Chapter 9

Ethernet

Part I

Ethernet

Standards and Implementation

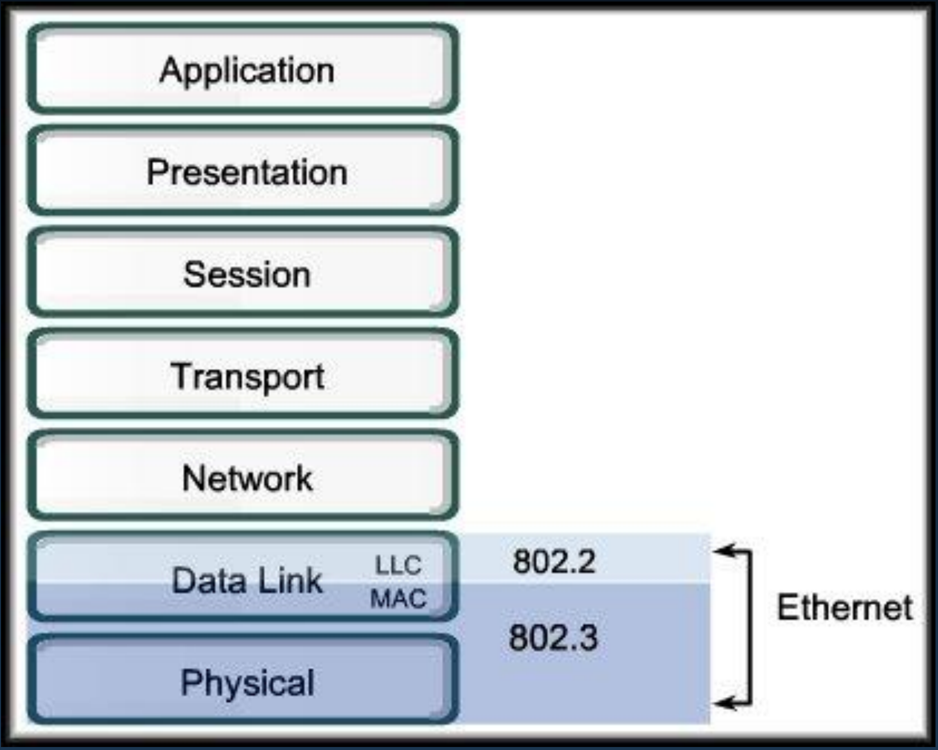

Standards and Implementation

- The most common LAN technology.
- Different media (copper cable, optical fibre)
- Different bandwidths (10, 100Mbps Gigabit, 10Gigabit)

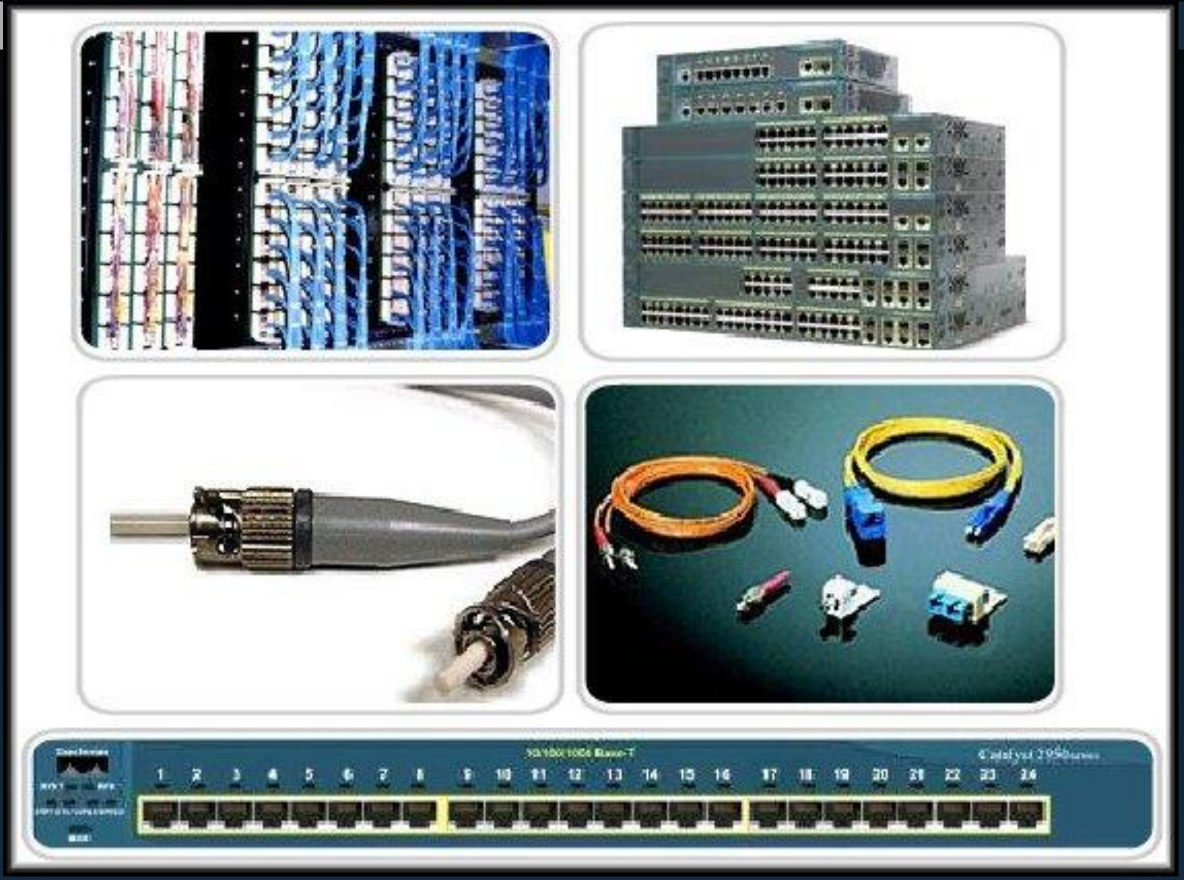

- Same addressing scheme
- Same basic frame format

Standards and Implementation

- History:
  - First LAN was Ethernet, designed at Xerox.
  - 1980: First Ethernet standard published by DIX (Digital, Intel, Xerox).
  - 1985: IEEE modified the Ethernet standard and published as 802.3.

Standards and Implementation

- 802.3 OSI Model Compatibility:
  - Needs of Layer 1.
  - The lower portion of Layer 2 of the OSI model.

