Name of Course : E1-E2 CFA

Chapter 4

Topic : Introduction to IPv6

Date of Creation : 19.03.2011

INTRODUCTION TO IPV6

The current version of IP (known as Version 4 or IPv4) has proven to be robust, easily implemented and interoperable, and has stood the test of scaling an internetwork to a global utility the size of today's Internet. This is a tribute to its initial design.

However, the initial design did not anticipate the following:

- The recent exponential growth of the Internet and the impending exhaustion of the IPv4 address space.

  IPv4 addresses have become relatively scarce, forcing some organizations to use a Network Address Translator (NAT) to map multiple private addresses to a single public IP address. While NATs promote reuse of the private address space, they do not support standards-based network layer security or the correct mapping of all higher layer protocols and can create problems when connecting two organizations that use the private address space.

- The growth of the Internet and the ability of Internet backbone routers to maintain large routing tables.

  Because of the way that IPv4 network IDs have been and are currently allocated, there are routinely over 85,000 routes in the routing tables of Internet backbone routers. The current IPv4 Internet routing infrastructure is a combination of both flat and hierarchical routing.

- The need for simpler configuration.

  Most current IPv4 implementations must be either manually configured or use a state full address configuration protocol such as Dynamic Host Configuration Protocol (DHCP). With more computers and devices using IP, there is a need for a simpler and more automatic configuration of addresses and other configuration settings that do not rely on the administration of a DHCP infrastructure.

- The requirement for security at the IP level.

  Private communication over a public medium like the Internet requires encryption services that protect the data being sent from being viewed or modified in transit. Although a standard now exists for providing security for IPv4 packets (known as Internet Protocol security or IPsec), this standard is optional and proprietary solutions are prevalent.

- The need for better support for real-time delivery of data—also called quality of service (QoS).

  While standards for QoS exist for IPv4, real-time traffic support relies on the IPv4 Type of Service (TOS) field and the identification of the payload, typically using a

UDP or TCP port. Unfortunately, the IPv4 TOS field has limited functionality and over time there were various local interpretations. In addition, payload identification using a TCP and UDP port is not possible when the IPv4 packet payload is encrypted.

To address these and other concerns, the Internet Engineering Task Force (IETF) has developed a suite of protocols and standards known as IP version 6 (IPv6). This new version, previously called IP-The Next Generation (IPng), incorporates the concepts of many proposed methods for updating the IPv4 protocol. The design of IPv6 is intentionally targeted for minimal impact on upper and lower layer protocols by avoiding the random addition of new features.

IPv6 Features

The following are the features of the IPv6 protocol:

- New header format
- Large address space
- Efficient and hierarchical addressing and routing infrastructure
- Stateless and stateful address configuration
- Built-in security
- Better support for QoS
- New protocol for neighboring node interaction
- Extensibility

The following sections discuss each of these new features in detail.

New Header Format

The IPv6 header has a new format that is designed to keep header overhead to a minimum. This is achieved by moving both non-essential fields and optional fields to extension headers that are placed after the IPv6 header. The streamlined IPv6 header is more efficiently processed at intermediate routers.

IPv4 headers and IPv6 headers are not interoperable. A host or router must use an implementation of both IPv4 and IPv6 in order to recognize and process both header formats. The new IPv6 header is only twice as large as the IPv4 header, even though IPv6 addresses are four times as large as IPv4 addresses.

Large Address Space

IPv6 has 128-bit (16-byte) source and destination IP addresses. Although 128 bits can express over $3.4?10^{38}$ possible combinations, the large address space of IPv6 has been designed to allow for multiple levels of subnetting and address allocation

from the Internet backbone to the individual subnets within an organization.

Even though only a small number of the possible addresses are currently allocated for use by hosts, there are plenty of addresses available for future use. With a much larger number of available addresses, address-conservation techniques, such as the deployment of NATs, are no longer necessary.

Efficient and Hierarchical Addressing and Routing Infrastructure

IPv6 global addresses used on the IPv6 portion of the Internet are designed to create an efficient, hierarchical, and summarizable routing infrastructure that is based on the common occurrence of multiple levels of Internet service providers. On the IPv6 Internet, backbone routers have much smaller routing tables, corresponding to the routing infrastructure of global ISPs.

