REDDIG II – Computer Networking Training

IP Addressing and Subnetting

IP Addressing and Subnetting

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, 172.16.254.1)

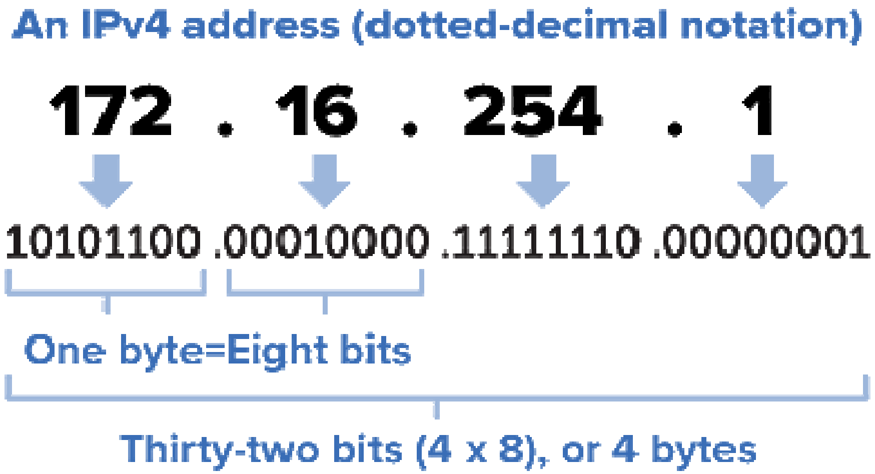
An IPv4 address (dotted-decimal notation)
172 . 16 . 254 . 1
10101100 .00010000 .11111110 .00000001
One byte=Eight bits
Thirty-two bits (4 x 8), or 4 bytes

IP Addressing and Subnetting

Binary and Decimal Conversion

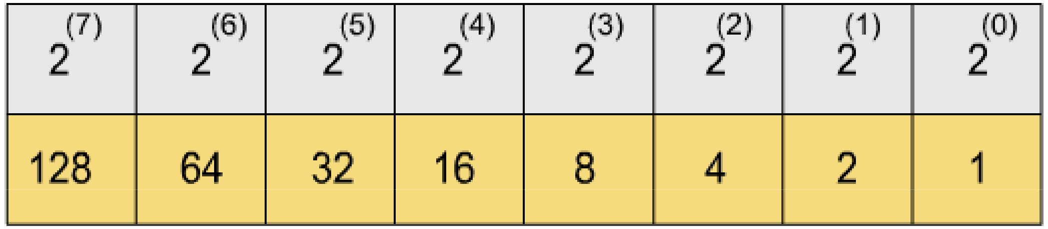

$2^{(7)}$	$2^{(6)}$	$2^{(5)}$	$2^{(4)}$	$2^{(3)}$	$2^{(2)}$	$2^{(1)}$	$2^{(0)}$
128	64	32	16	8	4	2	1

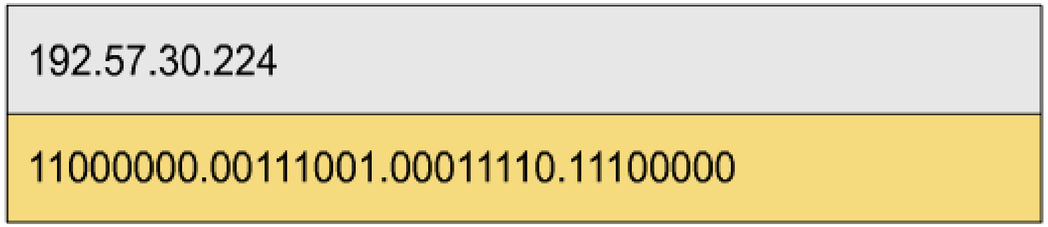

192.57.30.224
11000000.00111001.00011110.11100000

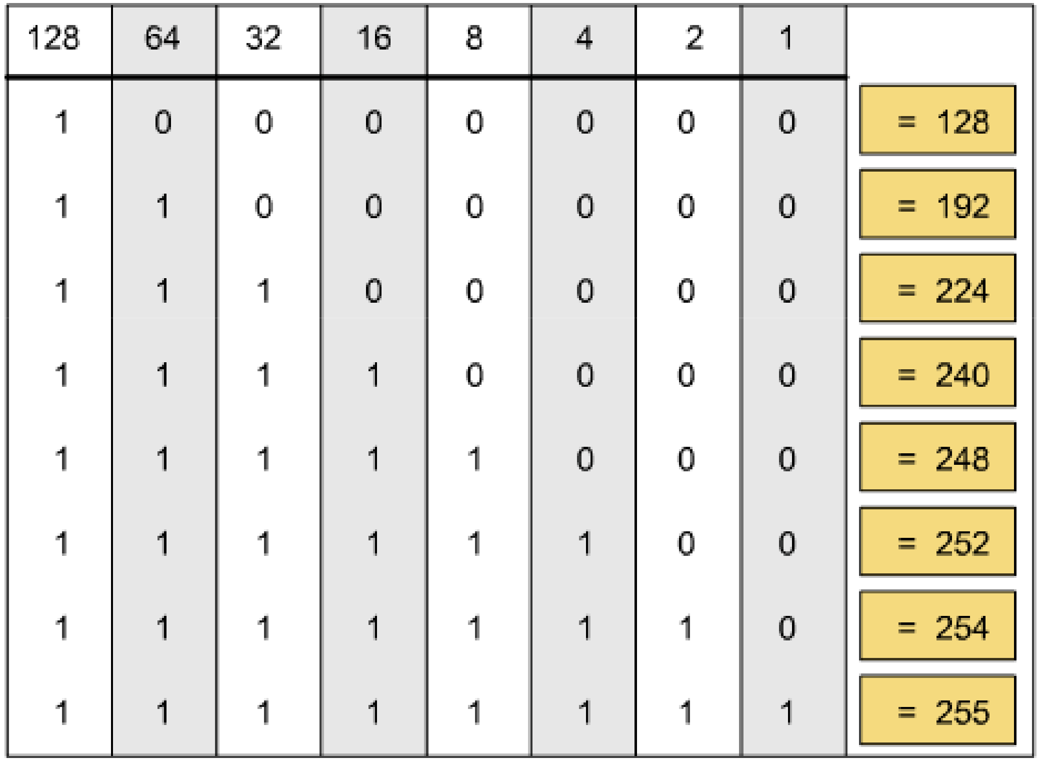

128	64	32	16	8	4	2	1	
1	0	0	0	0	0	0	0	= 128
1	1	0	0	0	0	0	0	= 192
1	1	1	0	0	0	0	0	= 224
1	1	1	1	0	0	0	0	= 240
1	1	1	1	1	0	0	0	= 248
1	1	1	1	1	1	0	0	= 252
1	1	1	1	1	1	1	0	= 254
1	1	1	1	1	1	1	1	= 255

IP Addressing and Subnetting

IP Address Classes

- IP classes are used to assist in assigning IP addresses to networks with different size requirements.
- Classful addressing:

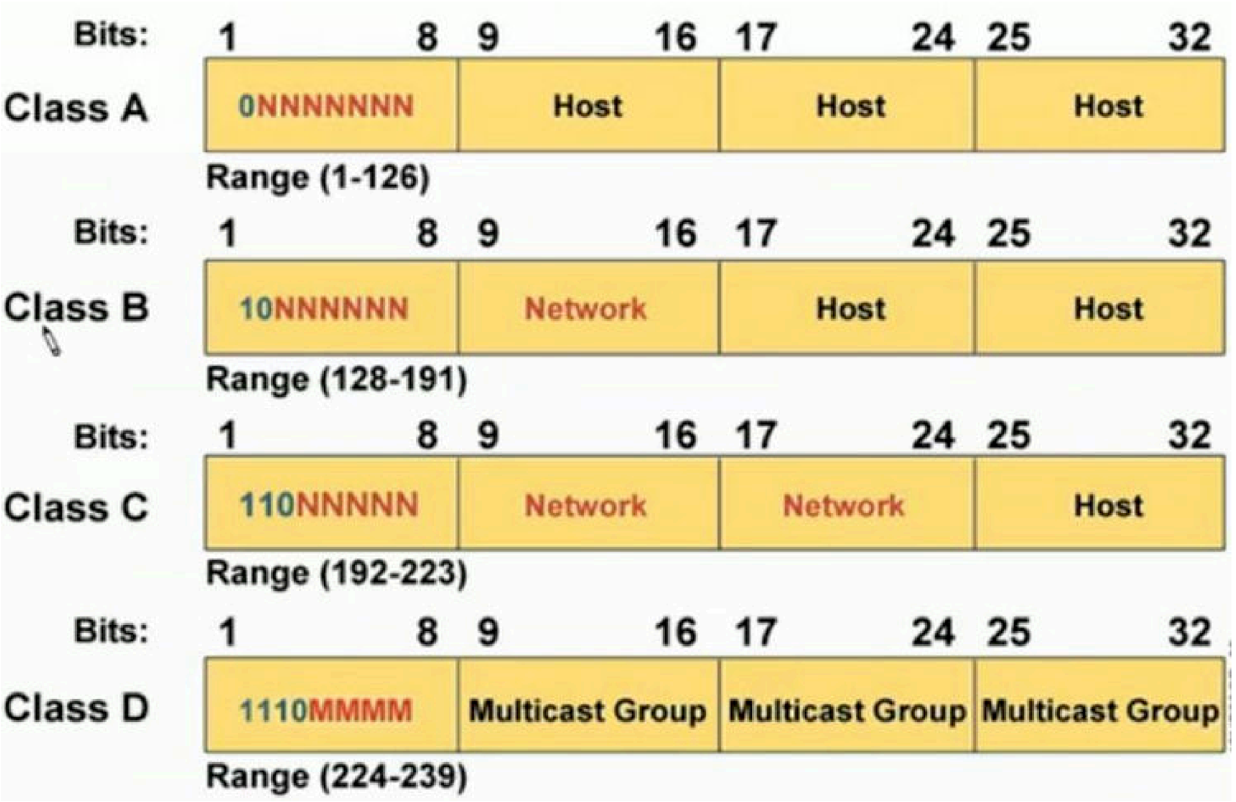

Bits:	1 8	9 16	17 24	25 32
Class A	0NNNNNNN	Host	Host	Host
	Range (1-126)			
Class B	10NNNNNN	Network	Host	Host
	Range (128-191)			
Class C	110NNNNN	Network	Network	Host
	Range (192-223)			
Class D	1110MMMM	Multicast Group	Multicast Group	Multicast Group
	Range (224-239)			

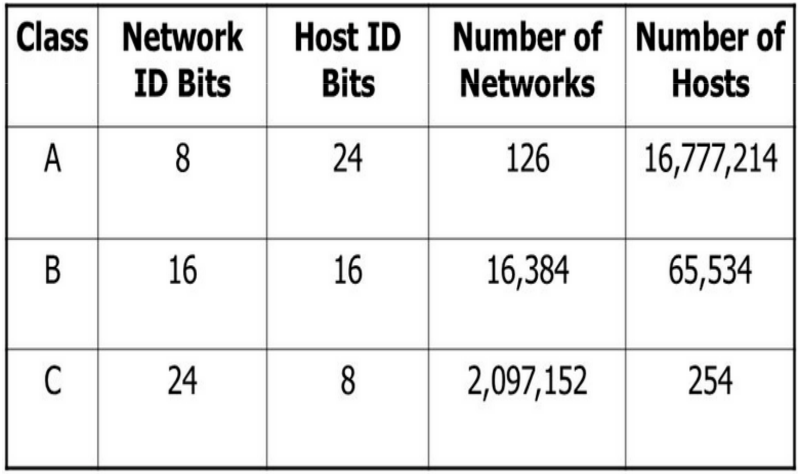

Class	Network ID Bits	Host ID Bits	Number of Networks	Number of Hosts
A	8	24	126	16,777,214
B	16	16	16,384	65,534
C	24	8	2,097,152	254

IP Addressing and Subnetting

Private Address Range

- Private IP addresses provide an entirely separate set of addresses that still allow access on a network but without taking up a public IP address space.
- Private addresses are not allowed to be routed out to the Internet, so devices using private addresses cannot communicate directly with devices on the Internet.

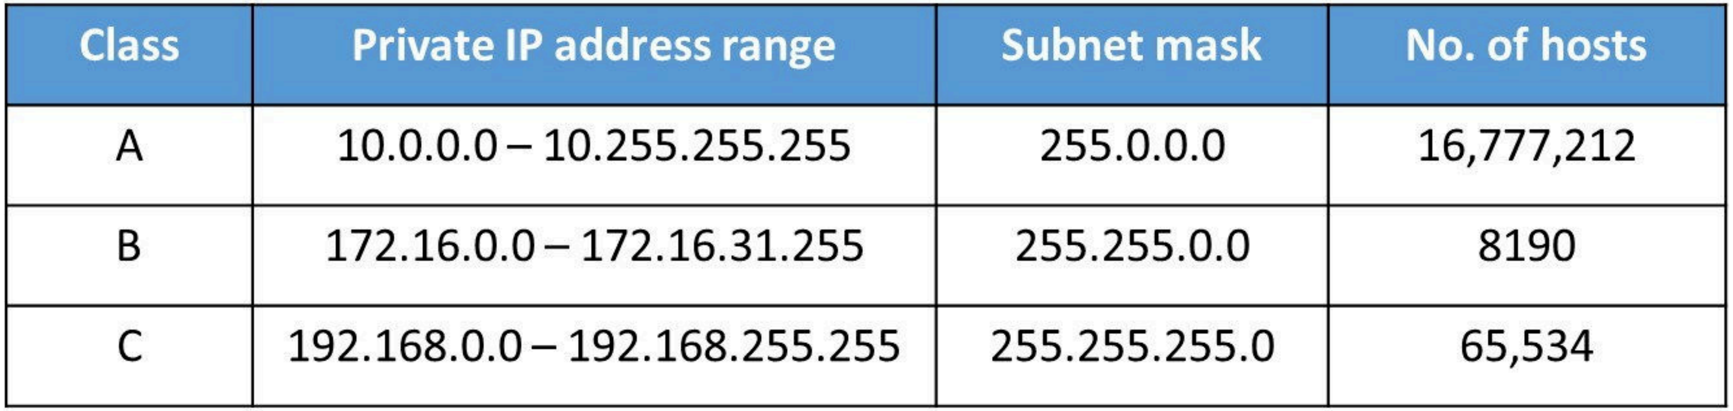

Class	Private IP address range	Subnet mask	No. of hosts
A	10.0.0.0 – 10.255.255.255	255.0.0.0	16,777,212
B	172.16.0.0 – 172.16.31.255	255.255.0.0	8190
C	192.168.0.0 – 192.168.255.255	255.255.255.0	65,534

Private IP Addresses

IP Addressing and Subnetting

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
- Once you have the address and the mask represented in binary, then identification of the network and host ID is easier.

•Example:

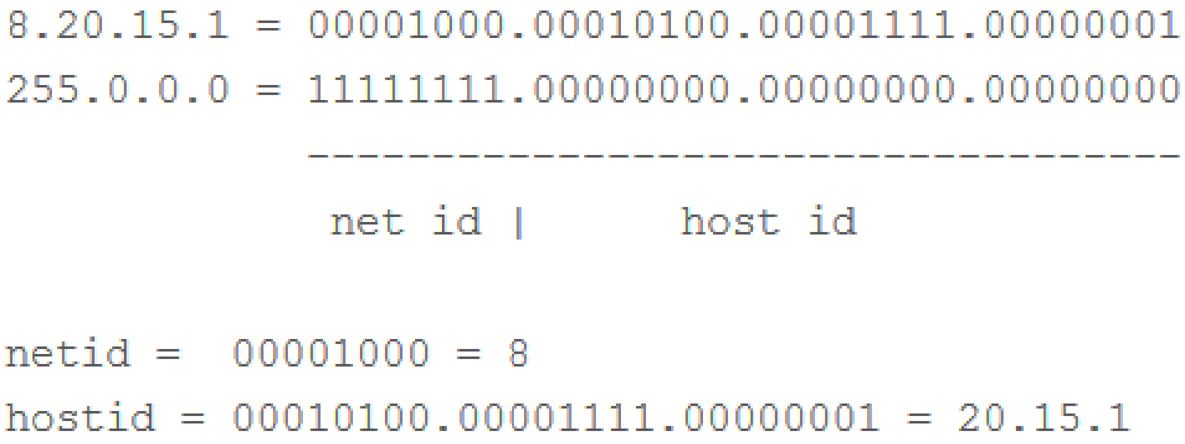

```
8.20.15.1 = 00001000.00010100.00001111.00000001
255.0.0.0 = 11111111.00000000.00000000.00000000
            -----------------------------------
             net id |      host id


netid =   00001000 = 8
hostid = 00010100.00001111.00000001 = 20.15.1
```

IP Addressing and Subnetting

Subnet and subnetting

- A subnet is a small network inside a larger network. It is a logical grouping of connected network devices that tend to be located in close physical proximity to each other on a local area network—a LAN.

- Subnetting allows you to create multiple logical networks that exist within a single Class A, B, or C network

- To create a subnet address, a network administrator <u>borrows</u> bits from the original host portion and designates them as the subnet field.

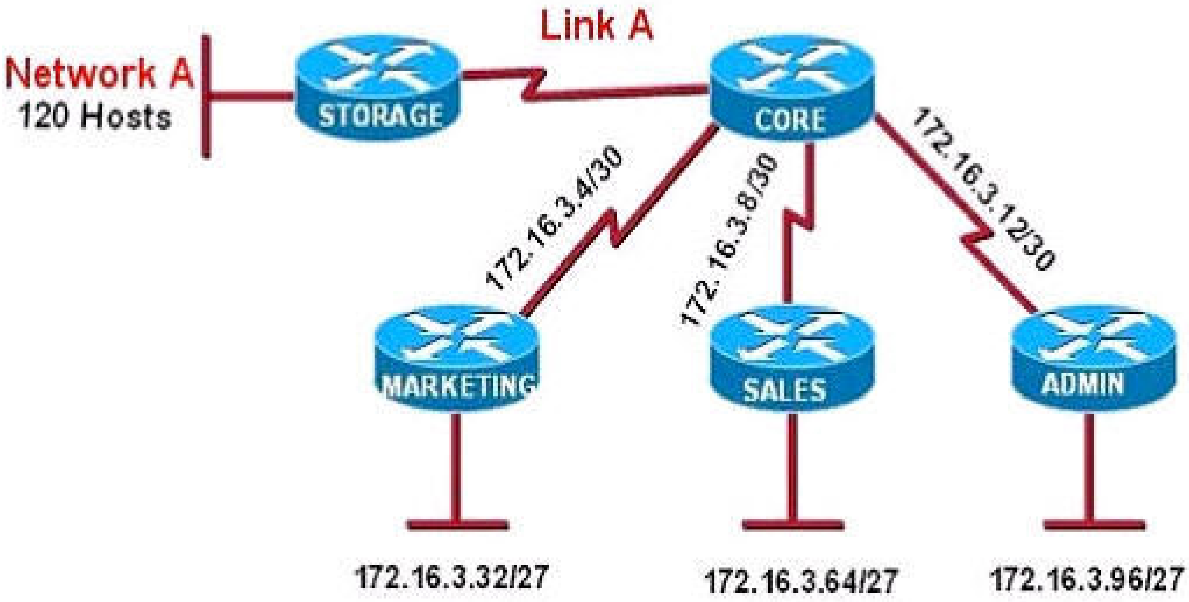

IP Addressing and Subnetting

Subnet and subnetting

- Subnetting does not give you more hosts, it only allows you to divide your larger network into smaller networks. Limit layer 2 and layer 3 broadcasts to their subnet.
- Example:

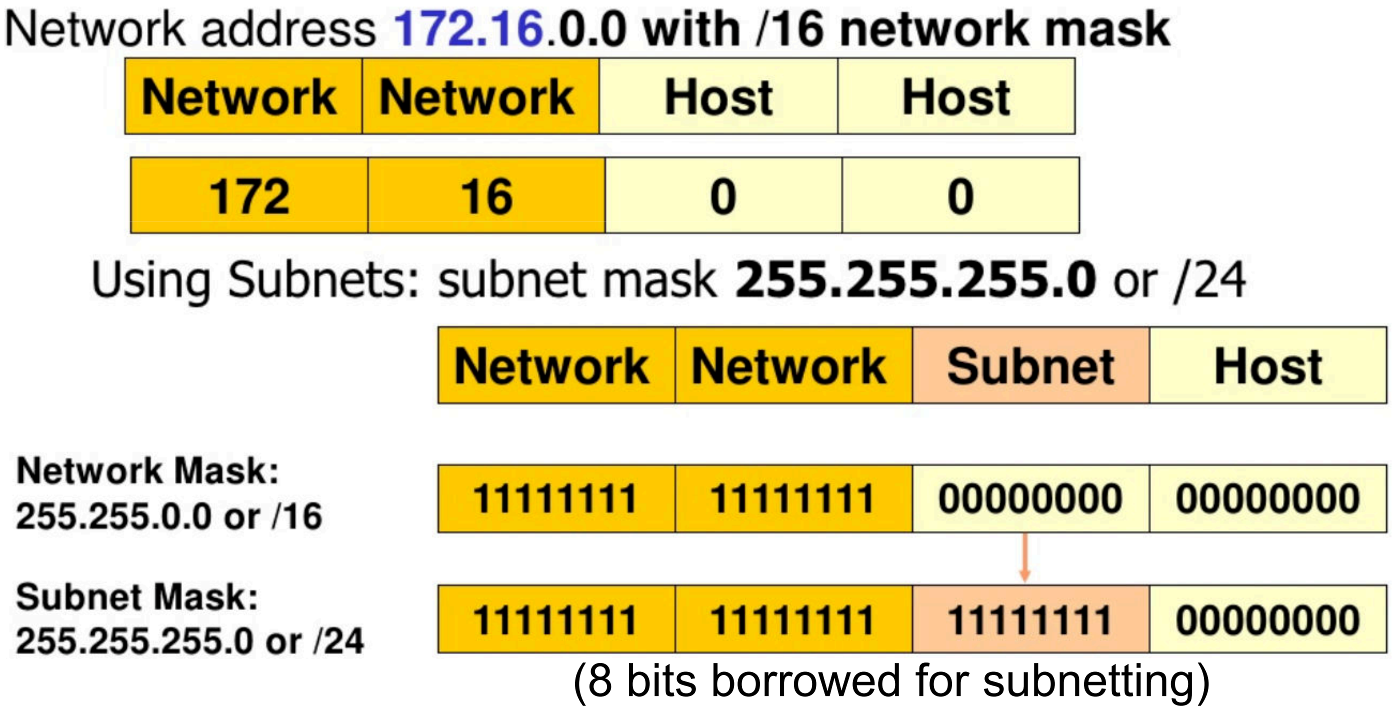

Network address **172.16.0.0 with /16 network mask**

Network	Network	Host	Host
172	16	0	0

Using Subnets: subnet mask **255.255.255.0** or /24

	Network	Network	Subnet	Host
Network Mask: **255.255.0.0 or /16**	11111111	11111111	00000000	00000000
Subnet Mask: **255.255.255.0 or /24**	11111111	11111111	11111111	00000000

(8 bits borrowed for subnetting)

- Applying a mask which is larger than the default subnet mask, will divide your network into subnets.

