Chapter 2

Internet Protocol Version 6 (IPv6)

Internet Protocol version 6, IPv6 (Deering, S. and Hinden, R., 1998), the new version of the Internet Protocol is designed to be an evolutionary step from Internet Protocol version 4 (University of Southern California, 1981), the current version of Internet Protocol. The Internet Engineering Steering Group approved IPv6 on 17 November 1994 as a Proposed Standard.

2.1 Motivation

IPv6 has been designed to enable high-performance, scalable internetworks to remain viable well into the next century. A large part of this design process involved correcting the inadequacies of IPv4. The enhanced features of IPv6 are larger address space, streamlined packet design, well-structured and efficient routing hierarchy, ease of administration, better support of security and QoS (Quality of Service). The motivations for the development of a new Internet Protocol are listed below:

Address Space Depletion

The world is running out of IP addresses for networked devices resulting from the rapid growth of the Internet. Communications technologies need permanent connection to the Internet (Stallings, W., 1996; Microsoft Corporation, 2000).

Because of insufficient address space, some organizations are forced to use temporary technologies such as Network Address Translator (NAT) (Srisuresh, P. and Holdrege, M., 1999) to map multiple private addresses to a single public IP address. However, problems arise when connecting two organizations that use the private address space because they do not support standards-based network layer security or the correct mapping of all higher layer protocols.

Hierarchical Addressing System

Without an address hierarchy system, backbone routers would be forced to store routing table information on the reachability of every network in the world (Bay Networks, 1997). With an address hierarchy system, backbone routers can use IP address prefixes to determine how traffic should be routed through the backbone.

Currently, IPv4 uses Classless Inter-Domain Routing (CIDR) (Fuller, V. et al., 1993) to allow flexible use of variable-length network prefixes. CIDR permits considerable "route aggregation" at various levels of the Internet hierarchy so that backbone routers can store a single routing table entry that provides reachability to many lower-level networks. However, CIDR does not guarantee an efficient and scalable hierarchy. Legacy IPv4 address assignment originating before CIDR do not facilitate summarization. The lack of uniformity of the current hierarchical system coupled with the rationing of IPv4 addresses complicates the situation.

The hierarchical approach to IPv6 is believed to make automatic router configuration a much more viable proposition than the current Internetwork at the moment.

System Management

The current IPv4 implementation must be either configure manually or use a stateful address configuration protocol such as Dynamic Host Configuration Protocol (DHCP) (Droms, R., 1997). With more computers and devices using IP, there is a need for a simpler and more automatic configuration of addresses and other configuration settings to reduce address administration workloads.

Security

Encryption, authentication, and data integrity safeguards are increasingly a standard aspect of enterprise internetworking. Vendors in the IPv4 arena are not very successful in adding robust security features to Network Layer components largely due to the lack of interoperability caused by proprietary security extensions. In IPv4, Internet Protocol security (IPsec) is not mandatory.

The proponents of IPv6 claim that to achieve the same level of security with IPv4 as is available with IPv6 would need more work, and would thus cost more money, than upgrading to the improved protocol.

Quality of Service (QoS)

Real-time traffic support relies on the IPv4 Type of Service (TOS) field and the identification of the payload, typically using a User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) port to deliver Quality of Service. QoS is getting more important as one of the significant shifts in the future Internet traffic is a

huge growth in the use of the Internet as a broadcast medium carrying video and audio.

However, the IPv4 TOS field has limited functionality and there were various local interpretations. Furthermore, payload identification using a TCP and UDP port is not possible when the IPv4 packet payload is encrypted.

2.2 IPv6 Features

The important features of the IPv6 protocol are as follows:

New Header Format

The new IPv6 header format is designed to keep header overhead to a minimum. Both non-essential fields and option fields in IPv4 are moved to extension headers that are placed after the IPv6 header (Deering, S. and Hinden, R., 1998).

Most of these optional headers are not examined or processed by intermediate nodes on the packet's path. In this way, IPv6 exhibits more efficient forwarding than IPv4. It is also easier now to add additional options. Processing of the IPv6 packet header is simpler as the IPv6 packet header is fixed-length whereas the IPv4 header is variable-length. Furthermore, only the source but not the IPv6 routers can perform packet fragmentation that is always time consuming.

