HY-335 Computer Networks

ARP - ETHERNET - MAC Addresses

Workshop

22nd Dec 2023

Before we start…

… any questions/remarks
on the 5th assignment?

Switched LANs

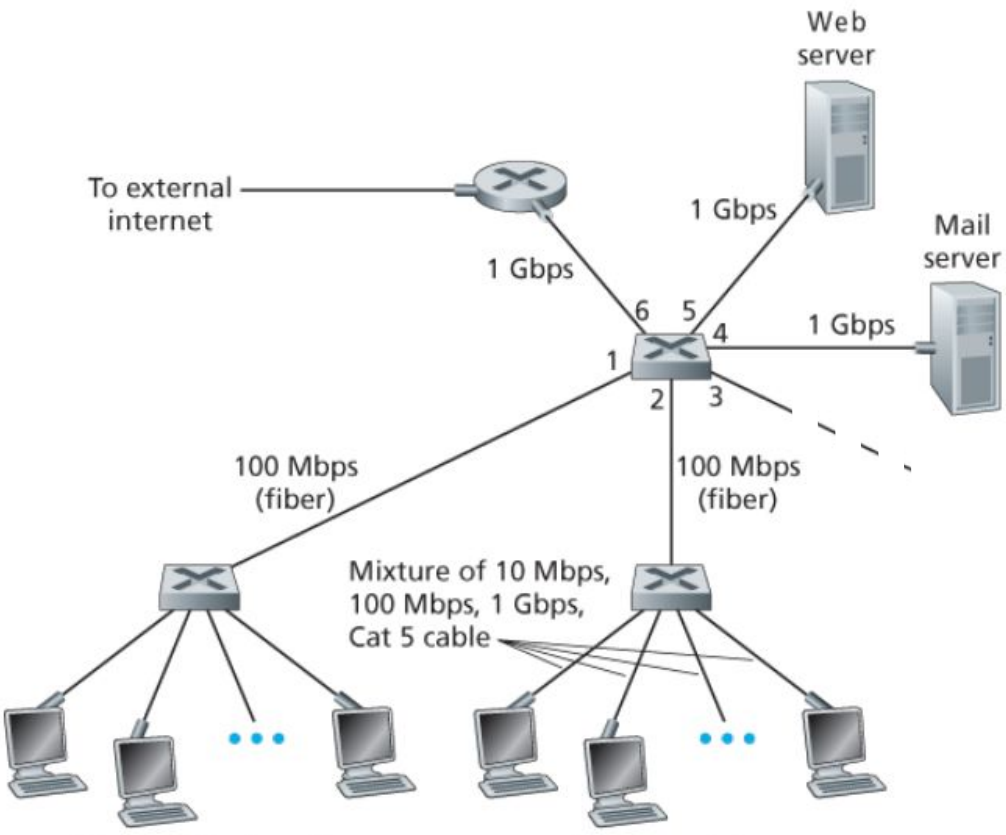

1) **Operate at the link layer**
2) **They switch link-layer frames** (rather than network-layer datagrams)
3) They *don't recognize network-layer addresses*, and *don't use routing algorithms like RIP or OSPF* to determine paths through the network of layer-2 switches.
4) **They use link-layer addresses to forward link-layer frames through the network of switches** (instead of using IP addresses).

Roadmap

I. **Link-Layer Addressing and ARP**
   A. MAC Addresses
   B. Address Resolution Protocol (ARP)
      1. *ARP table*
      2. *Sending a Datagram off the Subnet*
      3. *ARP queries*

II. **Ethernet**
   A. Ethernet Frame Structure
   B. Ethernet Technologies

III. **Link-Layer Switches**

MAC Addresses

- **Globally unique**
- **Flat structure** (not hierarchical, like IP)!
- **Doesn't change** no matter where the adapter goes.
- **Necessary for communication** on a local area network (LAN).
- Devices use the **Address Resolution Protocol** to acquire the MAC address of another device.

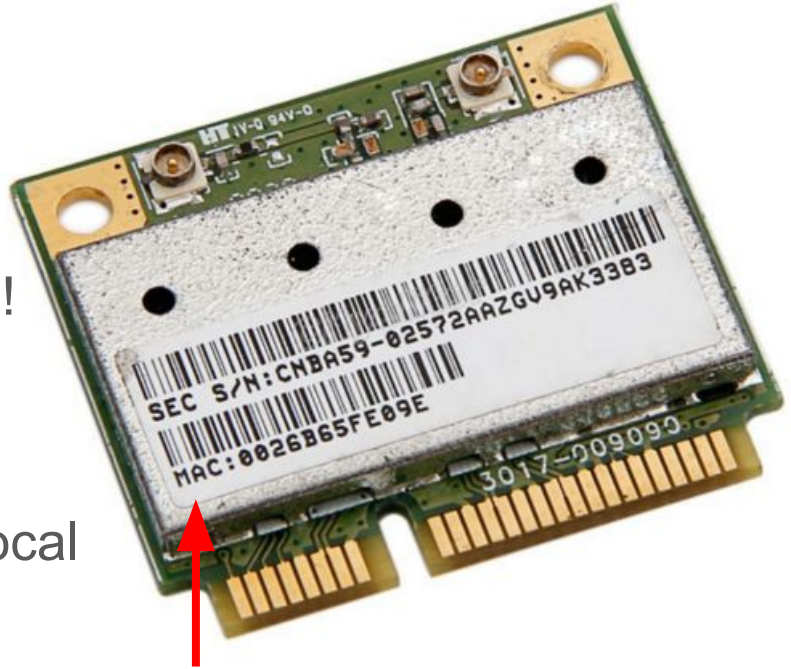

MAC ~ personal ID card
IP ~ postal address

When an adapter wants to send a frame to some destination adapter, the sending adapter:

1) Inserts the destination adapter's MAC address into the frame

2) Then sends the frame into the LAN.

When an adapter receives a frame:

1) Will check to see whether the destination MAC address in the frame matches its own MAC address

2.1) If there is a match, the adapter extracts the enclosed datagram and passes the datagram up the protocol stack.

2.2) If there isn't a match, the adapter discards the frame, without passing the network-layer datagram up.

