Internet Protocol (IP)

Part of TCP/IP, used by the Internet

Specifies interface with higher layer, e.g. TCP

Specifies protocol format and mechanisms

IP Services

Primitives

Functions to be performed

Form of primitive implementation dependent, e.g. subroutine call

Send

Request transmission of data unit

Deliver

Notify user of arrival of data unit

Parameters

Used to pass data and control info

Source address

Destination address

Protocol

Recipient, e.g. TCP

Type of Service

Specify treatment of data unit during transmission through networks

Identification of IP packet

Source, destination address and user protocol

Uniquely identifies PDU

Needed for re-assembly and error reporting

Send only

Don't fragment indicator

Can IP fragment data

If not, may not be possible to deliver

Send only

Time to live

Send only

Data length

Option data

User data

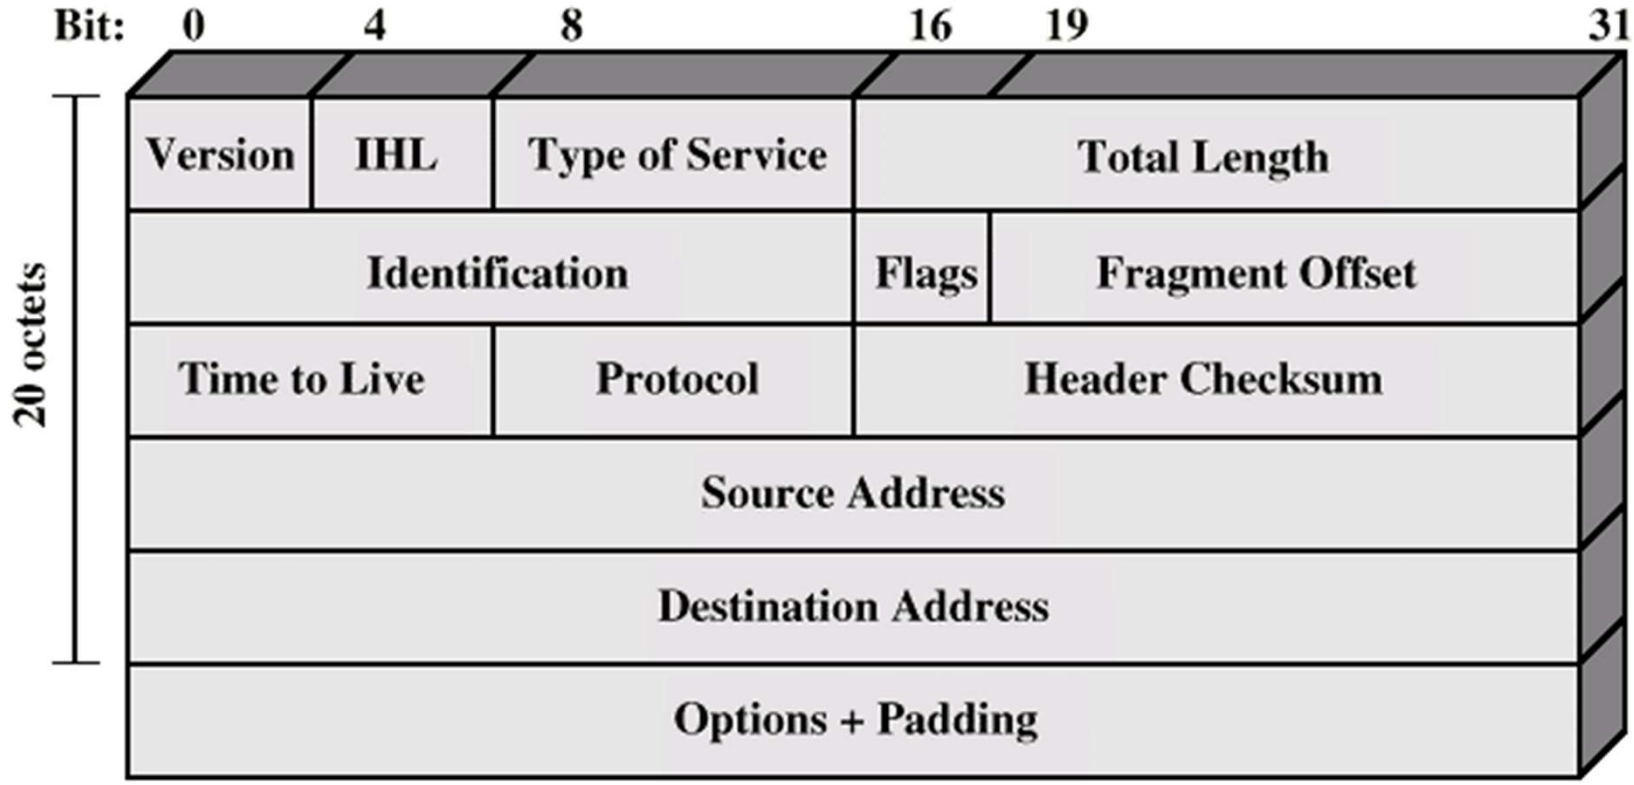

IP packet structure

Header Fields

Version

Currently 4

IP v6 - see later

Internet header length

In 32 bit words

Including options

Type of service

Total length

Of datagram, in octets

Identification

Sequence number

Used with addresses and user protocol to identify datagram uniquely

Flags

More bit

Don’t fragment

Fragmentation offset

Time to live

Protocol

Next higher layer to receive data field at destination

Header checksum

Reverified and recomputed at each router

16 bit ones complement sum of all 16 bit words in header

Set to zero during calculation

Source address

Destination address

Options

Padding

To fill to multiple of 32 bits long

Data Field

Carries user data from next layer up

Integer multiple of 8 bits long (octet)

Max length of datagram (header plus data) 65,535 octets

IP Addresses

32 bit global internet address

Network part and host part

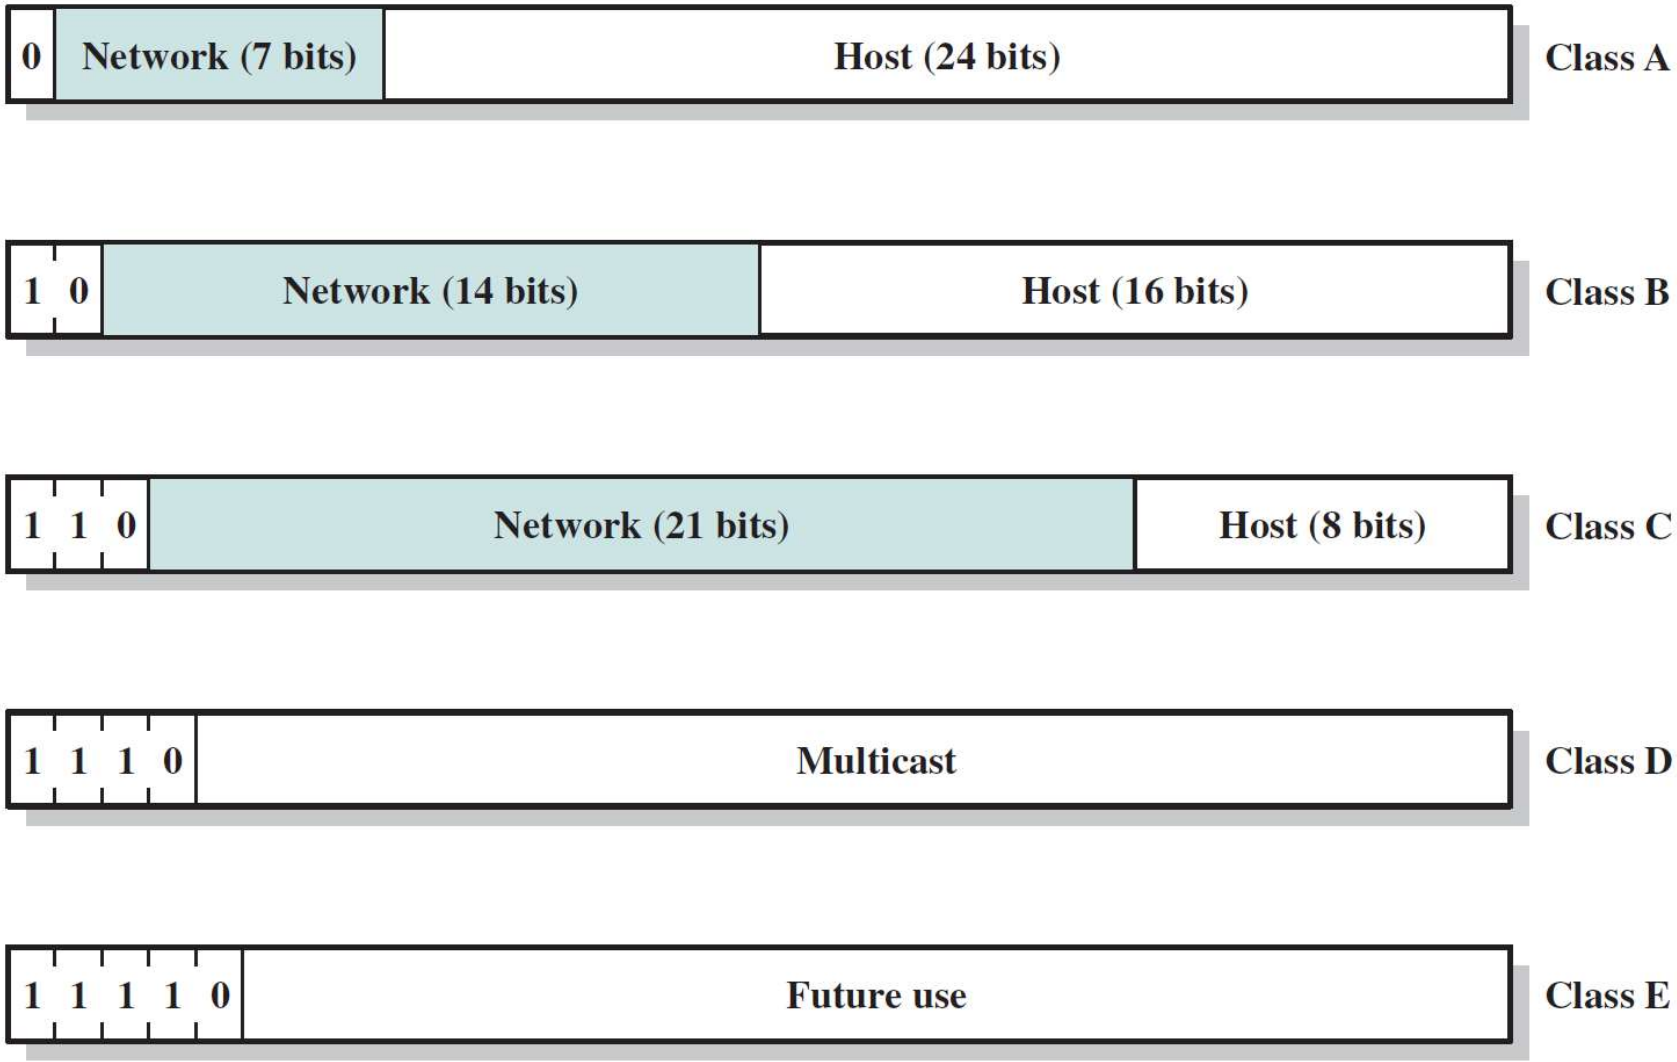
0
Network (7 bits)
Host (24 bits)
Class A
1 0
Network (14 bits)
Host (16 bits)
Class B
1 1 0
Network (21 bits)
Host (8 bits)
Class C
1 1 1 0
Multicast
Class D
1 1 1 1 0
Future use
Class E

