

Troubleshooting Output Drops on ATM Router Interfaces

Document ID: 10416

Introduction

Prerequisites

Requirements

Components Used

Conventions

Traditional Reasons for Output Drops

ATM-Specific Reasons for Output Queue Drops

Layer 3 Per-VC Queues

Understand Different Drop Counters

Troubleshoot

Tuning Queue Sizes

Output Drop Counters

Known Issue: VC Appears Stuck

Related Information

Introduction

This document provides the information you need to understand and troubleshoot output drops on ATM interfaces.

Prerequisites

Requirements

Readers of this document should have knowledge of these topics:

You can use the **show interface** command on any Cisco router interface to see several important values:

- Input and output rate in bits per second and packets per second (five minutes is the default period).
- Input and output queue size and the number of drops.
- Input error counters such as cyclic redundancy checks (CRCs), ignores, and no buffers.

In this output, an enhanced ATM port adapter (PA-A3) has experienced 11,184 output queue drops since the counters were last cleared on week and one day ago:

```
router#show interface atm 5/0/0
  ATM5/0/0 is up, line protocol is up
  Hardware is cyBus ENHANCED ATM PA
  MTU 4470 bytes, sub MTU 4470, BW 149760 Kbit, DLY 80 usec, rely 255/255,
  load 2/255
  Encapsulation ATM, loopback not set, keepalive set (10 sec)
  Encapsulation(s): AAL5 AAL3/4
  4096 maximum active VCs, 7 current VCCs
  VC idle disconnect time: 300 seconds
  Last input never, output 00:00:00, output hang never
```

```
  Last clearing of "show interface" counters 1w1d
  Queueing strategy: fifo
```

```
Output queue 0/40, 11184 drops; input queue 0/150, 675 drops
5 minute input rate 1854000 bits/sec, 382 packets/sec
5 minute output rate 1368000 bits/sec, 376 packets/sec
155080012 packets input, 3430455270 bytes, 0 no buffer
Received 0 broadcasts, 0 runts, 0 giants
313 input errors, 313 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
157107224 packets output, 1159429109 bytes, 0 underruns
0 output errors, 0 collisions, 0 interface resets
0 output buffers copied, 0 interrupts, 0 failures
```

On an ATM interface, the output of the **show interface atm** command sometimes displays a large number of output queue drops. All types of router interfaces, from serial to Ethernet, may experience output queue drops. This is due to the amount of traffic or the method in which the router switches packets from the ingress (incoming interface) to the egress (exiting interface). ATM interfaces also experience output drops due to ATM-layer traffic shaping on a virtual circuit.

Components Used

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

Conventions

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

Traditional Reasons for Output Drops

For information on traditional reasons for output drops, refer to Troubleshooting Input Queue Drops and Output Queue Drops

ATM-Specific Reasons for Output Queue Drops

On ATM interfaces, output drops can be interpreted as something other than buffer exhaustion for the interface.

Note: Any interface that is overdriven (that is, when the offered rate is greater than the line rate) presents output drops.

ATM interfaces typically use ATM-layer traffic shaping to limit the maximum amount of bandwidth used by a virtual connection. If you present more traffic to the virtual circuit (VC) than it is configured to transmit, the ATM interface tries to store the packet until it can be scheduled for transmission. However, the interface may need to drop some packets. This can particularly happen if you burst above the traffic-shaping parameters for a period of time longer than the virtual circuit is configured to handle. Traffic shaping is often implemented as part of a traffic contract with the circuit provider.

The ATM Forum defines five ATM service categories in its Traffic Management Specification Version 4.0. Each of these service categories supports a unique set of traffic parameters that may include peak cell rate (PCR), sustained cell rate (SCR), and maximum burst size (MBS):

- constant bit rate (CBR).
- variable bit rate – real time (VBR-rt).
- variable bit rate – non-real time (VBR-nrt).
- available bit rate (ABR).
- unspecified bit rate (UBR).

When you specify a peak cell rate, you can tell the ATM interface to shape the output rate and ensure that the bits per second rate for the VC does not exceed the maximum value.