IP Addressing and Subnetting

Variable Length Subnet Masking (VLSM)

- The process of sub-netting a subnet to fit your needs.

- Example:

- Subnet 10.1.0.0/16, 8 more bits are borrowed again, to create 256 subnets with a /24 mask.
  - -Mask allows for 254 host addresses per subnet
  - -Subnets range from: 10.1.0.0 / 24 to 10.1.255.0 / 24

- The same process can be done for other subnets (10.2.0.0/16, 10.3.0.0/16 …)

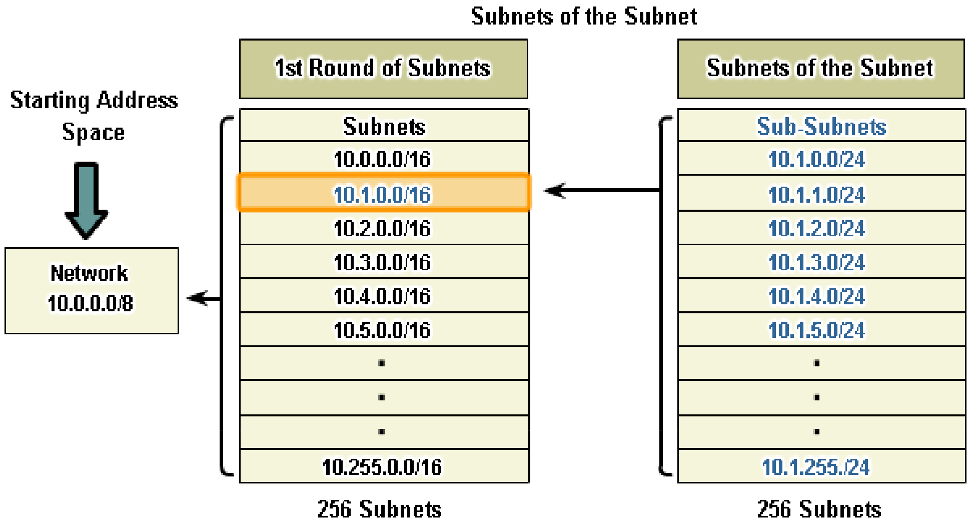

Subnets of the Subnet

Starting Address Space

Network 10.0.0.0/8

1st Round of Subnets
Subnets
10.0.0.0/16
10.1.0.0/16
10.2.0.0/16
10.3.0.0/16
10.4.0.0/16
10.5.0.0/16
.
.
.
10.255.0.0/16

256 Subnets

Subnets of the Subnet
Sub-Subnets
10.1.0.0/24
10.1.1.0/24
10.1.2.0/24
10.1.3.0/24
10.1.4.0/24
10.1.5.0/24
.
.
.
10.1.255./24

256 Subnets

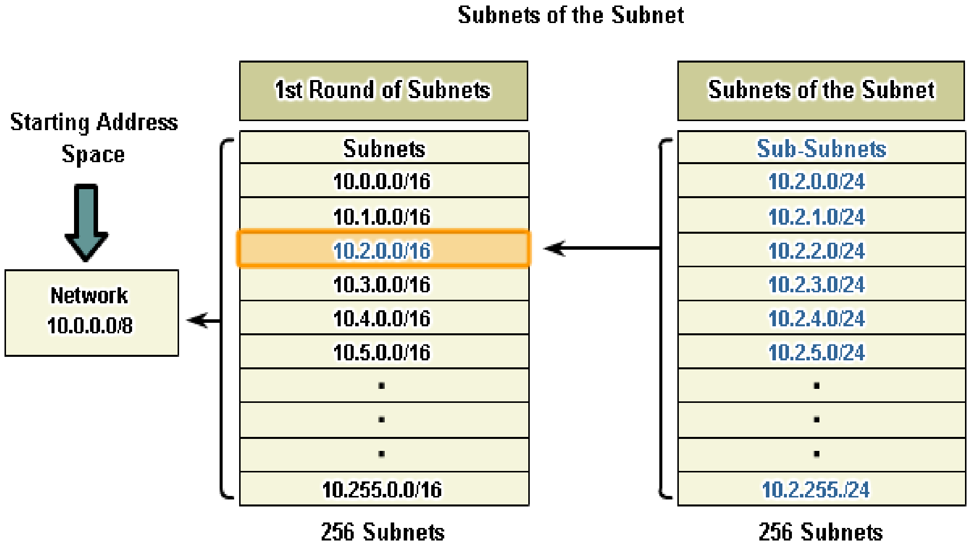

Subnets of the Subnet

Starting Address Space

Network 10.0.0.0/8

1st Round of Subnets
Subnets
10.0.0.0/16
10.1.0.0/16
10.2.0.0/16
10.3.0.0/16
10.4.0.0/16
10.5.0.0/16
.
.
.
10.255.0.0/16

256 Subnets

Subnets of the Subnet
Sub-Subnets
10.2.0.0/24
10.2.1.0/24
10.2.2.0/24
10.2.3.0/24
10.2.4.0/24
10.2.5.0/24
.
.
.
10.2.255./24

256 Subnets

IP Addressing and Subnetting

Variable Length Subnet Masking (VLSM)

- **Examples:**

- **Subnet 10.3.0.0/16, 12 more bits are borrowed again, to create 4,096 subnets with a /28 (255.255.255.240) mask.**
  - **–Mask allows for 14 host addresses per subnet**
  - **–Subnet bits: 4**
  - **–Subnets range from: 10.3.0.0 / 28 to 10.3.255.240 / 28**

- **Subnet 10.4.0.0/16, 4 more bits are borrowed again, to create 16 subnets with a /20 (255.255.240.0) mask.**
  - **–Mask allows for 4,094 host addresses per subnet**
  - **–Subnet bits: 12**
  - **–Subnets range from: 10.4.0.0 / 20 to 10.4.240.0 / 20**

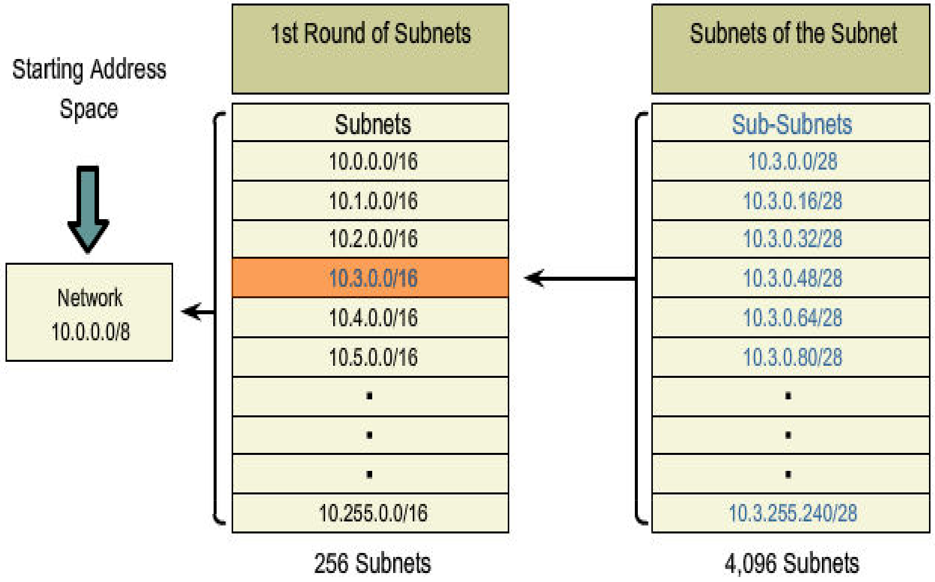

Starting Address Space

Network 10.0.0.0/8

1st Round of Subnets

Subnets
10.0.0.0/16
10.1.0.0/16
10.2.0.0/16
10.3.0.0/16
10.4.0.0/16
10.5.0.0/16
.
.
.
10.255.0.0/16

256 Subnets

Subnets of the Subnet

Sub-Subnets
10.3.0.0/28
10.3.0.16/28
10.3.0.32/28
10.3.0.48/28
10.3.0.64/28
10.3.0.80/28
.
.
.
10.3.255.240/28

4,096 Subnets

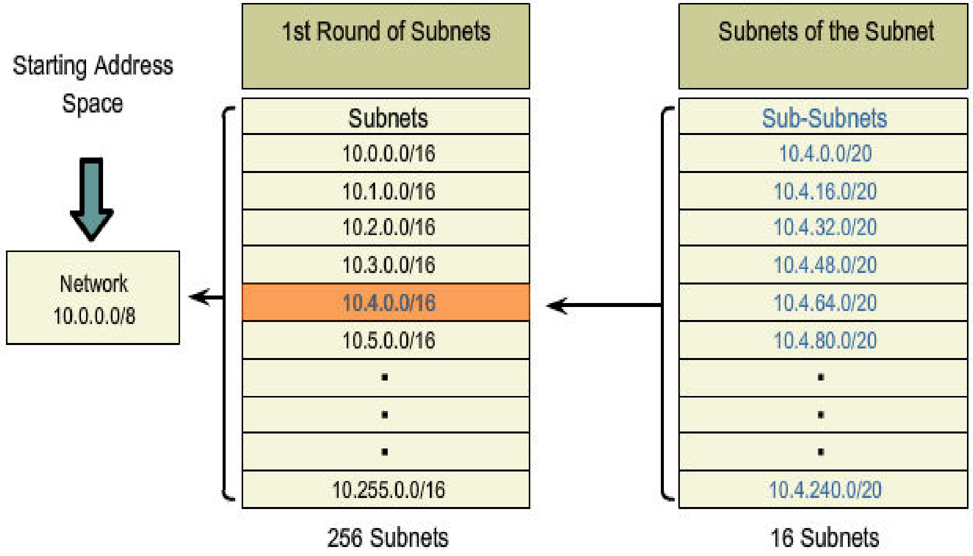

Starting Address Space

Network 10.0.0.0/8

1st Round of Subnets

Subnets
10.0.0.0/16
10.1.0.0/16
10.2.0.0/16
10.3.0.0/16
10.4.0.0/16
10.5.0.0/16
.
.
.
10.255.0.0/16

256 Subnets

Subnets of the Subnet

Sub-Subnets
10.4.0.0/20
10.4.16.0/20
10.4.32.0/20
10.4.48.0/20
10.4.64.0/20
10.4.80.0/20
.
.
.
10.4.240.0/20

16 Subnets

IP Addressing and Subnetting

Classless Inter-Domain Routing – CIDR

- The name is unfortunate because CIDR only specifies addressing and forwarding
- Designers wanted to make it easy for a human to specify a mask

- CIDR Notation:
  The number after the '/' is the number of bits that are 1s in the subnet mask

- CIDR allowed for more efficient use of IPv4 address space and prefix aggregation, known as route summarization or supernetting

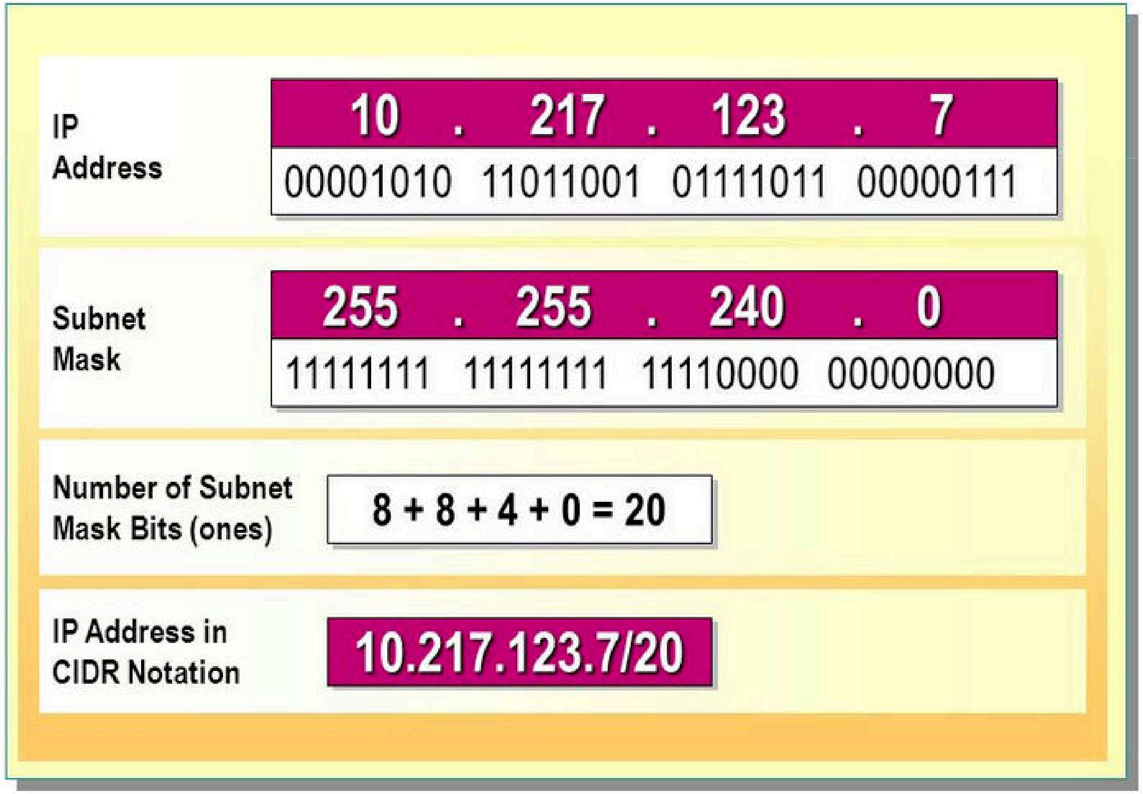

IP Address	10 .	217 .	123 .	7
	00001010	11011001	01111011	00000111
Subnet Mask	255 .	255 .	240 .	0
	11111111	11111111	11110000	00000000
Number of Subnet Mask Bits (ones)	8 + 8 + 4 + 0 = 20			
IP Address in CIDR Notation	10.217.123.7/20			

IP Addressing and Subnetting

CIDR - Supernetting

- Supernetting is the opposite of subnetting
- In subnetting you borrow bits from the host part
- Supernetting is done by borrowing bits from the network side.
- And combine a group of networks into one large supernetwork.

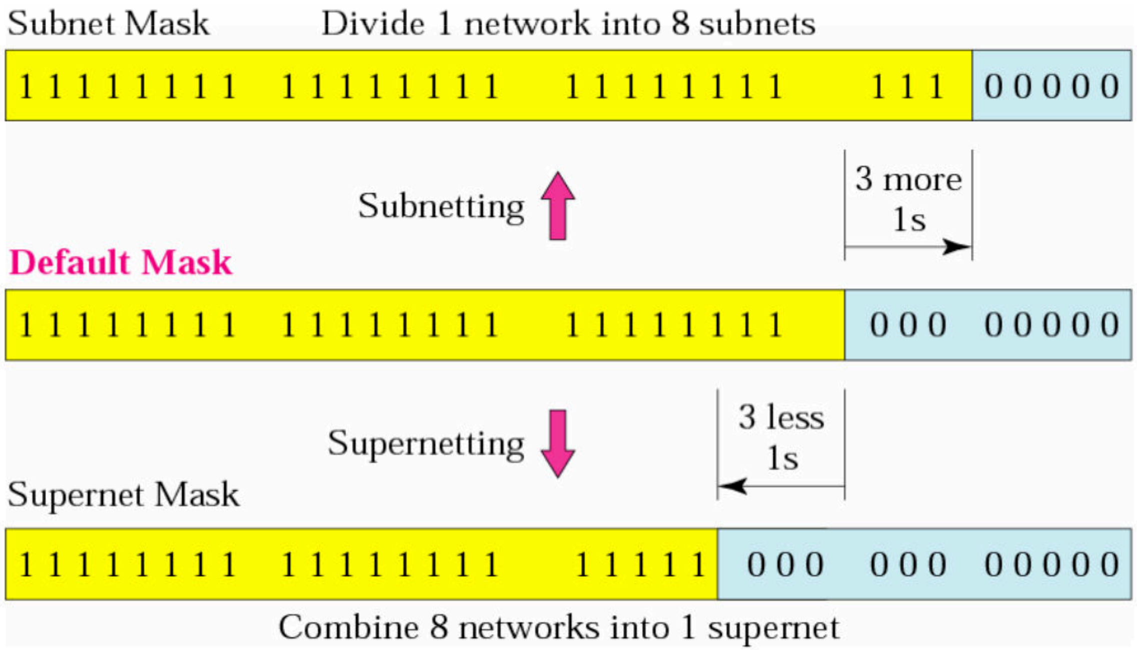

Subnet Mask

Divide 1 network into 8 subnets

11111111 11111111 11111111 111 00000

Subnetting

3 more 1s

Default Mask

11111111 11111111 11111111 000 00000

Supernetting

3 less 1s

Supernet Mask

11111111 11111111 11111 000 000 00000

Combine 8 networks into 1 supernet

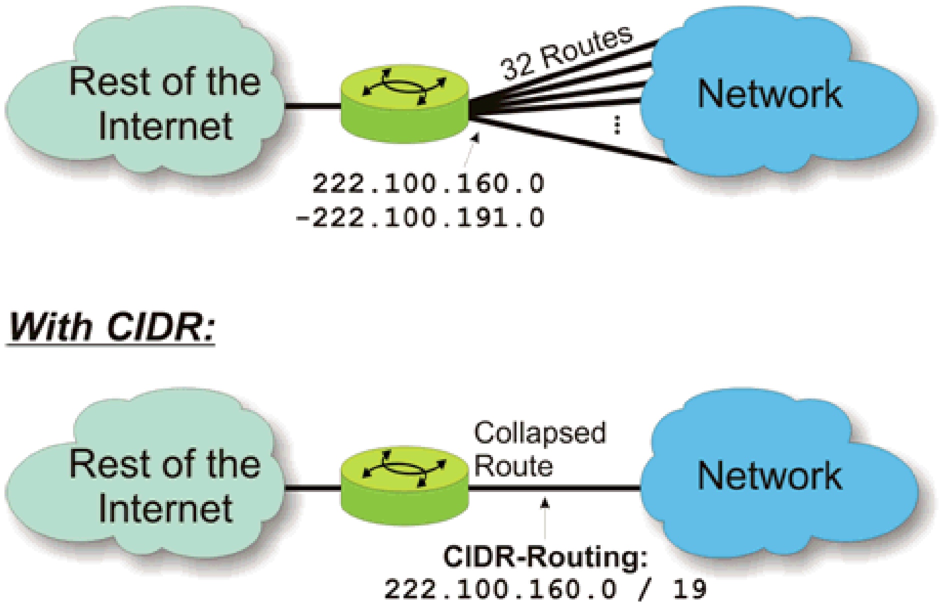

Rest of the Internet

32 Routes

Network

222.100.160.0
-222.100.191.0

With CIDR:

