


Lecture3-TCP/IP

The Internet Layer

Andrei.doncescu@laas.fr

TCP/IP Internet provides 3 layers of service




Internet Protocol (IP RFC-791)

- Layering allows one to replace one service without affecting others
- IP layer (basic unit of transfer in TCP/IP) provides:
 - *Best-effort* (does not discard capriciously), *unreliable* (no guarantees)
 - Packet may be lost, duplicated, out-of-order with no notification
 - *Connectionless* (each packet treated independently)
 - IP software provides routing

Addressing and Delivering

- This physical addressing scheme works well on an individual LAN segment.
- A network that consists of only a few computers on an uninterrupted medium can function with nothing more than physical addresses.
- Data can pass directly from network adapter to network adapter using the low-level protocols associated with the Network Access layer.



Addressing and Delivering

Unfortunately, on a routed network, it is not possible to deliver data by physical address.

The discovery procedures required for delivering by physical address do not work across a router interface.

Even if they did work, delivery by physical address would be cumbersome because the permanent physical address built in to a network card does not allow you to impose a logical structure on the address space.


Addressing and Delivering



TCP/IP therefore makes the physical address invisible and instead organizes the network around a logical, hierarchical addressing scheme.



This logical addressing scheme is maintained by the **IP protocol** at the Internet layer.



The logical address is called the **IP address**.

Another Internet layer protocol called **Address Resolution Protocol (ARP)** assembles a table that maps IP addresses to physical addresses.

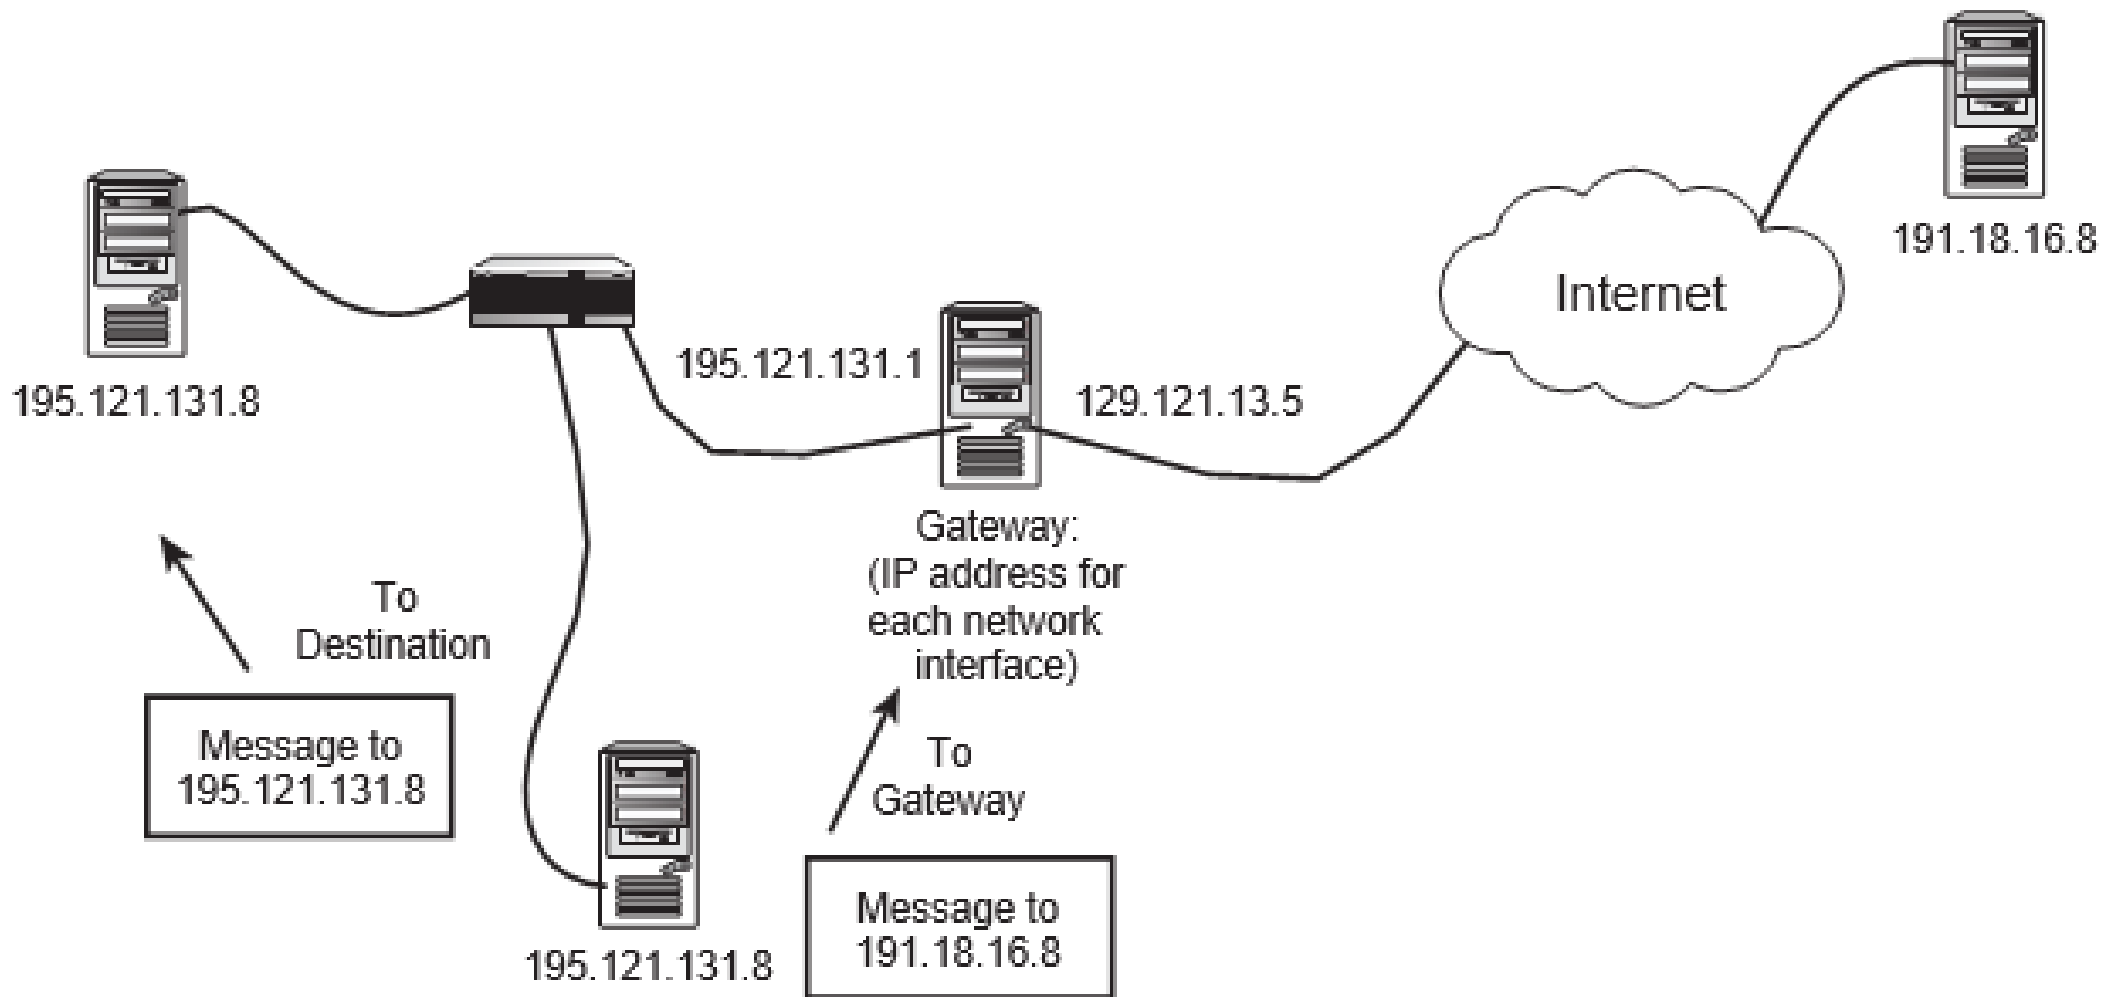
This ARP table is the link between the IP address and the physical address burned into the network adapter card.

Addressing and Delivering

- TCP/IP software uses the following strategy for sending data on the network:
 - If the destination address is on the same network segment as the source computer, the source computer sends the packet directly to the destination.
 - The IP address is resolved to a physical address using **ARP**, and the data is directed to the destination network adapter.

Addressing and Delivering

- If the destination address is on a different segment from the source computer, the following process begins:
 - The datagram is directed to a gateway. A gateway is a device on the local network segment that is capable of forwarding a datagram to other network segments.
The gateway address is resolved to a physical address using ARP, and the data is sent to the gateway's network adapter.
 - The datagram is routed through the gateway to a higher-level network segment where the process is repeated. If the destination address is on the new segment, the data is delivered to its destination. If not, the datagram is sent to another gateway.
 - The datagram passes through the chain of gateways to the destination segment, where the destination IP address is mapped to a physical address using ARP and the data is directed to the destination network adapter.



Addressing and Delivering

- To deliver data on a complex routed network, the Internet layer protocols must therefore be able to:
 - Identify any computer on the network
 - Provide a means for determining when a message must be sent through the gateway
 - Provide a hardware-independent means of identifying the destination network segment so that the datagram will pass efficiently through the routers to the correct segment
 - Provide a means for converting the logical IP address of the destination computer to a physical address so that the data can be delivered to the network adapter of the destination computer

Internet Protocol

The Internet Protocol (IP) provides a hierarchical, hardware-independent addressing system and offers the services necessary for delivering data on a complex, routed network.

Each network adapter on a TCP/IP network has a unique IP address.

Internet Protocol



The Host

Descriptions of TCP/IP often talk about a *computer* having an IP address.

A computer is sometimes said to have an IP address because most computers have only one network adapter.

However, computers with multiple network adapters are also common.


A computer that is acting as a router or a proxy server, for instance, must have more than one network adapter and, therefore, has more than one IP address.

The term *host* is often used for a network device associated with an IP address.

Under some operating systems, it is also possible to assign more than one IP address to a single network adapter.

Internet Protocol

- The IP address is therefore divided into two parts:
 - The network ID
 - The host ID



Internet Protocol

- The network must provide a means for determining which part of the IP address is the **network ID** and which part is the **host ID**.
- Unfortunately, the variety and complexity of networks in the real world precludes a simple, one-size-fits-all solution to this problem.
- Big networks must reserve a large number of host bits for their large number of hosts.
- Small networks do not need many bits to give each host a unique ID; however, the vast number of small networks means that more bits of the IP address are necessary for the network ID.



- Original solution to this problem was to divide the IP address space into a series of address classes.
- Class A networks used the first 8 bits of the address for the network ID;
- Class B used the first 16 bits;
- Class C networks used the first 24 bits. This system was extended through a feature called subnetting to provide greater control at the local level for structuring the network.

Internet Protocol

Internet Protocol

A more recent technique known as classless interdomain routing (CIDR) essentially renders the address class system unnecessary.

CIDR, which is now quite common on the Internet, offers a simple, flexible, and unambiguous notation for allocating blocks of IP addresses.

IP Header Fields

- Every IP datagram begins with an IP header.
- The TCP/IP software on the source computer constructs the IP header.
- The TCP/IP software at the destination uses the information enclosed in the IP header to process the datagram.
- The IP header contains a great deal of information, including the IP addresses of the source and destination computers, the length of the datagram, the IP version number, and special instructions to routers.

- Basic transfer unit



- Format of Internet datagram

0	4	8	16	19	24	31
Vers	Hlen	Type of serv.	Total length			
Identification			Flags	Fragment offset		
TTL		Protocol	Header Checksum			
Source IP address						
Destination IP address						
IP Options (if any)					Padding	
Data						
...						

IP datagram format (cont.)

- **Vers** (4 bits): version of IP protocol (IPv4=4)
- **Hlen** (4 bits): Header length in 32 bit words, without options (usual case) = 20
- **Type of Service – TOS** (8 bits): little used in past, now being used for QoS
- **Total length** (16 bits): length of datagram in bytes, includes header and data
- **Time to live – TTL** (8bits): specifies how long datagram is allowed to remain in internet
 - Routers decrement by 1
 - When TTL = 0 router discards datagram
 - Prevents infinite loops
- **Protocol** (8 bits): specifies the format of the data area
 - Protocol numbers administered by central authority to guarantee agreement, e.g. TCP=6, UDP=17 ...

IP Header Fields

Bit Position: 0 4 8 16 24 31

Version	IHL	Type of Service	Total Length	
Identification		Flags	Fragment Offset	
Time to Live	Protocol	Header Checksum		
Source IP Address				
Destination IP Address				
IP Options (optional)			Padding	
Data				
More Data...?				

- **Version:** This 4-bit field indicates which version of IP is being used. The current version of IP is 4. The binary pattern for 4 is 0100.
- **IHL (Internet Header Length):** This 4-bit field gives the length of the IP header in 32-bit words. The minimum header length is five 32-bit words. The binary pattern for 5 is 0101.
- **Type of Service:** The source IP can designate special routing information. Some routers ignore the Type of Service field, although this field recently has received more attention with the emergence of quality of service (QoS) technologies. The primary purpose of this 8-bit field is to provide a means of prioritizing datagrams that are waiting to pass through a router. Most implementations of IP today simply put all 0s in this field.
- **Total Length:** This 16-bit field identifies the length, in octets, of the IP data-gram. This length includes the IP header and the data payload.

IP Header Fields

Bit Position: 0	4	8	16	24	31
Version	IHL	Type of Service	Total Length		
Identification			Flags	Fragment Offset	
Time to Live	Protocol		Header Checksum		
Source IP Address					
Destination IP Address					
IP Options (optional)				Padding	
Data					
More Data...?					

Identification: This 16-bit field is an incrementing sequence number assigned to messages sent by the source IP.

When a message is sent to the IP layer and it is too large to fit in one datagram, IP fragments the message into multiple datagrams, giving all datagrams the same identification number.

This number is used on the receiving end to reassemble the original message.

Flags: The Flags field indicates fragmentation possibilities. The first bit is unused and should always have a value of 0. The next bit is called the DF (Don't Fragment) flag. The DF flag signifies whether fragmentation is allowed (value = 0) or not (value = 1).

The next bit is the MF (More Fragments) flag, which tells the receiver that more fragments are on the way. When MF is set to 0, no more fragments need to be sent or the datagram never was fragmented.

IP Header Fields

Bit Position: 0	4	8	16	24	31
Version	IHL	Type of Service	Total Length		
Identification			Flags	Fragment Offset	
Time to Live	Protocol		Header Checksum		
Source IP Address					
Destination IP Address					
IP Options (optional)				Padding	
Data					
More Data...?					

Fragment Offset: This 13-bit field is a numeric value assigned to each successive fragment. IP at the destination uses the fragment offset to reassemble the fragments into the proper order. The offset value found here expresses the offset as a number of 8-byte units.

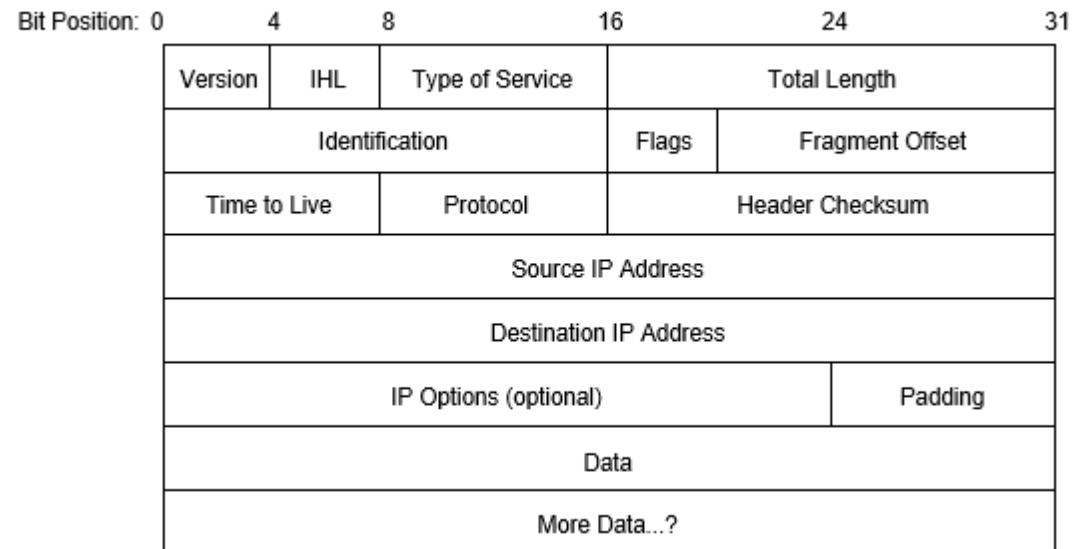
Time To Live (TTL): This bit field indicates the amount of time in seconds or router hops that the datagram can survive before being discarded. Every router examines and decrements this field by at least 1, or by the number of seconds the datagram is delayed inside the router. The datagram is discarded when this field reaches 0.

A **hop** represents the number of routers a datagram must cross on the way to its destination. If a datagram passes through five routers before arriving at its destination, the destination is said to be five hops, or five router hops, away.

IP Header Fields

- **Protocol:** The 8-bit Protocol field indicates the protocol that will receive the data payload. A datagram with the protocol identifier 6 (binary 00000110) is passed up the stack to the TCP module, for example. The following are some common protocol values:

Protocol Name	Protocol Identifier
ICMP	1
TCP	6
UDP	17



IP Datagram format (cont.)

Source & destination IP address (32 bits each):
contain IP address of sender and intended recipient

Options (variable length):
Mainly used to record a route, or timestamps, or specify routing

IP Header Fields

Bit Position: 0	4	8	16	24	31
Version	IHL	Type of Service	Total Length		
Identification			Flags	Fragment Offset	
Time to Live	Protocol		Header Checksum		
Source IP Address					
Destination IP Address					
IP Options (optional)				Padding	
Data					
More Data...?					

Header Checksum: This field holds a 16-bit calculated value to verify the validity of the header only. This field is recomputed in every router as the TTL field decrements.

Source IP Address: This 32-bit field holds the address of the source of the datagram.

Destination IP Address: This 32-bit field holds the destination address of the datagram and is used by the destination IP to verify correct delivery.

IP Options: This field supports a number of optional header settings primarily used for testing, debugging, and security. Options include Strict Source Route (a specific path router path that the datagram should follow), Internet Timestamp (a record of timestamps at each router), and security restrictions.

IP Header Fields

Bit Position: 0	4	8	16	24	31
Version	IHL	Type of Service	Total Length		
Identification			Flags	Fragment Offset	
Time to Live	Protocol		Header Checksum		
Source IP Address					
Destination IP Address					
IP Options (optional)				Padding	
Data					
More Data...?					

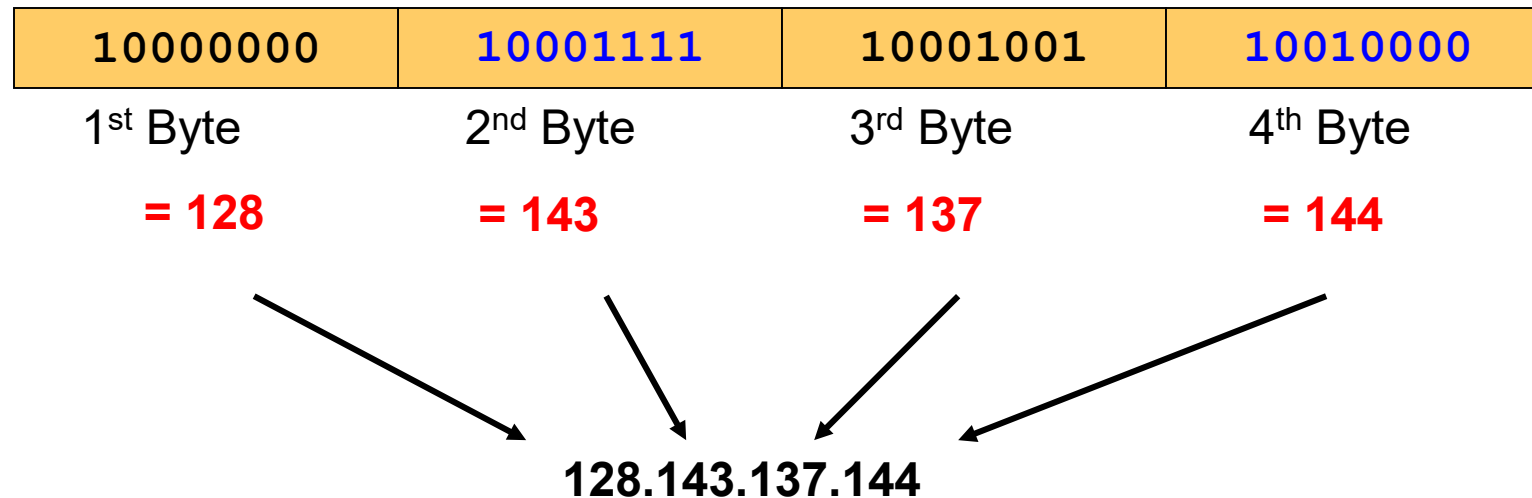
Padding: The IP Options field may vary in length. The Padding field provides additional 0 bits so that the total header length is an exact multiple of 32 bits. (The header must end after a 32-bit word because the IHL field measures the header length in 32-bit words.)

