The Internet Protocol

- IP Addresses
- Address Resolution Protocol:
- IP datagram format and forwarding:
- IP fragmentation and reassembly

IP Addresses

- IP Addresses are 32 bit.
- Written in dotted decimal format: X.X.X.X
  - Example: A Clemson   address : 130.127.48.4

- An address encodes the identification of the network as well as the host (network id, host id)

- What is a host?

- How does an address relate to a host ?

- Three types of addresses?

IP Addresses

- IP Addresses are 32 bit.
- Written in dotted decimal format: X.X.X.X
  - Example: A Clemson address : 130.127.48.4

- An address encodes the identification of the network as well as the host (network id, host id)

- What is a host? A computer that can communicate with another computer over a TCP/IP network.

- How does an address relate to a host ?
  - Not one –to-one !!
  - Multihomed hosts have more than one interface AND more than one IP address
  - So slightly better answer is one IP address per interface although….
    - NAT totally changes this - Network Address Port Translation allows any number of Hosts to share the same IP address

- Three types of addresses?
  - Unicast
  - Broadcast
  - Multicast

Original Classful Addressing Scheme

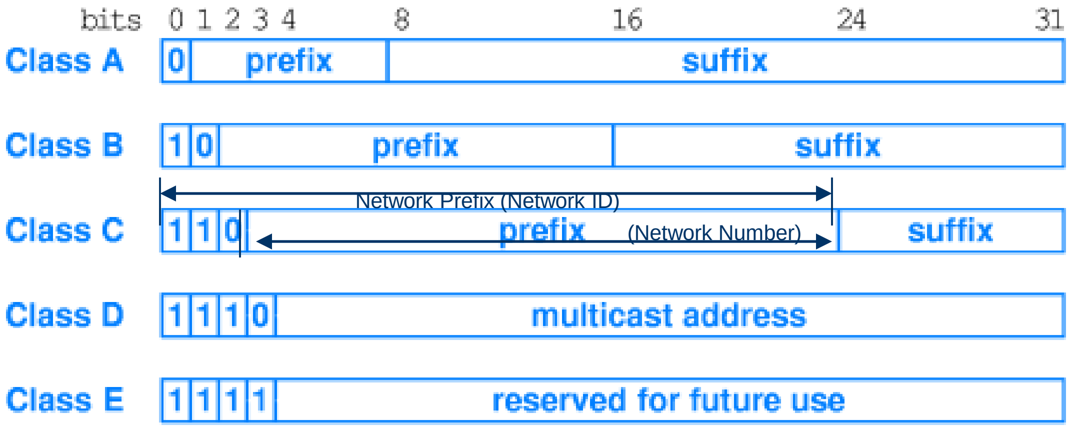

Original address scheme was classful:

- Class A for large networks (>64K hosts)
- Class B for medium networks (>256 hosts)
- Class C for small networks (<256 hosts)
- Class D for multicast
- Class E reserved

Original Classful Addressing Scheme

- What's the total address space?

- Example: Clemson address: 130.127.49.225
  - Class ?
  - Network ID?
  - Host ID?

- Example: 192.168.1.100

Broadcast Addresses: Two Types

- A network directed broadcast
  - contains a valid network id with an all 1's host id
  - class C example: 192.168.1.255
  - From a host on a different network (192.168.2.2), conceptually, if we issue 'ping -c 192.168.1.255' , ping should receive an echo response from all hosts on the 192.168.2 network.
    - There's not a router in the world that will allow this!!!!
- A local or limited broadcast does not require knowledge of the network address.
  - Referred to as the all 1's broadcast: 255.255.255.255
  - A limited broadcast useful for certain startup protocols:

Special Addresses

- Conventions…. The ‘this’ rule and the ‘all’ rule:
  - A netid or a host id of ‘0’ implies ‘this’
  - A netid or a host id of ‘1’ implies ‘all’

- Class C example: 192.168.1.0

  The ‘this’ rule: ‘This host’ on the network

- Class C example: 192.168.1.255

  The ‘1’s’ rule: ‘All’ hosts on the network

Addresses with Special Meaning

- All 0's and all 1's: has slightly different meanings for the entire address, just the network id or just the host id
- Loop Back : 127.x.x.x (e.g., 127.0.0.1)
  - What address class?
  - Conceptually, all 2exp24 addresses in this class "A" network are equivalent.
    - In reality, the Host OS probably will only accept a few. Try
      - Ping 127.0.0.1
      - Ping 127.22.22.22
      - Ping localhost //A defacto standard host name that gets resolved to 127.0.0.1
- Private address space:
  - RFC 1918 defines certain address ranges for private use.
    - 10.0.0.0 - 10.255.255.255 (Class A space)
    - 172.16.0.0 - 172.31.255.255(Class B space)
    - 192.168.0.0 - 192.168.255.255 (Class C space)

Special Addresses

So, how many valid host ids are available with the following address:

192.1.1.0/24

What happens if you ping 0.0.0.0 ?

Classful Address Ranges

Class	Lowest Address	Highest Address
A	1.0.0.0	126.0.0.0
B	128.1.0.0	191.255.0.0
C	192.0.1.0	223.255.255.0
D	224.0.0.0	239.255.255.255
E	240.0.0.0	247.255.255.255

Dotted Decimal Class Address Ranges

Classful Addresses

- Addresses do not specify computers, but rather connections to particular hosts.
- Multihomed: A host that has >1 physical connection.

Classful Addresses

Basic mechanism: A two level class hierarchy.

- requires a unique network prefix for each physical interface.
- Two additional schemes designed to conserve net addresses: subnet and classless addressing (CIDR).

Other issues:

- Mobility
- Flexibility
- Naming
- And of course the IP Address shortage problem...

Subnetworks

Subnet addressing: Allows an organization to create multiple logical networks internally while preserving the presence of a single autonomous network to the outside.

An autonomous network:

A collection of routers and networks that fall under one administrative entity.