Rest of the Internet

Collapsed Route

Network

CIDR-Routing:
222.100.160.0 / 19

IP Addressing and Subnetting

IP Header Format

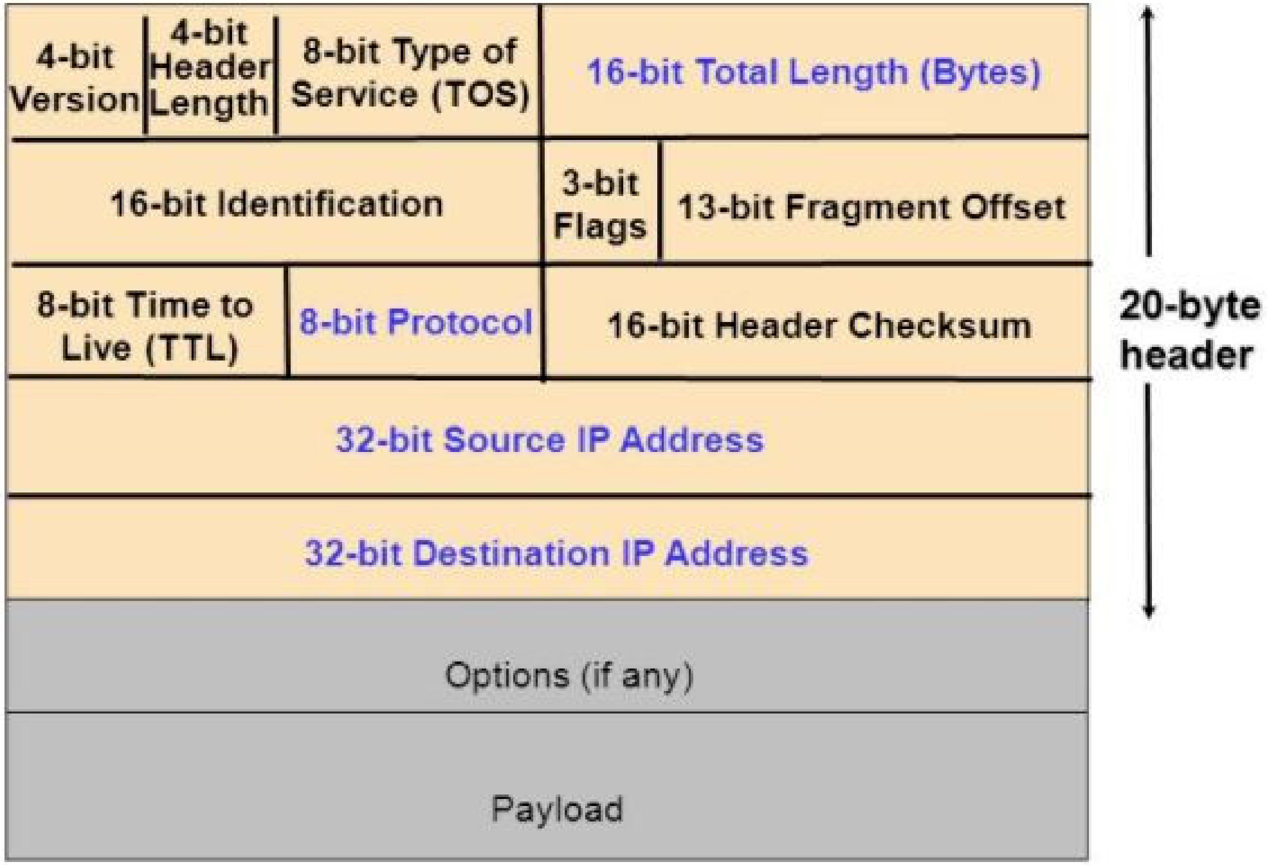

4-bit Version | 4-bit Header Length | 8-bit Type of Service (TOS) | 16-bit Total Length (Bytes)

16-bit Identification | 3-bit Flags | 13-bit Fragment Offset

8-bit Time to Live (TTL) | 8-bit Protocol | 16-bit Header Checksum

32-bit Source IP Address

32-bit Destination IP Address

Options (if any)

Payload

20-byte header

- **20 bytes < Header Size < 60 bytes**
- **20 bytes < Total Length < 65536 bytes**

IP Addressing and Subnetting

Maximum Transmission Unit (MTU)

- Maximum size of IP datagram is 65535, but the data link layer protocol generally imposes a limit that is much smaller. The limit is called maximum transmission unit (MTU).

- What if the size of an IP datagram exceeds the MTU?
- What if the route contains networks with different MTUs?

> IP datagram is fragmented into smaller units!

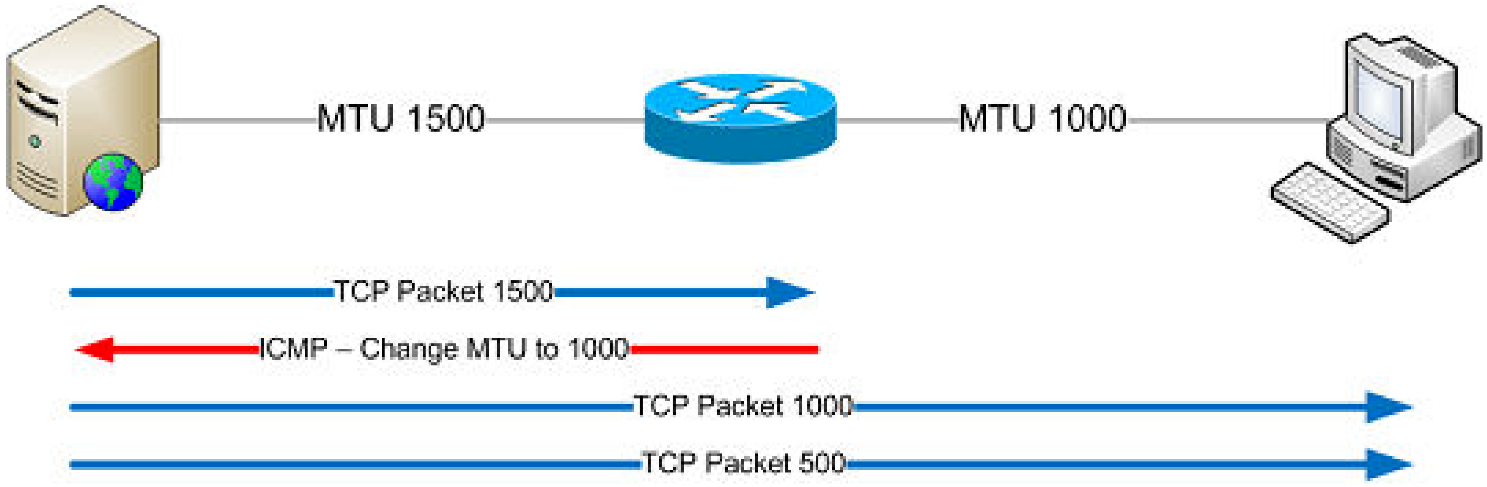

IP Addressing and Subnetting

IP Fragmentation

- IP router splits the datagram into several datagram. Fragments are reassembled at receiver.

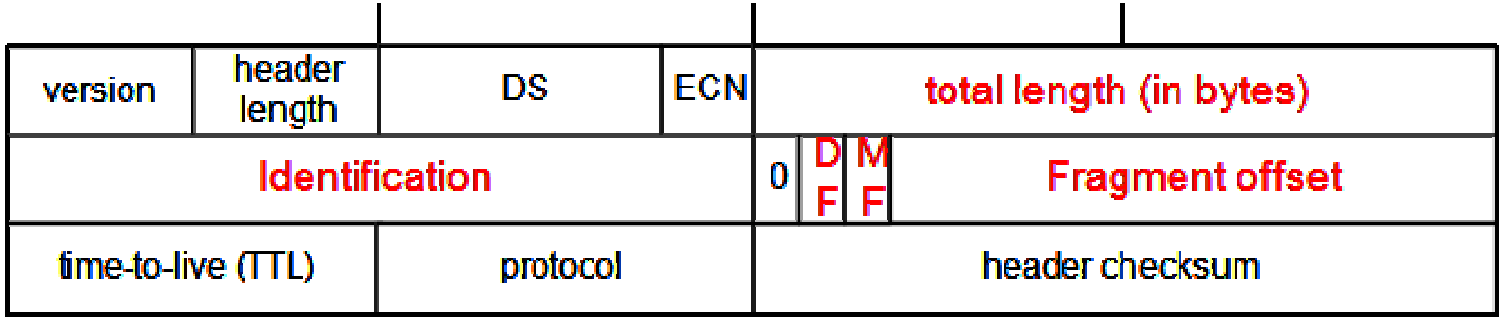

<table>
<tr><td>version</td><td>header length</td><td>DS</td><td>ECN</td><td colspan="4">total length (in bytes)</td></tr>
<tr><td colspan="4">Identification</td><td>0</td><td>DF</td><td>MF</td><td>Fragment offset</td></tr>
<tr><td colspan="2">time-to-live (TTL)</td><td colspan="2">protocol</td><td colspan="4">header checksum</td></tr>
</table>

- Identification: When a datagram is fragmented, the identification is the same in all fragments
- Flags:
  - DF bit is set: Datagram cannot be fragmented and must be discarded if MTU is too small
  - MF bit set: This datagram is part of a fragment and an additional fragment follows this one

IP Addressing and Subnetting

IP Fragmentation

- Fragment offset: Offset of the payload of the current fragment in the original datagram
- Total length: Total length of the current fragment

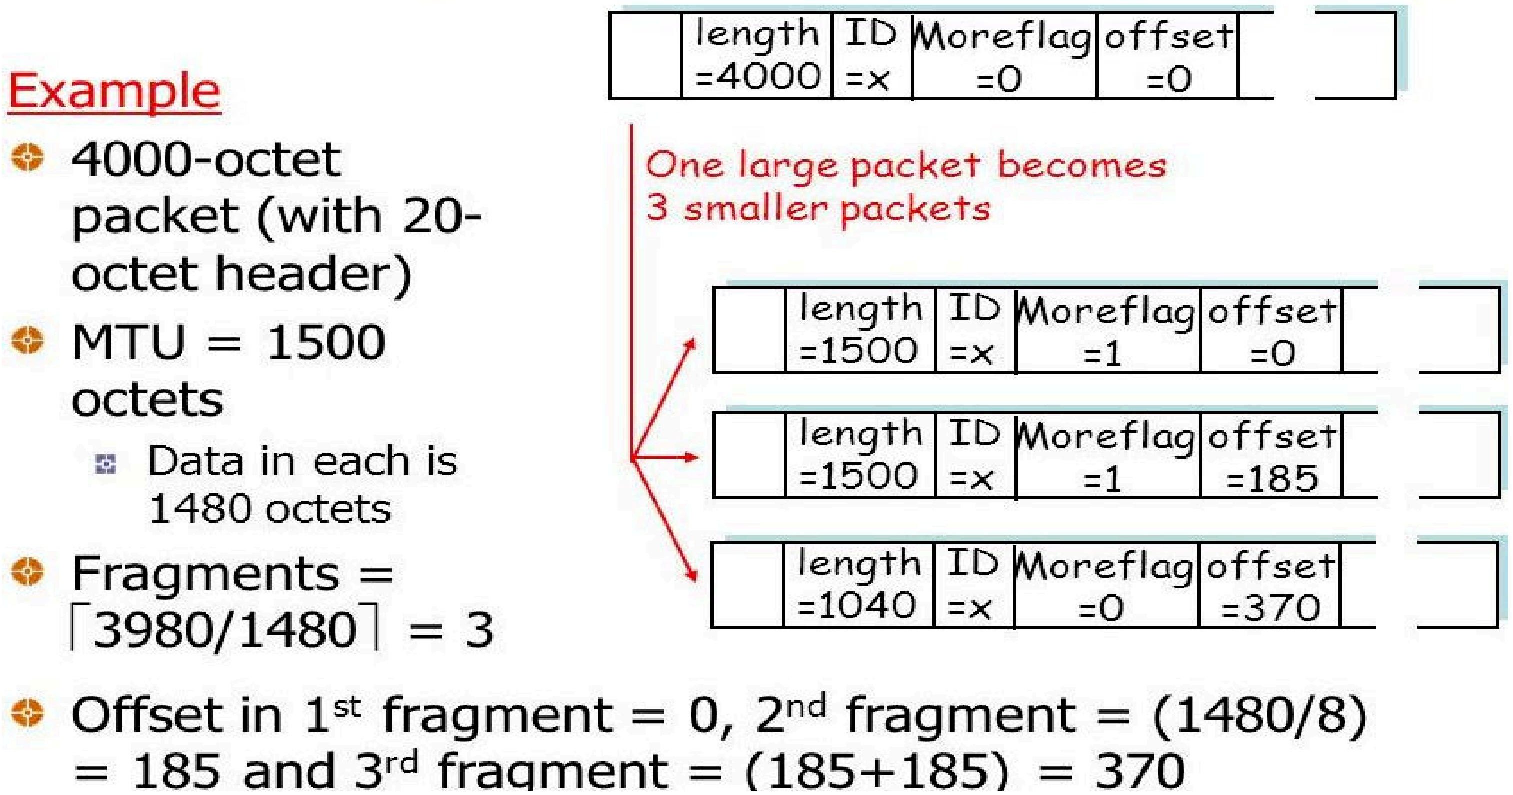

Example

- 4000-octet packet (with 20-octet header)
- MTU = 1500 octets
  - Data in each is 1480 octets
- Fragments = ⌈3980/1480⌉ = 3
- Offset in 1st fragment = 0, 2nd fragment = (1480/8) = 185 and 3rd fragment = (185+185) = 370

	length	ID	Moreflag	offset		
	=4000	=x	=0	=0		

One large packet becomes 3 smaller packets

	length	ID	Moreflag	offset		
	=1500	=x	=1	=0		
	=1500	=x	=1	=185		
	=1040	=x	=0	=370		

IP Addressing and Subnetting

•Address Resolution Protocol (ARP)

- IP Addresses are not recognized by hardware.
- The process of finding the hardware address of a host given the IP address is called Address Resolution.
- The Address Resolution Protocol is used by a sending host when it knows the IP address of the destination but needs the Ethernet address.

- ARP is a broadcast protocol - every host on the network receives the request.
- Each host checks the request against it’s IP address - the right one responds.

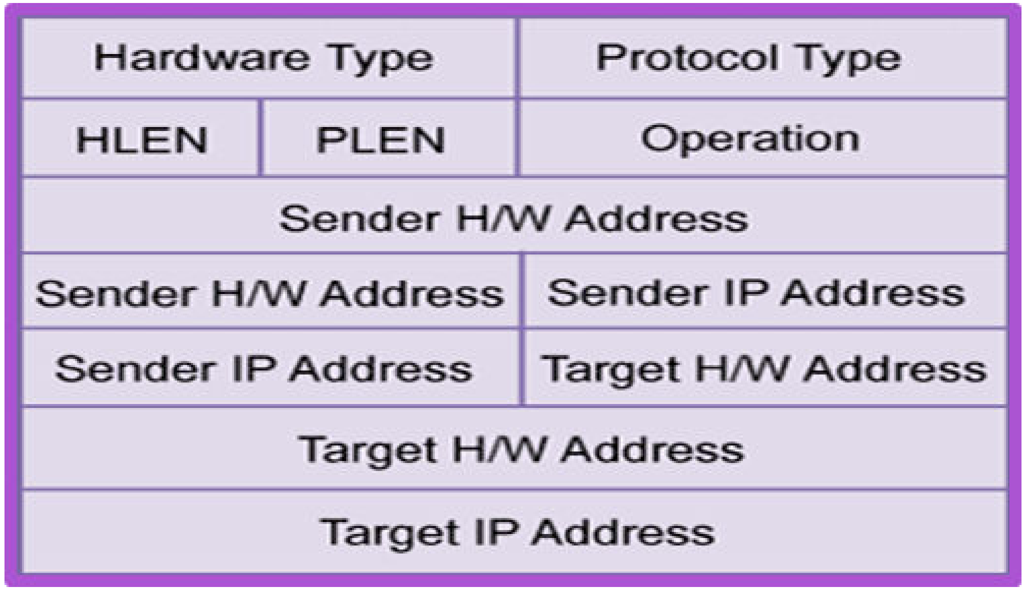

<table>
<tr><td colspan="2">Hardware Type</td><td>Protocol Type</td></tr>
<tr><td>HLEN</td><td>PLEN</td><td>Operation</td></tr>
<tr><td colspan="3">Sender H/W Address</td></tr>
<tr><td colspan="2">Sender H/W Address</td><td>Sender IP Address</td></tr>
<tr><td colspan="2">Sender IP Address</td><td>Target H/W Address</td></tr>
<tr><td colspan="3">Target H/W Address</td></tr>
<tr><td colspan="3">Target IP Address</td></tr>
</table>

IP Addressing and Subnetting

Address Resolution Protocol (ARP)

- **Example: Host 10.1.1.1 want to resolve MAC address of 10.1.1.10**

1) Host 10.1.1.1 sends broadcast ARP request

2) Host 10.1.1.1 gets unicast ARP reply from host 10.1.1.10

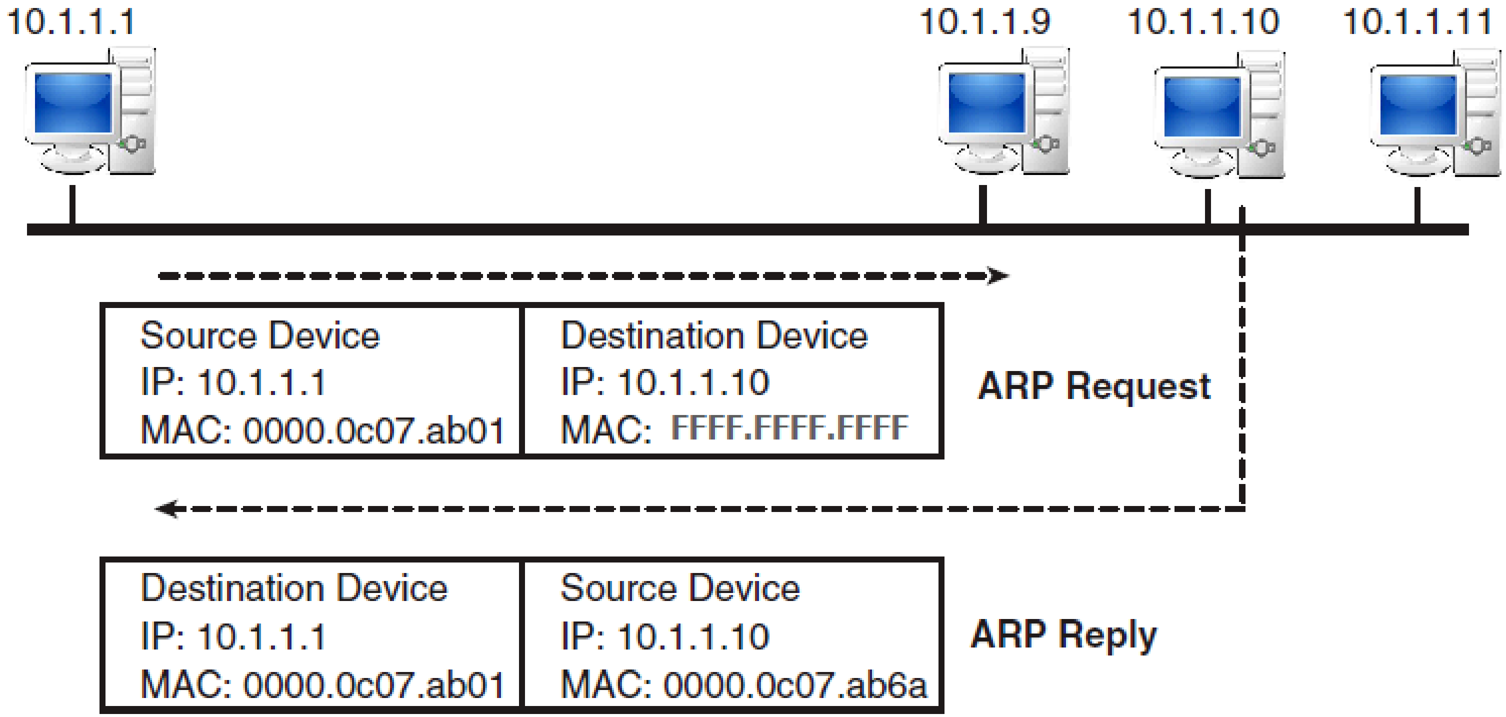

Source Device IP: 10.1.1.1 MAC: 0000.0c07.ab01	Destination Device IP: 10.1.1.10 MAC: FFFF.FFFF.FFFF	ARP Request

