Advance Networking

Unit- 1

-Madhavi Dave

Introduction

- The pre-requisite for Advanced Networking is the basic knowledge of TCP/IP reference model and Protocol Layering.
- TCP/IP internet as a virtual network built by interconnecting physical networks with routers.
- This chapter discusses addressing, an essential ingredient that helps TCP/IP software hide physical network details and makes the resulting internet appear to be a single, uniform entity.

Universal Host Identifiers

- TCP/IP uses tern *host* to refer to the end system that attaches to internet.
- A host can be small/large, general purpose/special purpose computer.
- A host may have user interface or embedded system for input/output process.

Universal Host Identifiers

- The internet divides all machines into two classes:
  - Router
  - Host
- To make the communication system universal, it needs global identifying system for each host.
- These identifiers are classified as names, addresses or routes.

Universal Host Identifiers

- In general, human prefers to use pronounceable names to identify computers, while software works more efficiently on binary identifiers.
- Thus the decision was made to standardize on compact binary address.
- The internet is considered as a large network; just the difference is that the internet is a virtual structure.

Universal Host Identifiers

- Because the internet is virtual, the designers are free to choose packet size, address, delivery techniques, etc.
- The designers of TCP/IP chose a scheme analogous to physical network addressing in which each host is assigned a unique number called *Internet Protocol Address*.

Universal Host Identifiers

- Usually, computers communicate through the Internet. The packet transmitted by the sending computer may pass through several LANs or WANs before reaching the destination computer.
- For this level of communication, we need a ***global addressing scheme;*** we called this ***logical addressing.***
- The term IP address to mean a logical address in the network layer of the TCP/IP protocol suite.

Universal Host Identifiers

- The Internet addresses are 32 bits in length; this gives us a maximum of 2^{32} addresses. These addresses are referred to as IPv4 (IP version 4) addresses or simply IP addresses.
- The need for more addresses, in addition to other concerns about the IP layer, motivated a new design of the IP layer called the new generation of IP or IPv6. In this version, the Internet uses 128-bit addresses that give much greater flexibility in address allocation

IPv4 ADDRESSES

- The identifier used in the IP layer of the TCP/IP protocol suite to identify the connection of each device to the Internet is called the Internet address or IP address.
- An IPv4 address is a 32-bit address that ***uniquely and universally*** defines the connection of a host or a router to the Internet.
- The IP address is the ***address of the connection, not the host or the router.***

Notations: Three different notations in IPv4 addressing. because each byte (octet) is 8 bits, each number in dotted-decimal notation is a value ranging from 0 to 255.

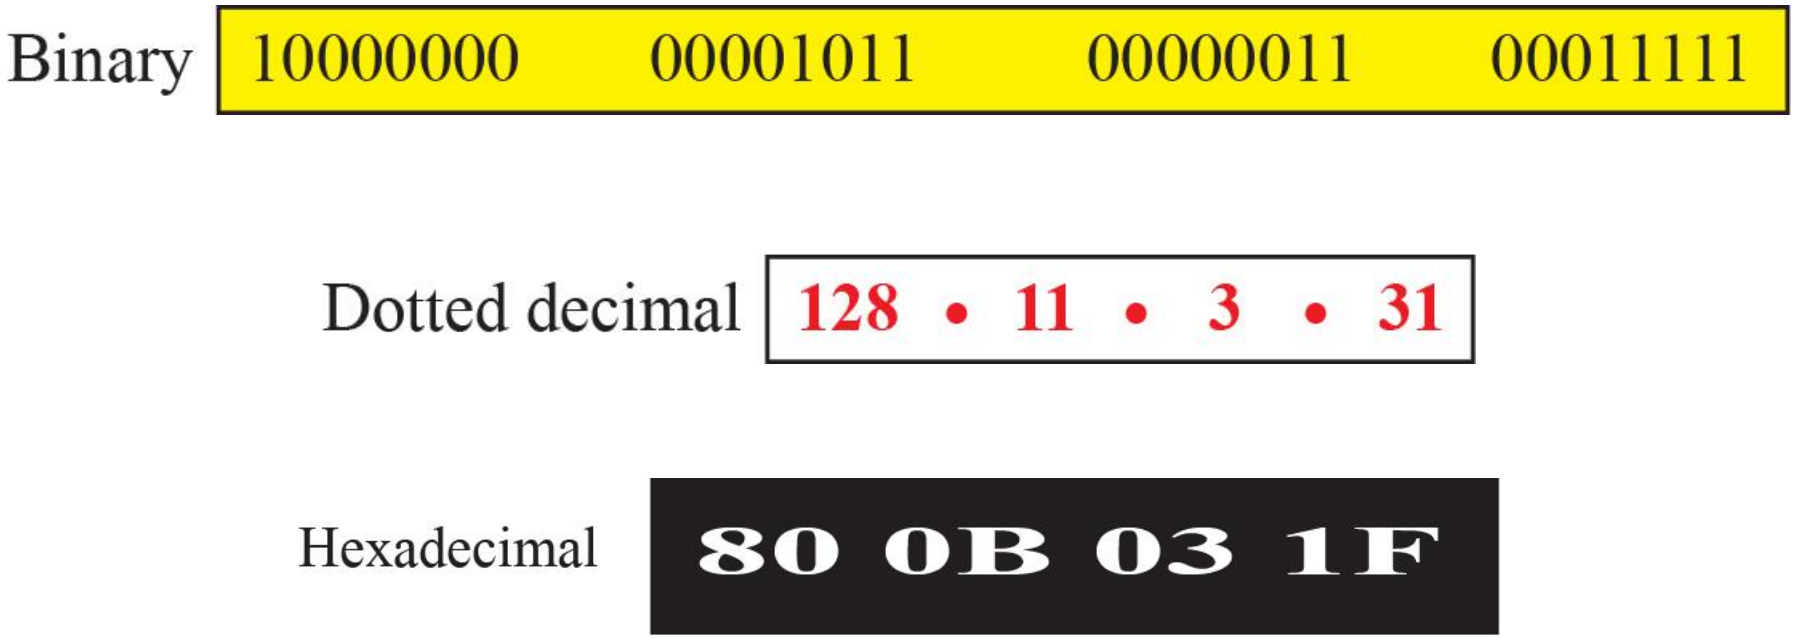

Classful Addressing

- When the Internet started, The whole address space was divided into five classes (class A, B, C, D, and E), as shown in.
- This scheme is referred to as classful addressing.
- Although classful addressing belongs to the past, it helps us to understand classless addressing

Classful Addressing

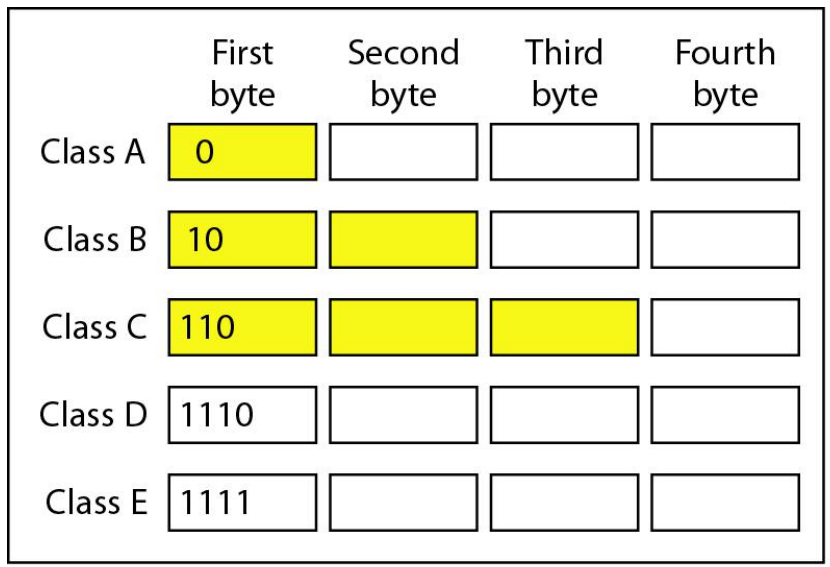

a. Binary notation

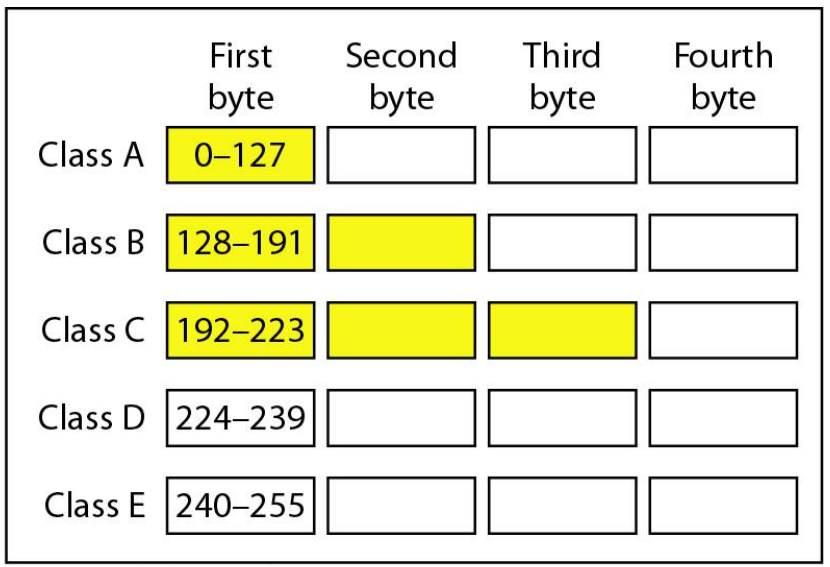

b. Dotted-decimal notation

Classful Addressing

- **Classes and Blocks:** One problem with classful addressing is that each class is divided into a fixed number of blocks with each block having a fixed size. Because of that a large part of the available addresses were wasted.

Class	*Number of Blocks*	*Block Size*	*Application*
A	128	16,777,216	Unicast
B	16,384	65,536	Unicast
C	2,097,152	256	Unicast
D	1	268,435,456	Multicast
E	1	268,435,456	Reserved

Netid and Hostid

- IP address in class A, B, or C is divided into netid and hostid.
- These parts are of varying lengths, depending on the class of the address.
- In class A, one byte defines the netid and three bytes define the hostid. In class B, two bytes define the netid and two bytes define the hostid. In class C, three bytes define the netid and one byte defines the hostid.