Subnetworks

Example : 130.127.49.225 mask 255.255.255.0

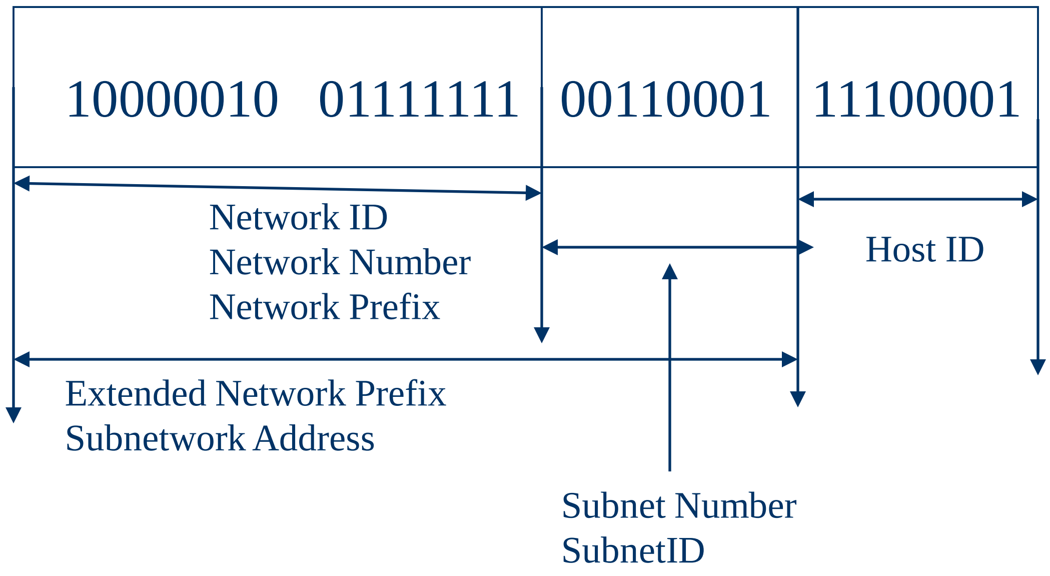

Classfull network address: 130.127.0.0
Classfull network prefix notation: 130.127/16
Subnetwork address: 130.127.49.0
Network Prefix Notation: 130.127.49/24

Address Resolution Protocol

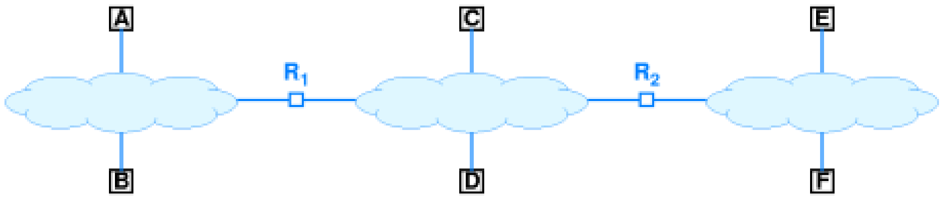

The Address Resolution Problem:
In order for two hosts on the same network to communicate, they need to know each others MAC address.

IP Header	IP Data

↓

Dest Address (6)	Source Address (6)	Length (2)	DSAP (1)	SSAP (1)	cntl (1)	code (3)	type (2)	Data (38 - 1492)	CRC (4)

Address Resolution Protocol

IP Address	Hardware Address
197.15.3.2	0A:07:4B:12:82:36
197.15.3.3	0A:9C:28:71:32:8D
197.15.3.4	0A:11:C3:68:01:99
197.15.3.5	0A:74:59:32:CC:1F
197.15.3.6	0A:04:BC:00:03:28
197.15.3.7	0A:77:81:0E:52:FA

One approach is through direct binding: Phys = f(IP)

- Objective is to encode the IP address in Physical

•Dynamic binding protocol is required for Ethernet networks where:

- Physical address larger than IP address (48 vs 32 bits)
- Physical address can change
- Potentially many hosts on the network

