Cisco 7200 Series Design Library: ATM Traffic Management

December, 2005

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

Text Part Number: OL-3274-01

Cisco 7200 Series Design Library: ATM Traffic Management

About the Cisco 7200 Series Design Library

This chapter provides an introduction to the *Cisco 7200 Series Design Library: ATM Traffic Management* book in the Cisco 7200 Series design library.

Objective

The *Cisco 7200 Series Design Library: ATM Traffic Management* book is the first publication in the Cisco 7200 series design library. The purpose of the book is to provide you with the architectural and design concepts and guidelines that are necessary to understand for effective management of your ATM traffic on a Cisco 7200 series router.

Audience and Scope

This book is intended for ATM network engineers, designers, and network support personnel who configure and maintain the Cisco 7200 series router in their ATM network. It is expected that the reader is knowledgeable about ATM technology.

The book includes a wide range of topics that involve the processing of ATM traffic on the Cisco 7200 series router, including both the hardware and software aspects of that processing. It attempts to provide both the "big picture" by discussing the complete flow of traffic on the router and the key architectural concepts involved in that flow, and also provide the in-depth picture in those areas where optimization might occur to support your ATM traffic requirements.

Although some information is provided about all of the ATM port adapters that are supported on the Cisco 7200 series routers, the book primarily focuses on permanent virtual circuit (PVC) configuration on the PA-A3 and PA-A6 ATM port adapters.

Document Revision History

Table 1 records technical changes to this document. The table shows the document revision number for the change, the date of the change, and a brief summary of the change.

Table 1 Document Revision History

Revision	Date	Change Summary
OL-3274-01	December 15, 2005	Second release, with the following revisions: • Revised statements about GTS and class-based shaping in Chapter 2 to read as follows: "In most cases you do not use class-based shaping to implement traffic shaping on the outbound ATM interface. However, in certain Cisco IOS software releases, you can use the **shape** command within an outbound service-policy with the PA-A3 or PA-A6 ATM port adapters to achieve class-based shaping at Layer 3." • Removed reference to use of GTS on PA-A1 for nrt-VBR service because the PA-A1 does not support nrt-VBR class of service. • Updated this chapter with new information about Obtaining Documentation.
OL-3274-01	June 23, 2003	First release.

Organization and Use

This book contains the following chapters:

- Chapter 1, "Introduction to ATM Traffic Management on the Cisco 7200 Series Routers"—Provides a brief introduction to ATM traffic management, and begins a discussion of some of the concepts associated with traffic management on the Cisco 7200 series router as an edge device on the User-Network Interface (UNI).
- Chapter 2, "Cisco 7200 Series Architecture and Design for ATM Traffic Management"—Discusses the overall flow of ATM traffic on the Cisco 7200 series router, and the different hardware and software architectures that are part of that flow. These hardware and software components work together to affect the overall performance of the flow of a packet through the router and onto the network as ATM cells.
- Chapter 3, "ATM Traffic Management Hardware and Software Planning"—Provides an introduction to the ATM port adapter hardware and software that is supported on the Cisco 7200 series routers. It includes hardware installation guidelines and verification information and a review of the Cisco IOS software releases supported by the Cisco 7200 series routers, including a summary of where certain key ATM features were introduced.
- Chapter 4, "Preparing to Configure ATM Traffic Management and QoS Features"—Includes guidelines for preparing to configure ATM traffic management and QoS features and identifies tasks that you should implement as a regular and ongoing assessment of your network.

- Chapter 5, "Configuring Traffic Shaping on the PA-A3 and PA-A6 ATM Port Adapters"—Provides a combination of design and configuration information to help you make informed decisions about implementing and optimizing traffic shaping on your ATM port adapters.
- Chapter 6, "Configuring QoS on the Layer 3 Queues for the PA-A3 and PA-A6 ATM Port Adapters"—Provides a brief introduction and some guidelines for configuring the IP to ATM Class of Service (CoS) features on the PA-A3 and PA-A6 ATM port adapters.
- Chapter 7, "Configuring the Ring Limits on the PA-A3 and PA-A6 ATM Port Adapters"—Describes how to optimize the ring limits on the PA-A3 and PA-A6 ATM port adapters.
- Chapter 8, "ATM Traffic Management Case Studies and Configuration Examples"—Provides case studies and configuration examples for enterprise networks using Cisco 7200 series routers in an ATM environment. All of the examples in this chapter represent actual lab-tested configurations from Cisco Systems proof-of-concept and solutions labs.
- Chapter 9, "Frequently Asked Questions"—Answers some of the frequently asked questions (FAQs) about traffic management on the PA-A3 and PA-A6 ATM port adapters.

Conventions

Command descriptions use the following conventions:

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.

Show command output examples use the following conventions:

`screen` font	Terminal sessions and information the system displays are in `screen` font.
`boldface screen` font	The **show** command that you must enter is in **`boldface screen`** font. Important areas of the display that are discussed within the text are highlighted in the output display using **`boldface screen`** font.
italic screen font	Arguments for which you supply values are in *italic screen* font.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

Notes and cautionary statements use these conventions:

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

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

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. Cisco also provides several ways to obtain technical assistance and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation at this URL:

http://www.cisco.com/techsupport

You can access the Cisco website at this URL:

http://www.cisco.com

You can access international Cisco websites at this URL:

http://www.cisco.com/public/countries_languages.shtml

Product Documentation DVD

Cisco documentation and additional literature are available in the Product Documentation DVD package, which may have shipped with your product. The Product Documentation DVD is updated regularly and may be more current than printed documentation.

The Product Documentation DVD is a comprehensive library of technical product documentation on portable media. The DVD enables you to access multiple versions of hardware and software installation, configuration, and command guides for Cisco products and to view technical documentation in HTML. With the DVD, you have access to the same documentation that is found on the Cisco website without being connected to the Internet. Certain products also have .pdf versions of the documentation available.

The Product Documentation DVD is available as a single unit or as a subscription. Registered Cisco.com users (Cisco direct customers) can order a Product Documentation DVD (product number DOC-DOCDVD=) from Cisco Marketplace at this URL:

http://www.cisco.com/go/marketplace/

Ordering Documentation

Beginning June 30, 2005, registered Cisco.com users may order Cisco documentation at the Product Documentation Store in the Cisco Marketplace at this URL:

http://www.cisco.com/go/marketplace/

Nonregistered Cisco.com users can order technical documentation from 8:00 a.m. to 5:00 p.m. (0800 to 1700) PDT by calling 1 866 463-3487 in the United States and Canada, or elsewhere by calling 011 408 519-5055. You can also order documentation by e-mail at tech-doc-store-mkpl@external.cisco.com or by fax at 1 408 519-5001 in the United States and Canada, or elsewhere at 011 408 519-5001.

Documentation Feedback

You can rate and provide feedback about Cisco technical documents by completing the online feedback form that appears with the technical documents on Cisco.com.

You can send comments about Cisco documentation to bug-doc@cisco.com.

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Cisco Product Security Overview

Cisco provides a free online Security Vulnerability Policy portal at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

From this site, you can perform these tasks:

- Report security vulnerabilities in Cisco products.
- Obtain assistance with security incidents that involve Cisco products.
- Register to receive security information from Cisco.

A current list of security advisories and notices for Cisco products is available at this URL:

http://www.cisco.com/go/psirt

If you prefer to see advisories and notices as they are updated in real time, you can access a Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed from this URL:

http://www.cisco.com/en/US/products/products_psirt_rss_feed.html

Reporting Security Problems in Cisco Products

Cisco is committed to delivering secure products. We test our products internally before we release them, and we strive to correct all vulnerabilities quickly. If you think that you might have identified a vulnerability in a Cisco product, contact PSIRT:

- Emergencies—security-alert@cisco.com

  An emergency is either a condition in which a system is under active attack or a condition for which a severe and urgent security vulnerability should be reported. All other conditions are considered nonemergencies.

- Nonemergencies—psirt@cisco.com

In an emergency, you can also reach PSIRT by telephone:

- 1 877 228-7302
- 1 408 525-6532

Tip We encourage you to use Pretty Good Privacy (PGP) or a compatible product to encrypt any sensitive information that you send to Cisco. PSIRT can work from encrypted information that is compatible with PGP versions 2.*x* through 8.*x*.

Never use a revoked or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one linked in the Contact Summary section of the Security Vulnerability Policy page at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

The link on this page has the current PGP key ID in use.

Obtaining Technical Assistance

Cisco Technical Support provides 24-hour-a-day award-winning technical assistance. The Cisco Technical Support & Documentation website on Cisco.com features extensive online support resources. In addition, if you have a valid Cisco service contract, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not have a valid Cisco service contract, contact your reseller.

Cisco Technical Support & Documentation Website

The Cisco Technical Support & Documentation website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day, at this URL:

http://www.cisco.com/techsupport

Access to all tools on the Cisco Technical Support & Documentation website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

http://tools.cisco.com/RPF/register/register.do

Note Use the Cisco Product Identification (CPI) tool to locate your product serial number before submitting a web or phone request for service. You can access the CPI tool from the Cisco Technical Support & Documentation website by clicking the **Tools & Resources** link under Documentation & Tools. Choose **Cisco Product Identification Tool** from the Alphabetical Index drop-down list, or click the **Cisco Product Identification Tool** link under Alerts & RMAs. The CPI tool offers three search options: by product ID or model name; by tree view; or for certain products, by copying and pasting **show** command output. Search results show an illustration of your product with the serial number label location highlighted. Locate the serial number label on your product and record the information before placing a service call.

Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool provides recommended solutions. If your issue is not resolved using the recommended resources, your service request is assigned to a Cisco engineer. The TAC Service Request Tool is located at this URL:

http://www.cisco.com/techsupport/servicerequest

For S1 or S2 service requests or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)
EMEA: +32 2 704 55 55
USA: 1 800 553-2447

For a complete list of Cisco TAC contacts, go to this URL:

http://www.cisco.com/techsupport/contacts

Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

Severity 1 (S1)—Your network is “down,” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Severity 2 (S2)—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Severity 3 (S3)—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Severity 4 (S4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- Cisco Marketplace provides a variety of Cisco books, reference guides, documentation, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:

  http://www.cisco.com/go/marketplace/

- *Cisco Press* publishes a wide range of general networking, training and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press at this URL:

  http://www.ciscopress.com

- *Packet* magazine is the Cisco Systems technical user magazine for maximizing Internet and networking investments. Each quarter, Packet delivers coverage of the latest industry trends, technology breakthroughs, and Cisco products and solutions, as well as network deployment and troubleshooting tips, configuration examples, customer case studies, certification and training information, and links to scores of in-depth online resources. You can access Packet magazine at this URL:

  http://www.cisco.com/packet

- *iQ Magazine* is the quarterly publication from Cisco Systems designed to help growing companies learn how they can use technology to increase revenue, streamline their business, and expand services. The publication identifies the challenges facing these companies and the technologies to help solve them, using real-world case studies and business strategies to help readers make sound technology investment decisions. You can access iQ Magazine at this URL:

  http://www.cisco.com/go/iqmagazine

  or view the digital edition at this URL:

  http://ciscoiq.texterity.com/ciscoiq/sample/

- *Internet Protocol Journal* is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

  http://www.cisco.com/ipj

- Networking products offered by Cisco Systems, as well as customer support services, can be obtained at this URL:

  http://www.cisco.com/en/US/products/index.html

- Networking Professionals Connection is an interactive website for networking professionals to share questions, suggestions, and information about networking products and technologies with Cisco experts and other networking professionals. Join a discussion at this URL:

  http://www.cisco.com/discuss/networking

- World-class networking training is available from Cisco. You can view current offerings at this URL:

  http://www.cisco.com/en/US/learning/index.html

Introduction to ATM Traffic Management on the Cisco 7200 Series Routers

In the latest generation of IP networks, with the growing implementation of Voice over IP (VoIP) and multimedia applications, the addition of voice and video traffic to the traditional IP data network has become increasingly common. Voice, video, and data traffic types have different transmission characteristics and service-level requirements.

The ATM technology is well-suited to transport mixed traffic because of its built-in ability to negotiate and guarantee a certain level of quality of service (QoS) from the source to the end device. This makes ATM a desirable transport method for mixed traffic through an IP network over a WAN.

This chapter provides a brief introduction to ATM traffic management, and begins a discussion of some of the concepts associated with traffic management on the Cisco 7200 series routers as an edge device on the User-Network Interface (UNI).

This chapter includes the following topics:

Traffic Characteristics

Voice, video, and data traffic are differentiated by the following transmission characteristics:

- Voice—Traffic flows with a regular pattern at a constant rate that is sensitive to delay and delay variation. When compression techniques are in use, voice traffic is more sensitive to error than uncompressed voice.
- Video—Real-time video traffic has similar transmission characteristics to voice traffic, but also requires high bandwidth. When compression techniques are in use, video traffic is more sensitive to error than uncompressed video.
- Data—Traffic flows with an irregular pattern that is often called *bursty* because of its variability in rate and amount of traffic. Data traffic is not sensitive to delay or delay variation, but it is sensitive to error.

Traffic management is vital to the performance and overall health of the ATM network. ATM uniquely satisfies the different transmission requirements of mixed traffic on a common network through its multiple service categories and QoS implementation.

Figure 1-1 Voice, Video, and Data Transmission Requirements

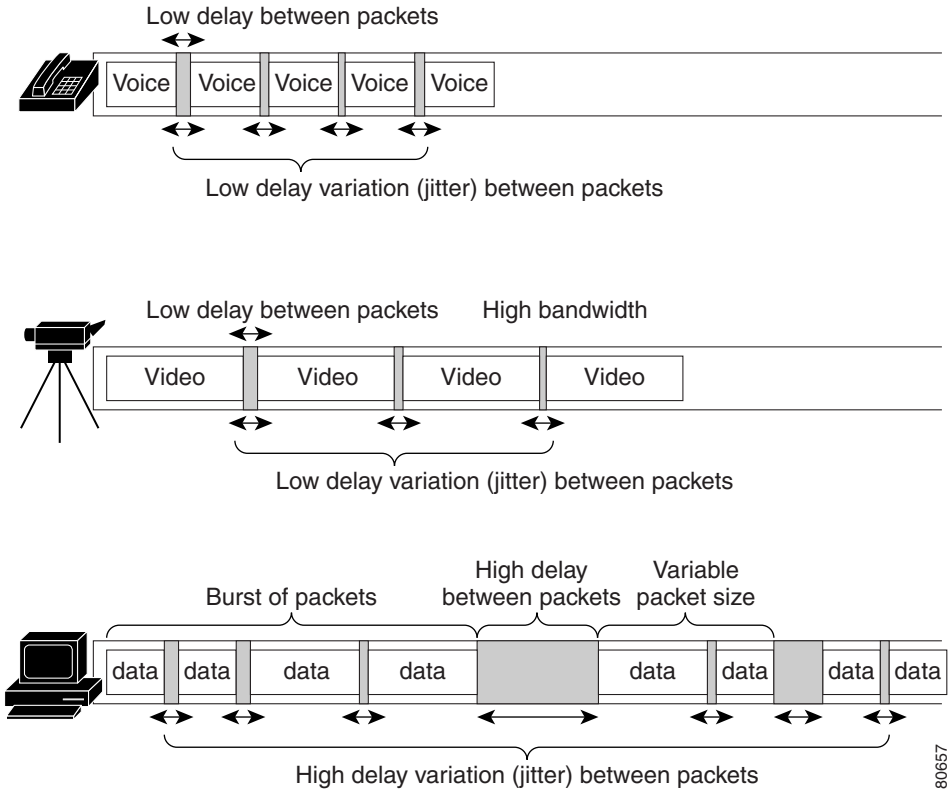

Traffic Contract

An ATM WAN is frequently a public network owned and managed by a service provider who supports multiple customers. These customers agree upon and pay for a certain level of bandwidth and performance from the service provider over that WAN. This agreement becomes the basis of the traffic contract, which defines the traffic parameters and the QoS that is negotiated for each virtual connection for that user on the network.

References to the traffic contract in an ATM network represent a couple of things. First, the traffic contract represents an actual service agreement between the user and the service provider for the expected network-level support. Second, the traffic contract refers to the specific traffic parameters and QoS values negotiated for an ATM virtual connection at call setup, which are implemented during data flow to support that service agreement.

The traffic contract also establishes the criteria for policing of ATM virtual connections on the network to ensure that violations of the agreed-upon service levels do not occur.

ATM Service Categories and Traffic Parameters

The ATM Forum Traffic Management specifications define several service categories to group traffic according to different transmission characteristics and performance needs. Each ATM service category is qualified by certain traffic parameters and QoS parameters that define the desired network performance for the permanent virtual circuit (PVC) or switched virtual circuit (SVC) on the ATM network.

The traffic parameters, sometimes called *descriptors*, are used to shape the flow of ATM cells. ATM service categories, and their corresponding traffic and QoS parameters, are the basis for differentiating services on the ATM network and for establishing the traffic contract for a particular connection.

Differences in Implementation of Traffic Parameters and QoS in PVCs and SVCs

All PVC and SVC traffic parameters and QoS parameters are established for the duration of a connection. The difference between PVCs and SVCs occurs in the implementation of these parameters.

On PVCs, traffic shaping parameters are based upon a manual configuration on both the edge device (router) and the switch. Therefore, no exchange of service-level information occurs between the edge device and the switch through signaling while a PVC connection is being established. Therefore, it is possible for configuration mismatches to occur between the router and the switch.

However, for SVCs, traffic parameters and QoS parameters are exchanged between the edge device and the switch through signaling. The edge device requests the required performance from the network, and the network responds with what it can provide. From there, the edge device can either accept or reject the connection. This is referred to as a *two-way handshake*.

ATM Service Categories

The following ATM service categories are defined by the ATM Forum specifications and are supported on the Cisco 7200 series router to perform traffic shaping. The ATM service categories can be subdivided by their support for real-time or non-real-time applications.

Note Cisco Systems does not support the Guaranteed Frame Rate (GFR) service category on the Cisco 7200 series router.

Table 1-1 on page 1-7 provides examples and summarizes the ATM traffic parameters and QoS parameters associated with each service category.

Real-Time Service Categories

There are two ATM service categories that are designed to support real-time applications, which require low cell delay and cell loss:

- **Constant bit rate (CBR)**—Supports real-time applications that request a static amount of bandwidth that is continuously available for the duration of the connection.
- **Real-time variable bit rate (rt-VBR)**—Supports real-time applications that have bursty transmission characteristics.

Non-Real-Time Service Categories

There are three ATM service categories that are designed to support non-real-time applications, which typically support data services:

- **Available bit rate (ABR)**—Supports non-real-time applications that tolerate high cell delay, and can adapt cell rates according to changing network resource availability to prevent cell loss. The ABR service category is characterized by reactive congestion control, where it uses flow control mechanisms to learn about the network conditions and adjust cell rates accordingly.
- **Non-real-time variable bit rate (nrt-VBR)**—Supports non-real-time applications with bursty transmission characteristics that tolerate high cell delay, but require low cell loss.
- **Unspecified bit rate (UBR)**—Supports non-real-time applications that tolerate both high cell delay and cell loss on the network. There are no network service-level guarantees for the UBR service category, and therefore it is a best-effort service.

Cisco-Specific UBR+ Service Category

Cisco Systems has also developed a second UBR service category called UBR+, which implements the Minimum Cell Rate (MCR) traffic parameter:

- **Unspecified bit rate plus (UBR+)** —Supports non-real-time applications that tolerate both high cell delay and cell loss on the network, but request a minimum guaranteed cell rate. As with the UBR service category, there are no network service-level guarantees for UBR+. However, the network can grant a service-level guarantee for the requested MCR.

For a description of MCR, see the "ATM Traffic Parameters" section on page 1-5.

A similar UBR service category is specified in an addendum to the ATM Forum Traffic Management specifications, which discusses implementation of an optional Minimum Desired Cell Rate (MDCR) parameter for the UBR service category.

However, the Cisco Systems implementation and the ATM Forum implementation vary in how the minimum rate is signaled to the ATM network. Cisco Systems uses the existing MCR information element (IE) that is used by the ABR service category, but the parameter has a different interpretation. For UBR+, the MCR parameter represents a desired cell rate; but in ABR, the MCR specifies the lowest acceptable cell rate.

The ATM Forum does not use the MCR IE, but implements a new IE for the MDCR traffic parameter.

Note Cisco Systems introduced support for UBR+ on SVCs only for the PA-A3 ATM port adapter in Cisco IOS Release 11.3 T. It is also available for SVC configuration on the PA-A6 ATM port adapter. However, the UBR+ service category is not supported for PVCs and is not available on the PA-A1 or PA-A2 ATM port adapters. For configuration guidelines and an example, see the "UBR+ Configuration Guidelines" section on page 5-28 and the "UBR+ Configuration Example" section on page 5-29.

ATM Traffic Parameters

The following traffic parameters are used to qualify the different ATM service categories:

- **Minimum Cell Rate (MCR)**—Cell rate (cells per second) at which the edge device is always allowed to transmit.

  For UBR+, the MCR is the minimum cell rate *requested* by the edge device as a guaranteed service-level for the SVC.
- **Peak Cell Rate (PCR)**—Cell rate (cells per second) that the edge device cannot exceed. Some service categories have a limit on the number of cells that can be sent at the PCR without penalty for violation of the traffic contract.
- **Cell Delay Variation Tolerance (CDVT)**—Allowable deviation in cell times for a PVC that is transmitting above the PCR. For a given cell interarrival time expected by the ATM switch, CDVT allows for some variance in the transmission rate. It allows a certain number of cells to arrive faster than the expected cell interarrival time without penalty for violation of the traffic contract.
- **Sustainable Cell Rate (SCR)**—Upper boundary for the average rate at which the edge device can transmit cells without loss.
- **Maximum Burst Size (MBS)**—Number of cells that the edge device can transmit up to the PCR for a limited period of time without penalty for violation of the traffic contract.

Note To configure traffic shaping parameters for the ATM port adapters, you typically specify a value in terms of bits per second, which uses the same unit of measure as the line rate. However, be aware that ATM transmission rates over the network actually are implemented according to a total number of cell time slots (or cells per second). Each time slot represents a cell time (in microseconds). Further, ATM switches frequently measure bandwidth according to cell times, not bits per second.

Table 1-1 on page 1-7 provides examples and summarizes the ATM traffic parameters and QoS parameters associated with each service category.

Figure 1-2 shows the relationships between the different ATM traffic parameters.

Figure 1-2 Relationships of ATM Traffic Parameters

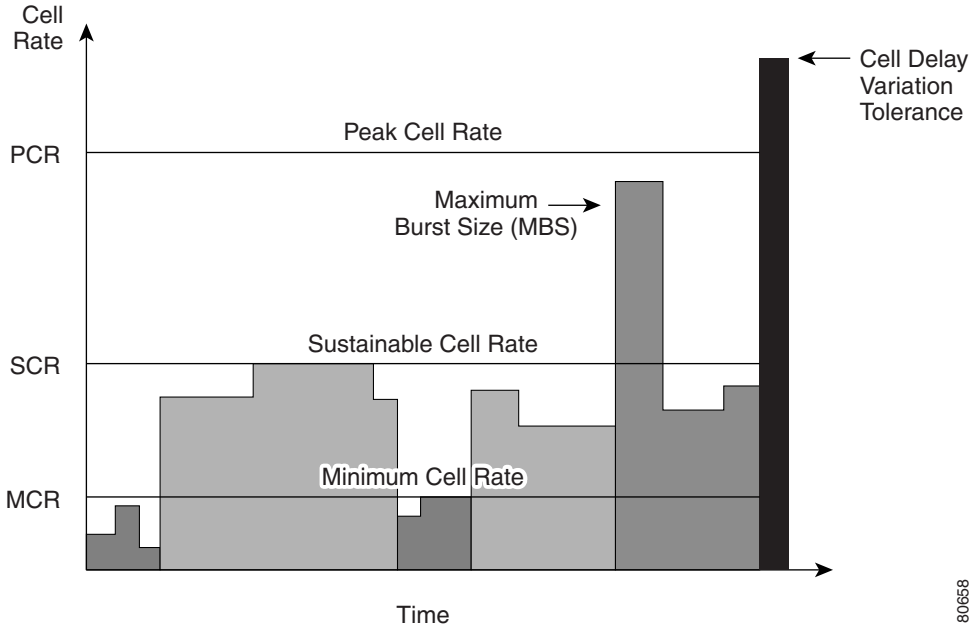

ATM QoS Parameters

The ATM Forum specifications define specific QoS parameters that are used to manage cell delay and cell loss over the ATM network for each of the different ATM service categories. Some of these QoS parameters are considered negotiable and some are not.

For SVCs, ATM switches evaluate the requested traffic parameters and QoS parameters using the Connection Admission Control (CAC) algorithm. CAC ensures that the requested QoS can be served throughout the duration of the connection over the network, from the source to the destination, without impacting other connections.

Negotiable QoS Parameters

The following cell delay and cell loss parameters are considered negotiable because the information is exchanged through signaling between the UNI edge device and the network-to-network interface (NNI) switch while an ATM connection is being established.

Cell Delay Parameters

The ATM Forum specifications support two negotiable parameters for cell delay:

- **Maximum cell transfer delay (maxCTD)**—Maximum length of time allowed for the network to transmit a cell from the source UNI device to the destination UNI device.
- **Peak-to-peak cell delay variation (peak-to-peak CDV)**—Maximum variation allowed from the fixed CTD for each cell transmitted from the source UNI device to the destination UNI device. Represents the allowable jitter, or distortion, between cell interarrival times over the network.

Cell Loss Parameters

The ATM Forum specifications support the following negotiable parameter for cell loss:

- **Cell loss ratio (CLR)**—Allowable percentage of cells (lost cells divided by total number of cells transmitted) that the network can discard due to congestion.

Non-Negotiable QoS Parameters

The following QoS parameters are not exchanged during connection setup on the ATM network:

- **Cell error ratio (CER)**—Allowable percentage of cells (errored cells divided by the total number of all transmitted cells) that can be in error.
- **Severely errored cell block ratio (SECBR)**—Allowable percentage of cell blocks (severely errored cell blocks divided by the total number of transmitted cell blocks) that can be severely in error. A cell block is a number of consecutively transmitted cells on a particular connection. A cell block is considered severely errored when more than a maximum numbe of errored cells, lost cells, or misinserted cells occur within that cell block.
- **Cell misinsertion rate (CMR)**—Allowable rate of misinserted cells (misinserted cells divided by the time period during which misinserted cells were collected). This rate does not include severely errored cell blocks. Misinserted cells are cells that are received with an incorrect VPI/VCI value.

Table 1-1 provides examples and summarizes the ATM traffic parameters and QoS parameters associated with each service category.

Table 1-1 ATM Traffic Parameters and QoS Parameters by Service Category

ATM Service Category	Application Examples	Traffic Parameters	ATM QoS Parameters
ABR	Critical data transfer, such as for defense information where rapid access to network bandwidth is important.	MCR, PCR	CLR (optional)
CBR	Telephone conversations, voice mail, or audio services (radio, or audio library). Videoconferencing, video on demand.	PCR, CDVT	Peak-to-peak CDV, maxCTD, CLR
nrt-VBR	Airline reservations, banking transactions.	PCR, CDVT, SCR, MBS	CLR
rt-VBR	Compressed or packetized voice or video including telephone conversations, voicemail, HDTV.	PCR, CDVT, SCR, MBS	Peak-to-peak CDV, maxCTD, CLR

Table 1-1 ATM Traffic Parameters and QoS Parameters by Service Category (continued)

ATM Service Category	Application Examples	Traffic Parameters	ATM QoS Parameters
UBR	File transfer and e-mail.	PCR (optional)[1]	None supported
UBR+[2]	Interconnecting IP routers with virtual channel connections (VCCs) or virtul path connections (VPCs).	PCR (optional), MCR	None supported

1. Cisco Systems supports specification of the PCR parameter; however, some ATM switches do not support enforcement of PCR and the value becomes informational.
2. UBR+ is a special ATM service category developed by Cisco Systems. It is similar to the ATM Forum's addendum specification for a Minimum Desired Cell Rate (MDCR) parameter for the UBR service category.

ATM QoS and Cisco IOS QoS Distinctions

The term QoS is a frequently used term that can represent many different aspects of data transmission and traffic management on a network. Therefore, it is important to be clear about what is meant by QoS.

It is helpful to your understanding of ATM traffic management on the Cisco 7200 series router if you do not confuse the QoS parameters associated with an ATM service category with the QoS features in the Cisco IOS software that can be implemented for IP over ATM. These are two very distinct areas of QoS.

ATM QoS

The ATM technology defines QoS in terms of the management of cell delay and cell loss over the ATM network. Therefore, ATM QoS represents the end-to-end network performance of cell transmission from the source to the destination, not including the edge router.

The ATM Forum specifications define specific QoS parameters that are used to manage cell delay and cell loss over the ATM network. Technically, the Cisco IOS QoS features supported by ATM have no relevance to these parameters or to the definition of QoS in the ATM Forum specifications.

Cisco IOS QoS Software

Cisco IOS QoS software features do not affect the performance of cell transmission once the cells leave the edge router and are transported over the ATM network. Cisco IOS QoS software affects packets in the Layer 3 queues on the router—not cells over the ATM network.

ATM Adaptation Layers and ATM Service Categories

Like ATM service categories, different ATM Adaptation Layers (AALs) also are defined to support different classifications of traffic types. The primary role of an AAL is to separate data into even, 48-byte chunks called a segmentation and reassembly (SAR) protocol data unit (PDU). Once the AAL creates a 48-byte chunk, the ATM Layer adds a 5-byte header to create a 53-byte ATM cell.

To optimize services for different traffic classifications, each AAL type (other than AAL5) provides different information within the 48-byte PDU. This extra information becomes overhead for the payload in the 48-byte PDU. All AAL types segment data into 48-byte SAR-PDUs, but the encapsulations within those 48 bytes vary by AAL type.

For example, AAL1 uses one byte out of the 48 bytes to support a sequence number field and a sequence number protection field. AAL5 is the most efficient of all of the AAL types because it does not use any of the 48-byte payload for extra information.

The ATM Forum identifies four different types of AALs (AAL3 and AAL4 have been combined), which correspond to the different types of traffic to be supported. Table 1-2 shows the relationship of AAL types to ATM service categories.

Table 1-2 Relationship of AAL Types to ATM Service Categories

AAL Type	ATM Service Category	Application
AAL0[1]	(Not applicable)	Cell relay over MPLS
AAL1	CBR	Voice
AAL2	rt-VBR	Compressed voice (allows for silent periods) or compressed video; Voice over ATM (VoATM).
AAL3/4[2]	ABR, UBR	No longer used
AAL5	ABR, CBR, nrt-VBR, rt-VBR, UBR	Data

1. AAL0 represents a null adaptation layer, which is used for cell relay over Multiprotocol Layer Switching (MPLS). Support for AAL0 was introduced for the PA-A3 ATM port adapters in Cisco IOS Release 12.0(25)S and Cisco IOS Release 12.2 S. As of Cisco IOS Release 12.0(25)S, it is not supported for the PA-A6 ATM port adapters.
2. AAL3/4 is no longer in use, but is included here to show the evolution of AAL types. AAL5 is a simplified version of AAL3/4.

Port Adapter Support for AAL on the Cisco 7200 Series Routers

AAL5 is one of the most widely used AAL types, and is the default encapsulation type for the ATM port adapters on the Cisco 7200 series router. AAL5 uses the entire 48-byte payload for data.

Different ATM port adapters on the Cisco 7200 series routers support different AAL encapsulation methods.

For more information about the differences in ATM port adapter support, see Chapter 3, “ATM Traffic Management Hardware and Software Planning.”

CBR for Voice and Data on the Cisco 7200 Series Routers

Cisco documentation differentiates between CBR for voice and CBR for data, depending on the AAL type supporting the CBR virtual connection:

- CBR for voice—Uses AAL1, and includes circuit emulation service (CES) and Voice over ATM applications. A 1-byte AAL1 header uses time stamps, sequence numbers, and other bits to help the ATM network deal with ATM-layer defects such as cell delay variation, cell misinsertion, and cell loss.
- CBR for data—Uses AAL5, but the same interface typically does not support CBR for voice. AAL5 adds an 8-byte trailer with a 4-byte cyclic redundancy check (CRC) for detecting errors in a PDU.

Congestion on an ATM Network

Well-behaved traffic that conforms to the agreed-upon service levels is critical to the performance of the public ATM WAN. Without the proper controls and management in place, there is the potential for certain customers to consume bandwidth above the agreed-upon rate. This can cause congestion, which not only prevents other user traffic from its right to access that bandwidth, but can cause significant degradation to the performance on the network.

The cost of congestion to ATM network performance is better understood when you consider what happens if one or more cells are marked and dropped during transmission of a packet. Consider an AAL5 PDU. It is important to recall that the cells are reassembled and the CRC of a packet is checked at the destination. This means that regardless of when or how many cells are dropped during transmission, all of the remaining cells associated with the packet are still transmitted across the ATM network. Then, when the destination receives the last cell with the end-of-message bit turned on, it reassembles the cells. When an application [such as the Transmission Control Protocol (TCP)] detects an error in the packet due to the lost cells, it requests that the source resend the entire packet. This results in more traffic being sent across the ATM network, creating even more congestion, which makes the problem worse. The congestion problem can grow exponentially out of control.

When congestion occurs, packets are marked and dropped, which causes retransmissions. A disruptive phenomenon called *global synchronization* can occur network wide, particularly with TCP applications. During a global synchronization event, the queues fill and retransmissions occur. If the backoff period (or window) for retransmissions is too close, then when the cells are retransmitted onto the network, the queues again quickly fill and the cells are dropped again.

Even with an ATM network that has been traffic engineered, congestion on the network can occur. The ATM public network also must be configured properly to manage all of the flows from the UNIs and NNIs that it supports. However, effective management of traffic on the ATM network begins with well-managed ATM traffic at the edge devices, such as the Cisco 7200 series router.

Therefore, the primary goal of ATM traffic management is congestion prevention at the UNI interface. If the UNI device can present cells to the public ATM network in a predictable way, then the ATM network can be more efficient and effectively managed.

Traffic Control Functions in ATM Traffic Management

Two of the most important aspects of ATM traffic management are the traffic control functions of shaping and policing. The Cisco 7200 series routers support both of these traffic control functions for ATM.

Traffic Shaping

Traffic shaping at the edge device of an ATM network is considered a preventive measure for the control of network congestion. Traffic shaping controls the flow of traffic onto the network to smoothe out peaks of traffic.

The concept of traffic shaping is particularly relevant for data transfer, which is characterized by variable bursts of traffic onto the network. These bursts create peaks of traffic, and can cause periodic violations to the traffic contract by exceeding the allowable rate of transfer. Bursty traffic patterns also make inefficient use of the network bandwidth.

Figure 1-3 shows the effect of shaping peaks of traffic on the Cisco 7200 series router to produce a smoother, more efficient flow of traffic outbound to an ATM switch.

Figure 1-3 Effect of Traffic Shaping on Traffic Peaks

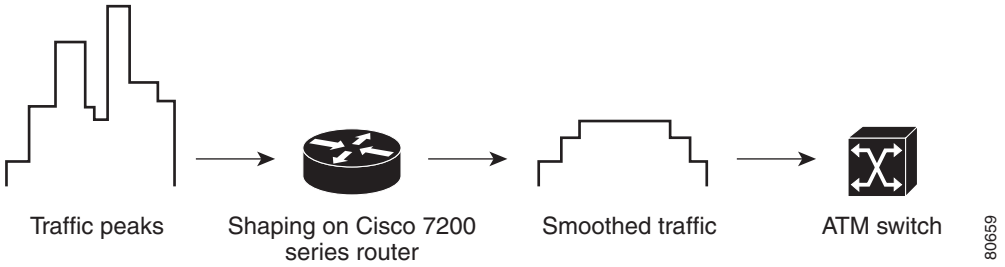

You can implement traffic shaping by configuring the set of traffic parameters associated with a particular ATM service category for a PVC or SVC.

Note Traffic shaping is sometimes referred to as *traffic conditioning*.

Traffic Shaping on the Cisco 7200 Series Router

The Cisco 7200 series router is normally an edge device located on the UNI side of the ATM network. It is very important to configure traffic shaping on the Cisco 7200 series router to effectively control the traffic going onto the ATM network to conform to the traffic contract—but it is only one aspect of the flow.

When you implement traffic shaping, cells are sent onto the network in consistent patterns of cells with fixed, minimum intercell gaps. This rate is based on the traffic shaping parameters that you configure for that PVC or SVC.

However, by shaping the traffic, and with the likely support of multiple service categories with competing transmission characteristics, you effectively create congestion on the router itself—this is where queueing comes in, and also the availability of certain Cisco IOS QoS software features to manage the performance of the queues.

You begin with traffic shaping to configure the performance levels that you want to support on the ATM network. From there, because traffic shaping produces congestion, you need to optimize the applicable hardware and software queues to increase overall performance of the flow of traffic through the router.

Port Adapter Support for Traffic Shaping on the Cisco 7200 Series Router

It is very important to understand that each ATM port adapter on the Cisco 7200 series routers supports different ATM service categories and also implements traffic shaping functions uniquely.

All ATM port adapters support traffic shaping on the Cisco 7200 series routers *except* the PA-A1 ATM port adapter. Although the PA-A1 does support the UBR service category, this is a best-effort service and technically does not perform the function of shaping the traffic over the PVC.

The PA-A3 ATM port adapter and PA-A6 ATM port adapter provides enhanced functionality to the PA-A1 port adapter, and are highly recommended for ATM traffic shaping. The PA-A6 ATM port adapter is an enhanced version of the PA-A3 ATM port adapter and supports twice as many virtual circuits.

For more information about how traffic shaping is implemented on ATM port adapters on the Cisco 7200 series routers, see Chapter 2, "Cisco 7200 Series Architecture and Design for ATM Traffic Management."

For more information about the differences in ATM port adapter support, see Chapter 3, "ATM Traffic Management Hardware and Software Planning."

Benefits of Traffic Shaping on the Cisco 7200 Series Router

Traffic shaping on the Cisco 7200 series router provides the following benefits:

- Smoothes rates of cell transmission to consistent interarrival times, which prevents ATM switches from marking and dropping traffic.
- Allows you to partition your T1/E1, T3/E3, and OC-3 links into smaller, logical channels.
- Helps prevent traffic from any particular VC from consuming the entire interface bandwidth.
- Allows you to match the router's interface transmission rate to the speed of a remote target interface so that you have an even, end-to-end connection.
- Allows any packet drops to occur closer to the source of the traffic, which is more efficient than if cell drops occur on the network side. When packet drops occur at the edge device, the retransmission can be handled much more efficiently by the source and without consuming as much network resource than if one or more cells of that packet are dropped on the ATM network. For further information, see the "Congestion on an ATM Network" section on page 1-10.
- Allows you to buffer some of the traffic waiting to be transmitted to help limit the number of packets that might ordinarily be dropped during any bursts of transmission.
- Allows you to optimize the traffic, rather than having the network side indiscriminately drop cells to force compliance with the traffic contract.

Traffic Policing

Another method used to control traffic on the ATM network is traffic policing. Traffic policing is typically performed by the ATM switch to monitor connections to ensure that they are in conformance with the traffic contract. Policing is important to maintain good performance on the network and to prevent misuse of network resources by users. This helps ensure that all network users get the service levels for which they are paying.

Traffic policing can be done at the UNI or NNI. Service providers typically implement policing at the UNI, on the first switch at the UNI reference point to the ATM network. The switch uses what is known as Usage Parameter Control (UPC) to police connections at the UNI. UPC applies a mathematical formula to determine whether the traffic over a virtual circuit (VC) complies with the contract.

In ATM, part of the policing function is to mark cells as low priority so that if congestion occurs, these cells are dropped. Cells are said to be *marked* as low priority when the Cell Loss Priority (CLP) bit is set to 1. Switches base this marking on cell arrival times and the traffic contract. If cells are found to be in violation of the traffic contract—that is, cells are arriving at a faster rate than agreed upon for the connection—then the cells are marked as low priority and can be dropped.

Figure 1-4 shows cells arriving at a rate above the CDVT to the first ATM switch. The first ATM switch is not congested, so all of the cells above CDVT are marked (the CLP bit is set to 1) and passed to the next switch on the network. The second ATM switch is experiencing high congestion, so it selectively drops any cells with CLP=1.

Figure 1-4 Traffic Policing and Marking on ATM Switches

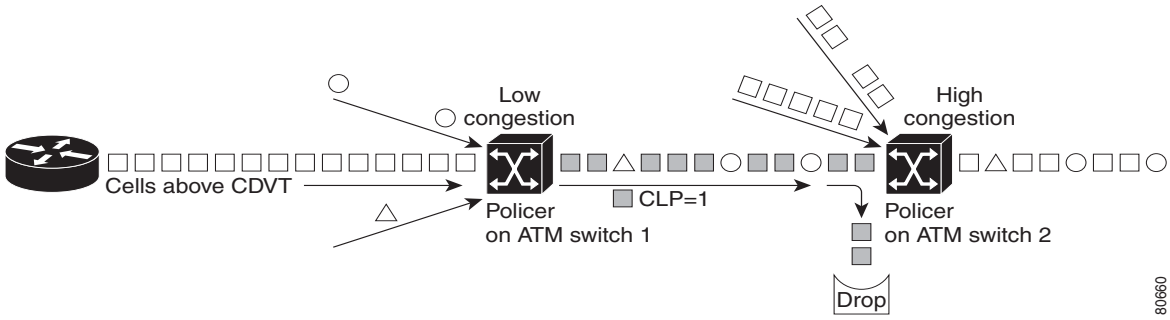

Be aware that some service providers can simply drop nonconforming cells (cells transmitting above the traffic contract), regardless of the level of congestion being experienced on the switch. In this situation, it could be that ATM switch 1 in Figure 1-4 drops the cells, rather than marking and passing along in the network.

Cisco 7200 series routers also support setting the CLP bit through simple marking and also through policing.

Traffic Policing on the Cisco 7200 Series Router

Although the switch on an ATM network commonly implements traffic policing by marking and dropping cells, you can also set the CLP bit using a QoS service policy on a Layer 3 queue on the Cisco 7200 series router. However, the Cisco 7200 series policer never drops ATM traffic based on the CLP bit. It merely marks the packet for CLP (the CLP bit is set in the ATM cell header) and continues.

The difference in this approach is that the switch implements marking at Layer 2, but on the Cisco 7200 series router, you can police and mark IP packets at Layer 3 using QoS service policies.

On the Cisco 7200 series router, you can set the CLP bit for ATM cells in a couple of ways:

- You can configure simple marking for all traffic matching a policy class—this is called *class-based packet marking*. When you implement class-based packet marking, all packets that match the class are marked. Congestion on the VC is not a consideration.
- You can police packets for certain rate criteria using a QoS service policy on a Layer 3 queue. When you implement traffic policing for a VC, the policer determines whether traffic conforms to configured contract values and then, according to your configuration, acts on violations by setting the CLP bit. The advantage over class-based packet marking is that the marking is performed on packets according to rate-conformance criteria, rather than on all packets in a class.

Note The Committed Access Rate (CAR) feature is considered a legacy form of policing and is no longer recommended for use on the Cisco 7200 series routers. Newer, class-based policing mechanisms are now available for some ATM port adapters using the modular QoS CLI (MQC) configuration method.

Benefits of Traffic Policing on the Cisco 7200 Series Router

By marking packets on the router, you can have some control over which traffic is marked and dropped on the network. In this way, cells are not randomly marked as low priority by the switch on the ATM network due to traffic violations.

Design Objectives for ATM Traffic Management

The result of successful ATM traffic management is the efficient transport of traffic through the network with minimization of congestion, while providing fair and sufficient bandwidth access for all service categories when needed.

To efficiently transport mixed traffic through an ATM network, the challenge lies in meeting the following design objectives over the network:

- Prevent congestion on the network by creating a more consistent flow of traffic at the edge device—this is known as *traffic shaping*.
- Control cell delay and cell loss while satisfying the transmission requirements of the different traffic types—this is the basis of QoS for ATM.
- Maximize the use of network bandwidth to fulfill the traffic contract, but prevent a particular application or location from monopolizing the bandwidth—this is part of queue management on the Cisco 7200 edge device; and, on the ATM network, the enforcement of bandwidth usage is known as *traffic policing*.

Related Documentation

The following table provides information about additional resources that you can read to learn more about some of the topics discussed in this chapter:

For more information about:	Refer to the following publications:
ATM technical standards	Approved ATM Forum Specifications
ATM technology and other Cisco Systems products	*Cisco ATM Solutions*, Cisco Press
Cisco IOS QoS software features	*Cisco IOS Quality of Service Solutions Configuration Guide*

Next Steps

This book focuses on how to implement traffic management functions and optimize the overall flow of ATM traffic on the Cisco 7200 series router. It emphasizes traffic management for the PA-A3 and PA-A6 ATM port adapters.

Before you implement traffic shaping to manage your ATM traffic, it is important that you understand how the hardware and software architectural concepts on the Cisco 7200 series router apply to the flow of ATM traffic.

Chapter 2, "Cisco 7200 Series Architecture and Design for ATM Traffic Management," provides detailed explanations of the hardware and software concepts that are applicable to ATM traffic management on the Cisco 7200 series router.

Cisco 7200 Series Architecture and Design for ATM Traffic Management

Revised: December 15, 2005, OL-3274-01

Traffic shaping and traffic policing are the two traffic control functions of ATM traffic management. Both forms of traffic control are supported on the Cisco 7200 series routers. Traffic shaping and the related queueing mechanisms are the primary focus of this book.

Traffic shaping is a form of preventive control and is highly recommended for managing ATM traffic on your edge router. You can consider traffic shaping as only a part—even the beginning—of ATM traffic management on the Cisco 7200 series router.

A result of traffic shaping is congestion on the router because the inbound flow of traffic to the router can be faster than the desired outbound flow of cells to the ATM network. You are effectively holding back congestion from the network, and increasing congestion on the router. Therefore, as part of ATM traffic management, you also want to optimize your router configuration to handle the packets that are waiting to be transmitted.

This chapter discusses the overall flow of ATM traffic on the Cisco 7200 series router, and the different hardware and software architectures that are part of that flow. These hardware and software components work together to affect the overall performance of the flow of a packet through the router, and onto the network as ATM cells. Once you understand how these different areas work together, you can better optimize the flow of traffic onto your ATM network.

This chapter includes the following topics:

Basic Traffic Flow on a Cisco 7200 Series Router

To better understand the scope of managing ATM traffic on a Cisco 7200 series router, it is worthwhile to review the basic flow of ATM traffic from the receipt of packets to their release as ATM cells.

There are certain dependencies and variances in the exact flow. There are aspects of the flow of traffic on the Cisco 7200 series router that apply to every packet, and then there are other aspects that are specific to your Cisco IOS software configuration, or your port adapter. Some variance also depends on how much congestion is being experienced on the router.

As packets are received by an interface and are inserted as ATM cells onto the network, they are processed through the following primary architectures on the Cisco 7200 series router:

- Receive ring
- Switching paths
- Layer 3 hold queue (either an interface hold queue, or a per-VC hold queue)
- Transmit ring
- Segmentation and reassembly (SAR) processor

Figure 2-1 shows the basic flow of traffic destined for an outbound ATM interface and where the different queues come into use on the Cisco 7200 series.

Figure 2-1 Basic Flow of a Packet Through the Cisco 7200 Series Router

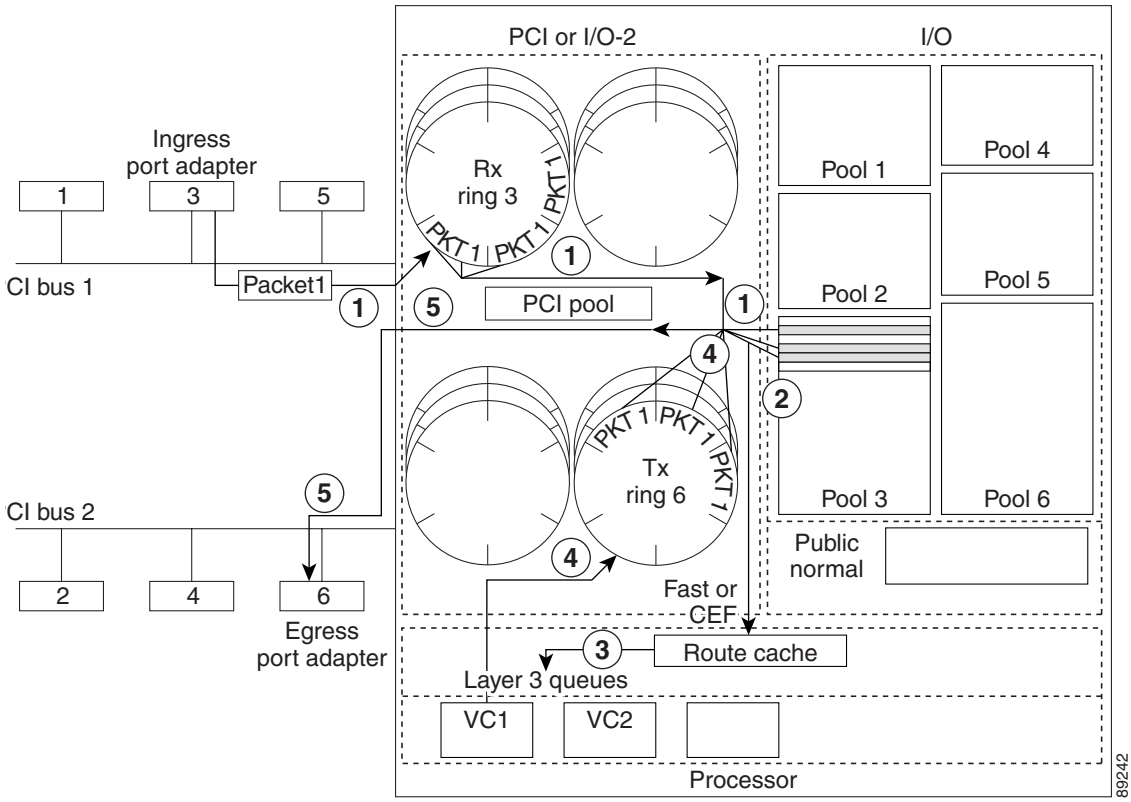

1	The Cisco 7200 series router receives packets on an ingress interface to a hardware queue called the receive ring.
2	The router processes the packets through a switching path: Cisco Express Forwarding (CEF), fast switching, or process switching. **Note** For optimal feature and forwarding performance, CEF switching is highly recommended for per-VC Class-Based Weighted Fair Queueing (CBWFQ) on a PA-A3 or PA-A6 ATM port adapter on the Cisco 7200 series routers. Distributed CEF (dCEF) is required on the Cisco 7500 series routers for per-VC CBWFQ.
3	If the Layer 3 queue is activated (that is, it already has packets in the queue), then the router enqueues the incoming packets to the corresponding Layer 3 queue: • For ATM port adapters other than the PA-A3 or PA-A6, this is an interface-level hold queue, which is common to all VCs supported by the interface. • For PA-A3 and PA-A6 ATM port adapters, this is a per-VC hold queue. **Note** All process-switched packets automatically enqueue to the corresponding Layer 3 queue whether or not there are packets currently enqueued there.
4	From the Layer 3 queue, the router dequeues the packets to a hardware queue called the transmit ring.
5	From the transmit ring, the router dequeues the packets to a SAR processor on the ATM port adapter for segmentation and scheduling of cells onto the network.

