

# Understanding and Tuning the tx-ring-limit Value

[TAC Notice: What's Changing on TAC Web](#)

## Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Conventions](#)

[Understanding Particles](#)

[Understanding Buffer Rings](#)

[PA-A3 Architecture Overview](#)

[Transmit Ring Allocation Scheme on the PA-A3](#)

[Displaying the Current Transmit Ring Values](#)

[When Should the Transmit Ring Be Tuned?](#)

[Impact of Very Small tx-ring-limit Values](#)

[Known Issues](#)

[Tuning the tx-ring-limit on 3600 and 2600 Routers](#)

[Related Information](#)

Help us help you.

Please rate this document.

Excellent

Good

Average

Fair

Poor

This document solved my problem.

Yes

No

Just browsing

Suggestions for improvement:

(256 character limit)

Send

## Introduction

This document discusses the function of a hardware transmit ring and the purpose of the **tx-ring-limit** command on ATM router interface hardware that supports per-virtual circuit (VC) queueing.

Cisco router interfaces configured with service policies store packets for an ATM VC in one of two sets of queues depending on the congestion level of the VC:

Queue	Location	Queueing Methods	Service Policies Apply	Command to Tune
Hardware queue or transmit ring	Port adapter or network module	FIFO only	No	<b>tx-ring-limit</b>
Layer-3 queue	Layer-3 processor system or	N/A	Yes	Varies with queueing method: <b>- vc-hold-</b>



## Prerequisites

### Requirements

There are no specific requirements for this document.

### Components Used

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

### Conventions

Refer to [Cisco Technical Tips Conventions](#) for more information on document conventions.

## Understanding Particles

Before discussing the transmit ring, we first need to understand what a particle is. A particle forms the basic building block of packet buffering on many platforms, including the Cisco 7200 router series and the versatile interface processor (VIP) on the Cisco 7500 router series.

Depending on the packet length, Cisco IOS® software uses one or more particles to store a packet. Let's look at an example. When receiving a 1200-byte packet, IOS retrieves the next free particle and copies the packet data into the particle. When the first particle is filled, IOS moves to the next free particle, links it to the first particle, and continues copying the data into this second particle. Upon completion, the 1200 bytes of the packet are stored in three discontinuous pieces of memory that IOS logically makes a part of a single packet buffers.

IOS particle size varies from platform to platform. All particles within a given pool are the same size. This uniformity simplifies the particle management algorithms and helps contribute to efficient memory use.

## Understanding Buffer Rings

Along with public and private interface pools, Cisco IOS creates special buffer control structures called rings. Cisco IOS and interface controllers use these rings to control which buffers are used to receive and transmit packets to the media. The rings themselves consist of media-controller-specific elements that point to individual packet buffers elsewhere in I/O memory.

Each interface has a pair of rings - a receive ring for receiving packets and a transmit ring for transmitting packets. The size of the rings can vary with the interface controller. In general, the size of the transmit ring is based on bandwidth of the interface or VC and is a power of two (Cisco Bug ID CSCdk17210).

Interface	Rings					
Line Rate(Mb/s) <	2	10	20	30	40	...

<b>txcount</b>	2	4	8	16	32	64
----------------	---	---	---	----	----	----

**Note:** On the 7200 series platform, the transmit ring packet buffers come from the receive ring of the originating interface for a switched packet or from a public pool if the packet was originated by IOS. They are deallocated from the transmit ring and returned to their original pool after the payload data is transmitted.

## PA-A3 Architecture Overview

To ensure high forwarding performance, the PA-A3 port adapter uses separate receive and transmit segmentation and reassembly (SAR) chips. Each SAR is supported by its own subsystem of onboard memory to store packets as well as key data structures like the VC table. This memory specifically includes 4 MB of SDRAM, which is chunked into particles.

The following table illustrates the number and size of particles on the receive and transmit paths on the PA-A3.

Ring	Particle Size	Number of Particles
Receive Ring	288 bytes	n/a
Transmit Ring	576* bytes	6000 (144 particles are reserved)

\* The transmit ring's particle size also is described as being 580 bytes. This value includes the 4-byte ATM core header that travels with the packet inside the router.

