CS-435

spring semester 2020

Network Technology & Programming Laboratory

University of Crete
Computer Science Department

Stefanos Papadakis

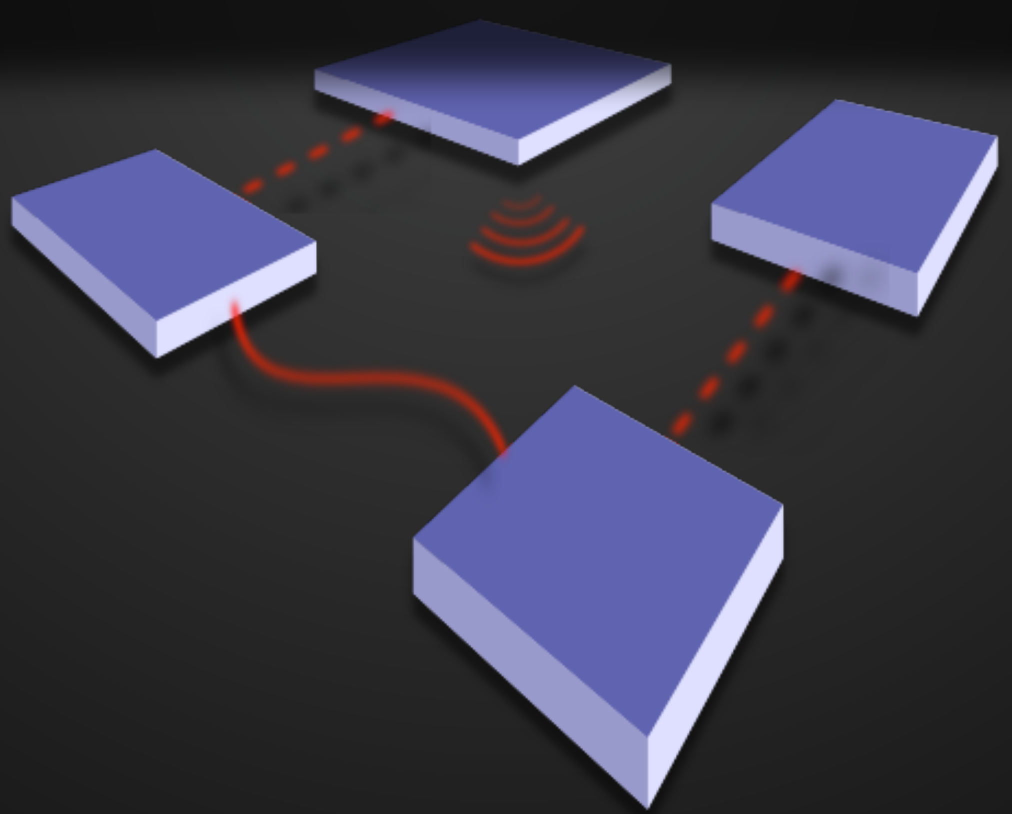

CS-435

Lecture #4 preview

- ICMP
- ARP
- DHCP
- NAT
- IPv6
  - Header
  - Addressing

ICMP

Internet Control Message Protocol

- core IP protocol
- ICMPv4 (RFC792) / ICMPv6 (RFC4443)
- used for signaling error messages
- 8 bytes header
- encapsulated in a single IP datagram
- used by:
  - ping
  - traceroute

The ARP

Address Resolution Protocol

- MAC address is **required** to send to LAN host
- How to find it:
  - manually
  - included in network address
  - use central directory
  - use address resolution protocol
- ARP (RFC 826) provides **dynamic** IP to ethernet address mapping
  - source **broadcasts** ARP request
  - destination replies with ARP response

DHCP

Dynamic Host Configuration Protocol

- “Plug & play”
- client-server protocol
- Requires at least one trusted server in the infrastructure
- Four step process

DHCP

- DHCP_Discover
  - Src 0.0.0.0
  - Dst 255.255.255.255
  - Must be forwarded across broadcast domains
- DHCP_Offer
  - Provides the parameters in unicast
- DHCP_Request
  - Src 0.0.0.0
  - Dst 255.255.255.255
- DHCP_ACK

