

Transmission of IPv6 over IPv4 Domains without Explicit Tunnels

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 (1999). All Rights Reserved.

Abstract

This memo specifies the frame format for transmission of IPv6 [IPV6] packets and the method of forming IPv6 link-local addresses over IPv4 domains. It also specifies the content of the Source/Target Link-layer Address option used in the Router Solicitation, Router Advertisement, Neighbor Solicitation, and Neighbor Advertisement and Redirect messages, when those messages are transmitted on an IPv4 multicast network.

The motivation for this method is to allow isolated IPv6 hosts, located on a physical link which has no directly connected IPv6 router, to become fully functional IPv6 hosts by using an IPv4 domain that supports IPv4 multicast as their virtual local link. It uses IPv4 multicast as a "virtual Ethernet".

Table of Contents

1. Introduction.....	2
2. Maximum Transmission Unit.....	2
3. Frame Format.....	3
4. Stateless Autoconfiguration and Link-Local Addresses.....	3
5. Address Mapping -- Unicast.....	4
6. Address Mapping -- Multicast.....	4
7. Scaling and Transition Issues.....	5
8. IANA Considerations.....	6
9. Security Considerations.....	6

Acknowledgements.....	7
References.....	7
APPENDIX A: IPv4 Multicast Addresses for Neighbor Discovery.....	8
Authors' Addresses.....	9
Full Copyright Notice.....	10

1. Introduction

This memo specifies the frame format for transmission of IPv6 [IPv6] packets and the method of forming IPv6 link-local addresses over IPv4 multicast "domains". For the purposes of this document, an IPv4 domain is a fully interconnected set of IPv4 subnets, within the same local multicast scope, on which there are at least two IPv6 nodes conforming to this specification. This IPv4 domain could form part of the globally-unique IPv4 address space, or could form part of a private IPv4 network [RFC 1918].

This memo also specifies the content of the Source/Target Link-layer Address option used in the Router Solicitation, Router Advertisement, Neighbor Solicitation, Neighbor Advertisement and Redirect messages described in [DISC], when those messages are transmitted on an IPv4 multicast domain.

The motivation for this method is to allow isolated IPv6 hosts, located on a physical link which has no directly connected IPv6 router, to become fully functional IPv6 hosts by using an IPv4 multicast domain as their virtual local link. Thus, at least one IPv6 router using the same method must be connected to the same IPv4 domain if IPv6 routing to other links is required.

IPv6 hosts connected using this method do not require IPv4-compatible addresses or configured tunnels. In this way IPv6 gains considerable independence of the underlying links and can step over many hops of IPv4 subnets. The mechanism is known formally as "IPv6 over IPv4" or "6over4" and colloquially as "virtual Ethernet".

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 [RFC2119].

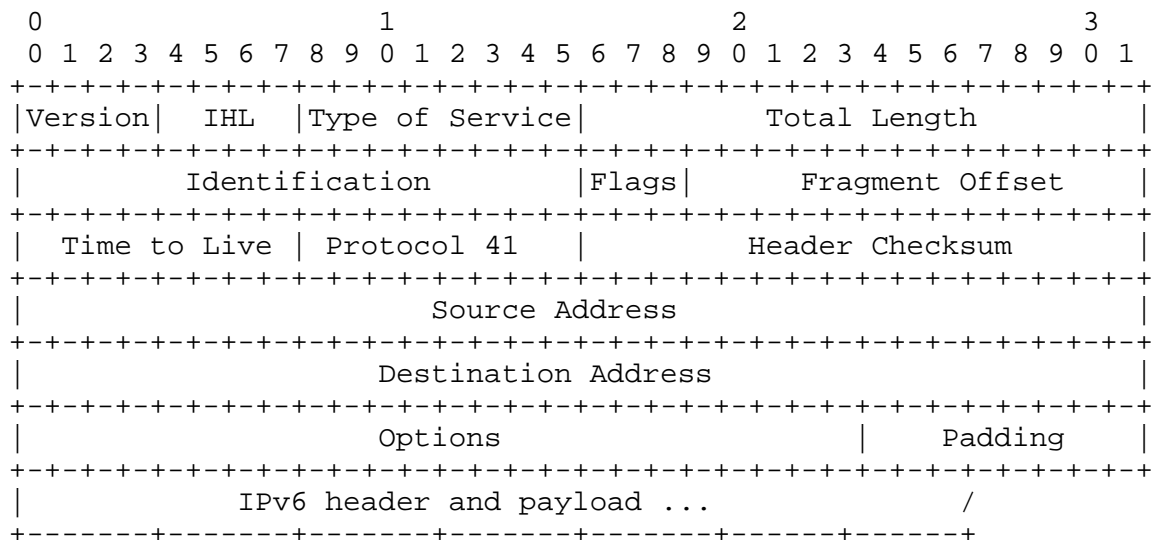
2. Maximum Transmission Unit

The default MTU size for IPv6 packets on an IPv4 domain is 1480 octets. This size may be varied by a Router Advertisement [DISC] containing an MTU option which specifies a different MTU, or by manual configuration of each node.