If you configure a Permanent Virtual Circuit (PVC) and do not specify the PCR or SCR, you create a PVC of the UBR service class. This PVC is automatically assigned a PCR equal to the line rate of the interface. Here is an example:

```
router(config)#interface atm 3/0
router(config-if)#pvc 5/200
router(config-if-atm-vc)#end
router#sh atm pvc 5/200
ATM3/0: VCD: 5, VPI: 5, VCI: 200
UBR, PeakRate: 44209
AAL5-LLC/SNAP, etype:0x0, Flags: 0xC20, VCmode: 0x0, Encapsize: 12
OAM frequency: 0 second(s), OAM retry frequency: 1 second(s)
OAM up retry count: 3, OAM down retry count: 5
OAM Loopback status: OAM Disabled
...
```

Similarly, if you configure a PVC with the same value for PCR and SCR, you create a UBR PVC. However, by doing this, you also shape this VC and limit the PCR. Here is an example:

```
router(config)#interface atm 6/0
  7200-1(config-if)#atm pvc 300 5 300 aal5snap ?
    <1-45000>      Peak rate(Kbps)
    abr           Available Bit Rate
    inarp         Inverse ARP enable
    oam           OAM loopback enable
    random-detect WRED enable
    tx-ring-limit Configure PA level transmit ring limit
    <cr>
router(config-if)#atm pvc 300 5 300 aal5snap 10000 ?
<1-10000> Average rate(Kbps)

router(config-if)#atm pvc 300 5 300 aal5snap 10000 10000
router(config-if)#end

router#show atm pvc 5/300
ATM3/0: VCD: 300, VPI: 5, VCI: 300
UBR, PeakRate: 10000
AAL5-LLC/SNAP, etype:0x0, Flags: 0x820, VCmode: 0x0, Encapsize: 12
OAM frequency: 0 second(s), OAM retry frequency: 0 second(s)
OAM up retry count: 0, OAM down retry count: 0
OAM Loopback status: OAM Disabled
OAM VC Status: Not Managed
ILMI VC status: Not Managed
...
```

The most common ATM service class for transmitting data (as opposed to voice or video traffic) is VBR-nrt. ATM interfaces are only capable of forwarding a limited amount of traffic. This amount is based on your traffic shaping parameters (PCR, SCR and MBS). SCR is a long-term rate average. The PCR and SCR bits per second values count the bits of an entire cell. This includes both the five-byte ATM header and the cell payload. On the following PVC, we have configured a PCR of 384 kbps, an SCR of 269 kbps, and an MBS of 250 cells. MBS is the number of cells that you can send at PCR.

Note: There are certain limitations on the PCR and SCR values. For more information on these limitations, refer to additional configuration documents on Traffic Management.

MBS is a low number relative to the output rate. For example, if your SCR is 269 kbps, and has 250 cells of MBS with 53 bytes per cell, it equals only a fraction of a second at which you send at PCR.

```

router#show atm pvc 1/59
  ATM4/1/0.8: VCD: 8, VPI: 1, VCI: 59

  VBR-NRT, PeakRate: 384, Average Rate: 269, Burst Cells: 250
  AAL5-NLPID, etype:0x2, Flags: 0x21, VCmode: 0x0
  OAM frequency: 0 second(s)
  InARP DISABLED
  Transmit priority 2
  InPkts: 302868, OutPkts: 386988, InBytes: 32380573, OutBytes: 199648072
  InPRoc: 79259, OutPRoc: 90978
  InFast: 222241, OutFast: 1931, InAS: 1368, OutAS: 294079

  InPktDrops: 0, OutPktDrops: 355
  CrcErrors: 0, SarTimeOuts: 0, OverSizedSDUs: 0
  OAM cells received: 0
  OAM cells sent: 0
  Status: UP

```

If you present more outbound traffic to the PVC than it can handle (or is configured to shape), the router tries to use queuing and drops mechanisms such as Weighted Random Early Detection (WRED) or another Quality of Service (QoS) method, to minimize packet loss. Some of these must be configured explicitly.

To determine if you are exceeding the PCR and SCR values of the PVC, look for the OutPktDrops counter in the output of the **show atm vc {vcd#}** or **show atm pvc <vpi>/<vci>** commands. These commands are only available, per VC, on the PA-A3, PA-A6 and on Cisco 2600 and 3600 routers (DS3, E3, OC3 and IMA interfaces). Observe the five minute input rate and output rate displayed by the **show interface atm** command. The traffic shaper should start to drop packets when the average traffic volume reaches SCR.

Even though it can cause the router to drop packets, traffic shaping is beneficial for multiple reasons:

- Drops happen closer to the source of the traffic (on the user side instead of on the network side).
- User equipment can usually buffer some traffic and reduce the amount of packets dropped during bursts.
- The key reason is that the network (that is, the service provider) may drop cells indiscriminately in order to force compliance to a contract. These drops may affect multiple packets, whereas the router has the intelligence to apply optimal shaping. For more information on this, refer to Troubleshooting ATM PVCs in a WAN Environment.

Note: It is important to understand that an ATM interface on a router only drops packets and never drops cells on the transmit side. Traffic shaping causes output queues to back up and may lead to drops if the congestion state is sustained.