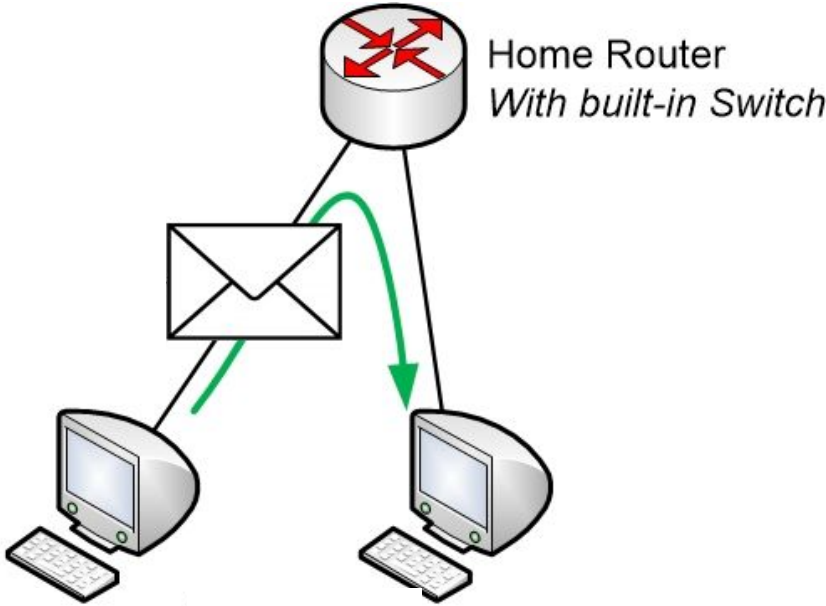

The ARP table

Each host and router has an ARP table in its memory, which contains:

- IP addresses
- MAC addresses
- Type
- a time to-live (TTL) value (typically 20 min)

IP Address	Physical Address	Type	Expiry
192.168.1.1	00:21:29:CC:A7:31	Static	-
192.168.1.24	00:17:31:BA:5C:A4	Dynamic	2m 27s
10.1.2.66	00:1A:70:3C:A6:3D	Dynamic	35s
210.210.3.1	90:6C:AC:2D:DE:9A	Dynamic	1m 14s

A table does not necessarily contain an entry for every host and router on the subnet:

- may have never been entered into the table
- may have expired

Let’s see my own ARP table!

(plz dont hacc)

Address Resolution Protocol (ARP)

The host or the router sends an ARP query packet, which includes the physical and IP addresses of the sender and the receiver.

The ARP response packet is unicast directly to the inquirer (host/router) by using the physical address received in the query packet.

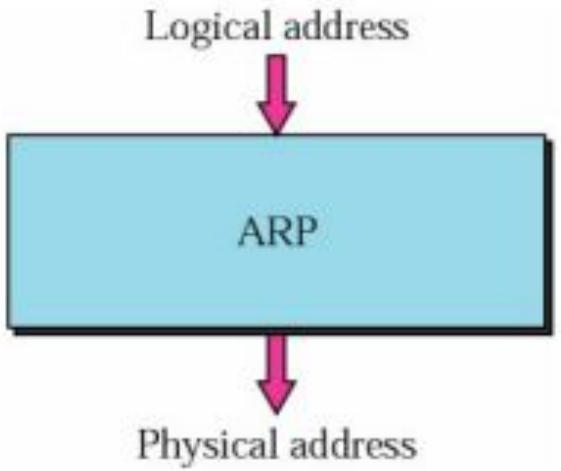

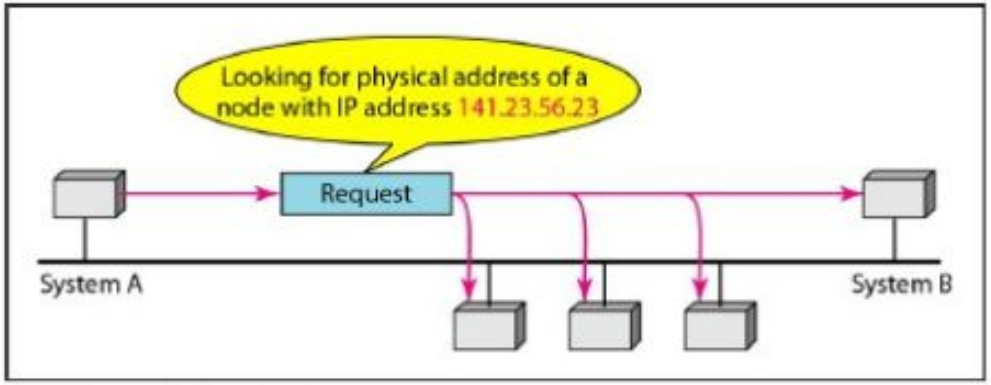

a. ARP request is broadcast

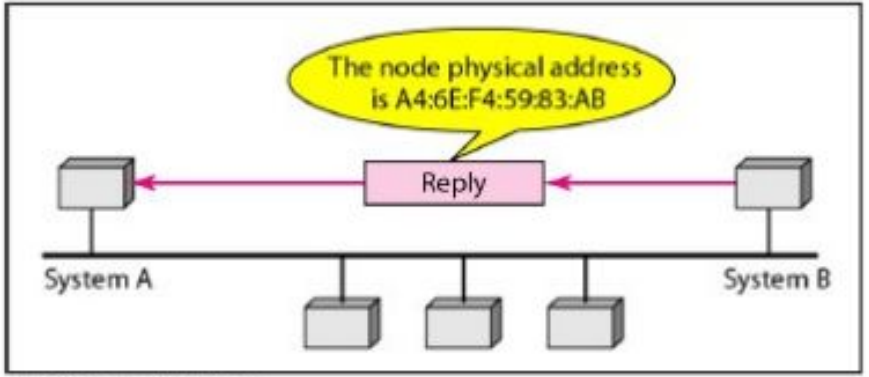

b. ARP reply is unicast

Address Resolution Protocol (ARP)

1. The sender knows the IP address of the target.
2. IP asks ARP to create an ARP request message.
3. The message is passed to the data link layer where it is encapsulated in a frame.
4. Every host or router receives the frame. All machines except the one targeted drop the packet. The target machine recognizes its IP address.
5. The target machine replies with an ARP reply message
6. The sender receives the reply message (and updates its ARP table).
7. The IP datagram, is now encapsulated in a link-layer frame and unicast to the destination.

