

Multi-Link Subnet Issues

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

There have been several proposals around the notion that a subnet may span multiple links connected by routers. This memo documents the issues and potential problems that have been raised with such an approach.

Table of Contents

1. Introduction	2
2. Issues	3
2.1. IP Model	3
2.2. TTL/Hop Limit Issues	4
2.3. Link-scoped Multicast and Broadcast	6
2.4. Duplicate Address Detection Issues	7
3. Security Considerations	8
4. Recommendations	9
4.1. IP Link Model	9
4.2. IPv6 Address Assignment	10
4.3. Duplicate Address Detection Optimizations	12
5. Normative References	12
6. Informative References	13

1. Introduction

The original IPv4 address definition [RFC791] consisted of a Network field, identifying a network number, and a Local Address field, identifying a host within that network. As organizations grew to want many links within their network, their choices were (from [RFC950]) to:

1. Acquire a distinct Internet network number for each cable; subnets are not used at all.
2. Use a single network number for the entire organization, but assign host numbers without regard to which LAN a host is on ("transparent subnets").
3. Use a single network number, and partition the host address space by assigning subnet numbers to the LANs ("explicit subnets").

[RFC925] was a proposal for option 2 that defined a specific type of Address Resolution Protocol (ARP) proxy behavior, where the forwarding plane had the properties of decrementing the Time To Live (TTL) to prevent loops when forwarding, not forwarding packets destined to 255.255.255.255, and supporting subnet broadcast by requiring that the ARP-based bridge maintain a list of recent broadcast packets. This approach was never standardized, although [RFC1027] later documented an implementation of a subset of [RFC925].

Instead, the IETF standardized option 3 with [RFC950], whereby hosts were required to learn a subnet mask, and this became the IPv4 model.

Over the recent past, there have been several newer protocols proposing to extend the notion of a subnet to be able to span multiple links, similar to [RFC925].

Early versions of the IPv6 scoped address architecture [SCOPID] proposed a subnet scope above the link scope, to allow for multi-link subnets. This notion was rejected by the WG due to the issues discussed in this memo, and as a result the final version [RFC4007] has no such notion.

There was also a proposal to define multi-link subnets [MLSR] for IPv6. However, this notion was abandoned by the IPv6 WG due to the issues discussed in this memo, and that proposal was replaced by a different mechanism that preserves the notion that a subnet spans only one link [RFC4389].

However, other WGs continued to allow for this concept even though it had been rejected in the IPv6 WG. Mobile IPv6 [RFC3775] allows tunnels to mobile nodes to use the same subnet as a home link, with the Home Agent doing layer 3 forwarding between them.

The notion also arises in Mobile Ad-hoc NETWORKS (MANETs) with proposals that an entire MANET is a subnet, with routers doing layer 3 forwarding within it.

The use of multi-link subnets has also been considered by other working groups, including NetLMM, 16ng, and Autoconf, and by other external organizations such as WiMax.

In this memo, we document the issues raised in the IPv6 WG which motivated the abandonment of the multi-link subnet concept, so that designers of other protocols can (and should) be aware of the issues.

The key words "MUST", "RECOMMENDED", and "SHOULD" in this document are to be interpreted as described in [RFC2119].

2. Issues

2.1. IP Model

The term "link" is generally used to refer to a topological area bounded by routers that decrement the IPv4 TTL or IPv6 Hop Limit when forwarding the packet. A link-local address prefix is defined in both IPv4 [RFC3927] and IPv6 [RFC4291].

The term "subnet" is generally used to refer to a topological area that uses the same address prefix, where that prefix is not further subdivided except into individual addresses.

In December 1995, the original IP Version 6 Addressing Architecture [RFC1884] was published, stating: "IPv6 continues the IPv4 model that a subnet is associated with one link. Multiple subnets may be assigned to the same link."

Thus, it explicitly acknowledges that the current IPv4 model has been that a subnet is associated with one link and that IPv6 does not change this model. Furthermore, a subnet is sometimes considered to be only a subset of a link, when multiple subnets are assigned to the same link.

The IPv6 addressing architecture has since been updated three times, first in July 1998 [RFC2373], then April 2003 [RFC3513], and finally in February 2006 [RFC4291]. All updates include the language: "Currently IPv6 continues the IPv4 model that a subnet prefix is

associated with one link. Multiple subnet prefixes may be assigned to the same link."

Clearly, the notion of a multi-link subnet would be a change to the existing IP model.

