Naming and Addressing

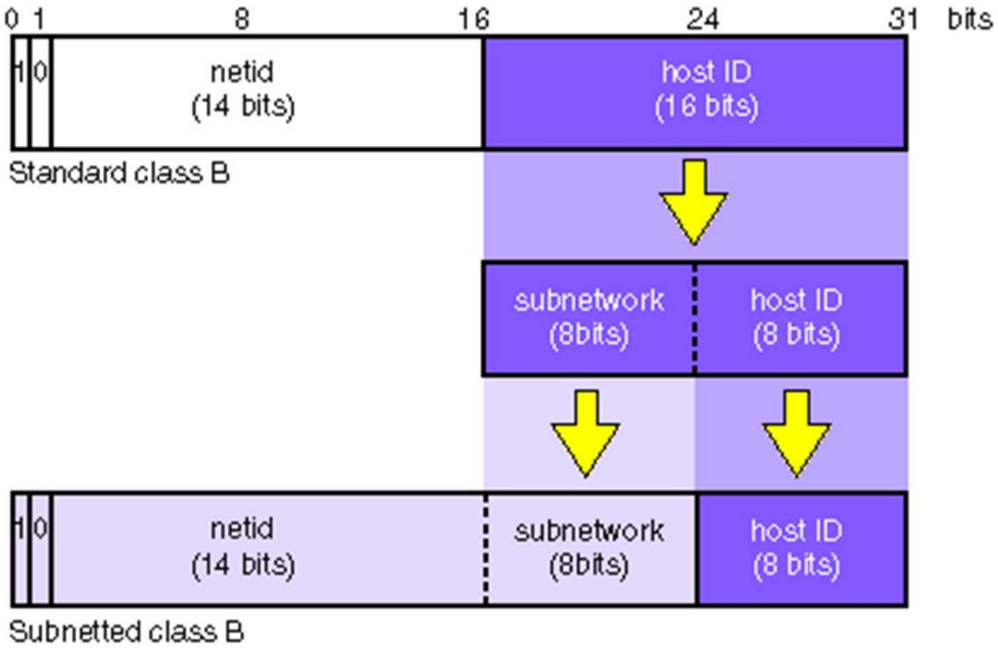

Contents

Review of Web

Naming

Addressing

Performance

Review of Web

What Happens?

Click on a link (http://foo.bar.com/xx)

Conversion from name to address

Open connection to remote machine

Pass arguments to process

Retrieve contents from server

Display locally

Review of Web Requirements

Name mapping service (DNS)

Addressing/routing (IP)

Reliable delivery (TCP)

Representation of content (HTTP)

Local display (application)

Naming

- Naming Computers
  - Flat (not divisible into components)
  - Hierarchical
- URLs

Naming Computers

- Need a way to locate services; easier for humans than numbers
- Flat Name Space:
  - Every computer has unstructured name
  - Must coordinate not to stomp on each other
  - Examples:, *portal*, *ie1*, *heuristic*, *ucbvax*
  - Didn’t scale very well

Hierarchical Naming

- First real growth problem of Internet
  - Rule of thumb: things break if they grow 2 orders of magnitude (5-7 years in today's Internet)
  - Common Idea: **hierarchies scale well**
- Divide up scale into "domains"
  - Examples: EDU, COM, MIL, ORG, NET
  - (ISO3166-based): KR, FI, JP, DK, US, ..

Benefits of Naming Hierarchy

Much better scaling

Decentralized administration

Redundant databases

Recursive, can subdivide each subdivision

URLs: New Names

Relatively New Formats on Internet

- Popularized by web browsers
- *Format: proto://host-name:port/args…*
- http://www.cs.berkeley.edu/~kfall
- gopher://gopher.umsl.edu
- ftp://ftp.microsoft.com
- telnet://rainmaker.wunderground.com:3000

Addressing

- IP (v4) Addresses
- Expressing IP Address
- Address Classes
  - Examples
  - Subnets
- CIDR

IP (v4) Addresses

Every interface has at least one IP address

IP addresses are 32-bit numbers (4.3 billion of them!)

Divided into parts: (network prefix, host number)

Classical structure uses net/subnet/host partitioning where hosts on same subnet share net and subnet number

The *prefix* is the concatenation of the two: [net/subnet]

Expressing IPv4 Addresses

4 decimal numbers, called "dotted quad"

Each (decimal) number is one byte

Example: 143.248.92.27

Can generally be used in place of names

Originally, parts of pre-defined addresses "Classes"

IP Address Classes (before CIDR)

Class	32bit			Highest Address
A	0	Network	Host	127.255.255.255
B	1 0	Network	Host	191.255.255.255
C	1 1 0	Network	Host	223.255.255.255
D	1 1 1 0	Multicast Group ID		239.255.255.255

Special IP Addresses

32 bits		Use
all zero		This host on this network (during boot) Default route (in tables)
all one		Limited broadcast (when a host needs subnet mask/its IP address)
netid	all one	Net-directed broadcast to netid
0 1 1 1 1 1 1 1	not 0, often 1	Loopback