Destination Device IP: 10.1.1.1 MAC: 0000.0c07.ab01	Source Device IP: 10.1.1.10 MAC: 0000.0c07.ab6a	ARP Reply

IP Addressing and Subnetting

Internet Control Message Protocol (ICMP)

- ICMP messages are encapsulated in IP datagrams

- Functions of ICMP:

- A node recognizing a transmission problem (TTL exceed, destination unreachable, etc.) generates ICMP messages

- ICMP provides some useful diagnostics about network operation (ping, traceroute)

- ICMP Echo Request/Reply

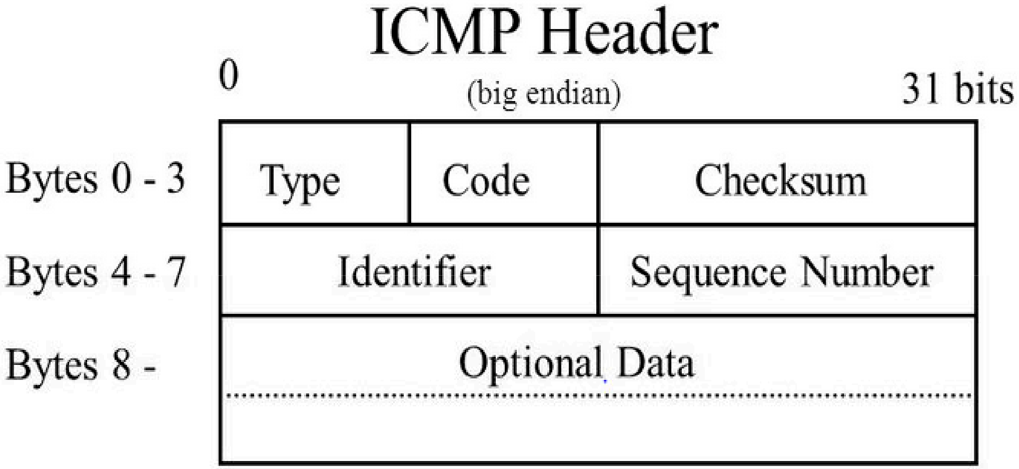

ICMP Header

(big endian)

0 … 31 bits

Bytes 0 - 3	Type	Code	Checksum
Bytes 4 - 7	Identifier		Sequence Number
Bytes 8 -	Optional Data		

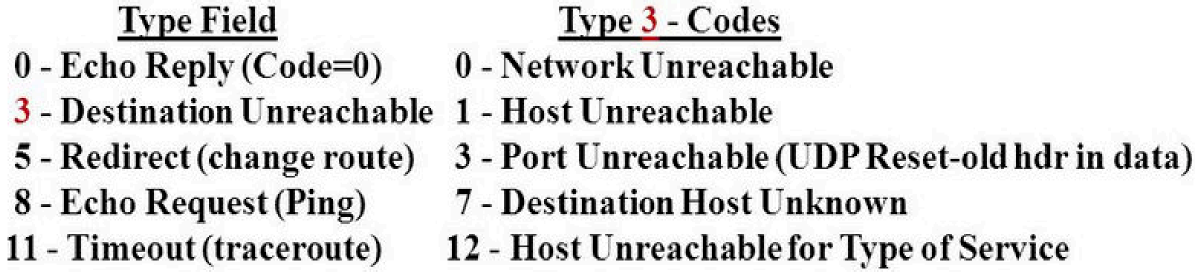

Type Field	Type 3 - Codes
0 - Echo Reply (Code=0)	0 - Network Unreachable
3 - Destination Unreachable	1 - Host Unreachable
5 - Redirect (change route)	3 - Port Unreachable (UDP Reset-old hdr in data)
8 - Echo Request (Ping)	7 - Destination Host Unknown
11 - Timeout (traceroute)	12 - Host Unreachable for Type of Service

IP Addressing and Subnetting

Internet Control Message Protocol (ICMP)

- Ping (Packet Internet groper) = ICMP echo request
- Why first ping fails? That is a result of the ARP resolution process:

```
PC>ping 192.168.1.12

Pinging 192.168.1.12 with 32 bytes of data:

Request timed out.

Reply from 192.168.1.12: bytes=32 time=156ms TTL=126

Reply from 192.168.1.12: bytes=32 time=156ms TTL=126

Reply from 192.168.1.12: bytes=32 time=156ms TTL=126
```

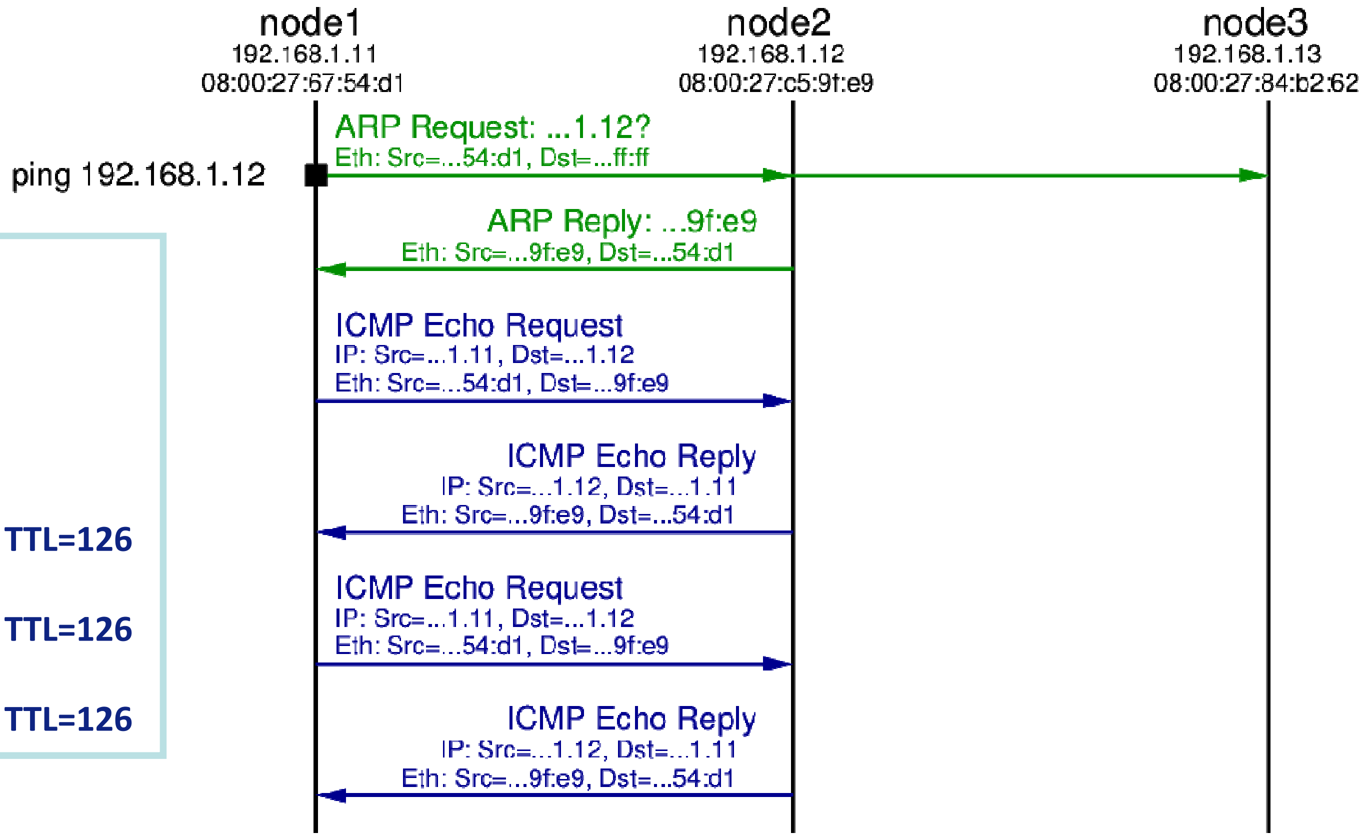

node1
192.168.1.11
08:00:27:67:54:d1

node2
192.168.1.12
08:00:27:c5:9f:e9

node3
192.168.1.13
08:00:27:84:b2:62

ping 192.168.1.12

ARP Request: ...1.12?
Eth: Src=...54:d1, Dst=...ff:ff

ARP Reply: ...9f:e9
Eth: Src=...9f:e9, Dst=...54:d1

ICMP Echo Request
IP: Src=...1.11, Dst=...1.12
Eth: Src=...54:d1, Dst=...9f:e9

ICMP Echo Reply
IP: Src=...1.12, Dst=...1.11
Eth: Src=...9f:e9, Dst=...54:d1

ICMP Echo Request
IP: Src=...1.11, Dst=...1.12
Eth: Src=...54:d1, Dst=...9f:e9

ICMP Echo Reply
IP: Src=...1.12, Dst=...1.11
Eth: Src=...9f:e9, Dst=...54:d1

IP Addressing and Subnetting

- **The Cisco IOS offers an "extended" mode of the ping command. This mode is entered by typing ping in privileged EXEC mode, without a destination IP address.**

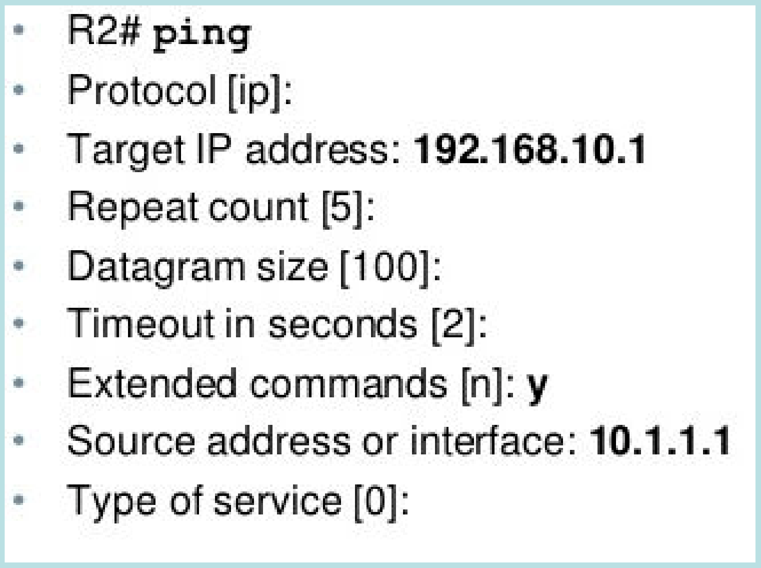

```
• R2# ping
• Protocol [ip]:
• Target IP address: 192.168.10.1
• Repeat count [5]:
• Datagram size [100]:
• Timeout in seconds [2]:
• Extended commands [n]: y
• Source address or interface: 10.1.1.1
• Type of service [0]:
```

- **Repeat count**: How many pings do you want to send? The default is five with standard ping.
- **Datagram size**: While the default is to send a 100-byte ping, with extended ping you could send very large ping packets
- **Timeout**: The default timeout is two seconds, but you could allow ping to wait much longer for a reply if you choose to do so.
- **Source interface**: You can specify the source of your ping because, otherwise, the receiving router may not be able to see all interfaces of your router and your standard ping may fail.
- **df-bit**: This option sets the Don’t Fragment bit in the IP header to indicate that routers should not fragment this packet.

IP Addressing and Subnetting

Traceroute

- Trace (Cisco = traceroute, Windows = tracert) is used to trace the path a packet takes between source and destination. Uses ICMP message within an IP Packet.
- Uses UDP as a the transport layer.

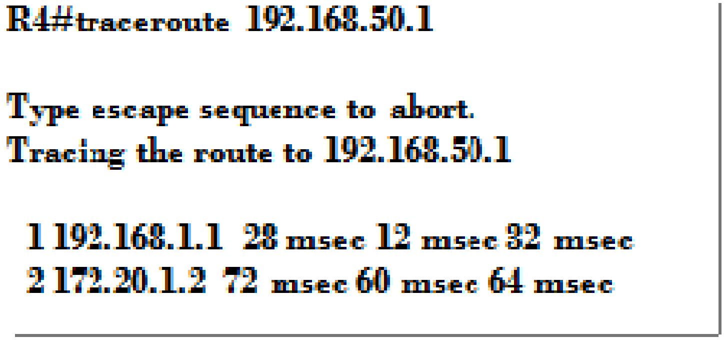

```
R4#traceroute  192.168.50.1

Type escape sequence to abort.
Tracing the route to 192.168.50.1

  1 192.168.1.1  28 msec 12 msec 32 msec
  2 172.20.1.2  72 msec 60 msec 64 msec
```

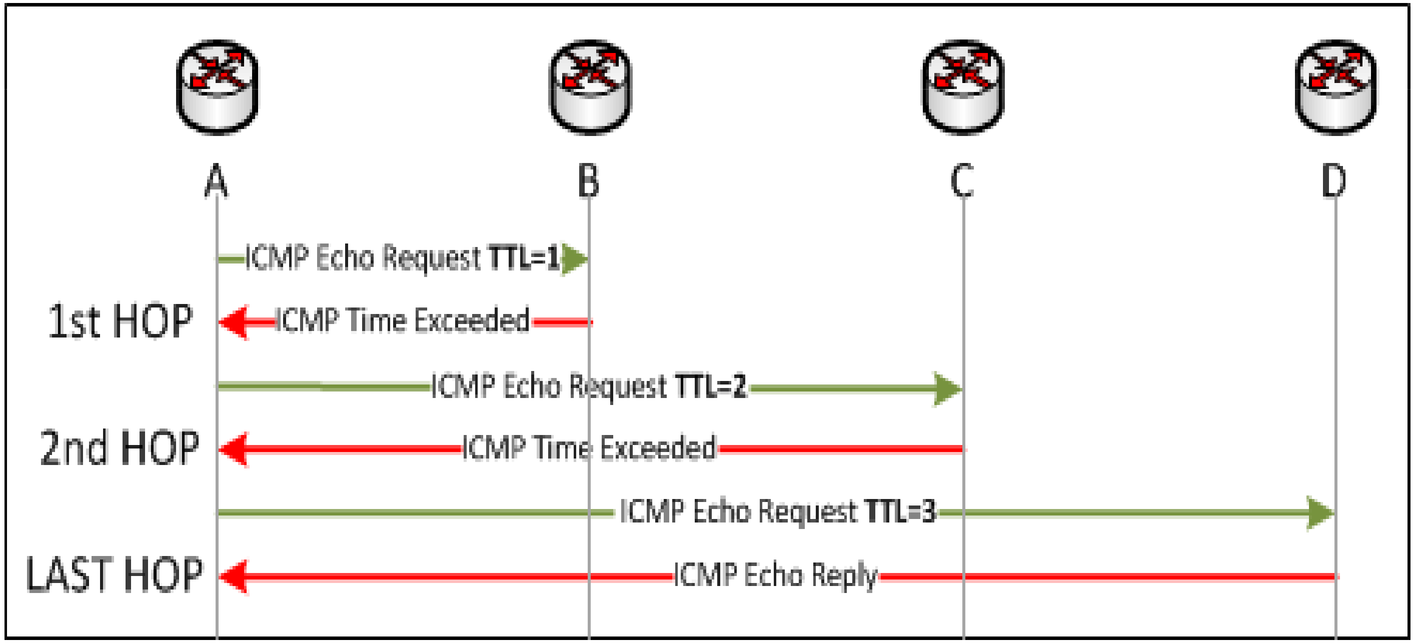

Cisco escape sequence: ctrl+shift+6

TRANPORT LAYER

TRANSPORT LAYER

Transport Layer Duties

- Packetizing: breaks application messages into segments
- Connection control: Connection-oriented or Connectionless
- Addressing: Port numbers to identify which network application
- Reliability: Flow control and Error Control

two transport layer protocols:

  - The Transport Control Protocol (TCP) for reliable service
  - The unreliable User Datagram Protocol (UDP)

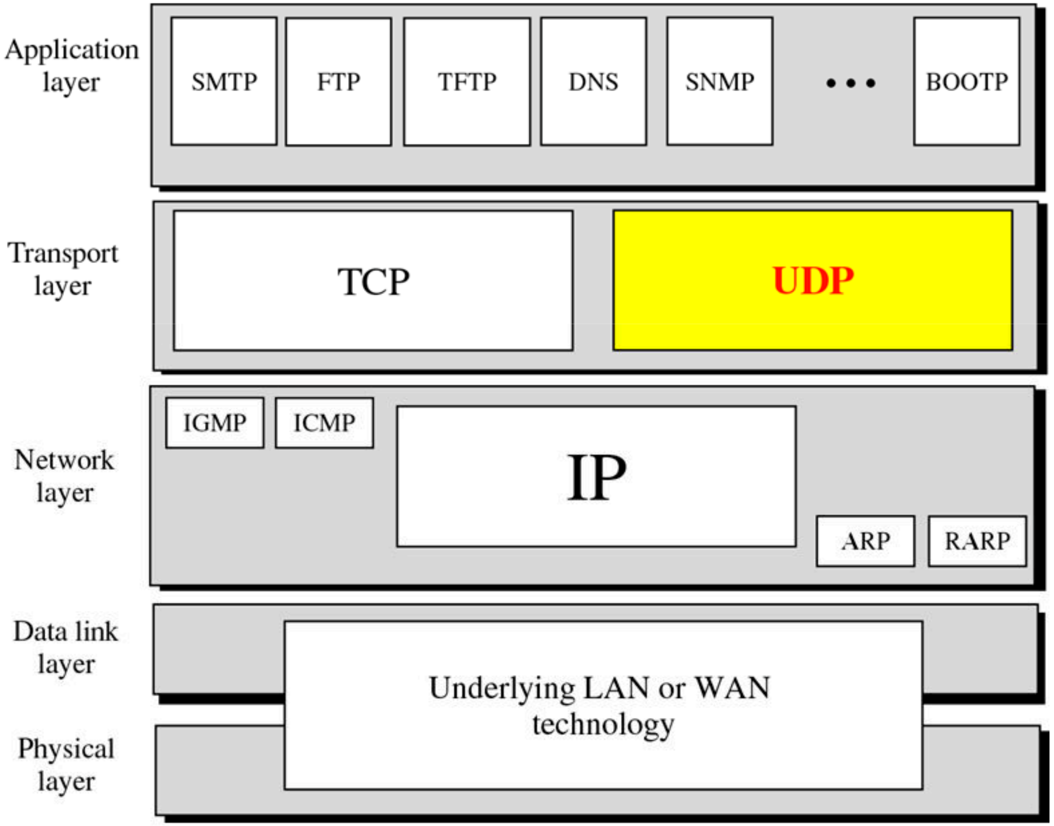

TRANSPORT LAYER

Transport layer - Sockets

- Transport layer at the receiving host delivers data to the socket
- There should be a unique identifier for each socket.
- Socket identifier is called socket address
- Socket address = IP address & Port number

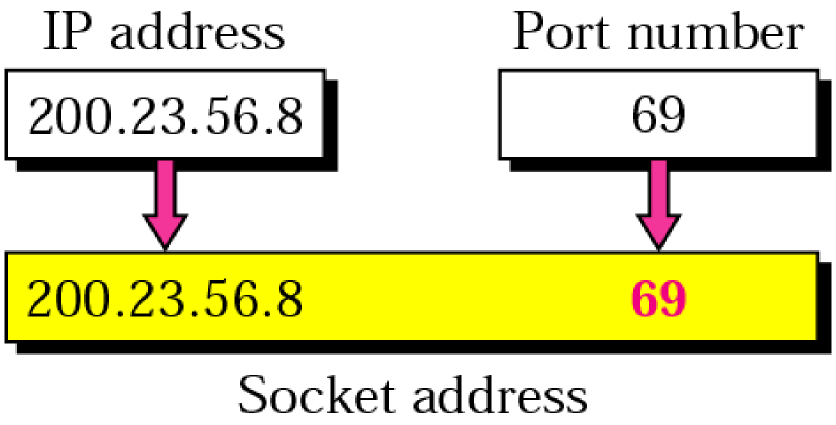

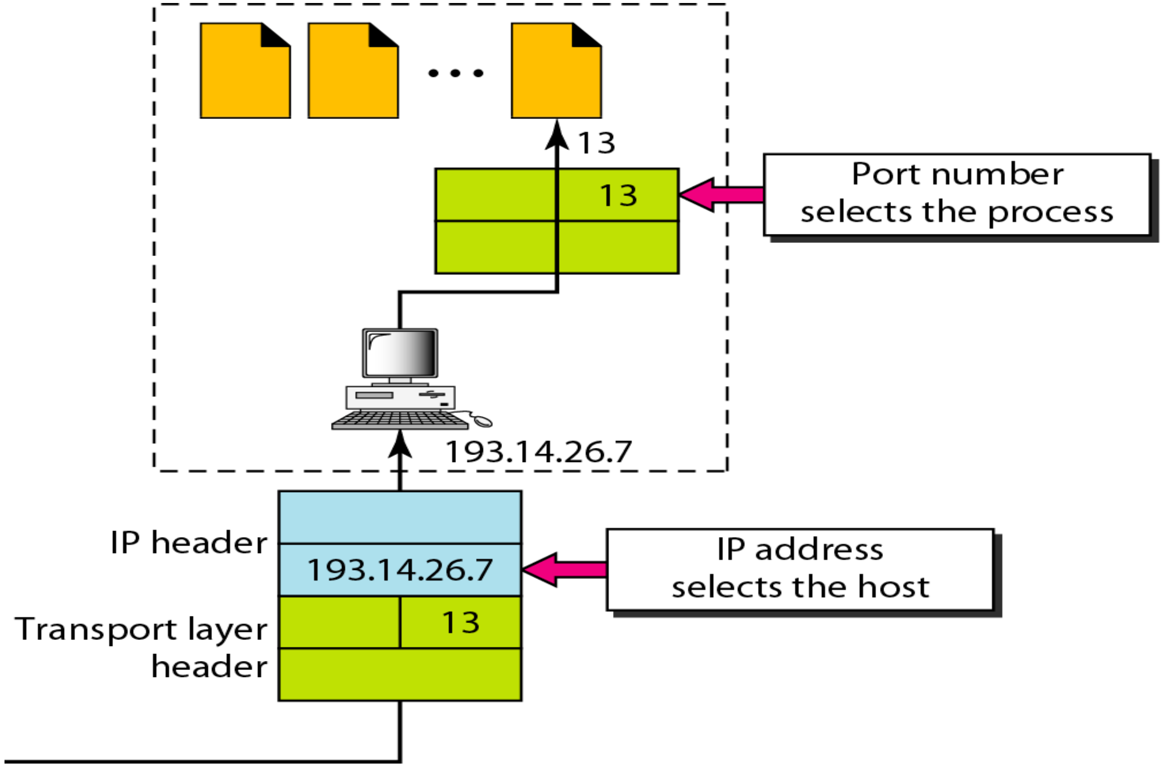

TRANSPORT LAYER

Port numbers Ranges

- Port numbers are 16-bit integers between 0 – 65535
- The Internet Assigned Numbers Authority (IANA) is responsible for maintaining the official assignments of port numbers for specific uses

<u>Well Known Ports (Numbers 0 to 1023)</u>
- These numbers are reserved for services and applications.

