IP Version 6 Addressing Architecture

Status of This Memo

Copyright Notice

Abstract

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

   This document obsoletes RFC 3513, "IP Version 6 Addressing
   Architecture".

Table of Contents

1. Introduction

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

2. IPv6 Addressing

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

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

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

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

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   superseded by multicast addresses.

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

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

2.1. Addressing Model

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

   All interfaces are required to have at least one Link-Local unicast
   address (see Section 2.8 for additional required addresses). A
   single interface may also have multiple IPv6 addresses of any type
   (unicast, anycast, and multicast) or scope. Unicast addresses with a
   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 in 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 one to
   four hexadecimal digits of the eight 16-bit pieces of the address.
   Examples:

       ABCD:EF01:2345:6789:ABCD:EF01:2345:6789

       2001:DB8:0:0:8:800:200C:417A

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

2. Due to some methods of allocating certain styles of IPv6
   addresses, it will be common for addresses to contain long strings
   of zero bits. In order to make writing addresses containing zero
   bits easier, a special syntax is available to compress the zeros.
   The use of "::" indicates one or more groups of 16 bits of zeros.
   The "::" can only appear once in an address. The "::" can also be
   used to compress leading or trailing zeros in an address.

   For example, the following addresses

       2001:DB8:0:0:8:800:200C: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 address

   may be represented as

       2001:DB8::8:800:200C:417A   a unicast address
       FF01::101   a multicast address
       ::1   the loopback address
       ::   the unspecified address

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 address prefixes are written in Classless Inter-Domain
   Routing (CIDR) notation [CIDR]. An IPv6 address prefix is
   represented by the notation:

       ipv6-address/prefix-length

   where

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

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

   For example, the following are legal representations of the 60-bit
   prefix 20010DB80000CD3 (hexadecimal):

       2001:0DB8:0000:CD30:0000:0000:0000:0000/60
       2001:0DB8::CD30:0:0:0:0/60
       2001:0DB8:0:CD30::/60

   The following are NOT legal representations of the above prefix:

       2001:0DB8:0:CD3/60   may drop leading zeros, but not trailing
                              zeros, within any 16-bit chunk of the address

       2001:0DB8::CD30/60   address to left of "/" expands to
                              2001:0DB8:0000:0000:0000:0000:0000:CD30

       2001:0DB8::CD3/60   address to left of "/" expands to
                              2001:0DB8:0000:0000:0000:0000:0000:0CD3

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   2001:0DB8:0:CD30:123:4567:89AB:CDEF
    and its subnet number 2001:0DB8:0:CD30::/60

    can be abbreviated as 2001:0DB8:0:CD30:123:4567:89AB:CDEF/60

2.4. Address Type Identification

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

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

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

   The general format of Global Unicast addresses is described in
   Section 2.5.4. Some special-purpose subtypes of Global Unicast
   addresses that contain embedded IPv4 addresses (for the purposes of
   IPv4-IPv6 interoperation) are described in Section 2.5.5.

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

2.5. Unicast Addresses

   IPv6 unicast addresses are aggregatable with prefixes of arbitrary
   bit-length, similar to IPv4 addresses under Classless Inter-Domain
   Routing.

   There are several types of unicast addresses in IPv6, in particular,
   Global Unicast, site-local unicast (deprecated, see Section 2.5.7),
   and Link-Local unicast. There are also some special-purpose subtypes
   of Global Unicast, such as IPv6 addresses with embedded IPv4
   addresses. Additional address types or subtypes can be defined in
   the future.

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

```
|                          128 bits                                 |
+-------------------------------------------------------------------+
|                         node address                              |
+-------------------------------------------------------------------+
```
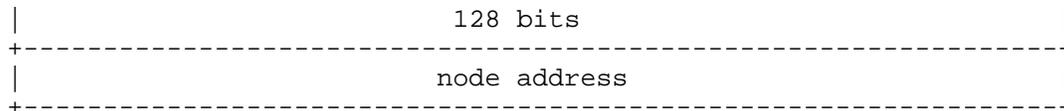
A slightly sophisticated host (but still rather simple) may
additionally be aware of subnet prefix(es) for the link(s) it is
attached to, where different addresses may have different values for
n:

```
|          n bits               |           128-n bits              |
+-------------------------------+-----------------------------------+
|        subnet prefix          |            interface ID           |
+-------------------------------+-----------------------------------+
```
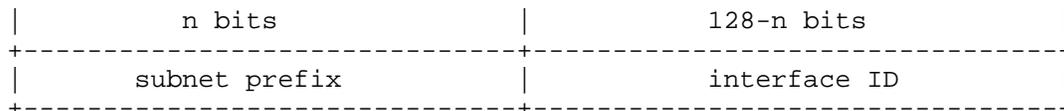
Though a very simple router may have no knowledge of the internal
structure of IPv6 unicast addresses, routers will more generally have
knowledge of one or more of the hierarchical boundaries for the
operation of routing protocols. The known boundaries will differ
from router to router, depending on what positions the router holds
in the routing hierarchy.

