

Network Working Group
Request for Comments: 3513
Obsoletes: 2373
Category: Standards Track

R. Hinden
Nokia
S. Deering
Cisco Systems
April 2003

Internet Protocol Version 6 (IPv6) Addressing Architecture

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 (2003). All Rights Reserved.

Abstract

This specification defines the addressing architecture of the IP Version 6 (IPv6) protocol. The document includes the IPv6 addressing model, text representations of IPv6 addresses, definition of IPv6 unicast addresses, anycast addresses, and multicast addresses, and an IPv6 node's required addresses.

Table of Contents

1. Introduction.....	3
2. IPv6 Addressing.....	3
2.1 Addressing Model.....	4
2.2 Text Representation of Addresses.....	4
2.3 Text Representation of Address Prefixes.....	5
2.4 Address Type Identification.....	6
2.5 Unicast Addresses.....	7
2.5.1 Interface Identifiers.....	8
2.5.2 The Unspecified Address.....	9
2.5.3 The Loopback Address.....	9
2.5.4 Global Unicast Addresses.....	10
2.5.5 IPv6 Addresses with Embedded IPv4 Addresses.....	10
2.5.6 Local-use IPv6 Unicast Addresses.....	11
2.6 Anycast Addresses.....	12
2.6.1 Required Anycast Address.....	13
2.7 Multicast Addresses.....	13
2.7.1 Pre-Defined Multicast Addresses.....	15
2.8 A Node's Required Addresses.....	17
3. Security Considerations.....	17
4. IANA Considerations.....	18
5. References.....	19
5.1 Normative References.....	19
5.2 Informative References.....	19
APPENDIX A: Creating Modified EUI-64 format Interface IDs.....	21
APPENDIX B: Changes from RFC-2373.....	24
Authors' Addresses.....	25
Full Copyright Statement.....	26

1. Introduction

This specification defines the addressing architecture of the IP Version 6 (IPv6) protocol. It includes the basic formats for the various types of IPv6 addresses (unicast, anycast, and multicast).

The authors would like to acknowledge the contributions of Paul Francis, Scott Bradner, Jim Bound, Brian Carpenter, Matt Crawford, Deborah Estrin, Roger Fajman, Bob Fink, Peter Ford, Bob Gilligan, Dimitry Haskin, Tom Harsch, Christian Huitema, Tony Li, Greg Minshall, Thomas Narten, Erik Nordmark, Yakov Rekhter, Bill Simpson, Sue Thomson, Markku Savela, and Larry Masinter.

2. IPv6 Addressing

IPv6 addresses are 128-bit identifiers for interfaces and sets of interfaces (where "interface" is as defined in section 2 of [IPV6]). There are three types of addresses:

Unicast: An identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address.

Anycast: An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the "nearest" one, according to the routing protocols' measure of distance).

Multicast: An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to a multicast address is delivered to all interfaces identified by that address.

There are no broadcast addresses in IPv6, their function being superseded by multicast addresses.

In this document, fields in addresses are given a specific name, for example "subnet". When this name is used with the term "ID" for identifier after the name (e.g., "subnet ID"), it refers to the contents of the named field. When it is used with the term "prefix" (e.g., "subnet prefix") it refers to all of the address from the left up to and including this field.

In IPv6, all zeros and all ones are legal values for any field, unless specifically excluded. Specifically, prefixes may contain, or end with, zero-valued fields.

2.1 Addressing Model

IPv6 addresses of all types are assigned to interfaces, not nodes. An IPv6 unicast address refers to a single interface. Since each interface belongs to a single node, any of that node's interfaces' unicast addresses may be used as an identifier for the node.

All interfaces are required to have at least one link-local unicast address (see section 2.8 for additional required addresses). A single interface may also have multiple IPv6 addresses of any type (unicast, anycast, and multicast) or scope. Unicast addresses with scope greater than link-scope are not needed for interfaces that are not used as the origin or destination of any IPv6 packets to or from non-neighbors. This is sometimes convenient for point-to-point interfaces. There is one exception to this addressing model:

A unicast address or a set of unicast addresses may be assigned to multiple physical interfaces if the implementation treats the multiple physical interfaces as one interface when presenting it to the internet layer. This is useful for load-sharing over multiple physical interfaces.