Class A

Start with binary 0

All 0 reserved

01111111 (127) reserved for loopback

Range 1.x.x.x to 126.x.x.x

All allocated

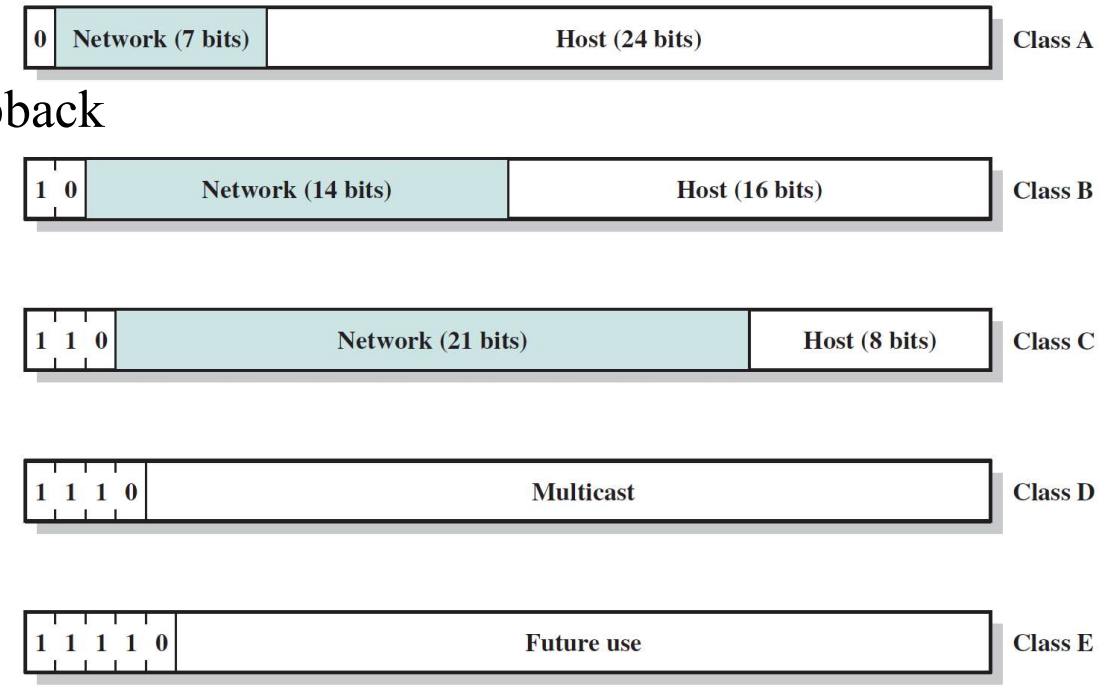

Class B

Start 10

Range 128.x.x.x to 191.x.x.x

Second Octet also included in network address

2^{14} = 16,384 class B addresses

All allocated

Class C

Start 110

Range 192.x.x.x to 223.x.x.x

Second and third octet also part of network address

2^{21} = 2,097,152 addresses

Nearly all allocated

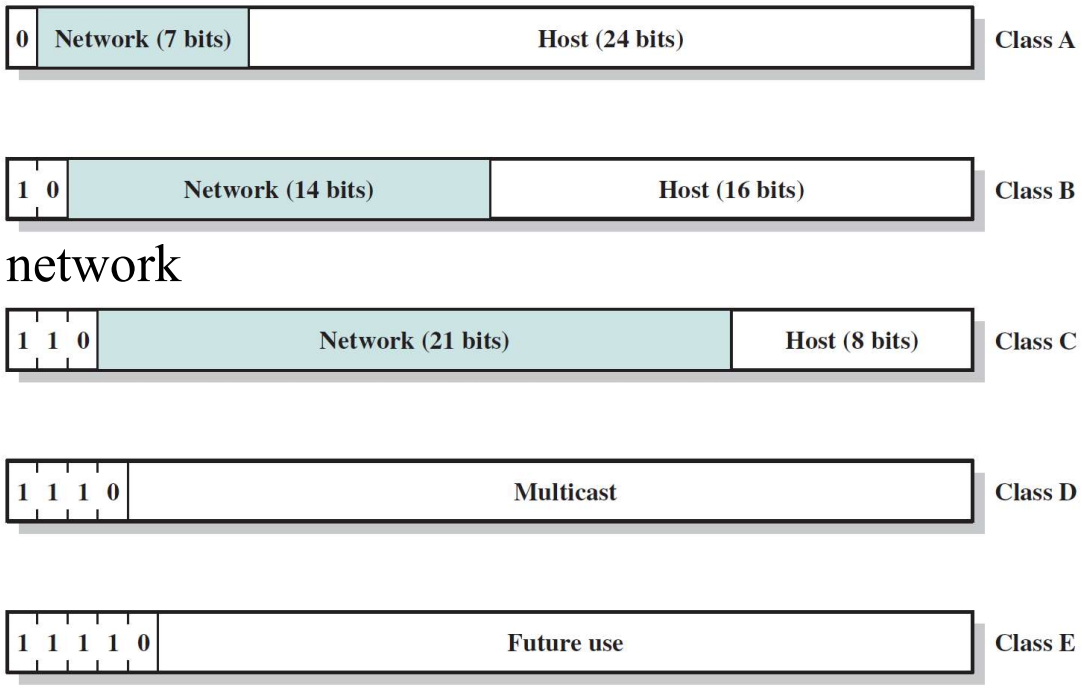

Class D

Contain multicast addresses for group users

first decimal field is between 224 and 239

Class E

Rreserved for research and future developments

First decimal field between 240 and 255

Class	1st Octet Decimal Range	1st Octet High Order Bits	Network/Host ID (N=Network, H=Host)	Default Subnet Mask	
A	1 – 126*	0	N.H.H.H	255.0.0.0	/8
B	128 – 191	10	N.N.H.H	255.255.0.0	/16
C	192 – 223	110	N.N.N.H	255.255.255.0	/24
D	224 – 239	1110	Reserved for Multicasting		
E	240 – 254	1111	Experimental; used for research		

Note: Class A addresses 127.0.0.0 to 127.255.255.255 cannot be used and is reserved for loopback and diagnostic functions.

- *Private IP addresses*

Class A: 10.0.0.0 - 10.255.255.255 /8
Class B: 172.16.0.0 - 172.31.255.255 /16
Class C: 192.168.0.0 - 192.168.255.255 /24

	172	25	114	250
IP Address (B class)	10101100	00011001	01110010	11111010
Network Mask	11111111	11111111	00000000	00000000
	255	255	0	0
Network Address	10101100	00011001	00000000	00000000
	172	25	0	0
Broadcast Address	10101100	00011001	11111111	11111111
	172	25	255	255

AND

AND

0	0	0
0	1	0
1	0	0
1	1	1

- Network address: all host bits = 0
- Network broadcast address: all host bits = 1
- Total number of host bits: 2^{16} = 65,536
- Number of hosts: 2^{16} -2 = 65,536 – 2 = 65,534

	Binary Representation	Dotted Decimal
IP address	11000000.11100100.00010001.00111001	192.228.17.57
Subnet mask	11111111.11111111.11111111.11100000	255.255.255.224
Bitwise AND of address and mask (resultant network/subnet number)	11000000.11100100.00010001.00100000	192.228.17.32
Subnet number	11000000.11100100.00010001.001	1
Host number	00000000.00000000.00000000.00011001	25

(b) Default subnet masks

	Binary Representation	Dotted Decimal
Class A default mask	11111111.00000000.00000000.00000000	255.0.0.0
Example Class A mask	11111111.11000000.00000000.00000000	255.192.0.0
Class B default mask	11111111.11111111.00000000.00000000	255.255.0.0
Example Class B mask	11111111.11111111.11111000.00000000	255.255.248.0
Class C default mask	11111111.11111111.11111111.00000000	255. 255. 255.0
Example Class C mask	11111111.11111111.11111111.11111100	255. 255. 255.252

Subnets and Subnet Masks

Allow arbitrary complexity of internetworked LANs within organization

Insulate overall internet from growth of network numbers and routing complexity

Site looks to rest of internet like single network

Each LAN assigned subnet number

Host portion of address partitioned into subnet number and host number

Local routers route within subnetted network

Subnet mask indicates which bits are subnet number (1s) and which are host number (0s)

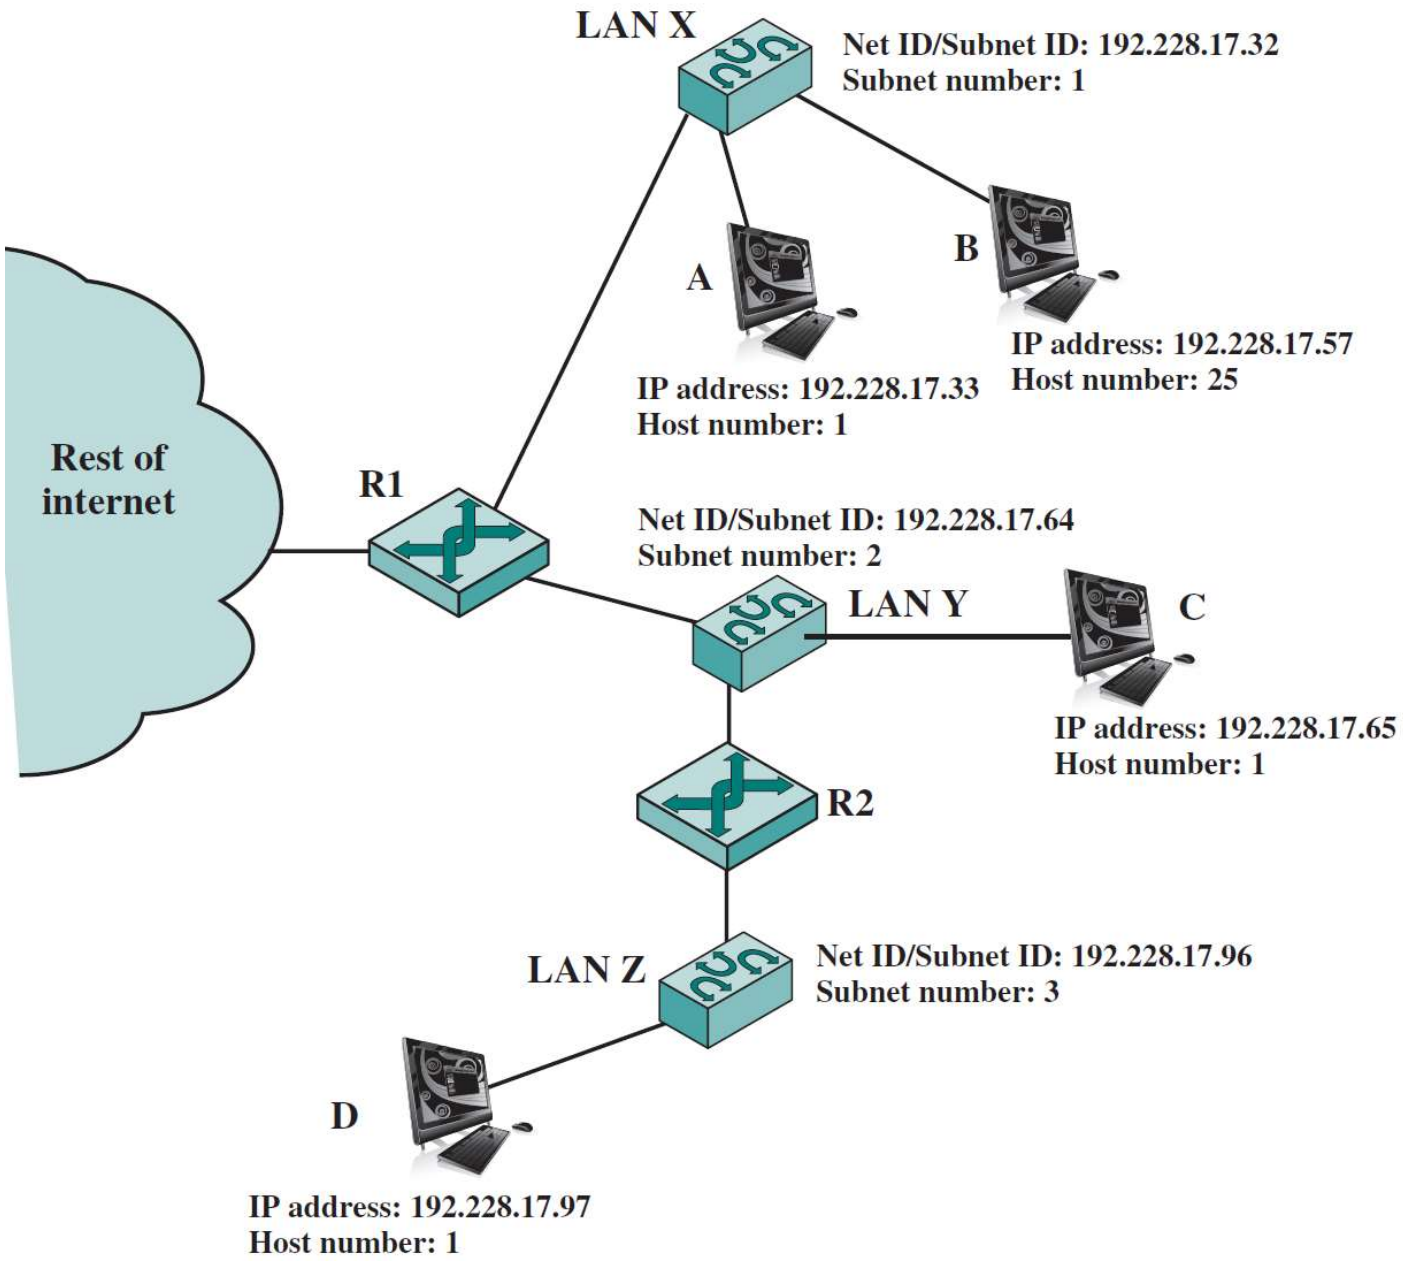
LAN X
Net ID/Subnet ID: 192.228.17.32
Subnet number: 1
A
B
IP address: 192.228.17.33
Host number: 1
IP address: 192.228.17.57
Host number: 25
Rest of internet
R1
Net ID/Subnet ID: 192.228.17.64
Subnet number: 2
LAN Y
C
IP address: 192.228.17.65
Host number: 1
R2
LAN Z
Net ID/Subnet ID: 192.228.17.96
Subnet number: 3
D
IP address: 192.228.17.97
Host number: 1