Hierarchy in addressing

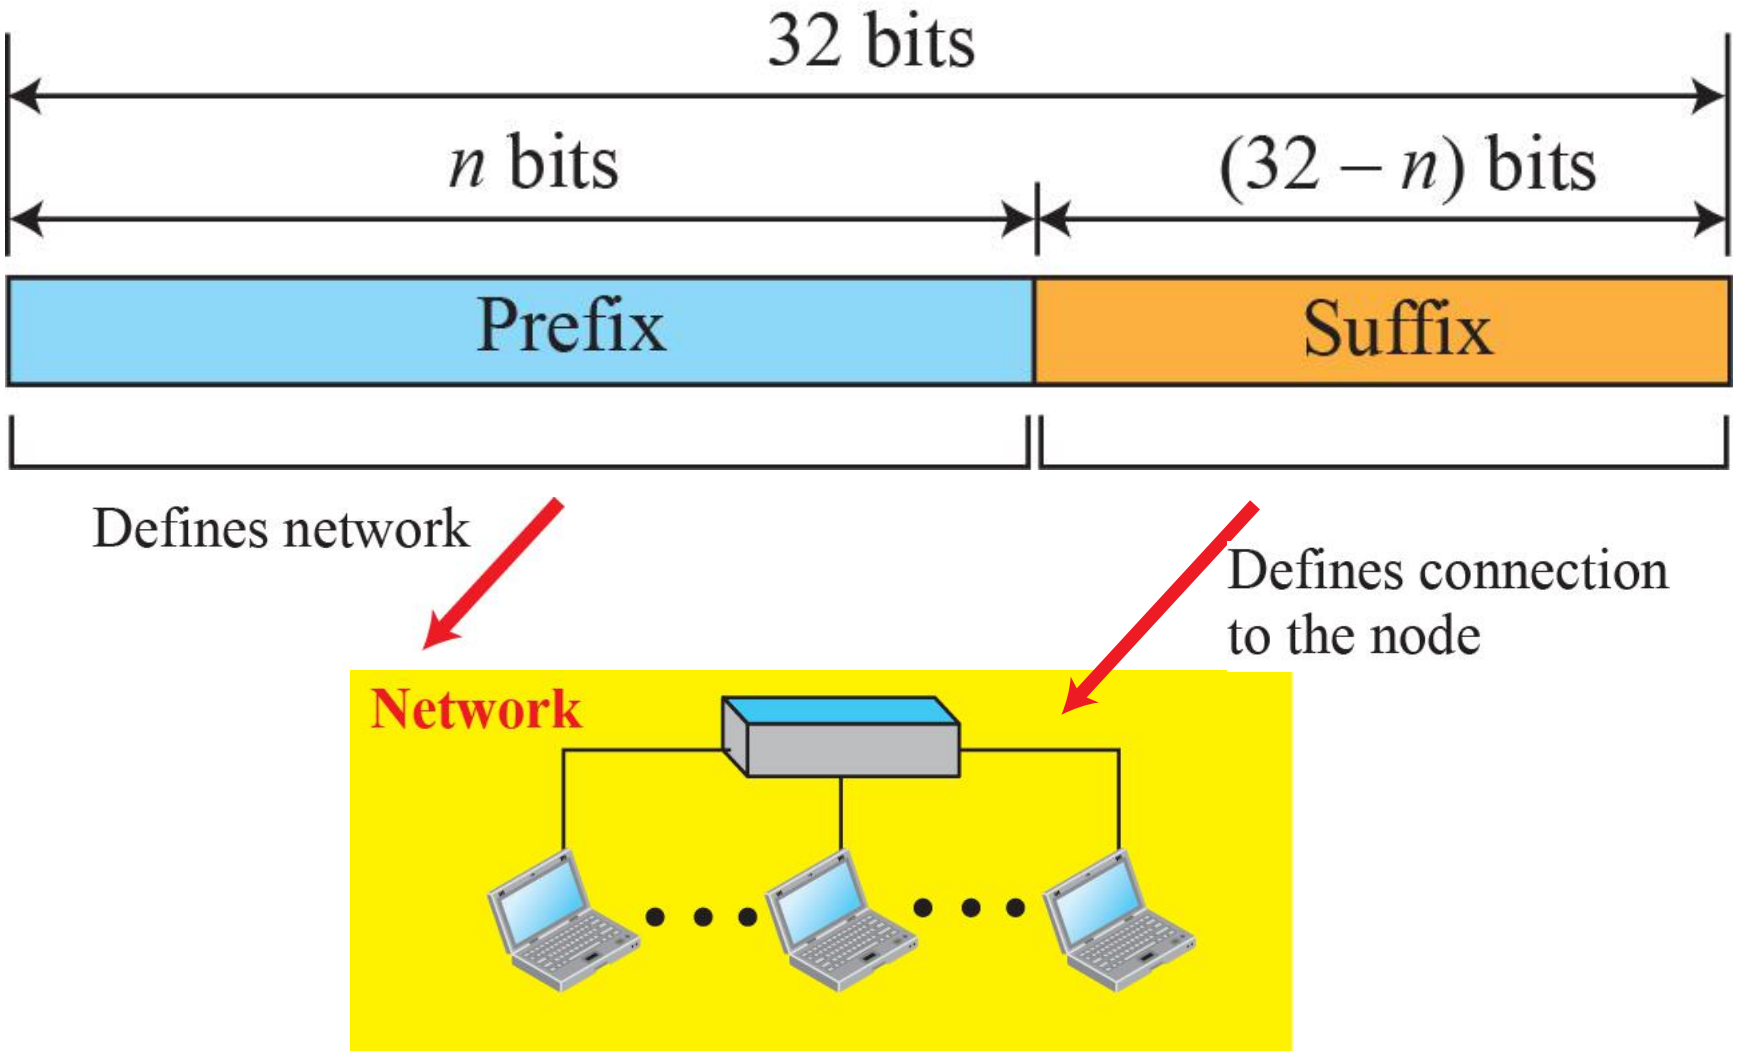

Occupation of the address space in classfull addressing

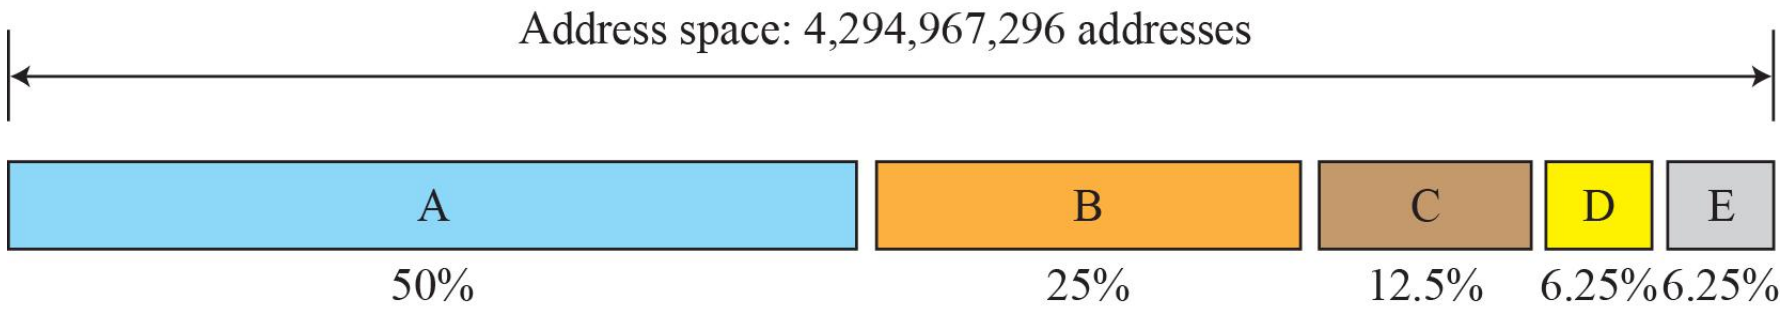

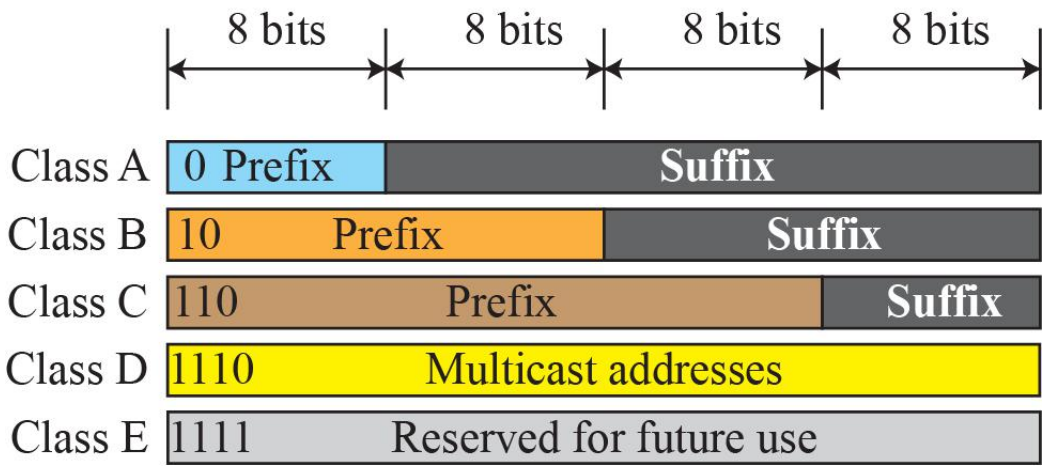

Class	Prefixes	First byte
A	n = 8 bits	0 to 127
B	n = 16 bits	128 to 191
C	n = 24 bits	192 to 223
D	Not applicable	224 to 239
E	Not applicable	240 to 255

Subnetting

- During the era of classfull addressing, subnetting was introduced. If an organization was granted a large block in class A or B, it could divide the addresses into several contiguous groups and assign each group to smaller networks (called subnets) or, in rare cases, share part of the addresses with neighbors.
- Subnetting allows a single network prefix to be used for multiple physical network.
- It allows the admin to divide host portion of their address into multiple networks.

Subnetting

- A 32 bit IPv4 address can be considered as internet portion and local portion.
- Internet address can be classified as netID and local address as hostID.
- For subnetting, local address will be divided into physical network and host.
- Following are two types of subnetting
  - Fixed length IPv4 subnet
  - Variable length IPv4 subnet

Masking

- ***Mask:*** we can also use a mask (also called the default mask), a 32-bit number made of contiguous 1s followed by contiguous 0s.
- Last column is also called slash notation or Classless Inter Domain Routing (CIDR) notation. The notation is used in classless addressing

Class	*Binary*	*Dotted-Decimal*	*CIDR*
A	**11111111** 00000000 00000000 00000000	**255**.0.0.0	/8
B	**11111111 11111111** 00000000 00000000	**255.255**.0.0	/16
C	**11111111 11111111 11111111** 00000000	**255.255.255**.0	/24

Subnet mask

- A 32 bit mask is used to divide internet, physical and host address of subnetting.
- The mask covers internet and physical address of local portion.
- The bits in subnet mask are set to 1 if the corresponding bit in IP address are subnet prefix; set to 0 otherwise.

Supernetting

- The time came when most of the class A and class B addresses were depleted; however, there was still a huge demand for midsize blocks. The size of a class C block with a maximum number of 256 addresses did not satisfy the needs of most organizations.
- Even a midsize organization needed more addresses. One solution was supernetting.

Supernetting

- In supernetting, an organization can combine several class C blocks to create a larger range of addresses. In other words, several networks are combined to create a supernetwork or a supemet.
- An organization can apply for a set of class C blocks instead of just one. For example, an organization that needs 1000 addresses can be granted four contiguous class C blocks.
- The organization can then use these addresses to create one supernetwork. Supernetting decreases the number of Is in the mask. For example, if an organization is given four class C addresses, the mask changes from /24 to /22.

Supernetting

- In supernetting, an organization can combine several class C blocks to create a larger range of addresses. In other words, several networks are combined to create a supernetwork or a supemet.
- An organization can apply for a set of class C blocks instead of just one. For example, an organization that needs 1000 addresses can be granted four contiguous class C blocks.
- The organization can then use these addresses to create one supernetwork. Supernetting decreases the number of Is in the mask.
- For example, if an organization is given four class C addresses, the mask changes from /24 to /22.

Address Depletion

- The flaws in classful addressing scheme combined with the fast growth of the Internet led to the near depletion of the available addresses. Yet the number of devices on the Internet is much less than the 2^{32} address space.
- We have run out of class A and B addresses, and a class C block is too small for most midsize organizations. One solution that has alleviated the problem is the idea of classless addressing.

Classless Addressing

- To overcome address depletion and give more organizations access to the Internet, classless addressing was designed and implemented.
- In this scheme, there are no classes, but the addresses are still granted in blocks.
- When an entity, small or large, needs to be connected to the Internet, it is granted a block (range) of addresses.

Classless Addressing

- The size of the block (the number of addresses) varies based on the nature and size of the entity.
- For example, a household may be given only two addresses; a large organization may be given thousands of addresses. An ISP, as the Internet service provider, may be given thousands or hundreds of thousands based on the number of customers it may serve.

Classless Addressing

- **Restriction:** To simplify the handling of addresses, the Internet authorities impose three restrictions on classless address blocks:

1. The addresses in a block must be contiguous, one after another.
2. The number of addresses in a block must be a power of 2 (I, 2, 4, 8, ...).
3. The first address must be evenly divisible by the number of addresses.

Classless Addressing

- Figure shows a block of addresses, in both binary and dotted-decimal notation, granted to a small business that needs 16 addresses.
- We can see that the restrictions are applied to this block. The addresses are contiguous. The number of addresses is a power of 2 ($16 = 2^4$), and the first address is divisible by 16. The first address, when converted to a decimal number, is 3,440,387,360, which when divided by 16 results in 215,024,210.