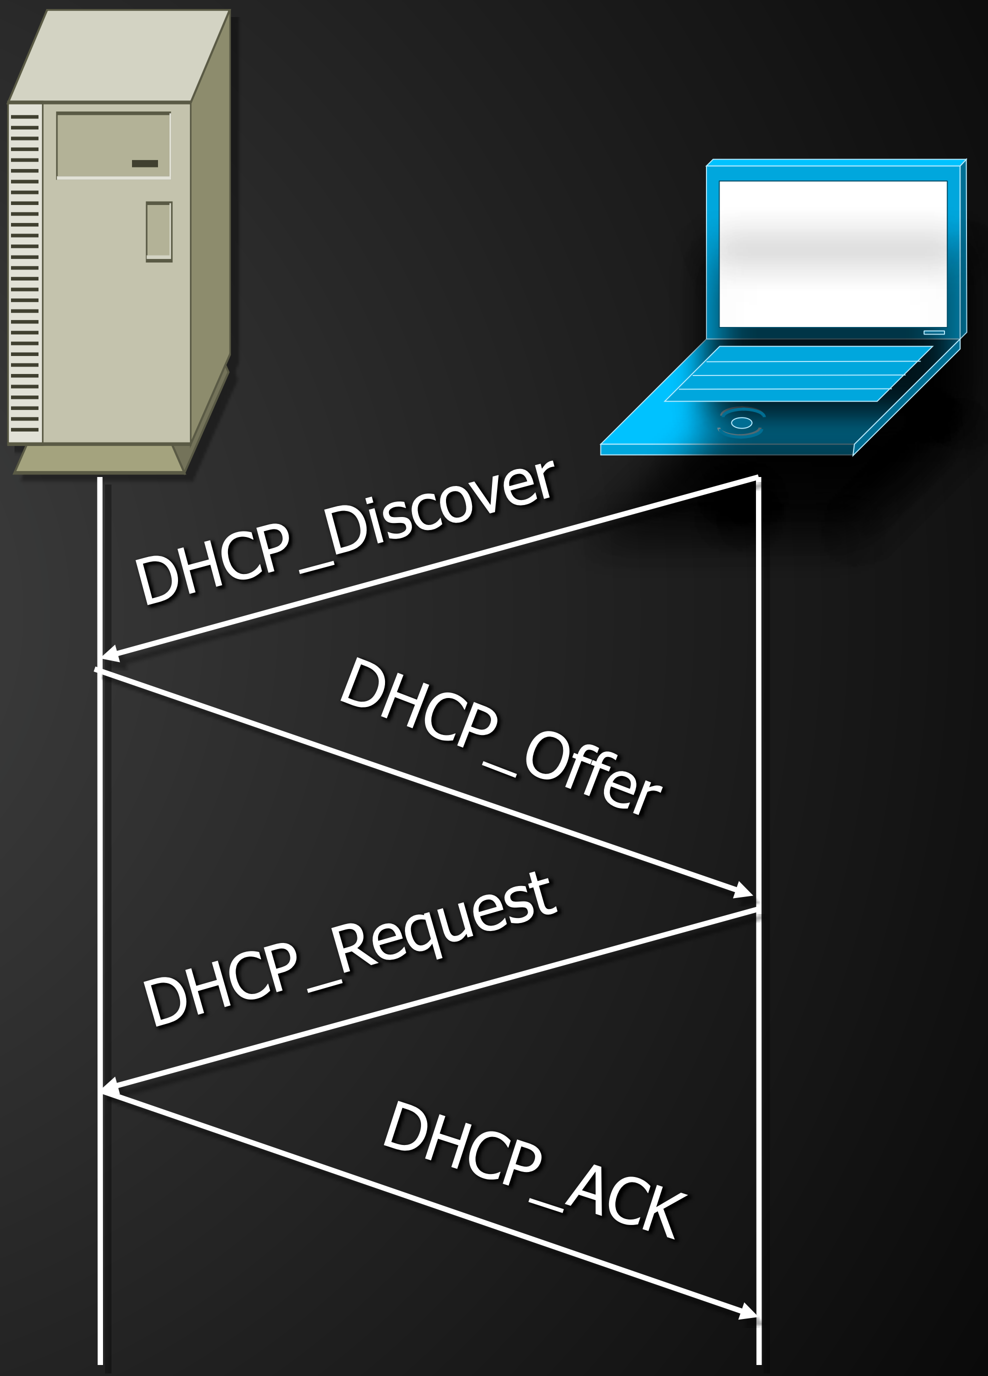

NAT

Network Address Translation

- IP address management within organizations should be **easy** & **safe**
  - Flexible w.r.t. growing number of machines
  - Not encumbered by “global” addressing problems
- Solution: NAT