The sizes in the above table were selected because they are divisible by 48 (the size of a cell's payload field) and by the cache line size (32 bytes) for maximum performance. They are designed to prevent the SAR from introducing inter-buffer delay when a packet requires multiple buffers. The transmit particle size of 576 bytes also was selected to cover about 90 percent of Internet packets.

## Transmit Ring Allocation Scheme on the PA-A3

The PA-A3 driver assigns a default transmit-ring value to each VC. This value varies with the ATM service category assigned to the VC. The following table lists the default values.

VC Service Category	PA-A3-OC3, T3, E3 Default Transmit Ring Value	PA-A3-IMA Default Transmit Ring Value	PA-A3-OC12 Default Transmit Ring Value	Time of Enforcement
	Based on formula**: $(48 \times \text{SCR}) / (\text{Particle\_size} \times 5)$ Minimum value is 40,	Based on formula: $(48 \times \text{SCR}) / (\text{Particle\_size} \times 5)$ Minimum value is 40,	Based on the following formula: $\text{Average rate (SCR)} * 2 * \text{TOTAL\_CREDITS} / \text{VISIBLE\_BANDWIDTH}$	

VBR-nrt	and overrides any calculated value less than 40 with a very low SCR.  <b>Note:</b> SCR is the cell rate with ATM overhead included.	and overrides any calculated value less than 40 with a very low SCR.  <b>Note:</b> SCR is the cell rate with ATM overhead included.	TOTAL_CREDITS = 8192  VISIBLE_BANDWIDTH = 599040 Note: If this formula calculates a value which is less than the default of 128, then the VC's transmit ring limit is set to 128.	Always
ABR	128	128	N/A	Always*
UBR	40	128	128	Only when total credit utilization exceeds 75 percent or the tx_threshold value, as shown in <b>show controller atm</b> .

\* Originally, the PA-A3-OC12 did not implement always-active limiting of VBR-nrt PVCs to the current transmit ring value. Bug ID CSCdx11084 resolves this issue. .

\*\* SCR should be expressed in cells/sec.

## Displaying the Current Transmit Ring Values

Originally, the value of the transmit ring was only visible via a hidden command. The **show atm vc {vcd}** command now displays the current value.

You also can use the **debug atm events** command to view the VC setup messages between the PA-A3 driver and the host CPU. The following sets of output were captured on a PA-A3 in a 7200 series router. The transmit ring value is displayed as the tx\_limit value, which implements the particle buffer quota allocated for a specific VC in the transmit direction.

PVC 1/100 is configured as VBR-nrt. Based on an SCR of 3500 kbps, the PA-A3 assigns a tx\_limit of 137. To see how this calculation is made, we need to convert an SCR of 3500 kbps to cells/sec. Notice that  $(3,500,000 \text{ bits /sec}) * (1 \text{ byte} / 8 \text{ bits}) * (1 \text{ cell} / 53 \text{ bytes}) = (3,500,000 \text{ cells}) / (8 * 53 \text{ sec}) = 8254 \text{ cells / sec}$ . Once we have the SCR value in cells / sec, we can apply the formula above to get tx\_limit = 137.

```
7200-17(config)#interface atm 4/0
7200-17(config-if)#pvc 1/100
7200-17(config-if-atm-vc)#vbr-nrt 4000 3500 94
7200-17(config-if-atm-vc)#
*Oct 14 17:56:06.886: Reserved bw for 1/100 Available bw = 141500
7200-17(config-if-atm-vc)#exit
```

```

7200-17(config-if)#logging
*Oct 14 17:56:16.370: atmdx_setup_vc(ATM4/0): vc:6 vpi:1 vci:100 state:2 con
*Oct 14 17:56:16.370: atmdx_setup_cos(ATM4/0): vc:6 wred_name:- max_q:0
*Oct 14 17:56:16.370: atmdx_pas_vc_setup(ATM4/0): vcd 6, atm_hdr 0x00100640,
*Oct 14 17:56:16.370: VBR: pcr 9433, scr 8254, mbs 94
*Oct 14 17:56:16.370:   vc tx_limit=137, rx_limit=47
*Oct 14 17:56:16.374:   Created 64-bit VC count