IPv4 headers and IPv6 headers are not interoperable. A host or router must implement both IPv4 and IPv6 in order to recognize and process both headers formats. The new

IPv6 header is only twice as large as the IPv4 header although IPv6 addresses are four times larger than IPv4 addresses.

Large Address Space

IPv6 has 128-bit (16-byte) source and destination IP addresses that can express over 3.4 * 10^{38} possible combinations (Hinden, R. and Deering, S., 1998). The large address space of IPv6 allows multiple levels of subnetting and address allocation from the Internet backbone to the individual subnets within an organization. Address-conservation techniques such as NAT are no longer necessary.

Efficient and hierarchical addressing and routing infrastructure

IPv6 implements address hierarchy where backbone routers use IP address prefixes to determine how traffic should be routed through the backbone, thus improve the routing efficiency. The IPv6 routing is simplified, as the IPv6 packet header is fixed-

Extensibility

IPv6 can easily be extended for new features by adding extension headers after the IPv6 header. The size of IPv6 extension headers is only limited by the size of the IPv6 packet, compare to IPv4 which can only support 40 bytes of options in the IPv4 header (Deering, S. and Hinden, R., 1998).

Multicast and Anycast

IPv6 extends IP multicasting capabilities in IPv4 by defining a very large multicast address space and a scope identifier that is used to limit the degree to which multicast routing information is propagated throughout an enterprise. Multicasting is an important feature in IPv6, as multicasting will eventually replace the IPv4 broadcast feature.

Conceptually anycast is a cross between unicast and multicast. Two or more interfaces on an arbitrary number of nodes are designated as an anycast group. A packet addressed to the group's anycast address is delivered to at least one of the interfaces in the group, typically "nearest" interface in the group, according to the routing protocols' measure of distance.

2.3 Differences between IPv4 and IPv6

Table 2.1 demonstrates the few main differences between IPv4 and IPv6.

Table 2.1 Differences between IPv4 and IPv6

IPv4	IPv6
Source and destination addresses are 32 bits (4 bytes) in length.	Source and destination addresses are 128 bits (16 bytes) in length.
Fragmentation is supported at both routers and the sending host.	Fragmentation is not supported at routers. Fragmentation is only supported at the sending host.
Header includes a checksum.	Header does not include a checksum.
Header includes options.	All optional data is moved to IPv6 extension headers.
IPSec support is optional.	IPSec support is mandatory.
No identification of payload for QoS handling by routers is present with the IPv4 header.	Payload identification for QoS handling by routers is included in the Flow Label inside the IPv6 header.
Address Resolution Protocol (ARP) broadcast ARP Request frames to resolve an IPv4 address to a link layer address.	ARP Request frames are replaced with multicast Neighbor Solicitation messages.
Uses Internet Group Management Protocol (IGMP) to manage local subnet group membership.	Replace IGMP with Multicast Listener Discovery (MLD).
Must be configured either manually or through DHCP.	Do not need manual configuration or DHCP.
Uses host address (A) resource records in the Domain Name System (DNS) to map host name to IPv4 addresses.	Uses host address (AAAA) resource records in DNS to map host names to IPv6 addresses.
Uses pointer (PTR) resource records in the IN-ADDR.ARPA DNS domain to map IPv4 addresses to host names.	Uses PTR resource records in the IP6.INT DNS domain to map IPv6 addresses to host names.

2.4 IPv6 Formats and Functions

2.4.1 IPv6 Packet

An IPv6 protocol data unit (known as a packet) (Deering, S. and Hinden, R., 1998) has the general form show in Figure 2.1.

