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.