



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, page 2-2](#)
- [Memory Architecture on a Cisco 7200 Series Router, page 2-4](#)
- [Layer 3 Software Queues and QoS Processing, page 2-16](#)
- [Summary of Hardware and Software Queues on the Cisco 7200 Series Router, page 2-21](#)
- [SAR Processors, page 2-22](#)
- [Summary of Traffic Flow Through the ATM Port Adapter, page 2-36](#)
- [Related Documentation, page 2-36](#)
- [Next Steps, page 2-37](#)

Basic Traffic Flow on a Cisco 7200 Series Router

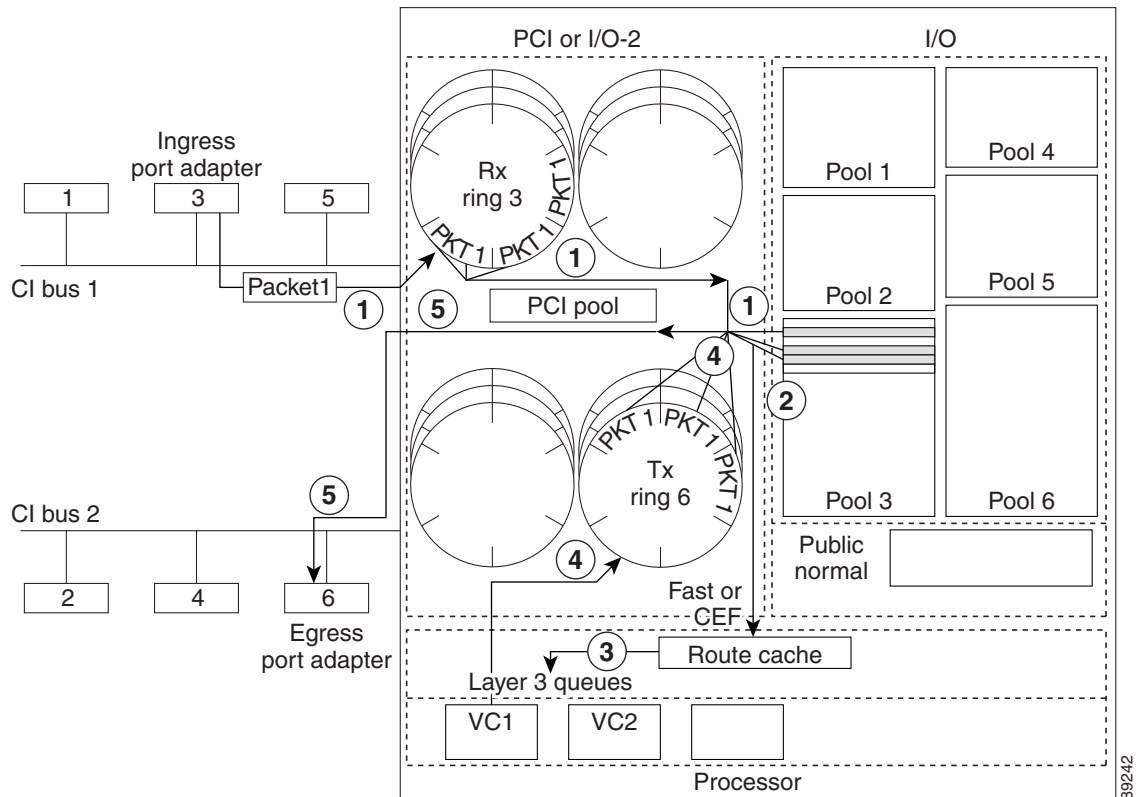
To better understand the scope of managing ATM traffic on a Cisco 7200 series router, it is worthwhile to review the basic flow of ATM traffic from the receipt of packets to their release as ATM cells.