Classless Addressing

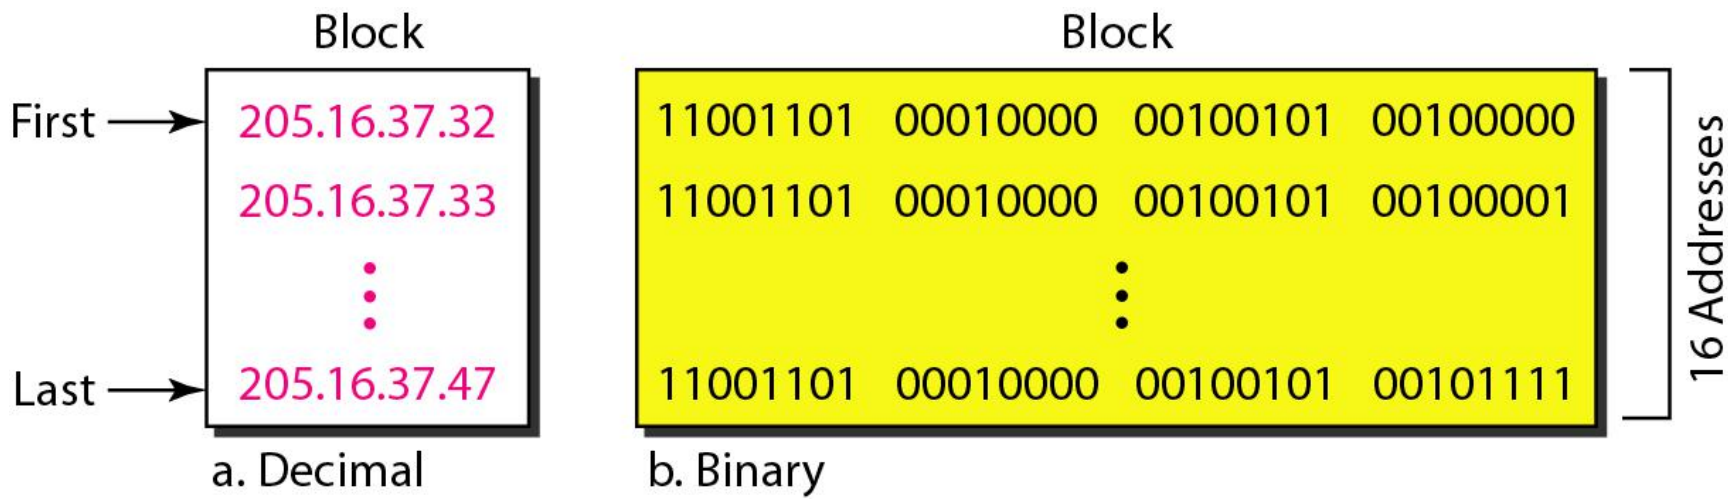

CIDR Notation

- Mask: A better way to define a block of addresses is to select any address in the block and the mask.
- A mask is a 32-bit number in which the n leftmost bits are 1s and the 32 - n rightmost bits are 0s.
- In classless addressing the mask for a block can take any value from 0 to 32. It is very convenient to give just the value of n preceded by a slash (CIDR notation).

CIDR Notation

- The address and the /n notation completely define the whole block. This is how you can find out first address, last address and total number of addresses in block.
  - The first address in the block can be found by setting the rightmost 32 - *n bits to 0s.*
  - The last address in the block can be found by setting the rightmost 32 - *n bits to 1s.*
  - The number of addresses in the block can be found by using the formula 2^{32-n}

Example

- A block of addresses is granted to a small organization. We know that one of the addresses is 205.16.37.12/27. Find out first address, last address and no. of addresses.
- The binary representation of the given address is
  - 11001101 00010000 00100101 00001100
- The value of n is 27, which means that number of addresses is 2^{32-27} or 32.

- Another way to find the first address, the last address, and the number of addresses is to represent the mask as a 32-bit binary (or 8-digit hexadecimal) number. In previous example the /28 can be represented as
- 11111111 11111111 11111111 11110000
- (twenty-eight 1s and four 0s). Find

a. The first address: The first address can be found by ANDing the given addresses with the mask.

b. The last address: The last address can be found by ORing the given addresses with the complement of the mask.

c. The number of addresses: The number of addresses can be found by complementing the mask, interpreting it as a decimal

number, and adding 1 to it.

Address:	11001101	00010000	00100101	00100111
Mask:	**11111111**	**11111111**	**11111111**	**11110000**
First address:	11001101	00010000	00100101	00100000

Address:	11001101	00010000	00100101	00100111
Mask complement:	**00000000**	**00000000**	**00000000**	**00001111**
Last address:	11001101	00010000	00100101	00101111

Mask complement:	**00000000**	**00000000**	**00000000**	**00001111**
Number of addresses:	15 + 1 = 16			

IPv6

- Despite all short-term solutions, such as classless addressing, Dynamic Host Configuration Protocol (DHCP), address depletion is still a long-term problem for the Internet.
- This and other problems in the IP protocol itself such as lack of accommodation for real-time audio and video transmission, and encryption and authentication of data for some applications, have been the motivation for IPv6.
- **Structure:** An IPv6 address consists of 16 bytes (octets); it is 128 bits long

IPv6

- To make addresses more readable, IPv6 specifies hexadecimal colon notation. In this notation, 128 bits is divided into eight sections, each 2 bytes in length.
- Two bytes in hexadecimal notation requires four hexadecimal digits. Therefore, the address consists of 32 hexadecimal digits, with every four digits separated by a colon.

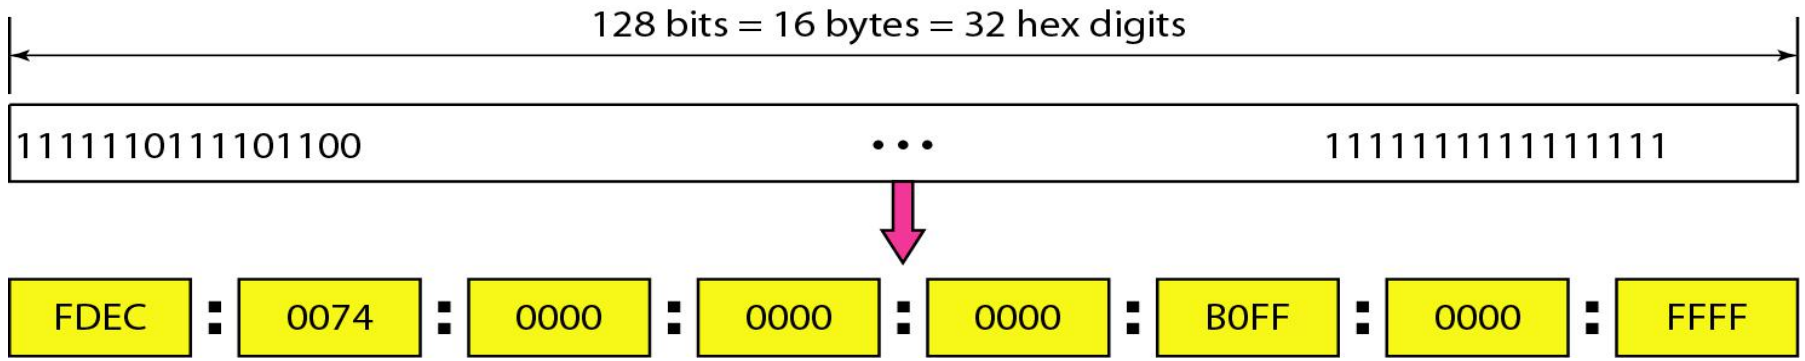
128 bits = 16 bytes = 32 hex digits
1111110111101100 • • • 1111111111111111
FDEC : 0074 : 0000 : 0000 : 0000 : B0FF : 0000 : FFFF

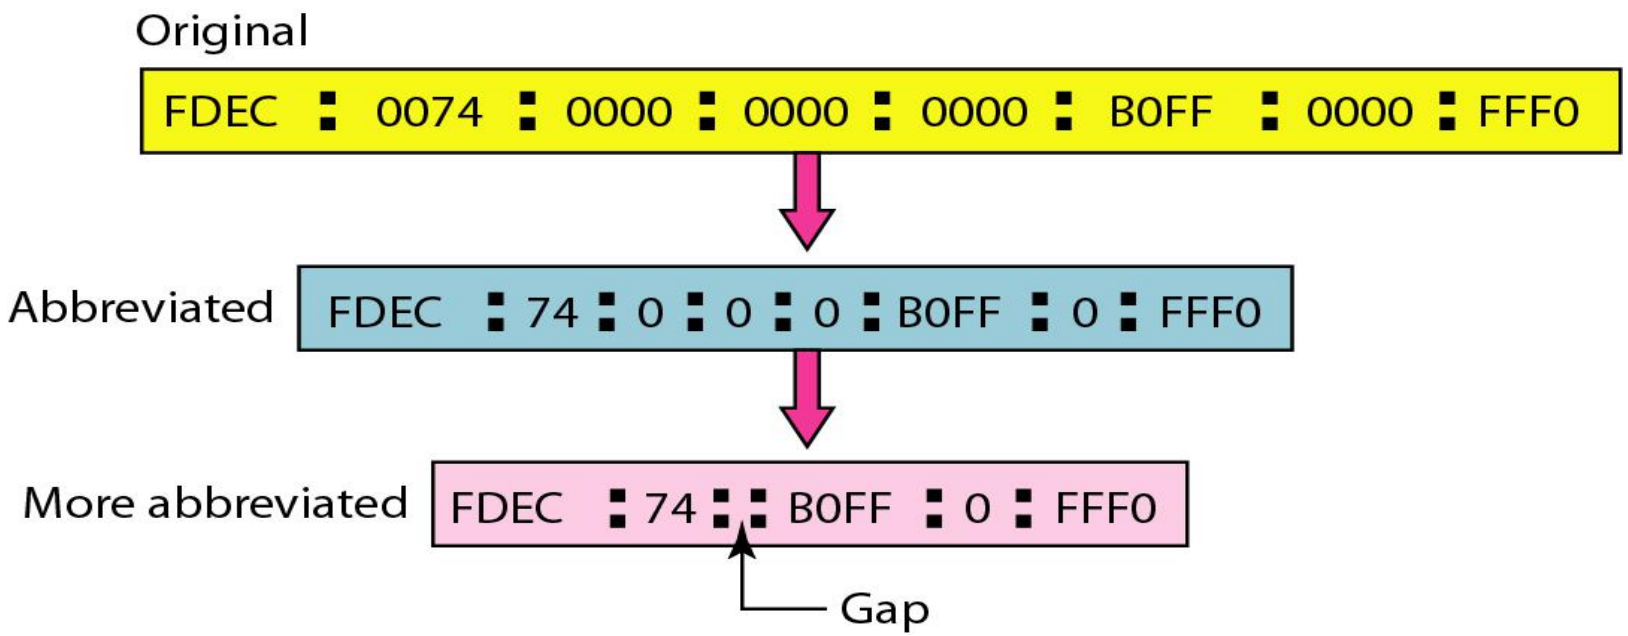
Original
FDEC : 0074 : 0000 : 0000 : 0000 : B0FF : 0000 : FFF0
Abbreviated
FDEC : 74 : 0 : 0 : 0 : B0FF : 0 : FFF0
More abbreviated
FDEC : 74 : : B0FF : 0 : FFF0
Gap

