

IP Version 6 Addressing Architecture

<[draft-ietf-ipngwg-addr-arch-v3-05.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of \[RFC2026\]](#).

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 view the list Internet-Draft Shadow Directories, see <http://www.ietf.org/shadow.html>.

This Internet Draft expires September 2, 2001.

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 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 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.....11](#)
    - [2.5.4 IPv6 Addresses with Embedded IPv4 Addresses.....11](#)
    - [2.5.5 NSAP Addresses.....12](#)
    - [2.5.6 Aggregatable Global Unicast Addresses.....12](#)
    - [2.5.7 Local-use IPv6 Unicast Addresses.....12](#)
  - [2.6 Anycast Addresses.....13](#)
    - [2.6.1 Required Anycast Address.....14](#)
  - [2.7 Multicast Addresses.....15](#)
    - [2.7.1 Pre-Defined Multicast Addresses.....17](#)
    - [2.7.2 Assignment of New IPv6 Multicast Addresses.....18](#)
  - [2.8 A Node's Required Addresses.....19](#)
- [3. Security Considerations.....19](#)
- [APPENDIX A: Creating EUI-64 based Interface Identifiers.....20](#)
- [APPENDIX B: CHANGES FROM \[RFC-2373\]\(#\).....23](#)
- [REFERENCES.....24](#)
- [AUTHOR'S ADDRESSES.....25](#)



## **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, 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, Roger Fajman, Markku Savela, and Larry Masinter.

## **2.0 IPv6 ADDRESSING**

IPv6 addresses are 128-bit identifiers for interfaces and sets of interfaces. There are three types of addresses:

- a 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 zero-valued fields or end in zeros.



## 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 be assigned 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 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
FF01:0:0:0:0:0:0:101	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: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:0000:0000:0000:0000:0000:CD30

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

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 Space	Prefix (binary)	Fraction of Address
----- -----	-----	
Reserved	0000 0000	1/256
Unassigned	0000 0001	1/256
Reserved for NSAP Allocation	0000 001	1/128
Unassigned	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 2.5.2](#)), the loopback address (see [section 2.5.3](#)), and the IPv6 Addresses with Embedded IPv4 Addresses (see [section 2.5.4](#)), are assigned out of the 0000 0000 format prefix space.
- (2) The format prefixes 001 through 111, except for Multicast Addresses (1111 1111), are all required to have 64-bit interface identifiers in EUI-64 format. See [section 2.5.1](#) for definitions.
- (3) All format prefixes labelled "Unassigned" must be treated as if they were global unicast addresses, unless and until a future standards document defines an exception for a particular prefix. Thus, an implementation must not have "wired-in" knowledge that a particular format prefix



all identifies global unicast addresses, but rather must treat

addresses as global unicast unless they fall under one of the defined exceptions or are explicitly configured otherwise (as in the case of anycast addresses). The currently defined exceptions are as follows:

::/128	the unspecified address
::1/128	the loopback address
::/96	IPv4-compatible address prefix
::FFFF:0:0/96	IPv4-mapped address prefix
FE80::/10	link-local unicast address prefix
FEC0::/10	site-local unicast address prefix
FF00::/8	multicast address prefix

This allocation supports the direct allocation of aggregation addresses, local use addresses, and multicast addresses. Space is reserved for NSAP 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**

IPv6 unicast addresses are aggregatable with contiguous bit-wise masks similar to IPv4 addresses under Class-less Interdomain Routing [[CIDR](#)].


There are several forms of unicast address assignment in IPv6, including the aggregatable global unicast address, the NSAP 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:





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.

### 2.5.1 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 derived directly from that interface's link-layer address. The same interface identifier may be used on multiple interfaces on a single node.


Note that the use of the same interface identifier on multiple interfaces of a single node does not affect the interface identifier's global uniqueness or each IPv6 address's global uniqueness created using that interface identifier.

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 [[EUI64](#)]. EUI-64 based Interface identifiers may have global scope when a global token is available (e.g., IEEE 48bit MAC) or may have local scope where a global token is not available (e.g., serial links, tunnel end-points, etc.). It is required that the "u" bit (universal/local bit in IEEE EUI-64 terminology) be inverted when





forming the interface identifier from the EUI-64. 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 EUI-64 identifier are as follows:



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 EUI-64 based Interface

Identifiers" provides examples on the creation of different EUI-64 based interface identifiers.

The motivation for inverting the "u" bit when forming the interface identifier is to make it easy for system administrators to hand configure local scope identifiers when hardware tokens are not available. This is expected to be case for serial links, tunnel end-points, 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 IEEE EUI-64 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 may be thought of as being associated with a virtual interface (e.g., the loopback interface).


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 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 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 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:






**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 Aggregatable Global Unicast Addresses**


The aggregatable global unicast address is defined in [AGGR]. It is a specific example of the following general format for global unicast addresses:



where FP is a format prefix (see section 2.4), and the format of the routing prefix, subnet ID, and interface ID are specific to a given FP. Note that, as specified in section 2.4, for all FPs whose high-order three bits have binary value 001 or greater, excluding the multicast FP 11111111, the interface ID is a 64-bit field in EUI-64 format.