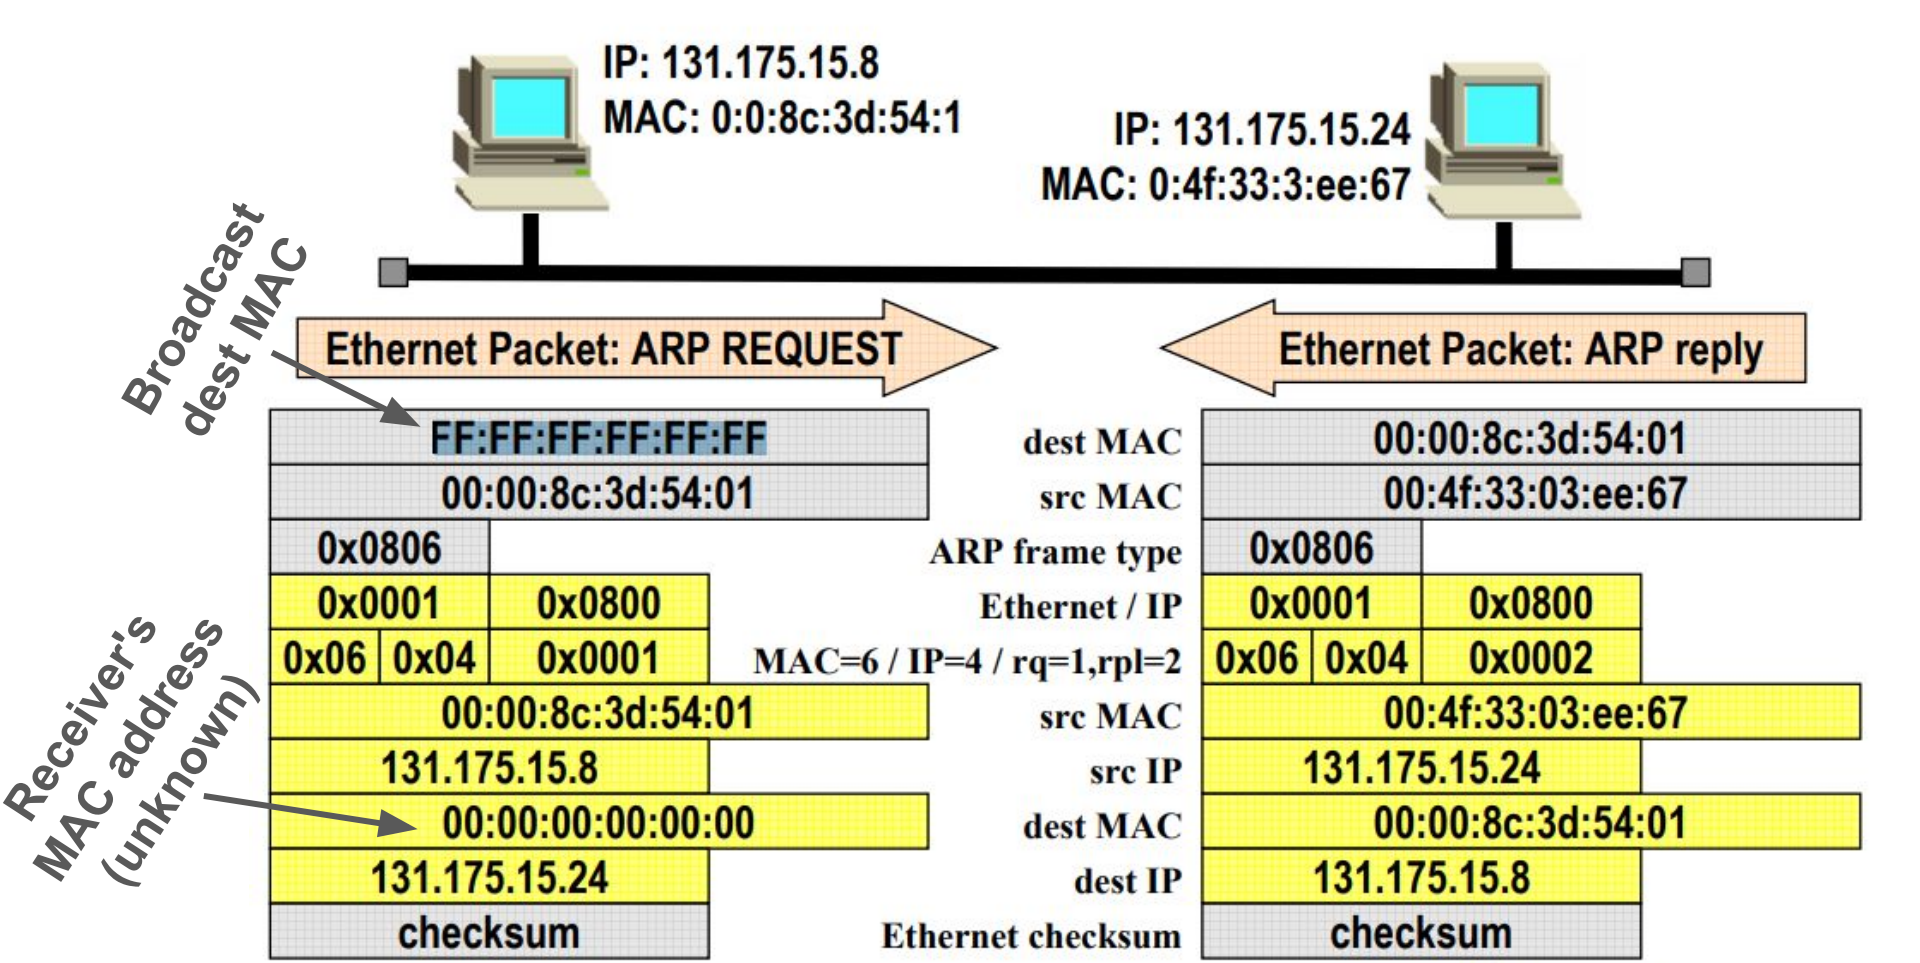
IP: 131.175.15.8
MAC: 0:0:8c:3d:54:1
IP: 131.175.15.24
MAC: 0:4f:33:3:ee:67
Broadcast dest MAC
Ethernet Packet: ARP REQUEST
Ethernet Packet: ARP reply
FF:FF:FF:FF:FF:FF
00:00:8c:3d:54:01
dest MAC
00:00:8c:3d:54:01
00:00:8c:3d:54:01
src MAC
00:4f:33:03:ee:67
0x0806
ARP frame type
0x0806
0x0001
0x0800
Ethernet / IP
0x0001
0x0800
0x06
0x04
0x0001
MAC=6 / IP=4 / rq=1,rpl=2
0x06
0x04
0x0002
00:00:8c:3d:54:01
src MAC
00:4f:33:03:ee:67
131.175.15.8
src IP
131.175.15.24
Receiver's MAC address (unknown)
00:00:00:00:00:00
dest MAC
00:00:8c:3d:54:01
131.175.15.24
dest IP
131.175.15.8
checksum
Ethernet checksum
checksum

A couple closing comments on ARP

ARP is plug-and-play: the ARP table gets built automatically— no need to be configured.