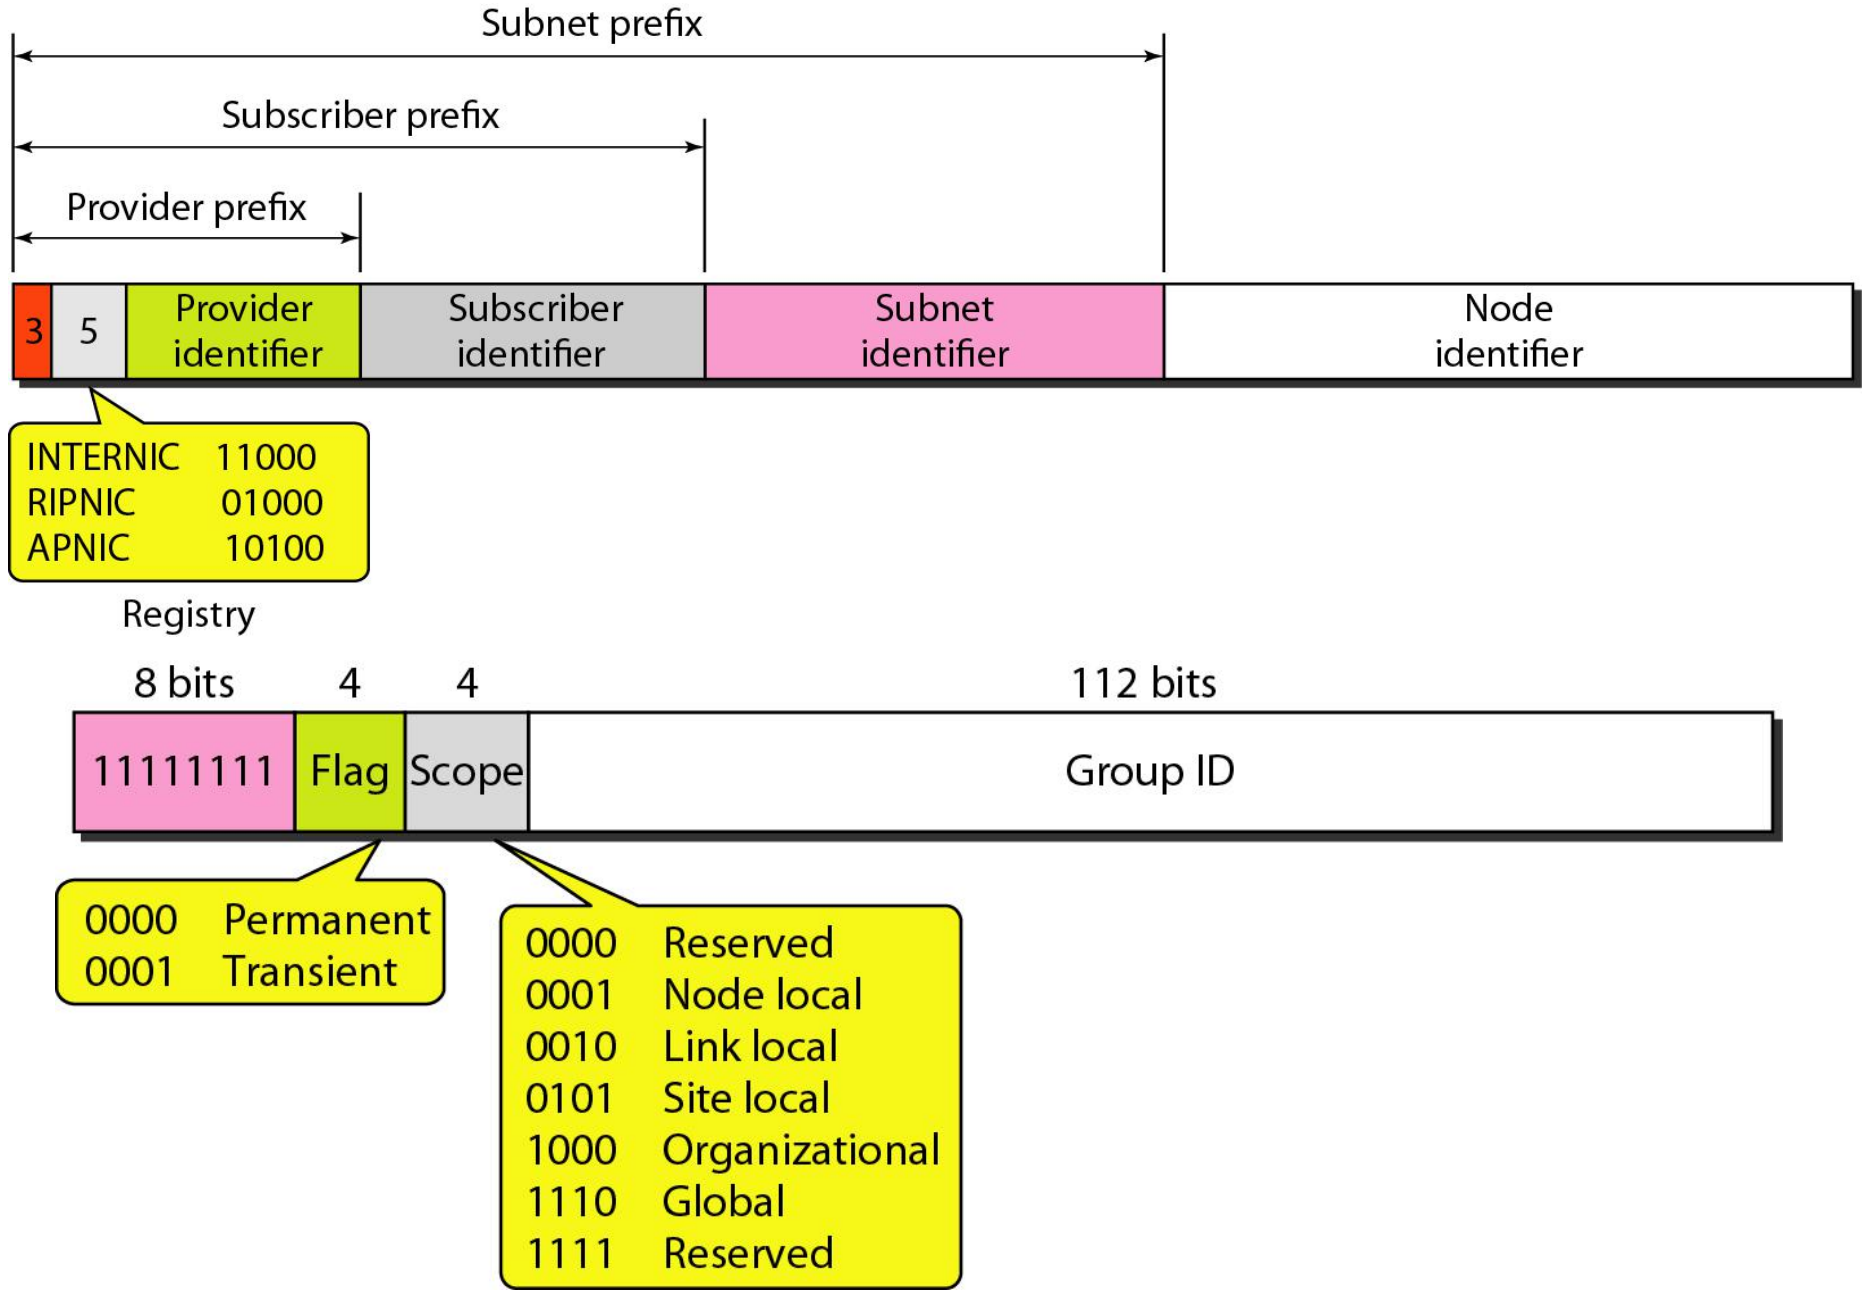
Subnet prefix
Subscriber prefix
Provider prefix
3
5
Provider identifier
Subscriber identifier
Subnet identifier
Node identifier
INTERNIC 11000
RIPNIC 01000
APNIC 10100
Registry
8 bits
4
4
112 bits
11111111
Flag
Scope
Group ID
0000 Permanent
0001 Transient
0000 Reserved
0001 Node local
0010 Link local
0101 Site local
1000 Organizational
1110 Global
1111 Reserved

Reserved Addresses in IPv6

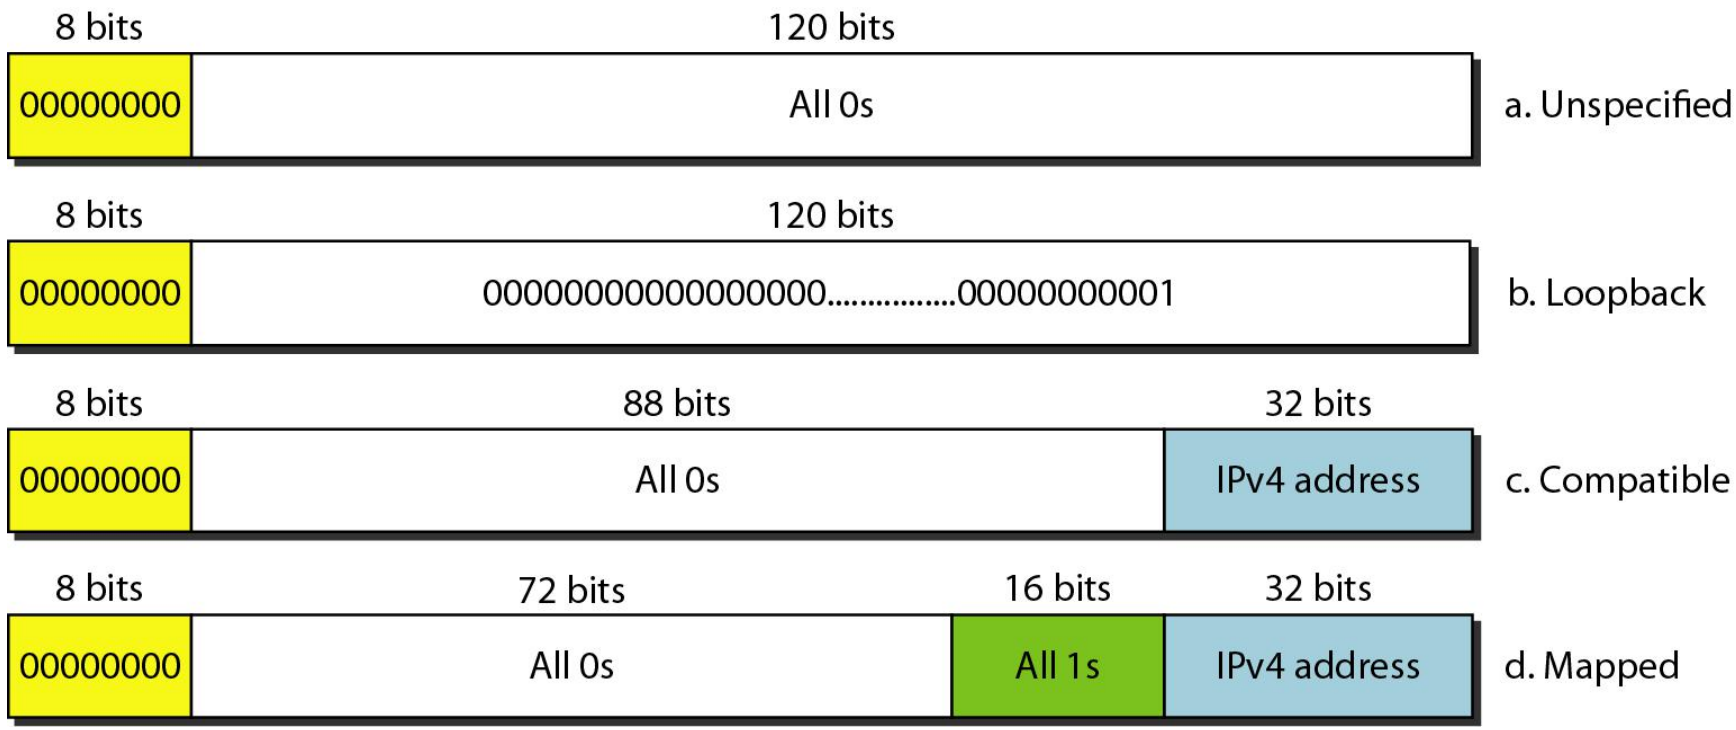

Special Addresses

- Both IPv4 and IPv6 have special interpretation for some addresses
  - IPv4 Network Address
    - An IPv4 address that has a hostID of 0s refer to as network address
  - IPv4 Direct Broadcast address
    - It is also called network broadcast address.
    - When it is used as a destination address, message sent to all host of the network.
    - hostID of all 1s is reserved for direct broadcast.