Addressing Operations

Example 1

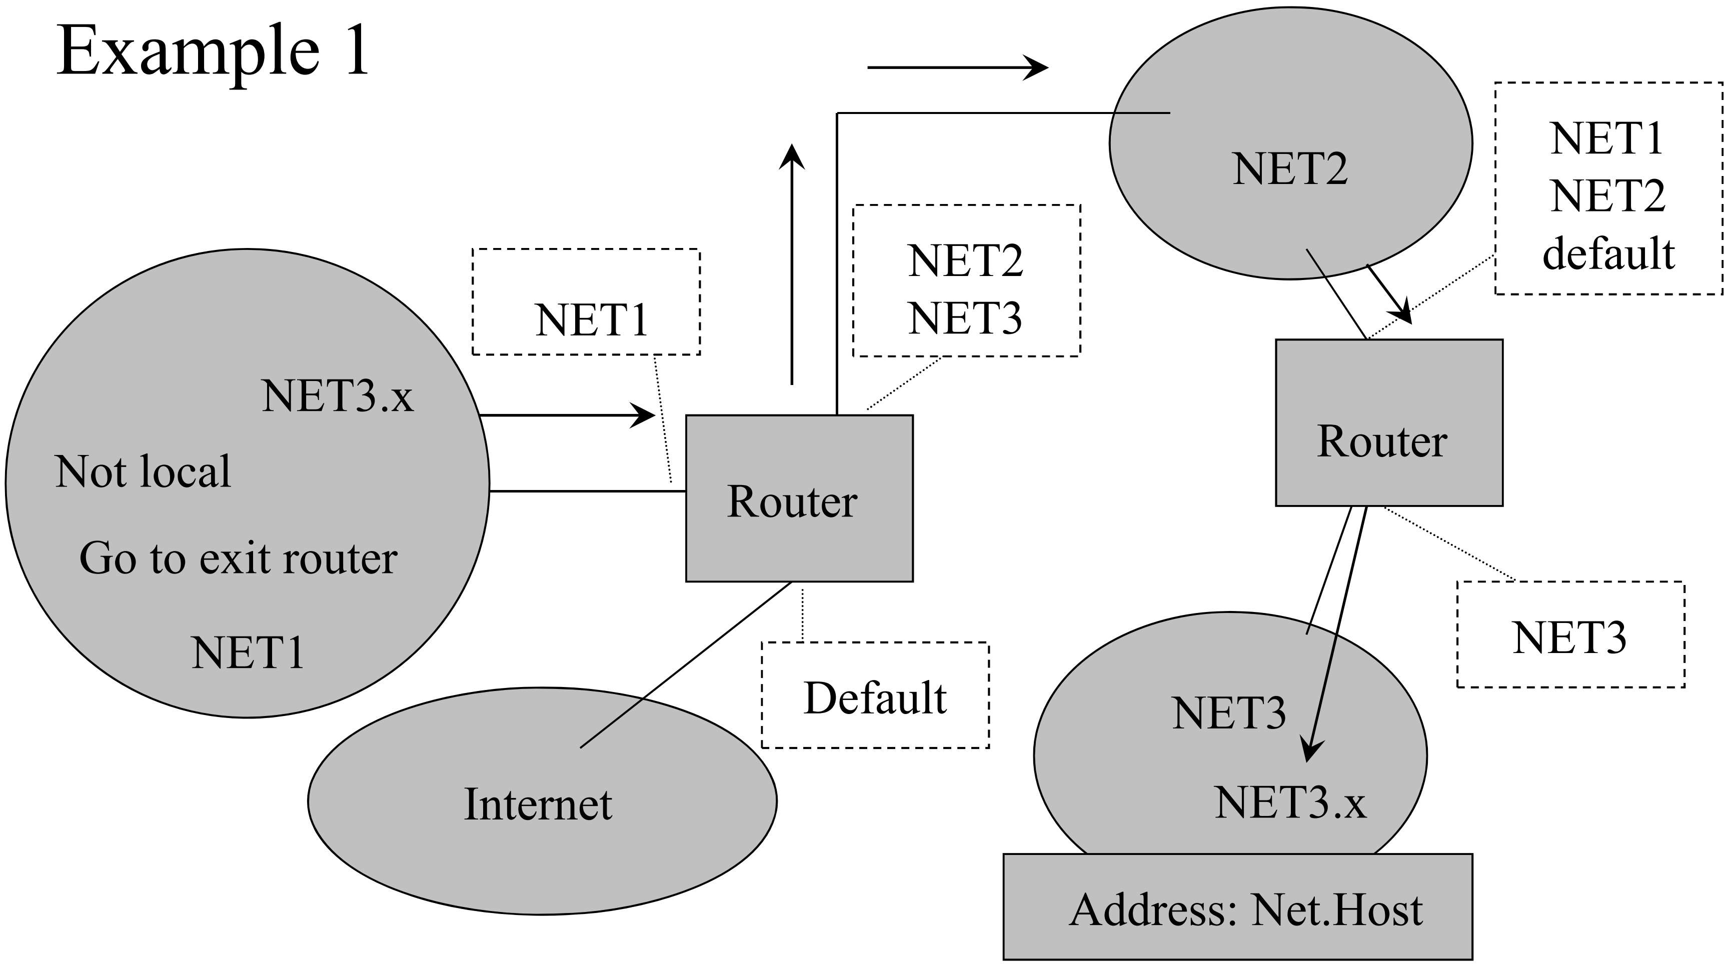

Addressing Operations

Example 2

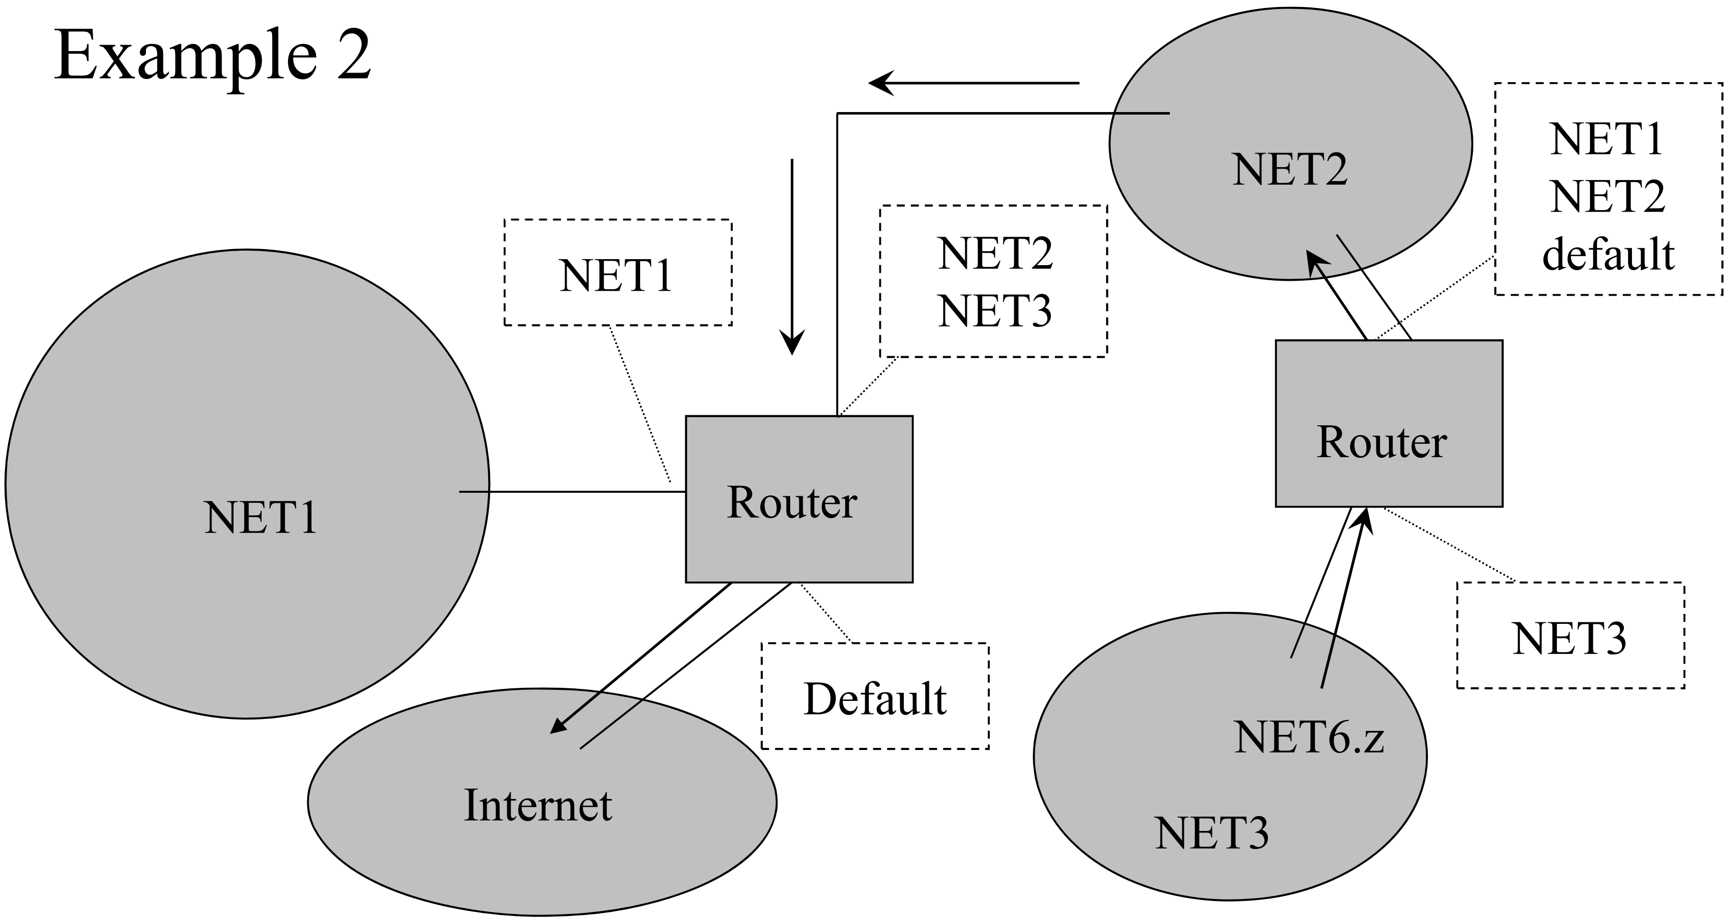