Memory Architecture on a Cisco 7200 Series Router

As the Cisco 7200 series router processes packets through the primary structures identified in the "Basic Traffic Flow on a Cisco 7200 Series Router" section on page 2-2, it uses various forms of memory architecture.

This section provides an overview of the types of memory found on the Cisco 7200 series router and describes where some of that memory can be optimized to increase performance through an ATM port adapter.

This section includes the following topics:

Areas of Memory and Types of Storage on a Cisco 7200 Series Router

The Cisco 7200 series router uses several different areas of memory as it processes packets:

- Processor memory—Stores Cisco IOS code, the routing table, and system buffers.
- I/O memory—Stores private interface particle pools and the public pool called *normal*.
- Peripheral Component Interconnect (PCI) memory (also called I/O-2 on the network processing engines [NPEs]-175, NPE-225, NPE-300, NPE-400, and the network services engine [NSE]-1)—Generally a smaller pool of memory that is used for interface receive and transmit rings. Sometimes it also allocates private interface pools for high-speed interfaces.

There are a variety of types of storage used for these three memory areas and memory located on the ATM port adapter hardware, including dynamic random-access memory (DRAM), synchronous static random-access memory (SSRAM), and synchronous dynamic random-access memory (SDRAM).

The exact architecture supported by the router and the size of these storage areas depends upon the type of NPE or NSE processor in use, and also the type of port adapters in use. Some port adapters, such as the PA-A3 and PA-A6 ATM port adapters, support additional memory located on the physical interface for other specialized functions such as SAR processing. For more information about memory located on the PA-A3 and PA-A6 ATM port adapters, see the "What is the SDRAM and SSRAM used for in the PA-A3 and PA-A6 ATM port adapters and why is it important?" section on page 9-4 in Chapter 9, "Frequently Asked Questions."

For additional details about Cisco IOS software architecture and packet processing, refer to the *Inside Cisco IOS Software Architecture* book by Cisco Press.

Particle-Based Memory

The Cisco 7200 series routers use a form of memory architecture based on a unit of storage called a *particle*. Particles are fixed-size segments of storage within an area of memory. Areas of memory that are particle-based consist of a collection of particles of a certain fixed size. For the Cisco 7200 series routers, a particle size of 512 bytes is typical.

A particle buffer, or collection of particles in an area of memory, is also known as a *particle pool*, as shown in Figure 2-2.

Figure 2-2 Particle Pool on the Cisco 7200 Series Router

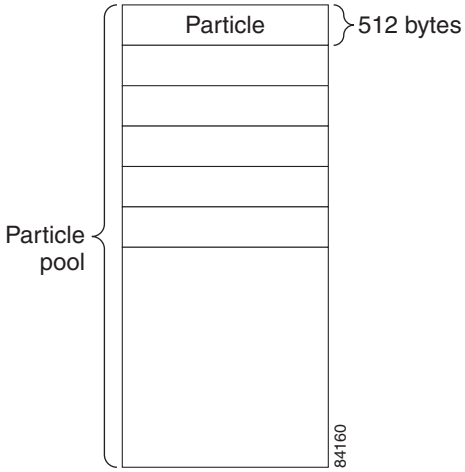

Depending on packet length, packets are stored within one or more particles. Within the particle pool, the location of particles used to store a particular packet can be discontiguous, or nonadjacent. Discontiguous particles that comprise a particular packet are linked together to form a logical packet buffer as shown in Figure 2-3.

CEF and fastswitching methods support discontiguous particle storage for packets. However, for process-switched packets, all packets are collected, or *coalesced*, into a single contiguous buffer in the public normal pool. The Cisco IOS software copies a process-switched packet from its original particle pool (whether the particles are contiguous or discontiguous there), into a single buffer within the public normal pool that is large enough to hold the entire packet.

Note The Cisco 7200 series routers use a Direct Memory Access (DMA) engine to transfer content during coalescing.

Figure 2-3 shows the difference in packet handling based on the switching path when a packet is stored in discontiguous particles.

Figure 2-3 Handling of Discontiguous Particles in Private Interface Pool for Different Switching Paths

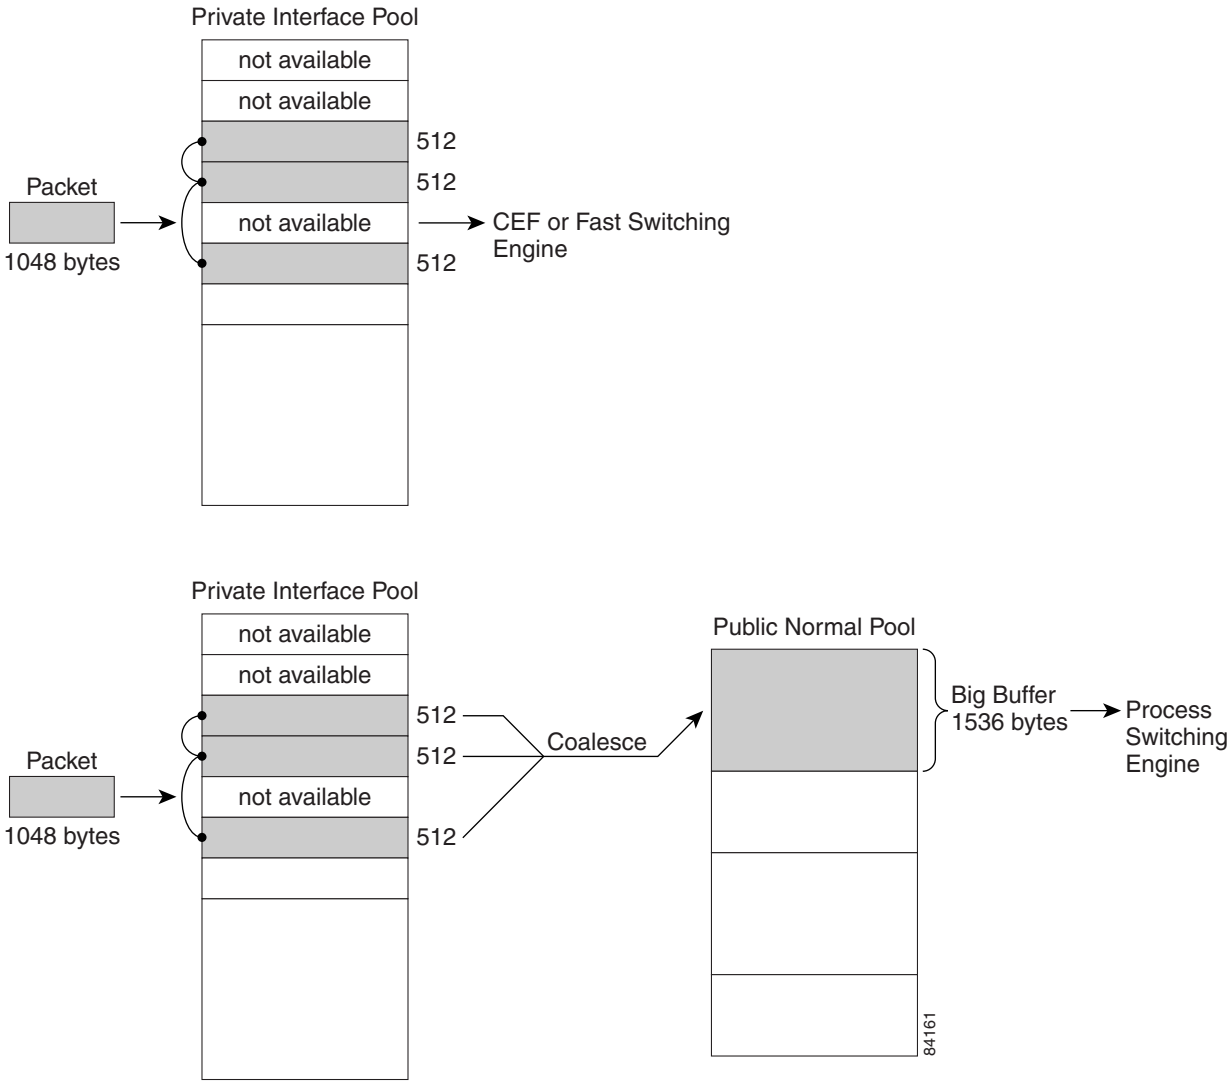

Significance of Particles and Memory Allocation for ATM Port Adapters

It is important for you to understand particles and how they are allocated to store packets so that you can better analyze and interpret possible performance issues if packet drops occur, and to optimize resources for the receipt and transmission of ATM traffic.

To better support certain types of ATM traffic (such as voice, which requires low latencies) and to prevent certain VCs from monopolizing memory resources, it can be necessary to tune the receive ring or transmit ring. The receive ring and transmit ring control structures have a direct relationship with particle allocation.

For more information about what the receive and transmit rings are and how they work, see the "Receive Rings and Transmit Rings" section on page 2-10. For more information about tuning the receive and transmit rings, see the "Per-VC Limits on the Receive and Transmit Rings" section on page 2-15.

Private Interface Pools and Public Pools

Within I/O memory, the Cisco IOS software creates private particle pools for each interface and a public dynamic pool, called the *normal pool*, that all interfaces and processes share. During normal system operation, there are the following two types of pools:

- Private interface pools
- Public pools

Private Interface Pools

Interface pools are considered private because they are not shared by other interfaces or processes. They are available only for storage of packets from a particular physical hardware interface. One interface pool exists for each port adapter on the Cisco 7200 series router, and the size of the pool varies by the type of port adapter and the NPE or NSE.

Private interface pools normally contain a fixed number of particles and are referred to as *static pools*. However, some of the high-speed interfaces supported by the Cisco 7200 series router now have the ability to dynamically allocate more particles for private interface pools.

On the Cisco 7200 series routers, PA-A1 and PA-A2 ATM port adapters have a default interface pool size of 400 particles. Table 2-1 shows the default number of particles within the private interface pools for the PA-A3 and PA-A6 ATM port adapters, which varies by the type of NPE or NSE.

Table 2-1 Default Number of Particles in Private Interface Pool for PA-A3 and PA-A6 ATM Port Adapters

NPE or NSE Model	Particle Size	Default Particles for PA-A3 and PA-A6 ATM PAs
NPE-225 and below	512 bytes	1200
NPE-300 and NSE-1	512 bytes	2400
NPE-400	512 bytes	5200
NPE-G1	512 bytes	5200

Public Pools

The public pool is sometimes referred to as the *normal pool* or the *global pool*. The buffers in the public pools grow and shrink based upon demand. Some public pools are temporary and are created and destroyed as needed. Other public pools are permanently allocated and cannot be destroyed.

The public pools are also used for process switching. This is shown at the bottom of Figure 2-3 on page 2-6.

Cisco IOS software uses five different public buffer pool sizes, or categories, as shown in Table 2-2:

Table 2-2 Default Buffer Sizes in Public Normal Pool

Public Normal Pool Categories	Default Size (in bytes)
Small	104
Middle	600
Big	1536
Large	4520
Very big	5024
Huge	18024

For certain types of interfaces, such as the PA-A1 and PA-A2 ATM port adapters, public pools are used for fallback. Fallback occurs when private interface pools are full and can no longer store incoming packets. When this happens, if the port adapter supports fallback, the router uses available buffers in the public pool to store packets.

Figure 2-4 Fallback to Public Normal Pool for PA-A1 and PA-A2 ATM Port Adapters

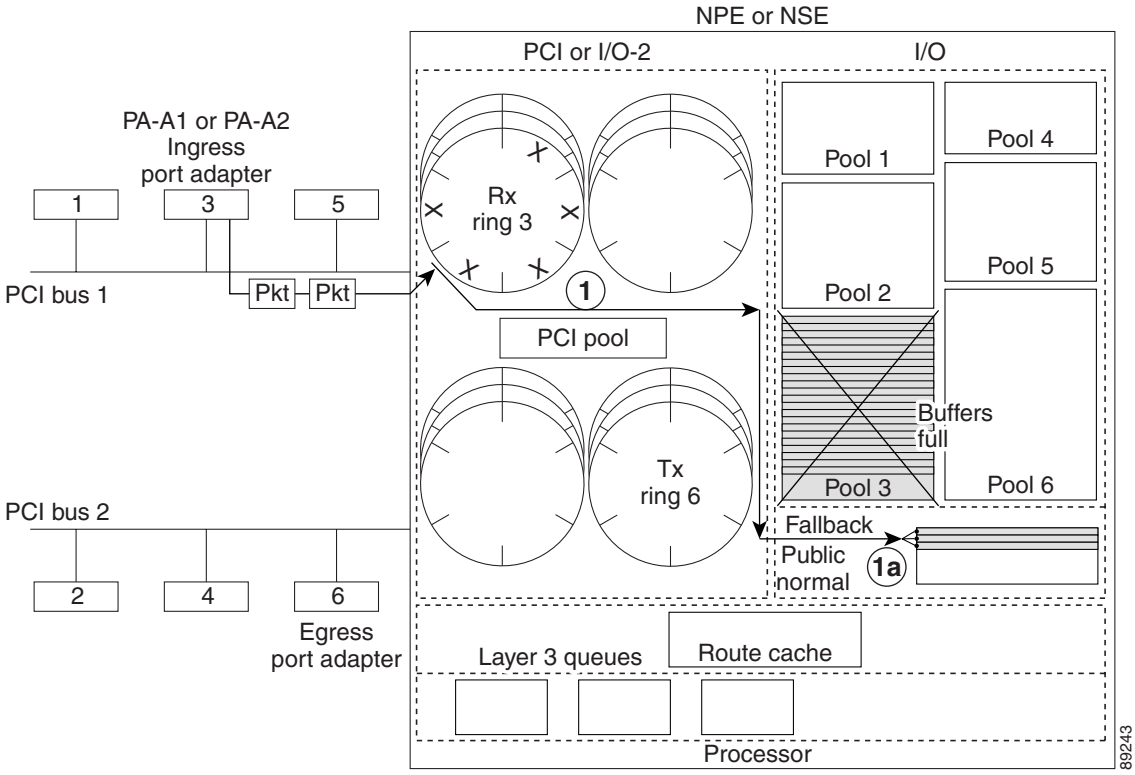

1	The Cisco 7200 series router receives packets on a PA-A1 or PA-A2 ATM port adapter to a hardware queue called the receive ring.
1a	When the private interface pool is full for PA-A1 and PA-A2 ATM port adapters, fallback occurs and the packet is stored in the public normal pool.

The PA-A3 and PA-A6 ATM port adapters do not support fallback to public pools. When the private interface pool is full, or the receive ring limit is reached for a particular PVC on a PA-A3 or PA-A6 ATM port adapter, then packets are dropped (see Figure 2-5). These drops are recorded in the Ignored error field of the **show interface atm** command.

Figure 2-5 No Fallback Support on PA-A3 and PA-A6 ATM Port Adapters

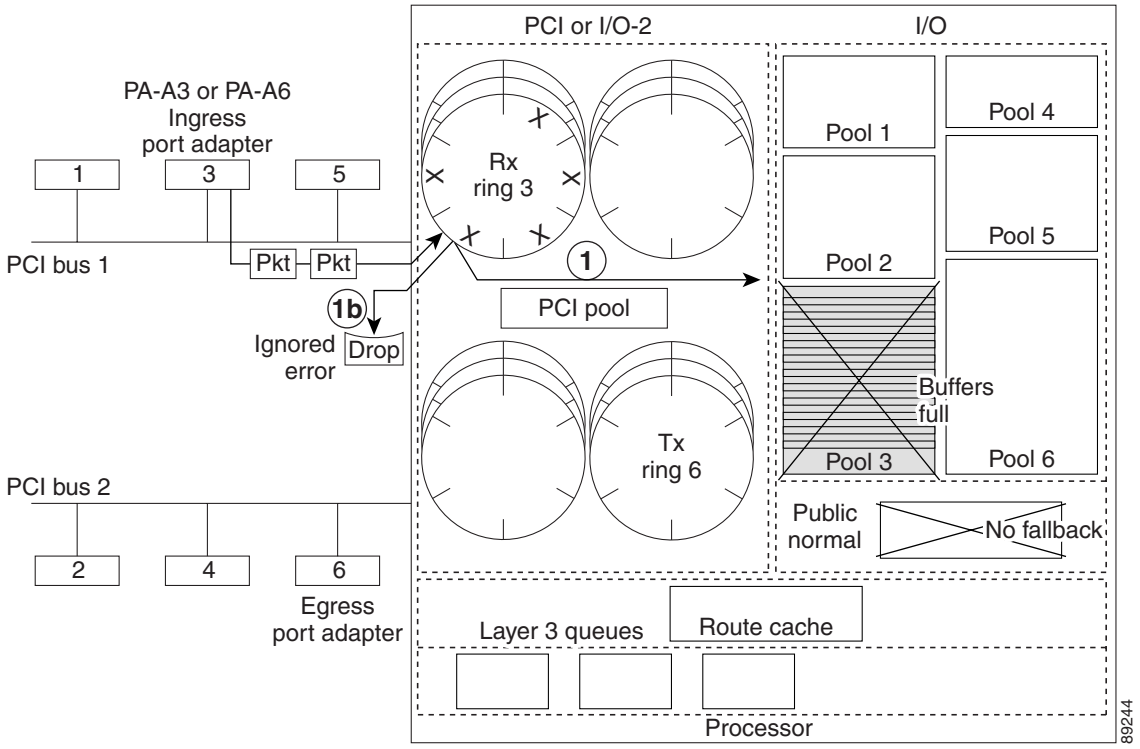

1	The Cisco 7200 series router receives packets on a PA-A3 or PA-A6 ATM port adapter to a hardware queue called the receive ring.
1b	When the private interface pool is full for PA-A3 and PA-A6 ATM port adapters, fallback *does not* occur and packets are dropped.

Monitoring the Buffer Pools

To view the size and usage of the public pools and private interface pools, use the **show buffers** command:

```
Router# show buffers

Public particle pools:
F/S buffers, 128 bytes (total 512, permanent 512):
     0 in free list (0 min, 512 max allowed)
     512 hits, 0 misses
     512 max cache size, 512 in cache
     0 hits in cache, 0 misses in cache
Normal buffers, 512 bytes (total 2048, permanent 2048):
     2048 in free list (1024 min, 4096 max allowed)
     0 hits, 0 misses, 0 trims, 0 created
     0 failures (0 no memory)
```

```
Private particle pools:
     ATM1/0 buffers, 512 bytes (total 1200, permanent 1200):
     0 in free list (0 min, 1200 max allowed)
     1200 hits, 1 misses
```

Note You normally do not need to adjust the size of the buffer pools, and improper settings can adversely impact system performance. Only modify the size of the pools after careful consideration or recommendation by technical support personnel. You can tune the private interface pools and public pools using the **buffers** command for some port adapters, but not on the PA-A3 and PA-A6 ATM port adapters.

Receive Rings and Transmit Rings

Along with public pools and private interface pools, the Cisco IOS software makes use of two packet-control structures called the *receive ring* and the *transmit ring*, also known collectively as *buffer rings*. As shown in Figure 2-1 on page 2-3, these buffer rings reside in PCI or I/O-2 memory depending upon the type of processor.

A unique receive ring and transmit ring structure exists for each port adapter on the Cisco 7200 series router. So, for six port adapters, there are six corresponding sets of receive and transmit rings that reside on the NPE or NSE.

Cisco IOS software and the interface controllers use these rings to maintain the location of particle buffers where incoming packets are temporarily stored for route processing and transmission to the network. The rings consist of media controller-specific elements that point to individual packet buffers that are located elsewhere in PCI (and I/O-2) or I/O memory. Therefore, the rings themselves do not store packets. The rings keep track of the locations in memory where those packets that are under the control of the ring are stored.

Relationship of Buffer Rings to Interface Pools

Each port adapter on the Cisco 7200 series router has a corresponding receive ring and transmit ring, and a private interface pool. As the router receives packets, it stores the physical packet contents in the private interface pool that corresponds to the ingress port adapter.

Some port adapters use only a single entry on the receive ring, as shown in Example 2 in Figure 2-6. This entry links to one or more particles where the owned packet is stored in the private interface pool. Other port adapters, including all of the ATM port adapters on the Cisco 7200 series router, use the same number of ring entries as the number of particles required to store the packet. This is shown in Example 1 in Figure 2-6.

Figure 2-6 Ratio of Ring Entries to Particles in the Private Interface Pool

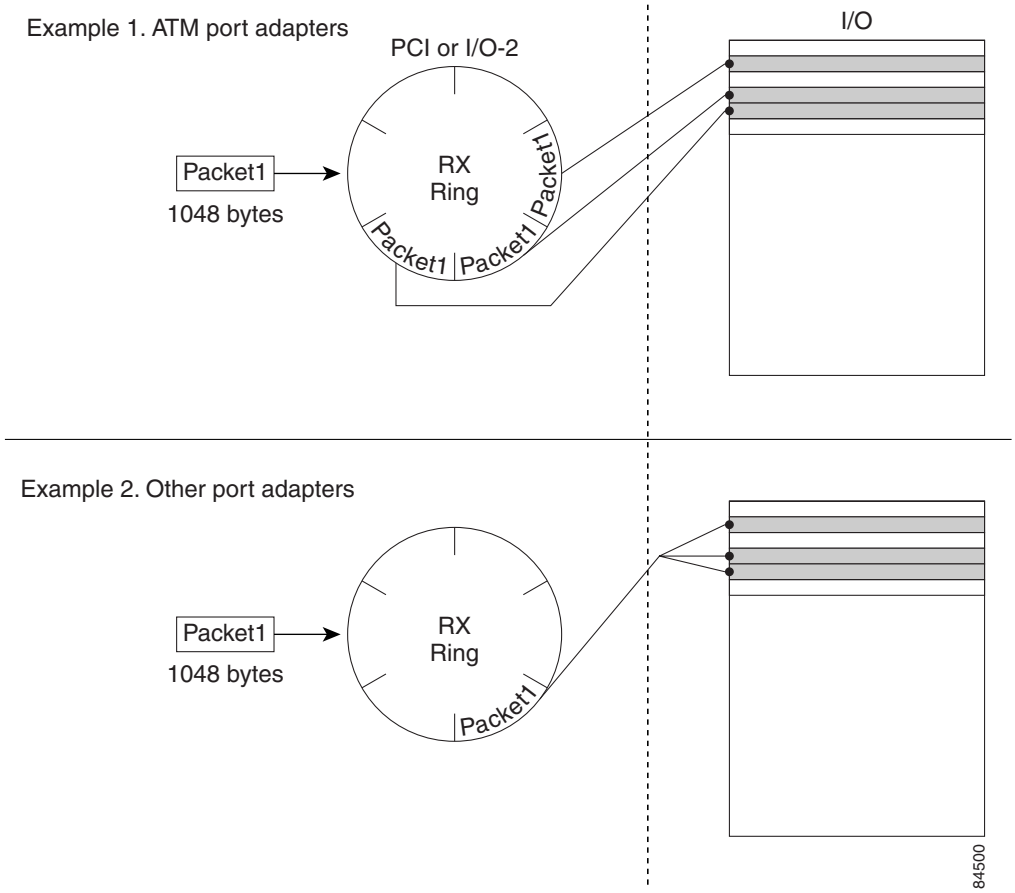

As the router processes the packet all of the way through to the transmit ring, the packet remains in the private interface pool of the ingress interface—unless the received packet coalesces to the public pool (for process switching), or it originally resides in the public pool due to fallback (recall that fallback is not supported by the PA-A3 or PA-A6 ATM port adapters).

In Figure 2-7, port adapter 3 represents an ATM interface that receives a 1048-byte packet. Because the packet is 1048 bytes, and particles are 512 bytes on the Cisco 7200 series routers, this packet requires three 512-byte particles in the private interface pool. Because there is a one-to-one relationship of ring entries to particles, three entries are reserved in Rx Ring3 for the packet. Each ring entry points to a particle location in private interface pool 3 for the corresponding packet content.

Figure 2-7 Receive Ring Entries Link to Private Interface Pool Particles

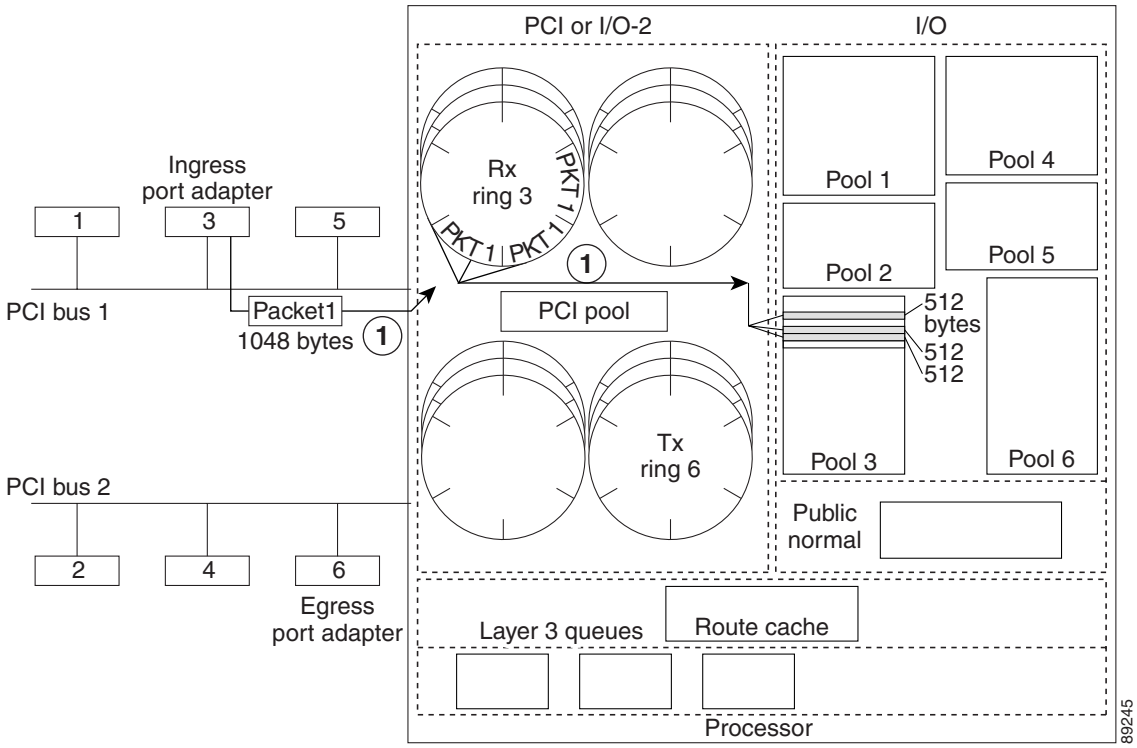

1	The Cisco 7200 series router receives a 1048-byte packet on an ATM port adapter to a hardware queue called the receive ring. The packet requires three 512-byte particles for storage in private interface pool 3. The receive ring creates three entries (one for each particle of the packet) that point to the location of a particle in the private interface pool.

The transmit ring of the egress interface, which corresponds to the interface for the outbound network destination of the packet, ultimately gains control of the packet. It creates one or more ring entries for the packet that link back to the same particles in the interface pool of the ingress interface, where the packet was originally received. The egress port adapter never uses particles associated with its own private interface pool for storing packets to be transmitted.

Figure 2-8 shows the case of a fast-switched or CEF-switched packet, which does not require the particles to be coalesced. After the packet has been switched and processed through any Layer 3 queues, the transmit ring of the destination ATM port adapter (Tx Ring6) creates three ring entries. Each of these ring entries point back to the particles in the same location of the private interface pool (Pool 3) where the original content for the packet resides.

Notice that at this point, the ring entries in Rx Ring3 have been freed and are available for the receipt of new packets over port adapter 3. However, until the transmit ring transfers the contents of the packet to the outbound port adapter, the particles for the packet still reside in private interface pool 3, and are owned by Tx Ring6.

Figure 2-8 Transmit Ring Entries Link to Private Interface Pool Particles of the Inbound Interface

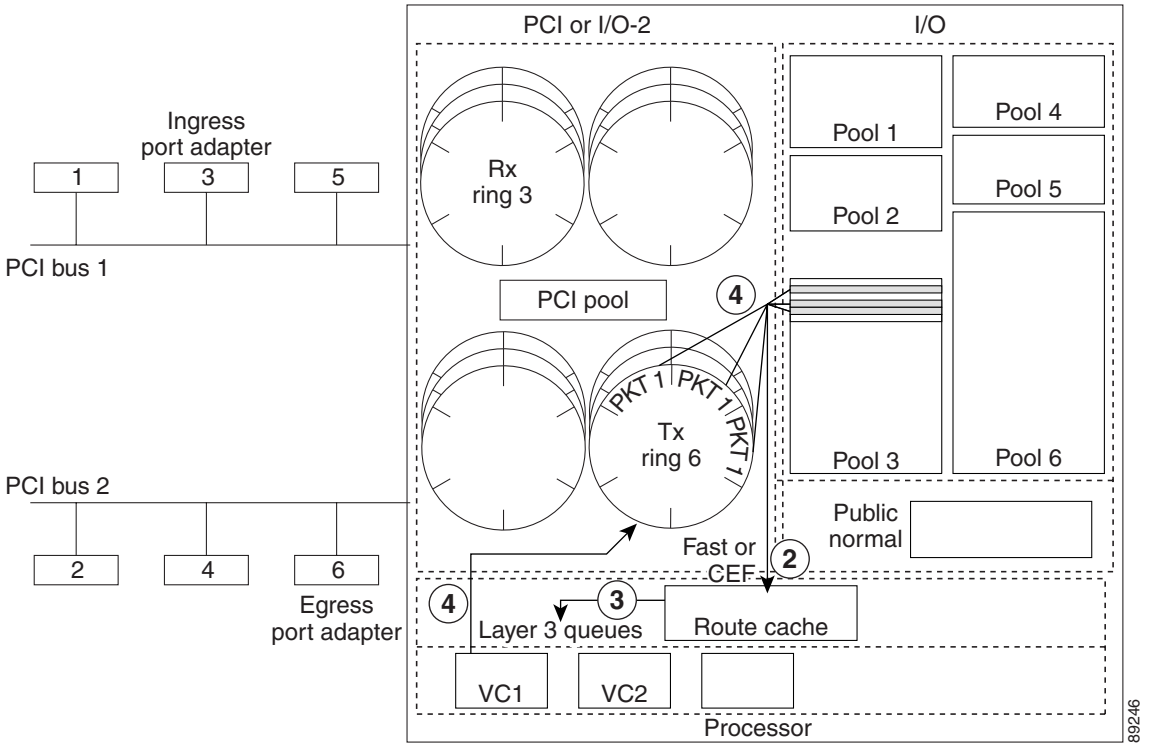

2	The router processes the packets using CEF or fast switching.
3	If the Layer 3 queue is activated (that is, it already has packets in the queue or the transmit ring is full), then the router enqueues the incoming packets to the corresponding per-VC queue.
4	From the per-VC queue, the router dequeues the packets to the transmit ring. The transmit ring creates three ring entries that point to the location of each particle for the packet in private interface pool 3. The receive ring entries are freed and no longer point to the particles in the private interface pool.

In summary, private interface pools store incoming packets. Both receive rings and transmit rings provide links to the ingress interface pool when they have control of a packet.

The concept of a one-to-one relationship of ring entries to particles for the PA-A3 and PA-A6 ATM port adapters becomes relevant if you need to customize the receive ring or transmit ring limits for these port adapters. It is helpful to understand the relationship of ring entries to particles to better understand the methods used by these port adapters to control ring consumption by a VC. For more information about controlling receive ring and transmit ring limits on the PA-A3 and PA-A6 ATM port adapters, see "Per-VC Limits on the Receive and Transmit Rings" section on page 2-15.

PA-A3 and PA-A6 ATM Port Adapter Architecture

The PA-A3 and PA-A6 ATM port adapters are the most advanced port adapters developed for ATM processing on the Cisco 7200 series router. This section discusses the following architectural areas supported by these port adapters:

- Receive Buffer and Transmit Buffer Located on the PA-A3 and PA-A6 ATM Port Adapters
- Per-VC Limits on the Receive and Transmit Rings

Receive Buffer and Transmit Buffer Located on the PA-A3 and PA-A6 ATM Port Adapters

In addition to storage on the NPE or NSE, the PA-A3 and PA-A6 ATM port adapters themselves provide storage for receive and transmit processing. The buffers located on these ATM port adapters receive and store packets or cells for SAR processing. The receive buffer and transmit buffer located on the ATM port adapters work in addition to the standard receive and transmit ring structures and private interface pools on the NPE or NSE.

For the PA-A3 and PA-A6 ATM port adapters, a DMA transfer occurs between the memory located on the ATM port adapter and the private interface pool. This occurs on both the receive and transmit side for these models of ATM port adapters as shown in Figure 2-9.

Figure 2-9 DMA Transfer of Packets Between the Private Interface Pool and the PA-A3 and PA-A6 ATM Port Adapters

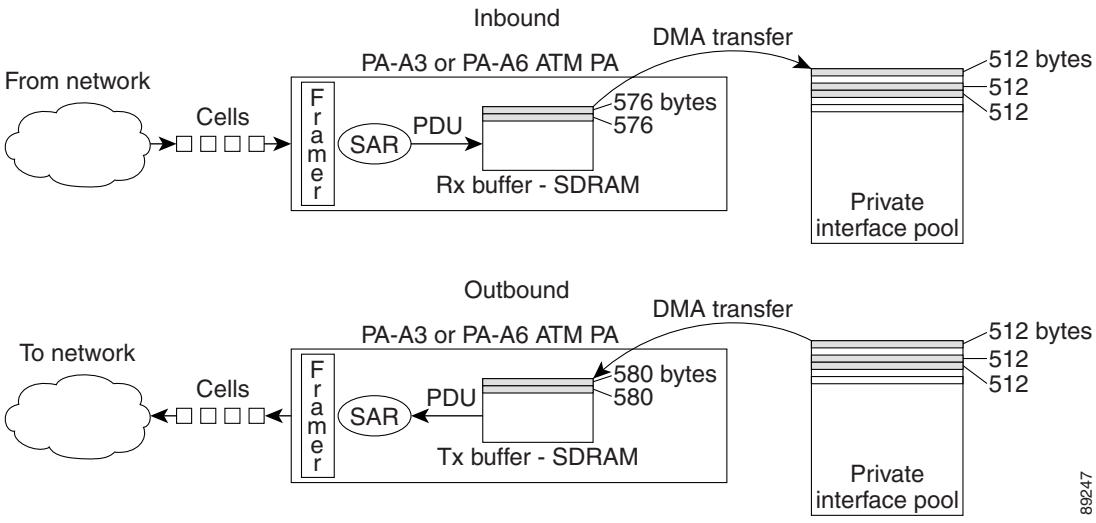

- On the receive side, the port adapter receives cells and reassembles them into packets. It transfers the packets to storage in the private interface pool on the NPE or NSE.
- On the transmit side, the NPE or NSE transfers the packet content from the private interface pool back to storage on the port adapter. The SAR processor on the port adapter segments the packet into cells and, based on the traffic shaping parameters, schedules the cells for transmission onto the network.

Note For efficiency in receiving full cells on the PA-A3 and PA-A6 ATM port adapters, the particle size in the receive buffer located on the PA-A3 and PA-A6 ATM port adapter is 576 bytes, as compared with 512 bytes in the private interface pool on the NPE or NSE. The particle size in the transmit buffer located on the PA-A3 and PA-A6 ATM port adapters is 580 bytes.

Both the receive buffer and the transmit buffer located on the PA-A3 and PA-A6 ATM port adapters work on a first-in first-out (FIFO) basis. On the transmit side, the scheduling of cells onto the network is performed by the SAR processor. The SAR processor implements the appropriate transmission slots based on the configured traffic shaping parameters for the VC.

Per-VC Limits on the Receive and Transmit Rings

The PA-A3 and PA-A6 ATM port adapters provide a way for you to limit the consumption of the receive and transmit ring resources on the NPE or NSE on a per-VC basis. An important concept to understand is that there is still physically only a single receive ring and a single transmit ring. However, the effect of these per-VC limits on the rings is a division of the hardware queue into logical, per-VC queues.

To implement these logical per-VC queues, these port adapters use a method of credits on each VC against a threshold value (the available amount of credit) to prevent any single VC from consuming all of the available resources. On both the transmit and receive rings, the PA-A3 and PA-A6 ATM port adapters use a method of credits that accounts for particles in use. The default limits for both rings is calculated using internal logic based upon configured parameters (such as traffic shaping values) for the VC. However, the PA-A3 and PA-A6 port adapters use a slightly different method of accounting for particles in use based upon whether the limit is for the receive or transmit ring.

This methodology is important to understand if you tune these limits:

- Receive ring—The per-VC receive credits are based on particles in use, not ring entries. This is appropriate because receive ring entries free up before the particles do, and you want the credit check to be based on actual consumption of resource by the VC. The most accurate way to do this on the receive side is to perform a check on particles in use. Accordingly, you configure the **rx-limit** command as a percentage of the private interface pool. In addition, the number of particles is fixed for each NPE or NSE, so with the receive limit as a percentage you can change the processor and still maintain a valid configuration.
- Transmit ring—The per-VC transmit credits are based on the number of ring entries. This equates to particles in use because of the one-to-one relationship of ring entries to particles on the rings. But, for the transmit case, the NPE or NSE does not free a ring entry until the contents of the packet in the particles have been transferred to memory located on the port adapter. Therefore, checking ring entries on the transmit side is effectively the same as checking particles in use for packets awaiting transmission.

  The command-line interface (CLI) range (up to 6000) for the transmit ring limit is based on a credit check for private interface particles against the number of particles available in the transmit hardware buffer located on the PA-A3 and PA-A6 ATM port adapters. Multiple interface pools might be feeding into a single outbound VC on an ATM port adapter, which means that you might have more total particles in use than what a single private interface pool can store. Therefore, the

credit check needs to be against the upper limit of what the transmit hardware buffer can store on the ATM port adapter. Therefore, you configure the **tx-ring-limit** command as a number of ring entries, which are checked against the hardware buffer. The hardware buffer size does not change across NPEs or NSEs, so this value also allows you to maintain a valid configuration if you change processors.

Note Packets are queued to the transmit ring as soon as there is a free particle, even if the packet requires more than one particle to be stored.

Figure 2-7 on page 2-12 and Figure 2-8 on page 2-13 demonstrate these concepts of ring entries and particle allocation. On the receiving end, you can see that initially both the receive ring entries and the particle allocation reflect the resource consumption. Three receive ring entries are in use, which correspond to three particles in use. However, when the transmit ring assumes ownership of the packet, the receive ring entries are no longer allocated to the packet, even though particle resource is still being used.

Note It is important to consider that the ring limits for the receive and transmit side are effectively operating against the same resource—particles within the private interface pool. Therefore, you must be very careful if you plan to tune these limits. Just as with adjustments to the buffer pools, improper settings for the receive ring or transmit ring limits can adversely impact system performance. In this case, adjustments to either side of the ring limits can impact the performance of both receiving and transmitting packets. Only modify the ring limits after careful evaluation of network impact or when recommended by technical support personnel. For more information about tuning these limits, see Chapter 7, "Configuring the Ring Limits on the PA-A3 and PA-A6 ATM Port Adapters."

Layer 3 Software Queues and QoS Processing

Cisco IOS QoS service policies are distinct from the concepts of QoS for the ATM network. Cisco IOS QoS service policies apply to the Layer 3 queues on the NPE or NSE. These QoS service policies do not address cell delay or cell loss over the ATM network itself, which is what defines QoS in the ATM standards.

This section defines the different Layer 3 queues that apply to ATM interfaces and how they are activated. It discusses how Cisco IOS QoS service policies work on Layer 3 queues for ATM traffic and what their relationship is with the transmit ring.

This section includes the following topics:

Software Queueing Terminology

It is helpful to recognize the following terminology usages in Cisco Systems documentation for the Cisco 7200 series software queues:

- A *hold queue* refers to a Layer 3 queue.
- A Layer 3 queue is sometimes referred to by its type. There is more than one type of Layer 3 queue:
  - *Interface queue*—A single Layer 3 queue per ATM interface, which is used for all port adapters excluding the PA-A3 and PA-A6 ATM port adapters.
  - *Per-VC queue*—One of multiple Layer 3 queues for each PVC on the PA-A3 and PA-A6 ATM port adapters. With per-VC queues, the single Layer 3 interface queue is not used for the PA-A3 or PA-A6 ATM port adapters.
- *Fancy queueing* refers to any type of QoS service policy, other than the default, that is configured for a Layer 3 queue.

Activation of Layer 3 Queues

The Cisco 7200 series router activates per-VC Layer 3 queues on the PA-A3 and PA-A6 ATM port adapters whenever congestion builds on an egress interface and outbound traffic cannot be processed to the transmit ring, or onto the hardware queue located on the ATM interface. Traffic shaping is frequently a cause for congestion on the egress ATM interface.

Traffic shaping parameters determine the rate at which the ATM port adapter inserts cells onto the network. If the port adapter receives many large packets, or it receives packets at a rate greater than it can transmit as cells to support the traffic contract, then congestion occurs and queueing is activated.

For PA-A1 and PA-A2 ATM port adapters, there is a single interface hold queue that all PVCs share. In this environment, certain over-subscribed PVCs can use up the available space on the single interface hold queue and prevent other PVCs from their share of that resource. The PA-A2 ATM port adapter supports Layer 2 queues within the port adapter hardware to preserve fairness among VCs. The hold queue at Layer 3 for the interface is for process-switched packets only.

For PA-A3 and PA-A6 ATM port adapters, there is a hold queue for each PVC that is configured for that interface. This environment provides more control and prevents any single over-subscribed PVC from starving other PVCs for transmission resources.

Switching Paths and Layer 3 Queue Activation

With the exception of process-switched packets, whenever entries are available for packets on the transmit ring, packets go directly to the transmit ring on the NPE or NSE and onto the FIFO hardware queue located on the ATM port adapter. Process-switched packets always enqueue to the Layer 3 queue first, before being placed onto the transmit ring, regardless of availability on the ring.

An important thing to keep in mind when designing your network for CEF and fast-switched packets is that QoS service policies will only apply to packets when there is congestion on the ATM port adapter and the transmit ring becomes full. Without congestion, the Layer 3 queueing mechanisms are never activated for CEF and fast-switched packets. This can be a useful thing to remember when tuning the transmit ring limit. For more information, see the "Relationship of Layer 3 Queues to the Transmit Ring" section on page 2-18.

Also, even when a Layer 3 queue is activated, if a service policy is not configured, then the default congestion management mechanism is FIFO (just as it is in the corresponding hardware queue), along with the default of tail drop for congestion avoidance on that Layer 3 queue. Once a Layer 3 queue is activated, any configured methods for congestion avoidance or congestion management apply to the queue accordingly. Therefore, to get the full benefit of Layer 3 queueing, you should configure policies for congestion avoidance and congestion management.

Relationship of Layer 3 Queues to the Transmit Ring

Layer 3 queueing and its relationship to the transmit ring frequently causes confusion, and is an important area for you to understand when optimizing the flow of ATM traffic. The transmit ring capacity and the use of Layer 3 queueing are closely related.

Layer 3 queues are activated for CEF and fast-switched packets when the transmit ring becomes full. When the hold queue is activated (either a single interface hold queue or a per-VC hold queue), the service policies for that queue are applied to enqueue and dequeue packets.

For PA-A3 and PA-A6 ATM port adapters, the transmit ring is considered full for any PVC whenever that PVC reaches the threshold of particles that it is allowed to consume on the transmit ring. This does not necessarily indicate that the entire transmit ring is full for all PVCs. This means that the logical per-VC ring is full. When the transmit ring limit is reached for that PVC, then packets are enqueued to the corresponding per-VC queue. In the meantime, packets from other PVCs can still be placed on the transmit ring for the egress port adapter.

An important thing to consider in the relationship of the Layer 3 queues with the transmit ring is that the hardware queues (both the transmit ring and the buffers located on the ATM port adapter) operate on a FIFO basis. You can control the size of the transmit ring, but you cannot differentiate service levels for packets once they reach these hardware queues. FIFO is the only available queueing method for the hardware queues. Therefore, to achieve any packet differentiation, you need to activate the Layer 3 queue.

It might seem reasonable that the larger the transmit ring size, the more efficient ATM transmission will be. However, when you consider that the transmit ring operates on a FIFO basis, a large transmit ring does not always lead to optimal transmission characteristics for your network traffic. And, it can prevent the Layer 3 queues from activating.

If the transmit ring limit is too large, latency can occur as packets build up on the hardware queue. These packets cannot achieve any priority as they await transmission on a FIFO basis. However, with a smaller transmit ring limit, packets are more readily sent to the hold queue, where they can be differentiated according to configured service policies and gain priority to make it onto the transmit ring ahead of other packets with a lower priority.

Cisco IOS QoS Software

Cisco IOS software provides a comprehensive set of QoS features and solutions to address the diverse transmission needs of voice, video, and data applications and to provide end-to-end QoS services. Cisco IOS software allows you to configure policies to provide differentiated service levels for different classifications of traffic on a Layer 3 queue.

QoS Feature Categories

In Cisco IOS software, QoS features are classified into the following categories:

- Classification and Marking
- Congestion Avoidance
- Congestion Management
- Traffic Shaping and Policing
- Signaling
- Link Efficiency Mechanisms

In Cisco IOS software, there is a subset of QoS features that you also can apply at the ATM PVC level, and these QoS features are collectively referred to as IP to ATM Class of Service (CoS).

IP to ATM CoS

IP to ATM CoS refers to a subset of the overall QoS features available on the Cisco 7200 series router that enables you to specify queueing service policies on a per-VC basis. IP to ATM CoS identifies certain QoS features that can be specifically applied at a more discreet, per-VC level for PVCs on the PA-A3 and PA-A6 ATM port adapters.

When IP to ATM CoS was first introduced, it included the following QoS feature support:

- Weighted Random Early Detection (WRED)
- Class-Based Weighted Fair Queueing (CBWFQ)
- Low Latency Queueing (LLQ)

An important thing to consider is that IP to ATM CoS does not limit the use of other QoS features to support your particular QoS service model. You can still use other QoS features to classify and mark different IP traffic in combination with implementing IP to ATM CoS features at the PVC.

Table 2-3 provides a description of the QoS categories and lists some of the features that are available in the Cisco IOS software in that category. The table also indicates whether IP to ATM CoS support is available in that QoS category.

For guidelines about configuring IP to ATM CoS features on the PA-A3 and PA-A6 ATM port adapters, see Chapter 6, "Configuring QoS on the Layer 3 Queues for the PA-A3 and PA-A6 ATM Port Adapters." For more information about IP to ATM CoS features, refer to the "IP to ATM CoS Overview" chapter of the *Cisco IOS Quality of Service Solutions Configuration Guide*.

Table 2-3 QoS Category Descriptions and IP to ATM CoS Support

QoS Category	Description	Feature Examples	Per-VC Support
Classification and Marking	QoS features that provide packet classification so that you can differentiate traffic into multiple priority levels, or classes of service. It includes those features that allow you to mark IP packets.	Network-Based Application Recognition (NBAR) ATM CLP bit (Layer 3 marking) IP precedence (Layer 3 marking) Differentiated Services Code Point (DSCP) (Layer 3 marking)	No
Congestion Avoidance	QoS features that allow you to anticipate and avoid congestion on your Layer 3 buffers to prevent exceeding the capacity of the queue.	WRED	Yes
Congestion Management	QoS features that allow you to implement priorities for traffic on a congested Layer 3 queue, such as to provide low latencies for delay-sensitive traffic.	LLQ CBWFQ	Yes
Traffic Shaping and Policing	QoS features that allow you to enforce a rate limit (policing) or smooth traffic flow to a specified rate (shaping).	Committed Access Rate (CAR)[1] Generic Traffic Shaping (GTS)[2]	No
Signaling	QoS features that support a way for an end station or network node to signal neighboring nodes to request special handling of certain traffic.	Resource Reservation Protocol (RSVP)	No
Link Efficiency Mechanisms	QoS features that optimize bandwidth usage, such as compression of headers.	Frame-Relay Forum specification for frame fragmentation (FRF.12) Cisco Link Fragmentation and Interleaving (LFI)	No

1. CAR is considered a legacy form of policing on the Cisco 7200 series router. For more information about policing and ATM traffic management, see the "Traffic Policing" section on page 1-12.
2. GTS is generally not recommended for ATM traffic shaping. All ATM port adapters except the PA-A1 port adapters implement native traffic shaping. For more information, see the "Related Documentation" section on page 2-36.