```

PVC 1/101 is configured as ABR. The PA-A3 assigns the default ABR tx\_limit value of 128. (See the table [above](#).)

```

7200-17(config-if)#pvc 1/102
7200-17(config-if-atm-vc)#abr ?
  <1-155000>      Peak Cell Rate(PCR) in Kbps
  rate-factors   Specify rate increase and rate decrease factors (inverse)
7200-17(config-if-atm-vc)#abr 4000 1000
7200-17(config-if-atm-vc)#
*Oct 14 17:57:45.066:   Reserved bw for 1/102 Available bw = 140500
*Oct 14 18:00:11.662: atmdx_setup_vc(ATM4/0): vc:8 vpi:1 vci:102 state:2 con
*Oct 14 18:00:11.662: atmdx_setup_cos(ATM4/0): vc:8 wred_name:- max_q:0
*Oct 14 18:00:11.662: atmdx_pas_vc_setup(ATM4/0): vcd 8, atm_hdr 0x00100660,
*Oct 14 18:00:11.662: ABR: pcr 9433, mcr 2358, icr 9433
*Oct 14 18:00:11.662:   vc tx_limit=128, rx_limit=47
*Oct 14 18:00:11.666:   Created 64-bit VC counters

```

PVC 1/102 is configured as UBR. The PA-A3 assigns the default UBR tx\_limit value of 40. (See the [table](#) above.)

```

7200-17(config-if)#pvc 1/101
7200-17(config-if-atm-vc)#ubr 10000
7200-17(config-if-atm-vc)#
*Oct 14 17:56:49.466:   Reserved bw for 1/101 Available bw = 141500
*Oct 14 17:57:03.734: atmdx_setup_vc(ATM4/0): vc:7 vpi:1 vci:101 state:2 con
*Oct 14 17:57:03.734: atmdx_setup_cos(ATM4/0): vc:7 wred_name:- max_q:0
*Oct 14 17:57:03.734: atmdx_pas_vc_setup(ATM4/0): vcd 7, atm_hdr 0x00100650,
*Oct 14 17:57:03.734: UBR: pcr 23584
*Oct 14 17:57:03.734:   vc tx_limit=40, rx_limit=117
*Oct 14 17:57:03.738:   Created 64-bit VC counters

```

The purpose of the tx\_limit is to implement a per-VC transmit credit or memory allocation scheme that prevents any consistently oversubscribed VC from grabbing all the packet buffer resources and hindering other VCs from transmitting normal traffic within their traffic contracts.

The PA-A3 implements a memory credit check under two conditions:

- Individual quota on each VBR-nrt and ABR VC - Compares each VC's tx\_count and tx\_limit values. It discards subsequent packets when the tx\_count is greater than the tx\_limit on any one VC. It is important to note that a burst of packets can exceed the transmit ring of a VBR-nrt VC at an instant in time and lead to output drops.
- Overall quota - Considers the tx\_threshold value. The PA-A3 allows for larger bursts on UBR VCs by enforcing traffic policing on such VCs only when the total packet buffer usage on the PA-A3 reaches this preset threshold.

**Note:** If a packet requires multiple particles and the transmit ring is full, the PA-A3 allows a VC to exceed its quota if particles are available. This scheme is designed to accommodate a small burst of packets without output drops.