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.

2.2 Text Representation of Addresses

There are three conventional forms for representing IPv6 addresses as text strings:

1. The preferred form is `x:x:x:x:x:x:x:x`, where the 'x's are the hexadecimal values of the eight 16-bit pieces of the address.

Examples:

`FEDC:BA98:7654:3210:FEDC:BA98:7654:3210`

`1080:0:0:0:8:800:200C:417A`

Note that it is not necessary to write the leading zeros in an individual field, but there must be at least one numeral in every field (except for the case described in 2.).

2. Due to some methods of allocating certain styles of IPv6 addresses, it will be common for addresses to contain long strings of zero bits. In order to make writing addresses containing zero bits easier a special syntax is available to compress the zeros.

The use of "::" indicates one or more groups of 16 bits of zeros. The "::" can only appear once in an address. The "::" can also be used to compress leading or trailing zeros in an address.

For example, the following addresses:

1080:0:0:0:8:800:200C:417A	a unicast address
FF01:0:0:0:0:0:0:101	a multicast address
0:0:0:0:0:0:0:1	the loopback address
0:0:0:0:0:0:0:0	the unspecified addresses

may be represented as:

1080::8:800:200C:417A	a unicast address
FF01::101	a multicast address
::1	the loopback address
::	the unspecified addresses

3. An alternative form that is sometimes more convenient when dealing with a mixed environment of IPv4 and IPv6 nodes is x:x:x:x:x:x:d.d.d.d, where the 'x's are the hexadecimal values of the six high-order 16-bit pieces of the address, and the 'd's are the decimal values of the four low-order 8-bit pieces of the address (standard IPv4 representation). Examples:

0:0:0:0:0:0:13.1.68.3

0:0:0:0:0:FFFF:129.144.52.38

or in compressed form:

::13.1.68.3

::FFFF:129.144.52.38

2.3 Text Representation of Address Prefixes

The text representation of IPv6 address prefixes is similar to the way IPv4 addresses prefixes are written in CIDR notation [CIDR]. An IPv6 address prefix is represented by the notation:

ipv6-address/prefix-length

where

ipv6-address	is an IPv6 address in any of the notations listed in section 2.2.
--------------	---

prefix-length is a decimal value specifying how many of the leftmost contiguous bits of the address comprise the prefix.

For example, the following are legal representations of the 60-bit prefix 12AB00000000CD3 (hexadecimal):

```
12AB:0000:0000:CD30:0000:0000:0000:0000/60
12AB::CD30:0:0:0:0/60
12AB:0:0:CD30::/60
```

The following are NOT legal representations of the above prefix:

```
12AB:0:0:CD3/60    may drop leading zeros, but not trailing zeros,
                    within any 16-bit chunk of the address

12AB::CD30/60      address to left of "/" expands to
                    12AB:0000:0000:0000:0000:000:0000:CD30

12AB::CD3/60       address to left of "/" expands to
                    12AB:0000:0000:0000:0000:000:0000:0CD3
```

When writing both a node address and a prefix of that node address (e.g., the node's subnet prefix), the two can be combined as follows:

```
the node address      12AB:0:0:CD30:123:4567:89AB:CDEF
and its subnet number 12AB:0:0:CD30::/60
```

can be abbreviated as 12AB:0:0:CD30:123:4567:89AB:CDEF/60

2.4 Address Type Identification

The type of an IPv6 address is identified by the high-order bits of the address, as follows:

Address type	Binary prefix	IPv6 notation	Section
-----	-----	-----	-----
Unspecified	00...0 (128 bits)	::/128	2.5.2
Loopback	00...1 (128 bits)	::1/128	2.5.3
Multicast	11111111	FF00::/8	2.7
Link-local unicast	1111111010	FE80::/10	2.5.6
Site-local unicast	1111111011	FEC0::/10	2.5.6
Global unicast	(everything else)		

Anycast addresses are taken from the unicast address spaces (of any scope) and are not syntactically distinguishable from unicast addresses.