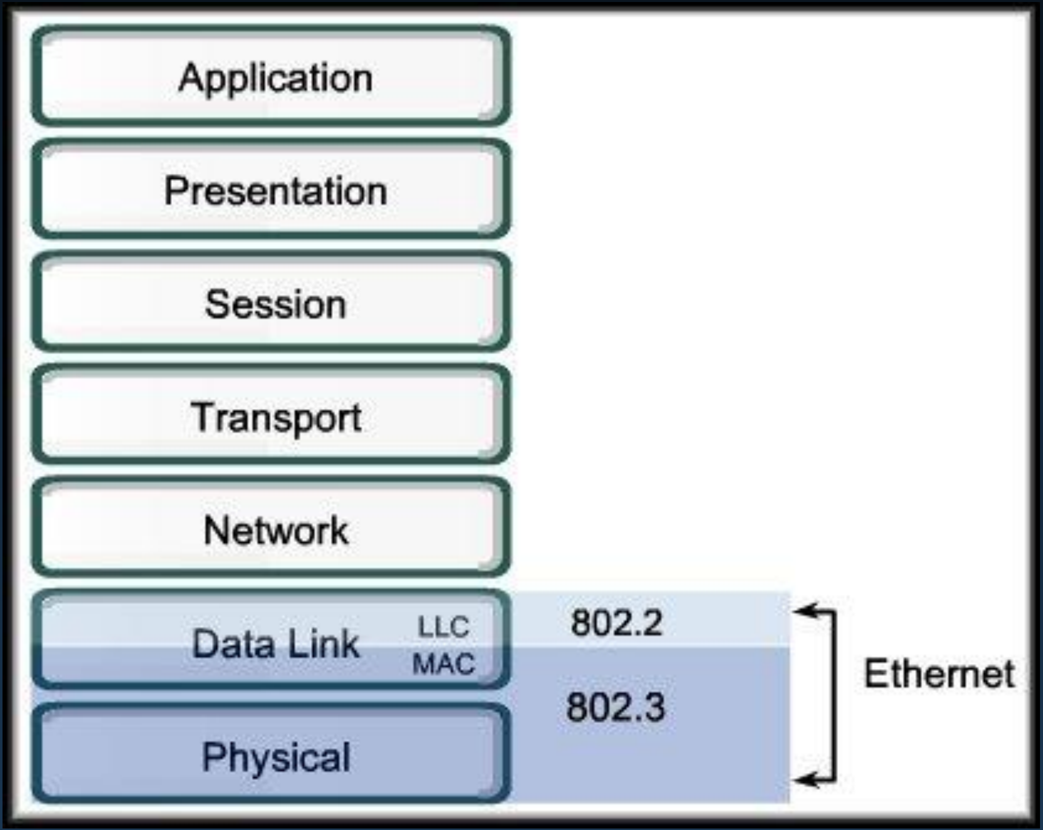

Standards and Implementation

- Layer 2 divided into two distinct areas of functionality or sub-layers.
  - Logical Link Control (LLC) – 802.2:
    - To communicate with the Network Layer.
  - Media Access Control (MAC) – 802.3:
    - To handle MAC addressing, framing and communication with the Physical Layer.

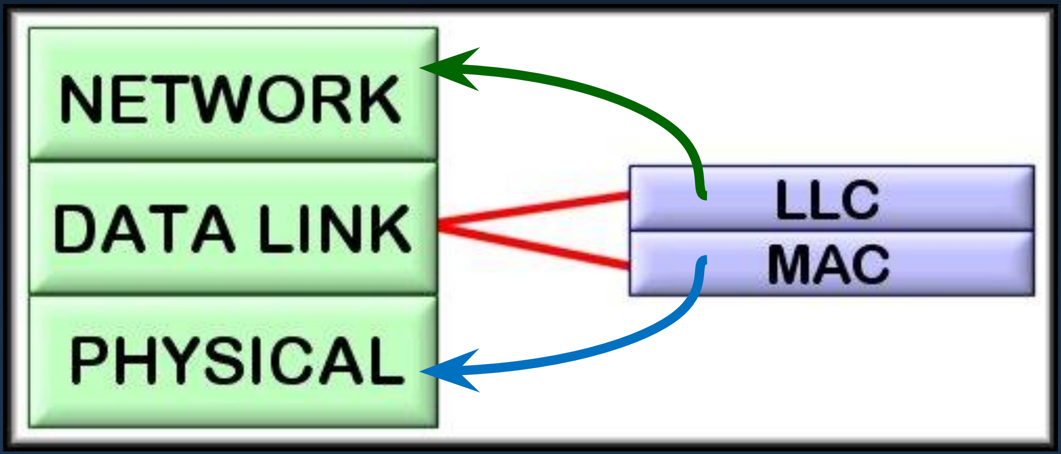

Standards and Implementation

- Layer 1 limitations were addressed at Layer 2.

Layer 1 Limitations	Layer 2 Functions
Cannot communicate with upper layers	Connects to upper layers via Logical Link Control (LLC)
Cannot identify devices	Uses addressing schemes to identify devices
Only recognizes streams of bits	Uses frames to organize bits into groups
Cannot determine the source of a transmission when multiple devices are transmitting	Uses Media Access Control (MAC) to identify transmission sources

Standards and Implementation

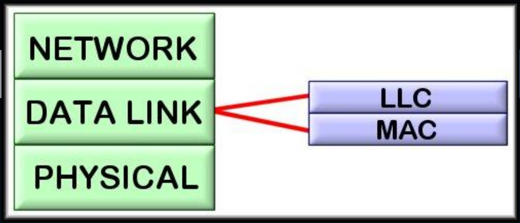

- Logical Link Control (LLC) – 802.2:
  - Communicates with the upper layers independent of the type of physical media.
- Media Access Control (MAC) – 802.3:
  - Controls access to the media.
  - Not the only standard that performs the same function.

802.2 Logical Link Control (LLC)								
MAC	802.3 Ethernet	802.5 Token Ring	802.6 MAN	802.7 Broadband	802.8 Fiber Optic	802.11 Wireless	802.15 Wireless PAN	802.16 Wireless Broadband

Standards and Implementation

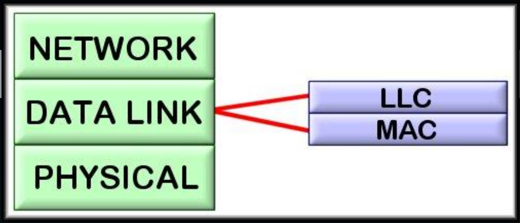

Ethernet 802.3

802.2 Logical Link Control (LLC)

802.3 Media Access Control (MAC)

Ethernet	Distance	Connector
10BASE5	500m	Coax
10BASE2	185m	Coax
10BASE-T	100m	UTP-RJ45
10BASE-TX	100m	UTP-RJ45
10BASE-CX	100m	STP-DB9
1000BASE-T	100m	UTP-RJ45
1000BASE-SX	550m	MM Fiber -SC
1000BASE-LX	5000m	MM/SM Fiber-SC

Standards and Implementation

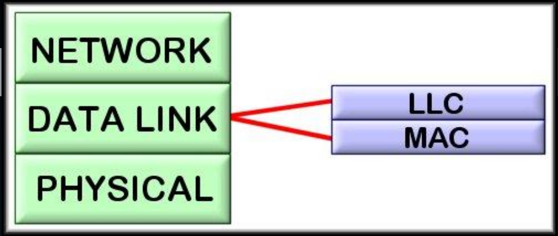

- Logical Link Control (LLC) – 802.2:
  - Prepares the data for the upper layers.
  - Allows running multiple network protocols on the same machine. Each protocol is assigned a specific ID.
  - Implemented mainly in software.
- Media Access Control (MAC) - 802.3:
  - Creates the frame and addresses the frame with the source and destination MAC address.
  - Checks for any errors using the FCS field.
  - Controls the assignment of frames on the media.
  - Controls the recovery of the media due to collisions.
  - Implemented mainly in hardware.