NAT

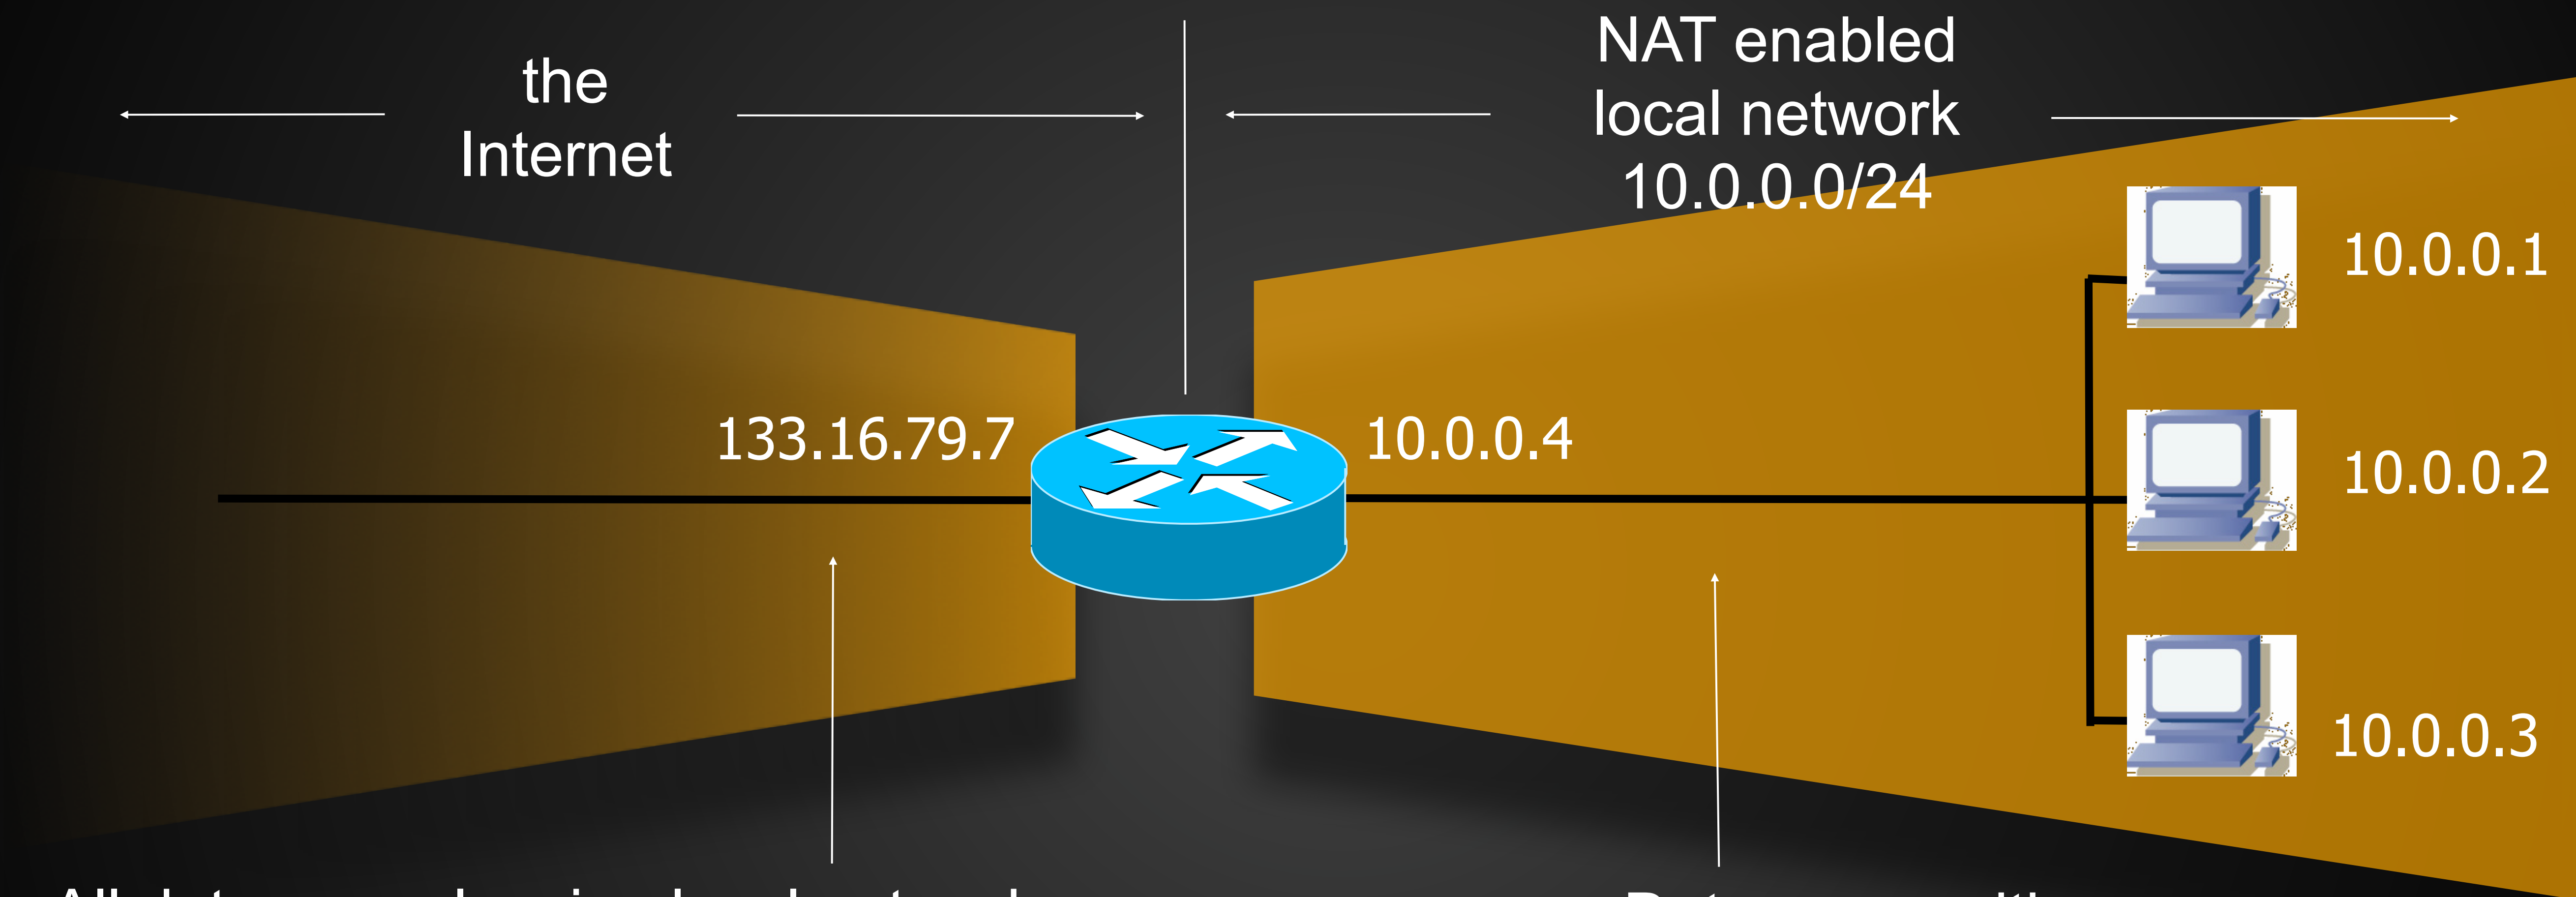

All datagrams leaving local network appear to have the same single source NAT IP address: 133.16.79.7, only with different **source port** numbers

Datagrams with source or destination in this network have 10.0.0.0/24 address for source, destination

NAT

Local network uses just **one** IP address as far as **outside** word is concerned:

- no need to be allocated range of addresses from ISP:

  just one IP address is used for all devices

- can change addresses of devices in local network without notifying outside world
- can change ISP without changing addresses of devices in local network
- devices inside local network not explicitly addressable, visible by outside world (security ++).

NAT

NAT router must:

- For the **outgoing** datagrams:
    - replace (source IP address, port #) of every outgoing datagram to (NAT IP address, new port #)
    - . . . remote clients/servers will respond using (NAT IP address, new port #) as destination address
- Keep for lookup a **NAT translation table** with every
    - (source IP address, port #) ⇔ (NAT IP address, new port #)
- For the **incoming** datagrams:
    - replace (NAT IP address, new port #) in dest fields of every incoming datagram with corresponding (source IP address, port #) stored in NAT table

IPv6

- The Internet Protocol (IP) v4 has been the foundation of the Internet and virtually all multivendor private internetworks.
- This protocol is reaching the end of its useful life and a new protocol, known as IPv6 (IP version 6, originally: IPng), has been defined to ultimately replace IP
- The Internet Engineering Task Force (IETF) issued a call for proposals for a next generation IP (IPng) in July 1992.
- By 1994 the final design for IPv6 had emerged.

Why Change IP versions?

- Address space exhaustion
  - **two** level addressing (**network** & **host**) is wasteful
  - network addresses used even if not connected
  - growth of networks and the Internet
  - extended use of TCP/IP
  - single/multi addresses per interface
- requirements for new types of service

Why Change IP versions?