•Address resolution protocol (ARP) is a table lookup approach

- Defined by RFC 826

ARP: Host A Sends to Host B

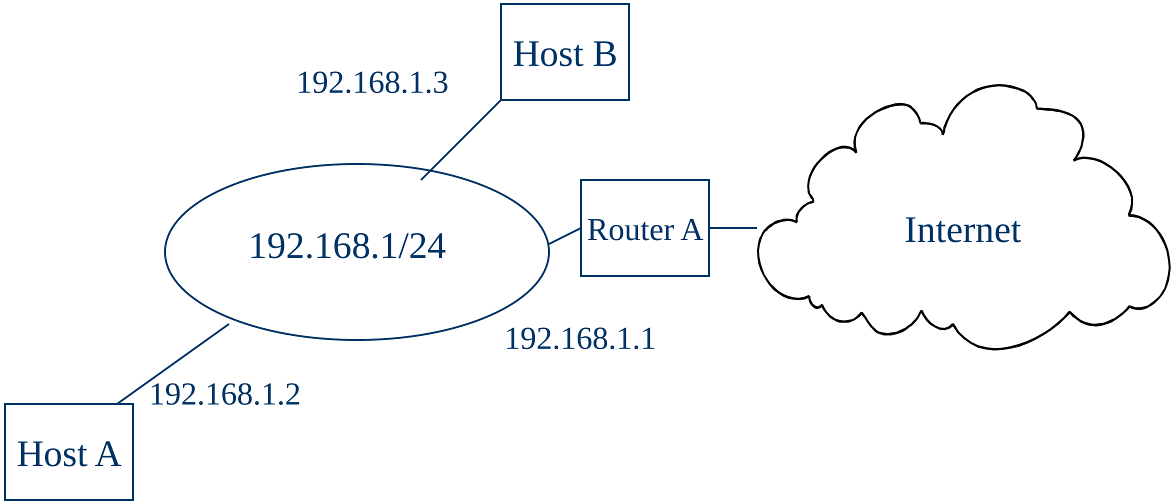

ARP: Host A Sends to Host B

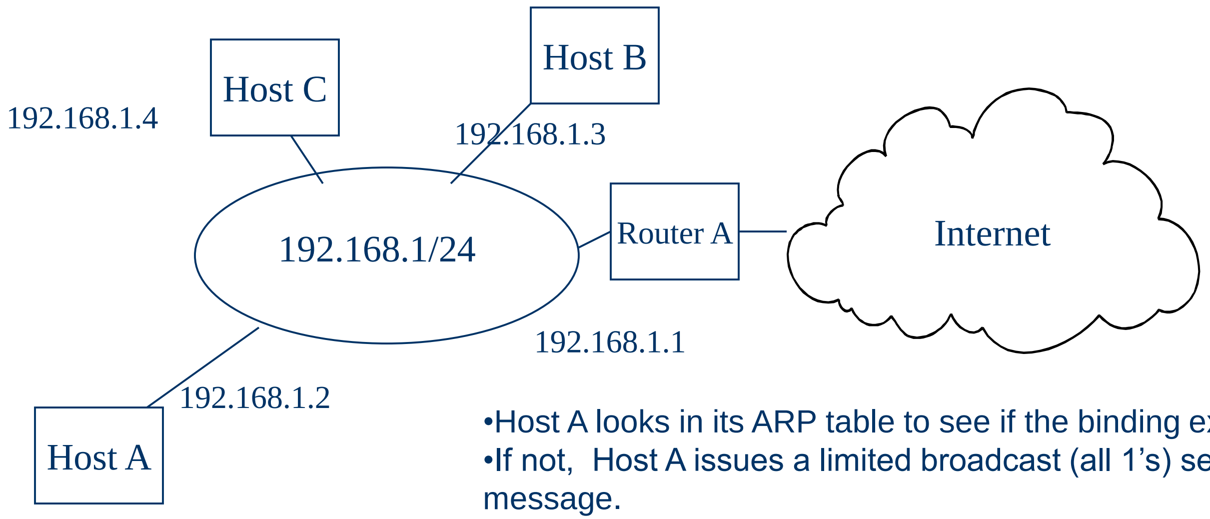

Host A's arp cache:

IP ADDR MAC ADDR

192.168.1.4 xxxxxx

- Host A looks in its ARP table to see if the binding exists
- If not, Host A issues a limited broadcast (all 1's) sending an ARP 'whois' message.
- All Hosts on the network receive the broadcast. Remember that on a LAN, the frames that are received must have a matching destination MAC address or the dest.Address must be the Ethernet broadcast address.
- Host C ignores the broadcast, Host B receives the message and passes it to its ARP program
- Host B replies with a response that says map my IP address to this MAC addr

Try an arp –a on a windows machine!!

The Internet Protocol (IP)

(a)

Destination	Next Hop
net 1	R_1
net 2	deliver direct
net 3	deliver direct
net 4	R_3

(b)

(a) An example internet with three routers connecting four physical networks, and (b) the conceptual routing table found in router R2. Each entry in the table lists a destination network and the next hop along a route to that network.

- The IP standard (RFC 791) defines:
  - The basic unit of transfer (the datagram)
  - Routing/data forwarding
  - Rules for how hosts/routers should deal with packets and errors

Routing/Forwarding

- **Routing** means process of selecting a path over which to send packets.
- **Router** is the computer which makes such a selection.
  - The Internet is composed of multiple physical networks inter-connected by computers called *routers*.
- ***Direct Delivery -*** transmission of a datagram from one machine across a physical network directly to another.
- ***Indirect Delivery -*** when the destination is not on a directly attached network; sender passes the datagram to a router for delivery.

Indirect Delivery: Host A sends to Host C

The forwarding problem when two Hosts on DIFFERENT networks want to communicate: The sending Host must learn the MAC address of the correct Router.

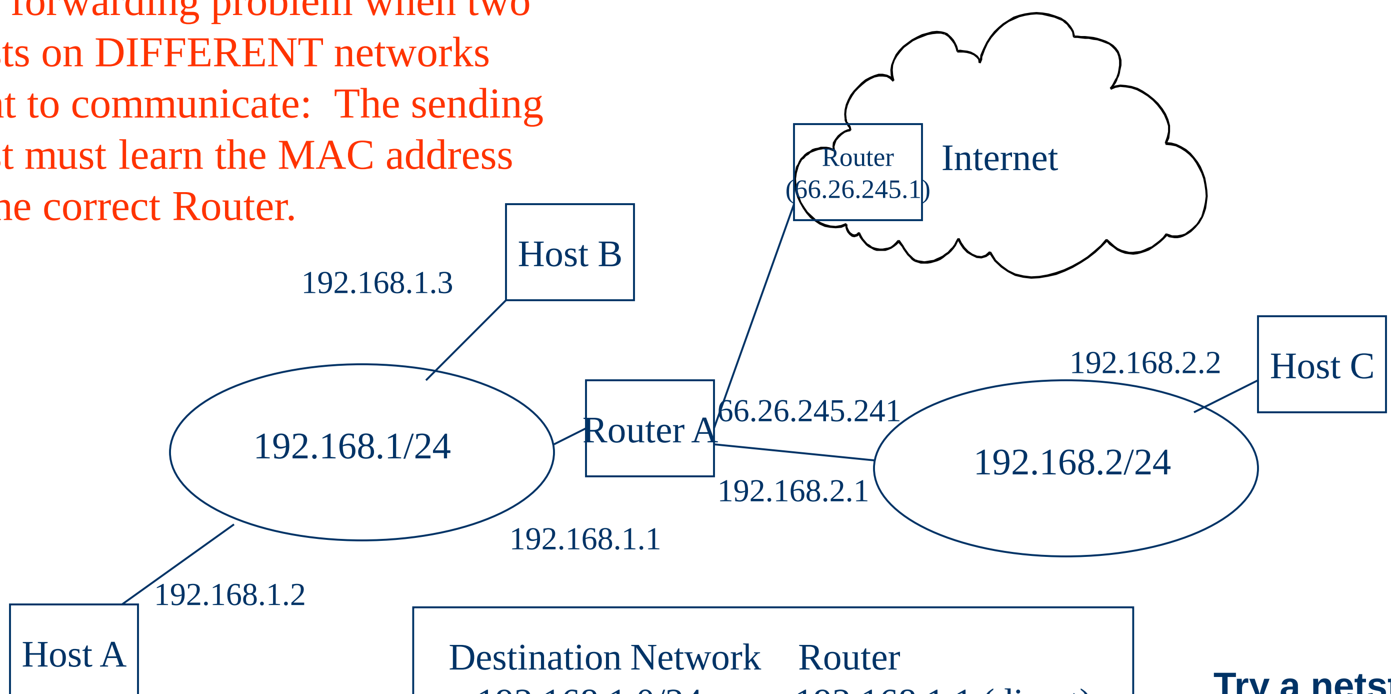

Destination Network	Router
192.168.1.0/24	192.168.1.1 (direct)
192.168.2.0/24	192.168.2.1 (direct)
66.26.245.1	66.26.245.241 (direct)
default	66.26.245.1 (indirect)

Try a netstat -r on a windows machine!!

Indirect Delivery: Host A sends to Host C

Example: Host A issues a ping to Host C.

