

Network Working Group
Request for Comments: 1188
Obsoletes: RFC 1103

D. Katz
Merit/NSFNET
October 1990

A Proposed Standard for the Transmission of IP Datagrams over FDDI Networks

Status of this Memo

This memo defines a method of encapsulating the Internet Protocol (IP) datagrams and Address Resolution Protocol (ARP) requests and replies on Fiber Distributed Data Interface (FDDI) Networks. This RFC specifies a protocol on the IAB Standards Track for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "IAB Official Protocol Standards" for the standardization state and status of this protocol.

This proposal is the product of the IP over FDDI Working Group of the Internet Engineering Task Force (IETF). Comments on this memo should be submitted to the IETF IP over FDDI Working Group Chair. Distribution of this memo is unlimited.

Abstract

This document specifies a method for the use of IP and ARP on FDDI networks. The encapsulation method used is described, as well as various media-specific issues.

Acknowledgments

This memo draws heavily in both concept and text from RFC 1042 [3], written by Jon Postel and Joyce K. Reynolds of USC/Information Sciences Institute. The author would also like to acknowledge the contributions of the IP Over FDDI Working Group of the IETF, members of ANSI ASC X3T9.5, and others in the FDDI community.

Conventions

The following language conventions are used in the items of specification in this document:

"Must," "Shall," or "Mandatory"--the item is an absolute requirement of the specification.

"Should" or "Recommended"--the item should generally be followed for all but exceptional circumstances.


"May" or "Optional"--the item is truly optional and may be followed or ignored according to the needs of the implementor.

Introduction

The goal of this specification is to allow compatible and interoperable implementations for transmitting IP datagrams [1] and ARP requests and replies [2].

The Fiber Distributed Data Interface (FDDI) specifications define a family of standards for Local Area Networks (LANs) that provides the Physical Layer and Media Access Control Sublayer of the Data Link Layer as defined by the ISO Open System Interconnection Reference Model (ISO/OSI). Documents are in various stages of progression toward International Standardization for Media Access Control (MAC) [4], Physical Layer Protocol (PHY) [5], Physical Layer Medium Dependent (PMD) [6], and Station Management (SMT) [7]. The family of FDDI standards corresponds to the IEEE 802 MAC layer standards [8, 9, 10].

The remainder of the Data Link Service is provided by the IEEE 802.2 Logical Link Control (LLC) service [11]. The resulting stack of services appears as follows:




This memo describes the use of IP and ARP in this environment. At this time, it is not necessary that the use of IP and ARP be consistent between FDDI and IEEE 802 networks, but it is the intent of this memo not to preclude Data Link Layer interoperability at such time as the standards define it.

The FDDI standards define both single and dual MAC stations. This document describes the use of IP and ARP on single MAC stations (single-attach or dual-attach) only. Operation on dual MAC stations will be described in a forthcoming document.

Packet Format

IP datagrams and ARP requests and replies sent on FDDI networks shall be encapsulated within the 802.2 LLC and Sub-Network Access Protocol (SNAP) [12] data link layers and the FDDI MAC and physical layers. The SNAP must be used with an Organization Code indicating that the SNAP header contains the EtherType code (as listed in Assigned Numbers [13]).

802.2 LLC Type 1 communication (which must be implemented by all conforming 802.2 stations) is used exclusively. All frames must be transmitted in standard 802.2 LLC Type 1 Unnumbered Information format, with the DSAP and the SSAP fields of the 802.2 header set to the assigned global SAP value for SNAP [11]. The 24-bit Organization Code in the SNAP must be zero, and the remaining 16 bits are the EtherType from Assigned Numbers [13] (IP = 2048, ARP = 2054).



The total length of the LLC Header and the SNAP header is 8 octets.

The K1 value is 170 (decimal).

The K2 value is 0 (zero).

The control value is 3 (Unnumbered Information).

Address Resolution

The mapping of 32-bit Internet addresses to 48-bit FDDI addresses shall be done via the dynamic discovery procedure of the Address Resolution Protocol (ARP) [2].

