IPv6

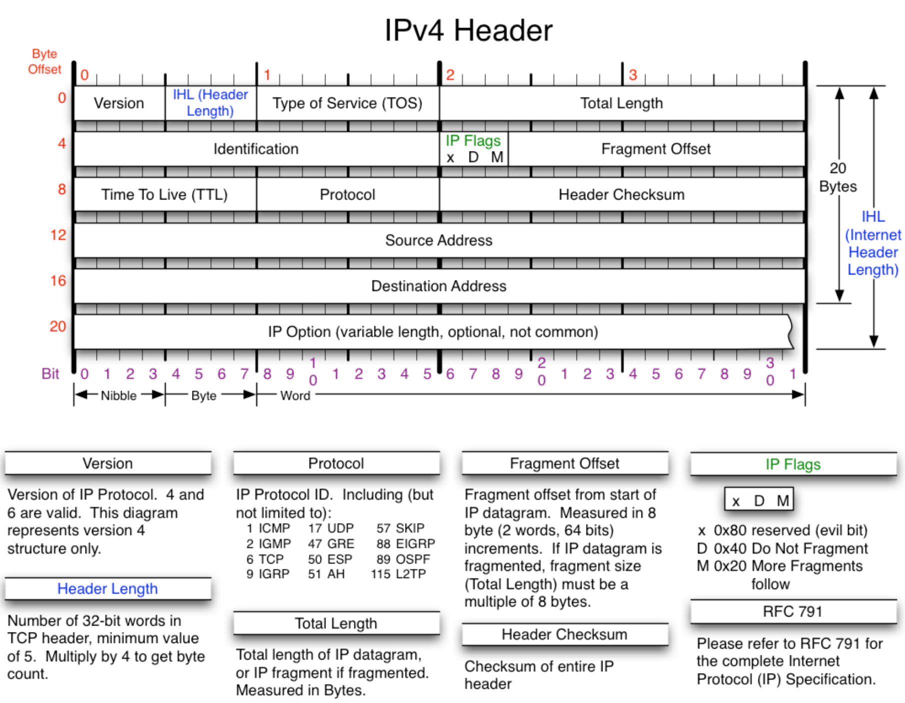
IPv4 Header
Byte Offset
0 1 2 3
0 Version | IHL (Header Length) | Type of Service (TOS) | Total Length
4 Identification | IP Flags x D M | Fragment Offset
8 Time To Live (TTL) | Protocol | Header Checksum
12 Source Address
16 Destination Address
20 IP Option (variable length, optional, not common)
20 Bytes
IHL (Internet Header Length)
Bit 0 1 2 3 4 5 6 7 8 9 10 1 2 3 4 5 6 7 8 9 20 1 2 3 4 5 6 7 8 9 30 1
Nibble
Byte
Word
Version
Version of IP Protocol. 4 and 6 are valid. This diagram represents version 4 structure only.
Header Length
Number of 32-bit words in TCP header, minimum value of 5. Multiply by 4 to get byte count.
Protocol
IP Protocol ID. Including (but not limited to):
1 ICMP 17 UDP 57 SKIP
2 IGMP 47 GRE 88 EIGRP
6 TCP 50 ESP 89 OSPF
9 IGRP 51 AH 115 L2TP
Total Length
Total length of IP datagram, or IP fragment if fragmented. Measured in Bytes.
Fragment Offset
Fragment offset from start of IP datagram. Measured in 8 byte (2 words, 64 bits) increments. If IP datagram is fragmented, fragment size (Total Length) must be a multiple of 8 bytes.
Header Checksum
Checksum of entire IP header
IP Flags
x D M
x 0x80 reserved (evil bit)
D 0x40 Do Not Fragment
M 0x20 More Fragments follow
RFC 791
Please refer to RFC 791 for the complete Internet Protocol (IP) Specification.