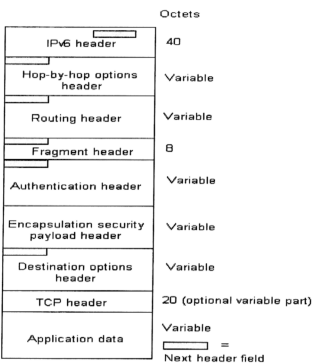

Figure 2.1 *IPv6 packet with all extension headers*

The only header that's required is IPv6 header. IPv6 header is a fixed size with a length of 40 octets, compare to 20 octets for the mandatory portion of the IPv4 header. The extension headers are as follows:

1) Hop-by-hop Options Header (Deering, S. and Hinden, R., 1998)
2) Routing Header (Deering, S. and Hinden, R., 1998)
3) Fragment Header (Deering, S. and Hinden, R., 1998)
4) Authentication Header (Kent, S. and Atkinson, R., 1998; Kent, S. and Atkinson, R., 1998c)

5) Encapsulating Security Payload Header (Kent, S. and Atkinson, R., 1998b)

6) Destination Options Header (Deering, S. and Hinden, R., 1998)

2.4.2 IPv6 Header

IPv6 header (Deering, S. and Hinden, R., 1998) is a streamlined version of the IPv4 header. Unneeded or rarely used fields in IPv4 are removed and some fields are added to provide better support of real-time traffic. The format of the IPv4 header and IPv6 header is shown in Figure 2.2.

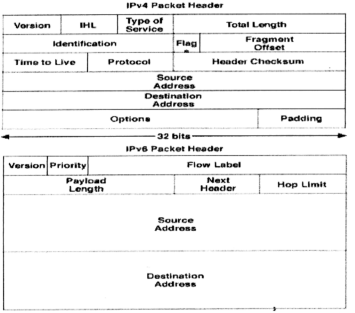

Figure 2.2 *IPv4 and IPv6 Header*

The IPv6 header has a fixed length of 40 octets and the fields in IPv6 header and the descriptions of IPv6 header are shown in Table 2.2.

Table 2.2 Fields of the IPv6 header and descriptions

Field	Description
Version (4 bits)	IP version number, 6 for IPv6
Traffic Class (8 bits)	Indicates the class or priority of the IPv6 packet.
Flow Label (20 bits)	Indicates that this packet belongs to a specific sequence of packets between a source and destination, requiring special handling by the intermediate IPv6 routers, such as non-default QoS or "real-time" services.
Payload Length (16 bits)	Length of the remainder of the IPv6 packet following the header in octets. The payload length field includes the extension headers and the upper layer Protocol Data Unit (PDU).
Next Header (8 bits)	Identifies the type of the header immediately following the IPv6 header.
Hop Limit (8 bits)	Decremented by 1 by each node that forwards the packet. The packet is discarded if Hop Limit is decremented to zero.
Source Address (128 bits)	Stores the IPv6 address of the originating host.
Destination Address (128 bits)	Stores the IPv6 address of the current destination host.

2.4.3 Differences between the IPv4 and IPv6 headers

The differences between IPv4 and IPv6 headers are shown in Table 2.3.

Table 2.3 Differences between the IPv4 Header fields and Corresponding IPv6 Equivalents

IPv4 Header Field	IPv6 Header Field
Version	Same field but with different version number.
Internet Header Length	Removed in IPv6, as IPv6 header is always a fixed size of 40 bytes. Each extension header is either a fixed size or indicates its own size.
Type of Services	Replaced by the IPv6 Traffic Class field.
Total Length	Replaced by the IPv6 Payload Length field, which only indicates the size of the payload.
Identification	Removed in IPv6 as fragmentation information is contained in a Fragment extension header.
Fragmentation Flags	Removed in IPv6 as fragmentation information is contained in a Fragment extension header.
Fragmentation Offset	Removed in IPv6 as fragmentation information is contained in a Fragment extension header.
Time to Live	Replaced by the IPv6 Hop Limit field.
Protocol	Replaced by the IPv6 Next Header field.

Header Checksum	Removed in IPv6 as bit-level error detection for the entire IPv6 packet is performed by the link layer.
Source Address	Same field but the IPv6 address is 128 bits in length.
Destination Address	Same field but the IPv6 address is 128 bits in length.
Options	Removed in IPv6 as the IPv6 extension headers replace IPv4 options.

2.4.4 Extension Headers

The IPv4 header includes all options. Therefore each intermediate router must check for their existence and process them when request and result in performance degradation in the forwarding of IPv4 packets. In IPv6 delivery and forwarding options are moved to extension headers and only the Hop-by-Hop Options extension header must be process at each intermediate router. This increases IPv6 header processing speed and improves forwarding process performance.