MQC Configuration Architecture

Modular QoS CLI (MQC) is a CLI structure that allows users to create traffic policies and attach these policies to interfaces. For ATM, the MQC architecture extends to application of service policies at the PVC level. MQC provides a more efficient and flexible way to configure QoS service models.

Using MQC to create QoS classes and configure policies involves the following steps:

1. Define a traffic class (**class-map** command).
2. Create a traffic policy and associate the traffic class with one or more QoS features (**policy-map** command).
3. Attach the traffic policy to the interface or PVC (**service-policy** command).

For more information about configuring traffic shaping, see Chapter 5, "Configuring Traffic Shaping on the PA-A3 and PA-A6 ATM Port Adapters." For more information about configuring QoS using MQC, refer to the *Cisco IOS Quality of Service Solutions Configuration Guide*.

Summary of Hardware and Software Queues on the Cisco 7200 Series Router

The Cisco 7200 series router uses both hardware and software queues to manage excess traffic and control its distribution to the physical media for transport onto the network. The hardware and software queues that the router supports, and their implementation, vary by the type of ATM port adapter.

Hardware Queues

Table 2-4 summarizes the hardware queues supported by the ATM port adapters on a Cisco 7200 series router. How packets flow through these queues and details about how these queues operate are discussed in the "Basic Traffic Flow on a Cisco 7200 Series Router" section on page 2-2 and the "Memory Architecture on a Cisco 7200 Series Router" section on page 2-4.

Table 2-4 Hardware Queues Supported by ATM Port Adapters on a Cisco 7200 Series Router

Type of Hardware Queue	Can it be optimized?	Queueing Method
Interface receive ring	Yes—**rx-limit** command (PA-A3 and PA-A6 only)[1]	FIFO
Interface transmit ring	Yes—**tx-ring-limit** command (PA-A3 and PA-A6 only)[1]	FIFO
Buffers local to the port adapter	No	FIFO

1. The **rx-limit** and the **tx-ring-limit** commands specify a particle limit on a per-VC basis. This limit determines the number of private interface pool particles that are available from that ring's FIFO queue for packets received or transmitted over that VC. The particle size itself is fixed and cannot be changed. For more information about using these commands, see Chapter 7, "Configuring the Ring Limits on the PA-A3 and PA-A6 ATM Port Adapters."

Software Queues

Table 2-5 summarizes the Layer 3 software queues that are supported by the ATM port adapters on a Cisco 7200 series router. You can optimize the size of the Layer 3 queues and also implement QoS policies for congestion avoidance and congestion management on those queues using the Cisco IOS software.

Table 2-5 Software Queues Supported by ATM Port Adapters on a Cisco 7200 Series Router

Type of Software Queue	Can it be optimized?	Supported Queueing Methods
Multiple Layer 3 hold queues, one for each VC (PA-A3 and PA-A6 port adapters only).	Yes—**vc-hold-queue** command to configure the size of each per-VC queue.	For congestion avoidance onto each per-VC queue: • Tail drop (default) or WRED For congestion management on the queue: • FIFO (default), CBWFQ, or LLQ
Single Layer 3 hold queue, one for each ATM interface (excluding PA-A3 or PA-A6 port adapters).	Yes—**hold-queue** command to configure the size of the single interface queue.	For congestion avoidance onto the queue: • Tail drop (default) or WRED For congestion management on the queue: • FIFO (default), CBWFQ, or LLQ

Note If you configure CBWFQ for congestion management, then you use the **queue-limit** command to specify the size of the Layer 3 queue.

SAR Processors

The SAR processors are responsible for reassembly of cells into packets on the receive side of the network. On the transmit side of the network, they are responsible for segmentation of packets into cells and scheduling them onto the network according to traffic shaping values.

Those SAR processors that support scheduling within the hardware based upon traffic shaping values that are configurable through the Cisco IOS software are said to support *native* ATM traffic shaping.

Different ATM port adapters implement different types of SAR processors. For more information on SAR processor types, see Chapter 3, "ATM Traffic Management Hardware and Software Planning."

This section includes the following topics:

Native ATM Traffic Shaping and Cisco IOS Traffic Shaping Distinctions

It is important to recognize that the Cisco 7200 series routers support a couple of different traffic shaping methods within their architecture: native ATM traffic shaping and Cisco IOS traffic shaping.

Cisco IOS traffic shaping is not generally used to implement traffic shaping for ATM.

Although the traffic parameters for both methods are configured in Cisco IOS software, native ATM traffic shaping is actually implemented within hardware, and Cisco IOS traffic shaping is implemented within the software, requiring more CPU resource.

Native ATM Traffic Shaping

Native ATM traffic shaping is the preferred method of providing shaping for outbound ATM traffic, and has the following characteristics and benefits:

- Native ATM traffic shaping is a hardware-based implementation that uses a SAR processor to perform the scheduling of cells according to the traffic parameters that you configure in the Cisco IOS software.
- On the Cisco 7200 series routers, native ATM traffic shaping is supported on the PA-A2, PA-A3 (T3, E3, OC-3, or IMA), and PA-A6 ATM port adapters.
- For the PA-A3 and PA-A6 ATM port adapters, native ATM traffic shaping is configured at the virtual circuit (VC) level (this can be per VC, for a VC class, bundle, or range of VCs).

Note You can configure CLI for traffic shaping at a VC bundle. However, the shaping itself does not occur at the bundle level. Shaping still occurs per VC on the PA-A3 and PA-A6 ATM port adapters. There is no shaping implemented for an entire VC bundle.

- For nrt-VBR traffic, the traffic descriptors for native ATM traffic shaping on a PA-A3 or PA-A6 ATM port adapter are better designed for nrt-VBR transmission requirements than are the traffic descriptors for Generic Traffic Shaping (GTS) in the Cisco IOS software.

Note Although the PA-A1 port adapter does implement a SAR processor, native traffic shaping is not available on the PA-A1 ATM port adapter.

Cisco IOS Traffic Shaping

Cisco IOS software supports a couple of forms of traffic shaping features that are part of its QoS feature set, including GTS and Class-Based Shaping:

- GTS—Implements traffic shaping at the interface using the **traffic-shape rate** command.
- Class-based shaping—Implements traffic shaping as part of a service policy using the modular QoS CLI (MQC) structure using the **shape (policy-map class)** command.

In most cases, you do not use class-based shaping to implement traffic shaping on the outbound ATM interface. However, in certain Cisco IOS software releases, you can use the **shape** command within an outbound service-policy with the PA-A3 or PA-A6 ATM port adapters to achieve class-based shaping at Layer 3.

Cisco IOS traffic shaping uses different traffic descriptors than native ATM traffic shaping, and the traffic descriptors for GTS are not as well-suited to support nrt-VBR services.

Note Be careful not to confuse the class-based shaping feature with Class-Based Weighted Fair Queueing (CBWFQ). Class-based shaping is used to configure GTS on a class and is a traffic conditioning feature. CBWFQ is supported and recommended for ATM on the PA-A3 and PA-A6 ATM port adapters for congestion management on a Layer 3 queue.

Understanding Line Rates and Cell Rates

When configuring traffic shaping and monitoring performance of your PVCs, you need to understand some other aspects about the physical interface, including the relationship of the line rates to cell rates and framing.

Each ATM port adapter for the Cisco 7200 series router is named according to its support for a certain line type, or physical interface. This line type represents a line rate (or port speed) that defines the maximum number of bits that can be transmitted and received over the physical interface.

For example, the PA-A3-T3 ATM port adapter supports a single T3 carrier, which uses the Digital Signal, Level 3 (DS-3) North American signaling standard supporting transmission rates of 44.736 Mbps. Table 2-6 shows the standard line rates supported by the ATM port adapters on the Cisco 7200 series routers and their corresponding cell rates. These rates include framing overhead.

Table 2-6 Standard Line Rates and Cell Rates (with Framing Overhead)

ATM Port Adapter	Physical Interface	Line Rate (Mbps)	Cell Rate (Cells per Second)
PA-A3-8T1 IMA	T1 (DS-1)	1.544	3641.51
PA-A3-8E1 IMA	E1	2.048	4830.19
PA-A2-4T1C-T3ATM PA-A3-T3	T3 (DS-3)	44.736	105509.43
PA-A2-4E1XC-E3ATM PA-A3-E3	E3	34.368	81056.60
PA-A1-OC3MM PA-A1-OC3SMI PA-A2-4E1XC-OC3SM PA-A2-4T1C-OC3SM PA-A3-OC3MM PA-A3-OC3SMI PA-A3-OC3SML	OC-3c (STM-1)	155.52	366792.45

Framing Types and Throughput

It is very important to understand that the theoretical line rate does not necessarily represent the actual data throughput that you will see over that interface. The type of framing used over the physical interface affects the maximum possible throughput due to variances in the overhead to implement that framing.

When you configure an interface on the Cisco 7200 series routers, a default framing type is implemented for all traffic over that interface. However, for some port adapters, you can override the default framing type. Each framing type supports a different maximum line rate.

For example, the ATM port adapters that support DS-3 and E3 framing allow you to specify several different framing types using the **atm framing** command. For C-bit ATM Direct Mapping (ADM) framing (the default) on DS-3 interfaces, the maximum line rate is 44.209 Mbps. For C-bit Physical Layer Convergence Protocol (PLCP) framing, the maximum line rate is 40.704 Mbps.

Note Cisco 7200 series routers always transmit traffic with framing overhead. Although you can enable or disable framing overhead on Cisco Systems *switches*, there is no command to enable or disable framing on Cisco Systems routers. It is important to be sure that the switching interface to the router is enabled for framing, and that the framing type corresponds to the framing configuration on the router.

Verifying the Framing Type on the Port Adapter

To verify the framing type on an ATM port adapter, you can use the **show controllers atm** privileged EXEC configuration command.

The following example shows the default framing type of C-bit ADM for the PA-A3 ATM port adapter:

```
Router# show controllers atm 1/0/0
ATM1/0/0: Port adaptor specific information
Hardware is DS3 (45Mbps) port adaptor
Framer is PMC PM7345 S/UNI-PDH, SAR is LSI ATMIZER II
Framing mode: DS3 C-bit ADM
No alarm detected
Facility statistics: current interval elapsed 796 seconds
 lcv      fbe      ezd       pe      ppe     febe     hcse
 ------------------------------------------------------------------------
 lcv: Line Code Violation
 be: Framing Bit Error
 ezd: Summed Excessive Zeros
 PE: Parity Error
 ppe: Path Parity Error
 febe: Far-end Block Error
 hcse: Rx Cell HCS Error
```

Note When an ATM port adapter is using the default framing type on the interface, you cannot verify the framing type using the **show running-configuration** command. However, if you override the default framing type using the **atm framing** command, then you will be able to see the framing configuration in that output.

For more information about framing formats, refer to the TAC Tech Note, "Framing Formats on DS-3 and E3 Interfaces."

PVC Performance

There are several factors that influence the actual performance of ATM traffic on a PVC over the line including the following:

- ATM overhead, which varies by encapsulation type and padding
- Number of VCs using the interface

- CDVT configuration on the ATM switch
- Operation, Administration, and Maintenance (OAM) cells and their interpretation for UPC on the switch

Traffic Shaping Considerations When Establishing Rates

To configure traffic shaping parameters for the ATM port adapters, you typically specify a value in terms of bits per second, which uses the same unit of measure as the line rate. However, be aware that ATM transmission rates over the network actually are implemented according to a total number of cell time slots (or cells per second). Each time slot represents a cell time (in microseconds). Further, ATM switches frequently measure bandwidth according to cell times, not bits per second.

So, it becomes important for you to understand the relationship of line rates to cell rates to understand how the SAR scheduler transmits cells and to be sure that your ATM connection between the router and switch are configured to support compatible rates.

Also, when you configure traffic shaping on the Cisco 7200 series router, you need to consider the effective line rate *without* framing overhead. This is because cell scheduling is established before the addition of framing overhead by the framer on the port adapter. If you were to base your traffic shaping parameters on the full line rate, then you might oversubscribe the line rate when the framing overhead is added.

Table 2-7 shows the corresponding cell rates without framing overhead by physical interface type.

Table 2-7 Cell Rates without Framing Overhead

ATM Port Adapter	Physical Interface	Cell Rate Without Framing Overhead (Cells per Second)
PA-A3-8T1 IMA	T1 (DS-1)	3622.64
PA-A3-8E1 IMA	E1	4528.30
PA-A2-4T1C-T3ATM PA-A3-T3	T3 (DS-3)	96000.00
PA-A2-4E1XC-E3ATM PA-A3-E3	E3	80000.00
PA-A1-OC3MM PA-A1-OC3SMI PA-A2-4E1XC-OC3SM PA-A2-4T1C-OC3SM PA-A3-OC3MM PA-A3-OC3SMI PA-A3-OC3SML	OC-3c (STM-1)	353207.55

On the Cisco 7200 series routers, some traffic parameters are configured in bits per second, but others, such as the maximum burst size (MBS) for the nrt-VBR service category, are configured in terms of the number of cells. A *cell time* represents the amount of time for the transmission of one cell over the line in a cell time slot. To appropriately configure the MBS, you need to consider cell times as well as the line rate. For more information, see the "Determining Cell Times" section on page 2-27.

When configuring traffic shaping on PA-A3 and PA-A6 ATM port adapters, it is also important to consider both the CDVT values on the switch and the use of OAM. Routers and switches treat OAM cells differently. When implementing shaping, the SAR on the PA-A3 and PA-A6 ATM port adapters considers data cells only. When enforcing rates using UPC, the switch typically counts both OAM cells and data cells.

You should also be sure that the router and the switch are basing their rate settings and policing on the same cell size. Some processors interpret rates based on 48-byte cells and others use 53-byte cells. When implementing PCR and SCR on ATM port adapters for Cisco Systems routers, the SAR accounts for the 5-byte ATM cell header, AAL5 padding, and an AAL5 trailer.

Note Transmission reporting in **show** command output varies between routers and switches. Routers typically provide ATM traffic counts in terms of packets (typically AAL5 packets) and sometimes rates (in bps), whereas switches frequently provide cell counts. You can use network management MIBs to analyze utilization. For more information about measuring PVC utilization, refer to the TAC Tech Note, "Measuring Utilization on ATM PVCs."

Converting Line Rates to Cell Rates

Every physical line rate can be represented by a number of cells per second, or cell time slots. To determine the number of cell time slots that a physical line can support, you need to divide the line rate by the size of each ATM cell.

The best way to do this conversion is to represent the ATM cell size as a number of bits per cell, because the physical line rates are defined in bits per second. A 53-byte ATM cell is equivalent to 424 bits per cell (53 bytes x 8 bits per byte). The bps unit of measure for the line rate, divided by bits per cell yields cells per second.

To convert line rates to cell rates, use the following formula:

Line rate (bits per second) / 424 (bits per cell) = Number of cells per second (cell time slots)

For example, the conversion of a T1 (DS-1) line rate at 1.544 Mbps is determined by the following equation:

154400 / 424 = 3641.51

Therefore, an ATM port adapter that supports a T1 physical line rate has approximately 3642 cell time slots as its maximum bandwidth. This calculation represents the number of cell time slots with framing overhead on the line. As Table 2-7 on page 2-26 shows, the cell rate is slightly less than this without framing overhead.

Determining Cell Times

It is useful to understand the concept of an ATM cell time. The amount of time that it takes for one ATM cell to transmit within a time slot over the interface is called the *cell time*. You can calculate this value as follows:

1 (cell) / *ATM cell rate* (cells per second) = ATM cell time

Here is a sample calculation of cell time for a DS-1 link using the ATM cell rate without framing:

1 / 3622.64 = .00027604 seconds, or 276.04 microseconds per ATM cell

Traffic Shaping Algorithms Used By the SAR Processors

The final piece of the architecture that you should understand in the transmission of ATM cells is the traffic shaping algorithms and scheduling used by the SAR processors. Every ATM port adapter uses a SAR processor and implements a certain scheduling algorithm for the transmission of cells onto the network. However, the SAR processor and scheduling algorithm varies by the type of ATM port adapter.

For this discussion, the focus is on the scheduling algorithms used by the most advanced ATM port adapters supported by the Cisco 7200 series routers—the PA-A3 and PA-A6 ATM port adapters. The PA-A3 ATM port adapters (*except* the PA-A3-OC12 model, which is not currently supported on the Cisco 7200 series router) and the PA-A6 ATM port adapters use the LSI ATMIZER II+ SAR processor. These port adapters support the Generic Cell Rate Algorithm (GCRA), commonly known as the Leaky Bucket algorithm.

Note The PA-A3-OC12 ATM port adapter does not use the same SAR processor and scheduling algorithm as the other models of the PA-A3 ATM port adapters, and it is not currently supported on the Cisco 7200 series routers. To find the latest information about port adapter support on different platforms, you can use the "Software Support for Hardware" feature of the Software Advisor tool on Cisco.com. For more information about hardware and software planning, see Chapter 3, "ATM Traffic Management Hardware and Software Planning."

GCRA (Leaky Bucket) on the PA-A3 and PA-A6 ATM Port Adapters

This section provides a brief introduction to GCRA, or the Leaky Bucket algorithm, and how it is used to control the transmission of cells for a PVC based on the traffic shaping parameters. The PA-A3 (*except* the PA-A3-OC12) and PA-A6 ATM port adapters use GCRA to implement the proper shaping for the PVC, as the PVC is scheduled for servicing within a time slot on the line.

For more information about the technical details of GCRA, refer to the ATM Forum Traffic Management specifications.

Before the SAR processor can place cells in a transmission slot and send them to the framer on the port adapter, it uses GCRA to control which cells are eligible for transmission on that PVC. This scheduling algorithm uses a token architecture to control access to the network for each PVC. Simply put, a PVC can only transmit a cell when that PVC has a token available in its transmission bucket.

The algorithm uses the concept of a bucket to represent an accumulation of tokens, or cell transmission credits, for a PVC. The traffic shaping parameters determine the rate at which tokens replenish the bucket, and the maximum number of tokens that can be used to burst cells onto the network.

With GCRA, the Sustainable Cell Rate (SCR) determines the rate at which tokens fill the bucket as shown in Figure 2-10. The maximum number of tokens that can be available at any time is determined by the Maximum Burst Size (MBS), and can be thought of as the size of the bucket.

Figure 2-10 Tokens Fill Bucket at SCR Up to Size of MBS

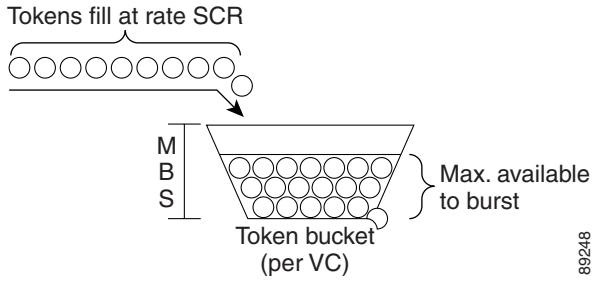

If a PVC is idle and does not transmit for a period of time, then tokens accumulate in the transmit bucket. When the PVC again has data to transmit, it can burst a number of cells less than the configured MBS.

Figure 2-11 shows that a PVC can use the accumulated tokens to burst up to the Peak Cell Rate (PCR) until the bucket is empty, at which point tokens are again replenished at the SCR.

Figure 2-11 PVC Can Use Available Tokens to Burst Up to the PCR

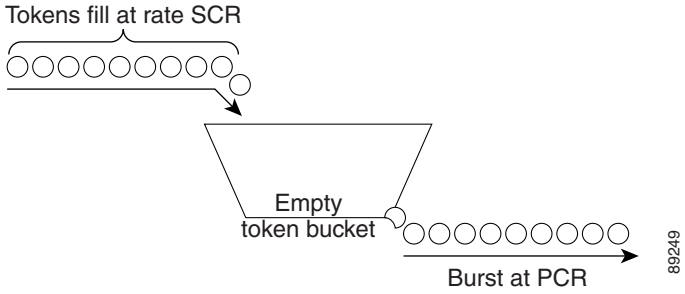

When the PVC has more cells to transmit than the allowable MBS, then the port adapter schedules the cells in an interval of time slots according to the traffic shaping parameters. For more details about how the PA-A3 and PA-A6 ATM port adapters implement scheduling, see the "Scheduling on the PA-A3 and PA-A6 ATM Port Adapters" section on page 2-29.

Scheduling on the PA-A3 and PA-A6 ATM Port Adapters

Based on the traffic shaping values, and the transmit priority for the PVC, the scheduler within the SAR processor determines which PVCs have access to the cell time slots on the physical interface, and it uses GCRA to enforce the shaping. The maximum line rate depends on the physical interface and its framing, which in turn determines the total number of cell time slots available. Each time slot division represents a cell time.

The PA-A3 and PA-A6 ATM port adapters implement these time slots using a calendar table (32K entries) and by keeping track of the list of PVCs to be serviced in each slot. If there is no traffic awaiting transmission from a particular PVC, then the SAR processor does not attach the virtual circuit descriptor (VCD) identifying that PVC to the calendar table for scheduling.

Note A VCD is used only internally by the router to uniquely identify a PVC.

Understanding Intercell Gaps

When considering scheduling, you should understand the concept of an intercell gap (ICG). An ATM port adapter can only send out cells at a fixed, minimum ICG according to the line rate (without framing) supported by the interface. When you configure SCR and PCR traffic shaping parameters on a PVC, the scheduler within the SAR processor translates these values into an ICG, which determines the interval of time slots that should be scheduled for that PVC to maintain the shaping configuration without bursting.

Figure 2-12 shows an example of time slots for a DS-1 physical interface. From Table 2-7 on page 2-26, you know that there are approximately 3622 time slots prior to framing that are available over the DS-1 physical interface. And, the corresponding cell time is 276.04 microseconds (from the "Determining Cell Times" section on page 2-27), which represents the minimum ICG for that interface.

Figure 2-12 Intercell Gap and Time Slots on a DS-1 Physical Interface

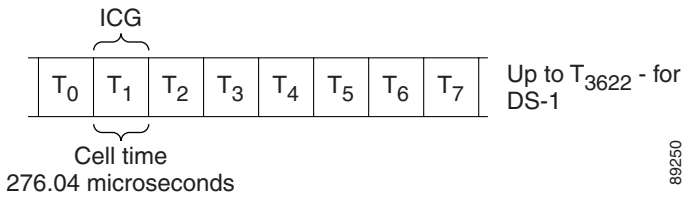

Transmission of cells with an ICG equal to 1/PCR is called bursting, and is characteristic of non-real-time service categories such as nrt-VBR. Real-time service categories generally are characterized by smaller and more evenly-distributed ICGs. However, real-time VBR can also burst cells in clumps.

VBR service categories are characterized by the ability to accommodate "bursty" traffic. To do this, the PA-A3 and PA-A6 ATM port adapters send out cells according to two different ICGs. When bursting, up to the MBS-number of cells can be sent with an ICG of 1/PCR. This duration is controlled by the concept of tokens available within a "leaky bucket." If no more tokens are available, cells are sent with an ICG of 1/SCR. When the offered traffic is below SCR, then the bucket is progressively replenished with tokens and the PVC is able to burst again.

In contrast, the CBR service category (often used for real-time services) is only characterized by one cell rate and it uses an ICG of 1/PCR. There is no concept of bursting in this service category because cells are sent at a constant rate.

Figure 2-13 shows an example of a time slot interval for VC1 with an ICG of 3. This ICG means that if a cell transmits in time slot T_1, then the next scheduled time slot would be T_4, T_7, and so on. Therefore, the scheduled time slots begin at T_n, and continue at an interval of $T_{n + ICG}$. The scheduler maintains this time slot interval for the PVC until there are no longer any cells to be transmitted.

Figure 2-13 Example of ICG=3 for VC1

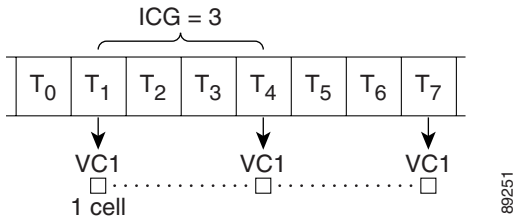

Scheduling Multiple PVCs

An ATM port adapter normally services multiple PVCs. Because the ATM port adapter breaks down the available bandwidth into evenly-spaced time slots, and these time slots are allotted to service the different PVCs, the scheduler within the SAR processor acts as a cell multiplexer by merging the traffic from several sources onto the line.

Some PVCs might share the same traffic shaping characteristics, and some might differ. In addition, the PVCs have a transmit priority (either by default, or configurable). So, when two or more PVCs compete for access to the same cell time slot, the port adapter considers the priority of the PVC and decides the order in which the cells are serviced.

Collision Handling

Overbooking an ATM port adapter can produce competition for cell time slots. However, cell collisions can happen at any time and are not only due to overbooking conditions. Overbooking increases the number of collisions and inevitably leads to drops, even on VCs that are not fully used. An ATM port adapter is overbooked for any type of service category when the sum of all SCRs for the PVCs is greater than the line rate. When the SAR processor attempts to schedule more than one PVC for a particular time slot, a collision occurs.

When a time slot collision occurs, the ATM port adapter considers the priority of the PVC and determines where to schedule the cells. Every ATM port adapter implements a prioritization scheme that varies by platform. When a collision occurs, the port adapter decides which cell transmits in the time slot, and bumps the deferred cell to the next adjacent time slot.

But, if another cell is already scheduled in the adjacent time slot, another collision occurs. Where does the bumped cell get placed, and how does the PVC priority affect the hierarchy? The prioritization scheme implemented by the port adapter determines how this is done.

The PA-A3 and PA-A6 ATM port adapters support one of two different possible algorithms to resolve time slot conflicts between two PVCs:

- Tail-insertion algorithm—Bumps cells to the bottom of the link list for a time slot according to PVC priorities.
- Head-insertion algorithm—Bumps cells to the top of the link list for a time slot according to PVC priorities.

The tail-insertion algorithm was the original scheduling algorithm used by the PA-A3 ATM port adapter, but has since been replaced by the head-insertion algorithm. With either algorithm, the PVC priority also affects the hierarchy of the competing cells.

Note The head-insertion algorithm was first implemented in the following Cisco IOS releases: 12.0(21)S, 12.0(21)ST, 12.1(11), 12.1(14)E, 12.2(6), 12.2(14)S, 12.2(8)T, 12.2(15)B.

The difference between the two collision algorithms is where the SAR processor places a cell in the time slot's link-list hierarchy (either at initial scheduling, or later due to bumping). And, the actual placement within a link list varies by whether the PVCs have the same or differing priorities.

The following list summarizes the SAR processor's behavior during collisions, by algorithm:

Note The software only uses one algorithm or the other. It does not support both head-insertion and tail-insertion at the same time.

- Using tail-insertion, the SAR processor performs the following actions:
  - Schedules a bumped PVC after previously scheduled PVCs of the same priority.
  - Schedules a bumped PVC ahead of previously scheduled PVCs if the priority of the bumped PVC is higher than the scheduled PVCs. But, the SAR processor continues to link the bumped PVC after any existing linked PVCs of the same priority.
- Using head-nsertion, the SAR processor performs the following actions:
  - Schedules a bumped PVC ahead of previously scheduled PVCs of the same priority.
  - Schedules a bumped PVC ahead of previously scheduled PVCs if the priority of the bumped PVC is higher than the scheduled PVCs. The SAR processor also links the bumped PVC ahead of any existing linked PVCs of the same priority.

The following examples illustrate these differences:

- Example 1: Collisions for PVCs with the Same Transmission Priority, page 2-32
- Example 2: Collisions with PVCs of Different Priorities, page 2-33

Example 1: Collisions for PVCs with the Same Transmission Priority

In this example, consider four VCs that each have an ICG of 2 and the same PVC priority. Begin with all four PVCs presenting a cell for transmission at the same time, and the port adapter initially schedules the first four time slots (T_1, T_2, T_3, and T_4).

Recall that only one cell can transmit in any time slot. Therefore, when all four VCs need to schedule a cell for transmission, the port adapter bumps the subsequent cells into adjacent time slots. In our example, we assume that the first four time slots, T_1, T_2, T_3, and T_4, are empty.

Figure 2-14 shows how tail-insertion scheduling occurs.

Figure 2-14 Tail-Insertion Time Slot Scheduling for PVCs with Same Transmission Priority

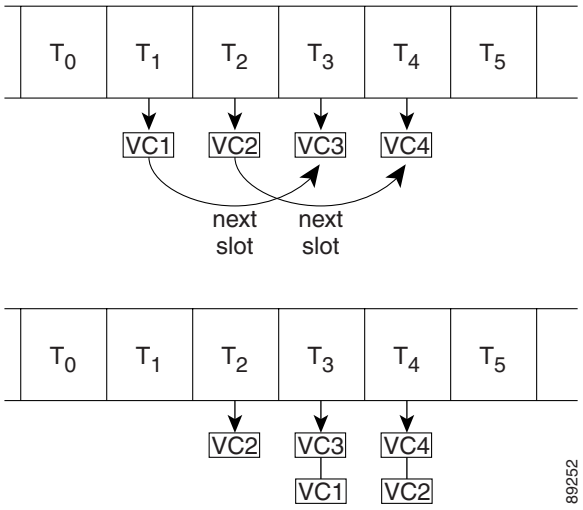

- At time T_1, the SAR processor takes the following actions:
  - Transmits a cell from VC1.
  - Reschedules VC1 to transmit at time slot T_3 ($T_{1 + 2 (ICG)}$).
  - Links VC1 at the tail-end of the T_3 link list (VC3 already has a cell scheduled at T_3).
- At time T_2, the SAR processor takes the following actions:

  - Transmits a cell from VC2.
  - Reschedules VC2 to transmit at time slot T_4 ($T_{2 + 2 (ICG)}$).
  - Links VC2 at the tail-end of the T_4 link list (VC4 already has a cell scheduled at T_4).
- At time T_3, the SAR processor takes the following actions:
  - Transmits a cell from VC3.
  - Reschedules VC3 to transmit at time slot T_5 ($T_{3 + 2 (ICG)}$).
  - Bumps the remaining entry in the link list for T_3 to the next time slot, T_4 (can only transmit a single cell in a time slot).

This means that the cell for VC1 now moves to T_4. With tail-insertion, VC1 again moves to the bottom of the link list for T_4, as shown in Figure 2-15.

Figure 2-15 Tail-Insertion Continues to Move VC1 to Bottom of Link List

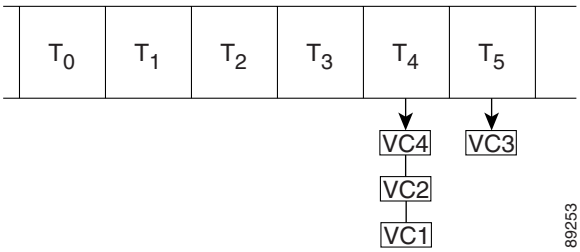

Figure 2-16 shows what happens if the head-insertion algorithm is used instead of tail insertion at the same point in time.

Because all of the PVCs have the same transmission priority, the SAR processor places the bumped cell ahead of the previously link-listed PVCs for that time slot. Therefore, instead of VC1 continuing to get bumped, it is first in the link list so that it will now have priority to transmit in the T_4 time slot.

Figure 2-16 Head-Insertion Moves VC1 to Top of Link List

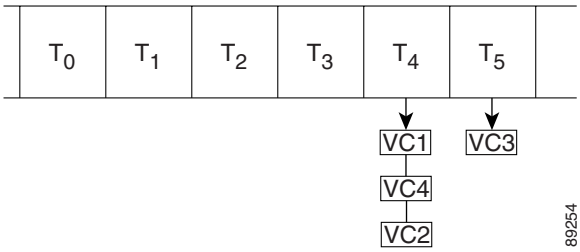

Example 2: Collisions with PVCs of Different Priorities

In this example, observe the affect of collisions for five PVCs of differing priorities. In the example, the PVCs are identified as A1, B2, C3, D1, and F2, where A, B, C, and so on identifies the PVC, and 1, 2, and 3 represents the priority of the PVC. As in the first example, each of the PVCs has an ICG of 2.

Also consider that with transmission priorities, the lower numbers have the higher priority. Therefore, A1 has a higher priority than B2 or C3.

Note For both the tail-insertion and head-insertion algorithms, the SAR processor never places a lower priority PVC ahead of a higher priority PVC in its link list for a time slot.

Figure 2-17 shows the initial link-list hierarchy of PVCs A1, B2, C3, D1, and F2 and their corresponding time slots.

Figure 2-17 Initial Link-List Hierarchy for PVCs with Different Transmission Priorities

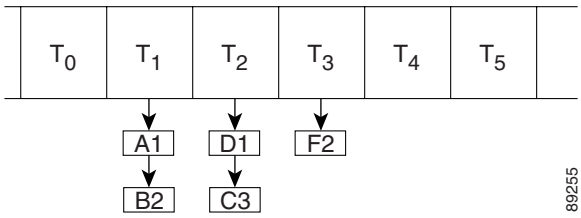

- At time T_1, the SAR processor takes the following actions:
  - Transmits a cell from A1.
  - Reschedules A1 to transmit at time slot T_3 ($T_{1 + 2 (ICG)}$).
  - Links B2 between D1 and C3 in the T_3 link list.

Figure 2-18 shows that the SAR processor places A1 ahead of F2 in time slot T_3 due to its priority. In similar fashion, it bumps B2 between D1 and C3 in time slot T_2. Notice that the SAR processor always schedules the higher priority PVCs ahead of the lower priority PVCs.

Figure 2-18 PVC Scheduling According to Priority

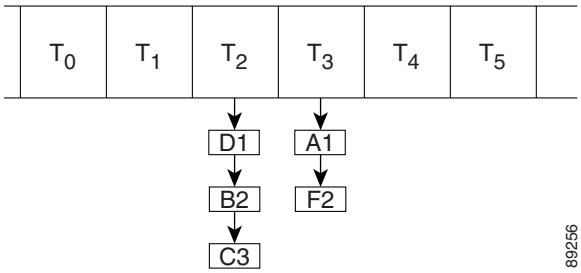

- At time T_2, the SAR takes the following actions:
  - Transmits a cell from D1.
  - Reschedules D1 to transmit at time slot T_4 ($T_{2 + 2 (ICG)}$).
  - Links B2 and C3 at the tail-end of the T_3 link list.

Figure 2-19 shows that when the SAR processor bumps B2 and C3 to time slot T_3 using tail-insertion, it places B2 below F2 (same priority) and C3 at the end.

Figure 2-19 Link-List Hierarchy of PVCs with Same Priority Using Tail-Insertion

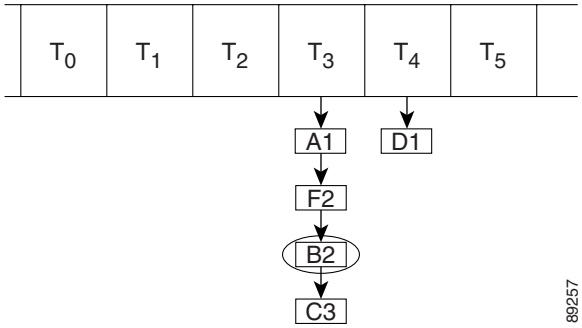

However, Figure 2-20 shows how the link-list hierarchy appears when the SAR processor uses head-insertion. With head-insertion, the SAR processor places B2 ahead of F2, and continues to place C3 at the end.

Figure 2-20 Link-List Hierarchy of PVCs with Same Priority Using Head-Insertion

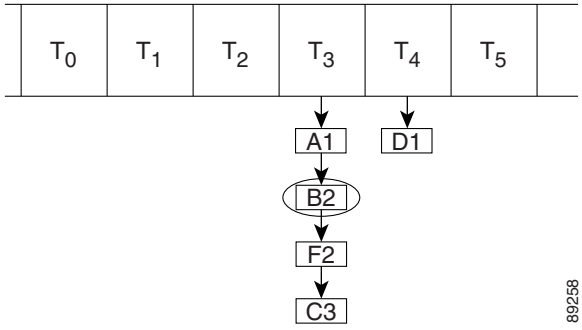

PVC Priorities

Based on the discussion of SAR scheduling, you can see that the PVC priority has similar significance in both collision algorithms. When collisions occur, the SAR processor always gives the PVC with the higher priority precedence over a PVC of lower priority in the link list. Therefore, if you need to increase the performance of a particular PVC, you might consider modifying its PVC priority.

The ATM port adapters originally supported four transmission priorities, but have now been enhanced to support six priorities. The default PVC priorities are established by the ATM port adapter according to the service category that you configure for the PVC. You can modify the default priorities using the **transmit-priority** interface ATM VC configuration command.

For more information about PVC priorities, see the "Configuring PVC Priorities" section on page 5-29 in Chapter 5, "Configuring Traffic Shaping on the PA-A3 and PA-A6 ATM Port Adapters."

Summary of Traffic Flow Through the ATM Port Adapter

After the discussion of the architectures and scheduling mechanisms on the PA-A3 and PA-A6 ATM port adapters, it is useful to summarize the flow of traffic from the router to the ATM port adapter and onto the network as cells:

Step 1 The router performs a DMA transfer of a packet from the transmit ring on the NPE or NSE to a FIFO ring in SDRAM local to the ATM port adapter.

Step 2 The SAR processor uses GCRA to determine when cells from each PVC are eligible for transmission based on the SCR, PCR, and MBS traffic parameters for that PVC.

Step 3 The SAR processor uses the traffic shaping parameters for a PVC to determine the appropriate ICG for cells to be transmitted. Recall that the SAR divides the total available bandwidth for the port adapter (prior to framing) into time slots with an even ICG.

Step 4 Using a calendar table, the SAR processor assigns time slots to each of the PVCs that it is configured to support. Those cells that are eligible for transmission according to GCRA for each PVC are scheduled and transmitted in the corresponding time slot for that PVC.

Step 5 When time slot collisions occur, the SAR processor uses PVC priorities and a head-insertion algorithm to bump cells to the next time slot.

Note The original collision algorithm for the PA-A3 ATM port adapter was tail insertion. For more information, see the "Collision Handling" section on page 2-31.

Step 6 On a FIFO basis, the SAR processor segments the packets that are scheduled for transmission into 52-byte cells [without the Header Error Check (HEC)] before sending them to the framer for the addition of physical-layer overhead and transmission onto the network.

Related Documentation

The following table provides information about additional resources that you can read to learn more about some of the topics discussed in this chapter:

For more information about:	Refer to the following publications:
ATM technical standards	Approved ATM Forum Specifications
ATM technology and other Cisco Systems products	*Cisco ATM Solutions*, Cisco Press
Cisco IOS QoS software features	*Cisco IOS Quality of Service Solutions Configuration Guide*
Framing formats on ATM Interfaces	Framing Formats on DS-3 and E3 Interfaces (TAC Tech Note)

For more information about:	Refer to the following publications:
Memory architecture and switching paths on the Cisco 7200 series	*Inside Cisco IOS Software Architecture,* Cisco Press
Network management variables and measuring rates and utilization for ATM PVCs	Measuring Utilization on ATM PVCs (TAC Tech Note)

Next Steps

The first two chapters of this book provide you with a foundation of ATM technology and concepts related to effectively designing and managing ATM traffic on your Cisco 7200 series router. They describe the architectures and the relationships that you should understand before configuring and optimizing your router to process ATM traffic.

The subsequent chapters in this book provide you with the necessary information to implement ATM traffic management, including hardware and software planning information and configuration guidelines and procedures.

Chapter 3, "ATM Traffic Management Hardware and Software Planning," provides you with additional information about the ATM port adapters supported by the Cisco 7200 series routers and describes some of the tools that you can use to find out more about Cisco IOS software releases and fixes, ATM features, and hardware and software compatibility.

ATM Traffic Management Hardware and Software Planning

This chapter provides an introduction to the ATM port adapter hardware and software that is supported on the Cisco 7200 series routers. It includes hardware installation guidelines and verification information and a review of the Cisco IOS software releases supported by the Cisco 7200 series routers, including a summary of where certain key ATM features were introduced.

This chapter includes the following topics:

Hardware Planning for ATM Traffic Management

This section provides an overview of the hardware guidelines for the Cisco 7200 series routers and an introduction to the ATM port adapters supported on the Cisco 7200 series routers.

This section includes the following topics:

Hardware Installation Guidelines on the Cisco 7200 Series Router

There are specific hardware installation and memory guidelines that you should observe to achieve the optimal operating results for your Cisco 7200 series router.

Cisco 7200 series routers have a finite data-carrying capacity, referred to as *bandwidth*, that affects the port adapter distribution in the chassis, as well as the number and types of port adapters that you can install. The Cisco 7200 series routers use a concept called *bandwidth points*, which allow you to

determine whether your port adapter and I/O controller configuration can be supported by your network processing engine (NPE) or network services engine (NSE) and memory configuration. Each port adapter and I/O controller is associated with a certain bandwidth point value, and the sum of all of these bandwidth points must not exceed allowable limits.

You should meet the following objectives during your hardware configuration:

- Find the total number of bandwidth points for each PCI bus by adding the bandwidth points for the port adapters and I/O controllers that correspond to that particular PCI bus.
- Be certain not to exceed the allowable number of total bandwidth points for each PCI bus (the allowable total varies by your NPE or NSE).
- Distribute the bandwidth evenly between the buses.

Refer to the *Cisco 7200 Series Port Adapter Hardware Configuration Guidelines* document for these important hardware configuration, memory, and operating guidelines.

ATM Port Adapter Support on the Cisco 7200 Series Router

The physical interfaces that support ATM on a Cisco 7200 series router are referred to as ATM port adapters (PAs). This section provides a brief introduction to the ATM port adapters that are supported on the Cisco 7200 series router, and includes the following topics:

For a list of documents that contain important details about the ATM port adapters and their installation, see the "Related Documentation" section on page 3-13.

Note Not all models of ATM port adapters are supported on the Cisco 7200 series router. For example, the PA-A3 OC-12 ATM port adapter is only available initially on the Cisco 7500 series routers.

PA-A1 ATM Port Adapter (OC-3)

The PA-A1 ATM port adapter is the original Cisco Systems ATM port adapter. It is designed for use as a local-area network (LAN) campus uplink supporting LAN Emulation (LANE).

The PA-A1 does not support traffic shaping. The peak cell rate (PCR) of all virtual circuits (VCs) is the line rate, which is 155 Mbps for the OC-3 (optical carrier) interface.

Therefore, the PA-A1 supports the unspecified bit rate (UBR) service class only. You cannot configure other ATM service classes such as non-real-time variable bit rate (nrt-VBR) or available bit rate (ABR) on permanent virtual circuits (PVCs) on the PA-A1 ATM port adapter. The UBR service class provides a high degree of multiplexing or bandwidth sharing, but does not provide any bounds on cell delay or cell loss.

PA-A2 ATM CES Port Adapter (T3, E3 , OC-3, and 4 CBR ports [T1 or E1])

The target application of the PA-A2 ATM CES port adapter is access to a broadband public or private ATM network where multiservice consolidation of voice, video, and data traffic over a single ATM link is a requirement.

The PA-A2 ATM CES port adapter is a dual-wide module. It has four T1 (1.544 Mbps) or four E1 (2.048 Mbps) 120-ohm constant bit rate (CBR) ports that can support both structured (N x 64 Kbps) and unstructured ATM Forum-compliant circuit emulation services (CES), and a single port of an OC-3 (155 Mbps) single-mode intermediate reach or a T3 (45 Mbps) or E3 (34 Mbps) standards-based ATM interface.

The PA-A2 ATM CES port adapter can be used only on a Cisco 7200 series router or a Cisco uBR7200 series router that has at least a 150-MHz network processing engine (NPE-150).

PA-A3 Enhanced ATM Port Adapter (T3, E3, OC-3, and T1/E1 Inverse Multiplexing Over ATM [IMA])

Cisco Systems introduced the enhanced PA-A3 ATM port adapter for WAN links that require native, hardware-based ATM traffic shaping to control bandwidth on the VCs. When connecting an ATM router interface to a switch network that requires traffic shaping, we recommend using the PA-A3 or the PA-A6 ATM port adapter.

The PA-A3 ATM port adapters are single-port, single- and dual-wide port adapters for the Cisco 7200, Cisco 7500, and Cisco 7600 OSR series routers using the FlexWAN module for the Catalyst 6000 family of switches. The ATM PA-A3 is designed with a high-performance, dual segmentation and reassembly (SAR) architecture with local buffer memory.

The PA-A3 ATM port adapter supports advanced ATM hardware features, such as per-VC and per-virtual path (VP) traffic shaping, and up to 4096 VCs. It supports most ATM service classes including ABR, CBR, nrt-VBR, rt-VBR, UBR and UBR+ (for SVCs only).

PA-A6 Enhanced ATM Port Adapter Plus (T3, E3, and OC-3)

With advanced ATM features, the enhanced PA-A6 ATM port adapter plus supports broadband aggregation, WAN aggregation, and campus/metropolitan-area network (MAN) aggregation, including the following types of applications:

- Broadband subscriber access aggregation
- High-speed customer premises equipment (CPE) WAN link
- High-speed WAN uplink
- High-speed enterprise backbone

The enhanced PA-A6 ATM port adapter plus supports many of the same features as the PA-A3 ATM port adapter, but provides additional hardware capacity, including support of up to 8191 VCs.

The enhanced PA-A6 ATM port adapter plus is a series of single-width, single-port ATM port adapters for Cisco 7200 series, Cisco 7401ASR, Cisco 7500 series, and Cisco 7600 series routers.

Like the PA-A3 ATM port adapter, the PA-A6 ATM port adapter also supports advanced ATM hardware features such as per-VC and per-VP traffic shaping. It supports most ATM service classes including ABR, CBR, nrt-VBR, rt-VBR, UBR, and UBR+ (for SVCs only).

ATM Port Adapter Summary

Table 3-1 provides a summary of the ATM port adapter support on the Cisco 7200 series routers.

Table 3-1 Summary of ATM Port Adapter Support on the Cisco 7200 Series Router

ATM Port Adapter	Description in show interfaces atm Command	Service Category Support	Virtual Circuit Support	Priority Scheduling
PA-A1 (OC-3)	"Hardware is TI1570 ATM"	UBR	2048 VCs	One level. Uses round robin scheduling among all VCs.
PA-A2 (T3/E3/OC-3 and 4 CBR [T1 or E1] ports)	"Hardware is ATM-CES"	• ABR • CBR—for voice • nrt-VBR • UBR	2046 VCs and 124 CBR VCs	2 levels based on CBR for voice or data VC. Strict priority is given to CBR and a weighted round robin is used for data VCs. Contention between VCs is handled using the fairness algorithm.[1]
PA-A3 (T3/E3/OC-3)	"Hardware is ENHANCED ATM PA"	• ABR • CBR—for data • nrt-VBR • rt-VBR • UBR • UBR+—SVCs only	4096 VCs	6 levels (configurable). For more information, see the "Configuring PVC Priorities" section on page 5-29.
PA-A3 IMA (T1/E1)	"Hardware is ENHANCED ATM PA"	• ABR • CBR—for data • nrt-VBR • rt-VBR • UBR • UBR+—SVCs only	• 512 VCs on UNI • 512 interface VCs per link on each IMA interface • Cell-based inverse multiplexing that allows Operation, Administration, and Maintenance (OAM) cells to provide management and monitoring information (including connectivity, alarm indication signals [AIS] and loopback) across the inverse multiplexed links.	6 levels (configurable). For more information, see the "Configuring PVC Priorities" section on page 5-29.

Table 3-1 Summary of ATM Port Adapter Support on the Cisco 7200 Series Router (continued)

ATM Port Adapter	Description in show interfaces atm Command	Service Category Support	Virtual Circuit Support	Priority Scheduling
PA-A6 (T3/E3/OC-3)	"Hardware is ENHANCED ATM PA Plus"	• ABR • CBR • nrt-VBR • rt-VBR • UBR • UBR+—SVCs only	8191 VCs (one VC is reserved for OAM processing)	6 levels (configurable). For more information, see the "Configuring PVC Priorities" section on page 5-29.

1. Uses a fairness algorithm (as defined in Appendix I.3 of Traffic Management Spec. 4.0). Specifically, the TI1585 ASIC uses the max-min fairness criteria. It divides the available bandwidth for bottlenecked connections among all connections bottlenecked on this link. The bandwidth available for bottlenecked connections is defined as the available link bandwidth minus the sum of all bandwidths of connections bottlenecked elsewhere.

Software Planning for ATM Traffic Management

The Cisco 7200 series routers and ATM port adapters are supported by a variety of Cisco IOS software releases. The currently available Cisco IOS releases include Cisco IOS Release 12.0, Cisco IOS Release 12.1, and Cisco IOS Release 12.2, and several early deployment releases including Cisco IOS Release 12.0 S, Cisco IOS Release 12.1 E, Cisco IOS Release 12.2 S, Cisco IOS Release 12.2 T, and Cisco IOS Release 12.2 B.

Note If you are new to Cisco IOS software and its release structure, you might find it useful to refer to the ABCs of Cisco IOS Software site located on Cisco.com. Use the ABCs Site Map to access specific topics, including information about how the Cisco IOS software is packaged and its release trains. You can also go to the Cisco IOS Software home page to locate other information about the Cisco IOS software, its product documentation, and tools.

This section provides background information about some of the latest Cisco IOS software releases supported by the Cisco 7200 series routers and describes some of the relationships between the software releases. It also includes information about where new feature support is introduced and about code stability and maturity that can help you to determine which Cisco IOS software release best suits your environment and objectives.

This section includes the following topics:

Cisco IOS Software Releases 12.0 T and 12.1 Release History

Cisco IOS Release 12.1 is based directly on Cisco IOS Release 12.0 T. It offers all of the advanced features of the 12.0 T release, but with some additional fixes. Therefore, it provides better code maturity and stability.