Physical Implementations

- Most of the traffic on the Internet originates and ends with Ethernet connections.
- When optical fiber media was introduced, Ethernet adapted to this new technology.
- The success of Ethernet is due to the following factors:
  - Simplicity and ease of maintenance
  - Ability to incorporate new technologies
  - Reliability
  - Low cost of installation and upgrade

Communication Through the LAN

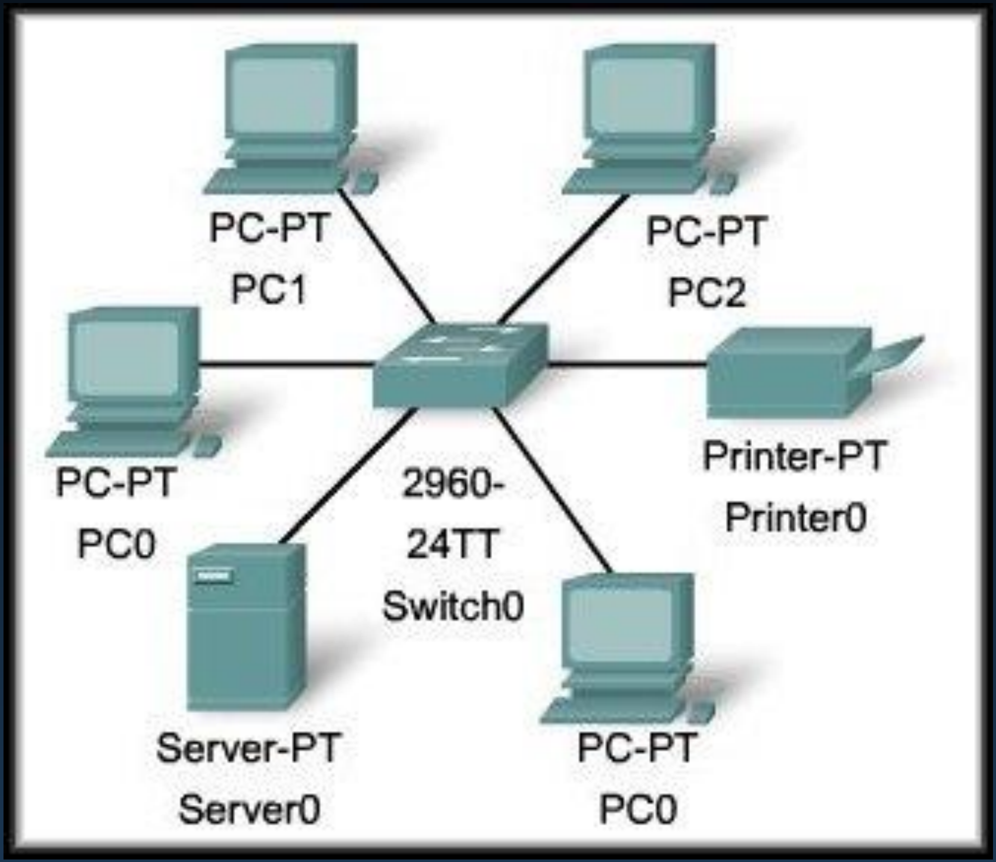

Historical Ethernet

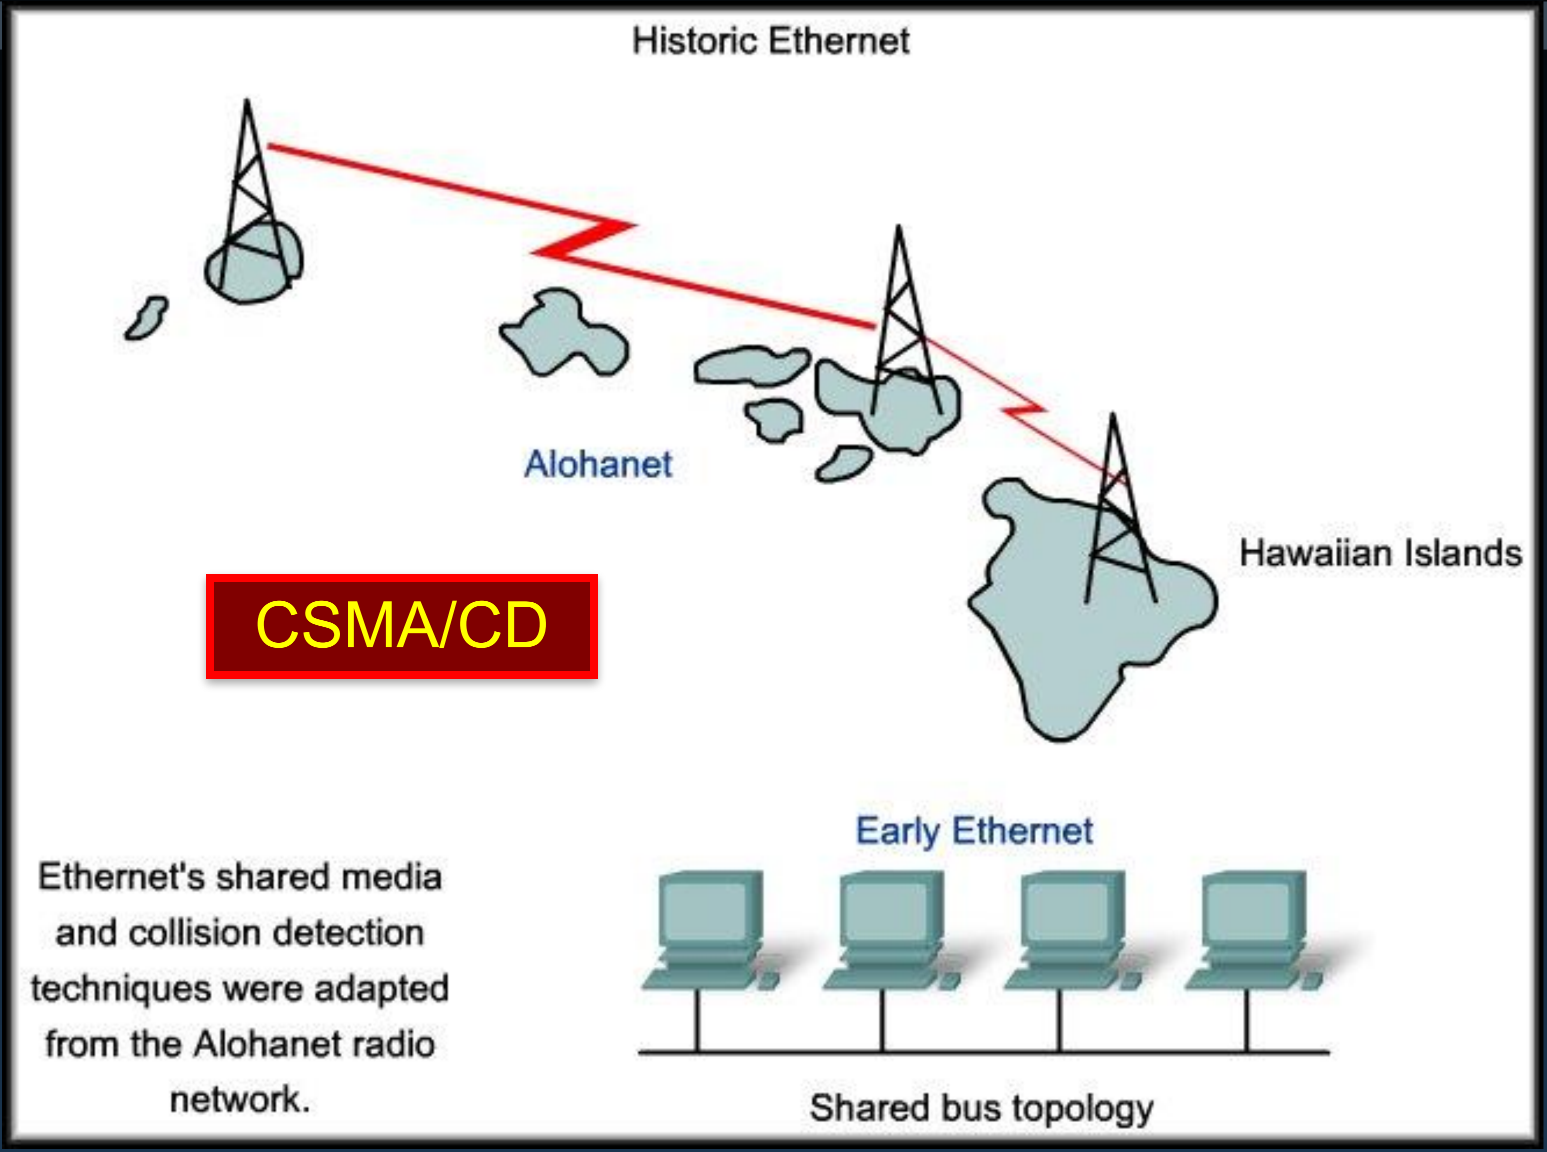

Early Media

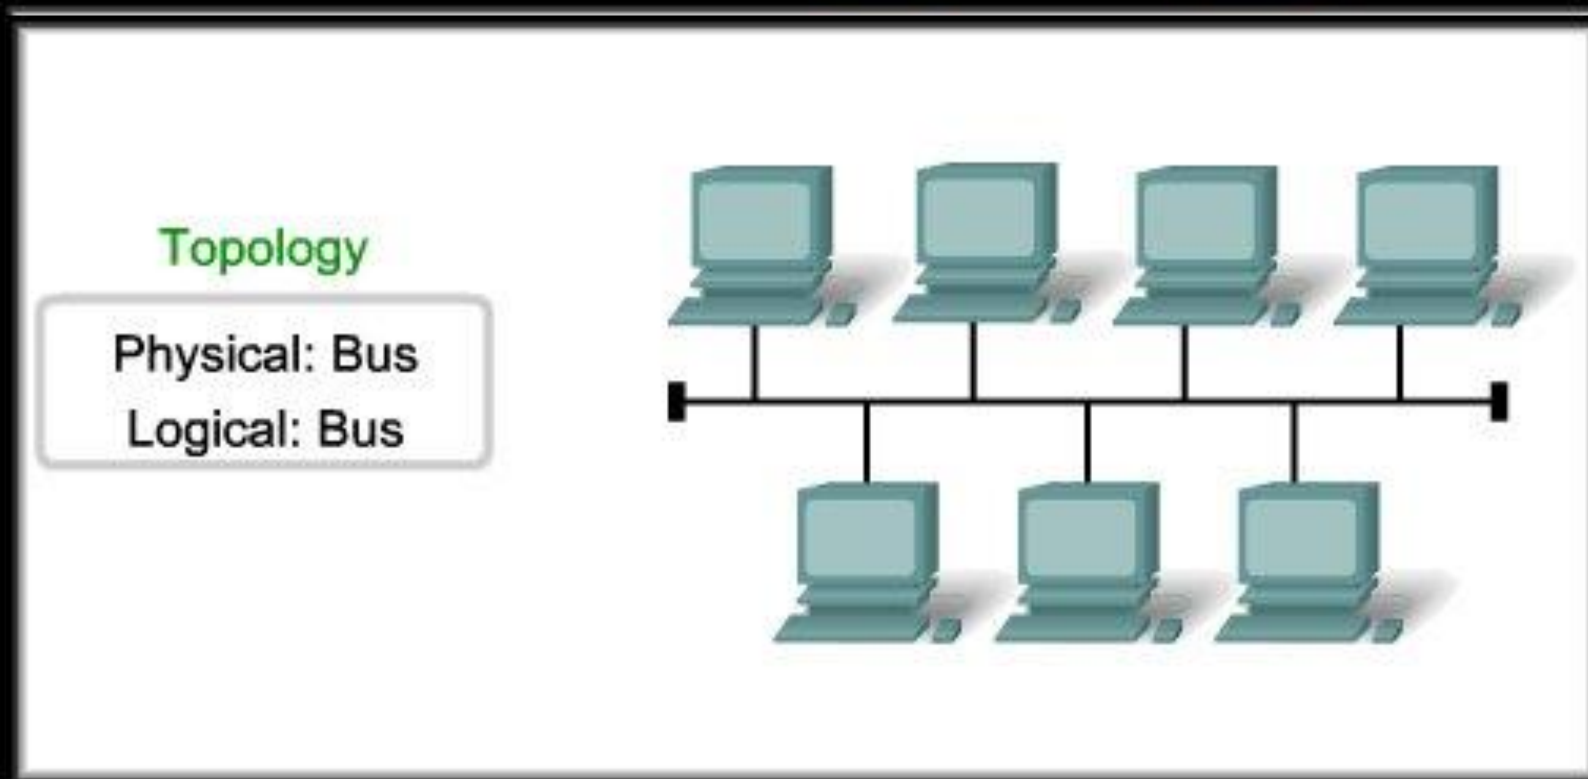

UT Coaxial Cable 45

Legacy Ethernet

Hub

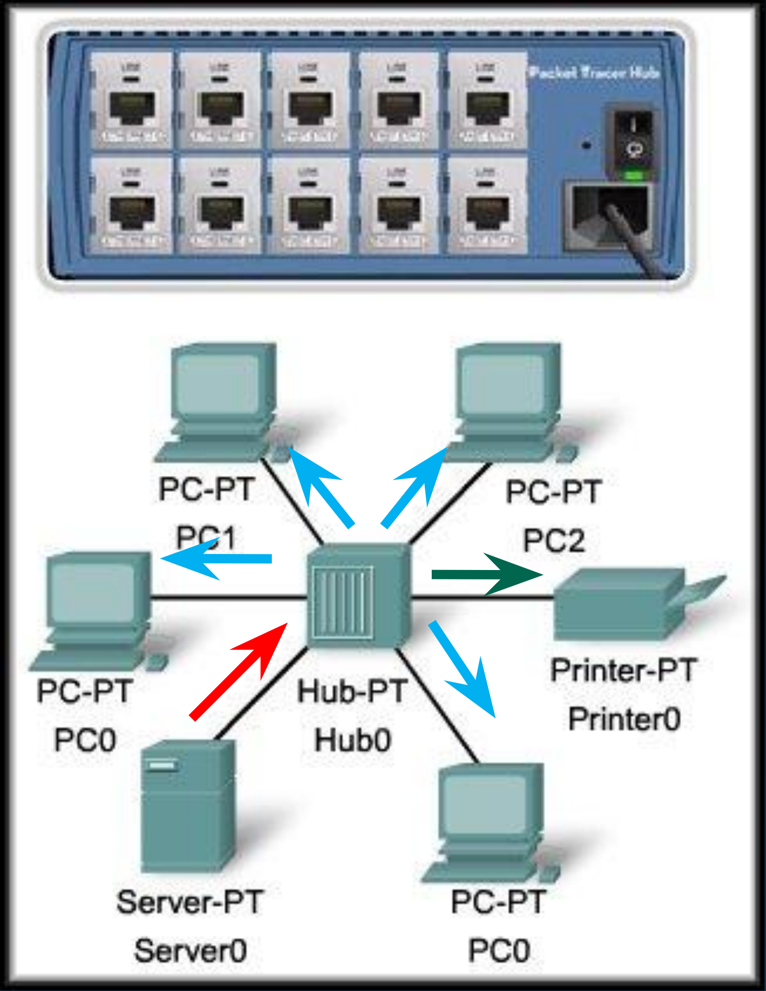

Switch

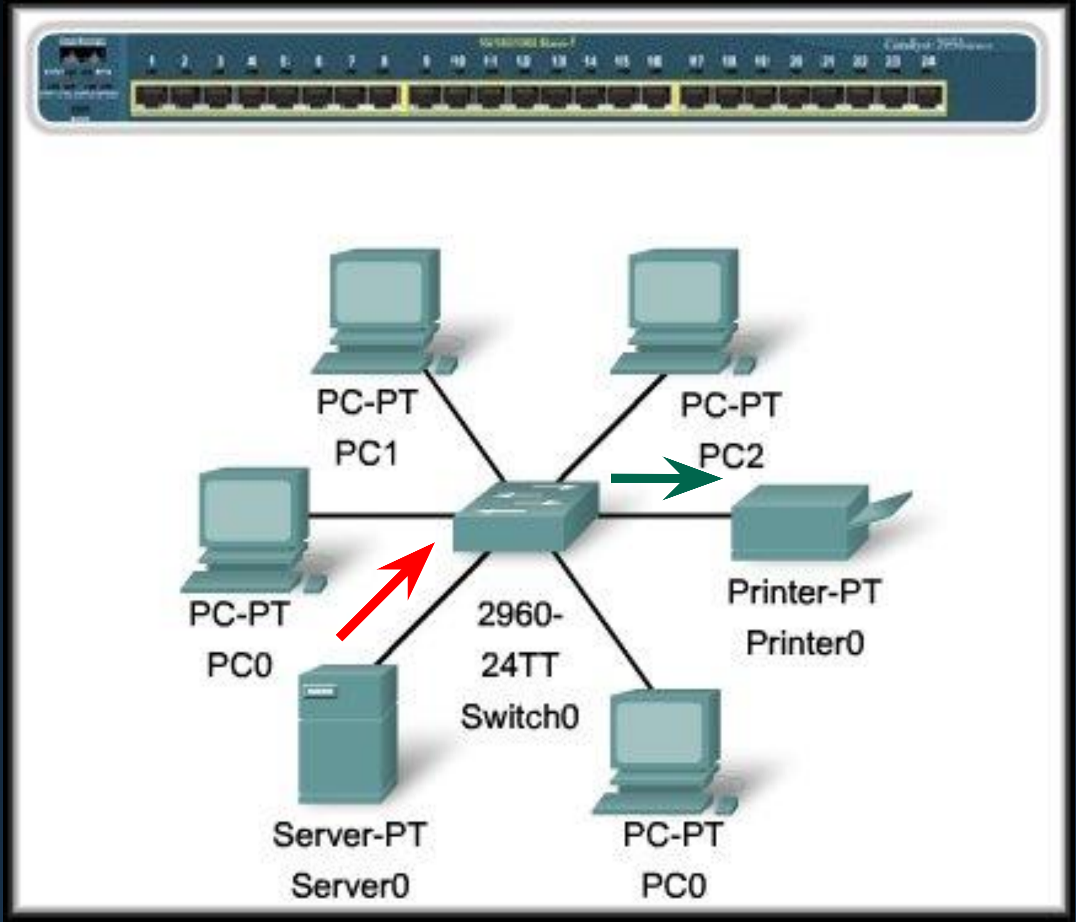

Legacy Ethernet