Stateless and Stateful Address Configuration

To simplify host configuration, IPv6 supports both stateful address configuration, such as address configuration in the presence of a DHCP server, and stateless address configuration (address configuration in the absence of a DHCP server). With stateless address configuration, hosts on a link automatically configure themselves with IPv6 addresses for the link (called link-local addresses) and with addresses derived from prefixes advertised by local routers. Even in the absence of a router, hosts on the same link can automatically configure themselves with link-local addresses and communicate without manual configuration.

Built-in Security

Support for IPsec is an IPv6 protocol suite requirement. This requirement provides a standards-based solution for network security needs and promotes interoperability between different IPv6 implementations.

Better Support for QoS

New fields in the IPv6 header define how traffic is handled and identified. Traffic identification using a Flow Label field in the IPv6 header allows routers to identify and provide special handling for packets belonging to a flow, a series of packets between a source and destination. Because the traffic is identified in the IPv6 header, support for QoS can be achieved even when the packet payload is encrypted through IPsec.

New Protocol for Neighboring Node Interaction

The Neighbor Discovery protocol for IPv6 is a series of Internet Control Message Protocol for IPv6 (ICMPv6) messages that manage the interaction of

neighboring nodes (nodes on the same link). Neighbor Discovery replaces the broadcast-based Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and ICMPv4 Redirect messages with efficient multicast and unicast Neighbor Discovery messages.

Extensibility

IPv6 can easily be extended for new features by adding extension headers after the IPv6 header. Unlike options in the IPv4 header, which can only support 40 bytes of options, the size of IPv6 extension headers is only constrained by the size of the IPv6 packet.

Differences Between IPv4 and IPv6

Table-1 highlights some of the key differences between IPv4 and IPv6.

Table-1: Difference 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
IPsec support is optional.	IPsec support is required
No identification of packet flow for QoS handling by routers is present within the IPv4 header.	Packet flow identification for QoS handling by routers is included in the IPv6 header using the Flow Label field.
Fragmentation is done by both routers and the sending host.	Fragmentation is not done by routers, only by 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
Address Resolution Protocol (ARP) uses broadcast ARP Request frames to resolve an IPv4 address to a link layer address.	ARP Request frames are replaced with multicast Neighbor Solicitation messages
Internet Group Management Protocol (IGMP) is used to manage local subnet group membership.	IGMP is replaced with Multicast Listener Discovery (MLD) messages.
ICMP Router Discovery is used to determine the IPv4 address of the best default gateway and is optional.	ICMP Router Discovery is replaced with ICMPv6 Router Solicitation and Router Advertisement messages and is required.

Broadcast addresses are used to send traffic to all nodes on a subnet.	There are no IPv6 broadcast addresses. Instead, a link-local scope all-nodes multicast address is used
Must be configured either manually or through DHCP.	Does not require manual configuration or DHCP.
Uses host address (A) resource records in the Domain Name System (DNS) to map host names to IPv4 addresses.	Uses host address (AAAA) resource records in the Domain Name System (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 pointer (PTR) resource records in the IP6.INT DNS domain to map IPv6 addresses to host names
Must support a 576-byte packet size (possibly fragmented).	Must support a 1280-byte packet size (without fragmentation

IPv6 Packets over LAN Media

A link layer frame containing an IPv6 packet consists of the following structure:

- Link Layer Header and Trailer – The encapsulation placed on the IPv6 packet at the link layer.
- IPv6 Header – The new IPv6 header. For more information, see "IPv6 Header."
- Payload –The payload of the IPv6 packet. For more information, see "IPv6 Header."

Figure 1 shows the structure of a link layer frame containing an IPv6 packet.

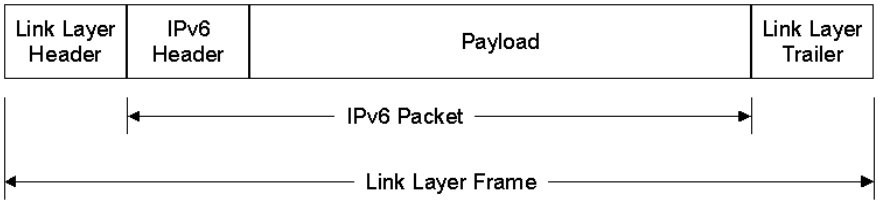

Figure 1 IPv6 packets at the link layer

For typical LAN technologies such as Ethernet, Token Ring, and Fiber Distributed Data Interface (FDDI), IPv6 packets are encapsulated in one of two

ways—with either the Ethernet II header or a Sub-Network Access Protocol (SNAP) header used by IEEE 802.3 (Ethernet), IEEE 802.5 (Token Ring), and FDDI.