IP Data Payload: This field typically contains data destined for delivery to TCP or UDP (in the Transport layer), ICMP, or IGMP. The amount of data is variable but could include thousands of bytes.

Dotted Decimal Notation

- IP addresses are written in a so-called *dotted decimal notation*
- Each byte is identified by a decimal number in the range [0..255]:

-



IP Addressing

An IP address is a 32-bit binary address. This 32-bit address is subdivided into four 8-bit segments called **octets**.

Humans do not work well with 32-bit binary addresses or even 8-bit binary octets, so the IP address is almost always expressed in what is called **dotted-decimal** format.

In dotted-decimal format, each octet is given as an equivalent decimal number. The four decimal values (4 x 8 = 32 bits) are then separated with periods.

Eight binary bits can represent any whole number from 0 to 255, so the segments of a dotted-decimal address are decimal numbers from 0 to 255.

You have probably seen examples of dotted-decimal IP addresses on your computer. A dotted-decimal IP address looks like this: 209.121.131.14.

Network prefix and Host number

- The network prefix identifies a network and the host number identifies a specific host (actually, interface on the network).



- **How do we know how long the network prefix is?**
 1. The network prefix is implicitly defined (see **class-based addressing**)
 2. The network prefix is indicated by a **netmask**.

IP Addressing

- Part of the IP address is used for the network ID, and part of the address is used for the host ID through a system of **address classes**.
- Although the more recent CIDR classless addressing has reduced the importance of the class system, address classes are still important enough to describe here as a starting point for understanding addressing in TCP/IP.
- The address class system divides the IP address space into address classes. Most IP addresses fall into the following classes:
 - **Class A addresses:** The first 8 bits of the IP address are used for the network ID. The final 24 bits are used for the host ID.
 - **Class B addresses:** The first 16 bits of the IP address are used for the network ID. The final 16 bits are used for the host ID.
 - **Class C addresses:** The first 24 bits of the IP address are used for the network ID. The final 8 bits are used for the host ID.

IP Addressing

More bits lead to more bit combinations. As you might guess, the Class A format provides a small number of possible network IDs and a huge number of possible host IDs for each network.

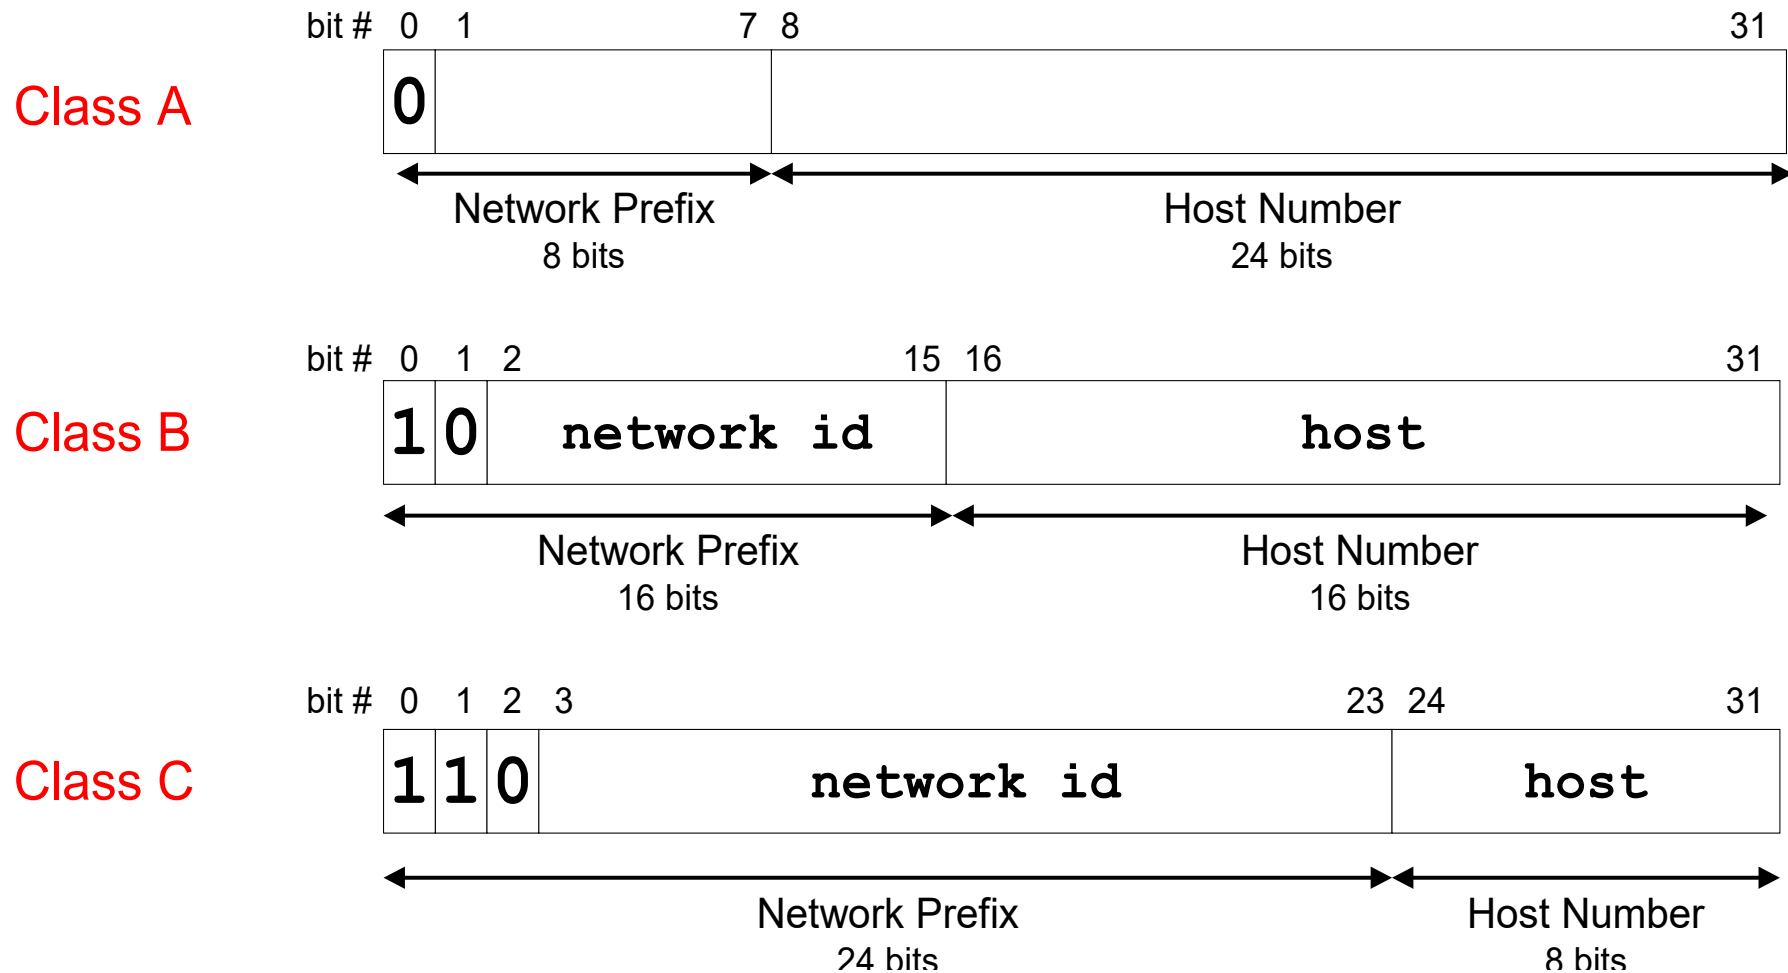
A Class A network can support approximately 224, or 16,777,216 hosts.

A Class C network, on the other hand, can provide host IDs for only a small number of hosts (254, which is 28, or 256, minus the unusable all 0s and all 1s addresses), but many more combinations of network IDs are available in the Class C format.



IP Addressing

- The designers of TCP/IP wrote the address rules such that the class of an address is obvious from the address itself. The first few bits of the binary address specify whether the address should be interpreted as a Class A, Class B, or Class C address.
- The rules for interpreting addresses are as follows:
 - If the 32-bit binary address starts with a 0 bit, the address is a Class A address.
 - If the 32-bit binary address starts with the bits 10, the address is a Class B address.
 - If the 32-bit binary address starts with the bits 110, the address is a Class C address.

The old way: Internet Address Classes



The old way: Internet Address Classes



IP Addressing

Address Class	Binary Address Must Begin With	First Term of Dotted-Decimal Address Must Be	Excluded Addresses
A	0	0 to 127	10.0.0.0 to 10.255.255.255 127.0.0.0 to 127.255.255.255
B	10	128 to 191	172.16.0.0 to 172.31.255.255
C	110	192 to 223	192.168.0.0 to 192.168.255.255

IP Addressing

Classes D and E

The Internet specifications also define special-purpose Class D and Class E addresses.

Class D addresses are used for multicasting. A multicast is a single message sent to a subset of the network, as opposed to a broadcast, which is processed by all nodes on the local net. The four leftmost bits of a Class D network address always start with the binary pattern 1110, which corresponds to decimal numbers 224 through 239.

Class E networks are considered experimental and are not normally used in production environments. The five leftmost bits of a Class E network always start with the binary pattern 11110, which corresponds to decimal numbers 240 through 247.

Special IP Addresses

A few IP addresses have special meanings and are not assigned to specific hosts. An all-0 host ID refers to the network itself.

For instance, the IP address 129.152.0.0 refers to the Class B network with the network ID 129.152.

An all-1s host ID signifies a broadcast. A broadcast is a message sent to all hosts on the network.

The IP address 129.152.255.255 is the broadcast address for the Class B network with the network ID 129.152. (Note that the dotted-decimal term 255 corresponds to the all-ones binary octet 11111111.)

The address 255.255.255.255 can also be used for broadcast on the network.

Special IP Addresses

Addresses beginning with the decimal number 127 are loopback addresses.

A message addressed to a loopback address is sent by the local TCP/IP software to itself. The loopback address is used to verify that the TCP/IP software is functioning.

The loopback address 127.0.0.1 is commonly used.

Because the private address ranges don't have to be synchronized with the rest of the world, the complete address range is available for any network.

A network administrator using these private addresses has more room for subnetting, and many more assignable addresses.

The address range 169.254.0.0 to 169.255.255.255 is reserved for autoconfiguration.

Problems with Classful IP Addresses



The original classful address scheme had a number of problems



Problem 1. Too few network addresses for large networks

Class A and Class B addresses are gone



Problem 2. Two-layer hierarchy is not appropriate for large networks with Class A and Class B addresses.

Fix #1: Subnetting

Problems with Classful IP Addresses

**Problem 3. The Internet is going to
outgrow the 32-bit addresses**

- **Fix #3: IP Version 6**

Problems with Classful IP Addresses

**Problem 3. The Internet is going to
outgrow the 32-bit addresses**

- **Fix #3: IP Version 6**

CIDR Example

- CIDR notation of a network address:
192.0.2.0/18
 - "18" says that the first 18 bits are the network part of the address (and 14 bits are available for specific host addresses)
- The network part is called the prefix
- Assume that a site requires a network address with 1000 addresses
- With CIDR, the network is assigned a continuous block of 1024 addresses with a 22-bit long prefix

CIDR: Prefix Size vs. Network Size

CIDR Block Prefix	# of Host Addresses
/27	32 hosts
/26	64 hosts
/25	128 hosts
/24	256 hosts
/23	512 hosts
/22	1,024 hosts
/21	2,048 hosts
/20	4,096 hosts
/19	8,192 hosts
/18	16,384 hosts
/17	32,768 hosts
/16	65,536 hosts
/15	131,072 hosts
/14	262,144 hosts
/13	524,288 hosts

Examples

- Example 1:
 - A classless address is given as 167.199.170.82/27
 - What is the number of addresses, the first address and the last address in the network?
 - Solution: $n=27$
Number of addresses = $2^{32-n} = 2^5 = 32$ addresses

Example 1

- First address
 - Address 167.199.170.82/27
10100111 11000111 10101010
01010010
 - First Address 167.199.170.64
10100111 11000111 10101010
01000000
 - Last Address 167.199.170.127
10100111 11000111 10101010
01011111

Example 2

- In classless addressing, an address cannot per se define the block the address belongs to.
- For example, the address 230.8.24.56 can belong to many blocks.
 - Prefix length 16: block: 230.8.0.0 - 230.8.255.255
 - Prefix length 20: block: 230.8.16.0 – 230.8.31.255
 - Prefix length 26: block: 230.8.24.0 – 230.8.24.63
 -

Block Allocation


- Block allocation is given to a global authority called Internet Corporation of Assigned Names and Numbers (ICANN).
- ICANN assigns a large block of addresses to an ISP.
- In CIDR, two restrictions need to be applied to the allocated block.
 - The number of requested addresses, N , needs to be a power of 2. The reason that $N = 2^{32-n}$, where N the number of addresses and n the network prefix. Since $n = 32 - \log_2 N$, which should be integer.
 - The block addresses should be contiguous

Example 3

- An ISP has requested a block of 1000 addresses. What is the prefix length?
 - Since 1000 is not a power of 2, 1024 addresses are granted.
 - The prefix length $n = 32 - \log_2 1024 = 22$. i.e. An available block 18.14.12.0/22 could be granted to the ISP.

IP Fragmentation

- How do we send a datagram of say 1400 bytes through a link that has a *Maximum Transfer Unit (MTU)* of say 620 bytes?
- Answer the datagram is broken into fragments



- Router fragments 1400 byte datagrams
 - Into 600 bytes, 600 bytes, 200bytes (note 20 bytes for IP header)
 - Routers do NOT reassemble, up to end host

Fragmentation Control

Identification: copied into fragment, allows destination to know which fragments belong to which datagram

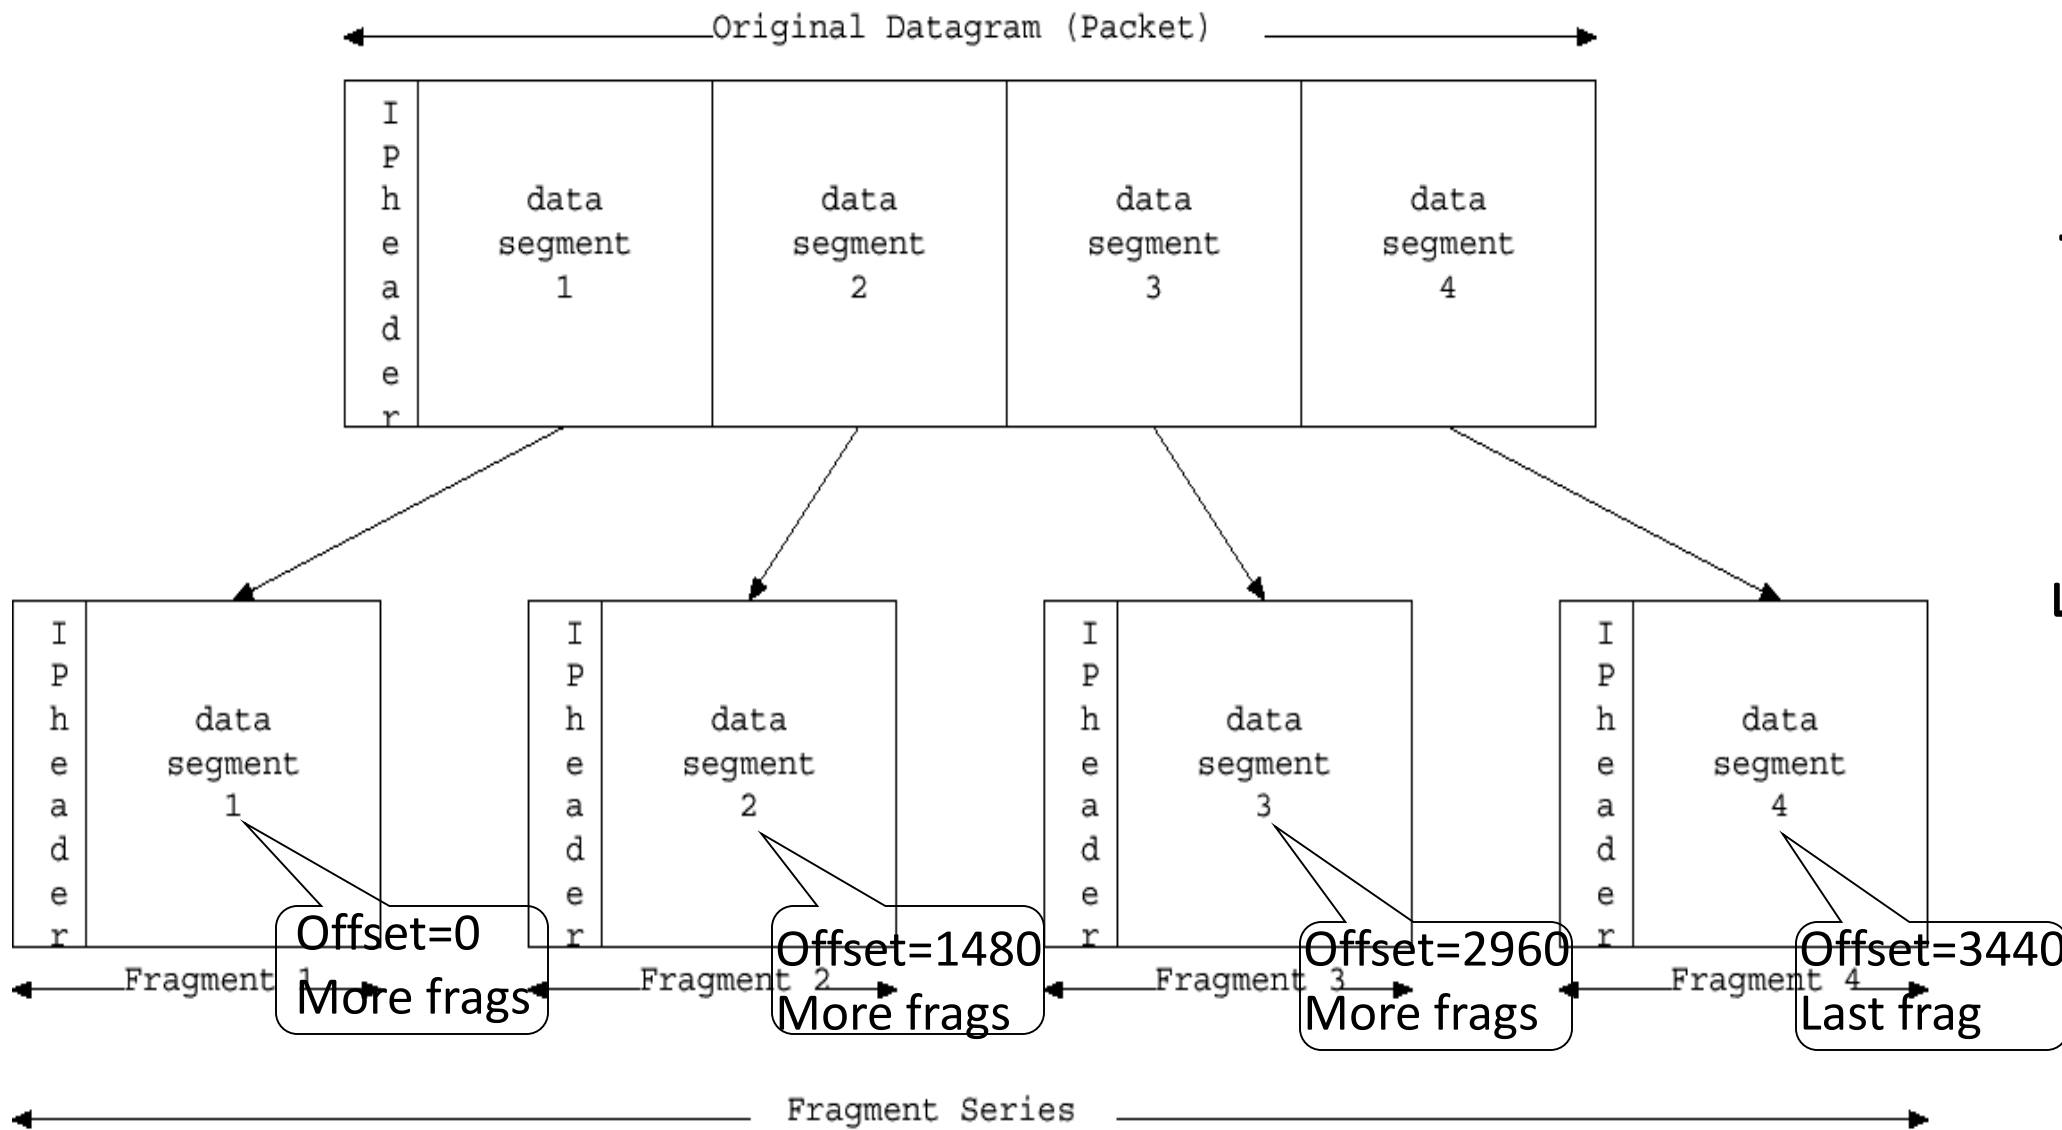
Fragment Offset (12 bits): specifies the offset in the original datagram of the data being carried in the fragment

- Measured in units of 8 bytes starting at 0

Flags (3 bits): control fragmentation

- Reserved (0-th bit)
- Don't Fragment – DF (1st bit):
 - useful for simple (computer bootstrap) application that can't handle
 - also used for MTU discovery (see later)
 - if need to fragment and can't router discards & sends error to source
- More Fragments (least sig bit): tells receiver it has got last fragment

TCP traffic is hardly ever fragmented (due to use of MTU discovery). About 0.5% - 0.1% of TCP packets are fragmented .