The general format of global unicast addresses is described in section 2.5.4. Some special-purpose subtypes of global unicast addresses which contain embedded IPv4 addresses (for the purposes of IPv4-IPv6 interoperation) are described in section 2.5.5.


Future specifications may redefine one or more sub-ranges of the global unicast space for other purposes, but unless and until that happens, implementations must treat all addresses that do not start with any of the above-listed prefixes as global unicast addresses.

2.5 Unicast Addresses


IPv6 unicast addresses are aggregable with prefixes of arbitrary bit-length similar to IPv4 addresses under Classless Interdomain Routing.

There are several types of unicast addresses in IPv6, in particular global unicast, site-local unicast, and link-local unicast. There are also some special-purpose subtypes of global unicast, such as IPv6 addresses with embedded IPv4 addresses or encoded NSAP addresses. Additional address types or subtypes can be defined in the future.

IPv6 nodes may have considerable or little knowledge of the internal structure of the IPv6 address, depending on the role the node plays (for instance, host versus router). At a minimum, a node may consider that unicast addresses (including its own) have no internal structure:



A slightly sophisticated host (but still rather simple) may additionally be aware of subnet prefix(es) for the link(s) it is attached to, where different addresses may have different values for n:



Though a very simple router may have no knowledge of the internal structure of IPv6 unicast addresses, routers will more generally have knowledge of one or more of the hierarchical boundaries for the operation of routing protocols. The known boundaries will differ

from router to router, depending on what positions the router holds in the routing hierarchy.

2.5.1 Interface Identifiers

Interface identifiers in IPv6 unicast addresses are used to identify interfaces on a link. They are required to be unique within a subnet prefix. It is recommended that the same interface identifier not be assigned to different nodes on a link. They may also be unique over a broader scope. In some cases an interface's identifier will be derived directly from that interface's link-layer address. The same interface identifier may be used on multiple interfaces on a single node, as long as they are attached to different subnets.

Note that the uniqueness of interface identifiers is independent of the uniqueness of IPv6 addresses. For example, a global unicast address may be created with a non-global scope interface identifier and a site-local address may be created with a global scope interface identifier.

For all unicast addresses, except those that start with binary value 000, Interface IDs are required to be 64 bits long and to be constructed in Modified EUI-64 format.

Modified EUI-64 format based Interface identifiers may have global scope when derived from a global token (e.g., IEEE 802 48-bit MAC or IEEE EUI-64 identifiers [EUI64]) or may have local scope where a global token is not available (e.g., serial links, tunnel end-points, etc.) or where global tokens are undesirable (e.g., temporary tokens for privacy [PRIV]).

Modified EUI-64 format interface identifiers are formed by inverting the "u" bit (universal/local bit in IEEE EUI-64 terminology) when forming the interface identifier from IEEE EUI-64 identifiers. In the resulting Modified EUI-64 format the "u" bit is set to one (1) to indicate global scope, and it is set to zero (0) to indicate local scope. The first three octets in binary of an IEEE EUI-64 identifier are as follows:

```

      0      0 0      1 1      2
    |0      7 8      5 6      3|
+---+---+---+---+---+---+
|cccc|ccug|cccc|cccc|cccc|cccc|
+---+---+---+---+---+---+

```

written in Internet standard bit-order, where "u" is the universal/local bit, "g" is the individual/group bit, and "c" are the bits of the company_id. Appendix A: "Creating Modified EUI-64 format

Interface Identifiers" provides examples on the creation of Modified EUI-64 format based interface identifiers.

The motivation for inverting the "u" bit when forming an interface identifier is to make it easy for system administrators to hand configure non-global identifiers when hardware tokens are not available. This is expected to be case for serial links, tunnel endpoints, etc. The alternative would have been for these to be of the form 0200:0:0:1, 0200:0:0:2, etc., instead of the much simpler 1, 2, etc.

The use of the universal/local bit in the Modified EUI-64 format identifier is to allow development of future technology that can take advantage of interface identifiers with global scope.

The details of forming interface identifiers are defined in the appropriate "IPv6 over <link>" specification such as "IPv6 over Ethernet" [ETHER], "IPv6 over FDDI" [FDDI], etc.

