

IP Version 6 Addressing Architecture

[<draft-ietf-ipngwg-addr-arch-v2-00.txt>](#)

Status of this Memo

This document is an Internet-Draft. 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.''

To learn the current status of any Internet-Draft, please check the ``1id-abstracts.txt' listing contained in the Internet- Drafts Shadow Directories on ds.internic.net (US East Coast), nic.nordu.net (Europe), ftp.isi.edu (US West Coast), or munnari.oz.au (Pacific Rim).

This Internet Draft expires November 16, 1997.

Abstract

This specification defines the addressing architecture of the IP Version 6 protocol [[IPv6](#)]. 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 nodes 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 Representation.....	6
2.5 Unicast Addresses.....	8
2.5.1 Interface Identifiers.....	9
2.5.2 The Unspecified Address.....	10
2.5.3 The Loopback Address.....	10
2.5.4 IPv6 Addresses with Embedded IPv4 Addresses.....	10
2.5.5 NSAP Addresses.....	11
2.5.6 IPX Addresses.....	11
2.5.7 Aggregatable Global Unicast Addresses.....	11
2.5.8 Local-use IPv6 Unicast Addresses.....	12
2.6 Anycast Addresses.....	13
2.6.1 Required Anycast Address.....	14
2.7 Multicast Addresses.....	14
2.7.1 Pre-Defined Multicast Addresses.....	16
2.8 A Node's Required Addresses.....	17
REFERENCES.....	19
SECURITY CONSIDERATIONS.....	20
AUTHOR'S ADDRESSES.....	20
CHANGES FROM RFC-1884	21

1.0 INTRODUCTION

This specification defines the addressing architecture of the IP Version 6 protocol. It includes a detailed description of the currently defined address formats for IPv6 [[IPV6](#)].

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

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](#)].

2.0 IPV6 ADDRESSING

IPv6 addresses are 128-bit identifiers for interfaces and sets of interfaces. 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 "subscriber". When this name is used with the term "ID" for identifier after the name (e.g., "subscriber ID"), it refers to the contents of the named field. When it is used with the term "prefix" (e.g. "subscriber prefix") it refers to all of the address 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 zero-valued fields or end in zeros.

2.1 Addressing Model

IPv6 Addresses of all types are assigned to interfaces, not nodes. 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.

An IPv6 unicast address refers to a single interface. A single interface may be assigned multiple IPv6 addresses of any type (unicast, anycast, and multicast). There are two exceptions to this model. These are:

- 1) A single address 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.
- 2) Routers may have unnumbered interfaces (i.e., no IPv6 address assigned to the interface) on point-to-point links to eliminate the necessity to manually configure and advertise the addresses. Addresses are not needed for point-to-point interfaces on routers if those interfaces are not to be used as the origins or destinations of any IPv6 datagrams.

Currently IPv6 continues the IPv4 model that a subnet is associated with one link. Multiple subnets 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 the method 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 multiple groups of 16-bits of zeros. The "::" can only appear once in an address. The "::" can also be used to compress the leading and/or trailing zeros in an address.

For example the following addresses:

1080:0:0:0:8:800:200C:417A	a unicast address
FF01:0:0:0:0:0:0:43	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::43	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. 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 Representation

The specific type of an IPv6 address is indicated by the leading bits in the address. The variable-length field comprising these leading bits is called the Format Prefix (FP). The initial allocation of these prefixes is as follows:

Allocation	Prefix (binary)	Fraction of Address Space
-----	-----	-----
Reserved	0000 0000	1/256
Unassigned	0000 0001	1/256
Reserved for NSAP Allocation	0000 001	1/128
Reserved for IPX Allocation	0000 010	1/128
Unassigned	0000 011	1/128
Unassigned	0000 1	1/32
Unassigned	0001	1/16
Aggregatable Global Unicast Addresses	001	1/8
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" (see [section 2.4.2](#)), the loopback address (see [section 2.4.3](#)), and the IPv6 Addresses with Embedded IPv4 Addresses (see [section 2.4.4](#)), are assigned out of the 0000 0000 format prefix space.
- (2) The format prefixes 001 and higher, except for Multicast Addresses (1111 1111), are all required to have to have 64-bit interface identifiers in EUI-64 format. See [section 2.5.1](#) for definitions.

This allocation supports the direct allocation of aggregation addresses, local use addresses, and multicast addresses. Space is reserved for NSAP addresses and IPX addresses. The remainder of the address space is unassigned for future use. This can be used for

expansion of existing use (e.g., additional aggregatable addresses, etc.) or new uses (e.g., separate locators and identifiers). Fifteen percent of the address space is initially allocated. The remaining 85% is reserved for future use.

Unicast addresses are distinguished from multicast addresses by the value of the high-order octet of the addresses: a value of FF (11111111) identifies an address as a multicast address; any other value identifies an address as a unicast address. Anycast addresses are taken from the unicast address space, and are not syntactically distinguishable from unicast addresses.

