual-template** *number* | **virtual-tokenring** *number* | **vlan** *number*}	Enables logging from a specific source interface.
logging trap [**0 - 7** | **alerts** | **critical** | **debugging** | **emergencies** | **errors** | **informational** | **notifications** | **warnings**]	Enables logging of SNMP[1] traps.

[1] SNMP = Simple Network Management Protocol

Examples

To enable message logging to all supported destinations other than the console, enter the **logging on** command as in the following example:

```
Switch(config)# logging on
```

To direct logging to the console terminal only and disable logging output to other destinations, enter the **no logging on** command as in the following example:

```
Switch(config)# no logging on
```

Setting the Message Logging Levels

You can set the logging levels when you log messages to the following devices:

- Console
- Monitor
- Syslog server

Table A-1 lists and briefly describes the logging levels and corresponding keywords you can use to set the logging levels for these types of messages. The highest level of message is 0, *emergencies*. The lowest level is 7, *debugging*, which also displays the largest number of messages.

Table A-1 Message Logging Keywords and Levels

Level	Keyword	Description	Syslog Definition
0	**emergencies**	System is unusable.	LOG_EMERG
1	**alerts**	Immediate action is needed.	LOG_ALERT
2	**critical**	Critical conditions exist.	LOG_CRIT
3	**errors**	Error conditions exist.	LOG_ERR
4	**warnings**	Warning conditions exist.	LOG_WARNING
5	**notifications**	Normal, but significant, conditions exist.	LOG_NOTICE
6	**informational**	Informational messages.	LOG_INFO
7	**debugging**	Debugging messages.	LOG_DEBUG

For information about limiting these messages, refer to the sections that follow.

Limiting the Types of Logging Messages Sent to the Console

The **logging console** command limits the logging messages displayed on the console terminal to messages up to and including the specified severity level, which is specified by the *level* argument.

The *level* argument is one of the keywords listed in Table A-1. Keywords are listed in order from the most severe level to the least severe level.

The **no logging console** command disables logging to the console terminal.

Example

The following example sets console logging of messages at the **debugging** level, which is the least severe level and displays all logging messages:

```
Switch(config)# logging console debugging
```

Logging Messages to an Internal Buffer

The **logging buffered** command copies logging messages to an internal buffer instead of writing them to the console terminal. The buffer is circular in nature, so newer messages overwrite older messages. To display the messages that are logged in the buffer, use the **show logging** privileged EXEC command. The first message displayed is the oldest message in the buffer. The **no logging buffered** command cancels the use of the buffer and writes messages to the console terminal, which is the default setting.

Example

The following example copies logging messages to buffered memory and sets the memory buffer space as 10,000 bytes:

```
Switch(config)# logging buffered 10000
```

Limiting the Types of Logging Messages Sent to Another Monitor

The **logging monitor** command limits the logging messages that are displayed on terminal lines—other than the console line—to messages with a level up to and including the specified *level* argument. The *level* argument is one of the keywords listed in Table A-1. To display logging messages on a terminal (virtual console), use the **terminal monitor** privileged EXEC command.

The **no logging monitor** command disables logging to terminal lines other than the console line.

Example

The following example sets the level of messages displayed on monitors other than the console to **notification** level:

```
Switch(config)# logging monitor notification
```

Logging Messages to a UNIX Syslog Server

The **logging** command identifies a syslog server host to receive logging messages. The *ip-address* argument is the IP address of the host. By issuing this command more than once, you build a list of syslog servers that receive logging messages.

The **no logging** command deletes the syslog server with the specified address from the list of syslogs.

Example

The following example configures syslog server host diablo.cisco.com to receive the logging messages:

```
Switch(config)# logging diablo.cisco.com
```

Example of Setting Up a UNIX Syslog Daemon

To set up the syslog daemon on a 4.3 BSD UNIX system, include a line such as the following in the file */etc/syslog.conf*:

```
local7.debugging /usr/adm/logs/tiplog
```

The **local7** keyword specifies the logging facility to be used. The **debugging** keyword specifies the syslog level. See Table A-1 for other keywords that can be listed.

The UNIX system sends messages at or above this level to the specified file, in this case */usr/adm/logs/tiplog*. The file must already exist, and the syslog daemon must have permission to write to the specific file. For System V UNIX systems, the line should read as follows:

```
local7.debug /usr/admin/logs/cisco.log
```

Limiting Messages to a Syslog Server

The **logging trap** command limits the logging messages sent to syslog servers to messages with a level up to and including the specified *level* argument. The *level* argument is one of the keywords listed in Table A-1.

To send logging messages to a syslog server, specify its host address with the **logging** command. The default trap level is the **informational** level. The **no logging trap** command disables logging to syslog servers.

Example

The following example configures traps sent to the syslog server as informational:

```
Switch(config)# logging trap informational
```

Displaying the Logging Configuration

To display the logging configuration, use the following command in user EXEC mode:

Command	Purpose
show logging	Displays the logging configuration.