<u>Registered Ports (Numbers 1024 to 49151)</u>
- These port numbers are assigned to user processes or applications

<u>Dynamic or Private Ports (Numbers 49152 to 65535)</u> - Also known as Ephemeral Ports, these are usually assigned dynamically to client applications when initiating a connection.

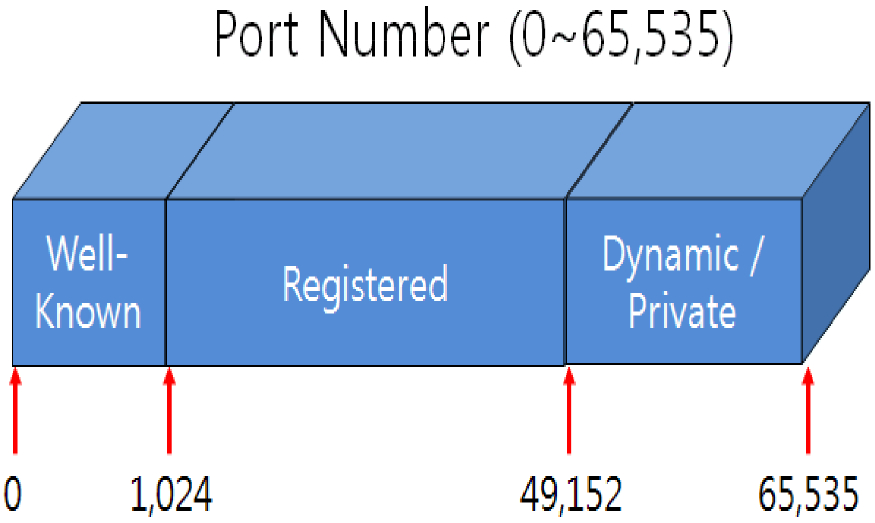

TRANSPORT LAYER

Port Numbers

- The *well known ports* are assigned by *IANA (Internet Assigned Numbers Authority)* in the range of 0 to 1023.
- Some well-known ports used by TCP and UDP:

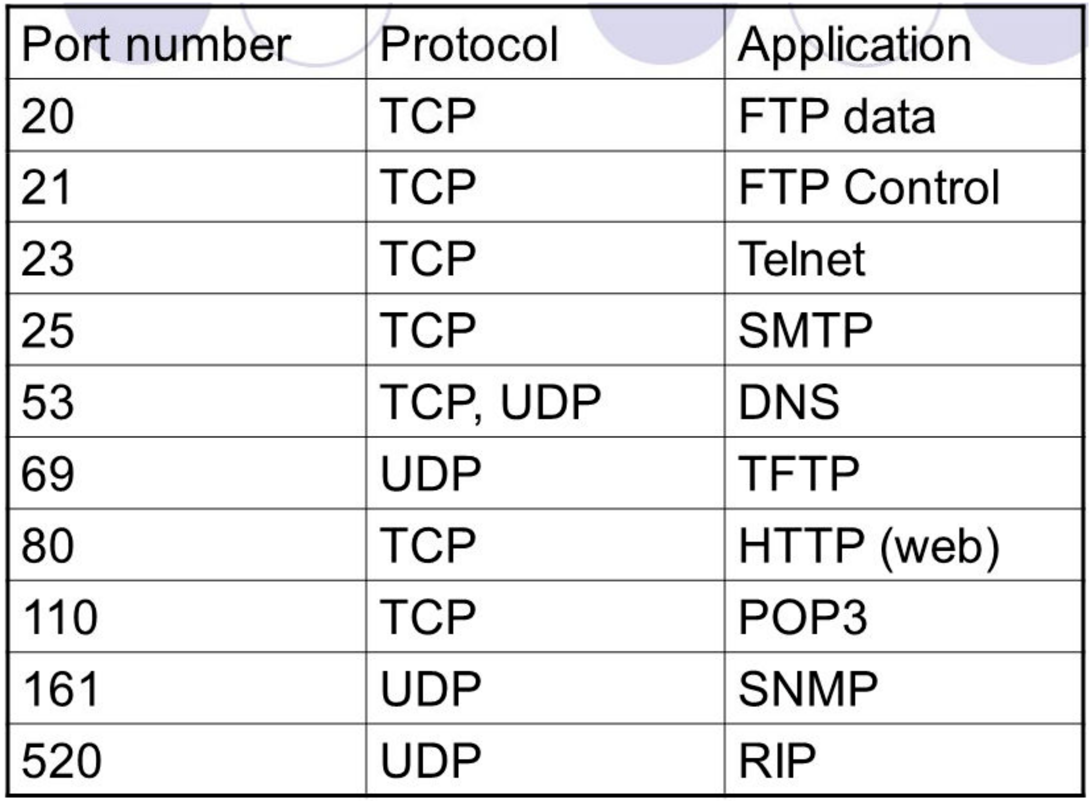

Port number	Protocol	Application
20	TCP	FTP data
21	TCP	FTP Control
23	TCP	Telnet
25	TCP	SMTP
53	TCP, UDP	DNS
69	UDP	TFTP
80	TCP	HTTP (web)
110	TCP	POP3
161	UDP	SNMP
520	UDP	RIP

TRANSPORT LAYER

Types of Connection

- Connection-Oriented or Connectionless:
  - Some protocols are connection-oriented: once things are set up, you always talk to a single endpoint
  - Connection oriented means that a virtual connection is established before any data is transferred.
  - Example: TCP
- Connectionless:
  - Each packet can go to or come from a different place
  - No handshaking between sender and receiver
  - Each UDP segment handled independently of others
  - Example: UDP, as used in the DNS, TFTP, etc

TRANSPORT LAYER

Multiplexing and demultiplexing

- Multiplexing: (at the sending node) The process of encapsulating data messages from different applications sockets with the header information and pass the segments to the network layer

- DeMultiplexing: (at the receiving node) The process of delivering the received data segment to the correct application

- Example:

  Suppose that the following is running on the same computer:

  - Downloading a web page while transferring data through FTP
  - Two telnet sessions are also running
  - Transport layer receives TPDUs from network layer for all four processes

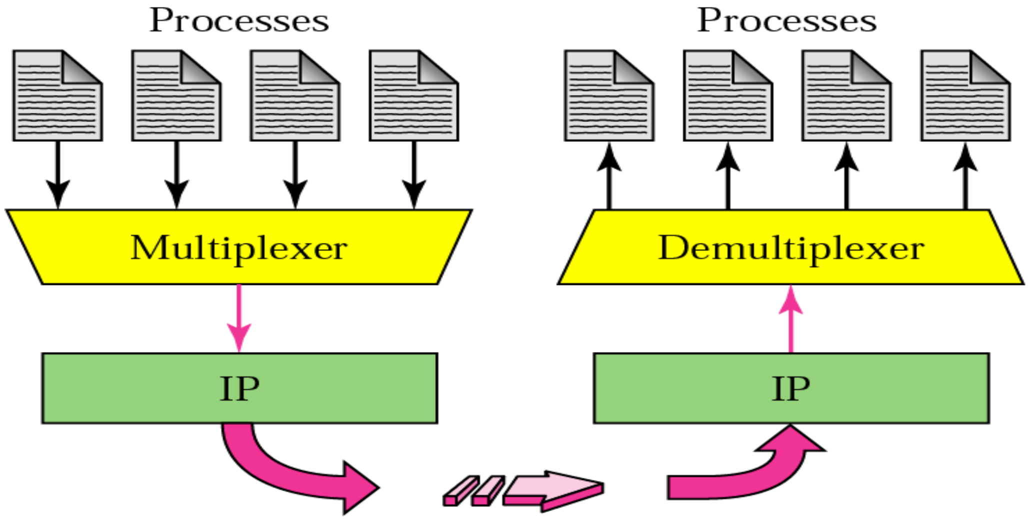

TRANSPORT LAYER

Layer 4 Protocols

- There are two main protocols at this layer; the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP).
- TCP e UDP Headers:

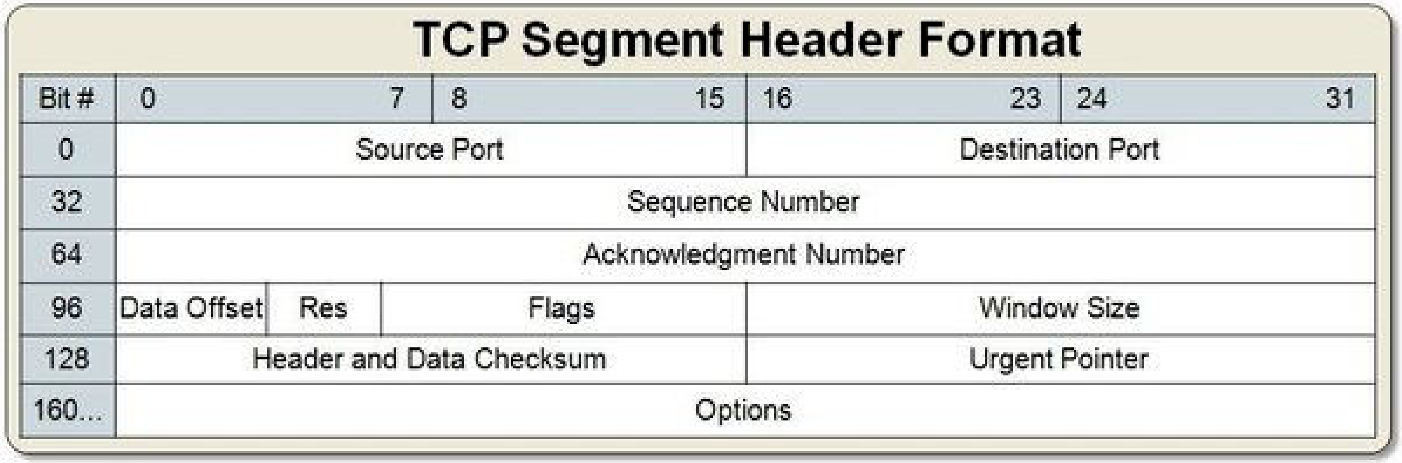

TCP Segment Header Format

Bit #	0 7	8 15	16 23	24 31
0	Source Port		Destination Port	
32	Sequence Number			
64	Acknowledgment Number			
96	Data Offset Res	Flags	Window Size	
128	Header and Data Checksum		Urgent Pointer	
160…	Options			

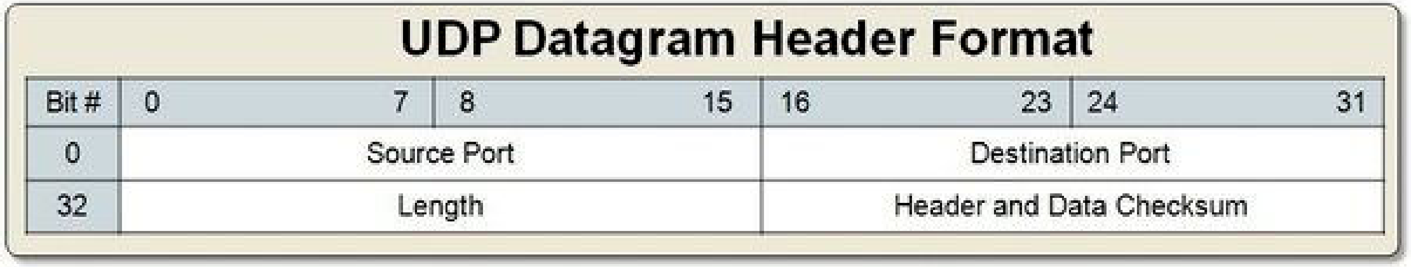

UDP Datagram Header Format

Bit #	0 7	8 15	16 23	24 31
0	Source Port		Destination Port	
32	Length		Header and Data Checksum	

Header TCP = 20bytes

Header UDP = 8 Btyes

TRANSPORT LAYER

- **Main differences between TCP and UDP:**

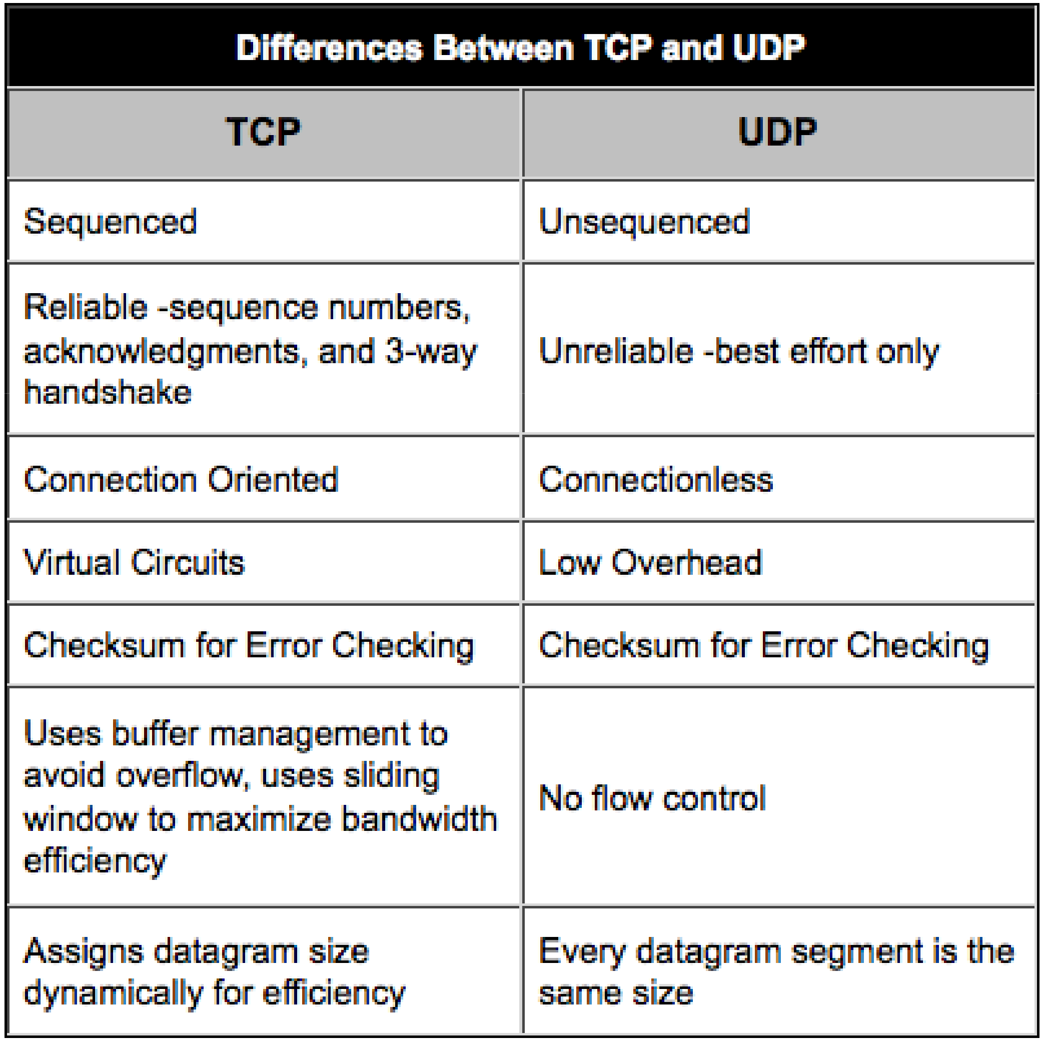

Differences Between TCP and UDP

TCP	UDP
Sequenced	Unsequenced
Reliable -sequence numbers, acknowledgments, and 3-way handshake	Unreliable -best effort only
Connection Oriented	Connectionless
Virtual Circuits	Low Overhead
Checksum for Error Checking	Checksum for Error Checking
Uses buffer management to avoid overflow, uses sliding window to maximize bandwidth efficiency	No flow control
Assigns datagram size dynamically for efficiency	Every datagram segment is the same size

TCP Flow-control

- TCP is a sliding window protocol
  - For window size n, can send up to n bytes without receiving an acknowledgement
  - When the data is acknowledged then the window slides forward

- Each packet advertises a window size
  Indicates number of bytes the receiver has space for

- Original TCP always sent entire window
  Congestion control now limits this

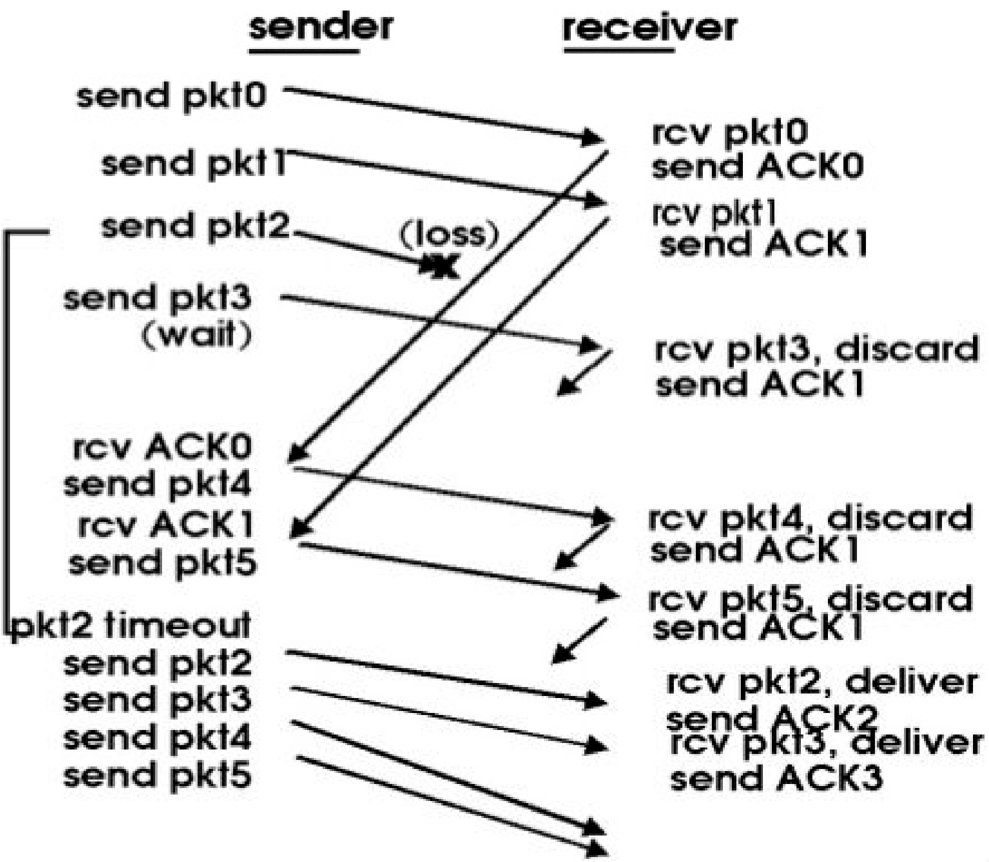

TRANSPORT LAYER

TCP Three-Way Handshake

- To establish or terminate connections reliably, TCP uses a 3-way handshake in which three messages are exchanged
- During the process to start a connection, each side sends a control message that specifies an initial buffer size (for flow control) and a sequence number.

