Planning for Information Network

Lecture 8 and 9:

Introduction to IPv6

IPv6 Features

The ability to scale networks for future demands requires a limitless supply of IP addresses; IPv6 combines expanded addressing with a more efficient and feature-rich header to meet these demands. IPv6 satisfies the increasingly complex requirements of hierarchical addressing that IPv4 does not support.

IPv6 Features

**The main benefits of IPv6 include the following:

■ **Larger address space:** IPv6 addresses are 128 bits, compared to IPv4’s 32 bits. This larger addressing space allows more support for addressing hierarchy levels, a much greater number of addressable nodes, and simpler auto configuration of addresses.

■ **Globally unique IP addresses:** Every node can have a unique global IPv6 address, which eliminates the need for NAT.

■ **Header format efficiency:** A simplified header with a fixed header size makes processing more efficient.

IPv6 Features

- **Improved privacy and security:** IPsec is the IETF standard for IP network security, available for both IPv4 and IPv6. Although the functions are essentially identical in both environments, IPsec is mandatory in IPv6. IPv6 also has optional security headers.

- **Flow labeling capability:** A new capability enables the labeling of packets belonging to particular traffic flows for which the sender requests special handling, such as nondefault quality of service (QoS) or real-time service.

- **Increased mobility and multicast capabilities:** Mobile IPv6 allows an IPv6 node to change its location on an IPv6 network and still maintain its existing connections. With Mobile IPv6, the mobile node is always reachable through one permanent address. A connection is established with a specific permanent address assigned to the mobile node, and the node remains connected no matter how many times it changes locations and addresses.

IPv6 Address Format

Rather than using dotted-decimal format, IPv6 addresses are written as hexadecimal numbers with colons between each set of four hexadecimal digits (which is 16 bits); we like to call this the “coloned hex” format. The format is x:x:x:x:x:x:x:x, where x is a 16-bit hexadecimal field. A sample address is as follows:

2035:0001:2BC5:0000:0000:087C:0000:000A

IPv6 Address Format

Note:

We can shorten the written form of IPv6 addresses. Leading 0s within each set of four hexadecimal digits can be omitted, and a pair of colons (::) can be used, once within an address, to represent any number of successive 0s.

For example, the previous address can be shortened to the following:

2035:1:2BC5::87C:0:A

An all-0s address can be written as :: .

IPv6 Address Format

Note:

A pair of colons (::) can be used only once within an IPv6 address. This is because an address parser identifies the number of missing 0s by separating the two parts and entering 0 until the 128 bits are complete. If two :: notations were to be placed in the address, there would be no way to identify the size of each block of 0s.

IPv6 Address Format

Example:

3FFE:**0**501:**0008**:**0000**:**0**260:97FF:FE40:EFAB

= 3FFE**:5**01**:8:0:2**60:97FF:FE40:EFAB

= 3FFE:501:8**::**260:97FF:FE40:EFAB

IPv6 Addressing in an Enterprise Network

An IPv6 address consists of two parts:

- **A subnet prefix** representing the network to which the interface is connected. Usually 64-bits in length.

- **An interface ID**, sometimes called a local identifier or a token. Usually 64-bits in length.

Subnet Prefix

IPv6 uses the "/prefix-length" to denote how many bits in the IPv6 address represent the subnet.

The syntax is **ipv6-address/prefix-length**

- ipv6-address is the 128-bit IPv6 address.
- /prefix-length is a decimal value representing how many of the left most contiguous bits of the address comprise the prefix.

For example:

fec0:0:0:1::1234/64

is really

fec0:0000:0000:0001:**0000:0000:0000:1234**/64

- The first 64-bits (**fec0:0000:0000:0001**) forms the address prefix.
- The last 64-bits (**0000:0000:0000:1234**) forms the Interface ID.