Layer 3 Per-VC Queues

On the PA-A3 and PA-A6, starting with Cisco IOS® Software Releases 11.1(22)CC and 12.0(3)T, a VIP2-50 and above builds a separate pool of buffers dedicated to the storage of packets for each VC. Each Layer 3 per-VC queue matches to a Layer 2 VC queue in the port adaptor. These two queues per VC ensure that a direct relationship exists between the outgoing ATM VC and the IP packets to be forwarded on that queue. When the PA per-VC queues become congested, they signal back pressure to the Layer 3 processor. The Layer 3 processor can then continue to buffer packets for that VC in the corresponding Layer 3 queue. Also, because the Layer 3 queues are accessible by the Layer 3 processor, a user can run advanced software scheduling and drop algorithms on those queues.

The number of buffers available for per-VC queuing on the VIP depends on the amount of static random-access memory (SRAM) (also known as MEMD) installed on the versatile interface processor (VIP). With 8 MB of SRAM on board, up to 1085 packets worth of buffers may be available to the IP to ATM Class of Service (CoS) feature for per-VC queuing. A per-VC queue only develops on the VIP for the ATM PVCs

on which there is temporary congestion. That is, there is more incoming IP traffic than the egress ATM shaping rate of the corresponding ATM PVC. This queue only remains on the VIP for the duration of the burst.

The VIP and the PA-A3/PA-A6 collaborate in these ways:

1. The port adaptor transmits ATM cells on each ATM PVC according to the ATM shaping rate.
2. The port adaptor maintains a per-VC first-in, first-out (FIFO) queue for each VC where it stores the packets waiting for transmission onto that VC.
3. Should this per-VC queue fill up, the port adaptor provides explicit back pressure to the VIP. This is so that the VIP only transmits packets for that VC to the PA when the PA has sufficient buffers available to store the packets. This ensures that the PA-A3 never needs to discard any packet, regardless of the level of congestion on the ATM VCs.
4. When the VIP has packets to transfer to the port adaptor but is throttled by the port adaptor back pressure, the VIP stores the packets into per-VC queues. That is, one logical queue for each ATM PVC configured on the ATM interface. The per-VC queue is a FIFO queue that stores all packets, in order of arrival, that are to be transmitted onto the corresponding VC. For more information, go to Detailed IP to ATM CoS Phase 1 Operations.

The VIP then monitors the level of congestion independently on each of its per-VC queues. If it is also configured, it performs a WRED selective congestion avoidance algorithm independently on each of these queues that enforces service differentiation across the IP Classes of Service. For each instance of the per-VC WRED algorithm, the IP to ATM CoS feature computes a separate moving average queue occupancy (expressed in number of packets and takes into account packets of all precedences). It also supports a separate set of configurable WRED drop profiles with one profile per precedence.

In summary, the ATM layer functions such as ATM shaping are handled by the PA-A3, while the IP-level service differentiation is performed by the VIP. Through explicit back pressure from the PA to the VIP, the PA operates in a lossless environment and all congestion management and selective drops are performed on the VIP.

The drops shown in the output of the **show interface** command includes VC drops, due to traffic shaping and lack of buffers. It is not necessary that the sum of VC drops matches to that of the interface. The output drops on the VC increases only when the packets are dropped by the driver. There can be two reasons behind the large output drops on interface and not on the VCs:

- The packets are dropped from the output hold queue of the interface.
- The packets are dropped by the queuing mechanism on the Route Processor Module (RPM) itself before passing the traffic to the driver.

Starting with Cisco IOS Software Release 11.1(22)CC and 12.0(3)T, Cisco IOS builds a separate pool of buffers dedicated to the storage of packets for each VC in the Layer 3 processor system. Each Layer 3 per-VC queue matches to a Layer 2 VC queue on the ATM interface. When the ATM per-VC queues become congested, the ATM interface signals back pressure to the Layer 3 processor. The Layer 3 processor can then continue to buffer packets for that VC in the corresponding Layer 3 queue. Furthermore, because the Layer 3 queues are accessible by the Layer 3 processor, you can run flexible software scheduling algorithms on those queues.

When you configure IP to ATM CoS, you apply policies to a class of traffic. These use the class-based weighted fair queuing (CBWFQ) feature to define match-on traffic via access lists, matching input interfaces or protocols such as IP and IPX. One of these policies is the **queue-limit** command. This command specifies the maximum number of packets that can be placed in the class queue (that is, the number of packets that can be queued or waiting in the queue). This number varies depending on the type of queuing you have configured. For more information on CBWFQ, refer to Per-VC CBWFQ on Cisco 7200, 3600 and 2600

Routers.