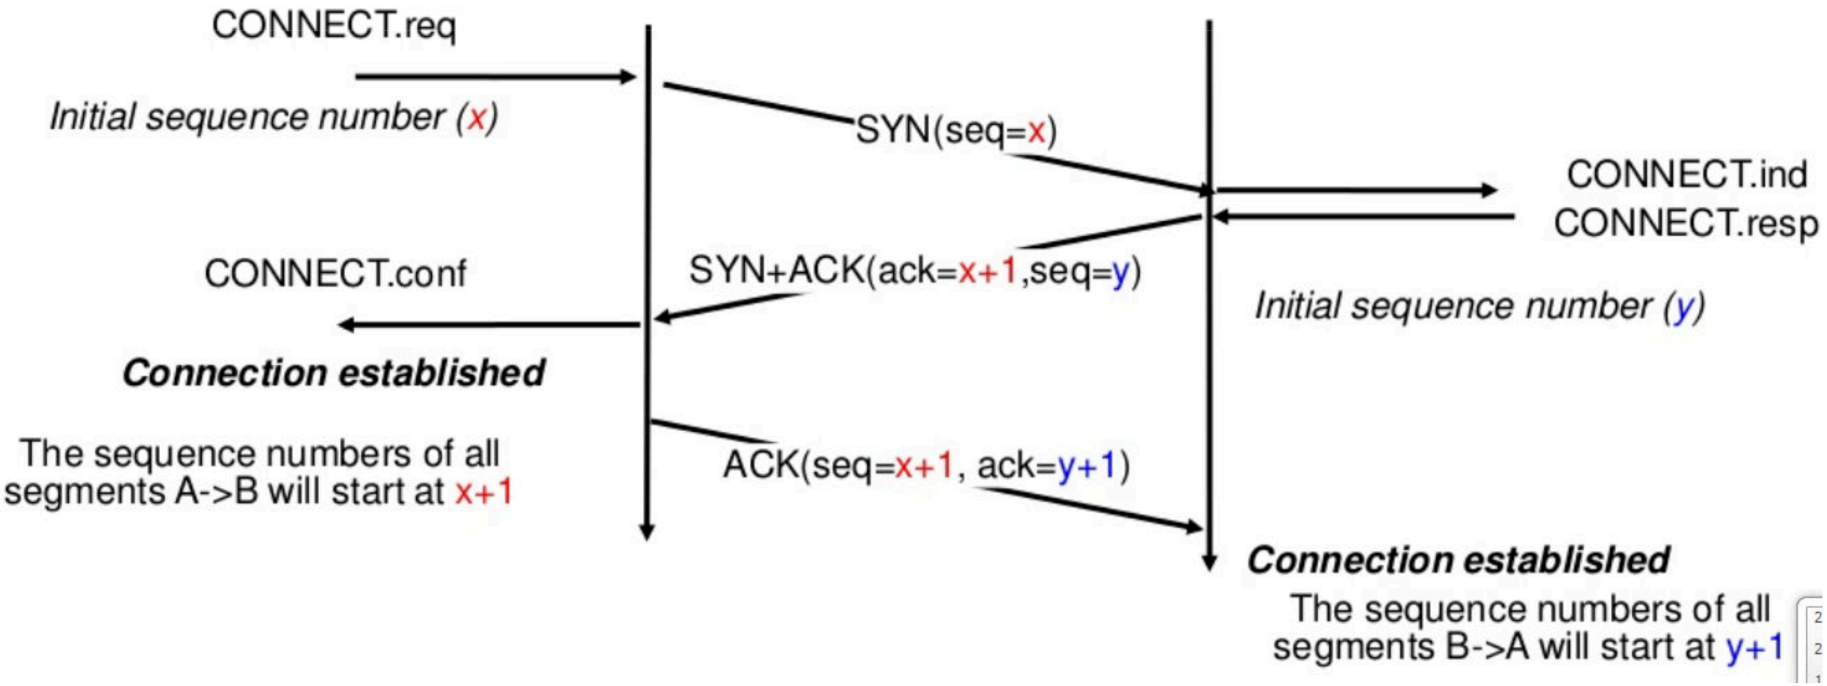
CONNECT.req
Initial sequence number (x)
SYN(seq=x)
CONNECT.ind
CONNECT.resp
CONNECT.conf
SYN+ACK(ack=x+1,seq=y)
Initial sequence number (y)
Connection established
The sequence numbers of all segments A->B will start at x+1
ACK(seq=x+1, ack=y+1)
Connection established
The sequence numbers of all segments B->A will start at y+1

TRANSPORT LAYER

TCP Client-Server Port Allocation

- Unless a client program explicitly requests a specific port number, the port number used is an dynamic port number (from 49152 through 65535).
- The allocations are temporary and only valid for the duration of the communication session. After completion (or timeout) of the communication session, the ports become available for reuse

Clients Sending TCP Requests

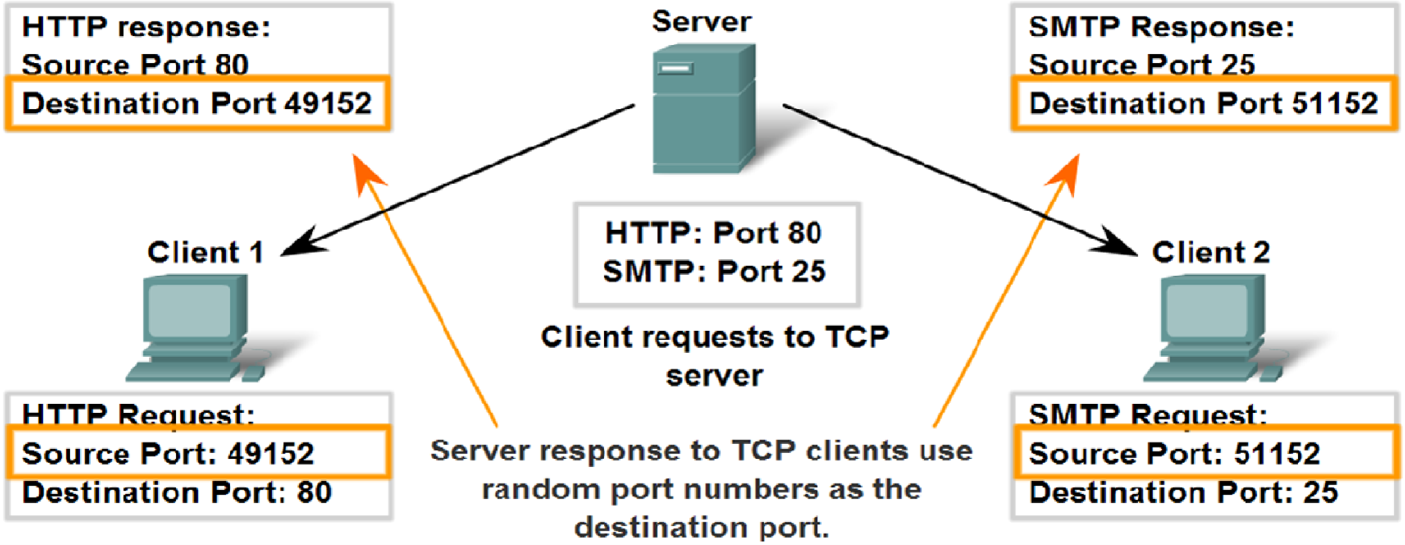

HTTP response:
Source Port 80
Destination Port 49152

Server

SMTP Response:
Source Port 25
Destination Port 51152

HTTP: Port 80
SMTP: Port 25

Client 1

Client 2

Client requests to TCP server

HTTP Request:
Source Port: 49152
Destination Port: 80

Server response to TCP clients use random port numbers as the destination port.

SMTP Request:
Source Port: 51152
Destination Port: 25

TRANSPORT LAYER

REDDIG Applications that use UDP or TCP:

- TCP:
  - AMHS
  - AFTN
  - AIDC (ATS Interfacility Data Communications)
  - RADAR

- UDP:
  - Voice Services

IP ROUTING

IP ROUTING

Primary Functions of the Router:

1) Packet forwarding
- the process used to switch a packet from an incoming interface to an outgoing interface on the same router.

2) Path selection
– determines the best path to the destination network

- Router de-encapsulates the frame
- Remaining packet passed up to layer 3
    - Routing decision made at this layer by examining destination IP address
- Packet is then re-encapsulated & sent out outbound interface

IP ROUTING

Router - Planes of Operation

- The control plane: The control plane is the brain of the router. It consists of routing protocols, routing updates, protocols such as IGMP, ICMP, ARP, BFD, LACP, and so on.
- The data plane: It is the forwarding plane, which is responsible for the switching of packets through the router.
- The management plane: It is used to manage a device through its connection to the network. Examples of protocols include SNMP, Telnet, FTP and SSH.

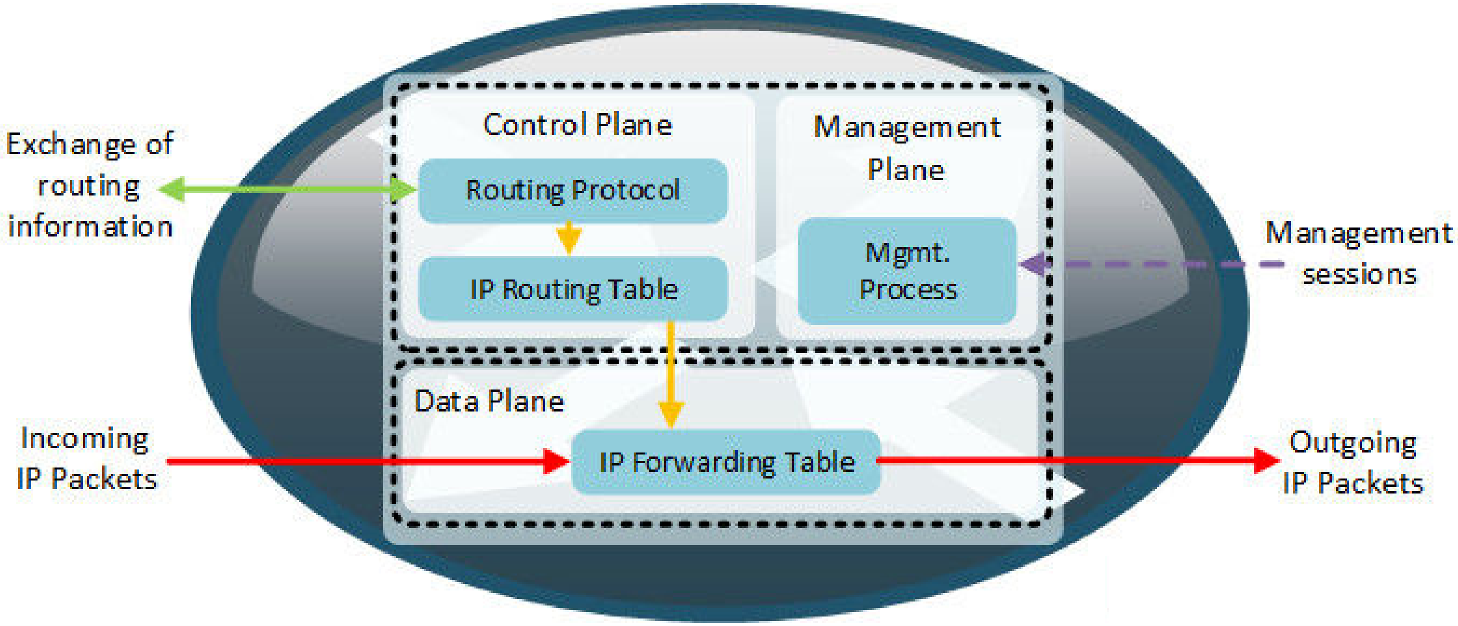

IP ROUTING

Cisco Express Forwarding (CEF)

- CEF is an optimized Layer 3 forwarding path through a router or switch. CEF optimizes routing table lookup by creating a special, easily searched tree structure based on the IP routing table. The forwarding information is called the Forwarding Information Base (FIB), and the cached adjacency information is called the Adjacency Table.

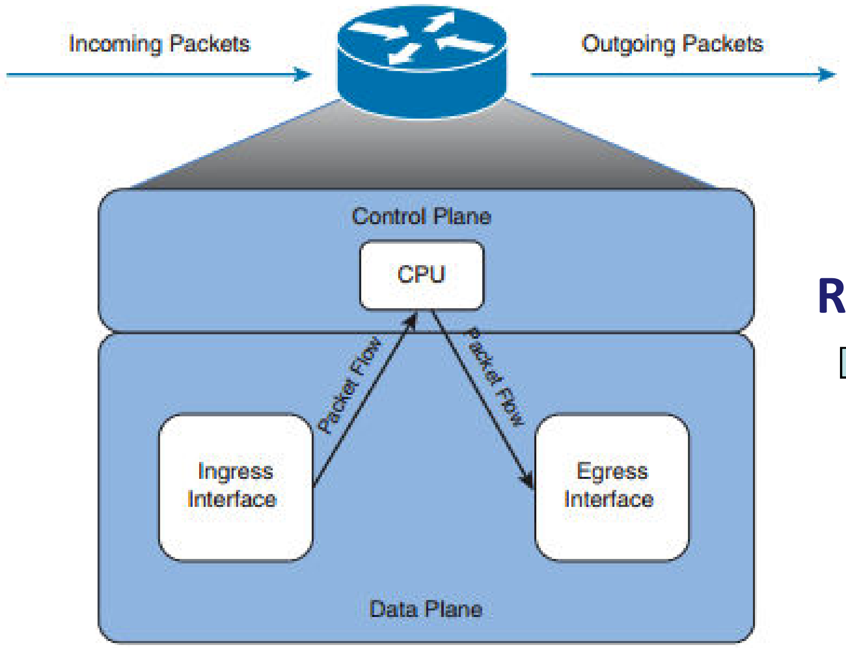

Router(config)# ip cef

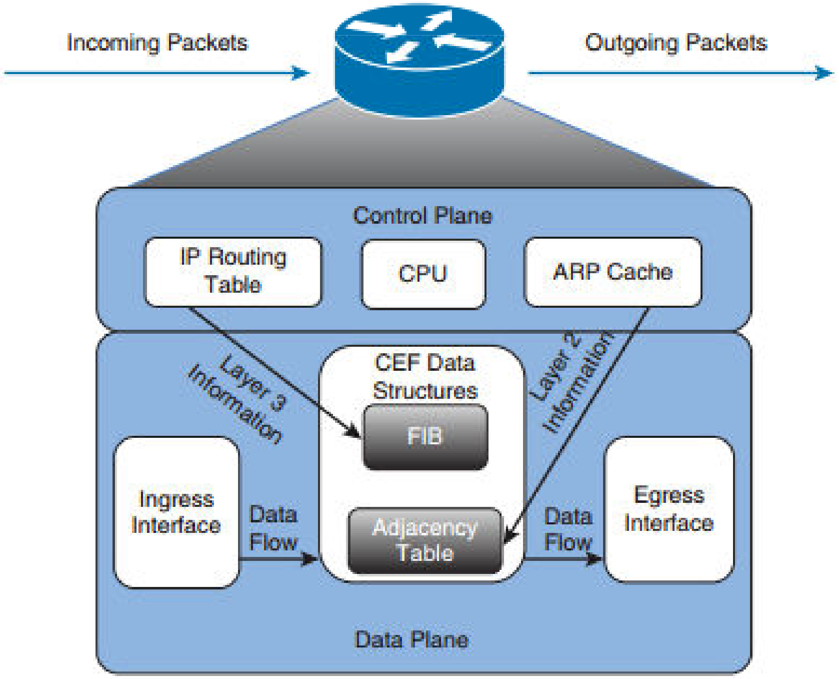

IP ROUTING

Routing Componentes

- Routing = building maps and giving directions
- Forwarding = moving packets between interfaces according to the "directions"

- RIB – Routing Information Base

This is the route table. When you do a show ip route, the RIB is what you see

- FIB – Forwarding Information Base

The FIB is an optimised version of the RIB. Or more correctly it's the table a router looks at when deciding where to actually forward traffic. In Cisco land, the CEF table is a FIB.

IP ROUTING

Routing Table

- A routing table lists all networks for which routes are known. The routing table is stored in the RAM of the device.
- When a router receives a packet that needs to be forwarded to a host on another network, it examines its destination IP address and looks for the routing information stored in the routing table.
- Each entry in the routing table consists of the following entries:
  - the network and the subnet mask – specifies a range of IP addresses.
  - the remote router – the IP address of the router used to reach that network.
  - the outgoing interface – the outgoing interface the packet should go out to reach the destination network.

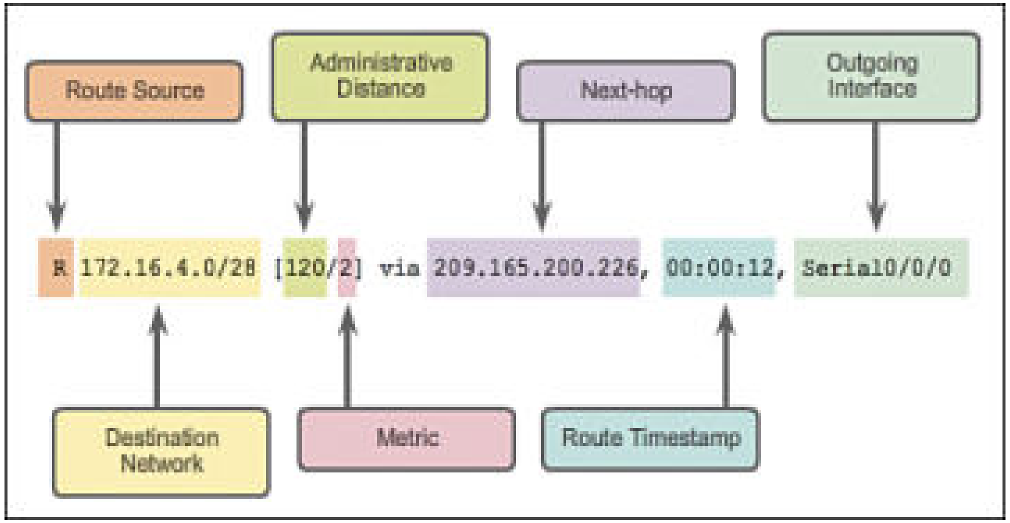

```
R 172.16.4.0/28 [120/2] via 209.165.200.226, 00:00:12, Serial0/0/0
```

IP ROUTING

Routing table lookup

Consider the following example:

Host A wants to communicate with Host B. Because hosts are on different subnets, Host A sends its packet to the default gateway (the router)

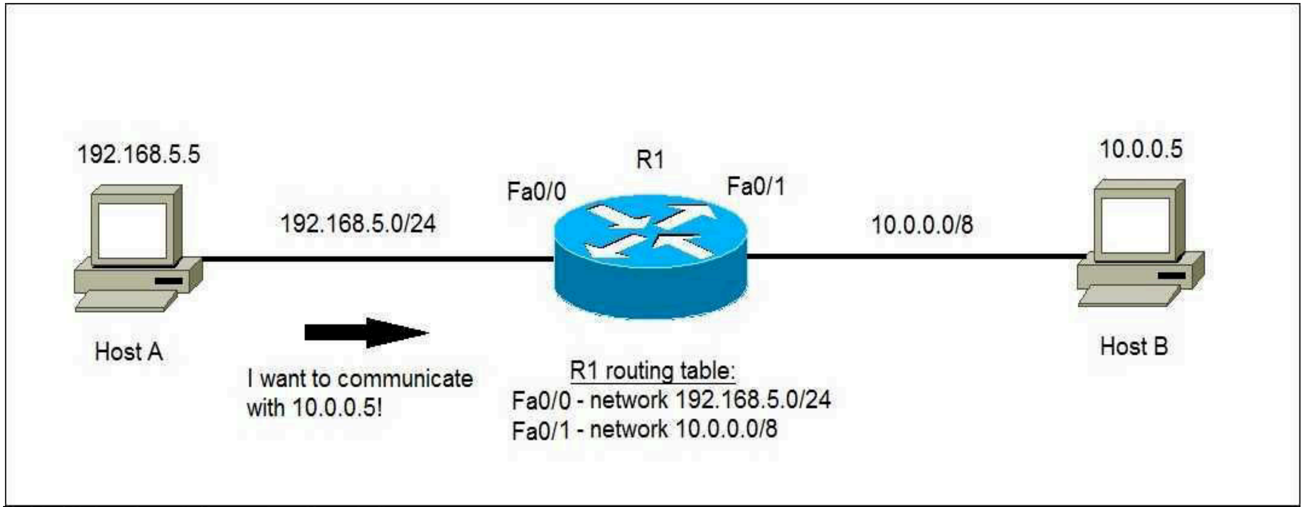

I want to communicate with 10.0.0.5!