2.5.2 The Unspecified Address

The address 0:0:0:0:0:0:0:0 is called the unspecified address. It must never be assigned to any node. It indicates the absence of an address. One example of its use is in the Source Address field of any IPv6 packets sent by an initializing host before it has learned its own address.

The unspecified address must not be used as the destination address of IPv6 packets or in IPv6 Routing Headers. An IPv6 packet with a source address of unspecified must never be forwarded by an IPv6 router.

2.5.3 The Loopback Address

The unicast address 0:0:0:0:0:0:0:1 is called the loopback address. It may be used by a node to send an IPv6 packet to itself. It may never be assigned to any physical interface. It is treated as having link-local scope, and may be thought of as the link-local unicast address of a virtual interface (typically called "the loopback interface") to an imaginary link that goes nowhere.

The loopback address must not be used as the source address in IPv6 packets that are sent outside of a single node. An IPv6 packet with a destination address of loopback must never be sent outside of a single node and must never be forwarded by an IPv6 router. A packet received on an interface with destination address of loopback must be dropped.

2.5.4 Global Unicast Addresses

The general format for IPv6 global unicast addresses is as follows:

	n bits		m bits		128-n-m bits	
+-----+		+-----+		+-----+		+-----+
	global routing prefix		subnet ID		interface ID	
+-----+		+-----+		+-----+		+-----+

where the global routing prefix is a (typically hierarchically-structured) value assigned to a site (a cluster of subnets/links), the subnet ID is an identifier of a link within the site, and the interface ID is as defined in section 2.5.1.

All global unicast addresses other than those that start with binary 000 have a 64-bit interface ID field (i.e., $n + m = 64$), formatted as described in section 2.5.1. Global unicast addresses that start with binary 000 have no such constraint on the size or structure of the interface ID field.

Examples of global unicast addresses that start with binary 000 are the IPv6 address with embedded IPv4 addresses described in section 2.5.5 and the IPv6 address containing encoded NSAP addresses specified in [NSAP]. An example of global addresses starting with a binary value other than 000 (and therefore having a 64-bit interface ID field) can be found in [AGGR].

2.5.5 IPv6 Addresses with Embedded IPv4 Addresses

The IPv6 transition mechanisms [TRAN] include a technique for hosts and routers to dynamically tunnel IPv6 packets over IPv4 routing infrastructure. IPv6 nodes that use this technique are assigned special IPv6 unicast addresses that carry a global IPv4 address in the low-order 32 bits. This type of address is termed an "IPv4-compatible IPv6 address" and has the format:

	80 bits		16		32 bits	
+-----+		+-----+		+-----+		+-----+
	0000.....0000		0000		IPv4 address	
+-----+		+-----+		+-----+		+-----+

Note: The IPv4 address used in the "IPv4-compatible IPv6 address" must be a globally-unique IPv4 unicast address.

A second type of IPv6 address which holds an embedded IPv4 address is also defined. This address type is used to represent the addresses of IPv4 nodes as IPv6 addresses. This type of address is termed an "IPv4-mapped IPv6 address" and has the format:

	80 bits	16	32 bits	
+-----+		+-----+		+-----+
	0000.....0000	FFFF	IPv4 address	
+-----+		+-----+		+-----+

2.5.6 Local-Use IPv6 Unicast Addresses

There are two types of local-use unicast addresses defined. These are Link-Local and Site-Local. The Link-Local is for use on a single link and the Site-Local is for use in a single site. Link-Local addresses have the following format:

	10 bits		54 bits		64 bits	
+-----+		+-----+		+-----+		+-----+
	1111111010		0		interface ID	
+-----+		+-----+		+-----+		+-----+

Link-Local addresses are designed to be used for addressing on a single link for purposes such as automatic address configuration, neighbor discovery, or when no routers are present.

Routers must not forward any packets with link-local source or destination addresses to other links.