Example Assignments of Address

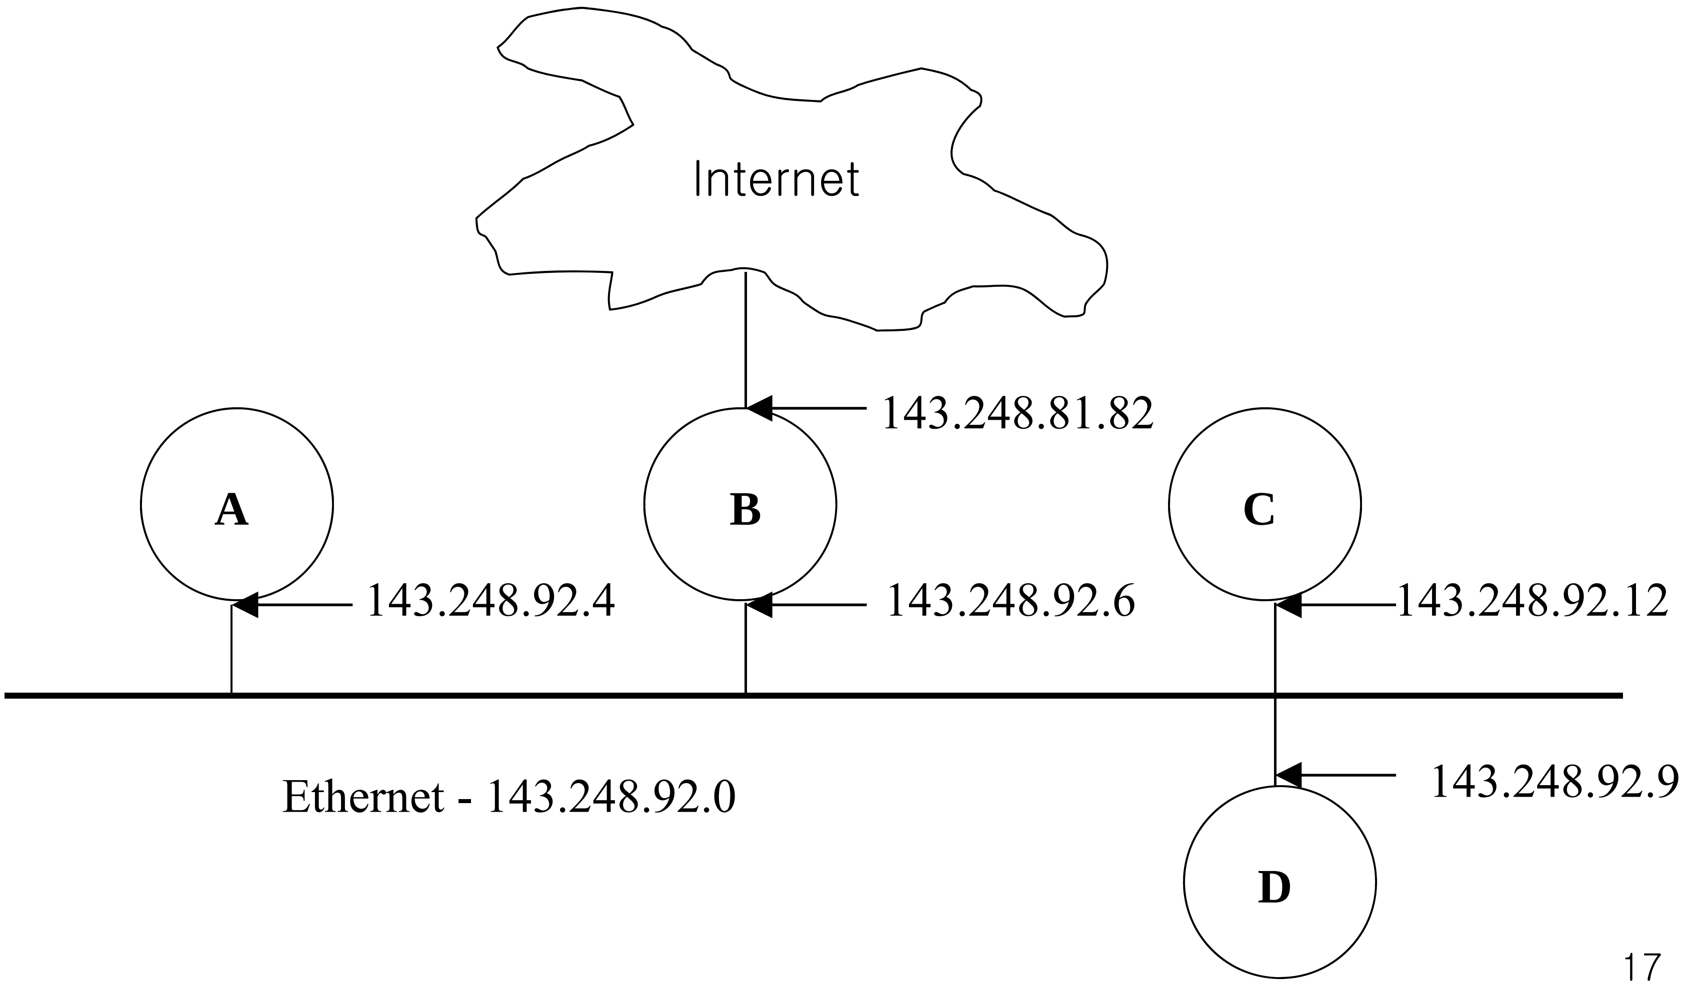

Subnet Addressing

- Somewhat historical, but terminology is consistent and still used
- Allows one site to have multiple *subnetworks* of their main network
  - Practical result: multiple segments
- Subnetting scheme is a **local** decision
- Requires a "Subnet mask"

