

# Understanding Per-VC Transmit Queuing on the PA-A3 and NM-1A ATM Interfaces

Document ID: 6187

---

## Introduction

### Before You Begin

Conventions

Prerequisites

Components Used

### Approaches to Per-VC Queues

### PA-A3 Architecture Overview

### ATM Network Modules Architecture Overview

### Before IP to ATM CoS

### After IP to ATM CoS

### hold-queue and vc-hold-queue Commands

### Related Information

---

## Introduction

Newer router ATM hardware, including the PA-A3 port adapter, NM-1A-OC3/DS3, and Inverse Multiplexing for ATM (IMA) adapter, creates a separate packet queue for each virtual circuit (VC) in the interface hardware buffers. These buffers are also known as the transmit ring. The purpose of per-VC queues is to ensure that one congested VC does not consume all the memory resources and starve other VCs.

This document reviews the per-VC queuing approach taken by the PA-A3 and by the NM-1A network modules. It also reviews how per-VC queuing changes when the ATM interface is configured with IP to ATM Class of Service (CoS) features.

## Before You Begin

### Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.

### Prerequisites

There are no specific prerequisites for this document.

### Components Used

This document is not restricted to specific software and hardware versions.

## Approaches to Per-VC Queues

Every ATM interface must ensure fair access to packet buffers by all VCs. In general, there are two approaches to allocating packet buffers:

- Allocate a fixed number of buffers to each VC. In other words, divide the buffers evenly among all

the configured VCs.

- Allow any one VC to use as many packet buffers as needed.

The PA-A3 and the ATM network modules implement a combination of the two approaches. During VC setup, the interface driver grants a transmit credit to the VC. The Cisco IOS® configuration guides refer to the transmit credit as the transmit ring. You can tune the value of the transmit ring with the tx-ring-limit command.

Here is a brief description of how each hardware type computes the transmit credit:

- **PA-A3** – Computes the transmit credit based on the number of buffers needed to meet the sustained cell rate (SCR) of a variable bit rate, non-real time (VBR-nrt) VC. Available bit rate (ABR) and unspecified bit rate (UBR) VCs are assigned default values of 128 and 40, respectively.
- **NM-1A** – Computes the transmit credit based on the number of buffers needed to meet the peak cell rate (PCR) of the VC or a maximum transmission unit (MTU)-size frame. The network modules select the larger value.

More will be explained about each of these approaches in the following sections of this technical note.

## PA-A3 Architecture Overview

The PA-A3 uses separate receive and transmit segmentation and reassembly (SAR) chips. Each SAR is supported by its own subsystem of local memory on the PA-A3 itself. This memory stores packets as well as key data structures like the VC table. On the transmit path, the local memory consists of 6144 particles of 576 bytes (or 580 with an internal 4-byte header that travels with the packet inside the router). Out of these, the PA-A3 reserves 144 particles for system packets like operations, administration, and maintenance (OAM) cells. Use the **show controller atm** command to view the 6144 particles.

```
[...snip...]  
BFD Cache status:  
  base=0x62931AA0, size=6144, read=143  
Rx Cache status:  
[...snip...]
```

The PA-A3 assigns the following default values for the per-VC transmit credit:

VC Service Category	Default Transmit Credit	Time of Enforcement
VBR-nrt	Based on the formula: $(48 * SCR) / (Particle\_size * 5)$ <b>Note:</b> SCR is the cell rate with ATM overhead included.	Always
ABR	128	Always
UBR	40	Only when total credit utilization exceeds 75 percent or the tx_threshold value, as shown in <b>show controller atm</b> .

# ATM Network Modules Architecture Overview

The ATM network modules support 2048 transmit buffer descriptors (TBDs).

**Note:** If we took the first approach to allocating transmit credits and evenly divided these 2048 TBDs among the 1024 possible active VCs, each VC would have only two TBDs.

Use the **show controller atm** command to view the 2048 TBDs.

```
3640-2.2#show controller atm 3/0
Interface ATM3/0 is up
Hardware is RS8234 ATMOC3
LANE client MAC address is 0030.9475.10d0
hwidb=0x61FDA664, ds=0x61FDC31C
RS8234 base 3D800000, ds 61FDC31C, PM5346 base 3DC00000, slave base 3DC00000
SBDs - avail 2048, guaranteed 22, unguaranteed 2026, starved 0 rbds 3588
```

The interface driver grants one guaranteed TBD to each VC. This TBD ensures that every VC has a minimum instantaneous transmit opportunity. All other TBDs go into the unguaranteed pool, from which the active VCs can pull more buffers up to their transmit credit.

The interface driver computes a transmit credit for each VC at setup time. It chooses the larger value resulting from one of the two following formulas:

- MTU-based credit = maximum transmission unit (MTU) of the interface or subinterface divided by particle size (in other words, MTU / particle size)
- PCR-based credit = 100 microseconds \* PCR / typical frame size (in bits), where the typical frame size is 512 bytes.

The interface driver considers MTU because it needs to transmit one MTU-size frame with a single command to the SAR to ensure maximum performance. This requirement means the VC needs at least enough TBDs equal to the number of particles needed to transmit one MTU-size frame.

## Before IP to ATM CoS

This section discusses the per-VC queuing architecture before IP to ATM Class of Service (CoS).

Traditionally, in addition to the per-VC transmit credit on the transmit ring, an ATM interface (and all Cisco router interfaces) supported an output hold queue. This queue held packets that were locally generated by the router and any other packets that followed the process switched path. Process switching defines a method of forwarding packets through a router. See Cisco IOS Switching Paths.