In a typical IPv6 packet, no extension headers are present. If either the intermediate routers or the destination requires special handling, the sending host adds one or more extension headers.

Each extension header must fall on 64-bits (8-bytes) boundaries. Extension headers of variable size contain a Header Extension Header Length field and padding must be used to ensure that their size is a multiple of 8 bytes.

A variable number of the type-length-value (TLV) encoded "options", carrying by the Hop-by-Hop Options header and the Destination Options header is defined by Hinden, R. and Deering, S. (1998).

Figure 2.3 shows the Next Header field in the IPv6 header and zero or more extension headers that form a chain of pointers. Each pointer indicates the type of header that comes after the immediate header until the upper layer protocol is ultimately identified.

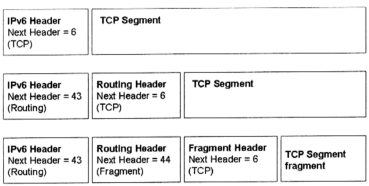

Figure 2.3 *IPv6 Extension headers*

List of extension headers are as follows:

1) Hop-by-hop Options Header (Deering, S. and Hinden, R., 1998)
2) Routing Header (Deering, S. and Hinden, R., 1998)
3) Fragment Header (Deering, S. and Hinden, R., 1998)
4) Authentication Header (Kent, S. and Atkinson, R., 1998; Kent, S. and Atkinson, R., 1998c)
5) Encapsulating Security Payload Header (Kent, S. and Atkinson, R., 1998b)
6) Destination Options Header (Deering, S. and Hinden, R., 1998)

Values of the Next Header Field

Table 2.4 shows the value of the Next Header field for an IPv6 header or an extension header.

Table 2.4 Values of the Next Header Field

Value (in decimal)	Header
0	Hop-by-Hop Options Header
6	Transmission Control Protocol (TCP)
17	User Datagram Protocol (UDP)
41	Encapsulated IPv6 Header
43	Routing Header
44	Fragmentation Header
46	Resource ReSerVation Protocol (RSVP)
50	Encapsulating Security Payload
51	Authentication Header
58	Internet Control Message Protocol version 6 (ICMPv6)
59	No next header
60	Destination Options Header

Extension Header Order

Referring to Deering, S. and Hinden, R. (1998), the extension headers are recommended to be organized in the following order when more than one extension header is used in the same packet:

1) IPv6 Header
2) Hop-by-Hop Options Header
3) Destination Options Header (for options to be processed by the first destination appears in the IPv6 Destination Address field plus subsequent destinations listed in the Routing header)
4) Routing Header
5) Fragment Header
6) Authentication Header

7) Encapsulating Security Payload header

8) Destination Options Header (for options to be processed only by the final destination of the packet)

9) Upper-layer header

Each extension header should occurs at most once, except for the Destination Options header which should be occur at most twice (one before a Routing header and one before the upper-layer header).

2.5 IPv6 Addressing Architecture

The most obvious distinguish feature of IPv6 is the use of much larger addresses. The size of an address in IPv6 is 128 bits, which is four times the larger than an address in IPv4. A 32-bit address space allows for 4,294,967,296 possible addresses. A 128-bit address space allows for 340,282,266,920,938,463,463,374,607,431,768,211,465 (or $3.4 * 10^{38}$) possible addresses.

2.5.1 Issues in current IPv4 Addressing Architecture

IPv6 provides an advanced hierarchical address space that facilitates efficient routing architectures. IPv4 was initially design with a class-based address scheme (Figure 2.4), which divided address bits between network and host but did not create a hierarchy that would allow a single high-level address to represent many lower-level addresses.

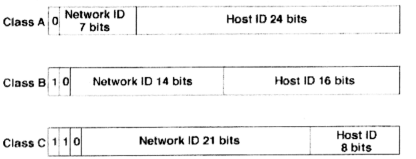

Figure 2.4 *IPv4 address classes*

The limitations of IPv4 addresses are hampering both the local and global levels of internetworking. Although techniques like NAT and CIDR are implemented to overcome the IPv4 deficiencies at the level of local area network (LAN) and Internet backbone, they are not able overcome all the problems due to the legacy design of the IPv4.