The last Cisco IOS 12.0 T release was Cisco IOS Release 12.0(7)T, after which the release numbering was incremented—the next maintenance release of this code is called Cisco IOS Release 12.1(1). The 12.1(1) and subsequent 12.1 releases are focused on code maturity and stability, and they are intended to achieve general deployment (GD) certification. To reach this goal, the feature set of the 12.1 release was frozen with Cisco IOS Release 12.1(1). At that time, new features and hardware support were directed into the 12.1 E and 12.1 T releases.

When the12.0 T release became the 12.1 release, the following two new branches, called Cisco IOS Release 12.1 E and Cisco IOS Release 12.1 T, were created for new feature integration:

- Cisco IOS Release 12.1 T—Integrates features and hardware support for platforms across the Cisco product line.
- Cisco IOS Release 12.1 E—Focuses exclusively on feature and hardware support for the Cisco 7500, Cisco 7200, and Cisco 7100, and Catalyst 6000 families.

Both the 12.1 E and 12.1 T branches are synchronized to the 12.1 release, receiving the same bug fixes that are committed into the 12.1 branch.

Cisco IOS Software Releases 12.1 T and 12.2 Release History

The last 12.1 T release was Cisco IOS Release 12.1(5)T, after which the release numbering was again incremented. The next maintenance release of this code is called Cisco IOS Release 12.2(1). Similar to the release history and objectives of Cisco IOS Release 12.1 described in the previous topic, the 12.2(1) and subsequent 12.2 releases are focused on code maturity and stability to achieve general deployment (GD) certification. To reach this goal, the feature set of Cisco IOS Release 12.2 was frozen with Cisco IOS Release 12.2(1).

When the 12.1 T release became the 12.2 release, the following three new branches, called Cisco IOS Release 12.2 B, Cisco IOS Release 12.2 S, and Cisco IOS Release 12.2 T, were created for new feature integration:

- Cisco IOS Release 12.2 B—Focuses primarily on broadband solutions and supports the Cisco 7400, Cisco 7200, and Cisco 6400 product families.
- Cisco IOS Release 12.2 S—Focuses exclusively on feature and hardware support for the Cisco 7600, Cisco 7500, Cisco 7400, Cisco 7200, and Cisco 7100, and Catalyst 6000 product families.
- Cisco IOS Release 12.2 T—Integrates features and hardware support for platforms across the Cisco product line.

Cisco IOS Release Summary

Table 3-2 provides a list of some of the latest available Cisco IOS releases and their descriptions for the Cisco 7200 series routers.

Table 3-2 Summary of Cisco IOS Releases Supported on the Cisco 7200 Series Routers

Cisco IOS Release	Description
12.0	Offers the most code maturity, extensive field exposure, and is the preferred GD release for the Cisco 7200 series routers.
12.0 S	Based on Cisco IOS Release 12.0. Release that is most suitable for service providers. Contains additional features and hardware support for service providers, but feature sets are limited to IP protocols.
12.0 ST	Based on Cisco IOS Release 12.0 S. Technology train for Cisco IOS Release 12.0 S that contains enhanced features suitable for service providers.
12.1	Based on Cisco IOS Release 12.0 T. Offers code maturity and is focused on stability. Latest GD-available release for the Cisco 7200 series routers.
12.1 E	Based on Cisco IOS Release 12.1. Focused on Cisco 7x00 and Catalyst 6000 platforms. Offers the latest Cisco 7200 series hardware and feature support and has extensive field exposure. Includes support for all 12.0 T, 12.0 XE, and 12.1 features.
12.1 T	Based on Cisco IOS Release 12.1. Offers all 12.0 T, 12.1, and some 12.0 XE features. Offers additional features for all Cisco platforms over time.
12.2	Based on Cisco IOS Release 12.1 T. Feature-rich and supports wide variety of hardware. Focused on stability and GD certification.
12.2 B	Based on Cisco IOS Release 12.2. Primarily focused on broadband solutions, supporting the Cisco 7200, Cisco 7400, and Cisco 6400 product families.
12.2 S	Based on Cisco IOS Release12.2. Offers all 12.0 S and 12.1 E features. Offers additional features for all Cisco 7x00, Cisco 10000, Cisco 12000, and Catalyst 6000 platforms over time.
12.2 T	Based on Cisco IOS Release 12.2. Offers additional hardware and software feature support for all Cisco platforms over time.

Additional Software Planning Information for the PA-A3 and PA-A6 ATM Port Adapters

This section provides information about some of the Cisco IOS software releases in which some specific software feature additions for the PA-A3 and PA-A6 ATM port adapters were introduced.

Table 3-3 List of Key Feature Support Additions for PA-A3 and PA-A6 ATM Port Adapters

Feature Description	Cisco IOS Software Release Where Introduced
ABR service category	12.0(4)T, 12.0(5)S **Note** The following minimum Cisco IOS software releases are recommended for ABR support—12.0(7)T and later, 12.0(8)S and later, 12.1(5) and later.
CBR service category	12.2(5), 12.2(8)T, 12.2(14)S, 12.2(15)B
Class-map groups	12.1(5)E, 12.1(5)T
Configurable per-VC hold queue	12.1(5)T
Head-insertion scheduling algorithm	12.0(21)S, 12.0(21)ST, 12.1(11), 12.1(11b)E, 12.2(6), 12.2(8)T, 12.2(14)S, 12.2(15)B
Per-VC CBWFQ (for the PA-A3 ATM port adapter on the Cisco 7200 series router)	12.0(5)T, 12.0(5)XE, 12.1(1), 12.1(1)T, 12.1(1)E
Per-VC CBWFQ (NSE support)	12.1(7)E
Per-VC CBWFQ (NSE-1 support)	12.2(4)B1
PVC transmit priority increase from 4 to 6 levels	12.2(5), 12.2(14)S, 12.2(8)T, 12.2(15)B
Real-time VBR service category	12.2(5), 12.2(8)T, 12.2(14)S, 12.2(15)B
Receive buffer default size increase	12.0 S, 12.1 E, 12.2 T
UBR+ service category (SVCs only)	11.3 T, enhanced in 12.0(3)T
Voice over ATM with AAL2 trunking	12.(2)T
VP shaping	12.0(4)T, 12.0(5)S, 12.0(7)XE

Cisco Systems Tools Overview

Cisco Systems maintains several tools that help you to configure and maintain your Cisco Systems router hardware and software. Table 3-4 provides a brief description of some of the helpful tools that you can use for hardware and software planning.

Note All of the tools shown in Table 3-4 require that you log in to your Cisco.com account. If you do not have an account or have forgotten your username or password, click **Cancel** at the Login dialog box and follow the instructions that appear.

Table 3-4 Helpful Tools for Hardware and Software Planning

Tool Name	URL	Description
Bug Toolkit	http://www.cisco.com/cgi-bin/Support/Bugtool/launch_bugtool.pl	Use the Bug Toolkit to find the latest information about Cisco IOS software defects.
Dynamic Configuration Tool	http://www.cisco.com/order/apollo/configureHome.html	Use the Dynamic Configuration Tool to help you select the appropriate hardware and software components for new Cisco Systems equipment that you want to purchase.
Feature Navigator	http://tools.cisco.com/ITDIT/CFN/jsp/index.jsp	Use Feature Navigator to find information about platform and software image support, including available features. **Note** Feature Navigator now supports all of the major Cisco IOS software releases.
Software Advisor	http://www.cisco.com/cgi-bin/front.x/Support/HWSWmatrix/hwswmatrix.cgi	Use the Software Advisor to find the minimum software requirements for your Cisco Systems hardware.

Verifying Software Support for Hardware

To find the minimum Cisco IOS software requirements for your Cisco 7200 series hardware, use the Software Advisor tool on Cisco.com. This tool does not verify whether hardware modules within a system are compatible, but it does provide the minimum Cisco IOS requirements for individual hardware modules or components.

Note You need a Cisco.com account to access Software Advisor. If you do not have an account or have forgotten your username or password, click **Cancel** at the Login dialog box and follow the instructions that appear.

To access Software Advisor, perform the following steps:

Step 1 Go to Cisco.com and click **Login** at the top of the Cisco.com home page.

Step 2 Enter your username and password.

Step 3 Point your browser directly to http://www.cisco.com/cgi-bin/front.x/Support/HWSWmatrix/hwswmatrix.cgi.

The Software Advisor page appears.

Step 4 Click **Software Support for Hardware**.

Step 5 From the Software Support for Hardware page, do one of the following to search for the minimum supported software release needed for your hardware:

- Choose a product family.
- Type a specific product number.

A list of the compatible hardware product numbers and minimum software releases is provided.

Step 6 To find a list of the common software releases for your hardware, select the checkbox beside the hardware products for which you want software information and click the **Display Intersection** button at the bottom of the page.

The products that you selected are displayed with the intersecting software releases shown in bold.

Verifying Feature Support

Both Feature Navigator and Software Advisor provide information about Cisco IOS software features and releases. Each tool allows you to find the following information about features in the Cisco IOS software:

- Which Cisco IOS software release supports one or more specified features.
- A list of features that are supported by a particular Cisco IOS software release, image name, or product number.
- The list of features that are shared by, and the list of features that are unique to, any two specified Cisco IOS software releases.

Note You can also find image names and product numbers, view MIBs, view release notes, and download images from the output of the **Compare Images** link.

Using Feature Navigator to Search for Features

To access Feature Navigator, perform the following steps:

Step 1 Go to Cisco.com and click **Login** at the top of the Cisco.com home page.

Step 2 Enter your username and password.

Step 3 Point your browser directly to http://tools.cisco.com/ITDIT/CFN/jsp/index.jsp.

The Cisco Feature Navigator II page appears.

Step 4 From the Cisco Feature Navigator II page, choose one of the following links:

- **Search by Feature—**To find which Cisco IOS software release supports the features that you want.
- **Search by Release—**To find a list of features supported by a particular Cisco IOS software release, image name, product number, or platform.
- **Compare Releases—**To compare two Cisco IOS software releases to find out what features the software has in common and which features are unique.

Verifying the Hardware and Software Installation

To obtain information about your Cisco 7200 series router hardware and software, you can use several different **show** commands, including the following commands:

- **show diag**—Displays the types of port adapters installed in your system and specific information about each one.
- **show interfaces atm**—Displays status information, including the physical slot and interface address, for the interfaces that you specify.
- **show version**—Displays the configuration of the system hardware, the number of each interface type installed, the Cisco IOS software version, the names and sources of configuration files, and boot image information.

This section includes the following examples:

- Example of the show diag Command, page 3-11
- Example of the show interfaces atm Command, page 3-12
- Example of the show version Command, page 3-12

Note The sample output that appears in this document might not match the output that you receive when running these commands. The sample output in this document is intended only as an example.

Example of the show diag Command

To display the types of port adapters installed in your system (and specific information about each), use the **show diag** command.

The following example shows output from the **show diag** command on a Cisco 7200 series router with a PA-A6 ATM port adapter installed in slot 6:

```
Router# show diag 6
Slot 6:
ATM WAN OC3+ (MM) Port adapter, 1 port
Port adapter is analyzed
Port adapter insertion time 00:44:57 ago
EEPROM contents at hardware discovery:
Hardware Revision: 1.0
PCB Serial Number:-----H
Part Number: 73-7981-0
Board Revision: 1A
RMA Test History: 00
RMA Number: 00-00-00
RMA History: 00
Unknown Field (type 0088): 00 00 00 01
Product Number: PA-ATM-DBL-DLX-OC3MM
Top Assy. Part Number: 800-20782-01
EEPROM format version 4
EEPROM contents (hex):
0x00: 04 FF 40 03 A7 41 01 00 C1 8B 2D 20 2D 20 2D 20
0x10: 2D 20 2D 20 48 82 49 1F 2D 01 42 31 41 03 00 81
0x20: 00 00 00 00 04 00 88 00 00 00 01 CB 94 50 41 2D
0x30: 41 54 4D 2D 44 42 4C 2D 44 4C 58 2D 4F 43 33 4D
```

```
0x40: 4D C0 46 03 20 00 51 2E 01 FF FF FF FF FF FF FF
0x50: FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF
0x60: FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF
0x70: FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF
```

Example of the show interfaces atm Command

To display status information (including the physical slot and interface address) for the interfaces that you specify, use the **show interfaces atm** command.

The following example shows output from the **show interfaces atm** command on a Cisco 7200 series router with a PA-A6 ATM port adapter installed in slot 6:

```
Router# show interfaces atm 6/0
ATM6/0 is up, line protocol is up
Hardware is ENHANCED ATM PA Plus
MTU 4470 bytes, sub MTU 4470, BW 149760 Kbit, DLY 80 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ATM, loopback not set
Encapsulation(s): AAL5
8191 maximum active VCs, 1 current VCCs
VC idle disconnect time: 300 seconds
8 carrier transitions
Last input 00:37:33, output 00:26:16, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: None
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
54033 packets input, 81670740 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
8763526 packets output, 296268354 bytes, 0 underruns
0 output errors, 0 collisions, 0 interface resets
0 output buffer failures, 0 output buffers swapped out
```

Example of the show version Command

To display the configuration of the system hardware, the number of each interface type installed, the Cisco IOS software version, the names and sources of configuration files, and boot image information use the **show version** command.

The following example shows output from the **show version** command on a Cisco 7200 series router:

```
Router# show version
Cisco Internetwork Operating System Software
IOS (tm) 7200 Software (C7200-IS-M), Version 12.2(8)B, RELEASE SOFTWARE (fc1)
TAC Support: http://www.cisco.com/tac
Copyright (c) 1986-2002 by cisco Systems, Inc.
Compiled Fri 02-Aug-02 10:51 by ccai
Image text-base: 0x60008940, data-base: 0x61850000
ROM: System Bootstrap, Version 11.1(13)CA, EARLY DEPLOYMENT RELEASE SOFTWARE (f)
BOOTLDR: 7200 Software (C7200-BOOT-M), Version 11.1(16)CA, EARLY DEPLOYMENT REL
d11-5-7206-15 uptime is 1 week, 31 minutes
System returned to ROM by reload at 14:08:34 UTC Tue Aug 20 2002
System image file is "slot0:c7200-is-mz.122-11.T"
cisco 7206 (NPE200) processor (revision B) with 114688K/16384K bytes of memory.
Processor board ID 15455885
R5000 CPU at 200Mhz, Implementation 35, Rev 2.1, 512KB L2 Cache
6 slot midplane, Version 1.3
```

```
Last reset from power-on
Bridging software.
X.25 software, Version 3.0.0.
Primary Rate ISDN software, Version 1.1.
1 FastEthernet/IEEE 802.3 interface(s)
2 Serial network interface(s)
2 HSSI network interface(s)
11 ATM network interface(s)
4 Channelized T1/PRI port(s)
125K bytes of non-volatile configuration memory.
4096K bytes of packet SRAM memory.
20480K bytes of Flash PCMCIA card at slot 0 (Sector size 128K).
4096K bytes of Flash internal SIMM (Sector size 256K).
Configuration register is 0x2102
```

Related Documentation

The following table provides information about additional resources that you can read to learn more about some of the topics discussed in this chapter:

For more information about:	Refer to the following publications:
Cisco IOS software overview and release structure	ABCs of Cisco IOS Software
Hardware guidelines, bandwidth points, and memory requirements	*Cisco 7200 Series Port Adapter Hardware Configuration Guidelines*
Hardware installation	• *PA-A1 ATM Port Adapter Installation and Configuration* • *PA-A2 ATM CES Port Adapter Installation and Configuration* • *PA-A3 Enhanced ATM Port Adapter Installation and Configuration* • *Inverse Multiplexing over ATM Port Adapter Installation and Configuration* • *PA-A6 Port Adapter Installation and Configuration*
Port adapter specifications and feature summaries	• *ATM Circuit Emulation Services Port Adapter for Cisco 7200 Routers* (brochure) • *PA-A1 ATM Port Adapter* (data sheet) • *Enhanced ATM Port Adapter (ATM PA-A3)* (data sheet) • *Enhanced ATM Port Adapter (ATM PA-A3)* (brochure) • *Enhanced ATM Port Adapter for Cisco 7200, 7400, 7500, and 7600 Series Routers* (ATM PA-A6 data sheet)

Next Steps

This chapter provides you with hardware and software planning information for your ATM port adapter on the Cisco 7200 series router.

Chapter 4, "Preparing to Configure ATM Traffic Management and QoS Features," provides you with some specific guidelines on gathering the necessary network information and defining your service models in preparation for configuring traffic shaping and QoS on your router.

Preparing to Configure ATM Traffic Management and QoS Features

Before you begin to configure traffic shaping and implement QoS, there are several areas of the network design that you should consider.

Several of these tasks should be implemented as a regular and ongoing assessment of your network. Most networks and their traffic usage patterns are dynamic. Network designs that work initially often need to be adjusted as usage increases and loads and requirements change.

This chapter includes the following guidelines for preparing to configure ATM traffic management and QoS features:

Defining the Service Model

To properly design your network, you need to define the business model, goals and requirements that the network needs to support. This serves as a basis for establishing the proper criteria for the traffic shaping and QoS.

Basically, you need to know the objectives of the network so that you can measure whether or not you are meeting those objectives, and assess where design changes might need to occur.

Analyzing the Network Traffic

Analyzing your network traffic is a key task throughout the design and maintenance of your network.

When you analyze your network traffic, obtain the following information:

- Assess the traffic mix—Classify the different types of traffic being handled and determine service levels for the different types of traffic to be supported. Be sure to identify any traffic that is vital to the support of your business applications. This is called *mission critical traffic*.
- Assess the load—Determine how much of the bandwidth is used, and by what types of traffic, over the network links.
- Assess the performance—Does the current network design optimize the flow of different traffic types and minimize loss or latency? During periods of congestion, is mission critical traffic adversely affected?
- Assess CPU utilization—During busy periods, monitor the CPU utilization on network devices. Be aware of any impact that QoS configuration might have on CPU utilization.

In addition to analyzing your network, you should gather other information about your network, including network topology diagrams, device configurations, and software versions.

Considering the Traffic Contract

The traffic contract establishes the criteria for policing of ATM virtual connections on the network to ensure that violations of the agreed-upon service levels do not occur.

Therefore, if you already have established an agreement with your service provider, you want to be sure that your ATM network design conforms to the service levels for which you are paying, and also so that you can understand where adjustments to the contract might need to be made.

Be sure that you understand how your service provider is policing your ATM network connections.

For more information about the traffic contract, see the "Traffic Contract" section on page 1-3.

Evaluating the PVC Configuration Over the Physical Interface

Consider the following areas when evaluating your PVC configuration over the physical interface:

- Verify the appropriate framing between the router and the switch.
- Verify whether you are using payload scrambling on the PVC. Both ends of a VC must have the same scrambling setting. For more information, refer to the TAC Tech Note, When Should Scrambling Be Enabled on ATM Virtual Circuits?.
- Determine the rate supported by the physical interface and evaluate PVC performance.
- Determine the Layer 2 and Layer 3 overhead for the classes of traffic to be supported over the physical interface when evaluating bandwidth.
- Determine the maximum delay tolerable for mission critical traffic and adjust burst parameters for traffic shaping accordingly.
- Determine if the PVC requires any bandwidth guarantees for its applications.

- If you are limited to a certain type of PVC, configure the closest service category supported by the ATM port adapter.
- Anticipate PVC growth in both load (amount of bandwidth consumption) and number of PVCs that you might need to support. For multiple PVCs, consider the following:
  - Project the maximum number of PVCs that you need to support.
  - Consider whether you need to implement individual PVCs or a VC bundle.

Related Documentation

The following table provides information about additional resources that you can read to learn more about some of the topics discussed in this chapter:

For more information about:	Refer to the following publications:
Payload Scrambling	When Should Scrambling Be Enabled on ATM Virtual Circuits? (TAC Tech Note)

Next Steps

After you evaluate your network by assessing your traffic and defining your business and service models, you are ready to begin shaping your traffic and configuring your QoS service policies to meet your network and business objectives.

For guidelines on configuring traffic shaping to effectively manager your ATM traffic, read Chapter 5, "Configuring Traffic Shaping on the PA-A3 and PA-A6 ATM Port Adapters."

For guidelines on configuring IP to ATM Class of Service (CoS) features and information about queue limits, read Chapter 6, "Configuring QoS on the Layer 3 Queues for the PA-A3 and PA-A6 ATM Port Adapters."

Configuring Traffic Shaping on the PA-A3 and PA-A6 ATM Port Adapters

The PA-A3 and PA-A6 ATM port adapters provide the most advanced set of features for ATM traffic management on the Cisco 7200 series routers. This chapter provides details about some of the different service categories supported by the PA-A3 and PA-A6 ATM port adapters for traffic shaping, and provides some guidelines about configuring those service categories.

The purpose of this chapter is to provide you with a combination of design and configuration information to help you make informed decisions about implementing and optimizing traffic shaping on your ATM port adapters. This chapter does not provide information about all applicable configuration procedures or commands that you can configure for a virtual circuit (VC). Therefore, you should also refer to the Cisco IOS software publications applicable to the Cisco IOS software release that you are running on your Cisco 7200 series router.

For more information about configuring ATM on the Cisco 7200 series routers, refer to the *Cisco IOS Wide-Area Networking Configuration Guide* and *Cisco IOS Wide-Area Networking Command Reference* publications available on Cisco.com.

This chapter includes the following sections:

Preparing to Configure Traffic Shaping

This section includes some basic information that you should consider prior to configuring traffic shaping on the PA-A3 and PA-A6 ATM port adapters. Before you begin to configure traffic shaping, be sure you have performed the preparatory tasks recommended in Chapter 4, "Preparing to Configure ATM Traffic Management and QoS Features."

Once you have determined your service model and have analyzed your network traffic needs, you need to decide how you are going to implement your VCs and the type of traffic shaping that those VCs should support.

Note As you plan to implement ATM traffic shaping, be aware that Operation, Administration, and Maintenance (OAM) cells are not shaped.

This section includes the following topics to help you begin to configure traffic shaping on your VC:

Determining the PVC Configuration Method

The PA-A3 and PA-A6 ATM port adapters support both switched virtual circuit (SVC) and permanent virtual circuit (PVC) implementation. This book focuses on PVC configuration and the associated traffic shaping and QoS configuration associated with PVCs.

Before you begin, you should think about how you want to configure your PVCs. You can implement PVC bundles and ranges of PVCs to simplify the configuration of multiple PVCs.

The PA-A3 and PA-A6 ATM port adapters support several different methods of configuring PVCs:

- Single PVCs using ATM VC configuration mode under an interface
- Multiple PVCs in VC bundles using ATM VC bundle configuration mode
- Multiple PVCs in ranges using ATM PVC range or ATM PVC-in-range configuration modes

When you configure PVCs, you also configure certain attributes for those PVCs, including the traffic shaping. To configure traffic shaping and other attributes for a PVC, you can use a couple of methods:

- Use explicit commands to configure traffic shaping directly at the PVC, PVC bundle, or PVC range.
- Define a VC class that contains the explicit traffic shaping commands (and other attributes). Then, attach the VC class where you want the configured attributes to be applied.

You can apply a VC class at many levels, including at an ATM interface or subinterface, an individual PVC, PVC bundle, or PVC range. VC classes are useful when you want to categorize different configurations and apply them to multiple PVCs.

In addition, you can use VC classes to define certain areas of PVC configuration (such as traffic shaping) once, but apply that configuration to multiple PVCs.

Configuring VC Classes

VC classes provide a simplified method of configuring and applying one or more parameters to individual PVCs, PVC bundles, and PVC ranges. To implement VC classes, complete the following steps:

Step 1 Create the VC class.

From global configuration mode, use the **vc-class atm** command.

Step 2 Configure the VC parameters.

From ATM VC class configuration mode, configure the VC parameters using one or more of the following commands: **abr**, **broadcast**, **bump**, **cbr**, **encapsulation**, **idle-timeout**, **ilmi manage**, **inarp**, **oam-bundle**, **oam-pvc**, **oam retry**, **oam-svc**, **protocol**, **ubr**, **ubr+**, **vbr-nrt**, and **vbr-rt**.

Note The **ubr+** command was introduced in Cisco IOS Release 11.3 T for SVCs. UBR+ is not applicable to PVCs. However, in releases prior to Cisco IOS Release 12.0(7)T, the command is still available within the CLI for PVC configuration—but it should not be used. If you apply a VC class with the **ubr+** command to a PVC, the Cisco IOS software assigns the UBR class to the PVCs.

Step 3 Apply the VC class.

There are several different commands to apply VC classes depending on where in the configuration you want to apply the VC class. You can apply the VC class at the ATM interface or subinterface, individual PVC or SVC, bundle or bundle member, range of PVCs or individual PVC in a range.

To apply a VC class, use one of the commands shown in Table 5-1.

Table 5-1 List of Commands Used to Apply VC Classes

Command	Configuration Mode	Description
class-bundle	ATM VC bundle (config-atm-bundle)	Applies the attributes defined within a VC class to all members of an ATM VC bundle.
class-int	Interface (config-if) or Subinterface (config-subif)	Applies the attributes defined within a VC class to an ATM main interface or subinterface.
class-range	ATM PVC range (config-if-atm-range)	Applies the attributes defined within a VC class to all members of an ATM VC range.
class-vc	ATM VC (config-if-atm-vc) or ATM VC bundle-member (config-if-atm-member) or ATM PVC-in-range (cfg-if-atm-range-pvc)	Applies the attributes defined within a VC class to an individual PVC, SVC, PVC bundle member, or PVC in a range.

VC Class Configuration Guidelines

When you configure VC classes, consider the following guidelines:

- When you create a VC class for a VC bundle *member*, you can use the following commands to define your parameters: **abr**, **bump**, **cbr**, **precedence**, **protect**, **ubr**, **vbr-nrt**, and **vbr-rt**.
- You cannot use the following commands in ATM VC class configuration mode to configure a VC bundle *member*: **encapsulation**, **protocol**, **inarp**, and **broadcast**. These commands are useful only at the bundle level, not the bundle member level.
- If an SVC command (for example, **idle-timeout** or **oam-svc**) is configured in a VC class, but the VC class is applied on a PVC, the SVC command is ignored. This is also true if a PVC command is applied to an SVC.

Inheritance Rules

If you combine methods of configuring traffic shaping for PVCs—that is, you might configure certain commands directly at the VC and then also implement certain attributes using a VC class—it is important to understand that there is a hierarchy in which the resulting values are set. The parameters that you set in certain levels of configuration take precedence over other levels of configuration where the same values might be set.

Generally speaking, the values of parameters that you configure directly at any level using explicit configuration commands take precedence over the same parameters that you might apply to a VC through the application of a VC class. This allows you to maximize the organization of your configuration so that you can combine common attributes shared by multiple VCs into a VC class; but, it also provides the flexibility for you to override certain attributes directly at the VC level where exceptions might occur.

The PA-A3 and PA-A6 ATM port adapters support the following rules of inheritance:

- Parameters that you set directly for a VC using discrete commands (in ATM VC configuration mode, ATM VC bundle configuration mode, or ATM PVC-in-range configuration mode) supersede VC class parameters applied using the **class-int** command at an ATM main interface or subinterface or using the **class-vc** command.
- Parameters that you set directly for a VC bundle member using discrete commands in ATM VC bundle configuration mode supersede values for the same parameters set for the VC at any other level, including VC class parameters applied using the **class-bundle** command for the entire VC bundle.

  VCs within a VC bundle are subject to the following configuration inheritance rules (listed in order of highest precedence):

  - VC configuration in ATM VC bundle configuration mode
  - VC class configuration in ATM VC class configuration mode
  - Subinterface configuration in subinterface configuration mode

- Parameters that you set directly for a VC range member using discrete commands in PVC-in-range configuration mode supersede values for the same parameters set for the VC at any other level, including VC class parameters applied using the **class-range** command for the entire PVC range.
- If you do not explicitly configure an ATM service category (using the **abr**, **cbr**, **vbr-nrt**, **vbr-rt**, **ubr**, or **ubr+** commands) on an ATM PVC, SVC, or VC bundle member, the VC inherits the configuration according to the following hierarchy (listed in order of highest precedence):
  - Configuration within a VC class applied to the PVC or SVC itself.
  - Configuration within a VC class applied to the PVC's or SVC's ATM subinterface.
  - Configuration within a VC class applied to the PVC's or SVC's ATM main interface.
  - Global default configuration: Traffic shaping at the maximum line rate of the PVC or SVC.

PVC Configuration Method Examples

This section provides some examples showing the available configuration methods to configure a PVC:

For more information on PVC configuration and the commands used in these examples, refer to the *Cisco IOS Wide-Area Network Configuration Guide* and *Cisco IOS Wide-Area Network Command Reference*, and the *Cisco IOS Quality of Service Solutions Configuration Guide* and *Cisco IOS Quality of Service Solutions Command Reference* publications.

Single PVC Configuration Example

To create a single PVC, use the **pvc** interface configuration command. From there you can configure traffic shaping and other parameters directly at the VC using explicit commands or by applying a VC class.

The following example shows a single PVC configuration on an ATM main interface with AAL5/MUX encapsulation and the nrt-VBR service category configured directly at the VC:

```
interface atm 2/0
 pvc cisco 1/40
 encapsulation aal5mux ip
 vbr-nrt 100000 50000 20
 exit
```

VC Bundle Configuration Example

To create an ATM PVC bundle, use the **bundle** interface configuration command.

The following example shows creation of a bundle named *new-york*. It specifies the IP address of the subinterface and the router protocol—the router uses Intermediate System to Intermediate System (IS-IS) as an IP routing protocol:

```
interface a1/0.1 multipoint
 ip address 10.0.0.1 255.255.255.0
 ip router isis
 bundle new-york
```

From ATM VC bundle configuration mode, you can configure traffic shaping and other parameters directly at the bundle using explicit commands or by applying a VC class.

VC Range Configuration Example

Another way to configure multiple PVCs is to create a range of PVCs. To create a range of PVCs, use the **range pvc** subinterface configuration command. The **range pvc** command enters you into PVC range configuration mode. From there you can configure traffic shaping and other parameters directly at the VC range using explicit commands or by applying a VC class.

The number of PVCs in a range can be calculated using the following formula:

Number of PVCs = (*end-vpi* – *start-vpi* + 1) x (*end-vci* – *start-vci* +1).

The *start-vpi* argument may be omitted if it is zero. The *end-vpi* argument may be omitted, but if it is omitted, it is assigned the value of *start-vpi*. The *end-vpi* and *end-vci* arguments are always greater than or equal to *start-vpi* and *start-vci* respectively.

When applied to multipoint subinterfaces, the **range pvc** command creates a range of ATM PVCs.

When applied to point-to-point subinterfaces, the **range pvc** command creates a range of PVCs and a corresponding range of point-to-point subinterfaces. For point-to-point subinterfaces, subinterface numbering begins with the subinterface on which the PVC range is configured and increases sequentially through the range.

In the following example, 100 PVCs with virtual channel identifier (VCI) values from 100 to 199 for each virtual path identifier (VPI) value from 0 to 4 are created for a PVC range called *range-pppoa-1*. This configuration creates a total of 500 PVCs in the range. PVC parameters are then configured for the range.

```
interface atm 6/0.110 multipoint
 range range-pppoa-1 pvc 100 4/199
  class-range class-pppoa-1
  ubr 1000
  encapsulation aal5snap
  protocol ppp virtual-Template 2
```

PVC-in-Range Configuration Example

To configure an individual PVC within a PVC range, use the **pvc-in-range** PVC range configuration command. From there you can configure traffic shaping and other parameters directly at the VC using explicit commands or by applying a VC class.

The **pvc-in-range** command defines an individual PVC within a PVC range and enables PVC-in-range configuration mode.

In the following example, a PVC called *pppoa* is deactivated. The PVC called *pppoa* is an individual PVC within a configured PVC range:

```
pvc-in-range pppoa 0/130
 shutdown
```

VC Class at an ATM Main Interface Configuration Example

To configure traffic shaping using a VC class, use the **vc-class atm** global configuration command to create the class and configure the VC parameters. You can apply VC classes at different levels of the configuration. To apply a VC class at the ATM main interface, use the **class-int** interface configuration command.

In the following example, a class called *classA* is first created and then applied to the ATM main interface 2/0:

```
! The following commands create the class named classA:
vc-class atm classA
 ubr 10000
 encapsulation aal5mux ip
! The following commands apply classA to ATM main interface 2/0:
interface atm 2/0
 class-int classA
```

VC Class at a Single PVC Configuration Example

To apply a VC class at the PVC, use the **class-vc** ATM VC configuration command. The following example shows creation of a class called *classA* and its application to an ATM PVC called *router5*:

```
! The following commands create the class named classA:
vc-class atm classA
 ubr 10000
 encapsulation aal5mux ip
! The following commands apply classA to an ATM PVC:
interface atm 2/0
 pvc router5 1/32
 class-vc classA
```

VC Class at an ATM Bundle Configuration Example

To apply a VC class at the PVC bundle, use the **class-bundle** ATM VC bundel configuration command. The following example shows creation of a class called *class1* and its application to the bundle called *bundle1*:

```
! The following commands create the class named class1:
vc-class atm class1
 encapsulation aal5snap
 broadcast
 protocol ip inarp
 oam-bundle manage 3
 oam 4 3 10
 vbr-nrt 100000 50000 20
! The following commands apply class1 to the bundle named bundle1:
bundle bundle1
 class-bundle class1
```

Taking into account hierarchy precedence rules, VCs belonging to the *bundle1* bundle will be characterized by these parameters: AAL5SNAP encapsulation, broadcast enabled, use of Inverse Address Resolution Protocol (Inverse ARP) to resolve IP addresses, Operation, Administration, and Maintenance (OAM) enabled, and non-real-time variable bit rate traffic shaping.

This example shows how you can apply VC attributes to all members of a VC bundle using the **class-bundle** command.

VC Class at a VC Bundle Member Configuration Example

You can also apply VC attributes to a particular bundle member using the **class-vc** command, which is shown in this example.

The following example shows creation of a class called *classA* and its application to the bundle member called *vcmember*, which is a member of *bundle1*:

```
! The following commands create the class named classA:
vc-class atm classA
 precedence 6-5
 no bump traffic
 protect group
 bump explicitly 7
 vbr-nrt 20000 10000 32
! The following commands create bundle1, add member named vcmember to
! bundle1, and then apply classA to vcmember:
bundle bundle1
 pvc-bundle vcmember
  class-vc classA
```

Taking into account hierarchy precedence rules, the VC bundle member *vcmember* will be characterized by these parameters:

- It carries traffic whose IP Precedence level is 6 and 5.
- It does not allow other traffic to be bumped onto it. When the VC goes down, its bumped traffic will be redirected to a VC whose IP Precedence level is 7.
- It is a member of the protected group of the bundle. When all members of a protected group go down, the bundle goes down.
- It supports the non-real-time variable bit rate service category.

VC Class at a VC Range Configuration Example

You can also define VC attributes by applying a VC class to a range of PVCs. To apply a VC class to an ATM PVC range, use the **class-range** PVC range configuration command.

The following example shows creation of a class called *classA* and its application to an ATM PVC range called *range-pppoa-1*:

```
! The following commands create the class named classA:
vc-class atm classA
 ubr 10000
 encapsulation aal5snap
! The following commands apply classA to an ATM PVC range:
interface atm 6/0.110 multipoint
 range range-pppoa-1 pvc 0/102 0/199
  class-range classA
```

VC Class at a PVC Within a Range Configuration Example

You can apply a VC class to an individual PVC within a range of PVCs. To apply a VC class to a PVC within a range, use the **class-vc** command in ATM PVC-in-range configuration mode.

The following example shows creation of a class called *classA* and its application to a PVC called *ppoa* within a range of PVCs:

```
! The following commands create the class named classA:
vc-class atm classA
 ubr 10000
 encapsulation aal5snap
! The following commands apply classA to the ATM PVC named pppoa
! within the range:
interface atm 6/0.110 multipoint
 range range-pppoa-1 pvc 0/102 0/199
  pvc-in-range pppoa 0/130
   class-vc classA
```

Choosing a Service Category

Once you have determined how you want to organize the configuration of your PVCs, as single or multiple PVCs in bundles or ranges, you need to choose the service categories to be supported by those PVCs for traffic shaping.

When determining the type of service category that you are going to support for one or more VCs, you need to determine the type of characteristics that are most important for the applications and types of traffic that those VCs support. You should also consider the traffic contract that you have established with your service provider.

You can reference Table 5-2 for some of the common types of applications that are typically represented by a service category supported on the PA-A3 and PA-A6 ATM port adapters. You can make a comparable selection according to your traffic needs. In addition, when choosing the type of service category for a PVC, it is important to assess whether any of the network applications require bandwidth guarantees or bounds on cell delay and cell loss.

Table 5-2 Common Applications by Service Category

Service Category	Application Examples
ABR	File transfer
CBR (for data, using AAL5)	Audio library, videoconferencing, video on demand
nrt-VBR	Interactive, bursty applications such as airline reservations or banking transactions
rt-VBR	Voice over ATM (VoATM), compressed voice over IP, video conferencing
UBR	File transfer, e-mail, library browsing, fax transmission, Telnet, LAN and remote office interconnections
UBR+ (SVCs only)	Same as UBR, but seek possible minimum bandwidth guarantee

It is important for you to realize that you are not necessarily restricted to a particular type of service category to carry your ATM traffic, nor are you bound to establish the same service categories on both ends of a link. However, each service category uses certain traffic parameters that best define the transmission characteristics of a type of traffic, and will help you to optimize your flow and to meet the requirements of your traffic contract.

In addition, each service category has a default transmission priority associated with it. Therefore, you will achieve the best use of bandwidth and optimize the performance for your traffic if you choose the service category that best represents the type of traffic and applications that will be carried over the PVC. The primary thing to keep in mind is that the ATM service category defines how the ATM network devices and the router treat the cells of the VC with respect to bandwidth guarantees, cell delay, and cell loss.

Table 5-3 provides a list of the traffic parameters by service category that you can reference to determine which parameters provide cell rate guarantees over the network.

Table 5-3 Cell Rate Guarantees by Service Category

Service Category	Traffic Parameter Used to Guarantee Cell Rate
ABR	Non-zero Minimum Cell Rate (MCR) if specified
CBR	PCR
nrt-VBR	SCR
rt-VBR	SCR
UBR	none
UBR+	Non-zero MCR if signalled by the router; applies to switched virtual circuits (SVCs) only on the PA-A3 and PA-A6 ATM port adapters

The remaining sections in this chapter provide information and guidelines for configuring the real-time and non-real-time service categories on the Cisco 7200 series router.

Configuring Real-Time Service Categories

There are two ATM service categories that are designed to support real-time applications, which require low cell delay and cell loss:

- **Constant bit rate (CBR)**—Supports real-time applications that request a static amount of bandwidth that is continuously available for the duration of the connection.
- **Real-time variable bit rate (rt-VBR)**—Supports real-time applications that have bursty transmission characteristics.

The following sections provide an overview of the real-time service categories that are supported on the PA-A3 and PA-A6 ATM port adapters, and include configuration guidelines and examples:

- Overview of CBR and Real-Time VBR Service Categories, page 5-11
- Configuring the CBR Service Category on a PVC, page 5-12
- Configuring the Real-Time VBR Service Category on a PVC, page 5-12

Overview of CBR and Real-Time VBR Service Categories

Both CBR and rt-VBR are typically used for voice and video applications. The rt-VBR service category is frequently used to support Voice over ATM (VoATM). To support these types of applications, both service categories place a limit on the cell delay variation, which is the variability in the arrival of adjacent cells. The primary difference between these two real-time service categories is the way in which the SAR processor services the PVCs with bandwidth.

CBR relies on the regular availability of scheduled time slots with small intercell gaps. PVCs using CBR are characterized by a steady interval of cell transmission over the line. They do not burst cells in clumps. Therefore, unless your PVC is transmitting at the line rate, it will only use some of the total bandwidth that is available.

The CBR service class is designed for ATM VCs needing a static amount of bandwidth that is continuously available for the duration of the active connection. An ATM VC configured as CBR can send cells at the peak cell rate (PCR) at any time and for any duration. It can also send cells at a rate less than the PCR or even emit no cells.

In contrast, rt-VBR makes better use of bandwidth when the traffic is bursty. Like CBR, rt-VBR also relies on the regular availability of scheduled time slots, but also allows the PVC to transmit cells in time slots using an ICG=1/PCR for limited periods of time.

Because of these differences, you do not configure the same traffic parameters for CBR and rt-VBR service categories. You only specify the PCR to implement the CBR service category. The rt-VBR service category implements the PCR, SCR, and MBS traffic parameters (as does the nrt-VBR service category).

CBR for Voice and CBR for Data

ATM defines a protocol stack consisting of three layers. The ATM adaptation layer (AAL) supports the QoS needs of an ATM service class such as CBR or nrt-VBR, and better enables an ATM network to carry different traffic types. AAL1 and AAL5 are the two most commonly used AAL types.

Cisco Systems differentiates between CBR for voice and CBR for data, depending on the AAL type supporting the CBR virtual connection:

- CBR for voice, which includes circuit emulation services (CES) and Voice over ATM (VoATM) applications, uses AAL1. A 1-byte AAL1 header uses time stamps, sequence numbers and other bits to help the ATM network handle ATM-layer defects like cell delay variation, cell misinsertion, and cell loss.
- CBR for data uses AAL5, and the same interface typically does not support CBR for voice. AAL5 adds an 8-byte trailer with a 4-byte cyclic redundancy check (CRC) for detecting errors in a protocol data unit (PDU).

The AAL sublayer functions, which include segmentation and reassembly (SAR), are performed only at the user side of a User-Network Interface (UNI) between a router or Catalyst ATM module and an ATM switch.

Distinguishing CBR and CES

It is important to clarify the differences between CBR and CES. CBR defines an ATM class of user traffic. In contrast, CES defines a method of carrying traffic from non-ATM telephony devices over an ATM cloud. In fact, CES provides an interworking function (IWF) that allows the two protocols to communicate.

To do this, Cisco ATM port adapters that support CES or voice CBR provide two interface types:

- One or more CBR interfaces (typically physical T1s or E1s)—Connects to non-ATM telephony devices, such as a private branch exchange (PBX) or time-division multiplexer (TDM). Use the **interface cbr** command on the PA-A2 to identify the CBR port.
- One ATM interface—Connects to the ATM cloud. Use the **interface atm** command to identify the ATM port.

With CES applications, the source router accepts standard T1 or E1 frames on the CBR port, converts these frames into ATM cells, and transmits the cells out the ATM interface through the ATM cloud. The destination router reassembles the ATM cells and sends them back through the interworking function to a CBR port. The CES specification mandates the transmission of voice traffic on CBR VCs.

To implement CES on the Cisco 7200 series router, use the PA-A2 ATM CES port adapters. To implement CBR for data on the Cisco 7200 series router, use the PA-A3 or PA-A6 ATM port adapters.

Configuring the CBR Service Category on a PVC

As of Cisco IOS Release 12.2, support was added for explicit configuration of CBR on the PA-A3 and PA-A6 ATM port adapters.

To configure a CBR PVC, use the **cbr** ATM VC configuration command as shown in the following example:

```
Router(config)# interface atm 1/0
Router(config-if)# pvc 0/100
Router(config-if-atm-vc)# cbr ?
  <1-155000>  Peak Cell Rate(PCR) in Kbps
```

Note CBR support was introduced for the PA-A3 ATM port adapter in the following Cisco IOS software releases: 12.2(5), 12.2(14)S, 12.2(8)T, and 12.2(15)B.

Prior to Cisco IOS Release 12.2, an explicit CLI to configure CBR on the PA-A3 ATM port adapter was not supported. To achieve an equivalent of this real-time service class on the PA-A3 for CBR for data, configure an nrt-VBR PVC and set the PCR and SCR to the same value, while lowering the transmission priority of the VC.

Note You can also configure the **cbr** command as part of a VC class. For more information, see the "Configuring VC Classes" section on page 5-2.

Configuring the Real-Time VBR Service Category on a PVC

The rt-VBR service category is intended for real-time applications, such as compressed Voice over IP (VoIP) and video conferencing, that require tightly constrained delays (cell transfer delay, or CTD) and cell delay variation (CDV). In some cases, the cells on a permanent virtual circuit (PVC) experience CDV when two or more VCs share a single ATM interface. You most commonly use the rt-VBR service category to support Voice over ATM (VoATM).

Delay can occur when the SAR processor needs to schedule cells for multiple PVCs, or when the SAR processor must schedule Operation, Administration, and Maintenance (OAM) cells along with servicing cells from one or more PVCs. As a result, the interarrival time between consecutive cells of a connection may vary. This phenomenon is known as *jitter*.

This section includes the following topics:

Real-Time VBR Configuration Guidelines

With rt-VBR traffic, you can expect the router to transmit in bursts at a rate that varies with time. To accomplish this, the rt-VBR service category implements the PCR, SCR, and MBS traffic parameters.

When configuring the rt-VBR service category, consider the following guidelines:

- When configuring VoATM, take care when calculating sufficient peak, average, and burst values, and ensure that the PVC can effectively handle the bandwidth for the number of voice calls.
- Use the following formulas as guidelines when establishing the traffic parameters for VoATM:

Traffic Parameter	Guideline for VoATM
PCR	(2 x maximum number of calls) x 16 Kbps
SCR	(1 x maximum number of calls) x 16 Kbps
MBS	4 x maximum number of calls

- The PCR range depends on the line rate of the ATM port adapter.
- In the CLI, the Average Cell Rate represents the SCR and is upwardly bounded by the PCR that you configure. See the "Real-Time VBR Configuration Examples" section on page 5-13.
- There are no default values, and the acceptable ranges vary by the type of ATM port adapter.
- The real-time VBR service category uses the same traffic parameters as the non-real-time VBR service category. For MBS, PCR, and SCR guidelines for both service categories, see the "MBS Configuration Guidelines" section on page 5-23 and the "PCR and SCR Configuration Guidelines" section on page 5-23.

Real-Time VBR Configuration Examples

To configure a rt-VBR PVC, use the **vbr-rt** ATM VC configuration command. The available values for traffic parameters vary by the type of ATM port adapter that you are configuring.

Note You can also configure the **vbr-rt** command as part of a VC class. For more information, see the "Configuring VC Classes" section on page 5-2.

PA-A3 OC-3 ATM Port Adapter rt-VBR Configuration Example

The following is an example of a rt-VBR configuration on a PA-A3 OC-3 ATM port adapter:

```
Router(config)# interface atm 1/0
Router(config-if)# pvc 0/100
Router(config-if-atm-vc)# vbr-rt ?
  <64-155000> Peak Cell Rate(PCR) in Kbps
Router(config-if-atm-vc)# vbr-rt 600 ?
  <64-600> Average Cell Rate in Kbps
Router(config-if-atm-vc)# vbr-rt 600 300 ?
  <1-64000>  Burst cell size in number of cells
Router(config-if-atm-vc)# vbr-rt 600 300 32 ?
```

PA-A3 E3 ATM Port Adapter rt-VBR Configuration Example

The following is an example of a rt-VBR configuration on a PA-A3 E3 ATM port adapter:

```
Router(config-if-atm-vc)# vbr-rt ?
  <1-34000>  Peak Cell Rate(PCR) in Kbps
Router(config-if-atm-vc)# vbr-rt 600 ?
  <1-600>  Average Cell Rate in Kbps
Router(config-if-atm-vc)# vbr-rt 600 300 ?
  <1-65535>  Burst cell size in number of cells
```

PA-A3 T3 ATM Port Adapter rt-VBR Configuration Example

The following is an example of a rt-VBR configuration on a PA-A3 T3 ATM port adapter:

```
Router(config-if-atm-vc)# vbr-rt ?
  <1-45000>  Peak Cell Rate(PCR) in Kbps
Router(config-if-atm-vc)# vbr-rt 1000 ?
  <1-1000>  Average Cell Rate in Kbps
Router(config-if-atm-vc)# vbr-rt 1000 1000 ?
  <1-65535>  Burst cell size in number of cells
```

Configuring Non-Real-Time Service Categories

There are three ATM service categories that are designed to support non-real-time applications, which typically support data services:

- **Available bit rate (ABR)**—Supports non-real-time applications that tolerate high cell delay, and can adapt cell rates according to changing network resource availability to prevent cell loss. The ABR service category is characterized by reactive congestion control, where it uses flow control mechanisms to learn about the network conditions and adjust cell rates accordingly.
- **Non-real time variable bit rate (nrt-VBR)**—Supports non-real-time applications with bursty transmission characteristics that tolerate high cell delay, but require low cell loss.
- **Unspecified bit rate (UBR)**—Supports non-real-time applications that tolerate both high cell delay and cell loss on the network. There are no network service-level guarantees for the UBR service category, and therefore it is a best-effort service.

Note Cisco Systems has also developed a second UBR service category called UBR+, which implements the MCR traffic parameter. UBR+ supports non-real-time applications that tolerate both high cell delay and cell loss on the network, but request a minimum guaranteed cell rate. As with the UBR service category, there are no network service-level guarantees for UBR+. UBR+ is available for SVCs only on the PA-A3 and PA-A6 ATM port adapters.

The following sections provide guidelines and describe how to configure these non-real-time service categories on the Cisco 7200 series router:

Configuring the ABR Service Category on a PVC

The ABR service category is designed for VCs that carry file transfer data and other bursty non-real-time traffic that simply requires some minimum amount of bandwidth (specified as a minimum cell rate) to be available while the VC is configured and active. With ABR, the delay or variation in delay from source to destination router can vary and can be a large value, making ABR unsuitable for real-time applications. As discussed previously, the CBR and VBR-rt service categories address applications that require tight boundaries on throughput and delay.