Example

The following example shows the addresses and levels associated with the current logging setup. The command output also includes ancillary statistics.

```
Switch# show logging
Syslog logging: disabled (0 messages dropped, 3 flushes, 0 overruns)
    Console logging: level critical, 130 messages logged
    Monitor logging: level notifications, 0 messages logged
    Trap logging: level informational, 132 message lines logged
    Buffer logging: level debugging, 127 messages logged
Log Buffer (10000 bytes):
Switch#
```

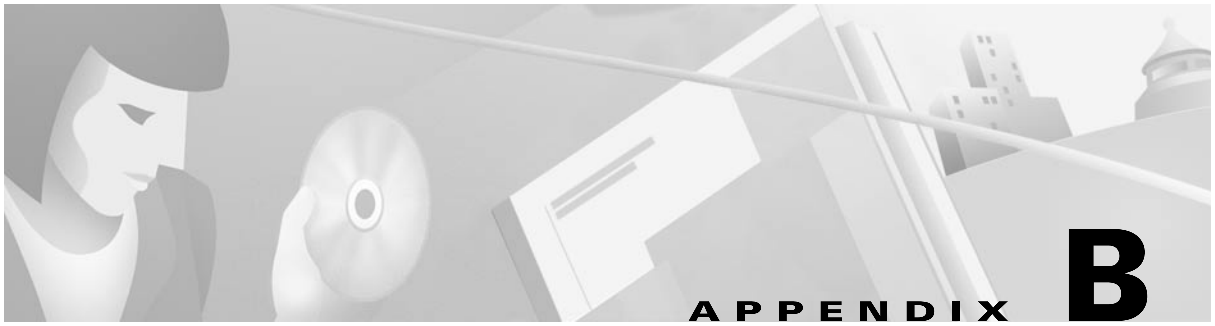

APPENDIX B

Troubleshooting TACACS+ and Recovering Passwords

This chapter describes troubleshooting information relating to security implementations and contains the following sections:

- Troubleshooting TACACS+ Problems, page B-1
- Recovering a Lost Password, page B-5

If you want detailed information about configuring and using TACACS+, refer to the *ATM Switch Router Software Configuration Guide* and *ATM Switch Router Command Reference* publications. For additional information about TACACS+, refer to the *Cisco IOS Security Configuration Guide* and *Cisco IOS Security Command Reference*.

Troubleshooting TACACS+ Problems

The following sections describe problems with TACACS+ operation and possible solutions.

Errors Unarchiving Source File

Symptom: Errors are generated when unarchiving the TACACS+ archive file (tac_plus.2.1.tar).

Table B-1 outlines possible problems and describes solutions.

Table B-1 TACACS+: Errors Unarchiving Source File

Possible Problem	Solution
Archive file was not transferred by using FTP[1] binary (image) mode	The TACACS+ archive file must be transferred by using FTP binary (image) mode. FTP the tac_plus.2.1.tar file again, using FTP binary transfer mode. From the FTP command line, enter the **image** command to set the image mode. For other FTP software, refer to your documentation for instructions on setting the image mode.
Insufficient disk space	Make sure there is sufficient disk space for the expanded tac_plus.2.1.tar file. If there is not enough space on your UNIX system, create enough free disk space to accommodate decompression of the file. TACACS+ requires about 900 KB.

[1] FTP = File Transfer Protocol

Cannot Compile Daemon

Symptom: Attempts to compile the TACACS+ daemon result in errors.

Table B-2 outlines possible problems and describes solutions.

Table B-2 TACACS+: Cannot Compile Daemon

Possible Problem	Solution
make is not in $PATH or is not installed on the UNIX machine	1. Enter the **which make** command at the UNIX prompt. If the output says “No make in $PATH...,” **make** is not in the specified path or is not installed. 2. If **make** is already installed, modify the $PATH variable to include the directory in which **make** is located. If **make** is not installed, see your system administrator for help installing it. 3. Compile the TACACS+ daemon again.
gcc is not in $PATH or is not installed correctly	1. Enter the **which gcc** command at the UNIX prompt. If the output says “No gcc in $PATH...,” **gcc** is not in the specified path or is not installed. 2. If **gcc** is already installed, modify the $PATH variable to include the directory in which **gcc** is located. If **gcc** is not installed, ask your system administrator to install it. 3. Compile the TACACS+ daemon again.
UNIX platform is commented out or is not in the Makefile	Your UNIX platform must be listed and uncommented in the Makefile for **make** to compile the TACACS+ source code properly. The Makefile is located in the tac_plus.2.1 directory. 1. Make sure that your UNIX platform is not commented out in the Makefile. 2. If your platform is not listed at all, see your system administrator for help with compiling the source code. The only supported platforms are those listed in the Makefile. 3. Compile the TACACS+ daemon again.

Daemon Is Not Up and Running

Symptom: The TACACS+ daemon is not running.

Table B-3 outlines possible problems and describes solutions.

Table B-3 TACACS+: Daemon Is Not Up and Running

Possible Problem	Solution
TACACS+ has not been launched	Launch TACACS+ with the **tac_plus -C** *configuration filename* command.
TACACS+ is not specified in the /etc/services file	1. Check the /etc/services file for the following line: `tacacs 49/tcp` 2. This line must be included in the file. If the line is not present, add the line to the file.
The **tac_plus** executable does not exist	The TACACS+ daemon cannot run if the **tac_plus** executable does not exist. 1. Check the directory where you installed tac_plus.2.1 to see if the **tac_plus** file exists. 2. If the file does not exist, use the **make tac_plus** command to compile **tac_plus**.

