IPv6 Addressing

In this section, we examine:

- The IPv6 address space
- IPv6 address syntax
- IPv6 prefixes
- Types of IPv6 addresses
- Unicast IPv6 addresses
- Multicast IPv6 addresses
- Anycast IPv6 addresses
- IPv6 addresses for a host
- IPv6 addresses for a router
- IPv6 interface identifiers

The IPv6 Address Space

The most obvious distinguishing feature of IPv6 is its use of much larger addresses. The size of an address in IPv6 is 128 bits, which is four times the larger than an IPv4 address. A 32-bit address space allows for 2^{32} or 4,294,967,296 possible addresses. A 128-bit address space allows for 2^{128} or 340,282,366,920,938,463,463,374,607,431,768,211,456 (or 3.4×10^{38} or 340 undecillion) possible addresses.

In the late 1970s when the IPv4 address space was designed, it was unimaginable that it could be exhausted. However, due to changes in technology and an allocation practice that did not anticipate the recent explosion of hosts on the Internet, the IPv4 address space was consumed to the point that by 1992 it was clear a replacement would be necessary.

With IPv6, it is even harder to conceive that the IPv6 address space will be consumed. To help put this number in perspective, a 128-bit address space provides 655,570,793,348,866,943,898,599 (6.5×10^{23}) addresses for every square meter of the Earth's surface.

It is important to remember that the decision to make the IPv6 address 128 bits in length was not so that every square meter of the Earth could have 6.5×10^{23} addresses. Rather, the relatively large size of the IPv6 address is designed to be subdivided into hierarchical routing domains that reflect the topology of the modern-day Internet. The use of 128 bits allows for multiple levels of hierarchy and flexibility in designing hierarchical addressing and routing that is currently lacking on the IPv4-based Internet.

The IPv6 addressing architecture is described in RFC 4291.

IPv6 Address Syntax

IPv4 addresses are represented in dotted-decimal format. This 32-bit address is divided along 8-bit boundaries. Each set of 8 bits is converted to its decimal equivalent and separated by periods. For IPv6, the 128-bit address is divided along 16-bit boundaries, and each 16-bit block is converted to a 4-digit hexadecimal number and separated by colons. The resulting representation is called colon-hexadecimal.

The following is an IPv6 address in binary form:

```
0010000000000001000011011011100000000000000000000010111100111011
0000001010101010000000001111111111111110001010001001110001011010
```

The 128-bit address is divided along 16-bit boundaries:

```
0010000000000001   0000110110111000   0000000000000000   0010111100111011   0000001010101010
0000000011111111   1111111000101000   1001110001011010
```

Each 16-bit block is converted to hexadecimal and delimited with colons. The result is:

2001:0DB8:0000:2F3B:02AA:00FF:FE28:9C5A

IPv6 representation can be further simplified by removing the leading zeros within each 16-bit block. However, each block must have at least a single digit. With leading zero suppression, the address representation becomes:

2001:DB8:0:2F3B:2AA:FF:FE28:9C5A

Compressing Zeros

Some types of addresses contain long sequences of zeros. To further simplify the representation of IPv6 addresses, a contiguous sequence of 16-bit blocks set to 0 in the colon hexadecimal format can be compressed to "::", known as *double-colon*.

For example, the link-local address of FE80:0:0:0:2AA:FF:FE9A:4CA2 can be compressed to FE80::2AA:FF:FE9A:4CA2. The multicast address FF02:0:0:0:0:0:0:2 can be compressed to FF02::2.

Zero compression can only be used to compress a single contiguous series of 16-bit blocks expressed in colon hexadecimal notation. You cannot use zero compression to include part of a 16-bit block. For example, you cannot express FF02:30:0:0:0:0:0:5 as FF02:3::5. The correct representation is FF02:30::5.

To determine how many 0 bits are represented by the "::", you can count the number of blocks in the compressed address, subtract this number from 8, and then multiply the result by 16. For example, in the address FF02::2, there are two blocks (the "FF02" block and the "2" block.) The number of bits expressed by the "::" is 96 (96 = (8 – 2)× 16).

Zero compression can only be used once in a given address. Otherwise, you could not determine the number of 0 bits represented by each instance of "::".

IPv6 Prefixes

The prefix is the part of the address that indicates the bits that have fixed values or are the bits of the subnet prefix. Prefixes for IPv6 subnets, routes, and address ranges are expressed in the same way as Classless Inter-Domain Routing (CIDR) notation for IPv4. An IPv6 prefix is written in *address/prefix-length* notation. For example, 21DA:D3::/48 and 21DA:D3:0:2F3B::/64 are IPv6 address prefixes.

Note IPv4 implementations commonly use a dotted decimal representation of the network prefix known as the subnet mask. A subnet mask is not used for IPv6. Only the prefix length notation is supported.

Types of IPv6 Addresses

There are three types of IPv6 addresses:

1. Unicast

A unicast address identifies a single interface within the scope of the type of unicast address. With the appropriate unicast routing topology, packets addressed to a unicast address are delivered to a single interface.

2. Multicast
A multicast address identifies multiple interfaces. With the appropriate multicast routing topology, packets addressed to a multicast address are delivered to all interfaces that are identified by the address. A multicast address is used for one-to-many communication, with delivery to multiple interfaces.
3. Anycast
An anycast address identifies multiple interfaces. With the appropriate routing topology, packets addressed to an anycast address are delivered to a single interface, the nearest interface that is identified by the address. The "nearest" interface is defined as being closest in terms of routing distance. An anycast address is used for one-to-one-of-many communication, with delivery to a single interface.

In all cases, IPv6 addresses identify interfaces, not nodes. A node is identified by any unicast address assigned to one of its interfaces.

Note RFC 4291 does not define a broadcast address. All types of IPv4 broadcast addressing are performed in IPv6 using multicast addresses. For example, the subnet and limited broadcast addresses from IPv4 are replaced with the link-local scope all-nodes multicast address of FF02::1.

Links and Subnets

Similar to IPv4, an IPv6 subnet prefix is assigned to a single link. Multiple subnet prefixes can be assigned to the same link. This technique is called *multinetting*.

Unicast IPv6 Addresses

The following types of addresses are unicast IPv6 addresses:

- Global unicast addresses
- Link-local addresses
- Site-local addresses
- Unique local IPv6 unicast addresses
- Special addresses

Global Unicast Addresses

Global unicast addresses are equivalent to public IPv4 addresses. They are globally routable and reachable on the IPv6 portion of the Internet. Unlike the current IPv4-based Internet, which is a mixture of both flat and hierarchical routing, the IPv6-based Internet has been designed from its foundation to support efficient, hierarchical addressing and routing. The scope, the portion of the IPv6 internetwork over which the address is unique, of a global unicast address is the entire IPv6 Internet.