Running out of 32-bit Addresses

- Not all that many unique addresses
  - 2^{32} = 4,294,967,296 (just over four billion)
  - Plus, some are reserved for special purposes
  - And, addresses are allocated in larger blocks
- Solution: a larger address space
  - IPv6 has 128-bit addresses ($2^{128} = 3.403 \times 10^{38}$)
  - 1500 addresses/square foot of Earth's surface ☺

3

IPv6 vs. IPv4: Address Comparison

- **IPv4** has a maximum of

  $2^{32} \approx 4$ billion addresses

- **IPv6** has a maximum of

  $2^{128} = (2^{32})^4 \approx$ 4 billion x 4 billion x 4 billion x 4 billion

  addresses

4

IPv6 Address Format

- Base format (16-byte)

  2001:0660:3003:0001:0000:0000:6543:210F

- Compact Format:

  2001:660:3003:1::6543:210F

  – Leading zeros in a group can be omitted
  – Only one set of zeros can be truncated

IPv6 Address Notation: Example

128.91.45.157.220.40.0.0.0.0.252.87.212.200.31.255

Binary	1000000001011011001011011001110111011100001010000000000000000000 0000000000000000111111000101011111010100110010000001111111111111

Dotted Decimal	128	91	45	157	220	40	0	0	0	0	252	87	212	200	31	255

Hexadecimal	0	32		64		96		128
Straight Hex	805B	2D9D	DC28	0000	0000	FC57	D4C8	1FFF
Leading-Zero Suppressed	805B	2D9D	DC28	0	0	FC57	D4C8	1FFF
Zero-Compressed	805B	2D9D	DC28	::		FC57	D4C8	1FFF
Mixed Notation	805B	2D9D	DC28	::		FC57	212 200	31 255

IPv6 Zero Compression

- To determine the number of 0 bits represented by the "::"
  - count the number of blocks in the compressed address
  - (-) subtract this number from 8
  - (*) multiply the result by 16.

- For example
  - FF02::2
  - two blocks - "FF02" block and "2" block.
  - The number of bits expressed by the "::" is 96
    (96 = (8 – 2)×16).

7

IPv6 Address Categories

- Unicast
  - Uniquely identifies an interface of an IPv6 node. A packet sent to a unicast address is delivered to the interface identified by that address
- Multicast
  - Identifies a group of IPv6 interfaces. A packet sent to a multicast address is processed by all members of the multicast group
- Anycast
  - Assigned to multiple interfaces (on multiple nodes). A packet sent to an anycast address is delivered to only one of these interfaces, usually the nearest one.
- There is no broadcast address in IPv6.

8

IPv6 Address Types

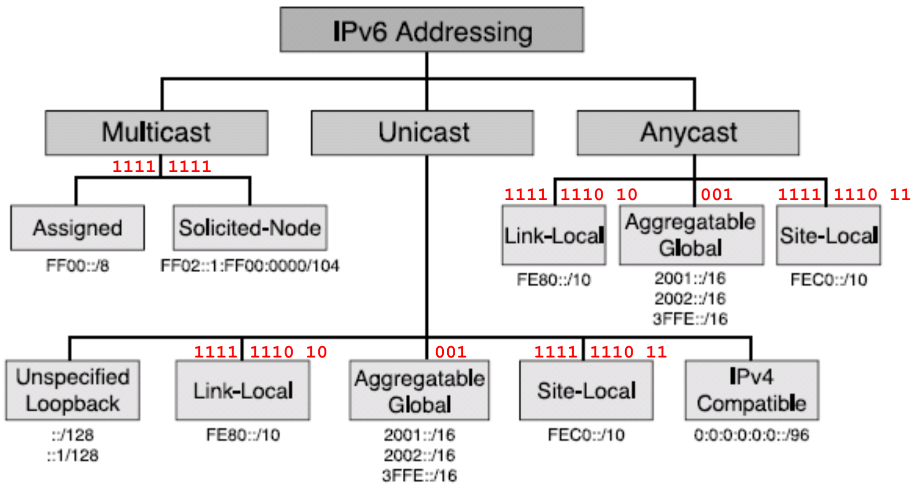

9

IPv6 Addressing Model

- Addresses are assigned to interfaces
- Change from the IPv4 model :
  - Interface 'expected' to have multiple addresses
- Addresses have scope
  - Link local
  - Site local
  - Global