With weighted fair queuing (WFQ), the default queue-limit is 64, as per the value specified for threshold. This is shown in this output:

```
core-1.msp#show queueing interface atm 2/0.100032
Interface ATM2/0.100032 VC 10/32
Queueing strategy: weighted fair
Total output drops per VC: 1539
Output queue: 0/512/64/1539 (size/max total/threshold/drops)
Conversations 0/37/128 (active/max active/max total)
Reserved Conversations 0/0 (allocated/max allocated)
```

The **queue-limit** command takes a number of packets from 1 to 64 as its argument.

With FIFO, the queue limit is 40, as shown in this output:

```
core-1.msp#show queueing interface atm 2/0.100032
Interface ATM2/0.100032 VC 10/32
Queueing strategy: FIFO
Output queue 0/40, 244 drops per VC
```

A new feature called Configurable per-VC Hold Queue Support lets you significantly increase the FIFO queue limit up to 1024 packets. The command to change the FIFO hold queue is **vc-hold-queue** in global configuration mode. This command was introduced in Cisco IOS Software Release 12.1(5)T. For more information, see Configurable per-VC Hold Queue Support for ATM Adapters.

You can enable flow-based WFQ using the **fair-queue** command. The **fair-queue** command also takes an argument that specifies the number of hashed queues for the class-default default class. The **queue-limit** command specifies the maximum number of packets each of these queues can hold. After this, any further enqueued packets are subject to tail drop. The router uses tail drop or (if you have configured it) WRED to manage the queue when packets are enqueued to it that exceed the configured limit.

In this example, a policy map is configured with a class class-default default class. The **fair-queue 32** command reserves 32 hashed queues that are created as traffic traverses the interface. WFQ queues are based on Layer 3 and Layer 4 header information. A queue limit of 20 is also configured. This command means that each hashed queue can hold 20 packets. When the 21st packet arrives, the router drops it using either tail drop or WRED as the dropping decision mechanism. This means that 20 packets accumulate in the queue reserved for the class before the tail drop or WRED packet drop is enacted.

```
class class-default
    fair-queue 32
    queue-limit 20
```

You can see in this output that there are 65 packets in the output queue. The threshold per conversation is 64. Conversation number 15 reaches a maximum of 64. On conversation number eleven, there have been 1,505,776 drops due to discards. This is the total number of drops for this queue. Trail drops counts the number of drops from this queue only when another queue has an incoming packet with a lower WFQ sequence number and the WFQ system reaches the max-queue-limit number of packets.

```
router2#show queue atm 4/0.102
Interface ATM4/0.102 VC 0/102
Queueing strategy: weighted fair
Total output drops per VC: 1505772

Output queue: 65/512/64/1505772 (size/max total/threshold/drops)
Conversations 2/3/16 (active/max active/max total)
```

```

Reserved Conversations 0/0 (allocated/max allocated)
(depth/weight/discards/tail drops/interleaves) 1/32384/0/0/0

Conversation 2, linktype: ip, length: 58
source: 8.0.0.1, destination: 6.6.6.6, id: 0x2DA1, ttl: 254, prot: 1
(depth/weight/discards/tail drops/interleaves) 64/32384/1505776/0/0
Conversation 15, linktype: ip, length: 1494
source: 7.0.0.1, destination: 6.6.6.6, id: 0x0000, ttl: 63, prot: 255

```

In addition to the **queue-limit** command, you can also apply the **bandwidth** command to a service policy. The **bandwidth** statement is used only with CBWFQ to give a minimum guarantee in times of congestion. In times of non-congestion, the class is free to use as much bandwidth as is available on the VC, even up to the maximum value of the VC.

The equivalent command with low latency queueing is the **priority** command. The **priority** command provides both a maximum and a guarantee. In periods of congestion, the class is guaranteed a certain amount of bandwidth. At the same time, it also is limited to this bandwidth, and drops occur if more packets above the priority kbps value are presented to the VC via the class. In times of non-congestion, the class is free to use as much bandwidth as possible up to the maximum value of the VC.

More specifically, policing is used to drop packets in times of congestion when the bandwidth is exceeded. Policing is used to ensure the class's traffic does not go above its configured priority value in kbps. Because of policing, you do not need the **queue-limit** command to police or put a limit on the priority queue. When congestion occurs, traffic destined for the priority queue is metered to ensure that the bandwidth allocation configured for the class to which the traffic belongs is not exceeded.

