

Hardware Troubleshooting for Catalyst 8540/8510 MSRs and LightStream 1010 ATM Switch

Document ID: 22013

Interface Cell Loss Issues

Contents

Invalid Cells on the LightStream 1010 and 8510
Invalid Cells on the Catalyst 8540
Invalid Cells on Ethernet Interfaces
Rejected and Discarded Cells
Understanding Switch Architecture
Inadequate Buffer Space
Exceeding Maximum Queue Limits
Usage Parameter Control (UPC) Violations
Cell Loss Priority (CLP) Discarding
Intelligent Tail Packet Discard/Early Packet Discard (ITPD/EPD)
Related Information

[<<<Previous Section](#) [Next Section>>>](#)

Invalid Cells on the LightStream 1010 and 8510

Cisco ATM switch routers, which include the LightStream 1010 and the Catalyst 8500 series switch routers, use a switching fabric with a shared-memory architecture. In some cases, the switch drops cells and increments the invalid cells counter, as shown in the output of one of the following commands and their respective platforms:

- **show switch fabric** – Catalyst 8540
- **show controller atm 2/0/0** – LightStream 1010
- **show controller controller0** – Catalyst 8510

These switches increment the invalid cells counter when they discard an ATM cell that has a valid header error checksum (HEC) value, but arrives on a non-existent virtual circuit (VC). Possible reasons include the following:

- Corrupted cell header
- Incomplete or no configuration of that VC in the switch fabric. For example, if you configure a virtual path identifier/virtual channel identifier (VPI/VCI) pair only on a router and not on the attached ATM switch, cells transmitted on this VC from the router are considered invalid by the switch.

This document explains the invalid cells counter on Cisco ATM campus switches and provides tips on how to troubleshoot incrementing invalid cells.

The output of **show controllers atm 2/0/0 (or 13/0/0) or atm 0** (depending on software version and chassis)

on the LightStream 1010 or Catalyst 8510 prints a table of the most recently received invalid cells. The **show controllers atm** command is cleared on read, meaning that the invalid cells counter is cleared when the show command is executed. Thus, if you are not receiving invalid cells continuously on an interface, the invalid cells counter will show as zero when you read a subsequent time.

```
cisco# show controllers atm 2/0/0
MMC Switch Fabric (idb=0x607F7DE0)
```

```
Key: discarded cells - # cells discarded due to lack of
resources
or policing (16-bit)
invalid cells - # good cells that came in on a non-existent
conn.
memory buffer - # cell buffers currently in use
RXcells - # rx cells (16-bit)
TXcells - # tx cells (16-bit)
RHEC - # cells with HEC errors
TPE - # cells with memory parity errors
```

```
discarded cells = 0
invalid cells = 132
memory buffer = 0
port type status RXcells TXcells RHEC TPE PACE_I PACE_M PACE_X
PACE_Y
0/0/0 155MBPS xytrpm 0xD00D 0x2420 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
0/0/1 155MBPS xytrpm 0x969D 0x2DDE 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
0/0/2 155MBPS xytrpm 0x43CF 0x6D9B 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
0/0/3 155MBPS xytrpm 0x0000 0x0000 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
0/1/0 155MBPS xytrpm 0xF7AC 0xE115 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
0/1/1 155MBPS xytrpm 0x0000 0x0000 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
0/1/2 155MBPS xytrpm 0x0000 0x0000 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
0/1/3 155MBPS xytrpm 0x7969 0x3575 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
1/0/0 622MBPS xytrPm 0xB54F 0x8B73 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
2/0/0 CPU 0x9496A8 0x5EAA4D
3/0/0 155MBPS xytrpm 0x0000 0x0000 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
3/0/1 155MBPS xytrpm 0xFB23 0xB8FB 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
3/0/2 155MBPS xytrpm 0xC5F9 0x2319 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
3/0/3 155MBPS xytrpm 0x0000 0x0000 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
3/1/0 155MBPS xytrpm 0x9B0A 0x52F0 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
3/1/1 155MBPS xytrpm 0x6B08 0x2342 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
```

```

0x0000 0x0000
3/1/2 155MBPS xytrpm 0x7467 0x0737 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000
3/1/3 155MBPS xytrpm 0x0000 0x0000 0x0000 0x0000 0x0000 0x0000
0x0000 0x0000

```

Invalid Cell Log

```

time stamp port pt clp gfc vpi vci
41 0xBDC15C5C.0x851EB690 3/0/2 0x1 0x0 0x0 0x0 0x5
42 0xBDC15C5D.0x851EB568 3/1/1 0x1 0x0 0x0 0x0 0x5
43 0xBDC15C64.0x851EAD50 3/1/1 0x1 0x0 0x0 0x0 0x11
44 0xBDC15C65.0x851EAC28 3/1/1 0x1 0x0 0x0 0x0 0x11
45 0xBDC15C66.0x851EAB00 3/1/1 0x1 0x0 0x0 0x0 0x11
46 0xBDC15C68.0x851EA8B0 3/1/1 0x1 0x0 0x0 0x0 0x11
47 0xBDC15C69.0x851EA788 0/0/2 0x1 0x0 0x0 0x0 0x11
48 0xBDC15C6B.0x851EA538 0/0/2 0x1 0x0 0x0 0x0 0x11
49 0xBDC15C6D.0x851EA2E8 3/0/1 0x1 0x0 0x0 0x0 0x11
50 0xBDC15C6E.0x851EA1C0 3/0/1 0x1 0x0 0x0 0x0 0x11
51 0xBDC15C6F.0x851EA098 3/0/1 0x1 0x0 0x0 0x0 0x11

```

Step 1 Check the **invalid cells**. The displayed invalid cells value is cleared on read.