Daemon Does Not Run

Symptom: The TACACS+ daemon does not run when invoked.

Table B-4 outlines possible problems and describes solutions.

Table B-4 TACACS+: Daemon Does Not Run

Possible Problem	Solution

TACACS+ configuration file is not present	1. Check the directory in which you installed TACACS+ for a configuration file in the TACACS+ format.
	2. If there is no TACACS+ configuration file present and you are upgrading from XTACACS, convert your password file into a configuration file by issuing the following command:
	`unix_host% convert.pl /etc/passwd > configuration-file`
	The configuration file can have any name you want.
	3. If there is no TACACS+ configuration file present, create one by using a text editor. At a minimum, the configuration file must contain the following text:
	`user = userid {`
	`login = cleartext “passwd”`
	`}`
	The configuration file can be given any name.
	For more information, refer to the user guide located in the tac_plus.2.1 directory.

Users Cannot Connect Using TACACS+

Symptom: Users cannot log in using TACACS+. Either users cannot get the Username prompt or they get the prompt but authentication or authorization fails.

Table B-5 outlines possible problems and describes solutions.

Table B-5 TACACS+: Users Cannot Log in Using TACACS+

Possible Problem	Solution

Switch router missing minimum configuration	1. Use the **show running-config** privileged EXEC command to view the local switch router configuration. Look for the following commands:
	`aaa new-model`
	`aaa authentication login default tacacs+ enable`
	`[...]`
	`tacacs-server host name`
	`tacacs-server key key`
	where *name* is the IP address or DNS[1] host name of the TACACS+ server and *key* is the authentication and encryption key.
	2. If all of these commands are not present, add the missing commands to the configuration. If there is no key configured on the TACACS+ daemon, the **tacacs-server key** command is not necessary.

aaa authorization command is present	1. Use the **show running-config** privileged EXEC command to view the local switch router configuration. Look for an **aaa authorization exec tacacs+** global configuration command entry. 2. If the command is present, remove it from the configuration by using the **no** version of the command.
PPP[2] not functioning correctly	If PPP is not functioning properly, problems will occur when using TACACS+. Use the **debug ppp negotiation** privileged EXEC command to see if both sides are communicating. For information on configuring PPP, refer to the *Cisco IOS Dial Solutions Configuration Guide: Terminal Services* and *Cisco IOS Dial Solutions Command Reference* publications.

PAP[3] is misconfigured	1. Use the **show running-config** privileged EXEC command to make sure your configuration includes the following global configuration command:
	`aaa authentication ppp default if-needed tacacs+`
	2. If the command is not present, add it to the configuration.
	3. In addition, check the configuration of the async interface being used. The interface must have the following commands configured:
	`encapsulation ppp`
	`ppp authentication pap`
	4. If these commands are not present, add them to the interface configuration.

Table B-5 TACACS+: Users Cannot Log in Using TACACS+ (continued)

Possible Problem	Solution

CHAP[4] is misconfigured	**1.** Use the **show running-config** privileged EXEC command to make sure your configuration includes the following global configuration command:
	`aaa authentication ppp default if-needed tacacs+`
	2. If the command is not present, add it to the configuration.
	3. In addition, check the configuration of the async interface being used. The interface must have the following commands configured:
	`encapsulation ppp`
	`ppp authentication chap`
	4. If these commands are not present, add them to the interface configuration.
	5. Make sure your daemon configuration file, located in the tac_plus.2.1 directory, includes one of the following lines, as appropriate:
	`chap = cleartext` *`password`*
	or
	`global = cleartext` *`password`*

Username and password are not in the /etc/passwd file	**1.** Check to make sure that the appropriate username and password pairs are contained in the /etc/passwd file. **2.** If the appropriate users are not specified, generate a new user with the correct username and password, using the **add user** command.
There is no TCP connection to the TACACS+ daemon	**1.** From the switch router, try to connect to port 49 by using Telnet on the TACACS+ daemon. **2.** If the attempt to connect via Telnet is unsuccessful, make sure the daemon is running. For more information, see the "Daemon Is Not Up and Running" section on page B-3. **3.** If the daemon is running but the Telnet connection times out, check the IP connectivity.

[1] DNS = Domain Naming System

[2] PPP = Point-to-Point Protocol

[3] PAP = Password Authentication Protocol

[4] CHAP = Challenge Handshake Authentication Protocol

Recovering a Lost Password

This section describes the procedure to recover a lost login or to enable a password. The procedure differs depending on the platform and the software used, but in all cases, password recovery requires that the switch router be taken out of operation and powered down.

If you need to perform the following procedure, make certain that there are secondary systems that can temporarily serve the functions of the switch router undergoing the procedure. If this is not possible, advise all potential users and, if possible, perform the procedure during low-use hours.

Note Make a note of your password, and store it in a secure place.

All of the procedures for recovering lost passwords depend on changing the configuration register of the switch router. This is done by reconfiguring the switch router software.

More recent Cisco platforms run from Flash memory or are netbooted from a network server and can ignore the contents of nonvolatile random-access memory (NVRAM) when booting. By ignoring the contents of NVRAM, you can bypass the configuration file (which contains the passwords) and gain complete access to the switch router. You can then recover the lost password or configure a new one.