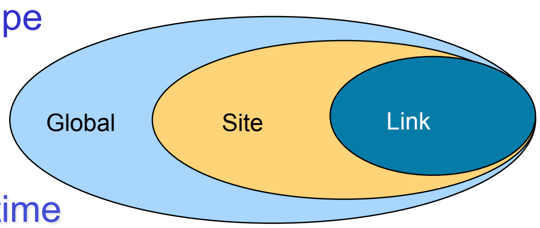

- Addresses have lifetime
  - Valid and Preferred lifetime

10

Global Unicast Addresses

- Equivalent to public IPv4 addresses, globally routable
- Global scoped communication are identified by high-level 3 bits set to 001 (2000::/3)

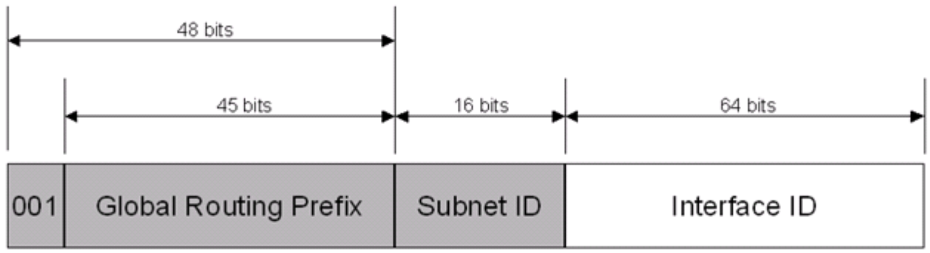

11

Global Unicast Address

- Global Routing Prefix – Site Prefix
  – Site prefix assigned to an organization by a provider should be at least a /48 prefix = 45 + high-order bits (001)
- Subnet ID - Site
  – With the /48 prefix, it is possible for that organization to enable up to 65,535 subnets
  – The organization can use bits 49 to 64 (16-bit) of the prefix received for subnetting
- Interface ID – Host
  – The host part uses each node's interface identifier
  – This part of the IPv6 address represented by the low-order 64-bit is called the *interface ID*

12

Global Unicast Address Example

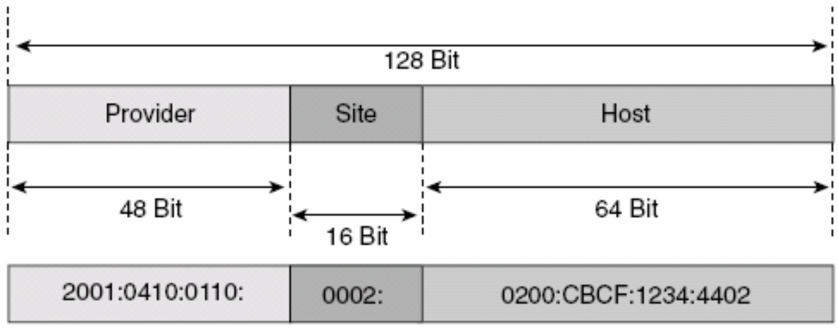

- 2001:0410:0110::/48 is assigned by a provider
- 2001:0410:0110:0002::/64 subnet within the organization
- 2001:0410:0110:0002:0200:CBCF:1234:4402 – node address within the subnet

