

Understanding Maximum Transmission Unit (MTU) on ATM Interfaces

Document ID: 10479

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Why are MTUs 4470 Bytes?

- AAL5 Oversized SDUs and Length Violations
- Benefits of Large and Same-Size MTUs

Relevant RFCs

- IP Fragmentation

Jumbo Frame Support

Troubleshoot

Known Issue – MTU and Bridging

Related Information

Introduction

Maximum transmission unit (MTU) defines the largest size of packets that an interface can transmit without the need to fragment. IP packets larger than the MTU must go through IP fragmentation procedures.

Cisco ATM router interfaces support an MTU between 64 and 17966 bytes. Each interface supports a default maximum packet size. For example, the maximum value is 9288 bytes on both the ATM interface processor (AIP) and network processor module (NP), and 4470 bytes on the PA-A3 and PA-A2 port adapters.

This document reviews the default MTU values for ATM interfaces and clarifies when a router increments the AAL5 Oversized SDUs and the AAL5 length violation counters.

Prerequisites

Requirements

There are no specific requirements for this document.

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.

Why are MTUs 4470 Bytes?

Most Cisco ATM router interfaces use a default MTU size of 4470 bytes. This number was chosen to match exactly Fiber Distributed Data Interface (FDDI) and High-Speed Serial Interface (HSSI) interfaces for autonomous switching.

Use the **mtu** command in interface configuration mode to configure a non-default value. Note that subinterfaces support a value that is different from the main interface as long as the main interface's value is as large as, or larger than the largest subinterface MTU.

```
7200#show interface atm 3/0
ATM3/0 is up, line protocol is up
Hardware is ENHANCED ATM PA
Internet address is 1.1.1.1/8
MTU 4470 bytes, sub MTU 1500, BW 149760 Kbit, DLY 80 usec,
reliability 255/255, txload 1/255, rxload 1/255
```

Use the **show atm interface atm** command to view the currently configured value.

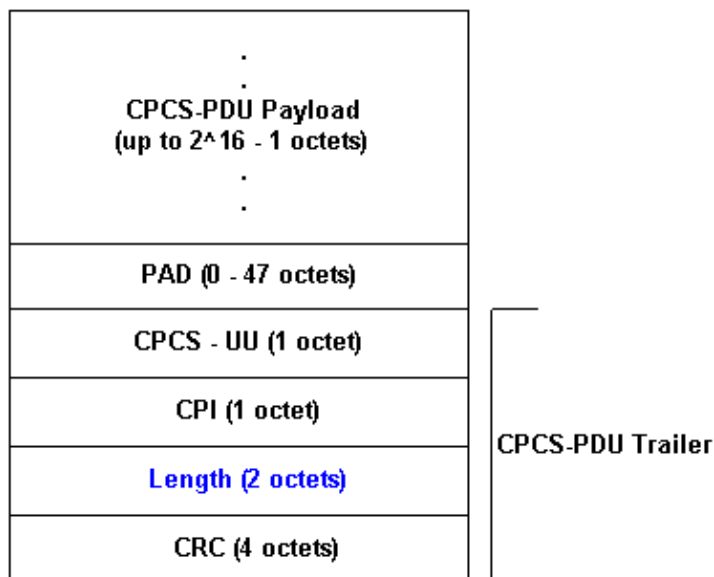
```
7200#show atm interface atm 3/0
Interface ATM3/0:
AAL enabled: AAL5 , Maximum VCs: 4096, Current VCCs: 2
Maximum Transmit Channels: 0
Max. Datagram Size: 4528
PLIM Type: SONET - 155000Kbps, TX clocking: LINE
Cell-payload scrambling: ON
sts-stream scrambling: ON
8359 input, 8495 output, 0 IN fast, 0 OUT fast, 0 out drop
Avail bw = 155000
Config. is ACTIVE
```

AAL5 Oversized SDUs and Length Violations

The **show interface atm** command reports two counters highlighted in bold and relevant to a discussion of packet size.

```
7200#show interface atm1/ima0
ATM1/IMA0.1 is up, line protocol is up
Hardware is ATM IMA
MTU 4470 bytes, BW 6000 Kbit, DLY 20000 usec,
reliability 255/255, txload 1/255, rxload 2/255
Encapsulation ATM
1382 packets input, 399282 bytes
1558 packets output, 205883 bytes
0 OAM cells input, 0 OAM cells output
AAL5 CRC errors : 280
AAL5 SAR Timeouts : 0
AAL5 Oversized SDUs : 0
AAL5 length violation : 210285
AAL5 CPI Error : 302
```

Both counters refer to ATM adaptation layer 5 (AAL5). They encapsulate routed or bridged protocol data units (PDUs) at the ATM stack's common part convergence sublayer (CPCS). RFC 1483 [\[1\]](#) defines the format of the AAL5 trailer, as illustrated in this diagram.



