

IPNGWG Working Group
Internet Draft
[draft-haberman-ipngwg-mcast-arch-01.txt](#)
March 2000
Expires September 2000

B. Haberman
Nortel Networks
D. Thaler
Microsoft

IP Version 6 Multicast Addressing Architecture

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) [[RFC 2026](#)].

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts. Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Abstract


This specification defines the multicast addressing architecture of the IP Version 6 protocol. The updated multicast address architecture presented in this document allows for unicast prefix-based allocation of multicast addresses. It is an update of [section 2.7](#) of the [RFC 2373](#).

[1.](#) Introduction

This document specifies an update to the IPv6 multicast addressing architecture [[RFC 2373](#)]. The current architecture does not contain any built-in support for dynamic address allocation. This proposal introduces encoded information in the multicast address to allow for dynamic, network prefix-based allocation of IPv6 multicast addresses, as well as allocation of single-source multicast addresses.

[2.](#) Multicast Address Format

An IPv6 multicast address is an identifier for a group of nodes. A node may belong to any number of multicast groups. Multicast addresses have the following generic format:



Haberman, Thaler

1

Internet Draft IPv6 Multicast Address Architecture

March 2000

11111111 at the start of the address identifies the address as being a multicast address.

flgs is a set of 4 flags:

+	-	+	-	+	+
0	S	P	T		
+	+	+	+	+	+

- o The high-order flag is reserved, and must be initialized to 0.
- o P = 0 indicates a multicast address that is not assigned based on the network prefix.
- o P = 1 indicates a multicast address that is assigned based on the network prefix.
- o T = 0 indicates a permanently assigned ("well-known") multicast address, assigned by the global Internet numbering authority.
- o T = 1 indicates a non-permanently-assigned ("transient") multicast address.
- o S = 0 indicates a multicast address that is not a single-source group address.
- o S = 1 indicates a multicast address that is a single-source group address.

scop is a 4-bit multicast scope value used to limit the scope of the multicast group. The values are:

- 0 reserved
- 1 node-local scope
- 2 link-local scope
- 3 local scope
- 4 (unassigned)
- 5 site-local scope
- 6 allocation scope

- 7 (unassigned)
- 8 organization-local scope
- 9 (unassigned)
- A (unassigned)
- B (unassigned)
- C (unassigned)
- D (unassigned)
- E global scope
- F reserved

The format of the next 80 bits is dependent on the value of the flags, and is discussed in [Section 4](#) below.

group ID identifies the multicast group, either permanent or transient, within the given scope.

The "meaning" of a permanently assigned multicast address is independent of the scope value. For example, if the "NTP servers group" is assigned a permanent multicast address with a group ID of 101 (hex), then:

FF01::101 means all NTP servers on the same node as the sender.

FF02::101 means all NTP servers on the same link as the sender.

FF05::101 means all NTP servers in the same site as the sender.

Haberman, Thaler

2

Internet Draft IPv6 Multicast Address Architecture

March 2000

FF0E::101 means all NTP servers in the Internet.

Non-permanently-assigned multicast addresses are meaningful only within a given scope. For example, a group identified by the non-permanent, site-local multicast address FF15::101 at one site bears no relationship to a group using the same address at a different site, or to a non-permanent group using the same group ID with a different scope, nor to a permanent group with the same group ID.

Multicast addresses must not be used as source addresses in IPv6 packets or appear in any routing header.

[3.](#) Pre-Defined Multicast Addresses

Since the meaning of a permanently assigned multicast address is independent of the scope value, such addresses are valid over all scope

ranges. This is shown by an "x" in the scope field of the address that means any legal scope value.

Note that multicast addresses that differ only in scope represent different groups. Nodes must join each group individually.

The following well-known multicast addresses are pre-defined:

Reserved Multicast Addresses: FF0x:0:0:0:0:0:0:0

The above multicast addresses are reserved and shall never be assigned to any multicast group.

All Nodes Addresses: FF0x:0:0:0:0:0:0:1

The above multicast addresses identify the group of all IPv6 nodes within a given scope. The All-Nodes-Address for scopes larger than scope 2 (link-local) SHOULD NOT be used or joined by nodes.

All Routers Addresses: FF0x:0:0:0:0:0:0:2

The above multicast addresses identify the group of all IPv6 routers within a given scope. The All-Routers-Address for scopes other than scope 1 (node-local), 2 (link-local), and 5 (site-local) SHOULD NOT be used or joined by routers.