R1 routing table:
Fa0/0 - network 192.168.5.0/24
Fa0/1 - network 10.0.0.0/8

The router receives the packet, examines the destination IP address, and looks up into its routing table to figure out which interface the packet will be sent out.

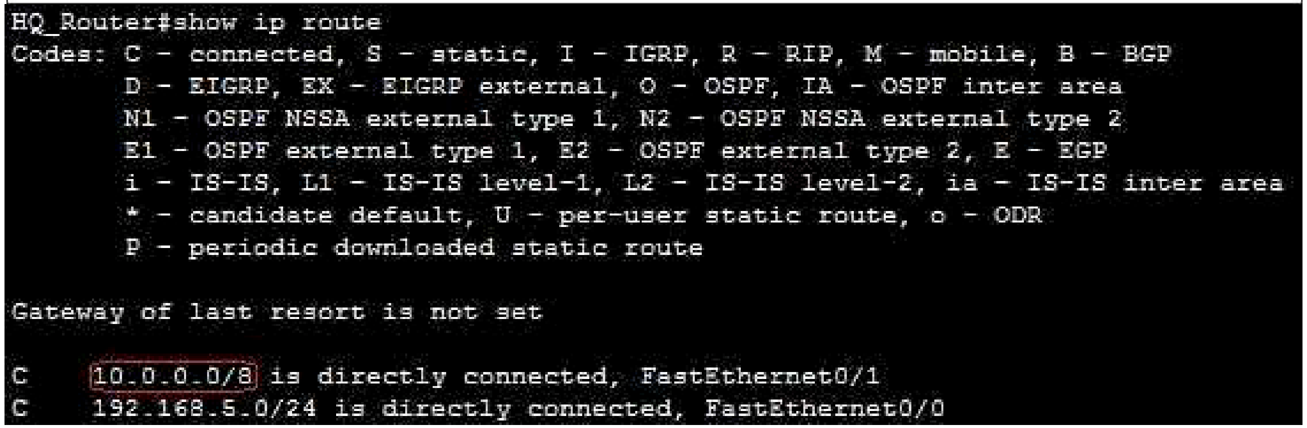

```
HQ_Router#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

C    10.0.0.0/8 is directly connected, FastEthernet0/1
C    192.168.5.0/24 is directly connected, FastEthernet0/0
```

This is the entry that will be used to route the packet:

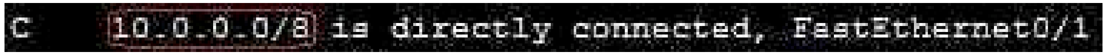

```
C    10.0.0.0/8 is directly connected, FastEthernet0/1
```

IP ROUTING

Routing Table Sources

The show ip route command is used to display the contents of the routing table. Entries in the routing table can be added as:

- Local Route interfaces - Added when an interface is configured and active. This entry is only displayed in IOS 15 or newer for IPv4 routes and all IOS releases for IPv6 routes.
- Directly connected interfaces - Added to the routing table when an interface is configured and active.
- Static routes - Added when a route is manually configured and the exit interface is active.
- Dynamic routing protocol - Added when routing protocols that dynamically learn about the network, such as RIP, EIGRP or OSPF, are implemented and networks are identified.

IP ROUTING

Packet forwarding Decision Process in a Router

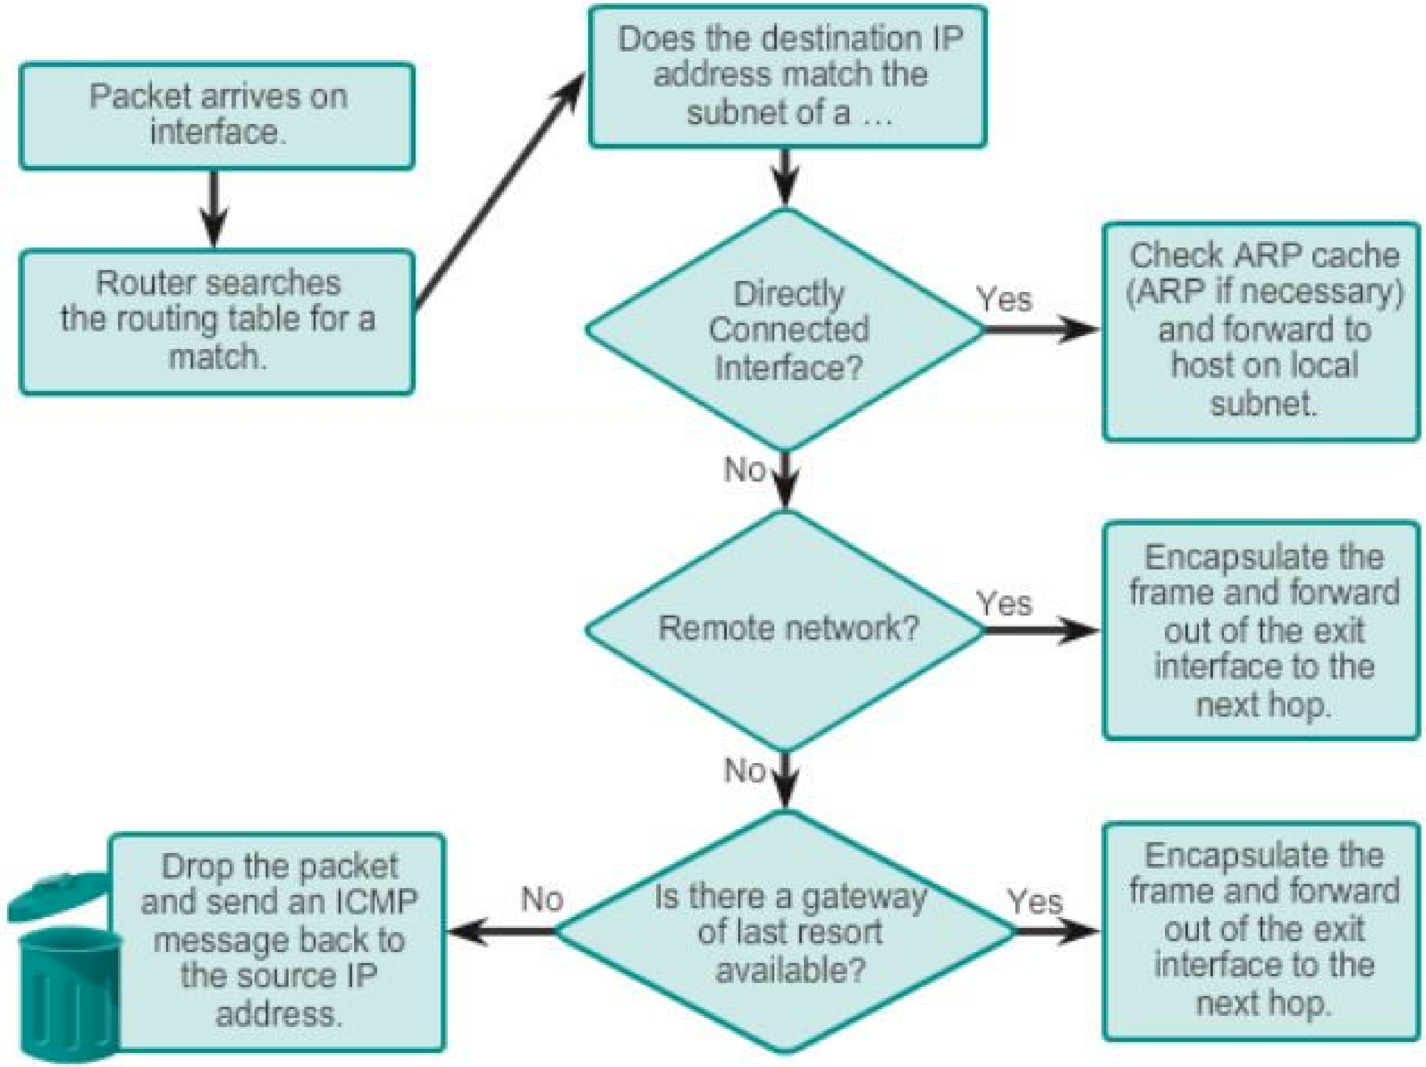

Packet arrives on interface.

Router searches the routing table for a match.

Does the destination IP address match the subnet of a …

Directly Connected Interface?

Yes

Check ARP cache (ARP if necessary) and forward to host on local subnet.

No

Remote network?

Yes

Encapsulate the frame and forward out of the exit interface to the next hop.

No

Is there a gateway of last resort available?

Yes

Encapsulate the frame and forward out of the exit interface to the next hop.

No

Drop the packet and send an ICMP message back to the source IP address.

IP ROUTING

Router Path selection

- Routers select best routes and build the routing table based on the following criteria:
  - Administrative distance
  - Metric
  - Longest prefix match

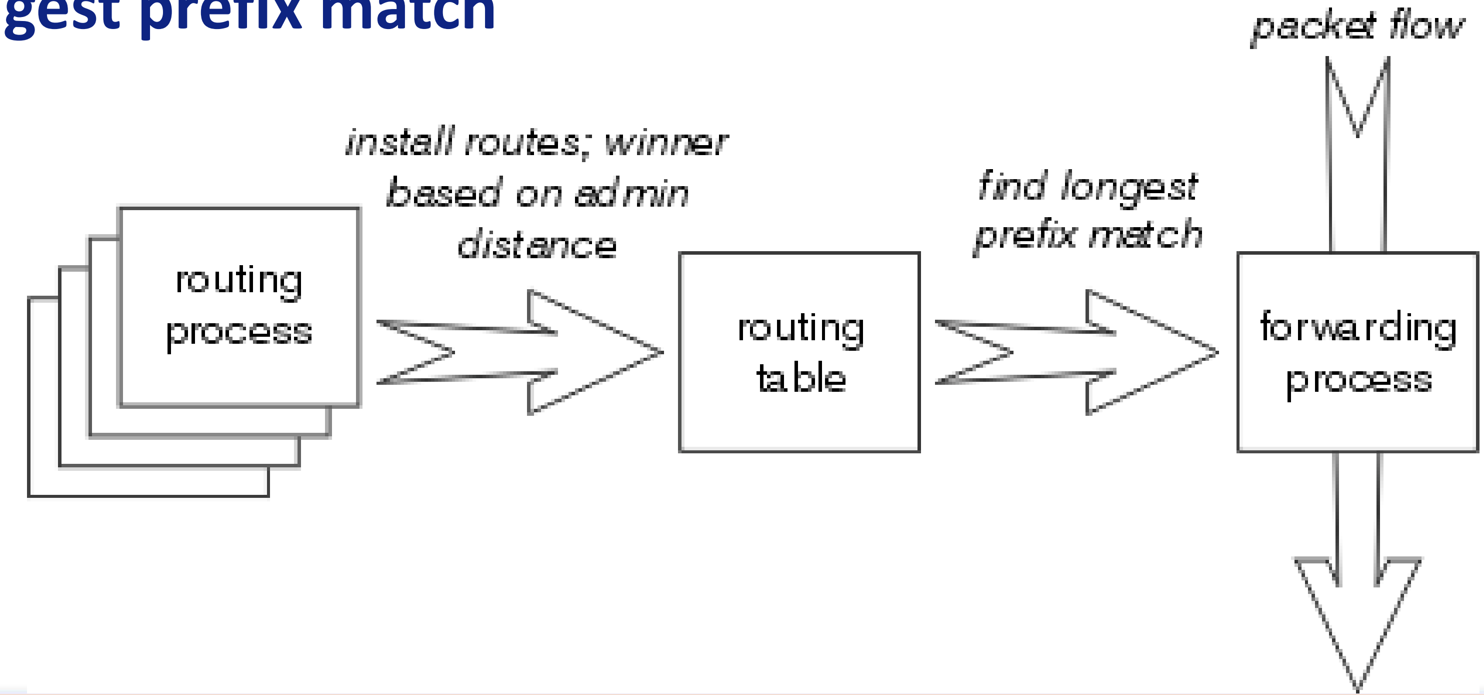
packet flow
install routes; winner based on admin distance
find longest prefix match
routing process
routing table
forwarding process

IP ROUTING

Longest Prefix Match

- Search for the routing table entry that has the longest prefix match with the destination IP address. WHY?
  - The longer the prefix the closer you are to destination….

1. Search for a match on all 32 bits
2. Search for a match for 31 bits

   ……..

   ……..

32. Search for a match on 0 bits

- Host route

  → 32-bit prefix match
- Default route is represented as 0.0.0.0/0

  → 0-bit prefix match

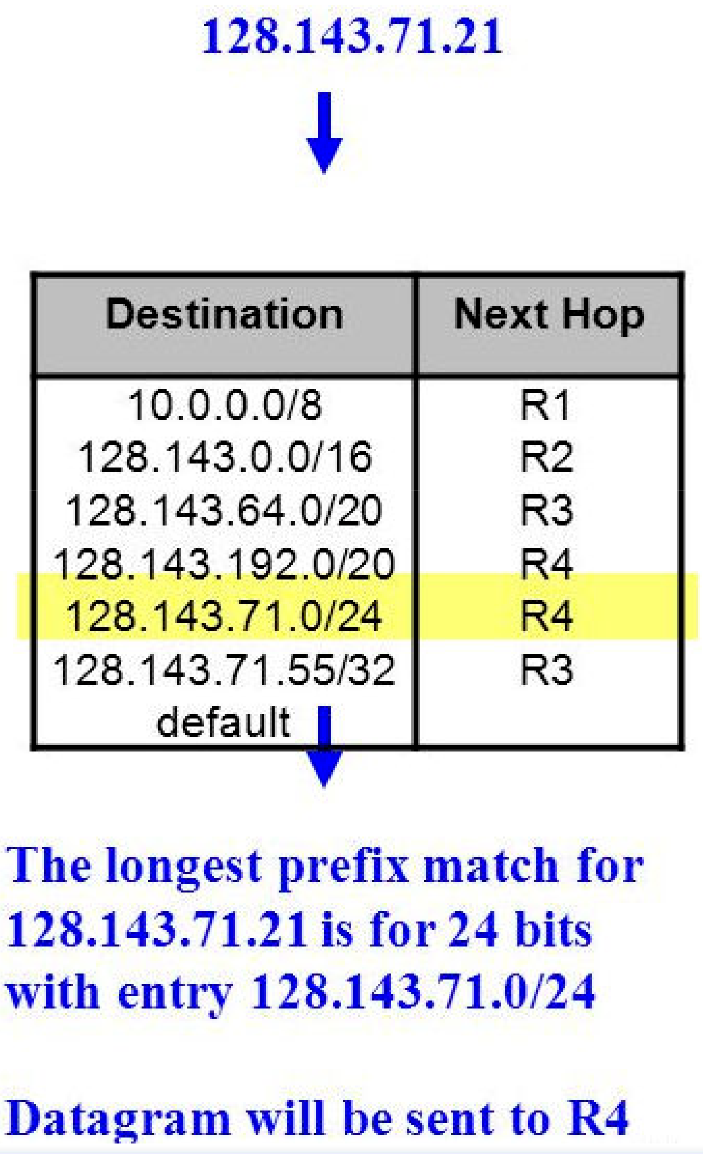

128.143.71.21

Destination	Next Hop
10.0.0.0/8	R1
128.143.0.0/16	R2
128.143.64.0/20	R3
128.143.192.0/20	R4
128.143.71.0/24	R4
128.143.71.55/32	R3
default	

The longest prefix match for 128.143.71.21 is for 24 bits with entry 128.143.71.0/24

Datagram will be sent to R4

IP ROUTING

Metrics

- Multiple path to same destination
- Best path is selected by the routing protocol, based on a specific value (metric)
- Each protocol uses its own metrics to build and update routing tables
- Metric is used to measure the distance to the destination network
- Lowest metric = best path, placed in routing table
- The following are metrics used in determining the best path:
  - Hop cont
  - Bandwidth
  - Load
  - Delay
  - Reliability
  - Cost

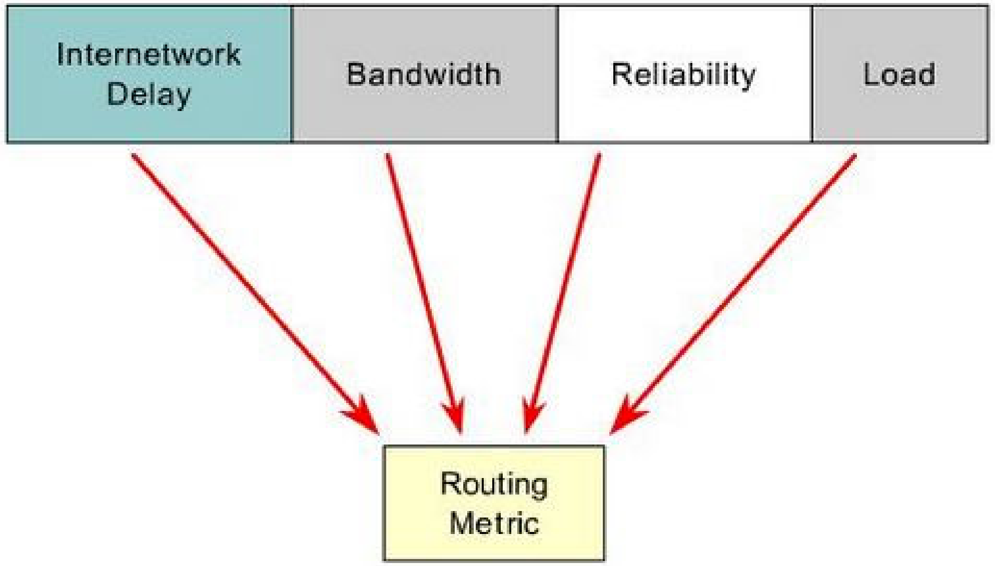

IP ROUTING

Administrative distance (AD)

- This is the measure of trustworthiness of the source of the route.
- Routes are chosen and built in the routing table based on the routing protocol's administrative distance. The routes learned from the routing protocol with the lowest administrative distance are installed in the routing table.

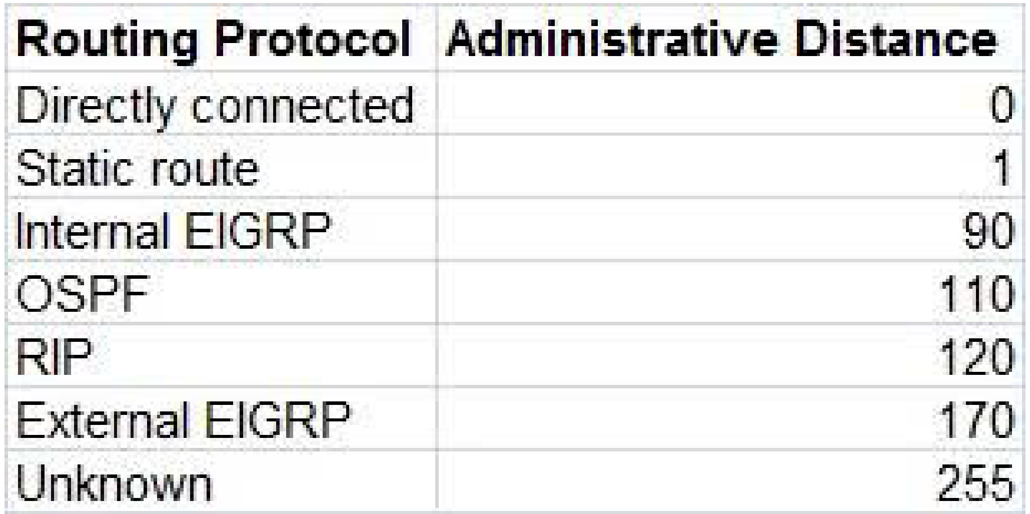

Routing Protocol	Administrative Distance
Directly connected	0
Static route	1
Internal EIGRP	90
OSPF	110
RIP	120
External EIGRP	170
Unknown	255

Static Routing

Static Routing

Static Route

- By adding static routes, a router can learn a route to a remote network that is not directly connected to one of its interfaces.
- A static route is created, maintained, and updated by a network administrator, manually.
- A static route to every network must be configured on every router for full connectivity.
- Administrative distance (AD) value of Static Routing is 1 that means it is most reliable protocol available for routing after Directly Connected networks (AD=0).