Site-Local addresses have the following format:

	10 bits		54 bits		64 bits	
+-----+		+-----+		+-----+		+-----+
	1111111011		subnet ID		interface ID	
+-----+		+-----+		+-----+		+-----+

Site-local addresses are designed to be used for addressing inside of a site without the need for a global prefix. Although a subnet ID may be up to 54-bits long, it is expected that globally-connected sites will use the same subnet IDs for site-local and global prefixes.

Routers must not forward any packets with site-local source or destination addresses outside of the site.

2.6 Anycast Addresses

An IPv6 anycast address is an address that is assigned to more than one interface (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.

Anycast addresses are allocated from the unicast address space, using any of the defined unicast address formats. Thus, anycast addresses are syntactically indistinguishable from unicast addresses. When a unicast address is assigned to more than one interface, thus turning it into an anycast address, the nodes to which the address is assigned must be explicitly configured to know that it is an anycast address.

For any assigned anycast address, there is a longest prefix P of that address that identifies the topological region in which all interfaces belonging to that anycast address reside. Within the region identified by P, the anycast address must be maintained as a separate entry in the routing system (commonly referred to as a "host route"); outside the region identified by P, the anycast address may be aggregated into the routing entry for prefix P.

Note that in the worst case, the prefix P of an anycast set may be the null prefix, i.e., the members of the set may have no topological locality. In that case, the anycast address must be maintained as a separate routing entry throughout the entire internet, which presents a severe scaling limit on how many such "global" anycast sets may be supported. Therefore, it is expected that support for global anycast sets may be unavailable or very restricted.

One expected use of anycast addresses is to identify the set of routers belonging to an organization providing internet service. Such addresses could be used as intermediate addresses in an IPv6 Routing header, to cause a packet to be delivered via a particular service provider or sequence of service providers.

Some other possible uses are to identify the set of routers attached to a particular subnet, or the set of routers providing entry into a particular routing domain.

There is little experience with widespread, arbitrary use of internet anycast addresses, and some known complications and hazards when using them in their full generality [ANYCST]. Until more experience has been gained and solutions are specified, the following restrictions are imposed on IPv6 anycast addresses:

- o An anycast address must not be used as the source address of an IPv6 packet.
- o An anycast address must not be assigned to an IPv6 host, that is, it may be assigned to an IPv6 router only.

2.6.1 Required Anycast Address

The Subnet-Router anycast address is predefined. Its format is as follows:

	n bits		128-n bits	
+-----+		+-----+		
	subnet prefix		0000000000000000	
+-----+		+-----+		

The "subnet prefix" in an anycast address is the prefix which identifies a specific link. This anycast address is syntactically the same as a unicast address for an interface on the link with the interface identifier set to zero.

Packets sent to the Subnet-Router anycast address will be delivered to one router on the subnet. All routers are required to support the Subnet-Router anycast addresses for the subnets to which they have interfaces.

The subnet-router anycast address is intended to be used for applications where a node needs to communicate with any one of the set of routers.

2.7 Multicast Addresses

An IPv6 multicast address is an identifier for a group of interfaces (typically on different link). An interface may belong to any number of multicast groups. Multicast addresses have the following format:

	8		4		4		112 bits	
+-----+		+-----+		+-----+		+-----+		
	11111111		flgs		scop		group ID	
+-----+		+-----+		+-----+		+-----+		

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

flgs is a set of 4 flags:

+-+-+--+
0 0 0 T
+-+-+--+

The high-order 3 flags are reserved, and must be initialized to 0.

T = 0 indicates a permanently-assigned ("well-known") multicast address, assigned by the Internet Assigned Number Authority (IANA).

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 interface-local scope
- 2 link-local scope
- 3 reserved
- 4 admin-local scope
- 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

interface-local scope spans only a single interface on a node, and is useful only for loopback transmission of multicast.

link-local and site-local multicast scopes span the same topological regions as the corresponding unicast scopes.

admin-local scope is the smallest scope that must be administratively configured, i.e., not automatically derived from physical connectivity or other, non-multicast-related configuration.

organization-local scope is intended to span multiple sites belonging to a single organization.

scopes labeled "(unassigned)" are available for administrators to define additional multicast regions.

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:

FF01:0:0:0:0:0:0:101 means all NTP servers on the same interface (i.e., the same node) as the sender.