This section includes the following topics:

Overview of the ABR Service Category

ABR allows the ATM port adapter to transmit at a rate that varies with the amount of bandwidth available in the network. When the network is congested and other source devices are transmitting, there is little available bandwidth. However, when the network is not congested, additional bandwidth is available for use by active devices. ABR allows the router to take advantage of this extra bandwidth and increase its transmission rates. To accomplish this, ABR uses resource management (RM) cells to carry information about network congestion. The router uses the information contained in the RM cells to adjust its transmission rates.

An ABR VC binds a source router to a contract with the ATM switch network. As part of this contract, a source router agrees to examine information that indicates whether or not the network is congested and, in turn, adapt the source transmission rate if required. In return, the ATM switch network agrees to drop no more than a maximum number of cells when congestion occurs. The ratio of dropped cells to transmitted cells is known as the *cell loss ratio* (CLR).

Resource Management Cells

RM cells are standard 53-byte ATM cells with the payload type field in the header set to a binary value of 110. Forward RM cells are sent to the destination end-system on the same VC as data cells and at an interval defined by the Number of RM Cells (NRM) parameter. By default, a source ABR device sends one forward RM cell for every 32 data cells.

Table 5-4 describes fields in an RM cell:

Table 5-4 RM Cell Field Descriptions

Field	Octet	Bits	Description or Value
Header	1-5	all	ATM header
ID	6	all	Protocol identifier
Reserved	7	1-3	0
RA	7	4	Request Acknowledge; 0 or according to ITU-T, I.371
NI	7	5	Non-Increase: 0—Not congested 1—Non increase
CI	7	6	Congestion Indication: 0—Not congested 1—Congestion indication
BN	7	7	BECN cell: 0—Source generated
DIR	7	8	Direction: 0—Forward (source generated) 1—Backward (destination generated)
ER	8-9	all	Explicit cell rate
CCR	10-11	all	Commitment, concurrency, and recovery; current cell rate; the allowed cell rate by the source
MCR	12-13	all	Minimum cell rate; MCR parameter
QL	14-17	all	Queue length; 0 or according to ITU-T, I.371
SN	18-21	all	Sequence number; 0 or according to ITU-T, I.371
Reserved	22-51	all	6A (hex) for each octet
Reserved	52	3-8	0
CRC-10	53	all	Cyclic redundancy check; according to ATM Forum Traffic Management specifications

ABR Flow Control Mechanisms

ABR supports the following three rate-based methods of communicating congestion information from ATM switches and destination end-systems back to a source device:

- Binary—Uses the explicit forward congestion indication (EFCI) bit in ATM data cells to indicate congestion.
- Relative Rate—Uses the non-increase (NI) and congestion indication (CI) bits in either forward (to the destination) or backward (to the source) RM cells to indicate presence or absence of congestion. No actual rate is set in any RM cell rate fields.
- Explicit Rate—Uses the explicit rate field in backward RM cells to indicate at which rate the source router can transmit. More specifically, with the explicit rate flow control method, a source router places its current transmission rate in the commitment, concurrency, and recovery (CCR) field. Intermediate switches explicitly communicate the rate at which the source is allowed to send at that given moment by placing a value in the explicit rate (ER) field. The source router reads the ER field and adjusts its CCR to match the ER as long as the calculated rate is not less than the minimum cell rate.

These flow control methods are rate-based mechanisms, in which the ATM switch network communicates the rate at which the source can transmit. Rate-based mechanisms contrast with credit-based mechanisms, in which the network communicates the amount of buffer space available for a given VC. The source device transmits only if it knows that the network can buffer the data.

Binary Flow Control

Binary flow control was the first mechanism implemented for ATM networks. ATM switches set the EFCI bit in the headers of forward data cells to indicate congestion. When a destination router receives a data cell with the EFCI bit set, it marks the congestion indication bit in RM cells to indicate congestion and sends the RM cells back to the source.

A standard ATM cell header consists of five bytes. The payload type identifier (PTI) field consists of three bits, each of which defines a different parameter. The first bit indicates whether the cell contains user data or control data. If the cell contains user data, the second bit indicates whether the cell experienced congestion as it moved through the network. This second bit is known as the explicit forward congestion indication (EFCI) bit, and is used in the binary flow control method.

Relative Rate Flow Control

Congestion control schemes operate best when the latency of the feedback path is minimized. Relative rate mode can greatly reduce feedback delays and deliver better performance than the binary mode. Relative rate mode provides the ability for switches to source backward-RM cells to send a congestion indicator rather than relying on the destination end-system to turn around forward RM cells and map the EFCI bit to the CI bit in the backward RM cells.

Explicit Rate Flow Control

Explicit rate ABR is typically deployed in ATM WAN switches, and is used in products like the Cisco 8400 IGX and 8800 MGX ATM switches. Relative rate ABR is more effectively deployed in the campus and is supported by the Cisco Lightstream 1010 and Catalyst 8510 ATM switch routers. The Catalyst 8540 supports EFCI marking only. EFCI is typically used for backward compatibility with legacy ATM switches that support neither explicit rate nor relative rate ABR.

Cisco ATM port adapters implement all three ABR flow-control mechanisms. However, there is no option to select a specific mechanism. Instead, the router adapts to the format and indications received in the incoming RM cells. The mechanism used depends on the configuration of the ATM switches.

ABR Traffic Parameters

Table 5-5 provides information about some of the traffic parameters that the network uses to manage ABR traffic flows.

Note Only the PCR, MCR, RIF, and RDF traffic parameters are configurable on the Cisco 7200 series router.

Table 5-5 ABR Traffic Parameter Descriptions

ABR Traffic Parameter	Description
Peak Cell Rate (PCR)	Specifies the maximum cell rate at which the source can transmit. This traffic parameter is configurable on the PA-A3 and PA-A6 ATM port adapters.
Minimum Cell Rate (MCR)	Specifies the rate at which a source router can always send. This traffic parameter is configurable on the PA-A3 and PA-A6 ATM port adapters.
Initial Cell Rate (ICR)	Specifies the rate at which a source router should send when the interface first becomes active and when it begins transmitting again after an idle period.
Available or Allowed Cell Rate (ACR)	Specifies the current permitted rate at which the source router can send, based on dynamic feedback from the network.
Rate Increase Factor (RIF)	Specifies the amount by which the transmission rate increases after the source interface receives a resource management (RM) cell with NI and CI set to zero. This traffic parameter is configurable on the PA-A3 and PA-A6 ATM port adapters.
Rate Decrease Factor (RDF)	Specifies the amount by which the transmission rate decreases after the source interface receives an RM cell with the CI bit set to one. This traffic parameter is configurable on the PA-A3 and PA-A6 ATM port adapters.
Number of RM Cells (NRM)	Specifies the number of data cells sent between RM cells. By default, the source sends one RM cell for every 32 data cells. This value is not currently configurable on the PA-A3 and PA-A6 ATM port adapters.
Transient Buffer Exposure (TBE)	Specifies the number of cells that a source can transmit before receiving feedback from the network via a returned RM cell.
Fixed Round Trip Time (FRTT)	Indicates an estimate of the round trip time or the amount of time it takes for an RM cell to be transmitted from the source to the destination and back.

ABR Operation

When permitted, the source device begins transmitting cells at the ACR. The ACR is originally set to the ICR value, and is bounded by the MCR and PCR. The MCR and PCR traffic parameters are established at connection setup.

The router transmits an RM cell prior to the first data cell, and thereafter transmits an RM cell every 32 data cells. Within the RM cell, the router indicates its ACR by placing the ACR value in the CCR field of the RM cell. The ER field reflects the PCR, which is the maximum rate at which the source wants to transmit.

RM cells move forward through the network from the source to the destination. The destination returns the RM cells back to the source, logging information along the way to indicate congestion and rate status. If the destination is incurring internal congestion, it can alter the contents of the RM cell to reflect that congestion. In addition, while the RM cell travels through the network, the switches also can modify the RM cells to set the CI and NI bits of the RM cell.

Upon receipt of an RM cell, a source router first looks at the CI bit:

- If the CI bit is set, the source reduces its ACR by at least ACR x RDF, but no lower than the MCR value.
- If the CI bit is not set, the source increases its ACR by no more than RIF x PCR to a maximum of the PCR value.

Next, the source looks at the NI bit. If the NI bit equals zero, the source does not increase the ACR.

Finally, if the source router is using explicit rate flow control, it looks at the ER field (after calculating the new ACR based on the CI bit) and adjusts its rate to whichever is lower: the new ACR or the ER.

ABR Configuration Guidelines

When configuring the ABR service category, consider the following guidelines:

- Use the **abr** command to configure the PCR and MCR for the PVC. The PCR is the maximum rate at which the source router is allowed to transmit. The MCR can be set to zero, or it can be used to guarantee a minimum amount of bandwidth to the source router even during periods of congestion.
- Use the **atm abr rate-factor** command to specify the rate at which the ACR is increased or decreased. The default value for both RIF and RDF is 16, which results in a 1/16 factor. We recommend that you use the default values.

Software Guidelines

Support for the ABR service category on the PA-A3 ATM port adapter was introduced in Cisco IOS Release 12.0(4)T and Cisco IOS Release 12.0(5)S. ABR is now available in the Cisco IOS Release 12.1 mainline, Cisco IOS Release 12.1 T and Cisco IOS Release 12.1 E trains.

Note ABR is not available in the Cisco IOS Release 12.0 mainline.

For the PA-A3 ATM port adapter, we recommend the following minimum Cisco IOS Releases to support ABR:

- Cisco IOS Release 12.0(7)T and later
- Cisco IOS Release 12.0(8)S and later
- Cisco IOS Release 12.1(5) and later

ABR is available in all releases for the PA-A6 ATM port adapter. The PA-A6 ATM port adapter was introduced in Cisco IOS Release 12.2(15)B and Cisco IOS Release 12.2(15)T. It is also supported in Cisco IOS Release 12.3.

ABR service on the PA-A3 and PA-A6 ATM port adapters implement all three modes of rate control. This mode is selected automatically as the PA-A3 or PA-A6 ATM port adapter adapts to the format and indications received in the incoming RM cells.

ABR Configuration Example

To configure a PVC for the ABR service category, use the **abr** ATM VC configuration command as shown in the following example:

```
Router (config)# interface atm 1/0
Router(config-if)# pvc 1/32
Router(config-if-atm-vc)# abr 10000 3000
```

To change the default ABR rate-increase and rate-decrease factors, use the **atm abr rate-factor** interface configuration command as shown in the following example:

```
Router(config)# interface atm 1/0
Router(config-if)# atm abr rate-factor 32 32
```

Note You can also configure the **abr** command as part of a VC class. For more information, see the "Configuring VC Classes" section on page 5-2.

Configuring the Non-Real-Time VBR Service Category on a PVC

The non-real-time VBR service category is designed for VCs that carry bursty traffic and are not as sensitive to cell delay, but do not respond well to dropping cells. It is best suited to traffic that can benefit from short bursts up to the PCR, but not large bulk data transfers.

This section includes the following topics:

Overview of the Non-Real-Time VBR Service Category

The PA-A3 and PA-A6 ATM port adapters schedule nrt-VBR PVCs according to the "Leaky Bucket" algorithm. With this algorithm, an ATM VC needs to have a token in the bucket to transmit a cell. The algorithm replenishes tokens in the bucket at the rate of the SCR. If a source is idle and does not transmit for a period of time, tokens accumulate in the bucket. An ATM VC can use the accumulated tokens to burst at the rate of the PCR until the bucket is empty, at which point tokens are again replenished at the rate of the SCR. For more details about the Leaky Bucket algorithm and cell scheduling on the PA-A3 and PA-A6 ATM port adapters, see the "GCRA (Leaky Bucket) on the PA-A3 and PA-A6 ATM Port Adapters" section on page 2-28 and the "Scheduling on the PA-A3 and PA-A6 ATM Port Adapters" section on page 2-29.

It is important to understand that the PCR is a temporary burst. You derive the duration at which the VC can send at the PCR from the MBS. This translates to a "time on the wire."

Because the PCR burst is temporary, configure a VC for nrt-VBR if your traffic is bursty and can benefit from the short bursts at the PCR. Otherwise, if your traffic pattern consists of bulk data transfers, the PCR brings virtually no benefit. The reason is that to burst at the PCR, the ATM VC must send for some duration below the SCR.

How Bursting Works—Some Examples

This section describes some examples that show how bursting works and when it is most useful. To best understand these examples, it is helpful for you to understand how the GCRA works. For details about the GCRA, see the "GCRA (Leaky Bucket) on the PA-A3 and PA-A6 ATM Port Adapters" section on page 2-28.

How PCR Bursts Reduce Latency

To begin, assume that you need to transmit interactive traffic that consists of one 1500-byte packet (12000 bits) every second for a total of 12 Kbps. (Ignore ATM overhead in this example.) Configure nrt-VBR on the VC using the following values for the traffic parameters:

- PCR = 800 Kbps
- SCR = 64 Kbps
- MBS = 32 cells

A PCR of 800 Kbps means the first packet is sent in 15 microseconds (12-Kb packet / 800-Kbps PCR). It then takes 187.5 milliseconds (12000-byte packet / 64000-bps SCR) for the token bucket to replenish. The next packet is sent in 15 microseconds.

This example illustrates how PCR bursts reduce latency. Without the PCR, on a VC with only an SCR of 64 Kbps, it would take 187.5 milliseconds to send the first and the second packet.

Bursting and Large File Transfers

This example demonstrates how bursting has little benefit for large file transfers. In this example, use the same traffic parameters as in the previous example and assume that you need to transmit a large file.

Some applications, such as certain video devices, send very large IP packets up to 64 KB. In the previous example, the MBS value matches the size of a single 1500-byte packet. Using the 64-KB packet size, only the first packet is likely sent at the PCR. The average transfer rate will peak at the SCR because the tokens cannot accumulate. Therefore, nrt-VBR bursting offers little benefit for large file transfers.

When large packets exceed the maximum transmission unit (MTU) of the link, such as with a 64-KB packet, it can be useful to send the entire packet as a burst. In these cases, you might want to specify an MBS that accommodates the entire packet. In the example of a 64-KB packet, you should specify an MBS of 1334 cells (64 KB / 48 payload bytes per cell).

Understanding CDVT and Non-Real-Time VBR PVCs

Cell delay variation tolerance (CDVT) is a traffic parameter supported by ATM switches that is closely related to the performance of nrt-VBR PVCs over the network. CDVT establishes a tolerance level on the network switch to ATM interfaces whose PVCs temporarily exceed their traffic contract and aggressively send cells (back-to-back or very closely spaced) over the network. CDVT basically implements a "forgiveness factor" (measured in seconds) for PVCs that exceed the parameters of their traffic contract and it delays implementation of a Usage Parameter Control (UPC) penalty for the PVC.

UPC is the policing mechanism used by switches to control traffic contract violations on the network. UPC applies a mathematical formula to determine whether the traffic being sent by a router on a VC complies with the contract. Providers typically implement UPC policing on the first switch into the network at a point referred to as the User-Network Interface (UNI).

The per-VC UPC policy on Cisco ATM switches (such as the Cisco Catalyst 8500 or Cisco LightStream1010) specifies one of three actions to take with cells that it deems noncompliant:

- Drop the cells.
- Tag the cells by setting the Cell Loss Priority (CLP) bit in the ATM header.
- Pass the cells.

By default, UPC passes any noncompliant cells.

Here is a typical example of a set of rules that a UPC policy enforces at the ATM switch for a nrt-VBR VC:

- For cells that the switch receives that are at or below the SCR, the switch transmits those cells unchanged through the network.
- For cell bursts that the switch receives at rates above the SCR but below the PCR, the switch transmits the cells unchanged for burst sizes smaller than the MBS.
- For cells that the switch receives that are above the PCR and are deemed noncompliant, the switch either tags or discards the cells according to the configured UPC action on the switch.
- For cell bursts up to the PCR that exceed the MBS number of cells and are deemed noncompliant by the switch, the switch either tags or discards the cells according to the configured UPC action on the switch.

Non-Real-Time VBR Configuration Guidelines

This section provides guidelines for configuring the traffic parameters associated with the nrt-VBR service category. Table 5-6 provides a description of the traffic parameters that the network uses to manage nrt-VBR traffic flows.

Table 5-6 Non-Real-Time VBR Traffic Parameter Descriptions

NRT-VBR Traffic Parameter	Description
Maximum Burst Size (MBS)	Specifies the number of cells that can be transmitted at the PCR. Although MBS is specified as a number of cells, it translates to an amount of time, or a duration at which the router sends at PCR.
Peak Cell Rate (PCR)	Specifies the maximum rate at which you expect to transmit cells.
Sustainable Cell Rate (SCR)	Specifies the rate at which you expect to continuously transmit cells.

MBS Configuration Guidelines

When configuring MBS, specify a number of cells that accommodates the typical packet size you expect for bursty traffic. Consider the MBS along with PCR as a means of reducing latency, not increasing bandwidth.

To appropriately configure the MBS, you need to understand cell times. You should understand how the MBS value translates to a PCR duration when provisioning nrt-VBR VCs. For example, recall the formula for calculating the cell time over a DS-1 link:

1 cell / 3622 cells per second = 276.04 microseconds per ATM cell

On a DS-1 link, an MBS value of 100 equates to a PCR duration of 2.8 milliseconds.

For more information about cell rates supported by other types of ATM interfaces, see Table 2-7 on page 2-26.

You can also calculate this time in seconds using the following formula:

T = (burst cells x 424 bits per cell) / (PCR – SCR)

MBS yields the duration at which the router sends at the PCR. MBS accommodates temporary bursts or short spikes in the traffic pattern. An MBS of 100 cells allows a burst of three MTU-size Ethernet frames or one MTU-size Fiber Distributed Data Interface (FDDI) frame. It is important that you consider longer duration bursts when you specify the SCR.

When large packets exceed the maximum transmission unit (MTU) of the link, it can be useful to send the entire packet as a burst. In these cases, you might want to specify an MBS that accommodates the entire packet.

There is no official definition of a burst. Think of a burst in terms of MTU-sized frames or whatever size frame the traffic pattern presents. This frame will then break down into some number of cells. To maximize the use of bursting, you should understand when you use the MBS and follow the recommendations for establishing its size.

PCR and SCR Configuration Guidelines

When specifying the PCR and SCR traffic parameters for nrt-VBR VCs, consider the following guidelines:

- PCR—Derive this rate in combination with MBS in order to achieve the desired latency for bursty traffic. Consider the PCR as a means of decreasing the latency of a VC rather than increasing its bandwidth.
- SCR—Specify this rate as though your traffic is constrained to it, as if it were a constant bit-rate circuit, and without concern for latency. In other words, consider SCR as the true bandwidth of the VC and not the long-term average rate.
- If you configure PCR=SCR, the burst calculation for MBS is ignored and the credit is set to 1, regardless of the burst size.

Note Configuring nrt-VBR on a PVC with the PCR and SCR set to the same value, while lowering the transmit priority of the VC, provides equivalent real-time service class performance on the PA-A3 and PA-A6 ATM port adapters for CBR for data. However, Cisco IOS Release 12.2 introduces two new SAR priority levels to support proper prioritization for CBR and rt-VBR when competition for cell time slots arises. It also introduces the ability to configure CBR and rt-VBR at the command line, which makes the configuration of nrt-VBR to model a real-time service category unnecessary.

- Be sure that the router and the switch are basing their rate setting and policing on the same cell size. Some processors interpret rates based on 48-byte cells, and others use 53-byte cells. When implementing the PCR and SCR on ATM port adapters for Cisco Systems routers, the SAR processor accounts for the 5-byte ATM cell header, AAL5 padding, and an AAL5 trailer.
- When configuring multiple PVCs, be sure that the sum of the SCR values for all PVCs is less than the supported line rate.

Cisco ATM router interfaces do not support any Kbps values in the range from zero to the line rate. Instead, they support a set of incremental values.

It is important to be aware that the available values in Kbps include the bandwidth consumed by user data as well as by all ATM overhead, including the 5-byte cell header, cell padding, and AAL5 overhead.

When selecting PCR and SCR values, refer to Table 5-7 on page 5-24 and Table 5-8 on page 5-24, which describe the officially supported values for each interface hardware type.

Table 5-7 Supported Values for PCR, SCR, and MBS on the PA-A3 ATM Port Adapters

ATM Port Adapter	PCR and SCR	MBS
PA-A3-OC3	Supports increments of 4.57 Kbps	Supports increments of 1 cell
PA-A3-T3/E3	• For DS-3, supports increments of 1.33 Kbps • For E3, supports increments of 1.03 Kbps	Supports increments of 1 cell

Table 5-8 Supported Values for PCR, SCR, and MBS on the PA-A6 ATM Port Adapters

ATM Port Adapter	PCR and SCR	MBS
PA-A6-OC3	Supports increments of 2.28 Kbps	Supports increments of 1 cell
PA-A6-T3/E3	• For DS-3, supports increments of 0.665 Kbps • For E3, supports increments of 0.515 Kbps	Supports increments of 1 cell

Non-Real-Time VBR Configuration Examples

To configure a PVC for the nrt-VBR service category, use the **vbr-nrt** ATM VC configuration command. The following example configures a nrt-VBR PVC with a PCR of 55 MB, an SCR of 50 MB, and an MBS of 100 cells:

```
Router(config)# interface atm 5/0
Router(config-if)# pvc 1/200
Router(config-if-atm-vc)# ?
ATM virtual circuit configuration commands:
  abr                     Enter Available Bit Rate (pcr)(mcr)
  broadcast               Pseudo-broadcast
  cbr                     Enter Average Cell Rate in Kbps.
  class-vc                Configure default vc-class name
  create                  Configure VC Auto-creation Type
  dbs                     Dynamic Bandwidth Selection
  default                 Set a command to its defaults
  dialer                  set dialer pool this pvc belongs to
  encapsulation           Select ATM Encapsulation for VC
```

```
  exit-vc                 Exit from ATM VC configuration mode
  idle-timeout            Set idle time for disconnecting this SVC/AutoVC
  ilmi                    Configure ILMI management
  inarp                   Change the inverse arp timer on the PVC
  ip addr inarp           Assign an ip address to the atm interface through
                          ATMInarp
  max-reserved-bandwidth  Maximum Reservable Bandwidth on a vc
  no                      Negate a command or set its defaults
  oam                     Configure oam parameters
  oam-pvc                 Send oam cells on this pvc
  pppoe                   PPPoE options
  protocol                Map an upper layer protocol to this connection.
  random-detect           Configure WRED
  service-policy          Attach a policy-map to a VC
  transmit-priority       set the transmit priority for this VC
  tx-ring-limit           Configure PA level transmit ring limit
  ubr                     Configure Unspecified Bit Rate (UBR) for this
                          interface
  vbr-nrt                 Enter Variable Bit Rate (pcr)(scr)(bcs)
  vbr-rt                  Enter Variable Bit Rate (pcr)(average)
  vc-hold-queue           Configure hold queue size
  vcci                    VCC Identifier
  vpn                     Configure VPN parameters
Router(config-if-atm-vc)# vbr-nrt 55000 50000 100
```

Note You can also configure the **vbr-nrt** command as part of a VC class. For more information, see the "Configuring VC Classes" section on page 5-2.

MBS Configuration Example

The following example configures a PVC for nrt-VBR service with an MBS of 100 cells:

```
Router(config)# int atm 6/0.2
Router(config-subif)# pvc 1/100
Router(config-if-atm-vc)# vbr-nrt 55000 50000 ?
 <1-65535>  Maximum Burst Size(MBS) in Cells
Router(config-if-atm-vc)# vbr-nrt 55000 50000 100
```

Verifying a Burst

It can be tricky to detect a burst on an ATM VC.

It is important to understand that the ATM interface only bursts when the ATM VC has transmitted for a duration below the SCR. If the ATM VC has always transmitted at SCR, then no burst credits have accumulated and no bursting occurs.

To actually "see" a burst, we recommend using the following test procedure:

Note You need a traffic generator and an ATM cell tester for this procedure.

Step 1 Configure the PVC with a PCR that is two times the Kbps rate of the SCR.

Step 2 Start the cell tester.

Step 3 Start the traffic generator and transmit at a rate above the PCR.

Step 4 On the cell tester, look at the measured intercell gap. The cell tester reports a smaller intercell gap, which indicates the burst.

Step 5 Stop the cell tester and continue sending at PCR on the traffic generator.

Step 6 Start the cell tester again. You will not see the burst. This is because the traffic generator has always sent cells above the PCR (and/or above the SCR). The ATM VC has never sent cells below the SCR and therefore has never accumulated enough credits to send above the SCR again.

When configuring the traffic shaping values for a nrt-VBR VC, factor any sustained bursts into the SCR. As illustrated with the test procedure above, the MBS is not designed for sustained transmission above the SCR.

Configuring the UBR Service Category on a PVC

UBR is intended for non-real-time applications that do not require any maximum bound on the transfer delay or on the cell loss ratio. It is a best-effort service category. UBR is typically used for data communications applications such as file transfer and e-mail. Other applications include fax transmission, Telnet, and LAN and remote office interconnections. Such applications are not sensitive to cell delay, but they are sensitive to cell loss.

ATM switches, such as the Cisco Catalyst 8500 series, allocate larger maximum per-VC queue limits for UBR PVCs. Queueing minimizes loss at the expense of greater delay.

Overview of the UBR Service Category

Neither an ATM-attached router nor an ATM switch provides traffic shaping or quality of service guarantees to a UBR VC. As a result, UBR VCs can experience a large number of cell drops or a high cell transfer delay as cells move from the source to the destination device.

The only traffic parameter that you can specify on a Cisco Systems router for UBR is the PCR. Some ATM switches do not enforce the PCR, and the value of PCR becomes informational only. On switched virtual circuits (SVCs) defined for UBR service, a router communicates to the network that a virtual circuit is configured for UBR by using the best effort indicator field in the ATM User Cell Rate Information Element (IE) of a signaling packet.

The PA-A3 and PA-A6 ATM port adapters allocate bandwidth based on the priorities of the PVCs. Each service category that is available on these ATM port adapters has a a default transmission priority associated with it. For information on the default priorities, see the "Default PVC Priorities" section on page 5-29.

If an ATM port adapter supports multiple PVCs implementing different service categories, there are priorities in how the interface allocates bandwidth. An ATM port adapter first allocates bandwidth for the PCR of a CBR VC followed by AAL5 and AAL2 VoATM VCs. Next, the port adapter allocates bandwidth for the rt-VBR and nrt-VBR classes according to their PCRs and SCRs. Finally, the adapter allocates bandwidth for the MCR of the ABR VCs. Any remaining bandwidth is available for the VCs of the other service classes such as UBR. However, the amount of remaining bandwidth and when it appears is not guaranteed.

For more details about how the PA-A3 and PA-A6 ATM port adapters schedule cells, see the "Scheduling on the PA-A3 and PA-A6 ATM Port Adapters" section on page 2-29, the "Collision Handling" section on page 2-31, and the "PVC Priorities" section on page 2-35.

Advantages of UBR

The UBR service category provides the following advantages:

- Allows for a high degree of statistical multiplexing by not reserving any minimum bandwidth per VC. The VCs use the bandwidth up to the configured PCR when available.
- Models the best-effort service normally provided by the Internet. Best-effort service is suitable for applications tolerant to delay and not requiring real-time response.

Disadvantages of UBR

The UBR service category has the following disadvantages:

- The PCR only provides an indication of a physical bandwidth limitation within a VC.
- VCs of other ATM service categories have a higher transmission priority by the ATM port adapter's segmentation and reassembly (SAR) scheduler. When competition for a cell time slot arises, the scheduler will give the time slot to a VC of service classes with a higher priority.
- UBR does not place any bounds with respect to the cell loss ratio (CLR) or to the cell transfer delay (CTD). The end-system is expected to handle and adjust for any cell loss or cell delay.
- Retransmission occurs at higher layers because UBR does not guarantee cell delivery.

UBR Configuration Guidelines

If you configure a PVC without specifying any shaping parameters, the router implements a UBR PVC with a PCR equal to the line rate of the physical ATM interface.

UBR Configuration Example

To configure a PVC for the UBR service category, use the **ubr** ATM VC configuration command as shown in the following example:

```
Router(config)# interface atm 5/0
Router(config-if)# pvc 1/300
Router(config-if-atm-vc)# ubr 50000
```

Note You can also configure the **ubr** command as part of a VC class. For more information, see the "Configuring VC Classes" section on page 5-2.

Configuring the UBR+ Service Category—SVCs Only

Cisco Systems has developed a second UBR service category called UBR+, which implements the Minimum Cell Rate (MCR) traffic parameter for SVCs only.

UBR+ supports non-real-time applications that tolerate both high cell delay and cell loss on the network, but request a minimum guaranteed cell rate. As with the UBR service category, there are no network service-level guarantees for UBR+. However, the network can grant a service-level guarantee for the requested MCR.

Overview of the UBR+ Service Category

While UBR defines only an optional PCR traffic parameter, UBR+ adds the minimum cell rate (MCR) traffic parameter. On an ATM switch, UBR+ also defines a cell delay variation tolerance (CDVT).

It is important to understand that UBR+ implements MCR as a "soft guarantee" of minimum bandwidth. A router signals the MCR value at call setup time when creating a switched VC. It is up to the ATM switch to grant the bandwidth requested in the MCR parameter. In other words, a UBR+ VC is a UBR VC for which the MCR is signaled by the router and guaranteed by the ATM switch. Thus, UBR+ affects connection admission control and resource allocation on ATM switches.

In other words, UBR+ offers ATM interfaces the ability to communicate both the desired minimum and maximum cell rates to the ATM network. From a traffic shaping perspective, the SVC continues to be created as a standard UBR VC on the router.

Difference Between UBR+ and ATM Forum's MDCR Implementation

UBR+ is a special ATM service category developed by Cisco Systems. A similar UBR service category is specified in an addendum to the ATM Forum Traffic Management specifications, which discusses implementation of an optional Minimum Desired Cell Rate (MDCR) parameter for the UBR service category.

However, the Cisco Systems and the ATM Forum implementation vary in how the minimum rate is signalled to the ATM network. Cisco Systems uses the existing MCR information element (IE) that is used by the ABR service category, but the parameter has a different interpretation. For UBR+, the MCR parameter represents a desired cell rate; but in ABR, the MCR specifies the lowest acceptable cell rate.

The ATM Forum does not use the MCR IE, but implements a new IE for the MDCR traffic parameter.

UBR+ Support on the PA-A3 and PA-A6 ATM Port Adapters

UBR+ is not applicable to PVCs. However, in releases prior to Cisco IOS Release 12.0(7)T, the command is still available within the CLI for PVC configuration—but it should not be used. If you apply a VC class with the **ubr+** command to a PVC, the Cisco IOS software assigns the UBR class to the PVCs. The router rejects a VC class on a PVC if the PCR and MCR defined in the **ubr+** command are higher than the line rate of the underlying physical interface.

UBR+ Configuration Guidelines

When configuring the UBR+ service category for SVCs, consider the following guidelines:

- The UBR+ service category is only applicable to SVCs on the PA-A3 and PA-A6 ATM port adapters. UBR+ is not supported on the PA-A1 or PA-A2 ATM port adapters.
- The **ubr+** command was introduced in Cisco IOS Release 11.3 T for SVCs on the PA-A3 ATM port adapter. In Cisco IOS Release 12.0(3)T, the **ubr+** command was enhanced to support configuration of UBR+ and configuration of output PCR and output MCR for VC bundles.
- You can optionally specify *input-pcr* and *input-mcr* parameters for a UBR+ SVC. You only need to specify the input parameters when your output and input parameters differ.
- If you omit the input parameters on the UBR+ SVC, the router automatically assigns the same values to these parameters as the output parameters.

UBR+ Configuration Example

To configure an SVC for the UBR+ service category, use the **ubr+** ATM VC configuration command as shown in the example. The following example specifies the *output-pcr*, *output-mcr*, *input-pcr*, and *input-mcr* arguments for an ATM SVC to be 10,000 Kbps, 3000 Kbps, 9000 Kbps, and 1000 Kbps, respectively:

```
Router(config)# interface atm 1/0
Router(config-if)# svc lion nsap 47.0091.81.000000.0040.0B0A.2501.ABC1.3333.3333.05
Router(config-if-atm-vc)# ubr+ 10000 3000 9000 1000
```

Configuring PVC Priorities

The PVC priority is used by the SAR scheduler to determine a transmission hierarchy when PVCs compete for the same scheduling time slot. When collisions occur, the SAR scheduler always gives the PVC with the higher priority precedence over a PVC of lower priority in the link list.

If you are experiencing a lot of collisions and need to increase the performance of a particular PVC, you can modify the default PVC priority. Remember that the transmit priority only affects the likelihood that the VC will be given priority access to a particular cell time slot. The transmit priority does not change the behavior of the SAR scheduler and does not implement a minimum bandwidth guarantee.

For more information about how the PVC priority affects collision handling by the SAR scheduler, see the "Collision Handling" section on page 2-31 and "PVC Priorities" section on page 2-35.

Default PVC Priorities

The ATM port adapter implements a default set of PVC priorities based upon the service category that you configure for the PVC. However, you can override the default priorities for a PVC using the **transmit-priority** command.

Note Prior to Cisco IOS Release 12.2(4), the number of PVC priorities was four. The range of PVC priorities was increased from four to six priorities beginning in the following Cisco IOS software releases: 12.2(5), 12.2(14)S, 12.2(8)T, and 12.2(15)B.

As of Cisco IOS Release 12.2, the PA-A3 and PA-A6 ATM port adapters support the following default classes of PVC priorities:

- Priority 1—CBR, OAM cells, and signaling (Integrated Local Management Interface [ILMI], Q.2931 Signaling ATM Adaptation Layer [QSAAL])
- Priority 2—AAL5 or AAL2 VoATM VC (any service category)
- Priority 3—Real-time VBR
- Priority 4—Non-real-time VBR
- Priority 5—ABR
- Priority 6—UBR, UBR+

Transmit Priority Guidelines

When configuring the transmit priority, take care to avoid using a transmit priority of 1, which should be reserved for control traffic like OAM and signaling.

Changing the Default PVC Priority

To change the default transmit priority, use the **transmit-priority** command in ATM VC configuration mode as shown in the following example:

```
Router(config)# int atm 1/0.1
Router(config-subif)# pvc 0/100
Router(config-if-atm-vc)# transmit-priority ?
  <1-6>  priority level
Router(config-if-atm-vc)# transmit-priority 3
Router(config-if-atm-vc)# end
```

Verifying the PVC Priority

To verify the PVC priority, use the **show atm vc** privileged EXEC command. The following example displays a transmit priority of 3 for the PVC:

```
Router# show atm vc 1
ATM1/0.1:VCD:1, VPI:0, VCI:100
VBR-NRT, PeakRate:10000, Average Rate:5000, Burst Cells:0
AAL5-LLC/SNAP, etype:0x0, Flags:0x20, VCmode:0x0
OAM frequency:0 second(s)
InARP frequency:15 minutes(s)
Transmit priority 3
InPkts:0, OutPkts:0, InBytes:0, OutBytes:0
InPRoc:0, OutPRoc:0
InFast:0, OutFast:0, InAS:0, OutAS:0
InPktDrops:0,  OutPktDrops:0
InByteDrops:0, OutByteDrops:0
CrcErrors:0, SarTimeOuts:0, OverSizedSDUs:0, LengthViolation:0, CPIErrors:0
Out CLP=1 Pkts:0
OAM cells received:0
OAM cells sent:0
Status:INACTIVE
```

Verifying Traffic Shaping Configuration

To verify the traffic shaping configuration on your PVC, use the **show atm vc** and **show interfaces atm** privileged EXEC commands. For more information about these commands, refer to the *Cisco IOS Wide-Area Networking Command Reference* publication.

Consider the following guidelines when interpreting the output from these commands:

- The shaping engine does not report the ATM cell header overhead, AAL5 padding, or the AAL5 trailer in the **show atm vc** or **show interfaces atm** command output fields.
- The shaping engine does not differentiate between actual data bytes and padding or filler payload. An ATM cell must contain 48 bytes in the payload field. An ATM interface uses two cells to transmit a 64-byte IP packet. In the second cell, "wasted" payload in the form of padding is counted by the ATM switch, but ignored by the router. Thus, unused cell payload can prevent the actual bit rate from reaching the SCR.

Understanding ATM and AAL Overhead

When using ATM, overhead to the payload data occurs in a couple of forms:

- ATM Layer overhead
- AAL Layer overhead

ATM Layer Overhead

ATM Layer overhead consists of the 5-byte (40-bit) cell header that is added to the 48-byte payload data to make a 53-byte cell. This ATM header is sometimes called the ATM cell "tax." The ATM cell header is not included in the output of the **show atm vc** or **show interfaces atm** commands.

Figure 5-1 shows the parts of an ATM cell header.

Figure 5-1 Format of an ATM Cell Header

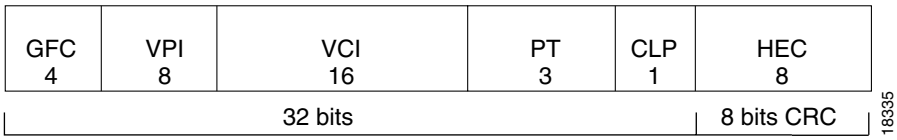

AAL Overhead

The AAL provides additional information to the protocol data unit (PDU) payload to support the QoS needs of a particular ATM service category. The amount of overhead varies by the type of AAL being used. AAL5 is the most common type and its format is described in RFC 1483.

AAL5 adds the following types of overhead:

- 8-byte trailer
- 8-byte encapsulation header for Logical Link Control/Subnetwork Access Protocol (LLC/SNAP) (this is the default header for ATM PVCs on the Cisco 7200 series router)
- Variable amount of padding (up to 47 bytes) to make the AAL5 PDU an even 48-byte multiple.

The AAL overhead is not included in the output of the **show atm vc** or **show interfaces atm** commands.

Using the show atm vc Command

The **show atm vc** command provides information about your PVC configuration, including the service category, encapsulation, ring limits, and transmit priority.

Note When you want to reference the VPI/VCI value, use the **show atm pvc** command.

The following example shows that the UBR service category has been configured with a PCR of 10000 Kbps for the PVC identified by virtual circuit descriptor (VCD) 2:

```
Router# show atm vc 2
VC 2 doesn't exist on interface ATM2/0
ATM5/0: VCD: 2, VPI: 1, VCI: 100
UBR, PeakRate: 10000
AAL5-LLC/SNAP, etype:0x0, Flags: 0x20, VCmode: 0x0
OAM frequency: 0 second(s)
PA TxRingLimit: 0 particles
PA Rx Limit: 0 particles
InARP frequency: 15 minutes(s)
Transmit priority 4
InPkts: 0, OutPkts: 0, InBytes: 0, OutBytes: 0
InPRoc: 0, OutPRoc: 0, Broadcasts: 0
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
InPktDrops: 0, OutPktDrops: 0
CrcErrors: 0, SarTimeOuts: 0, OverSizedSDUs: 0
OAM cells received: 0
OAM cells sent: 0
Status: ACTIVE
```

Note The number of bytes reported in the OutBytes field does not include the size of the AAL5 trailer or padding, or ATM cell header.

Displaying the VCD of a PVC

To display the VCD of a PVC, use the **show atm pvc** privileged EXEC command:

```
Router# show atm pvc
           VCD/                                        Peak       Avg/Min   Burst
Interface  Name   VPI   VCI   Type   Encaps      Kbps          Kbps   Cells       Sts
2/0        1        0     5    PVC     SAAL    155000        155000              UP
2/0        2        0    16    PVC     ILMI    155000        155000              UP
2/0.2      101      0    50    PVC     SNAP    155000        155000              UP
2/0.2      102      0    60    PVC     SNAP    155000        155000            DOWN
2/0.2      104      0    80    PVC     SNAP    155000        155000              UP
2/0        hello    0    99    PVC     SNAP      1000                            UP
```

Using the show interfaces atm Command

You can use the **show interfaces atm** privileged EXEC command to see several important values related to your ATM configuration and performance. The following statistics in particular provide some key information:

- Input and output rate in bits per second and packets per second (five minutes is the default period).
- Input and output queue size and the number of drops.
- Input error counters such as cyclic redundancy checks (CRCs), ignores, and no buffers.

The following sample output from a PA-A3 ATM port adapter shows that the port adapter has experienced 11,184 output queue drops since the counters were last cleared one week and one day ago:

```
Router# show interfaces atm 5/0/0
   ATM5/0/0 is up, line protocol is up
   Hardware is cyBus ENHANCED ATM PA
   MTU 4470 bytes, sub MTU 4470, BW 149760 Kbit, DLY 80 usec, rely 255/255,
   load 2/255
   Encapsulation ATM, loopback not set, keepalive set (10 sec)
   Encapsulation(s): AAL5 AAL3/4
   4096 maximum active VCs, 7 current VCCs
   VC idle disconnect time: 300 seconds
   Last input never, output 00:00:00, output hang never

   Last clearing of "show interface" counters 1w1d
   Queueing strategy: fifo

   Output queue 0/40, 11184 drops; input queue 0/150, 675 drops
   5 minute input rate 1854000 bits/sec, 382 packets/sec
   5 minute output rate 1368000 bits/sec, 376 packets/sec
   155080012 packets input, 3430455270 bytes, 0 no buffer
   Received 0 broadcasts, 0 runts, 0 giants
   313 input errors, 313 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
   157107224 packets output, 1159429109 bytes, 0 underruns
   0 output errors, 0 collisions, 0 interface resets
   0 output buffers copied, 0 interrupts, 0 failures
```

Note The statistics do not include the size of the AAL5 trailer or padding, or ATM cell header.

Measuring Traffic Shaping Accuracy

To measure traffic shaping accuracy, we strongly recommend that you use an ATM traffic analyzer.

If you want to use the **show interfaces atm** output, we recommend that you translate the SCR that you configured into a packets-per-second unit of measure. Then, refer to the packets/sec field in the **show interfaces atm** command output. You might also notice that a larger packet size normally produces a bit rate that is closer to the configured SCR.

You should avoid using the 5 minute output rate field of the **show interfaces atm** command output, which is reported as an average bits-per-second rate based on a default load interval of 5 minutes. To be more reactive to short bursts of traffic, you can adjust the load interval down to a minimum of 30 seconds using the **load-interval** interface configuration command.

Related Documentation

The following table provides information about additional resources that you can read to learn more about some of the topics discussed in this chapter:

For more information about:	Refer to the following publications:
ATM technical standards	Approved ATM Forum Specifications
Configuring VCs and VC parameters	*Cisco IOS Wide-Area Network Configuration Guide* and *Cisco IOS Wide-Area Network Command Reference*
Configuring VC bundles	*Cisco IOS Quality of Service Solutions Configuration Guide* and *Cisco IOS Quality of Service Solutions Command Reference*
Configuring VC classes	*Cisco IOS Wide-Area Network Configuration Guide* *Cisco IOS Quality of Service Solutions Configuration Guide* and *Cisco IOS Quality of Service Solutions Command Reference*
Configuring VC ranges	*Cisco IOS Wide-Area Network Configuration Guide* and *Cisco IOS Wide-Area Network Command Reference*

Next Steps

This chapter provides detailed information and guidelines for you to configure traffic shaping on the PA-A3 and PA-A6 ATM port adapters. As we describe in Chapter 2, "Cisco 7200 Series Architecture and Design for ATM Traffic Management," traffic shaping produces congestion on the router. Therefore, you need to manage your hardware and Layer 3 queues on the Cisco 7200 series router to optimize the complete flow of ATM traffic.

You can find information about configuring Layer 3 queues on the Cisco 7200 series routers in Chapter 6, "Configuring QoS on the Layer 3 Queues for the PA-A3 and PA-A6 ATM Port Adapters."

Chapter 7, "Configuring the Ring Limits on the PA-A3 and PA-A6 ATM Port Adapters" provides information about optimizing the hardware queues on the PA-A3 and PA-A6 ATM port adapters.

Configuring QoS on the Layer 3 Queues for the PA-A3 and PA-A6 ATM Port Adapters

In Chapter 2, "Cisco 7200 Series Architecture and Design for ATM Traffic Management," you learned that the PA-A3 and PA-A6 ATM port adapters use both hardware and software queues to support the receive and transmit processing of ATM traffic on a Cisco 7200 series router. The hardware queues, in the form of transmit and receive buffers on the ATM port adapter itself, operate on a first-in first-out (FIFO) basis and are not configurable. The hardware queues operate with the private interface pools (and their receive rings and transmit rings) on the network processing engine (NPE) or network services engine (NSE) for the storage of content during receive and transmit processing.

The software queues operate at Layer 3 and activate when congestion builds on the router and outbound traffic cannot be processed on the transmit ring. Control of the packets awaiting transmit processing passes to the Layer 3 queues according to the quality of service (QoS) policies that you configure. If you do not configure any service policies, then the default behaviors apply. Process-switched packets automatically enqueue to the Layer 3 queues regardless of the state of congestion on the router.

The PA-A3 and PA-A6 ATM port adapters support per-VC queueing at Layer 3. With this design, you can prevent any single VC from starving other VCs for resources.

This chapter provides a brief introduction and some guidelines for configuring the IP to ATM class of service (CoS) features on the PA-A3 and PA-A6 ATM port adapters. The chapter does not describe the full details about the supported QoS implementations in the Cisco IOS software, or all of the possible configuration options for QoS on the PA-A3 and PA-A6 ATM port adapters.

For in-depth information about configuring QoS, refer to the *Cisco IOS Quality of Service Solutions Command Reference* and the *Cisco IOS Quality of Service Solutions Configuration Guide* publications for your software release. Refer to the "Related Documentation" section on page 6-14 for additional references about configuring QoS.

This chapter includes the following sections:

Preparing to Configure QoS

Before you begin to configure QoS for the Layer 3 queues, you should have a good understanding of the Cisco 7200 series architecture and how the router processes ATM traffic. In particular, you need to understand when Layer 3 queues are activated. This information is provided in Chapter 2, "Cisco 7200 Series Architecture and Design for ATM Traffic Management."

You also need to evaluate your network traffic and assess your business needs to establish the criteria for appropriate service policies before you configure QoS on your network. A description of the kinds of information that you should acquire and the tasks that you should perform are shown in Chapter 4, "Preparing to Configure ATM Traffic Management and QoS Features."

Architecture Overview

This section describes some of the important characteristics about the architecture used by the Cisco 7200 router during ATM processing that might be helpful for you to review as you prepare to configure QoS service policies:

- The Cisco 7200 series router activates Layer 3 queues whenever congestion builds on an egress interface and outbound traffic can not be processed to the transmit ring, or onto the hardware queue located on the ATM port adapter. Traffic shaping is frequently a cause for congestion on the egress ATM interface.
- For PA-A3 and PA-A6 ATM port adapters, there is a hold queue for each PVC that is configured on that interface. This environment provides more control and prevents any single over-subscribed PVC from starving other PVCs for transmission resources. This queue is called the per-VC queue.
- With the exception of process-switched packets, whenever entries are available for packets on the transmit ring, packets go directly to the transmit ring on the NPE and onto the FIFO hardware queue located on the ATM port adapter. Process-switched packets always enqueue to the Layer 3 queue first, before being placed onto the transmit ring, regardless of availability on the ring.
- QoS service policies only begin to apply to CEF-switched and fastswitched packets when there is congestion on the ATM port adapter and the transmit ring is full.

IP to ATM CoS Overview

The Cisco IOS software classifies QoS features into the following categories:

- Classification and Marking
- Congestion Avoidance
- Congestion Management
- Traffic Shaping and Policing
- Signaling
- Link Efficiency Mechanisms

IP to ATM Class of Service (CoS) refers to a subset of the overall QoS features available on the Cisco 7200 series router that enable you to specify queueing service policies on a per-VC basis. IP to ATM CoS identifies certain QoS features that can be specifically applied at a more discreet, per-VC level for PVCs on the PA-A3 and PA-A6 ATM port adapters:

- Congestion avoidance policies determine the drop behavior in a queue. For congestion avoidance, you can configure Weighted Random Early Detection (WRED) on a per-VC basis on the PA-A3 and PA-A6 ATM port adapters. The default behavior is tail drop.
- Congestion management policies determine the order of the queue. For congestion management, you can configure Low Latency Queueing (LLQ) or Class-Based Weighted Fair Queueing (CBWFQ) on a per-VC basis on the PA-A3 and PA-A6 ATM port adapters. The default congestion management policy is FIFO.

It is important to remember that queueing occurs on the outbound path only. Therefore, you can configure a service policy that specifies a non-default queueing strategy, such as WRED, CBWFQ, or LLQ, for outbound traffic only on a PVC. If you attempt to configure WRED, CBWFQ, or LLQ as an inbound policy, the command will be rejected and a message similar to the following appears:

```
*Jun  4 07:27:17.210:  CBWFQ : Can be enabled as an output feature only
```

Although only WRED, CBWFQ, and LLQ are supported on a per-VC, outbound basis, you can still use other QoS features to classify and mark different IP traffic on the inbound path in combination with implementing IP to ATM CoS features at the PVC. This chapter focuses only on guidelines for the IP to ATM CoS features.

Understanding the Queue Limit

The queue limit is an important concept to understand when configuring IP to ATM CoS features, and there are several aspects to consider. The hold queue limit specifies the maximum number of packets that can be held in the Layer 3 queue. The queue limit is called the *queue depth* also. For some port adapters, this is an interface level queue that all VCs share. For the PA-A3 and PA-A6 ATM port adapters, it is a per-VC queue, and the interface queue is not used.

The type of queueing strategy implemented on the VC determines the default queue limit size, and also determines which command you might use to customize the queue depth. When the number of packets in the queue reaches the queue limit for that VC, then the router initiates a drop policy. By default, there technically is not a congestion avoidance policy. The default behavior is tail drop, which occurs when the queue is full. When the VC uses a tail drop policy, the last packet in is the first packet dropped, as long as there is not any available space in the queue.

To minimize tail drop, you can configure WRED as an alternative congestion avoidance policy to implement drop probabilities among different classified flows of traffic. WRED provides intelligent dropping. For more information, see the "Configuring WRED" section on page 6-6.