- Assume the ARP cache in all nodes is empty

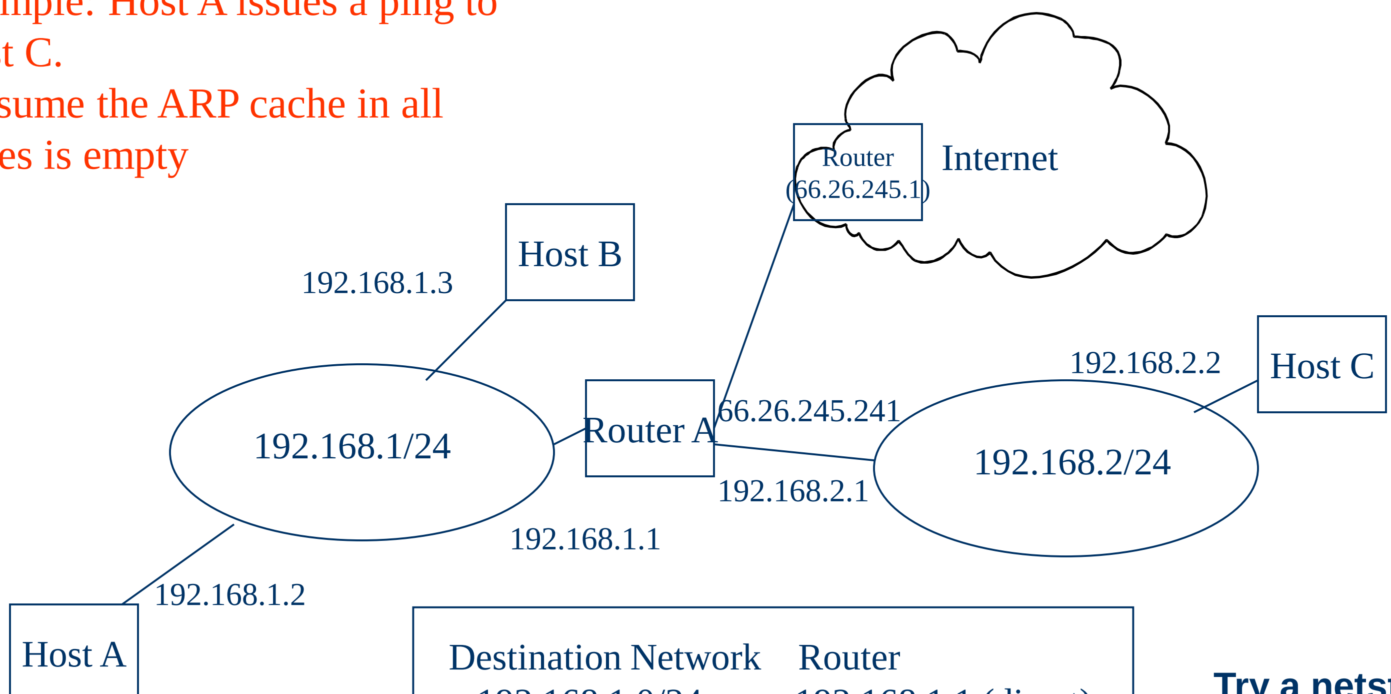

Destination Network	Router
192.168.1.0/24	192.168.1.1 (direct)
192.168.2.0/24	192.168.2.1 (direct)
66.26.245.1	66.26.245.241 (direct)
default	66.26.245.1 (indirect)

Try a netstat -r on a windows machine!!

Indirect Delivery: Host A sends to Host C

Example: Host A issues a ping to Host C.

- Assume the ARP cache in all nodes is empty
- Step 1: ARP 'whois' from Host A to Router A

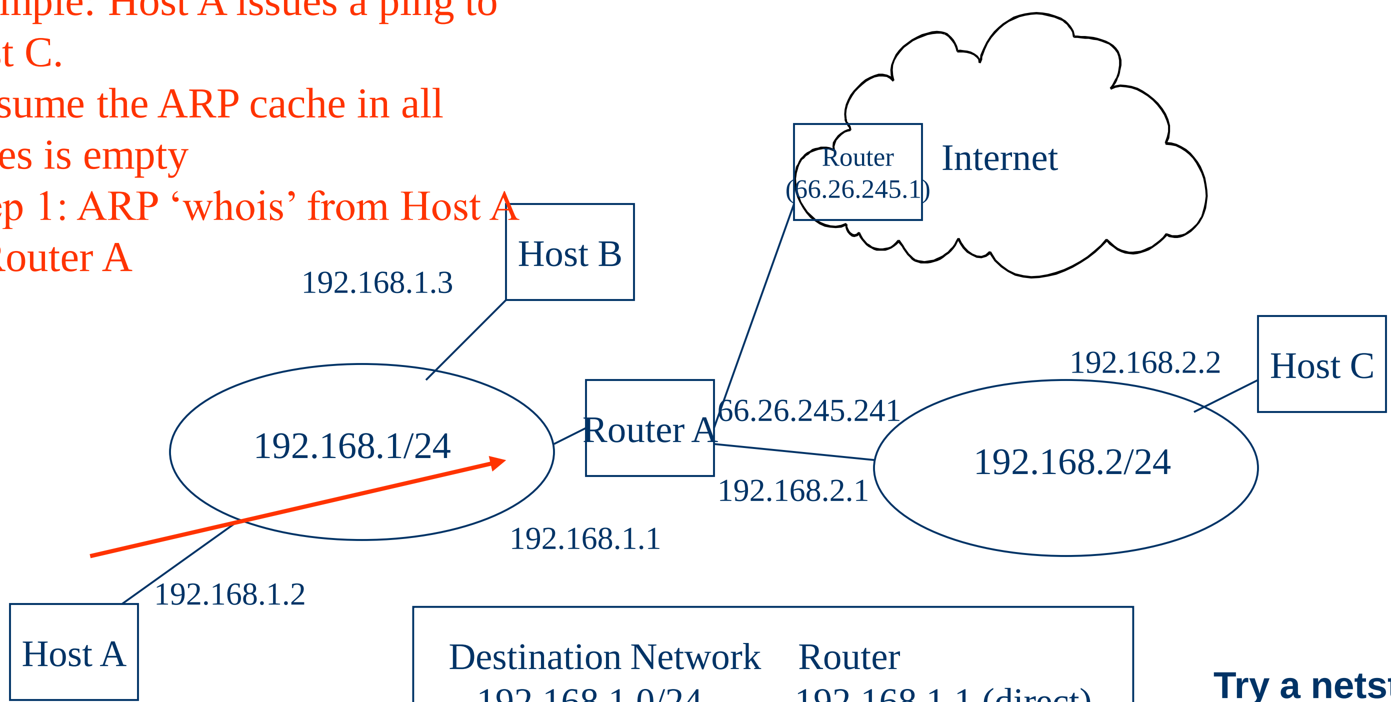

Destination Network	Router
192.168.1.0/24	192.168.1.1 (direct)
192.168.2.0/24	192.168.2.1 (direct)
66.26.245.1	66.26.245.241 (direct)
default	66.26.245.1 (indirect)

Try a netstat -r on a windows machine!!