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	<p>The router processes the packets through a switching path: Cisco Express Forwarding (CEF), fast switching, or process switching.</p> <p>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.</p>
3	<p>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:</p> <ul style="list-style-type: none"> 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. <p>Note All process-switched packets automatically enqueue to the corresponding Layer 3 queue whether or not there are packets currently enqueued there.</p>
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](#), page 2-4
- [Particle-Based Memory](#), page 2-4
- [Private Interface Pools and Public Pools](#), page 2-7
- [Receive Rings and Transmit Rings](#), page 2-10

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 discontinuous, or nonadjacent. Discontinuous 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 discontinuous 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 discontinuous there), into a single buffer within the public normal pool that is large enough to hold the entire packet.

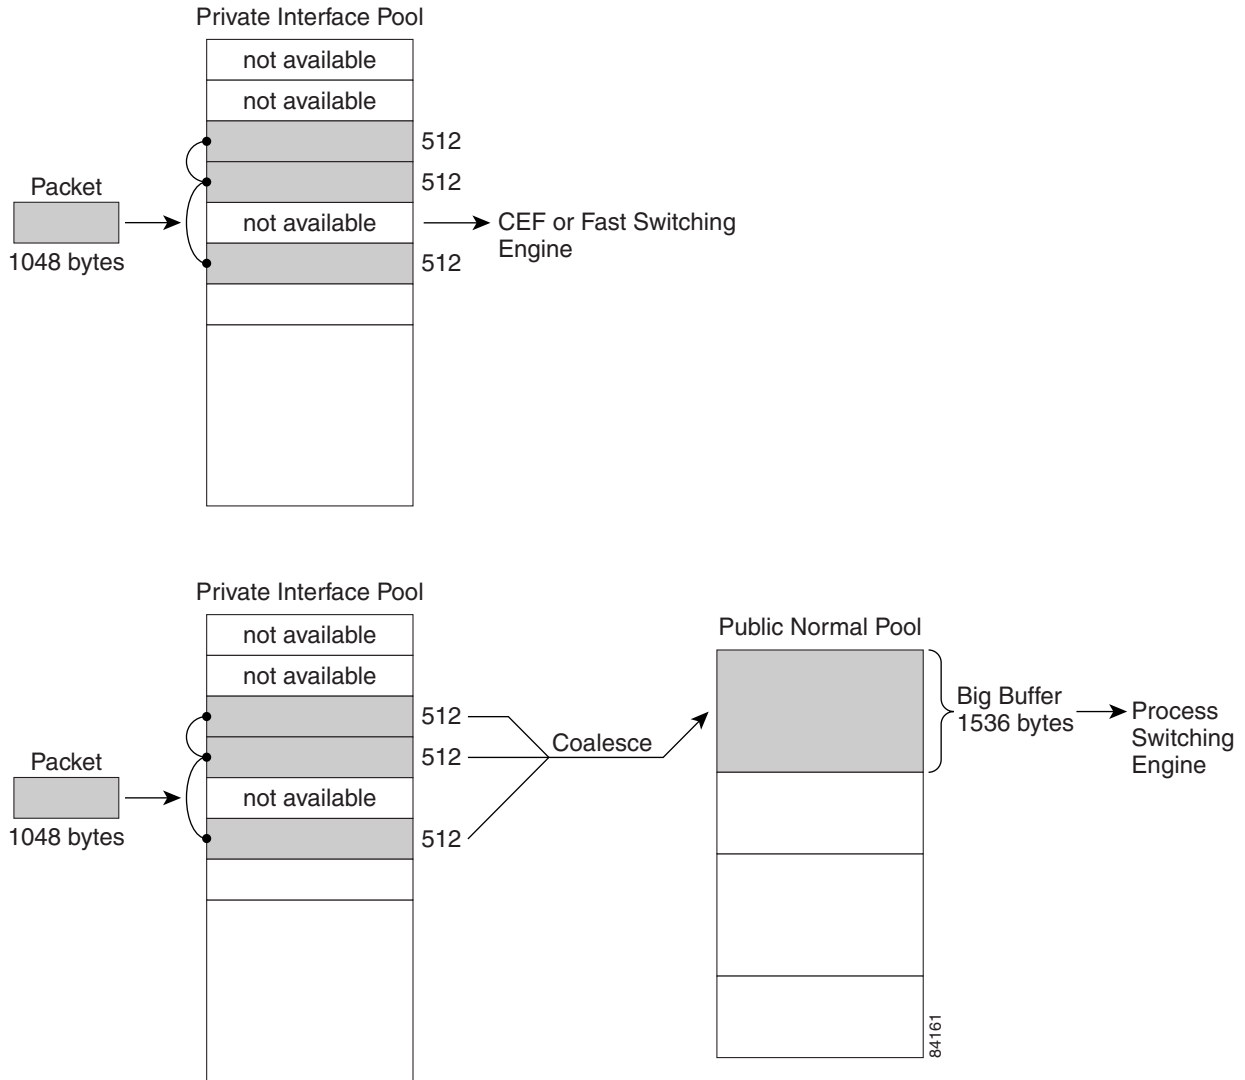


Note

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 discontinuous particles.

Figure 2-3 Handling of Discontinuous 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](#).

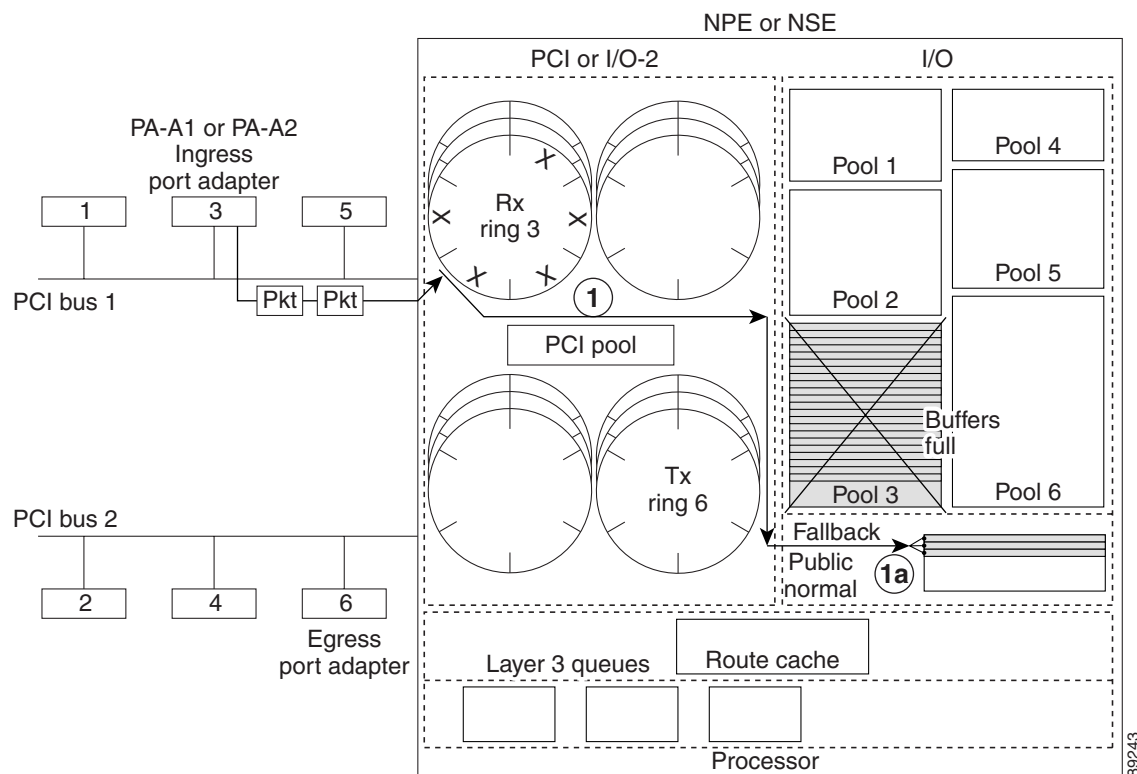
Cisco IOS software uses five different public buffer pool sizes, or categories, as shown in [Table 2-2](#):

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.