Step 2 Check the **Invalid Cell Log**.

You can use the following objects of the CISCO-RHINO-MIB to poll your ATM switch router for the number of invalid cells:

```

ciscoAtmSwitchInvalidCells OBJECT-TYPE
SYNTAX Counter32
ACCESS read-only
STATUS mandatory
DESCRIPTION
"The total invalid cells of the switch."
 ::= { ciscoLS1010CpuSwitchGroup 5 }

```

```

ciscoAtmSwitchInvalidCellHeaderTable OBJECT-TYPE
SYNTAX SEQUENCE OF CiscoAtmSwitchInvalidCellHeaderEntry
ACCESS not-accessible
STATUS mandatory
DESCRIPTION
"A list of invalid cell header entries."
 ::= { ciscoLS1010CpuSwitchGroup 6 }

```

```

ciscoAtmSwitchInvalidCellHeaderEntry OBJECT-TYPE
SYNTAX CiscoAtmSwitchInvalidCellHeaderEntry
ACCESS not-accessible
STATUS mandatory
DESCRIPTION
"a entry of invalid cell header."
INDEX { ciscoAtmSwitchInvalidCellHeaderIndex }
 ::= { ciscoAtmSwitchInvalidCellHeaderTable 1 }

```

```

CiscoAtmSwitchInvalidCellHeaderEntry ::= SEQUENCE {
ciscoAtmSwitchInvalidCellHeaderIndex INTEGER,
ciscoAtmSwitchInvalidCellHeader OCTET STRING

```

```

}

ciscoAtmSwitchInvalidCellHeaderIndex OBJECT-TYPE
SYNTAX INTEGER(1..64)
ACCESS not-accessible
STATUS mandatory
DESCRIPTION
"A sequence number that identifies a invalid cell header
entry in the table."
 ::= { ciscoAtmSwitchInvalidCellHeaderEntry 1 }

ciscoAtmSwitchInvalidCellHeader OBJECT-TYPE
SYNTAX OCTET STRING(SIZE(5))
ACCESS read-only
STATUS mandatory
DESCRIPTION
"The most recently received invalid cells header.
octet 0 is port number (0-32),
octet 1 bit 7-5 is PTI,
octet 1 bit 4 is CLP,
octet 1 bit 3-0 is GFC,
octet 2 is VPI,
octet 3 is high byte of VCI
octet 4 is low byte of VCI."
 ::= { ciscoAtmSwitchInvalidCellHeaderEntry 2 }

```

Invalid Cells on the Catalyst 8540

The **show switch fabric** command on the Catalyst 8540 does not print a log of the most recent invalid cells. However, you can use the following commands to determine on what VPI and VCI the invalid cell arrived.

Step 1 Use the **show switch fabric** command to determine the MSC# with incrementing invalid cells. Each of the two required switch processors in the Catalyst 8540 contains four MSC ASICs. Each ASIC forms the switch fabric for several ports.

```

8540# show switch fabric
swc_presence_mask: 0x5
Switch mode: NR_20G
Number of Switch Cards present in the Chassis: 2

```

SWC_SLOT	SWC_TYPE	SWC_STATUS
5	EVEN	ACTIVE
6	NOT-PRESENT	NOT-PRESENT
7	ODD	ACTIVE

```

MMC Switch Fabric (idb=0x6244FE24)

```

```

Key: Rej. Cells - # cells rejected due to lack of resources
                or policing (16-bit)
Inv. Cells     - # good cells that came in on a non-existent conn.
Mem Buffs      - # cell buffers currently in use
RX Cells       - # rx cells (16-bit)
TX Cells       - # tx cells (16-bit)
Rx HEC         - # cells Received with HEC errors
Tx PERR        - # cells with memory parity errors

```

MSC#	Rej. Cells	Inv. Cells	Mem. Buffs	Rx Cells	Tx Cells	r
MSC 0:	0	0	0	75085	37787	0
MSC 1:	0	0	0	0	0	0
MSC 2:	0	0	0	0	0	0
MSC 3:	0	0	0	0	0	0
MSC 4:	0	0	0	0	5	0
MSC 5:	0	0	0	987	989	0
MSC 6:	0	0	0	220	220	0
MSC 7:	0	0	0	2677	23606	0

Switch Fabric Statistics

```

Rejected Cells: 0
Invalid Cells: 0
Memory Buffers: 0
Rx Cells: 78969
Tx Cells: 62607
RHEC: 0
TPE: 0

```

[Information Deleted]

Step 2 Check the **MSC#**, **Rej. Cells**, and **Inv. Cells** fields. They indicate the rejected cells per MSC# or matching set of physical ports.

Step 3 Check the **Switch Fabric Statistics** section for the **Rejected Cells** and **Invalid Cells** fields. They indicate the total number of rejected cells.

Step 4 Use the **show mmc ports** command to determine which physical ports use the particular MSC.

```

8540# show mmc ports
int a0/0/0: msc#: 0 port#: 12
int a0/0/1: msc#: 0 port#: 8
int a0/0/2: msc#: 0 port#: 4
int a0/0/3: msc#: 0 port#: 0
int a0/0/4: msc#: 0 port#: 14
int a0/0/5: msc#: 0 port#: 10
int a0/0/6: msc#: 0 port#: 6
int a0/0/7: msc#: 0 port#: 2
int a0/0/8: msc#: 1 port#: 12
int a0/0/9: msc#: 1 port#: 8
int a0/0/10: msc#: 1 port#: 4
int a0/0/11: msc#: 1 port#: 0
int a0/0/12: msc#: 1 port#: 14
int a0/0/13: msc#: 1 port#: 10
int a0/0/14: msc#: 1 port#: 6
int a0/0/15: msc#: 1 port#: 2
[output omitted]

```

Step 3 Use the **show mmc msc_reg all** command to view details of the invalid cells. The value for "m" is the MSC number. The following sample output was taken from a switch with invalid cells on MSC# 1:

```

Switch# show mmc msc_reg all 1

gcr0[1] = 0x0000A112
...