13

Local-Use Unicast Addresses

1. Link-local addresses
   - Used between neighbours on the same link
   - Always automatically configured

   `1111 1110 10 (FE80::/10)`

2. Site-local addresses
   - Used between nodes communicating with other nodes in the same site.

   `1111 1110 11 (FEC0::/10)`

14

Link-Local Unicast Address

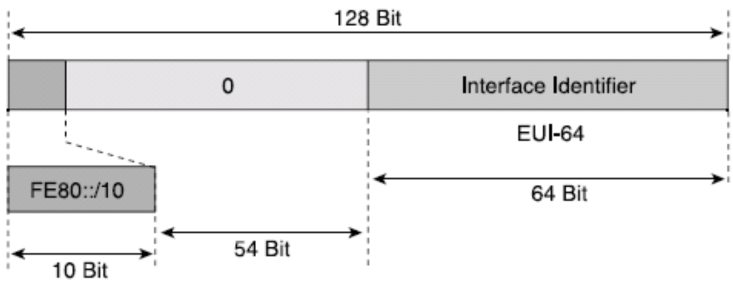

- When an IPv6 stack is enabled on a node, one link-local address is automatically assigned to each interface of the node at boot time.
- Link-local addresses are only for local-link scope and must never be routed between subnets within a site.

15

MAC to EUI-64 example

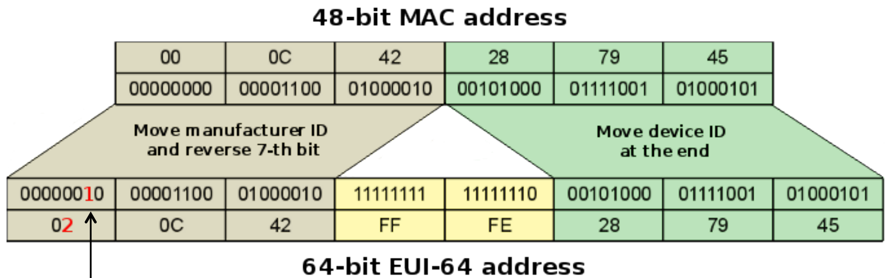

0: globally unique
1: locally administered

Exercise

- What is the link-local unicast address of the interface with MAC address 00:2c:66:8e:fa:11?

17

Site-Local Address

- A site is an organization network or portion of an organization's network that has a defined geographical location (such as an office, an office complex, or a campus).
- Site-local addresses are equivalent to the IPv4 private addresses (10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16).
- Site-local addresses are not reachable from other sites, and routers must not forward site-local traffic outside the site.
- Site-local addresses can be used in addition to global unicast addresses.

18

Site-Local Address

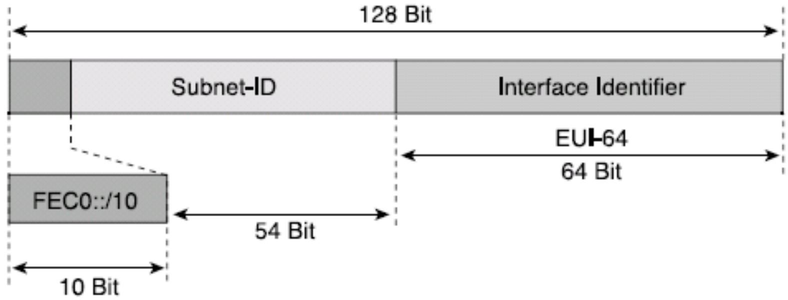

- Unlike link-local addresses, site-local addresses are not automatically configured and must be assigned either through stateless or stateful address configuration processes.
- May be assigned to any nodes and routers within a site

Site-Local Address Example

- Assign site-local prefixes to a site with 10 subnets:

Subnet	Site-Local Prefix (::/64)
1	
2	
3	
4	
5	
6	
7	
8	
9	
10	

Special IPv6 Addresses (1)

- Unspecified address
  - 0:0:0:0:0:0:0:0 or ::
  - Used to indicate the absence of an address
  - Equivalent to the IPv4 address 0.0.0.0
- Loopback address
  - 0:0:0:0:0:0:0:1 or ::1
  - Used to identify a loopback interface, enabling a node to send packets to itself.
  - Equivalent to IPv4 loopback address 127.0.0.1
  - Never sent out on a link

Special IPv6 Addresses (2)