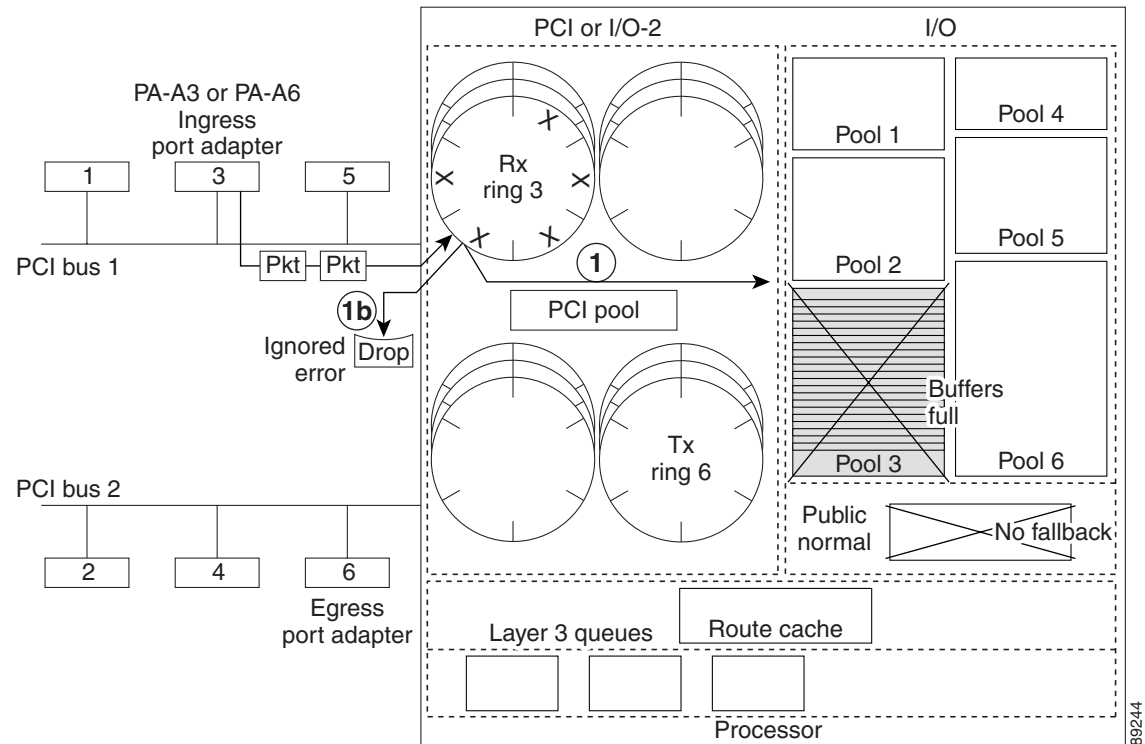
Figure 2-4 Fallback to Public Normal Pool for PA-A1 and PA-A2 ATM Port Adapters



- | | |
|-----------|--|
| 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 <i>does not</i> 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.