The two-byte length field in the AAL5 trailer indicates the size of the CPCS-PDU payload field. Two bytes is 16 bits or a maximum length value of 65,535 (2¹⁶) octets.

MTU defines the size of the Layer 3 datagram. An AAL5 service data unit (SDU) is defined as the Layer 3 datagram plus the optional Logical Link Control/Subnetwork Access Protocol (LLC/SNAP) header. An AAL5 PDU is defined as the combined AAL5 SDU plus the eight-byte AAL5 trailer. Therefore, an MTU of 9180 can produce an AAL5 SDU of 9180 bytes and an AAL5 PDU of 9188 bytes with the eight-byte AAL5 trailer.

When an ATM interface receives a packet larger than the MTU, the router increments the Oversized SDUs counter. The Oversize SDUs counter is defined in RFC 1695 [\[1\]](#).

```

aal5VccOverSizedSDUs OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The number of AAL5 CPCS PDUs discarded
        on this AAL5 VCC at the interface
        associated with an AAL5 entity because the
        AAL5 SDUs were too large."
    ::= { aal5VccEntry 5 }
  
```

RFC 1695 also supports the ability to set separate transmit and receive SDU sizes using these object IDs:

```

atmVccAal5CpcsTransmitSduSize OBJECT-TYPE
    SYNTAX INTEGER (1..65535)
    MAX-ACCESS read-create
    STATUS current
    DESCRIPTION
        "An instance of this object only exists when the
        local VCL end-point is also the VCC end-point,
        and AAL5 is in use.
        The maximum AAL5 CPCS SDU size in octets that is
        supported on the transmit direction of this VCC."
    DEFVAL { 9188 }
    ::= { atmVclEntry 9 }

atmVccAal5CpcsReceiveSduSize OBJECT-TYPE
    SYNTAX INTEGER (1..65535)
    MAX-ACCESS read-create
    STATUS current
  
```

DESCRIPTION

"An instance of this object only exists when the local VCL end-point is also the VCC end-point, and AAL5 is in use.
The maximum AAL5 CPCS SDU size in octets that is supported on the receive direction of this VCC."
DEFVAL { 9188 }
 ::= { atmVclEntry 10 }

ATM interfaces that follow RFC 1695 also increment the ifInErrors counter upon detecting oversized SDU errors. This is in addition to CRC-32 and SAR timeout errors, which are two counters also defined in the RFC.

A router increments the AAL5 length violation counter when the calculated size of a reassembled packet fails to match the received value of the AAL5 length field regardless of the MTU. To understand how these violations can occur, you need to understand how a receiving ATM interface recognizes the last cell of a frame.

A cell header includes a three-bit payload type identifier (PTI) field. These three bits signify:

- **Bit 1** Indicates whether the cell contains user data or management data.
- **Bit 2** Indicates whether the cell experiences congestion during transmission.
- **Bit 3** Indicates whether the cell is the final cell of a higher-layer data frame. When set to 1, this bit is called the end of marker (EOM).

PTI values of 001 or 011 mark the last cell of an AAL5 PDU and tell the receiving ATM interface to start reassembly. During periods of congestion or error conditions, an ATM link may drop the last cell. As a result, the receiving interface does not start reassembly until receiving the end of marker cell of the second AAL5 packet, producing a length violation.

In some cases, your router reports a large value for the AAL5 length violations counter and a much smaller value for the AAL5 CRC errors counter. This condition occurs when the ATM interface declares a length violation and drops a reassembled packet without bothering to check the CRC. An ATM interface checks the CRC only after it confirms that the packet size matches the AAL5 length field.

Benefits of Large and Same-Size MTUs

Using a consistent and max-sized MTU across multiple interfaces in your network offers these benefits:

- Reduces or eliminates fragmentation. Larger MTUs can enhance TCP performance by eliminating fragmentation. Therefore applications like Network File System (NFS) can take greater advantage of their large native MTUs of around 8 kB.
- Optimizes the size of the packet buffer pools carved in Packet memory (MEMD) on the route switch processor (RSP) on a Cisco 7500 series platform. On this platform, MTU plays an important role in buffer carving. Specifically, this platform uses a buffer-carving algorithm that creates four buffer pools based on MTU. If all interfaces use the same MTU, the router creates a large pool of same-sized buffers. Using large and widely varying MTUs on this platform forces the Cisco IOS® software to carve a small number of large buffers, possibly impacting other interfaces. On the 7500 series platform, adjusting the MTU can lead to a smaller number of ignored input errors. Refer to What Causes a "%RSP-3-RESTART: cbus complex"?