Indirect Delivery: Host A sends to Host C

Example: Host A issues a ping to Host C.

- Assume the ARP cache in all nodes is empty
- Step 1: ARP 'whois' from Host A to Router A
- Step 2: ARP 'I am' from Router A to Host A

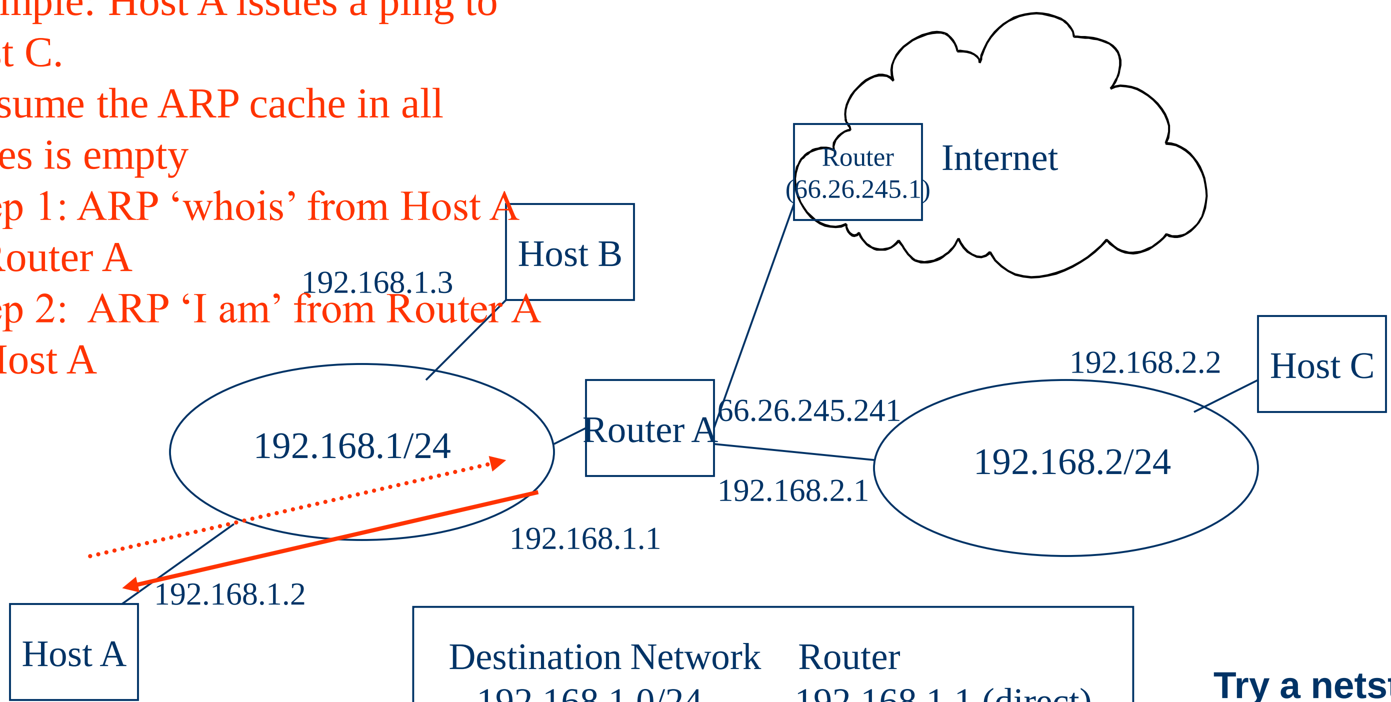

Destination Network	Router
192.168.1.0/24	192.168.1.1 (direct)
192.168.2.0/24	192.168.2.1 (direct)
66.26.245.1	66.26.245.241 (direct)
default	66.26.245.1 (indirect)

Try a netstat -r on a windows machine!!

Indirect Delivery: Host A sends to Host C

- Step 3: Host A forwards IP packet to Router A

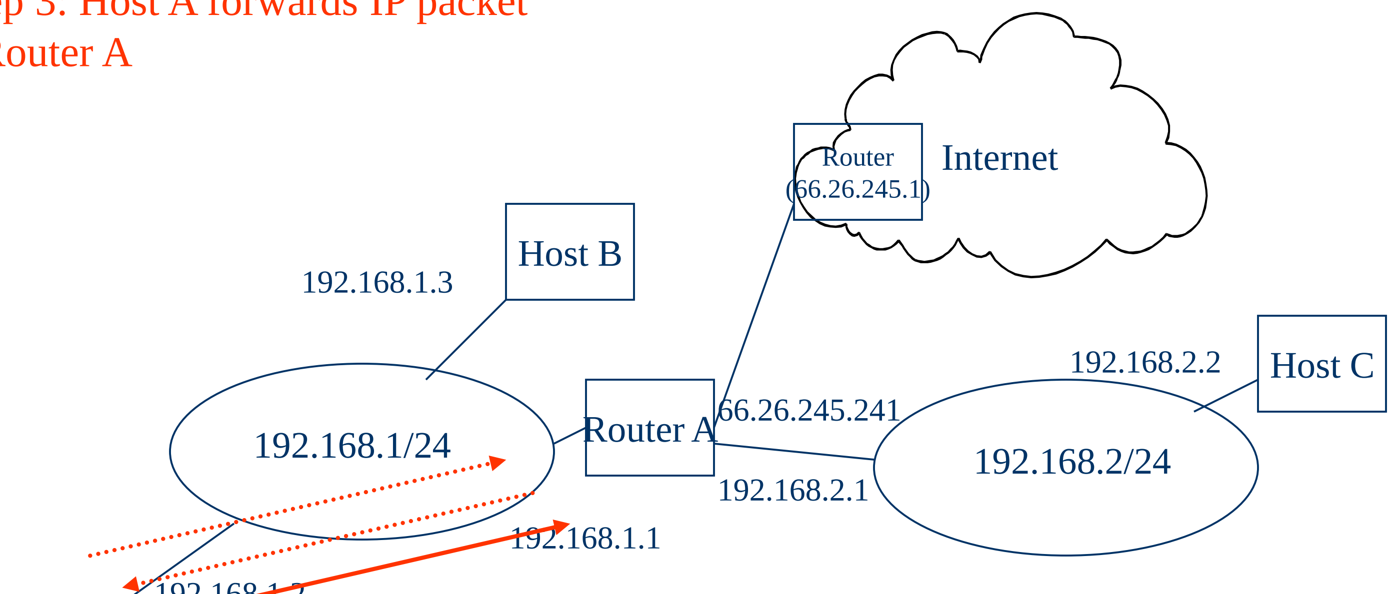

Destination Network	Router
192.168.1.0/24	192.168.1.1 (direct)
192.168.2.0/24	192.168.2.1 (direct)
66.26.245.1	66.26.245.241 (direct)
default	66.26.245.1 (indirect)

Try a netstat -r on a windows machine!!