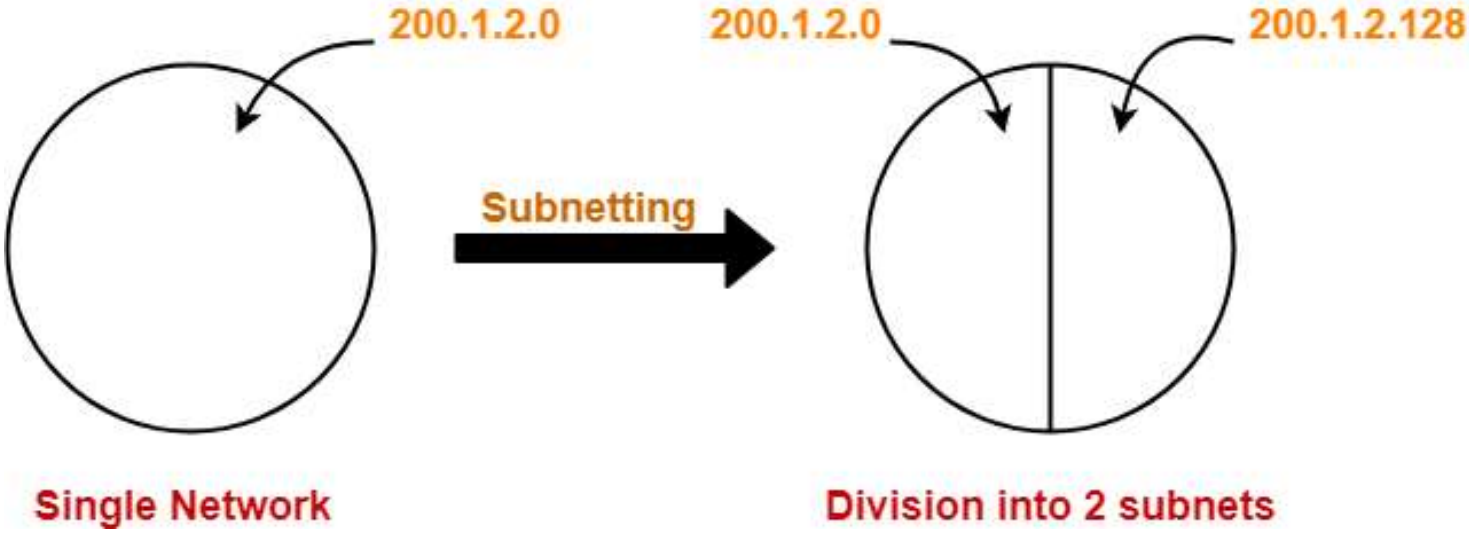
200.1.2.0
200.1.2.0
200.1.2.128
Subnetting
Single Network
Division into 2 subnets

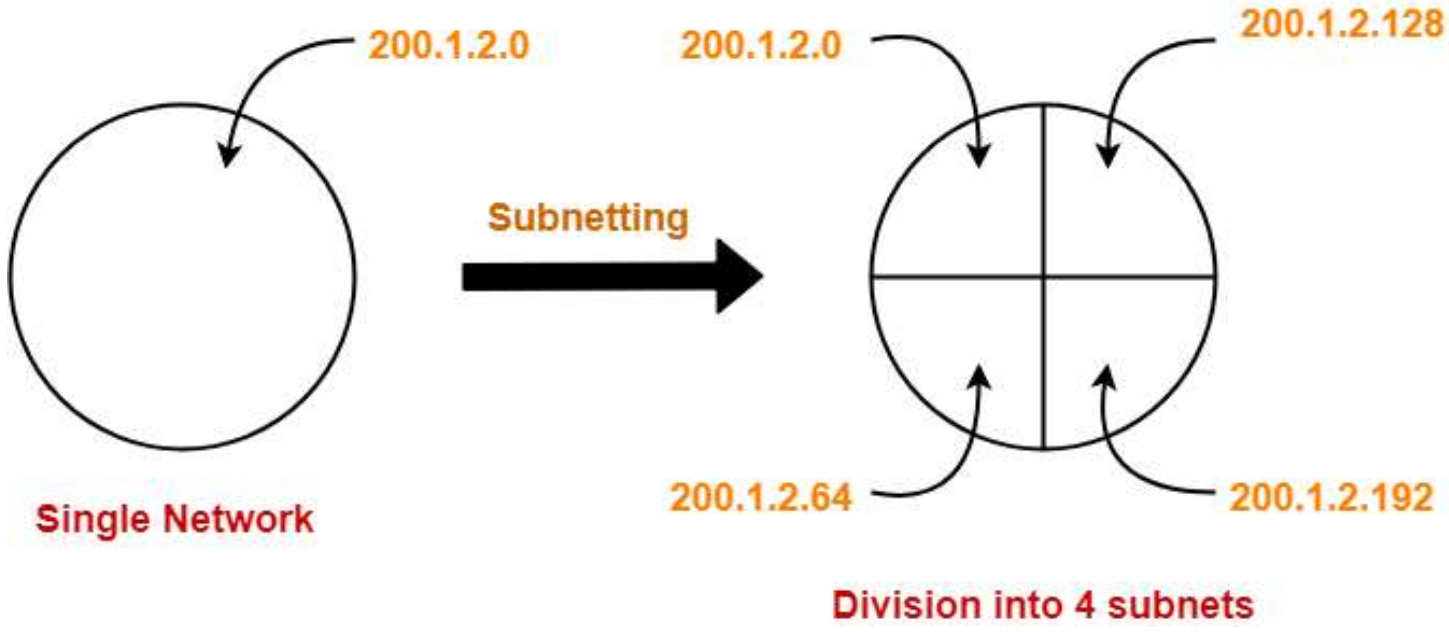
200.1.2.0
200.1.2.0
200.1.2.128
Subnetting
Single Network
200.1.2.64
200.1.2.192
Division into 4 subnets

- Creating subnets:
  - borrow bits from the host ID
  - create a new Network Mask to show the new structure of the IPv4

	172	25	114	250		AND
IP Address (B class)	10101100	00011001	01110010	11111010		0 0 0
Network Mask	255	255	0	0	/16	0 1 0
Subnet Mask	255	255	255	192	/26	1 0 0
	11111111	11111111	11111111	11000000	AND	1 1 1
Subnet Address	10101100	00011001	01110010	11000000		
	172	25	114	192		
Subnet Broadcast	10101100	00011001	01110010	11111111		
	172	25	114	255		

Total number of host bits: 2^6	
Number of hosts: 2^6 -2 = 64 – 2 = 62	First host IP on subnet: 172.25.114.193
Total number of subnet bits: 2^{10}	Last host IP on subnet: 172.25.114.254
Number of subnets: 2^{10} = 1024	

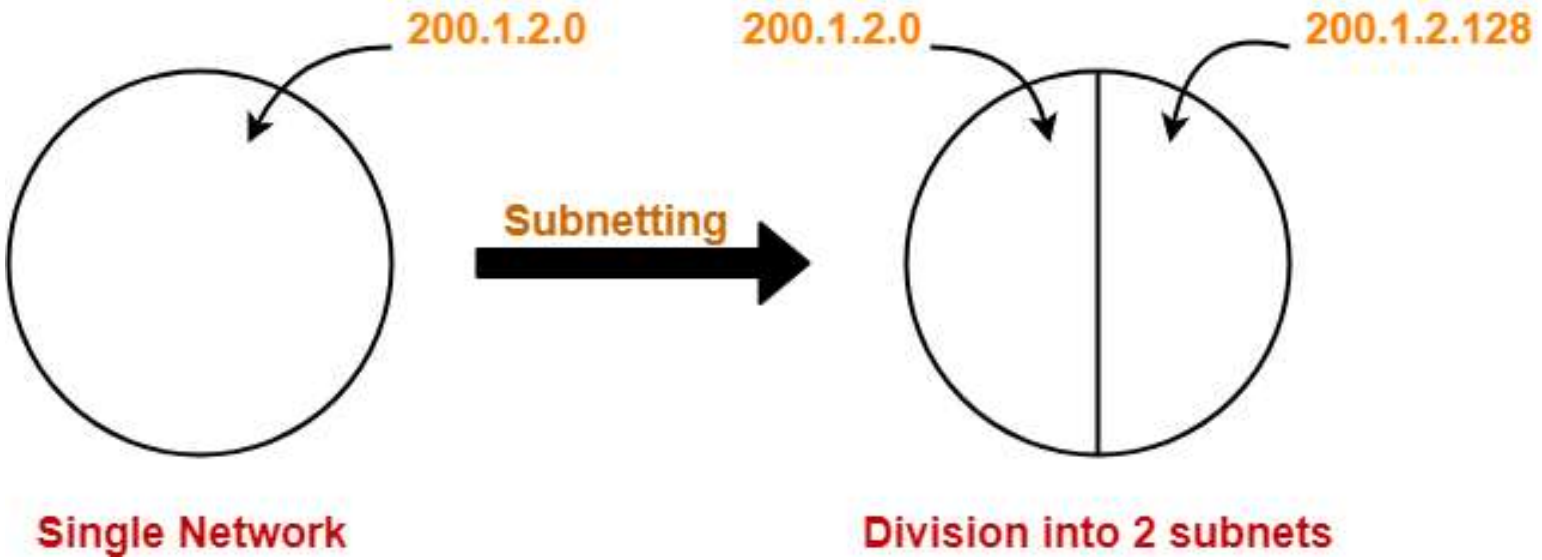
200.1.2.0
Subnetting
Single Network
200.1.2.0
200.1.2.128
Division into 2 subnets

