



POLYTECH[®]
PARIS-SUD

IPv6

Réseaux Avancés

IPv6

(Internet Protocol version 6)

IPv6

- **IP version 6 is the new version of the Internet Protocol (IP)**
- **The standardization process started in the '90s**
- **The main elements of IPv4 are still there, but there are some important changes**
- **... and, more important, the IP address length is now 128 bits, instead of only 32**

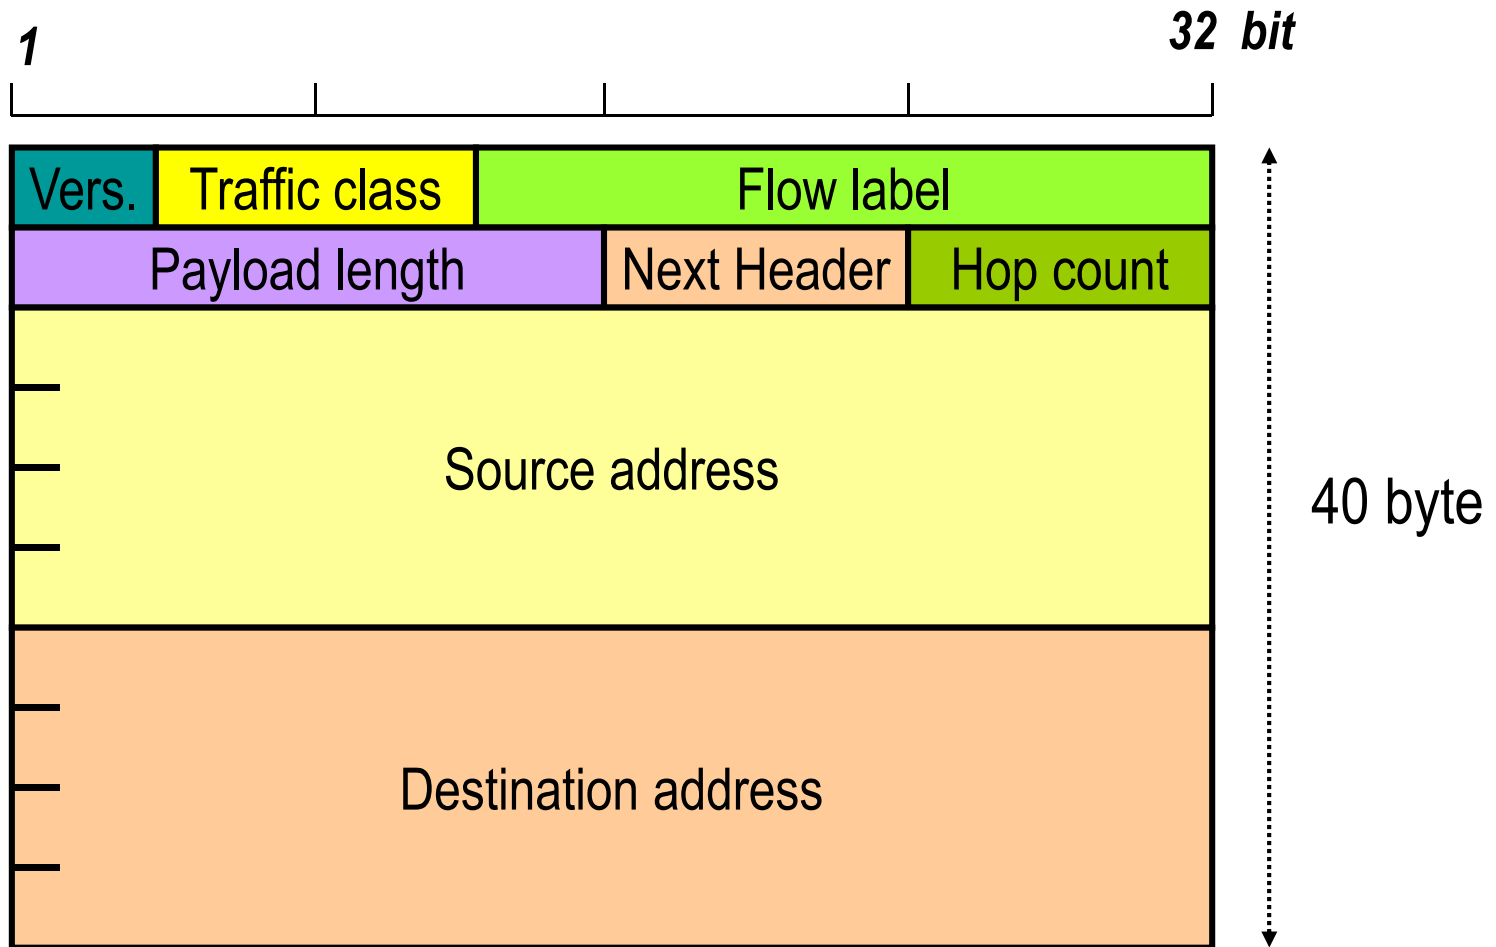
Motivations to introduce IPv6

- **IPv4 address space exhaustion**
- **Introduction of supports to Real-Time applications**
- **Supports for authentication (necessary, for example, for E-Commerce applications)**