Subnet Addressing

Example

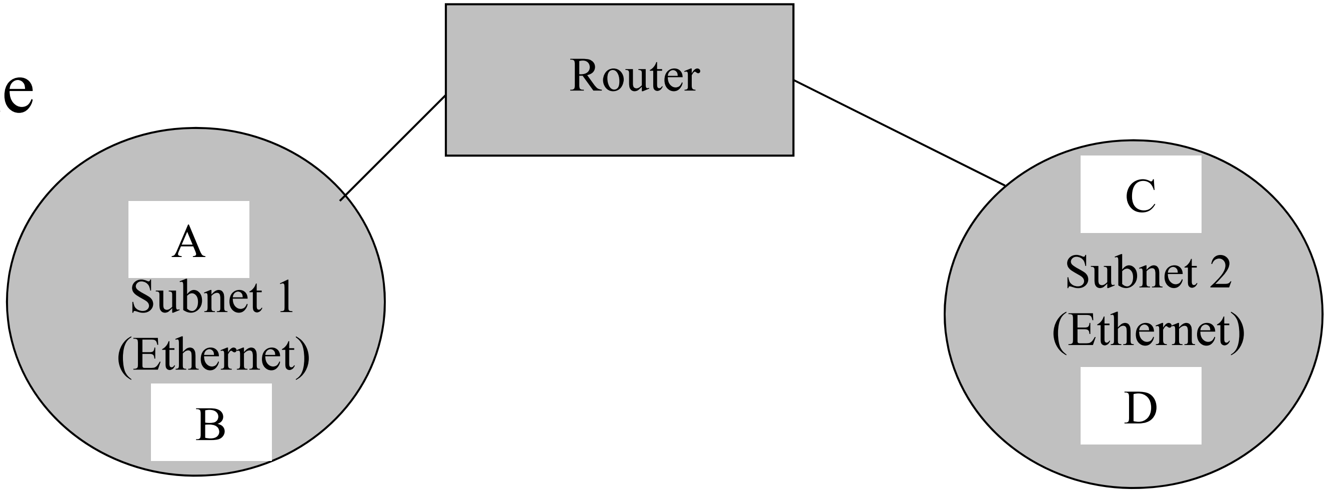

From A to B: Same Subnet (prefix) ☜ **Direct**

Ethernet packet from A to B

From A to C: Different Subnets ☜ **Via Router**

Ethernet packet 1 from A to Router

Ethernet packet 2 from Router to C

Subnet Addressing

Example

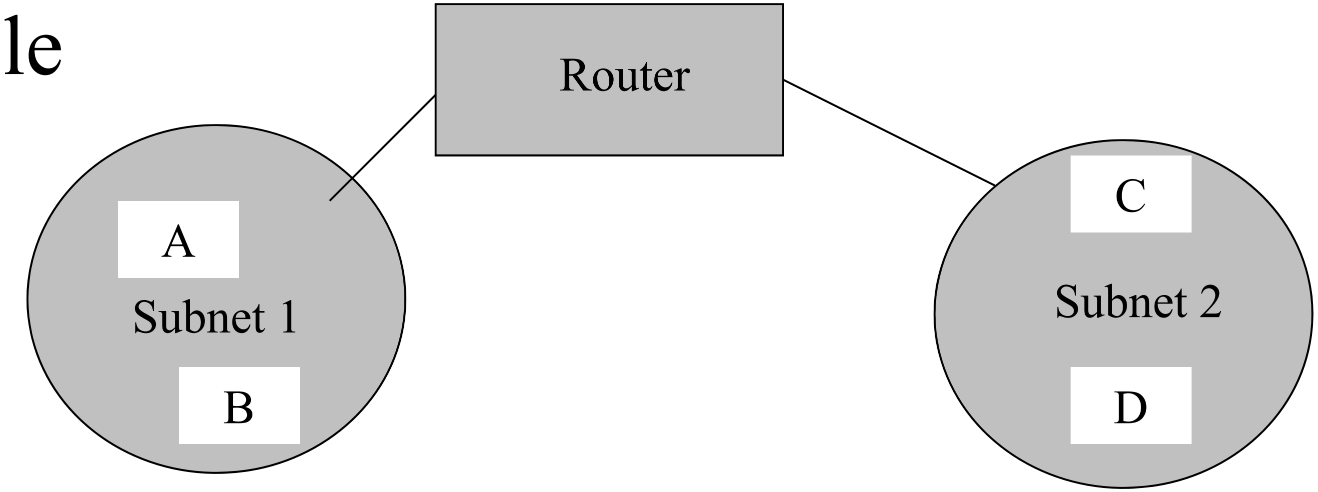

By looking at IP address (prefix) of B, A can tell whether B is on the same subnet as A or not!

(This determines whether to use a router or not)

Subnet Structure

Idea is to "steal" classful host bits and use them for numbering subnets

Rest of Internet only sees classes

A mask indicates which bits are [network/subnet] part, and which are host part

Subnet Example

143.248.92.27 is a “Class B” address

16 bits of network, 16 bits of host

Locally, want a 200 “Subnets”

So, need 8 bits to indicate subnet

Use a *subnet mask* of (16+8=24) bits

The networks utilize a Subnet Mask of:

1111 1111. 1111 1111. 1111 1111. 0000 0000

Converted to decimal: 255.255.255.0

Subnet Example

26 bit mask: 128.32.25.12/26 is:

10000000 00100000	00011001 00	001100
11111111 11111111	11111111 11	000000
NET	SUBNET	HOST

Subnet 100 of net 128.32, host 12

Subnet Partitioning

128.32.0.0/26 gives 2^(26-16) = 1024 Subnet of 2^(32-26)-2=62 hosts each

First usable address: 128.32.0.1 (see RFC1812, page 48)

Last usable address: 128.32.255.254

Any address with all "1" bits in host part is a (subnet) broadcast

Subnet Partitioning

128.32.25.12/26 is:

10000000 00100000	00011001 00	001100

128.32.0.65/26 is:

10000000 00100000	00000000 01	000001

128.32.255.190/26 is:

10000000 00100000	11111111 10	111110

Common Subnet?

Are 128.32.25.12 and 128.32.25.85 on the same subnet using a /26 mask?

128.32.25.12 is:

10000000 00100000 00011001 00 | 001100

128.32.25.85 is:

10000000 00100000 00011001 01 | 010101

Prefixes differ, so not on same subnet (need router to reach)

Observe: these would be on the same subnet using a mask shorter than /26!

Subnet Masking

The target router will implement a mathematical function that performs a logical binary AND function between the stored Subnet Mask and a particular received datagram IP address to determine whether the datagram can be delivered on the same subnet, or whether it must go through an IP router to another subnet.

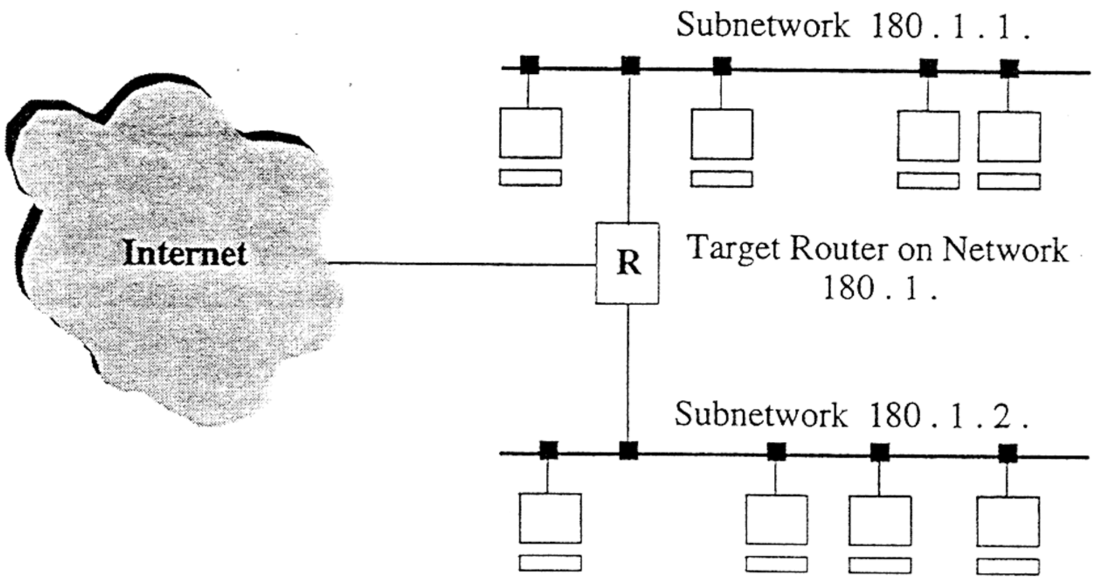

Subnet Masking Process

Assume the destination address is D and a router's is M.

Assume that both D and M contain a netid, subnetid, and hostid.

Assume the Subnetwork Mask is [255.255.255.0].

Two AND calculations are made and the answers compared:

[255.255.255.0] AND [D]
[255.255.255.0] AND [M]

1 AND 1 = 1
1 AND 0 = 0
0 AND 1 = 0
0 AND 0 = 0

Any binary number ANDed with 0 will be 0. Therefore, the hostid is "stripped off" leaving only the netid and subnetid.

Subnet Masking Process

The two results are compared for a routing decision.

If they are identical, Addresses M and D are on the same network.

If the answers are not equal, the two addresses are on different subnetworks, and the datagram must go to another router for delivery.