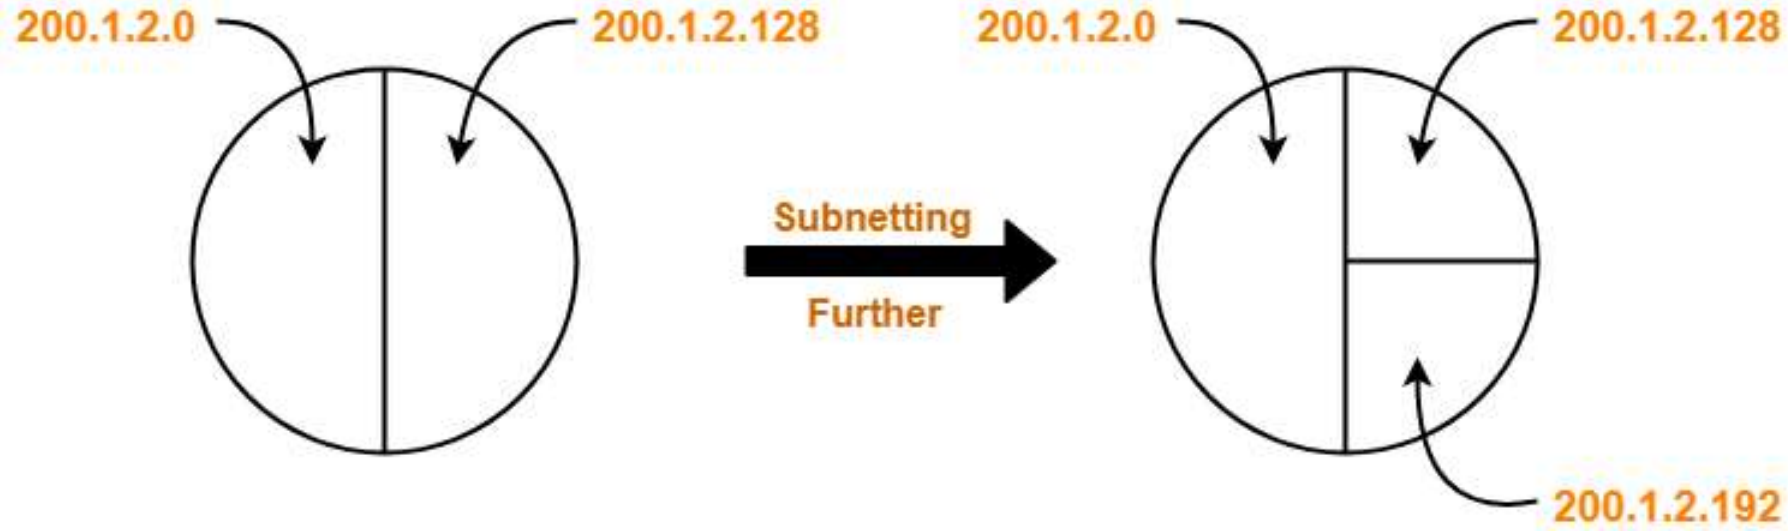
200.1.2.0
200.1.2.128
Subnetting
Further
200.1.2.0
200.1.2.128
200.1.2.192

ICMP

Need for appending to IP (used only for data transfer) of some **control protocols,** e.g. ICMP, ARP, RARP, BOOTP

Internet Control Message Protocol (ICMP) provides error-reporting mechanisms

RFC 792

Transfer of (control) messages from routers and hosts, to other hosts

Feedback about problems

e.g. time to live expired

ICMP packet encapsulated in IP datagram

Not a very reliable protocol, because (see next slide):

IP provides a *best-effort delivery* (not a secure one)
Internet layer can detect a small variety of errors:

- •Checksum (header only!)
- •TTL expires
- •No route to destination network
- •Can't deliver to destination host (e.g., no ARP reply)

Internet layer discards datagrams with problems

Some - e.g., checksum error - can't trigger error messages (ICMP message)

Some errors can be reported

Remember:
ICMP encapsulated in IP (see drawing)
ICMP messages sent in response to incoming datagrams with problems
ICMP message **not** sent for ICMP message

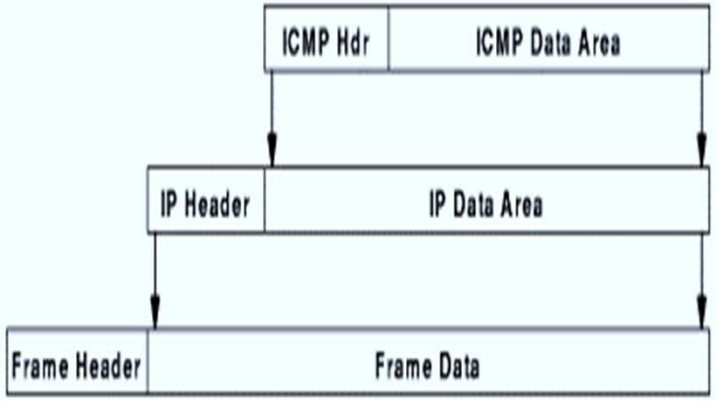

Message type	Description
Destination unreachable	Packet could not be delivered
Time exceeded	Time to live field hit 0
Parameter problem	Invalid header field
Source quench	Choke packet
Redirect	Teach a router about geography
Echo request	Ask a machine if it is alive
Echo reply	Yes, I am alive
Timestamp request	Same as Echo request, but with timestamp
Timestamp reply	Same as Echo reply, but with timestamp

The principal ICMP message types.

ICMP and reachability

An internet host, *A*, is *reachable* from another host *B*, if datagrams can be delivered from *A* to *B*

TCP/IP *ping* program tests reachability - sends datagram from *B* to *A,* that *A echoes* back to B

Uses ICMP *echo request* and *echo reply* messages, e.g. Internet layer includes code to reply to incoming ICMP *echo request* messages

ICMP and internet routes

List of all routers on path from *A* to *B* is called the *route* from *A* to *B*

TCP/IP *traceroute* program uses UDP to non-existent port and TTL field to find route via *expanding ring* search

traceroute must accommodate varying network delays & dynamically changing routes

Sends ICMP echo messages with increasing TTL

- Router that decrements TTL to 0, sends ICMP *time exceeded* message, with router's address as source address
- First, with TTL 1, gets to first router, which discards and sends time exceeded message
- Next, with TTL 1, gets through first router to second router
- Continue until message from destination received

ICMP and path MTU (smallest accepted probe) discovery

Fragmentation should be avoided
How can source configure outgoing datagrams to avoid fragmentation?
Source determines *path MTU* - smallest network MTU (Minimum Transmission Unit) on path from source to destination
Source *probes* path using IP datagrams with *don't fragment* flag
Router responds with ICMP *fragmentation required* message
Source sends smaller probes until destination reached

ICMP and router discovery

Router can fail, causing "black-hole" or isolating host from internet

- •ICMP *router discovery* used to find new route
- •Host can broadcast request for router announcements to auto-configure default route
- •Host can broadcast request if router fails
- •Router can broadcast advertisement of existence when first connected

ICMP redirect

Default route may cause *extra hop*

- •Router that forwards datagram on same interface sends ICMP *redirect*
- •Host installs new route with correct router as next hop

ICMP packet format

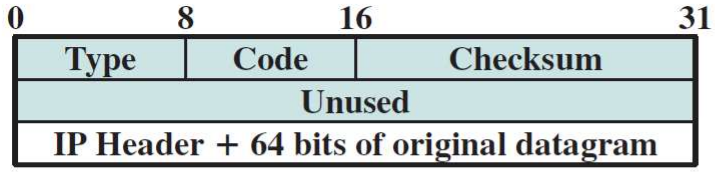

(a) Destination unreachable; time exceeded; source quench

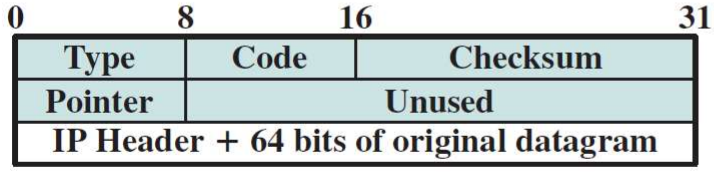

(b) Parameter problem

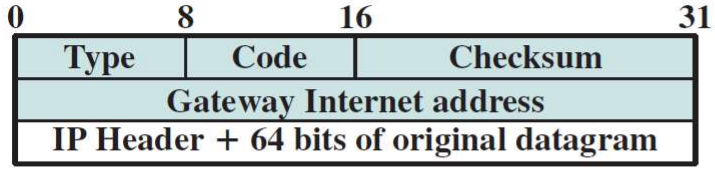

(c) Redirect

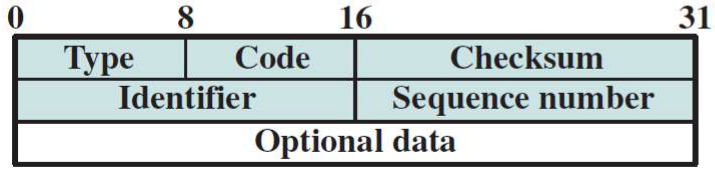

(d) Echo, echo reply

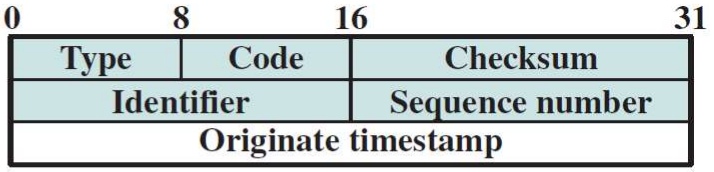

(e) Timestamp

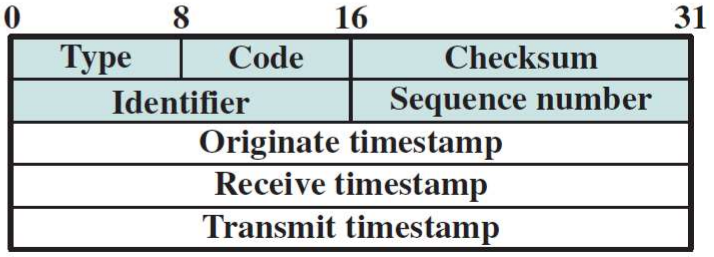

(f) Timestamp reply

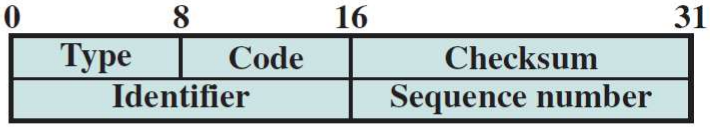

(g) Address mask request

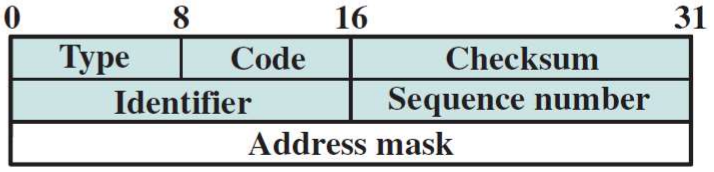

(h) Address mask reply

Proposed Problem

A company owns a building with one floor, having the following structure: **CC**-Communication Center, **O1**, **O2**-offices, **MS**-management, **A**- administration (see drawing). In room **MS** there are 2 computers, in room **A** are 5 computers, and the rooms **O1**, **O2** have 10 computers each.

The company gets for use the addresses 198.188.77.64 with the netmask 255.255.255.224.

Establish the address configuration of each subnetwork and computer located at this floor, knowing that in the both rooms **O** will be **one** subnetwork, the rooms **MS** and **A** belong to **one** other subnetwork , and **another** subnet will be setup for future applications in **CC**.