If a host becomes disconnected from the subnet, its entry is eventually deleted from the other ARP tables in the subnet.

ARP packet is encapsulated in a link-layer frame → architecturally **above the link layer.**

However, an ARP packet **has fields containing** link-layer addresses and thus is arguably **a link-layer protocol.**

But it also **contains network-layer addresses** and thus is also arguably **a network-layer protocol.**

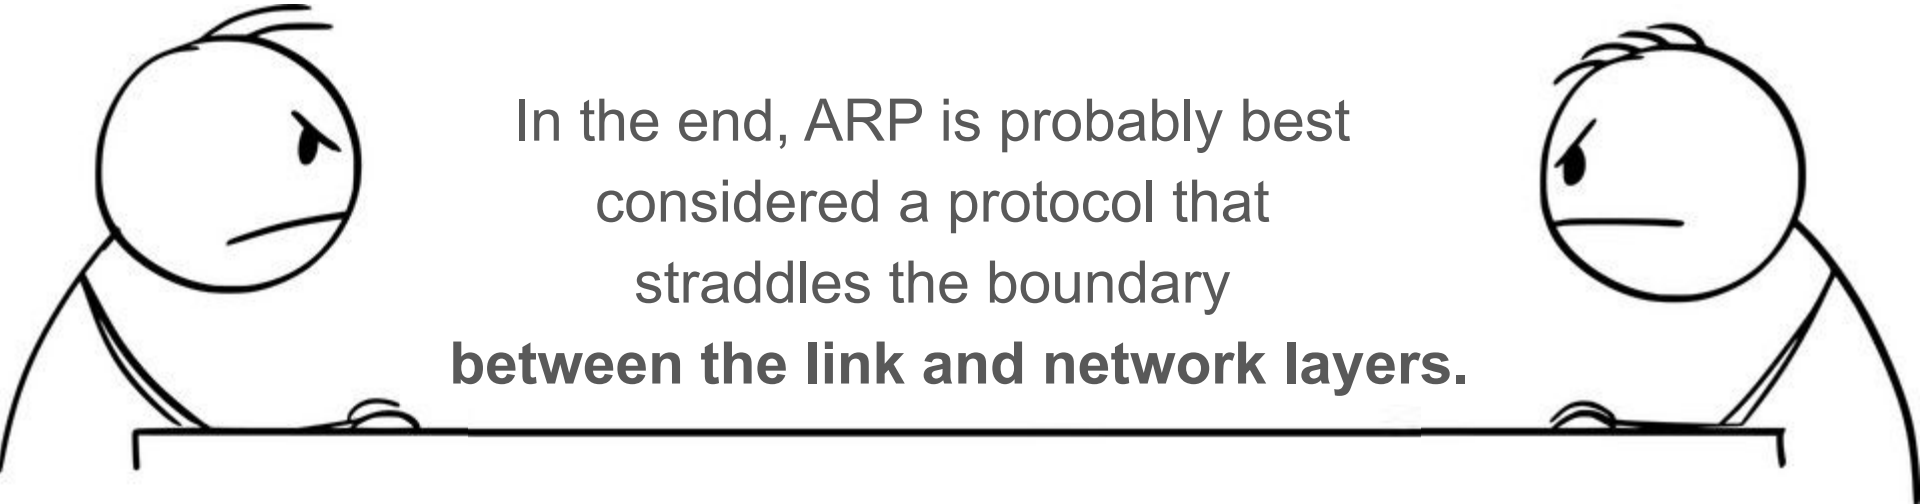

In the end, ARP is probably best considered a protocol that straddles the boundary **between the link and network layers.**

Ethernet

the first widely deployed high-speed LAN

deployed early, network administrators became intimately familiar with

others were more complex and expensive

1980s to mid-1990s: Bus topology .

By the late 1990s: (mostly) hub-based star topology.

In the early 2000s: the hub at the center was replaced with a switch.

Ethernet Frame Structure

We can learn a lot about Ethernet by examining the Ethernet frame.

- Data field (46 to 1,500 bytes). This field carries the IP datagram
- Destination address (6 bytes). This field contains the MAC address of the destination adapter
- Source address (6 bytes). This field contains the MAC address of the adapter that transmits
- Type field (2 bytes). The type field permits Ethernet to multiplex network-layer protocols
- Cyclic redundancy check (CRC) (4 bytes)
- Preamble (8 bytes).

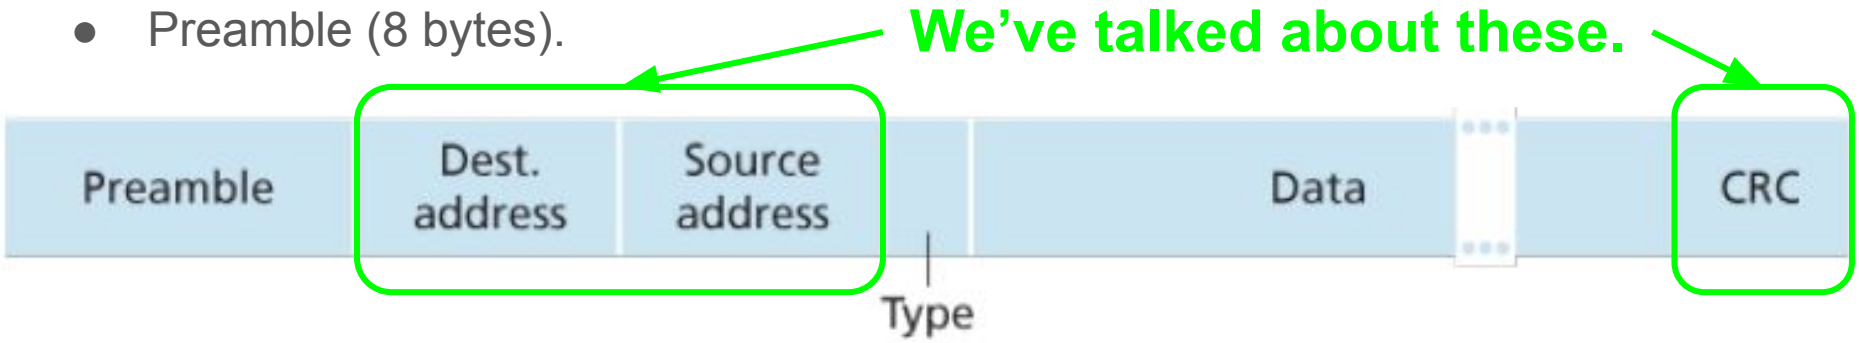

Data field (46 to 1,500 bytes)