Similarly, the Mobility Related Terminology [RFC3753] defines a Foreign subnet prefix as "a bit string that consists of some number of initial bits of an IP address which identifies a node's foreign link within the Internet topology" with a similar definition for a Home subnet prefix. These both state that the subnet prefix identifies a (singular) link.

2.2. TTL/Hop Limit Issues

Since a link is bounded by routers that decrement the IPv4 TTL or IPv6 Hop Limit, there may be issues with applications and protocols that make any assumption about the relationship between TTL/Hop Limit and subnet prefix.

There are two main cases that may arise. Some applications and protocols may send packets with a TTL/Hop Limit of 1. Other applications and protocols may send packets with a TTL/Hop Limit of 255 and verify that the value is 255 on receipt. Both are ways of limiting communication to within a single link, although the effects of these two approaches are quite different. Setting TTL/Hop Limit to 1 ensures that packets that are sent do not leave the link, but it does not prevent an off-link attacker from sending a packet that can reach the link. Checking that TTL/Hop Limit is 255 on receipt prevents a receiver from accepting packets from an off-link sender, but it doesn't prevent a sent packet from being forwarded off-link.

As for assumptions about the relationship between TTL/Hop Limit and subnet, let's look at some example references familiar to many protocol and application developers.

Stevens' "Unix Network Programming", 2nd ed. [UNP], states on page 490, "A TTL of 0 means node-local, 1 means link-local" (this of course being true by the definition of link). Then page 498 states, regarding IP_MULTICAST_TTL and IPV6_MULTICAST_HOPS, "If this is not specified, both default to 1, which restricts the datagram to the local subnet." Here, Unix programmers learn that TTL=1 packets are restricted to a subnet (as opposed to a link). This is typical of many documents that use the terms interchangeably due to the IP model described earlier.

Similarly, "TCP/IP Illustrated", Volume 1 [TCPILL], states on page 182: "By default, multicast datagrams are sent with a TTL of 1. This restricts the datagram to the same subnet."

Steve Deering's original multicast README file [DEERING] contained the statement "multicast datagrams with initial TTL 1 are restricted to the same subnet", and similar statements now appear in many vendors' documentation, including documentation for Windows (e.g., [TCPIP2K]) and Linux (e.g., [LINUX] says a TTL of 1 is "restricted to the same subnet. Won't be forwarded by a router.")

The above are only some examples. There is no shortage of places where application developers are being taught that a subnet is confined to a single link, and so we must expect that arbitrary applications may embed such assumptions.

Some examples of protocols today that are known to embed some assumption about the relationship between TTL and subnet prefix are the following:

- o Neighbor Discovery (ND) [RFC2461] uses messages with Hop Limit 255 checked on receipt, to resolve the link-layer address of any IP address in the subnet.
- o Older clients of Apple's Bonjour [MDNS] use messages with TTL 255 checked on receipt, and only respond to queries from addresses in the same subnet. (Note that multi-link subnets do not necessarily break this, as this behavior is to constrain communication to within a subnet, where a subnet is only a subset of a link. However, it will not work across a multi-link subnet.)

Some other examples of protocols today that are known to use a TTL 1 or 255, but do not appear to explicitly have any assumption about the relationship to subnet prefixes (other than the well-known link-local prefix) include the following:

- o Link-Local Multicast Name Resolution [LLMNR] uses a TTL/Hop Limit of 1 for TCP.
- o Multicast Listener Discovery (MLD) [RFC3810] uses a Hop Limit of 1.
- o Reverse tunneling for Mobile IPv4 [RFC3024] uses TTL 255 checked on receipt for Registration Requests sent to foreign agents.

- o [RFC3927] discusses the use of TTL=1 and TTL=255 within the IPv4 link-local address prefix.

It is unknown whether any implementations of such protocols exist that add such assumptions about the relationship to subnet prefixes for other reasons.

2.3. Link-scoped Multicast and Broadcast

Because multicast routing is not ubiquitous, the notion of a subnet that spans multiple links tends to result in cases where multicast does not work across the subnet. Per [RFC2644], the default behavior is that routers do not forward directed broadcast packets either, nor do they forward limited broadcasts (see [RFC1812], Section 4.2.2.11).