Priority traffic metering has these qualities:

- It is similar to Committed Access Rate (CAR) limits, except that priority traffic metering is only performed under congestion conditions. When the device is not congested, the priority class traffic is allowed to exceed its allocated bandwidth. When the device is congested, the priority class traffic above the allocated bandwidth is discarded.
- It is performed on a per-packet basis, and tokens are replenished as packets are sent. If not enough tokens are available to send the packet, it is dropped.
- It restrains priority traffic to its allocated bandwidth to ensure that non-priority traffic, such as routing packets and other data, is not starved.
- With metering, the classes are policed and rate-limited individually. That is, they are each treated as separate flows with separate bandwidth allocations and constraints. This is still the case even though a single policy map might contain four priority classes, all of which are enqueued in a single priority queue.

On the PA-A3 on 7200 routers, queueing does not happen in the interface queue, and you should not display interface queue at all in the **show interface** command. The **hold-queue** command does not make any changes. The driver takes the packet directly from the per-VC queue. The locally-generated process switched packets are also queued directly on the per-VC queue. There is also back-pressure and congestion on a per-VC basis.

Most drivers drop the packet when there is congestion along the Cisco Express Forwarding (CEF) or fast switching path. The interface queue is only used for locally-generated packets. Only a few ATM drivers support fancy queueing, which does not scale.

By default, the FIFO queueing method is enabled on the interface. Execute the command **show queueing interface atm/imay** to see the per-VC queues and the drops due to per-VC queueing. Here is an example:

```

7200#show queueing interface atm 2/0.1
Interface ATM2/0.1 VC 1/100

```

```
Queueing strategy: FIFO
Output queue 0/40, 244 drops per VC
```

Compare the value in the output of **show queueing interface atm** with the number in the **show interface atm** output. Are these numbers the same? Is the **show interface** number higher? If it is higher, then the drops can be due to a high number of process-switched packets that are sent to the system buffers.

Optionally, to see the drops due to IP flows, you can enable WFQ or weighted fair queueing on the ATM interface. WFQ creates queues for IP flows, which are defined based on source and destination IP addresses and port numbers. For more information, refer to Per-VC Class-Based, Weighted Fair Queuing (Per-VC CBWFQ) on the Cisco 7200, 3600, and 2600 Routers. Configure this:

```
policy-map mypol
  class class-default
    fair-queue
  !
interface ATM2/0.130 point-to-point
ip address 14.0.0.2 255.0.0.0
no ip directed-broadcast
PVC 1/130
vbr-nrt 100000 75000 100
  service-policy output mypol
broadcast
encapsulation aal5mux ip
```

Once you have configured WFQ, the output of the **show queueing** command changes:

```
core-1.msp#show queueing interface atm 2/0.100032
Interface ATM2/0.100032 VC 10/32
Queueing strategy: weighted fair
Total output drops per VC: 1539
Output queue: 0/512/64/1539 (size/max total/threshold/drops)
Conversations 0/37/128 (active/max active/max total)
Reserved Conversations 0/0 (allocated/max allocated)
```

There are currently 65 packets in the output queue. The threshold per conversation is 64. Conversation 15 reaches its maximum at 64. On conversation 11, there have been 1,505,776 drops due to discards, which is the total number of drops for this queue. Trail drops counts the number of drops from this queue only when another queue has an incoming packet with a lower WFQ sequence number, and the WFQ system reaches the max-queue-limit number of packets.

```
router2#show queue atm 4/0.102
Interface ATM4/0.102 VC 0/102
Queueing strategy: weighted fair
Total output drops per VC: 1505772
Output queue: 65/512/64/1505772 (size/max total/threshold/drops)
Conversations 2/3/16 (active/max active/max total)
Reserved Conversations 0/0 (allocated/max allocated)
(depth/weight/discards/tail drops/interleaves) 1/32384/0/0/0
Conversation 2, linktype: ip, length: 58
source: 8.0.0.1, destination: 6.6.6.6, id: 0x2DA1, ttl: 254, prot: 1
(depth/weight/discards/tail drops/interleaves) 64/32384/1505776/0/0
Conversation 15, linktype: ip, length: 1494
source: 7.0.0.1, destination: 6.6.6.6, id: 0x0000, ttl: 63, prot: 255
```

Understand Different Drop Counters

The important point to understand about interfaces that run per-VC queueing is that drops are seen under the output of **show queueing interface atm** and not under the output of the **show atm vc vcd#** command.

Troubleshoot

Complete these steps if you are having a problem.