```

```
icc[0] = 0x00000026
...
ich[1] = 0x00000D00 0x00640064
vci:64 pti:0 clp:0 vpi:64 ssp:0 sp:D
...
```

Step 4 Check the **icc** field. Any non-zero value indicates invalid cells.

Step 5 Check the **ich** field and the following **vci**, **vpi**, and **sp** values:

- **vci** -- identifies the VCI receiving the last invalid cell.
- **vpi** -- identifies the VPI receiving the last invalid cell.
- **sp** -- identifies the port (p) receiving the invalid cell.

Step 6 Use the **show atm vc interface atm** command to confirm the connection exists in the system.

Note: The invalid cells log may refer to port numbers and VPI/VCI values that do not correspond to actual port numbers and VCs. The reason is that the port interface (PIF) chip on some modules considers idle cells on some ports as invalid cells. The interface chip actually changes the default VPI of 0 for idle cells to a different value. For example, VPI/VCI 0/16 on port #1 will be changed to VPI/VCI 4/16 on the PIF chip. The interface chip drops idle cells from these ports and increments the invalid cells counter. On ports 0 and 6, the interface chip does not change the default idle cell value of VPI/VCI 0/0 since the physical VPI is the same as VPI on the interface chip.

Invalid Cells on Ethernet Interfaces

Ethernet interfaces can also experience cell drops due to invalid cells in an ATM switch router. Ethernet interfaces derive much of their local intelligence from a PIF ASIC, which segments Ethernet frames for transmission across the switch router's internal ATM switch fabric. A PIF checks if a received frame or packet is received for a protocol for which the interface is configured. It then searches the content-addressable memory (CAM) table and determines the egress VPI/VCI value for the destination port. Finally, the PIF segments the frame into cells, applies a five-byte header with the appropriate egress VPI/VCI information, and sends the cells out the switch fabric. If the PIF needs to drop a frame, it tags all cells in that frame with a VPI/VCI =0/0, and the switch fabric then drops these cells.

An Ethernet PIF drops cells and increments the invalid cells counter due to one of the following reasons:

- Layer-2 filtering of MAC addresses reachable out the same interface as the received Ethernet frame. The Catalyst 8500 filters such "local" MAC addresses by sending received frames into the switch fabric on the 0/0 VPI/VCI. These drops are equivalent to the in-discards counter on the Catalyst 5000.
- Packet discard on a multiplexing-based (MUX-based) Gigabit Ethernet line card. On such cards, when a received packet has a non-forwardable or un-recognizable protocol and must be discarded, the switch puts the packet on the 0/0 VPI/VCI.
- Microcode put a packet on a VPI/VCI that is not actually established in the switch core.

Rejected and Discarded Cells

In some cases, the switch discards cells and reports these drops in the output of one of the following commands, depending on the platform:

- **show switch fabric** – Catalyst 8540 MSR
- **show controller atm 2/0/0 or atm0** – LightStream 1010 or Catalyst 8510 in standalone chassis
- **show controller atm 13/0/0** – LightStream 1010 or Catalyst 8510 in bottom five slots of Catalyst 5500

A Cisco ATM switch router increments the discarded or rejected cells counter when it drops a cell due to one of the following reasons:

- Inadequate buffer space
- Exceeding of maximum queue limits
- Usage parameter control (UPC) violations
- Cell loss priority (CLP) discarding
- Intelligent Tail packet discard (ITPD) / early packet discard (EPD)

The purpose of these sections is to review each of the above reasons and provide tips on how to troubleshoot why you are seeing non-zero values for rejected cells. But before you start, look at Understanding Switch Architecture.

Understanding Switch Architecture

The LightStream 1010 and Catalyst 8510 use an architecture that differs from the Catalyst 8540.

On the LightStream 1010 and 8510, use the **show controller atm 2/0/0** (or **13/0/0** if used in the Catalyst 5500) or **show controller atm0** to view statistics for the CPU and ATM switch processor:

```
ls1010# show controller atm 2/0/0
MMC Switch Fabric (idb=0x60AD7B20)
```

```
Key: discarded cells - # cells discarded due to lack of
resources
or policing (16-bit)
invalid cells - # good cells that came in on a non-existent
conn.
memory buffer - # cell buffers currently in use
RXcells - # rx cells (16-bit)
TXcells - # tx cells (16-bit)
RHEC - # cells with HEC errors
TPE - # cells with memory parity errors
```

```
discarded cells = 0
invalid cells = 184027
memory buffer = 0
garbage cells to cpu = 0
unexpected marker intrs = 0
```

The **discarded cells** field indicates the total count of discarded cells across all ports.