FF02:0:0:0:0:0:0:101 means all NTP servers on the same link as the sender.

FF05:0:0:0:0:0:0:101 means all NTP servers in the same site as the sender.

FF0E:0:0:0:0:0:0:101 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 FF15:0:0:0:0:0:0:101 at one site bears no relationship to a group using the same address at a different site, nor to a non-permanent group using the same group ID with 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.

Routers must not forward any multicast packets beyond of the scope indicated by the scop field in the destination multicast address.

Nodes must not originate a packet to a multicast address whose scop field contains the reserved value 0; if such a packet is received, it must be silently dropped. Nodes should not originate a packet to a multicast address whose scop field contains the reserved value F; if such a packet is sent or received, it must be treated the same as packets destined to a global (scop E) multicast address.

2.7.1 Pre-Defined Multicast Addresses

The following well-known multicast addresses are pre-defined. The group ID's defined in this section are defined for explicit scope values.

Use of these group IDs for any other scope values, with the T flag equal to 0, is not allowed.

```

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 (interface-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 (interface-local), 2 (link-local), or 5 (site-local).

```

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

```

Solicited-node multicast address are computed as a function of a node's unicast and anycast addresses. A solicited-node multicast address is formed by taking the low-order 24 bits of an 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.

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

2.8 A Node's Required Addresses

A host is required to recognize the following addresses as identifying itself:

- o Its required Link-Local Address for each interface.
- o Any additional Unicast and Anycast Addresses that have been configured for the node's interfaces (manually or automatically).
- o The loopback address.
- o The All-Nodes Multicast Addresses defined in section 2.7.1.
- o The Solicited-Node Multicast Address for each of its unicast and anycast addresses.
- o Multicast Addresses of all other groups to which the node belongs.

A router is required to recognize all addresses that a host is required to recognize, plus the following addresses as identifying itself:

- o The Subnet-Router Anycast Addresses for all interfaces for which it is configured to act as a router.
- o All other Anycast Addresses with which the router has been configured.
- o The All-Routers Multicast Addresses defined in section 2.7.1.

3. Security Considerations

IPv6 addressing documents do not have any direct impact on Internet infrastructure security. Authentication of IPv6 packets is defined in [AUTH].

4. IANA Considerations

The table and notes at <http://www.isi.edu/in-notes/iana/assignments/ipv6-address-space.txt> should be replaced with the following:

INTERNET PROTOCOL VERSION 6 ADDRESS SPACE

The initial assignment of IPv6 address space is as follows:

Allocation	Prefix (binary)	Fraction of Address Space
-----	-----	-----
Unassigned (see Note 1 below)	0000 0000	1/256
Unassigned	0000 0001	1/256
Reserved for NSAP Allocation	0000 001	1/128 [RFC1888]
Unassigned	0000 01	1/64
Unassigned	0000 1	1/32
Unassigned	0001	1/16
Global Unicast	001	1/8 [RFC2374]
Unassigned	010	1/8
Unassigned	011	1/8
Unassigned	100	1/8
Unassigned	101	1/8
Unassigned	110	1/8
Unassigned	1110	1/16
Unassigned	1111 0	1/32
Unassigned	1111 10	1/64
Unassigned	1111 110	1/128
Unassigned	1111 1110 0	1/512
Link-Local Unicast Addresses	1111 1110 10	1/1024
Site-Local Unicast Addresses	1111 1110 11	1/1024
Multicast Addresses	1111 1111	1/256

Notes:

1. The "unspecified address", the "loopback address", and the IPv6 Addresses with Embedded IPv4 Addresses are assigned out of the 0000 0000 binary prefix space.
2. For now, IANA should limit its allocation of IPv6 unicast address space to the range of addresses that start with binary value 001. The rest of the global unicast address space (approximately 85% of the IPv6 address space) is reserved for future definition and use, and is not to be assigned by IANA at this time.

5. References

5.1 Normative References

- [IPV6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9 , RFC 2026, October 1996.