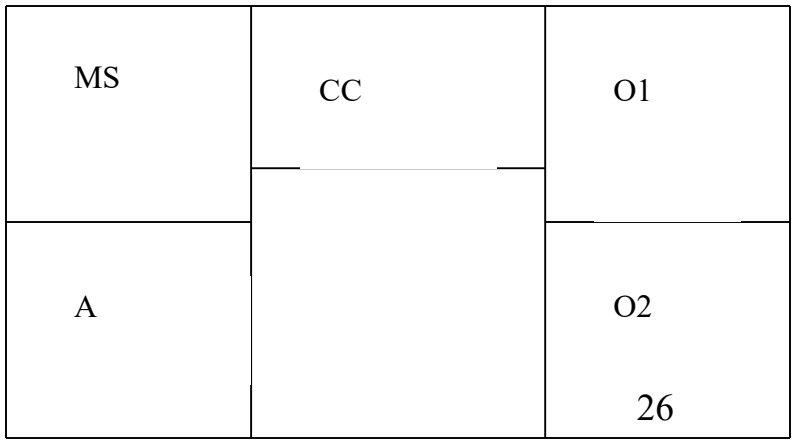

A *class C network* will be implemented.
Total number of *required subnetworks* will be *3*.

Bitwise AND of address and mask values (giving the resultant network/ subnet number) will be:

198.188.77.64

So the first subnet will have address: 198.188.77.64 and may serve rooms O1 and O2, with 20 stations together, filling address scheme from 198.188.77.65 to 198.188.77.84

The second subnet will have address: 198.188.77.128 and may serve rooms A+MS, with 7 computers, with addresses from 198.188.77.129 to 198.188.77.135

The third subnet with address 198.188.77.192 may serve the rest of the rooms.

Introduction to IPv6

IP v 1-3: defined and replaced

IP v4 - current version; 20+ years old

IP v5 - streams protocol

IP v6 - replacement for IP v4

During developments it was called IPng - Next Generation

Why Change IP?

32 bit Address space exhaustion

32 bit address space = millions of networks (could be enough?), BUT:

Two level addressing (network and host) wastes space: one network address used, even if not all possible associated hosts connected to Internet, or network connected to Internet

Growth of networks and the Internet (LANs, wireless LANs …)

2^{14} Class B network addresses already almost exhausted; class C networks too low size for most companies

Extended use of TCP/IP (new applications => requests for new IP addresses)

Requirements for new types of services

Different applications have different requirements for delivery, reliability and speed

Current IP has *type of service* that's not often implemented

Multicast transmissions

IPv6 RFCs

1752 - Recommendations for the IP Next Generation Protocol

2460 - Overall specification

4291 - addressing structure

….. Others

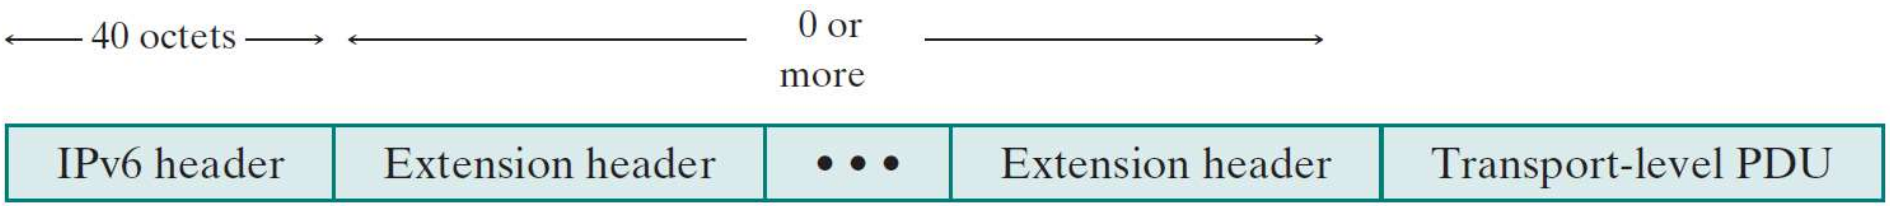

Enhancements over IPv4

Expanded address space: 128 bit

Improved option mechanism

- Separate optional headers between IPv6 header and Transport layer header
- Most are not examined by intermediate routers
  - Improved speed and simplified router processing of IPv6 datagrams
  - Easier to extend options

Address auto-configuration

- Dynamic assignment of addresses

Increased address flexibility & scalability

- Anycast address : packet delivered to a set of hosts

Support for resource allocation

- Allow packet labeling (those belonging to a traffic flow)

General considerations

Not generally compatible with IPv4

- But compatible with higher-level protocols

Longer addresses: expanded address space, 128 bit

Address auto-configuration; dynamic assignment of addresses

Traffic Priorities: 0 – 7 for variable flow rate, 8 – 15 for real time traffic

Improved option mechanism

- Separate optional headers between IPv6 header and transport layer header
- Most are not examined by intermediate routers
  - Improved speed and simplified router processing
  - Easier to extend options

Increased addressing flexibility

- Anycast - delivered to one of a set of nodes
- Improved scalability of multicast addresses

Support for resource allocation

Replaces *type of service* IPv4 field (most unused)

Labeling of packets to particular traffic flow

Allows special handling, e.g. real time video

IPv6 Header General format

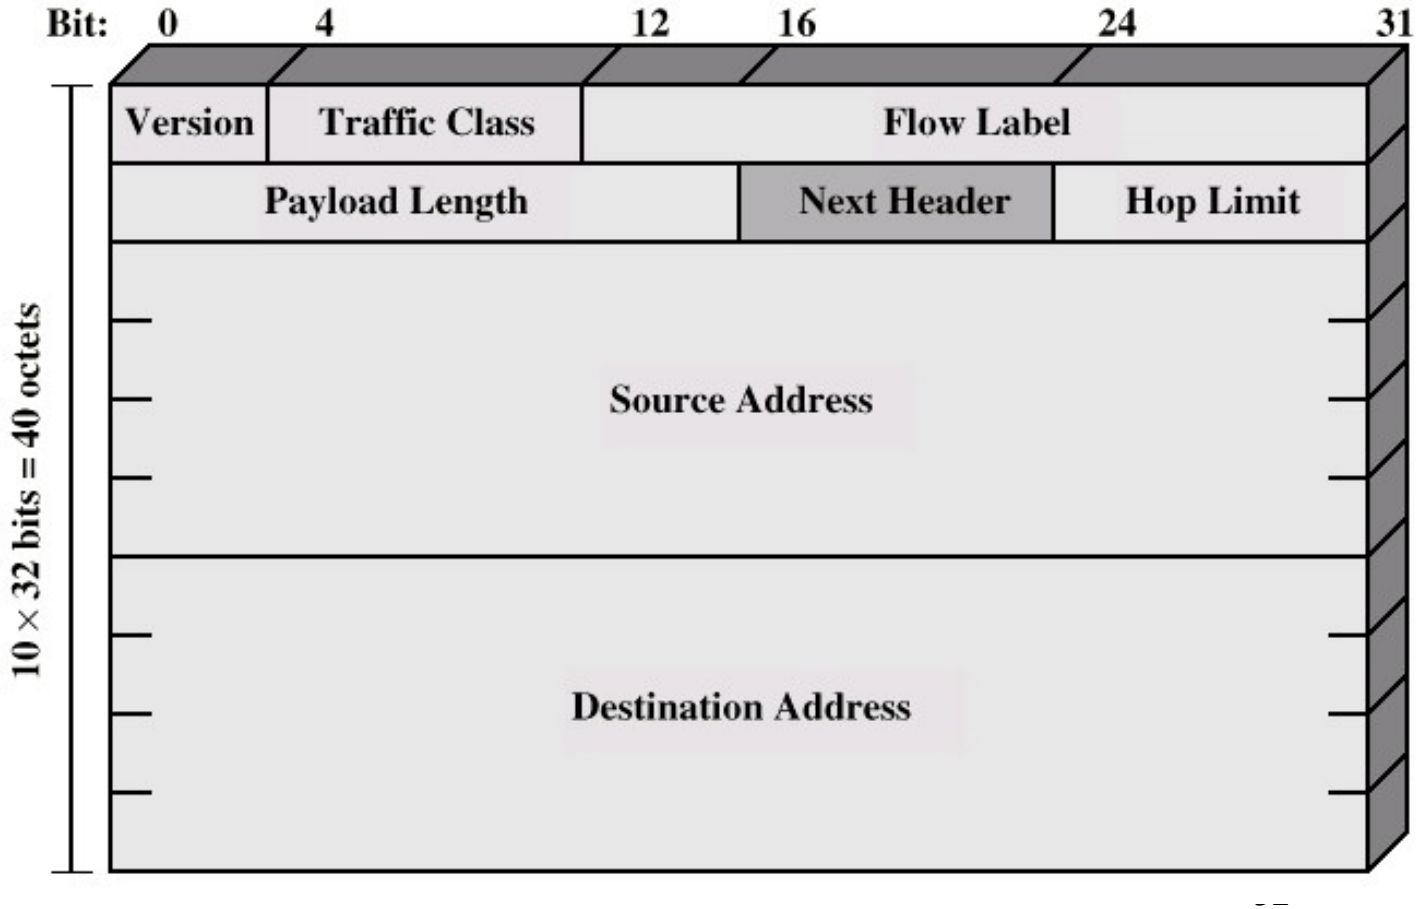

IPv6 base header format

Contains less information than IPv4 header; header format - entirely different

Next Header field points to first extension header

Flow Label field used to associate datagrams belonging to a *flow* or communication between two applications (support for audio-video connections, with appropriate QoS)

Traffic class

Classes or priorities of current packet

Still under development, see RFC 2460

Advantages of the header structure

Efficiency - header only as large as necessary

Flexibility & extension - can add new headers for new features

Incremental development - can add processing for new features to testbed; other routers will skip those headers