- Half Duplex:
  - One way traffic.
  - Necessary on a shared media.
  - Only one device can transmit at a time.
  - Collisions occur.

Hub

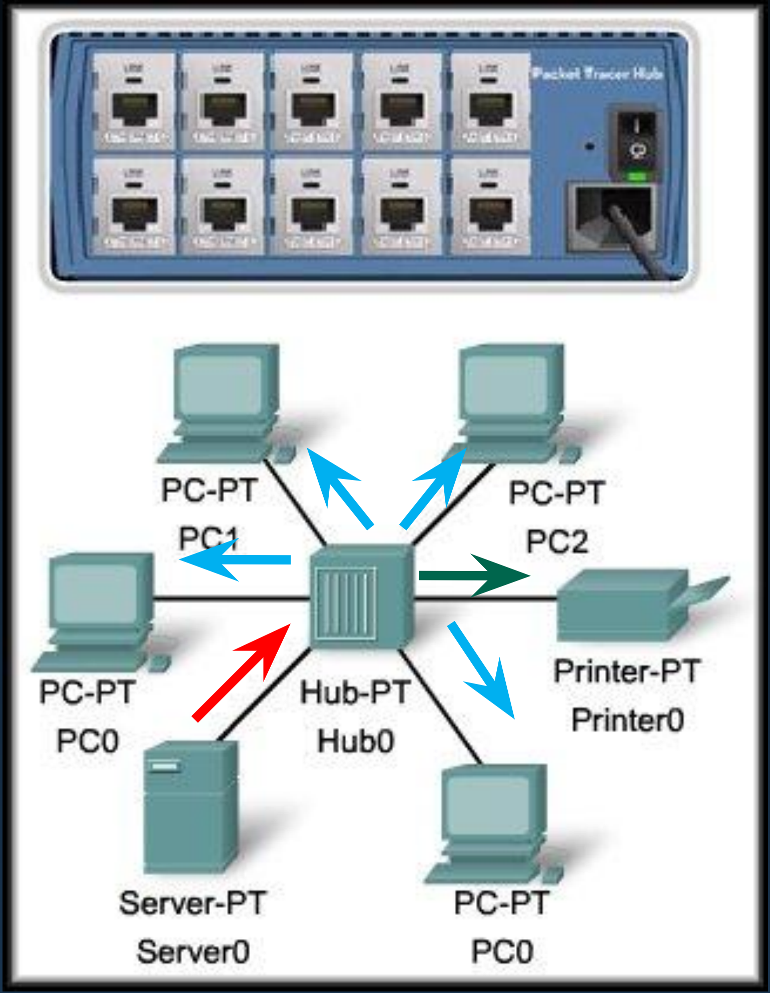

Legacy Ethernet

Switch

- Full Duplex:
  - Two way traffic.
  - Not a shared media.
  - Dedicated switch connection.
  - A device can transmit and receive at the same time.
  - No Collisions.