IPv4 Packet Header

20 Bytes + Options

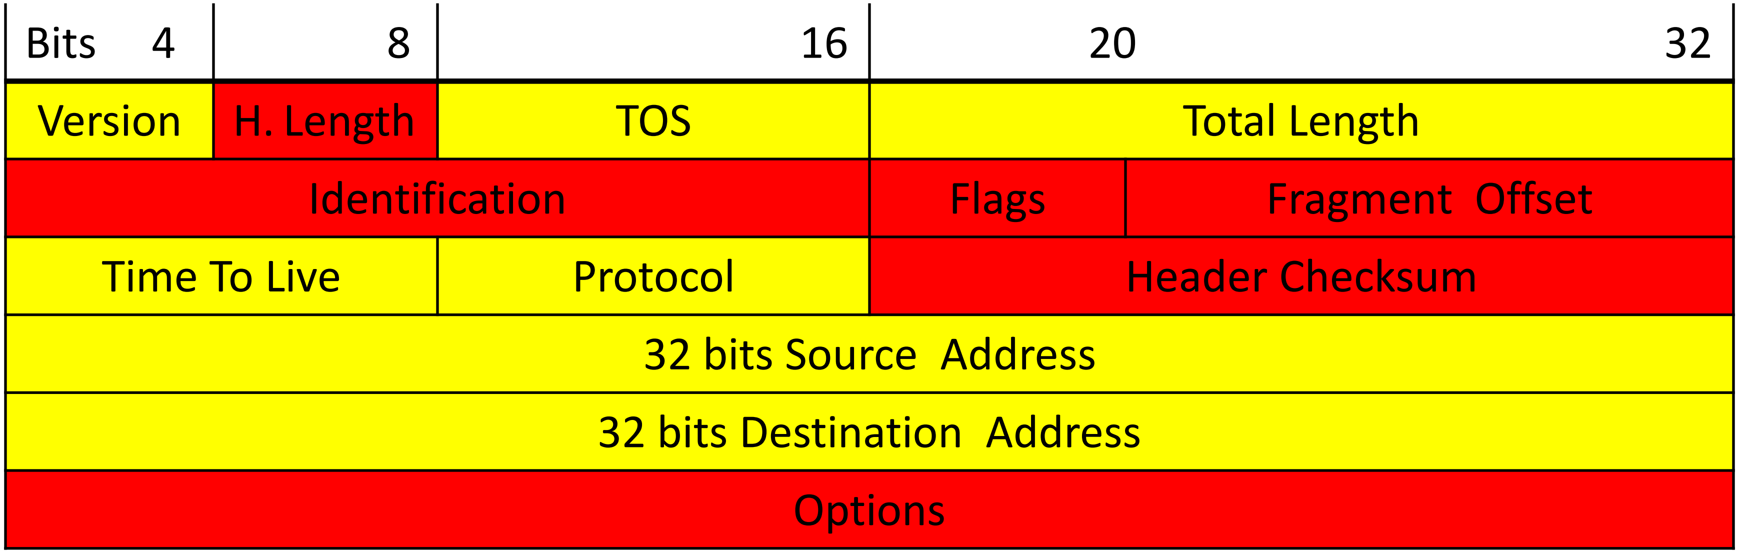

Bits 4	8	16	20	32
🟨 Version	🟥 H. Length	🟨 TOS	🟨 Total Length	
🟥 Identification			🟥 Flags	🟥 Fragment Offset
🟨 Time To Live		🟨 Protocol	🟥 Header Checksum	
🟨 32 bits Source Address				
🟨 32 bits Destination Address				
🟥 Options				

🟨 **Modified Fields**

🟥 **Deleted Fields**

IPv6 Packet Header

From 12 to 8 fields (40 Bytes)

<table>
<tr><th>Bits 4</th><th>12</th><th>16</th><th>24</th><th>32</th></tr>
<tr><td>Version</td><td>Class of Traffic</td><td colspan="3">Flow Label</td></tr>
<tr><td colspan="3">Payload Length</td><td>Next Header</td><td>Hop Limit</td></tr>
<tr><td colspan="5">128 bits Source Address</td></tr>
<tr><td colspan="5">128 bits Destination Address</td></tr>
</table>

IPv6 Packet Header

The IPv6 header has 40 octets, in contrast to the 20 octets in the IPv4 header. IPv6 has fewer fields, and the header is 64-bit-aligned to enable fast, efficient, hardware-based processing. The IPv6 address fields are four times larger than in IPv4.

The IPv4 header contains 12 basic header fields, followed by an options field and a data portion (which usually includes a transport layer segment). The basic IPv4 header has a fixed size of 20 octets; the variable-length options field increases the size of the total IPv4 header. IPv6 contains fields similar to 7 of the 12 IPv4 basic header fields (5 plus the source and destination address fields) but does not require the other fields.

IPv6 Packet Header

The IPv6 header contains the following fields:

■ **Version:** A 4-bit field, the same as in IPv4. For IPv6, this field contains the number 6; for IPv4, this field contains the number 4.

■ **Traffic class:** An 8-bit field similar to the type of service (ToS) field in IPv4. This field tags the packet with a traffic class that it uses in differentiated services (DiffServ) QoS. These functions are the same for IPv6 and IPv4.

■ **Flow label:** This 20-bit field is new in IPv6. It can be used by the source of the packet to tag the packet as being part of a specific flow, allowing multilayer switches and routers to handle traffic on a per-flow basis rather than per-packet, for faster packet-switching performance. This field can also be used to provide QoS.

■ **Payload length:** This 16-bit field is similar to the IPv4 total length field.

IPv6 Packet Header

- **Next header:** The value of this 8-bit field determines the type of information that follows the basic IPv6 header. It can be transport-layer information, such as Transmission Control Protocol (TCP) or User Datagram Protocol (UDP), or it can be an extension header. The next header field is similar to the protocol field of IPv4.

- **Hop limit:** This 8-bit field specifies the maximum number of hops that an IPv6 packet can traverse. Similar to the time to live (TTL) field in IPv4, each router decreases this field by 1. Because there is no checksum in the IPv6 header, an IPv6 router can decrease the field without re-computing the checksum; in IPv4 routers, the re-computation costs processing time. If this field ever reaches 0, a message is sent back to the source of the packet, and the packet is discarded.

- **Source address:** This field has 16 octets (128 bits). It identifies the source of the packet.

- **Destination address:** This field has 16 octets (128 bits). It identifies the destination of the packet.

IPv6 Packet Header

Note:

The IPv6 header does not have a header checksum field. Because link-layer technologies perform checksum and error control and are considered relatively reliable, an IPv6 header checksum is considered redundant. Without the IPv6 header checksum, upper-layer checksums, such as within UDP, are mandatory with IPv6.

Special IPv6 Addresses

IPv6 Address	Description
::/0	• All routes and used when specifying a default static route. • It is equivalent to the IPv4 quad-zero (0.0.0.0).
::/128	• Unspecified address and is initially assigned to a host when it first resolves its local link address.
::1/128	• Loopback address of local host. • Equivalent to 127.0.0.1 in IPv4.
FE80::/10	• Link-local unicast address. • Similar to the Windows autoconfiguration IP address of 169.254.x.x.
FF00::/8	Multicast addresses.
All other addresses	Global unicast address.

IPv6 Address Scope Types

Similar to IPv4, a single source can address datagrams to either one or many destinations at the same time in IPv6.

Following are the types of IPv6 addresses:

■ **Unicast (one-to-one):** Similar to an IPv4 unicast address, an IPv6 unicast address is for a single source to send data to a single destination. A packet sent to a unicast IPv6 address goes to the interface identified by that address. The IPv6 unicast address space encompasses the entire IPv6 address range, with the exception of the FF00::/8 range (addresses starting with binary 1111 1111), which is used for multicast addresses.