if the IP datagram exceeds 1,500 bytes, then the host has to fragment the datagram,

if the IP datagram is less than 46 bytes, the data field has to be “stuffed” to fill it out to 46 bytes (the network layer uses the length field in the IP datagram header to remove the stuffing.)

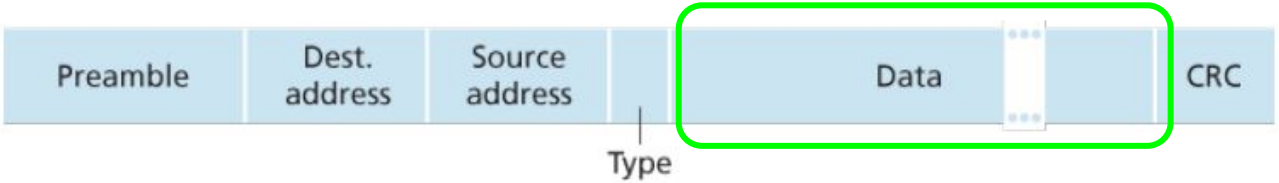

Type field (2 bytes)

hosts can use other network-layer protocols besides IP

when the Ethernet frame arrives at adapter B, adapter B needs to know to which network-layer protocol it should pass (that is, demultiplex) the contents of the data field

ARP protocol (discussed in the previous section) has its own type number, and if the arriving frame contains an ARP packet (i.e., has a type field of 0806 hexadecimal)

analogous to the protocol field in the network-layer datagram and the port-number fields in the transport-layer segment; all of these fields serve to glue a protocol at one layer to a protocol at the layer above

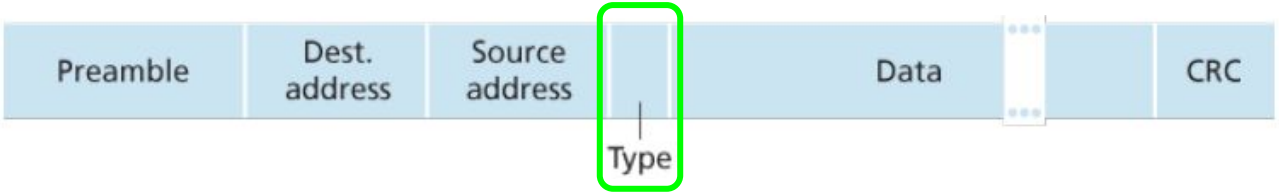

Preamble (8 bytes)

First 7 bytes: 10101010

Last byte: 10101011

The first 7 bytes of the preamble serve to “wake up” the receiving adapters and to synchronize their clocks to that of the sender’s clock.

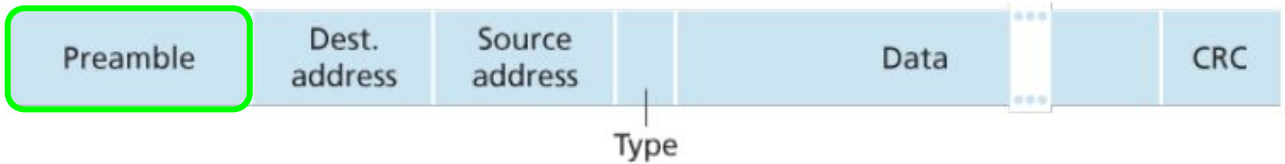

Ethernet Technologies

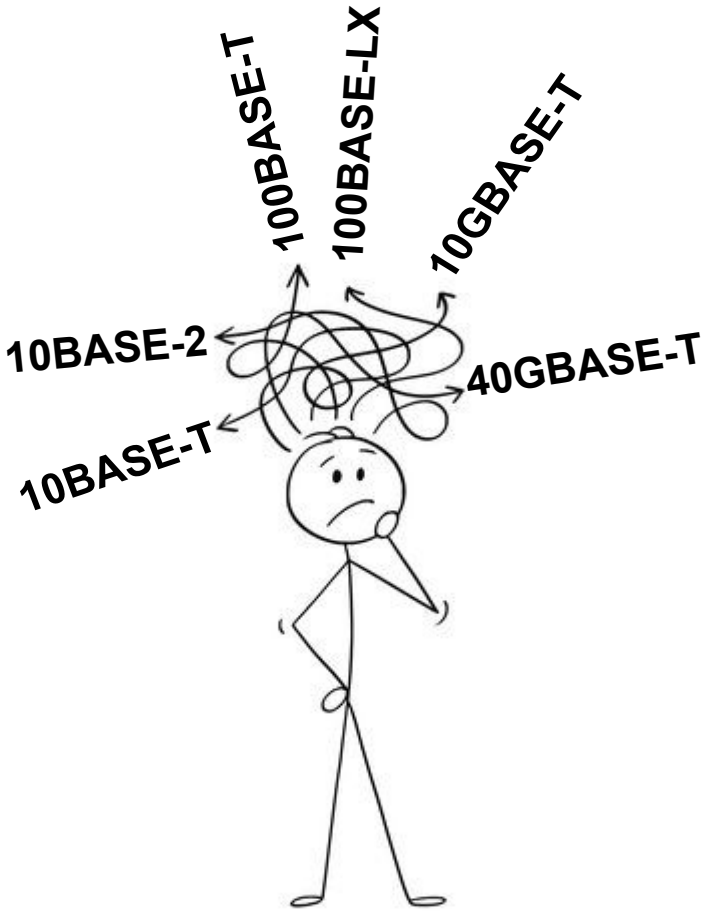

Standardized acronyms

10, 100, 1000:
10 Megabit (per second), 100 Megabit

10G, 40G:
10 Gigabit,40 Gigibit Ethernet

"BASE" → baseband Ethernet

T → twisted-pair copper wires

2 → Maximum length = 185 m

5 → Maximum length = 500 m

LX → long wavelength (1270–1355 nm, optic fibers)