Ethernet II Encapsulation

With Ethernet II encapsulation, IPv6 packets are indicated by setting the EtherType field in the Ethernet II header to 0x86DD (IPv4 is indicated by setting the EtherType field to 0x800). With Ethernet II encapsulation, IPv6 packets can have a minimum size of 46 bytes and a maximum size of 1,500 bytes. Figure 2 shows Ethernet II encapsulation for IPv6 packets.

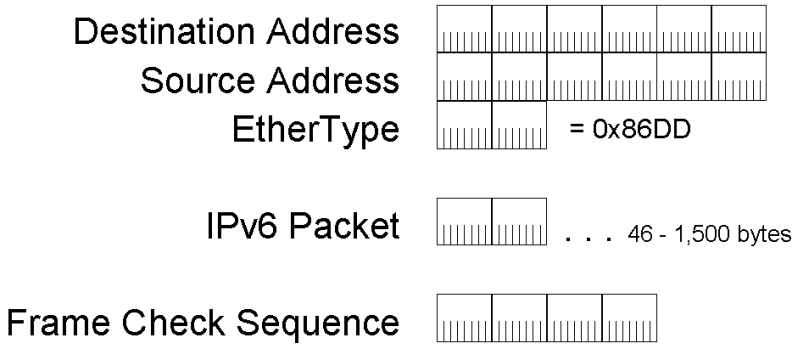

Figure 2: Ethernet II encapsulation

IPv6 Addressing

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}) 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 ($6.5x10^{23}$) 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 $6.5x10^{23}$ 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 2373.

Current Allocation

Similar to the same way in which the IPv4 address space is divided, the IPv6 address space is divided based on the value of high order bits. The high order bits and their fixed values are known as a Format Prefix (FP).

Table-2 shows the allocation of the IPv6 address space by FPs.

Table-2: 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 NSAP allocation	0000 001	1/128
Unassigned	0000 010	1/128
Unassigned	0000 011	1/128
Unassigned	0000 1	1/32
Unassigned	0001	1/16
Aggregatable global unicast addresses	001	1/8
Unassigned	010	1/8
Unassigned	011	1/8
Unassigned	100	1/8
Unassigned	101	1/8

Unassigned	110	1/8
Unassigned	1110	1/16
Unassigned	1111 0	1/32
Unassigned	1111 10	1/64
Unassigned	1111 110	1/128
Unassigned	1111 1110 0	1/512
Link-local unicast addresses	1111 1110 10	1/1024
Site-local unicast addresses	1111 1110 11	1/1024
Multicast addresses	1111 1111	1/256

The current set of unicast addresses that can be used with IPv6 nodes consists of aggregatable global unicast addresses, link-local unicast addresses, and site-local unicast addresses. These represent only 15 percent of the entire IPv6 address space.

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:

00100001110110100000000011010011000000000000000000101111001110110000001010101010000000001111111111111110001010001001110001011010

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

0010000111011010 0000000011010011 0000000000000000 0010111100111011 0000001010101010 0000000011111111 1111111000101000 1001110001011010

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

21DA:00D3:0000:2F3B: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:

21DA:D3:0:2F3B: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.

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 network identifier. Prefixes for IPv6 subnet identifiers, 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 is a route prefix and 21DA:D3:0:2F3B::/64 is a subnet prefix.

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. To accommodate load-balancing systems, RFC 2373 allows for multiple interfaces to use the same address as long as they appear as a single interface to the IPv6 implementation on the host.

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.

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. A multicast address is used for one-to-many communication, with delivery to multiple interfaces. 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.

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.

- Loop-back address

  The loop-back address (0:0:0:0:0:0:0:1 or ::1) is used to identify a loop-back interface, enabling a node to send packets to itself. It is equivalent to the IPv4 loop-back address of 127.0.0.1. Packets addressed to the loop-back 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.

- **6 to 4 address**

  The 6 to 4 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 of the node, forming a 48-bit prefix. 6to4 is a tunneling technique described in RFC 3056.

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 FP of 11111111. 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 header.

Beyond the FP, multicast addresses include additional structure to identify their flags, scope, and multicast group. Figure 8 shows the IPv6 multicast addresses.