There are many protocols and applications today that use link-scoped multicast. The list of such applications and protocols that have been assigned their own link-scoped multicast group address (and may also have assumptions about the TTL/Hop Limit as noted above) can be found at:

<http://www.iana.org/assignments/multicast-addresses>

<http://www.iana.org/assignments/ipv6-multicast-addresses>

In addition, an arbitrarily large number of other applications may be using the all-1's broadcast address, or the all-hosts link-scoped multicast address, rather than their own group address.

The well-known examples of protocols using link-scoped multicast or broadcast generally fall into one of the following groups:

- o Routing protocols: Distance Vector Multicast Routing Protocol (DVMRP) [RFC1075], OSPF [RFC2328], RIP [RFC2453][RFC2080], Enhanced Interior Gateway Routing Protocol (EIGRP) [EIGRP], etc. These protocols exchange routes to subnet prefixes.
- o Address management protocols: Neighbor Discovery, DHCPv4 [RFC2131], Dynamic Host Configuration Protocol for IPv6 (DHCPv6) [RFC3315], Teredo [RFC4380], etc. By their nature, this group tends to embed assumptions about the relationship between a link and a subnet prefix. For example, ND uses link-scoped multicast to resolve the link-layer address of an IP address in the same subnet prefix, and to do duplicate address detection (see Section 2.4 below) within the subnet. DHCP uses link-scoped multicast or broadcast to obtain an address in the subnet. Teredo states that the Teredo IPv4 Discovery Address is "an IPv4 multicast address used to

discover other Teredo clients on the same IPv4 subnet. The value of this address is 224.0.0.253", which is a link-scoped multicast address. It also says that "the client MUST silently discard all local discovery bubbles [...] whose IPv4 source address does not belong to the local IPv4 subnet".

- o Service discovery protocols: Simple Service Discovery Protocol (SSDP) [SSDP], Bonjour, WS-Discovery [WSDISC], etc. These often do not define any explicit assumption about the relationship to subnet prefix.
- o Name resolution protocols: NetBios [RFC1001], Bonjour, LLMNR, etc. Most often these do not define any explicit assumption about the relationship to subnet prefix, but Bonjour only responds to queries from addresses within the same subnet prefix.

Note that protocols such as Bonjour and Teredo that drop packets that don't come from an address within the subnet are not necessarily broken by multi-link subnets, as this behavior is meant to constrain the behavior to within a subnet, when a link is larger than a single subnet.

However, regardless of whether any assumption about the relationship to subnet prefixes exists, all protocols mentioned above or on the IANA assignments lists will not work across a multi-link subnet without protocol-specific proxying functionality in routers, and adding proxying for an arbitrary number of protocols and applications does not scale. Furthermore, it may hinder the development and use of future protocols using link-scoped multicast.

2.4. Duplicate Address Detection Issues

Duplicate Address Detection (DAD) uses link-scoped multicast in IPv6 and link-scoped broadcast in IPv4 and so has the issues mentioned in Section 2.3 above.

In addition, [RFC2462] contains the statement:

"Thus, for a set of addresses formed from the same interface identifier, it is sufficient to check that the link-local address generated from the identifier is unique on the link. In such cases, the link-local address MUST be tested for uniqueness, and if no duplicate address is detected, an implementation MAY choose to skip Duplicate Address Detection for additional addresses derived from the same interface identifier."

The last possibility, sometimes referred to as Duplicate Interface Identifier Detection (DIID), has been a matter of much debate, and the current work in progress [2462BIS] states:

Each individual unicast address SHOULD be tested for uniqueness. Note that there are implementations deployed that only perform Duplicate Address Detection for the link-local address and skip the test for the global address using the same interface identifier as that of the link-local address. Whereas this document does not invalidate such implementations, this kind of "optimization" is NOT RECOMMENDED, and new implementations MUST NOT do that optimization.

The existence of such implementations also causes problems with multi-link subnets. Specifically, a link-local address is only valid within a link, and hence is only tested for uniqueness within a single link. If the same interface identifier is then assumed to be unique across all links within a multi-link subnet, address conflicts can occur.

3. Security Considerations

The notion of multi-link subnets can cause problems with any security protocols that either rely on the assumption that a subnet only spans a single link or can leave gaps in the security solution where protocols are only defined for use on a single link.

Secure Neighbor Discovery (SEND) [RFC3971], in particular, is currently only defined within a single link. If a subnet were to span multiple links, SEND would not work as currently specified, since it secures Neighbor Discovery messages that include link-layer addresses, and if forwarded to other links, the link-layer address of the sender will be different. This same problem also exists in cases where a subnet does not span multiple links but where Neighbor Discovery is proxied within a link. Section 9 of [RFC4389] discusses some possible future directions in this regard.