IPv6: the main novelties

- **IPv6**
 - Address length, options management, fragmentation management, introduction of flow identifiers, traffic classes, no header checksum, etc.
- **ICMPv6:**
 - New version of ICMP with enhanced functionalities
- **ARP:**
 - Eliminated. At his place we have ICMPv6 Address Resolution
- **DHCPv6**
 - Modified for the new protocol (some functionalities are performed by ICMPv6)
- **Routing**
 - RIPng and OSPFv6

Base Header IPv6

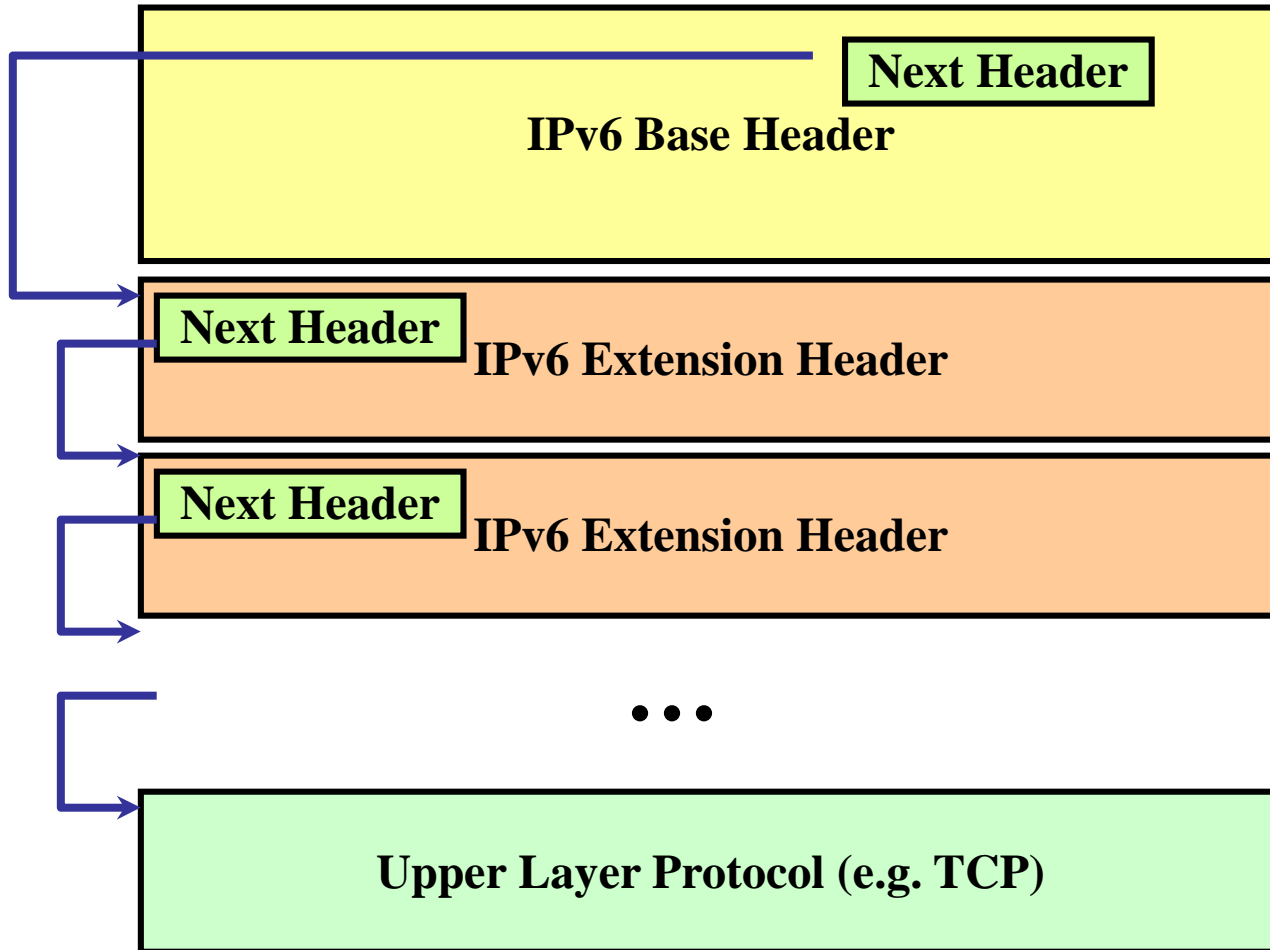


- The Base Header length is **FIXED** (**40 bytes**)