There are two different types of queue limits that you can tune for the PA-A3 and PA-A6 ATM port adapters, which are used for different types of queueing strategies:

- Per-VC queue—This Layer 3 queue applies at the VC level when FIFO queueing is being used by the VC. The default size is 40 packets, but you can tune the size of the queue using the **vc-hold-queue** ATM VC configuration command.
- Class queue—This Layer 3 queue applies at the class level when CBWFQ is being used by the VC. The default size is 64 packets, but you can tune the size of the queue using the **queue-limit** policy-map class configuration command.

  CBWFQ creates a queue for every class for which a class map is defined. Each class has a queue limit associated with it, which specifies the maximum number of packets that can enqueue there. Packets that satisfy the match criteria for a class accumulate in the queue reserved for the class until

they are sent, which occurs when the fair queueing process services the queue. When the class queue reaches the maximum packet threshold, enqueueing of any further packets to the class queue causes tail drop. If Weighted Random Early Detection (WRED) is configured for the class policy, intelligent packet drop takes effect before the queue limit is reached, while the average queue depth is between the minimum and maximum thresholds.

Configuring the Queue Limits

You can configure a per-VC queue limit when you use the FIFO queueing strategy at a VC, or you can configure a class queue limit at the VC if you are configuring a class-based policy. The default queue limit varies by the type of queueing policy that you configure.

For LLQ, the original default limit was 64. The queue limit is not configurable for LLQ. However, as of Cisco IOS release 12.1(3)T, the queue limit automatically adjusts to the configured bandwidth to accommodate packet bursts. For more information, see the *Configuring Burst Size in Low Latency Queueing* feature documentation.

Configuring the FIFO Per-VC Hold Queue Limit

For FIFO, the default maximum number of packets that can be in the per-VC queue is 40 packets. When using the default FIFO policy on a PVC, you can modify the queue limit associated with each VC using the **vc-hold-queue** ATM VC configuration command. The possible queue limits are 5 to 1024 packets.

FIFO Per-VC Hold Queue Limit Configuration Example

The following example specifies that PVC 0/100 can hold up to 50 packets before activating its drop policy during FIFO queueing:

```
Router(config)# interface atm3/0.1
Router(config-if)# pvc 0/100
Router(config-if-atm-vc)# vc-hold-queue 50
```

Verifying the Per-VC Hold Queue Limit

To verify the size of the per-VC hold queue on an ATM interface, use the **show queueing interface atm** command and observe the value of the "Output queue" field. The following example shows that the total possible per-VC queue depth is 50 packets, and no packets are currently in the queue (shown by "Output queue 0/50"):

```
Router# show queueing interface atm 3/0
  Interface ATM3/0 VC 0/100
  Queueing strategy: fifo
  Output queue 0/50, 0 drops per VC
[text omitted]
```

Configuring the Class Queue Limit

For CBWFQ, the default limit is 64 packets. When using a CBWFQ policy on a PVC, you can modify the queue limit associated with a class using the **queue-limit** policy-map class configuration command. The possible queue limits are 1 to 64 packets.

When you configure CBWFQ, the per-VC hold queue limit does not apply to the VC.

Class Queue Limit Configuration Example

The following example configures a policy map called "fairq" for the default class. The queue limit for this class is 50 packets:

```
Router(config)# policy-map fairq
Router(config-pmap)# class class-default
Router(config-pmap-c)# fair-queue
Router(config-pmap-c)# queue-limit 50
```

Verifying the Class Queue Limit

Once you have configured CBWFQ and attached the policy to a VC, you can use the **show queueing interface atm** command to verify the queueing strategy and observe the value of the "threshold" field to verify the limit for the class queue. The following example shows that the WFQ class limit on VC 10/32 is 50 packets:

```
Router# show queueing interface atm 2/0.100032
     Interface ATM2/0.100032 VC 10/32
     Queueing strategy: weighted fair
     Total output drops per VC: 1539
     Output queue: 0/512/50/1539 (size/max total/threshold/drops)
        Conversations  0/37/128 (active/max active/max total)
        Reserved Conversations 0/0 (allocated/max allocated)
```

Using MQC to Configure and Apply QoS Service Policies

The Modular QoS CLI (MQC) is a command-line interface (CLI) structure that allows you to create service polices and attach these policies to interfaces, subinterfaces, and ATM or Frame Relay virtual circuits (VCs). A traffic policy contains a traffic class and one or more QoS features. A traffic class is used to classify traffic, while the QoS features in the service policy determine how to treat the classified traffic.

For ATM port adapters, you can use MQC to create traffic policies and attach these policies at the PVC level, or at the interface level. For the PA-A3 and PA-A6 ATM port adapters, you should attach the policies to the VC, not the interface.

To create QoS classes and configure policies using MQC on the PA-A3 and PA-A6 ATM port adapters, complete the following basic steps:

Step 1 Define a traffic class using the **class-map** command.

Step 2 Create a traffic policy and associate the traffic class with one or more QoS features using the **policy-map** command.

Step 3 Attach the traffic policy to the PVC using the **service-policy** command.

For examples using MQC configuration, see Chapter 8, "ATM Traffic Management Case Studies and Configuration Examples." You can also find MQC configuration examples in the *Cisco IOS Quality of Service Solutions Configuration Guide* and *Cisco IOS Quality of Service Solutions Command Reference*.

Configuring WRED

Weighted Random Early Detection (WRED) is the Cisco Systems implementation of the Random Early Detection (RED) algorithm, which combines the capabilities of the RED algorithm with consideration of IP precedence. WRED flows are classified by different IP precedence levels.

WRED generally drops packets selectively based on IP precedence. Packets with a higher IP precedence are less likely to be dropped than packets with a lower precedence. Thus, the higher the priority of a packet, the higher the probability that the packet will be delivered.

WRED makes early detection of congestion possible and provides for multiple classes of traffic. It also protects against global synchronization and is therefore most useful in networks that transmit a large amount of TCP traffic. WRED reduces the chances of tail drop by selectively dropping packets when the output interface begins to show signs of congestion. By dropping some packets early rather than waiting until the queue is full, WRED avoids dropping large numbers of packets at once and minimizes the occurrence of global synchronization. Thus, WRED allows the bandwidth to be more efficiently used at all times.

Although WRED is an alternative method of congestion avoidance, it does not necessarily eliminate tail drop from occurring. WRED is a more proactive implementation than tail drop, because the WRED algorithm applies before the queue becomes full. If the queue limit is reached on the VC, then the tail drop strategy always applies. Therefore, you need to be careful about how you configure your queue limits and the maximum threshold for WRED to be sure that WRED can be properly activated before the queue limit is reached. If you configure WRED within a class policy, then the queue limit for the class applies.

For more information about global synchronization, see the "Congestion on an ATM Network" section on page 1-10. For more information about WRED, refer to the "Congestion Avoidance" section of the *Cisco IOS Quality of Service Solutions Configuration Guide*. You also can refer to the *IP to ATM Class of Service Phase 1 Design Guide* for more in-depth analysis of WRED implementation considerations. However, be aware that much of the discussion in the *IP to ATM Class of Service Phase I Design Guide* is based on the Cisco 7500 series router.

WRED Configuration Guidelines

When you configure WRED on a VC for the PA-A3 and PA-A6 ATM port adapters, consider the following guidelines:

- Be sure to enable Cisco Express Forwarding (CEF) switching, which is a requirement for WRED.
- Use WRED on VCs where you expect high congestion to occur, or that will be transmitting a high volume of Transmission Control Protocol (TCP) traffic.
- You can configure both WRED and a fancy queueing mechanism like CBWFQ or LLQ.
- Be aware that by default, non-IP traffic is treated with a precedence of zero, which means that your non-IP traffic will be dropped more often than your TCP/IP traffic.
- To transport IP best-effort traffic on the IP backbone, you need to implement a consistent policy for use of precedence values for different types of traffic throughout the network. In particular, you should perform precedence marking of IP traffic on the edge of the network (for example, marking of incoming IP precedence through the Cisco IOS Committed Access Rate [CAR] feature or through policy routing).
- WRED is supported at the interface level. However, for the PA-A3 and PA-A6 ATM port adapters, you should configure WRED at the per-VC level using the **random-detect** command. You can use a QoS policy map to apply WRED to the VC.
- The queue limit defines the maximum number of packets that the Layer 3 queues can store at any time. When the mean queue depth is between the minimum and maximum thresholds, WRED applies.
- If you change the queue limit but you are using the default WRED settings, then the maximum threshold automatically adjusts to the configured queue limit. However, if you manually configure the max-threshold, then you will lose the benefit of dynamic adjustment.

  The queue limit should be equal to or larger than the WRED max-threshold. If the per-VC queue limit is smaller, then the WRED mechanism can not be fully implemented on the VC because tail drop is enforced when the number of packets in the queue reaches the queue limit.
- Use the default parameters for WRED. The default WRED settings are very robust and automatically implement the following considerations:
  - The experience developed in the Internet research community on RED parameter setting
  - Configuration of related parameters (such as shaping parameters of the ATM VC on which WRED is run)
  - A different discard profile per precedence (the higher the precedence, the better the default corresponding service)
- The default values allocate the same max-thresholds and the same mark-probability to all the precedences. However, the default min-threshold is different for every precedence. The higher the precedence, the higher the min-thresholds. Consequently the default WRED configuration offers an increasingly better service to higher precedences.

Caution Because of the dynamic nature of RED and WRED and their complex interactions with transport-level flow control mechanisms (such as TCP flow control), fine-tuning WRED to achieve specific IP service differentiation objectives in particular operating conditions is a delicate exercise and great caution is recommended. We recommend that you start operations or testing with the default WRED settings (or

from configurations close to the WRED default settings) and fine-tune from there. Modifications to WRED parameters should be tested and validated under a vast range of network conditions before being deployed in a large network.

WRED Configuration Example on an ATM PVC

To configure WRED using MQC, complete the following steps:

Step 1 From global configuration mode, enable IP CEF:

```
Router(config)# ip cef
```

Step 2 From global configuration mode, create the policy and configure WRED. The following example creates a policy named "atm_wred." The **bandwidth** command implements WFQ, and the **random-detect** command without any other parameters enables WRED using the default weights and precedence for the class named "mytest" in the policy:

```
Router(config)# policy-map atm_wred
Router(config-pmap)# class mytest
Router(config-pmap-c)# bandwidth 64
Router(config-pmap-c)# random-detect
Router(config-pmap-c)# exit
```

Step 3 From global configuration mode, create the PVC at the interface and apply the service policy. The following example configures PVC 1/120 at an ATM point-to-point interface and applies the service policy named "atm_wred" for traffic outbound on the PVC:

```
Router(config)# interface ATM1/0.20 point-to-point
Router(config-if)# ip address 10.20.20.21 255.255.255.0
Router(config-if)# pvc 1/120
Router(config-if-atm-vc)# vbr-nrt 150 100 120
Router(config-if-atm-vc)# service-policy output atm_wred
```

For more information about configuring WRED, refer to the *Cisco IOS Quality of Service Solutions Configuration Guide*.

Monitoring WRED Status on a VC

To see how many packets have been selectively dropped by the WRED algorithm on a per-VC basis, you can use the **show policy-map interface atm** command and observe the value of the "Random drop" counters. The following example shows that no drops are present on VC 1/120:

```
Router# show policy-map interface atm 1/0.20 out
 ATM1/0.20: VC 1/120 -

  Service-policy output: atm_wred

    Class-map: mytest (match-all)
      169 packets, 191676 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group 121
      Queueing
        Output Queue: Conversation 25
        Bandwidth 64 (kbps)
        (pkts matched/bytes matched) 112/187344
```

```
  (depth/total drops/no-buffer drops) 0/0/0
   exponential weight: 9
   mean queue depth: 0

class    Transmitted       Random drop      Tail drop    Minimum Maximum  Mark
         pkts/bytes        pkts/bytes       pkts/bytes    thresh  thresh prob
    0     112/187344           0/0              0/0          20      40 1/10
    1       0/0                0/0              0/0          22      40 1/10
    2       0/0                0/0              0/0          24      40 1/10
    3       0/0                0/0              0/0          26      40 1/10
    4       0/0                0/0              0/0          28      40 1/10
    5       0/0                0/0              0/0          30      40 1/10
    6      63/4788             0/0              0/0          32      40 1/10
    7       0/0                0/0              0/0          34      40 1/10
 rsvp       0/0                0/0              0/0          36      40 1/10
```

Configuring CBWFQ

Native (flow-based) WFQ assigns a weight to each conversation, and then schedules the transmit time for each packet of the different flows. The weight is a function of the IP precedence of each flow, and the scheduling time depends on the packet size.

CBWFQ assigns a weight to each configured class instead of each flow. The bandwidth you assign to a class is used to calculate the weight of that class. More precisely, the weight is a function of the interface bandwidth divided by the class bandwidth. Therefore, the bigger the bandwidth parameter, the smaller the weight.

Without LLQ, CBWFQ provides weighted fair queueing based on defined classes with no strict priority queue available for real-time traffic. CBWFQ allows you to define traffic classes and then assign characteristics to those classes. For example, you can designate the minimum bandwidth delivered to a class during congestion.

For CBWFQ, the weight for a packet belonging to a specific class is derived from the bandwidth you assign to the class when you configure it. Therefore, the bandwidth assigned to the packets of a class determines the order in which packets are sent. All packets are serviced fairly based on weight; no class of packets may be granted strict priority.

This scheme poses problems for voice traffic that is largely intolerant of delay, especially variation in delay. For voice traffic, variations in delay introduce irregularities of transmission manifesting as jitter in the heard conversation. To overcome this limitation for voice traffic, you can use LLQ with CBWFQ.

CBWFQ Configuration Guidelines

When you configure CBWFQ on a VC for the PA-A3 and PA-A6 ATM port adapters, consider the following guidelines:

- The size of the transmit ring limit determines how quickly the Layer 3 queue is activated. Therefore, when you plan to implement WFQ, you should reduce the transmit ring limit. For more information, see Chapter 7, “Configuring the Ring Limits on the PA-A3 and PA-A6 ATM Port Adapters.”
- Implement CBWFQ on slow PVCs if you do not want bulk traffic to impact the transmission of smaller packet traffic.
- Modify the queue limit for each class using the **queue-limit** policy-map class configuration command. The possible queue limits are 1 to 64 packets.
- You can configure WRED with CBWFQ as an alternative drop strategy.

- Because CBWFQ provides a minimum bandwidth guarantee, you can only apply CBWFQ to VCs with classes of service other than UBR and UBR+. These service categories are best-effort classes that do not guarantee a minimum bandwidth.
- The PA-A3 and PA-A6 ATM port adapters do not support native, flow-based WFQ configured directly on an interface using the **fair-queue** command. You need to configure WFQ within the default class using a policy map to implement CBWFQ for the PA-A3 and PA-A6 ATM port adapters.
- Use class maps to classify and assign weights to traffic. Classification parameters and class maps are defined at the same place.
- After defining the classification parameters, configure a policy map to apply traffic parameters to these classified flows.
- After you configure the traffic parameters for each class, apply CBWFQ on a VC-basis using the **service-policy output** ATM VC configuration command.
- Traffic that does not match one of the defined class maps is assigned a default class map (class default) that you define in the policy map. The parameters configured under this default class apply to all non-classified traffic.

CBWFQ Configuration Example

To configure CBWFQ using MQC, complete the following steps:

Step 1 From global configuration mode, create the policy and configure fair queueing. The following example creates a policy named “cbwfq” for the class called “mytest.” The **bandwidth** command implements CBWFQ for the class:

```
Router(config)# policy-map cbwfq
Router(config-pmap)# class mytest
Router(config-pmap-c)# bandwidth 256
Router(config-pmap-c)# end
```

Step 2 Beginning in global configuration mode, create the interface, create the PVC, and apply the service policy. The following example configures PVC 0/101 at an ATM point-to-point interface and applies the service policy named “cbwfq” for traffic outbound on the PVC:

```
Router(config)# interface atm 4/0.11 point-to-point
Router(config-subif)# ip address 10.10.10.1 255.255.255.0
Router(config-subif)# pvc 0/101
Router(config-if-atm-vc)# vbr-nrt 2048 1024 96
Router(config-if-atm-vc)# service-policy output cbwfq
```

For more information about configuring CBWFQ, refer to the *Cisco IOS Quality of Service Solutions Configuration Guide*.

Monitoring CBWFQ Status on a VC

To monitor the status of CBWFQ on a VC, complete the following steps:

Step 1 From privileged EXEC mode, run the **show queue** command for the VC that you want to see:

```
Router# show queue atm 4/0.11 vc 0/101
  Interface ATM4/0.11 VC 0/101
  Queueing strategy: weighted fair
  Output queue: 56/512/64/38858 (size/max total/threshold/drops)
     Conversations  2/2/64 (active/max active/max total)
     Reserved Conversations 1/1 (allocated/max allocated)
     Available Bandwidth 512 kilobits/sec

  (depth/weight/total drops/no-buffer drops/interleaves) 1/228/9284/0/0
  Conversation 73, linktype: ip, length: 994
  source: 10.0.0.2, destination: 10.10.10.2, id: 0x0000, ttl: 63, prot: 255

  (depth/weight/total drops/no-buffer drops/interleaves) 55/32384/29574/0/0
  Conversation 44, linktype: ip, length: 994
  source: 10.0.0.2, destination: 10.10.10.3, id: 0x0000, ttl: 63, prot: 255
```

Step 2 To verify how many packets are currently in the hold queue across all conversations, observe the value of the "Output queue" size field. The following example shows that the per-VC queue is 56 packets for VC 0/101:

```
Router# show queue atm 4/0.11 vc 0/101
  Interface ATM4/0.11 VC 0/101
  Queueing strategy: weighted fair
  Output queue: 56/512/64/38858 (size/max total/threshold/drops)
     Conversations  2/2/64 (active/max active/max total)
[text omitted]
```

Step 3 To verify how many packets are in the queue for each conversation flow, observe the value of the "depth" fields, as shown in the following output excerpt:

```
[text omitted]
(depth/weight/total drops/no-buffer drops/interleaves) 1/228/9284/0/0
  Conversation 73, linktype: ip, length: 994
  source: 10.0.0.2, destination: 10.10.10.2, id: 0x0000, ttl: 63, prot: 255

  (depth/weight/total drops/no-buffer drops/interleaves) 55/32384/29574/0/0
  Conversation 44, linktype: ip, length: 994
  source: 10.0.0.2, destination: 10.10.10.3, id: 0x0000, ttl: 63, prot: 255
[text omitted]
```

Step 4 To verify how many packets you can queue for each conversation, observe the value of the "Output queue" threshold field as shown in the following output excerpt. Notice that for the WFQ queueing strategy, the number of possible packets in the hold queue is 64 (the default). The total number of packets that all conversations can queue is 512:

```
Router# show queue atm 4/0.11 vc 0/101
  Interface ATM4/0.11 VC 0/101
  Queueing strategy: weighted fair
  Output queue: 56/512/64/38858 (size/max total/threshold/drops)
[text omitted]
```

Step 5 To verify information about the classes for the output policy on a PVC, run the **show policy-map interface** command:

```
Router# show policy-map interface atm 4/0.11 vc 0/101
 ATM4/0.11: VC 0/101 -

  Service-policy output: cbwfq

    Class-map: mytest (match-all)
      153656 packets, 152734064 bytes
      30 second offered rate 230000 bps, drop rate 0 bps
      Match:  precedence 6
      Queueing
        Output Queue: Conversation 73
        Bandwidth 256 (kbps) Max Threshold 64 (packets)
        (pkts matched/bytes matched) 146307/145429158
        (depth/total drops/no-buffer drops) 0/9284/0

    Class-map: class-default (match-any)
      257250 packets, 255704576 bytes
      30 second offered rate 224000 bps, drop rate 93000 bps
      Match: any
```

Configuring LLQ

The Low Latency Queueing feature brings strict priority queueing to CBWFQ. Configured by the **priority** command, strict priority queueing gives delay-sensitive data, such as voice, preferential treatment over other traffic. With this feature, delay-sensitive data is sent first, before packets in other queues are serviced.

LLQ Configuration Guidelines

When you configure LLQ on a VC for the PA-A3 and PA-A6 ATM port adapters, consider the following guidelines:

- The size of the transmit ring limit determines how quickly the Layer 3 queue is activated. Therefore, when you plan to implement LLQ you should reduce the transmit ring limit. For more information, see Chapter 7, “Configuring the Ring Limits on the PA-A3 and PA-A6 ATM Port Adapters.”
- The **priority** command allows you to set up classes based on a variety of criteria (not just User Datagram Protocol [UDP] ports) and assign priority to them, and is available for use on serial interfaces and ATM permanent virtual circuits (PVCs). A similar command, **ip rtp priority**, allows you to define priority flows based only on UDP port numbers, but it is not available for ATM PVCs.
- Layer 2 encapsulations are accounted for in the amount of bandwidth specified with the **priority** command. However, the amount of bandwidth does not include other headers such as ATM cell tax overheads. You must also allow bandwidth for possible jitter introduced by the routers in the voice path.
- Use the **priority** command for Voice over IP (VoIP) on serial links and ATM PVCs.

Note In Cisco IOS Release 12.0(7)T and Cisco IOS Release 12.1, the **priority** command does not support VoIP over Frame Relay links. As of Cisco IOS Release 12.2 and later, the **priority** command is supported on Frame Relay links.

- You cannot configure the **priority** command with the **random-detect** command (for WRED), or with the the **queue-limit** command (to configure class queue depth). The **bandwidth** command and **priority** command are mutually exclusive.
- You can configure the **priority** command in multiple classes, but you should only use it for voice-like, constant bit rate (CBR) traffic.
- The functionality of LLQ has been extended to allow a configurable Committed Burst (Bc) size using the *Configuring Burst Size in Low Latency Queueing* feature. With this new functionality, the network can now accommodate temporary bursts of traffic and handle network traffic more efficiently.

LLQ Configuration Examples

The following example configures strict priority queueing with a guaranteed bandwidth of 50 kbps for the policy map named “llq”:

```
Router(config)# policy-map llq
Router(config-pmap)# class voice
Router(config-pmap-c)# priority 50
```

In this example, traffic classes named “voice” and “video” go into the high priority queue and get strict priority queueing over data traffic. However, voice traffic will be rate-limited to 50 Kbps and video traffic will be rate-limited to 100 Kbps. The classes will be individually rate-limited even if they go into the same queue:

```
Router(config)# policy-map llq
Router(config-pmap)# class voice
Router(config-pmap-c)# priority 50
Router(config-pmap-c)# exit
Router(config-pmap)# class video
Router(config-pmap-c)# priority 100
Router(config-pmap-c)# exit
Router(config-pmap)# class data
Router(config-pmap-c)# bandwidth 500
```

Monitoring LLQ Status on a VC

To monitor LLQ status on a VC, use the **show policy-map interface atm** command.

The following sample output was obtained from an ATM PVC with an SCR of 1024 Kbps. For LLQ, the queueing system adjusts the burst size as the value of the **priority** command changes:

```
Router# show policy-map interface atm 4/0.11 vc 0/101
 ATM4/0.11: VC 0/101 -

  Service-policy output: llq

    Class-map: data (match-all)
      79793 packets, 79314242 bytes
      30 second offered rate 254000 bps, drop rate 0 bps
      Match: ip precedence 0
```

```
  Queueing
    Output Queue: Conversation 73
    Bandwidth 500 (kbps) Max Threshold 64 (packets)
    (pkts matched/bytes matched) 38326/38096044
    (depth/total drops/no-buffer drops) 0/1693/0

Class-map: voice (match-all)
  996 packets, 93624 bytes
  30 second offered rate 5000 bps, drop rate 0 bps
  Match: ip precedence 5
  Queueing
    Strict Priority
    Output Queue: Conversation 72
    Bandwidth 50 (kbps) Burst 1250 (Bytes)
    (pkts matched/bytes matched) 0/0
    (total drops/bytes drops) 0/0

Class-map: video (match-all)
  1029 packets, 96726 bytes
  30 second offered rate 5000 bps, drop rate 0 bps
  Match: ip precedence 6
  Queueing
    Strict Priority
    Output Queue: Conversation 72
    Bandwidth 100 (kbps) Burst 2500 (Bytes)
    (pkts matched/bytes matched) 503/47282
    (total drops/bytes drops) 0/0

Class-map: class-default (match-any)
  1 packets, 32 bytes
  30 second offered rate 0 bps, drop rate 0 bps
  Match: any
```

Monitoring QoS on the PA-A3 and PA-A6 ATM Port Adapters

To monitor per-VC drop counters on the PA-A3 and PA-A6 ATM port adapters, you need to use the **show queueing interface atm** command. Do not use the **show atm vc** command.

Related Documentation

The following table provides information about additional resources that you can read to learn more about some of the topics discussed in this chapter:

For more information about:	Refer to the following publications:
Cisco IOS QoS software commands	*Cisco IOS Quality of Service Solutions Command Reference*
Cisco IOS QoS software features	*Cisco IOS Quality of Service Solutions Configuration Guide*
Applying Service Policies for ATM	Where Do I Apply a QoS Service Policy on an ATM Interface? (TAC Tech Note)
Burst sizes and LLQ	Configuring Burst Size in Low Latency Queueing (Cisco IOS feature module)

For more information about:	Refer to the following publications:
CBWFQ an transmit ring limit relationship	Understanding Class Based Weighted Fair Queueing on ATM (TAC Tech Note)
Per-VC CBWFQ	Per-VC Class-Based, Weighted Fair Queueing (Per-VC CBWFQ) on the Cisco 7200, 3600, and 2600 Routers (TAC Tech Note)
QoS FAQs	QoS Frequently Asked Questions
WRED implementation and fine-tuning	*IP to ATM Class of Service Phase 1 Design Guide*

Next Steps

This chapter provides guidelines and information about queue limits and how to configure the IP to ATM CoS features that are supported on the PA-A3 and PA-A6 ATM port adapters, including WRED, CBWFQ, and LLQ.

To activate Layer 3 queues, you might need to optimize the size of the transmit ring for the ATM port adapter on the NPE/NSE. For more information, see Chapter 7, "Configuring the Ring Limits on the PA-A3 and PA-A6 ATM Port Adapters."

For in-depth examples and case studies of QoS configuration in an ATM network, see Chapter 8, "ATM Traffic Management Case Studies and Configuration Examples."

Configuring the Ring Limits on the PA-A3 and PA-A6 ATM Port Adapters

The PA-A3 and the PA-A6 ATM port adapters provide a way for you to limit the consumption of the receive ring and transmit ring resources on the NPE or NSE on a per-VC basis. The effect of these per-VC limits on the rings is a division of the resource into logical, per-VC queues. The default limits for both rings is calculated using internal logic based upon configured parameters (such as traffic shaping values) for the VC.

Chapter 2, "Cisco 7200 Series Architecture and Design for ATM Traffic Management," provides an in-depth discussion about how the receive ring and transmit ring structures work on the NPE or NSE in the overall processing of ATM traffic on the Cisco 7200 series router. It describes the relationship of the receive rings and transmit rings to the private interface pools on the NPE or NSE and also to the hardware buffers located on the ATM port adapters. It also discusses the relationship of the Layer 3 queues to the transmit ring.

This chapter describes how to optimize the ring limits on the PA-A3 and PA-A6 ATM port adapters. It includes the following sections:

Preparing to Configure the Ring Limits

Before you begin to configure the ring limits, you should have a good understanding of the Cisco 7200 series architecture and how the router processes ATM traffic.

Caution It is important to consider that the ring limits for the receive and transmit side are effectively operating against the same resource—particles within the private interface pool. Therefore, you must be very careful if you plan to tune these limits. Just as with adjustments to the buffer pools, improper settings for the receive ring or transmit ring limits can adversely impact system performance. In this case, adjustments to either side of the ring limits can impact the performance of both receiving and transmitting packets. Only modify the ring limits after careful evaluation of network impact or when recommended by technical support personnel.

Architecture Overview

This section reviews some of the important characteristics about the memory architecture during ATM processing that you should understand as you prepare to configure the ring limits. For more details, see Chapter 2, "Cisco 7200 Series Architecture and Design for ATM Traffic Management."

There are several memory structures that are used in the processing of ATM traffic on a Cisco 7200 series router. Some of these structures are located on the ATM port adapter hardware itself, while others are located on the NPE or NSE:

- Transmit and receive buffers—Particle-based hardware buffers located on the PA-A3 and PA-A6 ATM port adapters where actual content is stored upon receipt of ATM cells or in preparation of cell transmission. These buffers are fixed and you cannot optimize them. A DMA engine transfers content between the hardware buffers on the ATM port adapter and the private interface pool on the NPE or NSE.
- Private interface pool—Particle-based memory located on the NPE or NSE where actual packet content also is stored for receive and transmit processing. The private interface pool that is associated with the ingress interface (where the data is received from the network) stores the content. The private interface pool is typically a static pool that you cannot optimize. For the PA-A3 and PA-A6 ATM port adapters, the default number of particles in the private interface pool varies by the type of NPE or NSE in use. For more information, see Table 2-1 on page 2-7.
- Transmit and receive rings—Control structures located on the NPE or NSE that point to the stored location of packets in the private interface pool for receive or transmit processing. A corresponding receive ring and a transmit ring control structure is associated with each physical interface on the Cisco 7200 series router.

  You can optimize the size of the receive and transmit rings on a per-VC basis, which effectively determines how much of the private interface pool can be used for receive and transmit processing over that VC.

Ring Limit Overview

Operation of the receive rings and transmit rings is described in detail in the "Receive Rings and Transmit Rings" section on page 2-10. This section provides an overview of some of the important concepts about the ring limits for PA-A3 and PA-A6 ATM port adapters:

- There is a single receive ring and a single transmit ring corresponding to any physical interface on the Cisco 7200 series router.
- For the ATM port adapters, each entry in the receive ring or transmit ring shares a one-to-one correspondence with a particle stored in the private interface pool for a packet.
- The ring limit creates a logical division of a ring's total resource into per-VC queues. In other words, when an ATM port adapter supports traffic from multiple VCs, the ring limit places a boundary on the number of private interface particles that can be used by that VC.
- The implementation of ring limits affects the number of particles available in the private interface pool to store data for a given VC. You can interpret the ring limit as a threshold for particle usage by a VC.
- You specify the receive ring limit as a percentage of the private interface pool, and the transmit ring limit as a number of ring entries. Each of these limits translates to a number of particles in the private interface pool. It is important to recognize that although there are two different limits whose functions are based on whether you are in the receive stage or transmit stage of processing, both the receive ring limit and the transmit ring limit affect the same resource—particles in the private interface pool of the ingress interface.

Configuring the Receive Ring Limit

This section provides guidelines for configuring and verifying the receive ring limit. It includes the following topics:

Receive Ring Limit Configuration Guidelines

Use the following guidelines when planning to tune the receive ring limit:

Caution It is important to consider that the ring limits for the receive and transmit side are effectively operating against the same resource—particles within the private interface pool. Therefore, you must be very careful if you plan to tune these limits. Just as with adjustments to the buffer pools, improper settings for the receive ring or transmit ring limits can adversely impact system performance. In this case, adjustments to either side of the ring limits can impact the performance of both receiving and transmitting packets. Only modify the ring limits after careful evaluation of network impact or when recommended by technical support personnel.

- Consider tuning the per-VC receive ring in the case where a higher throughput ingress ATM interface might be feeding several slower egress serial interfaces.

  Private interface particles are not freed until the packet contents are transferred to the outbound port adapter or they are transmitted. This means that slow egress interfaces can tie up particle resource in the private interface pool, making the number of particles to receive data limited or unavailable for the receive ring. To help minimize this condition, you can increase the value of the **rx-limit** command.
- An indication that a VC is exceeding its receive ring limit is an increase in the number of ignored errors on the ATM interface, or input drops on the VC. For more information, see the "Monitoring Ring Limits and Resource Usage" section on page 7-12.
- Configure the **rx-limit** command as a percentage of the total particles in the private interface pool. Not all VCs need to add up to 100 percent. For example, it is possible to allow every VC to use up to 50 percent of the available resource. In this case, no more than two VCs would be able to use the full 50 percent at any given time or resource would not be available.
- You can configure the receive ring limit for a VC to be 100 percent, or the entire particle pool.

Default Values for the Receive Ring Limit

The default value for the receive ring limit is calculated internally based on the traffic shaping configuration on the VC using the following logic:

If PCR > 200

A = Min (PCR/200, Rx-Threshold)

Else

A = 0

Rx-Limit = Max (A, (MTU/particle_size) x 2)

The following further describes some of the variables used in this logic:

- The ATM port adapter sets the Rx-Threshold as the maximum Rx-limit that any VC on a given ATM interface can have. The Rx-Threshold is determined as 2/3 of the total number of available particles in the private interface pool for that ATM interface. For example, if the private interface buffer is 1200 particles, then the Rx-Threshold is 2/3 x 1200, which is 900. You can view the Rx-Threshold value using the **show controllers atm** command.

Note The Rx-Threshold is used only for computing the default receive ring limit. You can manually configure the **rx-limit** command on a VC to be 100 percent of the particle pool.

- The default maximum transmission unit (MTU) on the PA-A3 and PA-A6 ATM port adapters is 4470 bytes.

  The MTU defines the largest size of packets that an interface can transmit without needing to fragment. IP packets larger than the MTU must go through IP fragmentation procedures. Most Cisco ATM router interfaces use a default MTU size of 4470 bytes. This number was chosen to match Fiber Distributed Data Interface (FDDI) and High-Speed Serial Interface (HSSI) interfaces for autonomous switching. Cisco ATM router interfaces support an MTU between 64 and 17966 bytes. You can use the **mtu** interface configuration command to modify the default value.
- The particle_size value is 512 bytes, which is the size of the private interface pool particles.

Receive Ring Limit Configuration Example

Caution The **rx-limit** ATM VC configuration command is an internal command. Therefore, you must run the **service internal** global configuration command to access the **rx-limit** ATM VC configuration command. Any commands revealed by the use of the **service internal** command are unsupported.

The following example specifies that PVC 1/100 can use 25% of the private interface particle pool when receiving traffic:

```
Router# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Router(config)# service internal
Router(config)# int atm 1/0.1
Router(config-subif)# pvc 1/100
Router(config-if-atm-vc)# rx-limit 25
Router(config-if-atm-vc)# end
```

Verifying the Receive Ring Limit and Particle Buffers

To verify the number of particles in the PA-A3 or PA-A6 private interface pool and how much of the pool is available to a particular PVC for receive processing, complete the following steps:

Step 1 To view the total number of particles in the PA-A3 or PA-A6 private interface pool, which are allocated when data is received on that interface, use the **show controllers atm** command and observe the "rx buffers" field. In this example, there are 1200 total particles of size 512 bytes available in the ATM private interface pool:

```
Router# show controllers atm 5/0
Interface ATM5/0 is up
Hardware is ENHANCED ATM PA - OC3 (155000Kbps)
Framer is PMC PM5346 S/UNI-155-LITE, SAR is LSI ATMIZER II
Firmware rev: G127, Framer rev: 0, ATMIZER II rev: 3
  idb=0x62948598, ds=0x6294FEA0, vc=0x6297F940
  slot 5, unit 2, subunit 0, fci_type 0x0056, ticks 120012
  1200 rx buffers: size=512, encap=64, trailer=28, magic=4
[text omitted]
```

Step 2 To monitor how much of the private interface pool is currently allocated, use the **show buffers** command. The following example shows that all 1200 particles are in use (shown as "1200 hits") and no more particles are available (shown as "0 in free list"):

```
Router# show buffers
[text omitted]
Private particle pools:
Serial4/0 buffers, 512 bytes (total 192, permanent 192):
     0 in free list (0 min, 192 max allowed)
     192 hits, 0 fallbacks
     192 max cache size, 128 in cache
     10 buffer threshold, 0 threshold transitions
Serial4/1 buffers, 512 bytes (total 192, permanent 192):
     0 in free list (0 min, 192 max allowed)
     192 hits, 0 fallbacks
     192 max cache size, 128 in cache
     10 buffer threshold, 0 threshold transitions
Serial4/2 buffers, 512 bytes (total 192, permanent 192):
     0 in free list (0 min, 192 max allowed)
```

```
     192 hits, 0 fallbacks
     192 max cache size, 128 in cache
     10 buffer threshold, 0 threshold transitions
Serial4/3 buffers, 512 bytes (total 192, permanent 192):
     0 in free list (0 min, 192 max allowed)
     192 hits, 0 fallbacks
     192 max cache size, 128 in cache
     10 buffer threshold, 0 threshold transitions
ATM5/0 buffers, 512 bytes (total 1200, permanent 1200):
    0 in free list (0 min, 1200 max allowed)
    1200 hits, 1 misses
```

Step 3 To verify the percentage of the particle pool that is configured for possible allocation by the receive ring for a particular PVC, use the **show atm pvc** command. The following example shows that the receive ring limit is 25 percent for PVC 1/100:

```
Router# show atm pvc 1/100
 ATM1/0.1: VCD: 14, VPI: 1, VCI: 100
 UBR, PeakRate: 149760
 AAL5-LLC/SNAP, etype:0x0, Flags: 0xC20, VCmode: 0x0
 OAM frequency: 0 second(s), OAM retry frequency: 1 second(s),
 OAM retry frequency: 1 second(s)
 OAM up retry count: 3, OAM down retry count: 5
 OAM Loopback status: OAM Disabled
 OAM VC state: Not Managed
 ILMI VC state: Not Managed
 Rx Limit: 25 percent
 InARP frequency: 15 minutes(s)
 Transmit priority 4
 InPkts: 0, OutPkts: 0, InBytes: 0, OutBytes: 0
 InPRoc: 0, OutPRoc: 0
 InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
 InPktDrops: 0, OutPktDrops: 0
 CrcErrors: 0, SarTimeOuts: 0, OverSizedSDUs: 0,
 LengthViolation: 0, CPIErrors: 0
 Out CLP=1 Pkts: 0
 OAM cells received: 0
 F5 InEndloop: 0, F5 InSegloop: 0, F5 InAIS: 0, F5 InRDI: 0
 F4 InEndloop: 0, F4 InSegloop: 0, F4 InAIS: 0, F4 InRDI: 0
 OAM cells sent: 0
 F5 OutEndloop: 0, F5 OutSegloop: 0, F5 OutRDI: 0
 F4 OutEndloop: 0, F4 OutSegloop: 0, F4 OutRDI: 0
 OAM cell drops: 0
 Status: UP
```

Note In later releases of the Cisco IOS software, the **show atm pvc** command displays the receive ring limit in the "VC Rx Limit" field as a number of private interface particles. The receive limit output was modified from percentages to particles in some of the following Cisco IOS software releases: 12.2(4)T, 12.2(9)S, 12.2(4)B, 12.2(3), 12.1(9)E, and 12.1(10).

The following example displays **show** output using Cisco IOS Release 12.2(10):

```
Router# show atm pvc 1/101
ATM6/0: VCD: 2, VPI: 1, VCI: 101
UBR, PeakRate: 149760
AAL5-LLC/SNAP, etype:0x0, Flags: 0xC20, VCmode: 0x0
OAM frequency: 0 second(s), OAM retry frequency: 1 second(s), OAM retry
frequency: 1 second(s)
OAM up retry count: 3, OAM down retry count: 5
OAM Loopback status: OAM Disabled
OAM VC state: Not Managed
```

```
ILMI VC state: Not Managed
VC TxRingLimit: 40 particles
VC Rx Limit: 800 particles
[text omitted]
```

Note On the PA-A3 and PA-A6 ATM port adapters for the Cisco 7200 series router, the **show** output always displays the value of the receive ring limit whether it is set by default or it is manually configured. However, on the Cisco 7500 series router, you can only show the value of the receive ring limit when you have manually configured the **rx-limit** command.

Configuring the Transmit Ring Limit

This section provides guidleines for configuring and verifying the transmit ring limit. It includes the following topics:

Transmit Ring Limit Configuration Guidelines

It is important for you to understand that QoS policies for ATM traffic are applied in the Layer 3 queues prior to the receipt of packets on the transmit ring. These QoS policies can affect which packets arrive to the transmit queue soonest based on the service policies for the class. However, the transmit ring itself uses a FIFO-based queue. Therefore, you do not want to optimize flow at the Layer 3 level, only to introduce high latencies at the transmit ring.

The primary consideration for tuning the transmit ring is finding the appropriate size of the transmit ring that is small enough to avoid latency in the ring's FIFO queue, but large enough to avoid drops that have a significant impact on TCP-based flows.

In general, low values are recommended for voice VCs and higher values are recommended for data VCs. Low values reduce jitter and delay due to queueing, and high values accommodate bursts.

Consider the following general guidelines when planning to tune the transmit ring limit:

- Configure the **tx-ring-limit** command as a number of ring entries, which equates to a number of particles in the private interface pool.
- Configure higher values for high-speed VCs.
- Lower the size of the transmit ring when you want to achieve packet differentiation using Layer 3 service policies and to avoid latency on the hardware queue.
- Configure a low value for low-bandwidth VCs, such as VCs with an SCR of 128 Kbps.

Note Packets are queued to the transmit ring as soon as there is a free particle, even if the packet requires more than one particle to be stored.

Transmit Ring Guidelines for Voice VCs

For VCs carrying voice traffic, reduce the size of the transmit ring limit. Select a value based on the amount of serialization delay, expressed in seconds, that is introduced by the transmit ring. To determine the amount of delay, you can use the following formula:

[(P x 8) x D] / S

where

P = Packet size in bytes (Multiply by eight to convert to bits.)

D = Transmit ring depth

S = Speed of the VC in bps

Transmit Ring Guidelines for Data VCs

For VCs carrying data, use the following guidelines:

- Consider the packet size and configure the **tx-ring-limit** command to accommodate 4 packets. Be aware that 64 bytes are always reserved in the first particle for header rewrites. Therefore, as an example, a 1500-byte packet requires 4 particles (of size 512 bytes). Multiplying 4 particles x 4 packets yields 16 for the the **tx-ring-limit** value.
- Be sure that the transmit ring limit is large enough to support one MTU-sized packet or the number of cells equal to the maximum burst size (MBS) for a nrt-VBR PVC.
- Use Table 7-1 for suggested transmit ring limits by link speed, when other specific guidelines do not apply:

Table 7-1 Suggested Transmit Ring Limit Values by Link Speed

Link Speed	Transmit Ring Limit
Less than or equal to 128 Kbps	5
192 Kbps	6
256 Kbps	7
512 Kbps	14
768 Kbps	21

- Tune the size of the queue when you think that the VC is experiencing unnecessary delay. On any network interface, queueing forces a choice between latency and the amount of burst that the interface can sustain. Larger queue sizes sustain longer bursts, but increase delay.

Default Values for the Transmit Ring Limit

The PA-A3 and PA-A6 ATM port adapters assign a default transmit ring limit for every VC. The way that this default value is determined varies by the service category that is configured for the VC. Table 7-2 shows how the default values for the transmit ring are implemented.

Table 7-2 Default Values for the Transmit Ring Limit by ATM Service Category

Service Category	Default Value	Time of Enforcement
ABR	128	Always
CBR	Calculated using the following formula: (48 x PCR) / (particle_size x 5) The minimum default value is 40. • The PCR includes ATM overhead, and is translated to cells per second for this formula. • The particle_size is 580 bytes, reflecting the particle size of the transmit hardware buffer on the ATM port adapter.	Always
UBR	• PA-A3 and PA-A6 (T3/E3/OC-3)—40 • PA-A3 IMA (T1/E1)—128	When the total credit utilization exceeds 75 percent of the tx_threshold value shown in the **show controllers atm** command output[1]
nrt-VBR	Calculated using the following formula: (48 x SCR) / (particle_size x 5) The minimum default value is 40. • The SCR includes ATM overhead, and is translated to cells per second for this formula. • The particle_size is 580 bytes, reflecting the particle size of the transmit hardware buffer on the ATM port adapter.	Always
rt-VBR	Same as nrt-VBR calculation.	Always

1. The tx_threshold value is used as an upper boundary for the PA-A3 or PA-A6 ATM port adapter to use during processing of UBR VCs. The PA-A3 and PA-A6 ATM port adapters allow for larger bursts on UBR VCs by enforcing the transmit limit on such VCs only when the total packet buffer usage on the PA-A3 or PA-A6 reaches 75 percent of this preset threshold.

Verifying the Default Transmit Ring Limit

To verify the default number of particles that can be used by a particular PVC for transmit processing during setup of a VC, you can use the **debug atm events** command. The following steps show the default transmit ring limit value assigned to a nrt-VBR PVC as it is configured:

Step 1 From global configuration mode, enable the **debug atm events** command:

```
Router(config)# debug atm events
```

Step 2 Configure a nrt-VBR PVC on an ATM interface and enable logging to display the debug messages on the console as shown in the following example. The default transmit limit is shown in the "vc tx_limit" field as 137:

```
Router(config)# interface atm 4/0
Router(config-if)# pvc 1/100
Router(config-if-atm-vc)# vbr-nrt 4000 3500 94
Router(config-if-atm-vc)#
*Oct 14 17:56:06.886:  Reserved bw for 1/100 Available bw = 141500
Router(config-if-atm-vc)# exit
Router(config-if)# logging
*Oct 14 17:56:16.370: atmdx_setup_vc(ATM4/0): vc:6 vpi:1 vci:100 state:2 config_status:0
*Oct 14 17:56:16.370: atmdx_setup_cos(ATM4/0): vc:6 wred_name:- max_q:0
*Oct 14 17:56:16.370: atmdx_pas_vc_setup(ATM4/0): vcd 6, atm hdr 0x00100640, mtu 4482
*Oct 14 17:56:16.370: VBR: pcr 9433, scr 8254, mbs 94
*Oct 14 17:56:16.370:  vc tx_limit=137, rx_limit=47
*Oct 14 17:56:16.374:  Created 64-bit VC count
```

Step 3 Based on an SCR of 3500 Kbps, the PA-A3 assigns a default tx_limit of 137. To see how this calculation is made, convert the SCR of 3500 Kbps to cells/sec.

- First change the SCR to a number of bytes/sec, (3,500,000 bits/sec) / 8 bits/byte = 437,500 bytes/sec.
- Then, divide by 53 bytes to determine the number of cells per second. (437,500 bytes/sec) / 53 bytes = 8254 cells/sec.

Step 4 Now, you can apply the formula for the default transmit ring limit of (48 x SCR) / (particle_size x 5). (48 x 8254) / (580 x 5) = 136.6, rounded to 137.

Transmit Ring Configuration Example

The following example specifies that PVC 1/121 can use 10 particles for data awaiting transmission:

Note The available range within the CLI for the transmit ring limit reflects the number of particles available in the transmit hardware buffer located on the ATM port adapter. This makes sense because the PVC might support multiple ingress interfaces. Therefore, the credit check on the number of private interface particles in use must not exceed the number of particles available in the hardware buffer on the ATM port adapter.

```
Router# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Router(config)# interface atm 5/0.2
Router(config-if)# pvc 1/121
Router(config-if-atm-vc)# tx-ring-limit ?
  <3-6000> Number (ring limit)
```

```
Router(config-if-atm-vc)# tx-ring-limit 10
Router(config-if-atm-vc)# end
```

Verifying the Transmit Ring Limit

Note On the PA-A3 and PA-A6 ATM port adapters for the Cisco 7200 series router, the **show** output always displays the value of the transmit ring limit whether it is set by default or it is manually configured. However, on the Cisco 7500 series router, you can only show the value of the transmit ring limit when you have manually configured the **tx-ring-limit** command.

To verify the number of particles that can be used by a particular PVC for transmit processing, you can use the **show atm vc** or the **show atm pvc** commands. The following **show atm vc** command example shows that the transmit ring limit is 10 particles for VC 1/121:

```
Router# show atm vc 5
ATM3/0.21: VCD: 5, VPI: 1, VCI: 121
UBR, PeakRate: 149760
AAL5-LLC/SNAP, etype:0x0, Flags: 0xC20, VCmode: 0x0
OAM frequency: 0 second(s)
VC TxRingLimit: 10 particles
VC Rx Limit: 120 particles
InARP frequency: 15 minutes(s)
Transmit priority 6
InPkts: 5, OutPkts: 379, InBytes: 540, OutBytes: 27380
InCells: 0, OutCells: 761
InPRoc: 5, OutPRoc: 379
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
InPktDrops: 0, OutPktDrops: 0/0/0 (holdq/outputq/total)
InCellDrops: 0, OutCellDrops: 0
InByteDrops: 0, OutByteDrops: 0
CrcErrors: 0, SarTimeOuts: 0, OverSizedSDUs: 0, LengthViolation: 0, CPIErrors: 0
Out CLP=1 Pkts: 0, Cells: 0
OAM cells received: 0
OAM cells sent: 0
Status: UP
```

The following **show atm pvc** command example for another PVC shows that a transmit ring limit of 40 particles:

```
Router# show atm pvc 1/101
ATM6/0: VCD: 2, VPI: 1, VCI: 101
UBR, PeakRate: 149760
AAL5-LLC/SNAP, etype:0x0, Flags: 0xC20, VCmode: 0x0
OAM frequency: 0 second(s), OAM retry frequency: 1 second(s), OAM retry
frequency: 1 second(s)
OAM up retry count: 3, OAM down retry count: 5
OAM Loopback status: OAM Disabled
OAM VC state: Not Managed
ILMI VC state: Not Managed
VC TxRingLimit: 40 particles
VC Rx Limit: 800 particles
[text omitted]
```

Monitoring Ring Limits and Resource Usage

It is important to understand the architecture involved in the processing of ATM traffic on the Cisco 7200 series router. Drops can occur due to a limitation of buffer resources either on board the ATM port adapter, or in the private interface pool on the NPE or NSE. Therefore, you should understand how to properly monitor the performance of your system's resources so that you can properly interpret where performance tuning might be appropriate.