Furthermore, as noted above some applications and protocols (ND, Bonjour, Mobile IPv4, etc.) mitigate against off-link spoofing attempts by requiring a TTL or Hop Limit of 255 on receipt. If this restriction were removed, or if alternative protocols were used, then off-link spoofing attempts would become easier, and some alternative way to mitigate such attacks would be needed.

4. Recommendations

4.1. IP Link Model

There are two models that do not have the issues pointed out in the rest of the document.

The IAB recommends that protocol designers use one of the following two models:

- o Multi-access link model: In this model, there can be multiple nodes on the same link, including zero or more routers. Data packets sent to the IPv4 link-local broadcast address (255.255.255.255) or to a link-local multicast address can be received by all other interested nodes on the link. Two nodes on the link are able to communicate without any IPv4 TTL or IPv6 Hop Limit decrement. There can be any number of layer 2 devices (bridges, switches, access points, etc.) in the middle of the link.
- o Point-to-point link model: In this model, there are exactly two nodes on the same link. Data packets sent to the IPv4 link-local broadcast address or to a link-local multicast address can be received by the other node on the link. The two nodes are able to communicate without any IPv4 TTL or IPv6 Hop Limit decrement. There can be any number of layer 2 devices (bridges, switches, access points, etc.) in the middle of the link.

A variant of the multi-access link model, which has fewer issues, but still some, is the following:

- o Non-broadcast multi-access (NBMA) model: Same as the multi-access link model, except that no broadcast or multicast packets can be sent, even between two nodes on the same link. As a result, no protocols or applications that make use of broadcast or multicast will work.

Links that appear as NBMA links at layer 3 are problematic. Instead, if a link is an NBMA link at layer 2, then protocol designers should define some mechanism such that it appears as either the multi-access link model or point-to-point link model at layer 3.

One use of an NBMA link is when the link itself is intended as a wide-area link (e.g., a tunnel such as 6to4 [RFC3056]) where none of the groups of functionality in Section 2.3 are required across the wide area. Admittedly, the definition of wide-area is somewhat subjective. Support for multicast on a wide-area link would be

analogous to supporting multicast routing across a series of local-area links. The issues discussed in Section 2.3 will arise, but may be acceptable over a wide area until multicast routing is also supported.

Note that the distinction of whether or not a link is a tunnel is orthogonal to the choice of model; there exist tunnel links for all link models mentioned above.

A multi-link subnet model should be avoided. IETF working groups using, or considering using, multi-link subnets today should investigate moving to one of the other models. For example, the Mobile IPv6 WG should investigate having the Home Agent not decrement the Hop Limit, and forward multicast traffic.

When considering changing an existing multi-link subnet solution to another model, the following issues should be considered:

Loop prevention: If physical loops cannot exist within the subnet, then removing the TTL/Hop Limit decrement is not an issue. Otherwise, protocol designers can (for example) retain the decrement but use a separate prefix per link, or use some form of bridging protocol instead (e.g., [BRIDGE] or [RBRIDGE]).

Limiting broadcast (including all-hosts multicast): If there is no efficiency requirement to prevent broadcast from going to other on-link hosts, then flooding it within the subnet is not an issue. Otherwise, protocol designers can (for example) use a separate prefix per link, or flood broadcast other than ARP within the subnet (ARP is covered below in Section 4.3).

Limiting the scope of other multicast (including IPv6 Neighbor Discovery): If there is no efficiency requirement to prevent multicast from going to other on-link hosts, then flooding multicast within the subnet is not an issue. Otherwise, protocol designers can (for example) use a separate prefix per link, or use Internet Group Management Protocol (IGMP)/MLD snooping [RFC4541] instead.

4.2. IPv6 Address Assignment

In IPv6, the Prefix Information Option in a Router Advertisement (RA) is defined for use by a router to advertise an on-link prefix. That is, it indicates that a prefix is assigned to the link over which the RA is sent/received. That is, the router and the node both have an on-link route in their routing table (or on-link Prefix List, in the conceptual model of a host in [RFC2461]), and any addresses used in

the prefix are assigned to an interface (on any node) attached to that.