Fragment series composition

NB. If data segment contains its own header that is not replicated

Address Resolution Protocol



- The computers on a local network use an Internet layer protocol called **Address Resolution Protocol (ARP)** to map IP addresses to physical addresses.
- A host must know the physical address of the destination network adapter to send any data to it.
- For this reason, ARP is an important protocol.
- TCP/IP is implemented in such a way that ARP and all the details of physical address translation are almost totally invisible to the user.
- As far as the user is concerned, a network adapter is identified by its IP address.
- The IP address must be mapped to a physical address for a message to reach its destination

Address Resolution Protocol

Each host on a network segment maintains a table in memory called the ARP table or ARP cache.

The ARP cache associates the IP addresses of other hosts on the network segment with physical addresses.

When a host needs to send data to another host on the segment, the host checks the ARP cache to determine the physical address of the recipient.

The ARP cache is assembled dynamically. If the address that is to receive the data is not currently listed in the ARP cache, the host sends a broadcast called an ARP request frame.

IP: 206.154.13.82
physical: 00-E0-98-07-8E-39



IP: 206.154.13.83
physical: 35-00-21-01-31



IP: 206.154.13.84
physical: 44-45-53-54-00-00



IP: 206.154.13.85
physical: 91-03-20-51-09-26

00-E0-98-07-8E-39	206.154.13.82
35-00-21-01-3B-14	206.154.13.83
44-45-53-54-00-00	206.154.13.84
•	•
•	•
•	•

Address Resolution Protocol

The ARP request frame contains the unresolved IP address.

The ARP request frame also contains the IP address and physical address of the host that sent the request.

The other hosts on the network segment receive the ARP request, and the host that owns the unresolved IP address responds by sending its physical address to the host that sent the request.

The newly resolved IP address-to-physical address mapping is then added to the ARP cache of the requesting host

Reverse ARP

RARP stands for Reverse ARP. RARP is the opposite of ARP. ARP is used when the IP address is known but the physical address is not known. RARP is used when the physical address is known but the IP address is not known.

RARP is often used in conjunction with the **BOOTP** protocol to boot diskless workstations.

Many network adapters contain an empty socket for insertion of an integrated circuit known as a boot PROM.

The boot PROM firmware starts as soon as the computer is powered on. It loads an operating system into the computer by reading it from a network server instead of a local disk drive.

The operating system downloaded to the BOOTP device is preconfigured for a specific IP address.

Internet Control Message Protocol

Data sent to a remote computer often travels through one or more routers; these routers can encounter a number of problems in sending the message to its ultimate destination.

Routers use **Internet Control Message Protocol (ICMP)** messages to notify the source IP of these problems.

ICMP is also used for other diagnosis and troubleshooting functions.

Internet Control Message Protocol

- **Echo Request and Echo Reply:** ICMP is often used during testing.
 - A technician who uses the ping command to check connectivity with another host is using ICMP.
 - The ping command sends a datagram to an IP address and requests the destination computer to return the data sent in a response datagram.
 - The commands actually used by ping are the ICMP Echo Request and Echo Reply.
- **Source Quench:** If a fast computer is sending large amounts of data to a remote computer, the volume can overwhelm the router.
 - The router might use ICMP to send a Source Quench message to the source IP to ask it to slow down the rate at which it is shipping data.
 - If necessary, additional source quenches can be sent to the source IP.

Internet Control Message Protocol

- **Destination Unreachable:** If a router receives a datagram that cannot be delivered, ICMP returns a Destination Unreachable message to the source IP.
 - One reason that a router cannot deliver a message is a network that is down because of equipment failure or maintenance.
- **Time Exceeded:** ICMP sends this message to the source IP if a datagram is discarded because TTL reaches 0.
 - This indicates that the destination is too many router hops away to reach with the current TTL value, or it indicates router table problems that cause the datagram to loop through the same routers continuously

Internet Control Message Protocol

- A **routing loop** occurs when a datagram circulates endlessly and never reaches its destination.
 - Suppose three routers are located in Los Angeles, San Francisco, and Denver. The Los Angeles router sends datagrams to San Francisco, which sends them to Denver, which sends them back to Los Angeles again. The datagram becomes trapped and will circulate continuously through these three routers until the TTL reaches 0. A routing loop should not occur, but occasionally it does. Routing loops sometimes occur when a network administrator places static routing entries in a routing table.
- **Fragmentation Needed:** ICMP sends this message if it receives a datagram with the Don't Fragment bit set and if the router needs to fragment the datagram to forward it to the next router or the destination.

Other Internet Layer Protocols

- A number of other protocols also inhabit the Internet layer. Some of these other protocols, such as Border Gateway Protocol (BGP) and Routing Information Protocol (RIP), facilitate the routing process.
- The IPsec protocols, which are optional in IPv4 but are an integral part of IPv6, operate at the Internet layer to provide secure encrypted communication.
- Other Internet layer protocols assist with tasks such as multicasting. As mentioned earlier, the Internet protocol layer is known in OSI shorthand as Layer 3.
- Any protocol referred to as a Layer 3 protocol is operating at the Internet layer.

Q&A

Q. What common address notation is used to simplify a 32-bit binary address?

A. Dotted-decimal notation.

Q. ARP returns what type of information when given an IP address?

A. The corresponding physical (or MAC) address.

Q. If a router is unable to keep up with the volume of traffic, what type of ICMP message is sent to the source IP?

A. A Source Quench message.

Q. What class does an IP address belong to that starts with the binary pattern 110 as the 3 leftmost bits?

A Class C network.

Quiz

- What is the purpose of the TTL field in the IP header?
- How big are the network and host ID fields for a Class A address?
- What is an octet?
- What is the IP address an address of?
- What is the difference between ARP and RARP?

Exercises

11001111 00001110 00100001 01011100

Answer = 207.14.33.92

00001010 00001101 01011001 01001101

Answer = 10.13.89.77

10111101 10010011 01010101 01100001

Answer = 189.147.85.97

- Convert the following 32-bit IP addresses into dotted-decimal notation:

Key Terms

- **Address Class:** A classification system for IP addresses. The network class determines how the address is subdivided into a network ID and host ID.
- **Address Resolution Protocol (ARP):** A key Internet layer protocol used to obtain the physical address associated with an IP address. ARP maintains a cache of recently resolved physical address-to-IP address pairs.
- **BOOTP:** A protocol used to boot a computer or other network device from a remote location.
- **Dotted Decimal:** Base 10 representation of a binary IP address using 4 numerals representing the 4 octets of the original address, separated by periods (209.121.131.14).
- **Host ID:** A portion of the IP address that refers to a node on the network. Each node within a network should have an IP address that contains a unique host ID.


Key Terms

- **Internet Control Message Protocol (ICMP):** A key Internet layer protocol used by routers to send messages that inform the source IP of routing problems. ICMP is also used by the ping command to determine the status of other hosts on the network.
- **Internet Protocol (IP):** A key Internet layer protocol used for addressing, delivering, and routing datagrams.
- **Multicast:** A technique that allows datagrams to be delivered to a group of hosts simultaneously.
- **Network ID:** A portion of the IP address that identifies the network.
- **Octet:** An eight-digit binary number.
- **Reverse Address Resolution Protocol (RARP):** A TCP/IP protocol that returns an IP address if given a physical address. This protocol is typically used by a diskless workstation that has a remote boot PROM installed in its network adapter.
- **Subnet:** A logical division of a TCP/IP address space.

IP Addressing & Subnetting

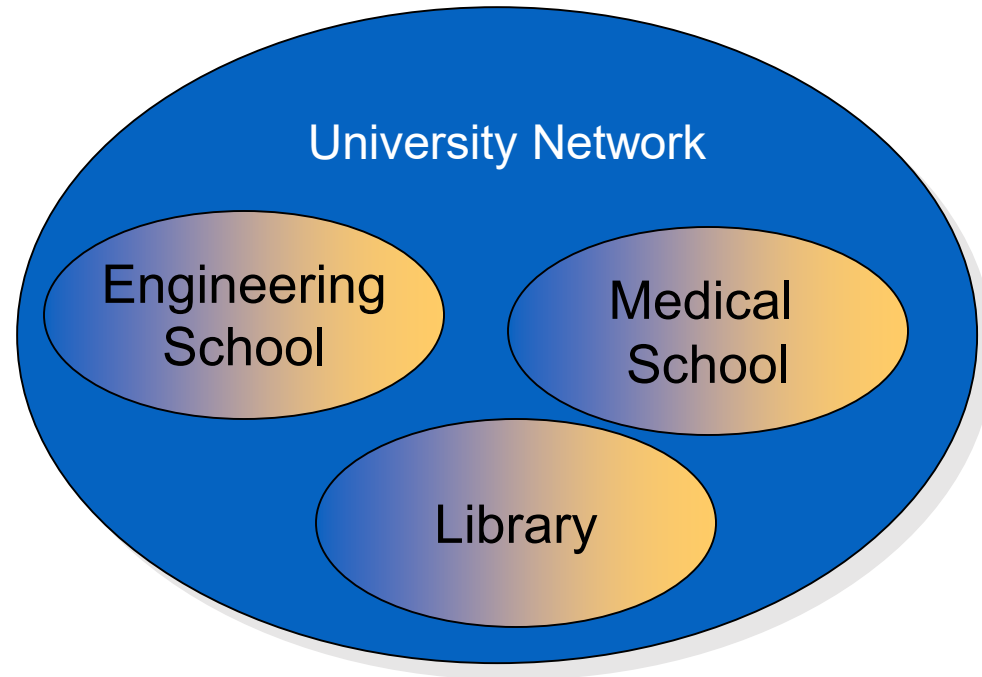


- IP Addressing
- Subnetting
- VLSM
- CIDR



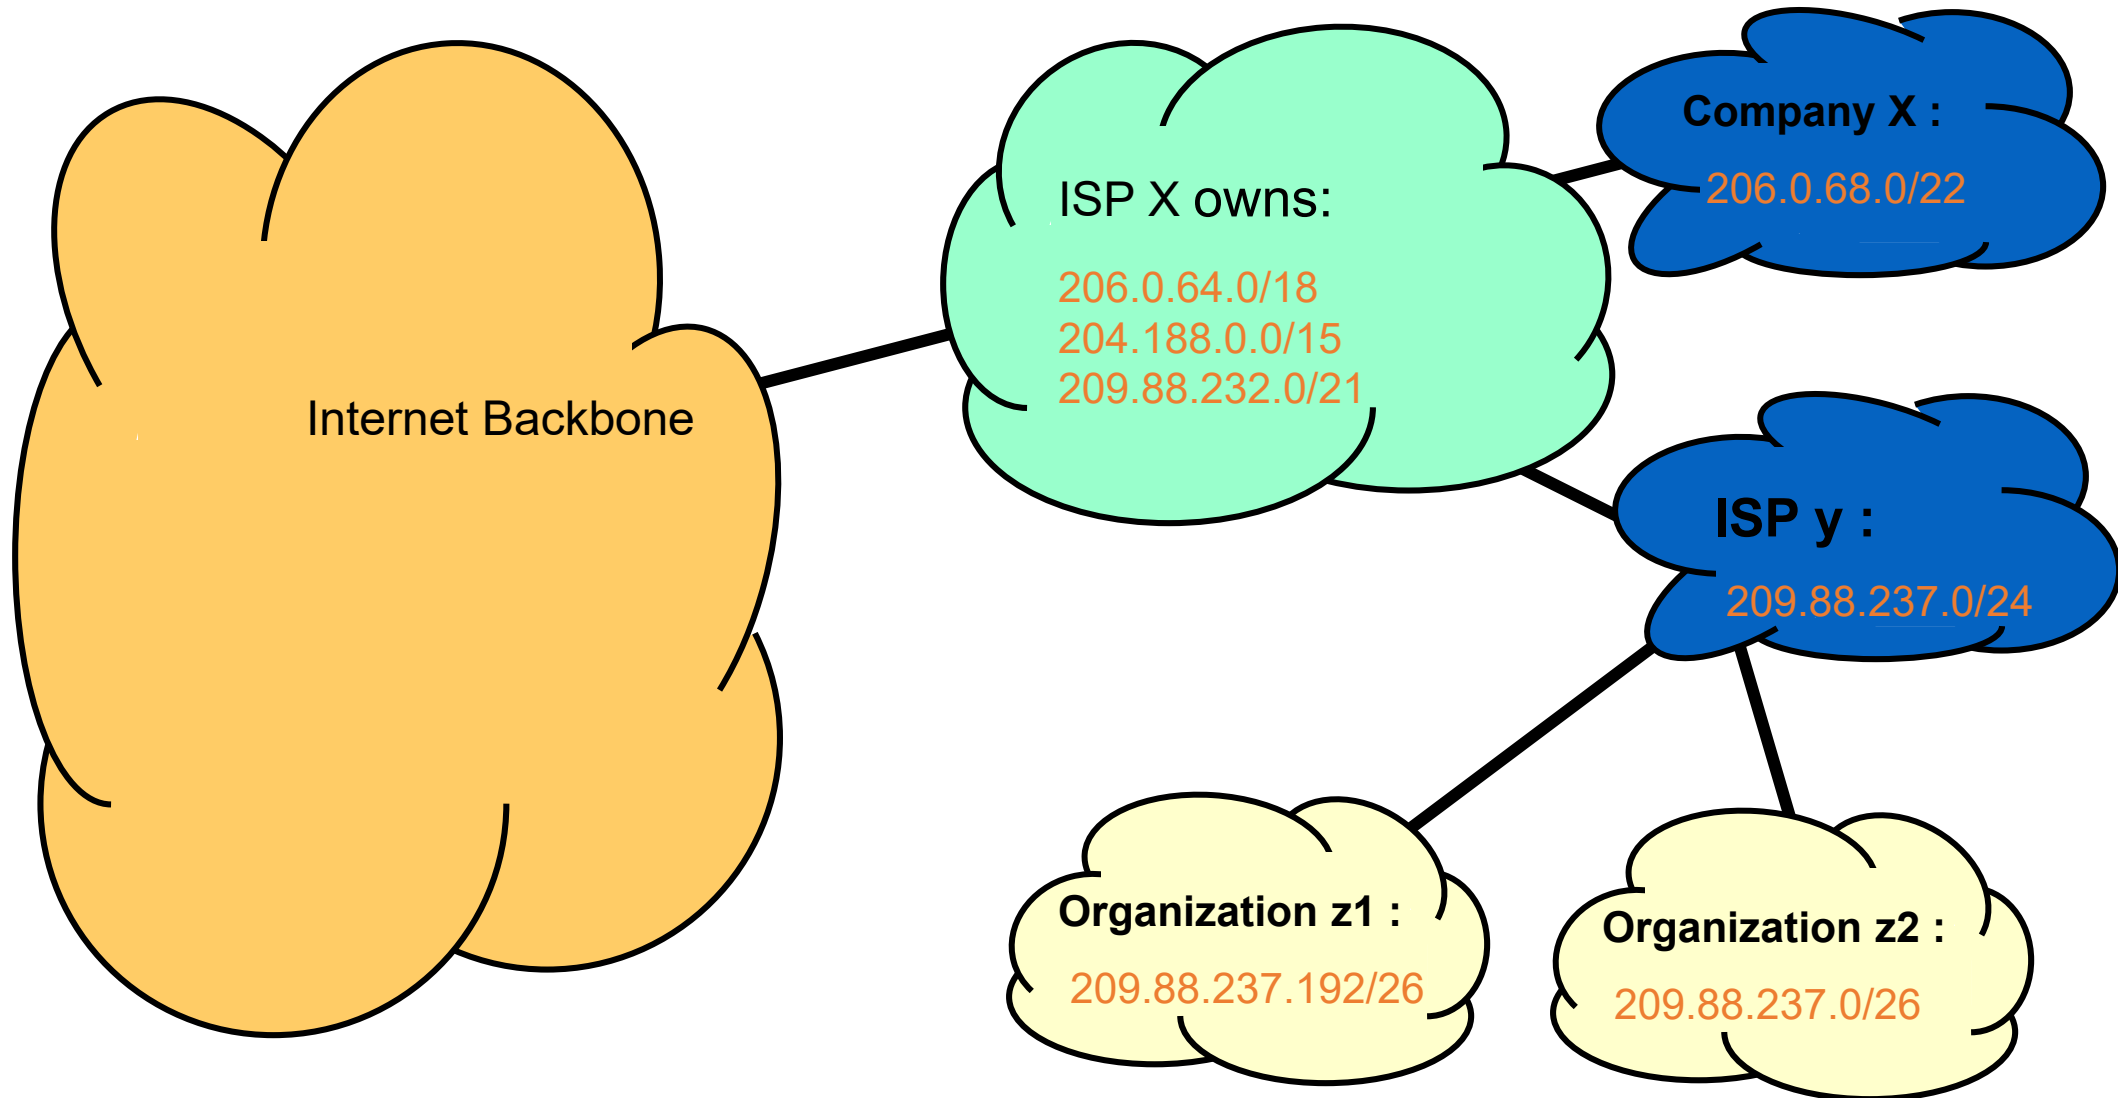
Subnetting

- **Problem:** Organizations have multiple networks which are independently managed
 - **Solution 1:** Allocate one or more addresses for each network
 - Difficult to manage
 - From the outside of the organization, each network must be addressable.
 - **Solution 2:** Add another level of hierarchy to the IP addressing structure