IPv6 Header

Field	Length (bits)	Description
Version	4	Protocol Version (6)
Traffic Class	8	Used to identify different types of traffics in Differentiated Services (DS) networks
Flow Label	20	Used to identify a flow of packets (like <i>labels</i> in MPLS)
Payload Length	16	Payload length in bytes (except headers, Base+eventually extension headers)
Next Header	8	Identifies the type of header that follows the Base header
Hop Limit	8	Same function as Time To Live in IPv4
Source Address	128	Source IPv6 address
Destination Address	128	Destination IPv6 address

Next Header



IPv6 Extension Headers

- **Routing Header:**
 - IPv6 can implement, exactly like IPv4, a loose source routing
 - Differently from IPv4, where we had an option for this, in IPv6 there is an Extension Header for this functionality
 - The Routing Header is used by the IPv6 source node (the packet sender) to specify one or more intermediate nodes (routers) that must be traversed by a packet in its way towards the destination
 - Very similar to the Loose Source and Record Route options in IPv4
 - The Routing header is identified with a Next Header value of 43

IPv6 Extension Headers

■ Fragment Extension Header:

- It is used to manage fragmentation. In IPv6 only the sender can fragment a packet, differently from IPv4 where *any router* along the path can fragment a packet
- This permits to reduce the overhead caused by such operation in routers
- To know if fragmentation is needed or not, the sender must know the MTU of the whole path (this is obtained using the MTU Path discovery message of ICMPv6)
- ➔ ■ When a route changes, the MTU can diminish, so that fragmentation would still be needed
- In such case, ICMPv6 has been extended: the router can tell the packet sender (using an appropriate ICMPv6 message) that fragmentation is needed.
- IPv6 standard does not encourage the utilisation of fragmentation. It encourages applications to send packets with length inferior to the path MTU
- The Fragment Header is identified with a Next Header value of 44 10

IPv6 Extension Headers

- **Authentication Header**
 - To authenticate the sender
- **Encapsulating Security Payload (ESP)**
 - To cypher the payload (another IP packet or upper layer TCP/UDP messages)

IPv6 Addresses

Synthetic Notation:

- The classic IPv4 notation (dotted decimal) is not efficient:

104.230.140.100.255.255.255.255.0.0.17.128.150.10.255.255

- The Colon Hexadecimal Notation is introduced:

68E6:8C64:FFFF:FFFF:0:1180:96A:FFFF

- Consecutive Zeros can be omitted (only one time in the address!):

FF05:0:0:0:0:0:0:B3 → FF05::B3

- Special notation to support transition from IPv4

0:0:0:0:0:0:128.10.2.1 → ::128.10.2.1



last 32 bits

Using Synthetic Notation

- The following addresses...

3080:0:0:0:8:800:200C:417A
FF01:0:0:0:0:0:0:43
0:0:0:0:0:0:0:1
0:0:0:0:0:0:0:0

unicast-address
multicast-address
loopback-address
unspecified-address

- ...can be expressed as follows:

3080::8:800:200C:4170
FF01::43
::1
::

unicast-address
multicast-address
loopback-address
unspecified-address

IPv6 Addreses

- **Addressing space in IPv6:**

✓ $2^{128}=3.4*10^{38}$

- **Number of addresses for square meter on the Earth:**

$7*10^{23}$ (greater than Avogadro number!!!)

- **If we assign 1 million addresses for each nanosecond**

✓ **We need 10^{16} years to assign them all**

IPv6 Address Types

- IPv6 has a rich set of addresses. Each interface normally has more than 1 address assigned
- Address types:
 - Unicast (*only 1 destination*)
 - Multicast (*all those who belong to a group*)
 - Anycast
 - ✓ Like multicast addresses, they identify a group of nodes
 - ✓ Differently from multicast, a packet destined to an anycast address will be delivered to the node (belonging to the anycast group) CLOSER to the source node (based on routing metrics)
 - ✓ In this way it is possible to identify the closer router, the closer DNS server ...