Internet addresses are assigned arbitrarily on Internet networks. Each host's implementation must know its own Internet address and respond to Address Resolution requests appropriately. It must also use ARP to translate Internet addresses to FDDI addresses when

needed.

The ARP protocol has several fields that parameterize its use in any specific context [2]. These fields are:

hrd	16 - bits	The Hardware Type Code
pro	16 - bits	The Protocol Type Code
hln	8 - bits	Octets in each hardware address
pln	8 - bits	Octets in each protocol address
op	16 - bits	Operation Code

The hardware type code assigned for IEEE 802 networks is 6 [13]. The hardware type code assigned for Ethernet networks is 1 [13]. Unfortunately, differing values between Ethernet and IEEE 802 networks cause interoperability problems in bridged environments. In order to not preclude interoperability with Ethernets in a bridged environment, ARP packets shall be transmitted with a hardware type code of 1. Furthermore, ARP packets shall be accepted if received with hardware type codes of either 1 or 6.

The protocol type code for IP is 2048 [13].

The hardware address length is 6.

The protocol address length (for IP) is 4.

The operation code is 1 for request and 2 for reply.

In order to not preclude interoperability in a bridged environment, the hardware addresses in ARP packets (ar\$sha, ar\$tha) must be carried in "canonical" bit order, with the Group bit positioned as the low order bit of the first octet. As FDDI addresses are normally expressed with the Group bit in the high order bit position, the addresses must be bit-reversed within each octet.

Although outside the scope of this document, it is recommended that MAC addresses be represented in canonical order in all Network Layer protocols carried within the information field of an FDDI frame.

Broadcast Address

The broadcast Internet address (the address on that network with a host part of all binary ones) must be mapped to the broadcast FDDI address (of all binary ones) (see [14]).

Multicast Support

A method of supporting IP multicasting is specified in [15]. This

method shall be used in FDDI networks if IP multicasting is to be supported. The use of this method may require the ability to copy frames addressed to any one of an arbitrary number of multicast (group) addresses.

An IP multicast address is mapped to an FDDI group address by placing the low order 23 bits of the IP address into the low order 23 bits of the FDDI group address 01-00-5E-00-00-00 (in "canonical" order). [See 13, page 20.]

For example, the IP multicast address:

224.255.0.2

maps to the FDDI group address:

01-00-5E-7F-00-02

in which the multicast (group) bit is the low order bit of the first octet (canonical order). When bit-reversed for transmission in the destination MAC address field of an FDDI frame (native order), it becomes:

80-00-7A-FE-00-40

that is, with the multicast (group) bit as the high order bit of the first octet, that being the first bit transmitted on the medium.

Trailer Formats

Some versions of Unix 4.x bsd use a different encapsulation method in order to get better network performance with the VAX virtual memory architecture. Hosts directly connected to FDDI networks shall not use trailers.

Byte Order

As described in Appendix B of the Internet Protocol specification [1], the IP datagram is transmitted over FDDI networks as a series of 8-bit bytes. This byte transmission order has been called "big-endian" [16].

MAC Layer Details

Packet Size

The FDDI MAC specification [4] defines a maximum frame size of 9000 symbols (4500 octets) for all frame fields, including four

symbols (two octets) of preamble. This leaves roughly 4470 octets for data after the LLC/SNAP header is taken into account.

However, in order to allow future extensions to the MAC header and frame status fields, it is desirable to reserve additional space for MAC overhead.

Therefore, the MTU of FDDI networks shall be 4352 octets. This provides for 4096 octets of data and 256 octets of headers at the network layer and above. Implementations must not send packets larger than the MTU.

Gateway implementations must be prepared to accept packets as large as the MTU and fragment them when necessary. Gateway implementations should be able to accept packets as large as can be carried within a maximum size FDDI frame and fragment them.

