

Transmission of IPv6 Packets over IEEE 1394 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.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

This document describes the frame format for transmission of IPv6 packets and the method of forming IPv6 link-local addresses and statelessly autoconfigured addresses on IEEE1394 networks.

1. INTRODUCTION

IEEE Std 1394-1995 (and its amendment) is a standard for a High Performance Serial Bus. IETF IP1394 Working Group has standardized the method to carry IPv4 datagrams and ARP packets over IEEE1394 subnetwork [IP1394].

This document describes the frame format for transmission of IPv6 [IPv6] packets and the method of forming IPv6 link-local addresses and statelessly autoconfigured addresses on IEEE1394 networks. It also describes the content of the Source/Target Link-layer Address option used in Neighbor Discovery [DISC] when the messages are transmitted on an IEEE1394 network.

2. SPECIFICATION TERMINOLOGY

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

3. IPv6-CAPABLE NODES

An IPv6-capable node MUST fulfill the following minimum requirements:

- it MUST implement configuration ROM in the general format specified by ISO/IEC 13213:1994 and MUST implement the bus information block specified by IEEE Std 1394a-2000 [1394a] and a unit directory specified by this document;
- the max_rec field in its bus information block MUST be at least 8; this indicates an ability to accept block write requests and asynchronous stream packets with data payload of 512 octets. The same ability MUST also apply to read requests; that is, the node MUST be able to transmit a block response packet with a data payload of 512 octets;
- it MUST be isochronous resource manager capable, as specified by IEEE Std 1394a-2000;
- it MUST support both reception and transmission of asynchronous streams as specified by IEEE Std 1394a-2000.

4. LINK ENCAPSULATION AND FRAGMENTATION

The encapsulation and fragmentation mechanism MUST be the same as "4. LINK ENCAPSULATION AND FRAGMENTATION" of [IP1394].

Note: Since there is an ether_type field to discriminate protocols and MCAP (multicast channel allocation protocol) is used for both IPv4 and IPv6, the version field in GASP (global asynchronous stream packet) header of IPv6 datagrams is the same value (one) as [IP1394].

The ether_type value for IPv6 is 0x86dd.

The default MTU size for IPv6 packets on an IEEE1394 network is 1500 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 each node. If a Router Advertisement received on an IEEE1394 interface has an MTU option specifying an MTU larger than 1500, or larger than a manually configured value, that MTU option may be logged to system management but MUST be otherwise ignored. The mechanism to extend MTU size between particular two nodes is for further study.

5. CONFIGURATION ROM

Configuration ROM for IPv6-capable nodes MUST contain a unit directory in the format specified by [IP1394] except following rules.

- The value for Unit_SW_Version is 0x000002.
- The textual descriptor for the Unit_SW_Version MUST be "IPv6".

Note: A dual-stack (IPv4 and IPv6) node will have two unit directories for IPv4 and IPv6 respectively.

6. STATELESS AUTOCONFIGURATION

The Interface Identifier [AARCH] for an IEEE1394 interface is formed from the interface's built-in EUI-64 identifier by complementing the "Universal/Local" (U/L) bit, which is the next-to-lowest order bit of the first octet of the EUI-64 identifier. Complementing this bit will generally change a 0 value to a 1, since an interface's built-in EUI-64 identifier is expected to be from a universally administered address space and hence have a globally unique value. A universally administered EUI-64 identifier is signified by a 0 in the U/L bit position, while a globally unique IPv6 Interface Identifier is signified by a 1 in the corresponding position. For further discussion on this point, see [AARCH].

An IPv6 address prefix used for stateless autoconfiguration [ACONF] of an IEEE1394 interface MUST have a length of 64 bits.

7. LINK-LOCAL ADDRESSES

The IPv6 link-local address [AARCH] for an IEEE1394 interface is formed by appending the Interface Identifier, as defined above, to the prefix FE80::/64.

10 bits	54 bits	64 bits
1111111010	(zeros)	Interface Identifier


8. ADDRESS MAPPING FOR UNICAST

The procedure for mapping IPv6 unicast addresses into IEEE1394 link-layer addresses uses the Neighbor Discovery [DISC]. Since 1394 link address (node_ID) will not be constant across a 1394 bridge, we have chosen not to put it in the Link-layer Address option. The recipient of the Neighbor Discovery SHOULD use the source_ID (obtained from either the asynchronous packet header or the GASP header) in

conjunction with the content of the Source link-layer address. An implementation MAY use some other methods to obtain a node_ID of the sender utilizing a mapping table between node_unique_ID (EUI-64 identifier) and node_ID. The mechanism to make such mapping table is out of scope of this document.

The recipient of an Neighbor Discovery packet MUST ignore it unless the most significant ten bits of the source_ID are equal to either 0x3FF or the most significant ten bits of the recipient's NODE_IDS register.