Except for the knowledge of the subnet boundary discussed in the
previous paragraphs, nodes should not make any assumptions about the
structure of an IPv6 address.

2.5.1. Interface Identifiers

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

Note that the uniqueness of interface identifiers is independent of
the uniqueness of IPv6 addresses. For example, a Global Unicast
address may be created with a local scope interface identifier and a
Link-Local address may be created with a universal scope interface
identifier.

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

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

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

```
      0        0 0       1 1       2
     |0        7 8       5 6       3|
     +----+----+----+----+----+----+
     |cccc|ccug|cccc|cccc|cccc|cccc|
     +----+----+----+----+----+----+
```
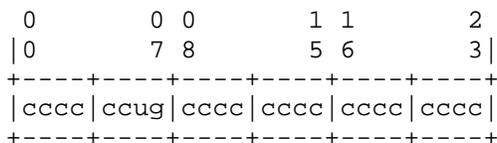
written in Internet standard bit-order, where "u" is the
universal/local bit, "g" is the individual/group bit, and "c" is the
bits of the company_id. Appendix A, "Creating Modified EUI-64 Format
Interface Identifiers", provides examples on the creation of Modified
EUI-64 format-based interface identifiers.

The motivation for inverting the "u" bit when forming an interface
identifier is to make it easy for system administrators to hand
configure non-global identifiers when hardware tokens are not
available. This is expected to be the case for serial links and
tunnel end-points, for example. 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 0:0:0:1, 0:0:0:2, etc.

IPv6 nodes are not required to validate that interface identifiers
created with modified EUI-64 tokens with the "u" bit set to universal
are unique.

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

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

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 must
not be assigned to any physical interface. It is treated as having
Link-Local scope, and may be thought of as the Link-Local unicast
address of a virtual interface (typically called the "loopback
interface") to an imaginary link that goes nowhere.

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

2.5.4. Global Unicast Addresses

The general format for IPv6 Global Unicast addresses is as follows:

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

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

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

   Examples of Global Unicast addresses that start with binary 000 are
   the IPv6 address with embedded IPv4 addresses described in Section
   2.5.5. An example of global addresses starting with a binary value
   other than 000 (and therefore having a 64-bit interface ID field) can
   be found in [GLOBAL].

2.5.5. IPv6 Addresses with Embedded IPv4 Addresses

   Two types of IPv6 addresses are defined that carry an IPv4 address in
   the low-order 32 bits of the address. These are the "IPv4-Compatible
   IPv6 address" and the "IPv4-mapped IPv6 address".

2.5.5.1. IPv4-Compatible IPv6 Address

   The "IPv4-Compatible IPv6 address" was defined to assist in the IPv6
   transition. The format of the "IPv4-Compatible IPv6 address" is as
   follows:

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

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

   The "IPv4-Compatible IPv6 address" is now deprecated because the
   current IPv6 transition mechanisms no longer use these addresses.
   New or updated implementations are not required to support this
   address type.

2.5.5.2. IPv4-Mapped IPv6 Address

   A second type of IPv6 address that holds an embedded IPv4 address is
   defined. This address type is used to represent the addresses of
   IPv4 nodes as IPv6 addresses. The format of the "IPv4-mapped IPv6
   address" is as follows:

```
   |                  80 bits               | 16 |      32 bits        |
   +--------------------------------------+--------------------------+
   |0000..............................0000|FFFF|    IPv4 address      |
   +--------------------------------------+----+--------------------+
```
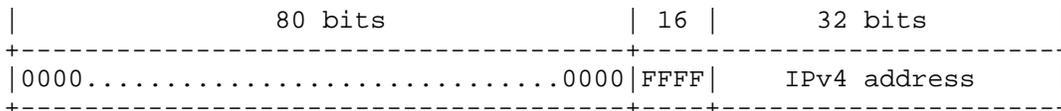
See [RFC4038] for background on the usage of the "IPv4-mapped IPv6
address".