Figure 4 shows the structure of global unicast addresses currently being allocated by IANA, as defined in RFC 3587.

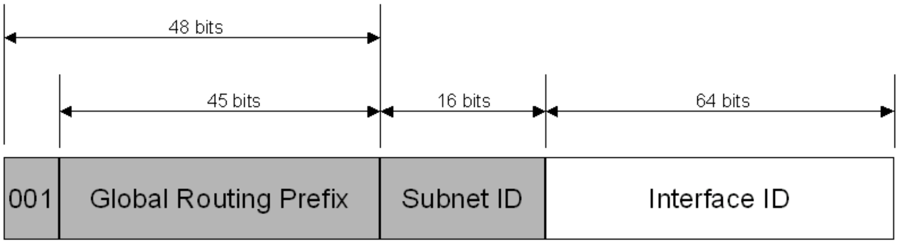

Figure 4 The global unicast address as defined in RFC 3587

The fields in the global unicast address are the following:

Fixed portion set to 001 – The three high-order bits are set to 001. The address prefix for currently assigned global addresses is 2000::/3.

Global Routing Prefix – Indicates the global routing prefix for a specific organization's site. The combination of the three fixed bits and the 45-bit Global Routing Prefix is used to create a 48-bit site prefix, which is assigned to an individual site of an organization. Once assigned, routers on the IPv6 Internet forward IPv6 traffic matching the 48-bit prefix to the routers of the organization's site.

Subnet ID – The Subnet ID is used within an organization's site to identify subnets. The size of this field is 16 bits. The organization's site can use these 16 bits within its site to create 65,536 subnets or multiple levels of addressing hierarchy and an efficient routing infrastructure.

Interface ID – Indicates the interface on a specific subnet within the site. The size of this field is 64 bits.

The fields within the global unicast address create a three-level structure shown in Figure 5.

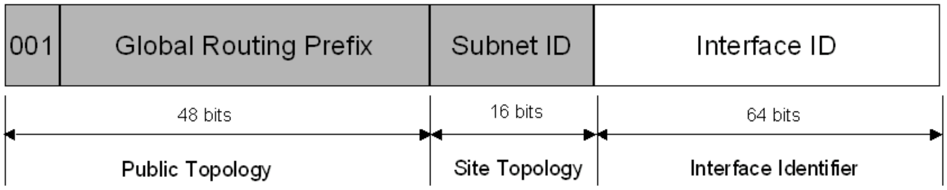

Figure 5 The three-level structure of the global unicast address

The public topology is the collection of larger and smaller ISPs that provide access to the IPv6 Internet. The site topology is the collection of subnets within an organization's site. The interface identifier identifies a specific interface on a subnet within an organization's site. For more information about global unicast addresses, see RFC 3587.

Local-Use Unicast Addresses

There are two types of local-use unicast addresses:

1. Link-local addresses are used between on-link neighbors and for Neighbor Discovery processes.
2. Site-local addresses are used between nodes communicating with other nodes in the same site.

Link-Local Addresses

Link-local addresses are used by nodes when communicating with neighboring nodes on the same link. For example, on a single link IPv6 network with no router, link-local addresses are used to communicate between hosts on the link. IPv6 link-local addresses are equivalent to

IPv4 link-local addresses defined in RFC 3927 that use the 169.254.0.0/16 prefix. IPv4 link-local addresses are known as Automatic Private IP Addressing (APIPA) addresses for computers running current Microsoft Windows operating systems. The scope of a link-local address is the local link.

A link-local address is required for Neighbor Discovery processes and is always automatically configured, even in the absence of all other unicast addresses. For more information on the address autoconfiguration process for link-local addresses, see “Address Autoconfiguration.”

Figure 6 shows the structure of the link-local address.

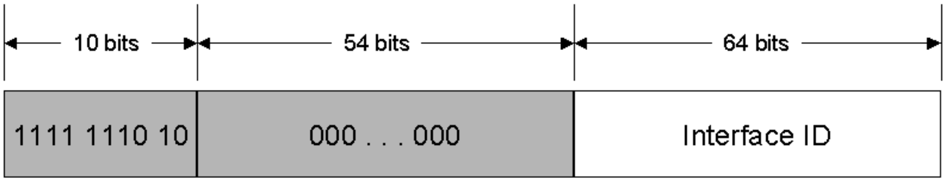

Figure 6 The link-local address

Link-local addresses always begin with FE80. With the 64-bit interface identifier, the prefix for link-local addresses is always FE80::/64. An IPv6 router never forwards link-local traffic beyond the link.

Site-Local Addresses

Site-local addresses are equivalent to the IPv4 private address space (10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16). For example, private intranets that do not have a direct, routed connection to the IPv6 Internet can use site-local addresses without conflicting with global unicast addresses. Site-local addresses are not reachable from other sites, and routers must not forward site-local traffic outside the site. Site-local addresses can be used in addition to global unicast addresses. The scope of a site-local address is the site. A site is an organization network or portion of an organization's network that has a defined geographical location (such as an office, an office complex, or a campus).

Unlike link-local addresses, site-local addresses are not automatically configured and must be assigned either through stateless or stateful address configuration processes. For more information, see “Address Autoconfiguration.”

Figure 7 shows the structure of the site-local address.

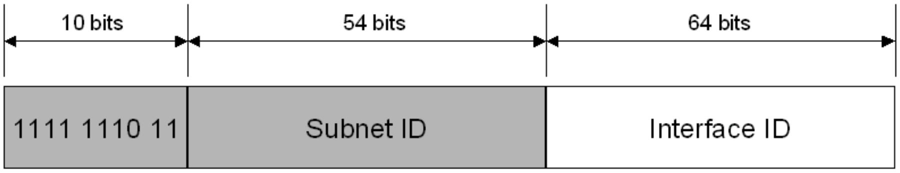

Figure 7 The site-local address

The first 10-bits are always fixed for site-local addresses (FEC0::/10). After the 10 fixed bits is a Subnet ID field that provides 54 bits with which you can create a hierarchical and summarizable routing infrastructure within the site. After the Subnet ID field is a 64-bit Interface ID field that identifies a specific interface on a subnet.

Note RFC 3879 formally deprecates the use of site-local addresses for future IPv6 implementations. Existing implementations of IPv6 can continue to use site-local addresses.

