IPNGWG Working Group
Internet Draft
draft-haberman-ipngwg-mcast-arch-00.txt
December 1999
Expires June 2000

B. Haberma
Nortel Network

IP Version 6 Multicast Addressing Architecture

Status of this Memo

Abstract

This specification defines the multicast addressing architecture of the IP Version 6 protocol [RFC 2460]. The updated multicast address architecture presented in this document allows for prefix-based allocation of multicast addresses. It is an update of section 2.7 of the RFC 2373 [RFC 2373].

1. 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].

2. Introduction

This document specifies an update to the IPv6 multicast addressing architecture. 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.

3. 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 format:

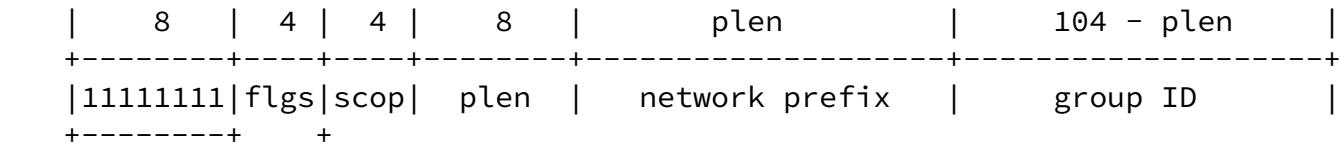

11111111 at the start of the address identifies the address as being a multicast address.

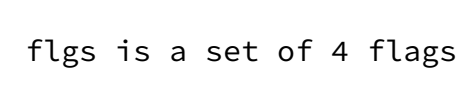

- The high-order 2 flags are reserved, and must be initialized to 0.
- P = 0 indicates a multicast address that is not assigned based on the network prefix. When P = 0, the plen field and the network prefix portion of the address are a part of the group ID.
- P = 1 indicates a multicast address that is assigned based on the network prefix.
- T = 0 indicates a permanently assigned (_well-known_) multicast

address, assigned by the global Internet numbering authority.

- T = 1 indicates a non-permanently-assigned (_transient_) multicast 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 (unassigned)
4 (unassigned)
5 site-local scope
6 (unassigned)
7 (unassigned)
8 organization-local scope
9 (unassigned)
A (unassigned)
B (unassigned)
C (unassigned)
D (unassigned)
E global scope
F reserved

plen indicates the length of the network prefix embedded in the address when P = 1. When P = 0, this field is considered a part of the group ID.

network prefix identifies the network prefix of the unicast subnet owning the multicast address. If P = 0, this field is considered a part of the group ID. If P = 1, this field contains the unicast network prefix defined in [RFC 2374] and assigned to the domain owning the multicast address.

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.

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.

4. Pre-Defined Multicast Addresses

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

Reserved Multicast Addresses:	FF00:0:0:0:0:0:0:0
	FF01:0:0:0:0:0:0:0
	FF02:0:0:0:0:0:0:0
	FF03:0:0:0:0:0:0:0
	FF04:0:0:0:0:0:0:0
	FF05:0:0:0:0:0:0:0
	FF06:0:0:0:0:0:0:0
	FF07:0:0:0:0:0:0:0
	FF08:0:0:0:0:0:0:0
	FF09:0:0:0:0:0:0:0
	FF0A:0:0:0:0:0:0:0
	FF0B:0:0:0:0:0:0:0
	FF0C:0:0:0:0:0:0:0
	FF0D:0:0:0:0:0:0:0
	FF0E:0:0:0:0:0:0:0
	FF0F: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:	FF01:0:0:0:0:0:0:1
	FF02:0:0:0:0:0:0:1

The above multicast addresses identify the group of all IPv6 nodes, within scope 1 (node-local) or 2 (link-local).

All Routers Addresses:	FF01:0:0:0:0:0:0:2
	FF02:0:0:0:0:0:0:2

FF05:0:0:0:0:0:0:2

The above multicast addresses identify the group of all IPv6 routers, within scope 1 (node-local), 2 (link-local), or 5 (site-local).

Solicited-Node Address: FF02:0:0:0:0:1:FFXX:XXXX

The above multicast address is computed as a function of a node's unicast and anycast addresses. 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 FF02:0:0:0:01:FFFF:FFFF.

For example, the solicited node multicast address corresponding to the IPv6 address 4037::01:800:200E:8C6C is FF02::1:FF0E:8C:6C. 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.

A node is required to compute and join the associated Solicited-Node multicast addresses for every unicast and anycast address it is assigned.

5. Assignment of New IPv6 Multicast Addresses

The current approach [RFC 2464] to map IPv6 multicast addresses into IEEE 802 MAC addresses takes that 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. This is defined in [RFC 2470]. Group ID's less than or equal to 32 bits will generate unique MAC addresses.

Due to this, new IPv6 multicast addresses that are not network prefix-based should be assigned so that the group identifier is always in the low order 32 bits as shown in the following:

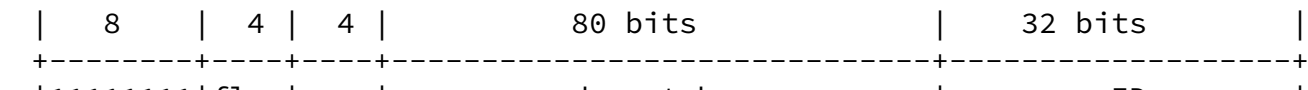

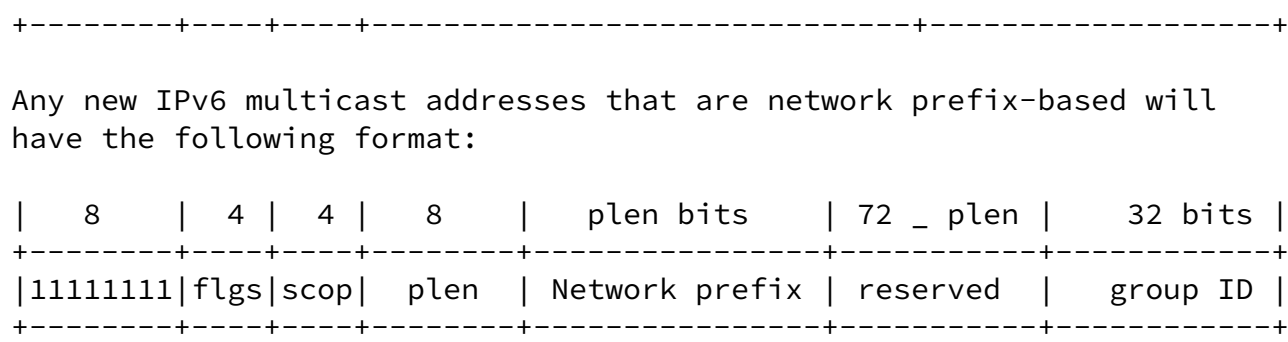

Any new IPv6 multicast addresses that are network prefix-based will have the following format:

While this limits the number of permanent IPv6 multicast groups to 2^32 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 40 bits.

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

6. Security Considerations

This document does not have any direct impact on Internet infrastructure security.

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.

Author's Address

Brian Haberman
Nortel Networks
4309 Emperor Blvd.
Suite 200
Durham, NC 27703
1-919-992-4439

Email : haberman@nortelnetworks.com