Special Addresses

- IPv4 Limited Broadcast Address
  - It is also known as local network broadcasting.
  - The packet is broadcasted locally with all 1s (for learning addresses of connected host)
  - It can be used as startup before computer knows its IP address.
- IPv4 Subnet Broadcast Address
  - It can be used when subnetting is used.
  - A subnet broadcast address consist of network prefix, a subnet number and all 1 in host field.

Special Addresses

- IPv4 all-0s Source Address
  - An address consist of thirty two 0s is reserved for a host which wants to communicate but does not know its own IP.
  - It is used as a temporary source address at startup.
- IPv4 Multicast Address
  - It is known as special form of multipoint delivery of a message.
  - Here the packet will be delivered to specific subnet host only.

Special Addresses

- IPv4 Loopback Address
  - It is used for testing of TCP/IP and inter-process communication of local computer.
  - The network prefix is 127.0.0.0/8
  - When application send packet using this address, the protocol software accepts outgoing packet and feeds it back to module that handles incoming packets.
  - The packet with address 127 does not travel through network and stays within computer.

Special Addresses

- IPv6 Link-Local Addresses
  - Link-Local address provide a way for computer to talk to its neighbor.
  - Any IPv6 address begins with a prefix of following 10 bits: 1111 1110 10
  - Computers connected to isolated network can use this address to find out neighbor router address as a startup.

Weaknesses in Internet Addressing

- The disadvantage of using embedded network is that address refer to network connection and not host computer.
- If a host moves from one network to another, its address must be changed.
- It is cumbersome to change address in some cases (IP conflict)

Weaknesses in Internet Addressing

- In IPv4 the size on network and host is fixed
- If the network grow beyond the size, a new prefix must be allocated.
- Though IPv6 solves the problem of network growth, it implies some other issues of forwarding.
- Some host and router carries more than one IP address, but humans think each must be give single name for identification.

Weaknesses in Internet Addressing

- When the network is down, it is impossible to reach to the host uses specific address.
- In case of congestion alternate path will not be used if the alternate address is not mentioned.

Introduction

- TCP/IP address scheme in which each host is assigned a 32-bit address.
- internet behaves like a virtual network, using only the assigned addresses when sending and receiving packets.
- Any two machines on a given physical network can communicate ***only*** if ***they know each other's physical network address.***
- But how a host or a router maps an **IP** address to the correct physical address when it needs to send a packet across a physical network?

Address Resolution Problem

- Consider two machines A and B that connect to the same physical network.
- Each has an assigned **IP** address ***ZA*** and ***ZB*** and a physical address ***PA*** and ***PB.***
- The goal is to devise low-level software that hides physical addresses and allows higher-level programs to work only with internet addresses.
- Communication must be carried out by physical networks using whatever physical address scheme the underlying network hardware supplies.

Address Resolution Problem

- Suppose machine A wants to send a packet to machine B across a physical network to which they both attach, but A has only B's internet address IB.
- The question arises: how does A map that address to B's physical address, PB?
- Address mapping must be performed at each step along a path from the original source to the ultimate destination.

Address Resolution Problem

- The problem of mapping high-level addresses to physical addresses is known as the ***address resolution problem*** and has been solved in several ways.
- Some protocol suites keep tables in each machine that contain pairs of high-level and physical addresses.
- Others solve the problem by encoding hardware addresses in high-level addresses.
- Using either approach exclusively makes high-level addressing at best.
- Here we will discuss two techniques for address resolution used by TCP/IP protocols and shows when each is appropriate.

Address Mapping

- An internet is made of a combination of physical networks connected by internetworking devices such as routers.
- A packet starting from a source host may pass through several different physical networks before finally reaching the destination host. The hosts and routers are recognized at the network level by their logical (IP) addresses.
- Packets pass through physical networks to reach these hosts and routers. At the physical level, the hosts and routers are recognized by their physical addresses.

Address Mapping

- A physical address is a local address. It must be unique locally, but is not necessarily unique universally. It is called a *physical address* because it is usually (but not always) implemented in hardware. It is also called as **MAC address**.
- The physical address and the logical address are two different identifiers and we need both to transfer data.
- This means that delivery of a packet to a host or a router requires two levels of addressing: logical and physical. We need to be able to map a logical address to its corresponding physical address and vice versa. This can be done by static or dynamic mapping:

Address Mapping

- Static mapping involves in the creation of a table that associates a logical address with a physical address. This table is stored in each machine on the network. Each machine that knows, the IP address of another machine but not its physical address can look it up in the table. This has some limitations:
  1. A machine could change its NIC, resulting in a new physical address.
  2. In some LANs, the physical address changes every time the computer is turned on.
  3. A mobile computer can move from one physical network to another, resulting in a change in its physical address.
- To implement these changes, a static mapping table must be updated periodically. This overhead could affect network performance.
- In dynamic mapping each time a machine knows one of the two addresses (logical or physical), it can use a protocol to find the other one

Resolution Through Direct Mapping

- In such network one or more computers (servers) store pairs of addresses, where each pair contains an Internet address and the corresponding physical address.
- Such servers store the pairs in a table in memory to speed searching.
- To guarantee efficient address resolution in such cases, software can use a conventional hash function to search the table.

Resolution Through Direct Mapping

- Designers of TCP/IP protocols found a creative solution to the address resolution problem for networks like the Ethernet that have broadcast capability.
- The solution allows new hosts or routers to be added to the network without recompiling code, and does not require maintenance of a centralized database.
- To avoid maintaining a table of mappings, the designers chose to use a low-level protocol to bind addresses dynamically.
- Termed the ***Address Resolution Protocol (ARP),*** the protocol provides a mechanism that is both reasonably efficient and easy to maintain.

ARP

- Anytime a host or a router has an IP datagram to send to another host or router, it has the logical (IP) address of the receiver.
- The IP datagram must be encapsulated in a frame to be able to pass through the physical network. This means that the sender needs the physical address of the receiver. The host or the router sends an ARP query packet. The packet includes the physical and IP addresses of the sender and the IP address of the receiver. Because the sender does not know the physical address of the receiver, the query is broadcast over the network.
- Every host or router on the network receives and processes the ARP query packet, but only the intended recipient recognizes its IP address and sends back an ARP response packet.
- The response packet contains the recipient's IP and physical addresses. The packet is unicast directly to the inquirer by using the physical address received in the query packet.