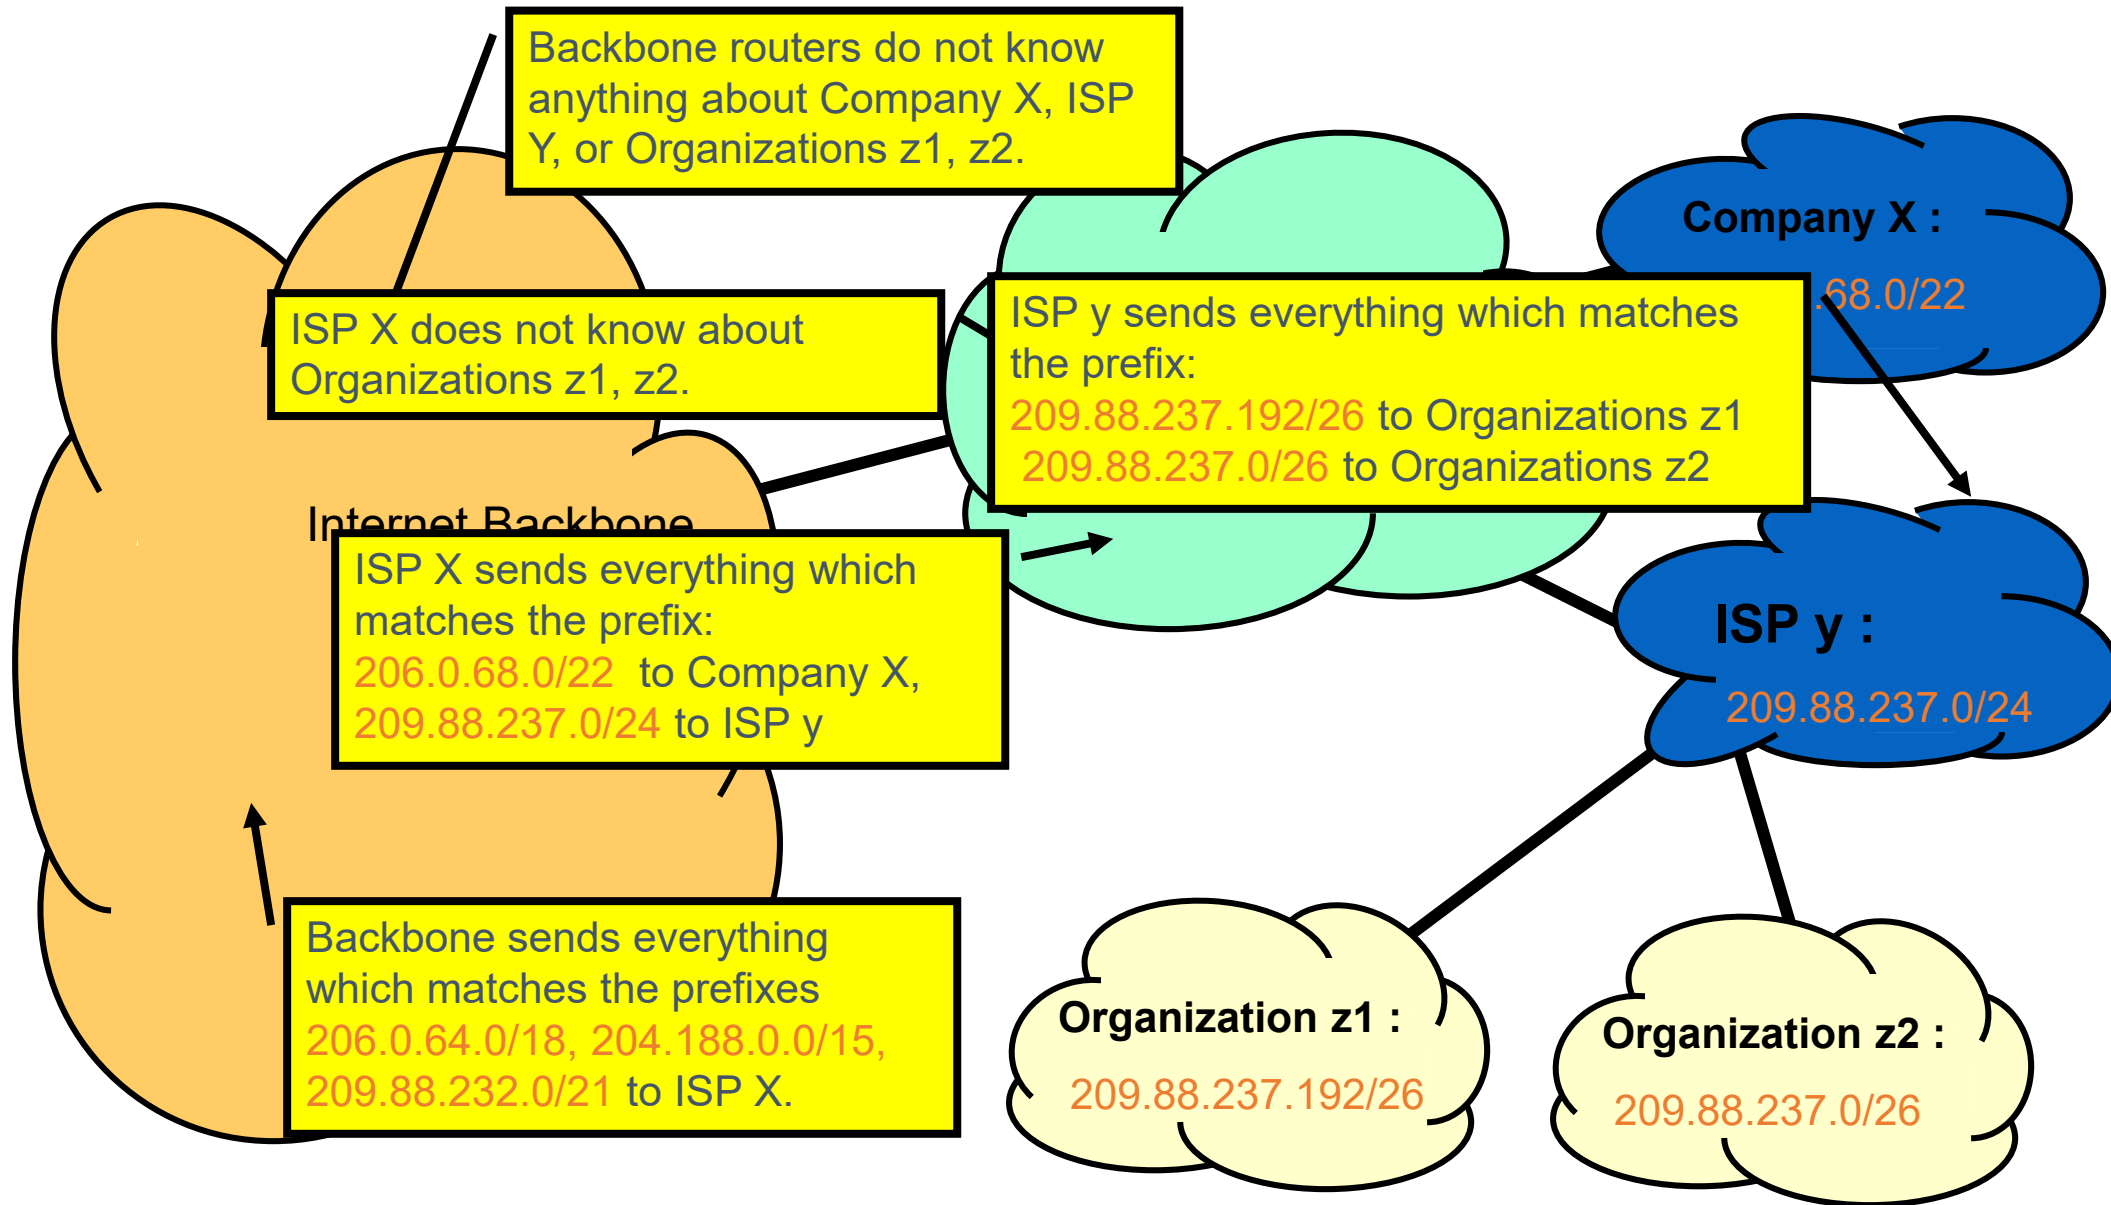


 **Subnetting**

CIDR and Routing Information



CIDR and Routing Information




What is it ?

Subnetting evolved as a means for using IP addressing to break up a physical network into smaller logical entities called subnets.

Later developments, such as classless interdomain routing and IPv6 have reduced the need for the classical approach to subnetting, but these later techniques borrow from the basic subnetting principles, and no discussion of TCP/IP is complete without a description of subnetting.

Dividing the Network



Datagrams arrive efficiently at the gateway and pass into the 99.0.0.0 address space.

However, the picture gets more complicated when you consider how to deliver the datagram after it passes into the 99.0.0.0 address space.

A Class A network has room for over 16 million host IDs.

Advantages of Subnetting

- With subnetting, IP addresses use a 3-layer hierarchy:
 - Network
 - Subnet
 - Host
- Reduces router complexity. Since external routers do not know about subnetting, the complexity of routing tables at external routers is reduced.

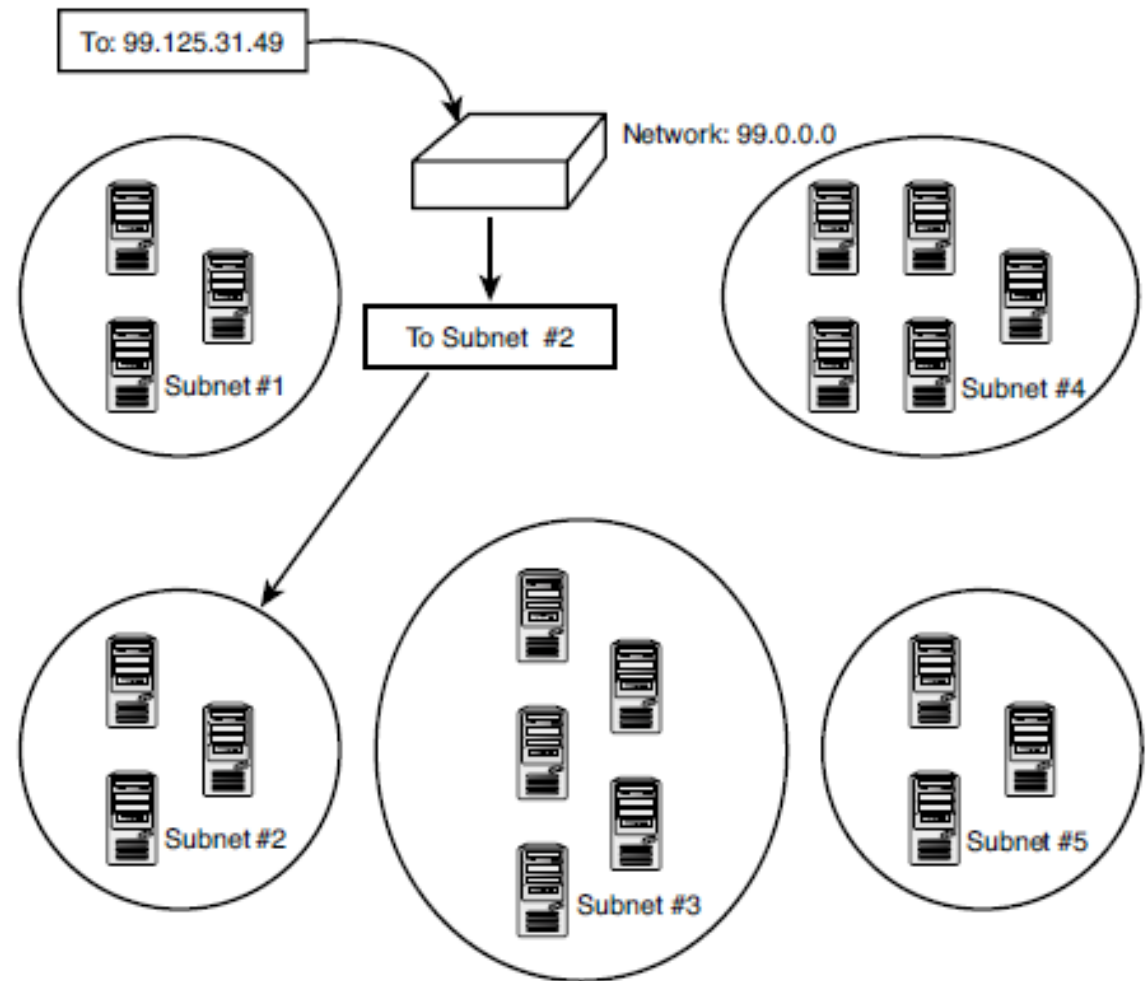
Designing Subnets

An organization of ISP that is granted a range of addresses may divide the range into several subranges and assign each subrange to a subnetwork(or subnet).

A subnetwork can be divided into several sub-subnetworks.

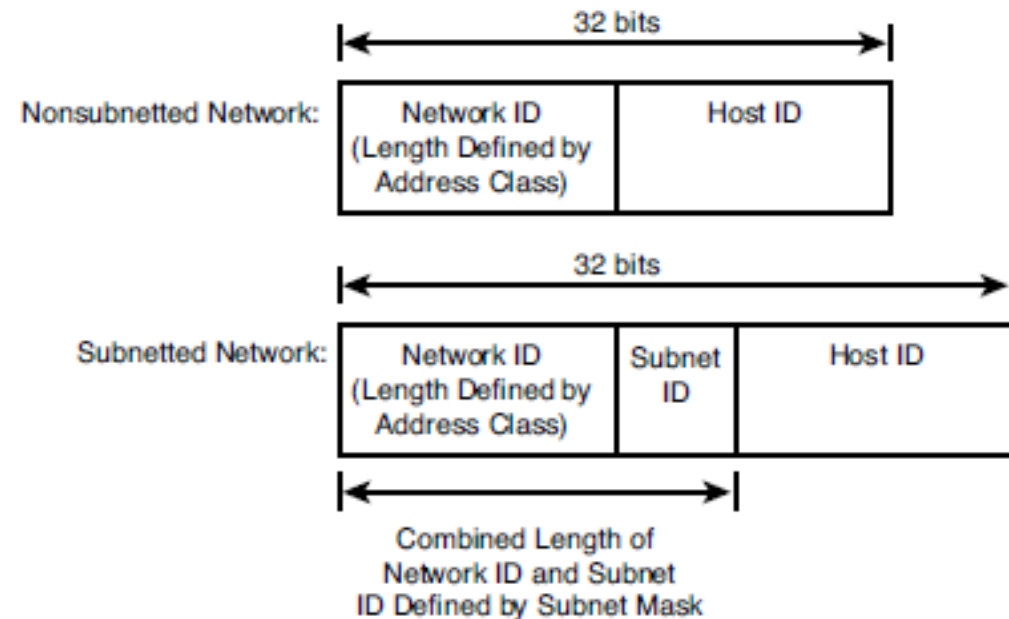
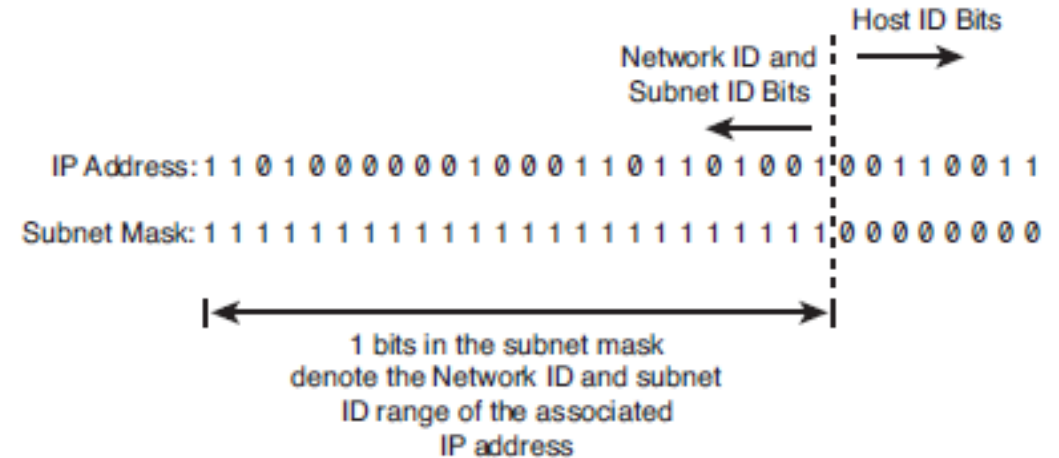
The subnetwork in a network should be carefully designed to enable the routing of packets.

- The routers can deliver a datagram to a subnet address within the network (generally corresponding to a network segment),
- and when the datagram reaches the subnet, it can be resolved to a physical address using ARP



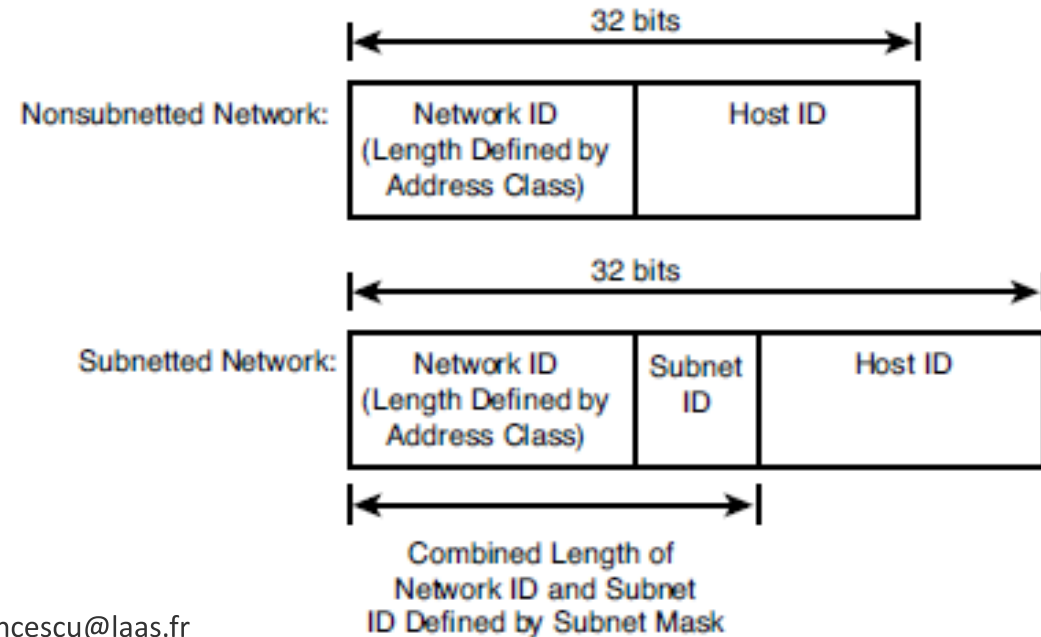
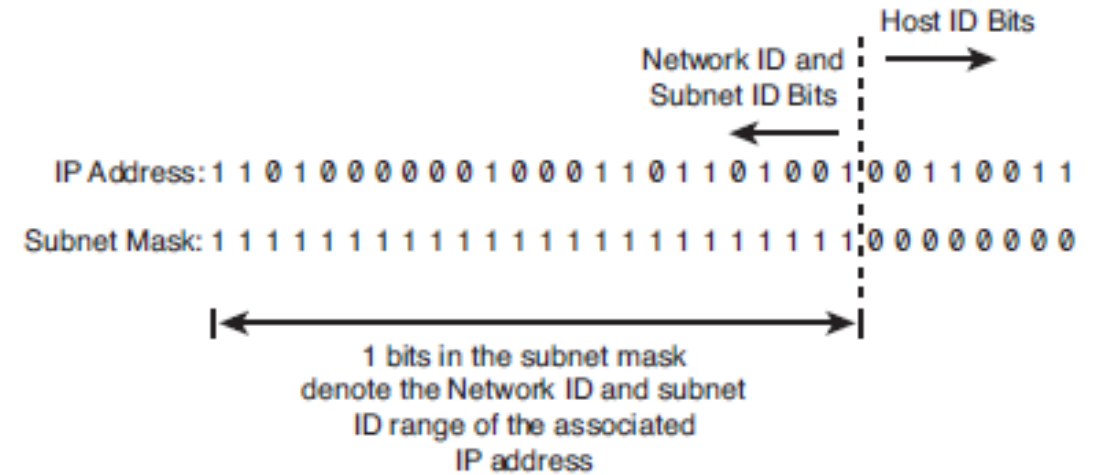
Organizing the network for efficient delivery.

- A parameter called the **subnet mask** tells how much of the address should be used for the subnet ID and how much is left for the actual host ID.



Subnet Mask

- The subnet mask uses a 1 for every bit in the IP address that is part of the network ID or subnet ID.
- The subnet mask uses a 0 to designate any bit in the IP address that is part of the host ID.



IP Addresses

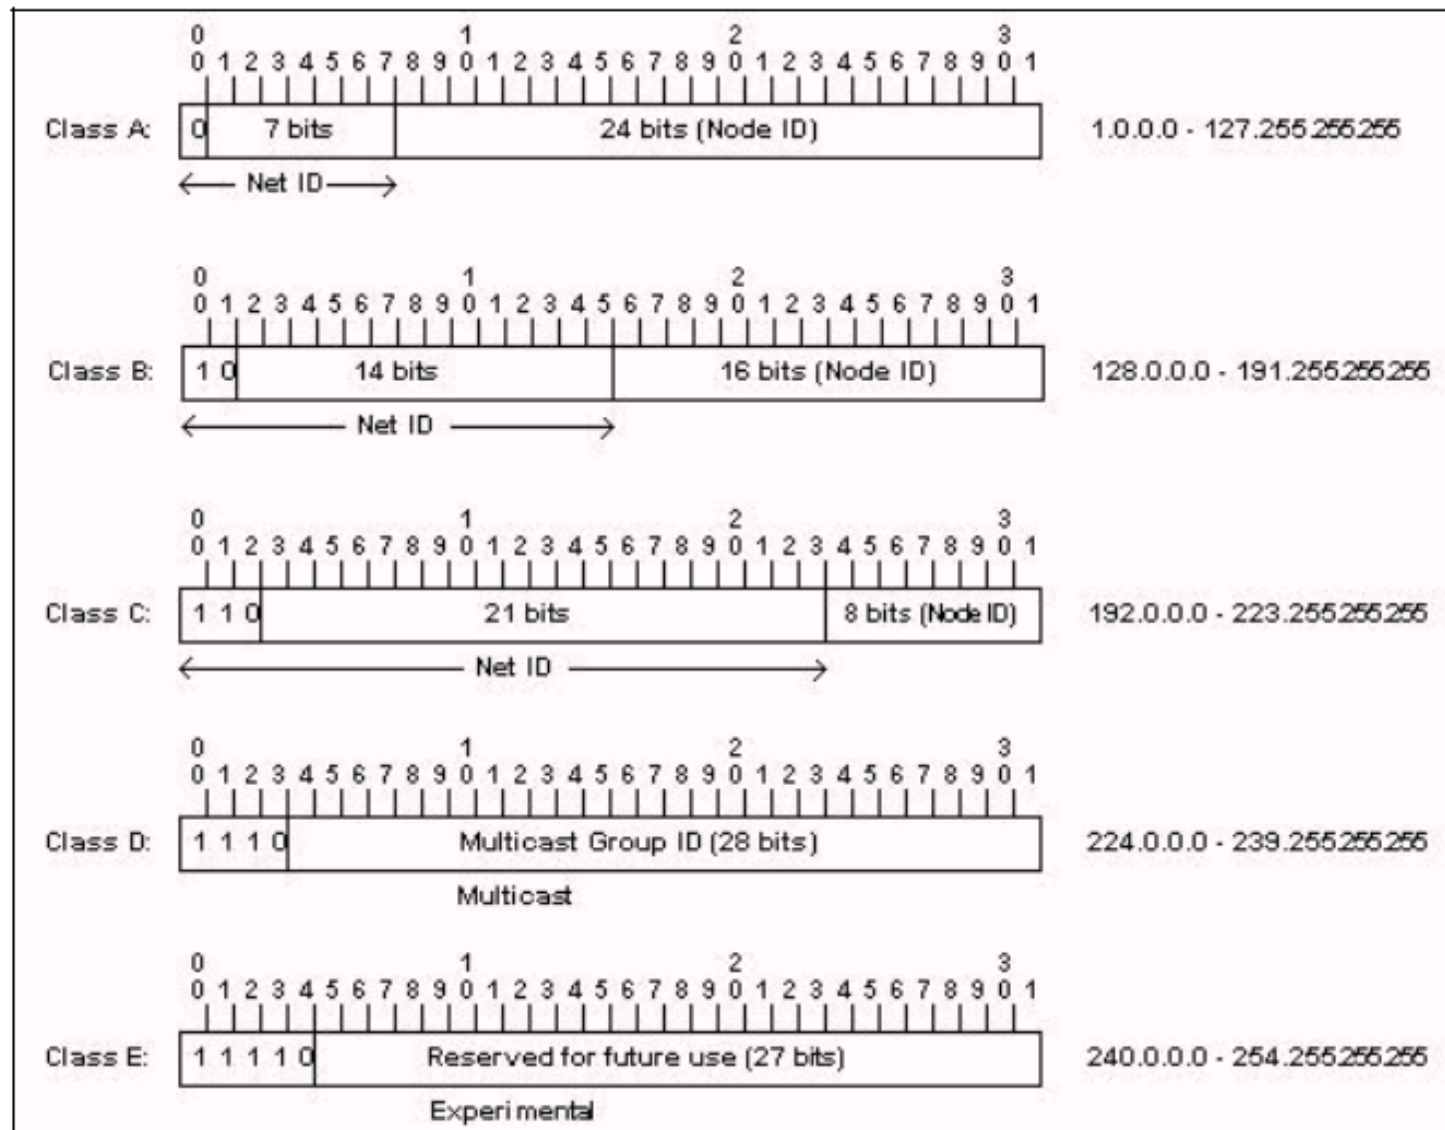
An IP address is an address used to uniquely identify a device on an IP network.

