Presentation 3

IPv6

Version 6 of the Internet Protocol (IP) was designed to solve various problems associated with version 4 and provides an essential basis for providing advanced networking services.

Slide 2 IPv4 & IPv6 Header Comparison

Several header fields in IPv4 have been replaced: IHL, type of service, total length, identification, flags, fragment offset, TTL, protocol, and header checksum.

Roughly speaking, type of service is replaced by **Traffic Class,** total length is replaced by **Payload Length**, TTL is replaced by **Hop Limit** (operating in the same way), and protocol is replaced by **Next Header**.

The **Next Header** field indicates the type of header following the IPv6 header — it may be another IP header, used for implementing options, or another protocol.

The **Flow Label** can used as a hint to routers to route packets in the same way.

Slides 3 – 8 IPv6 Addresses

Addresses in IPv6 depend on the scope of their use — link local, site local, or global.

A **link local** address is for use on a local hardware connection over the broadcast domain of a LAN. A **site local** address is for use in the autonomous system (AS) one organization. A **global** address must be legal in the entire Internet.

An address may be defined as **Unicast**, specifying delivery to one address on a single interface, **Multicast**, specifying delivery to a set of interfaces, or **Anycast**, specifying delivery to the closest single interface in set of interfaces.

The IPv6 address is **128 bits**, defining 340,282,366,920,938,463,463,374,607,431,768,211,456 separate addresses (50,000 addresses per square meter of land on Earth). The address is written as a colon-separated **16-bit hex number**, with leading zeros optional (see slide 5) in any 16-bit field. When there are several successive 16-bit 0 fields one (only) continuous field of zeros may be replaced with a double colon (::).

IPv4 compatible addresses are 32-bit IPv4 addresses prefixed with 96 zeros. These are used to identify IPv6 datagrams to be sent into an IPv4 tunnel. An **IPv4 mapped** address is used as the destination address in an IPv6 datagram to be sent to an IPv4 destination. It prefixes the IPv4 address with 0:0:0:0:0:FFFF.

IPv6 address types are determined by the initial bits called a **prefix**. The prefixes 00000000 (binary) and 0000001 are **reserved**. The prefix 11111111 indicates a **multicast address**.

The formats for **link local**, **site local**, and **global** addresses are shown on slide 7.

A link local address is formed by adding a fixed prefix to the hardware MAC address of the node.

A site local address is formed by adding a fixed prefix and a 54-bit administrator assigned subnet address to the MAC address.

A global address is formed by the 3-bit prefix, 45-bit ISP-assigned address, 16-bit locally assigned subnet address, and a 64-bit node address (assigned or derived from MAC address).

The format of a **global unicast address** is

3 bits	45 bits (ISP)	Subnet (16 bits)	64 bits
001	Routing Prefix	Subnet ID	Interface ID (host)

Slide 8 shows an example of IPv6 global addresses:

An **addressing authority** assigns the address range 2001:0410::/32 to the ISP.

3 bits	45 bits (ISP)	Subnet (16 bits)	64 bits
001	Routing Prefix	Subnet ID	Interface ID (host)
001	0 0000 0000 0001 0000 0100 0001 0000		
2001:0410			

The **ISP** assigns 2001:0410:1:/48 to customer 1.

3 bits	45 bits (ISP)	Subnet (16 bits)	64 bits
001	Routing Prefix	Subnet ID	Interface ID (host)
001	0 0000 0000 0001 0000 0100 0001 0000	0000 0000 0000 0001	
2001:0410		1	

Customer 1 can now assign any 64-bit address to each host in its AS.

Slides 9 – 11 Extension Headers

The **next header** field points to the type of header following IPv6 header. Slide 9 shows a list of upper layer protocol headers, such as TCP and UDP, and a list of **IP option headers** for specific services. Some option headers are:

Tunneling	An IPv6 datagram is in data field of this IPv6 datagram. This is used with fragmentation when router adds option field.
Hop-by-hop header	Options that must be checked by every router and destination node
Destination options header	Options that must be checked by the destination node
Routing header	The source sets the route to be taken through the network
Fragment header	Handles fragmentation
Authentication header **Encapsulating Security Payload**	For implementing IPsec

Slide 11 shows examples of chaining IPv6 headers.

Slide 12 **Traffic Class + Flow**

The traffic class field is 8 bits, consisting of 6 bits for the **Differentiated Services Code Point (DSCP)** and 2 bits for **Explicit Congestion Notification (ECN)**.

The DSCP field sets per-hop behavior according to the service required. The ECN field permits routers to indicate congestion to hosts.

