

Transmission of IPv6 Packets over Ethernet 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 (1998). All Rights Reserved.

1. Introduction

This document specifies the frame format for transmission of IPv6 packets and the method of forming IPv6 link-local addresses and statelessly autoconfigured addresses on Ethernet networks. It also specifies the content of the Source/Target Link-layer Address option used in Router Solicitation, Router Advertisement, Neighbor Solicitation, Neighbor Advertisement and Redirect messages when those messages are transmitted on an Ethernet.

This document replaces RFC 1972, "A Method for the Transmission of IPv6 Packets over Ethernet Networks", which will become historic.

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].


2. Maximum Transmission Unit

The default MTU size for IPv6 [IPV6] packets on an Ethernet 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 Ethernet 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.

For purposes of this document, information received from DHCP is considered "manually configured" and the term Ethernet includes CSMA/CD and full-duplex subnetworks based on ISO/IEC 8802-3, with various data rates.

3. Frame Format

IPv6 packets are transmitted in standard Ethernet frames. The Ethernet header contains the Destination and Source Ethernet addresses and the Ethernet type code, which must contain the value 86DD hexadecimal. The data field contains the IPv6 header followed immediately by the payload, and possibly padding octets to meet the minimum frame size for the Ethernet link.



(Each tic mark represents one bit.)

4. Stateless Autoconfiguration

The Interface Identifier [AARCH] for an Ethernet interface is based on the EUI-64 identifier [EUI64] derived from the interface's built-in 48-bit IEEE 802 address. The EUI-64 is formed as follows. (Canonical bit order is assumed throughout.)

The OUI of the Ethernet address (the first three octets) becomes the company_id of the EUI-64 (the first three octets). The fourth and fifth octets of the EUI are set to the fixed value FFFE hexadecimal. The last three octets of the Ethernet address become the last three octets of the EUI-64.

The Interface Identifier is then formed from the EUI-64 by complementing the "Universal/Local" (U/L) bit, which is the next-to-lowest order bit of the first octet of the EUI-64. Complementing this bit will generally change a 0 value to a 1, since an interface's built-in address is expected to be from a universally administered address space and hence have a globally unique value. A universally administered IEEE 802 address or an EUI-64 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].

For example, the Interface Identifier for an Ethernet interface whose built-in address is, in hexadecimal,

34-56-78-9A-BC-DE

would be


36-56-78-FF-FE-9A-BC-DE.

A different MAC address set manually or by software should not be used to derive the Interface Identifier. If such a MAC address must be used, its global uniqueness property should be reflected in the value of the U/L bit.

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


5. Link-Local Addresses

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



6. Address Mapping -- Unicast

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



Option fields:

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

Length 1 (in units of 8 octets).

Ethernet Address

The 48 bit Ethernet 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 to derive the Interface Identifier.

7. Address Mapping -- Multicast

An IPv6 packet with a multicast destination address DST, consisting of the sixteen octets DST[1] through DST[16], is transmitted to the Ethernet multicast address whose first two octets are the value 3333 hexadecimal and whose last four octets are the last four octets of DST.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0 0 1 1 0 0 1 1|0 0 1 1 0 0 1 1|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  DST[13]      |  DST[14]      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  DST[15]      |  DST[16]      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

8. Differences From RFC 1972

The following are the functional differences between this specification and RFC 1972.

The Address Token, which was a node's 48-bit MAC address, is replaced with the Interface Identifier, which is 64 bits in length and based on the EUI-64 format [EUI64]. An IEEE-defined mapping exists from 48-bit MAC addresses to EUI-64 form.

A prefix used for stateless autoconfiguration must now be 64 bits long rather than 80. The link-local prefix is also shortened to 64 bits.

9. Security Considerations

The method of derivation of Interface Identifiers from MAC addresses is intended to preserve global uniqueness when possible. However, there is no protection from duplication through accident or forgery.

10. References

- [AARCH] Hinden, R. and S. Deering "IP Version 6 Addressing Architecture", RFC 2373, July 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.

- [EUI64] "Guidelines For 64-bit Global Identifier (EUI-64)",
<http://standards.ieee.org/db/oui/tutorials/EUI64.html>
- [IPV6] Deering, S. and R. Hinden, "Internet Protocol, Version 6
(IPv6) Specification", RFC 2460, December 1998.
- [RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119, March 1997.

11. Author's Address

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

Phone: +1 630 840-3461
EMail: crawdada@fnal.gov

12. Full Copyright Statement

Copyright (C) The Internet Society (1998). 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.