On the Catalyst 8540, use the **show switch fabric** command to display the number of rejected cells. Note that this output differs from the LightStream 1010 output in that it shows rejected cells per Modular Switching Component (MSC) number. MSC application-specific integrated circuits (ASICs) form the switch fabric for a defined set of modules and ports.

```
8540MSR# show switch fabric
swc_presence_mask: 0x5
```

Switch mode: NR_20G
Number of Switch Cards present in the Chassis: 2

```
SWC_SLOT          SWC_TYPE          SWC_STATUS
=====
      5             EVEN             ACTIVE
      6             NOT-PRESENT       NOT-PRESENT
      7             ODD              ACTIVE
```

MMC Switch Fabric (idb=0x6244FE24)

Key: Rej. Cells - # cells rejected due to lack of resources
 or policing (16-bit)
Inv. Cells - # good cells that came in on a non-existent conn.
Mem Buffs - # cell buffers currently in use
RX Cells - # rx cells (16-bit)
TX Cells - # tx cells (16-bit)
Rx HEC - # cells Received with HEC errors
Tx PERR - # cells with memory parity errors

MSC#	Rej. Cells	Inv. Cells	Mem. Buffs	Rx Cells	Tx Cells	r
MSC 0:	0	0	0	82678	28733	0
MSC 1:	0	0	0	0	0	0
MSC 2:	0	0	0	0	0	0
MSC 3:	0	0	0	0	0	0
MSC 4:	0	0	0	0	5	0
MSC 5:	0	0	0	987	989	0
MSC 6:	0	0	0	220	220	0
MSC 7:	0	0	0	2677	28138	0

Switch Fabric Statistics

Rejected Cells: 0
Invalid Cells: 0
Memory Buffers: 0
Rx Cells: 86562
Tx Cells: 58085
RHEC: 0
TPE: 0

[Information Deleted]

Step 2 Check the **MSC#**, **Rej. Cells**, and **Inv. Cells** fields. They indicate the rejected cells per MSC# or matching set of physical ports.

Step 3 Check the **Switch Fabric Statistics** section for the **Rejected Cells** and **Invalid Cells** fields. They indicate the total number of rejected cells.

Each of the two required switch processors in the 8540 contain four MSC ASICs, which build the internal switch fabric for half of the ports in a system. Use the **show mmc ports** command to determine which physical ports use a particular MSC#.

```
8540#show mmc ports
int a0/0/0: msc#: 0 port#: 12
int a0/0/1: msc#: 0 port#: 8
int a0/0/2: msc#: 0 port#: 4
int a0/0/3: msc#: 0 port#: 0
int a0/0/4: msc#: 0 port#: 14
```

```

int a0/0/5: msc#: 0 port#: 10
int a0/0/6: msc#: 0 port#: 6
int a0/0/7: msc#: 0 port#: 2
int a0/0/8: msc#: 1 port#: 12
int a0/0/9: msc#: 1 port#: 8
int a0/0/10: msc#: 1 port#: 4
int a0/0/11: msc#: 1 port#: 0
int a0/0/12: msc#: 1 port#: 14
int a0/0/13: msc#: 1 port#: 10
int a0/0/14: msc#: 1 port#: 6
int a0/0/15: msc#: 1 port#: 2
[output omitted]

```

With each slot, the first half of the ports use an even MSC#, and the second half use an odd MSC#. However, when using the original LightStream 1010 port adapter modules (PAMs) with a super carrier module (SuperCAM) carrier access module, all ports on a single SuperCAM map to an even SP and MSC ASIC. For example, the first half of the ports in slot 0 normally connect to MSC0 of SP0, while the second half of the ports in slot 0 connect to MSC1 of SP1. However, with a SuperCAM, both sets of ports connect to MSC0 of SP0.

Inadequate Buffer Space

Cisco campus ATM switches use a shared-memory architecture that stores up to 65,536 cells. Using a shared-memory design provides the following benefits:

- Supports more connections or higher loads.
- Supports an increased amount of multicast traffic since only a single copy of any multicast cell is stored in the common cell memory.
- Provides complete sharing with a maximum level of statistical buffer sharing.

Since all ports can use the full memory, it is important that the buffer management process ensures fairness among ports by ensuring that one or a small subset of ports cannot occupy all of the buffers.

Note: The opposite of a shared-memory architecture is a per-port output buffer architecture, in which each port has dedicated memory buffers that cannot be accessed by other ports. The Catalyst 6000 and Catalyst 5000 are output-buffered switches.

On an LightStream 1010, use the **sh controller atm 2/0/0** command to view the number of memory buffers currently in use.

```

ls1010# show controller atm 2/0/0
MMC Switch Fabric (idb=0x60AD7B20)

```

```

Key: discarded cells - # cells discarded due to lack of
resources
or policing (16-bit)
invalid cells - # good cells that came in on a non-existent
conn.
memory buffer - # cell buffers currently in use
RXcells - # rx cells (16-bit)
TXcells - # tx cells (16-bit)

```

```
RHEC - # cells with HEC errors
TPE - # cells with memory parity errors
```

```
discarded cells = 0
invalid cells = 184027
memory buffer = 0
garbage cells to cpu = 0
unexpected marker intrs = 0
```

Check the **memory buffer** field. It should display a non-zero value on a busy production switch.

You can use the following managed objects of the CISCO-RHINO-MIB to poll your ATM switch router for the number of free buffers and discarded cells:

Managed Object	Description
ciscoAtmSwitchTotalBuffer	Total cell buffer count in the switch's shared memory.
ciscoAtmSwitchFreeBuffer	Free cell buffer count in switch's shared memory.
ciscoAtmSwitchDiscardCells	Total discarded cells of the switch.

Exceeding Maximum Queue Limits

ATM switch routers use configurable queue limits and thresholds to control queuing in the system. The queuing processes and configurable values vary with the feature card installed on the ATM switch processor (ASP) or multiservice switch processor (MSP):

	Feature card per-class queuing (FC-PCQ)	Feature card per-flow queuing (FC-PFQ) and 8540
Service category limit	Yes	No
Maximum queue size per interface	Yes	No
Threshold groups	No	Yes

The Catalyst 8510 and the LightStream 1010 with a FC-PCQ support service category limits, which restrict the number of cells admitted into the switch as determined by the type of output queues. Use the **show atm resource** command to display these limits. Use the **atm service-category-limit** command to configure a non-default value.

```
Switch# show atm resource
```

```
Resource configuration:
```

```
Over-subscription-factor 16Sustained-cell-rate-margin-factor
1%
Abr-mode: relative-rate
Atm service-category-limit (in cells):
64544 cbr 64544 vbr-rt 64544 vbr-nrt 64544 abr-ubr
```

```
Resource state:
Cells per service-category:
0 cbr 0 vbr-rt 0 vbr-nrt 0 abr-ubr
```

Note: In the previous example all ATM service classes have access to most of the shared memory by default.