Address Space in IPv6

Reserved for:
 ✓ Unspecified
 ✓ Loopback
 ✓ IPv6 addresses compatible with IPv4

Unicast addresses for local use:
 ✓ Link Local: used for autoconfiguration and neighbor discovery
 ✓ Site Local: similar to private addresses in IPv4 (class 10.0.0.0)


prefix (binary)	usage	fraction
0000 0000	Reserved	1/256
0000 0001	Unassigned	1/256
0000 001	Reserved for NSAP Allocation	1/128
0000 010	Reserved for IPX Allocation	1/128
0000 011	Unassigned	1/128
0000 1	Unassigned	1/32
0001	Unassigned	1/16
001	Aggregatable Global Unicast addr.	1/8
010	Unassigned	1/8
011	Unassigned	1/8
100	Unassigned	1/8
101	Unassigned	1/8
110	Unassigned	1/8
1110	Unassigned	1/16
1111 0	Unassigned	1/32
1111 10	Unassigned	1/64
1111 110	Unassigned	1/128
1111 1110 0	Unassigned	1/512
1111 1110 10	Link local Unicast Addresses	1/1024
1111 1110 11	Site local Unicast addresses	1/1024
1111 1111	Multicast	1/256

Special Addresses (Reserved)

- **Unspecified address (0:0:0:0:0:0:0:0)**
 - Used as *source* address during bootstrap phase (the node does not have other addresses)
 - It CANNOT be used as destination address
- **Loopback address (0:0:0:0:0:0:0:1)**
 - Analogous to 127.x.y.z in IPv4
- **IPv4-compatible IPv6 address (::IPv4_addr) (deprecated)**
 - The host with such address has ALSO a valid IPv6 address
 - Used to permit communications between IPv6 hosts when an IPv4 network is traversed
 - Disabled in Windows by default. To enable, use the command `netsh interface ipv6 set state v4compat=enabled`
 - Format: 80 zeros + 16 zeros + 32 bit IPv4_addr
- **IPv4-mapped IPv6 address (::FFFF:IPv4_addr)**
 - The host has ONLY an IPv4 address
 - Used to make IPv6 hosts communicate with IPv4 hosts
 - Format: 80 zeros + 16 ones + 32 bit IPv4_addr

Aggregatable Global Unicast Address (RFC 3587)

- Global unicast address format
- It has a hierarchical structure, to permit aggregation and reduce scalability problems with routing tables
- Global Routing Prefix (organized hierarchically): it is assigned to a Site/Organization
- Subnet ID: identifies a Subnet inside the Site.
 - This field is used by organizations to define a hierarchy inside the AS, defining the subnets
 - 16 bits are available: up to 65536 subnets can be defined
- Interface ID
 - 64 bits with format derived from IEEE EUI-64 (EUI: Extended Unique Identifier)




Link-Local Unicast Address

- **FP = 1111 1110 10**
- **They are used ONLY to address machines on a single link (which means, for example, on a single LAN segment)**
- **In IPv6 each interface has at least one link-local unicast address**
 - **Which is normally defined using autoconfiguration, starting from the physical address (IEEE EUI-64) (EUI= Extended Unique Identifier)**
- **These addresses are fundamental in the Neighbor Discovery process**



Site-Local Unicast Address

- **FP = 1111 1110 11**
- **Also these addresses are destined to a local use**
- **They define a private address space (equivalent, for example, to class 10 in IPv4)**




- **Note that: if the prefix of a Site-Local address is replaced by the Global Routing Prefix, we obtain a global address very easily**

EUI-64 format

- The “Interface ID” field identifies an interface; hence it must be unique
- The EUI-64 identifier is a 64 bit number that identifies both the manufacturer and the “serial number” of the card (in a similar way to Ethernet MAC addresses)
 1. If a card already has an EUI-64 identifier, then it almost has its Interface ID.
 2. For Ethernet cards, there exists a procedure which starts from the MAC-Address (48 bits long) and permits to obtain the EUI-64 format