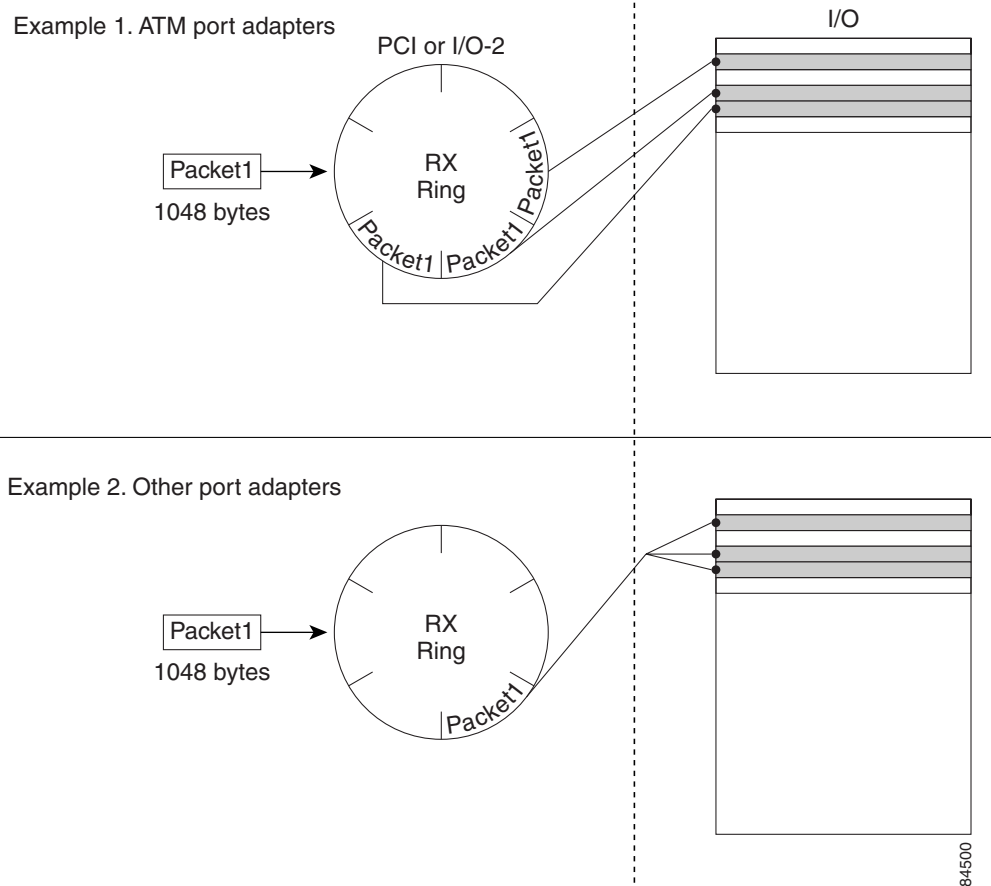
A unique receive ring and transmit ring structure exists for each port adapter on the Cisco 7200 series router. So, for six port adapters, there are six corresponding sets of receive and transmit rings that reside on the NPE or NSE.

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.

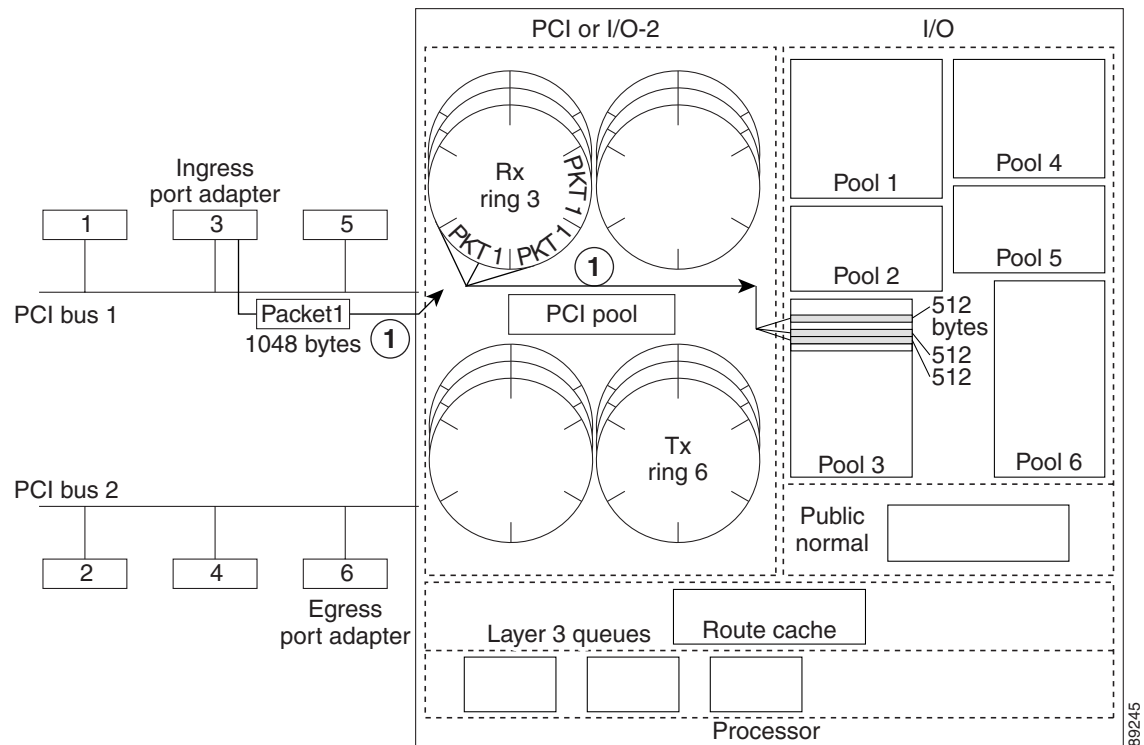
Some port adapters use only a single entry on the receive ring, as shown in Example 2 in [Figure 2-6](#). This entry links to one or more particles where the owned packet is stored in the private interface pool. Other port adapters, including all of the ATM port adapters on the Cisco 7200 series router,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