IPv6 Address Scope Types

Unicast Topology:

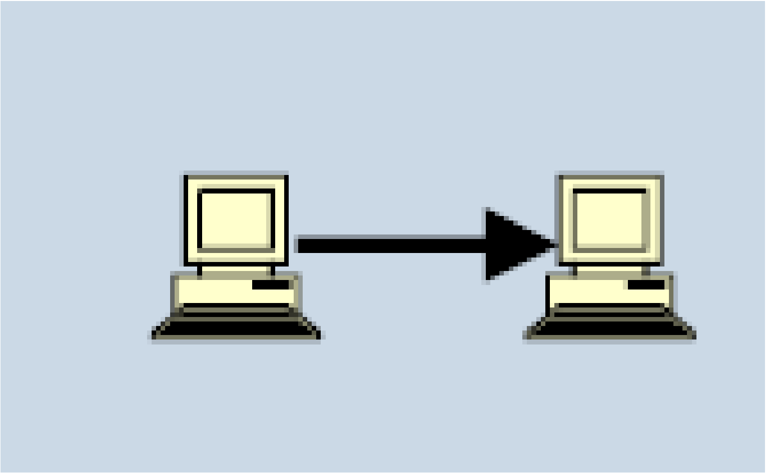

IPv6 Address Scope Types

■ **Anycast (one-to-nearest):** An IPv6 anycast address is a new type of address that is assigned to a set of interfaces on different devices; an anycast address identifies multiple interfaces. A packet that is sent to an anycast address goes to the closest interface (as determined by the routing protocol being used) identified by the anycast address.

Therefore, all nodes with the same anycast address should provide uniform service. Anycast addresses are syntactically indistinguishable from global unicast addresses because anycast addresses are allocated from the global unicast address space. Nodes to which the anycast address is assigned must be explicitly configured to recognize the anycast address.

Anycast addresses must not be used as the source address of an IPv6 packet. Examples of when anycast addresses could be used are load balancing, content delivery services, and service location. For example, an anycast address could be assigned to a set of replicated FTP servers. A user in China who wants to retrieve a file would be directed to the Chinese server, whereas a user in the Europe would be directed to the European server.

IPv6 Address Scope Types

Anycast topology:

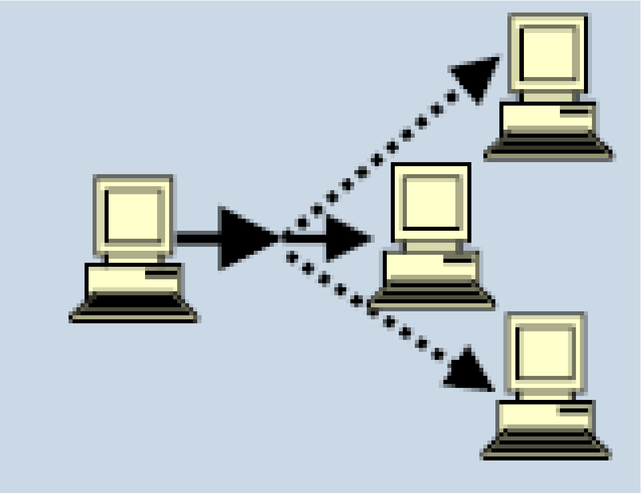

IPv6 Address Scope Types

■ **Multicast (one-to-many):** Similar to IPv4 multicast, an IPv6 multicast address identifies a set of interfaces (in a given scope), typically on different devices.

A packet sent to a multicast address is delivered to all interfaces identified by the multicast address (in a given scope). IPv6 multicast addresses have a 4-bit scope identifier (ID) to specify how far the multicast packet may travel.

Scope:

- 1 (0001) = Node
- 2 (0010) = Link
- 5 (0101) = Site
- 8 (1000) = Organization
- E (1110) = Global

IPv6 Address Scope Types

Multicast topology:

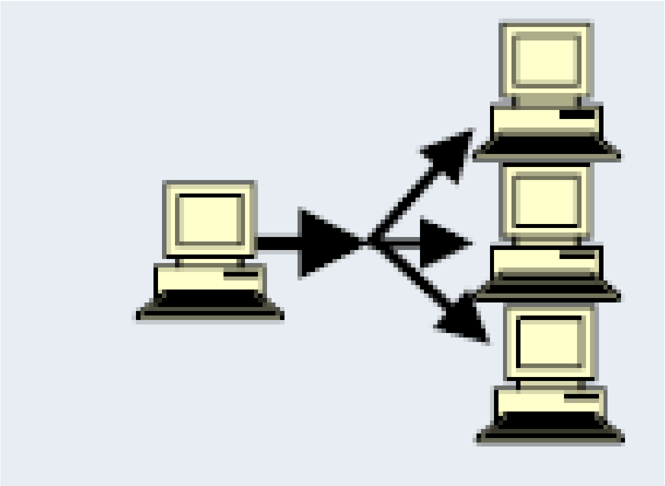

IPv6 Address Scope Types

Note:

- IPv6 has no concept of broadcast addresses; multicast addresses are used instead.

- A single interface may be assigned multiple IPv6 addresses of any type (unicast, anycast, and multicast).

Interface Identifiers in IPv6 Addresses

In IPv6, a link is a network medium over which network nodes communicate using the link layer. Interface IDs in IPv6 addresses are used to identify a unique interface on a link. They can also be thought of as the “host portion” of an IPv6 address. Interface IDs are required to be unique on a link and can also be unique over a broader scope.

When the interface identifier is derived directly from the data link layer address of the interface, the scope of that identifier is assumed to be universal (global). Interface identifiers are always 64 bits and are dynamically created based on the data link layer.

Interface Identifiers in IPv6 Addresses

For Ethernet, the interface ID used is based on the MAC address of the interface and is in an **E**xtended **U**niversal **I**dentifier 64-bit (EUI-64) format.

The EUI-64 format interface ID is derived from the 48-bit link-layer MAC address by inserting the hexadecimal number **FFFE** between the upper 3 bytes (the organizational unique identifier “OUI” field) and the lower 3 bytes (the vendor code or serial number field) of the link-layer address. The seventh bit in the high-order byte is set to 1 (equivalent to the IEEE G/L bit) to indicate the uniqueness of the 48-bit address.