Note: Originally, the AIP supported an MTU as large as 9180. The reason requires an understanding of architecture. The ability of ATM interfaces to support the advertised maximum number of active simultaneous virtual circuits (VCs) is based on statistical multiplexing and on having enough packet buffers to perform some number of simultaneous reassemblies. Cisco limits the MTU size to roughly

9000 bytes on the AIP to support the advertised maximum active VCs value of 2000.

- Increases router performance by minimizing the number of packets processed. Most of the performance costs in routers relate to "packets handled", rather than "bytes transferred". A router typically processes transit packets in interrupt mode. A large MTU can result in higher performance since faster CPUs do not necessarily result in fast interrupt-intensive operations.

Relevant RFCs

This table lists requests for comment (RFCs) related to datagram sizes.

Note: All links in the table are RFC1483 [↗](#).

Request for Comment	Description
RFC 791 ↗	Defines IP fragmentation procedures.
RFC 1191 ↗ and RFC 1435 ↗	Define Path MTU discovery, a key mechanism for reducing IP fragmentation in the Internet. This mechanism is important because ATM uses default MTU sizes that are significantly different from other technologies like Ethernet and FDDI.
RFC 1209 ↗	Specifies an IP MTU over SMDS of 9180 octets. The Internet Engineering Task Force (IETF) used this value and RFC to set an MTU of 9180 octets for IP over ATM AAL5, as defined in RFC 2225 ↗ .
RFC 1626 ↗ and RFC 2225 ↗	Specify among other items that ATM interfaces must attempt to negotiate the AAL CPCS-SDU size using the ATM signaling protocol for switched virtual circuits (SVCs).

IP Fragmentation

[RFC 791](#) [↗](#) defines IP fragmentation and describes the procedure as "If the total length is less than or equal the maximum transmission unit then submit this datagram to the next step in datagram processing; otherwise cut the datagram into two fragments, the first fragment being the maximum size, and the second fragment being the rest of the datagram."

The **debug ip packet** *{host access-list}* command output captures a ping between the two hosts 192.168.1.51 and 192.168.1.254. For each packet, the router reports that it receives two fragments: one 1500 bytes in length and one 48 bytes in length.



Caution: Before issuing debug commands, please see Important Information on Debug Commands.

```
*Mar 28 09:59:27.002: IP: s=192.168.1.51 (ATM4/0.3), d=192.168.1.254, len 1500, rcvd 4
*Mar 28 09:59:27.002: IP: recv fragment from 192.168.1.51 offset 0 bytes
*Mar 28 09:59:27.002: IP: s=192.168.1.51 (ATM4/0.3), d=192.168.1.254, len 48, rcvd 4
*Mar 28 09:59:27.002: IP: recv fragment from 192.168.1.51 offset 1480 bytes
```

The router responds with an echo reply and reports that it is sending two fragments.

```
*Mar 28 09:59:27.002: ICMP: echo reply sent, src 192.168.1.254, dst 192.168.1.51
*Mar 28 09:59:27.002: IP: s=192.168.1.254 (local), d=192.168.1.51 (ATM4/0.3),
len 1528, sending
*Mar 28 09:59:27.002: IP: s=192.168.1.254 (local), d=192.168.1.51 (ATM4/0.3),
len 1500, sending fragment
*Mar 28 09:59:27.006: IP: s=192.168.1.254 (local), d=192.168.1.51 (ATM4/0.3),
len 48, sending last fragment
```

Jumbo Frame Support

Gigabit Ethernet interfaces on Cisco Catalyst 5000 and 6000 switches support jumbo frames, which have an MTU of 9,216 bytes. Support for jumbo frames for the Catalyst 6000 family ATM module (WS-X6101) is available as of Cisco IOS Software Release 12.1(10)E, as per the Release Notes.

Configuring the MTU size on the subinterface does not affect the maximum frame size that can be transferred on a Catalyst 6000 family ATM module. The maximum frame size (9218 bytes) is initialized when the module comes up and does not change when the MTU size changes using the CLI.

To bridge the jumbo frames, the feature should be enabled for the ATM module on the supervisor engine by using the **set port jumbo mod/port** command.