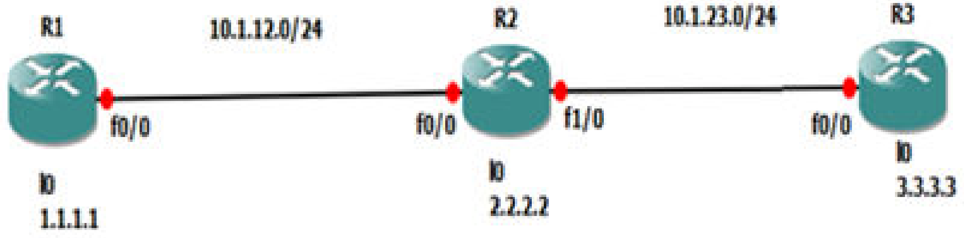

```
R1#show run | sec route
ip route 2.2.2.2 255.255.255.255 10.1.12.2
ip route 3.3.3.3 255.255.255.255 10.1.12.2
ip route 10.1.23.0 255.255.255.0 10.1.12.2
R1#
```

```
R2#show run | se route
ip route 1.1.1.1 255.255.255.255 10.1.12.1
ip route 3.3.3.3 255.255.255.255 10.1.23.3
R2#
```

```
R3#show run | sec route
ip route 1.1.1.1 255.255.255.255 10.1.23.2
ip route 2.2.2.2 255.255.255.255 10.1.23.2
ip route 10.1.12.0 255.255.255.0 10.1.23.2
R3#
```

Static Routing

Why Use Static Routing?

Advantages :

- Static routes are not advertised over the network, resulting in better security.
- Static routes use less bandwidth than dynamic routing protocols, no CPU cycles are used to calculate and communicate routes.
- The path a static route uses to send data is known

Disadvantages:

- Initial configuration and maintenance is time-consuming.
- Configuration is error-prone, especially in large networks.
- Administrator intervention is required to maintain changing route information.
- Does not scale well with growing networks; maintenance becomes cumbersome.
- Requires complete knowledge of the whole network for proper implementation.

Static Routing

Static Routing

Static routing has three primary uses:

- **Small networks:**
  Providing ease of routing table maintenance in smaller networks that are not expected to grow significantly.

- **Default route:**
  Using a single default route to represent a path to any network that does not have a more specific match with another route in the routing table. Default routes are used to send traffic to any destination beyond the next upstream router.

- **Routing to and from stub networks.**
  A stub network is a network accessed by a single route, and the router has no other neighbors (only one router).

Static Routing

Using static routing in stub networks

- For an example, here we see that any network attached to R1 would only have one way to reach other destinations, whether to networks attached to R2 or to destinations beyond R2.
- Therefore, network 172.16.3.0 is a stub network and R1 is a stub router.
- Running a routing protocol between R1 and R2 is a waste of resources

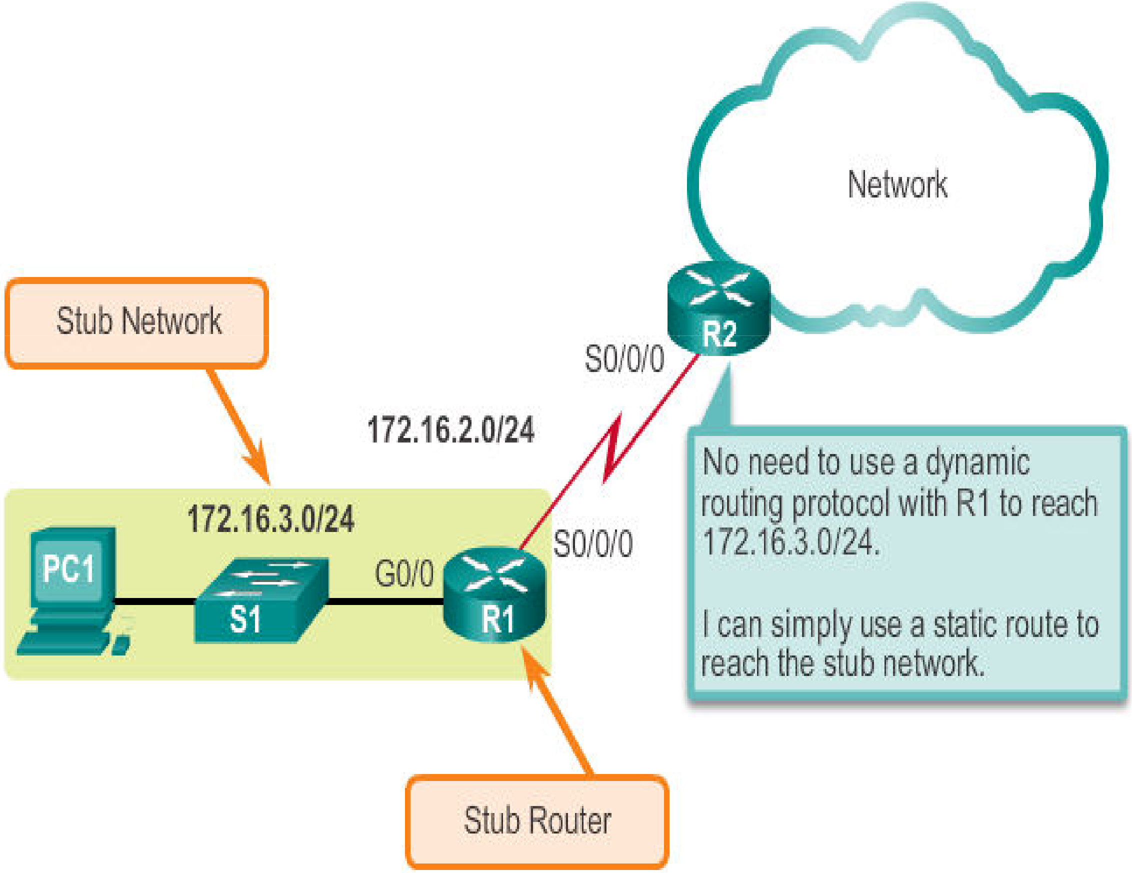

No need to use a dynamic routing protocol with R1 to reach 172.16.3.0/24.

I can simply use a static route to reach the stub network.

Static Routing

Static Route Configuration

- IP route command:
  - To configure a static route use the following command: ip route
  - Parameters:

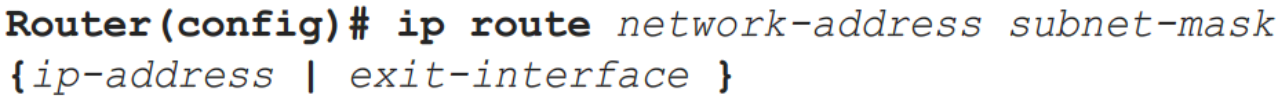

```
Router(config)# ip route network-address subnet-mask
{ip-address | exit-interface }
```

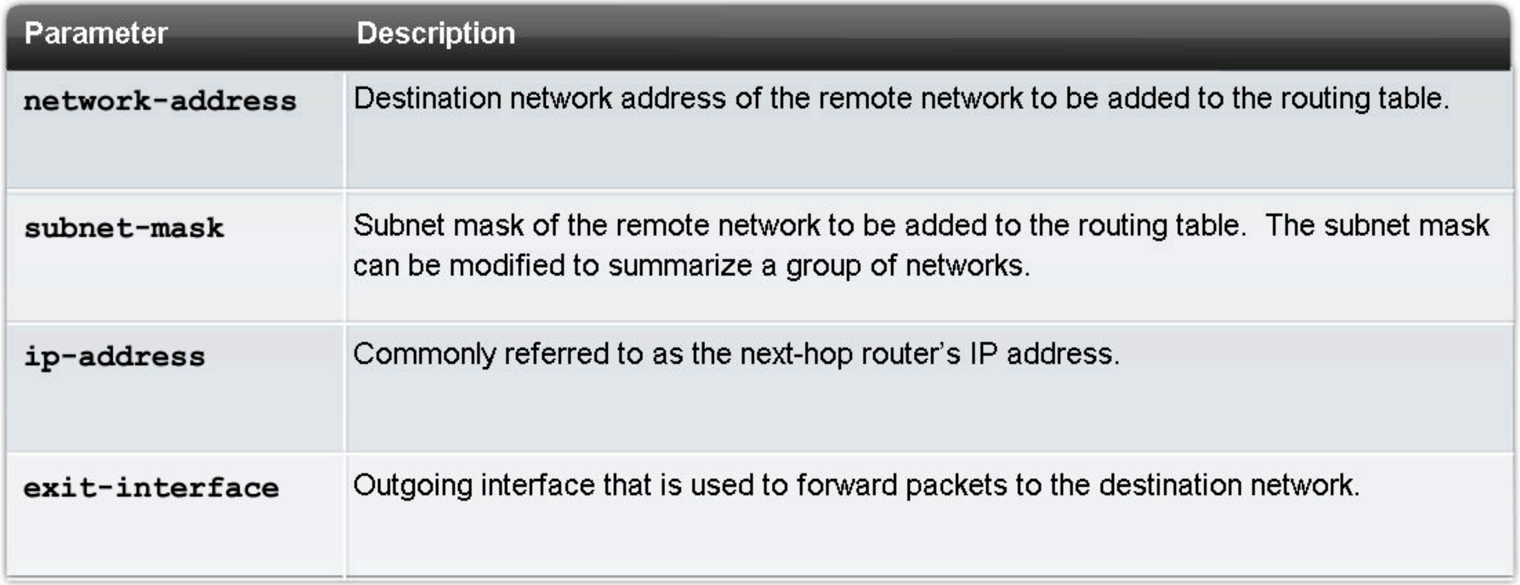

Parameter	Description
`network-address`	Destination network address of the remote network to be added to the routing table.
`subnet-mask`	Subnet mask of the remote network to be added to the routing table. The subnet mask can be modified to summarize a group of networks.
`ip-address`	Commonly referred to as the next-hop router's IP address.
`exit-interface`	Outgoing interface that is used to forward packets to the destination network.

Static Routing

Static Route Configuration - Next-Hop Options

- Configuring a Static route with an Exit Interface
  - A static route that forwards all packets to the next-hop IP address goes through the following process (reclusive route lookup)
    - The router first must match static route’s destination IP address with the Next hop address
    - The next hop address is then matched to an exit interface
  - For point-to-point interfaces, you can use static routes that point to the exit interface or to the next-hop address
- Configuring a Static route to the next-hop IP address
  - It is more efficient because the routing table can resolve the exit interface in a single search instead of 2 searches
  - For multipoint/broadcast interfaces, it is more suitable to use static routes that point to a next-hop address.

Static Routing

Modifying Static routes

- Existing static routes cannot be modified. The old static route must be deleted by placing no in front of the ip route
- Example:

  router# no ip route 192.168.2.0 255.255.255.0 172.16.2.2
- A new static route must be rewritten in the configuration

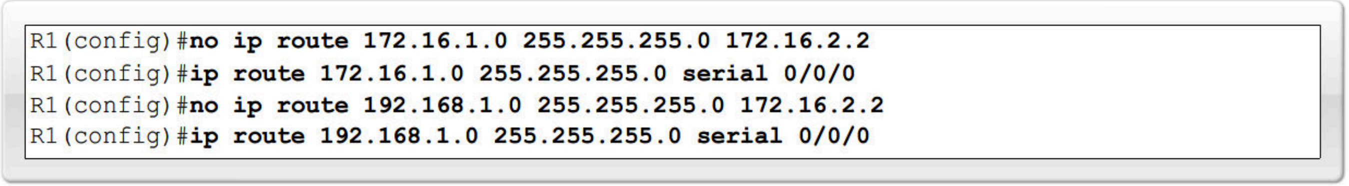

```
R1(config)#no ip route 172.16.1.0 255.255.255.0 172.16.2.2
R1(config)#ip route 172.16.1.0 255.255.255.0 serial 0/0/0
R1(config)#no ip route 192.168.1.0 255.255.255.0 172.16.2.2
R1(config)#ip route 192.168.1.0 255.255.255.0 serial 0/0/0
```

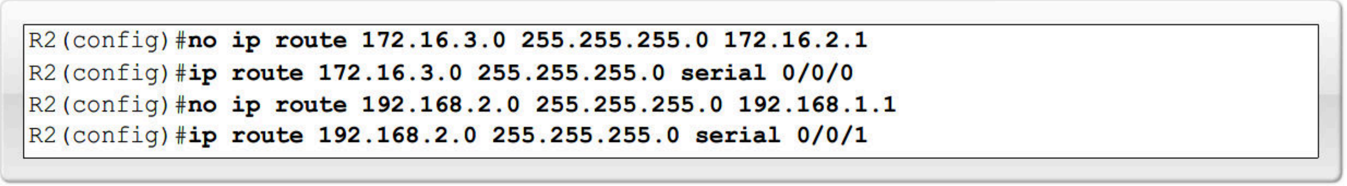

```
R2(config)#no ip route 172.16.3.0 255.255.255.0 172.16.2.1
R2(config)#ip route 172.16.3.0 255.255.255.0 serial 0/0/0
R2(config)#no ip route 192.168.2.0 255.255.255.0 192.168.1.1
R2(config)#ip route 192.168.2.0 255.255.255.0 serial 0/0/1
```

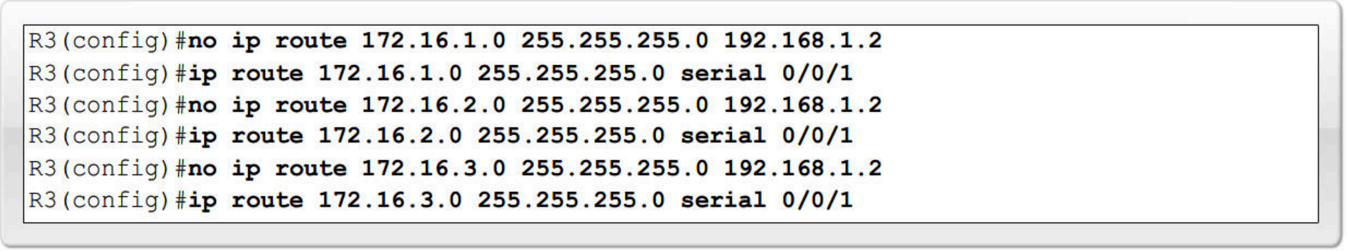

```
R3(config)#no ip route 172.16.1.0 255.255.255.0 192.168.1.2
R3(config)#ip route 172.16.1.0 255.255.255.0 serial 0/0/1
R3(config)#no ip route 172.16.2.0 255.255.255.0 192.168.1.2
R3(config)#ip route 172.16.2.0 255.255.255.0 serial 0/0/1
R3(config)#no ip route 172.16.3.0 255.255.255.0 192.168.1.2
R3(config)#ip route 172.16.3.0 255.255.255.0 serial 0/0/1
```

Static Routing

Default Static Route

- Static route can be used to configure default route.
- A default static route is a route that matches all packets.
- A default route identifies the gateway IP address to which the router sends all IP packets that it does not have a learned or static route.
- A default static route is simply a static route with 0.0.0.0/0 as the destination IPv4 address.
- Configuring a default static route creates a Gateway of Last Resort.

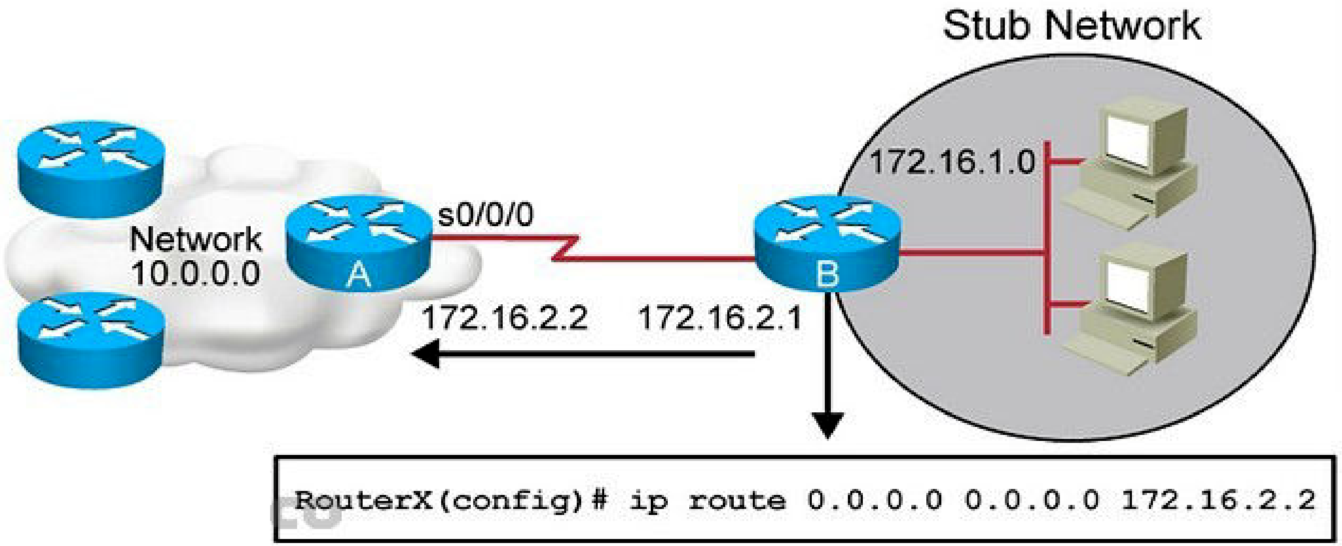

```
RouterX(config)# ip route 0.0.0.0 0.0.0.0 172.16.2.2
```

Static Routing

Floating Static Route

- Floating static routes are static routes that are used to provide a backup path to a primary static or dynamic route, in the event of a link failure.
- The floating static route is only used when the primary route is notavailable.
- To accomplish this, the floating static route is configured with a higher administrative distance than the primary route.

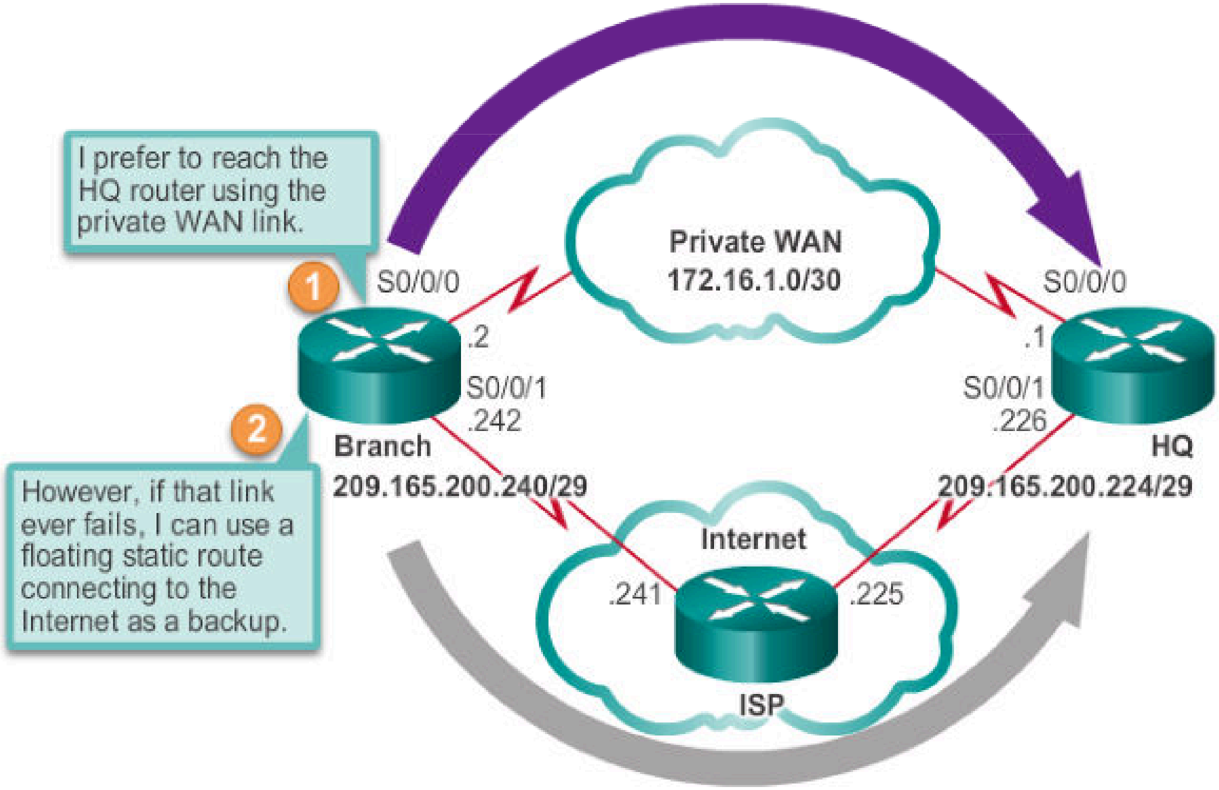

1 I prefer to reach the HQ router using the private WAN link.

2 However, if that link ever fails, I can use a floating static route connecting to the Internet as a backup.

Static Routing

Static Route Troubleshooting

- Troubleshooting a Missing Route
- Tools that can be used to isolate routing problems include:
  - Ping – tests end to end connectivity
  - Traceroute – used to discover all of the hops (routers) along the path between 2 points
  - Show IP route– used to display routing table & certain forwarding process
  - Show ip route static
  - Show ip route *network*
  - Show ip interface brief- used to show status of router interfaces
  - Show cdp neighbors detail– used to gather configuration information about directly connected neighbors