1. Determine the type of ATM router interface by looking at the description line in the **show interface atm** command.

| Part Number | Description in show interface | Per-VC |
|-------------|--------------------------------------|----------------|
| AIP | Output Hardware is cxBus ATM | Counters No |
| PA-A1 | Hardware is TI1570 ATM | No |
| PA-A2 | Hardware is ATM-CES | No |
| PA-A3 | Hardware is cyBus ENHANCED ATM PA | Yes |
| PA-A6 | Hardware is ENHANCED ATM PA Plus | Yes |

2. Consult the table in step 1 to determine whether your interface supports per-VC counters.

- ◆ If it does, use the **show atm vc {vcd#}** or **show atm pvc <vpi>/<vci>** command on all VCs configured for an interface or subinterface.
- ◆ Add up the OutPktDrops counters for all the VCs and compare this value with the number of output queue drops displayed in the **show interface atm** command. Are the two numbers nearly the same?

◇ If yes, then the output drops are due to traffic shaping at the ATM layer.

◇ If no, then the output drops are due to lack of buffer resources.

3. Determine whether the interface's buffers are full with the command **show controllers cbus** on a Cisco 7500 series router. Look for a txacc value at or near zero.

```
router#show controllers cbus
[snip]
slot5: VIP2 R5K, hw 2.00, SW 22.20, ccb 5800FF70, cmdq 480000A8, VPs      8192
software loaded from system
IOS (TM) VIP Software (SVIP-DW-M), Version 12.1(5), RELEASE      SOFTWARE (fc1)
ROM Monitor version 115.0
ATM5/0/0, applique is OC3 (155000Kbps)
  gfreeq 48000160, lfreeq 480001F0 (4544 bytes)
  rxlo 4, rxhi 305, rxcurr 305, maxrxcurr 305
  txq 48001A48, txacc 48001A4A (value 5), txlimit      203
```

4. Since **show controllers cbus** does not indicate per-VC statistics, use the command **show atm vc**, followed by the command **show atm vc {vcd#}** or **show atm pvc <vpi>/<vci>** to see per-VC drop counters.

```
router#show atm vc
ATM5/0/0.4      4      4      32 PVC AAL5-SNAP      1536 1536      32 ACTI
ATM5/0/0.6      6      4      34 PVC AAL5-SNAP      1024 1024      32 ACTI
ATM5/0/0.7      7      6      32 PVC AAL5-SNAP      1024 1024      32 ACTI
router#show atm vc 7
ATM5/0/0.7: VCD: 7, VPI: 6, VCI: 32, etype:0x0, AAL5 -
LLC/SNAP, Flags: 0x40030
PeakRate: 1024, Average Rate: 1024, Burst Cells: 32, VCmode: 0x0
OAM DISABLED, InARP DISABLED
InPkts: 31672500, OutPkts: 23342085, InBytes: 1592433047, OutBytes:
2557199223
InPRoc: 386157, OutPRoc: 9791, Broadcasts: 380352
```

```
InFast: 0, OutFast: 0, InAS: 31286343, OutAS: 22951942
```

```
InPktDrops: 3, OutPktDrops: 4476
```

```
CrcErrors: 308, SarTimeOuts: 0, OverSizedSDUs: 0
```

```
OAM F5 cells sent: 0, OAM cells received: 0
```

```
Status: ACTIVE
```

```
router# show atm pvc 6/32
```

```
ATM5/0/0.7: VCD: 7, VPI: 6, VCI: 32
```

```
...
```

```
InPkts: 31672500, OutPkts: 23342085, InBytes: 1592433047, OutBytes: 2557199223
```

```
InPRoc: 386157, OutPRoc: 9791, Broadcasts: 380352
```

```
InFast: 0, OutFast: 0, InAS: 31286343, OutAS: 22951942
```

```
InPktDrops: 3, OutPktDrops: 4476
```

```
...
```

5. If you use an ATM port adapter on a VIP, determine whether the distributed VIP memory resources are congested with the command **show controllers VIP <slot>tech-support**, where <slot> is the slot number where the ATM port adapter resides.

- ◆ Use a VIP2 with more SRAM. Determine the type of VIP and the amount of SRAM with the command **show diag {slot #}**. A VIP2-40 has 32 MB of dynamic random-access memory (DRAM) and 2 MB of SRAM that cannot be upgraded. The VIP2-50 names the VIP2 R5K controller.

```
Slot 5:
```

```
Physical slot 5, ~physical slot 0xA, logical slot 5, CBus 0
```

```
Microcode Status 0x4
```

```
Master Enable, LED, WCS Loaded
```

```
Board is analyzed
```

```
Pending I/O Status: None
```

```
EEPROM format version 1
```

```
VIP2 controller, HW rev 2.11, board revision C0
```

```
Serial number: 12313902 Part number: 73-1684-04
```

```
Test history: 0x00 RMA number: 00-00-00
```

```
Flags: cisco 7000 board; 7500 compatible
```

```
EEPROM contents (hex):
```

```
0x20: 01 15 02 0B 00 BB E5 2E 49 06 94 04 00 00 00 00
```

```
0x30: 60 00 00 01 00 00 00 00 00 00 00 00 00 00 00 00
```

```
Slot database information:
```

```
Flags: 0x4 Insertion time: 0x1484 (5w3d ago)
```

```
Controller Memory Size: 32 MBytes DRAM, 2048 KBytes SRAM
```