The Catalyst 8510 and LightStream 1010 with a FC-PCQ also support maximum queue sizes, which determine the number of cells that can be scheduled for transmission per ATM service class per interface. Use the **atm output-queue** command to configure a non-default value.

Command	Description
Switch(config-if)# atm output-queue [force] { cbr vbr-rt vbr-nrt abr-ubr } max-size <i>number</i>	Configures the the maximum queue size of the output queue.
Switch> show atm interface resourceatm { <i>card/subcard/port</i> }	Displays the resource management interface configuration status and statistics.
Switch(config-if)# Switch(config)# atm threshold-group service { cbr vbr-rt vbr-nrt abr ubr } <i>group</i> #	Configures a service category to a threshold group.

Because not all queue size values are supported by the switch fabric, the value installed is displayed, as well as the configuration value requested. The value installed is always greater than or equal to that requested. Use the **show atm interface resource atm** command to display both values.

```
Switch> show atm interface resource atm 3/0/0

Resource Management configuration:
Output queues:

Max sizes(explicit cfg): 30000 cbr, none vbr-rt, none vbr-nrt,
none abr-ubr
! -- Note the "explicit cfg" values.
Max sizes(installed): 30208 cbr, 256 vbr-rt, 4096 vbr-nrt,
12032 abr-ubr
! -- Note the "installed" values.
Efcf threshold: 25% cbr, 25% vbr-rt, 25% vbr-nrt, 25% abr, 25%
ubr
Discard threshold: 87% cbr, 87% vbr-rt, 87% vbr-nrt, 87% abr,
87% ubr
Abr-relative-rate threshold: 25% abr
[output omitted]
```

Note: Systems with a FC-PFQ queue cells on input, not output, so the **atm output-queue** commands do not apply.

Catalyst 8510 and LightStream 1010 systems with a FC-PFQ and Catalyst 8540s support the threshold groups feature. Each group consists of virtual paths (VPs) and virtual circuits (VCs) that belong to the same ATM service category, such as VBR-nrt or UBR. By default, one threshold group holds cells for one ATM service class. Use the **atm threshold-group** command to assign more than one service category to a threshold group and to assign a service category to a non-default group number. Use the **show atm resource** command to confirm your changes.

In the following sample output, the ATM switch router is using the default settings. One ATM service class is assigned to one threshold group.

```

Ls1010# show atm resource
Resource configuration:
  Over-subscription-factor 8  Sustained-cell-rate-margin-factor 1%
  Abr-mode:  EFCI
  Hierarchical Scheduling Mode : disabled
  Service Category to Threshold Group mapping:
    cbr 1 vbr-rt 2 vbr-nrt 3 abr 4 ubr 5
Threshold Groups:
Group Max      Max Q  Min Q  Q thresholds  Cell  Name
  cells limit  limit  Mark Discard  count
  instal instal instal
-----
  1   16447   767   767   25 %  62 %    0   cpu-switched-tg
  2   65535   127   127   25 %  87 %    0   vbr-rt-default-tg
  3   65535   511    31   25 %  87 %    0   vbr-nrt-default-tg
  4   65535   511   511   25 %  87 %    0   ipc-tg
  5   61439   511    31   25 %  62 %    0   switching-tg
  6   65535  4095  1023   25 %  87 %    0   well-known-vc-tg

Ls1010#

```

Note: Each group matches to one ATM service category by default.

Each threshold group consists of eight regions, with each region having a set of thresholds. A threshold group congests when its member VCs have a large number of cells stored in the shared cell memory. As the cumulative number of queued cells for the member VCs approaches the "Max cells instal" value, the maximum number of cells in each per-VC and per-VP queue shrinks from the max-queue-limit to min-queue-limit. Refer to the "Max Q limit instal" and "Min Q limit instal" columns in **show atm resource** output for the queue size values.

When congestion is in the range of 0 cells (uncongested) to 1/8th full, the connection queues are limited to max-queue-size. In general, as you move from one region to another, you make the new threshold max (previous-threshold/2, min-queue-threshold). When congestion is in the range of 7/8ths full to completely full, the connection queues are limited to min-queue-size. Note that the operation of the switch for threshold groups in upper regions only occurs if the group congests by going above 1/8th full. However, the maximum size and threshold position commands are effective even for threshold groups in the lowest region.

The following commands adjust the threshold group values.

Command	Description
atm threshold-group <i>group max-cells</i> <i>number</i>	Configures the maximum number of cells queued for <i>all</i> VCs in the group. See the "Max cells instal" value in show atm resource .
atm threshold-group	Configures the largest <i>per-VC</i> queue

<i>group max-queue-limit number</i>	limit applied to all VCs in the group. See the "Max Q limit install" value in show atm resource .
atm threshold-group <i>group min-queue-limit number</i>	Configures the smallest <i>per-VC</i> queue-limit applied to all VCs in the group. See the "Min Q limit install" value in show atm resource .
atm threshold-group <i>group marking-threshold percent</i>	Determines the point at which a per-VC queue is considered "full", and the switch begins to set the explicit forward congestion indication (EFCI) bit or implements available bit rate (ABR) relative-rate marking. See the "Q thresholds Mark" value in show atm resource .
atm threshold-group <i>group discard-threshold percent</i>	Determines the point at which a per-VC queue is considered "full", and the switch begins to discard cells with the cell loss priority (CLP) bit to one and implements early packet discard (EPD). See the "Q thresholds Discard" value in show atm resource .

