

Network Working Group
Request for Comments: 1884
Category: Standards Track

R. Hinden, Ipsilon Networks
S. Deering, Xerox PARC
Editors
December 1995

IP Version 6 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.

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 Address Type Representation.....	5
2.4 Unicast Addresses.....	7
2.4.1 Unicast Address Example.....	8
2.4.2 The Unspecified Address.....	9
2.4.3 The Loopback Address.....	9
2.4.4 IPv6 Addresses with Embedded IPv4 Addresses.....	9
2.4.5 NSAP Addresses.....	10
2.4.6 IPX Addresses.....	10
2.4.7 Provider-Based Global Unicast Addresses.....	10
2.4.8 Local-use IPv6 Unicast Addresses.....	11
2.5 Anycast Addresses.....	12
2.5.1 Required Anycast Address.....	13
2.6 Multicast Addresses.....	14
2.6.1 Pre-Defined Multicast Addresses.....	15
2.7 A Node's Required Addresses.....	17
REFERENCES.....	18
SECURITY CONSIDERATIONS.....	18
DOCUMENT EDITOR'S ADDRESSES.....	18

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 editors would like to acknowledge the contributions of Paul Francis, Jim Bound, Brian Carpenter, Deborah Estrin, Peter Ford, Bob Gilligan, Christian Huitema, Tony Li, Greg Minshall, Erik Nordmark, Yakov Rekhter, Bill Simpson, and Sue Thomson.

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.

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 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
Unassigned	001	1/8
Provider-Based Unicast Address	010	1/8
Unassigned	011	1/8
Reserved for Geographic- Based Unicast Addresses	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 Use Addresses	1111 1110 10	1/1024
Site Local Use Addresses	1111 1110 11	1/1024
Multicast Addresses	1111 1111	1/256

Note: 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.

This allocation supports the direct allocation of provider addresses, local use addresses, and multicast addresses. Space is reserved for NSAP addresses, IPX addresses, and geographic addresses. The remainder of the address space is unassigned for future use. This can be used for expansion of existing use (e.g., additional provider 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.

2.4 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 provider based unicast address, the geographic based unicast address, the NSAP address, the IPX hierarchical address, the site-local-use address, the link-local-use 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:



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:




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.

2.4.1 Unicast Address Examples


An example of a Unicast address format which will likely be common on LANs and other environments where IEEE 802 MAC addresses are available is:



Where the 48-bit Interface ID is an IEEE-802 MAC address. The use of IEEE 802 MAC addresses as a interface ID is expected to be very common in environments where nodes have an IEEE 802 MAC address. In other environments, where IEEE 802 MAC addresses are not available, other types of link layer addresses can be used, such as E.164 addresses, for the interface ID.

The inclusion of a unique global interface identifier, such as an IEEE MAC address, makes possible a very simple form of auto-configuration of addresses. A node may discover a subnet ID by listening to Router Advertisement messages sent by a router on its attached link(s), and then fabricating an IPv6 address for itself by using its IEEE MAC address as the interface ID on that subnet.

Another unicast address format example is where a site or organization requires additional layers of internal hierarchy. In this example the subnet ID is divided into an area ID and a subnet ID. Its format is:



This technique can be continued to allow a site or organization to add additional layers of internal hierarchy. It may be desirable to use an interface ID smaller than a 48-bit IEEE 802 MAC address to allow more space for the additional layers of internal hierarchy. These could be interface IDs which are administratively created by

the site or organization.

2.4.2 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.


2.4.3 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.

2.4.4 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:




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:



2.4.5 NSAP Addresses


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



The draft definition, motivation, and usage are under study [NSAP].

2.4.6 IPX Addresses


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



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

2.4.7 Provider-Based Global Unicast Addresses

The global provider-based unicast address is assigned as described in [ALLOC]. This initial assignment plan for these unicast addresses is similar to assignment of IPv4 addresses under the CIDR scheme [CIDR]. The IPv6 global provider-based unicast address format is as follows:



The high-order part of the address is assigned to registries, who then assign portions of the address space to providers, who then assign portions of the address space to subscribers, etc.