2.5.6. Link-Local IPv6 Unicast Addresses

Link-Local addresses are for use on a single link. Link-Local
addresses have the following format:

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

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

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

2.5.7. Site-Local IPv6 Unicast Addresses

Site-Local addresses were originally designed to be used for
addressing inside of a site without the need for a global prefix.
Site-local addresses are now deprecated as defined in [SLDEP].

Site-Local addresses have the following format:

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

The special behavior of this prefix defined in [RFC3513] must no
longer be supported in new implementations (i.e., new implementations
must treat this prefix as Global Unicast).

Existing implementations and deployments may continue to use this
prefix.

2.6. Anycast Addresses

   An IPv6 anycast address is an address that is assigned to more than
   one interface (typically belonging to different nodes), with the
   property that a packet sent to an anycast address is routed to the
   "nearest" interface having that address, according to the routing
   protocols' measure of distance.

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

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

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

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

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

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              | 00000000000000 |
+---------------------------------------------+----------------+
```
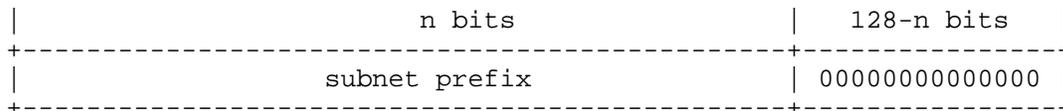
The "subnet prefix" in an anycast address is the prefix that
identifies a specific link. This anycast address is syntactically
the same as a unicast address for an interface on the link with the
interface identifier set to zero.

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

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

2.7. Multicast Addresses

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

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

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