ARP Operation

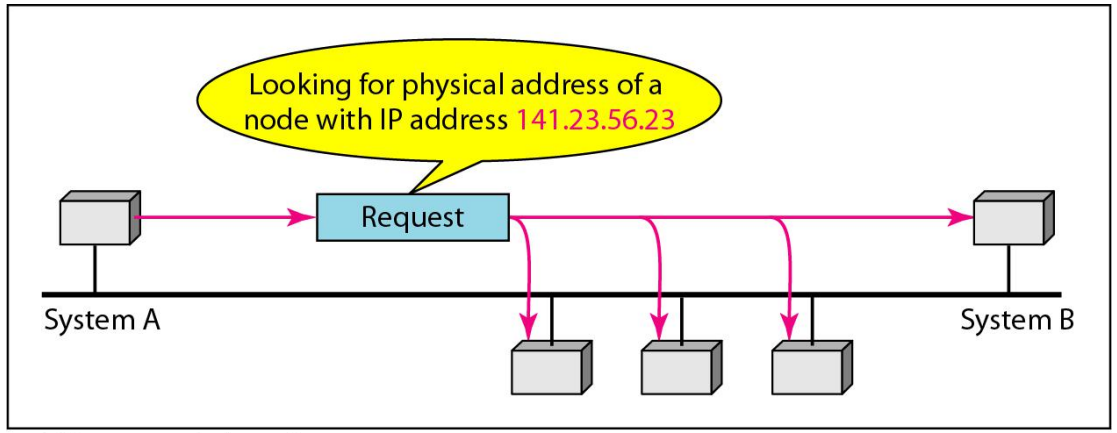

a. ARP request is broadcast

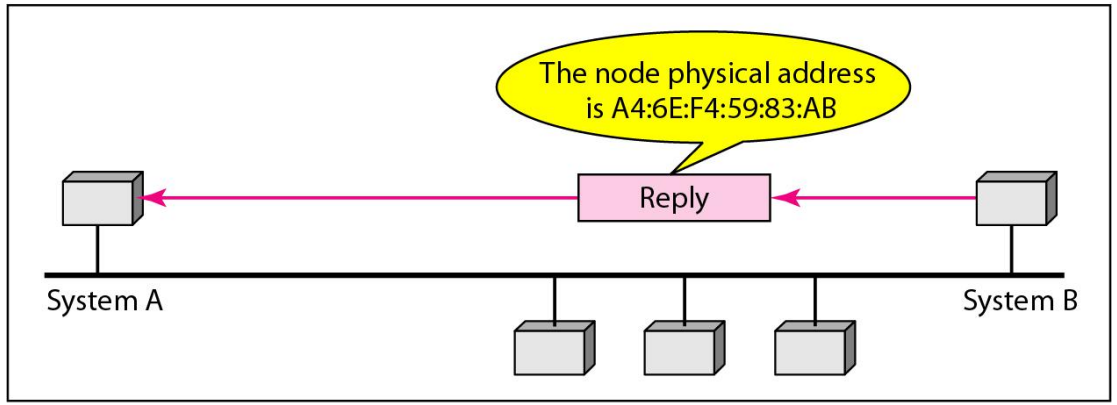

b. ARP reply is unicast

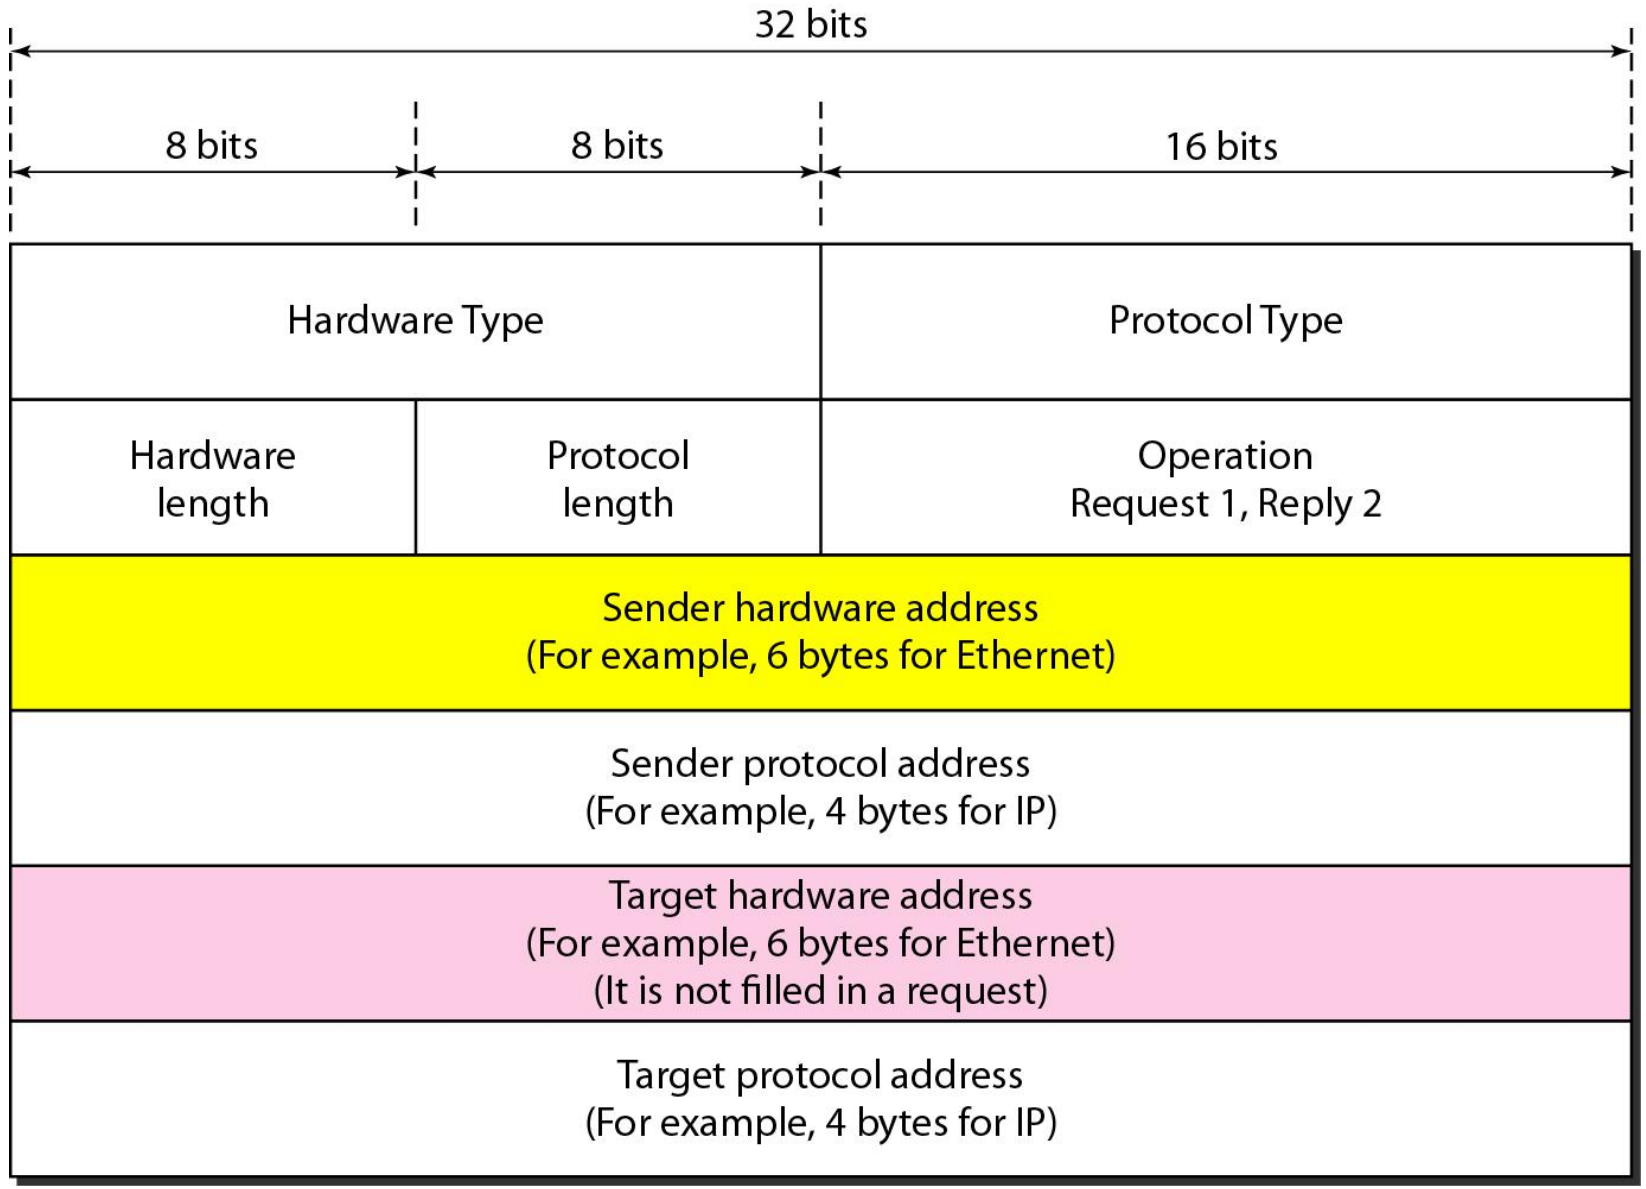
32 bits
8 bits
8 bits
16 bits
Hardware Type
Protocol Type
Hardware length
Protocol length
Operation
Request 1, Reply 2
Sender hardware address
(For example, 6 bytes for Ethernet)
Sender protocol address
(For example, 4 bytes for IP)
Target hardware address
(For example, 6 bytes for Ethernet)
(It is not filled in a request)
Target protocol address
(For example, 4 bytes for IP)

ARP Operation

- ARP can be useful if the ARP reply is ***cached*** because a system normally sends several packets to the same destination.
- A system that receives an ARP reply stores the mapping in the cache memory and keeps it for 20 to 30 minutes unless the space in the cache is exhausted.
- **ARP Packet Format:**
- Hardware type. 16-bit field defining the type of the network on which ARP is running. Each LAN has been assigned an integer based on its type
- Protocol type. 16-bit field defining the protocol
- Hardware length. 8-bit field defining the length of the physical address in bytes.
- Protocol length. 8-bit field defining the length of the logical address in bytes.

ARP Operation

- Operation. 16-bit field defining the type of packet. Two types are defined: ARP request (1) and ARP reply (2).
- Sender hardware address. variable-length field defining the physical address of the sender
- Sender protocol address. variable-length field defining the logical address of the sender.
- Target hardware address. variable-length field defining the physical address of the target.
- Target protocol address. variable-length field defining the logical (for example, IP) address of the target

Encapsulation of ARP Packet

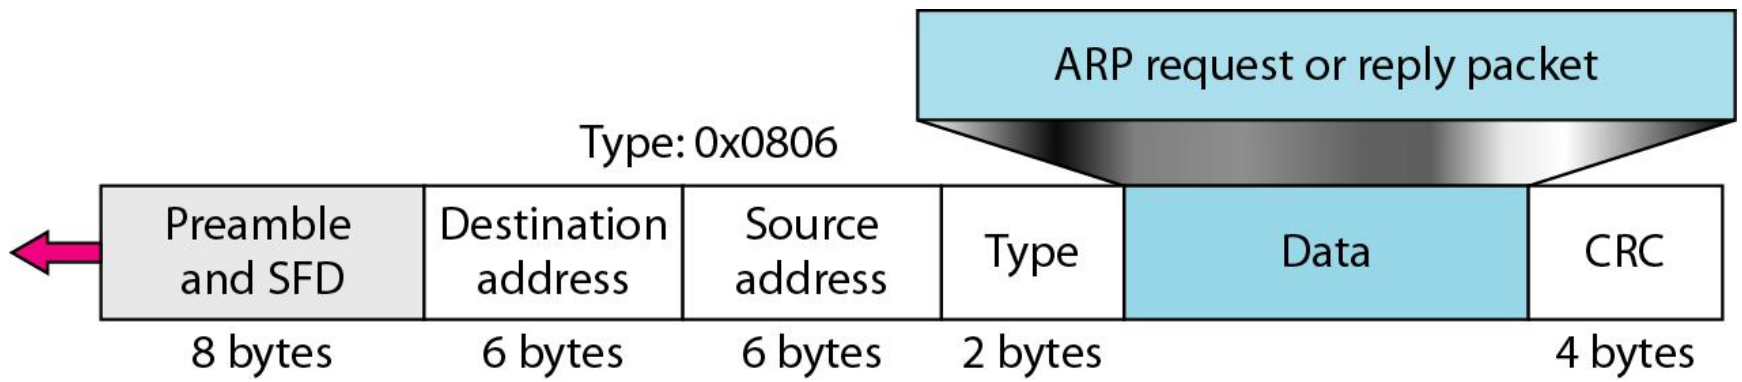

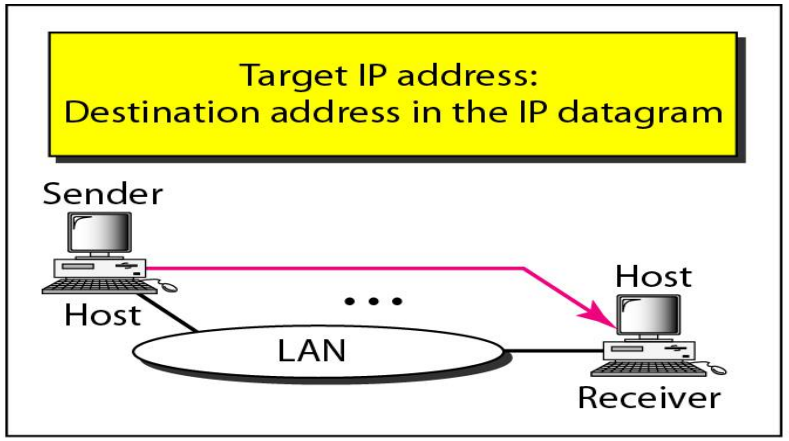

Case 1. A host has a packet to send to another host on the same network.

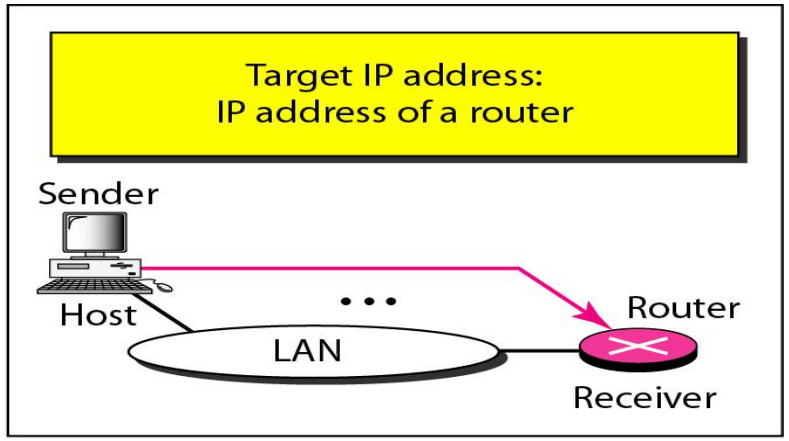

Case 2. A host wants to send a packet to another host on another network. It must first be delivered to a router.

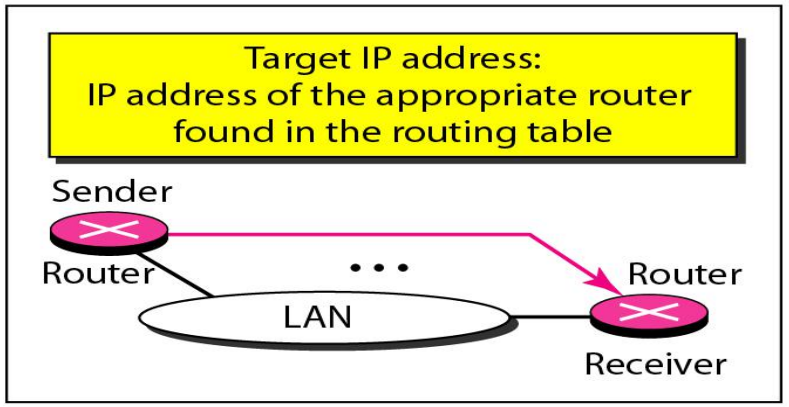

Case 3. A router receives a packet to be sent to a host on another network. It must first be delivered to the appropriate router.

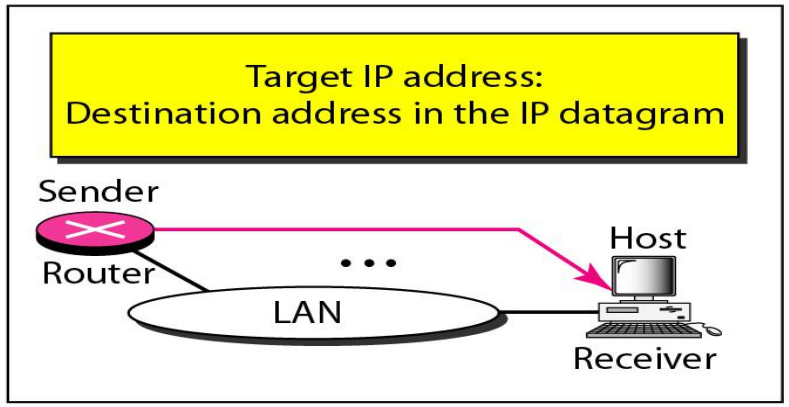

Case 4. A router receives a packet to be sent to a host on the same network.