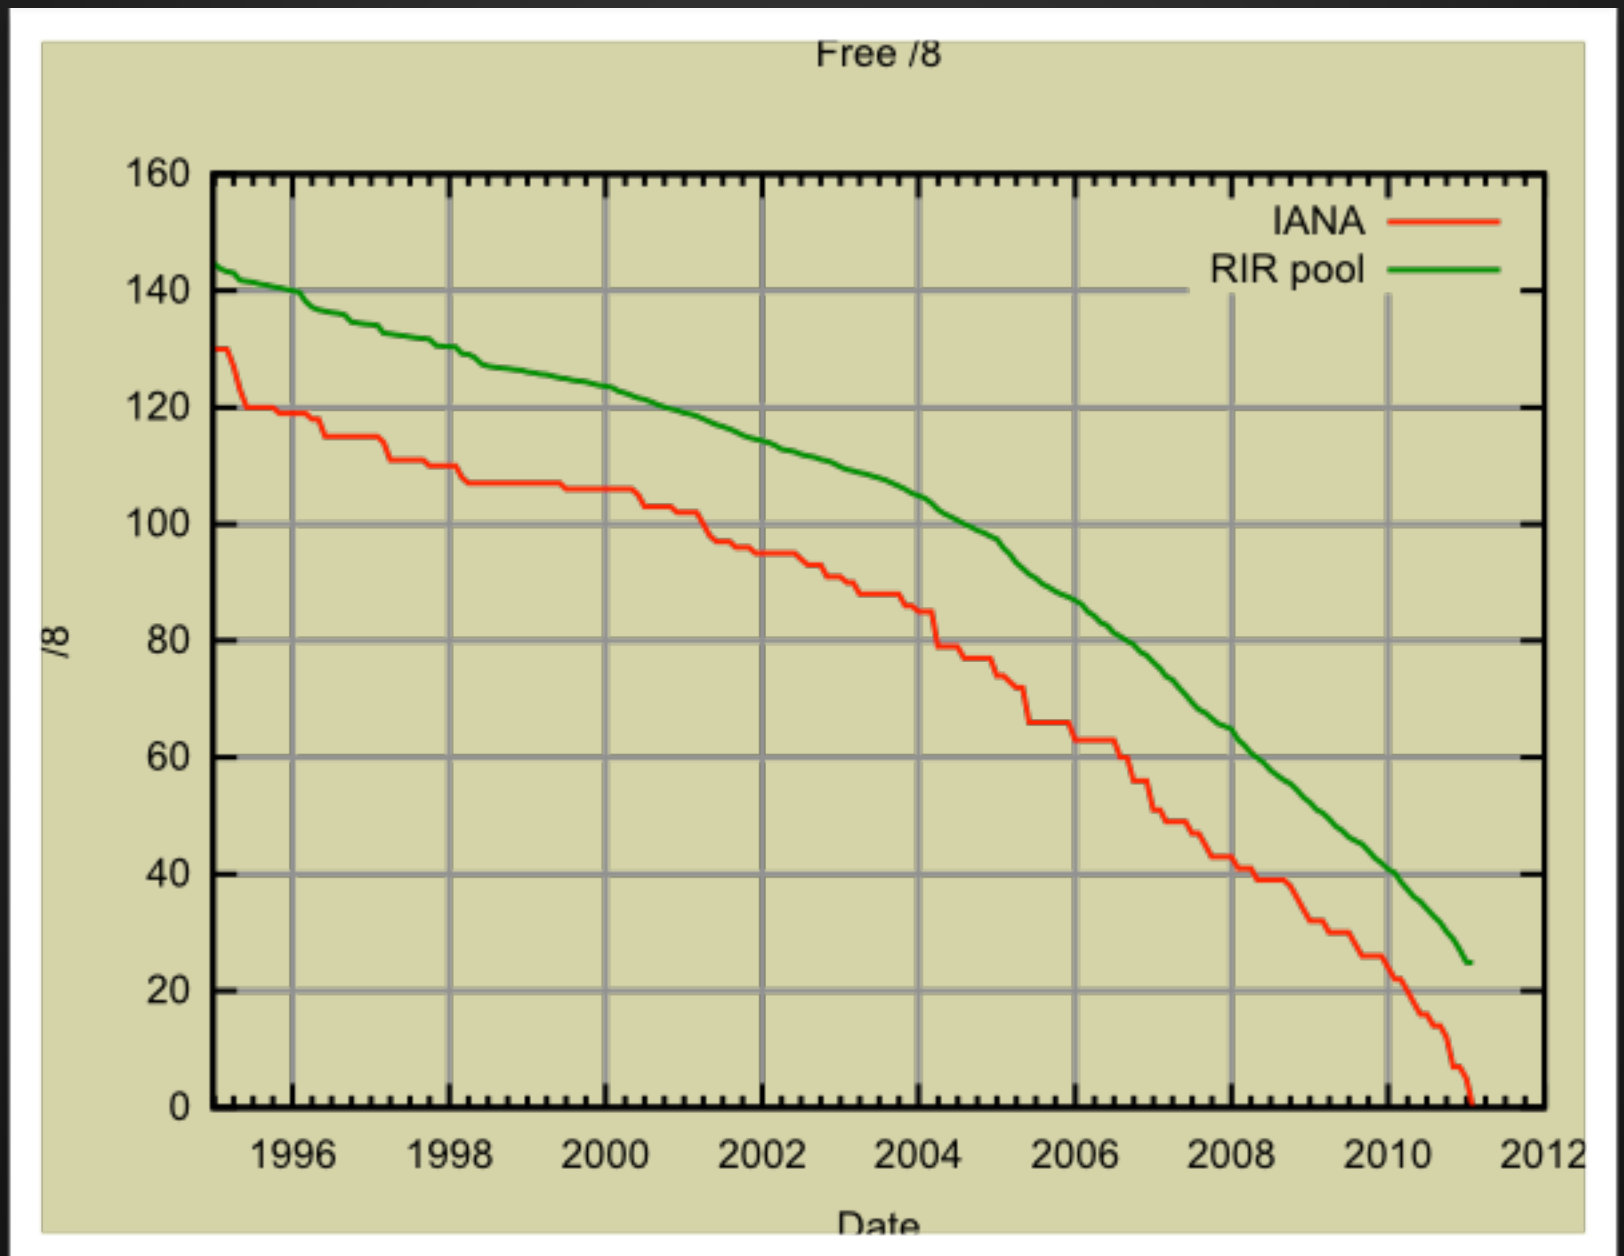

IPv6 Enhancements

- expanded 128 bit address space (vs 32bit of IPv4)
- improved option mechanism
  - most not examined by intermediate routes
- dynamic address assignment
- increased addressing flexibility
  - anycast & multicast (NO broadcast)
- support for resource allocation
  - labeled packet flows (**QoS**)
- integrates security (**IPsec**)

IPv6 packet structure

- <u>Hop-by-Hop Options header</u>: Defines special options that require hop-by-hop processing.
- <u>Routing header</u>: Provides extended routing, similar to IPv4 source routing.
- <u>Fragment header</u>: Contains fragmentation and reassembly information.
- <u>Authentication header</u>: Provides packet integrity and authentication.
- <u>Encapsulating Security Payload header</u>: Provides privacy.
- <u>Destination Options header</u>: Contains optional information to be examined by the destination node.