Interface Identifiers in IPv6 Addresses

EUI-64 Format IPv6 Interface Identifier

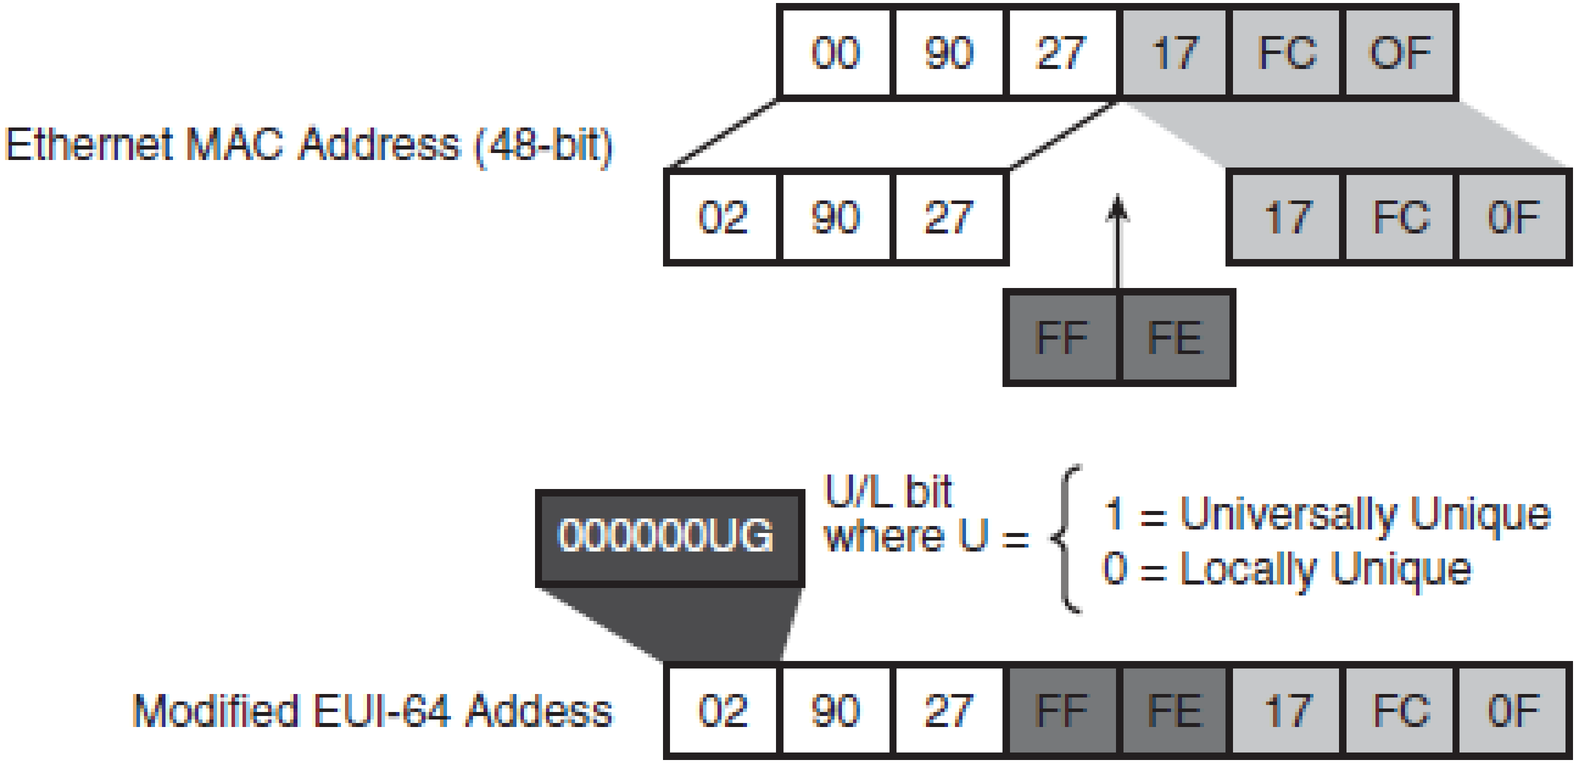

Interface Identifiers in IPv6 Addresses

The seventh bit in an IPv6 interface identifier is referred to as the **U**niversal/**L**ocal (**U/L**) bit. This bit identifies whether this interface identifier is locally unique on the link or whether it is universally unique.

When the interface identifier is created from an Ethernet MAC address, it is assumed that the MAC address is universally unique and, therefore, that the interface identifier is universally unique. The U/L bit is for future use by upper-layer protocols to uniquely identify a connection, even in the context of a change in the leftmost part of the address. However, this feature is not yet used. The eighth bit in an IPv6 interface identifier, also known as the "**G**" bit, is the **g**roup/individual bit for managing groups.

IPv6 Unicast Addresses

Following are the different unicast addresses that IPv6 supports:

- Global aggregatable address (also called global unicast address)
- Link-local address
- IPv4-compatible IPv6 address

Global Aggregatable Unicast Addresses

Note:

IPv6 global aggregatable unicast addresses are equivalent to IPv4 unicast addresses.

The structure of global aggregatable unicast addresses enables summarization (aggregation) of routing prefixes so that the number of routing table entries in the global routing table can be reduced. Global unicast addresses used on links are aggregated upward, through organizations, and then to intermediate-level ISPs, and eventually to top-level ISPs. A global unicast address typically consists of a 48-bit global routing prefix, a 16-bit subnet ID, and a 64-bit interface ID (typically in EUI-64 bit format), as illustrated in the figure bellow.

Global Aggregatable Unicast Addresses

The subnet ID can be used by individual organizations to create their own local addressing hierarchy using subnets. This field allows an organization to use up to 65,536 individual subnets.

A fixed prefix of binary 2000::/3 (binary 001) indicates a global aggregatable IPv6 address; this is the current range of IPv6 global unicast addresses assigned by the Internet Assigned Numbers Authority (IANA).