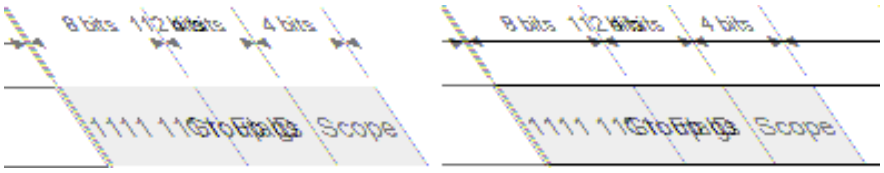

Figure 8: The IPv6 multicast address

The fields in the multicast address are:

Flags – Indicates flags set on the multicast address. The size of this field is 4 bits. As of RFC 2373, the only flag defined is the Transient (T) flag. The T flag uses the low-order bit of the Flags field. When set to 0, the T flag indicates that the multicast address is a permanently assigned (well-known) multicast address allocated by the Internet Assigned Numbers Authority (IANA). When set to 1, the T flag indicates that the multicast address is a transient (non-permanently-assigned) multicast address.

Scope – Indicates the scope of the IPv6 internet-work 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.

Table-3 lists the values for the Scope field defined in RFC 2373.

Table-3: 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

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 node-local and link-local scopes, the following addresses are defined:

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

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

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

With 112 bits for the Group ID, it is possible to have 2^{112} group IDs. However, because of the way in which IPv6 multicast addresses are mapped to Ethernet multicast MAC addresses, RFC 2373 recommends assigning the Group ID from the low order 32 bits of the IPv6 multicast address and setting the remaining original group ID bits to 0. By using only the low-order 32 bits, each group ID maps to a unique Ethernet multicast MAC address. Figure 9 shows the modified IPv6 multicast addresses.

Figure 9: The modified IPv6 multicast address using a 32-bit group ID

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.

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. Figure 10 shows the Subnet-Router anycast address.

Figure 10: The Subnet-Router anycast address

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 Header

The IPv6 header is a streamlined version of the IPv4 header. It eliminates fields that are unneeded or rarely used and adds fields that provide better support for real-time traffic.

Structure of an IPv6 Packet

Figure 17 shows the structure of an IPv6 packet.

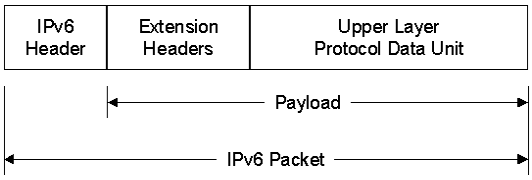

Figure 17: The structure of an IPv6 packet

IPv6 Header

The IPv6 header is always present and is a fixed size of 40 bytes. The fields in the IPv6 header are described in detail later.

Extension Headers

Zero or more extension headers can be present and are of varying lengths. A Next Header field in the IPv6 header indicates the next extension header. Within each extension header is a Next Header field that indicates the next extension header. The last extension header indicates the upper layer protocol (such as TCP, UDP, or ICMPv6) contained within the upper layer protocol data unit.

The IPv6 header and extension headers replace the existing IPv4 IP header with options. The new extension header format allows IPv6 to be augmented to support future needs and capabilities. Unlike options in the IPv4 header, IPv6 extension headers have no maximum size and can expand to accommodate all the extension data needed for IPv6 communication.

Upper Layer Protocol Data Unit

The upper layer protocol data unit (PDU) usually consists of an upper layer protocol header and its payload (for example, an ICMPv6 message, a UDP message, or a TCP segment).

The IPv6 packet payload is the combination of the IPv6 extension headers and the upper layer PDU. Normally, it can be up to 65,535 bytes long. Payloads greater than 65,535 bytes in length can be sent using the Jumbo Payload option in the Hop-by-Hop Options extension header.

Figure 19 shows the IPv6 header as defined in RFC 2460.

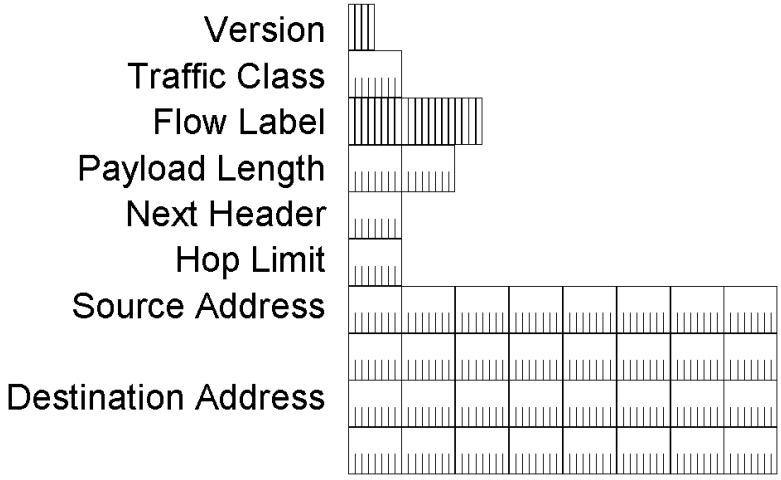

Figure 19: The IPv6 header

The fields in the IPv6 header are:

Version – 4 bits are used to indicate the version of IP and is set to 6.