Extension Headers

Additional information stored in optional extension headers, followed by data

Hop-by-Hop Options

Require processing at each router

Routing

Similar to IPv4 source routing

Fragment

Authentication

Encapsulating security payload

Destination options

For destination node

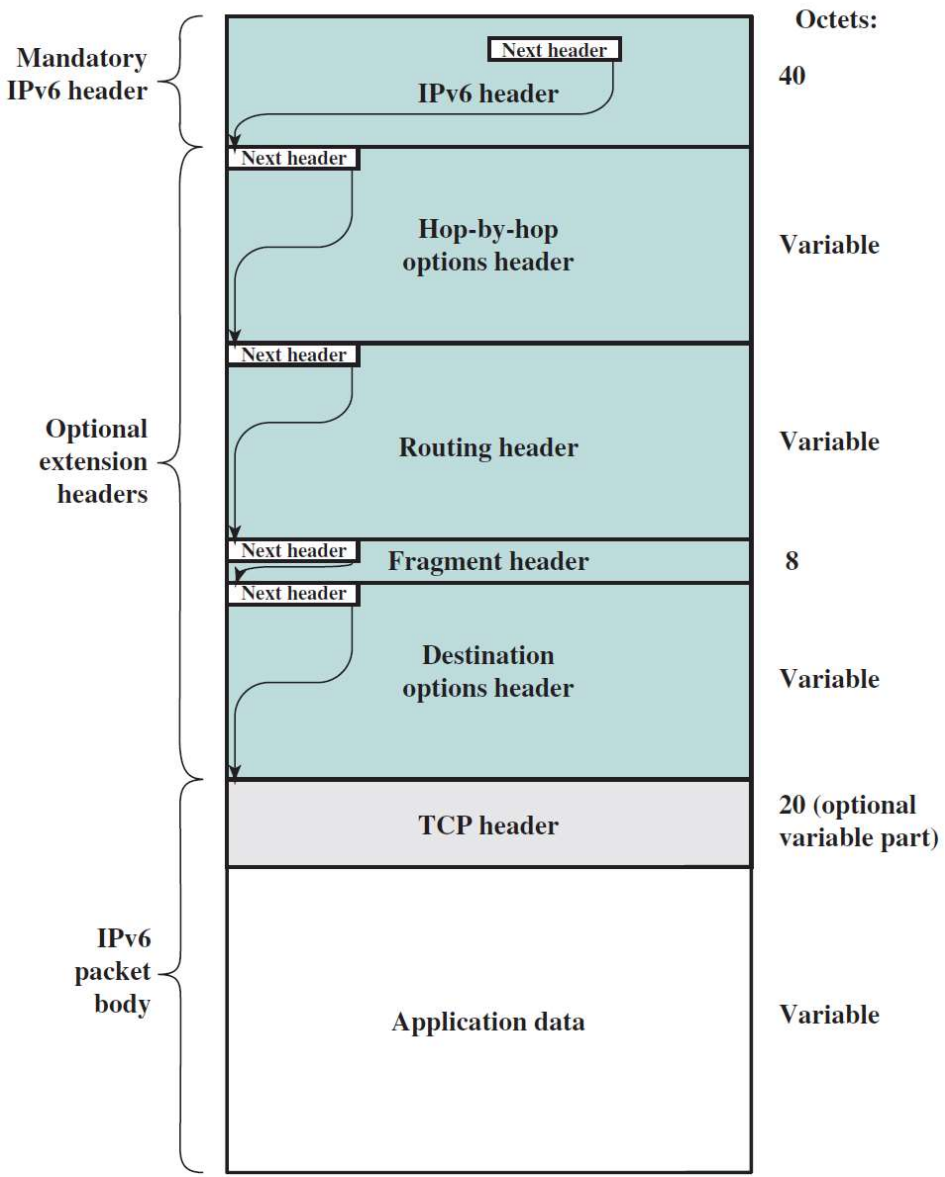

IPv6 Packet with Extension Headers

IP v6 Header Fields (Basic Header)

Version: 6

Traffic Class (DS/ECN)

Classes or priorities of packets

Used by originating nodes and/or forwarding routers for differentiated services and congestion functions

Flow Label

Used by routers to label packets requesting special handling within the network

Payload length

Includes all extension headers plus user (application) data

Next Header

Identifies type of header immediate following

May have another extension header or next layer up protocol header

Source Address & *Destination address* (128 bits)

IPv6 Addresses

128 bits long

Assigned to node's interface, not to node

One single interface may have multiple unique unicast addresses

Three types of address:

Unicast

Single interface; packet delivered there

Anycast

Set of interfaces (typically belong to different nodes)

Delivered to any **one** interface, usual the "nearest"

Multicast

Packets delivered to all interfaces identified

Prefix (binary)	Usage	Fraction
0000 0000	Reserved (including IPv4)	1/256
0000 0001	Unassigned	1/256
0000 001	OSI NSAP addresses	1/128
0000 010	Novell NetWare IPX addresses	1/128
0000 011	Unassigned	1/128
0000 1	Unassigned	1/32
0001	Unassigned	1/16
001	Unassigned	1/8
010	Provider-based addresses	1/8
011	Unassigned	1/8
100	Geographic-based addresses	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 use addresses	1/1024
1111 1110 11	Site local use addresses	1/1024
1111 1111	Multicast	1/256

IPv6 addresses

Address Type	Description
Unicast	One to One (Global, Link local, Site local) + An address destined for a single interface.
Multicast	One to Many + An address for a set of interfaces + Delivered to a group of interfaces identified by that address. + Replaces IPv4 “broadcast”
Anycast	One to Nearest (Allocated from Unicast) + Delivered to the closest interface as determined by the IGP

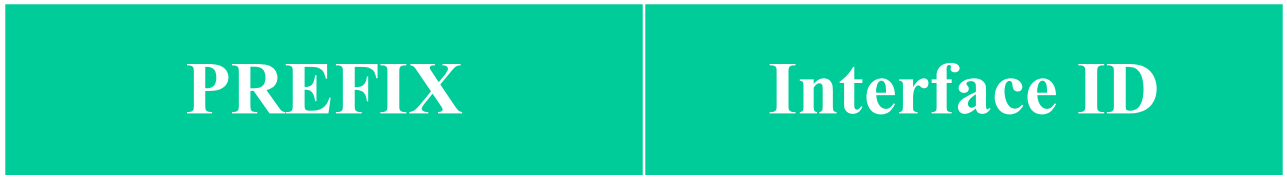

IPv6 Address Space Usage

Address Type	Binary Prefix	IPv6 Notation	Fraction of Address Space
Embedded IPv4 address	00…1111 1111 1111 1111 (96 bits)	::FFFF/96	2^{-96}
Loopback	00…1 (128 bits)	::1/128	2^{-128}
Link-local unicast	1111 1110 10	FE80::/10	1/1024
Multicast	1111 1111	FF00::/8	2/256
Global unicast	Everything else		

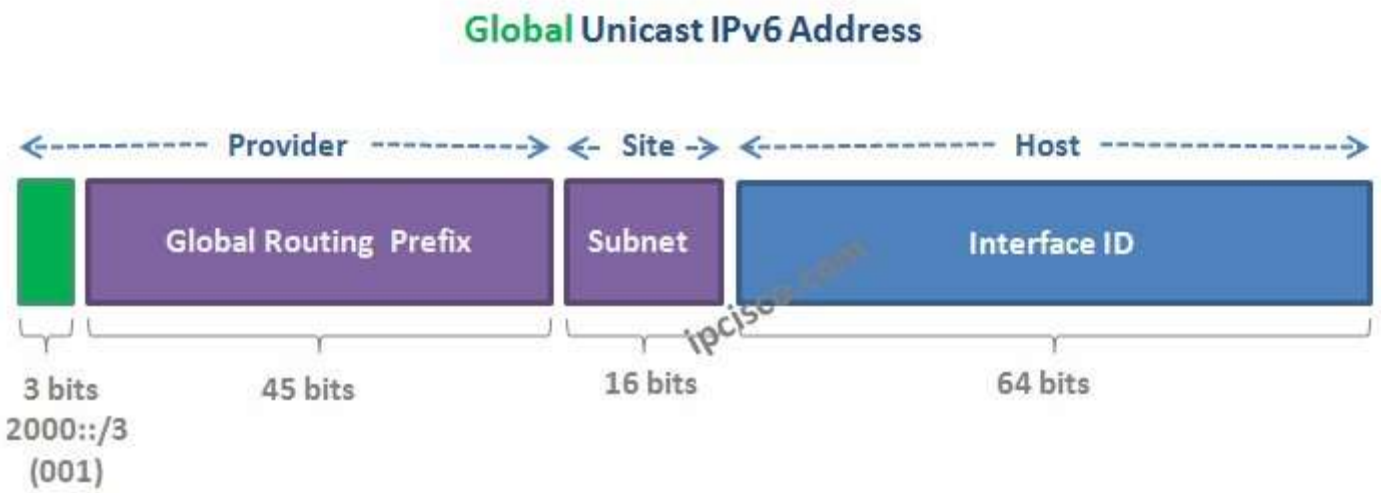

IPv6 addresses format:

•128-bit addresses, may use dotted decimal representation; requires 16 numbers

105.220.136.100.255.255.255.255.0.0.18.128.140.10.255.255

•Groups of 16-bit numbers in hex, separated by colons - *colon hexadecimal* (or *colon hex*) representation (not case sensitive)

69DC:8864:FFFF:FFFF:0:1280:8C0A:FFFF

•Zero-compression - series of zeroes indicated by two colons

FF0C:0:0:0:0:0:0:B1 is equivalent with: FF0C::B1

But once in an address

•IPv6 address with 96 leading zeros is interpreted to hold an IPv4 address

•Use of " / " notation to denote number of bits in address represent the subnet (prefix); rest of them represent interface ID);

• /64 is common prefix length

ICANN assigns a range of IP addresses to Regional Internet Registry (RIR) organizations. The size of address range assigned to the RIR may vary but with a minimum prefix of /12 and belong to the following range: 2000::/12 to 200F:FFFF:FFFF:FFFF::/64.

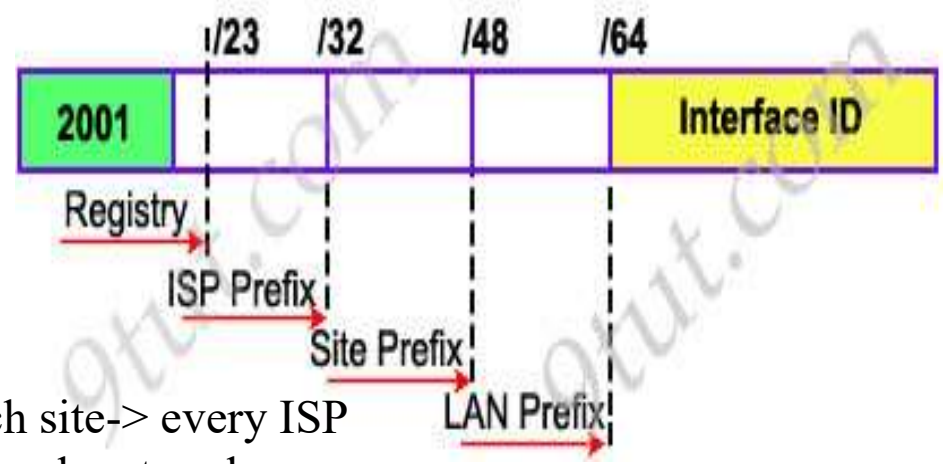

Each ISP receives a /32 and provides a /48 for each site-> every ISP can provide $2^{(48-32)}$ = 65,536 site addresses (note: each network organized by a single entity is often called a site).
Each site provides /64 for each LAN -> each site can provide $2^{(64-48)}$ = 65,536 LAN addresses for use in their private networks.
So each LAN can provide 2^{64} interface addresses for hosts.

EXAMPLE:

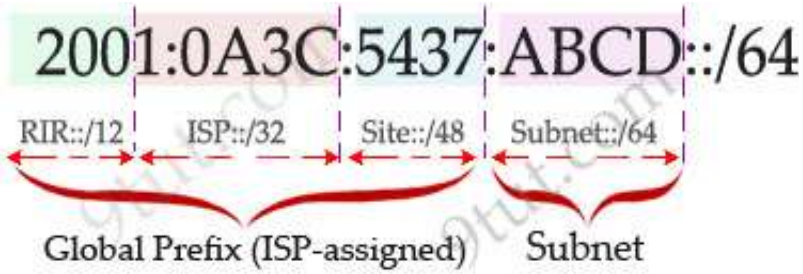

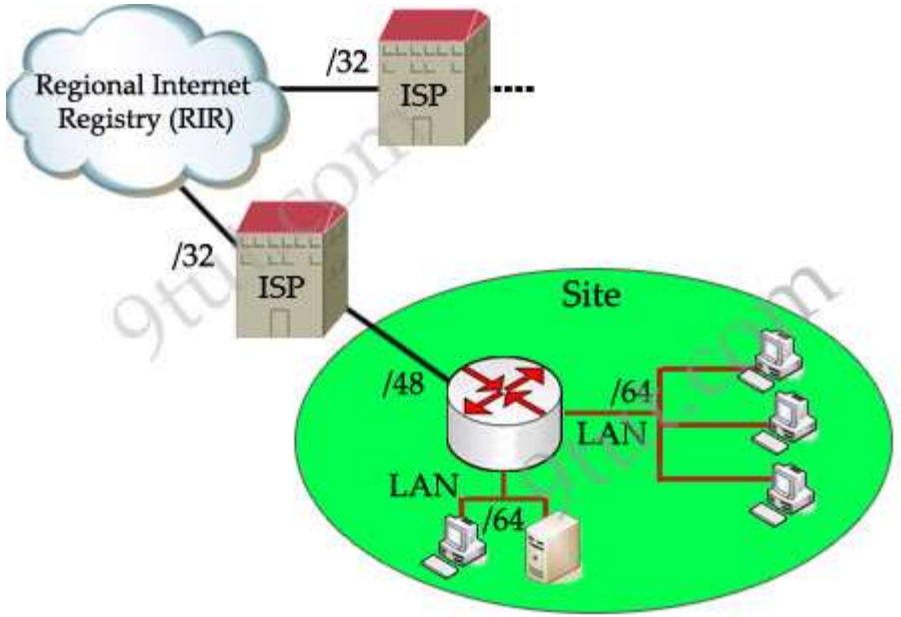

IPv6 Address Scopes	*Description*
Link-local address	+ only used for communications within the local subnetwork (automatic address configuration, neighbor discovery, router discovery, and by many routing protocols). Only valid on the current subnet. + routers do not forward packets with link-local addresses. + are allocated with the FE80::/64 prefix -> can be easily recognized by the prefix FE80. + is usually created dynamically using a link-local prefix of FE80::/10 and a 64-bit interface identifier (based on 48-bit MAC address).
Global unicast address	+ unicast packets sent through the public Internet + globally unique throughout the Internet + starts with a 2000::/3 prefix (this means any address beginning with 2 or 3). But in the future global unicast address might not have this limitation
Site-local address	+ allows devices in the same organization, or site, to exchange data. + starts with the prefix FEC0::/10. They are analogous to IPv4′s private address classes.