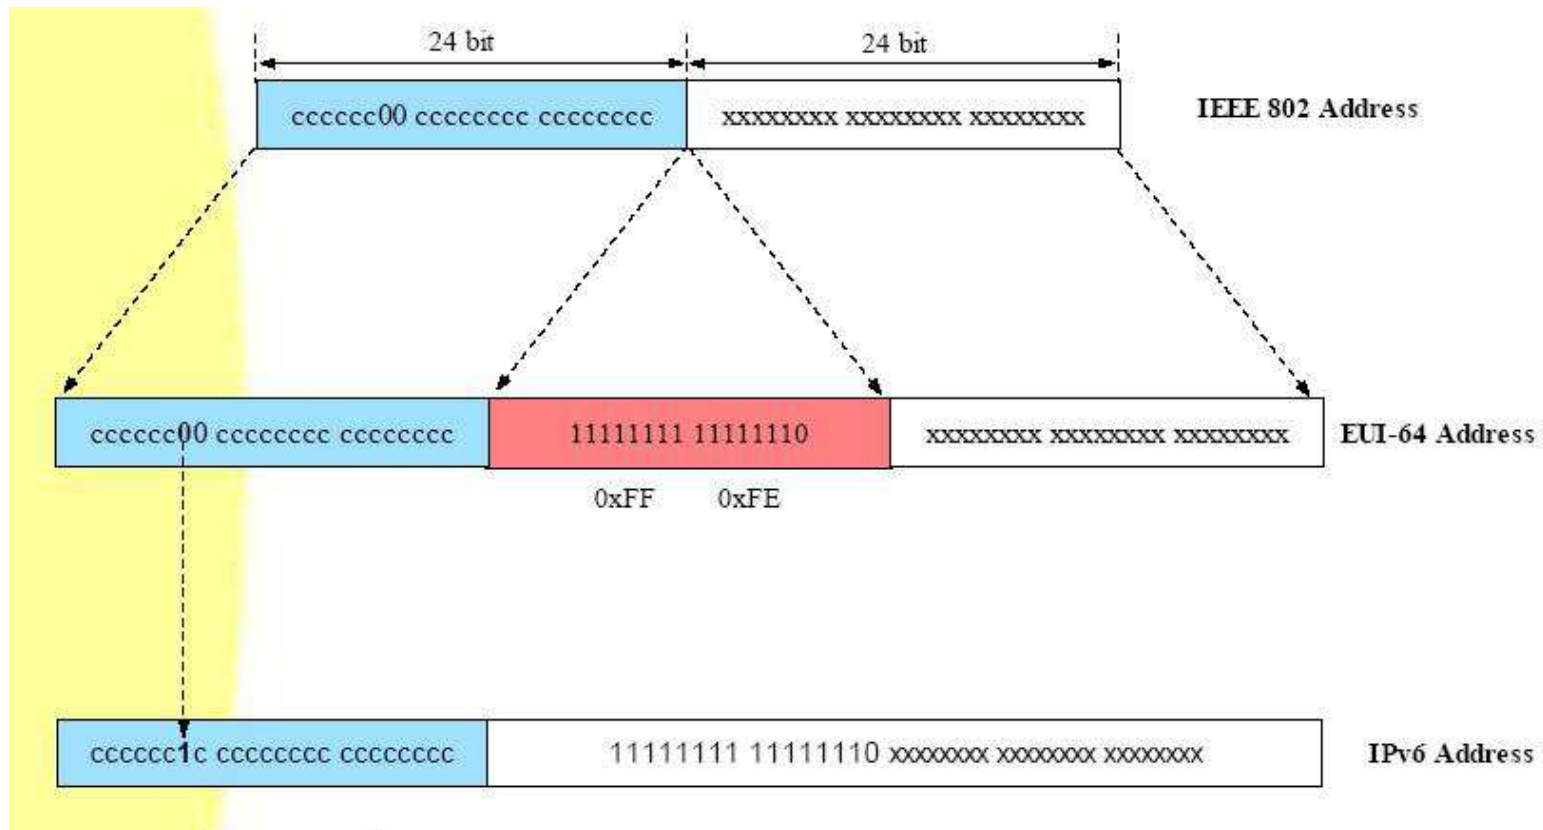
Building the Interface-ID (1)

- Very easy!
- The Interface-ID in IPv6 is obtained from the EUI-64 inverting the Universal/Local bit (the 7th bit)



Building the Interface-ID (2)

- Starting from a MAC address



Example

- **MAC Address:** **00-AA-00-3F-2A-1C**
- **EUI-64 Address:** **00-AA-00-FF-FE-3F-2A-1C**
- **Inverting U/L bit:** **02-AA-00-FF-FE-3F-2A-1C**
- **Using IPV6 notation:** **2AA:FF:FE3F:2A1C**

Multicast Address


- **FP = 1111 1111 (or equivalently, in hexadecimal: “FF”)**
- **Different sub-types of Multicast Addresses**
 - **Multicast global**
 - **Multicast link-local**
 - **Multicast site-local**
- **There exist several addresses reserved for special uses**



Multicast Address

- **Flags:**
 - **T=1 *non-permanent* (Transient) address**
 - **T=0 *permanent* address**
- **Scope (4 bits field):**
 - **0: reserved**
 - **1: interface-local scope** →
 - **2: link-local scope**
 - **5: site-local scope**
 - **8: organization-local scope**
 - **E: global scope**
 - **All others: unassigned**

Packets with this destination address may not be sent over any network link, but must remain within the current node. It is the multicast equivalent of the unicast loopback address.