- IPv4 compatible
  - 0:0:0:0:0:0:IPv4 = ::IPv4
  - 0:0:0:0:0:0:192.168.30.1 = ::192.168.30.1 = ::C0A8:1E01
  - Special unicast IPv6 address used by transition mechanisms on dual-stack hosts and routers to automatically create IPv4 tunnels to deliver IPv6 packets over IPv4 networks.
- IPv4 mapped
  - 0:0:0:0:0::FFFF:IPv4 = ::FFFF:IPv4
  - 0:0:0:0:0:FFFF:192.168.30.1 = ::FFFF:C0A8:1E01
  - Used only locally on nodes having both IPv4 and IPv6 stacks
  - These addresses are never known outside the node itself and should not go on the wire as IPv6 addresses
  - Rarely implemented

Questions

- What is the IPv4-compatible IPv6 address corresponding to the IPv4 address 2.15.17.20?

- What is the IPv4-mapped IPv6 address corresponding to the IPv4 address 2.15.17.20?

23

Multicast Addresses

- Always begin with FF
- Two types:
  - – Assigned FF00::/8
  - – Solicited-Node FF02::1:FF00:0000/104
- For each unicast and anycast address configured on an interface of a node or router, a corresponding solicited-node multicast address is automatically enabled.
- The solicited-node multicast address is local to the link.

24

Solicited-Node Multicast Address

- Consists of the prefix FF02::1:FF00:0000/104 + the low-order 24-bit of the unicast or anycast address.

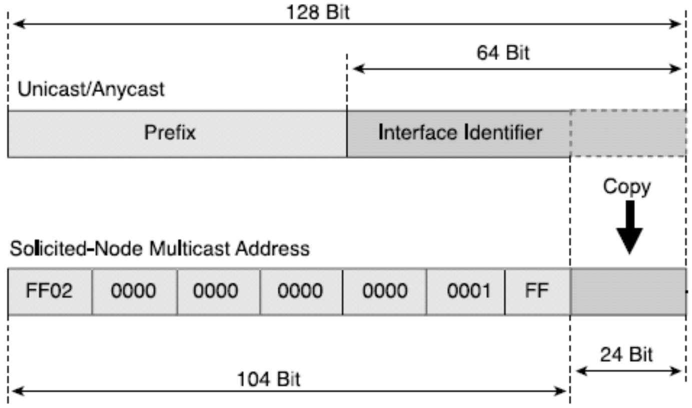

Solicited-Node Multicast Address Examples

Unicast Address	Solicited-Node Multicast Address
2001:425:0:1::45FF	
FEC0::1:1:1:777	

Multicast Addresses

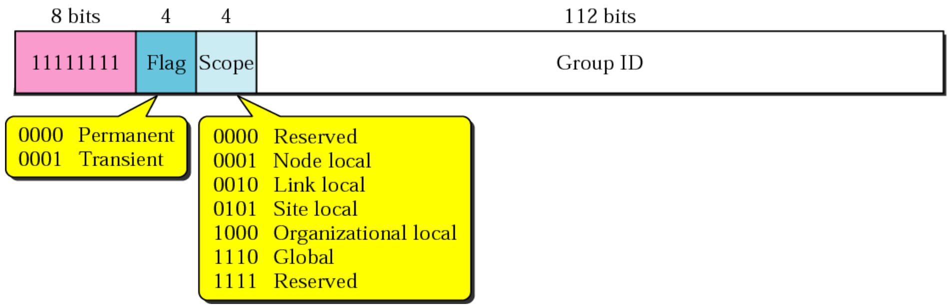

- Always begin with FF
- Special Multicast Addresses:
  - Link-local all-nodes (FF02::1)
  - Link-local all-routers(FF02::2)
  - Site-local all-routers (FF05::2)

27

Summary: IPv6 Address Space Subdivision

Prefix	Address	Compact	Use
			Unspecified
			Loopback
			Multicast
			Solicited-node multicast
			Link-local unicast
			Site-local unicast
Everything else	Global Unicast		

28

So how many addresses can an IPv6 host have?