100 Mbps Ethernet standards:

- Nodes are connected to a switch via point-to-point segments
- mid-1990s: Ethernet was standardized at 100 Mbps
- The original Ethernet MAC. protocol and frame format were preserved.
- higher-speed physical layers were defined for copper wire (100BASE-T) and fiber (100BASE-FX, 100BASE-SX, 100BASE-BX)

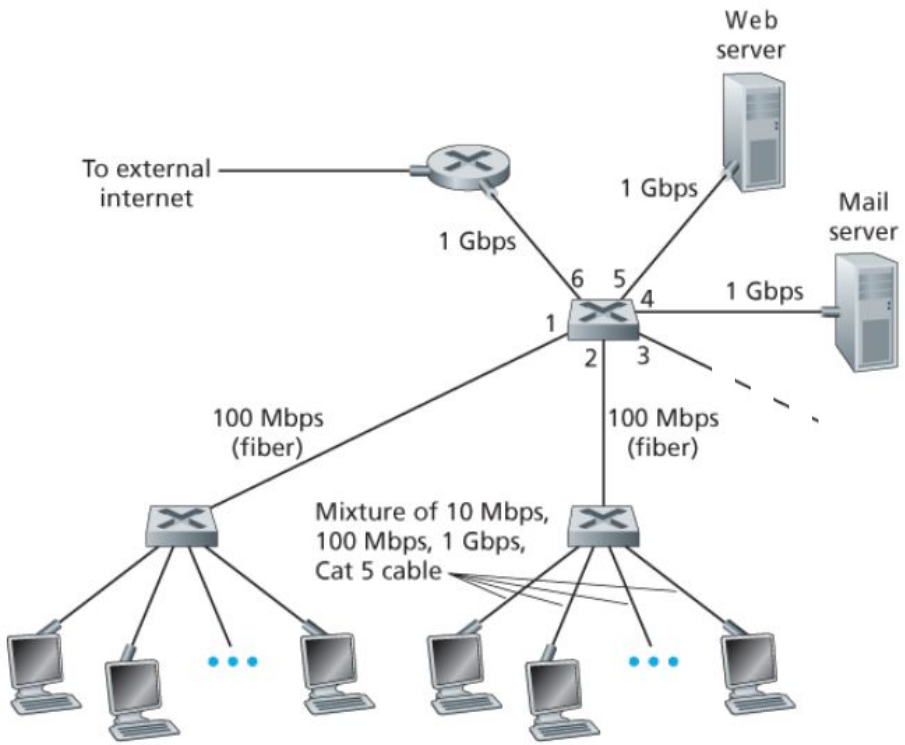

Gigabit Ethernet:

standard Ethernet frame format (backward compatible with 10BASE-T and 100BASE-T)

Allows for **point-to-point links** as well as shared **broadcast channels**

Uses CSMA/CD for shared broadcast channels.

Allows for full-duplex operation at 40 Gbps in both directions for point-to-point channels

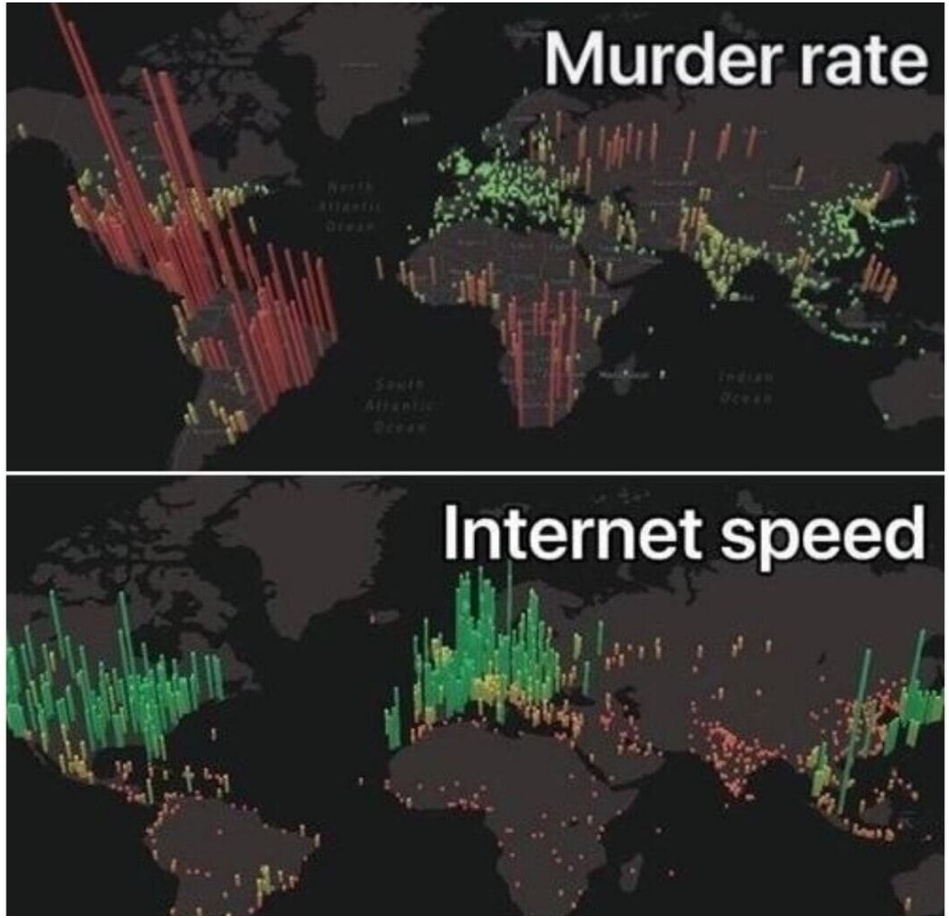

just sayin


Reverse Polarity
TNC Male
N connector
TNC connector
HD15 connector
Mini BNC connector
BNC connector

