

A Method for the Transmission of IPv6 Packets over FDDI Networks

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.

Introduction

This memo specifies the MTU and frame format for transmission of IPv6 [IPV6] packets on FDDI networks, including a method for MTU determination in the presence of 802.1d bridges to other media. It also specifies the method of forming IPv6 link-local addresses on FDDI networks and the content of the Source/Target Link-layer Address option used in the Router Solicitation, Router Advertisement, Neighbor Solicitation, and Neighbor Advertisement messages described in [DISC], when those messages are transmitted on an FDDI network.

Maximum Transmission Unit


FDDI permits a frame length of 4500 octets (9000 symbols), including at least 22 octets (44 symbols) of Data Link encapsulation when long-format addresses are used. Subtracting 8 octets of LLC/SNAP header, this would, in principle, allow the IPv6 packet in the Information field to be up to 4470 octets. However, it is desirable to allow for the variable sizes and possible future extensions to the MAC header and frame status fields. The default MTU size for IPv6 packets on an FDDI network is therefore 4352 octets. This size may be reduced by a Router Advertisement [DISC] containing an MTU option which specifies a smaller MTU, or by manual configuration of a smaller value on each node. If a Router Advertisement is received with an MTU option specifying an MTU larger than the default or the manually configured value, that MTU option may be logged to system management but must be otherwise ignored.

For purposes of this document, information received from DHCP is considered "manually configured".

Frame Format

FDDI provides both synchronous and asynchronous transmission, with the latter class further subdivided by the use of restricted and unrestricted tokens. Only asynchronous transmission with unrestricted tokens is required for FDDI interoperability. Accordingly, IPv6 packets shall be sent in asynchronous frames using unrestricted tokens. The robustness principle dictates that nodes should be able to receive synchronous frames and asynchronous frames sent using restricted tokens.

IPv6 packets are transmitted in LLC/SNAP frames, using long-format (48 bit) addresses. The data field contains the IPv6 header and payload and is followed by the FDDI Frame Check Sequence, Ending Delimiter, and Frame Status symbols.



FDDI Header Fields:

- FC** The Frame Code must be in the range 50 to 57 hexadecimal, inclusive, with the three low order bits indicating the frame priority. The Frame Code should be in the range 51 to 57 hexadecimal, inclusive, for reasons given in the next section.
- DSAP, SSAP** Both the DSAP and SSAP fields shall contain the value AA hexadecimal, indicating SNAP encapsulation.
- CTL** The Control field shall be set to 03 hexadecimal, indicating Unnumbered Information.
- OUI** The Organizationally Unique Identifier shall be set to 000000 hexadecimal.

Ethertype The ethernet protocol type ("ethertype") shall be set to the value 86DD hexadecimal.

Interaction with Bridges

802.1d MAC bridges which connect different media, for example Ethernet and FDDI, have become very widespread. Some of them do IPv4 packet fragmentation and/or support IPv4 Path MTU discovery [PMTU], many others do not, or do so incorrectly. Use of IPv6 in a bridged mixed-media environment should not depend on support from MAC bridges.

For correct operation when mixed media are bridged together, the smallest MTU of all the media must be advertised by routers in an MTU option. If there are no routers present, this MTU must be manually configured in each node which is connected to a medium with larger default MTU. Multicast packets on such a bridged network must not be larger than the smallest MTU of any of the bridged media. Often, the subnetwork topology will support larger unicast packets to be exchanged between certain pairs of nodes. To take advantage of high-MTU paths when possible, nodes transmitting IPv6 on FDDI should implement the following simple mechanism for "FDDI adjacency detection".

A node which implements FDDI adjacency detection and has it enabled on an FDDI interface must set a non-zero LLC priority in all Neighbor Advertisement, Neighbor Solicitation and, if applicable, Router Advertisement frames transmitted on that interface. (In IEEE 802 language, the user_priority parameter of the M_UNITDATA.request primitive must not be zero.) If FDDI adjacency detection has been disabled on an FDDI interface, the priority field of those frames must be zero.