```
                              +-+-+-+-+
   flgs is a set of 4 flags:  |0|R|P|T|
                              +-+-+-+-+
```

   The high-order flag is reserved, and must be initialized to 0.

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

   T = 1 indicates a non-permanently-assigned ("transient" or
   "dynamically" assigned) multicast address.

The P flag's definition and usage can be found in [RFC3306].

The R flag's definition and usage can be found in [RFC3956].

scop is a 4-bit multicast scope value used to limit the scope of
the multicast group. The values are as follows:

    0 reserved
    1 Interface-Local scope
    2 Link-Local scope
    3 reserved
    4 Admin-Local scope
    5 Site-Local scope
    6 (unassigned)
    7 (unassigned)
    8 Organization-Local scope
    9 (unassigned)
    A (unassigned)
    B (unassigned)
    C (unassigned)
    D (unassigned)
    E Global scope
    F reserved

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

Link-Local multicast scope spans the same topological region as
the corresponding unicast scope.

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

Site-Local scope is intended to span a single site.

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

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

group ID identifies the multicast group, either permanent or
transient, within the given scope. Additional definitions of the
multicast group ID field structure are provided in [RFC3306].

The "meaning" of a permanently-assigned multicast address is
independent of the scope value. For example, if the "NTP servers
group" is assigned a permanent multicast address with a group ID of
101 (hex), then

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

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

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

   FF0E:0:0:0:0:0:0:101 means all NTP servers in the Internet.

Non-permanently-assigned multicast addresses are meaningful only
within a given scope. For example, a group identified by the non-
permanent, site-local multicast address FF15:0:0:0:0:0:0:101 at one
site bears no relationship to a group using the same address at a
different site, nor to a non-permanent group using the same group ID
with a different scope, nor to a permanent group with the same group
ID.

Multicast addresses must not be used as source addresses in IPv6
packets or appear in any Routing header.

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

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

2.7.1. Pre-Defined Multicast Addresses

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

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

         Reserved Multicast Addresses: FF00:0:0:0:0:0:0:0
                                         FF01:0:0:0:0:0:0:0
                                         FF02:0:0:0:0:0:0:0
                                         FF03:0:0:0:0:0:0:0
                                         FF04:0:0:0:0:0:0:0
                                         FF05:0:0:0:0:0:0:0
                                         FF06:0:0:0:0:0:0:0
                                         FF07:0:0:0:0:0:0:0
                                         FF08:0:0:0:0:0:0:0
                                         FF09:0:0:0:0:0:0:0
                                         FF0A:0:0:0:0:0:0:0
                                         FF0B:0:0:0:0:0:0:0
                                         FF0C:0:0:0:0:0:0:0
                                         FF0D:0:0:0:0:0:0:0
                                         FF0E:0:0:0:0:0:0:0
                                         FF0F:0:0:0:0:0:0:0

   The above multicast addresses are reserved and shall never be
   assigned to any multicast group.

      All Nodes Addresses: FF01:0:0:0:0:0:0:1
                              FF02:0:0:0:0:0:0:1

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

      All Routers Addresses: FF01:0:0:0:0:0:0:2
                               FF02:0:0:0:0:0:0:2
                               FF05:0:0:0:0:0:0:2

   The above multicast addresses identify the group of all IPv6 routers,
   within scope 1 (interface-local), 2 (link-local), or 5 (site-local).

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

   Solicited-Node multicast address are computed as a function of a
   node's unicast and anycast addresses. A Solicited-Node multicast
   address is formed by taking the low-order 24 bits of an address
   (unicast or anycast) and appending those bits to the prefix
   FF02:0:0:0:0:1:FF00::/104 resulting in a multicast address in the
   range

         FF02:0:0:0:0:1:FF00:0000

   to

         FF02:0:0:0:0:1:FFFF:FFFF

For example, the Solicited-Node multicast address corresponding to
the IPv6 address 4037::01:800:200E:8C6C is FF02::1:FF0E:8C6C. IPv6
addresses that differ only in the high-order bits (e.g., due to
multiple high-order prefixes associated with different aggregations)
will map to the same Solicited-Node address, thereby reducing the
number of multicast addresses a node must join.

A node is required to compute and join (on the appropriate interface)
the associated Solicited-Node multicast addresses for all unicast and
anycast addresses that have been configured for the node's interfaces
(manually or automatically).

2.8. A Node's Required Addresses

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

     o Its required Link-Local address for each interface.

     o Any additional Unicast and Anycast addresses that have been
       configured for the node's interfaces (manually or
       automatically).

     o The loopback address.

     o The All-Nodes multicast addresses defined in Section 2.7.1.

     o The Solicited-Node multicast address for each of its unicast and
       anycast addresses.

     o Multicast addresses of all other groups to which the node
       belongs.

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

     o The Subnet-Router Anycast addresses for all interfaces for which
       it is configured to act as a router.

     o All other Anycast addresses with which the router has been
       configured.

     o The All-Routers multicast addresses defined in Section 2.7.1.

3. Security Considerations

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

4. IANA Considerations

   The "IPv4-Compatible IPv6 address" is deprecated by this document.
   The IANA should continue to list the address block containing these
   addresses at http://www.iana.org/assignments/ipv6-address-space as
   "Reserved by IETF" and not reassign it for any other purpose. For
   example:

      0000::/8   Reserved by IETF   [RFC3513]   [1]

   The IANA has added the following note and link to this address block.

      [5] 0000::/96 was previously defined as the "IPv4-Compatible IPv6
           address" prefix. This definition has been deprecated by RFC
           4291.

   The IANA has updated the references for the IPv6 Address Architecture
   in the IANA registries accordingly.

5. Acknowledgements

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

6. References

6.1. Normative References

   [IPV6]   Deering, S. and R. Hinden, "Internet Protocol, Version 6
              (IPv6) Specification", RFC 2460, December 1998.

6.2. Informative References

   [AUTH]   Kent, S. and R. Atkinson, "IP Authentication Header", RFC
   2402, November 1998.

   [CIDR]   Fuller, V., Li, T., Yu, J., and K. Varadhan, "Classless
               Inter-Domain Routing (CIDR): an Address Assignment and
               Aggregation Strategy", RFC 1519, September 1993.

   [ETHER]   Crawford, M., "Transmission of IPv6 Packets over Ethernet
               Networks", RFC 2464, December 1998.

   [EUI64]   IEEE, "Guidelines for 64-bit Global Identifier (EUI-64)
               Registration Authority",
               http://standards.ieee.org/regauth/oui/tutorials/EUI64.html,
               March 1997.

   [FDDI]   Crawford, M., "Transmission of IPv6 Packets over FDDI
               Networks", RFC 2467, December 1998.

   [GLOBAL]   Hinden, R., Deering, S., and E. Nordmark, "IPv6 Global
               Unicast Address Format", RFC 3587, August 2003.

   [PRIV]   Narten, T. and R. Draves, "Privacy Extensions for Stateless
               Address Autoconfiguration in IPv6", RFC 3041, January 2001.

   [RFC3513]   Hinden, R. and S. Deering, "Internet Protocol Version 6
               (IPv6) Addressing Architecture", RFC 3513, April 2005.

   [RFC3306]   Haberman, B. and D. Thaler, "Unicast-Prefix-based IPv6
               Multicast Addresses", RFC 3306, August 2002.

   [RFC3956]   Savola, P. and B. Haberman, "Embedding the Rendezvous Point
               (RP) Address in an IPv6 Multicast Address", RFC 3956,
               November 2004.

   [RFC4038]   Shin, M-K., Hong, Y-G., Hagino, J., Savola, P., and E.
               Castro, "Application Aspects of IPv6 Transition", RFC 4038,
               March 2005.

   [SLDEP]   Huitema, C. and B. Carpenter, "Deprecating Site Local
               Addresses", RFC 3879, September 2004.

Appendix A: Creating Modified EUI-64 Format Interface Identifiers

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

   Links or Nodes with IEEE EUI-64 Identifiers

   The only change needed to transform an IEEE EUI-64 identifier to an
   interface identifier is to invert the "u" (universal/local) bit. An
   example is a globally unique IEEE EUI-64 identifier of the form:
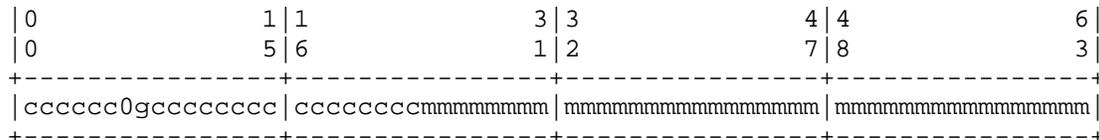
```
|0              1|1              3|3              4|4              6|
|0              5|6              1|2              7|8              3|
+----------------+----------------+----------------+----------------+
|cccccc0gcccccccc|ccccccccmmmmmmmm|mmmmmmmmmmmmmmmm|mmmmmmmmmmmmmmmm|
+----------------+----------------+----------------+----------------+
```

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