Traffic Class – Indicates the class or priority of the IPv6 packet. The size of this field is 8 bits. The Traffic Class field provides similar functionality to the IPv4 Type of Service field. In RFC 2460, the values of the Traffic Class field are not defined. However, an IPv6 implementation is required to provide a means for an application layer protocol to specify the value of the Traffic Class field for experimentation.

Flow Label – Indicates that this packet belongs to a specific sequence of packets between a source and destination, requiring special handling by intermediate IPv6 routers. The size of this field is 20 bits. The Flow Label is used for non-default quality of service connections, such as those needed by real-time data (voice and video). For default router handling, the Flow Label is set to 0. There can be multiple flows between a source and destination, as distinguished by separate non-zero Flow Labels.

Payload Length – Indicates the length of the IPv6 payload. The size of this field is 16 bits. The Payload Length field includes the extension headers and the upper layer PDU. With 16 bits, an IPv6 payload of up to 65,535 bytes can be indicated. For payload lengths greater than 65,535 bytes, the Payload Length field is set to 0 and the Jumbo Payload option is used in the Hop-by-Hop Options extension header.

Next Header – Indicates either the first extension header (if present) or the protocol in the upper layer PDU (such as TCP, UDP, or ICMPv6). The size of this field is 8 bits. When indicating an upper layer protocol above the Internet layer, the same values used in the IPv4 Protocol field are used here.

Hop Limit – Indicates the maximum number of links over which the IPv6 packet can travel before being discarded. The size of this field is 8 bits. The Hop Limit is similar to the IPv4 TTL field except that there is no historical relation to the amount of time (in seconds) that the packet is queued at the router. When the Hop Limit equals 0, an ICMPv6 Time Exceeded message is sent to the source address and the packet is discarded.

Source Address –Stores the IPv6 address of the originating host. The size of this field is 128 bits.

Destination Address – Stores the IPv6 address of the current destination host. The size of this field is 128 bits. In most cases the Destination Address is set to the final destination address. However, if a Routing extension header is present, the Destination Address might be set to the next router interface in the source route list.

Values of the Next Header Field

Table-5 shows the typical values of the Next Header field for an IPv6 header or an IPv6 extension header

Table-5: Values of the Next Header Field

Value (in decimal)	Header
0	Hop-by-Hop Options Header
6	TCP
17	UDP
41	Encapsulated IPv6 Header
43	Routing Header
44	Fragment Header
46	Resource ReSerVation Protocol
50	Encapsulating Security Payload

51	Authentication Header
58	ICMPv6
59	No next header
60	Destination Options Header

Comparing the IPv4 and IPv6 Headers

Table-6 shows the differences between the IPv4 and IPv6 header fields.

Table-6: IPv4 Header Fields and Corresponding IPv6 Equivalents

IPv4 Header Field	IPv6 Header Field
Version	Same field but with different version numbers.
Internet Header Length	Removed in IPv6. IPv6 does not include a Header Length field because the 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 Service	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 Fragmentation Flags Fragment Offset	Removed in IPv6. Fragmentation information is not included in the IPv6 header. It 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. In IPv6, bit-level error detection for the entire IPv6 packet is performed by the link layer.
Source Address	The field is the same except that IPv6 addresses are 128 bits in length.
Destination Address	The field is the same except that IPv6 addresses are 128 bits in length.
Options	Removed in IPv6. IPv4 options are replaced by IPv6 extension headers.

The one new field in the IPv6 header that is not included in the IPv4 header is the Flow Label field.

IPv6 Extension Headers

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

RFC 2460 defines the following IPv6 extension headers that must be supported by all IPv6 nodes:

- Hop-by-Hop Options header
- Destination Options header
- Routing header
- Fragment header
- Authentication header
- Encapsulating Security Payload header

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

xxxx