**2.5.7 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 or destination addresses to other links.

Site-Local addresses have the following format:



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

[draft-ietf-ipngwg-addr-arch-v3-05.txt](#)

13]

[Page

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 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 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:



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







0 reserved  
1 interface-local scope  
2 link-local scope  
3 reserved for subnet-local scope  
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

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 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 at 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 outside of the scope indicated by the scop field in the 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 ID's for any other scope values 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

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 join the associated Solicited-Node multicast addresses for every unicast and anycast address it is assigned.

### 2.7.2 Assignment of New IPv6 Multicast Addresses

The current approach [[ETHER](#)] to map IPv6 multicast addresses into IEEE 802 MAC addresses takes the low order 32 bits of the IPv6 multicast address and uses it to create a MAC address. Note that Token Ring networks are handled differently. This is defined in [[TOKEN](#)]. IPv6 multicast addresses, in which the 32-bit group ID portion of the address are unique will generate unique MAC addresses.

Due to this, new IPv6 multicast addresses should be assigned with an attempt at keeping the group ID portions of the addresses unique with

respect to all other IPv6 multicast addresses. Note that address assignments that result in MAC address collisions will still work correctly, but possibly with reduced performance in some cases. Thus, when assigning IPv6 multicast addresses, the performance costs of such collisions should be weighed against the cost of assigning addresses in a way that ensures the group ID maps into a unique MAC address.





While this limits the number of permanent IPv6 multicast groups to  $2^{32}$  this is unlikely to be a limitation in the future. If it becomes necessary to exceed this limit in the future multicast will still work but the processing will be slightly slower.

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 The All-Nodes Multicast Addresses defined in [section 2.7.1](#)
- o Solicited-Node Multicast Address for each of its assigned unicast and anycast addresses
- o Multicast Addresses of all other groups to which the host 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 the interfaces it is configured to act as a router on.
- o All other Anycast addresses with which the router has been configured.
- o The All-Routers Multicast Addresses defined in [section 2.7.1](#)
- o Multicast Addresses of all other groups to which the router belongs.

## **3. Security Considerations**

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




APPENDIX A : Creating EUI-64 based Interface Identifiers

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


Links or Nodes with EUI-64 Identifiers

The only change needed to transform an EUI-64 identifier to an interface identifier is to invert the "u" (universal/local) bit.

For example, a globally unique EUI-64 identifier of the form:



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:



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 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 MAC with global scope:


[draft-ietf-ipngwg-addr-arch-v3-05.txt](#)

20]

[Page




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:



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

### Links with Non-Global Identifiers

There are a number of types of links that, while multi-access, do not have globally unique link identifiers. Examples include LocalTalk and Arcnet. The method to create an EUI-64 formatted 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:



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 for

the link.

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. To use this approach no other interface connecting the same node to the same link 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 on the link. There are many possible approaches to select a link-unique interface identifier. They include:

- Manual Configuration
- Generated Random Number
- Node Serial Number (or other node-specific token)

The link-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":

- Added clarification that all unassigned FP's must be treated as global unicast addresses unless another standards document defines an exception.
- Added clarification that implementations must not have "wired-in" knowledge that a particular format prefix identifies global unicast addresses and that implementations must treat all addresses as global unicast addresses unless they fall under a defined exception.
- Moved the list of defined prefix exceptions from [section 2.5.6](#) to [section 2.4](#).
- Revised [section 2.5.6](#) "Aggregatable Global Unicast Addresses" to generalize the description in this section and clarify that the specifics are defined in the reference.
- 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.
- Changed the address block reserved for IPX addresses to unassigned.
- Multicast scope name changes:
  - o Changed name of scope value 1 from "node-local" to "interface-local"
  - o Reserved scope value 3 for "subnet-local" for multi-link subnets
  - o Defined scope value 4 as "admin-local"
- Clarified text describing Exchanges.
- Corrected reference to [RFC1933](#) and updated references.
- Several minor textual clarifications.



REFERENCES

- [AGGR] Hinden, R., S. Deering, M. O'Dell, "An Aggregatable Global Unicast Address Format", [RFC2374](#), July 1998.
- [AUTH] Kent, S., R. Atkinson, "IP Authentication Header", [RFC2402](#), November 1998.
- [ANYCST] Partridge, C., T. Mendez, and W. Milliken, "Host Anycasting Service", [RFC1546](#), November 1993.
- [CIDR] Fuller, V., Li, T., Yu, J., Varadhan, K., "Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy", [RFC1519](#), September 1993.
- [ETHER] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", [RFC2464](#), 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", [RFC2467](#), December 1998.
- [IPV6] Deering, S., R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC2460](#), December 1998.
- [MASGN] Hinden, R., "IPv6 Multicast Address Assignments", [RFC2375](#), July 1998.
- [NSAP] Bound, J., B. Carpenter, D. Harrington, J. Houldsworth, A. Lloyd, "OSI NSAPs and IPv6", [RFC1888](#), August 1996.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", [RFC2026](#), [BCP00009](#), October 1996.
- [TOKEN] Crawford, M., T. Narten, S. Thomas, "Transmission of IPv6 Packets over Token Ring Networks", [RFC2470](#), December 1998.
- [TRAN] Gilligan, R., E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", [RFC1933](#), April 1996.



AUTHOR'S 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