The Source/Target Link-layer Address option has the following form when the link layer is IEEE1394.



Type 1 for Source Link-layer address.

 2 for Target Link-layer address.

Length 3 (in units of 8 octets).

node_unique_ID This field contains the node unique ID of the node and MUST be equal to that specified in the node's configuration ROM.

max_rec This field MUST be equal to the value of max_rec in the node's configuration ROM.

spd This field MUST be set to the lesser of the node's link speed and PHY speed. The link speed is the maximum speed at which the link may send or receive packets; the PHY speed is the maximum speed at which the PHY may send, receive or repeat packets. The encoding used for spd is specified in the Table 2 of [IP1394].

unicast_FIFO	This field MUST specify the 48-bit offset of the node's FIFO available for the receipt of IPv6 datagrams. The offset of a node's unicast FIFO MUST NOT change, except as the result of a power reset.
reserved	This field MUST be set to all zeros by the sender and ignored by the receiver.

Note that node_ID may change when 1394 bus-reset occurs. The mapping cache held in the node SHOULD be cleared on 1394 bus-reset.

According to [1394], the maximum data payload and the transmission speed SHOULD be determined based on the sender's capability, the recipient's capability, and the PHYs of all intervening nodes.

9. IPv6 MULTICAST

By default, all best-effort IPv6 multicast MUST use asynchronous stream packets whose channel number is equal to the channel field from the BROADCAST_CHANNEL register. In particular, datagrams addressed to all-nodes multicast addresses, all-routers multicast addresses, and solicited-node multicast addresses [AARCH] MUST use the default channel specified by the BROADCAST_CHANNEL register.

Best-effort IPv6 multicast for other multicast group addresses may utilize a different channel number if such a channel number is allocated and advertised prior to use, by the multicast channel allocation protocol (MCAP), as described in [IP1394].

When a node wishes to receive multicast data addressed to other than all-nodes multicast addresses, all-routers multicast addresses, and solicited-node multicast addresses, it MUST confirm if the channel mapping between a multicast group address and a channel number exists using MCAP, as described in "9.3 Multicast Receive" in [IP1394].

The implementation of MCAP is optional for send-only nodes. A node MAY transmit multicast data addressed to any multicast addresses into the default broadcast channel regardless of the existing allocation of the channel. If a node wishes to transmit multicast data on other than the default channel, it MUST first confirm by MCAP whether or not a channel number for the group address has been already allocated. The implementors are encouraged to use this protocol when transmitting high-rate multicast streams.

The MCAP 'type' value for IPv6 group address descriptor is 2.

10. IANA CONSIDERATIONS

IANA has assigned a value of 0x000002 for "Unit_SW_Version for IPv6 over IEEE1394" out of the "CSR Protocol Identifiers" name space, as described in section 5. The details of the "CSR Protocol Identifiers" namespace is described in "10. IANA CONSIDERATIONS" of [IP1394].

Section 9.1 of [IP1394] defines MCAP group address descriptors, which include an 8-bit type name space. This document requests that IANA maintain a name space to manage MCAP group address descriptors. The initial assignments for that table are:

Value	Usage
0	reserved
1	IPv4 Multicast Address
2	IPv6 Multicast Address
255	reserved

Additional values from the range 3-254 can be assigned through Standards Action [RFC 2434].

11. Security Considerations

IPv6 over IEEE1394 does not introduce any additional security considerations over [IP1394]. The security concerns described in "11. SECURITY CONSIDERATIONS" in [IP1394] apply here as well.

12. Acknowledgment

The authors would like to acknowledge the authors of [IP1394] and [ETHER] since some part of this document has been derived from them.

13. References

- [1394] IEEE Std 1394-1995, Standard for a High Performance Serial Bus
- [1394a] IEEE Std 1394a-2000, Standard for a High Performance Serial Bus - Amendment 1
- [IP1394] Johansson, P., "IPv4 over IEEE 1394", RFC 2734, December 1999.
- [IPV6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

- [AARCH] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373 December 1998.
- [ACONF] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [DISC] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [ETHER] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.

14. Authors' Addresses

Kenji Fujisawa
Network & Software Technology Center, Sony Corporation
6-7-35 Kitashinagawa,
Shinagawa-ku, Tokyo 141-0001, JAPAN

Phone: +81-3-5795-8507
Fax: +81-3-5795-8977
EMail: fujisawa@sm.sony.co.jp

Atsushi Onoe
Internet Systems Laboratory,
Internet Laboratories, Sony Corporation
6-7-35 Kitashinagawa,
Shinagawa-ku, Tokyo 141-0001, JAPAN

Phone: +81-3-5448-4620
Fax: +81-3-5448-4622
EMail: onoe@sm.sony.co.jp

15. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