Indirect Delivery: Host A sends to Host C

- Step 3: Host A forwards IP packet to Router A
- Step 4: Router A issues a ‘who is’ for Host C

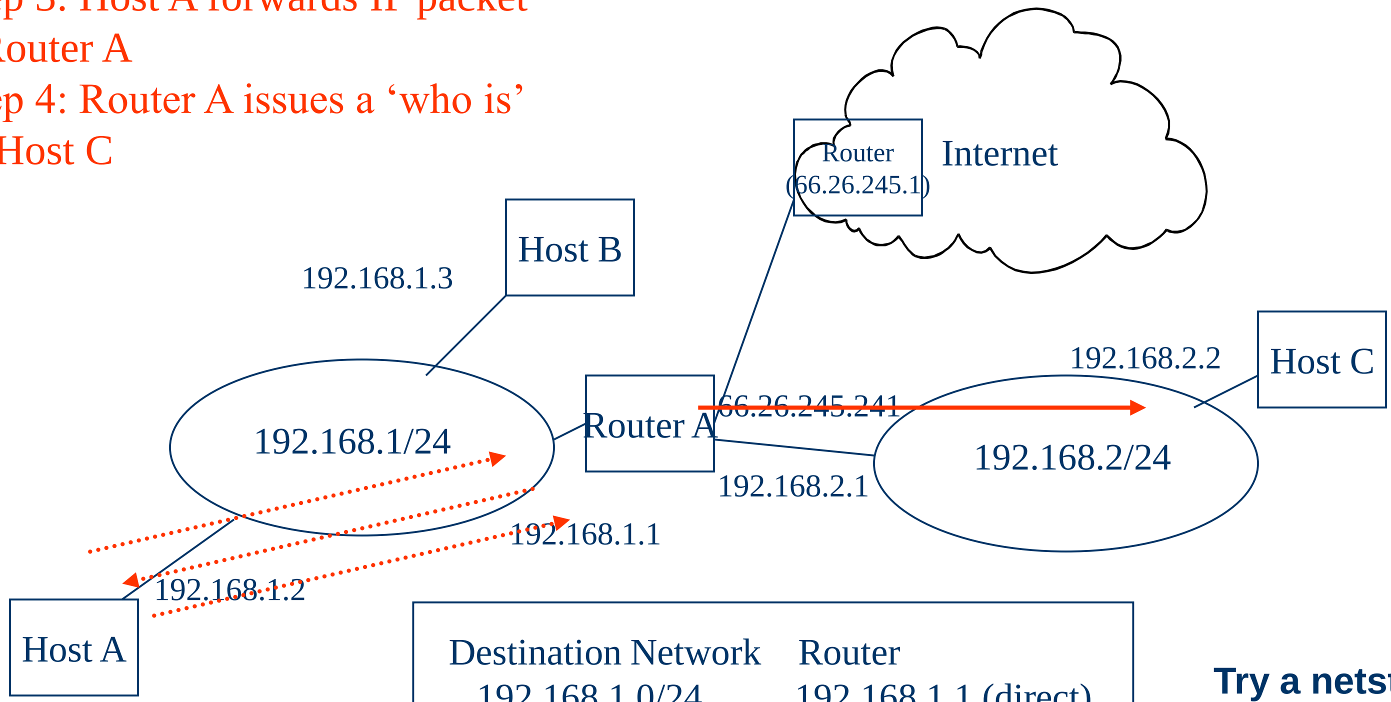

Destination Network	Router
192.168.1.0/24	192.168.1.1 (direct)
192.168.2.0/24	192.168.2.1 (direct)
66.26.245.1	66.26.245.241 (direct)
default	66.26.245.1 (indirect)

Try a netstat -r on a windows machine!!

Indirect Delivery: Host A sends to Host C

- Step 3: Host A forwards IP packet to Router A
- Step 4: Router A issues a ‘who is’ for Host C
- Step 5: Host C issues a ‘I am’ ARP message to Router A