Note that if by chance the IPv6 MTU size proves to be too large for some intermediate IPv4 subnet, IPv4 fragmentation will ensue. While undesirable, this is not disastrous. However, the IPv4 "do not fragment" bit MUST NOT be set in the encapsulating IPv4 header.

3. Frame Format

IPv6 packets are transmitted in IPv4 packets [RFC 791] with an IPv4 protocol type of 41, the same as has been assigned in [RFC 1933] for IPv6 packets that are tunneled inside of IPv4 frames. The IPv4 header contains the Destination and Source IPv4 addresses. The IPv4 packet body contains the IPv6 header followed immediately by the payload.



If there are IPv4 options, then padding SHOULD be added to the IPv4 header such that the IPv6 header starts on a boundary that is a 32-bit offset from the end of the datalink header.

The Time to Live field SHOULD be set to a low value, to prevent such packets accidentally leaking from the IPv4 domain. This MUST be a configurable parameter, with a recommended default of 8.

4. Stateless Autoconfiguration and Link-Local Addresses

The Interface Identifier [AARCH] of an IPv4 interface is the 32-bit IPv4 address of that interface, with the octets in the same order in which they would appear in the header of an IPv4 packet, padded at the left with zeros to a total of 64 bits. Note that the "Universal/Local" bit is zero, indicating that the Interface Identifier is not globally unique. When the host has more than one IPv4 address in use

on the physical interface concerned, an administrative choice of one of these IPv4 addresses is made.

An IPv6 address prefix used for stateless autoconfiguration [CONF] of an IPv4 interface MUST have a length of 64 bits except for a special case mentioned in Section 7.

The IPv6 Link-local address [AARCH] for an IPv4 virtual interface is formed by appending the Interface Identifier, as defined above, to the prefix FE80::/64.

FE	80	00	00	00	00	00	00
00	00	00	00	IPv4 Address			

5. Address Mapping -- Unicast

The procedure for mapping IPv6 addresses into IPv4 virtual link-layer addresses is described in [DISC]. The Source/Target Link-layer Address option has the following form when the link layer is IPv4. Since the length field is in units of 8 bytes, the value below is 1.

Type	Length	must be zero	IPv4 Address
------	--------	--------------	--------------

Type:

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

Length:

- 1 (in units of 8 octets).

IPv4 Address:

The 32 bit IPv4 address, in network byte order. This is the address the interface currently responds to, and may be different from the Interface Identifier for stateless autoconfiguration.

6. Address Mapping -- Multicast

IPv4 multicast MUST be available. An IPv6 packet with a multicast destination address DST MUST be transmitted to the IPv4 multicast address of Organization-Local Scope using the mapping below. These IPv4 multicast addresses SHOULD be taken from the block

239.192.0.0/16, a sub-block of the Organization-Local Scope address block, or, if all of those are not available, from the expansion blocks defined in [ADMIN]. Note that when they are formed using the expansion blocks, they use only a /16 sized block.

```

+-----+-----+-----+-----+
| 239   | OLS   | DST14  | DST15  |
+-----+-----+-----+-----+

```

DST14, DST15 last two bytes of IPv6 multicast address.

OLS from the configured Organization-Local
Scope address block. SHOULD be 192,
see [ADMIN] for details.

No new IANA registration procedures are required for the above. See appendix A. for a list of all the multicast groups that must be joined to support Neighbor Discovery.

7. Scaling and Transition Issues

The multicast mechanism described in Section 6 above appears to have essentially the same scaling properties as native IPv6 over most media, except for the slight reduction in MTU size which will slightly reduce bulk throughput. On an ATM network, where IPv4 multicast relies on relatively complex mechanisms, it is to be expected that IPv6 over IPv4 over ATM will perform less well than native IPv6 over ATM.

The "IPv6 over IPv4" mechanism is intended to take its place in the range of options available for transition from IPv4 to IPv6. In particular it allows a site to run both IPv4 and IPv6 in coexistence, without having to configure IPv6 hosts either with IPv4-compatible addresses or with tunnels. Interfaces of the IPv6 router and hosts will of course need to be enabled in "6over4" mode.

A site may choose to start its IPv6 transition by configuring one IPv6 router to support "6over4" on an interface connected to the site's IPv4 domain, and another IPv6 format on an interface connected to the IPv6 Internet. Any enabled "6over4" hosts in the IPv4 domain will then be able to communicate both with the router and with the IPv6 Internet, without manual configuration of a tunnel and without the need for an IPv4-compatible IPv6 address, either stateless or stateful address configuration providing the IPv6 address to the IPv6 host.