```
|0              1|1              3|3              4|4              6|
|0              5|6              1|2              7|8              3|
+----------------+----------------+----------------+----------------+
|cccccc1gcccccccc|ccccccccmmmmmmmm|mmmmmmmmmmmmmmmm|mmmmmmmmmmmmmmmm|
+----------------+----------------+----------------+----------------+
```
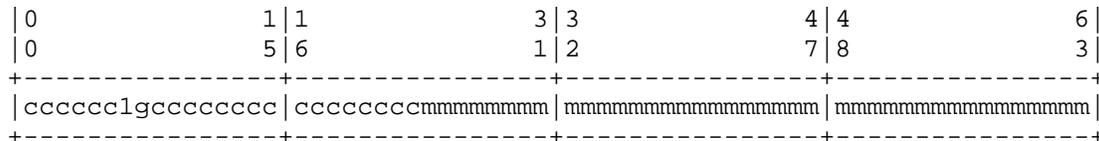
   The only change is inverting the value of the universal/local bit.

   Links or Nodes with IEEE 802 48-bit MACs

   [EUI64] defines a method to create an IEEE EUI-64 identifier from an
   IEEE 48-bit MAC identifier. This is to insert two octets, with
   hexadecimal values of 0xFF and 0xFE (see the Note at the end of
   appendix), in the middle of the 48-bit MAC (between the company_id
   and vendor-supplied id). An example is the 48-bit IEEE MAC with
   Global scope:
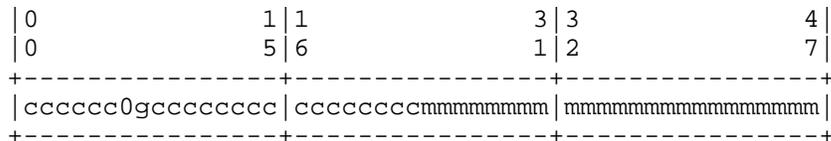
```
|0              1|1              3|3              4|
|0              5|6              1|2              7|
+----------------+----------------+----------------+
|cccccc0gcccccccc|ccccccccmmmmmmmm|mmmmmmmmmmmmmmmm|
+----------------+----------------+----------------+
```

where "c" is 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" is the bits of the manufacturer-
selected extension identifier. The interface identifier would be of
the form:

```
|0              1|1              3|3              4|4              6|
|0              5|6              1|2              7|8              3|
+---------------+---------------+---------------+---------------+
|cccccc1gcccccccc|cccccccc11111111|11111110mmmmmmmm|mmmmmmmmmmmmmmmm|
+---------------+---------------+---------------+---------------+
```
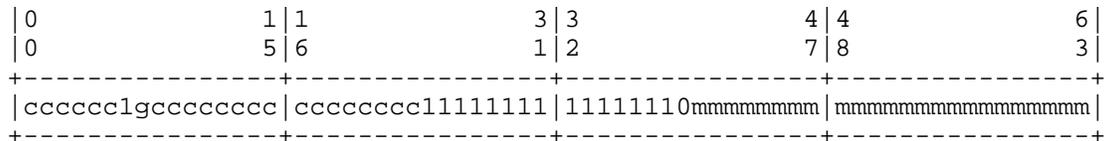
When IEEE 802 48-bit MAC addresses are available (on an interface or
a node), an implementation may use them to create interface
identifiers due to their availability and uniqueness properties.

Links with Other Kinds of Identifiers

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

```
|0              1|1              3|3              4|4              6|
|0              5|6              1|2              7|8              3|
+---------------+---------------+---------------+---------------+
|0000000000000000|0000000000000000|0000000000000000|0000000001001111|
+---------------+---------------+---------------+---------------+
```
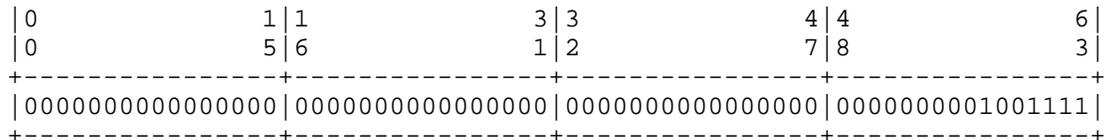
Note that this results in the universal/local bit set to "0" to
indicate local scope.

Links without Identifiers

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

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

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

    Manual Configuration
    Node Serial Number
    Other Node-Specific Token

The subnet-prefix-unique interface identifier should be generated in
a manner such 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.

Note: [EUI-64] actually defines 0xFF and 0xFF as the bits to be
      inserted to create an IEEE EUI-64 identifier from an IEEE MAC-
      48 identifier. The 0xFF and 0xFE values are used when starting
      with an IEEE EUI-48 identifier. The incorrect value was used
      in earlier versions of the specification due to a
      misunderstanding about the differences between IEEE MAC-48 and
      EUI-48 identifiers.

      This document purposely continues the use of 0xFF and 0xFE
      because it meets the requirements for IPv6 interface
      identifiers (i.e., that they must be unique on the link), IEEE
      EUI-48 and MAC-48 identifiers are syntactically equivalent, and
      that it doesn't cause any problems in practice.

Appendix B: Changes from RFC 3513

   The following changes were made from RFC 3513, "IP Version 6
   Addressing Architecture":

   o The restrictions on using IPv6 anycast addresses were removed
     because there is now sufficient experience with the use of anycast
     addresses, the issues are not specific to IPv6, and the GROW
     working group is working in this area.

   o Deprecated the Site-Local unicast prefix. Changes include the
     following:

- Removed Site-Local from special list of prefixes in Section
  2.4.

- Split section titled "Local-use IPv6 Unicast Addresses" into
  two sections, "Link-Local IPv6 Unicast Addresses" and "Site-
  Local IPv6 Unicast Addresses".

- Added text to new section describing Site-Local deprecation.

o Changes to resolve issues raised in IAB response to Robert Elz
  appeal. Changes include the following:

  - Added clarification to Section 2.5 that nodes should make no
    assumptions about the structure of an IPv6 address.

  - Changed the text in Section 2.5.1 and Appendix A to refer to
    the Modified EUI-64 format interface identifiers with the "u"
    bit set to one (1) as universal.

  - Added clarification to Section 2.5.1 that IPv6 nodes are not
    required to validate that interface identifiers created in
    Modified EUI-64 format with the "u" bit set to one are unique.

o Changed the reference indicated in Section 2.5.4 "Global Unicast
  Addresses" to RFC 3587.

o Removed mention of NSAP addresses in examples.

o Clarified that the "x" in the textual representation can be one to
  four digits.

o Deprecated the "IPv6 Compatible Address" because it is not being
  used in the IPv6 transition mechanisms.

o Added the "R" and "P" flags to Section 2.7 on multicast addresses,
  and pointers to the documents that define them.

o Editorial changes.

Authors' Addresses

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

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

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