2.5 Unicast Addresses

The IPv6 unicast address is contiguous bit-wise maskable, similar to IPv4 addresses under Class-less Interdomain Routing [[CIDR](#)].

There are several forms of unicast address assignment in IPv6, including the global aggregatable global unicast address, the NSAP address, the IPX hierarchical address, the site-local address, the link-local address, and the IPv4-capable host address. Additional address types 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:

128 bits
node address

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:

n bits	128-n bits
subnet prefix	interface ID

Still more sophisticated hosts may be aware of other hierarchical

boundaries in the unicast address. 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 on that link. They may also be unique over a broader scope. In many cases an interface's identifier will be the same as that interface's link-layer address.

In a number of the format prefixes (see [section 2.4](#)) Interface IDs are required to be 64 bits long and to be constructed in IEEE EUI-64 format [EUI-64]. EUI-64 based Interface identifiers may have global scope when a global token is available or may have local scope where a global token is not available (e.g., serial links, tunnel endpoints, etc.). It is required that the "u" bit (universal/local bit in IEEE EUI-64 terminology) be set correctly in the IEEE EUI-64 address. The "u" bit is set to zero (0) to indicate global scope, and it is set to one (1) to indicate local scope. The first three bytes in binary of an 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 group/individual bit, and "c" are the bits of the company_id.

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

The details on forming interface identifiers is 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 datagrams 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 datagrams or in IPv6 Routing Headers.

[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 datagram to itself. It may never be assigned to any interface.

The loopback address must not be used as the source address in IPv6 datagrams that are sent outside of a single node. An IPv6 datagram with a destination address of loopback must never be sent outside of a single node and must never be forwarded by an IPv6 router.

[2.5.4](#) IPv6 Addresses with Embedded IPv4 Addresses

The IPv6 transition mechanisms include a technique for hosts and routers to dynamically tunnel IPv6 packets over IPv4 routing infrastructure. IPv6 nodes that utilize this technique are assigned special IPv6 unicast addresses that carry an 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	
+-----+-----+-----+		

A second type of IPv6 address which holds an embedded IPv4 address is also defined. This address is used to represent the addresses of IPv4-only nodes (those that **do not** support IPv6) as IPv6 addresses. This type of address is termed an "IPv4-mapped IPv6 address" and has the format:




[2.5.5](#) NSAP Addresses

This mapping of NSAP address into IPv6 addresses is defined in [[NSAP](#)]. This document recommends that network implementors who have planned or deployed an OSI NSAP addressing plan, and who wish to deploy or transition to IPv6, should redesign a native IPv6 addressing plan to meet their needs. However, it also defines a set of mechanisms for the support of OSI NSAP addressing in an IPv6 network. These mechanisms are the ones that must be used if such support is required. This document also defines a mapping of IPv6 addresses within the OSI address format, should this be required.

[2.5.6](#) IPX Addresses

This mapping of IPX address into IPv6 addresses is as follows:



The draft definition, motivation, and usage are under study.

[2.5.7](#) Aggregatable Global Unicast Addresses

The global aggregatable global unicast address is defined in [[AGGR](#)]. This address format is designed to support both the current provider based aggregation and a new type of aggregation called exchanges. The combination will allow efficient routing aggregation for both sites which connect directly to providers and who connect to exchanges. Sites will have the choice to connect to either type of aggregation point.

The IPv6 aggregatable global unicast address format is as follows:


```

+---+-----+-----+-----+-----+
|001| TLA |   NLA*   |   SLA*   |           Interface ID           |
+---+-----+-----+-----+-----+

```

Where

001	Format Prefix (3 bit) for Aggregatable Global Unicast Addresses
TLA	Top Level Aggregator
NLA*	Next Level Aggregator(s)
SLA*	Site-Local Aggregator(s)
INTERFACE ID	Interface Identifier

The contents, field sizes, and assignment rules are defined in [\[AGGR\]](#).

2.5.8 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 auto-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 |   38 bits   |  16 bits  |           64 bits           |
+-----+-----+-----+-----+
|1111111011|   0   | 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.

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 address prefix P that identifies the topological region in which all interfaces belonging to that anycast address reside. Within the region identified by P, each member of the anycast set must be advertised 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 advertisement 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 advertised 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 internet service aggregation. Such addresses could be used as intermediate addresses in an IPv6 Routing header, to cause a packet to be delivered via a particular aggregation or sequence of aggregations. 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 agreed upon for those problems, 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 which they have interfaces.

The subnet-router anycast address is intended to be used for applications where a node needs to communicate with one of a set of routers on a remote subnet. For example when a mobile host needs to communicate with one of the mobile agents on its "home" subnet.

[2.7](#) Multicast Addresses

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:

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

43 (hex), then:

FF01:0:0:0:0:0:0:43 means all NTP servers on the same node as the sender.