Subnet Example

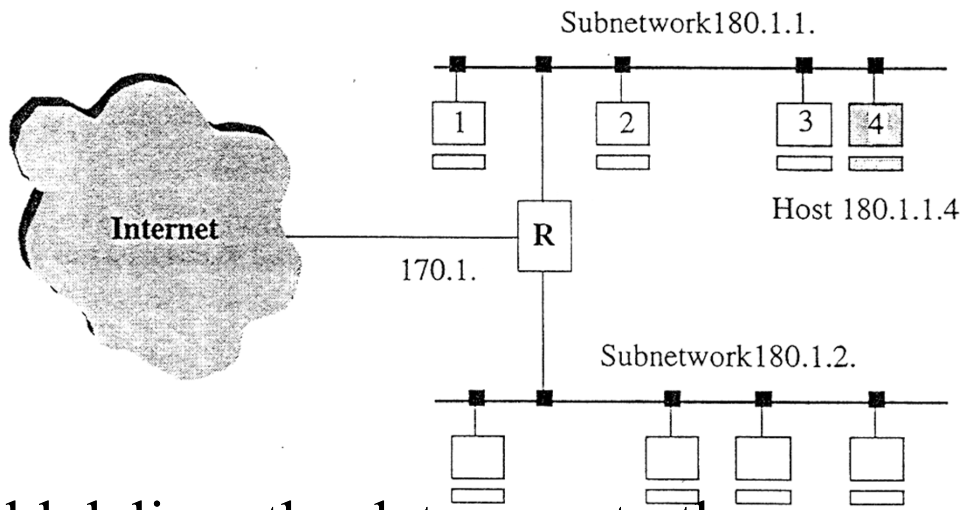

The first 16 bits (the netid) would deliver the datagram to the access point for the network (the target router).

Using the Mask, the router would then decide which of the 256 subnetworks for this datagram.

```
       Mask  = 11111111.11111111.11111111.00000000 [255.255.255.0]
Pkt Address  = 10110100.00000001.00000001.00000100 [180.1.1.4]
   AND calc  = 10110100.00000001.00000001.00000000 [180.1.1.0]
```

Classless Inter-domain Routing (CIDR)

About 1993, remove strict classes from architecture

To solve two of these: router memory exhaustion and a shortage of Class B addresses

Generalized notation of “network prefix”

Requires “longest prefix” match routing

Subsumes and generalizes subnetting

Classless Inter-domain Routing (CIDR)

About 1993, remove strict classes from architecture

To solve two of these: router memory exhaustion and a shortage of Class B addresses

Generalized notation of "network prefix"

Requires "longest prefix" match routing

Subsumes and generalizes subnetting

Classless Inter-domain Routing (CIDR)

Minimizes the number of routes that a router needs to store and increases the address efficiency

CIDR aggregates routes such that a single entry in a forwarding table reach a lot of different networks

A block of contiguous class C addresses are aggregated to a single network prefix

Classless Inter-domain Routing (CIDR)

192.4.16: 11000000 00000100 00010000

192.4.31: 11000000 00000100 00011111

First 20bits are the same:

20-bit prefix for all the networks 192.4.16 – 192.4.31 is represented as 192.4.16/20

Classless Inter-domain Routing (CIDR)

Similar to subnetting, but bits in CIDR are contiguous: bits in subnet mask may not

CIDR collapses multiple addresses onto one: subnet shares one address among multiple physical networks

Measuring Performance (Capacity and Utilization)

- Capacity
  - The rate (bits/second) of a communication channel
  - Typically fixed by oscillator rate, noise, coding, bandwidth
- Utilization
  - The fraction of capacity in actual use measured over some interval of time

Measuring Performance (Throughput and Delay)

Throughput

- The data rate (bits/second) available to a particular application over an interval of time

Delay/Latency (one-way)

- The time required to send a minimum-sized data unit from sender to receiver

Measuring Performance (Derived Values)

- Jitter: variability in delay
- Round-Trip Time (RTT)
  - Two-way delay from sender to receiver and back
- Bandwidth-Delay Product
  - Product of bandwidth and delay, indicates "storage" capacity of network

A Note on Units

"Mega" is sometimes 1,000,000 and sometimes 1,048,576 (similar for kilo-units)

1 MB is 8 times larger than 1 Mb

Data rates are often powers of 10 (100 Mb/s Ethernet), whereas messages are often powers of 2 (a 1KB message is 1024 bytes)

Transmission Time

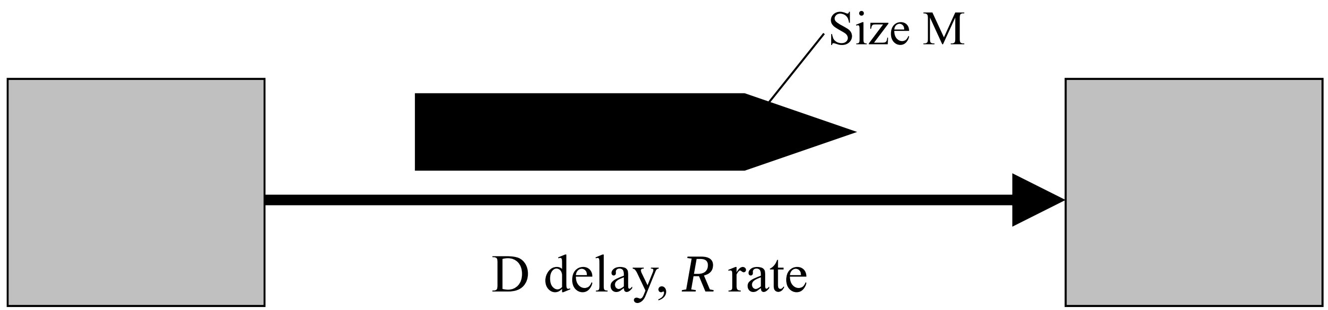

Tx delay = (M bits)/(R b/s) = M/R sec

Prop delay = D sec = Length/Prop_speed

Total Tx Time = D + M/R sec

Transmission Time

A Picture:

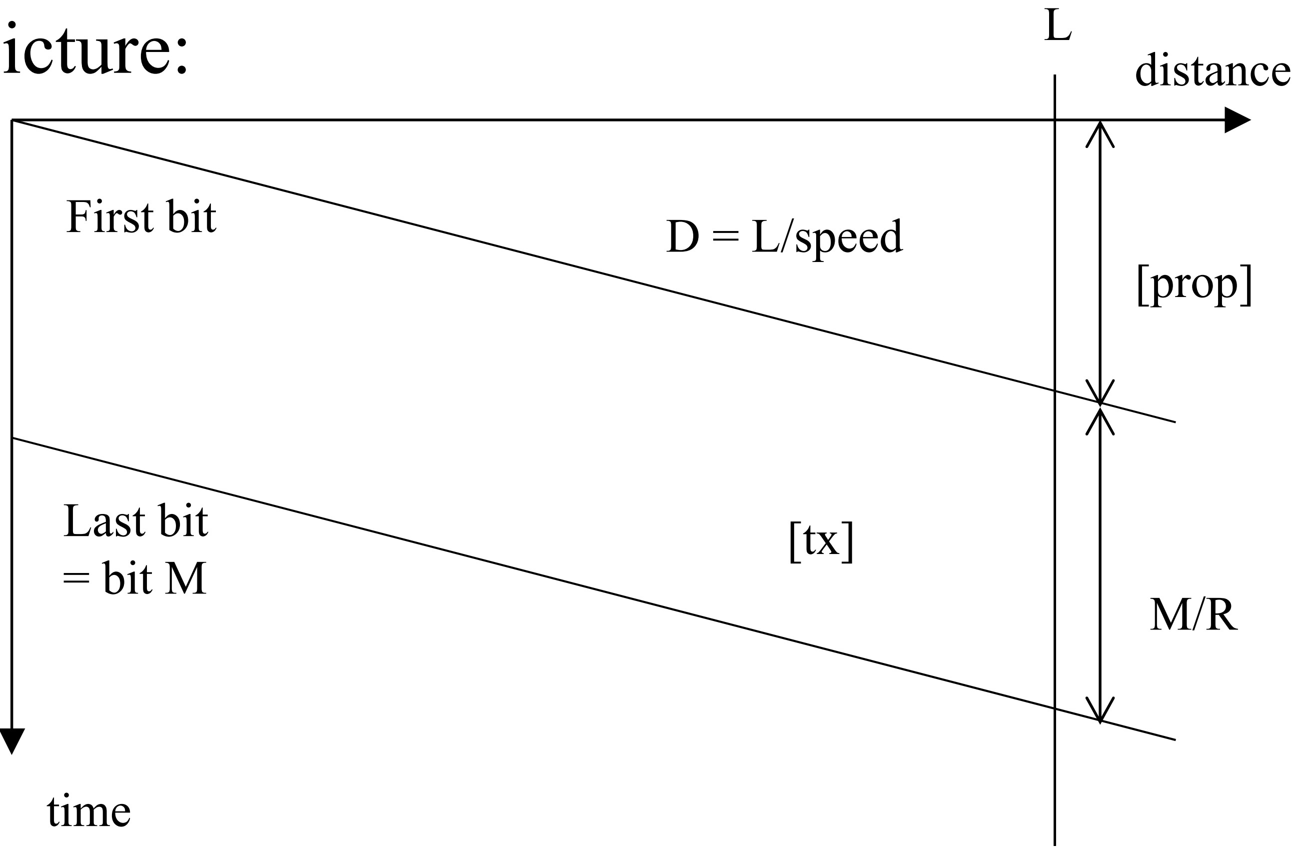

Each bit takes 1/R seconds to be sent

The signals propagate at some speed

Latency

Slower channels "Stretch out" bits in time:

- A bit on a 1Mb/s link is 1 μsec wide
- A bit on a 10Mb/s link is 0.1 μsec wide

Longer channels take a longer time to propagate

Low Speed Links

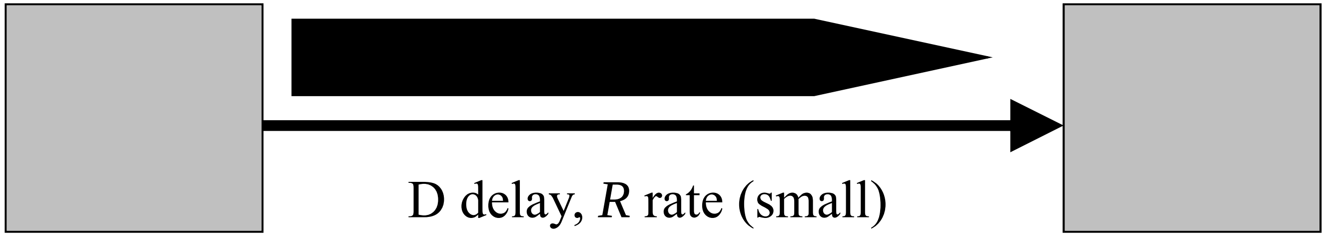

Small (slow) → large Tx Time (M/R)

Ex: Dialup (D = 10ms, R = 56Kb/s, M = 1KB)

Total Tx Time = .010 + ((1024 × 8)/(56 × 1024)) = 0.153 sec = 153 msec (1KB msg@56Kb/s)

(M/R dominates)

High Speed Link

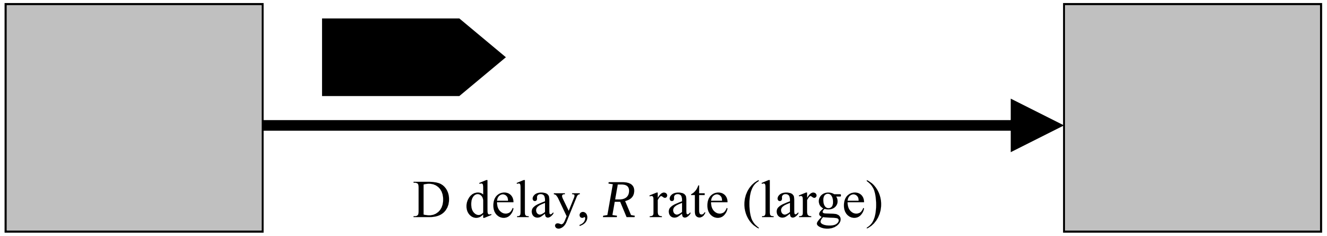

Large R (fast) → small Tx Time (M/R)

Ex: OC-3 (D = 10ms, R = 155Mb/s, M = 1KB)

Tx Time = .010 + ((1024 × 8)/(155 × 1024 × 1024)) = 0.01005 sec = 10.05 ms (D >> M/R)

(D dominates)

Queuing Delay

Total Latency = total tx time + queuing delay

transmit time = { last slides }

queuing delay = { depends! }

Example : If a packet arrives at the queue when there are already K bits waiting to be sent out at rate R, then

queuing delay = K/R

Statistically, one expects K to be some multiple of an average packet length. The multiple depends on the load and on the "bursty" nature of the traffic.