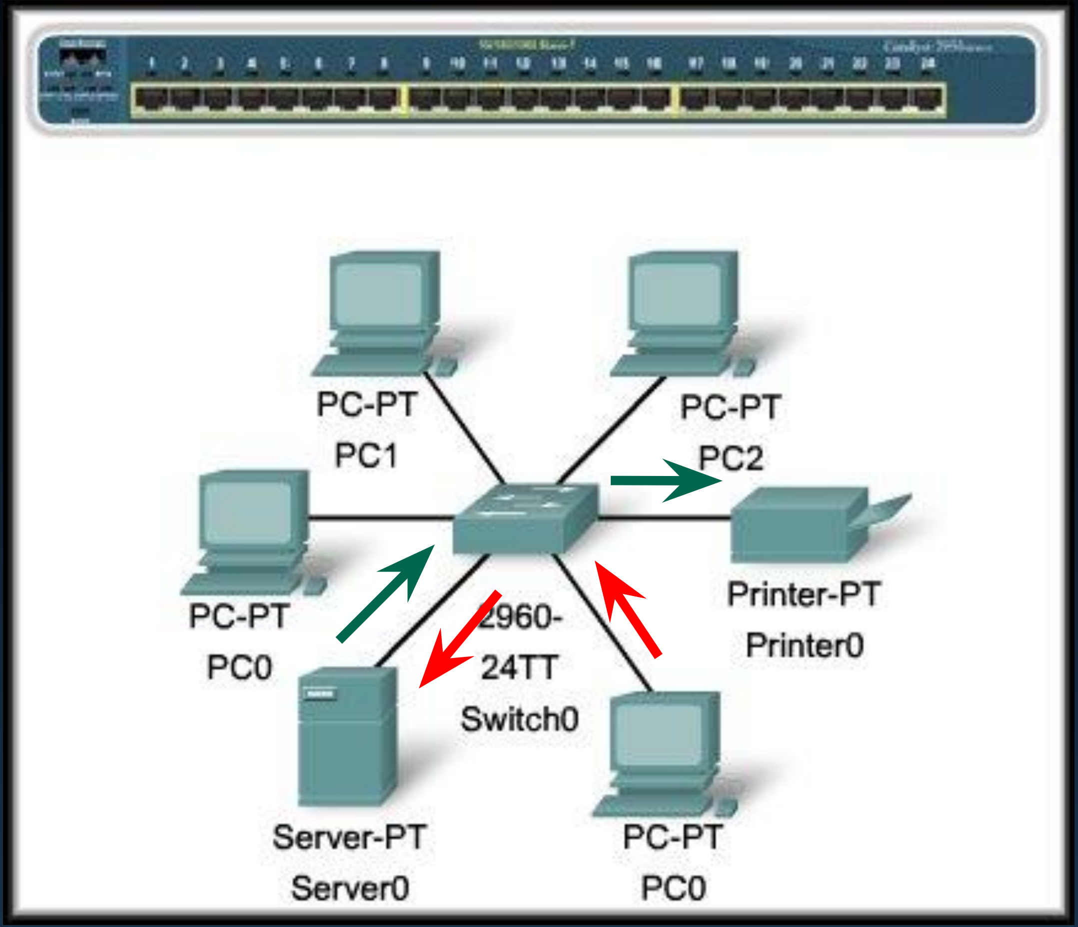

Legacy Ethernet

- Ethernet with hubs is designed to work with collisions.
  - Collisions occur when devices transmit at the same time.
  - Managed by CSMA/CD.
  - As more devices are added, more collisions occur.
  - As more collisions occur, network performance degrades.
  - Half Duplex communication.

- Ethernet with switches is designed to eliminate collisions.
  - Each device attached to switch only receives frames destined for that device.
  - Full Duplex communication.

Moving to 1 Gbps and Beyond

- Gigabit Ethernet is used to describe implementations that provide bandwidth of 1000 Mbps (1 Gbps) or greater.
- Built on the full-duplex capability and the UTP and fiber-optic media technologies of earlier Ethernet.

New networking services require high bandwidth LANs.

Does not always mean replacement of existing switches and cables .

Moving to 1 Gbps and Beyond

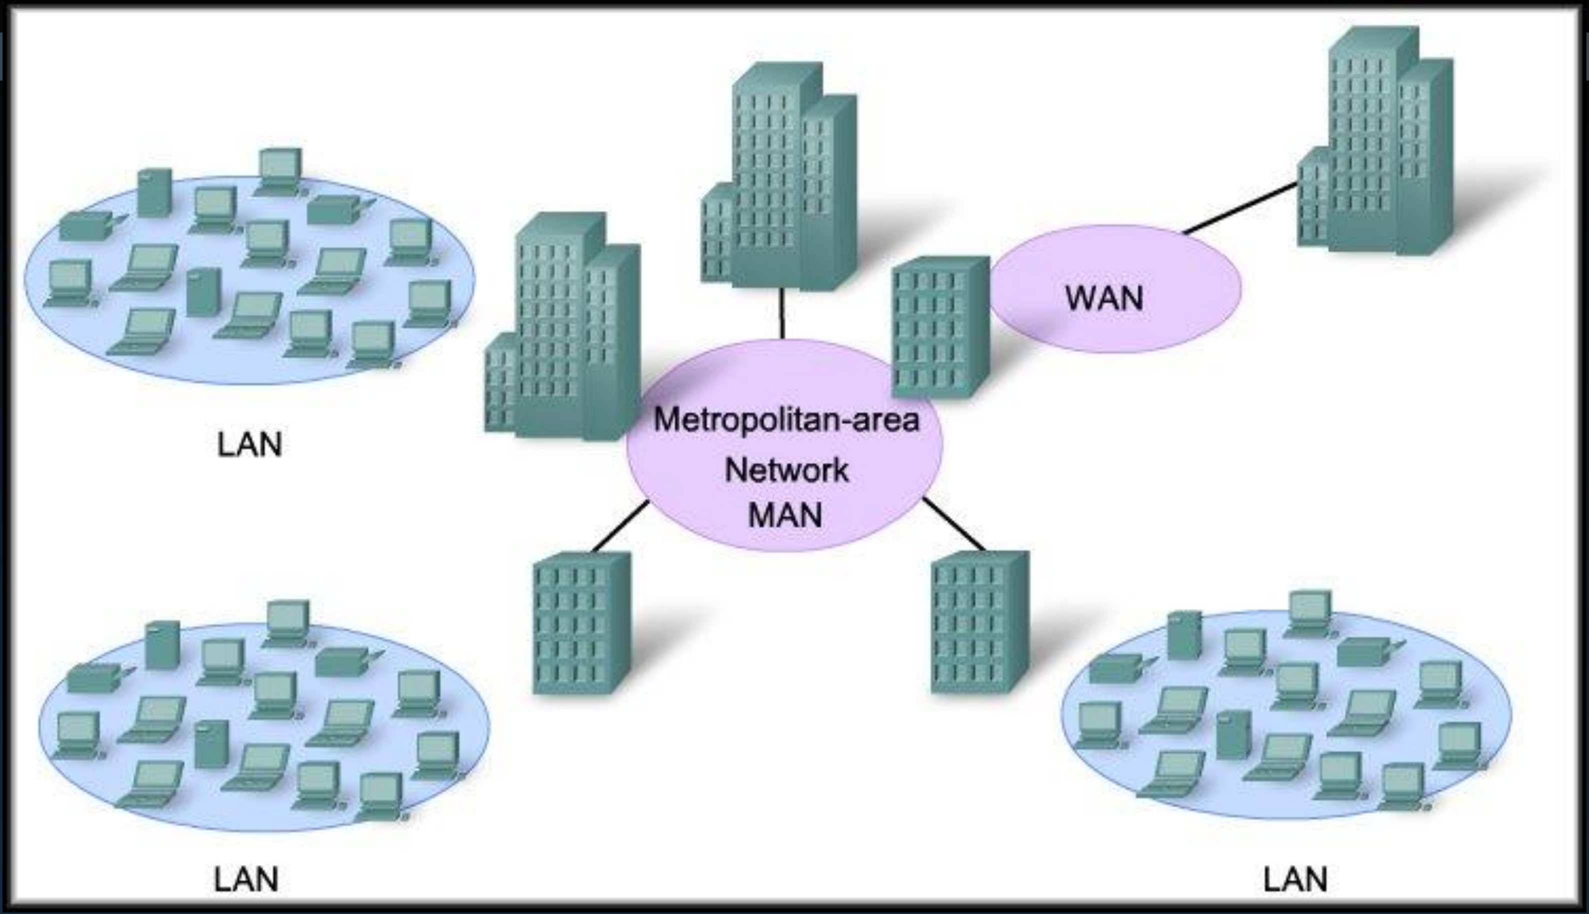

- Increased cabling distances enabled by the use of fiber-optic cable in Ethernet-based networks has resulted in a blurring of the distinction between LANs and WANs.

Ethernet Frame

Field size in bytes

IEEE 802.3

Preamble	Start of Frame delimiter	Destination Address	Source Address	Length/ Type	802.2 Header and Data	Frame Check Sequence
7	1	6	6	2	46 to 1500	4

Ethernet

Preamble	Destination Address	Source Address	Type	Data	Frame Check Sequence
8	6	6	2	46 to 1500	4

Encapsulating the Packet

- The Ethernet protocol defines the frame format.
  - Adds headers and trailers around the Layer 3 packet.