2.5.2 IPv6 Addressing

IPv6 addresses are 128-bit identifiers for interfaces and sets of interfaces. The text representations of the IPv6 addresses can be found in (Hinden, R. and Deering, S., 1998).

There are three types of addresses (Hinden, R. and Deering, S., 1998):

1) Unicast – An identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address.
2) 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.

3) 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 protocol's measure of distance).

There are no broadcast addresses in IPv6, their function being superseded by multicast addresses. In IPv6, all zeros and all ones are legal values for any field, unless specifically excluded. Specifically, prefixes may contain zero-valued fields or end in zeros.

2.5.3 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. A single interface may also be assigned multiple IPv6 addresses of any type (unicast, anycast, and multicast) or scope. Unicast addresses with scope greater than link-scope are not needed for interfaces that are not used as the origin or destination of any IPv6 packets to or from non-neighbors.

Currently IPv6 continues the IPv4 model that a subnet prefix is associated with one link. Multiple subnet prefixes may be assigned to the same link.

2.5.4 Address Type Specification

The leading bits in the address indicate the specific type of an IPv6 address. The high order bits and their fixed value are known as a Format Prefix (FP). Table 2.5 shows the allocation of the IPv6 address space by FPs.

Table 2.5 Current Allocation of the IPv6 Address Space

Allocation	Format Prefix (FP)	Fraction of the Address Space
Reserved	0000 0000	1/256
Unassigned	0000 0001	1/256
Reserved for Network Service Access Point (NSAP) allocation	0000 001	1/128
Reserved for Internetwork Packet Exchange (IPX) allocation	0000 010	1/128
Unassigned	0000 011	1/128
Unassigned	0000 1	1/32
Unassigned	0001	1/16
Aggregatable global unicast addresses	001	1/8
Unassigned	010	1/8
Unassigned	011	1/8
Unassigned	100	1/8
Unassigned	101	1/8
Unassigned	110	1/8
Unassigned	1110	1/16
Unassigned	1111 0	1/32
Unassigned	1111 10	1/64
Unassigned	1111 110	1/128
Unassigned	1111 1110 0	1/512
Link-local Unicast Addresses	1111 1110 10	1/1024
Site-local Unicast Addresses	1111 1110 11	1/1024
Multicast Addresses	1111 1111	1/256

2.5.5 IPv6 Prefixes

The prefix is the part of the address that indicates the bits that have fixed values or are the bits of the network identifier. IPv6 prefixes are expressed in the way similar to CIDR notation for IPv4. An IPv6 prefix is written in *address/prefix-length* notation.

For example, for an IPv6 address FE80:2222:3232:AE22:3355:DD33:EE22:3343 with prefix length 64, the prefix can be address as:

FE80:2222:3232:AE22:3355:DD33:EE22:3343/64 or

FE80:2222:3232:AE22::/64

2.5.6 Unicast Address

IPv6 unicast addresses are aggregatable with contiguous bit-wise masks similar to IPv4 addresses under CIDR. Several forms of unicast address assignment in IPv6 are as follows:

1) Aggregatable global unicast address
2) Link-local address
3) Site-local address
4) Unspecified address
5) Loopback address
6) NSAP address
7) IPX hierarchical address

Interface Identifiers

Interface identifiers in IPv6 unicast addresses are used to identify interfaces on a link. They are required to be unique on that link. They may also be unique over a broader scope. In many cases an interface's identifier will be the link-layer address. The same interface identifier may be used on multiple interfaces on a single node.

In a number of the format prefixes [see Table 2.5], Interface IDs are required to be 64 bits long and to be constructed in IEEE EUI-64 format (IEEE, 2001). EUI-64 based Interface identifiers may have global scope when a global token is available (e.g., IEEE 48 bit MAC) or may have local scope where a global token is not available (e.g., serial links).

Aggregatable Global Unicast Address

IPv6 has been designed from the ground up to provide a highly scalable address space that can be partitioned into a flexible and efficient global routing hierarchy. At the top of the hierarchy, several international registries assign blocks of addresses to top-level aggregators (TLA). These TLAs are essentially the public transit points (exchanges) where long-haul providers and telecom companies establish peer connections. TLAs allocate block of addresses to Next Level Aggregators (NLA), which represent large providers and global corporate networks. When a NLA is a provider, the NLA further allocates its addresses to its subscribers. More efficient routing is viable as the NLAs that are under the same TLA will have addresses with a common TLA prefix and subscribers with the same providers will have IP addresses with a common NLA prefix (R. Hinden et al., 1998).