Special Multicast Addresses

- **FF01::1 = all systems node-local scope**
- **FF02::1 = all systems link-local scope**
- **FF01::2 = all-routers node-local scope**
- **FF02::2 = all-routers on the local network segment**
- **FF05::2 = all-routers site-local scope**


All these addresses are used in a similar way as local broadcast addresses, distinguishing between all systems and all routers

Special Multicast Addresses

- Solicited-Node Multicast address
 - Each IPv6 node must have a “solicited-node multicast address” for each unicast or anycast address configured on the node itself
 - It is meant to facilitate efficient querying of network nodes during address resolution
 - Such address is built automatically concatenating the prefix

FF02::1:FF00:0/104

- ... with the last 24 bits of the corresponding unicast or anycast address



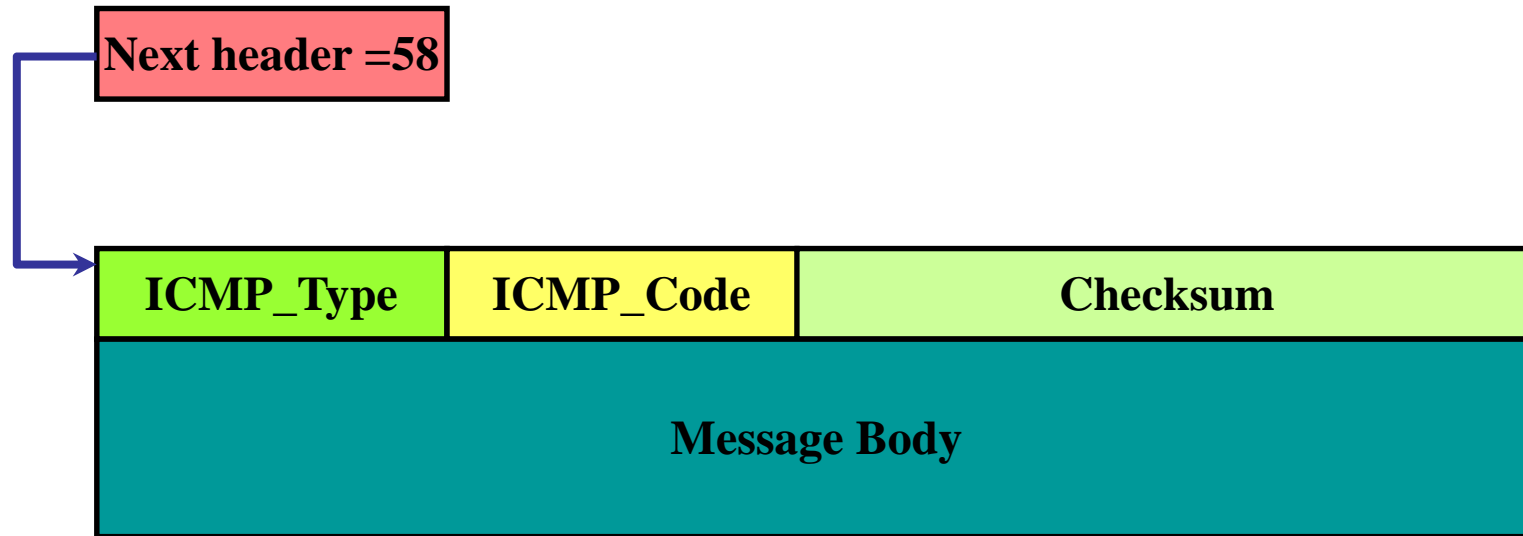
Several addresses for different functions

- **IPv6 has several self-configuration processes**
- **Normally, each node must:**
 - **Self-configure a link-local address starting from its physical address (Ethernet or EUI-64)**
 - **Self-configure a solicited-node multicast address for each of its unicast/anycast address**
 - **It can further self-configure other addresses using different procedures (ICMP/ DHCP)**

ICMP version 6

- **ICMP has a greater importance in IPv6 than in IPv4**
- **Some of its functions are:**
 - **Error reporting and network diagnostic**
 - **Address resolution**
 - **Self-configuration of IPv6 addresses**
 - **PATH-MTU computation (to manage fragmentation)**

ICMPv6: message structure



Some common types

- Type=1 – destination unreachable
- Type=2 – Packet too big
- Type=3 – Time exceeded
- Type=4 – Parameter problem,
- Type=128 – Echo request
- Type=129 – Echo reply

ICMPv6 Neighbor Discovery

- **There are several Neighbor Discovery procedures in IPv6**
 - **Address Resolution**
 - ✓ Analogous to ARP for IPv4
 - **Router Discovery**
 - ✓ To signal and discover the presence of routers on a network link
 - **Redirection**
 - ✓ Similar to redirection in IPv4
 - **Neighbor Unreachability Detection**