The registry ID identifies the registry which assigns the provider portion of the address. The term "registry prefix" refers to the high-order part of the address up to and including the registry ID.


The provider ID identifies a specific provider which assigns the subscriber portion of the address. The term "provider prefix" refers to the high-order part of the address up to and including the provider ID.

The subscriber ID distinguishes among multiple subscribers attached to the provider identified by the provider ID. The term "subscriber prefix" refers to the high-order part of the address up to and including the subscriber ID.

The intra-subscriber portion of the address is defined by an individual subscriber and is organized according to the subscribers local internet topology. It is likely that many subscribers will choose to divide the intra-subscriber portion of the address into a subnet ID and an interface ID. In this case the subnet ID identifies a specific physical link and the interface ID identifies a single interface on that subnet.

2.4.8 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:



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 addresses.

Site-Local addresses have the following format:

10 bits	n bits	m bits	118-n-m bits
1111111011	0	subnet ID	interface ID

Site-Local addresses may be used for sites or organizations that are not (yet) connected to the global Internet. They do not need to request or "steal" an address prefix from the global Internet address space. IPv6 site-local addresses can be used instead. When the organization connects to the global Internet, it can then form global addresses by replacing the site-local prefix with a subscriber prefix.

Routers MUST not forward any packets with site-local source addresses outside of the site.

2.5 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 provider. Such addresses could be used as intermediate addresses in an IPv6 Routing header, to cause a packet to be delivered via a particular provider or sequence of 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 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.

2.5.1 Required Anycast Address

The Subnet-Router anycast address is predefined. It's format is as follows:

n bits	128-n bits
subnet prefix	00000000000000


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 it's "home" subnet.

2.6 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:



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 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

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 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.

2.6.1 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

```

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

```

DHCP Server/Relay-Agent:  FF02:0:0:0:0:0:0:C

```

The above multicast addresses identify the group of all IPv6 DHCP Servers and Relay Agents within scope 2 (link-local).

```

Solicited-Node Address:  FF02:0:0:0:0:1:XXXX: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 32 bits of the address (unicast or anycast) and appending those bits to the 96-bit prefix FF02:0:0:0:0:1 resulting in a multicast address in the range

```

FF02:0:0:0:0:1:0000:0000

```

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::1:200E:8C6C. IPv6 addresses that differ only in the high-order bits, e.g., due to multiple high-order prefixes associated with different providers,

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.

2.7 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

- [ALLOC] Rekhter, Y., and T. Li, "An Architecture for IPv6 Unicast Address Allocation", RFC 1887, cisco Systems, December 1995.
- [ANYCST] Partridge, C., Mendez, T., and W. Milliken, "Host Anycasting Service", RFC 1546, BBN, November 1993.
- [CIDR] Fuller, V., Li, T., Varadhan, K., and J. Yu, "Supernetting: an Address Assignment and Aggregation Strategy", RFC 1338, BARRNet, cisco, Merit, OARnet, June 1992.
- [IPV6] Deering, S., and R. Hinden, Editors, "Internet Protocol, Version 6 (IPv6) Specification", RFC 1883, Xerox PARC, Ipsilon Networks, December 1995.
- [MULT] Deering, S., "Host Extensions for IP multicasting", STD 5, RFC 1112, Stanford University, August 1989.
- [NSAP] Carpenter, B., Editor, "Mechanisms for OSIN SAPs, CLNP and TP over IPv6", Work in Progress.

SECURITY CONSIDERATIONS

Security issues are not discussed in this document.

DOCUMENT EDITOR'S ADDRESSES

Robert M. Hinden
Ipsilon Networks, Inc.
2191 E. Bayshore Road, Suite 100
Palo Alto, CA 94303
USA

Phone: +1 415 846 4604
Fax: +1 415 855 1414
EMail: hinden@ipsilon.com

Stephen E. Deering
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
USA

Phone: +1 415 812 4839
Fax: +1 415 812 4471
EMail: deering@parc.xerox.com