Host implementations should be prepared to accept packets as large as the MTU; however, hosts must not send datagrams longer than 576 octets unless they have explicit knowledge that the destination is prepared to accept them. Host implementations may accept packets as large as can be carried within a maximum size FDDI frame. A host may communicate its size preference in TCP-based applications via the TCP Maximum Segment Size option [17].

Datagrams on FDDI networks may be longer than the general Internet default maximum packet size of 576 octets. Hosts connected to an FDDI network should keep this in mind when sending datagrams to hosts that are not on the same local network. It may be appropriate to send smaller datagrams to avoid unnecessary fragmentation at intermediate gateways. Please see [17] for further information.

There is no minimum packet size restriction on FDDI networks.

In order to not preclude interoperability with Ethernet in a bridged environment, FDDI implementations must be prepared to receive (and ignore) trailing pad octets.

Other MAC Layer Issues

The FDDI MAC specification does not require that 16-bit and 48-bit address stations be able to interwork fully. It does, however, require that 16-bit stations have full 48-bit functionality, and that both types of stations be able to receive frames sent to either size broadcast address. In order to avoid interoperability problems, only 48-bit addresses shall be used with IP and ARP.

The FDDI MAC specification defines two classes of LLC frames, Asynchronous and Synchronous. Asynchronous frames are further controlled by a priority mechanism and two classes of token, Restricted and Unrestricted. Only the use of Unrestricted tokens and Asynchronous frames are required by the standard for FDDI interoperability.

All IP and ARP frames shall be transmitted as Asynchronous LLC frames using Unrestricted tokens, and the Priority value is a matter of local convention. Implementations should make the priority a tunable parameter for future use. It is recommended that implementations provide for the reception of IP and ARP packets in Synchronous frames, as well as Restricted Asynchronous frames.

After packet transmission, FDDI provides Frame Copied (C) and Address Recognized (A) indicators. The use of these indicators is a local implementation decision. Implementations may choose to perform link-level retransmission, ARP cache entry invalidation, etc., based on the values of these indicators and other information. The semantics of these indicators, especially in the presence of bridges, are not well defined as of this writing. Implementors are urged to follow the work of ANSI ASC X3T9.5 in regard to this issue in order to avoid interoperability problems.

IEEE 802.2 Details

While not necessary for supporting IP and ARP, all implementations must support IEEE 802.2 standard Class I service in order to be compliant with 802.2. Described below is the minimum functionality necessary for a conformant station. Some of the functions are not related directly to the support of the SNAP SAP (e.g., responding to XID and TEST commands directed to the null or global SAP addresses), but are part of a general LLC implementation. Implementors should consult IEEE Std. 802.2 [11] for details.

802.2 Class I LLC requires the support of Unnumbered Information (UI) Commands, eXchange IDentification (XID) Commands and Responses, and TEST link (TEST) Commands and Responses. Stations need not be able to transmit XID and TEST commands, but must be able to transmit responses.

Encodings

Command frames are identified by having the low order bit of the SSAP address reset to zero. Response frames have the low order bit of the SSAP address set to one.

The UI command has an LLC control field value of 3.

The XID command/response has an LLC control field value of 175 (decimal) if the Poll/Final bit is off or 191 (decimal) if the Poll/Final bit is on.

The TEST command/response has an LLC control field value of 227 (decimal) if the Poll/Final bit is off or 243 (decimal) if the Poll/Final bit is on.

Elements of Procedure

UI responses and UI commands with the Poll bit set shall be ignored. UI commands having other than the SNAP SAP in the DSAP or SSAP fields shall not be processed as IP or ARP packets.