FF02:0:0:0:0:0:0:43 means all NTP servers on the same link as the sender.

FF05:0:0:0:0:0:0:43 means all NTP servers at the same site as the sender.

FF0E:0:0:0:0:0:0:43 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 FF15:0:0:0:0:0:0:43 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 datagrams or appear in any routing header.

2.7.1 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

FF02:0:0:0:0:1: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 support a Solicited-Node multicast addresses for every unicast and anycast address it is assigned.

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

2.8 A Node's Required Addresses

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

- o Its Link-Local Address for each interface
- o Assigned Unicast Addresses
- o Loopback Address
- o All-Nodes Multicast Address
- o Solicited-Node Multicast Address for each of its assigned unicast and anycast addresses
- o Multicast Addresses of all other groups which the host belongs.

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

- o Its Link-Local Address for each interface
- o Assigned Unicast Addresses
- o Loopback Address
- o The Subnet-Router anycast addresses for the links it has interfaces.
- o All other Anycast addresses with which the router has been configured.
- o All-Nodes Multicast Address
- o All-Router Multicast Address
- o Solicited-Node Multicast Address for each of its assigned unicast and anycast addresses
- o Multicast Addresses of all other groups which the router belongs.

The only address prefixes which should be predefined in an implementation are the:

- o Unspecified Address
- o Loopback Address
- o Multicast Prefix (FF)
- o Local-Use Prefixes (Link-Local and Site-Local)
- o Pre-Defined Multicast Addresses
- o IPv4-Compatible Prefixes

Implementations should assume all other addresses are unicast unless specifically configured (e.g., anycast addresses).

REFERENCES

- [AGGR] Hinden, R., Deering, S., O'Dell, M., "An Aggregatable Global Unicast Address Format", internet draft, <[draft-ietf-ipngwg-unicast-aggr-00.txt](#)>, May 1997.
- [ALLOC] Rekhter, Y., Li, T., "An Architecture for IPv6 Unicast Address Allocation", [RFC1887](#), December 1995.
- [ANYCST] C. Partridge, T. Mendez, and W. Milliken, "Host Anycasting Service", [RFC1546](#), November 1993.
- [CIDR] V. Fuller, T. Li, K. Varadhan, J. Yu, "Supernetting: an Address Assignment and Aggregation Strategy", [RFC1338](#).
- [ETHER] M. Crawford, "Transmission of IPv6 Packets over Ethernet Networks", Internet Draft, <[draft-ietf-ipngwg-trans-ethernet-00.txt](#)>, March 1997.
- [EUI64] IEEE, "Guidelines for 64-bit Global Identifier (EUI-64) Registration Authority", <http://standards.ieee.org/db/oui/tutorials/EUI64.html>, March 1997.
- [FDDI] M. Crawford, "Transmission of IPv6 Packets over FDDI Networks", Internet Draft, <[draft-ietf-ipngwg-trans-fddi-00.txt](#)>, March 1997.
- [IPv6] S. Deering, R. Hinden, Editors, "Internet Protocol, Version 6 (IPv6) Specification", [RFC1883](#), December 1995.
- [MASGN] R. Hinden, "IPv6 Multicast Address Assignments", Internet Draft, <[draft-ietf-ipngwg-multicast-assgn-02.txt](#)>, May 1997.
- [MULT] S. Deering, "Host Extensions for IP multicasting", [RFC1112](#).
- [NSAP] J. Bound, B. Carpenter, D. Harrington, J. Houldsworth, A. Lloyd, "OSI NSAPs and IPv6", [RFC1888](#), August 1996.
- [RFC2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", [RFC2119](#), [BCP14](#), March 1997.

SECURITY CONSIDERATIONS

Documents of this type do not directly impact the security of the Internet infrastructure or its applications.

AUTHOR'S ADDRESSES

Robert M. Hinden
Ipsilon Networks, Inc.
232 Java Drive
Sunnyvale, CA 94089
USA

phone: +1 408 990-2004
fax: +1 408 743-5677
email: hinden@ipsilon.com

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

CHANGES FROM [RFC-1884](#)

This draft has the following changes from [RFC-1884](#).

- Added notation for defining IPv6 prefixes.
- Changed solicited node multicast definition to use a longer prefix.
- Added site scope all routers multicast address.
- Defined Aggregatable Global Unicast Addresses to use "001" Format Prefix.
- Changed "010" (Provider-Based Unicast) and "100" (Reserved for Geographic) Format Prefixes to Unassigned.
- Added section on Interface ID definition for unicast addresses. Requires use of EUI-64 in range of format prefixes and rules for setting global/local scope bit in EUI-64.
- Updated NSAP text to reflect working in [RFC1888](#).
- Removed protocol specific IPv6 multicast addresses (e.g., DHCP) and referenced the IANA definitions.
- Removed section "Unicast Address Example". Had become OBE.
- Added new references.
- Minor text clarifications and improvements.