Zone IDs for Local-Use Addresses

Unlike global addresses, local-use addresses can be reused. Link-local addresses are reused on each link. Site-local addresses can be reused within each site of an organization. Because of this address reuse capability, link-local and site-local addresses are ambiguous. To specify which link on which an address is assigned or located or within which site an address is assigned or located, an additional identifier is needed. This additional identifier is a zone identifier (ID), also known as a scope ID, which identifies a connected portion of a network that has a specified scope. The syntax specified in RFC 4007 for identifying the zone associated with a local-use address is the following:

Address%zone_ID

Address is a local-use address and *zone_ID* is an integer value representing the zone. The values of the zone ID are defined relative to the sending host. Therefore, different hosts might determine different zone ID values for the same physical zone. For example, Host A might choose 3 to represent the zone ID of an attached link and Host B might choose 4 to represent the same link.

For Windows-based IPv6 hosts, the zone IDs for link-local and site-local addresses are defined as follows:

- For link-local addresses, the zone ID is typically the interface index of the interface either assigned the address or to be used as the sending interface for a link-local destination. The interface index is an integer starting at 1 that is assigned to IPv6 interfaces, which include a loopback and one or multiple tunnel or LAN interfaces. You can view the list of interface indexes from the display of the **netsh interface ipv6 show interface** command.
- For site-local addresses, the zone ID is the site ID, an integer assigned to the site of an organization. For organizations that do not reuse the site-local address prefix, the site ID is set to 1 by default and does not need to be specified. You can view the site ID from the display of the **netsh interface ipv6 show address level=verbose** command.

The following are examples of using Windows tools and the zone ID:

- **ping fe80::2b0:d0ff:fee9:4143%3** In this case, 3 is the interface index of the interface attached to the link containing the destination address.
- **tracert fec0::f282:2b0:d0ff:fee9:4143%2** In this case, 2 is the site ID of the organization site containing the destination address.

In Windows XP, Windows Server 2003, Windows Vista, and Windows Server "Longhorn," the Ipconfig.exe tool displays the zone ID of local-use IPv6 addresses. The following is an excerpt from the display of the **ipconfig** command:

```
Ethernet adapter Local Area Connection:

        Connection-specific DNS Suffix  . : wcoast.example.com
        IP Address. . . . . . . . . . . . : 157.60.14.219
        Subnet Mask . . . . . . . . . . . : 255.255.255.0
        IP Address. . . . . . . . . . . . : 3ffe:ffff:2a1c:2:1cc8:ef1d:1dd9:8066
        IP Address. . . . . . . . . . . . : 3ffe:ffff:2a1c:204:5aff:fe56:f5b
        IP Address. . . . . . . . . . . . : fe80::204:5aff:fe56:f5b%4
        Default Gateway . . . . . . . . . : 157.60.14.1
                                            fe80::20a:42ff:feb0:5400%4
```

For the link-local addresses in the display of the **ipconfig** command, the zone ID indicates the interface index of the interface either assigned the address (for IP Address) or the interface through which an address is reachable (for Default Gateway).

Unique Local IPv6 Unicast Addresses

Site-local addresses provide a private addressing alternative to using global addresses for intranet traffic. However, because the site-local address prefix can be used to address multiple sites within an organization, a site-local address prefix address can be duplicated. The ambiguity of site-local addresses in an organization adds complexity and difficulty for applications, routers, and network managers. For more information, see section 2 of RFC 3879. To replace site-local addresses with a new type of address that is private to an organization, yet unique across all of the sites of the organization, RFC 4193 defines unique local IPv6 unicast addresses. Figure 4 shows the structure of unique local addresses.

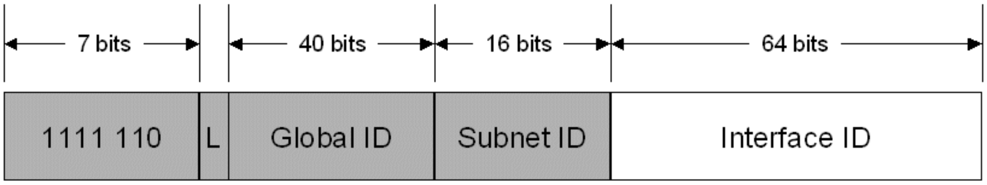

Figure 8 The unique local address

The first 7 bits have the fixed binary value of 1111110. All unique local addresses have the address prefix FC00::/7. The Local (L) flag is set 1 to indicate a local address. The L flag value set to 0 has not yet been defined. Therefore, unique local addresses with the L flag set to 1 have the address prefix of FD00::/8. The Global ID identifies a specific site within an organization and is set to a randomly derived 40-bit value. By deriving a random value for the Global ID, an organization can have statistically unique 48-bit prefixes assigned to the sites of their organizations. Additionally, two organizations that use unique local addresses that merge have a low probability of duplicating a 48-bit unique local address prefix, minimizing site renumbering. Unlike the Global Routing Prefix in global addresses, you should not assign Global IDs in unique local address prefixes so that they can be summarized.

The global address and unique local address share the same structure beyond the first 48 bits of the address. Figure 9 shows the structure of global and unique local addresses.

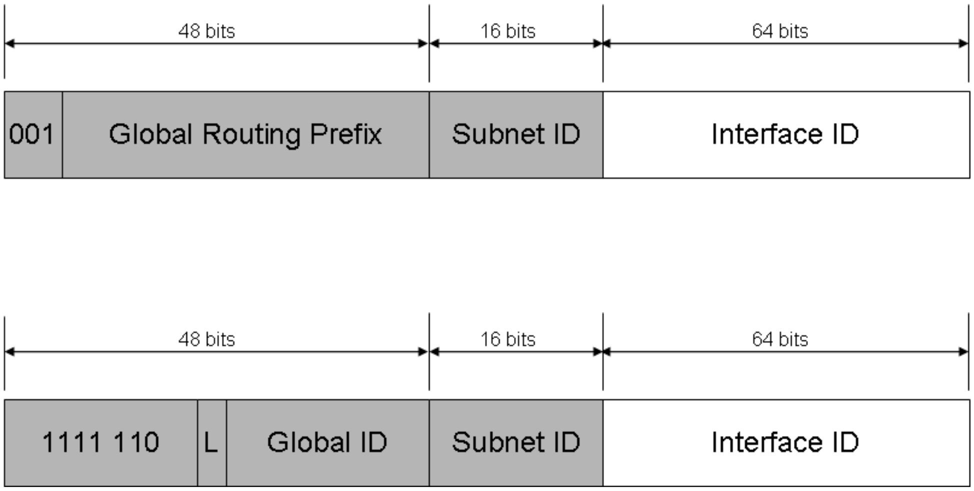

Figure 9 The structure of global and unique local addresses

In global addresses, the Subnet ID field identifies the subnet within an organization. For unique local addresses, the Subnet ID field can perform the same function. Therefore, you can create a subnet numbering scheme that can be used for both local and global unicast addresses. Unique local addresses have a global scope but their reachability is defined by routing topology. Organizations will not advertise their unique local address prefixes outside of their organizations or create DNS AAAA entries with unique local addresses in the Internet DNS.