The **show atm vc** command displays the following two counters related to rejected or discarded cells caused by queue threshold values being exceeded:

- Number of cells queued per threshold group
- Number of drops due to full queue via the counters "Rx Clp0 q full drops" and "Rx Clp1 qthresh drops"

Note: The output of the **show atm vc** command changes with respect to drop counters depending on whether packet discard is enabled on the VC.

```
switch# show atm vc int atm 12/0/3 0 100

Interface: ATM12/0/3, Type: oc3suni
VPI = 0 VCI = 100
Status: UP
Time-since-last-status-change: 00:18:09
Connection-type: PVC
Cast-type: point-to-point
Packet-discard-option: disabled
Usage-Parameter-Control (UPC): pass
Wrr weight: 2
Number of OAM-configured connections: 0
OAM-configuration: disabled
OAM-states: Not-applicable
Cross-connect-interface: ATM12/0/0, Type: oc3suni
Cross-connect-VPI = 0
Cross-connect-VCI = 100
Cross-connect-UPC: pass
```

```

Cross-connect OAM-configuration: disabled
Cross-connect OAM-state: Not-applicable
Threshold Group: 1, Cells queued: 63
Rx cells: 2010095, Tx cells: 0
Tx Clp0:0, Tx Clp1: 0
Rx Clp0:2010095, Rx Clp1: 0
Rx Upc Violations:0, Rx cell drops:148
Rx Clp0 q full drops:148, Rx Clp1 qthresh drops:0
[output omitted]

```

Note: The CBR service category is assigned to group 1 by default.

Check the numbers following the **Clp0 q full drops** and **Clp1 qthresh drops** fields.

You also can obtain these counts via SNMP polling.

Managed Object	Description
ciscoAtmVclClp0VcqFullCellDrops	Total number of cells received on this virtual channel link (VCL) with the CLP bit clear, discarded because the per-VC queue limit is exceeded. This counter is valid only if EPD is disabled on the VCL. On LightStream 1010s, this counter is valid only when the switch processor is equipped with an FC-PFQ.
ciscoAtmVclVcqClpThreshCellDrops	Total number of cells received on this VCL, discarded because the discard threshold (as opposed to the queue-limit) is exceeded on the per-VC queue, and the CLP bit is set. This counter is valid only if EPD is disabled on the VCL. On LightStream 1010s, this counter is valid only when the switch processor is equipped with an FC-PFQ.
ciscoAtmVclLsPerVcQThreshGrp	Threshold Group to which the cell-queue for cells received by this VC are queued. Note that this value is not valid until the VCL is in an active cross-connect. On LightStream 1010s, this counter is valid only when the switch processor is equipped with an FC-PFQ.

Usage Parameter Control (UPC) Violations

When configured, an ATM switch at the network side of a user-network interface (UNI) polices the flow of cells in the forward (into the network) direction of a virtual connection. These policing mechanisms are

known as usage parameter control (UPC). They determine whether received cells are complying with the negotiated traffic management values, and then take one of the following actions on violating cells, depending on the configuration:

- Pass the cell without changing the cell loss priority (CLP) bit in the cell header.
- Tags the cell with a CLP bit value of 1.
- Discard the cell. If you enable the discard option, we recommend that you also enable the tail packet discard (TPD) feature discussed in the section Intelligent Tail Packet Discard/Early Packet Discard later in this document.

Use the **upc** parameter in the **atm pvc** command to specify the violation action. The full syntax of the command is:

Command	Description
atm pvc vpi-A [vci-A any-vci] [upc upc-A] [pd pd] interface atm card-B/subcard-B/port-B[.vpt#] vpi-B [vci-B any-vci] [upcupc-B]	Configures a PVC. See the "Max cells install" value in show atm resource .

The **upc** parameter cannot be set to tag or drop on the processor port (ATM 0).

Normally, UPC polices only the source end of a soft VC. Use the **atm svc-upc-intent drop** command to enable UPC by default for all terminating VCs on the destination end of a soft VC.

Use the **show atm vc** command to view the configured UPC action and intelligent packet discard mechanisms, as well as the number of cells discarded due to UPC violations.

```
Switch# show atm vc interface atm 0/0/1.51 51 16

Interface: ATM0/0/1.51, Type: oc3suni
VPI = 51 VCI = 16
Status: DOWN
Time-since-last-status-change: 2w0d
Connection-type: PVC
Cast-type: point-to-point
Packet-discard-option: enabled
Usage-Parameter-Control (UPC): pass
Wrr weight: 32
Number of OAM-configured connections: 0
OAM-configuration: disabled
OAM-states: Not-applicable
Cross-connect-interface: ATM2/0/0, Type: ATM Swi/Proc
Cross-connect-VPI = 0
Cross-connect-VCI = 73
Cross-connect-UPC: pass
Cross-connect OAM-configuration: disabled
Cross-connect OAM-state: Not-applicable
Encapsulation: AAL5ILMI
```

```

Threshold Group: 6, Cells queued: 0
Rx cells: 0, Tx cells: 0
Tx Clp0:0, Tx Clp1: 0
Rx Clp0:0, Rx Clp1: 0
Rx Upc Violations:0, Rx cell drops:0
Rx pkts:0, Rx pkt drops:0
Rx connection-traffic-table-index: 6
Rx service-category: UBR (Unspecified Bit Rate)
Rx pcr-clp01: 424
Rx scr-clp01: none
Rx mcr-clp01: none
Rx cdvt: 1024 (from default for interface)
Rx mbs: none
Tx connection-traffic-table-index: 6
Tx service-category: UBR (Unspecified Bit Rate)
Tx pcr-clp01: 424
Tx scr-clp01: none
Tx mcr-clp01: none
Tx cdvt: none
Tx mbs: none
No AAL5 connection registered

```

Check the **Packet-discard-option** and **Usage-Parameter-Control** setting. Also check the **Upc Violations** field for the number of violations.