The address is made up of 32 binary bits which can be divisible into a network portion and host portion with the help of a subnet mask.

32 binary bits are broken into four octets (1 octet = 8 bits)

Dotted decimal format (for example, 172.16.81.100)

IP Address Classes



IP Address Classes



Class A: The first octet is the network portion. Octets 2, 3, and 4 are for subnets/hosts



Class B: The first two octets are the network portion. Octets 3 and 4 are for subnets/hosts



Class C: The first three octets are the network portion. Octet 4 is for subnets/hosts

Private Address Range

Address Class	Reserved Address Space
Class A	10.0.0.0 - 10.255.255.255
Class B	172.16.0.0 - 172.31.255.255
Class C	192.168.0.0 - 192.168.255.255



The number of addresses in each subnetwork should be a power of 2.



The prefix length for each subnetwork should be found using the following formula:



$$n_{\text{sub}} = 32 - \log_2 N_{\text{sub}}$$



Assign addresses to larger subnetworks first.

Network Masks

Distinguishes which portion of the address identifies the network and which portion of the address identifies the node.

Default masks:

- Class A: 255.0.0.0
- Class B: 255.255.0.0
- Class C: 255.255.255.0

Subnetting



Creates multiple logical networks that exist within a single Class A, B, or C network.



If you do not subnet, you will only be able to use one network from your Class A, B, or C network, which is unrealistic



Each data link on a network must have a unique network ID, with every node on that link being a member of the same network

Benefits of Subnetting

Reduced network traffic

Optimized network performance

Simplified management

Facilitated spanning of large geographical distances

IP Subnet-Zero

This command allows you to use the first and last subnet in your network design.

For example, the Class C mask of 192 provides subnets 64 and 128, but with the IP subnet-zero command, you now get to use subnets 0, 64, 128, and 192

How to create subnets

Determine the number of required network IDs:

- One for each subnet
- One for each wide area network connection

Determine the number of required host IDs per subnet:

- One for each TCP/IP host
- One for each router interface

Based on the above requirements, create the following:

- One subnet mask for your entire network
- A unique subnet ID for each physical segment
- A range of host IDs for each subnet

Subnetting a Class A/B/C Address

How many subnets does the chosen subnet mask produce?

How many valid hosts per subnet are available?

What are the valid subnets?

What's the broadcast address of each subnet?

What are the valid hosts in each subnet?

Practice Example #1C: 255.255.255.128 (/25) Network 192.168.10.0



How many subnets? Since 128 is 1 bit on (10000000), the answer would be $2^1 = 2$.



How many hosts per subnet? We have 7 host bits off (10000000), so the equation would be $2^7 - 2 = 126$ hosts.



What are the valid subnets? $256 - 128 = 128$. Remember, we'll start at zero and count in our block size, so our subnets are 0, 128.




What's the broadcast address for each subnet? The number right before the value of the next subnet is all host bits turned on and equals the broadcast address. For the zero subnet, the next subnet is 128, so the broadcast of the 0 subnet is 127.



What are the valid hosts? These are the numbers between the subnet and broadcast address

Logical Network Implementation



Example

- An organization is granted a block of addresses with the beginning address 14.24.74.0/24. The organization needs to have 3 subblocks to use in three subnets:
 - One subblock of 10 addresses
 - One subblock of 60 addresses,
 - One subblock of 120 addresses.

Cont.

There are $2^{32-24} = 256$ addresses.

First address is 14.24.74.0

Last address is 14.24.74.255

To assign subnets, start with the larger block.

120 is not a power of 2, we allocate 128 addresses.

(subnet mask) $n_1 = 32 - \log_2 128 = 25$,

- the first address in this block is 14.24.74.0/25
- The last address is 14.24.74.127/25

Cont.

Number of addresses in the second block, which requires 60 addresses, is 64 since 60 is not a power of 2.

$$n_2 = 32 - \log_2 64 = 26$$

The first address is 14.24.74.128/26

The last address is 14.24.74.191/26

Cont.

The number of addresses in the smallest block is 10, we allocate 16 since 10 is not a power of 2.

The subnet mask
 $n_3 = 32 - \log_2 16$
 $= 28$.

The first address is
14.24.74.192/28

The last address is
14.24.74.207/28

The sum of addresses in the subnets is 208 addresses, which means 48 addresses are left in reserve. The first address in the range is 14.24.74.208 and the last address is 14.24.74.255.

We don't know about the prefix length yet.

Practice Example #2C: 255.255.255.224 (/27) Network 192.168.10.0

- How many subnets? 224 is 11100000, so our equation would be $2^3 = 8$.
- How many hosts? $2^5 - 2 = 30$.
- What are the valid subnets? $256 - 224 = 32$. We just start at zero and count to the subnet mask value in blocks (increments) of 32: 0, 32, 64, 96, 128, 160, 192, and 224.
- What's the broadcast address for each subnet (always the number right before the next subnet)?
- What are the valid hosts (the numbers between the subnet number and the broadcast address)?

Subnet Address	0	32	192	224
First Host	1	33		193	225
Last Host	30	62		222	254
Broadcast Address	31	63		223	255

**Practice Example #2C: 255.255.255.224 (/27)
Network 192.168.10.0**

Practice Example #1B: 255.255.128.0 (/17)

Network 172.16.0.0

- Subnets? $2^1 = 2$
- Hosts? $2^{15} - 2 = 32,766$ (7 bits in the third octet, and 8 in the fourth)
- Valid subnets? $256 - 128 = 128$. 0, 128. Remember that subnetting is performed in the third octet, so the subnet numbers are really 0.0 and 128.0, as shown in the next table
- Broadcast address for each subnet?
- Valid hosts?

Practice Example #1B: 255.255.128.0 (/17)
Network 172.16.0.0

Subnet	0.0	128.0
First Host	0.1	128.1
Last Host	127.254	255.254
Broadcast	127.255	255.255

Practice
Example #2B:
255.255.240.0
(/20)
Network
172.16.0.0

Subnets? $2^4 = 16$.

Hosts? $2^{12} - 2 = 4094$.

Valid subnets? $256 - 240 = 0, 16, 32, 48,$
etc., up to 240.

Broadcast address for each subnet?

Valid hosts?


Practice Example #2B: 255.255.240.0 (/20)
Network 172.16.0.0

Subnet	0.0	16.0	240.0
First Host	0.1	16.1		240.1
Last Host	15.254	31.254		255.254
Broadcast	15.255	31.255		255.255

Variable Length
Subnet
Mask(VLSM)
Subnet with
requirements
shown?

netA: 14 hosts

netD: 7 hosts



netC: 2 hosts

netB: 28 hosts

netE: 28 hosts


Variable
Length
Subnet
Mask(VLSM)
Subnet with
requirements
shown?

5 subnets needed

Can be assigned as follows:

- netA: 204.15.5.0/27 host address range 1 to 30
- netB: 204.15.5.32/27 host address range 33 to 62
- netC: 204.15.5.64/27 host address range 65 to 94
- netD: 204.15.5.96/27 host address range 97 to 126
- netE: 204.15.5.128/27 host address range 129 to 158

Variable Length Subnet Mask(VLSM) Subnet with requirements shown?



Variable Length Subnet Mask(VLSM) Subnet with requirements shown?

Given the same network and requirements as in Sample Exercise 1 develop a subnetting scheme using VLSM, given:

- netA: must support 14 hosts
- netB: must support 28 hosts
- netC: must support 2 hosts
- netD: must support 7 hosts
- netE: must support 28 host


Determine what mask allows the required number of hosts.

- netA: requires a /28 (255.255.255.240) mask to support 14 hosts
- netB: requires a /27 (255.255.255.224) mask to support 28 hosts
- netC: requires a /30 (255.255.255.252) mask to support 2 hosts
- netD: requires a /28 (255.255.255.240) mask to support 7 hosts
- netE: requires a /27 (255.255.255.224) mask to support 28 hosts

Variable Length Subnet Mask(VLSM) Subnet with requirements shown?

- The easiest way to assign the subnets is to assign the largest first. For example, you can assign in this manner:
 - netB: 204.15.5.0/27 host address range 1 to 30
 - netE: 204.15.5.32/27 host address range 33 to 62
 - netA: 204.15.5.64/28 host address range 65 to 78
 - netD: 204.15.5.80/28 host address range 81 to 94
 - netC: 204.15.5.96/30 host address range 97 to 98

Variable Length
Subnet
Mask(VLSM)
Subnet with
requirements
shown?



CIDR

Classless Interdomain Routing

Improve address space utilization

Routing scalability in the Internet

For example, if an ISP owns network 172.16.0.0/16, then the ISP can offer 172.16.1.0/24, 172.16.2.0/24, and so on to customers. Yet, when advertising to other providers, the ISP only needs to advertise 172.16.0.0/16

The Address Resolution Protocol (ARP)

MAC Address vs. IP Address


- MAC addresses
 - Hard-coded in read-only memory when adaptor is built
 - Like a social security number
 - Flat name space of 48 bits (e.g., 00-0E-9B-6E-49-76)
 - Portable, and can stay the same as the host moves
 - Used to get packet between interfaces on same network [To guide delivery between two hosts, two routers, and a host and router within a single subnet]
- IP addresses
 - Configured, or learned dynamically
 - Like a postal mailing address
 - Hierarchical name space of 32 bits (e.g., 12.178.66.9)
 - Not portable, and depends on where the host is attached
 - Used to get a packet to destination IP subnet [To guide delivery to destination host across the Internet (across multiple networks)]

Address Resolution


- Problem

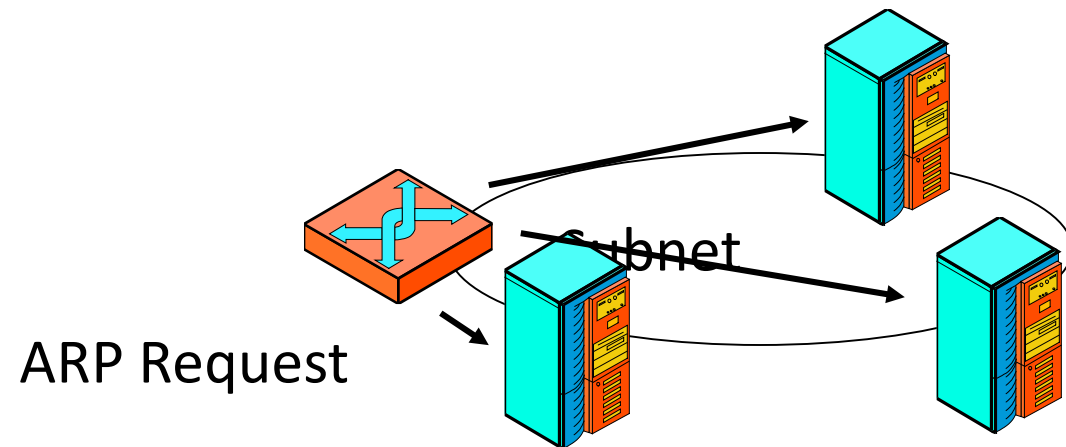


- Router knows that destination host is on its subnet based on the IP address of an arriving packet
- Does not know the destination host's subnet address, *so cannot deliver the packet across the subnet*



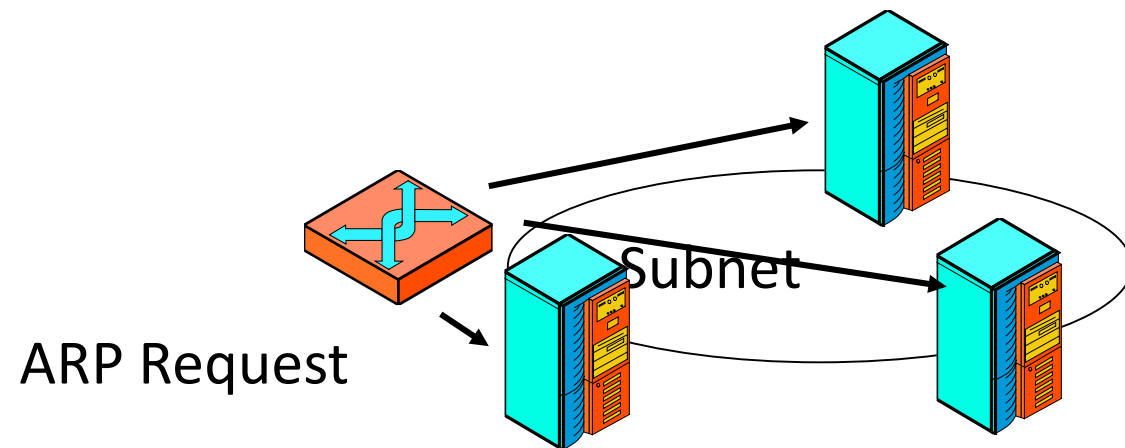
Address Resolution Protocol (ARP)

- Router creates an ARP Request message to be sent to all hosts on the subnet. 
- Address resolution protocol message asks “Who has IP address 128.171.17.13?”
- Passes ARP request to data link layer process for delivery




Address Resolution Protocol (ARP)

- Data link process of router broadcasts the ARP Request message to all hosts on the subnet.
 - On a LAN, MAC address of 48 ones tells all stations to pay attention to the frame



Address Resolution Protocol (ARP)

- Host with IP address 128.171.17.13 responds
 - Internet process creates an ARP response message
 - Contains the destination host's subnet address (48-bit MAC address on a LAN)

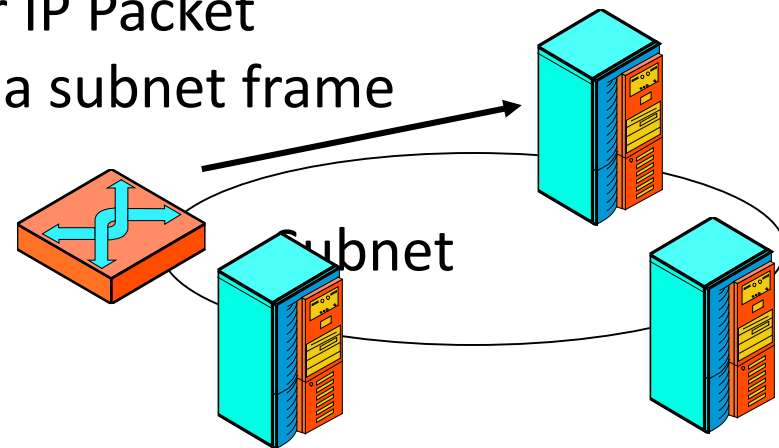


Address Resolution Protocol (ARP)

- Router delivers the IP packet to the destination host
 - Places the IP packet in the subnet frame
 - Puts the *destination host's subnet address* in the destination address field of the frame

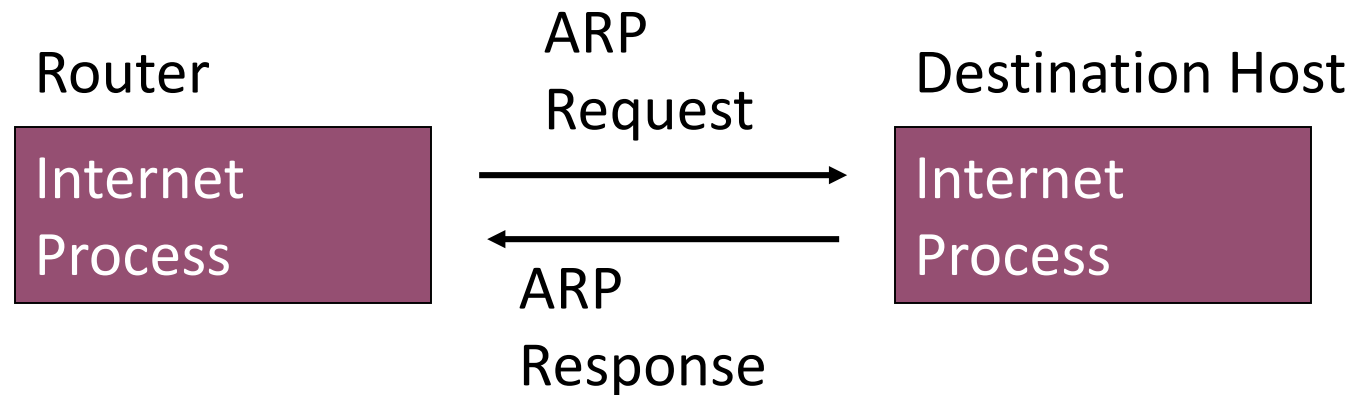


Deliver IP Packet
within a subnet frame



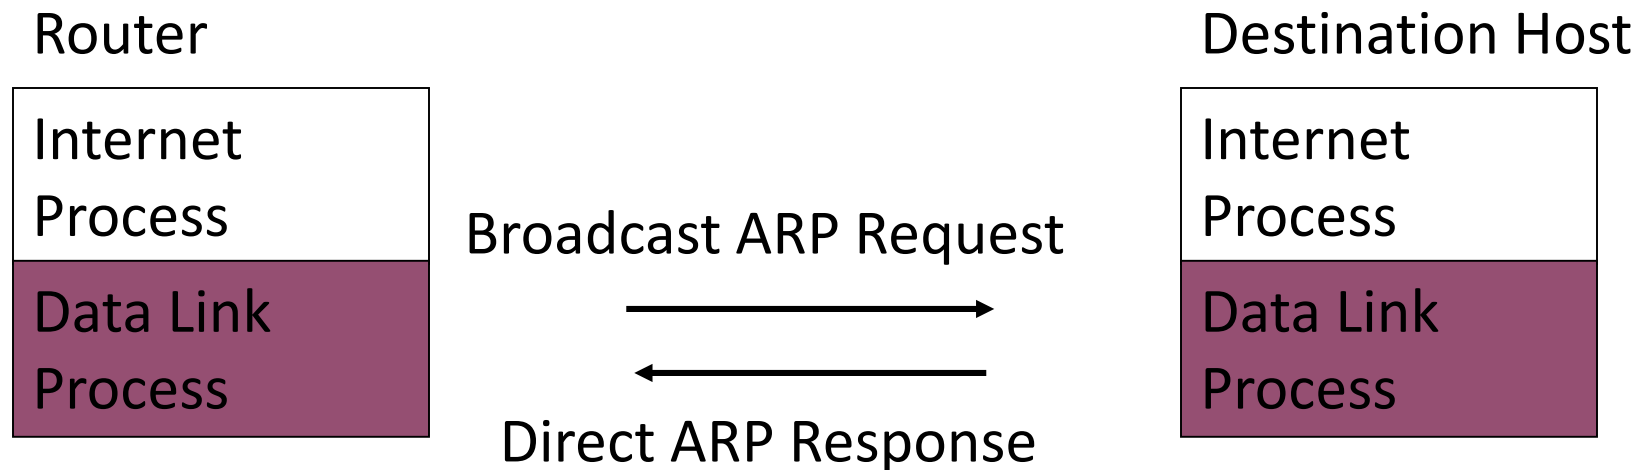
Address Resolution Protocol

- ARP Requests and Responses are sent between the internet layer processes on the router and the destination host




Address Resolution Protocol

- However, the data link processes deliver these ARP packets
 - Router broadcasts the ARP Request
 - Destination host sends ARP response to the subnet source address found in the broadcast frame



Example: A Sending a Packet to B


How does host A send an IP packet to host B?




A sends packet to R, and R sends packet to B.