In Cisco IOS Software Releases earlier than 12.1(10)E, Catalyst ATM modules accept the MTU command at the command line and a maximum value of 9218 bytes. However, without jumbo frame support, this configuration change is misleading. The original lack of support for jumbo frames comes from the maximum number of buffers supported for any VC.

```
ATM#show interface atm0
ATM0 is down, line protocol is down
Hardware is Catalyst 5000 ATM
MTU 1584 bytes, sub MTU 0, BW 156250 Kbit, DLY 80 usec, rely 255/255,
load 1/255
Encapsulation ATM, loopback not set, keepalive not supported
Encapsulation(s): AAL5, PVC mode
4096 maximum active VCs, 1024 VCs per VP, 0 current VCCs
VC idle disconnect time: 300 seconds
Signaling vc = 1, vpi = 0, vci = 5
UNI Version = 3.1, Link Side = user
PHY Type : SINGLE PHY; Link Status: DOWN
[snip]
```

The LANE version 1 specification requires that a SETUP message include the AAL parameters information element (IE). In this IE, the calling party or source ATM interface must specify the Forward Maximum CPCS-SDU Size and the Backward Maximum CPCS-SDU Size. The supported AAL5 SDU Max octet values are 1516, 4544, 9234, and 18190. As of Cisco IOS Software Release 12.1(10)E, LECs can transfer frames up to 9218 bytes.

Jumbo frames support is already on the roadmap for the 8540 enhanced Gigabit Ethernet line cards. Such support is being investigated for the Gigabit Ethernet cards for the 8510. The ATM router module 2 (ARM2) for the 8540 now supports a configurable MTU size.

Troubleshoot

Complete these steps to narrow your troubleshooting if your symptoms point to a problem with datagram sizes.

1. Confirm the correct MTU is on the main interface and on the subinterface.

2. If pings above a certain packet size fail, the problem may be related to traffic shaping. Refer to Understanding the VBR—nrt Service Category and Traffic Shaping for ATM VCs. Confirm the packets exit the source router and/or enter the destination router with these commands:

◆ **debug ip packet** (host access—list only)



Caution: This debug can produce a large amount of output on a production output.

Take extra precautions when you enable this debug.

◆ **debug atm packet interface atm mod/port vpi vci**

◆ **debug atm errors**

3. Check for a non—zero value for the giants counter in the output of **show interface atm**. Does the giants counter increment with your pings?
4. Execute the **show buffers** command and look for non—zero values for the misses and failures counters. Determine whether the counters are incrementing, particularly when you ping the router and use the system buffers. Refer to Buffer Tuning for more information.

```
7500#show buffers
Buffer elements:
  499 in free list (500 max allowed)
  913677 hits, 0 misses, 0 created
Public buffer pools:
Small buffers, 104 bytes (total 480, permanent 480):
  474 in free list (20 min, 1000 max allowed)
  1036212 hits, 0 misses, 0 trims, 0 created
  0 failures (0 no memory)
Middle buffers, 600 bytes (total 360, permanent 360):
  358 in free list (20 min, 800 max allowed)
  635809 hits, 0 misses, 0 trims, 0 created
  0 failures (0 no memory)
Big buffers, 1524 bytes (total 360, permanent 360):
  360 in free list (10 min, 1200 max allowed)
  23457 hits, 0 misses, 0 trims, 0 created
  0 failures (0 no memory)
VeryBig buffers, 4520 bytes (total 40, permanent 40):
  40 in free list (5 min, 1200 max allowed)
  8969 hits, 0 misses, 0 trims, 0 created
  0 failures (0 no memory)
Large buffers, 5024 bytes (total 40, permanent 40):
  40 in free list (3 min, 120 max allowed)
  0 hits, 0 misses, 0 trims, 0 created
  0 failures (0 no memory)
Huge buffers, 18024 bytes (total 4, permanent 0):
  3 in free list (3 min, 52 max allowed)
  0 hits, 1 misses, 427 trims, 431 created
  0 failures (0 no memory)
```

5. Execute the **show ip interface atm** command and determine whether Cisco express forwarding (CEF) is enabled. If so, check the MTU size referenced in the adjacency entry to the destination.

```
router#show adj atm 5/0.1 interface
Protocol Interface Address
IP ATM5/0.1 point2point(6)
  0 packets, 0 bytes
  00040000
  AAAA030000000800
  CEF expires: 00:02:49
  refresh: 00:00:49
  ATM-PVC never
  Fast adjacency enabled
  IP redirect enabled
  IP mtu 4470 (0x0)
  Fixup disabled
```

Known Issue – MTU and Bridging

Cisco bug ID CSCdv42095 (registered customers only) resolves a problem with failing pings for packets larger than 1498 bytes when the MTU is configured to be less than 1502 bytes on a bridged interface. The changes allow the maximum packet size to be equal to the MTU plus the maximum ATM encapsulation in bytes. Set the MTU to 1502 as a workaround.

Related Information

- [ATM Technology Support Pages](#)
 - [Cisco ATM Port Adapter](#)
 - [ATM Acronyms](#)
 - [Technical Support – Cisco Systems](#)
-

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

© 2014 – 2015 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Dec 14, 2007

Document ID: 10479