You also can obtain these counts via SNMP polling. Use the `ciscoAtmVclUpcViolations` managed object in the `CISCO-ATM-CONN-MIB`.

Note: In evaluating the arriving cell rate, the ATM switch router counts both Operation, Administration, and Maintenance (OAM) cells as well as with data cells since the current signaling protocol does not allow a user to explicitly specify traffic parameters for OAM flows.

Cell Loss Priority (CLP) Discarding

The standard ATM cell header includes the cell loss priority (CLP) bit, which explicitly indicates that a cell experienced congestion during transmission to the destination end. A CLP value of one means that the cell has a lower priority and thus is more likely to be dropped in times of congestion. Thus, you can use the CLP bit to generate different priority cell flows.

ATM switch routers use a threshold-based CLP selective discard mechanism that imposes a threshold on the number of cell buffers to be shared by CLP = 0 and CLP = 1 cells. When the switch port queue occupancy reaches a user-configurable threshold level, only CLP = 0 cells are allowed to enter the system, and CLP = 1 cells are discarded.

The **show atm resource** command displays the queue threshold percentage, after which cells are eligible for CLP discard or early packet discard. This value is the column labeled "Discard."

```

NewLs1010# show atm resource
Resource configuration:
  Over-subscription-factor 8 Sustained-cell-rate-margin-factor 1%
  Abr-mode: EFCI
  Hierarchical Scheduling Mode : disabled
  Service Category to Threshold Group mapping:
    cbr 1 vbr-rt 2 vbr-nrt 3 abr 4 ubr 5
  Threshold Groups:

```

Group	Max cells instal	Max Q limit instal	Min Q limit instal	Q thresholds		Cell count	Name
				Mark	Discard		
1	16447	767	767	25 %	62 %	0	cpu-switched-tg
2	65535	127	127	25 %	87 %	0	vbrnr-default-tg
3	65535	511	31	25 %	87 %	0	vbrnr-default-tg
4	65535	511	511	25 %	87 %	0	ipc-tg
5	61439	511	31	25 %	62 %	0	switching-tg
6	65535	4095	1023	25 %	87 %	0	well-known-vc-tg

NewLs1010#

Adjust the value of the discard threshold with the command **atm threshold-group** *[module-id module] group discard-threshold percent*.

Note too that there are two threshold values:

- Mark – Threshold at which the explicit forward congestion indication (EFCI) bit is set.
- Discard – Threshold at which cells are eligible for CLP discard or early packet discard (EPD).

Optionally, you can enable tail packet discard on each VC for use with CLP selective discard. With the TPD option, the system "goodput" (usable throughput) is enhanced. You enable tail packet discard (TPD) by specifying the "pd" or packet discard parameter in the **atm pvc** command. The "pd" parameter enables both tail packet discard and early packet discard.

The command syntax is as follows:

Command	Description
atm pvc <i>vpi vci</i> [pd pd] [rx-cttr index] [tx-cttr index]	Configures a PVC.
atm soft-vc <i>source-vpi source-vci</i> dest-address <i>atm-address dest-vpi dest-vci</i> [pd pd] [rx-cttr index] [tx-cttr index]	Configures a soft PVC on the switch router.

Use the **show atm interface resource** command to display the discard threshold percentages.

```
LS1010# show atm interface resource atm 4/1/0
Resource Management configuration:
  Output queues:
    Max sizes(explicit cfg): none cbr, none vbr-rt, none vbr-nrt, none abr-r
    Max sizes(installed): 256 cbr, 512 vbr-rt, 4096 vbr-nrt, 11776 abr-ubr
    EfcI threshold: 25% cbr, 25% vbr-rt, 25% vbr-nrt, 25% abr, 25% ubr
    Discard threshold: 87% cbr, 87% vbr-rt, 87% vbr-nrt, 87% abr, 87% ubr
    Abr-relative-rate threshold: 25% abr
  CAC Configuration to account for Framing Overhead : Disabled
  Pacing: disabled 0 Kbps rate configured, 0 Kbps rate installed
  overbooking : disabled
  Service Categories supported: cbr,vbr-rt,vbr-nrt,abr,ubr
  Link Distance: 0 kilometers
  Controlled Link sharing:
    [Information Deleted]
```

Check the **Discard threshold** values.

With UPC, you can implement either drop or tag as UPC policy. You cannot set up a tag-and-drop policy in which you tag above sustained cell rate (SCR) and drop above peak cell rate (PCR).