A **flow** is a stream of related packets from one source to one destination that requires particular handling by routers, for example real time priority. These requirements can be cached in routers. A 24-bit flow label is chosen randomly from 1 to FFFFF, where 0 indicates that the datagram is not part of a flow. A flow is identified by its label, source IP, and destination IP.

Slides 13 – 17 **Assistive Protocols**

Slide 13 shows the structure and some codes for the version 6 **Internet Control Message Protocol (ICMP)**. New message types include **Router Solicitation** and **Router Advertisement**, **Neighbor Solicitation**, and **Neighbor Advertisement**.

Neighbor solicitation and neighbor advertisement can replace the ARP protocol. To discover the MAC address of a local node, a host sends a neighbor solicitation addressed to the node IP address, with the global prefix replaced with the **link local multicast prefix** FF02:: (64 bits). The neighbor sends ICMP **neighbor advertisement** containing its MAC address as the node ID.

Router solicitation and router advertisement can be used to replace aspects of DHCP. In advertisements, routers send ICMP messages to hosts on some schedule. These are addressed to FF02::1 which addresses all systems on the local link. It provides service configuration parameters, such as maximum transition unit (MTU), the global IP prefix, and the preferred IP address configuration method (ICMP / stateless). In router solicitation, a host sends an ICMP message to request router service addressed to FF02::2, the multicast address for all routers on the local link. Routers respond with a router advertisement.

Slide 16 describes the DHCP operation in IPv6 for host configuration, similar to IPv4 operation.

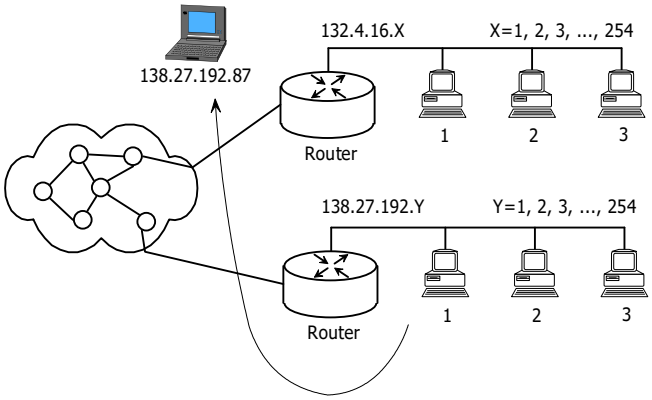
IPv6 also permits **stateless auto-configuration** by which a node can assign itself an IP address that enables communication with a configuration server (see slide 17). The steps are:

1. The node reads its 48-bit MAC address from hardware, splits it into 24-bit slices, and converts it to a 64-bit host ID given by: 24 bits FFEE 24 bits. For example, the MAC address **`11:22:33:44:55:66`** becomes the host ID **`1122 33FF EE44 5566`**
2. The node extends the link local prefix to 64 bits as **`FE80::/10`** $\longrightarrow$ **`FE80::/64`** and prefixes the node address to create a temporary link local unicast address. For example, the host ID becomes the IP address **`FE80::1122:33FF:EE44:5566`**
3. The host verifies that this is a unique address with neighbor solicitation. If there is no response then no other node has this address and it may be used by the host.
4. The host sends a router solicitation. If no routers respond with a router advertisement, the host attempts DHCP. The response message provides network parameters, such as the global prefix for the AS. The host replaces its link local prefix with the global prefix to form a global unicast address.

Slides 18 – 30 Mobile IP

IPv6 permits **management of mobile devices** using fixed IP addresses by implementing Mobile IP, a protocol developed as an optional addition to IPv4. This protocol is a good introduction to mobility management.

The problem to be solved in IPv4 begins with a mobile device, whose home network is a LAN subnet with network number 138.27.192.0. This user requires a fixed IP address to provide accessibility and network services for customers. But the user is mobile and wants to attach to a different LAN subnet with network number 132.4.16.0. The device will not be able to automatically obtain router connectivity, because the router and device are on different subnets with different subnet masks. Also, any datagrams addressed to the device will be routed through the system to the user's home subnet.



The basic idea of mobile IP is similar to call forwarding. The mobile node uses its home router when connected to its home network, but finds a mobile service router when roaming (connected to a non-home network). The home router and mobile service router coordinate between themselves, so that IP datagrams are sent to the usual home router, forwarded by home router to mobile service router, and then forwarded to the mobile node.

A mobile node (**MN**) is any host that can change its point of attachment. Its Home Address is its permanent IP address.