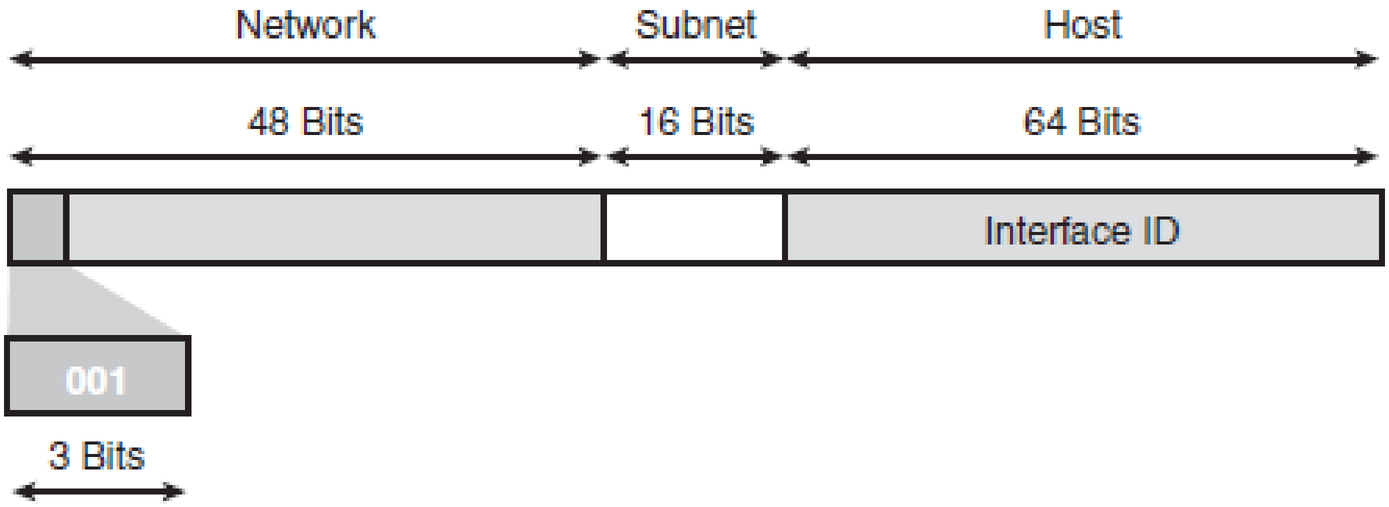

Link-Local Unicast Addresses

A link-local address is useful only in the context of the local link network; its scope limits its relevance to only one link. A link-local address is an IPv6 unicast address that can be automatically configured on any interface by using the link-local prefix FE80::/10 (1111 1110 10) and the 64-bit interface identifier. Link-local addresses are used in the neighbor discovery protocol and the dynamic address assignment process.

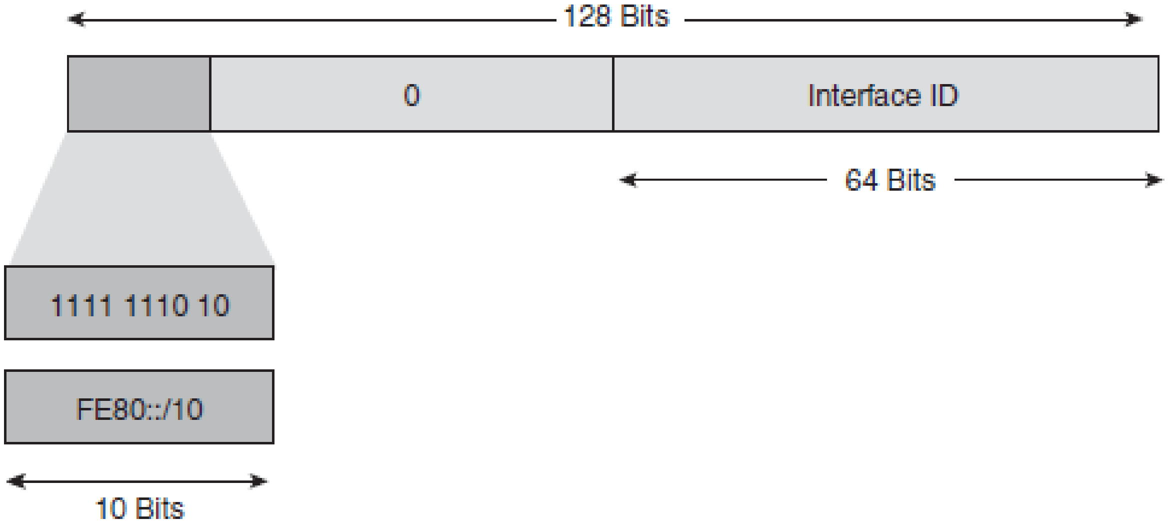

Link-Local Unicast Addresses

A link-local unicast address connects devices on the same local network without requiring globally unique addresses.

When communicating with a link-local address, the outgoing interface must be specified, because every interface is connected to FE80::/10.

An IPv6 router must not forward packets that have either link-local source or destination addresses to other links.

IPv6 Address Assignment Strategies

As with IPv4, IPv6 allows two address assignment strategies:

Static and Dynamic

Static IPv6 Address Assignment

Static address assignment in IPv6 is the same as in IPv4—the administrator must enter the IPv6 address configuration manually on every device in the network.

Dynamic IPv6 Address Assignment

IPv6 dynamic address assignment strategies allow dynamic assignment of IPv6 addresses, as follows:

■ **Link-local address:** The host configures its own link-local address autonomously, using the link-local prefix FE80::0/10 and a 64-bit identifier for the interface, in an EUI-64 format.

■ **Stateless autoconfiguration:** A router on the link advertises—either periodically or at the host's request—network information, such as the 64-bit prefix of the local network and its willingness to function as a default router for the link. Hosts can automatically generate their global IPv6 addresses by using the prefix in these router messages; the hosts do not need manual configuration or the help of a device such as a DHCP server. For example, the following figure shows a host using the prefix advertised by the router as the top 64 bits of its address; the remaining 64 bits contain the host's 48-bit MAC address in an EUI-64 format.