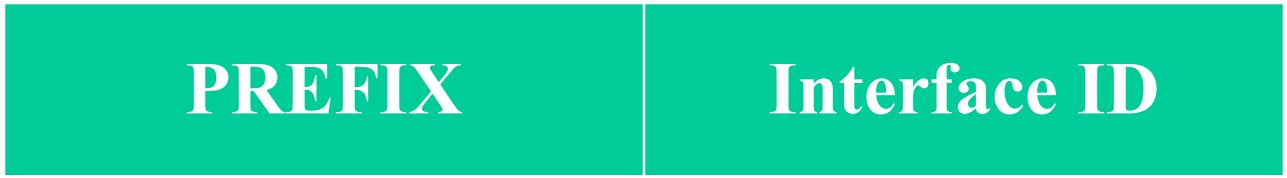

Options Headers

Carry optional information, not necessary examined by all routers or hosts

Hop-by-Hop Options Header

Consists of the following:

Next header

Header extension length

Options

One or more option definitions

Options definition contains the following sub-fields:

Option Type – identifies option

Length – length in octets of the option's data field

Option Data – option specification

Examples for such options:

Jumbo payload option

Over $2^{16} = 65{,}535$ octets in an IPv6 packet

Router alert option

Tells the router that the contents of this packet is of interest to the router (to handle data accordingly)

Provides support for RSVP (Reservation Protocol), used in multimedia transmissions, for flow control

Fragmentation Header

Fragmentation only allowed at source node, no fragmentation at intermediate routers

Node must perform path discovery operation, to find the smallest MTU (Maximum Transmission Unit) of intermediate networks

Source fragments IPv6 packets to match MTU

Otherwise limits to 1280 octets, that must be supported by any network

Fragmentation Header Fields:

Next Header – type of following header

Reserved

Fragmentation offset – any fragment data is multiple of 64bits; this field indicates where in the original packet this fragment's payload belongs

Reserved

More flag – more fragments or last fragment

Identification – identify the original packet (now fragmented)

Routing Header

List of one or more intermediate nodes to be visited

Structure (see next):

Next Header

Header extension length – length of this header

Routing type – identifies a routing protocol header variant

Segments left - number of nodes still to be visited

Destination Options Header

Information examined by the destination node

Same format as Hop-by-Hop options header

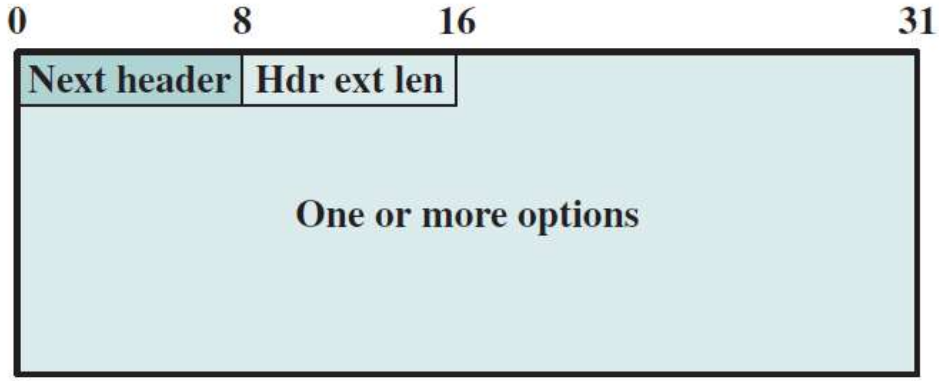

(a) Hop-by-Hop Options header;
Destination Options header

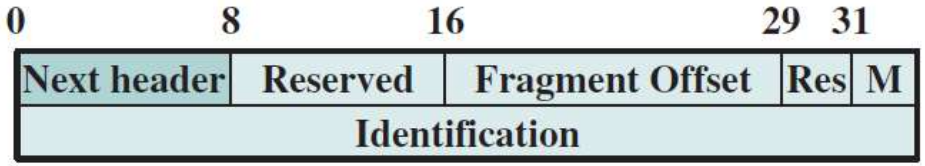

(b) Fragment header

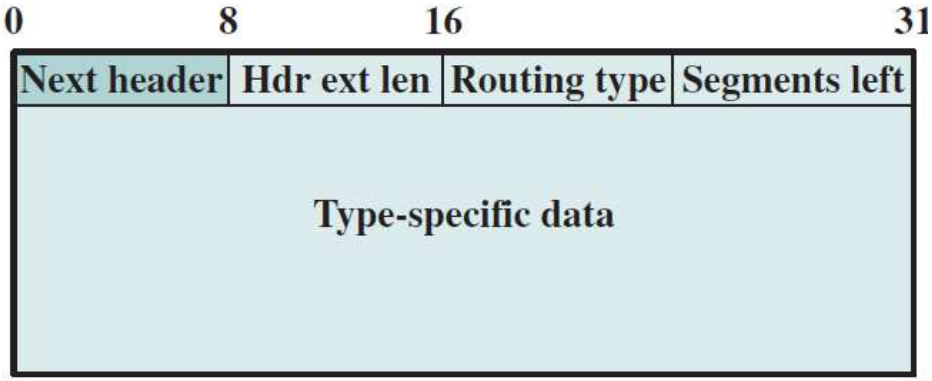

(c) Generic Routing header

IPv6 Extension Headers

Technologies can be used in transition from IPv4 to IPv6

Dual Stack Routers

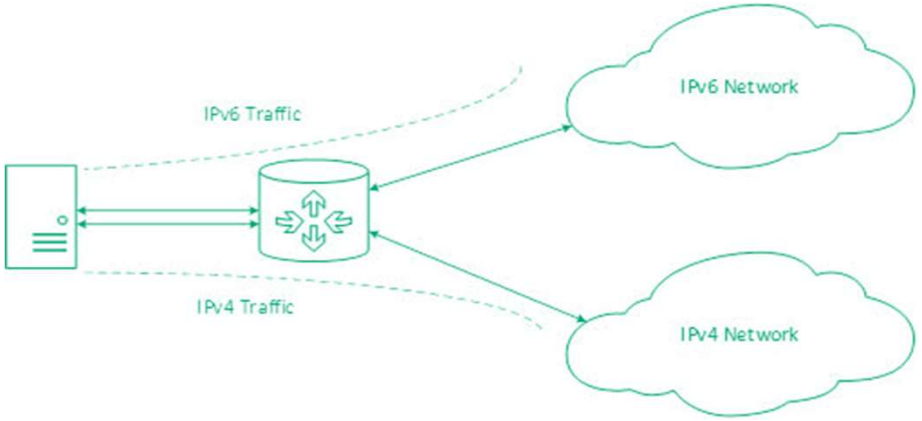

Tunneling

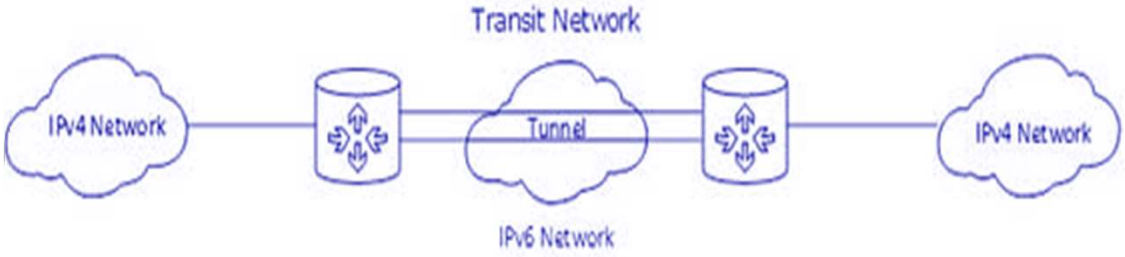

NAT Protocol Translation

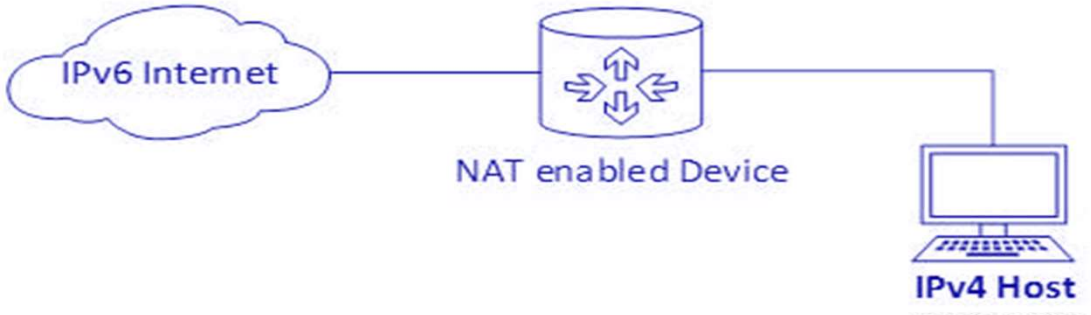

Multicasting

Addresses that refer to group of hosts on one or more networks

Used in:

Multimedia stream "broadcasts"