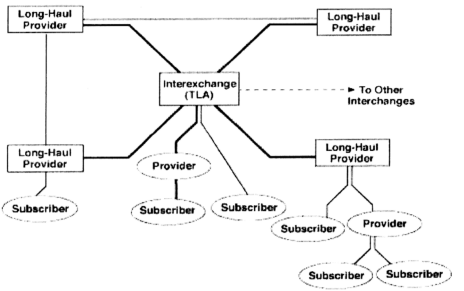

Figure 2.5 Aggregation-based allocation structures

Aggregation-based allocation is based on the existence today of a limited number of high-level exchange points, where large long-haul service providers and telecom companies networks interconnect. The use of these exchange points to divide IPv6 address hierarchy has a geographical component because exchanges are distributed around the globe. It also has a provider orientation because all large providers are represented at one or more exchange points.

In Figure 2.6, the first 3 address bits indicate what type of address follows (unicast, multicast, etc). The next 13 bits are allocated to the various TLAs around the world. The following 32 bits are allocated to the next lower level of providers and subscribers.

3	13	32 bits	16 bits	64 bits
001	TLA	NLA	SLA	Interface ID
	Public Topology		Site Topology	Local Interface

Figure 2.6 *Aggregation-based IP addresses*

Next level aggregators (NLA) can divide the NLA address field; create their own hierarchy that maps well to the current ISP industry, in which smaller ISPs subscribe to higher level IPSs, accomplish by the ongoing subdivision of the 32-bit NLA field.

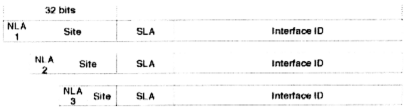

Figure 2.7 *Subdividing the NLA address space*

Following the NLA ID are fields for subscriber site networking information: Site Level Aggregator (SLA) and Interface ID. Service providers will typically supply subscribers with blocks of contiguous addresses, which are then used by individual organizations to create their own local addressing hierarchy and identify subnets and hosts. The 16-bit SLA field can supports up to 65,535 individual subnets. The 64-bit Interface ID will typically be derived from the installed IEEE LAN adapter address.

Link-local Address

Link-local addresses, identified by format prefix (FP) 1111 1110 10 are used by nodes when communicating with neighboring nodes on the same link. A link-local address is required for ND processes and is always automatically configured, even if all other unicast addresses are absent. The scope of a link-local address is the local link.

Site-local address

Site-local addresses, identified by FP of 1111 1110 11 are equivalent to the IPv4 private address space (10.0.0.0/8, 172.16.0.0/12 and 192.168.0.0/16). Private intranets that do not have direct, routed connection to the IPv6 Internet can use site-local addresses without conflicting with aggregatable global unicast addresses. Site-local addresses are not reachable from other sites, and routers must not forward site-local traffic outside the site. The scope of a site-local address is the site (the organization internetwork).

Unspecified address

The unspecified address (0:0:0:0:0:0:0:0 or ::) is only used to indicate the absence of the address. It can only used as the source address for packets attempting to verify the uniqueness of a tentative address. It is never assigned to an interface or used as the destination address of IPv6 packets.

Loopback address

The loopback address (0:0:0:0:0:0:0:1 or ::1) is used by a node to send an IPv6 packet to itself and must not be assigned to any physical interface. Packets addressed to the loopback address must never be sent on a link or forwarded by an IPv6 router.

NSAP address

NSAP addresses use the FP 0000001 and map the last 121 bits of the IPv6 address to an NSAP address. J. Bound et al. (1996) defined the NSAP address mapping.

IPX hierarchical address

IPX addresses use the FP 0000010 and map the last 121 bits of the IPv6 address to an IPX address. The IPX address mapping has not been defined yet.