Dynamic IPv6 Address Assignment

IPv6 Stateless Autoconfiguration Allows a Host to Automatically Configure Its IPv6 Address:

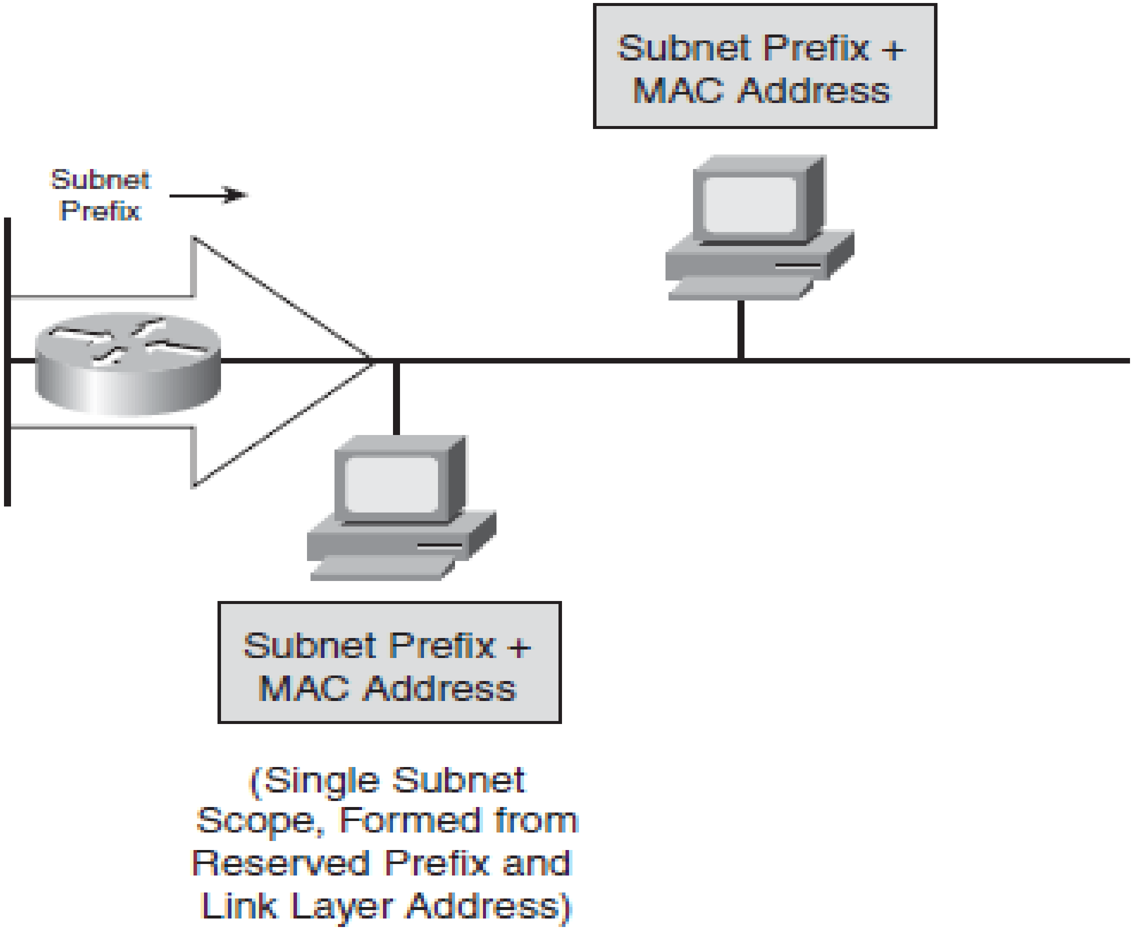

(Single Subnet Scope, Formed from Reserved Prefix and Link Layer Address)

Dynamic IPv6 Address Assignment

- **Stateful using DHCP for IPv6 (DHCPv6):** DHCPv6 is an updated version of DHCP for IPv4. DHCPv6 gives the network administrator more control than stateless autoconfiguration and can be used to distribute other information, including the address of the DNS server.

DHCPv6 can also be used for automatic domain name registration of hosts using a dynamic DNS server. **DHCPv6 uses multicast addresses.**

IPv4-to-IPv6 Transition Strategies and Deployments

Differences Between IPv4 and IPv6:

Regardless of which protocol is used, the communication between IPv4 and IPv6 domains must be transparent to end users. The major differences to consider between IPv4 and IPv6 include the following:

■ IPv4 addresses are 32 bits long, whereas IPv6 addresses are 128 bits long.

■ An IPv6 packet header is different from an IPv4 packet header. The IPv6 header is longer and simpler (new fields were added to the IPv6 header, and some old fields were removed).

■ IPv6 has no concept of broadcast addresses; instead, it uses multicast addresses.

■ Routing protocols must be changed to support native IPv6 routing.

Dual-Stack Transition Mechanism

a dual-stack node enables both IPv4 and IPv6 stacks. Applications communicate with both IPv4 and IPv6 stacks; the IP version choice is based on name lookup and application preference. This is the most appropriate method for campus and access networks during the transition period, and it is the preferred technique for transitioning to IPv6. A dual-stack approach supports the maximum number of applications. Operating systems that support the IPv6 stack include FreeBSD, Linux, Sun Solaris, and Windows 2000, XP, and Vista.