Special IPv6 Addresses

The following are special IPv6 addresses:

- Unspecified address
  The unspecified address (0:0:0:0:0:0:0:0 or ::) is only used to indicate the absence of an address. It is equivalent to the IPv4 unspecified address of 0.0.0.0. The unspecified address is typically used as a source address for packets attempting to verify the uniqueness of a tentative address. The unspecified address is never assigned to an interface or used as a destination address.
- Loopback address
  The loopback address (0:0:0:0:0:0:0:1 or ::1) is used to identify a loopback interface, enabling a node to send packets to itself. It is equivalent to the IPv4 loopback address of 127.0.0.1. Packets addressed to the loopback address must never be sent on a link or forwarded by an IPv6 router.

Compatibility Addresses

To aid in the migration from IPv4 to IPv6 and the coexistence of both types of hosts, the following addresses are defined:

- IPv4-compatible address
  The IPv4-compatible address, 0:0:0:0:0:0:*w.x.y.z* or ::*w.x.y.z* (where *w.x.y.z* is the dotted decimal representation of an IPv4 address), is used by IPv6/IPv4 nodes that are communicating using IPv6. IPv6/IPv4 nodes are nodes with both IPv4 and IPv6 protocols. When the IPv4-compatible address is used as an IPv6 destination, the IPv6 traffic is

automatically encapsulated with an IPv4 header and sent to the destination using the IPv4 infrastructure.

- IPv4-mapped address
  The IPv4-mapped address, 0:0:0:0:0:FFFF:*w.x.y.z* or ::FFFF:*w.x.y.z*, is used to represent an IPv4-only node to an IPv6 node. It is used only for internal representation. The IPv4-mapped address is never used as a source or destination address of an IPv6 packet.
- 6to4 address
  The 6to4 address is used for communicating between two nodes running both IPv4 and IPv6 over an IPv4 routing infrastructure. The 6to4 address is formed by combining the prefix 2002::/16 with the 32 bits of a public IPv4 address, forming a 48-bit prefix. 6to4 is a tunneling technique described in RFC 3056.

For more information on these address and IPv6 transition technologies, see IPv6 Transition Technologies at http://www.microsoft.com/technet/network/ipv6/ipv6coexist.mspx.

Multicast IPv6 Addresses

In IPv6, multicast traffic operates in the same way that it does in IPv4. Arbitrarily located IPv6 nodes can listen for multicast traffic on an arbitrary IPv6 multicast address. IPv6 nodes can listen to multiple multicast addresses at the same time. Nodes can join or leave a multicast group at any time.

IPv6 multicast addresses have the first eight bits set to 1111 1111. An IPv6 address is easy to classify as multicast because it always begins with “FF”. Multicast addresses cannot be used as source addresses or as intermediate destinations in a Routing extension header.

Beyond the first eight bits, multicast addresses include additional structure to identify their flags, scope, and multicast group. Figure 10 shows the IPv6 multicast address.

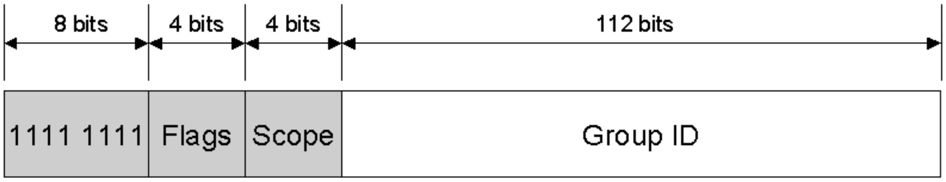

Figure 10 The IPv6 multicast address

The fields in the multicast address are:

- **Flags** – Indicates flags set on the multicast address. The size of this field is 4 bits. The first low-order bit is the Transient (T) flag. When set to 0, the T flag indicates that the multicast address is a permanently assigned (well-known) multicast address allocated by IANA. When set to 1, the T flag indicates that the multicast address is a transient (non-permanently-assigned) multicast address. The second low-order bit is for the Prefix (P) flag, which indicates whether the multicast address is based on a unicast address prefix. RFC 3306 describes the P flag. The third low-order bit is for the Rendezvous Point Address (R) flag, which indicates whether the multicast address contains an embedded rendezvous point address. RFC 3956 describes the R flag.
- **Scope** – Indicates the scope of the IPv6 internetwork for which the multicast traffic is intended. The size of this field is 4 bits. In addition to information provided by multicast routing protocols, routers use the multicast scope to determine whether multicast traffic can be

forwarded. The most prevalent values for the Scope field are 1 (interface-local scope), 2 (link-local scope), and 5 (site-local scope).
For example, traffic with the multicast address of FF02::2 has a link-local scope. An IPv6 router never forwards this traffic beyond the local link.

- **Group ID** – Identifies the multicast group and is unique within the scope. The size of this field is 112 bits. Permanently assigned group IDs are independent of the scope. Transient group IDs are only relevant to a specific scope. Multicast addresses from FF01:: through FF0F:: are reserved, well-known addresses.

To identify all nodes for the interface-local and link-local scopes, the following addresses are defined:

- FF01::1 (interface-local scope all-nodes multicast address)
- FF02::1 (link-local scope all-nodes multicast address)

To identify all routers for the interface-local, link-local, and site-local scopes, the following addresses are defined:

- FF01::2 (interface-local scope all-routers multicast address)
- FF02::2 (link-local scope all-routers multicast address)
- FF05::2 (site-local scope all-routers multicast address)

For the current list of permanently assigned IPv6 multicast addresses, see http://www.iana.org/assignments/ipv6-multicast-addresses.

Solicited-Node Address

The solicited-node address facilitates the efficient querying of network nodes during address resolution. In IPv4, the ARP Request frame is sent to the MAC-level broadcast, disturbing all nodes on the network segment, including those that are not running IPv4. IPv6 uses the Neighbor Solicitation message to perform address resolution. However, instead of using the local-link scope all-nodes multicast address as the Neighbor Solicitation message destination, which would disturb all IPv6 nodes on the local link, the solicited-node multicast address is used. The solicited-node multicast address is comprised of the prefix FF02::1:FF00:0/104 and the last 24-bits of the IPv6 address that is being resolved, as shown in Figure 11.