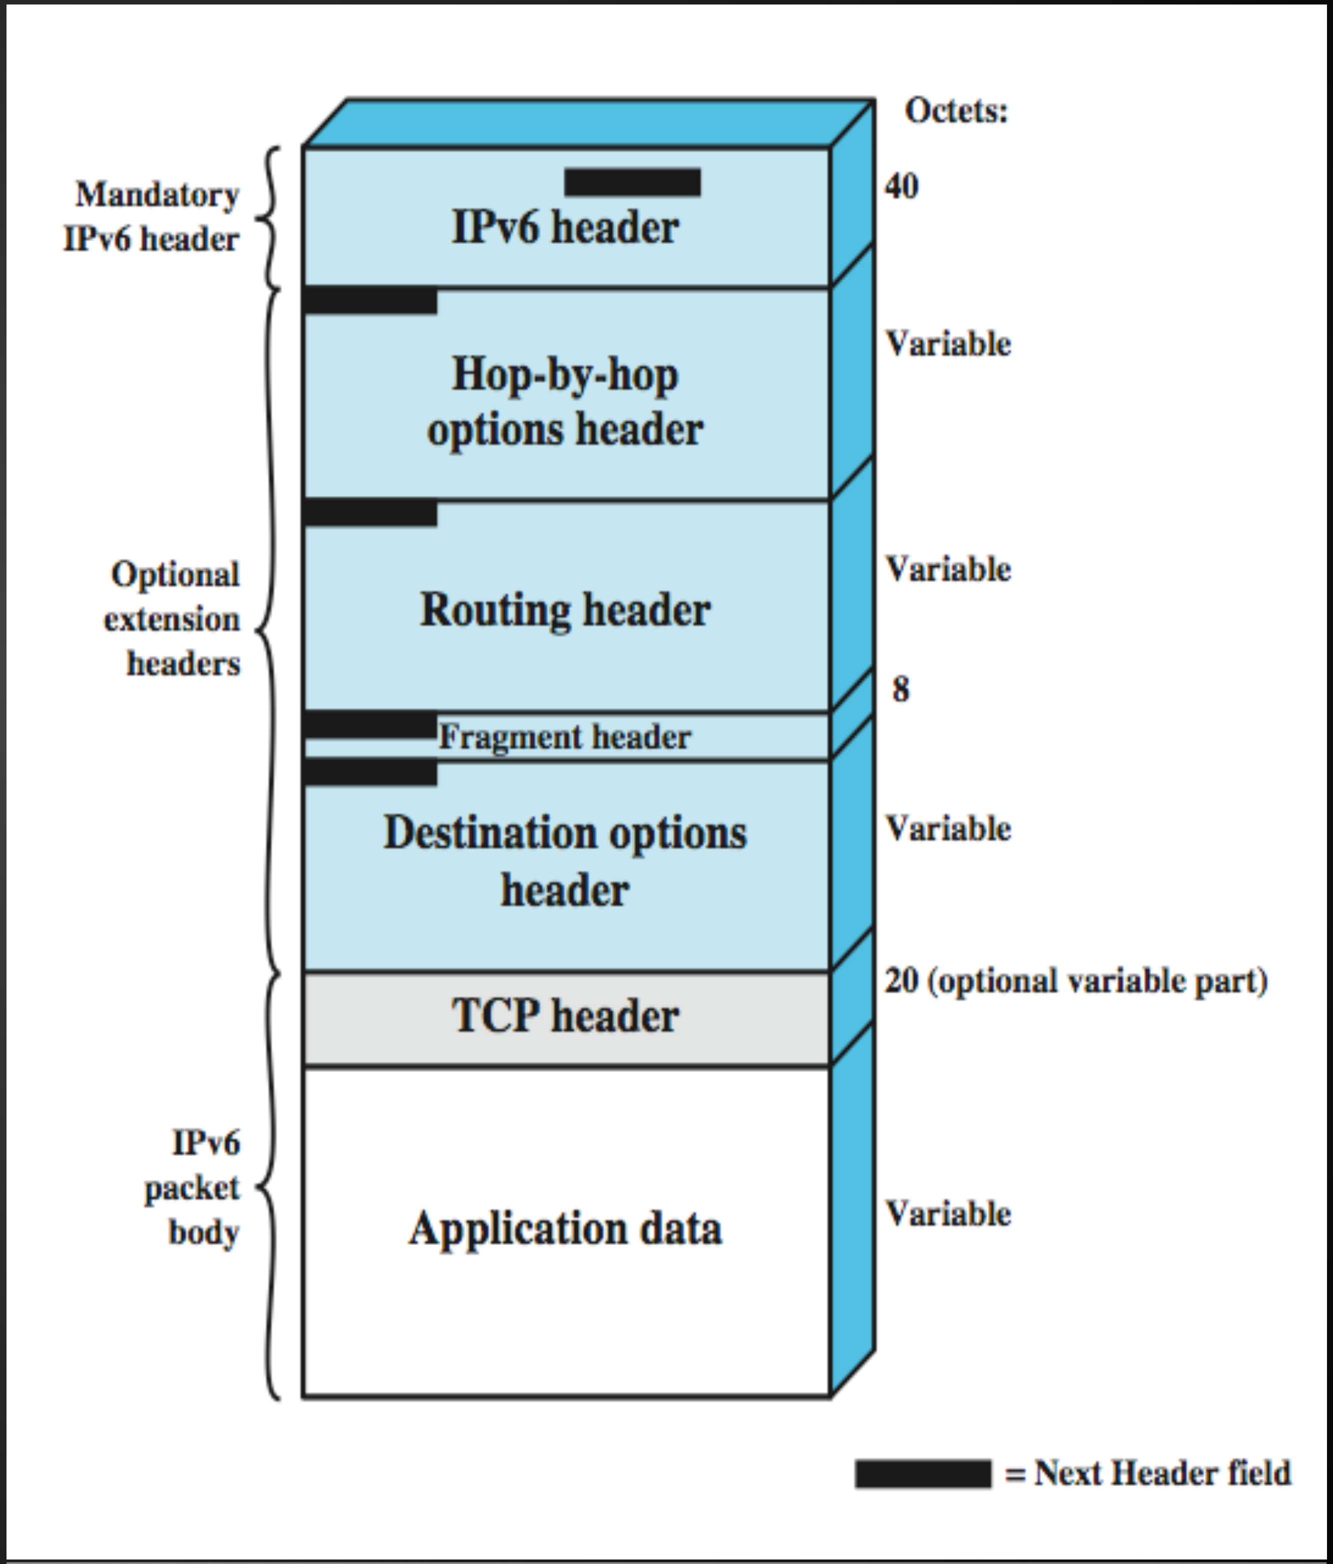

IPv4 vs IPv6 header

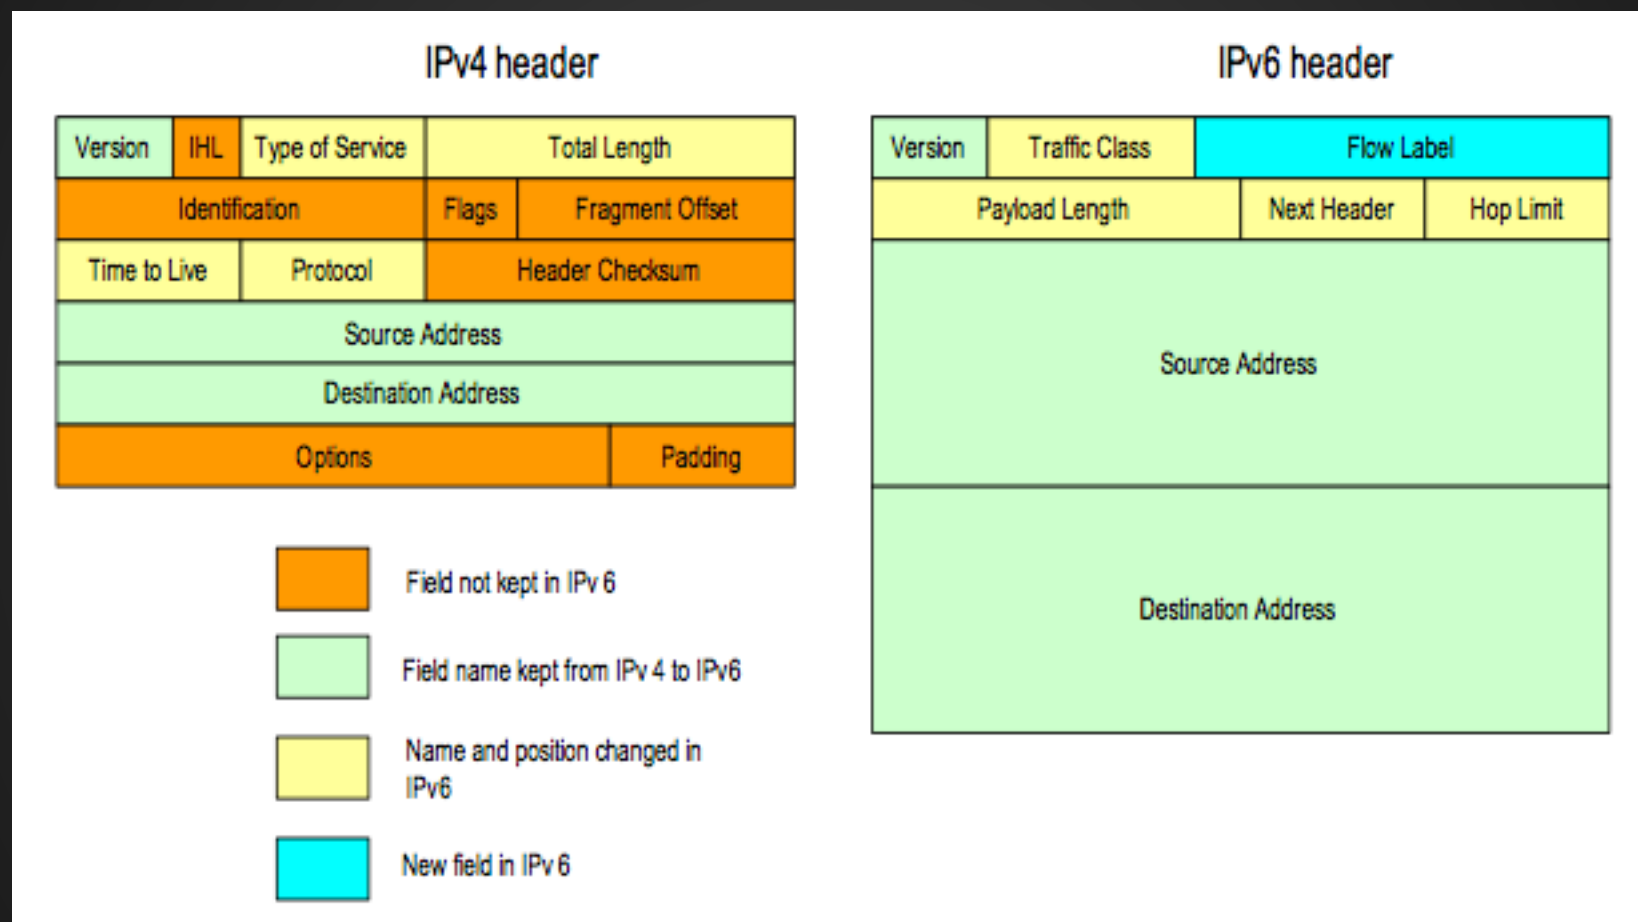

IPv6 header

- Version (4 bits): 6.
- DS/ECN (8 bits): used by originating nodes and/or forwarding routers for differentiated services and congestion functions, see IPv4 DS/ECN field.
- Flow Label (20 bits): used by a host to label those packets for which it is requesting special handling by routers within a network.
- Payload Length (16 bits): Length of the remainder of the IPv6 packet following the header, in octets: this is the total length of all of the extension headers plus the transport-level segment.
- Next Header (8 bits): Identifies the type of header immediately following the IPv6 header; this will either be an IPv6 extension header or a higher-layer header, such as TCP or UDP.