The **show controller atm** command displays several counters relevant to transmit credits.

```

7200-17#show controller atm 4/0
  Interface ATM4/0 is up
  Hardware is ENHANCED ATM PA - OC3 (155000Kbps)
  Framer is PMC PM5346 S/UNI-155-LITE, SAR is LSI ATMIZER II
  Firmware rev: G125, Framer rev: 0, ATMIZER II rev: 3
    idb=0x622105EC, ds=0x62217DE0, vc=0x62246A00
    slot 4, unit 9, subunit 0, fci_type 0x0059, ticks 190386
    1200 rx buffers: size=512, encap=64, trailer=28, magic=4
  Curr Stats:
    VCC count: current=7, peak=7
    SAR crashes: Rx SAR=0, Tx SAR=0
    rx_cell_lost=0, rx_no_buffer=0, rx_crc_10=0
    rx_cell_len=0, rx_no_vcd=0, rx_cell_throttle=0, tx_aci_err=0
  Rx Free Ring status:
    base=0x3E26E040, size=2048, write=176
  Rx Compl Ring status:
    base=0x7B162E60, size=2048, read=1200
  Tx Ring status:
    base=0x3E713540, size=8192, write=2157
  Tx Compl Ring status:
    base=0x4B166EA0, size=4096, read=1078
  BFD Cache status:
    base=0x62240980, size=6144, read=6142
  Rx Cache status:
    base=0x62237E80, size=16, write=0
  Tx Shadow status:
    base=0x62238900, size=8192, read=2143, write=2157
  Control data:
    rx_max_spins=3, max_tx_count=17, tx_count=14
    rx_threshold=800, rx_count=0, tx_threshold=4608
    tx_bfd_write_indx=0x4, rx_pool_info=0x62237F20

```

The following table describes the values used by the PA-A3 to enforce the overall transmit credit scheme:

Value	Description
max_tx_count	Histogram of the maximum number of transmit particles ever held by the PA-A3 microcode.
tx_count	Total number of transmit particles currently being held by the PA-A3 microcode.  <b>Note:</b> The PA-A3 microcode also tracks the tx_count of each VC. When a particle is sent to the PA-A3 microcode from the PA-A3 driver, the tx_count increments by one.
tx_threshold	When the total amount of free packet buffers falls below this threshold, the PA-A3 enforces the transmit credit on UBR VCs. Note that the PA-A3 always enforces the transmit credits of VBR and ABR VCs.

## When Should the Transmit Ring Be Tuned?

The transmit ring serves as a staging area for packets in line to be transmitted. The router needs to enqueue a sufficient number of packets on the transmit ring and ensure that the interface driver has packets with which to fill available cell timeslots.

Originally, the PA-A3 driver did not adjust the transmit ring size when a service policy with low latency queueing (LLQ) was applied. With current images, the PA-A3 tunes down the value from the above defaults (Cisco Bug ID CSCds63407) to minimize queueing-related delay.

The primary reason to tune the transmit ring is to reduce latency caused by queueing. When tuning the transmit ring, consider the following:

- On any network interface, queueing forces a choice between latency and the amount of burst that the interface can sustain. Larger queue sizes sustain longer bursts while increasing delay. Tune the size of a queue when you feel the VC's traffic is experiencing unnecessary delay.
- Consider the packet size. Configure a **tx-ring-limit** value that accommodates four packets. For example, if your packets are 1500 bytes, set a tx-ring-limit value of 16 = (4 packets) \* (4 particles).
- Ensure the transmit credit is large enough to support one MTU-sized packet and/or the number of cells equal to the maximum burst size (MBS) for a VBR-nrt PVC.
- Configure a low value with low-bandwidth VCs, such as a 128 kbps SCR. For example, on a low-speed VC with an SCR of 160 kbps, a tx-ring-limit of ten is relatively high and can lead to significant latency (for example, hundreds of milliseconds) in the driver-level queue. Tune the tx-ring-limit down to its minimum value in this configuration.
- Configure higher values for high-speed VCs. Selecting a value of less than four may inhibit the VC from transmitting at its configured rate if the PA-A3 implements back pressure too aggressively and the transmit ring does not have a ready supply of packets waiting to be transmitted. Ensure that a low value does not affect VC throughput. (See Cisco Bug ID CSCdk17210.)

In other words, the size of the transmit ring needs to be small enough to avoid introducing latency due to queueing, and it needs to be large enough to avoid drops and a resulting impact to TCP-based flows.

An interface first removes the packets from the layer-3 queueing system and then queues them on the transmit ring. Service policies apply only to packets in the layer-3 queues and are transparent to the transmit ring.