5.2 Informative References

- [ANYCST] Partridge, C., Mendez, T. and W. Milliken, "Host Anycasting Service", RFC 1546, November 1993.
- [AUTH] Kent, S. and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.
- [AGGR] Hinden, R., O'Dell, M. and S. Deering, "An Aggregatable Global Unicast Address Format", RFC 2374, July 1998.
- [CIDR] Fuller, V., Li, T., Yu, J. and K. Varadhan, "Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy", RFC 1519, September 1993.
- [ETHER] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.
- [EUI64] IEEE, "Guidelines for 64-bit Global Identifier (EUI-64) Registration Authority", <http://standards.ieee.org/regauth/oui/tutorials/EUI64.html>, March 1997.
- [FDDI] Crawford, M., "Transmission of IPv6 Packets over FDDI Networks", RFC 2467, December 1998.
- [MASGN] Hinden, R. and S. Deering, "IPv6 Multicast Address Assignments", RFC 2375, July 1998.
- [NSAP] Bound, J., Carpenter, B., Harrington, D., Houldsworth, J. and A. Lloyd, "OSI NSAPs and IPv6", RFC 1888, August 1996.
- [PRIV] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.
- [TOKEN] Crawford, M., Narten, T. and S. Thomas, "Transmission of IPv6 Packets over Token Ring Networks", RFC 2470, December 1998.

[TRAN] Gilligan, R. and E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", RFC 2893, August 2000.

APPENDIX A: Creating Modified EUI-64 format Interface Identifiers

Depending on the characteristics of a specific link or node there are a number of approaches for creating Modified EUI-64 format interface identifiers. This appendix describes some of these approaches.

Links or Nodes with IEEE EUI-64 Identifiers

The only change needed to transform an IEEE EUI-64 identifier to an interface identifier is to invert the "u" (universal/local) bit. For example, a globally unique IEEE EUI-64 identifier of the form:

0	1	3	4	6
0	5	1	7	3
-----+-----+-----+-----+-----				
cccccc0gcccccccc ccccccccmmmmmmmmmm mmmmmmmmmmmmmmmmmmmm mmmmmmmmmmmmmmmmmmmm				
+-----+-----+-----+-----+-----				

where "c" are the bits of the assigned company_id, "0" is the value of the universal/local bit to indicate global scope, "g" is individual/group bit, and "m" are the bits of the manufacturer-selected extension identifier. The IPv6 interface identifier would be of the form:

0	1	3	4	6
0	5	1	7	3
-----+-----+-----+-----+-----				
cccccc1gcccccccc ccccccccmmmmmmmmmm mmmmmmmmmmmmmmmmmmmm mmmmmmmmmmmmmmmmmmmm				
+-----+-----+-----+-----+-----				

The only change is inverting the value of the universal/local bit.

Links or Nodes with IEEE 802 48 bit MAC's

[EUI64] defines a method to create a IEEE EUI-64 identifier from an IEEE 48bit MAC identifier. This is to insert two octets, with hexadecimal values of 0xFF and 0xFE, in the middle of the 48 bit MAC (between the company_id and vendor supplied id). For example, the 48 bit IEEE MAC with global scope:

0	1	3	4
0	5	1	7

cccccc0gcccccccc	ccccccccmmmmmmmm	mmmmmmmmmmmmmmmm	

where "c" are the bits of the assigned company_id, "0" is the value of the universal/local bit to indicate global scope, "g" is individual/group bit, and "m" are the bits of the manufacturer-selected extension identifier. The interface identifier would be of the form:

0	1	3	4	6
0	5	1	7	3

cccccc1gcccccccc	cccccccc11111111	11111110mmmmmmmm	mmmmmmmmmmmmmmmm	

When IEEE 802 48bit MAC addresses are available (on an interface or a node), an implementation may use them to create interface identifiers due to their availability and uniqueness properties.

Links with Other Kinds of Identifiers

There are a number of types of links that have link-layer interface identifiers other than IEEE EUI-64 or IEEE 802 48-bit MACs. Examples include LocalTalk and Arcnet. The method to create an Modified EUI-64 format identifier is to take the link identifier (e.g., the LocalTalk 8 bit node identifier) and zero fill it to the left. For example, a LocalTalk 8 bit node identifier of hexadecimal value 0x4F results in the following interface identifier:

0	1	3	4	6
0	5	1	7	3

0000000000000000	0000000000000000	0000000000000000	0000000001001111	

Note that this results in the universal/local bit set to "0" to indicate local scope.

Links without Identifiers

There are a number of links that do not have any type of built-in identifier. The most common of these are serial links and configured tunnels. Interface identifiers must be chosen that are unique within a subnet-prefix.