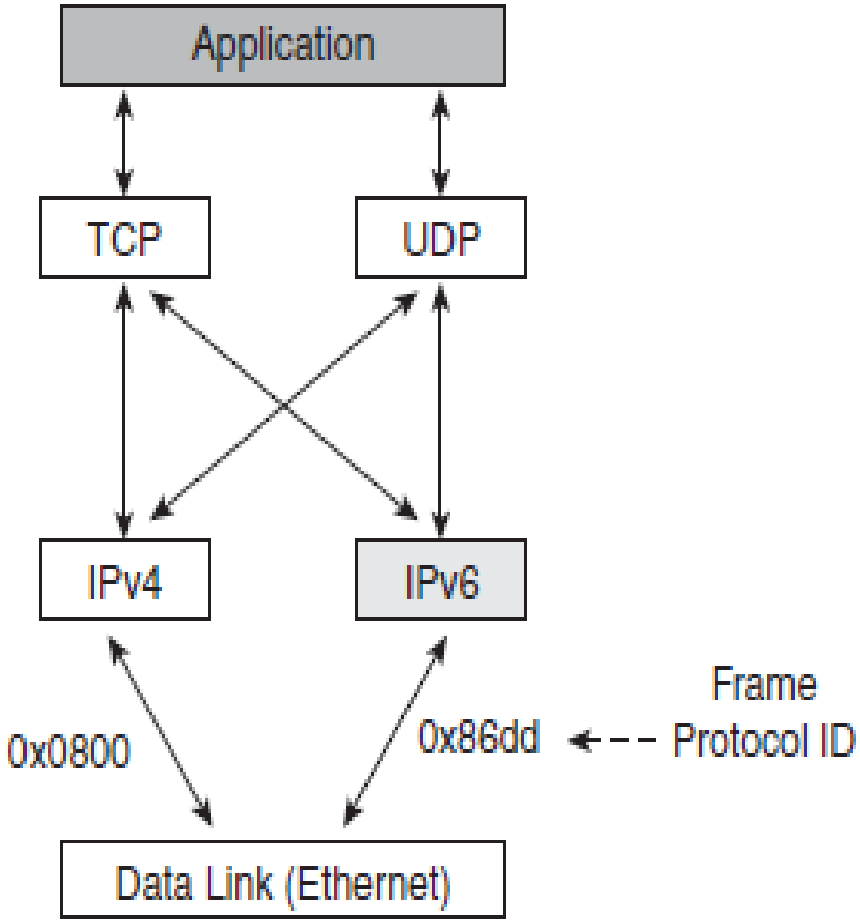

Tunneling Transition Mechanism

The purpose of tunneling is to encapsulate packets of one type in packets of another type. When transitioning to IPv6, tunneling encapsulates IPv6 packets in IPv4 packets, as shown in the following figure.

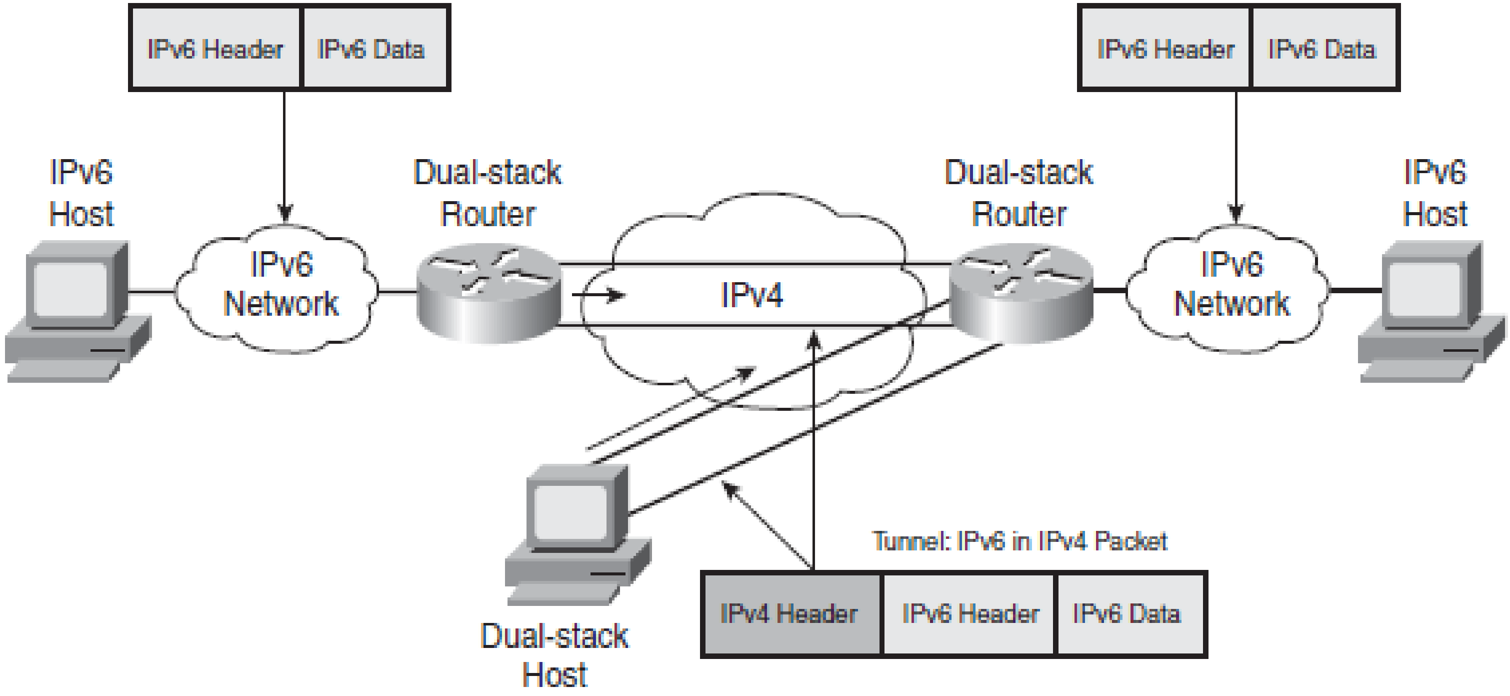

Tunneling Transition Mechanism

By using overlay tunnels, isolated IPv6 networks can communicate without having to upgrade the IPv4 infrastructure between them. Both routers and hosts can use tunneling. The following different techniques are available for establishing a tunnel:

■ **Manually configured:** For a manually configured tunnel, the tunnel source and tunnel destination are manually configured with static IPv4 and IPv6 addresses. Manual tunnels can be configured between border routers or between a border router and a host.

■ **Semi-automated:** Semi-automation is achieved by using a tunnel broker that uses a web based service to create a tunnel. A tunnel broker is a server on the IPv4 network that receives tunnel requests from dual-stack clients, configures the tunnel on the tunnel server or router, and associates the tunnel from the client to one of the tunnel servers or routers. A simpler model combines the tunnel broker and server onto one device.

Tunneling Transition Mechanism

■ **Automatic:** Various automatic mechanisms accomplish tunneling, including the following:

— IPv4-compatible: The tunnel is constructed dynamically using an IPv4-compatible IPv6 address (an IPv6 address that consists of 0s in the upper bits and an embedded IPv4 address in the lower 32 bits). Because it does not scale, this mechanism is appropriate only for testing.