IPv6 Addresses for a Host

- Unicast addresses:
  1. A link-local address for each interface
  2. Routable unicast addresses for each interface (which could be a site-local address and one or more global unicast addresses)
  3. The loopback address (::1) for the loopback interface

- Additionally, each host is listening for traffic on the following multicast addresses:
  1. The link-local all-nodes multicast address (FF02::1/128)
  2. The solicited-node address for each unicast address on each interface (FF02::1:FF00:0000/104)

And how many addresses can an IPv6 router have?

IPv6 Addresses for a Router

- Same as a host. In addition, each router listens for traffic on the following multicast addresses:
  1. The link-local all-routers multicast address (FF02::2)
  2. The site-local all-routers multicast address (FF05:2)

IPv6 Header

IPv6 Header

- The IPv6 header is redesigned.
- Goal is to minimize overhead and reduce the header processing for the majority of the packets.
- Less essential and optional fields are moved to extension headers

IPv4 (20 bytes+options)

4-bit Version	4-bit Header Len	8-bit Type of Service (TOS)	16-bit Total Length (Bytes)	
16-bit Identification			3-bit Flags	13-bit Fragment Offset
8-bit Time to Live (TTL)		8-bit Protocol	16-bit Header Checksum	
32-bit Source IP Address				
32-bit Destination IP Address				

IPv6 (40 bytes)

4-bit Version	8-bit Traffic Class	20-bit *Flow Label*	
16-bit Payload Length (Bytes)		8-bit Next Header	8-bit Hop Limit
128-bit Source IP Address			
128-bit Destination IP Address			

IPv6 Header

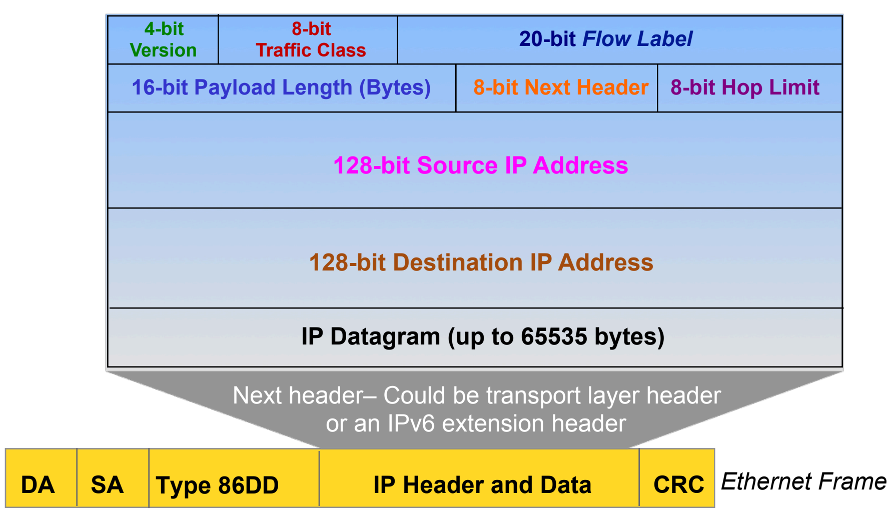

IPv6 Header Fields (1)

- Flow Label (20 bits)
  - Identifies specific flows needing special QoS
  - Each source chooses its own flow label values
  - Flow label value of 0 used when no special QoS requested (the common case today)
  - Routers use source address + flow label to identify distinct flows

IPv6 HeaderFields (2)

- Payload length:
  - similar to the total length in IPv4, except it does not include the 40-byte header
- Hop limit:
  - like TTL field, decrements by one for each router
- Next header:
  - similar to the protocol field in IPv4
  - tells you what type of information follows (TCP, UDP, extension header)