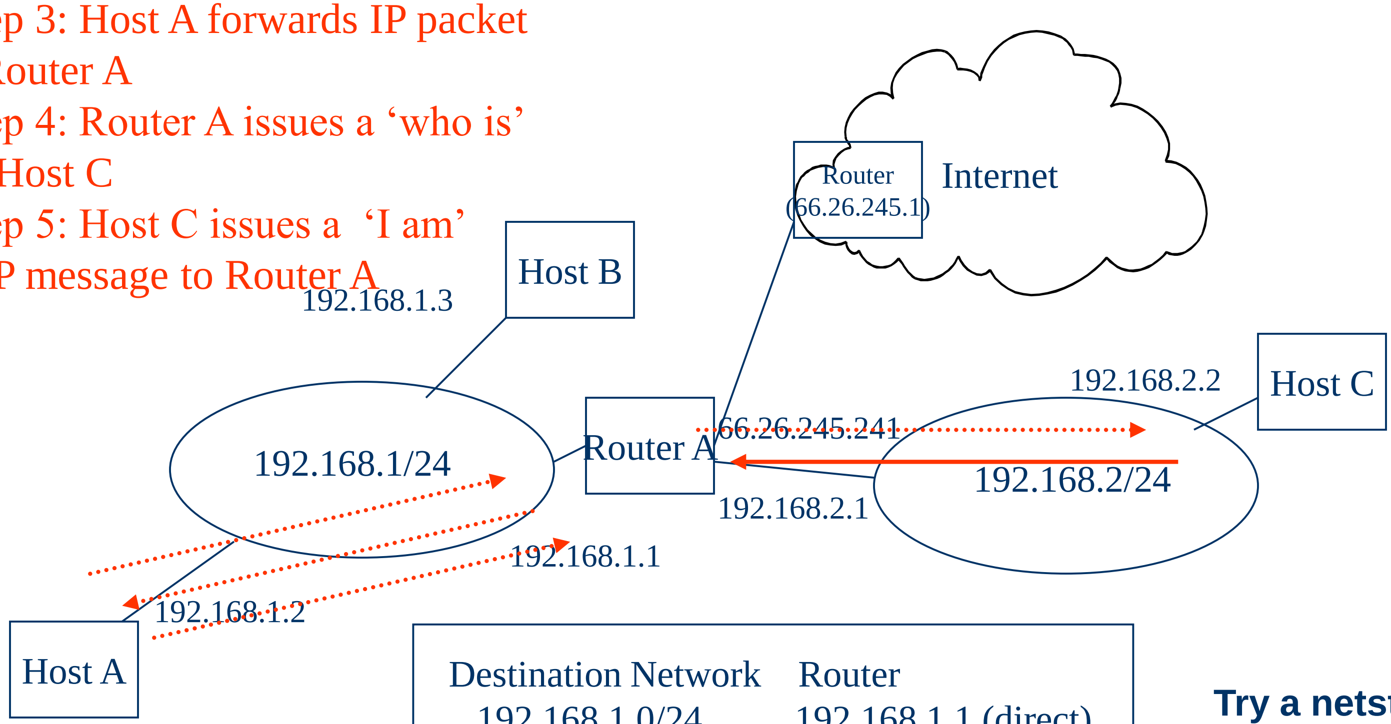

Destination Network	Router
192.168.1.0/24	192.168.1.1 (direct)
192.168.2.0/24	192.168.2.1 (direct)
66.26.245.1	66.26.245.241 (direct)
default	66.26.245.1 (indirect)

Try a netstat -r on a windows machine!!

Indirect Delivery: Host A sends to Host C

- Step 3: Host A forwards IP packet to Router A
- Step 4: Router A issues a ‘who is’ for Host C
- Step 5: Host C issues a ‘I am’ ARP message to Router A
- Step 6: Router A forwards the packet to Host C

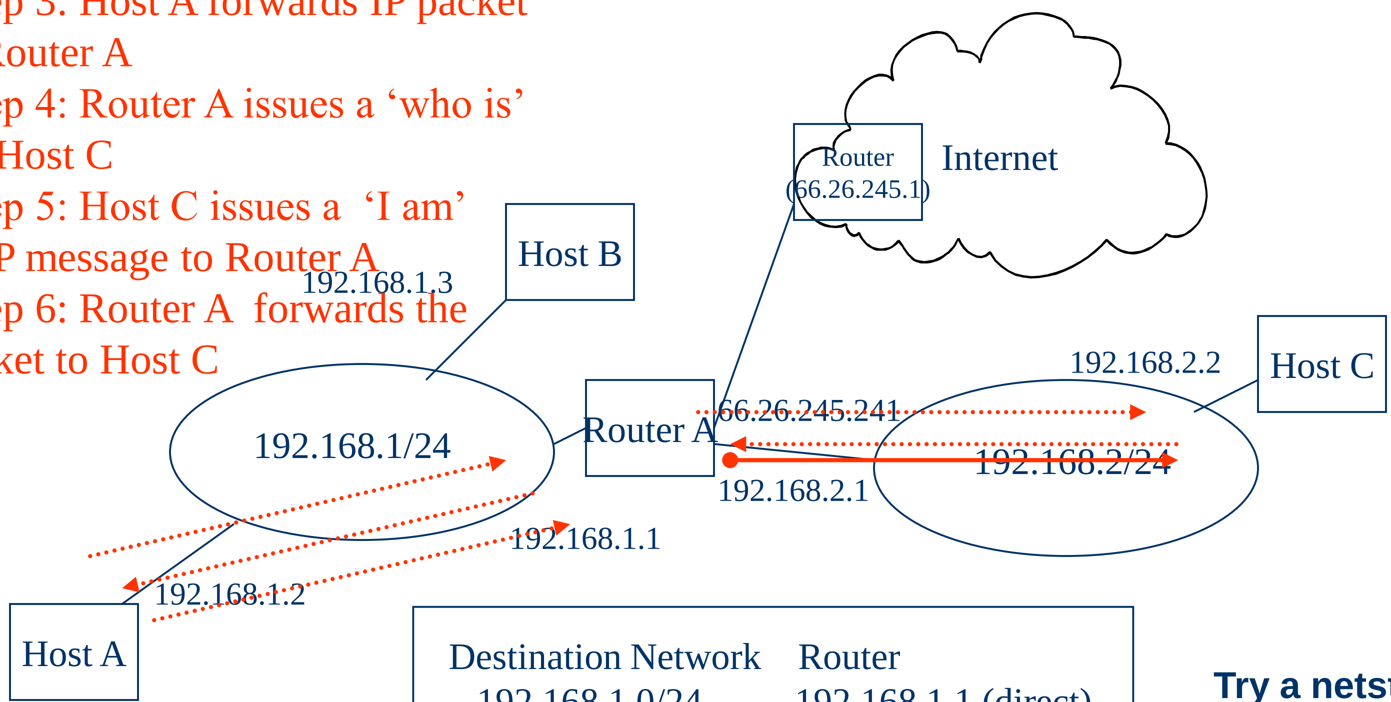

Destination Network	Router
192.168.1.0/24	192.168.1.1 (direct)
192.168.2.0/24	192.168.2.1 (direct)
66.26.245.1	66.26.245.241 (direct)
default	66.26.245.1 (indirect)

Try a netstat -r on a windows machine!!