When an XID or TEST command is received, an appropriate response must be returned. XID and TEST commands must be responded to only if the DSAP is the SNAP SAP (170 decimal), the Null SAP (0 decimal), or the Global SAP (255 decimal). XID and TEST commands received with other DSAP values must not be responded to unless the station supports the addressed service. Responses to XID and TEST frames shall be constructed as follows:

Destination MAC: Copied from Source MAC of the command
Source MAC: Set to the address of the MAC receiving the command
DSAP: Copied from SSAP of the command
SSAP: Set to 171 decimal (SNAP SAP + Response bit) if the DSAP in the command was the SNAP SAP or the Global SAP; set to 1 decimal (Null SAP + Response bit) if the DSAP in the command was the Null SAP

When responding to an XID or a TEST command, the value of the Final bit in the response must be copied from the value of the Poll bit in the command.

XID response frames must include an 802.2 XID Information field of 129.1.0 indicating Class I (connectionless) service.

TEST response frames must echo the information field received in the corresponding TEST command frame.

Appendix on Numbers

The IEEE specifies numbers as bit strings with the least significant bit first, or bit-wise little-endian order. The Internet protocols are documented in bit-wise big-endian order. This may cause some confusion about the proper values to use for numbers. Here are the conversions for some numbers of interest.

Number	IEEE Binary	Internet Binary	Internet Decimal
UI	11000000	00000011	3
SAP for SNAP	01010101	10101010	170
Global SAP	11111111	11111111	255
Null SAP	00000000	00000000	0
XID	11110101	10101111	175
XID Poll/Final	11111101	10111111	191
XID Info			129.1.0
TEST	11000111	11100011	227
TEST Poll/Final	11001111	11110011	243

References

- [1] Postel, J., "Internet Protocol", RFC 791, USC/Information Sciences Institute, September 1981.
- [2] Plummer, D., "An Ethernet Address Resolution Protocol - or - Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", RFC 826, MIT, November 1982.
- [3] Postel, J., and Reynolds, J., "A Standard for the Transmission of IP Datagrams over IEEE 802 Networks", RFC 1042, USC/Information Sciences Institute, February 1988.
- [4] ISO, "Fiber Distributed Data Interface (FDDI) - Media Access Control", ISO 9314-2, 1989. See also ANSI X3.139-1987.
- [5] ISO, "Fiber Distributed Data Interface (FDDI) - Token Ring Physical Layer Protocol", ISO 9314-1, 1989. See also ANSI X3.148-1988.
- [6] ISO, "Fiber Distributed Data Interface (FDDI) - Physical Layer Medium Dependent", ISO DIS 9314-3, 1989. See also ANSI X3.166-199x.
- [7] ANSI, "FDDI Station Management", ANSI X3T9.5/84-49 Rev 6.0, 1990.

- [8] IEEE, "IEEE Standards for Local Area Networks: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications", IEEE, New York, New York, 1985.
- [9] IEEE, "IEEE Standards for Local Area Networks: Token-Passing Bus Access Method and Physical Layer Specification", IEEE, New York, New York, 1985.
- [10] IEEE, "IEEE Standards for Local Area Networks: Token Ring Access Method and Physical Layer Specifications", IEEE, New York, New York, 1985.
- [11] IEEE, "IEEE Standards for Local Area Networks: Logical Link Control", IEEE, New York, New York, 1985.
- [12] IEEE, "Draft Standard P802.1A--Overview and Architecture", 1989.
- [13] Reynolds, J.K., and J. Postel, "Assigned Numbers", RFC 1060, USC/Information Sciences Institute, March 1990.
- [14] Braden, R., and J. Postel, "Requirements for Internet Gateways", RFC 1009, USC/Information Sciences Institute, June 1987.
- [15] Deering, S., "Host Extensions for IP Multicasting", RFC 1112, Stanford University, August 1989.
- [16] Cohen, D., "On Holy Wars and a Plea for Peace", Computer, IEEE, October 1981.
- [17] Postel, J., "The TCP Maximum Segment Size Option and Related Topics", RFC 879, USC/Information Sciences Institute, November 1983.

Security Considerations

Security issues are not discussed in this memo.

Author's Address

Dave Katz
Merit/NSFNET
1075 Beal Ave.
Ann Arbor, MI 48109

Phone: (313) 763-4898

EMail: dkatz@merit.edu