During transition, routers may need to advertise at least two IPv6 prefixes, one for the native LAN (e.g. Ethernet) and one for "6over4". As with any IPv6 prefix assigned to an IPv6 subnet, the latter MUST be unique within its scope, whether site-local or global addressing is used.

Also note that when a router is handling both native LAN and "6over4" on the same physical interface, during stateless autoconfiguration, there is a period when IPv6 link-local addresses are used, in both cases with the prefix FE80::/64. Note that the prefix-length for these link-local address MUST then be 128 so that the two cases can be distinguished.

As the site installs additional IPv6 routers, "6over4" hosts which become physically adjacent to IPv6 routers can be changed to run as native IPv6 hosts, with the only impact on IPv6 applications being a slight increase in MTU size. At some stage during transition, it might be convenient to dual home hosts in both native LAN and "6over4" mode, but this is not required.

8. IANA Considerations

No assignments by the IANA are required beyond those in [ADMIN].

9. Security Considerations

Implementors should be aware that, in addition to possible attacks against IPv6, security attacks against IPv4 must also be considered. Use of IP security at both IPv4 and IPv6 levels should nevertheless be avoided, for efficiency reasons. For example, if IPv6 is running encrypted, encryption of IPv4 would be redundant except if traffic analysis is felt to be a threat. If IPv6 is running authenticated, then authentication of IPv4 will add little. Conversely, IPv4 security will not protect IPv6 traffic once it leaves the IPv6-over-IPv4 domain. Therefore, implementing IPv6 security is required even if IPv4 security is available.

There is a possible spoofing attack in which spurious 6over4 packets are injected into a 6over4 domain from outside. Thus, boundary routers MUST discard multicast IPv4 packets with source or destination multicast addresses of organisation local scope as defined in section 6 above, if they arrive on physical interfaces outside that scope. To defend against spurious unicast 6over4 packets, boundary routers MUST discard incoming IPv4 packets with protocol type 41 from unknown sources, i.e. IPv6-in-IPv4 tunnels must only be accepted from trusted sources. Unless IPSEC

authentication is available, the RECOMMENDED technique for this is to configure the boundary router only to accept protocol type 41 packets from source addresses within a trusted range or ranges.

Acknowledgements

The basic idea presented above is not original, and we have had invaluable comments from Matt Crawford, Steve Deering, Dan Harrington, Rich Draves, Erik Nordmark, Quang Nguyen, Thomas Narten, and other members of the IPNG and NGTRANS working groups.

This document is seriously ripped off from RFC 1972 written by Matt Crawford. Brian Carpenter was at CERN when the work was started.

References

- [AARCH] Hinden, R., and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.
- [ADMIN] Meyer, D., "Administratively Scoped IP Multicast", BCP 23, RFC 2365, July 1998.
- [CONF] 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.
- [IPV6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC 791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [RFC 1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., de Groot, G. and E. Lear, "Address Allocation for Private Internets", RFC 1918, February 1996.
- [RFC 1933] Gilligan, R. and E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", RFC 1933, April 1996.
- [RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC 1972] Crawford, M., "A Method for the Transmission of IPv6 Packets over Ethernet Networks", RFC 1972, August 1996.

APPENDIX A: IPv4 Multicast Addresses for Neighbor Discovery

The following IPv4 multicast groups are used to support Neighbor Discovery with this specification. The IPv4 addresses listed in this section were obtained by looking at the IPv6 multicast addresses that Neighbor Discovery uses, and deriving the resulting IPv4 "virtual link-layer" addresses that are generated from them using the algorithm given in Section 6.

all-nodes multicast address

- the administratively-scoped IPv4 multicast address used to reach all nodes in the local IPv4 domain supporting this specification. 239.0LS.0.1

all-routers multicast address

- the administratively-scoped IPv4 multicast address to reach all routers in the local IPv4 domain supporting this specification. 239.0LS.0.2

solicited-node multicast address

- an administratively scoped multicast address that is computed as a function of the solicited target's address by taking the low-order 24 bits of the IPv4 address used to form the IPv6 address, and prepending the prefix FF02:0:0:0:0:1:FF00::/104 [AARCH]. This is then mapped to the IPv4 multicast address by the method described in this document. For example, if the IPv4 address used to form the IPv6 address is W.X.Y.Z, then the IPv6 solicited node multicast address is FF02::1:255.X.Y.Z and the corresponding IPv4 multicast address is 239.0LS.Y.Z

Authors' Addresses

Brian E. Carpenter
Internet Division
IBM United Kingdom Laboratories
MP 185, Hursley Park
Winchester, Hampshire S021 2JN, UK

EMail: brian@hursley.ibm.com

Cyndi Jung
3Com Corporation
5400 Bayfront Plaza, Mailstop 3219
Santa Clara, California 95052-8145

EMail: cmj@3Com.com

Full Copyright Statement

Copyright (C) The Internet Society (1999). 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.