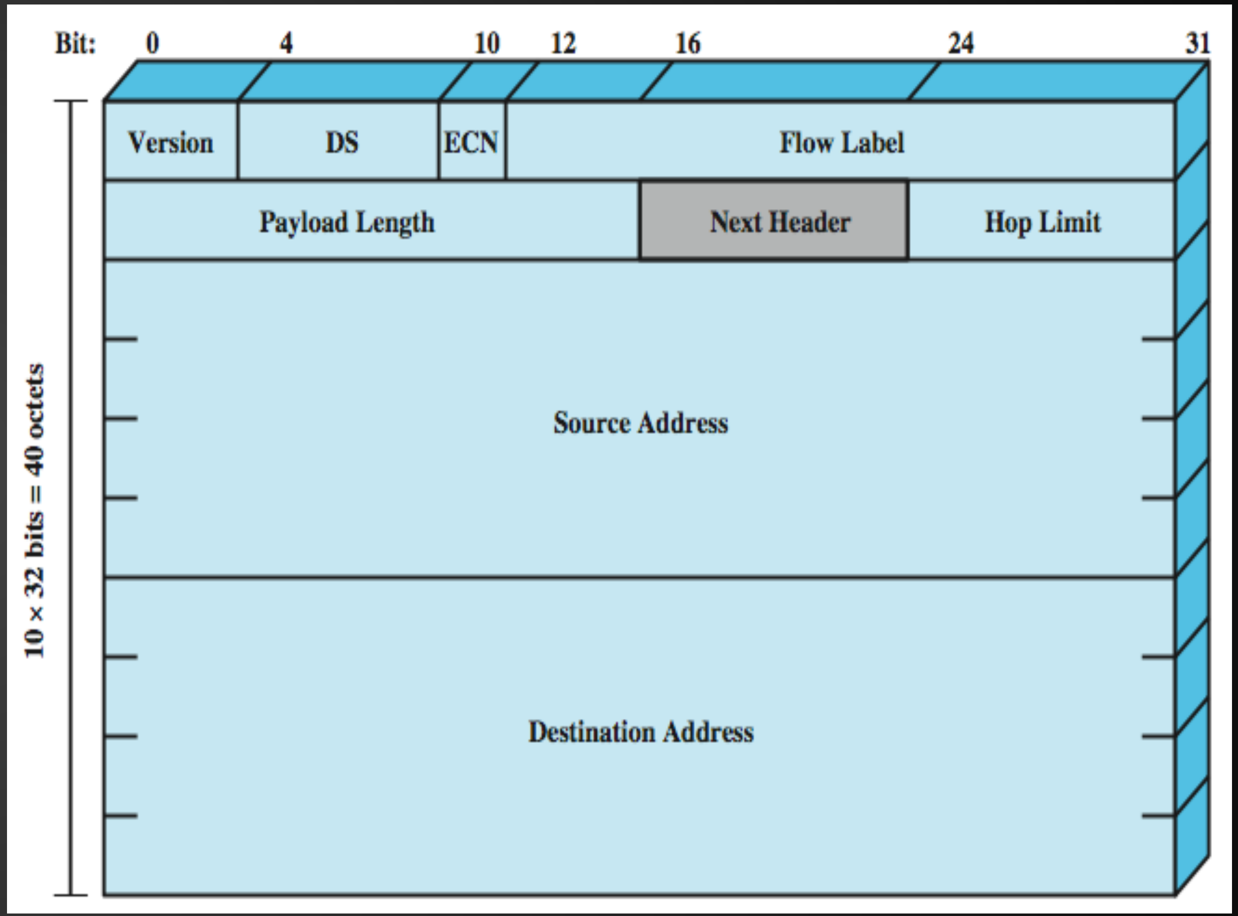

IPv6 header

- <u>Hop Limit (8 bits)</u>: The remaining number of allowable hops for this packet. The hop limit is set to some desired maximum value by the source and decremented by 1 by each node that forwards the packet. The packet is discarded if Hop Limit is decremented to zero.
- <u>Source Address (128 bits)</u>: address of originator of the packet.
- <u>Destination Address (128 bits)</u>: address of intended recipient of the packet.

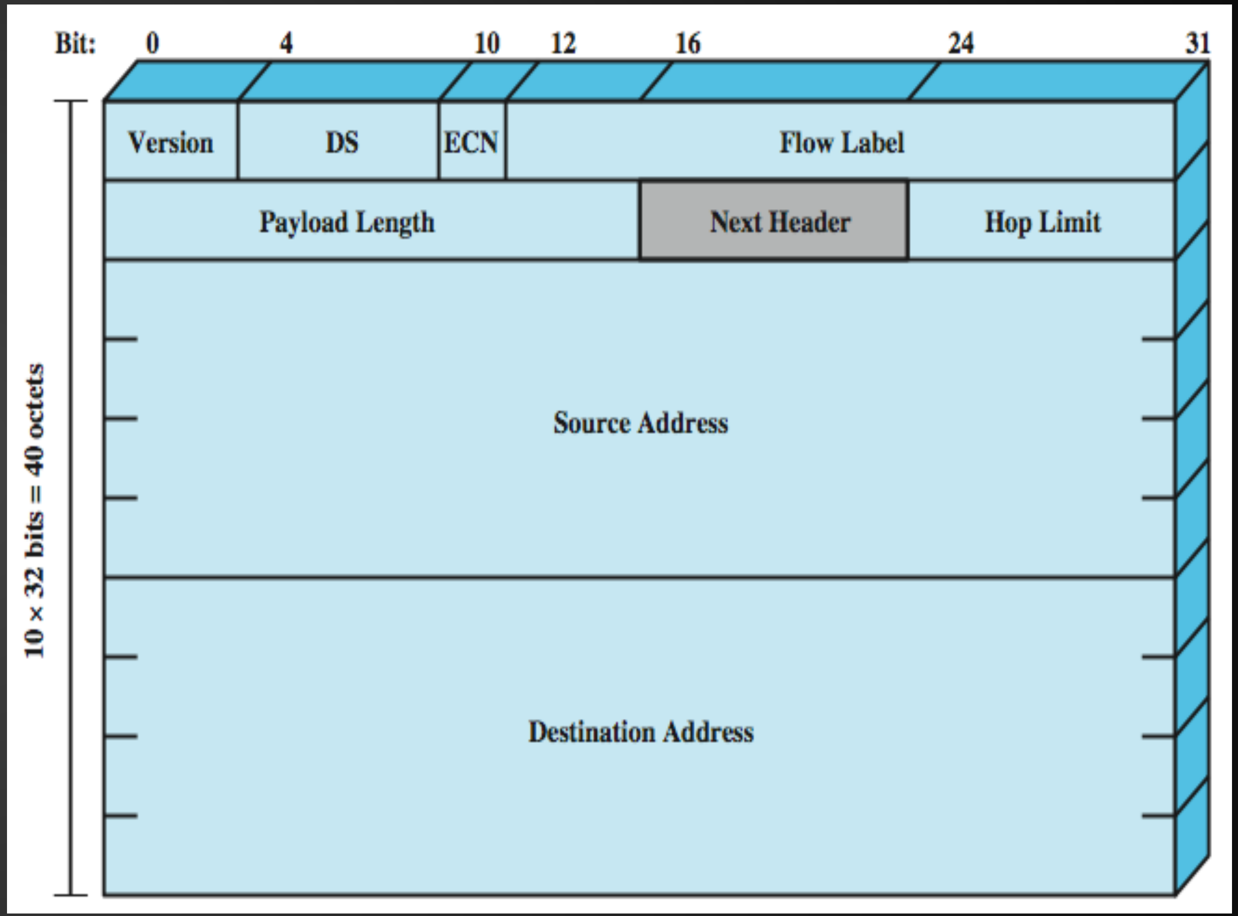

The flow label

- A flow is uniquely identified by the combination of a source address, destination address, and a nonzero 20-bit flow label.
- All packets that are to be part of the same flow are assigned the same flow label by the source.
- A “flow” may be a single TCP connection or even multiple TCP connections.

The flow label

- From the router's point of view:
  - a flow is a sequence of packets that share attributes that affect how these packets are handled by the router. These include:
    - path,
    - resource allocation,
    - discard requirements,
    - accounting, and
    - security attributes.
- The router may treat packets from different flows differently in a number of ways, including:
  - allocating different buffer sizes,
  - giving different precedence in terms of forwarding, and requesting different quality of service from networks.

The flow label

- In principle, all of a user's requirements for a particular flow **could be defined in an extension header** and included with each packet.
- alternative, **for IPv6**, using the flow label, in which the flow requirements are defined prior to flow commencement and a unique flow label is assigned to the flow.
- the router must save flow requirement information about each flow.

IPv6 Addressing

- Like IPv4…
  - Unicast
    - An identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address.
  - 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.
  - 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 protocols' measure of distance).
- Specified in the IPv6 address architecture RFC.

IPv6 Addressing

- Broadcast?
  - There is no broadcast in IPv6.
  - This functionality is taken over by multicast.
- A consequence of this is that the “highest” address may be used freely

IPv6 Addressing

- IPv6 addresses of all types are assigned to interfaces, not nodes.
  - An IPv6 unicast address refers to a single interface.
  - As each interface belongs to a single node, any of that node's interfaces' unicast addresses may be used as an identifier for the node.
- The same interface identifier may be used on multiple interfaces on a single node

IPv6 Addressing

- All addresses are 128 bits
- Write as sequence of eight sets of four hex digits (16 bits each) separated by colons
- Leading zeros in group may be omitted
  - Contiguous all-zero groups may be replaced by "::"
  - Only one such group can be replaced

IPv6 Addressing

- Example:
  - 3ffe:3700:0200:00ff:0000:0000:0000:0001
- This can be written as
  - 3ffe:3700:200:ff:0:0:0:1 or
  - 3ffe:3700:200:ff::1
- All three reduction methods are used here.