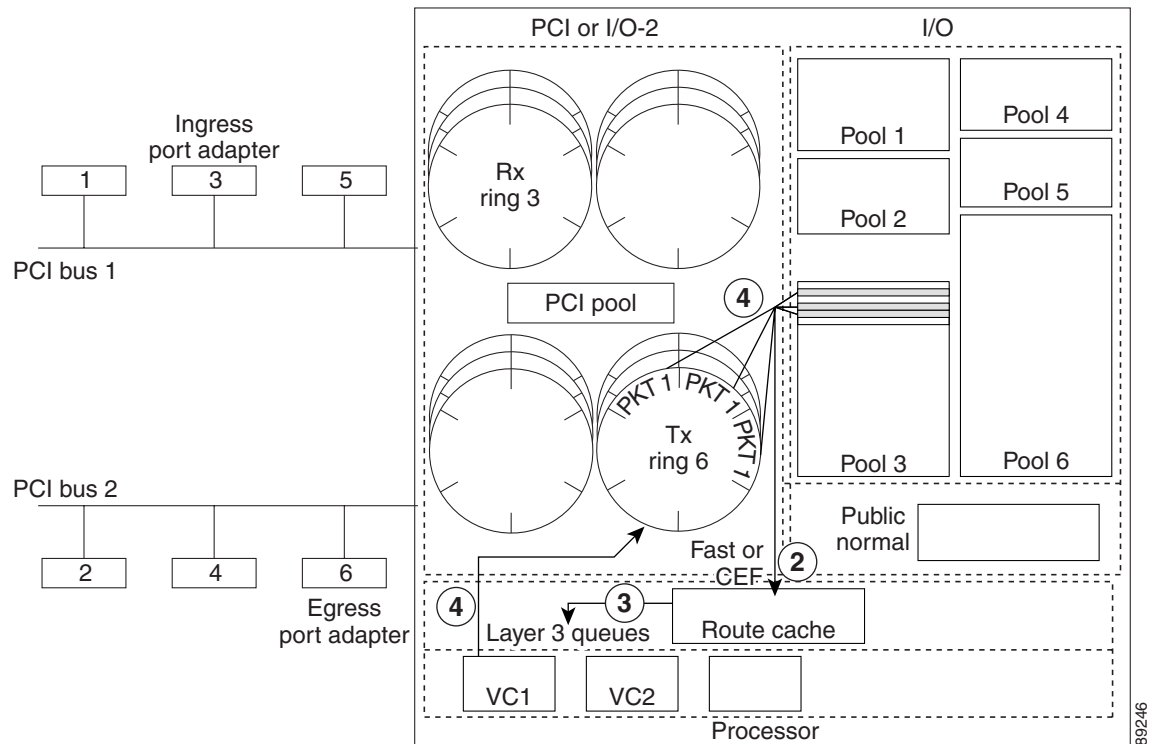
- 1** The Cisco 7200 series router receives a 1048-byte packet on an ATM port adapter to a hardware queue called the receive ring. The packet requires three 512-byte particles for storage in private interface pool 3. The receive ring creates three entries (one for each particle of the packet) that point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