Host A Decides to Send Through R

- Host A constructs an IP packet to send to B
 - Source 111.111.111.111, destination 222.222.222.222
- Host A has a gateway router R
 - Used to reach destinations outside of 111.111.111.0/24
 - Address 111.111.111.110 for R learned via DHCP



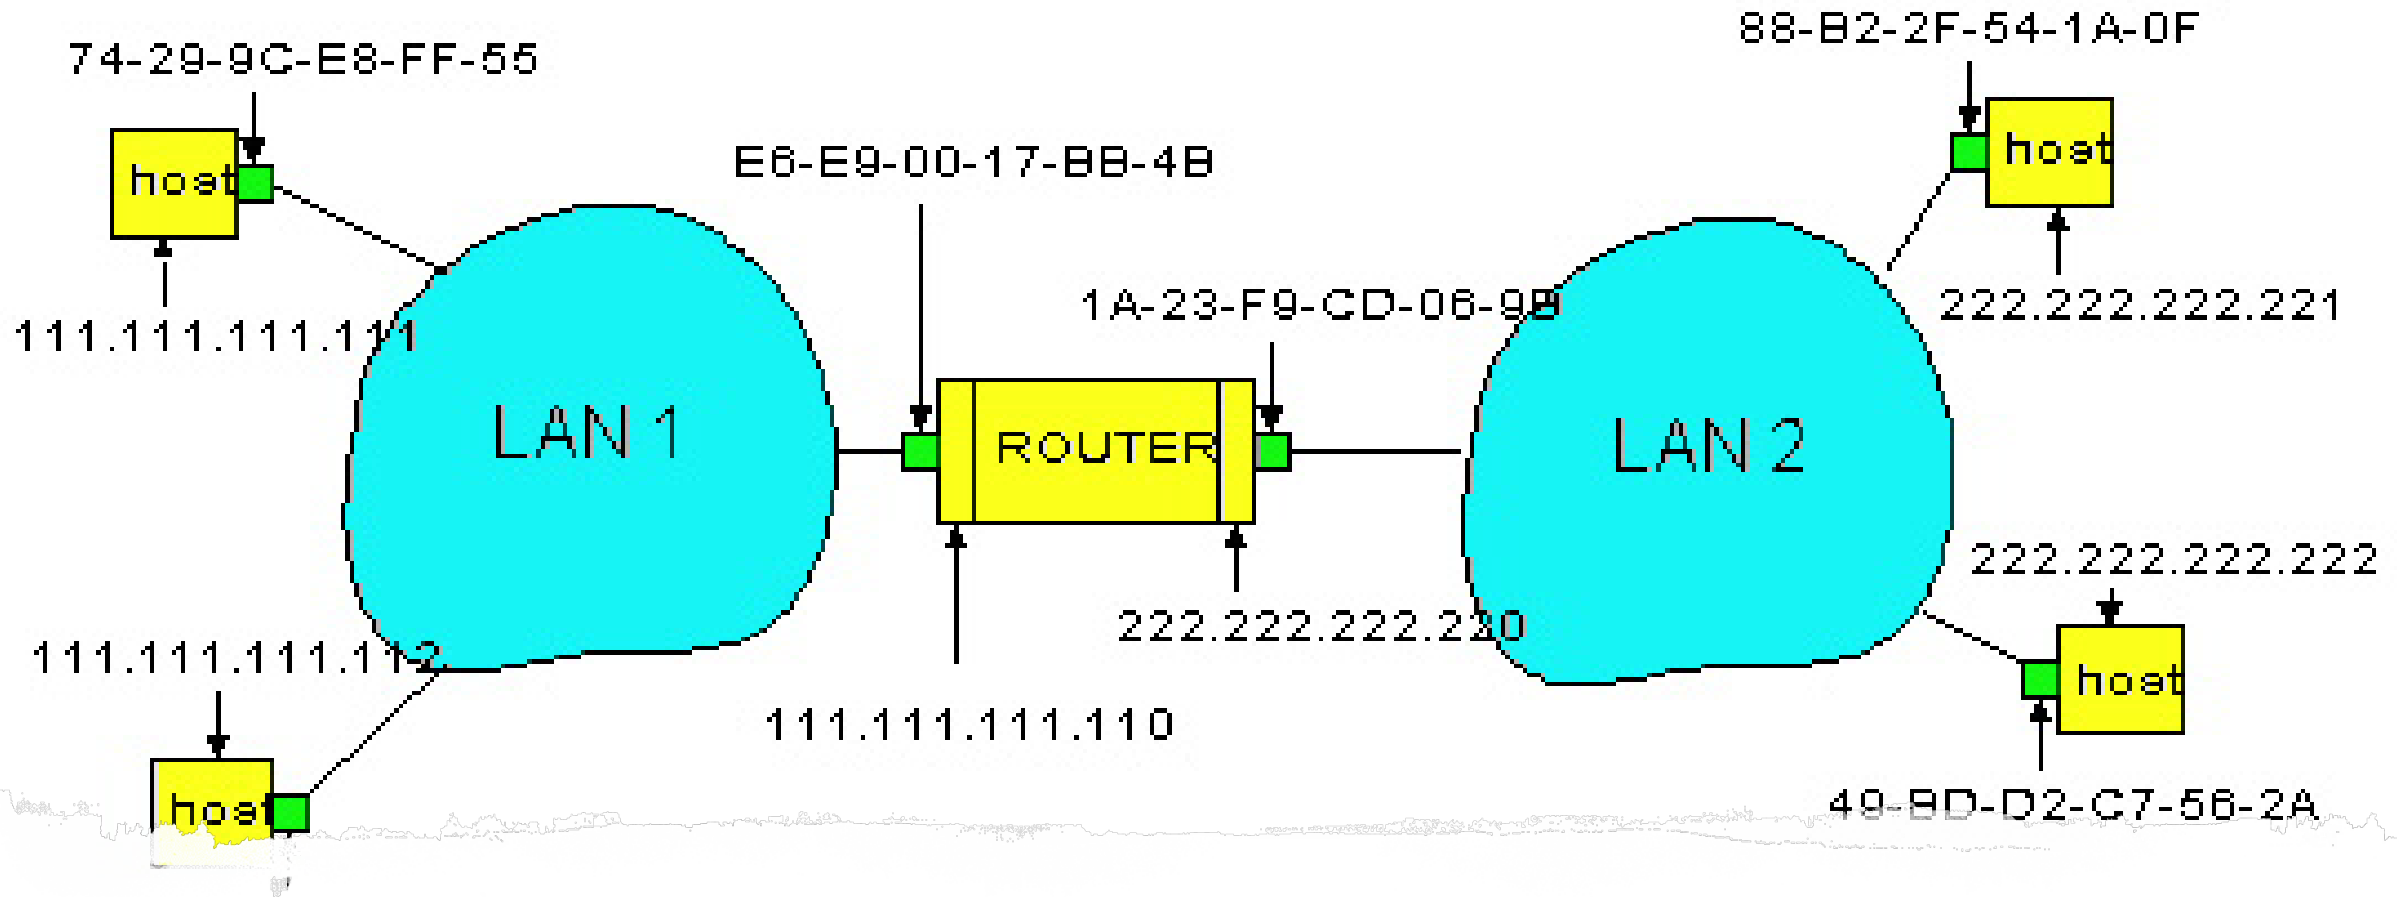


Host A ^A Sends Packet Through R

- Host A learns the MAC address of R's interface
 - ARP request: broadcast request for 111.111.111.110
 - ARP response: R responds with E6-E9-00-17-BB-4B
- Host A encapsulates the packet and sends to R

R

B



R Decides ^A how to Forward Packet


- Router R's adaptor receives the packet
 - R extracts the IP packet from the Ethernet frame
 - R sees the IP packet is destined to 222.222.222.222
- Router R consults its forwarding table
 - Packet matches 222.222.222.0/24 via other adaptor

R

B

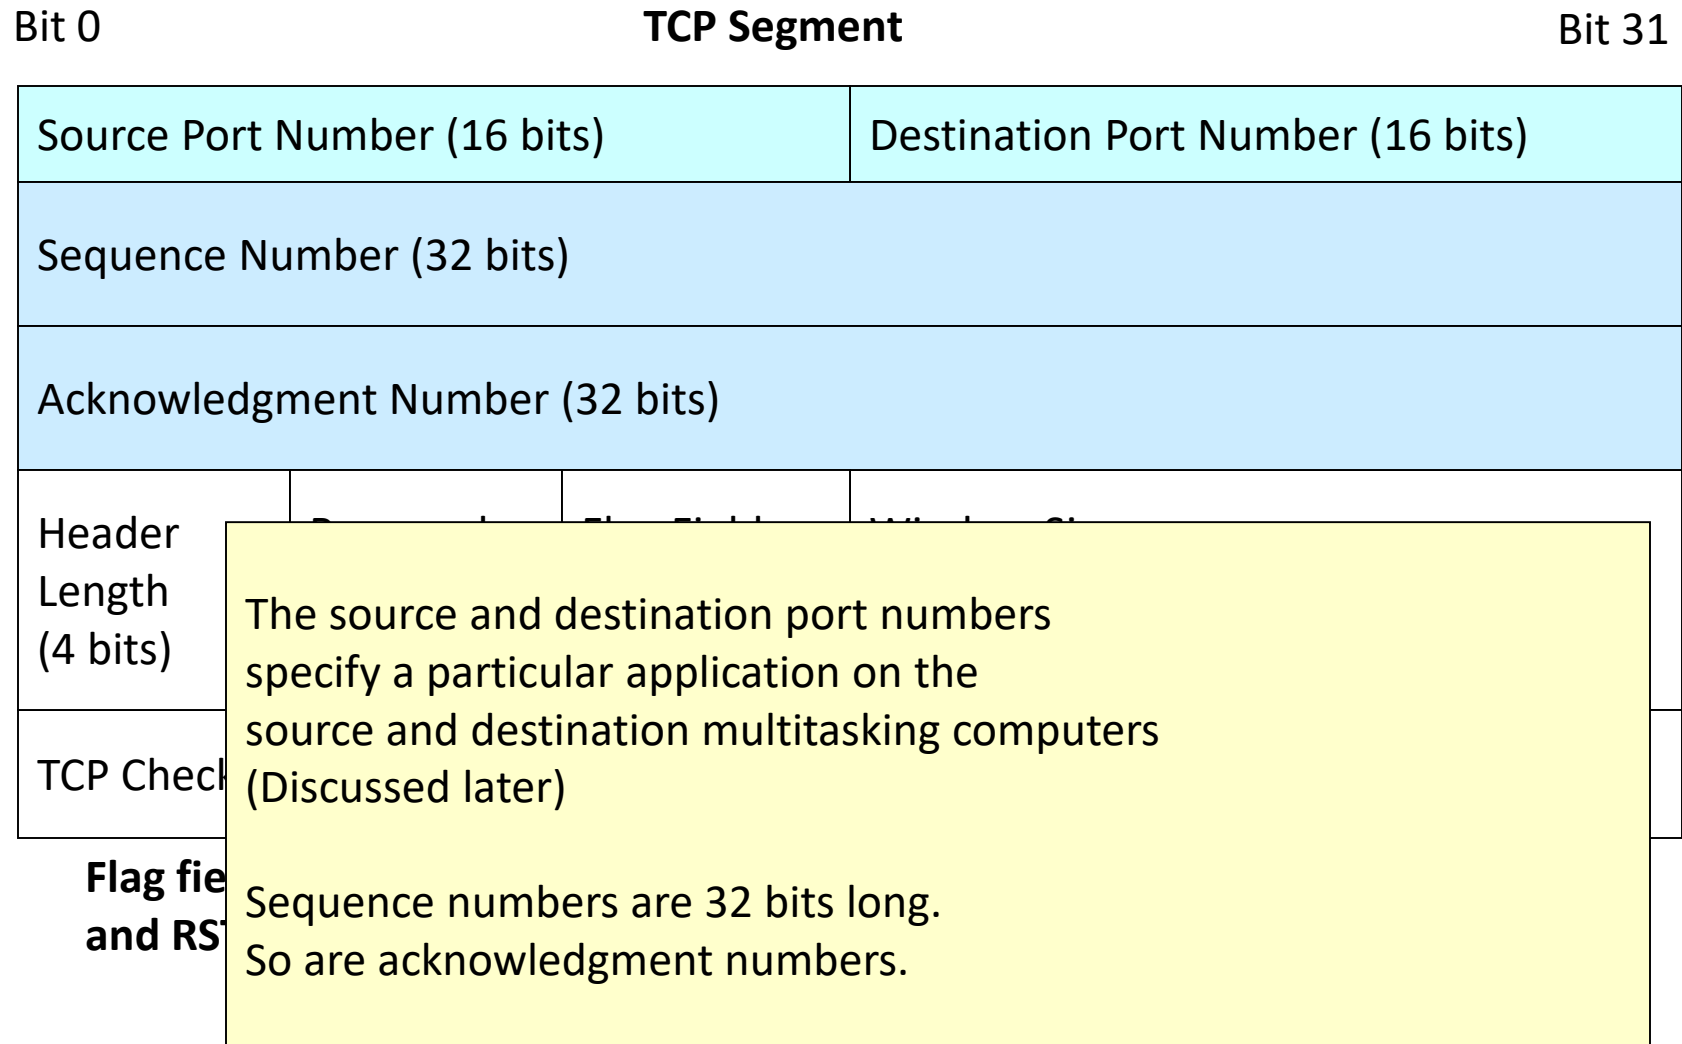
R Sends Packet to B

- Router R's learns the MAC address of host B
 - ARP request: broadcast request for 222.222.222.222
 - ARP response: B responds with 49-BD-D2-C7-56-2A
- Router R encapsulates the packet and sends to B



The Transmission Control Protocol (TCP)

TCP Segment and UDP Datagram



TCP Segment and UDP Datagram

Bit 0	<p>Flags are one-bit fields. If a flag's value is 1, it is "set". If a flag's value is 0, it is "not set." TCP has six flags</p>		
Source Port Number	<p>If the TCP Checksum field's value is correct, The receiving process sends back an acknowledgment.</p>		
Sequence Number			
Acknowledgment Number			
Header Length (4 bits)	Reserved (6 bits)	Flag Fields (6 bits)	Window Size (16 bits)
TCP Checksum (16 bits)			Urgent Pointer (16 bits)

Bit 0

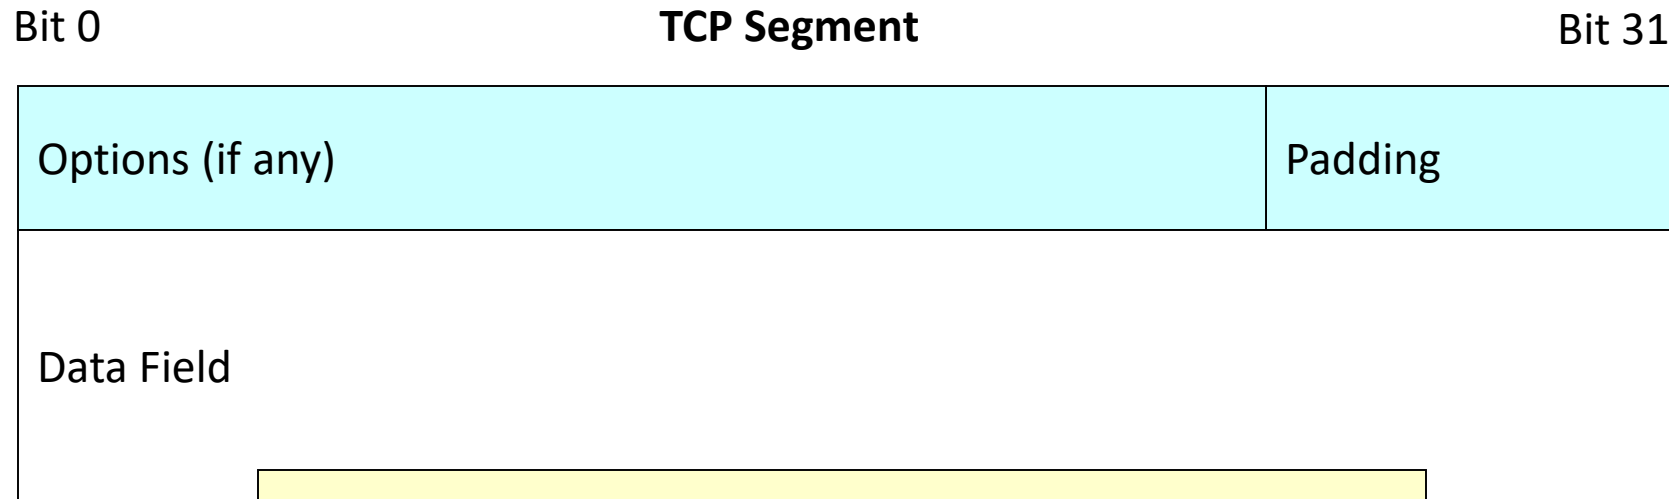
TCP Segment

Bit 31

Source Port Number (16 bits)		Destination Port Number (16 bits)	
Sequence Number (32 bits)			
Acknowledgment Number (32 bits)			
Header Length (4 bits)	Reserved (6 bits)	Flag Fields (6 bits)	Window Size (16 bits)
TCP Checksum (16 bits)		Urgent Pointer (16 bits)	

For flow control (to tell the other party to slow down),
The sender places a small value in the Window Size field.

TCP Segment and UDP Datagram

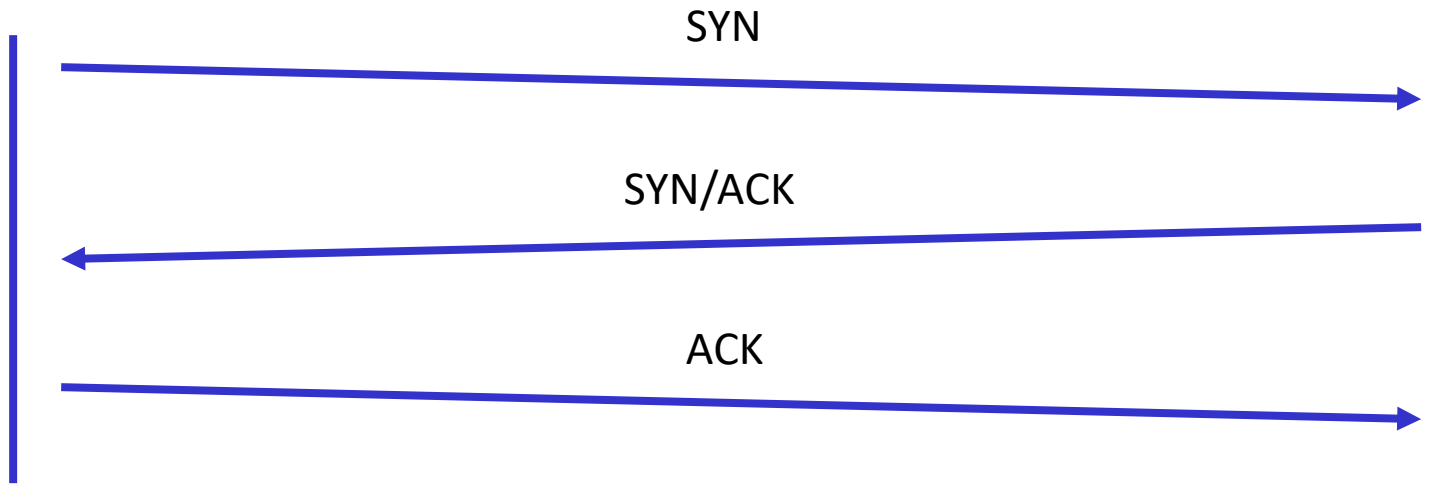


TCP segment headers can end with options.
Unlike IPv4 options,
TCP options are very common.

If an option does not end at a 32-bit boundary,
padding must be added.

TCP Session Openings and Closings


Normal Three-Way Opening



A SYN segment is a segment in which the SYN bit is set. One side sends a SYN segment requesting an opening. The other side sends a SYN/acknowledgment segment. Originating side acknowledges the SYN/ACK.

TCP Session Openings and Closings


Normal Four-Way Close



A FIN segment is a segment in which the FIN bit is set.
Like both sides saying “good bye” to end a conversation.