ICMPv6 Neighbor Discovery

- **Several special addresses are used (link-scope):**
 - **All-Systems Multicast Address (FF02::1)**
 - **All-Routers Multicast Address (FF02::2)**
 - **Solicited-node Multicast Address**
 - **Unicast Link-Local Address**
 - **Unspecified Address (0::0)**

- **5 new message types are introduced:**
 - **Router Solicitation message: type=133**
 - **Router Advertisement message: type=134**
 - **Neighbor Solicitation message: type=135**
 - **Neighbor Advertisement message: type=136**
 - **Redirect message: type=137**

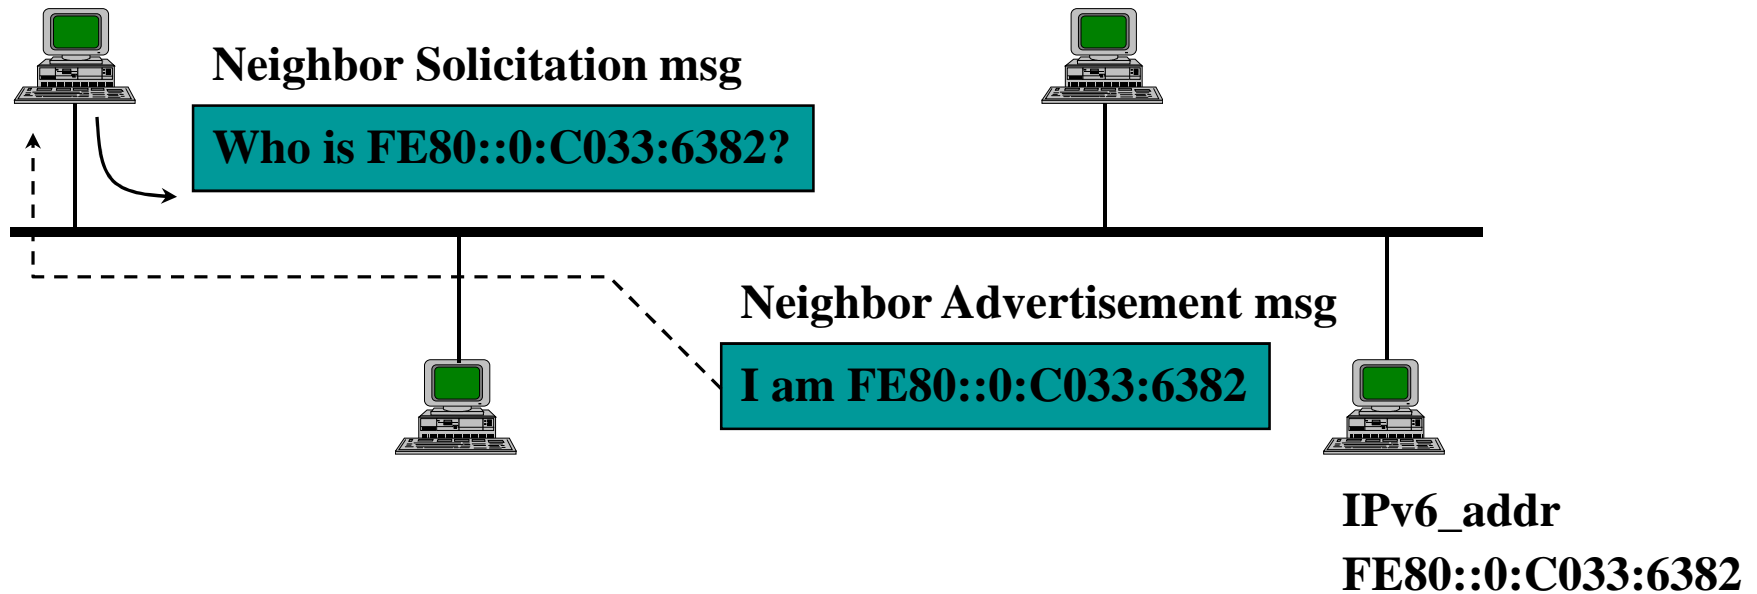
ICMPv6 Address Resolution

- Same function as ARP in IPv4
- Obviously, multicast/broadcast addresses are needed at the physical level
 - We suppose the existence of mapping between multicast IPv6 addresses and multicast/broadcast addresses at the link level
- “Neighbor Solicitation” and “Neighbor Advertisement” messages are used