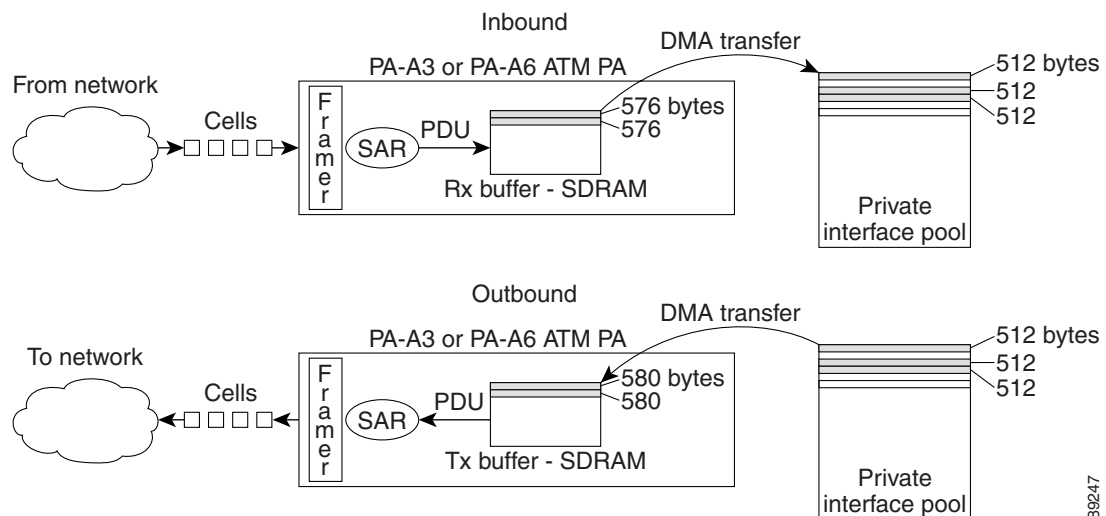
- [Receive Buffer and Transmit Buffer Located on the PA-A3 and PA-A6 ATM Port Adapters](#)
- [Per-VC Limits on the Receive and Transmit Rings](#)

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



89247

- 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, page 2-17](#)
- [Activation of Layer 3 Queues, page 2-17](#)
- [Relationship of Layer 3 Queues to the Transmit Ring, page 2-18](#)
- [Cisco IOS QoS Software, page 2-18](#)

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: <ul style="list-style-type: none"> Tail drop (default) or WRED For congestion management on the queue: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> Tail drop (default) or WRED For congestion management on the queue: <ul style="list-style-type: none"> 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, page 2-23](#)
- [Understanding Line Rates and Cell Rates, page 2-24](#)
- [Traffic Shaping Algorithms Used By the SAR Processors, page 2-28](#)

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:

$$1 \text{ (cell)} / \text{ATM cell rate (cells per second)} = \text{ATM cell time}$$

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

$$1 / 3622.64 = .00027604 \text{ seconds, or } 276.04 \text{ microseconds per ATM cell}$$

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 $1/\text{PCR}$ 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 $1/\text{PCR}$. 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 $1/\text{SCR}$. 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 $1/\text{PCR}$. 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 $T_n + \text{ICG}$. 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-insertion, 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](#)