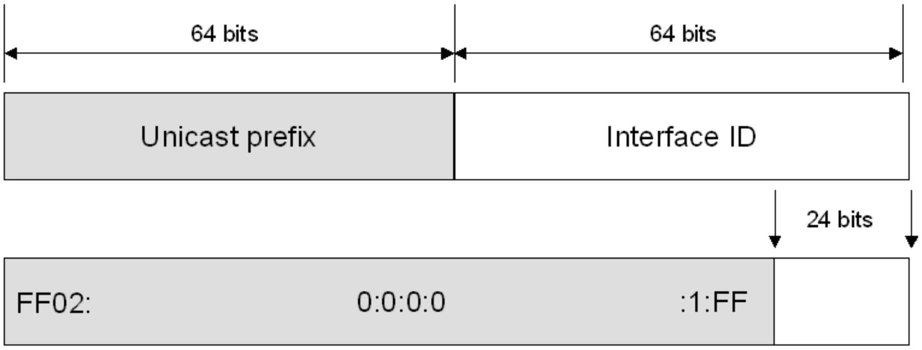

Figure 11 The solicited-node multicast address

For example, Node A is assigned the link-local address of FE80::2AA:FF:FE28:9C5A and is also listening on the corresponding solicited-node multicast address of FF02::1:FF28:9C5A (the underline highlights the correspondence of the last six hexadecimal digits). Node B on the local link must resolve Node A's link-local address FE80::2AA:FF:FE28:9C5A to its corresponding link-layer address. Node B sends a Neighbor Solicitation message to the solicited node multicast address of FF02::1:FF28:9C5A. Because Node A is listening on this multicast address, it processes the Neighbor Solicitation message and sends a unicast Neighbor Advertisement message in reply.

The result of using the solicited-node multicast address is that address resolutions, a common occurrence on a link, are not required to use a mechanism that disturbs all network nodes. By using the solicited-node address, very few nodes are disturbed during address resolution. In practice, due to the relationship between the Ethernet MAC address, the IPv6 interface ID, and the solicited-node address, the solicited-node address acts as a pseudo-unicast address for very efficient address resolution.

Anycast IPv6 Addresses

An anycast address is assigned to multiple interfaces. Packets addressed to an anycast address are forwarded by the routing infrastructure to the nearest interface to which the anycast address is assigned. In order to facilitate delivery, the routing infrastructure must be aware of the interfaces assigned anycast addresses and their “distance” in terms of routing metrics. At present, anycast addresses are only used as destination addresses and are only assigned to routers. Anycast addresses are assigned out of the unicast address space and the scope of an anycast address is the scope of the type of unicast address from which the anycast address is assigned.

The Subnet-Router anycast address is predefined and required. It is created from the subnet prefix for a given interface. To construct the Subnet-Router anycast address, the bits in the subnet prefix are fixed at their appropriate values and the remaining bits are set to 0. All router interfaces attached to a subnet are assigned the Subnet-Router anycast address for that subnet. The Subnet-Router anycast address is used for communication with one of multiple routers attached to a remote subnet.

IPv6 Addresses for a Host

An IPv4 host with a single network adapter typically has a single IPv4 address assigned to that adapter. An IPv6 host, however, usually has multiple IPv6 addresses—even with a single interface. An IPv6 host is assigned the following unicast addresses:

- A link-local address for each interface
- Unicast addresses for each interface (which could be a site-local address and one or multiple global unicast addresses)
- The loopback address (::1) for the loopback interface

Typical IPv6 hosts are logically multihomed because they have at least two addresses with which they can receive packets—a link-local address for local link traffic and a routable site-local or global address.

Additionally, each host is listening for traffic on the following multicast addresses:

- The interface-local scope all-nodes multicast address (FF01::1)
- The link-local scope all-nodes multicast address (FF02::1)
- The solicited-node address for each unicast address on each interface
- The multicast addresses of joined groups on each interface

IPv6 Addresses for a Router

An IPv6 router is assigned the following unicast addresses:

- A link-local address for each interface
- Unicast addresses for each interface (which could be a site-local address and one or multiple global unicast addresses)
- A Subnet-Router anycast address
- Additional anycast addresses (optional)
- The loopback address (::1) for the loopback interface

Additionally, each router is listening for traffic on the following multicast addresses:

- The interface-local scope all-nodes multicast address (FF01::1)
- The interface-local scope all-routers multicast address (FF01::2)
- The link-local scope all-nodes multicast address (FF02::1)
- The link-local scope all-routers multicast address (FF02::2)
- The site-local scope all-routers multicast address (FF05::2)
- The solicited-node address for each unicast address on each interface
- The multicast addresses of joined groups on each interface

IPv6 Interface Identifiers

The last 64 bits of an IPv6 address are the interface identifier that is unique to the 64-bit prefix of the IPv6 address. The following are the ways in which an IPv6 interface identifier is determined:

- A 64-bit interface identifier that is derived from the Extended Unique Identifier (EUI)-64 address. The 64-bit EUI-64 address is defined by the Institute of Electrical and Electronic Engineers (IEEE). EUI-64 addresses are either assigned to a network adapter or derived from IEEE 802 addresses. This is the default behavior for IPv6 in Windows XP and Windows Server 2003.
- As defined in RFC 3041, it might have a temporarily assigned, randomly generated interface identifier to provide a level of anonymity when acting as a client.
- As defined in RFC 2472, an interface identifier can be based on link-layer addresses or serial numbers, or randomly generated when configuring a Point-to-Point Protocol (PPP) interface and an EUI-64 address is not available.
- It is assigned during manual address configuration.
- It is a permanent interface identifier that is randomly generated to mitigate address scans of unicast IPv6 addresses on a subnet. This is the default behavior for IPv6 in Windows Vista and Windows Server “Longhorn.” You can disable this behavior with the **netsh interface ipv6 set global randomizeidentifiers=disabled** command.

EUI-64 address-based interface identifiers

RFC 4291 states that all unicast addresses that use the prefixes 001 through 111 must also use a 64-bit interface identifier that is derived from the EUI-64 address. The 64-bit EUI-64 address is defined by the Institute of Electrical and Electronic Engineers (IEEE). EUI-64 addresses are either assigned to a network adapter card or derived from IEEE 802 addresses.

IEEE 802 Addresses

Traditional interface identifiers for network adapters use a 48-bit address called an IEEE 802 address. It consists of a 24-bit company ID (also called the manufacturer ID), and a 24-bit extension ID (also called the board ID). The combination of the company ID, which is uniquely assigned to each manufacturer of network adapters, and the board ID, which is uniquely assigned to each network adapter at the time of assembly, produces a globally unique 48-bit address. This 48-bit address is also called the physical, hardware, or media access control (MAC) address.

Figure 12 shows the structure of the 48-bit IEEE 802 address.