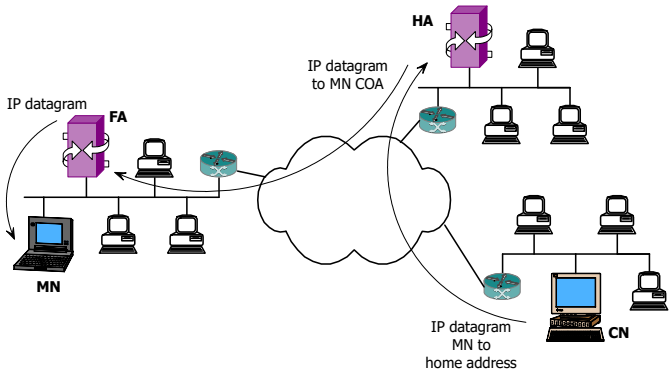
A correspondent node (**CN**) is any node that sends datagrams to the **MN** home address.

The home agent (**HA**) maintains a table of registered mobile nodes and forwards datagrams addressed to mobile nodes.

A foreign agent (**FA**) delivers datagrams between the **MN** and **HA**.

A mobility agent is any **HA** or **FA** supporting mobility.

A Care-of-Address (**COA**) is a FA address used to identify the current location of the MN.



The setup of a mobile connection is as follows:

- To permit discovery, mobility agents broadcast their availability, usually implemented within a router. The mobile node (MN) looks for local router services (HA) when connected to its home network
- When connected to a foreign network, the MN looks for a router advertising Foreign Agent (FA) services.
- The MN requests service from the FA, which assigns a Care-of-Address (COA) to the MN. The FA may have 1 or more available COAs within the IP address space of its subnet, but usually assigns the same COA to all MNs it serves.
- The MN must register to receive service. The FA requests mobile IP support from the MN's home agent (HA). If the HA agrees to provide Mobile IP support, the MN registers its COA with its HA, so that the HA can forward datagrams to the FA using the COA.
- Datagrams from a correspondent node (CN) arrive at the HA by normal routing (they are addressed to the MN).
- The HA encapsulates the arriving datagram in a tunnel datagram by adding a new IP header (the original datagram is treated by intermediate routers as meaningless data). The new IP datagram is addressed to the COA address from the HA address and arrives at the FA by normal routing.
- The FA recognizes the COA and removes the original datagram from the encapsulated datagram. This datagram must be handled carefully — it cannot be routed by the normal procedures of the routing table because the destination address is on the home network of the roaming MN. The FA places the datagram in a LAN data link (layer 2) frame addressed to the MN on the FA's network (even though the destination address in the MAC frame and the destination address in the IP datagram disagree).

Another complication is further roaming and **hand-off**. Suppose the roaming MN moves to new attachment point (network), so that it requires a change of FA. In this case, to avoid datagram loss (and to avoid very frequent registrations) the MN registers and maintains multiple bindings with multiple COAs, the old and new. The HA can now forward each packet to multiple COAs and the MN receives packet at one COA. The process is shown on slide 27. Route optimization is possible after a reconnection.

Triangle routing is the condition in which the CN sends datagrams to the HA, which forwards them to the FA for the MN, but the MN sends datagrams directly to the CN. **Route optimization** is possible if the CN is enabled to run mobile IP. In this case,

1. The CN sends data to the MN in its home network.
2. The HA advises the CN that the MN is mobile.
3. The CN requests a binding so that it can send datagrams directly to the MN at the COA.
4. The HA provides the COA to the CN.
5. The CN sends datagrams to the MN directly using the COA.

Slide 30 shows the process of roaming under route optimization.

Slides 31 – 32 IPv6 Mobile IP

With the growth in wireless mobile networking, mobile IP became a priority for IPv6.

Using the addressing features of IPv6, the mobile node (MN) can always form a **link local unicast address** from its MAC address and a **roaming address** (from the global prefix and MAC address) using the stateless auto-configuration procedure. No special Foreign Agent (FA) is required because the MN uses its roaming address as a care-of-address (COA).

The MN registers with its Home Agent (HA) by sending a binding update containing its roaming address. The HA forwards traffic for the registered MN by tunneling packets from the CN to the MN. The MN sends packets to CN directly and route optimization is achieved by the HA providing the CN with the COA.

IPv6 provides further support for mobility using option headers:

- Any node can write the home address for a MN in the **destination option header**. By examining the option header, the destination MN can identify datagrams by its home address.
- Tunneling can be accomplished by using the IPv6 routing extension headers instead of encapsulation, which reduces processing time.
- When an MN moves from one ESS to another in a WiFi network, there is a layer 2 handover (a change in the attached access point and associated ESS ID). A router node can detect this change in the subnet prefix and automatically updates the CoA.
- IPv6 Mobility header messages are provided for Home Test Init, Home Test, Care-of Test Init, and Care-of Test, binding update and acknowledgement.