Note that an IPv6 frame which originated on an Ethernet, or traversed an Ethernet, before being translated by an 802.1d bridge and delivered to a node's FDDI interface will have zero in the priority field, as required by [BRIDGE]. (There's a fine point here: a conforming bridge may provide a management-settable Outbound User Priority parameter for each port. However, the author is unaware of any product that provides this optional capability and, in any case, the default value for the parameter is zero.)

If a node N1 receives, in an FDDI frame with a non-zero LLC priority, a valid Router Advertisement, Neighbor Advertisement, or Neighbor Solicitation from a node N2, then N1 may send unicast IPv6 packets to N2 with sizes up to the default IPv6 FDDI MTU (4352 octets), regardless of any smaller MTU configured manually or received in a Router Advertisement MTU option. N2 may be the IPv6 destination or the next hop router to the destination.

Nodes implementing FDDI adjacency detection must provide a configuration option to disable the mechanism. This option may be used when a smaller MTU is desired for reasons other than mixed-media bridging. By default, FDDI adjacency detection should be enabled.

The only contemplated use of the LLC priority field of the FC octet is to aid in per-destination MTU determination. It would be sufficient for that purpose to require only that Router Advertisements, Neighbor Advertisements, and Neighbor Solicitations sent on FDDI always have non-zero priority. However, it may be simpler or more useful to transmit all IPv6 packets on FDDI with non-zero priority.

Stateless Autoconfiguration and Link-Local Addresses

The address token [CONF] for an FDDI interface is the interface's built-in 48-bit IEEE 802 address, in canonical bit order and with the octet in the same order in which they would appear in the header of an ethernet frame. (The individual/group bit is in the first octet and the OUI is in the first three octets.) A different MAC address set manually or by software should not be used as the address token.

An IPv6 address prefix used for stateless autoconfiguration of an FDDI interface must be 80 bits in length.

The IPv6 Link-local address [AARCH] for an FDDI interface is formed by appending the interface's IEEE 802 address to the 80-bit prefix FE80::.

```

+-----+-----+-----+-----+-----+-----+-----+-----+
|  FE      80      00      00      00      00      00      00  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|  00      00      |                               FDDI Address  |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Address Mapping -- Unicast

The procedure for mapping IPv6 addresses into FDDI link-layer addresses is described in [DISC]. The Source/Target Link-layer Address option has the following form when the link layer is FDDI.

```

+-----+-----+-----+-----+-----+-----+-----+
| Type  | Length |                               FDDI Address                               |
+-----+-----+-----+-----+-----+-----+-----+

```

Option fields:

Type 1 for Source Link-layer address.
 2 for Target Link-layer address.

Length 1 (in units of 8 octets).

FDDI Address

The 48 bit FDDI IEEE 802 address, in canonical bit order.
 This is the address the interface currently responds to, and
 may be different from the built-in address used as the
 address token.

Address Mapping -- Multicast

An IPv6 packet with a multicast destination address DST is
 transmitted to the FDDI multicast address whose first two octets are
 the value 3333 hexadecimal and whose last four octets are the last
 four octets of DST, ordered from more to least significant.

```

+-----+-----+-----+-----+-----+-----+
|  33   |  33   | DST13 | DST14 | DST15 | DST16 |
+-----+-----+-----+-----+-----+-----+

```

Security Considerations

Security considerations are not addressed in this memo.

Acknowledgments

Erik Nordmark contributed to the method for interaction with bridges.

References

- [AARCH] Hinden, and S. Deering, "IP Version 6 Addressing Architecture", RFC 1884, December 1995.
- [BRIDGE] ISO/IEC 10038 : 1993 [ANSI/IEEE Std 802.1D] Media access control (MAC) bridges.
- [CONF] Thomson, S., and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 1971, August 1996.
- [DISC] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 1970, August 1996.
- [IPV6] Deering, S., and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 1883, August 1996.
- [PMTU] Mogul, J., and S. Deering, "Path MTU Discovery", RFC 1191, November 1990.

Author's Address

Matt Crawford
Fermilab MS 368
PO Box 500
Batavia, IL 60510
USA

Phone: +1 708 840-3461

EMail: crawdad@fnal.gov