Note If your password is encrypted, you cannot recover it. You must configure a new password.

Follow these steps to recover a password:

Step 1 Beginning in the privileged executive mode, enter the **show version** command and the configuration register value. The default value is 0x2102.

Step 2 Power cycle the switch router.

Step 3 Within 60 seconds of turning the switch router On, press the **Break** key sequence or send a break signal, which is usually **^]**. If you do not see the `>` prompt with no switch router name, the terminal is not sending the correct **Break** signal. In that case, check the terminal or terminal emulation setup.

Step 4 Enter the **confreg** command at the `>` prompt.

Step 5 Answer **yes** to the `Do you wish to change configuration [y/n]?` prompt.

Step 6 Answer **no** to all the questions that appear until you reach the `Ignore system config info [y/n]` prompt. Answer **yes**.

Step 7 Answer **no** to the remaining questions until you reach the `Change boot characteristics [y/n]?` prompt. Answer **yes**.

Step 8 At the enter to boot: prompt, enter **2**.

Step 9 Answer **no** to the Do you wish to change configuration `[y/n]?` prompt.

Step 10 Enter the **reset** command at the `rommon>` prompt.

Step 11 Enter the **enable** command at the `Switch>` prompt. You'll be in enable mode and see the `Switch#` prompt.

Step 12 Enter the **show startup-config** command to view your password.

Step 13 If your password is clear text, proceed to Step 16.
or
If your password is encrypted, continue with Step 14.

Step 14 If your password is encrypted, enter the **configure memory** command to copy the NVRAM into memory.

Step 15 Enter the **copy running-config startup-config** command.

Step 16 Enter the **configure terminal** command.

Step 17 Enter the **enable secret** *password* command.

Step 18 Enter the **config-registe**r *value* command, where *value* is whatever value you entered in Step 1.

Step 19 Enter the **exit** command to exit configuration mode.

Step 20 Enter the **copy running-config startup-config** command.

Step 21 Enter the **reload** command at the prompt.

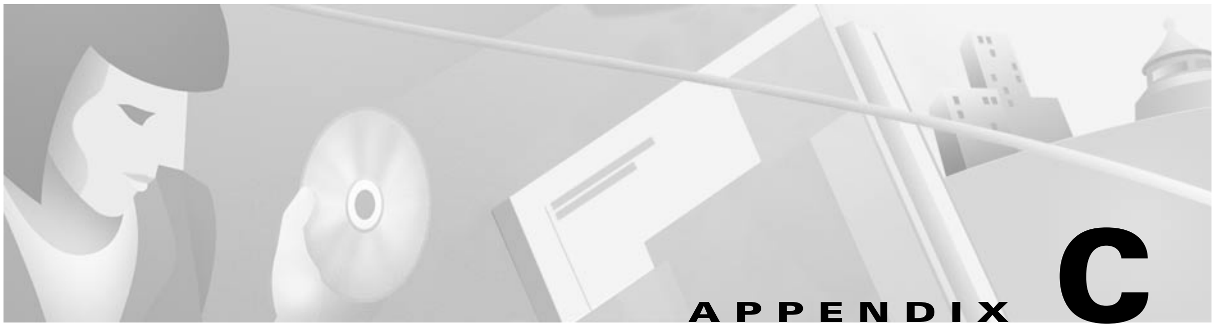

APPENDIX **C**

ATM Cell Structures

This appendix describes the various ATM cell types and their configurations and includes the following sections:

- Formats of the ATM Cell Header, page C-1
- OAM Cell Structure, page C-3
- Generic Identifier Transport IE Used by Signalling, page C-4
- LANE Data Frame, page C-5

Formats of the ATM Cell Header

The ATM standards groups have defined two header formats. The User-Network Interface (UNI) header format is defined by the UNI specification, and the Network-to-Network Interface (NNI) header format is defined by the NNI specification.

The UNI specification defines communications between ATM endpoints (such as workstations and routers) and switch routers in private ATM networks. The format of the UNI cell header is shown in Figure C-1.

Figure C-1 UNI Header Format

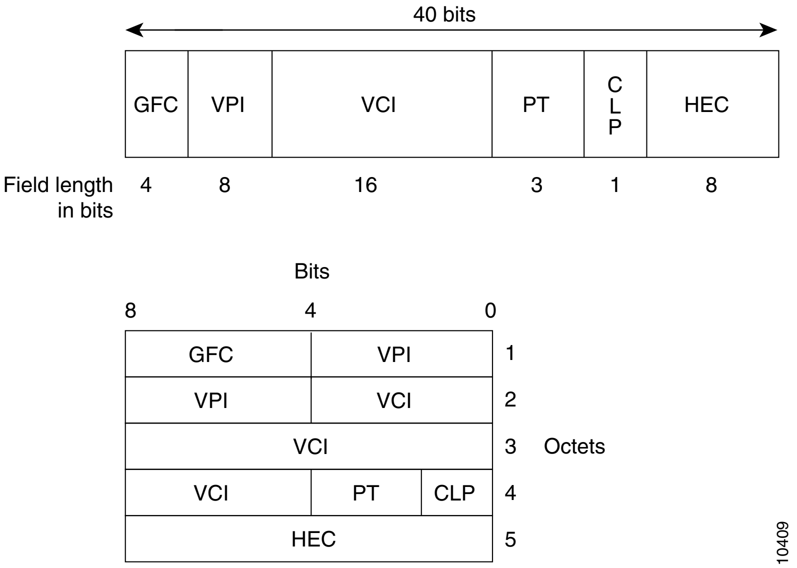

The UNI header consists of the following fields:

- GFC—4 bits of *generic flow control* that are used to provide local functions, such as identifying multiple stations that share a single ATM interface. The GFC field is typically not used and is set to a default value.
- VPI—8 bits of *virtual path identifier* that is used, in conjunction with the VCI, to identify the next destination of a cell as it passes through a series of switch routers on its way to its destination.
- VCI—16 bits of *virtual channel identifier* that is used, in conjunction with the VPI, to identify the next destination of a cell as it passes through a series of switch routers on its way to its destination.
- PT—3 bits of *payload type*. The first bit indicates whether the cell contains user data or control data. If the cell contains user data, the second bit indicates congestion, and the third bit indicates whether the cell is the last in a series of cells that represent a single AAL5 frame.
- CLP—1 bit of *congestion loss priority* that indicates whether the cell should be discarded if it encounters extreme congestion as it moves through the network.
- HEC—8 bits of *header error control* that are a checksum calculated only on the header itself.

The NNI specification defines communications between switch routers. The format of the NNI header is shown in Figure C-2.

Figure C-2 NNI Header Format

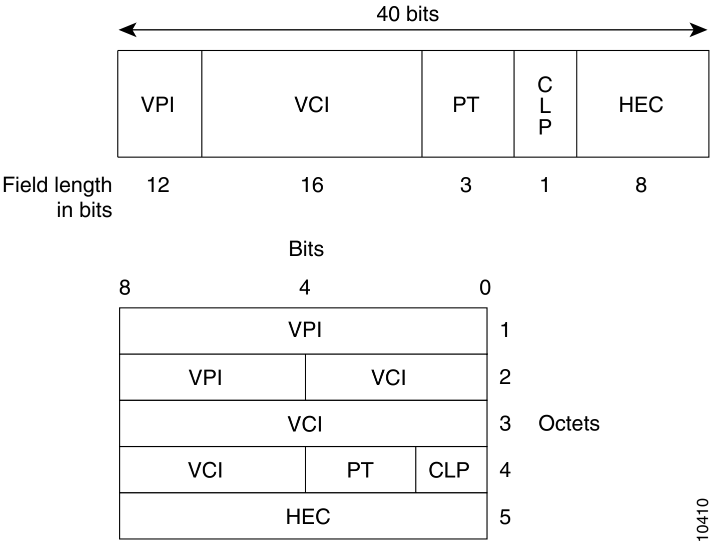

The GFC field is not present in the format of the NNI header. Instead, the VPI field occupies the first 12 bits, which allows switch routers to assign larger VPI values. With that exception, the format of the NNI header is identical to the format of the UNI header.

OAM Cell Structure

Operation, Administration, and Maintenance (OAM) performs standard loopback (end-to-end or segment) and fault detection and notification (alarm indication signal [AIS] and remote defect identification [RDI]) for each connection. It also maintains a group of timers for the OAM functions. When there is an OAM state change such as loopback failure, OAM software notifies the connection management software. You can enable or disable OAM operation for the following switch router components:

- The entire switch router
- A specific ATM interface
- Each ATM connection

Figure C-3 shows the format of the OAM loopback cell.

Figure C-3 OAM Cell Structure

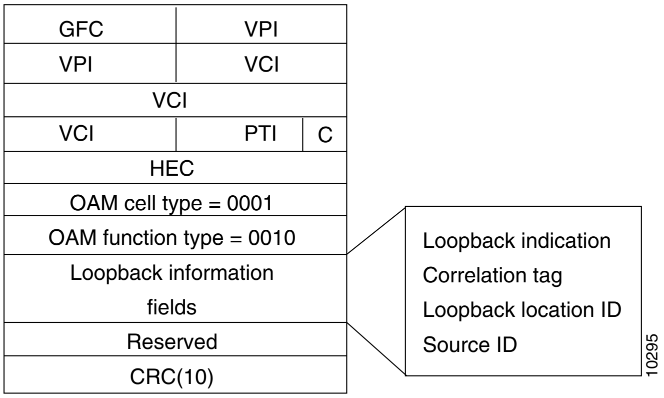

The OAM cell structure has the following features:~

- OAM cell type is coded as 0001.
- OAM function type is coded as 0010.
- 350 bits that are specific to the OAM type are divided into the following elements:
  - Loopback indicator—A bit that is set to 1 before the cell is looped back. The loopback node then sets the bit to 0, indicating it has been looped back.
  - Correlation tag—Identifies (correlates) related OAM cells within the same connection.
  - Loopback location ID—An optional field that identifies the site that is to loopback the cell.
  - Source ID—An optional field that identifies the site generating the cell.

Generic Identifier Transport IE Used by Signalling

The generic identifier transport information element (IE) is used by signalling to carry an identifier between two users.

Figure C-4 shows the format of the generic signalling IE.