The following example output was generated on a permanent virtual circuit (PVC) with packet discard (PD) enabled, UPC set to pass, and traffic shaping parameters set to 10 MB SCR and 20 MB PCR. Sending 25 MB through the PVC produces UPC violations on approximately 60% of the cells.

```
switch# show atm vc int a0/1/3 2 122
Interface: ATM0/1/3, Type: oc3suni
VPI = 2 VCI = 122
Status: UP
Time-since-last-status-change: 00:56:47
Connection-type: SoftVC
Cast-type: point-to-point
Soft vc location: Source
Remote ATM address:
39.840f.8011.4126.0002.fd98.0001.4000.0c80.1010.00
Remote VPI: 2
Remote VCI: 122
Soft vc call state: Active
Number of soft vc re-try attempts: 0
First-retry-interval: 5000 milliseconds
Maximum-retry-interval: 60000 milliseconds
Aggregate admin weight: 5040
TIME STAMPS:
Current Slot:2
Outgoing Setup March 12 11:45:31.180
Incoming Connect March 12 11:45:31.188
Packet-discard-option: enabled
Usage-Parameter-Control (UPC): tag
Wrr weight: 2
Number of OAM-configured connections: 0
OAM-configuration: disabled
OAM-states: Not-applicable
Cross-connect-interface: ATM0/1/2, Type: oc3suni
Cross-connect-VPI = 0
Cross-connect-VCI = 112
Cross-connect-UPC: pass
Cross-connect OAM-configuration: disabled
Cross-connect OAM-state: Not-applicable
Threshold Group: 2, Cells queued: 0
Rx cells: 3706784, Tx cells: 0
Tx Clp0:0, Tx Clp1: 0
Rx Clp0:3706784, Rx Clp1: 0
Rx Upc Violations:2257061, Rx cell drops:0
Rx pkts:115837, Rx pkt drops:0
Rx connection-traffic-table-index: 3020000
Rx service-category: VBR-RT (Realtime Variable Bit Rate)
Rx pcr-clp01: 20000
Rx scr-clp01: 10000
Rx mcr-clp01: none
Rx cdvt: 1024 (from default for interface)
Rx mbs: 1024 (from default for interface)
```

Check the **Packet-discard-option** and **Usage-Parameter-Control** setting.

Check the **Rx cells** and **Tx cells** fields plus the **Rx Upc Violations** and **Rx cell drops** fields.

With switched virtual circuits (SVCs), Cisco campus ATM switches use the AAL5 information element (IE) to indicate whether or not to enable packet discard; the presence of the AAL5 IE tells the switch to enable PD. With ATM interfaces using UNI 4.0 signaling, ATM switches may use the frame-discard bits in the traffic management options field of the ATM Traffic Descriptor IE.

Intelligent Tail Packet Discard/Early Packet Discard (ITPD/EPD)

Most data frames are segmented and transmitted across an ATM cloud as multiple cells. If one or more cells are dropped by the network, the resulting packet fails the CRC check at the receiving end and must be retransmitted. Such retransmissions lead to poor effective throughput or goodput, which is defined as the number of delivered cells that are not part of a retransmission or an incomplete packet.

To maximize the number of completely delivered packets, your ATM switch router implements a unique ITPD/EPD scheme that intelligently and selectively discards cells belonging to the same packets in order to minimize the effects of fragmentation. Working together, ITPD/EPD can prevent frequent buffer overflows by dumping corrupted or complete packets from rapidly filling buffers. By dropping a small number of packets instead of cells from a large number of packets, occasional buffer overflows do not have serious negative effects on the end-to-end system goodput.

ITPD works to minimize fragmentation as it is occurring. ITPD acts in response to cell drops due to one of the following reasons:

- UPC enforcement action upon violation
- Buffer overflow
- Exceeding any of the buffer limits
- CLP selective discard

When one cell of a packet has been discarded by the ATM switch router, ITPD discards all subsequent cells of the same packet. Depending on the feature card, the last cell (also known as the end of packet (EOP) cell) may be dropped as well.

ATM switches identify the EOP cell via a bit in the payload type identifier (PTI) field of a cell header. The FC-PCQ does not drop the last cell of frame when doing EPD, while the FC-PFQ does.

EPD works to prevent fragmentation before it occurs. With EPD, the ATM switch router begins to discard all cells except the EOP cell from newly arriving packets when the switch buffer queues reach a user-configurable threshold level. If the first cell of a packet has entered the buffer, all remaining cells of the packet are also allowed to enter if sufficient buffer space is available. Otherwise, TPD is enabled.

Use the **atm threshold-group group discard-threshold percent** command to configure the threshold at which point the queue is considered full and EPD starts dropping cells. See the "Q thresholds Discard" value in the output of **show atm resource** for the default discard percentage.

The placement of the EPD threshold determines how efficiently the buffer is used and how often cells are dropped. The EPD threshold essentially functions as the effective buffer size. The excess buffer capacity above the EPD threshold is used to accommodate cells from those packets that have already had cells in the buffer or in the transmission on the line.

Setting the threshold depends on many factors, including:

- Distribution of packet sizes
- Traffic distribution
- Duration of the congestion period
- Proportion of the incoming cells during the congestion period that belong to outstanding packets and as a result must be buffered.
- The interaction with other ATM-level or transport-level flow and congestion control mechanisms.

In addition, the amount of excess buffer capacity required depends on how the buffer is shared with non-TPD/EPD traffic. You can enable TPD by specifying the "pd" or packet discard parameter in the **atm pvc** command. The "pd" parameter enables both tail packet discard and early packet discard. Packet discard can only be enabled for AAL5 connections. Here is how the drop behavior changes with UPC and the PD option:

- If UPC is configured to drop and PD is disabled, then the switch drops violating cells only.
- If UPC is configured to drop and PD is enabled, then the switch performs ITPD and drops all the cells following the violating one (except the last cell).
- If PD is enabled and cells are queued to a buffer which has already exceeded its EPD threshold, then the switch drops the complete (AAL5) packet.

In other words, PD is applied as EPD whenever possible (high queue sizes, for example) and as ITPD in all other cases, including UPC drops and buffer overflow.

Related Information

- [Troubleshooting Switch Router ATM Interface Connections](#)
- [ATM and Layer 3 Module Installation Guide](#)
- [Traffic and Resource Management](#)
- [Configuring Resource Management](#)
- [ATM Technology Support Pages](#)

All contents are Copyright © 1992--2003 Cisco Systems Inc. All rights reserved. Important Notices and Privacy Statement.

Updated: Nov 02, 2005

Document ID: 22013