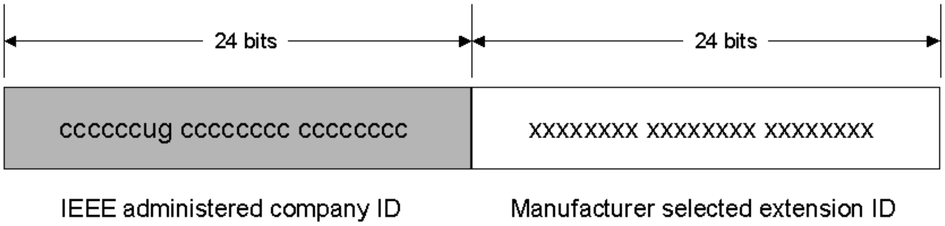
24 bits
24 bits
ccccccug cccccccc cccccccc
xxxxxxxx xxxxxxxx xxxxxxxx
IEEE administered company ID
Manufacturer selected extension ID

Figure 12 The 48-bit IEEE 802 address

Defined bits within the IEEE 802 address are:

Universal/Local (U/L) – The next-to-the low order bit in the first byte is used to indicate whether the address is universally or locally administered. If the U/L bit is set to 0, the IEEE (through the designation of a unique company ID) has administered the address. If the U/L bit is set to 1, the address is locally administered. The network administrator has overridden the manufactured address and specified a different address. The U/L bit is designated by the **u** in Figure 12.

Individual/Group (I/G) – The low order bit of the first byte is used to indicate whether the address is an individual address (unicast) or a group address (multicast). When set to 0, the address is a unicast address. When set to 1, the address is a multicast address. The I/G bit is designated by the **g** in Figure 12.

For a typical 802 network adapter address, both the U/L and I/G bits are set to 0, corresponding to a universally administered, unicast MAC address.

IEEE EUI-64 Addresses

The IEEE EUI-64 address represents a new standard for network interface addressing and is used for Gigabit Ethernet adapters. The company ID is still 24-bits long, but the extension ID is 40 bits, creating a much larger address space for a network adapter manufacturer. The EUI-64 address uses the U/L and I/G bits in the same way as the IEEE 802 address.

Figure 13 shows the structure of the EIU-64 address.

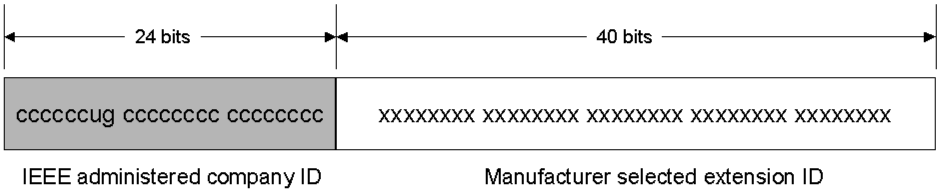
24 bits
40 bits
ccccccug cccccccc cccccccc
xxxxxxxx xxxxxxxx xxxxxxxx xxxxxxxx xxxxxxxx
IEEE administered company ID
Manufacturer selected extension ID

Figure 13 The EUI-64 address

Mapping IEEE 802 Addresses to EIU-64 Addresses

To create an EUI-64 address from an IEEE 802 address, the 16 bits of 11111111 11111110 (0xFFFE) are inserted into the IEEE 802 address between the company ID and the extension ID, as shown in Figure 14.

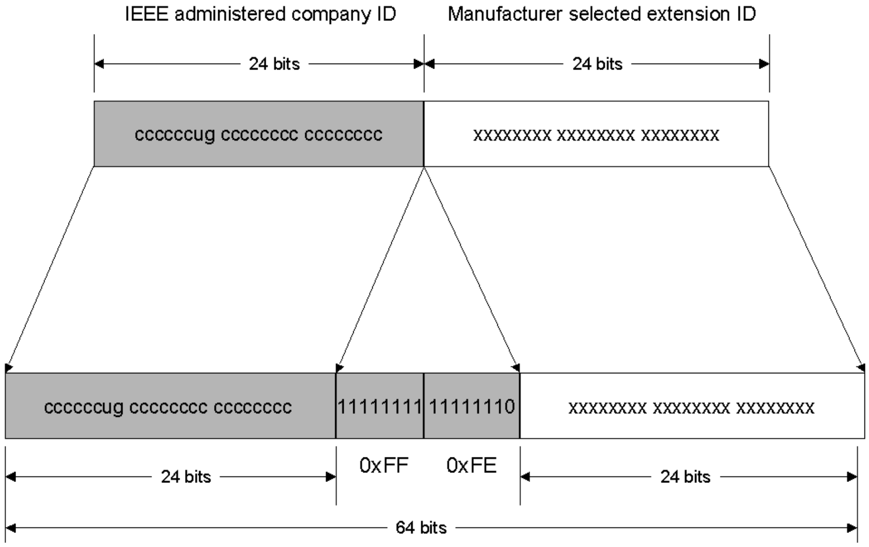
IEEE administered company ID
Manufacturer selected extension ID
24 bits
24 bits
ccccccug cccccccc cccccccc
xxxxxxxx xxxxxxxx xxxxxxxx
ccccccug cccccccc cccccccc
11111111
11111110
xxxxxxxx xxxxxxxx xxxxxxxx
24 bits
0xFF
0xFE
24 bits
64 bits

Figure 14 The conversion of an IEEE 802 address to an EUI-64 address

Mapping EUI-64 Addresses to IPv6 Interface Identifiers

To obtain the 64-bit interface identifier for IPv6 unicast addresses, the U/L bit in the EUI-64 address is complemented (if it is a 1, it is set to 0; and if it is a 0, it is set to 1). Figure 15 shows the conversion for a universally administered, unicast EUI-64 address.

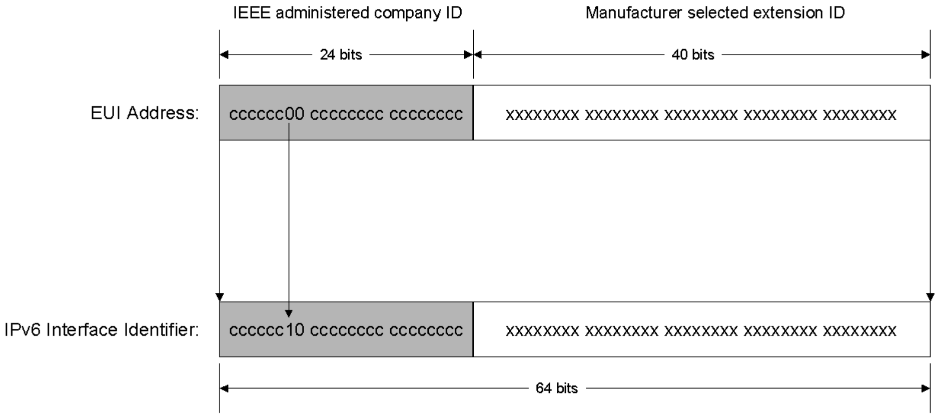
IEEE administered company ID
Manufacturer selected extension ID
24 bits
40 bits
EUI Address:
cccccc00 cccccccc cccccccc
xxxxxxxx xxxxxxxx xxxxxxxx xxxxxxxx xxxxxxxx
IPv6 Interface Identifier:
cccccc10 cccccccc cccccccc
xxxxxxxx xxxxxxxx xxxxxxxx xxxxxxxx xxxxxxxx
64 bits

Figure 15 The conversion of a universally administered, unicast EUI-64 address to an IPv6 interface identifier

To obtain an IPv6 interface identifier from an IEEE 802 address, you must first map the IEEE 802 address to an EUI-64 address, and then complement the U/L bit. Figure 16 shows this conversion process for a universally administered, unicast IEEE 802 address.

