Network Working Group
Request for Comments: 2526
Category: Standards Track

D. Johnson
Carnegie Mellon University
S. Deering
Cisco Systems, Inc.

March 1999

Reserved IPv6 Subnet Anycast Addresses

Status of this Memo

Copyright Notice

Abstract

The IP Version 6 addressing architecture defines an "anycast" address as an IPv6 address that is assigned to one or more network interfaces (typically belonging to different nodes), with the property that a packet sent to an anycast address is routed to the "nearest" interface having that address, according to the routing protocols' measure of distance. This document defines a set of reserved anycast addresses within each subnet prefix, and lists the initial allocation of these reserved subnet anycast addresses.

1. Introduction

IP Version 6 (IPv6) defines a new type of address, known as an "anycast" address, that allows a packet to be routed to one of a number of different nodes all responding to the same address [2, 3]. The anycast address may be assigned to one or more network interfaces (typically on different nodes), with the network delivering each packet addressed to this address to the "nearest" interface based on the notion of "distance" determined by the routing protocols in use.

The uses of anycast addresses are still evolving, but such addresses offer the potential for a number of important services [5, 6]. For example, an anycast address may be used to allow nodes to access one of a collection of servers providing a well-known service, without manual configuration in each node of the list of servers; or an

anycast address may be used in a source route to force routing through a specific internet service provider, without limiting routing to a single specific router providing access to that ISP.

IPv6 defines a required Subnet-Router anycast address [3] for all routers within a subnet prefix, and allows additional anycast addresses to be taken from the unicast address space. This document defines an additional set of reserved anycast addresses within each subnet prefix, and lists the initial allocation of these reserved subnet anycast addresses.

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 [1].

2. Format of Reserved Subnet Anycast Addresses

Within each subnet, the highest 128 interface identifier values are reserved for assignment as subnet anycast addresses.

The construction of a reserved subnet anycast address depends on the type of IPv6 addresses used within the subnet, as indicated by the format prefix in the addresses. In particular, for IPv6 address types required to have 64-bit interface identifiers in EUI-64 format, the universal/local bit MUST be set to 0 (local) in all reserved subnet anycast addresses, to indicate that the interface identifier in the address is not globally unique. IPv6 addresses of this type are currently specified to be those having format prefixes 001 through 111, except for Multicast Addresses (1111 1111) [3].

Specifically, for IPv6 address types required to have to have 64-bit interface identifiers in EUI-64 format, these reserved subnet anycast addresses are constructed as follows:

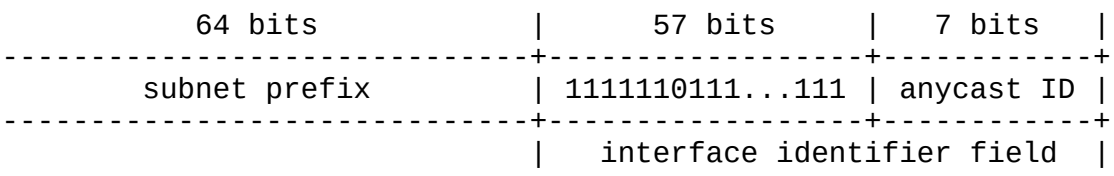

For other IPv6 address types (that is, with format prefixes other than those listed above), the interface identifier is not in EUI-64 format and may be other than 64 bits in length; these reserved subnet anycast addresses for such address types are constructed as follows:

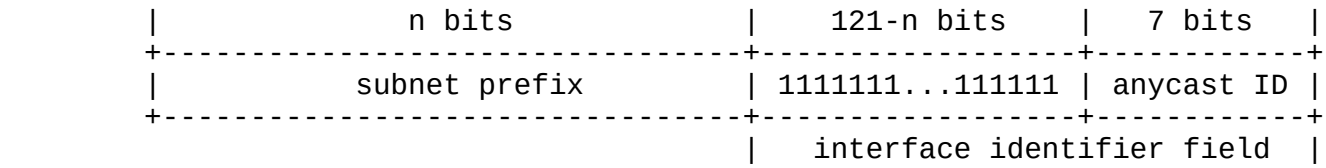

The subnet prefix here consists of all fields of the IPv6 address except the interface identifier field. The interface identifier field in these reserved subnet anycast addresses is formed from a 7-bit anycast identifier ("anycast ID"), with the remaining (highest order) bits filled with all one's; however, for interface identifiers in EUI-64 format, the universal/local bit in the interface identifier MUST be set to 0. The anycast identifier identifies a particular reserved anycast address within the subnet prefix, from the set of reserved subnet anycast addresses.