When no built-in identifier is available on a link the preferred approach is to use a global interface identifier from another interface or one which is assigned to the node itself. When using this approach no other interface connecting the same node to the same subnet-prefix may use the same identifier.

If there is no global interface identifier available for use on the link the implementation needs to create a local-scope interface identifier. The only requirement is that it be unique within a subnet prefix. There are many possible approaches to select a subnet-prefix-unique interface identifier. These include:

- Manual Configuration
- Node Serial Number
- Other node-specific token

The subnet-prefix-unique interface identifier should be generated in a manner that it does not change after a reboot of a node or if interfaces are added or deleted from the node.

The selection of the appropriate algorithm is link and implementation dependent. The details on forming interface identifiers are defined in the appropriate "IPv6 over <link>" specification. It is strongly recommended that a collision detection algorithm be implemented as part of any automatic algorithm.

APPENDIX B: Changes from RFC-2373

The following changes were made from RFC-2373 "IP Version 6 Addressing Architecture":

- Clarified text in section 2.2 to allow ":::" to represent one or more groups of 16 bits of zeros.
- Changed uniqueness requirement of Interface Identifiers from unique on a link to unique within a subnet prefix. Also added a recommendation that the same interface identifier not be assigned to different machines on a link.
- Change site-local format to make the subnet ID field 54-bit long and remove the 38-bit zero's field.
- Added description of multicast scop values and rules to handle the reserved scop value 0.
- Revised sections 2.4 and 2.5.6 to simplify and clarify how different address types are identified. This was done to insure that implementations do not build in any knowledge about global unicast format prefixes. Changes include:
 - o Removed Format Prefix (FP) terminology
 - o Revised list of address types to only include exceptions to global unicast and a single entry that identifies everything else as Global Unicast.
 - o Removed list of defined prefix exceptions from section 2.5.6 as it is now the main part of section 2.4.
- Clarified text relating to EUI-64 identifiers to distinguish between IPv6's "Modified EUI-64 format" identifiers and IEEE EUI-64 identifiers.
- Combined the sections on the Global Unicast Addresses and NSAP Addresses into a single section on Global Unicast Addresses, generalized the Global Unicast format, and cited [AGGR] and [NSAP] as examples.
- Reordered sections 2.5.4 and 2.5.5.
- Removed section 2.7.2 Assignment of New IPv6 Multicast Addresses because this is being redefined elsewhere.
- Added an IANA considerations section that updates the IANA IPv6 address allocations and documents the NSAP and AGGR allocations.
- Added clarification that the "IPv4-compatible IPv6 address" must use global IPv4 unicast addresses.
- Divided references in to normative and non-normative sections.
- Added reference to [PRIV] in section 2.5.1
- Added clarification that routers must not forward multicast packets outside of the scope indicated in the multicast address.
- Added clarification that routers must not forward packets with source address of the unspecified address.
- Added clarification that routers must drop packets received on an interface with destination address of loopback.
- Clarified the definition of IPv4-mapped addresses.

- Removed the ABNF Description of Text Representations Appendix.
- Removed the address block reserved for IPX addresses.
- Multicast scope changes:
 - o Changed name of scope value 1 from "node-local" to "interface-local"
 - o Defined scope value 4 as "admin-local"
- Corrected reference to RFC1933 and updated references.
- Many small changes to clarify and make the text more consistent.

Authors' Addresses

Robert M. Hinden
Nokia
313 Fairchild Drive
Mountain View, CA 94043
USA

Phone: +1 650 625-2004
EMail: hinden@iprg.nokia.com

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

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

Full Copyright Statement

Copyright (C) The Internet Society (2003). 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.

Acknowledgement

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