In contrast, DHCPv6 Prefix Delegation (DHCP-PD) [RFC3633] is defined for use by a client to request a prefix for use on a different link. Section 12.1 of RFC 3633 states:

Upon the receipt of a valid Reply message, for each IA_PD the requesting router assigns a subnet from each of the delegated prefixes to each of the links to which the associated interfaces are attached, with the following exception: the requesting router MUST NOT assign any delegated prefixes or subnets from the delegated prefix(es) to the link through which it received the DHCP message from the delegating router.

Hence, the upstream router has a route in its routing table that is not on-link, but points to the client; the prefix is assigned to a link other than the one over which DHCP-PD was done; and any addresses used in the prefix are assigned to an interface (on any node) attached to that other link.

The IAB believes that the distinction between these two cases (assigning a prefix to the same link vs. another link) is important, and that the IETF protocols noted above are appropriate for the two scenarios noted. The IAB recommends that other protocol designers remain consistent with the IETF-defined scopes of these protocols (e.g., not using DHCP-PD to assign a prefix to the same link, or using RAs to assign a prefix to another link).

In addition, the Prefix Information Option contains an L (on-link) flag. Normally, this flag is set, indicating that this prefix can be used for on-link determination. When not set, the advertisement makes no statement about on-link or off-link properties of the prefix. For instance, the prefix might be used for address configuration with some of the addresses belonging to the prefix being on-link and others being off-link. Care must be taken when the L flag is not set. Specifically, some platforms allow applications to retrieve the prefix length associated with each address of the node. If an implementation were to return the prefix length used for address configuration, then applications may incorrectly assume that TTL=1 is sufficient for communication, and that link-scoped multicast will reach other addresses in the prefix. As a result, the IAB recommends that designers and maintainers of APIs that provide a prefix length to applications address this issue. For example, they might indicate that no prefix length exists when the prefix is not on-link. If the API is not capable of reporting that one does not exist, then they might choose to report a value of 128 when the prefix is not on-link. This would result in such applications

believing they are on separate subnets, rather than on a multi-link subnet.

4.3. Duplicate Address Detection Optimizations

One of the reasons sometimes cited for wanting a multi-link subnet model (rather than a multi-access link model), is to minimize the ARP/ND traffic between end-nodes. This is primarily a concern in IPv4 where ARP results in a broadcast that would be seen by all nodes, not just the node with the IPv4 address being resolved. Even if this is a significant concern, the use of a multi-link subnet model is not necessary. The point-to-point link model is one way to avoid this issue entirely.

In the multi-access link model, IPv6 ND traffic can be reduced by using well-known multicast learning techniques (e.g., [RFC4541] at a layer 2 intermediate device (bridge, switch, access point, etc.)).

Some have suggested that a layer 2 device could maintain an ARP or ND cache and service requests from that cache. However, such a cache prevents any type of fast mobility between layer 2 ports, and breaks Secure Neighbor Discovery [RFC3971]. As a result, the IAB recommends to protocol designers that this approach be avoided, instead using an alternative such as layer 2 learning. For IPv4 (where no Secure ARP exists), the IAB recommends that protocol designers avoid having a device respond from its cache in cases where a node can legitimately move between layer 2 segments of the link without any layer 2 indications at the layer 2 intermediate device. Also, since currently there is no guarantee that any device other than the end-host knows all addresses of the end-host, protocol designers should avoid any dependency on such an assumption. For example, when no cache entry for a given request is found, protocol designers may specify that a node broadcast the request to all nodes.

5. Normative References

- [RFC791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [RFC950] Mogul, J. and J. Postel, "Internet Standard Subnetting Procedure", STD 5, RFC 950, August 1985.
- [RFC1812] Baker, F., Ed., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

- [RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [RFC2644] Senie, D., "Changing the Default for Directed Broadcasts in Routers", BCP 34, RFC 2644, August 1999.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", RFC 3927, May 2005.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.
- [RFC4007] Deering, S., Haberman, B., Jinmei, T., Nordmark, E., and B. Zill, "IPv6 Scoped Address Architecture", RFC 4007, March 2005.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4541] Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", RFC 4541, May 2006.