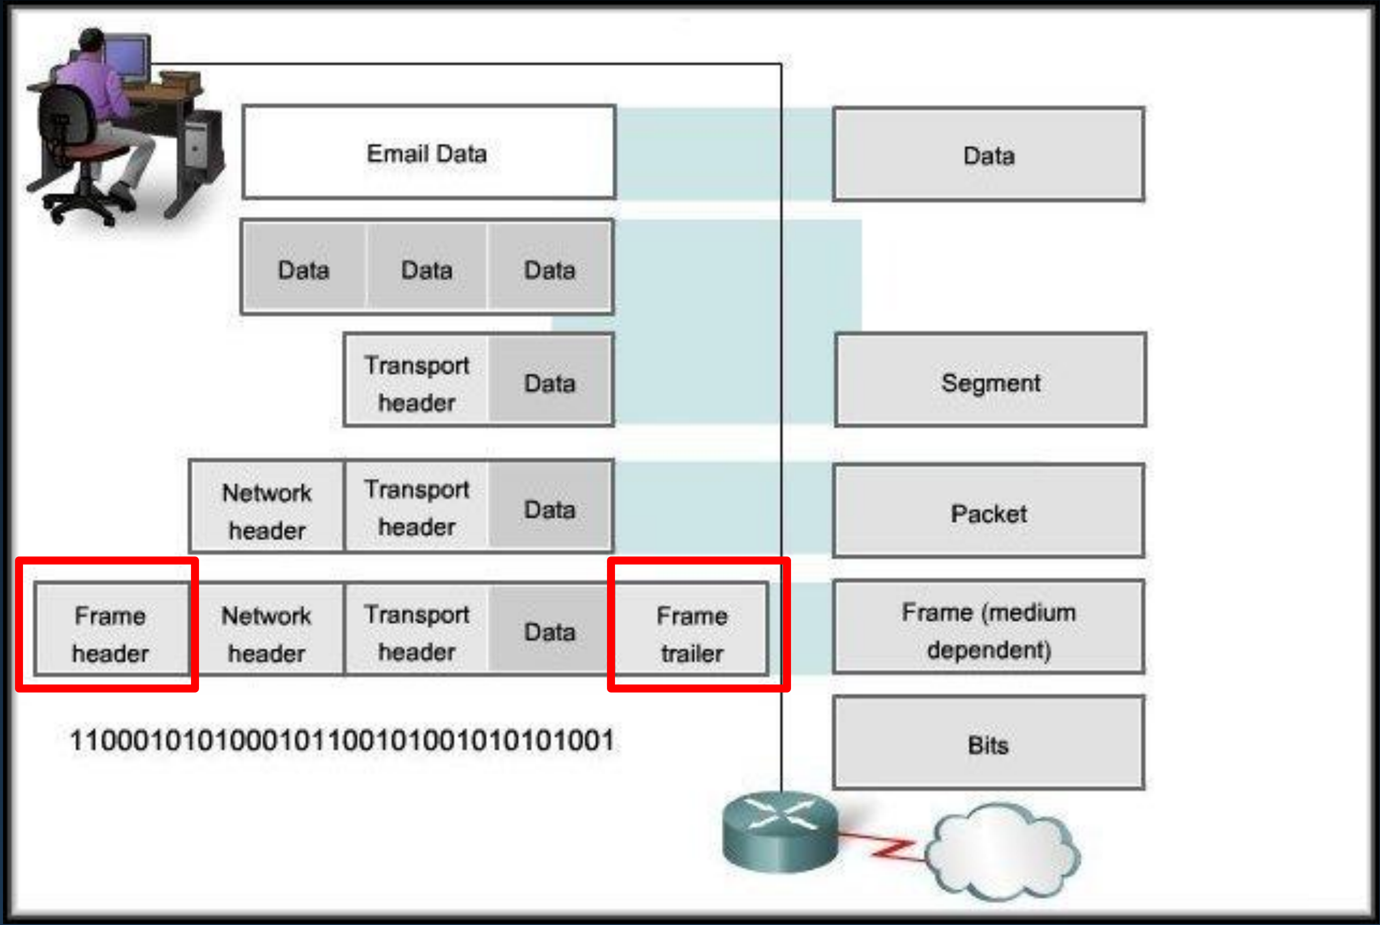

Encapsulating the Packet

- The IEEE 802.3 Ethernet Frame format:
  - Minimum Size: 64 Bytes
  - Maximum Size: 1518 Bytes
  - *If the frame is less than the minimum or greater than the maximum, it is considered corrupt and will be dropped.*

LENGTH OF FIELD IN BYTES						
7	1	6	6	2	46 – 1500	4
Preamble	Start of Frame Delimiter	Destination MAC Address	Source MAC Address	Length or Type	Data and Pad	FCS

Header (Preamble through Length or Type) — Trailer (FCS)

Encapsulating the Packet

LENGTH OF FIELD IN BYTES						
7	1	6	6	2	46 – 1500	4
Preamble	Start of Frame Delimiter	Destination MAC Address	Source MAC Address	Length or Type	Data and Pad	FCS

- Preamble and Start of Frame Delimiter (SFD) – 8 bytes:
  - Used to synchronize the NIC with the media in preparation for receiving a frame.
  - Is not considered part of the frame length.
  - Will not appear in any capture of the frame.

Encapsulating the Packet

LENGTH OF FIELD IN BYTES						
7	1	6	6	2	46 – 1500	4
Preamble	Start of Frame Delimiter	Destination MAC Address	Source MAC Address	Length or Type	Data and Pad	FCS

- Destination MAC Address – 6 bytes:
  - Identifies the node that is to receive the frame.
  - A receiving device compares its MAC address to the contents of this field.
  - If the addresses match, the frame is accepted.
  - Also used by switches to determine the interface to be used to forward the frame.

Encapsulating the Packet

LENGTH OF FIELD IN BYTES						
7	1	6	6	2	46 – 1500	4
Preamble	Start of Frame Delimiter	Destination MAC Address	Source MAC Address	Length or Type	Data and Pad	FCS

- Source MAC Address – 6 bytes:
  - Identifies the node that originated the frame.
  - Also used by switches to add addresses to their internal Port / MAC address tables.

Encapsulating the Packet

LENGTH OF FIELD IN BYTES						
7	1	6	6	2	46 – 1500	4
Preamble	Start of Frame Delimiter	Destination MAC Address	Source MAC Address	Length or Type	Data and Pad	FCS

- Length / Type – 2 bytes:
  - DIX used this for type, the original IEEE 802.3 standard used it for length. The later IEEE standard (Ethernet II) allows it to be used for either.
  - Ethernet II is the frame type used in TCP/IP networks.
  - If the value is greater than 1518 (0x600), it contains a code identifying the encapsulated upper layer protocol.
  - Any other value defines the length of the frame.

Encapsulating the Packet

LENGTH OF FIELD IN BYTES						
7	1	6	6	2	46 – 1500	4
Preamble	Start of Frame Delimiter	Destination MAC Address	Source MAC Address	Length or Type	Data and Pad	FCS

- Data and Pad – 46 to 1500 bytes:
  - The encapsulated data from Layer 3.
  - Most commonly an IPv4 packet.
  - If the total frame length is less than 64 bytes, the field is padded to the right with enough null characters to meet the minimum frame length.

Encapsulating the Packet

LENGTH OF FIELD IN BYTES						
7	1	6	6	2	46 – 1500	4
Preamble	Start of Frame Delimiter	Destination MAC Address	Source MAC Address	Length or Type	Data and Pad	FCS

- Frame Check Sequence (FCS)– 4 bytes:
  - Used to detect errors in a frame that may have occurred during transmission along the media.
  - The result of a Cyclic Redundancy Check (CRC) is placed in the frame by the sending node.
  - The receiving node performs the same CRC and compares the values….they should be equal.

Ethernet MAC Address

LENGTH OF FIELD IN BYTES						
7	1	6	6	2	46 – 1500	4
Preamble	Start of Frame Delimiter	Destination MAC Address	Source MAC Address	Length or Type	Data and Pad	FCS

- In order for a transmission to be received properly at the destination computer, there must be a method of uniquely identifying that host.
- A unique address is permanently programmed into ROM in each NIC *("burned in“)* when it is manufactured.
  - Because of this, the MAC Address is often referred to as the burned in (BIA) address or physical address of a machine.