NOTE The format of an IPv4-compatible IPv6 address is 0:0:0:0:0:0:A.B.C.D, or ::A.B.C.D, where A.B.C.D is the IPv4 address in dotted-decimal notation. The entire 128-bit IPv4- compatible IPv6 address is used as a node's IPv6 address, and the IPv4 address that is embedded in the low-order 32 bits is used as the node's IPv4 address. For example, the IPv4 address 192.168.30.1 would convert to the IPv4-compatible IPv6 address 0:0:0:0:0:0:192.168.30.1. Other acceptable representations for this address are ::192.168.30.1 and ::C0A8:1E01.

Tunneling Transition Mechanism

— IPv6-to-IPv4 (6-to-4): The 6-to-4 tunneling method automatically connects IPv6 islands through an IPv4 network. Each 6-to-4 edge router has an IPv6 address with a /48 prefix that is the concatenation of 2002::/16 and the IPv4 address of the edge router; 2002::/16 is a specially assigned address range for the purpose of 6-to-4. The edge routers automatically build the tunnel using the IPv4 addresses embedded in the IPv6 addresses. For example, if the IPv4 address of an edge router is 192.168.99.1, the prefix of its IPv6 address is 2002:C0A8:6301::/48 because 0xC0A86301 is the hexadecimal representation of 192.168.99.1.

When an edge router receives an IPv6 packet with a destination address in the range of 2002::/16, it determines from its routing table that the packet must traverse the tunnel. The router extracts the IPv4 address embedded in the third to sixth octets, inclusive, in the IPv6 next-hop address. This IPv4 address is the IPv4 address of the 6-to-4 router at the destination site—the router at the other end of the tunnel. The router encapsulates the IPv6 packet in an IPv4 packet with the destination edge router’s extracted IPv4 address. The packet passes through the IPv4 network. The destination edge router unencapsulates the IPv6 packet from the received IPv4 packet and forwards the IPv6 packet to its final destination. A 6-to-4 relay router, which offers traffic forwarding to the IPv6 Internet, is required for reaching a native IPv6 Internet.

Tunneling Transition Mechanism

— 6over4: A router connected to a native IPv6 network and with a 6over4-enabled interface can be used to forward IPv6 traffic between 6over4 hosts and native IPv6 hosts. IPv6 multicast addresses are mapped into the IPv4 multicast addresses. The IPv4 network becomes a virtual Ethernet for the IPv6 network; to achieve this, an IPv4 multicast-enabled network is required.

Translation Transition Mechanism

Dual-stack and tunneling techniques manage the interconnection of IPv6 domains. For legacy equipment that will not be upgraded to IPv6 and for some deployment scenarios, techniques are available for connecting IPv4-only nodes to IPv6-only nodes, using translation, an extension of NAT techniques.

As shown in the following figure, an IPv6 node behind a translation device has full connectivity to other IPv6 nodes and uses NAT functionality to communicate with IPv4 devices.

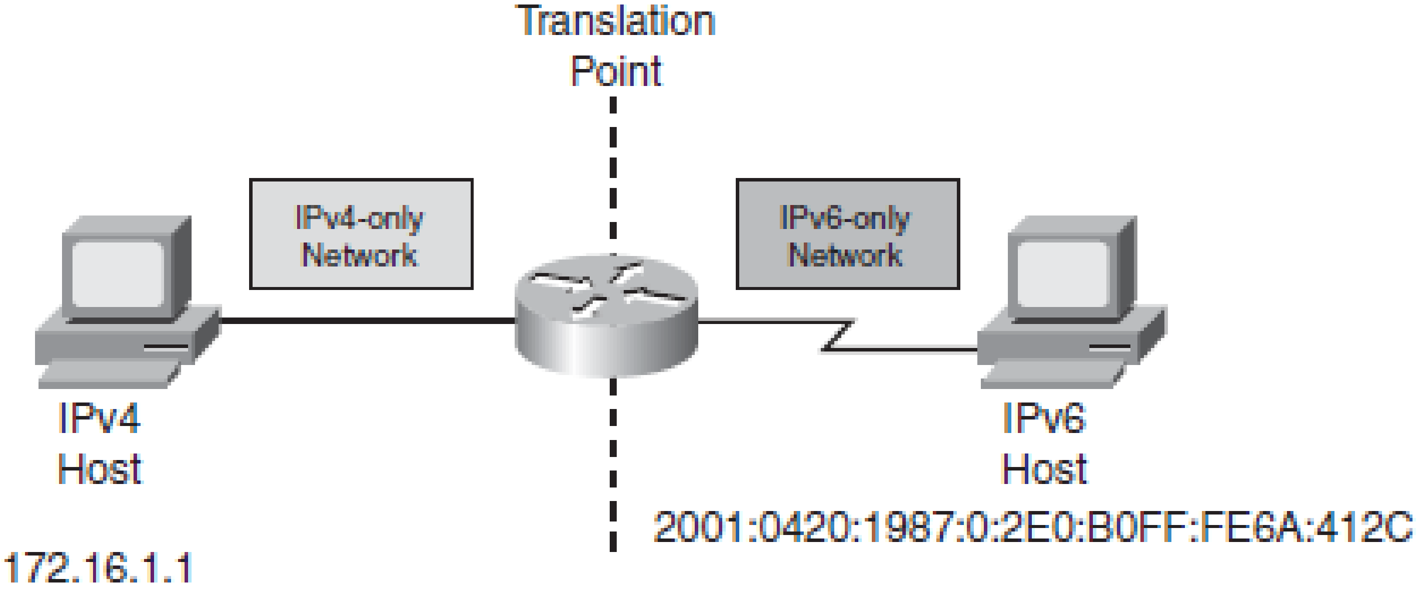

Translation Transition Mechanism

Translation techniques are available for translating IPv4 addresses to IPv6 addresses and vice versa. Similar to current NAT devices, translation is done at either the transport layer or the network layer.

NAT - Protocol Translation (NAT-PT) is the main translation technique; the Dual- Stack Transition Mechanism (DSTM) might also be available.

The NAT-PT translation mechanism translates at the network layer between IPv4 and IPv6 addresses and allows native IPv6 hosts and applications to communicate with native IPv4 hosts and applications. An application-level gateway (ALG) translates between the IPv4 and IPv6 DNS requests and responses.

Thank you