Four Different Cases

1. The sender is a host and wants to send a packet to another host on the same network.
2. The sender is a host and wants to send a packet to another host on another network. In this case, the host looks at its routing table and finds the IP address of the next hop (router) for this destination. If it does not have a routing table, it looks for the IP address of the default router. The IP address of the router becomes the logical address that must be mapped to a physical address.
3. The sender is a router that has received a datagram destined for a host on another network. It checks its routing table and finds the IP address of the next router. The IP address of the next router becomes the logical address that must be mapped to a physical address.
4. The sender is a router that has received a datagram destined for a host on the same network. The destination IP address of the datagram becomes the logical address that must be mapped to a physical address

Address Resolution Cache

- Broadcasting is far too expensive to be used every time one machine needs to transmit a packet to another because every machine on the network must receive and process the broadcast packet.

ARP Cache Timeout

- To reduce communication costs, computers that use ARP maintain a cache of recently acquired IP-to-physical address bindings.
- That is, whenever a computer sends an **ARP** request and receives an *ARP* reply, it saves the **IP** address and corresponding hardware address information in its cache for successive lookups.

ARP Cache Timeout

- When transmitting a packet, a computer always looks in its cache for a binding before sending an AFW request.
- If it finds the desired binding in its *ARP* cache, the computer need not broadcast on the network.
- Thus, when two computers on a network communicate, they begin with an ARP request and response, and then repeatedly transfer packets without using **ARP** for each one.

Proxy ARP

- A proxy ARP is an ARP that acts on behalf of a set of hosts. Whenever a router running a proxy ARP receives an ARP request looking for the IP address of one of these hosts, the router sends an ARP reply announcing its own hardware (physical) address. After the router receives the actual IP packet, it sends the packet to the appropriate host or router

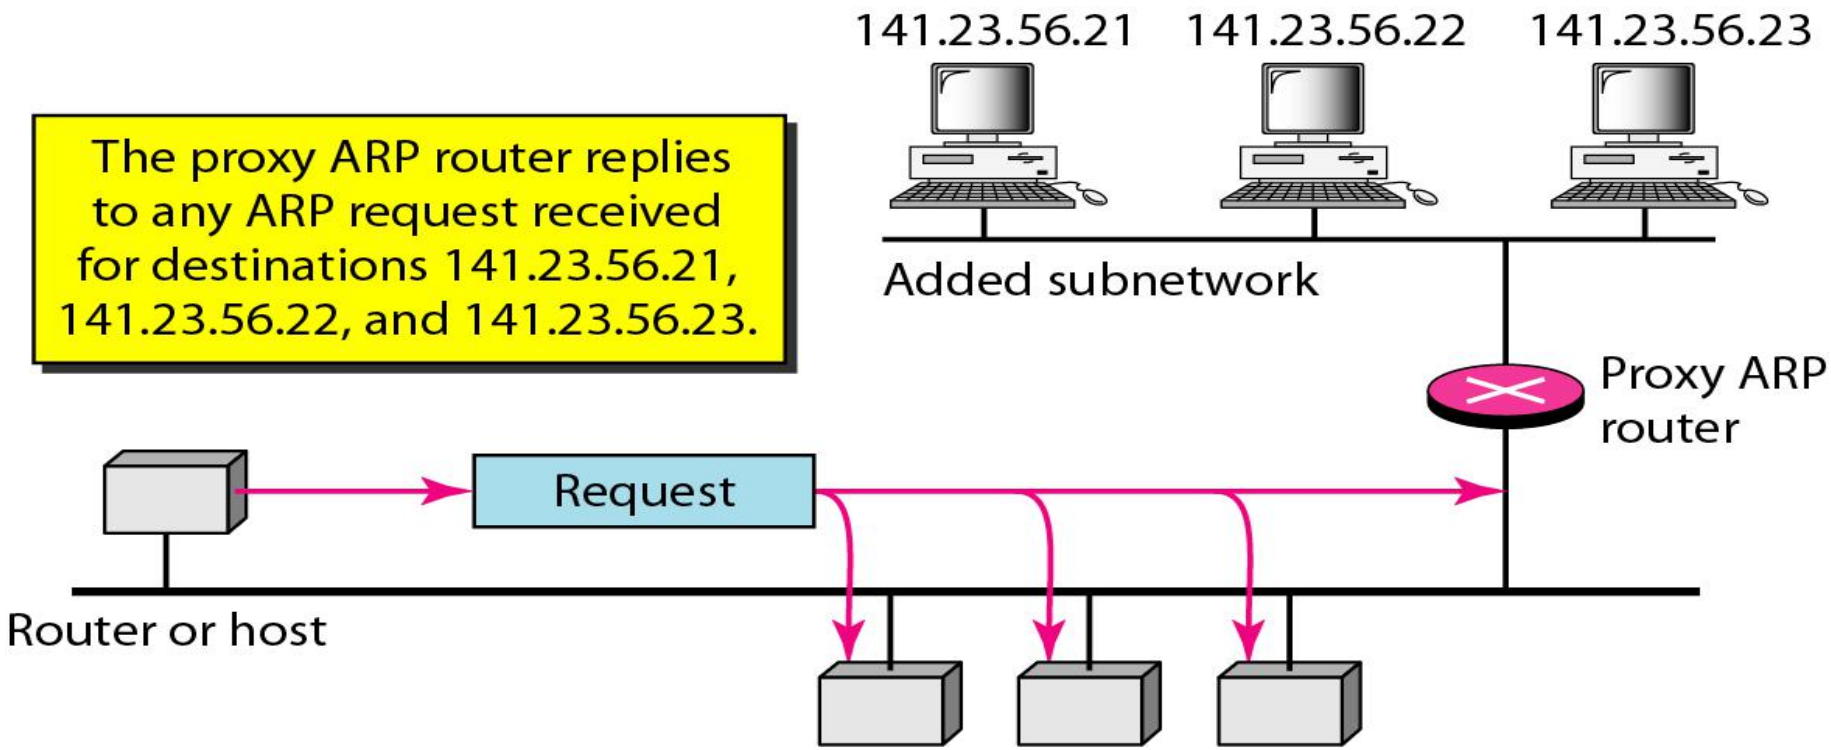
The proxy ARP router replies to any ARP request received for destinations 141.23.56.21, 141.23.56.22, and 141.23.56.23.
141.23.56.21
141.23.56.22
141.23.56.23
Added subnetwork
Proxy ARP router
Request
Router or host

ARP Refinements

- Several refinements of *ARP* have been included in the protocol.
- First, observe that if host A is about to use *ARP* because it needs to send to B, there is a high probability that host B will need to send to A in the near future.
- To anticipate B's need and avoid extra network traffic, A includes its IP-to-physical address binding when sending B a request.
- B extracts A's binding from the request, saves the binding in its **ARP** cache, and then sends a reply to A.

ARP Refinements

- Second, notice that because A broadcasts its initial request, **all** machines on the network receive it and can extract and update A's IP-to-physical address binding in their cache.
- Third, when a computer has its host interface replaced, (e.g., because the hardware has failed) its physical address changes.
- Other computers on the net that have stored a binding in their *ARP* cache need to be informed so they can change the entry.
- The computer can notify others of a new address by sending an *ARP* broadcast when it boots.

Relation of ARP to Other Protocols

- ARP is a low-level protocol that hides the underlying network physical addressing, permitting one to assign an arbitrary IP address to every machine.
- We think of ARP as part of the physical network system, and not as part of the internet protocols.

RARP

- Reverse Address Resolution Protocol (RARP) finds the logical address for a machine that knows only its physical address.
- The IP address of a machine is usually read from its configuration file stored on a disk file.
- However, a diskless machine is usually booted from ROM, which has minimum booting information.
- The ROM is installed by the manufacturer. It cannot include the IP address because the IP addresses on a network are assigned by the network administrator.

RARP

- The machine can get its physical address (by reading its NIC, for example), which is unique locally.
- It can then use the physical address to get the logical address by using the RARP protocol.
- A RARP request is created and broadcast on the local network. Another machine on the local network that knows all the IP addresses will respond with a RARP reply.
- There is a serious problem with RARP: Broadcasting is done at the data link layer.
- The physical broadcast address, a1s is in the case of Ethernet, does not pass the boundaries of a network.
- This means that if an administrator has several networks or several subnets, it needs to assign a RARP server for each network or subnet.

Timing RARP Transactions

- Some computers that rely on RARP to boot, choose to retry indefinitely until they receive a response.
- Other implementations announce failure after only a few tries to avoid flooding the network with unnecessary broadcast traffic (e.g., in case the server is unavailable).
- On an Ethernet, network failure is less likely than server overload.
- Making RARP software retransmit quickly may have the unwanted effect of flooding a congested server with more traffic.
- Using a large delay ensures that servers have ample time to satisfy the request and return an answer.

Primary And Backup RARP Servers

- The chief advantage of having several computers function as RARP servers is that it makes the system more reliable.
- If one server is down or too heavily loaded to respond, another answers the request.
- Thus, it is highly likely that the service will be available.
- The chief disadvantage of using many servers is that when a machine broadcasts a RARP request, the network becomes overloaded because all servers attempt to respond.
- On an Ethernet, for example, using multiple RARP servers makes the probability of collision high.

Internet Architecture

- An internet is an abstraction of physical networks because, at the lowest level, it provides the same functionality: accepting packets and delivering them.
- Higher levels of internet software add most of the rich functionality users perceive.
- **TCP/IP** internet provides three sets of services
  - Application Services
  - Reliable Transport Services
  - Connectionless Packet Delivery Services

Internet Architecture

- At the lowest level, a connectionless delivery service provides a foundation on which everything rests.
- At the next level, a reliable transport service provides a higher level platform on which applications depend.
- Internet software is designed around three conceptual networking services arranged in a hierarchy; much of its success has resulted because this architecture is surprisingly robust and adaptable.
- One of the most significant advantages of this conceptual separation is that it becomes possible to replace one service without disturbing others. Thus, research and development can proceed concurrently on all three.

Connectionless Delivery System

- The most fundamental internet service consists of a packet delivery system.
- Technically, the service is defined as an unreliable, best-effort, connectionless packet delivery system, analogous to the service provided by network hardware that operates on a best-effort delivery paradigm.
- The service is called ***unreliable*** because delivery is not guaranteed. The packet may be lost, duplicated, delayed, or delivered out of order, but the service will not detect such conditions, nor will it inform the sender or receiver.

Connectionless Delivery System

- The service is called ***connectionless*** because each packet is treated independently from all others.
- A sequence of packets sent from one computer to another may travel over different paths, or some may be lost while others are delivered.
- Finally, the service is said to use ***best-effort delivery*** because the internet software makes an earnest attempt to deliver packets.
- That is, the internet does not discard packets capriciously; unreliability arises only when resources are exhausted or underlying networks fail.

Fragmentation

- A datagram can travel through different networks. Each router decapsulates the IPv4 datagram from the frame it receives, processes it, and then encapsulates it in another frame.
- The format and size of the received frame depend on the protocol used by the physical network through which the frame has just traveled.
- The format and size of the sent frame depend on the protocol used by the physical network through which the frame is going to travel.

Fragmentation

- **Maximum Transfer Unit (MTU):** Each data link layer protocol has its own frame format in most protocols. One of the fields defined in the format is the maximum size of the data field. The value of the MTU depends on the physical network protocol.
- To make the IPv4 protocol independent of the physical network, the designers decided to make the maximum length of the IPv4 datagram equal to 65,535 bytes.
- This makes transmission more efficient if we use a protocol with an MTU of this size. However, for other physical networks, we must divide the datagram to make it possible to pass through these networks. This is called fragmentation.