IP Datagram Format: UDP example

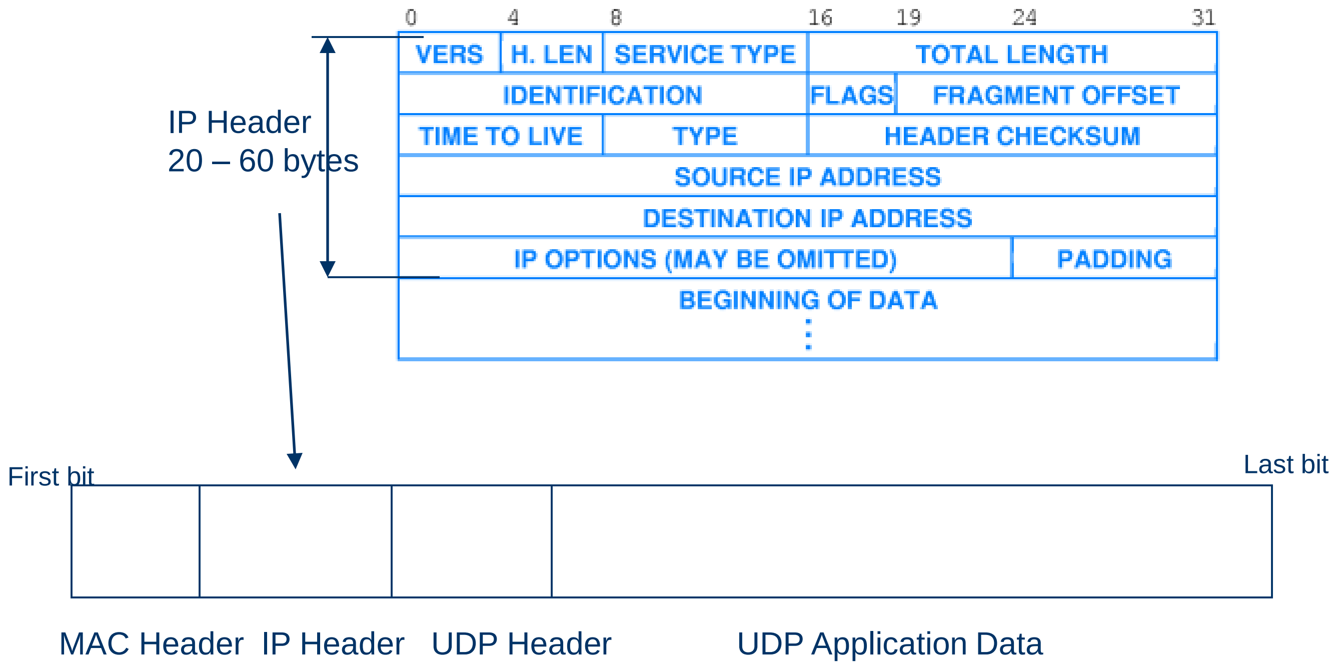

This is what gets sent 'on the wire': a frame which contains an IP Packet

IP Datagram Format: UDP example

Here is a 1500 byte UDP packet in hex:

4500 05dc   version 4, length 20bytes, TOS:0, total length:1500bytes
1234 4000   id:1234 Flags(3 bits):010 in binary (DF set), offset : 0
ff11 aaaa   TTL:255, protocol: 0x11(UDP), checksum: 0xaaaa
c0a8 0164   src:192.168.1.100
ca01 0101   dst: 202.1.1.1

5555 6666   src port: 0x5555, dst port 0x6666
05c8 bbbb   UDPlength:1480, UDP checksum:0xbbbb
…
1472 bytes of application data
…

IP Fragmentation: Ping example

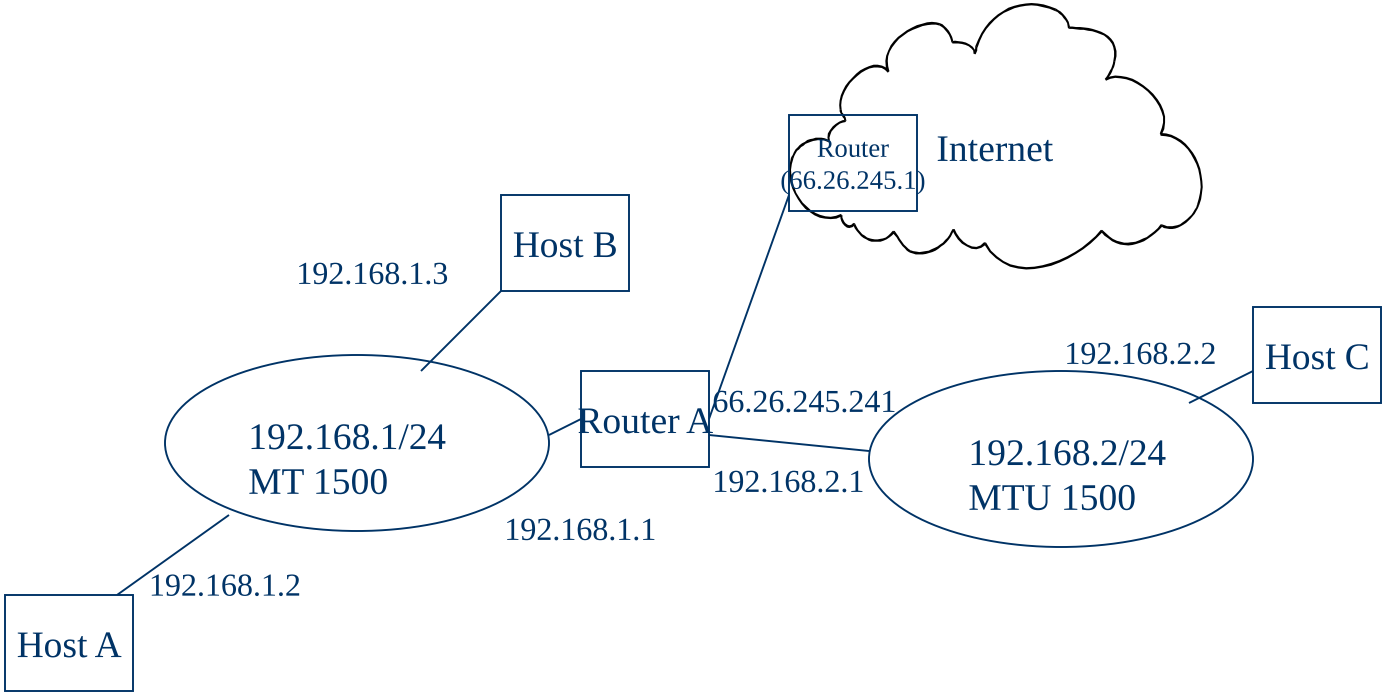

Ping –l 3000 192.168.2.2

IP Fragmentation control

{unused, don’t frag, more frags}	offset

•Identification field: uniquely ids a datagram
•Flags: DF (don’t frag) and MF (more frag)
•Fragment offset: specifies the offset in the original datagram of the data being carried in the fragment specified in units of 8 octets.
•The data portion of a fragment must be a multiple of 8 (except for the last fragment).

IP Fragmentation: Ping example