IPv6 Extension Headers

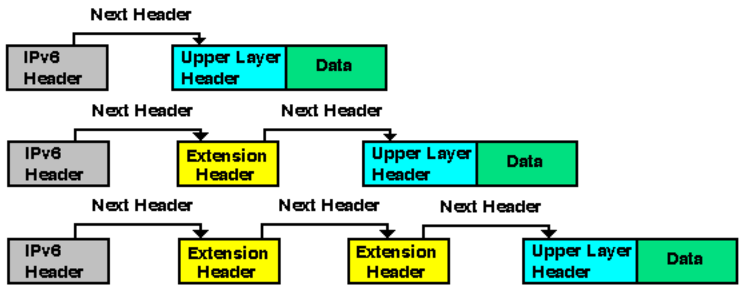

Common Next Headers:

1 (0x01)	ICMP header
43 (0x2B)	Routing header
44 (0x2C)	Fragmentation header
6 (0x06)	TCP protocol
17 (0x11)	UDP protocol
58 (0x3A)	ICMPv6 protocol
59 (0x3B)	IPv6 No Next Header (terminates a no upper layer frame)

39

IPv6 Fragmentation

- IPv6 datagrams include:

1. Unfragmentable Part
   - Includes the main header of the original datagram + any extension headers that need to be present in each fragment (such as the Routing header)

2. Fragmentable Part
   - Data portion of the datagram + other extension headers if present

- *Unfragmentable part* must be present in each fragment, while the *fragmentable* part is split up amongst the fragments.

40

IPv6 Fragment Sets

So to fragment a datagram, a device creates a set of fragment datagrams, each of which contains the following, in order:

- Unfragmentable Part
  - The full *Unfragmentable Part* of the original datagram, with its *Payload Length* changed to the length of the fragment datagram
- Fragment Header
  - A *Fragment* header with the *Fragment Offset*, *Identification* and *M* flags set in the same way they are used in IPv4
- Fragment
  - A fragment of the *Fragmentable Part* of the original datagram. Note that each fragment must have a length that is a multiple of 8 bytes, because the value in the *Fragment Offset* field is specified in multiples of 8 bytes.

IPv6 Fragmentation Example

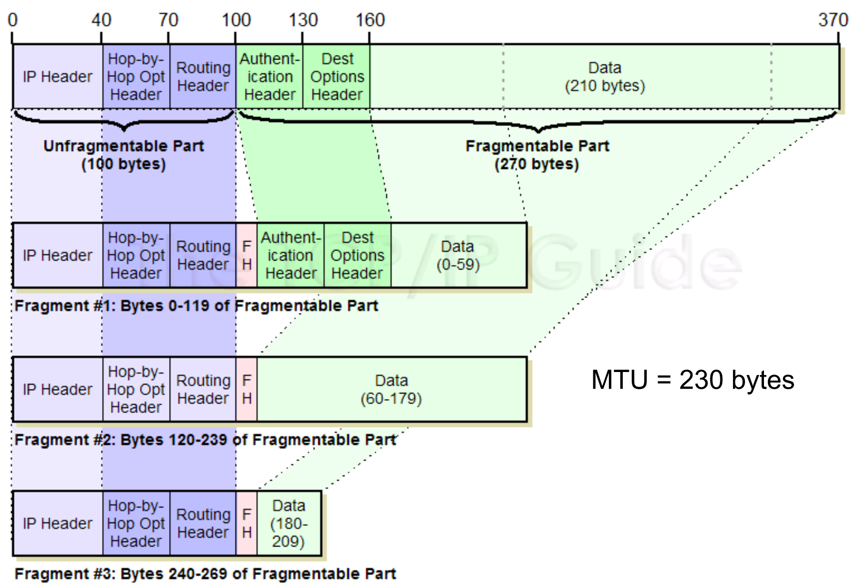

MTU = 230 bytes

IPv6 Neighbor Discovery

Neighbor Discovery (ND) Protocol

- Used to determine link-layer addresses, find routers
- Replaces ARP, ICMP Router Discovery, and ICMP Redirect used in IPv4 and provides additional functionality

- ICMPv6 is a "Super" Protocol that :
  – Covers ICMP (v4) features (error control, info, …)
  – Transports ND messages