Queueing on the transmit ring introduces a serialization delay that is directly proportional to the depth of the ring. An excessive serialization delay can impact latency budgets for delay-sensitive applications such as voice. Thus, Cisco recommends reducing the size of the transmit ring for VCs carrying voice. Select a value based on the amount of amount of serialization delay, expressed in seconds, introduced by the transmit ring. Use the following formula:

$$((P*8)*D)/S$$

P = Packet size in bytes. Multiply by eight to convert to bits.

D = Transmit-ring depth.

S = Speed of the VC in bps.

**Note:** IP packets on the Internet are typically one of three sizes: 64 bytes (for example, control messages), 1500 bytes (for example, file transfers), or 256 bytes (all other traffic). These values produce a typical overall Internet packet size of 250 bytes.

**Note:** The following table summarizes the advantages and disadvantages of larger or smaller transmit ring sizes:

Size of Transmit Ring	Advantage	Disadvantage
High Value	Recommended for data VCs to accommodate bursts.	Not recommended for voice VCs. Can introduce increased latency and jitter.
Low Value	Recommended for voice VCs to reduce delay due to queueing and jitter.	Not recommended for relatively high-speed VCs. Can introduce reduced throughput if tuned to such a low value that no packets are ready to be sent once the wire is free.

Use the **tx-ring-limit** command in VC configuration mode to tune the size of the transmit ring.

```

7200-1(config-subif)#pvc 2/2
  7200-1(config-if-atm-vc)#?
    ATM virtual circuit configuration commands:
abr          Enter Available Bit Rate (pcr)(mcr)
broadcast    Pseudo-broadcast
class-vc     Configure default vc-class name
default      Set a command to its defaults
encapsulation Select ATM Encapsulation for VC
exit-vc      Exit from ATM VC configuration mode
ilmi         Configure ILMI management
inarp        Change the inverse arp timer on the PVC
no           Negate a command or set its defaults
oam          Configure oam parameters
oam-pvc      Send oam cells on this pvc
protocol     Map an upper layer protocol to this connection.
random-detect Configure WRED
service-policy Attach a policy-map to a VC
transmit-priority set the transmit priority for this VC
tx-ring-limit Configure PA level transmit ring limit
ubr          Enter Unspecified Peak Cell Rate (pcr) in Kbps.
vbr-nrt      Enter Variable Bit Rate (pcr)(scr)(bcs)
7200-1(config-if-atm-vc)#tx-ring-limit ?
<3-6000>  Number (ring limit)
<cr>

```

Use the **show atm vc** command to display the currently configured value.

```

7200-1#show atm vc
VC 3 doesn't exist on interface ATM3/0
ATM5/0.2: VCD: 3, VPI: 2, VCI: 2
VBR-NRT, PeakRate: 30000, Average Rate: 20000, Burst Cells: 94
AAL5-LLC/SNAP, etype:0x0, Flags: 0x20, VCmode: 0x0

```

```
OAM frequency: 0 second(s)
PA TxRingLimit: 10
InARP frequency: 15 minutes(s)
Transmit priority 2
InPkts: 0, OutPkts: 0, InBytes: 0, OutBytes: 0
InPRoc: 0, OutPRoc: 0
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
InPktDrops: 0, OutPktDrops: 0
CrcErrors: 0, SarTimeOuts: 0, OverSizedSDUs: 0
OAM cells received: 0
OAM cells sent: 0
Status: UP
```

In addition, use the show **atm pvc vpi/vci** command to view both the current transmit and receive ring limits. The following output was captured on a 7200 Series router running Cisco IOS Software Release 12.2(10).

```
vikingshow    atm pvc 1/101
  ATM6/0: VCD: 2, VPI: 1, VCI: 101
UBR, PeakRate: 149760
AAL5-LLC/SNAP, etype:0x0, Flags: 0xC20, VCmode: 0x0
OAM frequency: 0 second(s), OAM retry frequency: 1 second(s), OAM retry
frequency: 1 second(s)
OAM up retry count: 3, OAM down retry count: 5
OAM Loopback status: OAM Disabled
OAM VC state: Not Managed
ILMI VC state: Not Managed
VC TxRingLimit: 40 particles
VC Rx Limit: 800 particles
```

## Impact of Very Small tx-ring-limit Values

On the transmit path, the host CPU transfers the payload from the host buffers to the local particle buffers on the PA-A3. The firmware running on the PA-A3 caches several buffer descriptors and frees them in a group. During the caching period, the PA-A3 does not accept new packets even though the contents of the local memory have been transmitted on the physical wire. The purpose of this scheme is to optimize overall performance. Thus, when configuring a non-default tx-ring-limit value, consider the buffer descriptor return delay.