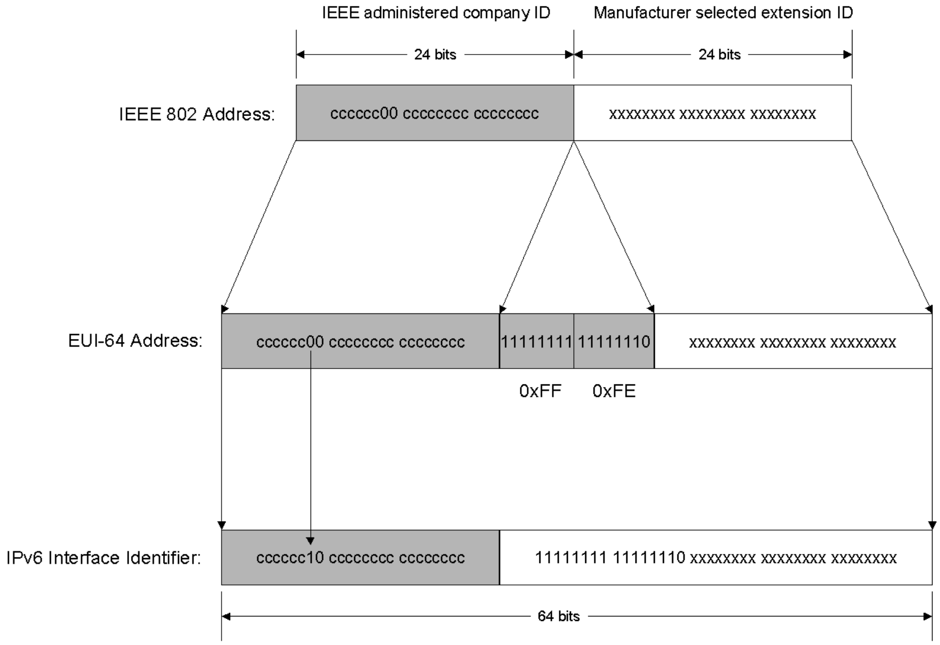
IEEE administered company ID
Manufacturer selected extension ID
24 bits
24 bits
IEEE 802 Address:
cccccc00 cccccccc cccccccc
xxxxxxxx xxxxxxxx xxxxxxxx
EUI-64 Address:
cccccc00 cccccccc cccccccc
11111111
11111110
xxxxxxxx xxxxxxxx xxxxxxxx
0xFF
0xFE
IPv6 Interface Identifier:
cccccc10 cccccccc cccccccc
11111111 11111110 xxxxxxxx xxxxxxxx xxxxxxxx
64 bits

Figure 16 The conversion of a universally administered, unicast IEEE 802 address to an IPv6 interface identifier

IEEE 802 Address Conversion Example

Host A has the Ethernet MAC address of 00-AA-00-3F-2A-1C. First, it is converted to EUI-64 format by inserting FF-FE between the third and fourth bytes, yielding 00-AA-00-FF-FE-3F-2A-1C. Then, the U/L bit, which is the seventh bit in the first byte, is complemented. The first byte in binary form is 00000000. When the seventh bit is complemented, it becomes 00000010 (0x02). The final result is 02-AA-00-FF-FE-3F-2A-1C which, when converted to colon hexadecimal notation, becomes the interface identifier 2AA:FF:FE3F:2A1C. As a result, the link-local address that corresponds to the network adapter with the MAC address of 00-AA-00-2A-1C is FE80::2AA:FF:FE3F:2A1C.

Note When complementing the U/L bit, add 0x2 to the first byte if the address is universally administered, and subtract 0x2 from the first byte if the address is locally administered.

Temporary Address Interface Identifiers

In today's IPv4-based Internet, a typical Internet user connects to an Internet service provider (ISP) and obtains an IPv4 address using the Point-to-Point Protocol (PPP) and the Internet Protocol Control Protocol (IPCP). Each time the user connects, a different IPv4 address might be obtained. Because of this, it is difficult to track a dial-up user's traffic on the Internet on the basis of IP address.

For IPv6-based dial-up connections, the user is assigned a 64-bit prefix after the connection is made through router discovery and stateless address autoconfiguration. If the interface identifier is always based on the EUI-64 address (as derived from the static IEEE 802 address), it is possible to identify the traffic of a specific node regardless of the prefix, making it easy to track a specific user and their use of the Internet. To address this concern and provide a level of anonymity, an alternative IPv6 interface identifier that is randomly generated and changes over time is described in RFC 3041.

The initial interface identifier is generated by using random numbers. For IPv6 systems that cannot store any historical information for generating future interface identifier values, a new random interface identifier is generated each time the IPv6 protocol is initialized. For IPv6 systems that have storage capabilities, a history value is stored and, when the IPv6 protocol is initialized, a new interface identifier is created through the following process:

1. Retrieve the history value from storage and append the interface identifier based on the EUI-64 address of the adapter.
2. Compute the Message Digest-5 (MD5) one-way encryption hash over the quantity in step 1.
3. Save the last 64 bits of the MD5 hash computed in step 2 as the history value for the next interface identifier computation.
4. Take the first 64 bits of the MD5 hash computed in Step 2 and set the seventh bit to zero. The seventh bit corresponds to the U/L bit which, when set to 0, indicates a locally administered interface identifier. The result is the interface identifier.

The resulting IPv6 address, based on this random interface identifier, is known as a temporary address. Temporary addresses are generated for public address prefixes that use stateless address autoconfiguration. Temporary addresses are used for the lower of the following values of the valid and preferred lifetimes:

- The lifetimes included in the Prefix Information option in the received Router Advertisement message.
- Local default values of 1 week for valid lifetime and 1 day for preferred lifetime.

After the temporary address valid lifetime expires, a new interface identifier and temporary address is generated.

Mapping IPv6 Multicast Addresses to Ethernet Addresses

When sending IPv6 multicast packets on an Ethernet link, the corresponding destination MAC address is 33-33-mm-mm-mm-mm where mm-mm-mm-mm is a direct mapping of the last 32 bits of the IPv6 multicast address, as shown in Figure 17.