Figure C-4 Generic Identifier Transport IE Used by Signalling

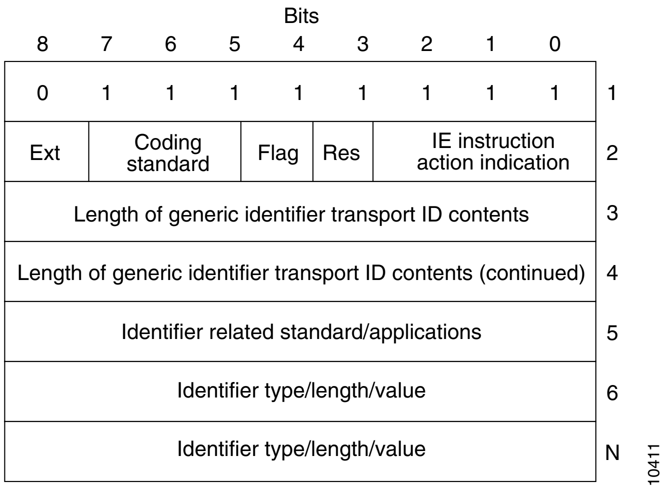

Bits 8	7	6	5	4	3	2	1	0	
0	1	1	1	1	1	1	1	1	1
Ext	Coding standard		Flag	Res	IE instruction action indication				2
Length of generic identifier transport ID contents									3
Length of generic identifier transport ID contents (continued)									4
Identifier related standard/applications									5
Identifier type/length/value									6
Identifier type/length/value									N

The generic identifier transport IE used by signalling has the following fields:

- Generic identifier transport information IE.
- Ext.
- Coding standard.
- Flag.
- Reserved.
- IE instruction action Indication.
- Length of generic indentifier transport IE.
- Identifier related standard/application—Each application requiring a different set or structure of identifiers (coded in octet 6 and possibly in subsequent octet groups) should use a different value of octet 5.
- Identifier type—This value is independent of the identifier related standard/application field, octet 5. The maximum length is 20 octets.
- Identifier length—A binary number indicating the length in octets of the identifier code in the subsequent octets of the octet group.
- Identifier value—Value of an identifier according to the recommendation or the standard identifier in octet 5.

LANE Data Frame

The LAN emulation data frame for Ethernet is based on ISO 8802.3/CSMA-CD (IEEE 802.3) and is used to provide connectivity between ATM attached end systems and LAN attached stations.

Figure C-5 shows the format of the LANE data frame.

Figure C-5 LANE Data Frame Format for IEEE 802.3/Ethernet

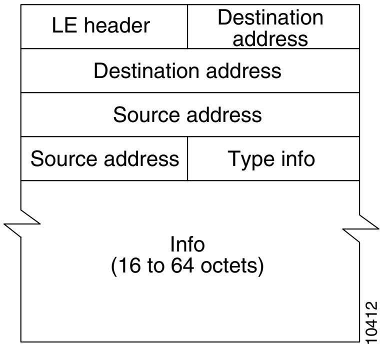

The LANE data frame has the following fields:

- LE header—Contains either the LAN emulation client identifier value, the sending client, or X'0000'.
- Destination address.
- Source address.
- Type information—Logical link control (LLC) data frames whose total length, including the LLC field and data, but not including padding required to meet minimum data frame length, is less than 1536 (X”0600”). It must be encoded by placing the length value in the type/length field. LLC data frames longer than the maximum must be encoded by placing the value 0 in the type/length field.
- Information—Encapsulated Ethernet data.

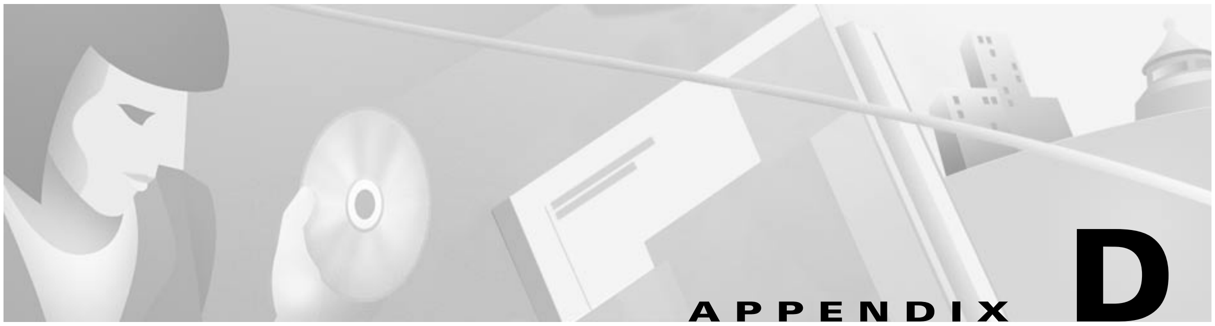

APPENDIX D

Creating a Core Dump

If the switch router fails, it is sometimes useful to get a full copy of the memory image, called a *core dump*, to identify the cause of the failure. Core dumps are generally only useful to your technical support representative.

⚠ **Caution** Use the commands discussed in this appendix only under the direction of a technical support representative. Creating a core dump while the switch router is functioning in a network can disrupt network operation. The resulting binary file, which is very large, must be transferred to a Trivial File Transfer Protocol (TFTP), File Transfer Protocol (FTP), or remote copy (rcp) server. It is subsequently interpreted by technical personnel who have access to source code and detailed memory maps.