- ◆ Remove a port adapter in the other bay of a VIP. The amount of SRAM that the IP to ATM CoS feature can use for per-VC queueing over the PA-A3/PA-A6 depends on whether another PA is supported on the same VIP. A VIP with a PA-A3 in one slot and the other slot left empty ensures that all of the VIP's SRAM buffers can be used by the PA-A3.
6. If your data gathering suggests that you are exceeding your traffic shaping parameters, then try increasing the PCR, SCR and MBS parameters on the VCs that record the highest number of drops.

Closely monitor the VC and determine if the drops are decreasing. Be sure to adjust these parameters in concert with your provider. Unilaterally increasing the values may lead to policing by the ingress switch to the ATM cloud.

7. Try an ATM interface that supports per-VC queueing, particularly if you see that one congested VC impacts other, non-congested VCs
8. Implement traffic management methods like fancy queueing and WRED. For more information, see Quality of Service Solutions.

- ◆ The output of **show interface atm** and **show queuing** indicates the type of queueing configured on the interface. If you have not explicitly configured fancy queueing, the ATM interface uses FIFO by default. Only when the VC becomes congested can you see the packets getting queued inside FIFO.

```

router#show queueing interface atm 1/0
  Interface ATM1/0 VC 1/35
  Queueing strategy: FIFO
  Output queue 0/40, 5161815 drops per VC
  Interface ATM1/0 VC 2/33
  Queueing strategy: FIFO
  Output queue 0/40, 0 drops per VC

```

9. Ensure that you use the newer PA-A3 (Revision 2.0), which is more stable in terms of drops and input errors. Refer to this field notice for more information.

Tuning Queue Sizes

The queue-limit keyword under the class-default is used to limit the queue depth of the congesting traffic. You can use the **TX-ring-limit** command to reduce the PA FIFO queue.

Output Drop Counters

You can obtain the number of output drops on your ATM VCs via a Cisco IOS command or via Simple Network Management Protocol (SNMP) polling (planned for Cisco IOS Software Release 12.2).

Originally, images without IP to ATM CoS displayed output packet drops by the ATM interface driver in the output of the **show atm pvc** command. In these images, the ATM interface driver made a random drop decision when a VC's transmit ring filled.

Originally, images with IP to ATM CoS displayed output packet drops by the Layer 3 processor in the output of the **show queueing int atm** command. In these images, the ATM interface throttles the receipt of new packets from the Layer 3 processor system until it has available space on the transmit ring of the VC. Therefore, IP to ATM CoS moves the drop decision from a random, last-in/first-dropped decision in the transmit ring's FIFO queue to a differentiated decision based on IP-level service policies implemented by the Layer 3 processor.

As of Cisco IOS Software Releases 12.1(9), 12.2(2), and 12.2(3)T (Cisco bug ID CSCdt44794 (registered customers only)), the **show atm pvc** command displays OutPktDrops by both the driver and by the Layer 3 processor.

- Without Layer 3 queueing enabled – Value displays as "OutPktDrops: 0".
- With Layer 3 queueing enabled – Valued displays as "OutPktDrops: 0/0/0 (holdq/outputq/total)."

This sample output shows that you can continue to use the **show queueing int atm** command to display drops by the Layer 3 processor.

```

router#show atm pvc 501
  Switch1.501: VCD: 10, VPI: 0, VCI: 501
  VBR-NRT, PeakRate: 128, Average Rate: 128, Burst Cells:    94
  AAL5-LLC/SNAP, etype:0x0, Flags: 0x8000020, VCmode:    0x0
  OAM frequency: 0 second(s), OAM retry frequency: 1 second(s), OAM retry frequency: 1 se
  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
  PA TxRingLimit: 3
  Rx Limit: 100 percent
  InARP frequency: 15 minutes(s)
  Transmit priority 2
  InPkts: 0, OutPkts: 2878, InBytes: 0, OutBytes: 816840
  InPRoc: 0, OutPRoc: 0
  InFast: 0, OutFast: 2876, InAS: 0, OutAS: 0