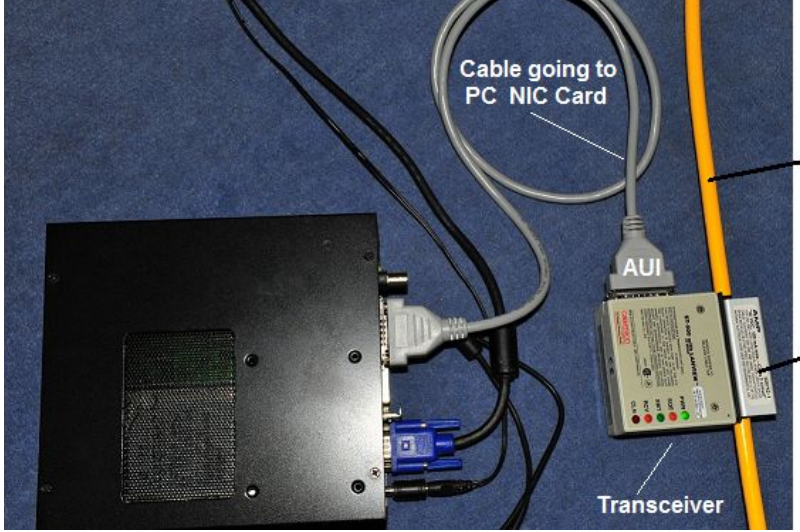
Cable going to
PC NIC Card
Thicknet
Coax Cable
AUI
Vampire Tap
Transceiver


another
innocent
victim

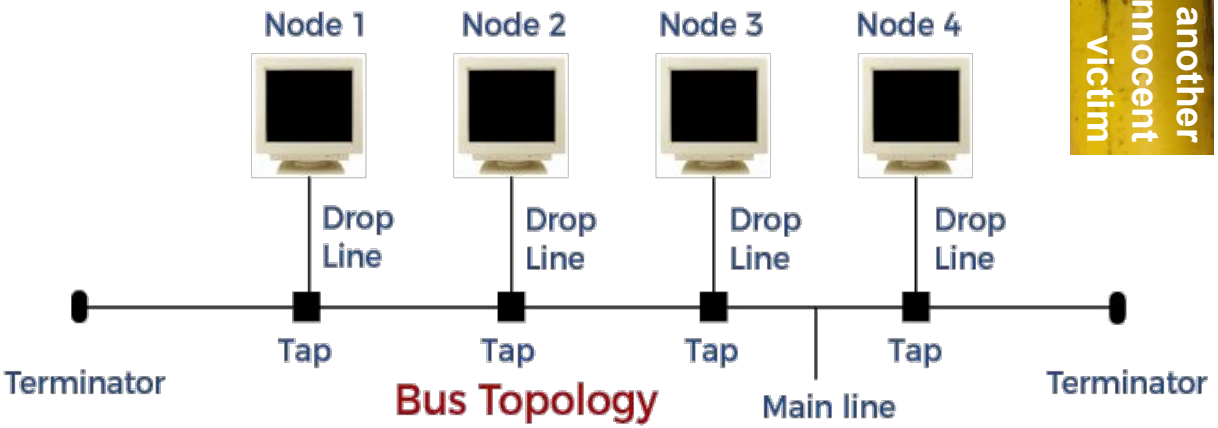
Node 1
Node 2
Node 3
Node 4
Drop
Line
Drop
Line
Drop
Line
Drop
Line
Tap
Tap
Tap
Tap
Terminator
Terminator
Bus Topology
Main line


Reverse Polarity
TNC Female
SMA connector
SMB connector
SMB HD connector
F connector
RCA connector

Link-Layer Switches

- They **receive incoming** link-layer frames and **forward them onto outgoing** links.
- They are **transparent to the hosts** and routers in the subnet.
- For scenarios of very high rate of frames arriva, switch output interfaces have buffers.
- (analogous to router output interface buffers for datagrams)

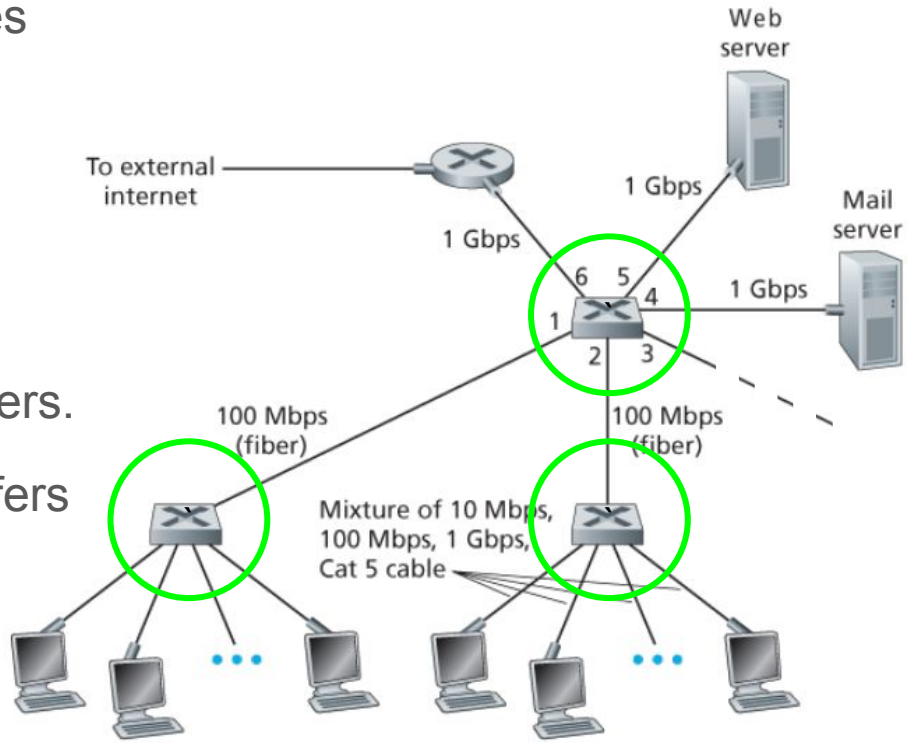

Switch Operation

Filtering: determines whether a frame should be forwarded to some interface or should just be dropped.

Forwarding: determines the interfaces to which a frame should be directed, and then moves the frame to those interfaces.