Fragmentation Example

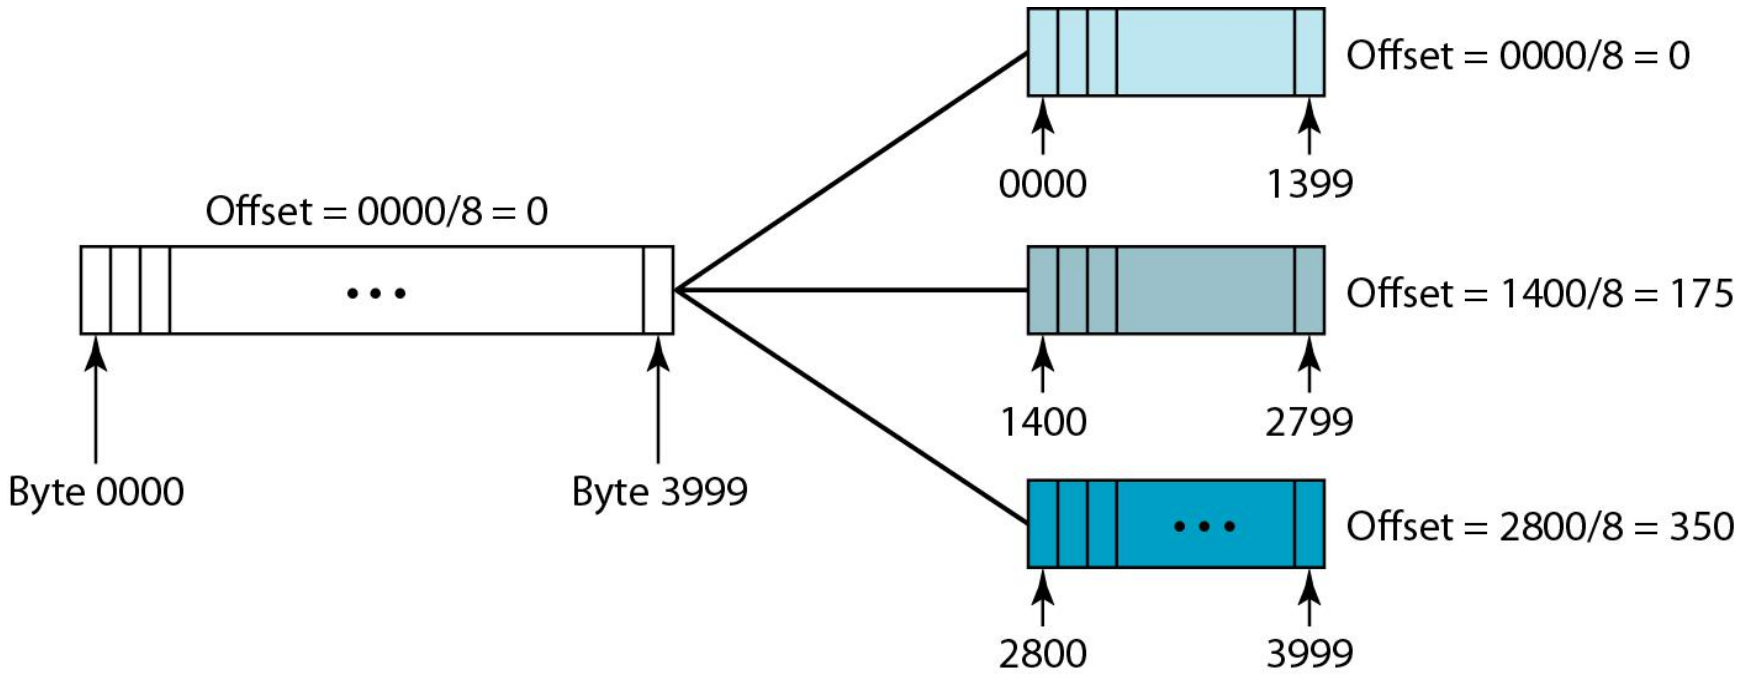

Fragmentation Example

- It is obvious that even if each fragment follows a different path and arrives out of order, the final destination host can reassemble the original datagram from the fragments received (if none of them is lost) by using the following strategy:

1. The first fragment has an offset field value of zero.
2. Divide the length of the first fragment by 8. The second fragment has an offset value equal to that result.
3. Divide the total length of the first and second fragments by 8. The third fragment has an offset value equal to that result.
4. Continue the process. The last fragment has a *more bit value of O.*

Internet Datagram

- The analogy between a physical network and a TCP/IP internet is strong.
- On a physical network, the unit of transfer is a frame that contains a header and data, where the header gives information such as the (physical) source and destination addresses.
- The internet call its basic transfer unit ***Internet datagram,*** sometimes referred to as ***an IP datagram***

Internet (IP) Datagram

- Like a typical physical network frame, a datagram is divided into header and data areas.
- Also like a frame, the datagram header contains the source and destination addresses and a type field that identifies the contents of the datagram.
- The difference, of course, is that the datagram header contains IP addresses whereas the frame header contains physical addresses.

Internet (IP) Datagram

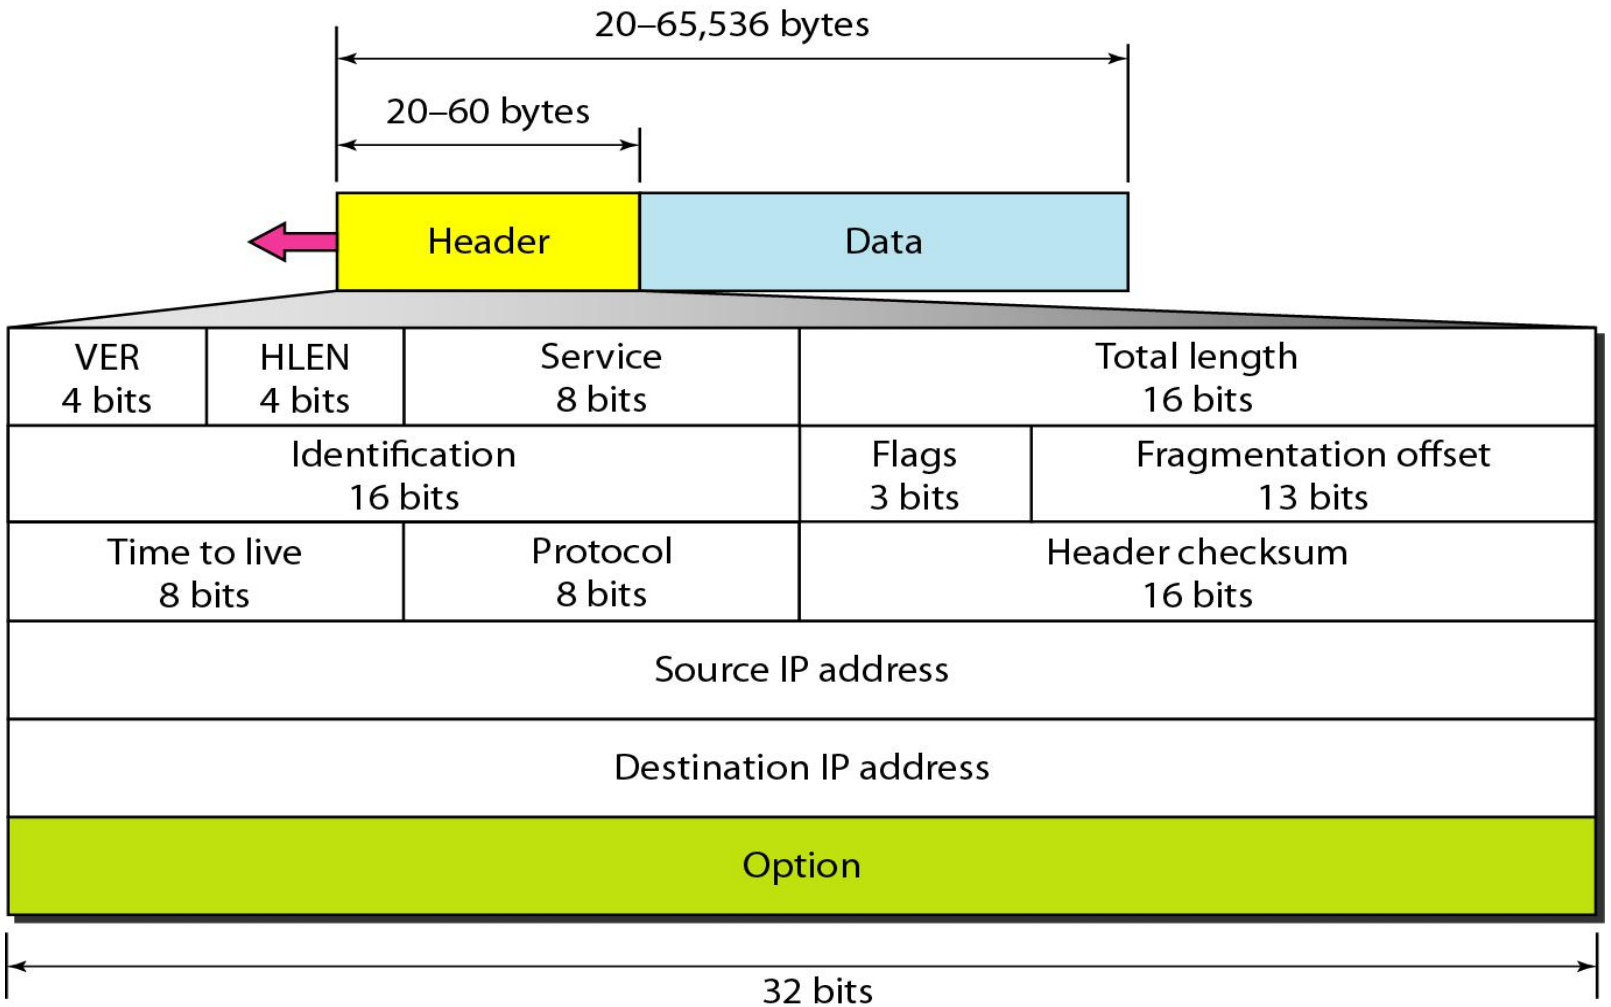

Internet (IP) Datagram

- An IP datagram consists of a header part and a text part. The header has a 20-byte fixed part and a variable length optional part.
- The Version field keeps track of which version of the protocol the datagram belongs to.
- Since the header length is not constant, a field in the header, HLEN, is provided to tell how long the header is.
- The Type of service field is the 6-bit field contained (from left to right), a three-bit Precedence field and three flags, D, T, and R. The Precedence field was a priority, from 0 (normal) to 7 (network control packet). The three flag bits allowed the host to specify what it cared most about from the set {Delay, Throughput, Reliability}.

Datagram Type Of Service

- The 8-bit SERVICE TYPE field specifies how the datagram should be handled.
- The field was originally divided into five subfields
- Three PRECEDENCE bits specify datagram precedence, with values ranging from ***0*** (normal precedence) through 7 (network control), allowing senders to indicate the importance of each datagram.
- Bits D, T, and R specify the **type** of transport desired for the datagram.
- When set, the D bit requests low delay, the T bit requests high throughput, and the R bit requests high reliability.

Internet (IP) Datagram

- The Total length includes everything in the datagram—both header and data. The maximum length is 65,535 bytes.
- The Identification field is needed to allow the destination host to determine which datagram a newly arrived fragment belongs to.
- All the fragments of a datagram contain the same Identification value.
- The Protocol field tells it which transport process to give it to.
- The Header checksum verifies the header only.
- The Source address and Destination address indicate the network number and host number.

Internet (IP) Datagram

- The Options field was designed to provide an escape to allow subsequent versions of the protocol to include information not present in the original design, to permit experimenters to try out new ideas, and to avoid allocating header bits to information that is rarely needed.
- Some options are followed by a 1-byte option length field, and then one or more data bytes. Originally, five options were defined

IP Option

- The ***IP OPTIONS*** field following the destination address is not required in every datagram; options are included primarily for network testing or debugging.
- Options processing is an integral part of the **IP** protocol, however, so all standard implementations must include it.
- The length of the ***IP OPTIONS*** field varies depending on which options are selected.
- Some options are one octet long; they consist of a single octet option code.
- Other options are variable length.
- When options are present in a datagram, they appear contiguously, with no special separators between them.