```

```

InPktDrops: 0, OutPktDrops: 6483/0/6483 (holdq/outputq/total)
CrcErrors: 0, SarTimeOuts: 0, OverSizedSDUs: 0, LengthViolation: 0, CPiErrors: 0
Out CLP=1 Pkts: 0
OAM cells received: 0
F5 InEndloop: 0, F5 InSegloop: 0, F5 InAIS: 0, F5 InRDI: 0
F4 InEndloop: 0, F4 InSegloop: 0, F4 InAIS: 0, F4 InRDI: 0
OAM cells sent: 0
router#show queueing int sw 1.501
    Interface Switch1.501 VC 0/501
    Queueing strategy: fifo
    Output queue 0/40, 6483 drops per VC

```

Cisco bug ID CSCdt26857 (registered customers only) defines a new MIB that augments the VC tables defined in RFC 1695, also known as the ATM MIB, and in the CISCO-AAL5-MIB. It counts AAL5 VC drops on Cisco ATM router interfaces, particularly the PA-A3.

Known Issue: VC Appears Stuck

In rare circumstances, incrementing output drops result from a problem with the transmit queue for a VC. During this condition, the VC appears "stuck".

Use these tips to determine whether you are experiencing a stuck VC condition:

- Execute several instances of the **show interface atm** command and look for a rapidly increasing value for output drops.
- If your image supports per-VC queueing, execute several instances of the **show queueing interface atm** command and look for a consistent value of "Output queue 40/40" if your VC uses Layer-3 FIFO queueing.
- Execute **shutdown** and then **no shutdown** on the interface or subinterface. These commands reset the transmit ring queues.
- Execute **show atm vc** and **show atm pvc** and analyze both the input and output packet counters. Are the input packet counters incrementing? Is the problem on the transmit side only?

This table lists known fixes in microcode version G.129. If you are a registered user, you can see the details of the bugs in the Bug Toolkit (registered customers only) page. Note that it is recommended to upgrade to the latest Cisco IOS software release (registered customers only) provided by Cisco.

| Cisco Bug ID | Fixed-In Versions |
|--------------|--|
| CSCdu09828 | Workaround provided. |
| CSCdt19788 | 12.2(2.2)T 12.0(16)S01 12.0(16.6)S 12.2(0.20)T 12.1(8.1) 12.0(16.6)S01 12.0(17.1)S 12.2(0.20)PI 12.2(0.21)T 12.0(15.6)ST03 12.2(1.1) 12.0(17.2) 12.2(0.21)S 12.0(16.6)ST 12.2(0.21)PI 12.0(17.1)ST 12.1(7.5)E 12.2(1.1)PI 12.0(17.3)ST 12.1(07a)E02 12.2(1.4)S 12.0(17.6)W05(21.16) 12.1(8.5)E 12.1(08a)E 12.1(7.5)EC 12.2(3.4)PB 12.2(3.4)B 12.1(4)XZ05 12.1(4)XY07 12.1(8.5)EC 12.2(2)DD01 |
| CSCdr22203 | 12.2(03.04)B 12.2(03.04)PB 12.2(02.02)T 12.2(01.04)S 12.2(01.01)PI 12.2(00.21)PI 12.2(00.21)S 12.2(00.21)T 012.002(001.001) 12.0(10.03)S 12.0(10.03)SC 12.1(02.03)E |

| | |
|---------------------------------|--|
| CSCds01236 and CSCds35103 | 12.1(4) 12.1(03a)E 12.1(4.1)T 12.0(12.6)S01 12.1(4)AA 12.1(4.2) 12.1(4.2)T 12.0(13.1)S 12.1(4.1) 12.1(4.3)PI 12.1(03a)EC 12.1(4.2)AA 12.1(4)DB 12.1(4)DC 12.0(12.6)SC01 12.0(13.6)ST 12.1(4.4)E 12.1(4)DC01 12.1(4.4)EC |
| CSCds57642 | 12.1(5.6)E01 12.2(0.05b) 12.2(0.9)T 12.2(0.10) 12.2(0.10)PI01 12.1(5.6)EC 12.2(0.18)S 12.2(3.4)PB 12.2(2)B |

On non-distributed platforms, the ATM VCs must use Layer 3 queueing if the Cisco IOS image supports it.

Related Information

- [Troubleshooting Input Queue Drops and Output Queue Drops](#)
 - [Troubleshooting Input Drops on ATM Router Interfaces](#)
 - [Configurable per-VC Hold Queue Support for ATM Adapters](#)
 - [CISCO ATM PORT ADAPTER](#)
 - [ATM Technology Support Pages](#)
 - [Performance Tuning Basics](#)
 - [Switching Paths Overview](#)
 - [Technical Support & Documentation – Cisco Systems](#)
-

[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: Nov 15, 2007

Document ID: 10416