Solicited-Node Address: FF0x:0:0:0:0:1:FFXX:XXXX

The above multicast address is computed as a function of a node's unicast and anycast addresses, and SHOULD NOT be used or joined for any scope value other than scope 2 (link-local). The solicited-node multicast address is formed by taking the low-order 24 bits of the address (unicast or anycast) and appending those bits to the prefix FF02:0:0:0:0:1:FF00::/104 resulting in a multicast address in the range FF02:0:0:0:0:1:FF00:0000 to FF02:0:0:0:0:1:FFFF:FFFF.


For example, the solicited node multicast address corresponding to the IPv6 address 4037::01:800:200E:8C6C is FF02::1:FF0E:8C6C. IPv6 addresses that differ only in the high-order bits, e.g. due to multiple high-order prefixes associated with different aggregations, will map to the same solicited-node address thereby reducing the number of multicast addresses a node must join.

multicast addresses for every unicast and anycast address it is assigned.


4. Assignment of New IPv6 Multicast Addresses

The current approach [[RFC 2464](#)] to map IPv6 multicast addresses into IEEE 802 MAC addresses takes the low order 32 bits of the IPv6 multicast address and uses it to create a MAC address. Note that Token Ring networks are handled differently. Token Ring support of IPv6 multicast is defined in [[RFC 2470](#)]. Group ID's less than or equal to 32 bits long will generate unique MAC addresses.

Due to this, the existing IPv6 multicast address architecture suggests that the group identifier always be in the low order 32 bits as shown in the following:



This document updates the above by specifying a different address format when P = 1. Any new IPv6 multicast addresses that are network prefix-based will have the following format:



plen indicates the length of the network prefix portion of the address when P = 1. This field is required in order to determine the number of bits to include as part of the unicast prefix.

network prefix identifies the network prefix of the unicast subnet owning the multicast address. If P = 1, this field contains the unicast network prefix defined in [[RFC 2374](#)] and assigned to the domain owning the multicast address.

The reserved field must be zero.

While this limits the number of unicast prefix-based IPv6 multicast groups to 2³² per prefix, this is unlikely to be a limitation in the future. If it becomes necessary to exceed this limit in the future, multicast will still work but the processing will be slightly slower.

With the network prefix-based architecture and the current unicast address architecture [[RFC 2374](#)], the network prefix portion of the multicast address will be at most 64 bits. This allows for the group ID field to be at least 40 bits.

Additional IPv6 multicast addresses are defined and registered by the IANA [[RFC 2375](#)].

Haberman, Thaler

4

Internet Draft IPv6 Multicast Address Architecture

March 2000

[5.](#) Open Issues

5.1 IPv4-compatible Multicast Addresses

Is ::224.1.2.24 considered a legal IPv6 multicast address? Is ::FFFF:224.1.2.24 considered a legal IPv6 multicast address?

5.2 Scope-relative Offsets

Scope-relative offsets as described in [RFC 2365](#) [[RFC 2365](#)] are defined with language implying the offset is the same between IPv6 and IPv4. IANA has assigned them from separate number spaces. Should these offsets be the same for IPv4 and IPv6 so that every protocol that wants a scope-relative offset only has to get one offset and not two?

[6.](#) Security Considerations

Using unicast network-prefix based multicast addresses can sometimes aid in identifying the allocation domain of a given multicast address, although no guarantee is provided.

Using single-source multicast addresses can sometimes aid in the prevention of denial-of-service attacks by arbitrary sources, although no guarantee is provided.

[7.](#) References

[RFC 2026] S. Bradner, "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.

[RFC 2460] S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.

- [RFC 2373] R. Hinden and S. Deering, "IP Version 6 Addressing Architecture", [RFC 2373](#), July 1998.
- [RFC 2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), [BCP14](#), March 1999.
- [RFC 2374] R. Hinden, M. O'Dell, and S. Deering, "An IPv6 Aggregatable Global Unicast Address Format", [RFC 2374](#), July 1998.
- [RFC 2464] M. Crawford, "Transmission of IPv6 Packets over Ethernet Networks", [RFC 2464](#), December 1998.
- [RFC 2470] M. Crawford, T. Narten, and S. Thomas, "Transmission of IPv6 Packets over Token Ring Networks", [RFC 2470](#), December 1998.
- [RFC 2375] R. Hinden and S. Deering, "IPv6 Multicast Address Assignments", [RFC 2375](#), July 1998.
- [RFC 2365] D. Meyer, "Administratively Scoped IP Multicast", [BCP 23](#), [RFC 2365](#), July 1998.

Haberman, Thaler

6

Author's Address

Brian Haberman
Nortel Networks
4309 Emperor Blvd.

Suite 200
Durham, NC 27703
1-919-992-4439
Email : haberman@nortelnetworks.com

Dave Thaler
Microsoft Corporation
One Microsoft Way
Redmond, WA 48105-6399
1-425-703-8835
Email: dthaler@microsoft.com