IPv6 ND Overview

- ND is used by hosts to:
  - Discover neighbouring routers
  - Discover addresses, address prefixes, and other configuration parameters
- ND is used by routers to:
  - Advertise their presence, host configuration parameters, and on-link prefixes
  - Inform hosts of a better next-hop address
- ND is used by nodes to:
  - Resolve the link-layer address of a neighbouring node
  - Determine whether a neighbour is still reachable

45

Neighbor Discovery (1)

- Most ND protocol functions are implemented using a set of five special ICMPv6 control messages.
  - Router Advertisement (RA) :
    - periodic advertisement (of the availability of a router) which contains:
      - IPv6 address prefix used on the link
      - a possible value for Max Hop Limit (TTL of IPv4)
      - value of MTU
  - Router Solicitation (RS) :
    - the host needs RA immediately (at boot time)

46

Neighbor Discovery (2)

- Neighbor Solicitation (NS):
  - to determine the link-layer address of a neighbor
  - also used to detect duplicate addresses (DAD)
- Neighbor Advertisement (NA):
  - answer to a NS packet
  - to advertise the change of physical address
- Redirect :
  - Used by a router to inform a host of a better route to a given destination

47

Router Advertisement Messages

- Router advertisement messages include configuration parameters such as link prefixes and their lifetimes, default hop limit, MTU
- Active IPv6 hosts on the local link receive the Router Advertisement messages and use the contents to maintain the default router list, the prefix list, and other configuration parameters.
- A host that is starting up sends a **Router Solicitation message** to the **link-local scope all-routers multicast address (FF02::2)**

51

Router Solicitation Messages

- Any node can send RS to all-routers multicast address FF02::2 on the local link
- When RS is received, a router responds with RA using all-node multicast FF02::1
- To avoid flooding of RS on the link, each node can send only three RS at boot time.

52

IPv6 Stateless Autoconfiguration

Stateless Autoconfiguration Procedure

- When booting, a host asks for network parameters:
  - IPv6 prefix (-es), default router address (-es), hop limit, (link local) MTU, prefix validity time, etc
- Host listens to Router Advertisements (RA)
  - RA are periodically transmitted by routers
- Hosts creates …
  - A link local address (fe80::/10)
  - A global IPv6 address using:
    - its interface identifier (EUI-64 address)
    - link prefix obtained via Router Advertisement

60

Stateless Autoconfiguration Procedure

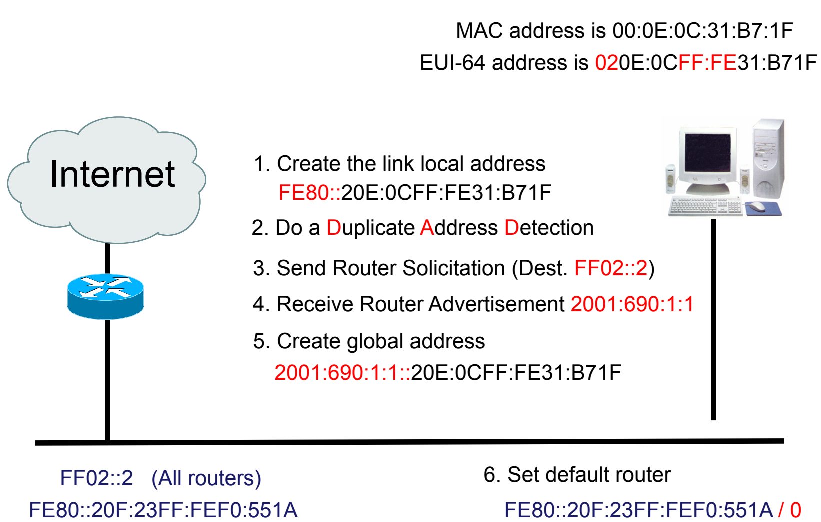

61

Conclusions

- IPv6 Addresses
- IPv6 Datagram Format
- IPv6 Neighbour Discovery
- IPv6 Stateless Autoconfiguration