ICMPv6 Address Resolution


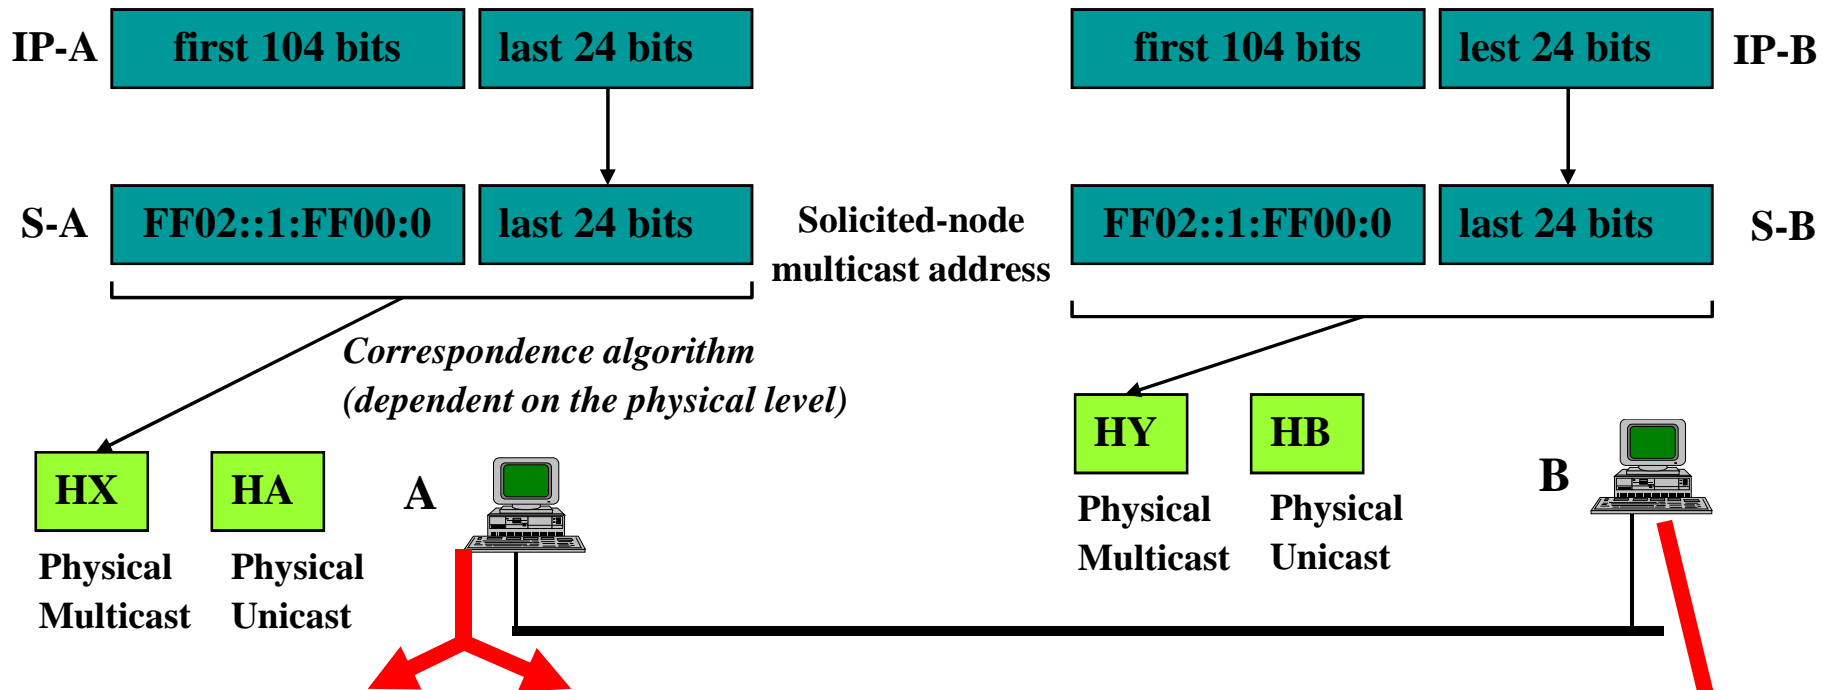
IPv6_addr

FE80::0800:2001:C782



- The Neighbor Solicitation message is sent at the Solicited-Node multicast address, which can be obtained also by the sender
- The Neighbor Advertisement message is sent to the IPv6 address of the requester

ICMPv6 Address Resolution



MTU Path Discovery

- **The sender must know what is the smallest MTU along the path**
- **It sends one packet with the same length as the MTU of the first link**
- **If the ICMP error message “Packet too big” comes back, then the MTU is reduced, and a new packet is sent out again ...**
- **... until no error message is received**

Transitioning from IPv4 to IPv6

- **It is based mainly on the following elements:**
 - **Dual stack:**
 - ✓ Systems with double stack, IPv4 and IPv6
 - **Tunneling:**
 - ✓ IPv4 networks are traversed using tunneling
 - **Header translation:**
 - ✓ Header translation between the two formats