In addition, if you configure a **tx-ring-limit** value of one with given a particle size of 576 bytes, a 1500-byte packet is removed from the queue as follows:

1. The PA-A3 driver queues the first particle in the transmit ring, and remembers that this packet is stored in two other memory particles.
2. During the next time that the transmit ring is empty, the second particle of the packet is put in the transmit ring.
3. During the next time that the transmit ring is empty again, the third particle is put in the transmit ring.

Even though the transmit ring consists of only one 576 byte particle, MTU/port-speed is still the worst-case latency through the transmit ring.

## Known Issues

When the **tx-ring-limit** command is applied to a VC through a vc-class statement, the PA-A3 does

not apply the configured value. Confirm this result by displaying the current value in the **show atm vc detail** command. Tuning the transmit ring using a vc-class was implemented in Cisco IOS Software Release 12.1 (Cisco Bug ID CSCdm93064). CSCdv59010 resolves a problem with the tx-ring-limit in certain versions of Cisco IOS Software Release 12.2. When you apply the **tx-ring-limit** command through the vc-class statement to an ATM PVC, the transmit ring size is not modified. Confirm this result using the **show atm vc detail** command, after applying the command through the vc-class and class-vc command pairs.

When added to a PVC on a PA-A3 in a Cisco 7200 series router running Cisco IOS Software Release 12.2(1), the **tx-ring-limit** command is duplicated, as shown below (Cisco Bug ID CSCdu19350).

```
interface ATM1/0.1 point-to-point
description dlci-101, cr3640
ip unnumbered Loopback0
pvc 0/101
tx-ring-limit 3
tx-ring-limit 3
```

The condition is harmless and does not affect the operation of the router.

Cisco bug ID CSCdv71623 resolves a problem with output drops on a multilink PPP bundle interface when the traffic rate is well below the line rate. This problem was seen in CSCdv89201 on an ATM interface with a tx-ring-limit value greater than five. The problem becomes particularly apparent when fragmentation is disabled or when the link weights (fragment size limits) are large -- common on higher speed links like T1s or E1s -- and the data traffic consists of a mix of small and large packets. Enabling fragmentation and using a small fragment size (set by the interface configuration command **ppp multilink fragment delay**) improves operation significantly. However, you should verify that your router has sufficient processing capacity to support these high levels of fragmentation without overloading the system CPU, before using this as a workaround.

Cisco bug ID CSCdw29890 resolves a problem with the tx-ring-limit command being accepted by the CLI for ATM PVC bundles, but not taking effect. However, you do not normally need to change the **tx-ring-limit** on ATM PVC bundles. The reason is that, reducing the ring size effectively moves all the transmit buffering to a QoS-controlled queue, so an arriving priority packet is transmitted immediately to minimize delay on low-speed interfaces. With ATM PVC bundles, cells from packets of all the member VCs are always sent simultaneously (and interleaved), so the delay is minimized automatically.

## Tuning the tx-ring-limit on 3600 and 2600 Routers

Current Cisco IOS software images support tuning the transmit ring on the ATM network modules for Cisco 2600 and 3600 series routers (Cisco Bug ID CSCdt73385). The current value appears in the **show atm vc** output.

---

## Related Information

- [More ATM Information](#)
- [Tools and Resources - Cisco Systems](#)
- [Technical Support & Documentation - Cisco Systems](#)

<a href="#">Home</a>	<a href="#">How to Buy</a>	<a href="#">Login</a>	<a href="#">Profile</a>	<a href="#">Feedback</a>	<a href="#">Site Map</a>	<a href="#">Help</a>
----------------------	----------------------------	-----------------------	-------------------------	--------------------------	--------------------------	----------------------

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

© 2007 - 2008 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)