6. Informative References

- [2462BIS] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", Work in Progress, May 2005.
- [BRIDGE] T. Jeffree, editor, "Media Access Control (MAC) Bridges", ANSI/IEEE Std 802.1D, 2004, <http://standards.ieee.org/getieee802/download/802.1D-2004.pdf>.
- [DEERING] Deering, S., "IP Multicast Extensions for 4.3BSD UNIX and related systems (MULTICAST 1.2 Release)", June 1989. <http://www.kohala.com/start/mcast.api.txt>

- [EIGRP] Cisco, "Enhanced Interior Gateway Routing Protocol", Cisco Document ID 16406, September 2005.
<http://www.cisco.com/warp/public/103/eigrp-toc.html>
- [LINUX] Juan-Mariano de Goyeneche, "Multicast over TCP/IP HOWTO", March 1998. <http://www.linux.com/howtos/Multicast-HOWTO-2.shtml>
- [LLMNR] Aboba, B., Thaler, D., and L. Esibov, "Link-local Multicast Name Resolution (LLMNR)", RFC 4795, January 2007.
- [MDNS] Cheshire, S. and M. Krochmal, "Multicast DNS", June 2005.
<http://files.multicastdns.org/draft-cheshire-dnsext-multicastdns.txt>
- [MLSR] Thaler, D. and C. Huitema, "Multi-link Subnet Support in IPv6", Proceedings of IETF 54, June 2002.
<http://www.ietf.org/proceedings/02jul/I-D/draft-ietf-ipv6-multilink-subnets-00.txt>
- [RBRIDGE] Perlman, R., Gai, S., and D. Dutt, "Rbridges: Base Protocol Specification", Work in Progress, March 2007.
- [RFC925] Postel, J., "Multi-LAN address resolution", RFC 925, October 1984.
- [RFC1001] NetBIOS Working Group in the Defense Advanced Research Projects Agency, Internet Activities Board, and End-to-End Services Task Force, "Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods", STD 19, RFC 1001, March 1987.
- [RFC1027] Carl-Mitchell, S. and J. Quarterman, "Using ARP to Implement Transparent Subnet Gateways", RFC 1027, October 1987.
- [RFC1075] Waitzman, D., Partridge, C., and S. Deering, "Distance Vector Multicast Routing Protocol", RFC 1075, November 1988.
- [RFC1884] Hinden, R., Ed., and S. Deering, Ed., "IP Version 6 Addressing Architecture", RFC 1884, December 1995.
- [RFC2080] Malkin, G. and R. Minnear, "RIPng for IPv6", RFC 2080, January 1997.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [RFC2373] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.
- [RFC2453] Malkin, G., "RIP Version 2", STD 56, RFC 2453, November 1998.
- [RFC3024] Montenegro, G., Ed., "Reverse Tunneling for Mobile IP, revised", RFC 3024, January 2001.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3513] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.
- [RFC3753] Manner, J., Ed., and M. Kojo, Ed., "Mobility Related Terminology", RFC 3753, June 2004.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [RFC3810] Vida, R., Ed., and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.
- [RFC4389] Thaler, D., Talwar, M., and C. Patel, "Neighbor Discovery Proxies (ND Proxy)", RFC 4389, April 2006.
- [SCOPID] Deering, S., Haberman, B., Jinmei, T., Nordmark, E., Onoe, A., and B. Zill, "IPv6 Scoped Address Architecture", Proceedings of IETF 54, July 2002.
<http://www.ietf.org/proceedings/02jul/I-D/draft-ietf-ipngwg-scoping-arch-04.txt>
- [SSDP] Goland, Yaron Y., Cai, T., Leach, P., Gu, Y., and S. Albright, "Simple Service Discovery Protocol (SSDP)", 1999.
<http://www.upnp.org/resources/specifications.asp>

- [TCPILL] Stevens, W. Richard, "TCP/IP Illustrated, Volume 1", Addison-Wesley, 1994.
- [TCPIP2K] MacDonald, D. and W. Barkley, "Microsoft Windows 2000 TCP/IP Implementation Details". <http://www.microsoft.com/technet/itsolutions/network/deploy/depovg/tcpip2k.msp>
- [UNP] Stevens, W. Richard, "Unix Network Programming, Volume 1, Second Edition", Prentice Hall, 1998.
- [WSDISC] Microsoft, "Web Services Dynamic Discovery (WS-Discovery)", 2005. <http://specs.xmlsoap.org/ws/2005/04/discovery/ws-discovery.pdf>

IAB Members at the time of this writing

Bernard Aboba
Loa Andersson
Brian Carpenter
Leslie Daigle
Elwyn Davies
Kevin Fall
Olaf Kolkman
Kurtis Lindqvist
David Meyer
David Oran
Eric Rescorla
Dave Thaler
Lixia Zhang

Author's Address

Dave Thaler
Microsoft
One Microsoft Way
Redmond, WA 98052

Phone: +1 425 703 8835
EMail: dthaler@microsoft.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

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