The motivation for reserving the highest addresses from each subnet rather than the lowest addresses, is to avoid conflicting with some existing official and unofficial uses of the low-numbered addresses in a subnet. For example, these low-numbered addresses are often used for the ends of a point-to-point link, for tunnel endpoints, for manually configured unicast addresses when a hardware token is not available for the network interface, and even for manually configured static addresses for the routers on a link. Reserving only 128 values for anycast identifiers (rather than perhaps 256) means that the minimum possible size of interface identifiers in an IPv6 address is 8 bits (including room in the subnet for unicast addresses as well as reserved subnet anycast addresses), allowing the division between subnet prefix and interface identifier in this case to be byte-aligned.

As with all IPv6 anycast addresses [3], these reserved subnet anycast addresses are allocated from the IPv6 unicast address space. All reserved subnet anycast addresses as defined in this document are reserved on all links, with all subnet prefixes. They MUST NOT be used for unicast addresses assigned to any interface.

3. List of Reserved Subnet Anycast Addresses

Currently, the following anycast identifiers for these reserved subnet anycast addresses are defined:

Decimal	Hexadecimal	Description
127	7F	Reserved
126	7E	Mobile IPv6 Home-Agents anycast [4]
0-125	00-7D	Reserved

Additional anycast identifiers are expected to be defined in the future.

4. Examples

To illustrate the construction of reserved subnet anycast addresses, this section details the construction of the reserved Mobile IPv6 Home-Agents subnet anycast address [4]. As noted in Section 3, the 7-bit anycast identifier for the Mobile IPv6 Home-Agents anycast address is 126 (decimal) or 7E (hexadecimal).

For IPv6 addresses containing a format prefix indicating that interface identifiers are required to be 64 bits in length and are required to be in EUI-64 format (currently format prefixes 001 through 111, except for 1111 1111 [3]), the reserved Mobile IPv6 Home-Agents subnet anycast address consists of the 64-bit subnet prefix followed by the 64-bit interface identifier shown below:

```
|0              1|1              3|3              4|4              6|
|0              5|6              1|2              7|8              3|
+----------------+----------------+----------------+----------------+
|1111110111111111|1111111111111111|1111111111111111|1111111111111110|
+----------------+----------------+----------------+----------------+
       ^                                                 ^^^^^^^
       +--- universal/local bit        anycast identifier ---+-----+
```

For other IPv6 address types, the interface identifier may be other than 64 bits in length and is not in EUI-64 format. In this example, assume that the length of the interface identifier is 64 bits, to allow clear comparison with the example given above (although interface identifiers of lengths other than 64 bits follow the same general construction of the interface identifier shown here). In this case, the reserved Mobile IPv6 Home-Agents subnet anycast address consists of the 64-bit subnet prefix followed by the 64-bit interface identifier shown below:

```
|0              1|1              3|3              4|4              6|
|0              5|6              1|2              7|8              3|
+----------------+----------------+----------------+----------------+
|1111111111111111|1111111111111111|1111111111111111|1111111111111110|
+----------------+----------------+----------------+----------------+
                                                        ^^^^^^^
                                       anycast identifier ---+-----+
```

5. IANA Considerations

This document defines a set of reserved subnet anycast addresses, based on a set of anycast identifiers within each subnet prefix in the IPv6 unicast address space. As future needs arise, new anycast identifiers may be defined. Such anycast identifiers MUST be reserved within all subnet prefixes, and so the assignment of these anycast identifiers requires centralized administration. New values SHOULD be assigned in descending numerical order and are expected to be assigned only with IESG approval.

6. Security Considerations

The use of any type of reserved anycast addresses poses a security concern only in allowing potential attackers a well-known address to attack. By designating certain services to be located at specific reserved anycast addresses, an attacker may more profitably focus an attack against such a specific service. Any such attack, however, is best dealt with in each service that uses a reserved anycast address.

RFC 1546, which originally proposed the idea of anycasting in IP, also points out a number of security considerations with the use of anycasting in general [6].

References

[1] Bradner, S., "Key words for use in RFCs to indicate requirement levels", BCP 14, RFC 2119, March 1997.

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

[3] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.

[4] David B. Johnson and Charles Perkins, "Mobility Support in IPv6", Work in Progress.

[5] Steve King et al, "The Case for IPv6", Work in Progress.

[6] Partridge, C., Mendez, T. and W. Milliken, "Host Anycasting Service", RFC 1546, November 1993.

Authors' Addresses

David B. Johnson
Carnegie Mellon University
Computer Science Department
5000 Forbes Avenue
Pittsburgh, PA 15213-3891
USA

Phone: +1 412 268-7399
Fax: +1 412 268-5576
EMail: dbj@cs.cmu.edu

Stephen E. Deering
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA

Phone: +1 408 527-8213
Fax: +1 408 527-8254
EMail: deering@cisco.com

Full Copyright Statement