This section describes some of the **show** commands that you can use to monitor the status of your input and output buffer resources located on the ATM port adapter hardware, and in the private interface pools on the the NPE or NSE. It includes the following topics:

Monitoring Hardware Buffers

The PA-A3 and PA-A6 ATM port adapters use memory on board the port adapter to store data for SAR processing. These ATM port adapters have a separate buffer for receive path processing, and another buffer for transmit path processing. The particle sizes vary in these two hardware buffers. For more information about the PA-A3 and PA-A6 ATM port adapter architecture, see the "Receive Buffer and Transmit Buffer Located on the PA-A3 and PA-A6 ATM Port Adapters" section on page 2-14.

The topics in this section describe how to determine whether you are experiencing a shortage of memory in the hardware buffers on the PA-A3 and PA-A6 ATM port adapters.

Monitoring the Status of Input Buffers Located on the PA-A3 and PA-A6 ATM Port Adapters

To monitor the status of the receive hardware buffer on the ATM port adapter, run the **show controllers atm** command and observe the "rx_no_buffer" counter. The following example indicates that there is not a shortage of input buffers because no packets have been dropped due to the VC reaching its transmit ring quota (shown by "rx_no_buffer=0"):

```
Router# show controllers atm 3/0
   Interface ATM3/0 is up
   Hardware is ENHANCED ATM PA - DS3 (45Mbps)
   Lane client mac address is 0030.7b1e.9054
   Framer is PMC PM7345 S/UNI-PDH, SAR is LSI ATMIZER II
   Firmware rev: G119, Framer rev: 1, ATMIZER II rev: 3
     idb=0x61499630, ds=0x6149E9C0, vc=0x614BE940
     slot 3, unit 2, subunit 0, fci_type 0x005B, ticks 73495
     400 rx buffers: size=512, encap=64, trailer=28, magic=4
   Curr Stats:
     rx_cell_lost=0, rx_no_buffer=0, rx_crc_10=0
     rx_cell_len=0, rx_no_vcd=0, rx_cell_throttle=0, tx_aci_err=0
[text omitted]
```

Monitoring the Status of Output Buffers Located on the PA-A3 and PA-A6 ATM Port Adapter

To determine when a shortage of output hardware buffers is occurring, complete the following steps:

Step 1 Run the **show interface atm** command and observe the "no buffer" counter. When the PA-A3 or PA-A6 ATM port adapter runs out of hardware storage in the transmit buffer on board the port adapter, the no buffer counter increments in the **show interface atm** command, as shown in the following example:

```
Router# show interface atm 4/0
    ATM4/0 is up, line protocol is up
    Hardware is ENHANCED ATM PA
    MTU 4470 bytes, sub MTU 4470, BW 149760 Kbit, DLY 80 usec,
    reliability 255/255, txload 136/255, rxload 1/255
    Encapsulation ATM, loopback not set
    Encapsulation(s): AAL5
    4095 maximum active VCs, 5 current VCCs
    VC idle disconnect time: 300 seconds
    Signalling vc = 4, vpi = 0, vci = 5
    UNI Version = 3.0, Link Side = user
    4 carrier transitions
    Last input 00:02:30, output 00:00:00, output hang never
    Last clearing of "show interface" counters never
    Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 103197668
    Queueing strategy: Per VC Queueing
    30 second input rate 0 bits/sec, 0 packets/sec
    30 second output rate 80210000 bits/sec, 6650 packets/sec
    308 packets input, 9856 bytes, 4138 no buffer
    Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    338179038 packets output, 3163620726 bytes, 0 underruns
    0 output errors, 0 collisions, 0 interface resets
    0 output buffer failures, 0 output buffers swapped out
```

Step 2 To further verify that the transmit hardware buffer is full, use the **show controllers atm** command and observe the value of the BFD cache status area of the output. The BFD cache "size" field indicates the total number of buffers in the local port adapter memory. The current number of free particles is shown by the "read" field. The PA-A3 reserves 144 particles for system packets like Operation, Administration, and Maintenance (OAM) cells. When the "read" value reaches 144, the PA-A3 driver starts dropping packets until a sufficient number of local memory particles becomes available. The following example shows that the free local memory particles for the hardware transmit buffer is 143, and therefore there is no available local hardware storage for transmit processing:

```
Router# show controllers atm 5/0
[text omitted]
   BFD Cache status:
     base=0x62931AA0, size=6144, read=143
   Rx Cache status:
[text omitted]
```

Monitoring Ring Limits for the Private Interface Pool

The topics in this section describe how to determine whether you are experiencing a shortage of memory in the private interface pool for receive or transmit processing on a PA-A3 or PA-A6 ATM port adapter.

Determining a Shortage of Private Interface Particles for Receive Processing

To determine if packets are being dropped due to a shortage of private interface particles for receive processing, use the **show interface atm** command and observe the "ignored" field. The following example shows that particles are still available for receive processing on ATM interface 4/0 because no packets have been dropped (shown by the "0 ignored" field):

```
Router# show interface atm 4/0
    ATM4/0 is up, line protocol is up
    Hardware is ENHANCED ATM PA
    MTU 4470 bytes, sub MTU 4470, BW 149760 Kbit, DLY 80 usec,
    reliability 255/255, txload 136/255, rxload 1/255
    Encapsulation ATM, loopback not set
    Encapsulation(s): AAL5
    4095 maximum active VCs, 5 current VCCs
    VC idle disconnect time: 300 seconds
    Signalling vc = 4, vpi = 0, vci = 5
    UNI Version = 3.0, Link Side = user
    4 carrier transitions
    Last input 00:02:30, output 00:00:00, output hang never
    Last clearing of "show interface" counters never
    Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 103197668
    Queueing strategy: Per VC Queueing
    30 second input rate 0 bits/sec, 0 packets/sec
    30 second output rate 80210000 bits/sec, 6650 packets/sec
    308 packets input, 9856 bytes, 4138 no buffer
    Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    338179038 packets output, 3163620726 bytes, 0 underruns
    0 output errors, 0 collisions, 0 interface resets
    0 output buffer failures, 0 output buffers swapped out
```

For more information about verifying the private interface particles, see the "Verifying the Receive Ring Limit and Particle Buffers" section on page 7-5.

Determining When the Transmit Ring Limit is Reached on a VC

To determine when a VC has reached its transmit ring limit quota, run the **show atm vc** command and observe the "OutPktDrops" counter. The following example indicates that the PVC 2/2 still has not reached its transmit quota because no packets have been dropped (shown by "OutPktDrops: 0"):

```
Router# show atm vc
VC 3 doesn't exist on interface ATM3/0
ATM5/0.2: VCD: 3, VPI: 2, VCI: 2
VBR-NRT, PeakRate: 30000, Average Rate: 20000, Burst Cells: 94
AAL5-LLC/SNAP, etype:0x0, Flags: 0x20, VCmode: 0x0
OAM frequency: 0 second(s)
PA TxRingLimit: 10
InARP frequency: 15 minutes(s)
Transmit priority 2
InPkts: 0, OutPkts: 0, InBytes: 0, OutBytes: 0
InPRoc: 0, OutPRoc: 0
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
InPktDrops: 0, OutPktDrops: 0
CrcErrors: 0, SarTimeOuts: 0, OverSizedSDUs: 0
OAM cells received: 0
OAM cells sent: 0
Status: UP
```

Note The **show controllers atm** command provides some transmit counter fields, such as "max_tx_count," "tx_count," and "tx_threshold." However, these output fields display the amount of transmit credits for the entire interface, not on an individual VC. For example, the "max_tx_count" field shows the maximum number of transmit particles held by the PA-A3 or PA-A6 ATM microcode. The "tx_count" field shows the total number of transmit particles currently being held by the port adapter microcode for all VCs. The "tx_threshold" field shows the preset limit used by the PA-A3 and PA-A6 ATM port adapters for enforcement of transmit limits on UBR PVCs.

Related Documentation

The following table provides information about additional resources that you can read to learn more about some of the topics discussed in this chapter:

For more information about :	Refer to the following publications:
MTUs on ATM port adapters	Understanding Maximum Transmission Unit (MTU) on ATM Interfaces (TAC Tech Note)
Troubleshooting input and output errors and interpreting **show** commands	• Troubleshooting Input and Output Errors on PA-A3 ATM Port Adapters (TAC Tech Note) • When Does the no buffer Error Counter Increment on the PA-A3?(TAC Tech Note) • Troubleshooting "Ignored" Errors on an ATM Port Adapter (TAC Tech Note)

Next Steps

This chapter describes how to optimize and verify the receive ring and transmit ring limits on the PA-A3 and PA-A6 ATM port adapters. If you want to know more about the memory architecture and the flow of ATM traffic processing using the receive rings, transmit rings, and the hardware buffers located on the PA-A3 and PA-A6 ATM port adapters, read Chapter 2, "Cisco 7200 Series Architecture and Design for ATM Traffic Management."

Chapter 8, "ATM Traffic Management Case Studies and Configuration Examples," provides case studies and configuration examples of some real-world implementations of ATM in enterprise networks, including tuning of transmit ring limits.

ATM Traffic Management Case Studies and Configuration Examples

This chapter provides case studies and configuration examples for enterprise networks using Cisco 7200 series routers in an ATM environment. All of the examples in this chapter represent actual lab-tested configurations from Cisco Systems proof-of-concept and solutions labs. These Cisco Systems labs validate and test proposed customer network designs according to specified objectives, and make recommendations about configuration. However, it is important to realize that every network has unique characteristics and requirements. Therefore, while the information contained in this chapter might be helpful for your network, there are no guarantees that the sample configurations apply to or will benefit your specific environment.

Due to their basis on real-world network configurations, the network case studies represent large and more complex network scenarios. The network configuration examples include other Cisco Systems products as well as many configuration elements that are beyond the scope of this document. Therefore, there is no intent to fully explain the network implementations shown. Rather, the case studies are provided so that you can see how ATM traffic management is applied on a Cisco 7200 series router as it fits into the wider scope of a real network environment to solve common customer objectives.

This chapter includes the following case studies and configuration examples:

High Density Aggregation Network Case Study

This case study provides information about designing large-scale, high-speed aggregation networks using the Cisco 7200 series routers in an ATM network.

The customer represented by this case study is a private Internet Service Provider (ISP) that offers Internet and intranet connectivity for the education and government sectors. Throughout this case, the customer will be referred to as Case SP.

This case includes the following sections:

Network Description

Case SP's network is primarily a T1-based service to subscribers who aggregate to a single Cisco 7200 router (with an NPE-200 and 128 MB of memory) at each point of presence (POP) in a hub-and-spoke fashion. The network consists of 10 POPs across the state. These 10 aggregation POPs connect as a hub-and-spoke back to the Case SP central site using a private Optical Carrier (OC)-3c ATM backbone comprised of BPX 8620 ATM switches. The central site houses two Cisco 12000 routers connected to the ATM backbone with OC-12 interfaces. Each Cisco 7200 POP router has an ATM permanent virtual circuit (PVC) back to each core Gigabit Switch Router (GSR). Each GSR has an OC-3 Synchronous Optical Network (SONET) connection to the Internet and routes all internal Case SP remote-POP-to-remote-POP data traffic. Full Internet route tables are carried to each POP's Cisco 7200 router using the Internal Border Gateway Protocol (IBGP) (100K routes).

Figure 8-1 High Density WAN Aggregation Network Configuration

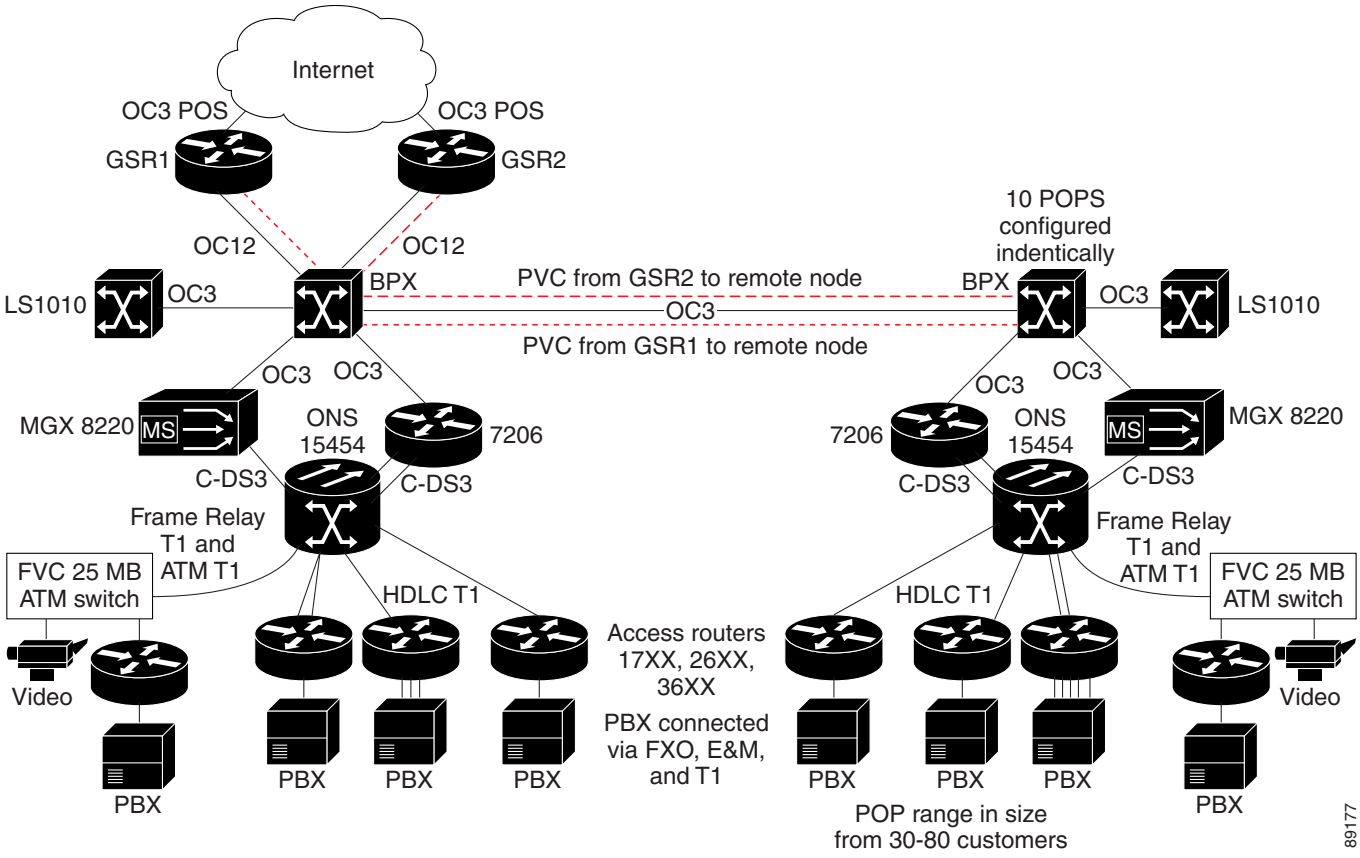

Initial Cisco 7206 Router Configuration

The following is a sample configuration representing the configuration for initial testing using the Cisco 7206 router. Portions of the configuration have been omitted where redundant configuration statements exist.

```
7200-A3# show run
Building configuration...

Current configuration : 48619 bytes
```

```
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service internal
!
hostname 7200-A3
!
boot system disk0:c7200-p-mz.121-7.E.bin
!
ip subnet-zero
!
no ip finger
no ip domain-lookup
!
controller T3 5/0
 framing m23
 t1 1 channel-group 0 timeslots 1-24
 t1 2 channel-group 0 timeslots 1-24
 t1 3 channel-group 0 timeslots 1-24
 t1 4 channel-group 0 timeslots 1-24
 t1 5 channel-group 0 timeslots 1-24
 t1 6 channel-group 0 timeslots 1-24
 t1 7 channel-group 0 timeslots 1-24
 t1 8 channel-group 0 timeslots 1-24
 t1 9 channel-group 0 timeslots 1-24
 t1 10 channel-group 0 timeslots 1-24
 t1 11 channel-group 0 timeslots 1-24
 t1 12 channel-group 0 timeslots 1-24
 t1 13 channel-group 0 timeslots 1-24
 t1 14 channel-group 0 timeslots 1-24
 t1 15 channel-group 0 timeslots 1-24
 t1 16 channel-group 0 timeslots 1-24
 t1 17 channel-group 0 timeslots 1-24
 t1 18 channel-group 0 timeslots 1-24
 t1 19 channel-group 0 timeslots 1-24
 t1 20 channel-group 0 timeslots 1-24
 t1 21 channel-group 0 timeslots 1-24
 t1 22 channel-group 0 timeslots 1-24
 t1 23 channel-group 0 timeslots 1-24
 t1 24 channel-group 0 timeslots 1-24
 t1 25 channel-group 0 timeslots 1-24
 t1 26 channel-group 0 timeslots 1-24
 t1 27 channel-group 0 timeslots 1-24
 t1 28 channel-group 0 timeslots 1-24
!
! Note: Two Channelized T3 controllers are configured similarly
! with 28 channel-groups, for a total of 112 Serial interfaces.
!
controller T3 5/1
 framing m23
 t1 1 channel-group 0 timeslots 1-24
!
[text omitted]
!
interface Loopback0
 ip address 10.15.1.1 255.255.255.255
!
interface FastEthernet1/0
 ip address 10.30.1.20 255.255.255.0
 no ip mroute-cache
 duplex full
!
```

```
interface ATM4/0
 no ip address
 load-interval 30
 no atm sonet ilmi-keepalive
 no atm ilmi-keepalive
!
interface ATM4/0.1 point-to-point
 ip address 10.10.10.2 255.255.255.0
 pvc 40/40
  vbr-nrt 155000 155000 1000
  encapsulation aal5snap
!
! Only a partial listing of the configured
! Channelized T3 controllers is shown.
!
interface Serial5/0/1:0
 ip address 10.13.1.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial5/0/2:0
 ip address 10.13.2.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial5/0/3:0
 ip address 10.13.3.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial5/0/4:0
 ip address 10.13.4.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
[text omitted]
!
interface Serial5/0/27:0
 ip address 10.13.27.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial5/0/28:0
 ip address 10.13.28.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
```

```
!
interface Serial5/1/1:0
 ip address 10.14.1.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial5/1/2:0
 ip address 10.14.2.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
[text omitted]
!
interface Serial5/1/27:0
 ip address 10.14.27.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial5/1/28:0
 ip address 10.14.28.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
router eigrp 100
 network 10.7.3.0 0.0.0.255
 network 10.13.0.0
 network 10.14.0.0
 network 10.15.0.0
 network 10.16.0.0
 network 10.15.1.1 0.0.0.0
 network 10.21.0.0
 network 10.200.1.0
 no auto-summary
 no eigrp log-neighbor-changes
!
router bgp 100
 no synchronization
 bgp log-neighbor-changes
 neighbor 10.26.1.1 remote-as 100
 neighbor 10.26.1.1 update-source Loopback0
!
! Note: A static route exists for each destination subnet on a Serial interface.
! Only the first four and last four are shown here as an example.
!
ip classless
ip route 10.7.3.0 255.255.255.0 Serial5/0/4:0
ip route 10.7.1.0 255.255.255.0 Serial5/0/2:0
ip route 10.22.1.0 255.255.255.0 ATM4/0.1
ip route 10.40.1.0 255.255.255.0 Serial5/0/1:0
!
[text omitted]
!
```

```
ip route 10.40.111.0 255.255.255.0 Serial6/1/27:0
ip route 10.40.112.0 255.255.255.0 Serial6/1/28:0
ip route 10.50.0.0 255.0.0.0 Serial5/0/1:0
ip route 10.70.0.0 255.255.0.0 10.10.10.1
no ip http server
!
access-list 100 permit udp any any range 1718 1720
access-list 100 permit udp any range 1718 1720 any
access-list 100 permit tcp any any range 1718 1720
access-list 100 permit tcp any range 1718 1720 any
access-list 100 permit udp any range 54000 56000 any
access-list 100 permit udp any any range 54000 56000
access-list 101 permit ip 10.50.5.0 0.0.0.15 host 10.6.6.6
access-list 101 permit ip 10.60.6.0 0.0.0.15 host 10.7.7.7
access-list 101 permit ip 10.60.6.0 0.0.0.15 host 10.7.7.8
access-list 101 permit ip 10.60.6.0 0.0.0.15 host 10.7.7.9
access-list 101 permit ip 10.60.6.0 0.0.0.15 host 10.7.7.0
access-list 101 permit ip host 10.1.1.1 host 10.7.7.0
access-list 101 permit ip host 10.1.1.1 host 10.7.7.6
access-list 101 permit ip host 10.1.1.1 host 10.7.6.6
access-list 101 permit ip host 10.1.1.1 host 10.3.6.6
access-list 101 deny   ip host 10.0.0.0 host 10.0.0.0
access-list 101 deny   ip 10.2.8.0 0.0.0.15 10.0.1.0 10.15.63.255
access-list 101 permit ip any any
route-map local permit 10
 match ip address 100
 set ip precedence flash-override
!
!
!
line con 0
 transport input none
line aux 0
line vty 0 4
 password [text omitted]
 login
line vty 5 15
 login
!
end
```

Recommended Cisco 7206VXR Router Configuration

The following is a sample of the recommended configuration for the Cisco 7206VXR router during this testing. Portions of the configuration have been omitted where redundant configuration statements exist.

```
7200-A3-VXR-RED# show run
Building configuration...

Current configuration : 48619 bytes
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service internal
!
hostname 7200-A3-VXR-RED
!
boot system disk0:c7200-p-mz.121-7.E.bin
boot system disk0:c7200-is-mz.121_7_E.dan.bin
!
```

```
ip subnet-zero
ip cef
!
!
no ip finger
no ip domain-lookup
!
!
class-map match-all call-control
  match access-group 100
class-map match-all voice
  match ip precedence 5
class-map match-all default
  match any
!
!
policy-map llq-atm
  class voice
    priority 596
  class call-control
   bandwidth 10
policy-map llq-serial-250
  class voice
    priority 250
  class call-control
   bandwidth 10
  class default
   bandwidth 892
   random-detect
   random-detect precedence 0   2     4     5
   random-detect precedence 1   2     4     5
   random-detect precedence 2   2     4     5
   random-detect precedence 3   2     4     5
policy-map llq-serial-548
  class voice
    priority 548
  class call-control
   bandwidth 10
  class default
   bandwidth 500
   random-detect
   random-detect precedence 0   2     4     5
   random-detect precedence 1   2     4     5
   random-detect precedence 2   2     4     5
   random-detect precedence 3   2     4     5
policy-map llq-serial-48
  class voice
    priority 48
  class call-control
   bandwidth 10
  class default
   bandwidth 1000
   random-detect
   random-detect precedence 0   2     4     5
   random-detect precedence 1   2     4     5
   random-detect precedence 2   2     4     5
   random-detect precedence 3   2     4     5
!
!
controller T3 5/0
 framing m23
 t1 1 channel-group 0 timeslots 1-24
 t1 2 channel-group 0 timeslots 1-24
 t1 3 channel-group 0 timeslots 1-24
```

```
 t1 4 channel-group 0 timeslots 1-24
 t1 5 channel-group 0 timeslots 1-24
 t1 6 channel-group 0 timeslots 1-24
 t1 7 channel-group 0 timeslots 1-24
 t1 8 channel-group 0 timeslots 1-24
 t1 9 channel-group 0 timeslots 1-24
 t1 10 channel-group 0 timeslots 1-24
 t1 11 channel-group 0 timeslots 1-24
 t1 12 channel-group 0 timeslots 1-24
 t1 13 channel-group 0 timeslots 1-24
 t1 14 channel-group 0 timeslots 1-24
 t1 15 channel-group 0 timeslots 1-24
 t1 16 channel-group 0 timeslots 1-24
 t1 17 channel-group 0 timeslots 1-24
 t1 18 channel-group 0 timeslots 1-24
 t1 19 channel-group 0 timeslots 1-24
 t1 20 channel-group 0 timeslots 1-24
 t1 21 channel-group 0 timeslots 1-24
 t1 22 channel-group 0 timeslots 1-24
 t1 23 channel-group 0 timeslots 1-24
 t1 24 channel-group 0 timeslots 1-24
 t1 25 channel-group 0 timeslots 1-24
 t1 26 channel-group 0 timeslots 1-24
 t1 27 channel-group 0 timeslots 1-24
 t1 28 channel-group 0 timeslots 1-24
!
! Note: Each of four Channelized T3 controllers is configured similarly
! with 28 channel-groups, for a total of 112 Serial interfaces.
!
controller T3 5/1
 framing m23
 t1 1 channel-group 0 timeslots 1-24
!
[text omitted]
!
controller T3 6/0
 framing m23
 t1 1 channel-group 0 timeslots 1-24
!
[text omitted]
!
controller T3 6/1
 framing m23
 t1 1 channel-group 0 timeslots 1-24
!
[text omitted]
!
interface Loopback0
 ip address 10.15.1.1 255.255.255.255
!
interface FastEthernet1/0
 ip address 10.30.1.20 255.255.255.0
 no ip mroute-cache
 duplex full
!
interface ATM4/0
 no ip address
 load-interval 30
 no atm sonet ilmi-keepalive
 no atm ilmi-keepalive
!
interface ATM4/0.1 point-to-point
 ip address 10.10.10.2 255.255.255.0
 pvc 40/40
```

```
  vbr-nrt 155000 155000 1000
  encapsulation aal5snap
  service-policy output llq-atm
 !
! Only a partial listing of the configured
! Channelized T3 controllers is shown.
!
interface Serial5/0/1:0
 ip address 10.13.1.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 service-policy output llq-serial-250
!
interface Serial5/0/2:0
 ip address 10.13.2.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 service-policy output llq-serial-250
!
interface Serial5/0/3:0
 ip address 10.13.3.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 service-policy output llq-serial-48
!
interface Serial5/0/4:0
 ip address 10.13.4.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 service-policy output llq-serial-548
!
[text omitted]
!
interface Serial5/0/27:0
 ip address 10.13.27.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial5/0/28:0
 ip address 10.13.28.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial5/1/1:0
 ip address 10.14.1.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial5/1/2:0
 ip address 10.14.2.1 255.255.255.0
 load-interval 30
```

```
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
.
.
.
interface Serial5/1/27:0
 ip address 10.14.27.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial5/1/28:0
 ip address 10.14.28.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial6/0/1:0
 ip address 10.15.1.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial6/0/2:0
 ip address 10.15.2.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
[text omitted]
!
interface Serial6/0/27:0
 ip address 10.15.27.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial6/0/28:0
 ip address 10.15.28.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial6/1/1:0
 ip address 10.16.1.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
```

```
!
interface Serial6/1/2:0
 ip address 10.16.2.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
[text omitted]
!
interface Serial6/1/27:0
 ip address 10.16.27.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
interface Serial6/1/28:0
 ip address 10.16.28.1 255.255.255.0
 load-interval 30
 no keepalive
 tx-ring-limit 12
 fair-queue 32 64 28
 hold-queue 30 out
!
router eigrp 100
 network 10.7.3.0 0.0.0.255
 network 10.13.0.0
 network 10.14.0.0
 network 10.15.0.0
 network 10.16.0.0
 network 10.15.1.1 0.0.0.0
 network 10.21.0.0
 network 10.200.1.0
 no auto-summary
 no eigrp log-neighbor-changes
!
router bgp 100
 no synchronization
 bgp log-neighbor-changes
 neighbor 10.26.1.1 remote-as 100
 neighbor 10.26.1.1 update-source Loopback0
!
! Note: A static route exists for each destination subnet on a Serial interface.
! Only the first four and last four are shown here as an example.
!
ip classless
ip route 10.7.3.0 255.255.255.0 Serial5/0/4:0
ip route 10.7.1.0 255.255.255.0 Serial5/0/2:0
ip route 10.22.1.0 255.255.255.0 ATM4/0.1
ip route 10.40.1.0 255.255.255.0 Serial5/0/1:0
!
[text omitted]
!
ip route 10.40.111.0 255.255.255.0 Serial6/1/27:0
ip route 10.40.112.0 255.255.255.0 Serial6/1/28:0
ip route 10.50.0.0 255.0.0.0 Serial5/0/1:0
ip route 10.70.0.0 255.255.0.0 10.10.10.1
no ip http server
!
access-list 100 permit udp any any range 1718 1720
access-list 100 permit udp any range 1718 1720 any
```

```
access-list 100 permit tcp any any range 1718 1720
access-list 100 permit tcp any range 1718 1720 any
access-list 100 permit udp any range 54000 56000 any
access-list 100 permit udp any any range 54000 56000
access-list 101 permit ip 10.50.5.0 0.0.0.15 host 10.6.6.6
access-list 101 permit ip 10.60.6.0 0.0.0.15 host 10.7.7.7
access-list 101 permit ip 10.60.6.0 0.0.0.15 host 10.7.7.8
access-list 101 permit ip 10.60.6.0 0.0.0.15 host 10.7.7.9
access-list 101 permit ip 10.60.6.0 0.0.0.15 host 10.7.7.0
access-list 101 permit ip host 10.1.1.1 host 10.7.7.0
access-list 101 permit ip host 10.1.1.1 host 10.7.7.6
access-list 101 permit ip host 10.1.1.1 host 10.7.6.6
access-list 101 permit ip host 10.1.1.1 host 10.3.6.6
access-list 101 deny   ip host 10.0.0.0 host 10.0.0.0
access-list 101 deny   ip 10.2.8.0 0.0.0.15 10.0.1.0 10.15.63.255
access-list 101 permit ip any any
route-map local permit 10
 match ip address 100
 set ip precedence flash-override
!
!
!
line con 0
 transport input none
line aux 0
line vty 0 4
 password [text omitted]
 login
line vty 5 15
 login
!
end
```

Network History and Problem Statement

The network was originally designed to carry data traffic. Since that time, it has grown significantly in terms of both the number of remote T1s at each aggregation POP, as well as the amount of data being transferred to and from each remote location. Many of the T1 or multiple T1 customers are schools with hundreds to thousands of users generating thousands of flows. The number of flows at a given site that are requesting data from the Internet has resulted in some subscriber links being constantly oversubscribed. The outbound T1 oversubscription has triggered ingress buffer starvation for the high-speed interface (which is the ATM Port Adapter PA-A3 OC-3) on the Cisco 7200 routers. This problem is seen in the form of ignores, flushes, and cyclic redundancy checks (CRCs) on the ATM interface. This change in network conditions has resulted in poor performance for customers even on lightly loaded links.

The constantly oversubscribed customers were aware of and accepted their performance issues, but the other subscribers were unhappy with the poor performance of their more lightly loaded links. Due to budget pressure, these oversubscribed accounts would not buy enough bandwidth to solve their oversubscription problem. Also, due to competitive pressure from other service providers, Case SP could not force them to do so.

In addition to facing performance problems on the data side, Case SP was also offering a trial for a Voice over Internet Protocol (VoIP) service to some customers. They wanted to be able to roll it out to more customers, plus add video over IP. Case SP's current video offering was in the form of H.320 and they were evaluating adding H.323 video on the network. It appeared that the platform was already becoming undersized for the applications it was supporting at that time, so it was highly unlikely that it would be

able to scale to the next level. Thus, the challenge was not only to find a solution to the existing performance issues but also to enable the network to support multiple differentiated services, such as VoIP and IP-based video conferencing.

Case Objectives

The following test objectives were established for this case study:

- Validate lab test criteria by reproducing ignored errors seen in the live environment during oversubscription of egress links.
- Identify the performance thresholds of the current platform in order to determine which POPs are affected and why.
- Test various methods for controlling the oversubscription problem.
- Identify a short-term solution to the existing platform that will stabilize the network, support the current rate of growth, and allow for added services such as voice and video conferencing.
- Test the performance of multiservice applications under congested network conditions in order to validate quality of service (QoS).
- Test a longer-term, next-generation solution to determine if a different platform is needed to provide significant scalability options.

Overview of the Testing

The Devices Under Test (DUT) included the Cisco 7200 router with an NPE-200, Cisco 7200VXR with an NPE-400, and the Cisco 7600 router.

After replicating the current customer environment, the Cisco lab team team was able to successfully reproduce the problem. In the test environment, the Cisco 7200 routers had a high-speed interface that consistently received an oversubscription of data destined for only a few of the lower-speed T1 interfaces. The ATM interface was far from being oversubscribed; it was the egress T1s that were driven beyond port speeds. This caused buffer starvation to occur, thus impacting even the non-oversubscribed links on the same router.

Each interface on the Cisco 7200 series router has a private interface pool in which all received packets over that interface are stored. The packet contents remain in the pool of the ingress interface until they are transmitted. While the packets are awaiting transmission out the low-speed (T1) interfaces (which were oversubscribed), the buffers cannot be used for incoming traffic. This caused the ATM interface to experience a lot of ignores in the oversubscribed environment. Ignores indicate that the data made it to the router, but was dropped prior to any processing or queueing. This occurred because there were not enough buffers to receive data from the ATM interface and process it through the router to the egress interface. Thus, with even just a few oversubscribed T1s, all connections were affected, as ingress data was dropped indiscriminately.

Testing Methodology

Two types of tests were run against the different router platforms being tested:

- The first test on the router provided a full load of data for each egress T1 in order to provide a maximum load reference. Traffic was generated in both directions for all T1s (1.536 Mbps per T1, TCP frames at 1024 bytes).

- The second test used a customer-supplied breakdown of the typical protocol and packet size distribution seen in the network, as well as a common interface load distribution in which only a percentage of the links were oversubscribed. This information was used to provide a customer mix load result. Table 8-1 shows the percentage of different types of traffic typically found on the network and the typical frame sizes and priorities.

Table 8-1 Case SP Traffic Mix

Traffic Type	Frame Size (bytes)	% Total Network Bandwidth	Priority	Type of Service (TOS)
Hypertext Transfer Protocol (HTTP)	256	57	4	2
Transmission Control Protocol (TCP)	1000	29	3	3
User Datagram Protocol (UDP)	80	5	1	5
File Transfer Protocol (FTP)	1500	9	2	4

Both tests were subjected to an injection of 120,000 BGP routes with 5,000 periodic flaps and 500 Enhanced Interior Gateway Protocol (EIGRP) routes with 60 periodic flaps. This route table load was used to simulate the customer's current environment with 20 percent growth. Table 8-2 shows the load that was used to test the T1s.

Table 8-2 Case SP Traffic Load

T1 Traffic Loads	Load per T1 (Mbps)	Example Load Using 56 T1s
25%	1.536	14 T1s will have a load of 1.53 Mbps
40%	1.200	22 T1s will have a load of 1.2 Mbps
25%	0.500	12 T1s will have a load of 500 Kbps
15%	0.100	8 T1s will have a load of 100 Kbps

Testing of Existing Cisco 7200 Series Router

The first set of tests were conducted using the existing Cisco 7200 series router with an NPE-200 and 128 MB DRAM.

The first set of tests were conducted to reproduce Case SP's oversubscription problem, identify the performance threshholds of the current platform, and test different ways of controlling the oversubscription problem using the current platform.

With a test bed set up for 56 T1s, four different approaches were tested in an attempt to overcome the issue of buffer starvation. Of the four optimizations attempted, configuring an output rate limit on the serial interfaces worked.

However, the following three approaches were also tried:

- Modifying the TX ring limit on the serial interface
- Implementing Weighted Fair Queueing (WFQ) on the serial interface

- Implementing Weighted Random Early Detect (WRED)

These three approaches did not fully address the problem, and are described further in the "Results of Existing 7200 Platform Testing Phase" section.

Results of Existing 7200 Platform Testing Phase

This section describes the configurations attempted and the results of the four approaches.

Modifying the TX Ring Limit

In this approach, there was an attempt to limit the number of buffers an interface can hold using the **tx-ring-limit** command. Although this approach yielded modest improvement, it did not drastically improve the discards.

The following is an example of the configuration:

```
interface Serial5/0/3:0
ip address 10.13.3.1 255.255.255.0
tx-ring-limit 12
```

Implementing WFQ

In this approach, queueing limits were configured on the serial interface. Configuring fair queueing with buffer limits helped reduce ATM ignores under light congestion, but it did not correct the problem under severe congestion.

The following is an example of the configuration:

```
interface Serial5/0/1:0
ip address 10.13.1.1 255.255.255.0
fair-queue 32 64 28
hold-queue 30 out
```

Implementing WRED

WRED requires enabling of Cisco Express Forwarding (CEF). IP CEF could not be used in this scenario due to the size of the route tables, as well as limited CPU and memory resources available on NPE-200.

Configuring an Output Rate Limit

In this approach, output rate limits were configured on the serial interfaces. Rate limiting resolved the issue of buffer starvation at the ingress interface. Packets were dropped at the offending output interface and non-oversubscribed ports performed normally. This was the result that Case SP desired. Be aware, however, that the use of rate limiting increases Central Processing Unit (CPU) utilization (by around 20 percent for the 14 T1s tested in this scenario). This approach also has the potential to cause the Transmission Control Protocol (TCP) window to shut, thereby reducing application throughput for those links.

The following is an example of the configuration:

```
interface Serial5/0/3:0
ip address 10.13.3.1 255.255.255.0
rate-limit output 1536000 1500 2000 conform-action
transmit exceed-action drop
```

Note The Committed Access Rate (CAR) feature is considered a legacy form of policing and is no longer recommended for use on the Cisco 7200 series routers. Newer, class-based policing mechanisms are now available using the modular QoS CLI (MQC) configuration method.

Conclusions of Existing 7200 Platform Testing Phase

With the current Cisco 7200 platform and using the full T1 load setup, a 100 percent load was determined to be 28 T1s. Using the customer traffic mix, the maximum load was 56 T1s. The NPE-200 could not intelligently discard traffic under heavy load conditions due to the lack of congestion avoidance techniques like WRED (which could not be turned on because it requires CEF, and CEF could not be enabled due to the high CPU usage for other services placed on the router).

Testing of Upgraded Cisco 7200 VXR Platform with NPE-400

In the second phase of testing, the router platform was upgraded to the NPE-400, with a faster processor. The goal of this testing was to identify a short-term solution that would stabilize the network, allow reuse of current router port adapters, and support the addition of VoIP and video services. The same testing methodology was used, but this time IP CEF was enabled on the router. In addition, rate limiting was not configured on the T1 serial interfaces to control discards, but WRED was configured.

Results of Upgraded 7200 Platform Testing Phase

Without using rate limiting, upgrading to the next processor improved control and scalability. The maximum full traffic load was 56 T1s. However, the problem recurred with ignores being received on the ATM PA-A3. To address this, IP CEF was enabled and WRED was used as a mechanism to control discards on the serial interface. Class-based queueing was also used to guarantee varying quality of service levels for different applications.

Implementing WRED

The Cisco WRED protocol combines IP Precedence and RED and provides differentiated drop thresholds for premium (high priority) traffic versus standard (lower priority) traffic. This allows Case SP to drop packets from the standard customer before dropping packets (if at all) from the premium customer or traffic type.

The following is an example of the WRED policy map configuration for the serial interface:

```
class-map match-all call-control
 match access-group 100
class-map match-all voice
 match ip precedence 5
class-map match-all default
 match any
!
policy-map llq-atm
 class voice
  priority 596
 class call-control
  bandwidth 10
!
policy-map llq-serial-250
 class voice
  priority 250
 class call-control
  bandwidth 10
 class default
  bandwidth 892
  random-detect
  random-detect precedence 0 2 4 5
  random-detect precedence 1 2 4 5
  random-detect precedence 2 2 4 5
  random-detect precedence 3 2 4 5
!
```

```
policy-map llq-serial-48
 class voice
  priority 48
 class call-control
  bandwidth 10
 class default
  bandwidth 1000
  random-detect
  random-detect precedence 0 2 4 5
  random-detect precedence 1 2 4 5
  random-detect precedence 2 2 4 5
  random-detect precedence 3 2 4 5
```

Configuring WRED with the above parameters allowed the oversubscribed interfaces to intelligently discard data without affecting the non-oversubscribed users. WRED did not have any detrimental effect on CPU utilization (in fact, there was only an increase of about 3 percent).

The configuration implements the **random-detect precedence** command using the following syntax:

random-detect precedence *precedence min-threshold max-threshold mark-prob-denominator*

The packet drop probability is based on the minimum threshold, maximum threshold, and mark probability denominator. When the average queue depth is above the minimum threshold of packets, RED starts dropping packets. The rate of packet drops increases linearly as the average queue size increases, until the average queue size reaches the maximum threshold.

The mark probability denominator is the fraction of packets dropped when the average queue depth is at the maximum threshold. As configured above, if the denominator is 5, one out of every five packets is dropped when the average queue is at the maximum threshold. When the average queue size is above the maximum threshold, all packets are dropped.

WRED Restrictions

At the time of this case study, WRED was characterized by the following restrictions:

- When WRED affects protocols other than TCP/IP, the packet sources might resend dropped packets at the same rate and therefore congestion is not decreased. WRED works best when most of the traffic is TCP/IP traffic. With TCP, dropped packets indicate congestion, so the packet source reduces its transmission rate.
- WRED treats non-IP traffic as precedence 0, the lowest precedence. Therefore, non-IP traffic is dropped more often than IP traffic.
- WRED does not support ATM encapsulations using AAL5-MUX.
- WRED requires CEF switching.

Conclusions of Upgraded 7200 Platform Testing Phase

With the upgraded platform and all the necessary tools available on the Cisco 7200 VXR, WRED was found to be the best way to control discards on a percentage of T1 serial links. The increase in CPU utilization from baseline records proved to be minimal (3 percent increase).

Additional tests were conducted with 56 T1s and 104 T1s and WRED rate limiting.

In the 56 T1 tests, 14 T1s were each overdriven at 2.0 Mbps/T1 while the remaining 42 T1s were loaded to their specific traffic rates. When the WRED rate limit was set at 1.475 Mbps per T1, no ATM ignores or errors occurred.

When generating traffic to 104 T1s (based on the customer's traffic load mix, with none over line rate), the processor was at an acceptable utilization percentage. No errors occurred (no ATM ignores, CRCs, flushes, or input errors), serial interfaces were left at default configurations, and fair queueing (with no prioritization or rate limiting) was configured.

When overdriving 20 T1s to 2.0 Mpbs/T1, with the other 84 T1s running at their normal rates with the same protocol mix and no rate control, ATM ignores did occur. CPU utilization increased by approximately 8 percent. When overdriving 20 T1s to 2.0 Mpbs/T1, with the other 84 T1s running at their normal rates with the same protocol mix but with WRED as the rate control mechanism, ATM ignores did not occur. This proved the capability of the upgraded, short-term solution for this customer scenario.

Multiservice Tests on the Cisco 7200 VXR Platform with NPE-400

Once the platform was configured to control the oversubscription issues, multiservice capabilities were tested. With 70 T1s being oversubscribed at 1.7 Mbps, WRED alleviated the buffer starvation problem that was previously seen as ignores on the ingress ATM OC-3 port adapter. WRED was applied to the default class of traffic (anything with a precedence of less than 3), and the buffers were tuned to reduce the queue size on the serial ports in order to prevent buffer starvation.

Buffer starvation was alleviated by applying WRED on the default policy, as shown in the following configuration example:

```
policy-map llq-serial-250
 class voice
  priority 250
 class call-control
  bandwidth 10
 class default
  bandwidth 892
  random-detect
  random-detect precedence 0 2 4 5
  random-detect precedence 1 2 4 5
  random-detect precedence 2 2 4 5
  random-detect precedence 3 2 4 5
```

Note When creating policies, be sure to consider that match criteria is applied on the first match of classes. Therefore, the default class (if it is the least important traffic classification) should be the last class that you configure.

By applying Low Latency Queueing (LLQ) for voice and Class-Based Weighted Fair Queueing (CBWFQ) for voice signaling, voice quality was consistently good, even in heavily congested links.

Use of H.323 gatekeeper allowed for centralized management of the dialing plan. No issues with call-setup delay were experienced. Channel Associated Signaling (CAS) was used between the private branch exchange (PBX) and the Cisco 2600 and Cisco 3600 routers. CAS typically can add call-setup delays of up to 3 to 5 seconds.

Testing of Cisco 7600 FlexWAN as Long-Term Solution

The Cisco 7600 configuration included a Supervisor II, with Multilayer Switch Feature Card 2 (MSFC2), 48-port 10/100 Ethernet card, and two FlexWAN modules containing two Multichannel Port Adapter (PA-MC)-T3+ and a PA-A3-OC3. The 48-port RJ-45 card was used to connect a router for injecting EIGRP routes and the gatekeeper for the local zone. The connectivity to the Cisco 7600 was achieved using an OC-3 link. The MSFC2 used Cisco IOS Release 12.1(7)E.

A total of 11 tests were performed with 120,000 BGP routes (5,000 routes flapping every 90 seconds) and 500 EIGRP routes (60 routes flapping every 90 seconds).

In this testing, the 7600 with FlexWAN was the DUT, and the same topology and traffic loads from the previous testing were used. Queueing parameters were set to match the previously tested configuration to provide quality of service.

The customer's typical traffic mix was used with all 70 T1s being oversubscribed at 1.7 Mbps. Border Gateway Protocol (BGP) and EIGRP routes were injected with 5,000 BGP flaps and 60 EIGRP flaps.

Because of SONET and ATM and ATM Adaption Layer Type-5 (AAL5) overhead, 123 Mbps is the maximum bandwidth that a single OC-3 link can sustain. Traffic was generated in both directions. Due to these overheads, the maximum number of T1s that could be tested was 80.

Multiservice Voice Test

By applying LLQ for voice and CBWFQ for voice signaling, voice quality was good even in heavily congested links. Digital PBX handsets and analog phone sets were tested by the customer for voice quality, delay in call setup, latency, echo, and distortion. A VoIP tester was used to quantitatively test voice. The VoIP tester uses prerecorded voice samples specified in the International Telecommunication Union (ITU) P861 specification. The tester analyzes the resultant voice waveforms and produces a Perceptual Speech Quality Measurement (PSQM) score for the resultant speech waveform. The VoIP tester was used to test PSQM speech quality, latency, and jitter. All tests produced very good voice quality.

QoS Testing for ATM in Airport Case Study

This case study provides information about QoS strategies to achieve priority for higher precedence data during times of network congestion using the Cisco 7200 series routers in an ATM network servicing voice, video, and data traffic. The customer represented by this case study is an airline. Throughout this case, the customer will be referred to as Case Airline.

This case includes the following sections:

Network Description

The test network was designed to closely emulate the two major classes of airports in the Case Airline network with both airports connected through a Cisco WAN switch carrier cloud back to the simulated City HQ core site as shown in Figure 8-2 on page 8-20. The data flow was to and from LANE LAN segments in City HQ and the remote airport sites.

The Cisco 7206 routers were running Cisco IOS release 12.1(3.4) using the PA-A3 OC-3 ATM port adapters.

Figure 8-2 Physical Network Diagram of Test Network Representing Case Airline

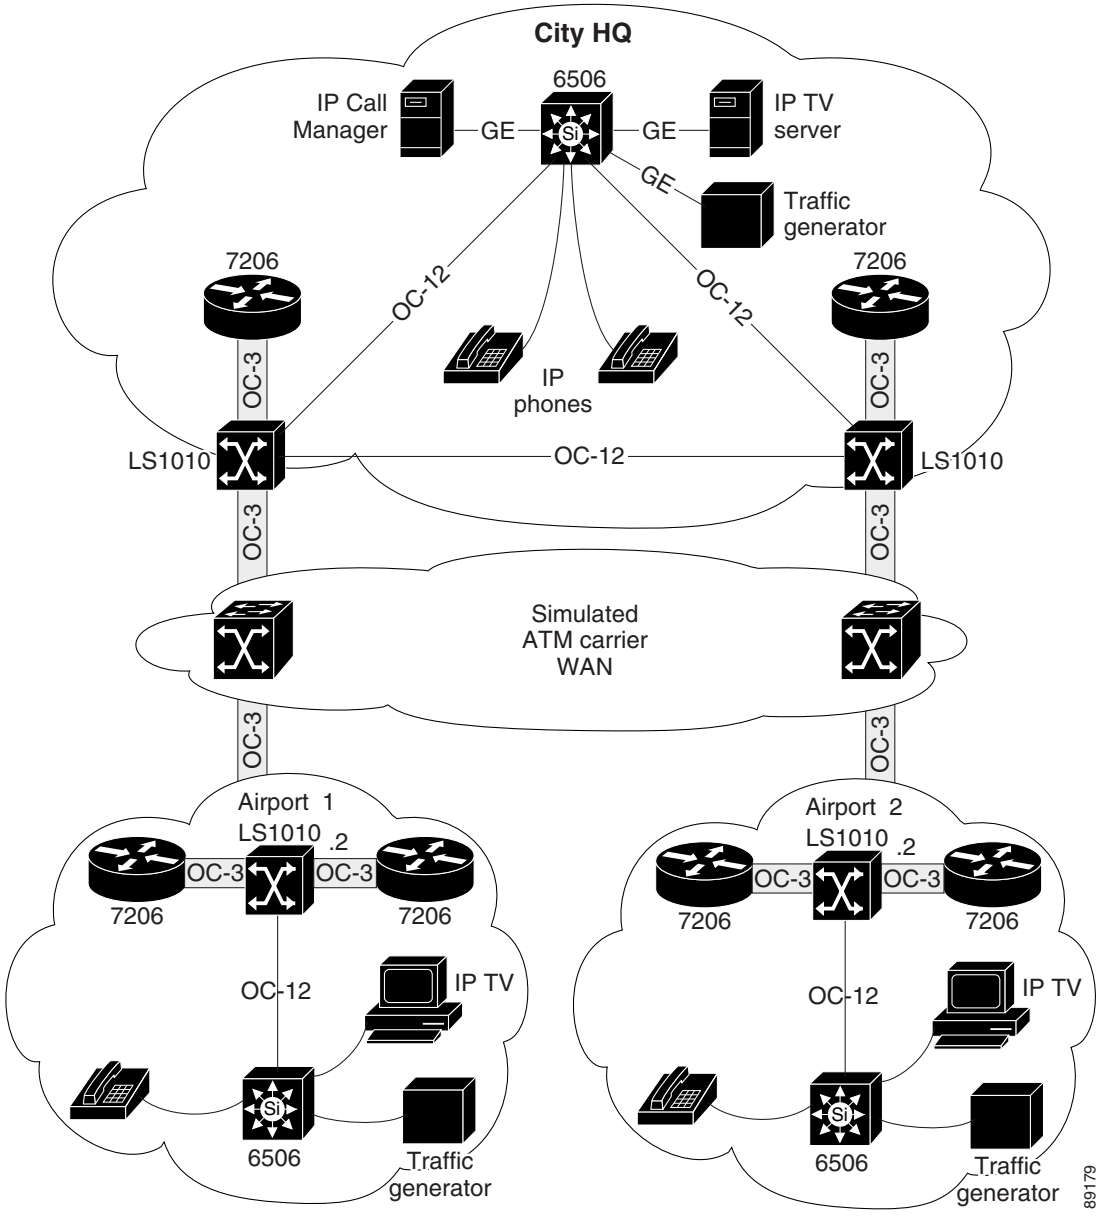

Figure 8-3 shows the logical view of the test network used to represent the Case Airline network.