Total (one way) Latency

Propagation Delay (D) = *distance/speed-of-light*

Transmission delay = (M / R)

Queuing delay (Q) (using statistical multiplexing) depends on utilization

Total Latency = D + (M/R) +Q

Beware of Overheads

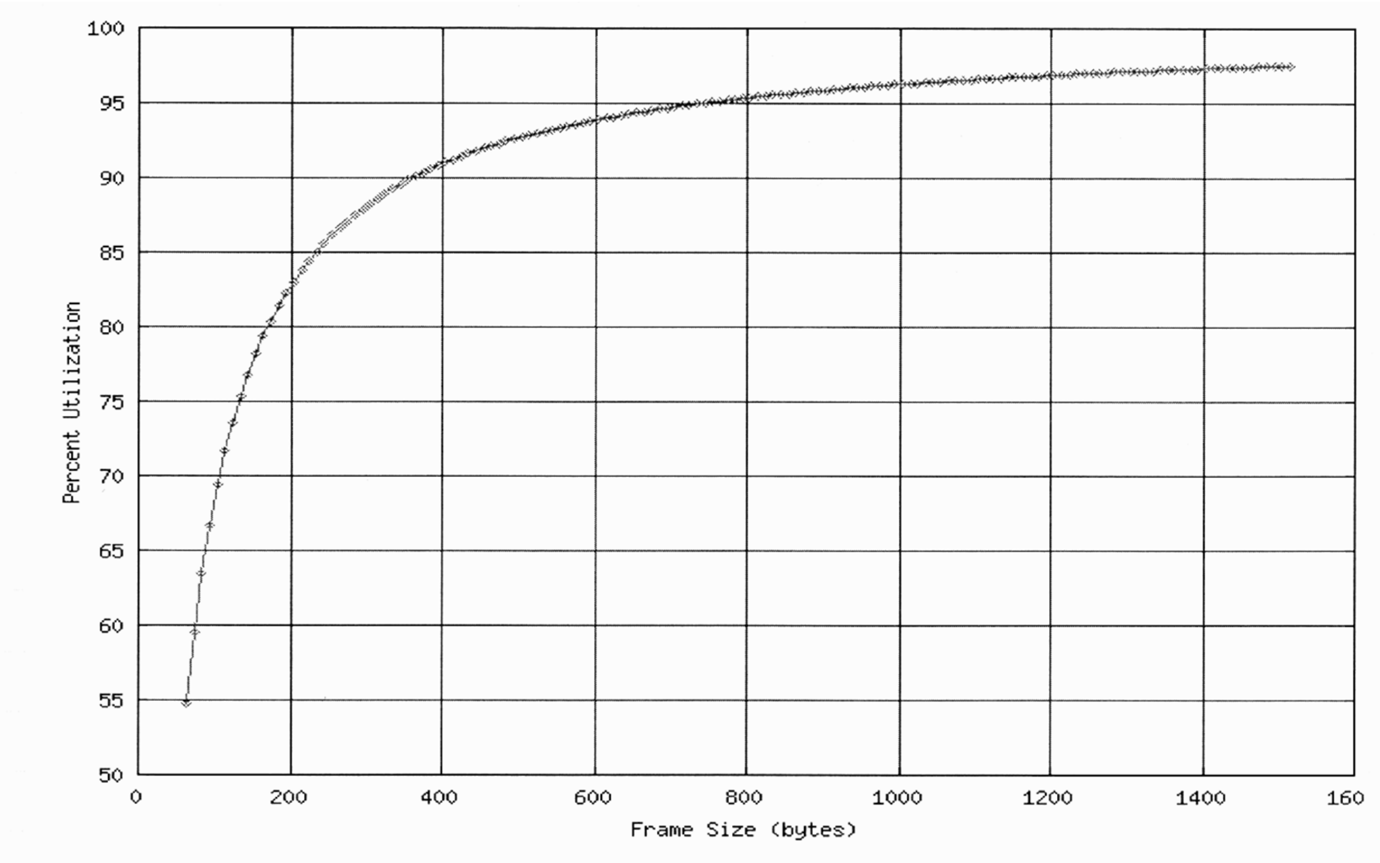

Measuring Latencies (1)

ie1% ping www.yahoo.com
    www.yahoo.com is alive

ie1% ping -s www.yahoo.com
    PING www.yahoo.com: 56 data bytes
    64 bytes from w4.scd.yahoo.com (66.218.71.83): icmp_seq=0. time=190. ms
    64 bytes from w4.scd.yahoo.com (66.218.71.83): icmp_seq=1. time=190. ms
    64 bytes from w4.scd.yahoo.com (66.218.71.83): icmp_seq=2. time=189. ms
    64 bytes from w4.scd.yahoo.com (66.218.71.83): icmp_seq=3. time=188. ms
    64 bytes from w4.scd.yahoo.com (66.218.71.83): icmp_seq=4. time=156. ms
    ^C

----www.yahoo.com PING Statistics----
    5 packets transmitted, 5 packets received, 0% packet loss
    round-trip (ms) min/avg/max = 156/182/190

Ping –s (Optional) Keyword to cause **ping** to send one datagram per second, printing one line of output for every response received. **Ping** command returns output only when a response is received.

Summary

Hierarchical naming:

heuristic.kaist.ac.kr (143.248.92.27)

IP addresses are 32-bit (4 bytes) numbers

Addressing at the router with the netid

A subnet mask indicates which bits are [network/subnet] part, and which are [host] part

143.248.92.27/24

Sunbet mask is used in routing to check common subnet