Teleconferencing – a transmission from a station sent to all members

Database updating – for all copies of replicated files or databases

Distributed computing – resource sharing

Real time workgroups – real time exchange of information

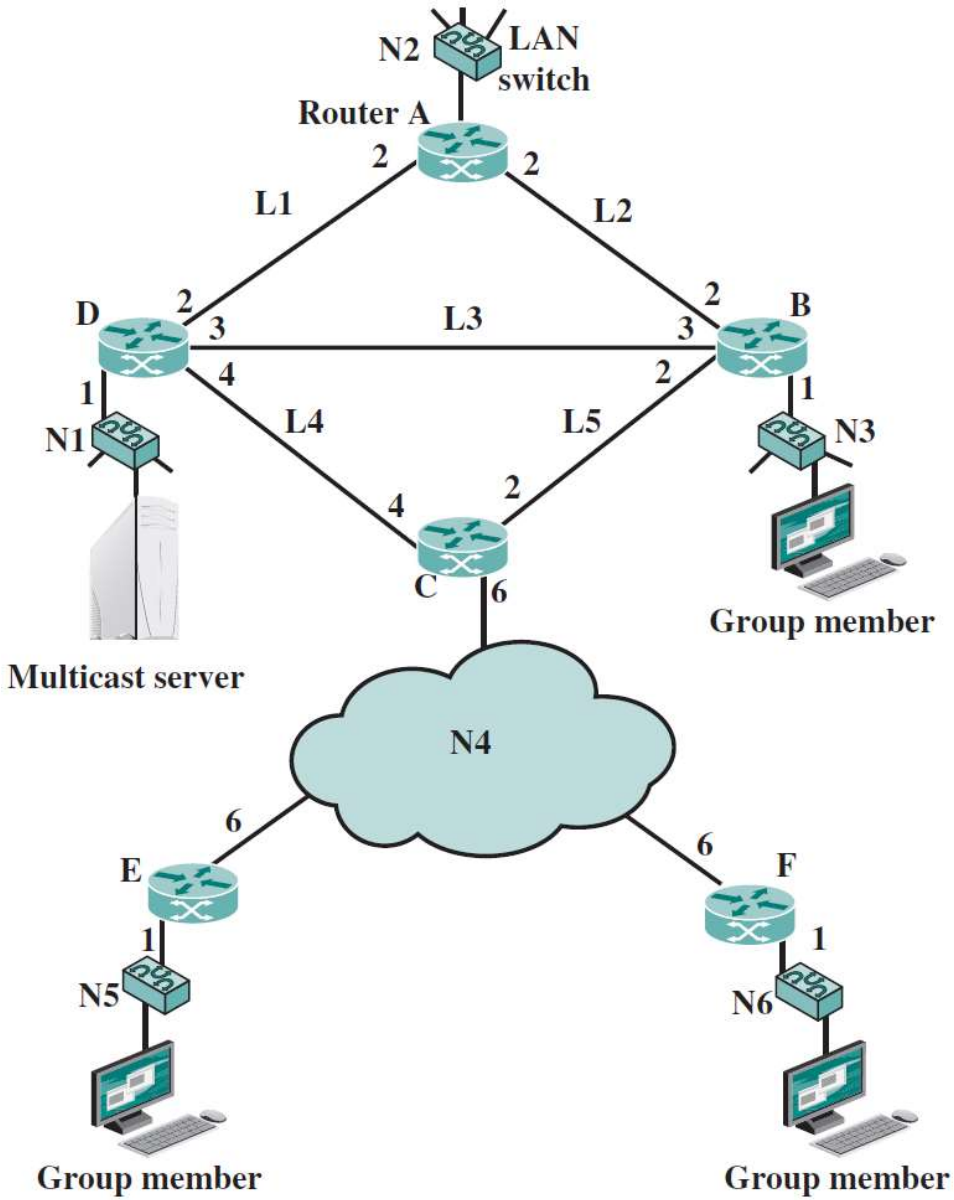

Multicast Configuration

Multicast over a single Ethernet LAN segment is straightforward; provision for MAC multicast addresses, due to the broadcast nature of LAN

For Internet environment, more approaches:

Broadcast and Multiple Unicast

Broadcast a copy of packet to each network, even if does not contain group members

> For figure behind, multicast server sends a packet to group hosts from networks N3, N5, N6: requires 13 copies of the packet

Multiple Unicast

> Send packet only to networks that have hosts in group
>
> Source knows location for each group member
>
> 11 packets

True Multicast

Use of following algorithm:

Determine least cost path to each network that has host in group

Gives spanning tree configuration containing networks with group members

Transmit single packet along spanning tree

Routers replicate packets at branch points of the spanning tree

8 packets required for above example

	Broadcast					Multiple Unicast				Multicast
	S → N2	S → N3	S → N5	S → N6	Total	S → N3	S → N5	S → N6	Total	
N1	1	1	1	1	4	1	1	1	3	1
N2										
N3	1	1	1					1	1	
N4			1	1	2		1	1	2	2
N5			1		1		1		1	1
N6			1	1			1	1	1	
L1	1			1						
L2										
L3		1			1	1			1	1
L4			1	1	2		1	1	2	1
L5										
Total	2	3	4	4	13	3	4	4	11	8

Multicast problems:

Router may have to forward more than one copy of packet (multiple output branches)

Convention needed to identify multicast addresses

IPv4 - Class D – starts with 1110 …

IPv6 - 8 bit prefix, all 1s, 4 bit flags field, 4 bit scope field, 112 bit group ID

Nodes (routers & source) must translate between IP multicast addresses and a list of networks containing group members; allows tree development

Multicast transmission example

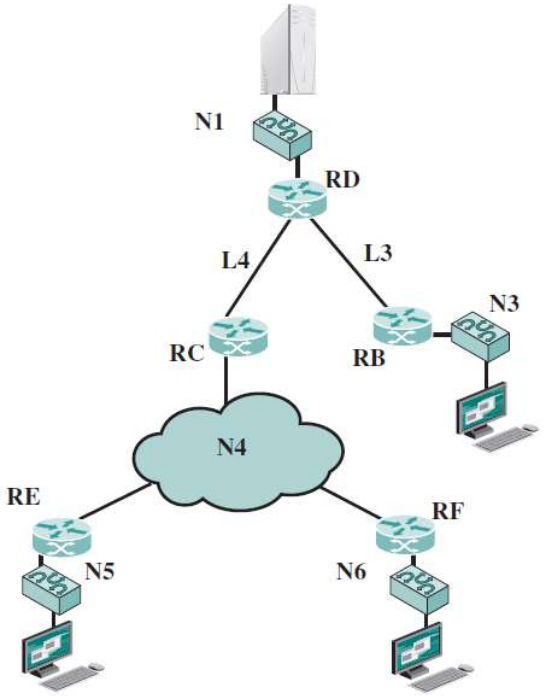

(a) Spanning tree from source to multicast group

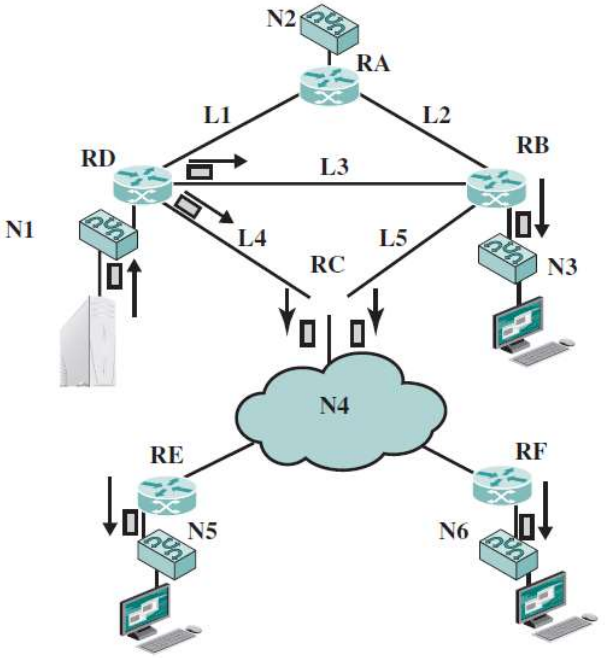

(b) Packets generated for multicast transmission

Router must translate between IP multicast address and a network LAN multicast address (at the MAC level) in order to deliver packet to that LAN

Mechanism required for hosts to dynamically join and leave multicast groups

Routers must exchange info

- Which networks include members of given group
- Sufficient info to work out shortest path to each network (spanning tree)
- Routing algorithm to work out shortest path
- Routers must determine routing paths based on source and destination addresses, for avoiding packet duplication

IGMP (Internet Group Management Protocol)

RFC 1112, initial developed for IPv4, but incorporated also in ICMPv6

Host and router exchange of **multicast group information**

Use broadcast LAN to transfer information among multiple hosts and routers

IGMP Fields

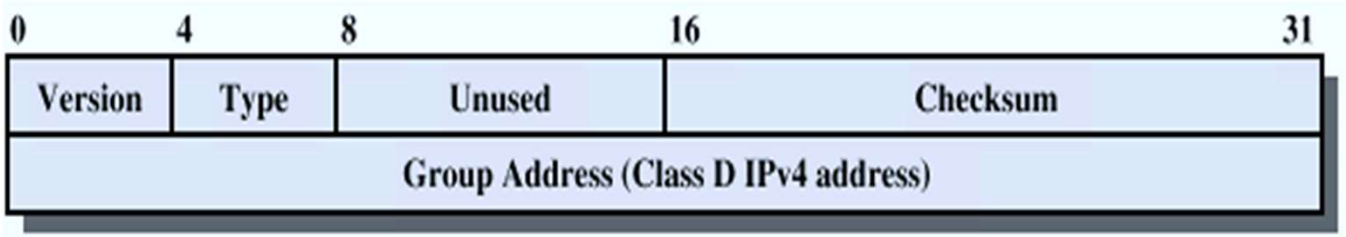

Version - 1

Type

1 - query sent by a multicast router

0 - report sent by a host

Checksum – 16 bit ones complement addition of all the 16-bit words in the message

Group address

Zero value in a request message

Valid group address in a report message

IGMP Operation

To join a group, hosts sends report message

- Group address of group to join
- Sent in a IP datagram with the same multicast destination address
- All hosts in group receive message and learn new member
- Routers listen to all multicast addresses to hear all reports

Routers periodically issue request messages (queries)

- Sent to *all-hosts* multicast address
- Host that want to stay in groups must read *all-hosts* messages and respond with report for each group it is in

Group Membership with IPv6

Function of IGMP included in ICMP v6; ICMPv6 contains a new type of message: group membership **termination** message, to allow host to leave the group

IGMP v3

RFC 3376

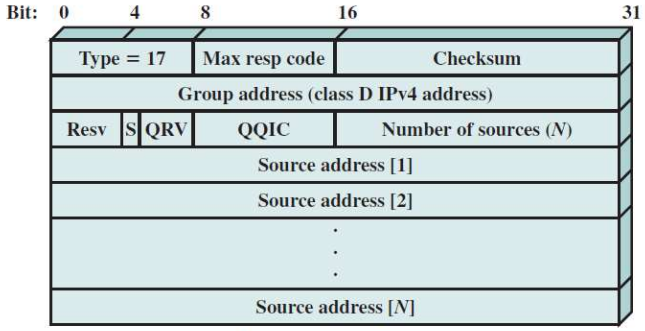

(a) Membership query message

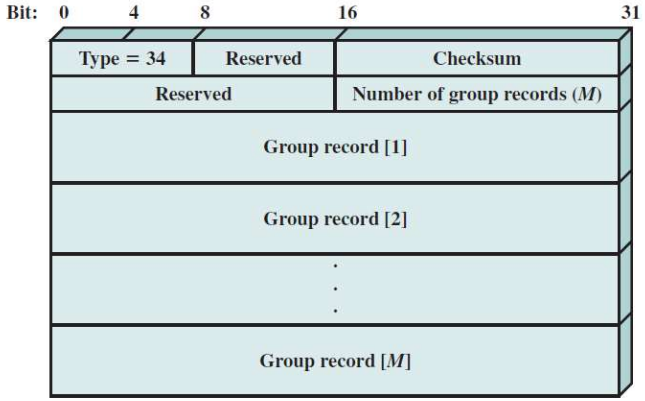

(b) Membership report message

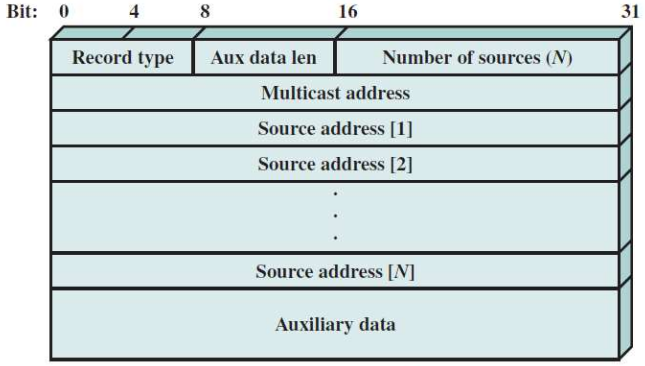

(c) Group record