Figure 8-3 Logical Network Diagram of Test Network Representing Case Airline

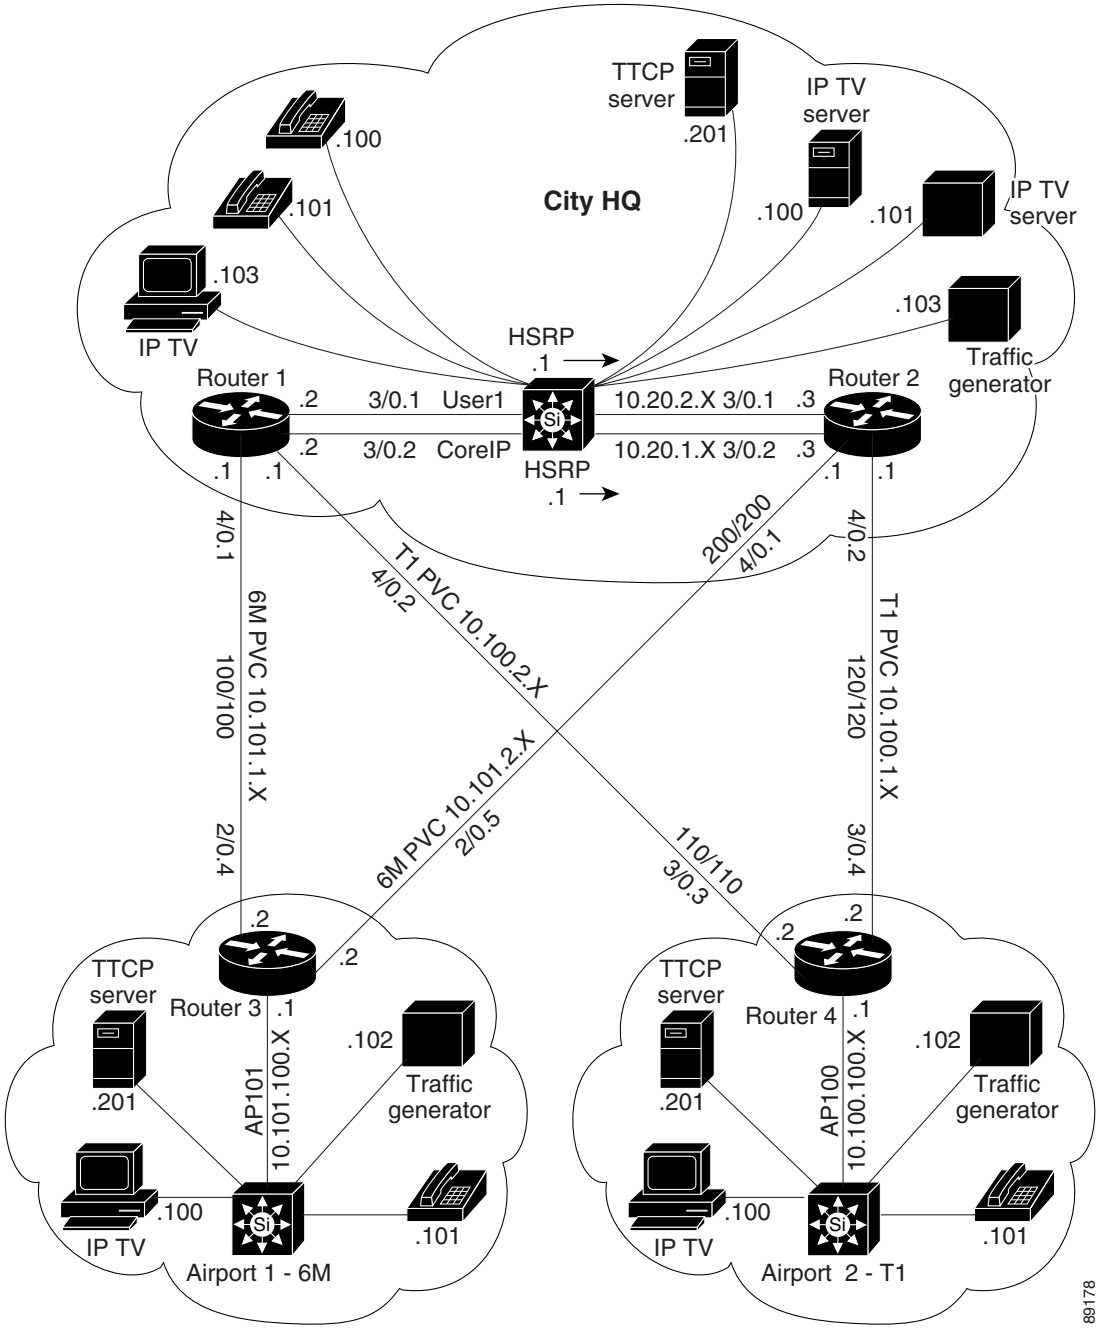

Baseline Cisco 7206 Router Configurations Using FIFO

The following configurations on Cisco 7206 routers represent those used during initial testing of voice, video, and data traffic under congestion with only FIFO queueing in place.

Router 1 Configuration

```
Current configuration:
!
version 12.1
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
service password-encryption
service udp-small-servers
service tcp-small-servers
!
hostname router1
!
!
!
ip subnet-zero
no ip domain-lookup
ip host router4 10.100.100.1
ip host router1 10.20.1.2
ip host router3 10.101.100.1
ip host router2 10.20.1.3
!
ip multicast-routing
ip cef
cns event-service server
!
!
!
!
!
vc-class atm 6Mbs
  vbr-nrt 5950 5900 32
  oam-pvc manage 3
  oam retry 5 3 3
  encapsulation aal5snap
!
interface Loopback0
 ip address 10.20.254.2 255.255.255.0
 ip pim sparse-dense-mode
!
interface ATM3/0
 no ip address
 no ip mroute-cache
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
 no atm ilmi-keepalive
!
interface ATM3/0.1 multipoint
 description User ELAN simulating Case Airline user ELAN
 ip address 10.20.1.2 255.255.255.0
 no ip redirects
 ip pim sparse-dense-mode
 no ip mroute-cache
 lane client ethernet elan1
 standby 2 timers 1 5
 standby 2 priority 110 preempt delay 60
 standby 2 ip 10.20.1.1
!
```

```
interface ATM3/0.2 multipoint
 description Server ELAN simulating Case Airline COPRIP
 ip address 10.20.2.2 255.255.255.0
 ip helper-address 10.20.1.110
 ip helper-address 10.20.1.101
 no ip redirects
 ip pim sparse-dense-mode
 no ip mroute-cache
 lane client ethernet elan2
 standby 1 timers 1 5
 standby 1 priority 100 preempt delay 60
 standby 1 ip 10.20.2.1
!
interface ATM4/0
 no ip address
 no atm ilmi-keepalive
!
interface ATM4/0.1 point-to-point
 bandwidth 6000
 ip address 10.101.1.1 255.255.255.0
 ip pim sparse-dense-mode
 pvc 100/100
  class-vc 6Mbs
 !
!
interface ATM4/0.2 point-to-point
 bandwidth 1000
 ip address 10.100.2.1 255.255.255.0
 ip pim sparse-dense-mode
 pvc 110/110
 !
!
router ospf 1
 log-adjacency-changes
 area 100 stub no-summary
 area 100 range 10.100.0.0 255.255.0.0
 area 101 stub no-summary
 area 101 range 10.101.0.0 255.255.0.0
 network 10.20.0.0 0.0.255.255 area 0
 network 10.100.0.0 0.0.255.255 area 100
 network 10.101.0.0 0.0.255.255 area 101
!
ip default-gateway 172.26.64.1
ip classless
ip route 172.26.0.0 255.255.0.0 FastEthernet0/0
ip route 172.26.11.0 255.255.255.0 ATM3/0.1
ip http server
ip pim send-rp-announce Loopback0 scope 16
ip pim send-rp-discovery Loopback0 scope 5
!
!
[text omitted]
!
end
```

Router 2 Configuration

```
version 12.1
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
service password-encryption
service udp-small-servers
service tcp-small-servers
!
```

```
hostname router2
!
ip subnet-zero
no ip domain-lookup
ip host router1 10.20.1.2
ip host router4 10.100.100.1
ip host router3 10.101.100.1
ip host router2 10.20.1.3
!
ip multicast-routing
ip cef
cns event-service server
!
vc-class atm 6Mbs
  vbr-nrt 5950 5900 32
  oam-pvc manage 3
  oam retry 5 3 3
  encapsulation aal5snap

interface Loopback0
 ip address 10.20.254.100 255.255.255.255
 ip pim sparse-dense-mode
!
interface ATM3/0
 no ip address
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
 no atm ilmi-keepalive
!
interface ATM3/0.1 multipoint
 description Client ELAN simulating Case Airline User ELAN
 ip address 10.20.1.3 255.255.255.0
 no ip redirects
 ip pim sparse-dense-mode
 lane client ethernet elan1
 standby 2 timers 1 5
 standby 2 priority 150 preempt delay 60
 standby 2 ip 10.20.1.1
!
interface ATM3/0.2 multipoint
 description server ELAN simulating Case Airline CORIP
 ip address 10.20.2.3 255.255.255.0
 ip helper-address 10.20.1.101
 no ip redirects
 ip pim sparse-dense-mode
 lane client ethernet elan2
 standby 1 timers 1 5
 standby 1 priority 110 preempt delay 60
 standby 1 ip 10.20.2.1
!
interface ATM4/0
 no ip address
 load-interval 30
 no atm ilmi-keepalive
!
interface ATM4/0.1 point-to-point
 bandwidth 6000
 ip address 10.101.2.1 255.255.255.0
 ip pim sparse-dense-mode
 pvc 200/200
  class-vc 6Mbs
 !
!
interface ATM4/0.2 point-to-point
```

```
 bandwidth 1000
 ip address 10.100.1.1 255.255.255.0
 ip pim sparse-dense-mode
 pvc 120/120
 !
!
router ospf 1
 log-adjacency-changes
 area 100 stub no-summary
 area 100 range 10.100.0.0 255.255.0.0
 area 101 stub no-summary
 area 101 range 10.101.0.0 255.255.0.0
 redistribute static subnets
 network 10.20.0.0 0.0.255.255 area 0
 network 10.100.0.0 0.0.255.255 area 100
 network 10.101.0.0 0.0.255.255 area 101
!
ip default-gateway 172.26.64.1
ip classless
ip route 172.26.11.0 255.255.255.0 10.20.1.4
ip http server
ip pim send-rp-announce Loopback0 scope 16
ip pim send-rp-discovery Loopback0 scope 5
!
!
!
[text omitted]
!
end
```

Router 3 Configuration

```
version 12.1
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
service password-encryption
service udp-small-servers
service tcp-small-servers
!
hostname router3
!
!
clock calendar-valid
ip subnet-zero
no ip domain-lookup
ip host router4 10.100.100.1
ip host router1 10.20.1.2
ip host router3 10.101.100.1
ip host router2 10.20.1.3
!
ip multicast-routing
ip cef
cns event-service server
!
!
!
!
!
vc-class atm 6Mbs
  vbr-nrt 5950 5900 32
  oam-pvc manage 3
  oam retry 5 3 3
  encapsulation aal5snap
!
```

```
interface Loopback0
 description "Router2 - Loopback address and Router ID"
 ip address 10.101.254.1 255.255.255.255
 ip pim sparse-dense-mode
!
interface ATM2/0
 no ip address
 no ip mroute-cache
 load-interval 30
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
 no atm ilmi-keepalive
!
interface ATM2/0.1 multipoint
 ip address 10.101.100.1 255.255.255.0
 ip helper-address 10.20.1.101
 ip pim sparse-dense-mode
 lane client ethernet elan1
!
interface ATM2/0.4 point-to-point
 bandwidth 6000
 ip address 10.101.1.2 255.255.255.0
 ip pim sparse-dense-mode
 shutdown
 pvc 100/100
  class-vc 6Mbs
 !
!
interface ATM2/0.5 point-to-point
 bandwidth 6000
 ip address 10.101.2.2 255.255.255.0
 ip pim sparse-dense-mode
 pvc 200/200
  class-vc 6Mbs
 !
!
router ospf 1
 log-adjacency-changes
 area 101 stub
 network 10.101.0.0 0.0.255.255 area 101
!
ip classless
no ip forward-protocol udp netbios-ns
no ip forward-protocol udp netbios-dgm
ip http server
ip pim send-rp-discovery Loopback0 scope 5
!
!
[text omitted]
!
end
```

Router 4 Configuration

```
!
version 12.1
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
service password-encryption
service udp-small-servers
service tcp-small-servers
!
hostname router4
!
```

```
!
!
ip subnet-zero
no ip domain-lookup
ip host router4 10.100.100.1
ip host router1 10.20.1.2
ip host router3 10.101.100.1
ip host router2 10.20.1.3
!
ip multicast-routing
ip cef
cns event-service server
!
!
!
vc-class atm 6Mbs
  vbr-nrt 5950 5900 32
  oam-pvc manage 3
  oam retry 5 3 3
  encapsulation aal5snap
!
interface Loopback0
 ip address 10.100.254.1 255.255.255.255
 ip pim sparse-dense-mode
!
interface ATM3/0
 no ip address
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
 no atm ilmi-keepalive
!
interface ATM3/0.1 multipoint
 ip address 10.100.100.1 255.255.255.0
 ip helper-address 10.20.1.101
 ip pim sparse-dense-mode
 lane client ethernet elan1
!
interface ATM3/0.3 point-to-point
 bandwidth 1000
 ip address 10.100.2.2 255.255.255.0
 ip pim sparse-dense-mode
 pvc 110/110
  !
!
interface ATM3/0.4 point-to-point
 bandwidth 6000
 ip address 10.100.1.2 255.255.255.0
 ip pim sparse-dense-mode
 pvc 120/120
  class-vc 6Mbs
!
router ospf 1
 log-adjacency-changes
 area 100 stub
 network 10.100.0.0 0.0.255.255 area 100
!
ip default-gateway 172.26.64.1
ip classless
ip http server
ip pim send-rp-discovery Loopback0 scope 5
!
[text omitted]
!
end
```

Case Objectives

The goal of this case study is to demonstrate that you can mark designated classes of traffic supporting voice, video, and data with a discreet precedence and then guarantee performance for higher precedence data during times of network congestion.

Overview of the Testing

During the tests, the quality and throughput of voice, video, and data traffic was monitored. Using a traffic generator, traffic was gradually increased across the same link to a level of three-to-one bandwidth oversubscription.

Testing Methodology

The testing was done in a series of incremental steps to demonstrate the use and effectiveness of each component of the Cisco IOS QoS solution. In each of the tests, the performance of voice and TCP data traffic over video traffic flows was observed for that QoS configuration. Then, additional data was injected onto the network to oversubscribe the links and observe the effects of congestion on voice and data traffic.

A total of seven tests were performed, using the following QoS configurations:

- First In First Out (FIFO)
- Weighted Fair Queueing (WFQ) without IP Precedence
- Weighted Random Early Detect (WRED)
- WFQ with IP Precedence
- WRED with IP Precedence
- Class-Based WFQ (CBWFQ) with IP Precedence
- Hybrid Using CBWFQ and WRED

Eight variables were recorded during each test to demonstrate the effectiveness of each configuration, as shown in the summary of results in Table 8-3 on page 8-29:

- TCP data throughput
- Voice quality (rating of 1–4, with 4 being toll quality)
- Video quality (1–4)
- Drop percentage of class 0 traffic
- Drop percentage of class 2 traffic
- Drop percentage of class 5 traffic
- Data latency of class 5 traffic during congestion
- CPU utilization

Results Summary

The case objectives were met and the results exceeded expectations. Even during three-to-one bandwidth oversubscription, priority traffic was relatively unaffected while best effort traffic was discarded. CPU utilization remained well within acceptable limits.

Table 8-3 Test Results by Test Type

Test No.	Test Type	TCP (Kb/s)	Voice (1-4)	Video (1-4)	Class 0 Drop%	Class 2 Drop %	Class 5 Drop %	Class 5 Latency (ms)
1	FIFO	209	1	1	45	–	–	60
2	FQ	189	2	1	45	–	–	60
3	RED	331	2	1	45	–	–	60
4	WFQ	248	3	2	45	–	–	55
5	WRED	313	3	2	45	–	–	60
6	CBWFQ	243	3.5	3	80	55	0	40
7	CBWFQ + WRED	333	3.5	3	80	53	0	40

QoS Test Configuration Examples

The following sections provide examples of the Cisco 7206 router configurations that were tested.

FIFO Configuration Example

The FIFO configurations used during the initial test are provided in the "Baseline Cisco 7206 Router Configurations Using FIFO" section on page 8-22.

WFQ Configuration Example

The configuration example in this section was used in tests 2 and 4 of the case study. During test 2, IP precedence was not tested.

During test 4, the effectiveness of WFQ was tested using IP precedence. In that test, the traffic generator created TOS-based flows where the precedence was set to 2, and the Catalyst switches marked precedence on the voice and video data. To oversubscribe the rate, data streams were injected with type of service (TOS) 0 on top of the prioritized voice and video traffic.

The following provides an example of the tested router configuration:

```
version 12.1
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
service password-encryption
service udp-small-servers
service tcp-small-servers
!
hostname router2
!
!
policy-map FAIR-Q
  class class-default
   fair-queue
!
ip subnet-zero
no ip domain-lookup
ip host router1 10.20.1.2
ip host router4 10.100.100.1
ip host router3 10.101.100.1
ip host router2 10.20.1.3
!
ip multicast-routing
ip cef
```

```
cns event-service server
!
!
!
vc-class atm 6Mbs
  vbr-nrt 5950 5900 32
  oam-pvc manage 3
  oam retry 5 3 3
  encapsulation aal5snap
!
interface Loopback0
 ip address 10.20.254.100 255.255.255.255
 ip pim sparse-dense-mode
!
interface FastEthernet0/0
 ip address 172.26.67.12 255.255.240.0
 no ip proxy-arp
 no ip mroute-cache
 shutdown
 half-duplex
!
interface ATM3/0
 no ip address
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
 no atm ilmi-keepalive
!
interface ATM3/0.1 multipoint
 description Client ELAN simulating Case Airline User ELAN
 ip address 10.20.1.3 255.255.255.0
 no ip redirects
 ip pim sparse-dense-mode
 lane client ethernet elan1
 standby 2 timers 1 5
 standby 2 priority 150 preempt delay 60
 standby 2 ip 10.20.1.1
!
interface ATM3/0.2 multipoint
 description server ELAN simulating Case Airline CORIP
 ip address 10.20.2.3 255.255.255.0
 ip helper-address 10.20.1.101
 no ip redirects
 ip pim sparse-dense-mode
 lane client ethernet elan2
 standby 1 timers 1 5
 standby 1 priority 110 preempt delay 60
 standby 1 ip 10.20.2.1
!
interface ATM4/0
 no ip address
 load-interval 30
 no atm ilmi-keepalive
!
interface ATM4/0.1 point-to-point
 bandwidth 6000
 ip address 10.101.2.1 255.255.255.0
 ip pim sparse-dense-mode
 pvc 200/200
  class-vc 6Mbs
  service-policy output FAIR-Q
 !
!
router ospf 1
 log-adjacency-changes
```

```
 area 100 stub no-summary
 area 100 range 10.100.0.0 255.255.0.0
 area 101 stub no-summary
 area 101 range 10.101.0.0 255.255.0.0
 redistribute static subnets
 network 10.20.0.0 0.0.255.255 area 0
 network 10.100.0.0 0.0.255.255 area 100
 network 10.101.0.0 0.0.255.255 area 101
!
ip default-gateway 172.26.64.1
ip classless
ip route 172.26.11.0 255.255.255.0 10.20.1.4
ip http server
ip pim send-rp-announce Loopback0 scope 16
ip pim send-rp-discovery Loopback0 scope 5
!
[text omitted]
!
end
```

WRED Configuration Example

The following configuration example was used in tests 3 and 5 of the case study. During test 3, IP precedence was not tested.

During test 5, the effectiveness of WFQ was tested using IP precedence. In that test, the traffic generator created TOS-based flows including prioritized voice, video, and TCP data. To oversubscribe the rate, data streams were injected with TOS 0 on top of the prioritized voice and video traffic.

The following provides an example of the tested router configuration:

```
version 12.1
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
service password-encryption
service udp-small-servers
service tcp-small-servers
!
hostname router2
!
policy-map RED
  class class-default
   bandwidth 4200
   random-detect
!
ip subnet-zero
no ip domain-lookup
ip host router1 10.20.1.2
ip host router4 10.100.100.1
ip host router3 10.101.100.1
ip host router2 10.20.1.3
!
ip multicast-routing
ip cef
cns event-service server
!
!
!
!
!
vc-class atm 6Mbs
  vbr-nrt 5950 5900 32
  oam-pvc manage 3
```

```
  oam retry 5 3 3
  encapsulation aal5snap
!
interface Loopback0
 ip address 10.20.254.100 255.255.255.255
 ip pim sparse-dense-mode
!
interface ATM3/0
 no ip address
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
 no atm ilmi-keepalive
!
interface ATM3/0.1 multipoint
 description Client ELAN simulating Case Airline User ELAN
 ip address 10.20.1.3 255.255.255.0
 no ip redirects
 ip pim sparse-dense-mode
 lane client ethernet elan1
 standby 2 timers 1 5
 standby 2 priority 150 preempt delay 60
 standby 2 ip 10.20.1.1
!
interface ATM3/0.2 multipoint
 description server ELAN simulating Case Airline CORIP
 ip address 10.20.2.3 255.255.255.0
 ip helper-address 10.20.1.101
 no ip redirects
 ip pim sparse-dense-mode
 lane client ethernet elan2
 standby 1 timers 1 5
 standby 1 priority 110 preempt delay 60
 standby 1 ip 10.20.2.1
!
interface ATM4/0
 no ip address
 load-interval 30
 no atm ilmi-keepalive
!
interface ATM4/0.1 point-to-point
 bandwidth 6000
 ip address 10.101.2.1 255.255.255.0
 ip pim sparse-dense-mode
 pvc 200/200
  class-vc 6Mbs
  service-policy output RED
 !
router ospf 1
 log-adjacency-changes
 area 100 stub no-summary
 area 100 range 10.100.0.0 255.255.0.0
 area 101 stub no-summary
 area 101 range 10.101.0.0 255.255.0.0
 redistribute static subnets
 network 10.20.0.0 0.0.255.255 area 0
 network 10.100.0.0 0.0.255.255 area 100
 network 10.101.0.0 0.0.255.255 area 101
!
ip default-gateway 172.26.64.1
ip classless
ip route 172.26.11.0 255.255.255.0 10.20.1.4
ip http server
ip pim send-rp-announce Loopback0 scope 16
ip pim send-rp-discovery Loopback0 scope 5
```

```
!
!
!
[text omitted]
!
end
```

Class-Based Weighted Fair Queueing (CBWFQ) Configuration with IP Precedence

The following configuration example was used in test 6 of the case study. During this test, the effectiveness of CBWFQ was measured with IP precedence. Three classes were configured to differentiate service levels based on IP precedence settings of 0, 2, and 5. Traffic was generated with the three IP precedence settings, and precedence was marked on voice and video traffic. To oversubscribe the rate, data streams were injected with TOS 0 on top of the prioritized voice, video, and TCP data traffic.

The following provides an example of the tested router configuration:

```
version 12.1
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
service password-encryption
service udp-small-servers
service tcp-small-servers
!
hostname router2
!
class-map CLASS5
  match access-group name IP-PRES5
class-map CLASS2
  match access-group name IP-PRES2
class-map CLASS0
  match access-group name IP-PRES0
!
!
policy-map TOS-WFQ
  class CLASS0
   bandwidth percent 10
  class CLASS2
   bandwidth percent 25
  class CLASS5
    priority 2500
!
ip subnet-zero
no ip domain-lookup
ip host router1 10.20.1.2
ip host router4 10.100.100.1
ip host router3 10.101.100.1
ip host router2 10.20.1.3
!
ip multicast-routing
ip cef
cns event-service server
!
!
!
!
!
vc-class atm 6Mbs
  vbr-nrt 5950 5900 32
  oam-pvc manage 3
  oam retry 5 3 3
  encapsulation aal5snap
```

```
!
interface Loopback0
 ip address 10.20.254.100 255.255.255.255
 ip pim sparse-dense-mode
!
interface ATM3/0
 no ip address
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
 no atm ilmi-keepalive
!
interface ATM3/0.1 multipoint
 description Client ELAN simulating Case Airline User ELAN
 ip address 10.20.1.3 255.255.255.0
 no ip redirects
 ip pim sparse-dense-mode
 lane client ethernet elan1
 standby 2 timers 1 5
 standby 2 priority 150 preempt delay 60
 standby 2 ip 10.20.1.1
!
interface ATM3/0.2 multipoint
 description server ELAN simulating Case Airline CORIP
 ip address 10.20.2.3 255.255.255.0
 ip helper-address 10.20.1.101
 no ip redirects
 ip pim sparse-dense-mode
 lane client ethernet elan2
 standby 1 timers 1 5
 standby 1 priority 110 preempt delay 60
 standby 1 ip 10.20.2.1
!
interface ATM4/0
 no ip address
 load-interval 30
 no atm ilmi-keepalive
!
interface ATM4/0.1 point-to-point
 bandwidth 6000
 ip address 10.101.2.1 255.255.255.0
 ip pim sparse-dense-mode
 pvc 200/200
  class-vc 6Mbs
  service-policy output TOS-WFQ
 !
!
router ospf 1
 log-adjacency-changes
 area 100 stub no-summary
 area 100 range 10.100.0.0 255.255.0.0
 area 101 stub no-summary
 area 101 range 10.101.0.0 255.255.0.0
 redistribute static subnets
 network 10.20.0.0 0.0.255.255 area 0
 network 10.100.0.0 0.0.255.255 area 100
 network 10.101.0.0 0.0.255.255 area 101
!
ip default-gateway 172.26.64.1
ip classless
ip route 172.26.11.0 255.255.255.0 10.20.1.4
ip http server
ip pim send-rp-announce Loopback0 scope 16
ip pim send-rp-discovery Loopback0 scope 5
!
```

```
!
ip access-list extended IP-PRES0
 permit ip any any precedence 0
ip access-list extended IP-PRES3
 permit ip any any precedence 3
ip access-list extended IP-PRES5
 permit ip any any precedence 5
!
[text omitted]
!
end
```

Hybrid Configuration Example Using WRED at City HQ and WFQ at Remote Airports

The final test (test 7) used WRED at City HQ and WFQ at the remote airports. RED was used for low priority TCP traffic with four classes used to define the three TOS-based classes, and a fourth class to match TOS 0 and the TCP protocol. Four traffic flows were tested across the different classes with voice, video, and data traffic. To oversubscribe the rate, data streams were injected with TOS 0 on top of the prioritized voice, video, and TCP data traffic.

The following provides an example of the tested router configuration:

```
Current configuration:
!
version 12.1
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
service password-encryption
service udp-small-servers
service tcp-small-servers
!
hostname router2
!
!
class-map CLASS0-TCP
  match access-group name TCP-CLASS0
class-map CLASS2-TCP
  match access-group name TCP-CLASS2
class-map CLASS5
  match access-group name IP-PRES5
class-map CLASS2
  match access-group name IP-PRES2
class-map CLASS0
  match access-group name IP-PRES0
!
policy-map TOS-HYBR
  class CLASS0-TCP
   bandwidth percent 5
   random-detect
  class CLASS0
   bandwidth percent 10
  class CLASS2
   bandwidth percent 25
  class CLASS5
    priority 2500
  class class-default
   fair-queue
!
ip subnet-zero
no ip domain-lookup
ip host router1 10.20.1.2
ip host router4 10.100.100.1
ip host router3 10.101.100.1
```

```
ip host router2 10.20.1.3
!
ip multicast-routing
ip cef
cns event-service server
!
!
!
!
!
vc-class atm 6Mbs
  vbr-nrt 5950 5900 32
  oam-pvc manage 3
  oam retry 5 3 3
  encapsulation aal5snap
!
interface Loopback0
 ip address 10.20.254.100 255.255.255.255
 ip pim sparse-dense-mode
!
interface ATM3/0
 no ip address
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
 no atm ilmi-keepalive
!
interface ATM3/0.1 multipoint
 description Client ELAN simulating Case Airline User ELAN
 ip address 10.20.1.3 255.255.255.0
 no ip redirects
 ip pim sparse-dense-mode
 lane client ethernet elan1
 standby 2 timers 1 5
 standby 2 priority 150 preempt delay 60
 standby 2 ip 10.20.1.1
!
interface ATM3/0.2 multipoint
 description server ELAN simulating Case Airline CORIP
 ip address 10.20.2.3 255.255.255.0
 ip helper-address 10.20.1.101
 no ip redirects
 ip pim sparse-dense-mode
 lane client ethernet elan2
 standby 1 timers 1 5
 standby 1 priority 110 preempt delay 60
 standby 1 ip 10.20.2.1
!
interface ATM4/0
 no ip address
 load-interval 30
 no atm ilmi-keepalive
!
interface ATM4/0.1 point-to-point
 bandwidth 6000
 ip address 10.101.2.1 255.255.255.0
 ip pim sparse-dense-mode
 pvc 200/200
  class-vc 6Mbs
  service-policy output TOS-HYBR
 !
!
router ospf 1
 log-adjacency-changes
 area 100 stub no-summary
```

```
 area 100 range 10.100.0.0 255.255.0.0
 area 101 stub no-summary
 area 101 range 10.101.0.0 255.255.0.0
 redistribute static subnets
 network 10.20.0.0 0.0.255.255 area 0
 network 10.100.0.0 0.0.255.255 area 100
 network 10.101.0.0 0.0.255.255 area 101
!
ip default-gateway 172.26.64.1
ip classless
ip route 172.26.11.0 255.255.255.0 10.20.1.4
ip http server
ip pim send-rp-announce Loopback0 scope 16
ip pim send-rp-discovery Loopback0 scope 5
!
!
ip access-list extended IP-PRES0
 permit ip any any precedence 0
ip access-list extended IP-PRES2
 permit ip any any precedence 2
ip access-list extended IP-PRES5
 permit ip any any precedence 5
ip access-list extended TCP-CLASS0
 permit tcp any any precedence 0
!
[text omitted]
^C
end
```

Case Conclusions

From the seven QoS test configurations explored in this case, the following conclusions were drawn:

- You can mark traffic with a higher precedence bit and implement QoS in a way that guarantees that preferred traffic preempts lower priority traffic even during periods of congestion, up to the total link bandwidth.
- Higher classes of traffic pass through the network with a lower latency than lower class traffic during periods of network congestion.
- Deploying RED for TCP data flows is an effective way to signal TCP end nodes to reduce the window size and to increase overall throughput of TCP during periods of congestion.
- CBWFQ is necessary to differentiate data flows from the same device that might have differing priorities.

QoS for AVVID Services over Low-Speed ATM VCs Configuration Example

In a VoIP network, voice quality is tightly coupled with packet loss, delay, and jitter. Therefore, network engineers must design enterprise networks with the objective of eliminating VoIP packet loss, and minimization of packet delay and jitter. To do this, you must use QoS in every part of the network. Specifically, it requires you to implement advance queueing and scheduling techniques in the distribution and core areas of the network, as well as in the WAN. Because WAN links are usually the lowest-speed circuits in an Architecture for Voice, Video, and Integrated Data (AVVID) network, you

need to give particular attention to reducing loss, delay, and jitter across the WAN links. WAN aggregation routers are often taxed by maintaining large routing tables for remote branches. Therefore, it is also important to consider the CPU overhead of queueing mechanisms at the WAN edge.

Figure 8-4 shows a typical AVVID deployment over low-speed ATM VCs running at or below 768 Kbps in a WAN aggregation, enterprise network environment.

Figure 8-4 Low-Speed WAN Aggregation Example

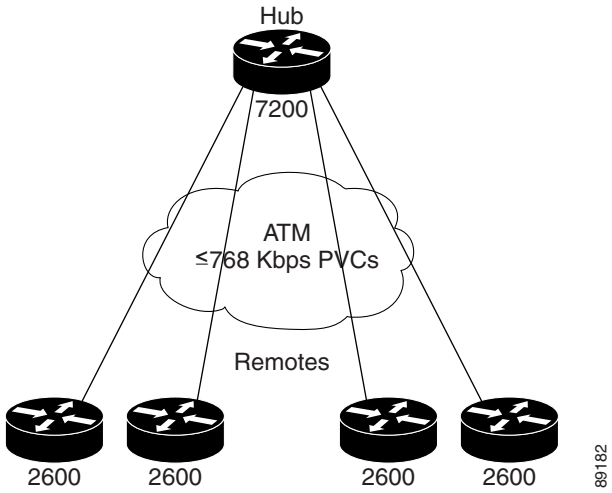

In this example, the hub is a Cisco 7200 series router, and the remote routers are Cisco 2600 multiservice access routers. To implement the configuration , all routers require a minimum of Cisco IOS Release 12.2 T.

Table 8-4 provides the recommended configuration guidelines in this environment:

Table 8-4 Configuration Guidelines for AVVID on Low-Speed ATM VCs

Configuration Area	Recommendation
IP Precedence	Use the following classifications: • Voice media—IP precedence 5 or DSCP EF • Voice control—IP precedence 3 or DSCP AF31 • Mission-critical traffic—IP precedence 2 or DSCP AF21
MLPPP over ATM	Use Multilink PPP (MLPPP) over ATM for fragmentation if you are using VCs at or below 768 Kbps.
Per-VC Low Latency Queueing (LLQ)	Enable priority queueing (PQ) for voice and Class-Based Weighted Fair Queueing (CBWFQ) for other classes of traffic on a per-VC basis.
Traffic Shaping	Shape the PVC to a rate that is appropriate for the service contract with your service provider and the ATM policing performed to enforce it.
Transmit Ring Size	Adjust the **tx-ring-limit** command for all ATM links below DS-3 speeds.

Note The following partial examples are abbreviated to show the most relevant pieces of the configuration.

Cisco 7200 Series Router Configuration

```
!
ip cef
!
class-map Voice
  match ip precedence 5
class-map Call-Control
  match ip precedence 3
class-map Mission-Critical
  match ip precedence 2
!
! Policy-map example assumes 768k link speed and 45% (of 768K shaped rate) for LLQ
!
policy-map QoS-Policy-768k
  class Voice
    priority 328
  class Call-Control
   bandwidth percent 10
  class Mission-Critical
   bandwidth percent 20
  class class-default
   fair-queue
!
interface ATM5/0
 no ip address
 no ip mroute-cache
 no atm ilmi-keepalive
!
! This example shows the use of abr for ATM VC; other possibilities (e.g. vbr-nrt) exist
!
interface ATM5/0.37 point-to-point
 pvc cisco37 100/37
  tx-ring-limit 3
  abr 768 768
  protocol ppp Virtual-Template37
 !
!
interface Virtual-Template37
 bandwidth 768
 ip address 10.1.37.2 255.255.255.252
 service-policy output QoS-Policy-768k
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
!
```

Cisco 2600 Series Router Configuration

```
!
ip cef
!
class-map Voice
  match ip precedence 5
class-map Call-Control
  match ip precedence 3
class-map Mission-Critical
  match ip precedence 2
!
! Policy-map example assumes 768K link speed and 45% (of shaped 768K rate) for LLQ
```

```
!
policy-map QoS-Policy-768k
  class Voice
    priority 328
  class Call-Control
   bandwidth percent 10
  class Mission-Critical
   bandwidth percent 20
  class class-default
   fair-queue
!
interface ATM2/0
 no ip address
 no ip mroute-cache
 no atm ilmi-keepalive
!
! This example shows the use of abr for ATM VC; other possibilities (e.g. vbr-nrt) exist
!
interface ATM2/0.37 point-to-point
 pvc cisco37 100/37
  tx-ring-limit 3
  abr 768 768
  protocol ppp Virtual-Template1
 !
!
interface Virtual-Template1
 bandwidth 768
 ip address 10.1.37.1 255.255.255.252
 service-policy output QoS-Policy-768k
 ppp multilink
 ppp multilink fragment-delay 10
 ppp multilink interleave
!
```

Frequently Asked Questions

This chapter answers some of the frequently asked questions (FAQs) about traffic management on the PA-A3 and PA-A6 ATM port adapters. Many of these questions are answered within the chapters of this document but are identified here also as a quick reference.

This chapter includes the following sections:

General FAQs

This section provides a list of general questions about the ATM processing on the Cisco 7200 series routers.

What types of queues are implemented on the Cisco 7200 series to support ATM traffic?

The Cisco 7200 series router implements both hardware and software queues. The ATM port adapter uses hardware queues located on the ATM port adapter itself, and also on the network processing engine (NPE) or network services engine (NSE). The PA-A3 and PA-A6 ATM port adapters have a separate receive buffer and transmit buffer for segmentation and reassembly (SAR) processing located on the port adapter. A direct memory access (DMA) transfer occurs between the hardware buffers on the port adapter and the private interface pool located on the NPE or NSE.

The software queue is one or more Layer 3 queues, whose implementation is dependent upon the type of ATM port adapter. The PA-A3 and PA-A6 ATM port adapters activate their Layer 3 queues when congestion builds on the router and the transmit ring is full. Packets awaiting transmission enqueue to the Layer 3 queue until they can be placed onto the transmit ring of the outbound port adapter. For the PA-A3 and PA-A6 ATM port adapters, the Layer 3 queue is per VC. For all other ATM PAs, the Layer 3 queue is per interface.

What is the transmit ring and how does it work?

The transmit ring is a control structure associated with the outbound port adapter. Each entry on the transmit ring is associated with one particle in the private interface pool. If a packet needs four private interface particles, it also uses four ring entries. The transmit ring is located on the NPE or NSE and points to packet content elsewhere in I/O memory that is awaiting transmission. The packet content is stored in particles of the inbound private interface pool. The transmit ring operates on a first-in first-out (FIFO) basis. When a packet is ready to be serviced by the transmit ring, a DMA transfer moves the packet contents from the private interface pool to the transmit buffer on the port adapter for SAR processing. Once the packet content has been transferred to the hardware buffer on the ATM port adapter, the ring entries are freed.

Note If there is only a single transmit ring entry available, but the packet is larger and requires additional particles and therefore ring entries, then the router still uses that one entry for the packet awaiting transmission.

What is the transmit ring limit and when should you tune it?

The transmit ring limit specifies an upper boundary on the number of ring entries that any one PVC can consume for outbound packets. The default ring limit varies by the type of service category that you configured for the PVC. You should customize the transmit ring limit when you need to adjust its size to allow activation of Layer 3 queues to reduce latency. The larger the ring limit, the greater the tolerance of bursts of traffic, but the longer the delay. Large ring limits can prevent Layer 3 queueing from activating.

Does the transmit ring store packets?

No. When the transmit ring receives control of a packet for transmit processing, the ring entries link back to the physical location of the particles within the ingress private interface pool where the packet content resides. Packet content generally is stored in the private interface pool of the interface on which the packet is received.

What are some of the differences between how process-switched packets and CEF or fastswitched packets are handled during ATM processing on the Cisco 7200 series router?

Process-switched packets are stored in the public normal pool on the NPE or NSE rather than the private interface pool, where CEF and fastswitched packets are stored. Process-switched packets automatically enqueue to the Layer 3 queue, regardless of whether the transmit ring has available entries. Therefore, Layer 3 queueing is always active for process-switched packets. The transmit ring must be full (no more ring entries available) for CEF and fastswitched packets to enqueue to their corresponding Layer 3 queues.

Is the Committed Access Rate (CAR) feature used for traffic policing on the Cisco 7200 series with ATM?

No. CAR is a legacy policing mechanism that is no longer recommended for any policing on the Cisco 7200 series router. Cisco Systems has a newer traffic policing mechanism that is class-based that you can implement using the modular QoS CLI (MQC) configuration method.

Does the Cisco 7200 series router support the Guaranteed Frame Rate (GFR) service category?

No.

What is native traffic shaping?

Those SAR processors that support scheduling within the hardware based upon traffic shaping values that are configurable through the Cisco IOS software are said to support native ATM traffic shaping. All ATM port adapters except the PA-A1 support native traffic shaping.

Note The PA-A1 ATM port adapter does support the UBR service category by default. However, UBR is considered a best-effort transmission method and therefore, is not technically considered to be traffic shaping.

What is the difference between native traffic shaping and Cisco IOS software shaping?

Native traffic shaping is implemented in the ATM port adapter hardware, while shaping within Cisco IOS is software based and requires more CPU resources. Native traffic shaping is the preferred method of implementing shaping on ATM port adapters and has better traffic descriptors for shaping variable bit rate non-real-time (nrt-VBR) traffic.

PA-A6 ATM Port Adapter FAQs

This section provides a list of questions about the PA-A6 ATM port adapter and provides information about how the PA-A6 ATM port adapter compares with the PA-A3 ATM port adapter.

What capabilities does the PA-A6 ATM port adapter provide over the PA-A3 ATM port adapter?

The PA-A6 ATM port adapter provides support for up to 8191 VCs compared to 4096 VCs for the PA-A3 ATM port adapter. The PA-A6 ATM port adapter also provides performance improvements over the PA-A3 ATM port adapter. The PA-A6 provides line rate performance using 128-byte packet sizes on the Cisco 7200 series routers using the NPE-400.

Are there any platforms that are currently supported on the PA-A3 ATM port adapter that are not supported on the PA-A6?

As of Cisco IOS Release 12.2(15)T, the PA-A6 ATM port adapter is not currently supported on the Cisco 7500 series routers. It is also not currently available on the Cisco 7600 FlexWAN.

Note For the latest information on the minimum supported Cisco IOS software releases and hardware compatibility, refer to the Software Advisor tool.

What kinds of applications does the PA-A6 ATM port adapter target support?

The PA-A6 ATM port adapter targets support for broadband aggregation applications on the Cisco 7200 series routers and Cisco 7401ASR router, where xDSL aggregation installations require large support for high numbers of VCs per interface. The PA-A6 ATM port adapter supports 8K connections (subscribers) per interface for features like Point-to-Point Protocol over ATM (PPPoA), Point-to-Point Protocol over Ethernet over ATM (PPPoEoA), and routed bridge encapsulation (RBE).

The PA-A6 ATM port adapter also supports WAN aggregation and campus/MAN networks that require high performance and can support greater than 4K VCs per interface.

What is the SDRAM and SSRAM used for in the PA-A3 and PA-A6 ATM port adapters and why is it important?

The synchronous dynamic random access memory (SDRAM) and the synchronous static random access memory (SSRAM) are used on the PA-A3 and PA-A6 ATM port adapter hardware to provide additional storage for the requirements of segmentation and reassembly (SAR) processing. This includes storage for the calendar scheduling table, and receive and transmit buffers.

ATM Port Adapter	SDRAM	SSRAM (per SAR processor)
PA-A3	4 MB	512 KB
PA-A6	32 MB	4 MB

What processing engines does the PA-A6 ATM port adapter currently support on the Cisco 7200 series routers?

The PA-A6 ATM port adapter is currently supported with the NPE-400 and NSE-1 processing engines.

Note For the latest information on hardware compatibility information, refer to the Software Advisor tool.

QoS FAQs

This section provides a list of questions about QoS support on the Cisco 7200 series routers for the PA-A3 and PA-A6 ATM port adapters. For additional questions about QoS, refer to the *QoS Frequently Asked Questions* publication.

What QoS features are supported on a per-VC basis for the PA-A3 and PA-A6 ATM port adapters?

You can configure Weighted Random Early Detection (WRED), Class-Based Weighted Fair Queueing (CBWFQ), and Low Latency Queueing (LLQ) at the VC level on the PA-A3 and PA-A6 ATM port adapters. The QoS features that are supported on a per-VC basis in ATM are referred to as IP to ATM Class of Service (CoS).

When is Cisco Express Forwarding (CEF) switching required for IP to ATM CoS features using the PA-A3 or PA-A6 ATM port adapters on the Cisco 7200 series router?

CEF switching is required to implement WRED on an ATM PVC on the Cisco 7200 series router. It is also required if you are using CBWFQ in a Network-Based Application Recognition (NBAR) environment. CEF switching is recommended when you are using CBWFQ and LLQ.

Note On the Cisco 7500 series router, dCEF is required for all IP to ATM CoS features using the PA-A3 and PA-A6 ATM port adapters.

Can WRED be configured at the same time with CBWFQ?

Yes. Be aware that the output from the **show queueing interface atm** command might only show "weighted fair" as the queueing strategy even if WRED is also configured. CBWFQ characterizes how packets are dequeued to the transmit ring. These queueing strategies define the order in which packets leave the Layer 3 queue for transmission. WRED provides an alternative method for congestion avoidance on the Layer 3 queue. WRED is a proactive drop policy to help manage congestion on a Layer 3 queue before the queue limit is reached. Therefore, WRED manages what packets are able to enqueue to the Layer 3 queue.

What is the difference between WFQ and CBWFQ?

Native Weighted Fair Queueing (WFQ) assigns a weight to each conversation, and then schedules the transmit time for each packet of the different flows. The weight is a function of the IP precedence of each flow, and the scheduling time depends on the packet size. WFQ was implemented for slow speed links (such as serial) to provide a fair treatment for each type of traffic. To do its job, WFQ classifies the traffic into different flows based on the associated Layer 3 and Layer 4 information (IP addresses, TCP ports, and so on).

You do not need to define access-lists in order for this to work. Therefore, with WFQ, low bandwidth traffic has effective priority over high bandwidth traffic. The high bandwidth traffic shares the transmission media proportionally to assigned weights. However, WFQ is not scalable if the flow amount increases considerably, and it is not available on high-speed interfaces such as ATM interfaces.

CBWFQ provides a solution to these limitations. CBWFQ assigns a weight to each configured class instead of each flow. The bandwidth you assign to a class is used to calculate the weight of that class. More precisely, the weight is a function of the interface bandwidth divided by the class bandwidth. Therefore, the bigger the bandwidth parameter, the smaller the weight. The weight of each packet that matches the class criteria is also calculated from this. WFQ is then applied to the classes (which can include several flows) rather than the flows themselves.

What is fancy queueing?

Fancy queueing is a general reference to any Layer 3 queueing policy that you configure to replace the default queueing behaviors.

What is a hold queue?

A hold queue is a Layer 3 software queue. The hold queue can be an interface queue or a per-VC queue. The PA-A3 and PA-A6 ATM port adapters support per-VC hold queues only.

What is the queue limit?

The queue limit is also called the queue depth, which specifies the maximum number of packets that can be placed onto the queue before activating a drop mechanism. The default queue limit varies by the type of queueing policy. The queue limit is configurable for FIFO, CBWFQ, and LLQ policies.

What is the default drop strategy?

Tail drop is the default drop strategy on a Layer 3 queue when the queue is full. With tail drop, no packets make it to the queue and all packets are dropped.

When is tail drop activated?

Tail drop occurs when the queue limit is reached for a per-VC queue on the PA-A3 ATM port adapter or PA-A6 ATM port adapter. Tail drop continues to be used even if you have configured WRED, when the average queue depth exceeds the maximum threshold value for WRED.

What is the default congestion management strategy?

FIFO is the default congestion management strategy on a Layer 3 queue. This is the same strategy implemented by the hardware buffer located on the ATM port adapter.

Why does tail drop occur when WRED is configured for congestion avoidance?

WRED does not replace tail drop. WRED is an intelligent drop policy that is implemented before the hold queue reaches its queue limit, or maximum threshold. Tail drop occurs after the queue is already full, when the mean queue depth for WRED exceeds the maximum threshold value, and when the queue limit is reached.

Where do you apply service policies on the PA-A3 and PA-A6 ATM port adapters?

The PA-A3 and PA-A6 ATM port adapters only support per-VC Layer 3 queues, not interface queues. Therefore, you should apply Layer 3 service policies at the VC for these port adapters.

Why would you want to activate Layer 3 queues instead of going directly to the transmit ring?

The transmit ring uses a FIFO queueing strategy, which means that higher priority traffic can be queued behind lower priority traffic. Therefore, you can use Layer 3 queueing policies to achieve differentiated levels of service and achieve priority for certain packets before they are sent to the transmit ring. However, if you activate the Layer 3 queue but do not configure a policy such as CBWFQ or LLQ, then FIFO is the default strategy and no benefit is achieved. Recall that FIFO is also the mechanism implemented in the hardware buffers on the PA-A3 and PA-A6 ATM port adapters. You cannot change the queueing mechanism within the hardware buffer. Therefore, to achieve any differentiation over which packets reach the FIFO-based hardware queue before other packets on the VC, you need to configure and activate a Layer 3 queueing policy.

A

B

C

D

F

G

H

I

J

M

N

O

P

Q

R

S

T

U

V