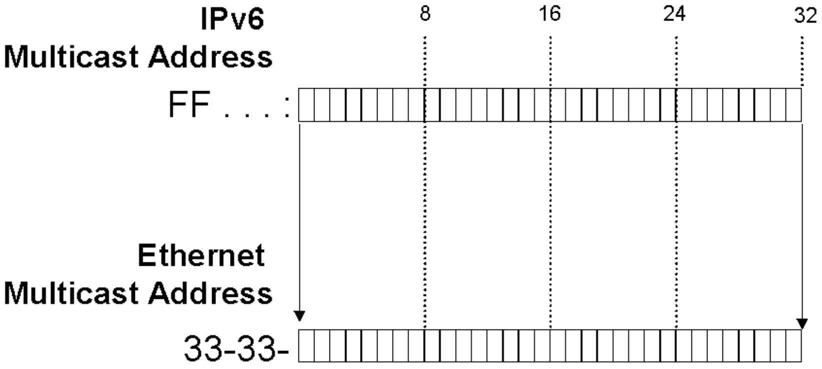

Figure 17 The mapping of an IPv6 multicast address to an Ethernet multicast MAC address

To efficiently receive IPv6 multicast packets on an Ethernet link, Ethernet network adapters can store additional interesting MAC addresses in a table on the network adapter. If an Ethernet frame with an interesting MAC address is received, it is passed to upper layers for additional processing. For every multicast address being listened to by the host, there is a corresponding entry in the table of interesting MAC address.

For example, a host with the Ethernet MAC address of 00-AA-00-3F-2A-1C (link-local address of FE80::2AA:FF:FE3F:2A1C) registers the following multicast MAC addresses with the Ethernet adapter:

- The address of 33-33-00-00-00-01, which corresponds to the link-local scope all-nodes multicast address of FF02::1.
- The address of 33-33-FF-3F-2A-1C, which corresponds to the solicited-node address of FF02::1:FF3F:2A1C. Remember that the solicited-node address is the prefix FF02::1:FF00:0/104 and the last 24-bits of the unicast IPv6 address.

Additional multicast addresses on which the host is listening are added and removed as needed from the table of interesting address on the Ethernet network adapter.

IPv6 and DNS

Enhancements to the Domain Name System (DNS) for IPv6 are described in RFC 1886 and consist of the following new elements:

- Host address (AAAA) resource record
- IP6.ARPA domain for reverse queries

Note According to RFC 3152, Internet Engineering Task Force (IETF) consensus has been reached that the IP6.ARPA domain be used, instead of IP6.INT as defined in RFC 1886. The IP6.ARPA domain is the domain used by IPv6 in Microsoft Windows.

The Host Address (AAAA) Resource Record

A new DNS resource record type, AAAA (called “quad A”), is used for resolving a fully qualified domain name to an IPv6 address. It is comparable to the host address (A) resource record used with IPv4. The resource record type is named AAAA (Type value of 28) because 128-bit IPv6 addresses are four times as large as 32-bit IPv4 addresses. The following is an example of a AAAA resource record:

host1.microsoft.com   IN   AAAA   2001:DB8:2F31:1A2D::2AA:FF:FE3F:2A1C

A host must specify either a AAAA query or a general query for a specific host name in order to receive IPv6 address resolution data in the DNS query answer sections.

The IP6.ARPA Domain

The IP6.ARPA domain has been created for IPv6 reverse queries. Also called pointer queries, reverse queries determine a host name based on the IP address. To create the namespace for reverse queries, each hexadecimal digit in the fully expressed 32-digit IPv6 address becomes a separate level in inverse order in the reverse domain hierarchy.

For example, the reverse lookup domain name for the address FEC0::2AA:FF:FE3F:2A1C (fully expressed as FEC0:0000:0000:0000:02AA:00FF:FE3F:2A1C) is:

C.1.A.2.F.3.E.F.F.F.0.0.A.A.2.0.0.0.0.0.0.0.0.0.0.0.0.0.C.E.F.IP6.ARPA.

The DNS support described in RFC 1886 represents a simple way to both map host names to IPv6 addresses and provide reverse name resolution.

Source and Destination Address Selection

For a typical IPv4-only host that has a single interface assigned one IPv4 address and resolves names using DNS, the choice of which IPv4 addresses to use as the source and destination when initiating communication is straightforward. The source IPv4 address is the address assigned to the interface of the host. The destination addresses to which connections are attempted are the IPv4 addresses returned in the DNS Name Query Response message.

For a typical IPv6 host that has multiple IPv6 addresses assigned to multiple interfaces and multiple IPv6 addresses are returned in the DNS Name Query Response message, the choice of the source and destination IPv6 address is more complex. The source and destination IPv6 addresses should be matched in scope and purpose. For example, an IPv6 host should not choose a link-local source address when communicating with a global destination address. Additionally, the possible destination address should be sorted by preference.

To provide a standardized method to choose source and destination IPv6 addresses with which to attempt connections, RFC 3484 defines the following required algorithms:

- A source address selection algorithm to choose the best source address to use with a destination address.
- A destination address selection algorithm to sort the list of possible destination addresses in order of preference.

For more information about the source and destination address selection algorithms defined in RFC 3484, see Source and Destination Address Selection for IPv6 at http://www.microsoft.com/technet/community/columns/cableguy/cg0206.mspx.

IPv4 Addresses and IPv6 Equivalents

Table 2 lists both IPv4 addresses and addressing concepts and their IPv6 equivalents.

Table 2 IPv4 Addressing Concepts and Their IPv6 Equivalents

IPv4 Address	IPv6 Address
Internet address classes	Not applicable in IPv6
Multicast addresses (224.0.0.0/4)	IPv6 multicast addresses (FF00::/8)
Broadcast addresses	Not applicable in IPv6
Unspecified address is 0.0.0.0	Unspecified address is ::
Loopback address is 127.0.0.1	Loopback address is ::1
Public IP addresses	Global unicast addresses
Private IP addresses (10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16)	Site-local addresses (FEC0::/10)
Autoconfigured addresses (169.254.0.0/16)	Link-local addresses (FE80::/64)
Text representation: Dotted decimal notation	Text representation: Colon hexadecimal format with suppression of leading zeros and zero compression. IPv4-compatible addresses are expressed in dotted decimal notation.
Network bits representation: Subnet mask in dotted decimal notation or prefix length	Network bits representation: Prefix length notation only
DNS name resolution: IPv4 host address (A) resource record	DNS name resolution: IPv6 host address (AAAA) resource record
DNS reverse resolution: IN-ADDR.ARPA domain	DNS reverse resolution: IP6.ARPA domain