The unicast addressing

- Unspecified address
  - All zeros (::)
  - Used as source address during initialization
  - Also used in representing default
- Loopback address
  - Low-order one bit (::1)
  - Same as 127.0.0.1 in IPv4

The unicast addressing

- Link-local address (FE80::/10)
  - Unique on a subnet
  - Auto configured
  - Routers must not forward any packets with link-local source or destination addresses.

The unicast addressing

- Site-local address
  - Unique to a “site”
  - Used when a network is isolated and no global address is available.

The unicast addressing

- Stateless Address Auto-Configuration (SLAAC)

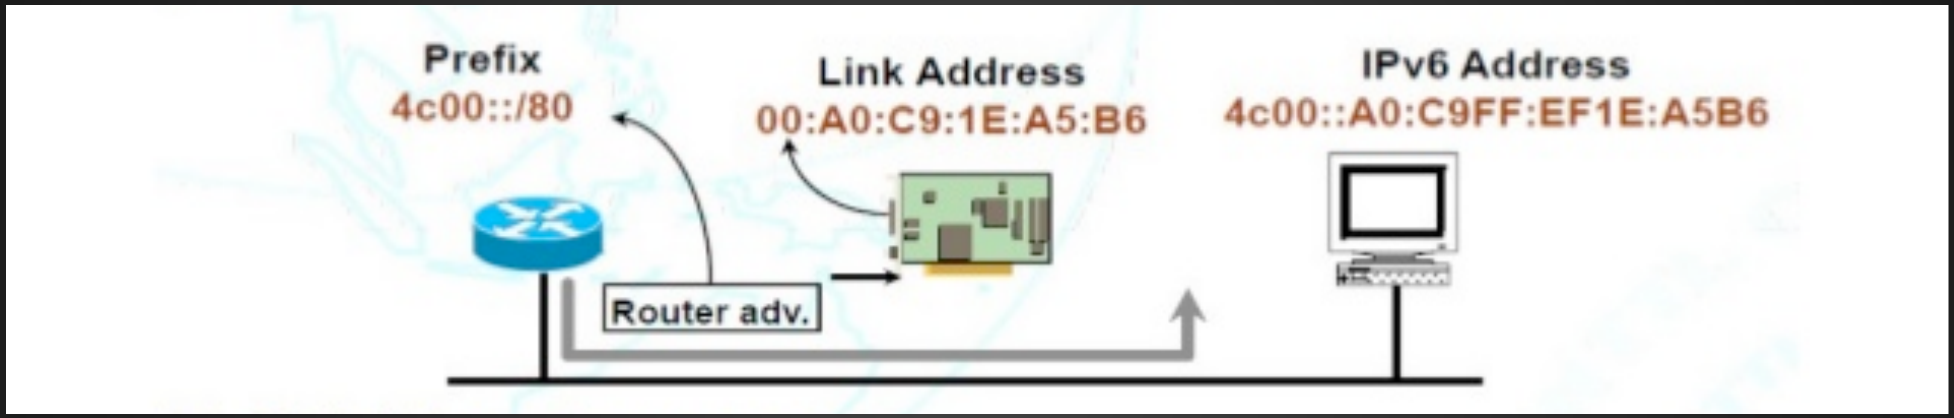

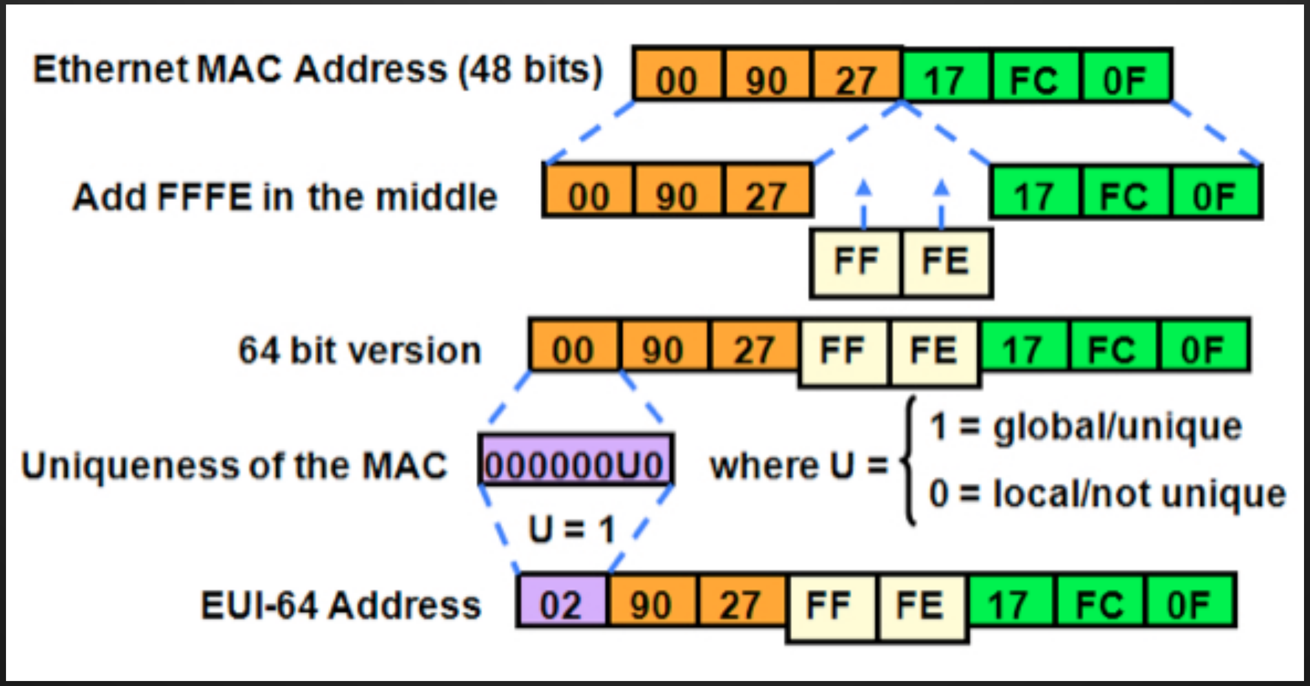

The unicast addressing

- Mapped IPv4 addresses
  - Of form ::FFFF:a.b.c.d
  - Used by dual-stack machines to communicate over IPv4 using IPv6 addressing
- Compatible IPv4 addresses (deprecated)
  - Of form ::a.b.c.d
  - Used by IPv6 hosts to communicate over automatic tunnels