TCP Session Openings and Closings

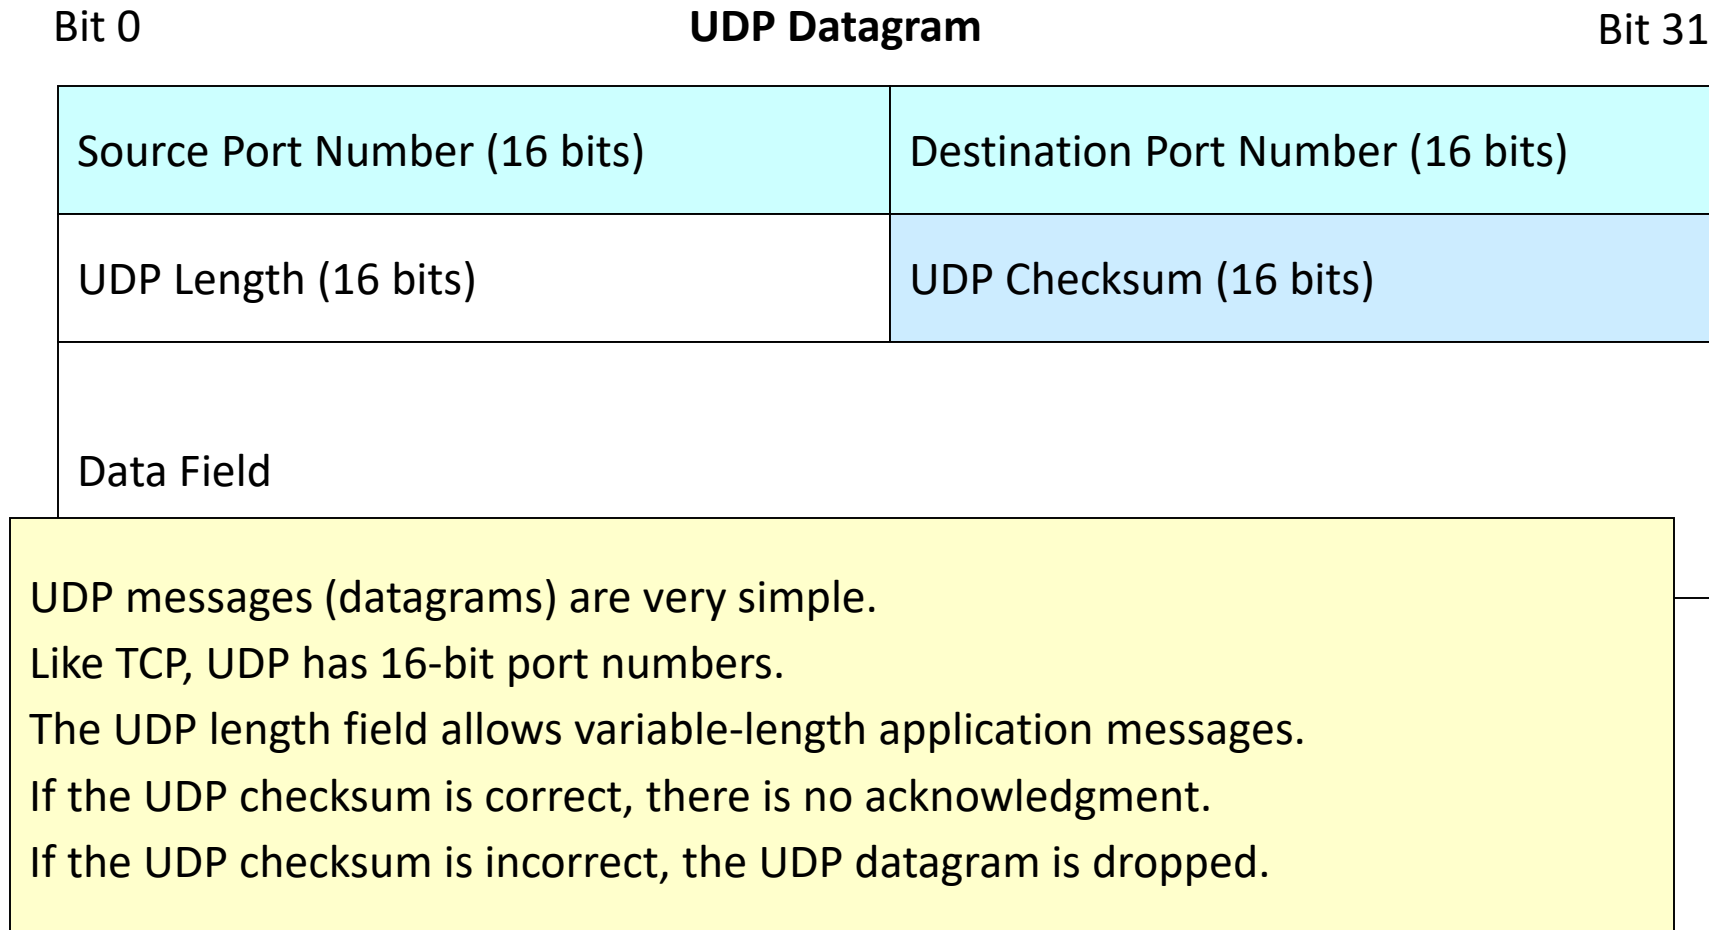
Abrupt Reset



An RST segment is a segment in which the RST bit is set.
A single RST segment breaks a connection.
Like hanging up during a phone call.
There is no acknowledgment.

The User Datagram Protocol (UDP)


TCP Segment and UDP Datagram



Port Numbers and Sockets in TCP and UDP


TCP and UDP Port Numbers

- Computers are multitasking devices
 - They run multiple applications at the same time
 - On a server, a port number designates a specific application




TCP and UDP Port Numbers

- Major Applications Have Well-Known Port Numbers between 0 to 1023, i.e.
 - HTTP is TCP Port 80
 - SMTP is TCP Port 25



TCP and UDP Port Numbers

- Clients Use Ephemeral Port Numbers
 - 1024 to 4999 for Windows Client PCs
 - A client has a separate port number for each connection to a program on a server



TCP and UDP Port Numbers

A socket is an IP address, a colon, and a port number.

1.33.17.3:80

123.30.17.120:25

128.171.17.13:2849

It represents a specific application (Port number) on a specific server (IP address)

Or a specific connection on a client.



Client PC
128.171.17.13
Port 2849




Webserver
1.33.17.13
Port 80




SMTP Server
123.30.17.120
Port 25

Use of TCP (and UDP) Port Numbers

Client
60.171.18.22




Source: 60.171.18.22:2707
Destination: 1.33.17.13:80



Webserver
1.33.17.13
Port 80


This shows sockets for a client packet sent to a webserver application on a webserver



SMTP Server
123.30.17.120
Port 25


\Use of TCP (and UDP) Port Numbers

Client
60.171.18.22




Source: 60.171.18.22:2707
Destination: 1.33.17.13:80

Source: 1.33.17.13:80
Destination: 60.171.18.22:2707



Webserver
1.33.17.13
Port 80

Sockets in
two-way
transmission




SMTP Server
123.30.17.120
Port 25

Use of TCP (and UDP) Port Numbers


Client

60.171.18.22




Source: 60.171.18.22:2707
Destination: 1.33.17.13:80

Source: 1.33.17.13:80
Destination: 60.171.18.22:2707



Webserver
1.33.17.13
Port 80


Source: 60.171.18.22:4400
Destination: 123.30.17.120:25



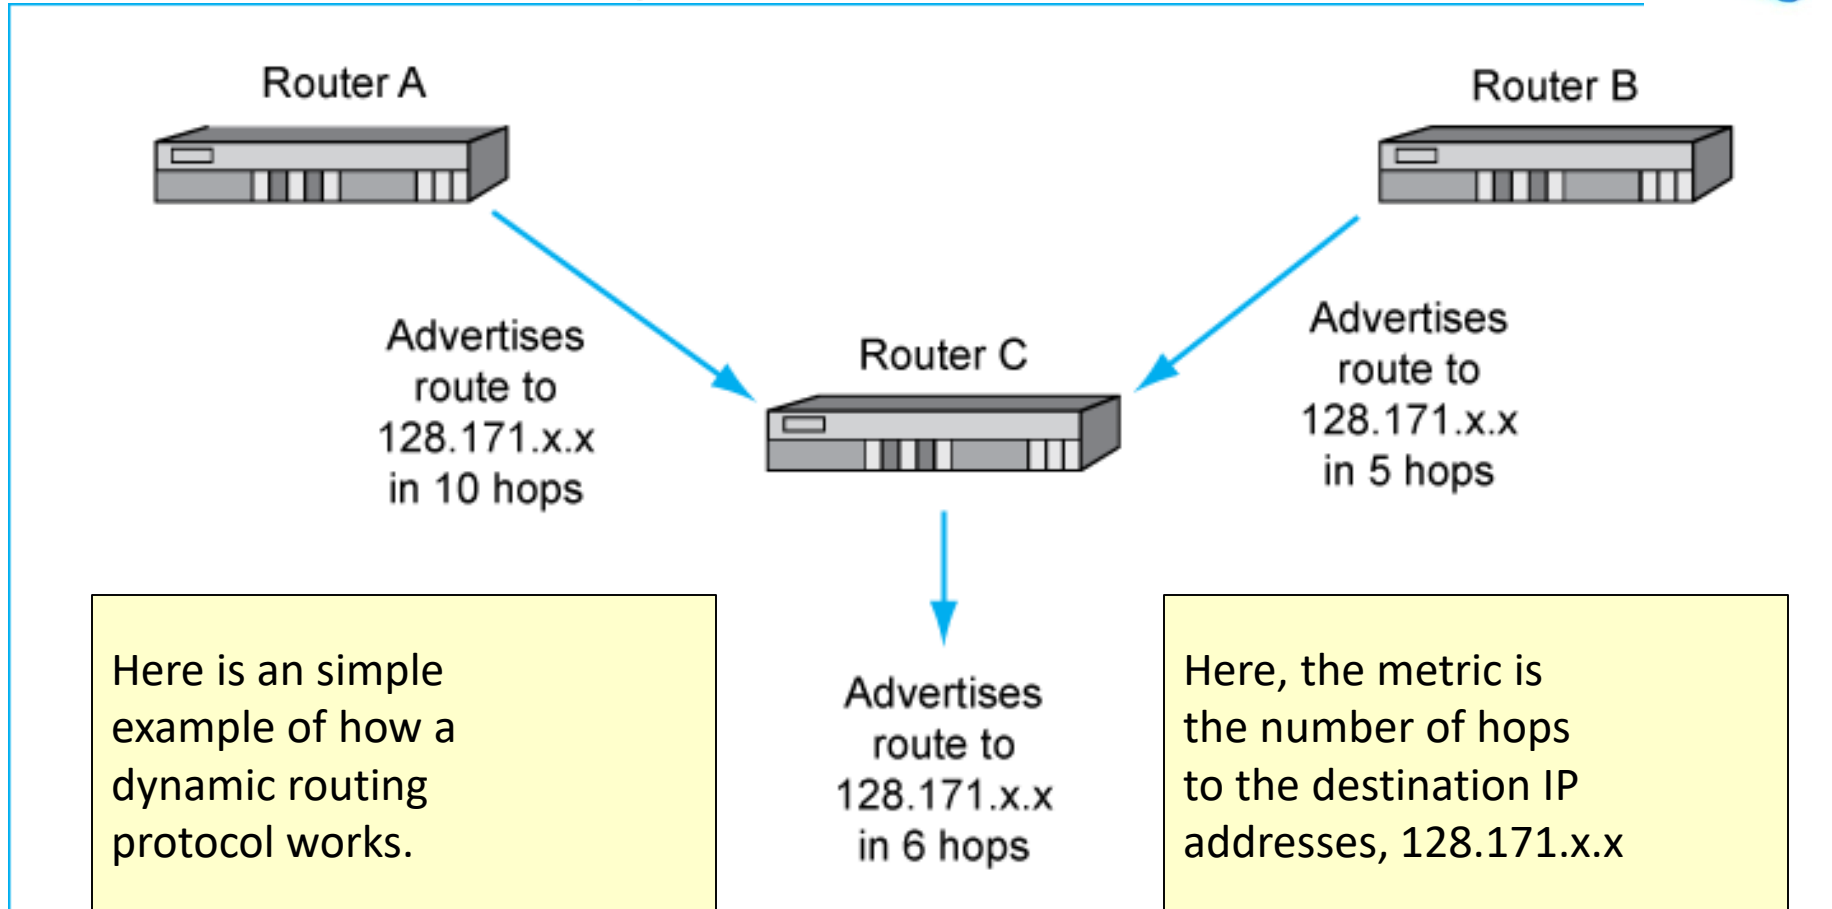
SMTP Server
123.30.17.120
Port 25

Clients use a different ephemeral port number for different connections

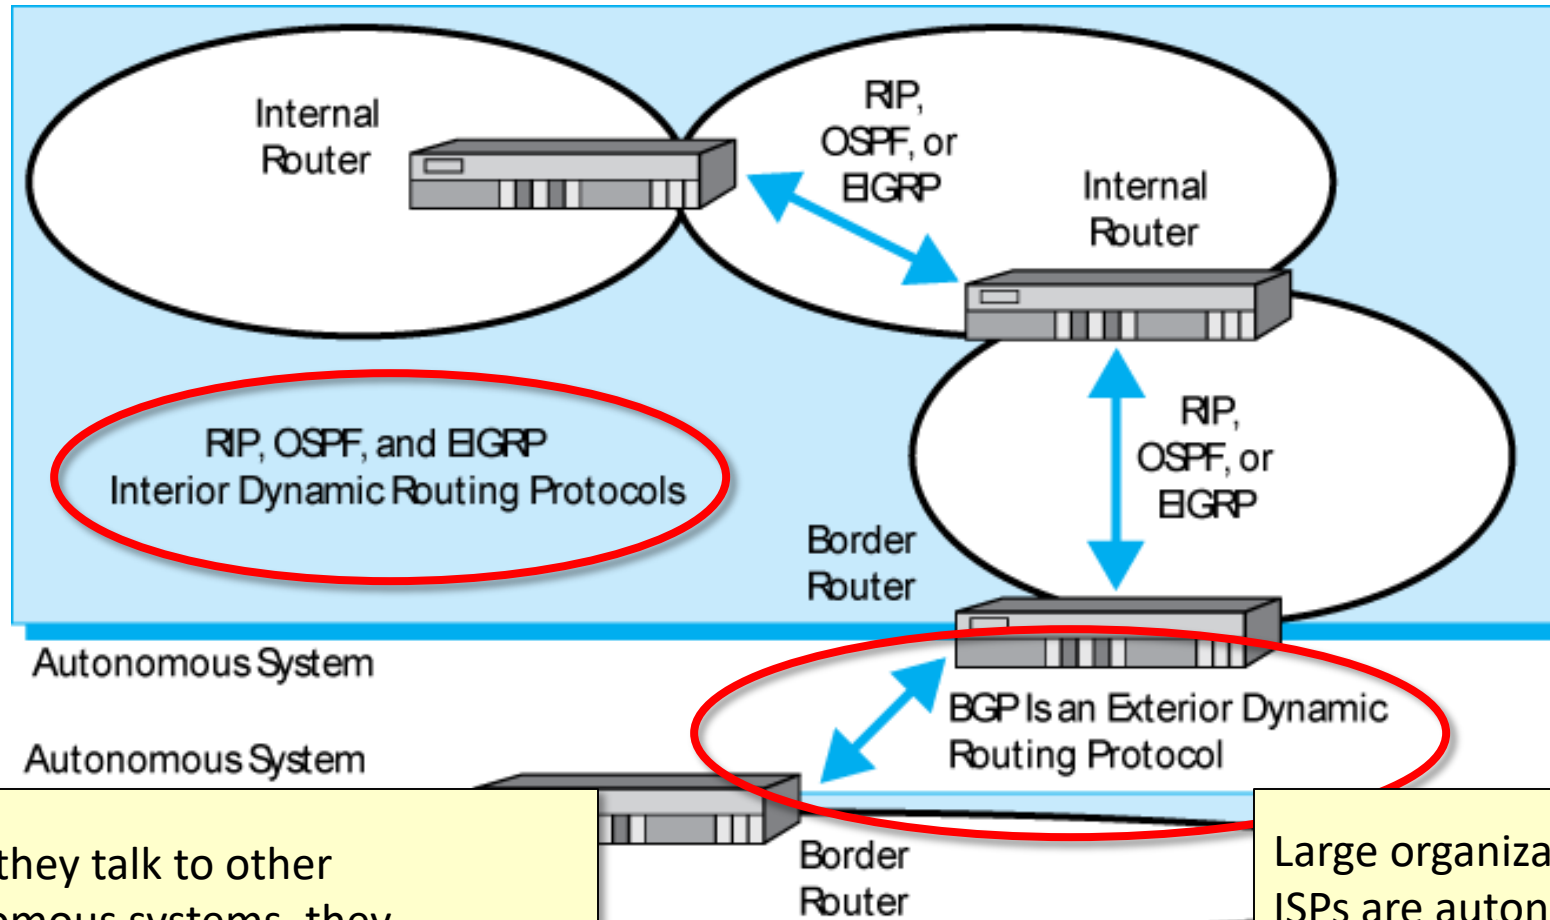
Dynamic Routing Protocols



Dynamic Routing Protocols



Dynamic Routing Protocols: Interior and Exterior



When they talk to other Autonomous systems, they Must negotiate which Exterior DRP they will use.

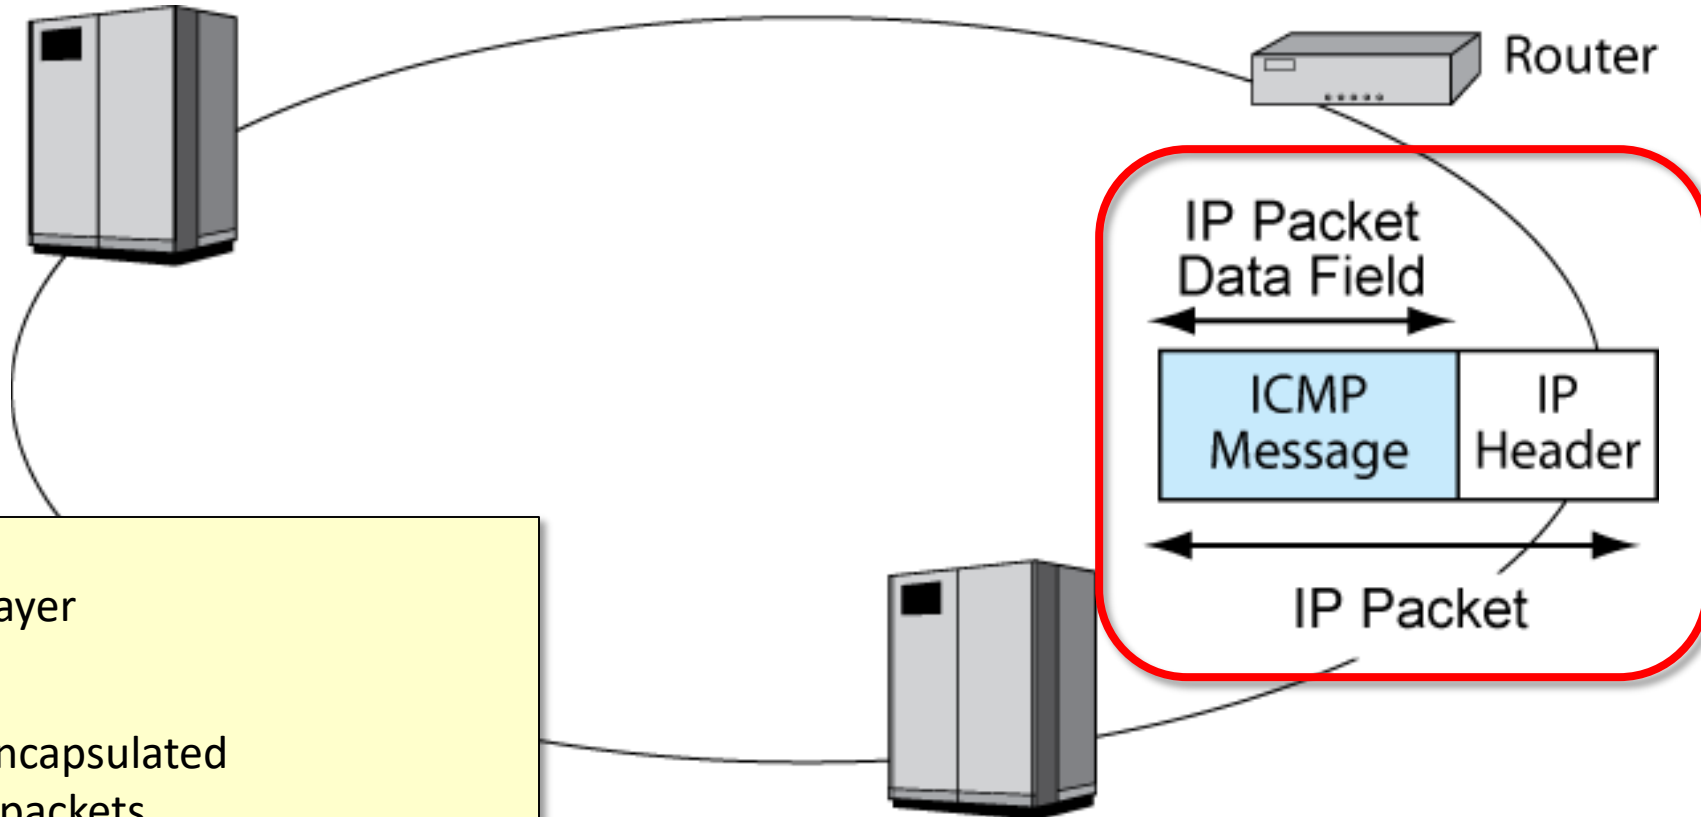
Large organizations and ISPs are autonomous systems. Autonomous systems can Select their interior Dynamic routing protocols.

