

Default IP MTU for use over ATM AAL5

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Default Value for IP MTU over ATM AAL5

Protocols in wide use throughout the Internet, such as the Network File System (NFS), currently use large frame sizes (e.g. 8 KB). Empirical evidence with various applications over the Transmission Control Protocol (TCP) indicates that larger Maximum Transmission Unit (MTU) sizes for the Internet Protocol (IP) tend to give better performance. Fragmentation of IP datagrams is known to be highly undesirable. [KM87] It is desirable to reduce fragmentation in the network and thereby enhance performance by having the IP Maximum Transmission Unit (MTU) for AAL5 be reasonably large. NFS defaults to an 8192 byte frame size. Allowing for RPC/XDR, UDP, IP, and LLC headers, NFS would prefer a default MTU of at least 8300 octets. Routers can sometimes perform better with larger packet sizes because most of the performance costs in routers relate to "packets handled" rather than "bytes transferred". So there are a number of good reasons to have a reasonably large default MTU value for IP over ATM AAL5.

RFC 1209 specifies the IP MTU over SMDS to be 9180 octets, which is larger than 8300 octets but still in the same range. [RFC-1209] There is no good reason for the default MTU of IP over ATM AAL5 to be different from IP over SMDS, given that they will be the same magnitude. Having the two be the same size will be helpful in interoperability and will also help reduce incidence of IP fragmentation.

Therefore, the default IP MTU for use with ATM AAL5 shall be 9180 octets. All implementations compliant and conformant with this specification shall support at least the default IP MTU value for use over ATM AAL5.

Permanent Virtual Circuits

Implementations which only support Permanent Virtual Circuits (PVCs) will (by definition) not implement any ATM signalling protocol. Such implementations shall use the default IP MTU value of 9180 octets unless both parties have agreed in advance to use some other IP MTU value via some mechanism not specified here.

Switched Virtual Circuits

Implementations that support Switched Virtual Circuits (SVCs) MUST attempt to negotiate the AAL CPCS-SDU size using the ATM signalling protocol. The industry standard ATM signalling protocol uses two different parts of the Information Element named "AAL Parameters" to exchange information on the MTU over the ATM circuit being setup [ATMF93a]. The Forward Maximum CPCS-SDU Size field contains the value over the path from the calling party to the called party. The Backwards Maximum CPCS-SDU Size Identifier field contains the value over the path from the called party to the calling party. The ATM Forum specifies the valid values of this identifier as 1 to 65535 inclusive. Note that the ATM Forum's User-to-Network-Interface (UNI) signalling permits the MTU in one direction to be different from the MTU in the opposite direction, so the Forward Maximum CPCS-SDU Size Identifier might have a different value from the Backwards Maximum CPCS-SDU Size Identifier on the same connection.

If the calling party wishes to use the default MTU it shall still include the "AAL Parameters" information element with the default values for the Maximum CPCS-SDU Size as part of the SETUP message of the ATM signalling protocol [ATMF93b]. If the calling party desires to use a different value than the default, it shall include the "AAL Parameters" information element with the desired value for the Maximum CPCS-SDU Size as part of the SETUP message of the ATM Signalling Protocol. The called party will respond using the same information elements and identifiers in its CONNECT message response [ATMF93c].

If the called party receives a SETUP message containing the "Maximum CPCS-SDU Size" in the AAL Parameters information element, it shall handle the Forward and Backward Maximum CPCS-SDU Size Identifier as follows:

- a) If it is able to accept the ATM MTU values proposed by the SETUP message, it shall include an AAL Parameters information element in its response. The Forward and Backwards Maximum CPCS-SDU Size fields shall be present and their values shall be equal to the corresponding values in the SETUP message.

- b) If it wishes a smaller ATM MTU size than that proposed, then it shall set the values of the Maximum CPCS-SDU Size in the AAL Parameters information elements equal to the desired value in the CONNECT message responding to the original SETUP message.
- c) If the calling endpoint receives a CONNECT message that does not contain the AAL Parameters Information Element, but the corresponding SETUP message did contain the AAL Parameters Information Telement (including the forward and backward CPCS-SDU Size fields), it shall clear the call with cause "AAL Parameters cannot be supported".
- d) If either endpoint receives a STATUS message with cause "Information Element Non-existent or Not Implemented" or cause "Access Information Discarded", and with a diagnostic field indicating the AAL Parameters Information Element identifier, it shall clear the call with cause "AAL Parameters cannot be supported."
- e) If either endpoint receives CPCS-SDUs in excess of the negotiated MTU size, it may use IP fragmentation or may clear the call with cause "AAL Parameters cannot be supported". In this case, an error has occurred either due to a fault in an end system or in the ATM network. The error should be noted by ATM network management for human examination and intervention.