2.5.7 Multicast Address

An IPv6 multicast address is an identifier for a group of nodes. A node may belong to any number of a multicast group. IPv6 multicast addresses have the FP of 1111 1111. An IPv6 address is easy to classify as multicast address because it always begins with "FF". Multicast addresses cannot be used as source addresses or as intermediate destinations in a Routing header. Format of a multicast group address is shown in Figure 2.8.

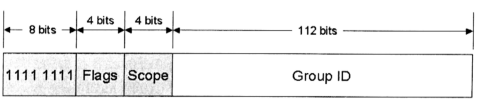

Figure 2.8 *The IPv6 multicast addresses*

Flags is a set of 4 flags:

0	0	0	T

The higher-order 3 flags are reserved, and must be initialized to 0.

T = 0 indicates a permanent-assigned ("well-known") multicast address, assigned by the global Internet numbering authority.

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

IPX hierarchical address

IPX addresses use the FP 0000010 and map the last 121 bits of the IPv6 address to an IPX address. The IPX address mapping has not been defined yet.

2.5.7 Multicast Address

An IPv6 multicast address is an identifier for a group of nodes. A node may belong to any number of a multicast group. IPv6 multicast addresses have the FP of 1111 1111. An IPv6 address is easy to classify as multicast address because it always begins with "FF". Multicast addresses cannot be used as source addresses or as intermediate destinations in a Routing header. Format of a multicast group address is shown in Figure 2.8.

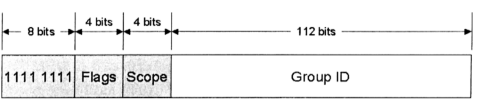

Figure 2.8 *The IPv6 multicast addresses*

Flags is a set of 4 flags:

0	0	0	T

The higher-order 3 flags are reserved, and must be initialized to 0.

T = 0 indicates a permanent-assigned ("well-known") multicast address, assigned by the global Internet numbering authority.

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

Scope is a 4-bit multicast scope value used to limit the scope of the multicast group. Table 2.6 lists the defined values for the Scope field.

Table 2.6 Defined Values for the Scope field

Value	Scope
0	Reserved
1	Node-local scope
2	Link-local scope
5	Site-local scope
8	Organization-local scope
E	Global scope
F	Reserved

Group ID identifies the multicast group, either permanent or transient, within the given scope.

Predefined Multicast Address

Multicast addresses from FF01:: through FF0F:: are reserved, well-known addresses (R. Hinden and S. Deering, 1998).

All Nodes Addresses: FF01:0:0:0:0:0:0:1 (node-local)

FF02:0:0:0:0:0:0:1 (link-local)

All Routers Addresses: FF01:0:0:0:0:0:0:2 (node-local)

FF01:0:0:0:0:0:0:2 (link-local)

FF01:0:0:0:0:0:0:2 (site-local)

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

The solicited-node multicast address is formed by taking the low-order 24 bits of the address (unicast or anycast) and appending these bits to the prefix

FF02:0:0:0:0:1:FF::/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.

Mapping IPv6 Multicast Address

The current approach to map IPv6 multicast addresses into IEEE 802 MAC addresses (Crawford, M., 1998) takes the low order 32 bits of the IPv6 multicast address and setting the remaining original group ID bits to 0. The modified IPv6 multicast address is shown in Figure 2.9.

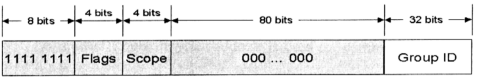

Figure 2.9 *The modified IPv6 multicast address using a 32-bit group ID*

2.5.8 Anycast Address

An IPv6 anycast address is an address that's 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. At present, anycast addresses are only used as destination address and are only assigned to routers.

Subnet-Router Anycast Address

The Subnet-Router anycast address is predefined and required. It is created from the subnet prefix for a given interface. The bits in the subnet prefix are fixed at their appropriate values and the remaining bits are set to 0. The Subnet-Router anycast address is shown in Figure 2.10.

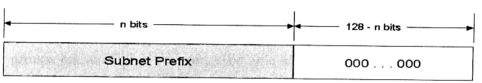

Figure 2.10 *The Subnet-Router anycast address*

All router interfaces attached to subnet are assigned the Subnet-Router anycast for that subnet. The Subnet-Router anycast address is used for communication with one of multiple routers attached to a remote subnet.