Dynamic Routing Protocols

Dynamic Routing Protocol	Interior or Exterior Routing Protocol?	Remarks
RIP (Routing Information Protocol)	Interior	Only for small autonomous TCP/IP systems with low needs for security
OSPF (Open Shortest Path First)	Interior	For large autonomous systems that only use TCP/IP
EIGRP (Enhanced Interior Gateway Routing Protocol)	Interior	Proprietary Cisco Systems protocol. Not limited to TCP/IP routing. Also handles IPX/SPX, SNA, and so forth
BGP (Border Gateway Protocol)	Exterior	Organization cannot choose what exterior routing protocol it will use. TCP/IP protocol

The Internet Control Message Protocol (ICMP)

Internet Control Message Protocol (ICMP) for Supervisory Messages

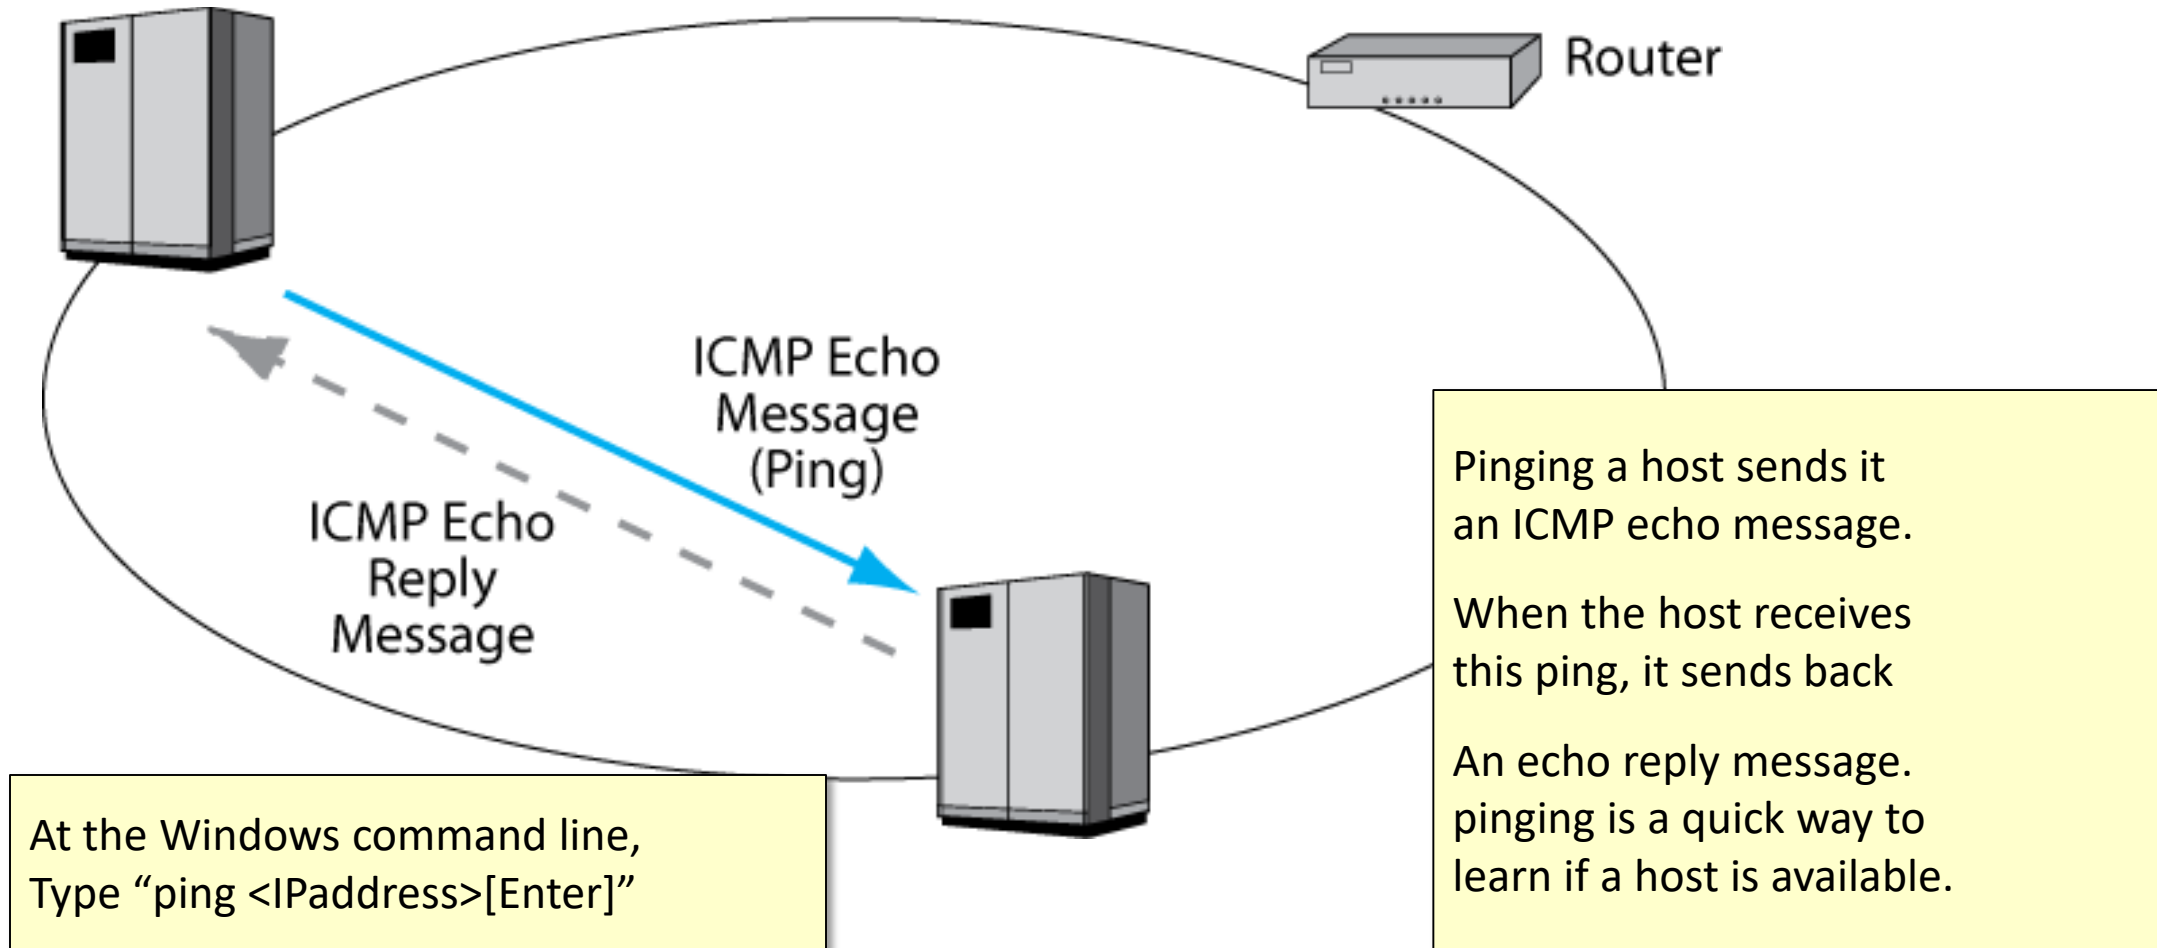


ICMP is the internet layer supervisory protocol.

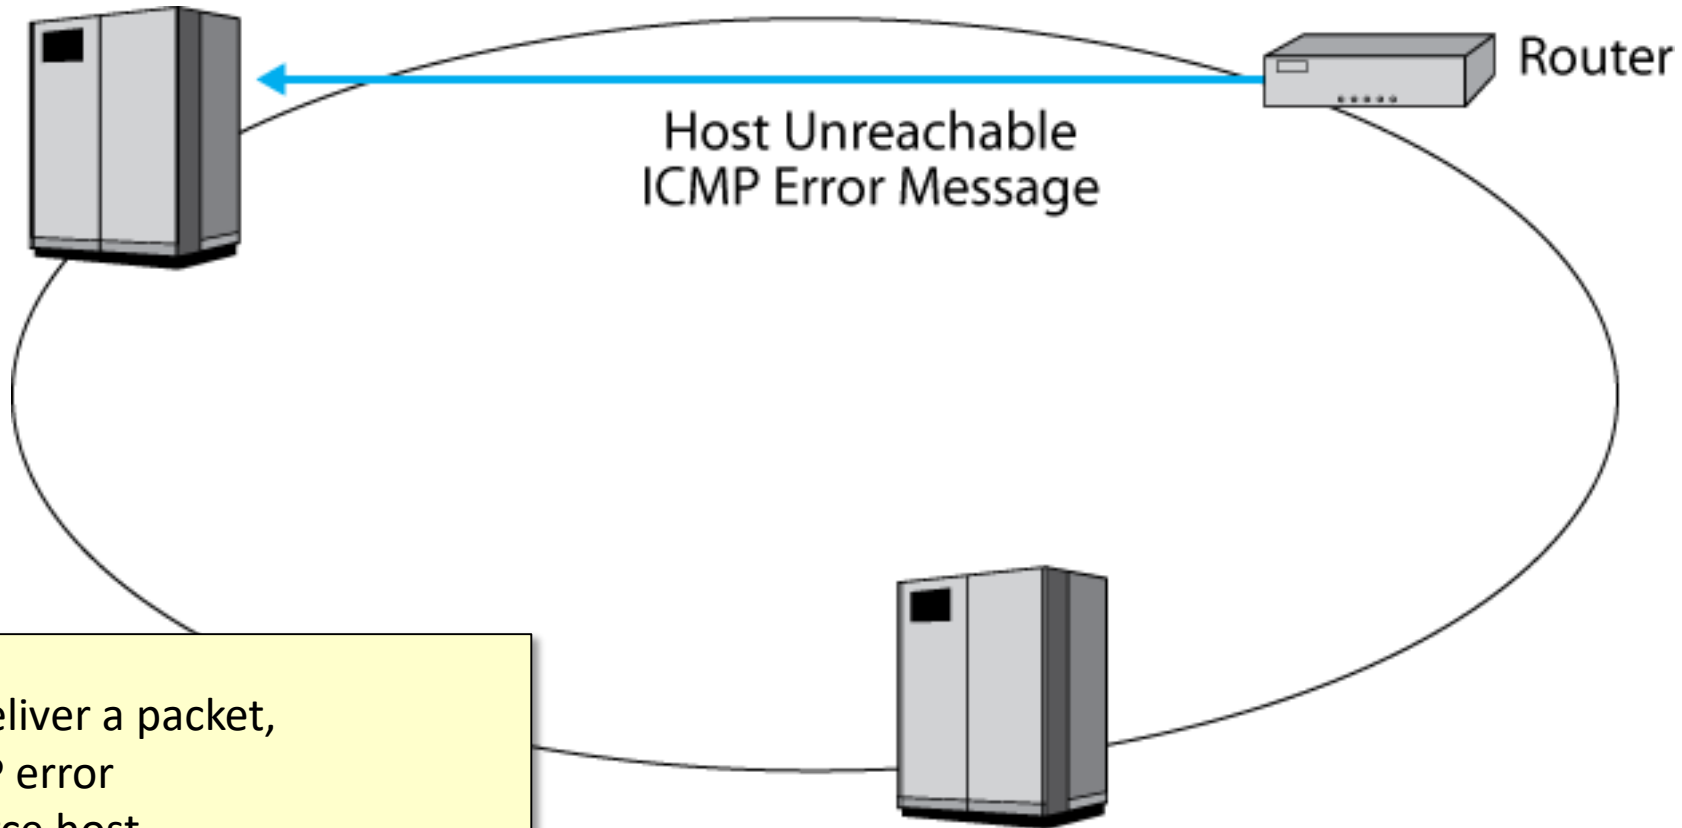
ICMP messages are encapsulated in the data field of IP packets.

These packets have no higher-layer contents

Internet Control Message Protocol (ICMP) for Supervisory Messages



8-16: Internet Control Message Protocol (ICMP) for Supervisory Messages



If a router cannot deliver a packet, it may send an ICMP error message to the source host.

There are several types of ICMP messages, for different types of error

Dynamic Host Configuration Protocol (DHCP)

From Chapter 1

Dynamic Host Configuration Protocol



Every Host Must Have a Unique IP address


Server hosts are given static IP addresses (unchanging)
Clients get dynamic (temporary) IP addresses that may be different each time they use an internet



Dynamic Host Configuration Protocol (DHCP)


Clients get these dynamic IP addresses from Dynamic Host Configuration Protocol (DHCP) servers

Dynamic Host Configuration Protocol (DHCP)



DHCP Request Message:
"My 48-bit Ethernet address is A3-4E-CD-59-28-7F".
Please give me a 32-bit IP address."

Dynamic Host Configuration Protocol (DHCP)



DHCP Response Message:

“Computer at A3-4E-CD-59-28-7F,
your 32-bit IP address is 1101000010111101010101100000010”.
(Usually other configuration parameters as well.)

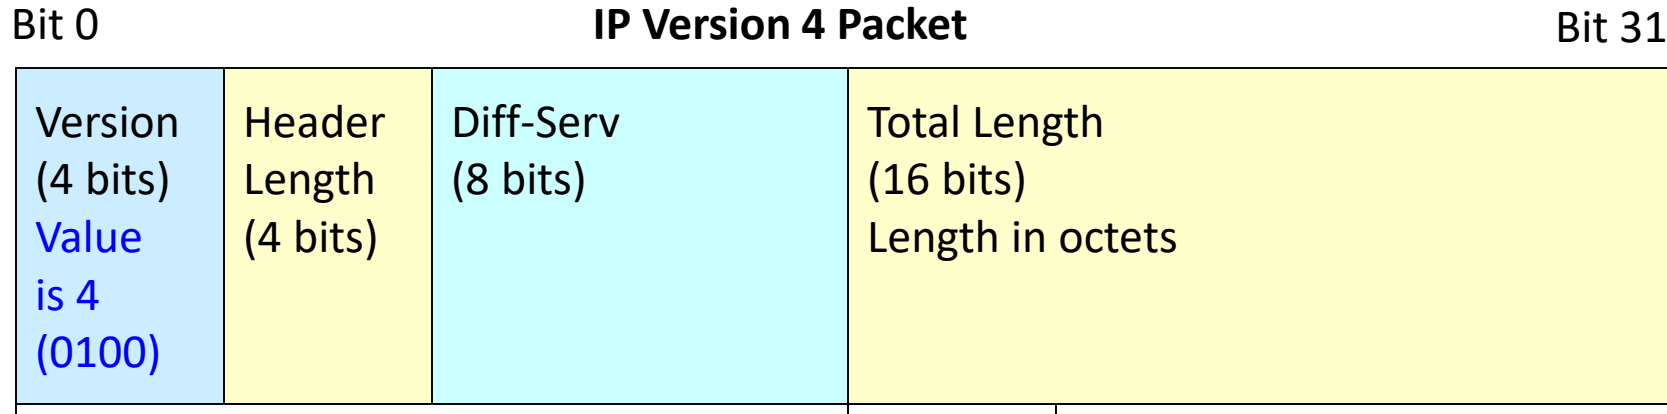
Why DHCP?

- If You Give PCs Static Information,
 - The cost of manual entry of configuration information (subnet mask, default router, DNS servers, etc.) is high
 - If something changes, such as the IP address of your DNS server, the cost of manually reconfiguring each PC is high
 - If something changes, your PCs may be inoperable until you make the manual changes
- With DHCP, users get hot fresh configuration data automatically

The Internet Protocol (IP)

Versions 4 and 6

IPv4 and IPv6 Packets



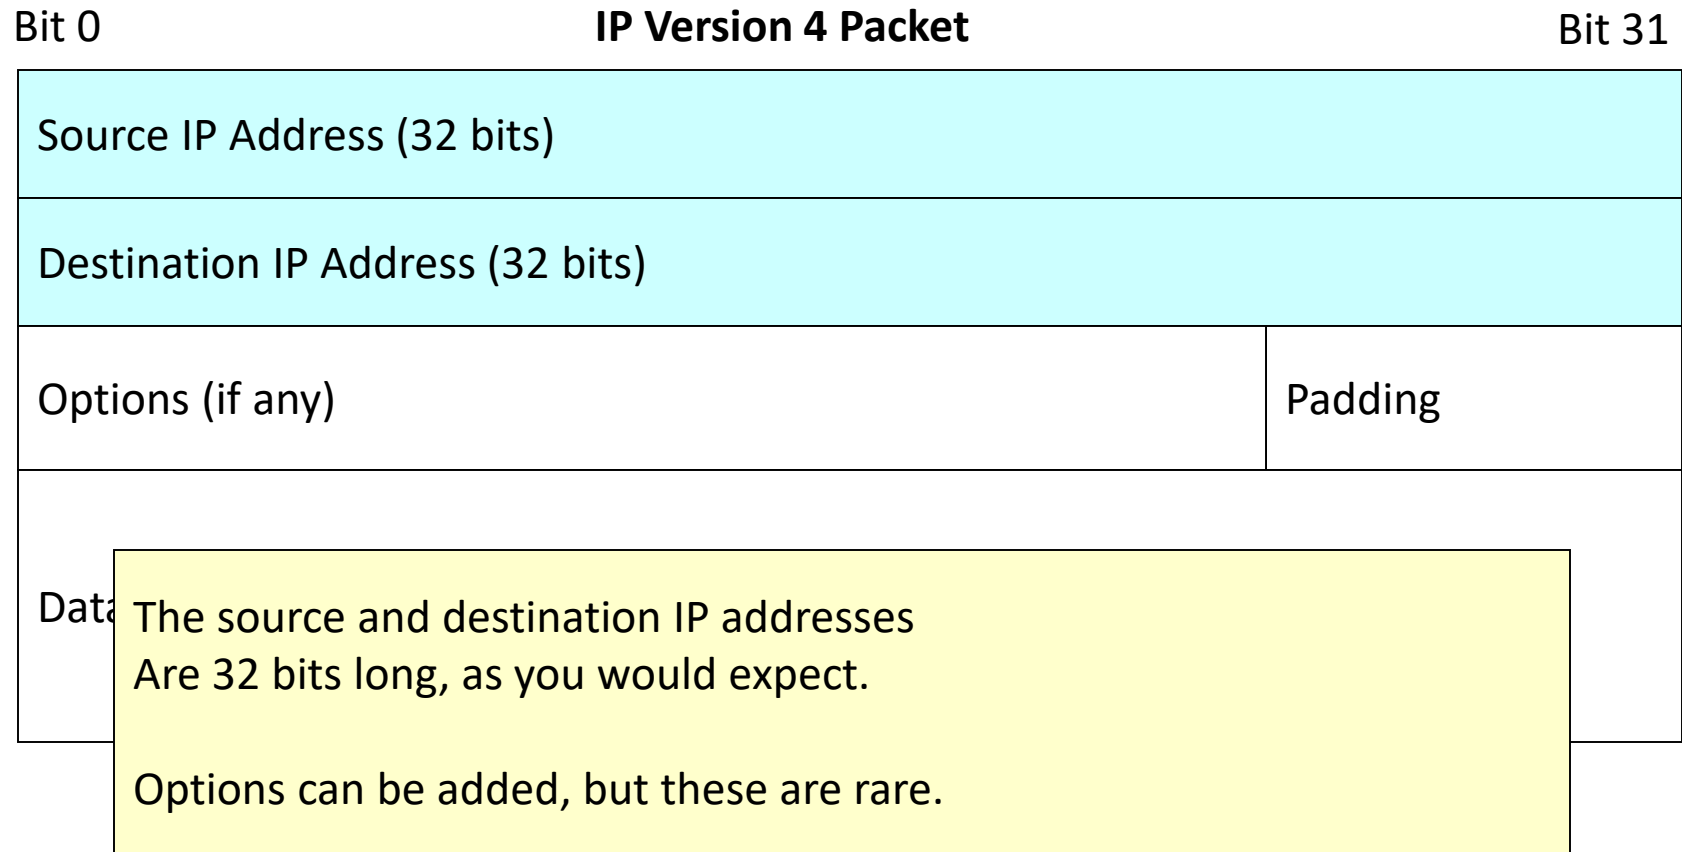
IPv4 is the dominant version of IP today.

The version number in its header is 4 (0100).

The header length and total length field tell the size of the packet.

The Diff-Serv field can be used for quality of service labeling.
(But MPLS is being used instead by most carriers)

IPv4 and IPv6 Packets



IPv4 and IPv6 Packets

Bit 0

Version (4 bits) Value is 6 (0110)	Diff-Serv (8 bits)
--	-----------------------

IP Version 6 is the emerging version of the Internet protocol.

Has 128 bit addresses for an almost unlimited number of IP addresses.

Needed because of rapid growth in Asia.

Also needed because of the exploding number of mobile devices

Payload Length (16 bits)	Next Header (8 bits) Name of next header	Hop Limit (8 bits)
Source IP Address (128 bits)		
Destination IP Address (128 bits)		
Next Header or Payload (Data Field)		

Layer 3 Switches


Layer 3 Switches

Traditionally, switches were fast and inexpensive while routers were slow and expensive

Using special-purpose hardware called application-specific integrated circuits (ASICs) allowed the creation of limited but fast and inexpensive routers

Marketing called these limited routers “Layer 3 switches” to indicate their speed, despite the fact that they are routers and operate at Layer 3, while switches operate at Layer 2

Layer 3 Switches and Routers in Site Internets



Layer 3 switches are routers.

Layer 3 switches are faster and cheaper to buy than traditional routers.

However, they are usually limited in functionality.


They also are expensive to manage.

They are typically used between workgroup switches and border routers.

Again, Layer 3 switches are true routers, Not switches.

However, they are faster and cheaper than traditional routers, at least to purchase.

Layer 3 Switches and Routers in Site Internets



Layer 3 switches are routers.

Layer 3 switches are faster and cheaper to buy than traditional routers.


However, they are usually limited in functionality.

They also are expensive to manage.

They are typically used between workgroup switches and border routers.

However, they have limited functionality that typically makes them unsuitable to being border routers to connect to different sites.

Layer 3 Switches and Routers in Site Internets



Layer 3 switches are routers.

Layer 3 switches are faster and cheaper to buy than traditional routers.

However, they are usually limited in functionality.


They also are expensive to manage.

As routers, however, they are expensive to manage (as we will see in Chapter 10).

After all, they really are routers, not switches.

They are typically used between workgroup switches and border routers.

Layer 3 Switches and Routers in Site Internets



Layer 3 switches are routers.

Layer 3 switches are faster and cheaper to buy than traditional routers.

However, they are usually limited in functionality.

They also are expensive to manage.

They are typically used between workgroup switches and border routers.

Too limited to be border routers and too expensive to manage to replace, Ethernet workgroup switches, L3 switches typically are used between the two.