Using exception Commands

Use the **exception** class of configuration commands only after consulting with a technical support representative. These commands are useful for debugging purposes, but they can result in unexpected behavior.

Creating a Core Dump

To obtain a core dump when a switch router fails, use the **exception dump** *ip-address* switch router configuration command (where *ip-address* is the address of your TFTP server).

If you include this command in your configuration, the switch router attempts a core dump when it crashes. The core dump is written to a file named *hostname*-core on your server, where *hostname* is the name of the switch router. You can change the name of the core file by entering the **exception core-file** *filename* command.

The default protocol for transferring the core dump is TFTP. However, TFTP transfers only 16 MB of the core dump file. If the switch router memory is over 16 MB, only the first 16 MB is transferred. To transfer the whole core dump, configure the switch router to use rcp or FTP for core dumps with the **exception protocol** command.

The following example configures a switch router to use rcp to dump the core file when it crashes:

```
Switch# configure terminal
Switch (config)# ip rcmd remote-username red
Switch (config)# exception protocol rcp
Switch (config)# exception dump 172.17.92.2
```

The following example configures a switch router to use FTP to dump the core file when it crashes:

```
Switch# configure terminal
Switch (config)# ip ftp username red
Switch (config)# ip ftp password blue
Switch (config)# exception protocol ftp
Switch (config)# exception dump 172.17.92.2
```

Note The remote machine must be configured to allow the switch router to write to it. For example, if you are using rcp with a UNIX system, the .rhosts file for the remote user must contain an entry for the switch router. Refer to the documentation for your FTP or rcp server for details.

This procedure can fail for certain types of system crashes. However, if successful, the core dump file will be the size of the memory available on the processor.

Creating an Exception Memory Core Dump

During the debugging process, you can cause the switch router to create a core dump and reboot when certain memory size parameters are violated. The **exception memory** commands define a minimum contiguous block of memory in the free pool and a minimum size for the free memory pool.

[**no**] **exception memory fragment** *size*

[**no**] **exception memory minimum** *size*

The value of *size* is in bytes and is checked every 60 seconds. If you enter a size that is greater than the free memory and the **exception dump** command has been configured, the switch router creates a core dump and reloads the Cisco IOS software after 60 seconds. If the **exception dump** command is not configured, the switch router reloads without generating a core dump.

The following example configures the switch router to monitor the free memory. If the free memory falls below 250,000 bytes, the switch router dumps the core and reloads.

```
Switch# configure terminal
Switch (config)# exception dump 131.108.92.2
Switch (config)# exception core-file memory.overrun
Switch (config)# exception memory minimum 250000
```

Using the write core Command

You can create test core dumps by using the **write core** privileged EXEC command. If you use this command, the switch router generates a core dump without reloading, which is useful if the switch router is malfunctioning but has not crashed.

Caution Use the **write core** command only under the direction of a technical support representative. Creating a core dump while the switch router is functioning in a network can disrupt network operation. The resulting binary file, which is very large, must be transferred to a TFTP, FTP, or rcp server and subsequently interpreted by technical personnel who have access to source code and detailed memory maps.

Depending on your TFTP server, you might need to create an empty target file to which the switch router can write the core dump.

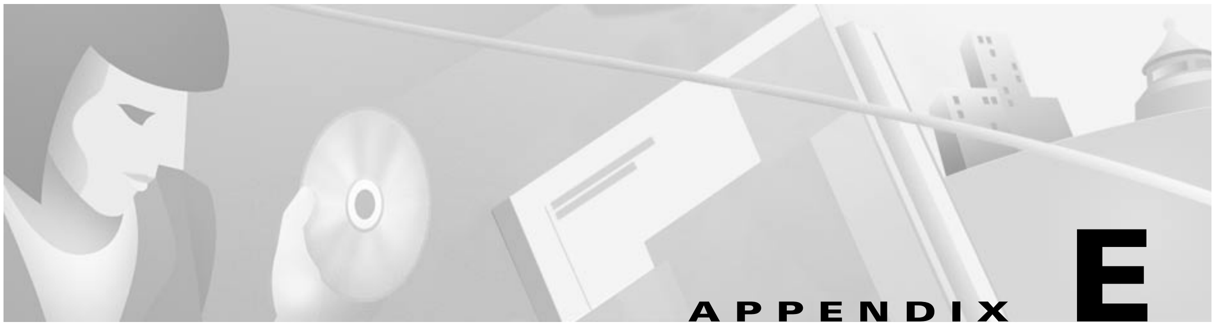

APPENDIX E

Technical Support

When you have a problem that you cannot resolve, contact customer service. To help resolve these problems, gather relevant information about your network prior to calling.

Gathering Information about Your Internetwork

Before gathering any specific data, compile a list of all symptoms users have reported on the internetwork (such as connections dropping or slow host response).

The next step is to gather specific information. Typical information needed to troubleshoot internetworking problems fall into two general categories: information required for any situation and information specific to the topology, technology, protocol, or problem.