By default, all interfaces use an output hold queue size of 40 packets. Use the **show interface atm** command to display the current value. Use the **hold-queue {value} out** command to configure a non-default value.

```
7206b(config)#interface atm 5/0
 7206b(config-if)#hold-queue ?
 <0-4096> Queue length
7206b(config-if)#hold-queue 75 out
7206b(config)#end
 7206b#show interface atm 5/0
 ATM5/0 is up, line protocol is up
 Hardware is ENHANCED ATM PA
 MTU 4470 bytes, sub MTU 4470, BW 44209 Kbit, DLY 190 usec, rely 255/255, load
 1/255
 Encapsulation ATM, loopback not set, keepalive not supported
 Encapsulation(s): AAL5
```

```

4096 maximum active VCs, 3 current VCCs
VC idle disconnect time: 300 seconds
22 carrier transitions
Last input 00:00:03, output 00:03:59, output hang never
Last clearing of "show interface" counters 2w0d
Queuing strategy: fifo
Output queue 0/75, 0 drops; input queue 0/75, 0 drops
[snip]

```

Now that we understand the interface output hold queue, we can discuss the path that a packet travels depending on whether it is following the fast path or the process path as it is forwarded by the router from ingress to egress interface.

A packet following the process-switched path goes through the following steps:

1. Packet is placed in the interface output hold queue.
2. The router catches the attention of the PA-A3 driver and announces that the queue holds data awaiting transmission.
3. The PA-A3 driver dequeues the packet from the output hold queue and copies the packet to local memory and to the appropriate VC's transmit ring.
4. If the VC's transmit credit or transmit ring is full, the interface driver continues to queue the packet on the common output hold queue.
5. If there is no room on the transmit ring after a short period, the driver discards the packet to avoid head of line blocking for packets destined from the common hold queue to other, non-congested VCs.

A packet following the fast-switched path goes through the following steps:

1. Fast-switched packets are sent directly to the transmit ring. Importantly, they initially bypass the output hold queue.
2. If the VC has filled its transmit credit, the packet is placed on the output hold queue.
3. If there is no room on the transmit ring after a short period, the driver discards the packet to avoid head of line blocking for packets destined from the common hold queue to other, non-congested VCs.

When the VC's transmit credit is reached, increasing the size of the hold queue will not prevent packet drops since the interface's priority is to avoid head-of-line-blocking.

## After IP to ATM CoS

This section explains how IP to ATM CoS changes output queuing on the ATM network modules and port adapters supporting per-VC queuing.

When running an image that automatically supports IP to ATM CoS (non-Route Switch Processor [RSP] platforms) or when VCs are configured explicitly with a service policy, the PA-A3 driver creates a hold queue per VC. In other words, since the hardware-level queue or transmit ring is maintained per VC, the PA-A3 extends the single output hold queue and creates a unique output hold queue per VC.

Importantly, this per-VC hold queue means that each VC has two sets of buffers, as illustrated below.

Queue	Location	Queuing Methods	Service Policies Apply	Command to Tune
Hardware queue or transmit	Port adapter or network	FIFO only	No	<b>tx-ring-limit</b>

ring	module			
Layer-3 queue	Layer-3 processor system or interface buffers	FIFO, flow-based weighted fair queuing (WFQ), class-based WFQ (CBWFQ), or low latency queuing (LLQ)	Yes	Varies with queuing method: – <b>vc-hold-queue</b> – <b>queue-limit</b>

In addition, this per-VC hold queue changes what happens to packets when a VC's transmit credit is full.

Packets on both the fast path and process path benefit from the per-VC hold queue.

- Process-switched packets – When the transmit credit is full, the packets are dequeued from the interface's hold-queue and then re-queued again in a per-VC hold queue.
- Fast-switched packets – When the transmit credit is full, the ATM driver simply suspends transmission for this VC until the hardware queues return one or more transmit credits for the VC. The packet is placed on the appropriate per-VC hold queue.

Since the router builds a hold queue per VC, the PA-A3 does not need to be concerned with avoiding head-of-line blocking and simply queues the excess packets instead of aggressively dropping them.

## hold-queue and vc-hold-queue Commands

It is important to understand the difference between the following two commands when tuning your router for optimal performance:

- **hold-queue {value} out** – Tunes the interface-level output hold queue on interfaces without per-VC queuing. You effectively can avoid using the **hold-queue {value} output** command when per-VC queuing is configured (via service policies on the 7500 series or by default starting with certain Cisco IOS software releases on non-RSP platforms).
- **vc-hold-queue {value}** – Tunes the per-VC output hold queue on interfaces with per-VC queuing. This command applies only to non-RSP platforms and only to VCs which are using the default FIFO queuing mechanism on packets inside the per-VC hold queue. The **vc-hold-queue** command determines how many packets that a VC can buffer after reaching its transmit credit. The **vc-hold-queue** command effectively set a limit on the aggregate size of the layer-3 queue. Use the **queue-limit** command to configure the size of the layer-3 queue per class.

**Note:** The **vc-hold-queue** command cannot be configured in the ATM PVC bundle CLI (CSCdw29901).

Removing the **vc-hold-queue** configuration under an ATM PVC using the **default vc-hold-queue** or **no vc-hold-queue** command still requires a value. The value can either be the current value or any number. This issue is a cosmetic issue and does not affect the performance of the router. It is resolved via Cisco bug ID CSCdx04931.

---

## Related Information

- [Configuring Per-VC Queuing and IP to ATM CoS on the P A-A3 and FlexWAN](#)
  - [Understanding Per-VC Queuing Options on the 4xOC3 ATM Line Card](#)
  - [Technical Support – Cisco Systems](#)
  - [Tools and Utilities – Cisco Systems](#)
  - [More ATM Information](#)
- 

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

---

Updated: Dec 14, 2007

Document ID: 6187

---