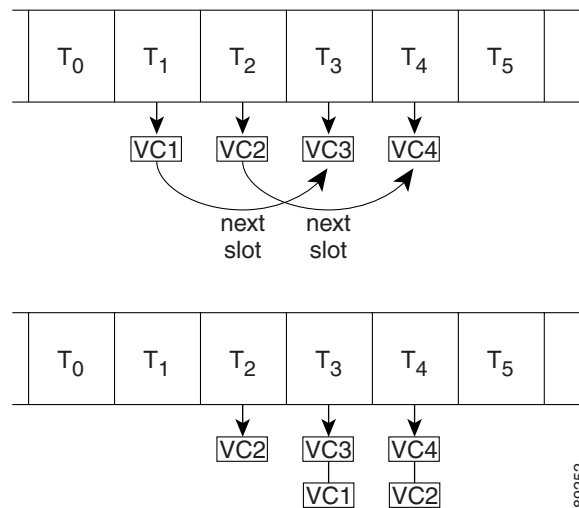
Example 1: Collisions for PVCs with the Same Transmission Priority

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 ($T_2 + 2 \text{ (ICG)}$).
- 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 ($T_3 + 2 \text{ (ICG)}$).
 - 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

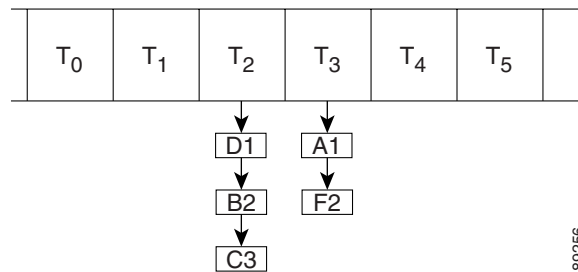


89255

- At time T₁, the SAR processor takes the following actions:
 - Transmits a cell from A1.
 - Reschedules A1 to transmit at time slot T₃ (T₁ + 2 (ICG)).
 - Links B2 between D1 and C3 in the T₃ link list.

Figure 2-18 shows that the SAR processor places A1 ahead of F2 in time slot T₃ due to its priority. In similar fashion, it bumps B2 between D1 and C3 in time slot T₂. 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

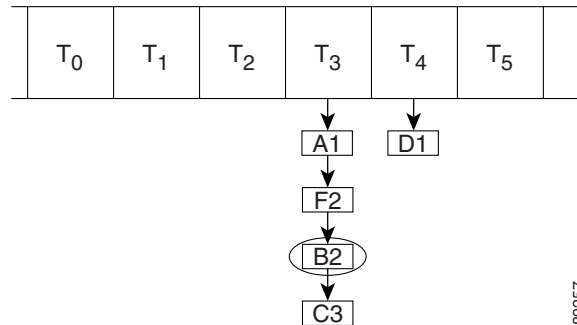


89256

- At time T₂, the SAR takes the following actions:
 - Transmits a cell from D1.
 - Reschedules D1 to transmit at time slot T₄ (T₂ + 2 (ICG)).
 - Links B2 and C3 at the tail-end of the T₃ link list.

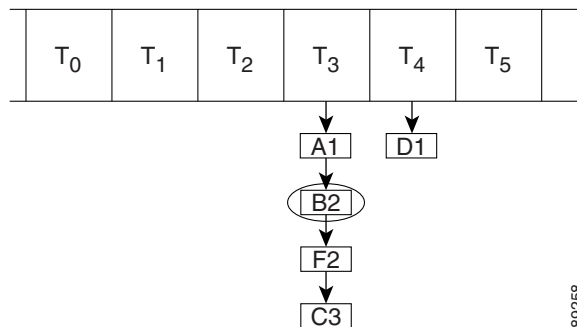
Figure 2-19 shows that when the SAR processor bumps B2 and C3 to time slot T₃ using tail-insertion, it places B2 below F2 (same priority) and C3 at the end.

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	<i>Cisco ATM Solutions</i> , 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	<i>Inside Cisco IOS Software Architecture</i> , 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.