Information that is always required by technical support engineers includes the following:

- Configuration listing of all switch routers involved
- Complete specifications of all switch routers involved
- Version numbers of software (obtained by using the **show version** command) and Flash code (obtained by using the **show controllers** command) on all relevant switch routers
- Network topology map
- List of hosts and servers (host and server type, number on network, description of host operating systems that are implemented)
- List of network layer protocols, versions, and vendors

To assist you in gathering this required data, the **show tech-support** EXEC command has been added in Cisco IOS Release 11.1(4) and later. This command provides general information about the switch router that you can provide to your technical support representative when you are reporting a problem.

The **show tech-support** command display includes the **show version**, **show hardware**, **show diag power-on**, **show running-config**, **show controllers**, **show stacks**, **show interfaces**, **show buffers**, **show process memory**, and **show process** EXEC commands.

Specific information that might be needed by technical support varies, depending on the situation, and include the following:

- Output from the following general **show** commands:

  show interfaces

  show controllers [**atm** | **serial** | **e1** | **ethernet**]

  show processes [**cpu** | **mem**]

  show buffers

  show memory summary
- Output from the following protocol-specific **show** commands:

  show *protocol* **route**

  show *protocol* **traffic**

  show *protocol* **interface**

  show *protocol* **arp**
- Output from relevant **debug** privileged EXEC commands
- Output from protocol-specific **ping** and **trace** command diagnostic tests, as applicable
- Network analyzer traces, as applicable
- Core dumps obtained by using the **exception dump** switch configuration command, or by using the **write core** switch configuration command if the system is operational, as appropriate

Getting the Data from Your Switch Router

When obtaining information from your switch router, tailor your method to the system that you are using to retrieve the information. Following are some hints for different platforms:

- PC and Macintosh—Connect a PC or Macintosh to the console port of the switch router and log all output to a disk file (using a terminal emulation program). The exact procedure varies depending on the communication package used with the system.
- Terminal connected to console port or remote terminal—The only way to get information with a terminal connected to the console port or with a remote terminal is to attach a printer to the AUX port on the terminal (if one exists) and to force all screen output to go to the printer. Using a terminal is undesirable because there is no way to capture the data to a file.
- UNIX workstation—At the UNIX prompt, enter the **script** *filename* command, then use Telnet to connect to the switch router. The UNIX **script** command captures all screen output to the specified filename. To stop capturing output and close the file, enter the end-of-file character (typically **^D**) for your UNIX system.

Note To get your system to automatically log specific error messages or operational information to a UNIX syslog server, use the **logging** *internet-address* switch configuration command. For more information about using the **logging** command and setting up a syslog server, refer to the Cisco IOS configuration guides and command reference publications.

Providing Data to Customer Service

If you need technical assistance with a Cisco product that is under warranty or covered by a maintenance contract, contact Cisco's Technical Assistance Center (TAC) to open a case. Contact the TAC with a phone call or an e-mail message:

- North America: 800-553-2447, e-mail: tac@cisco.com
- Europe: 32 2 778 4242, e-mail: euro-tac@cisco.com
- Asia-Pacific: 61 2 9935 4107, e-mail: asiapac-tac@cisco.com

When submitting information to your technical support representative, electronic data is preferred. Electronic data significantly eases the transfer of information between technical support personnel and development staff. Common electronic formats include data sent via e-mail and files sent using File Transfer Protocol (FTP).

If you are submitting data to your technical support representative, use the following list to determine the preferred method for submission:

1. The preferred method of information submission is via FTP service over the Internet. If your environment supports FTP, you can place your file in the *incoming* directory on the host *cco.cisco.com*.
2. The next best method is to send data by electronic mail. Before using this method, be sure to contact your technical support representative, especially when transferring binary core dumps or other large files.

   If you use e-mail, do not use encoding methods such as binhex or zip. Only MIME-compliant mail should be used.
3. Use a PC-based communications protocol, such as *Kermit*, to upload files to Cisco.com. Again, be sure to contact your technical support representative before attempting any transfer.
4. Transfer by disk or tape.
5. The least favorable method is hard-copy transfer by fax or physical mail.

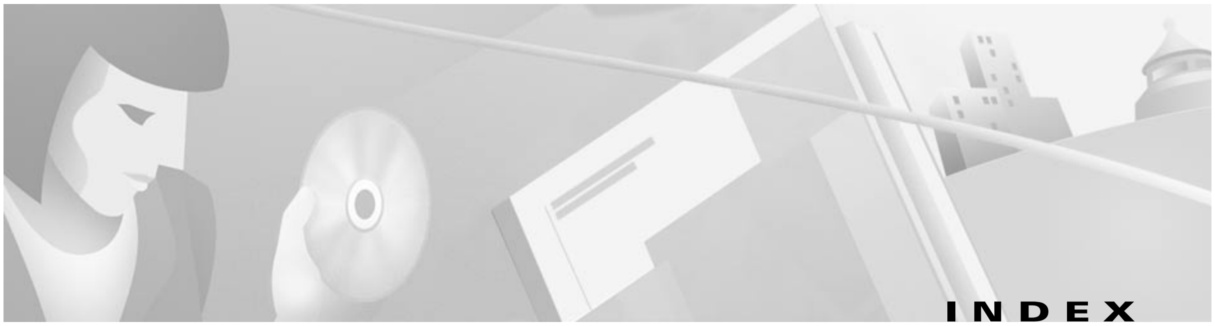

INDEX