If the called endpoint incorrectly includes the Forward and Backward Maximum CPCS-SDU Size fields in the CONNECT messages (e.g. because the original SETUP message did not include these fields) or it sets these fields to an invalid value, then the calling party shall clear the call with cause "Invalid Information Element Contents".

Path MTU Discovery Required

The Path MTU Discovery mechanism is an Internet Standard [RFC-1191] and is an important mechanism for reducing IP fragmentation in the Internet. This mechanism is particularly important because new subnet ATM uses a default MTU sizes significantly different from older subnet technologies such as Ethernet and FDDI.

In order to ensure good performance throughout the Internet and also to permit IP to take full advantage of the potentially larger IP datagram sizes supported by ATM, all routers implementations that comply or conform with this specification must also implement the IP Path MTU Discovery mechanism as defined in RFC-1191 and clarified by RFC-1435. Host implementations should implement the IP Path MTU Discovery mechanism as defined in RFC-1191.

Applicability Statement

The Multiprotocol Encapsulation over ATM AAL5 defined in RFC-1483 is not specific to any model of IP and ATM interaction. [RFC-1483] Similarly, this specification is general enough to apply to all models for use of IP over ATM AAL5. Use of this specification is recommended for all implementations of IP over ATM AAL5 in order to increase interoperability and performance. This specification does not conflict with the Classical IP over ATM specification and may be used as a conforming extension to that specification. [RFC-1577] Applicability of this draft is not limited to the Classical IP over ATM model.

Security Considerations

Security issues are not discussed in this memo.

References

[RFC-791] Postel, J., "Internet Protocol - DARPA Internet Program Protocol Specification", STD 5, RFC 791, DARPA, September 1981.

[RFC-793] Postel, J., "Transmission Control Protocol - DARPA Internet Program Protocol Specification", STD 7, RFC 793, DARPA, September 1981.

[RFC-1122] Braden, R., Editor, Requirements for Internet Hosts -- Communications Layers, STD 3, RFC 1122, USC/Information Sciences Institute, October 1989, pp.58-60.

[RFC-1191] Mogul, J., and S. Deering, "Path MTU Discovery", RFC 1191, DECWRL, Stanford University, November 1990.

[RFC-1209] Piscitello, D., and J. Lawrence, "The Transmission of IP Datagrams over the SMDS Service", RFC 1209, Bell Communications Research, March 1991.

[RFC-1435] Knowles, S., "IESG Advice from Experience with Path MTU Discovery", RFC-1435, IESG, March 1993.

[RFC-1483] Heinanen, J., "Multiprotocol Encapsulation over ATM Adaptation Layer 5", RFC 1483, Telecom Finland, July 1993.

[RFC-1577] Laubach, M., "Classical IP and ARP over ATM", RFC 1577, Hewlett-Packard Laboratories, January 1994.

[ATMF93a] Breault, R., Grace, J., Jaeger, J., and L. Wojnaroski, Editors, "ATM Forum User Network Interface Specification", Version 3.0, Section 5.4.5.5, p. 194-200, 10 September 1993, ATM Forum.

[ATMF93b] Breault, R., Grace, J., Jaeger, J., and L. Wojnaroski, Editors, "ATM Forum User Network Interface Specification", Version 3.0, Section 5.3.1.7, p. 171-172, 10 September 1993, ATM Forum.

[ATMF93c] Breault, R., Grace, J., Jaeger, J., and L. Wojnaroski, Editors, "ATM Forum User Network Interface Specification", Version 3.0, Section 5.3.1.3, p. 168, 10 September 1993, ATM Forum.

[KM87] Kent C., and J. Mogul, "Fragmentation Considered Harmful", Proceedings of the ACM SIGCOMM '87 Workshop on Frontiers in Computer Communications Technology, August 1987.

Acknowledgements

While all members of the IETF's IP over ATM Working Group have been helpful, Vern Schryver, Rob Warnock, Craig Partridge, Subbu Subramaniam, and Bryan Lyles have been especially helpful to the author in analysing the host and routing implications of the default IP MTU value. Similarly, Dan Grossman provided significant review and help in ensuring alignment of this text with the related work in the ATM Forum and ITU.

Disclaimer

Author's organisation provided for identification purposes only. This document presents the author's views and is not necessarily the official opinion of his employer.

Author's Address

Randall J. Atkinson
Information Technology Division
Naval Research Laboratory
Washington, DC 20375-5320
USA

EEmail: atkinson@itd.nrl.navy.mil