These are achieved with a switch table, which contain:

- MAC address (switches don't use IP!*)
- switch interface → MAC address
- time of entry

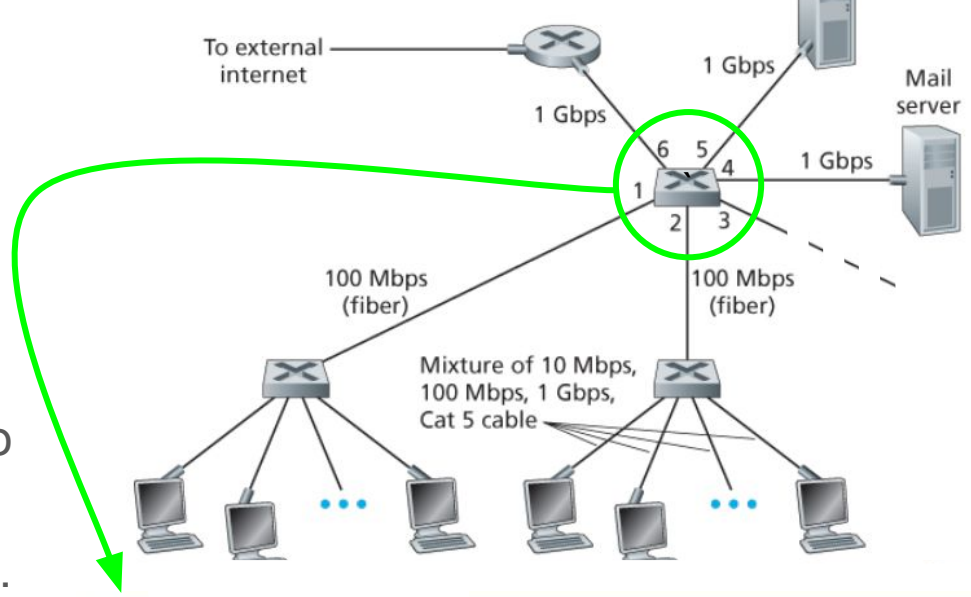

Address	Interface	Time
62-FE-F7-11-89-A3	1	9:32
7C-BA-B2-B4-91-10	3	9:36
....	....	....

Portion of a switch table for the **uppermost switch**

Switch Operation

Suppose a frame with **destination** address DDDD-DD-DD-DD-DD arrives at the switch on interface 1:

- There is **no entry** in the table for DD-DD-DD-DD-DD-DD:

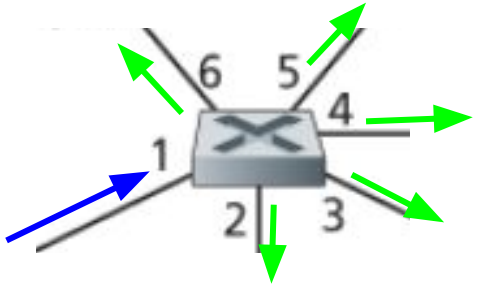

The switch **broadcasts** the frame.

- There is an entry in the table, associating DD-DD-DD-DD-DD-DD with interface 1:

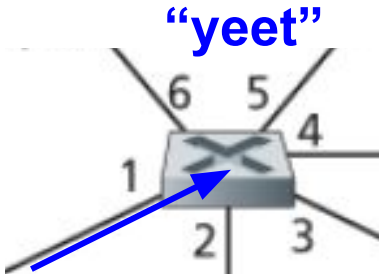

No need to forward, the switch discards the frame.

- There is an entry in the table, associating DD-DD-DD-DD-DD-DD with interface y≠1:

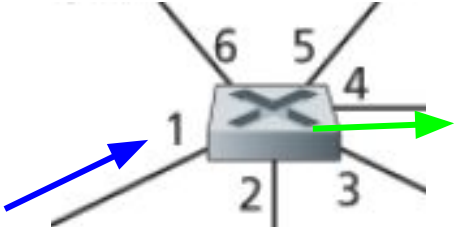

The switch puts the frame in an output buffer that precedes interface y.

Switch Self-Learning

The switch table is built automatically, dynamically, and autonomously!

1. The switch table is initially empty.
2. For each incoming frame received on an interface, the switch stores in its table:
   a. (1) the MAC address in the frame's source address field,
   b. (2) the interface from which the frame arrived, and
   c. (3) the current time.

3. The switch deletes an address in the table if no frames are received with that address as the source address after some period of time (the aging time).

Properties of Link-Layer Switching

- **Elimination of collisions:** no wasted bandwidth due to collisions!

- **Heterogeneous links:**

- **Management:**

  Example: An adapter malfunctions and continually sends Ethernet frames (jabbering adapter). The switch can detect the problem and internally disconnect the malfunctioning adapter.

Switches VS Routers

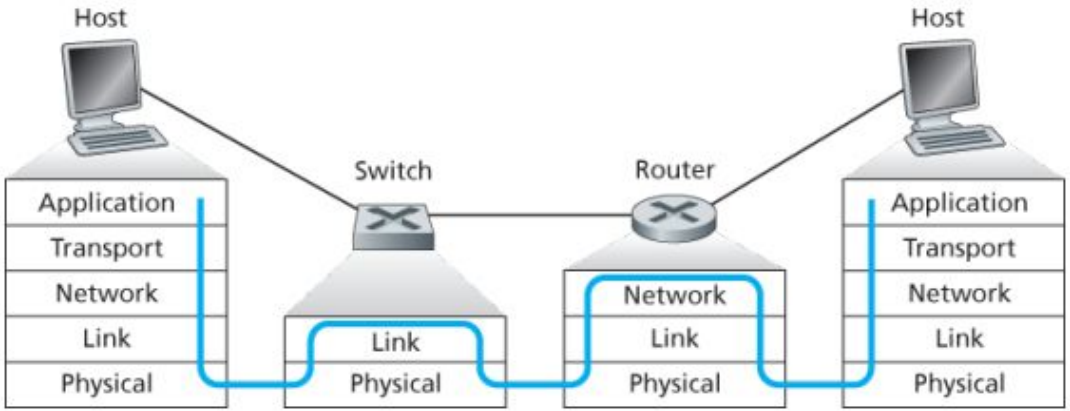

	Hubs	Routers	Switches
Traffic isolation	No	Yes	Yes
Plug and play	Yes	No	Yes
Optimal routing	No	Yes	No

Switches → small networks: localize traffic, increase aggregate throughput without requiring any configuration of IP addresses.

Routers → larger networks *(in addition to switches)*:
Routers provide a more robust isolation of traffic, control broadcast storms, and use more “intelligent” routes among the hosts in the network.

Thank you!

Sources:

- Computer Networking: A Top-Down Approach, 7th Edition
- https://www.csd.uoc.gr/~hy335a/material/Lectures/hy335a_mac_2020.pdf
- CS-204: COMPUTER NETWORKS (Dr. Vandana K.)