Ethernet MAC Address

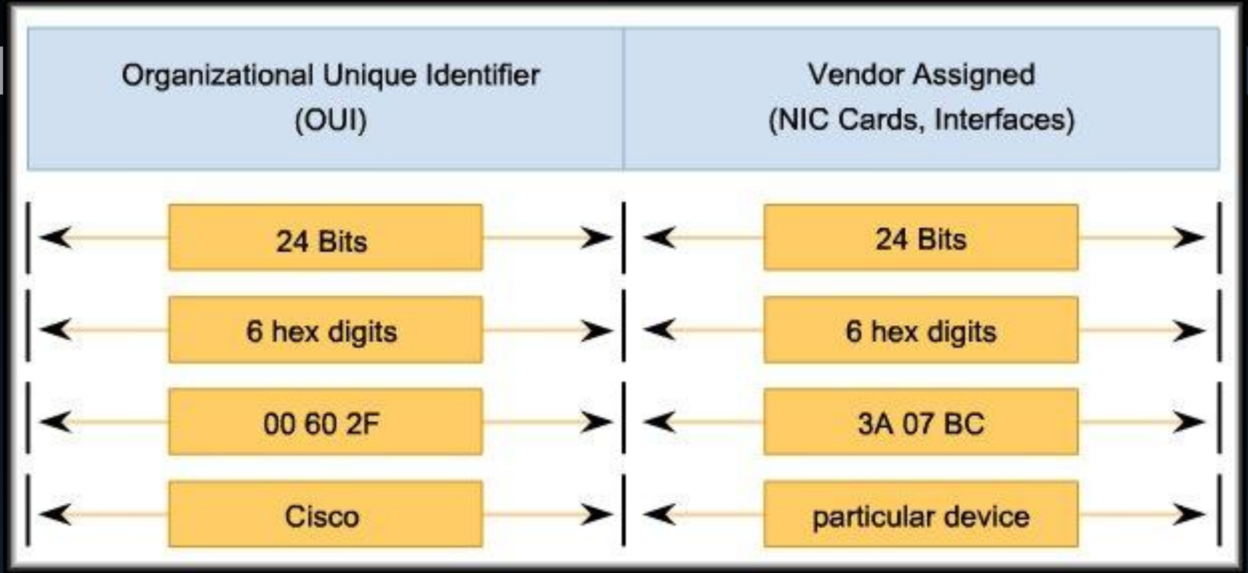

- 48 bits in length.
- Expressed as 12 hexadecimal digits.
- The first 6 hexadecimal digits, which are administered by the IEEE, identify the manufacturer or vendor and thus comprise the Organizational Unique Identifier (OUI).
- The remaining 6 hexadecimal digits comprise the interface serial number, or another value administered by the specific vendor.

Ethernet MAC Address

- The OUI and the sequential number ensure that the assigned MAC addresses remain unique.
- You will see them expressed in different ways.

Cisco MAC Address	Intel MAC Address
00-60-2F-3A-07-BC	00-20-E0-6B-17-62
00:60:2F:3A:07:BC	00:20:E0:6B:17:62
0060.2F3A.07BC	0020.E06B.1762

Ethernet MAC Address

LENGTH OF FIELD IN BYTES						
7	1	6	6	2	46 – 1500	4
Preamble	Start of Frame Delimiter	Destination MAC Address	Source MAC Address	Length or Type	Data and Pad	FCS

- When a network device matches the destination address to the address in the NIC, the NIC passes the frame up the OSI layers where the decapsulation process takes place.

The MAC address is essential to communications on a network. It is the only address that guarantees that the message will be accepted by the destination.

Hexadecimal Numbering and Addressing

- A big problem with the binary system was verbosity. In order to represent the number 202:
  - Requires 3 decimal digits (202).
  - Requires 8 bits (11001010).
- When representing large numbers, the binary system quickly becomes unwieldy.
- We can also convert from decimal to binary but the conversion is not a trivial task.

Hexadecimal Numbering and Addressing

- The hexadecimal numbering system addresses both of these issues:
  - It is compact.
  - It easy to convert from binary to hexadecimal and vice versa.
  - Because of this most of the computers in use today use the hexadecimal system.

Hexadecimal Numbering and Addressing

- You can expect to see hex numbers represented in documents and the web in different ways:
  - $23A9_{16}$ 2eb6H 0FCDh ‘7b’
  - 0xE0 0x23facb92 %0a000c834a >34ce
  - 10-00-5a-29-16-ab (NIC – e.g. ipconfig –all)
  - 00:00:0C:48:8C:11 (NIC – e.g. router MAC display
  - #FFFFFF (Web RGB Colour Code)
  - 1080:0:0:0:8:800:200C:417A (IP Version 6 Address)

Hexadecimal Numbering and Addressing

Binary Data Organization

BIT: 0 and 1

NIBBLE:

Bit No.	3	2	1	0
Value	8	4	2	1

BYTE:

Bit No.	7	6	5	4	3	2	1	0
Value	128	64	32	16	8	4	2	1

WORD:

Bit No.	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Value	32,768											16	8	4	2	1

Hexadecimal Numbering and Addressing

- Hexadecimal and Binary:
  - Hexadecimal numbering is base 16 and requires a way to represent the values 0 to 15:

Hex	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
Decimal	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15

  - Each hexadecimal digit is represented in binary by one nibble (4 bits).

Bit No.	3	2	1	0
Value	8	4	2	1

Hexadecimal Numbering and Addressing

- Hexadecimal and Binary:
  - A byte is 8 bits (2 nibbles).

Bit No.	7	6	5	4	3	2	1	0
Value	128	64	32	16	8	4	2	1

  - Each byte is represented by a group of 2 hexadecimal digits and each word by a group of 4 digits.
    - e.g. 0x1234, 0xBEEF, 0xDEAF, 0xDEAD, 0xFEED
      - Bytes are usually, but not always, separated by a colon (:), a dash (–) or a space.
      - 0x12:34 0xBE-EF 0xDE AF

Hexadecimal Numbering and Addressing

- Converting Hexadecimal to Binary:
  - Convert 0xCA to Binary.......
    - Convert each hexadecimal digit to its binary equivalent.
      - C = 1100 (12) A = 1010 (10)
    - Combine the nibbles in the same sequence to form the complete byte.
      - 11001010
  - 0xCA = 11001010 = 202

Hexadecimal Numbering and Addressing

- Converting Binary to Hexadecimal:
  - Convert 11001010 to Hexadecimal.....
    - Beginning at the left, divide the byte to form 4-bit nibbles.
      - 1100 1010
    - Convert each nibble to its hexadecimal equivalent.
      - 1100 = 12 = C   1010 = 10 = A
  - 11001010 = 0xCA

Hexadecimal Numbering and Addressing

Decimal and Binary equivalents of 0 to F Hexadecimal

Decimal	Binary	Hexadecimal
0	0000	0
1	0001	1
2	0010	2
3	0011	3
4	0100	4
5	0101	5
6	0110	6
7	0111	7
8	1000	8
9	1001	9
10	1010	A
11	1011	B
12	1100	C
13	1101	D
14	1110	E
15	1111	F

Selected Decimal, Binary and Hexadecimal equivalents

Decimal	Binary	Hexadecimal
0	0000 0000	00
1	0000 0001	01
2	0000 0010	02
3	0000 0011	03
4	0000 0100	04
5	0000 0101	05
6	0000 0110	06
7	0000 0111	07
8	0000 1000	08
10	0000 1010	0A
15	0000 1111	0F
16	0001 0000	10
32	0010 0000	20
64	0100 0000	40
128	1000 0000	80
192	1100 0000	C0
202	1100 1010	CA
240	1111 0000	F0
255	1111 1111	FF

Viewing the MAC Address

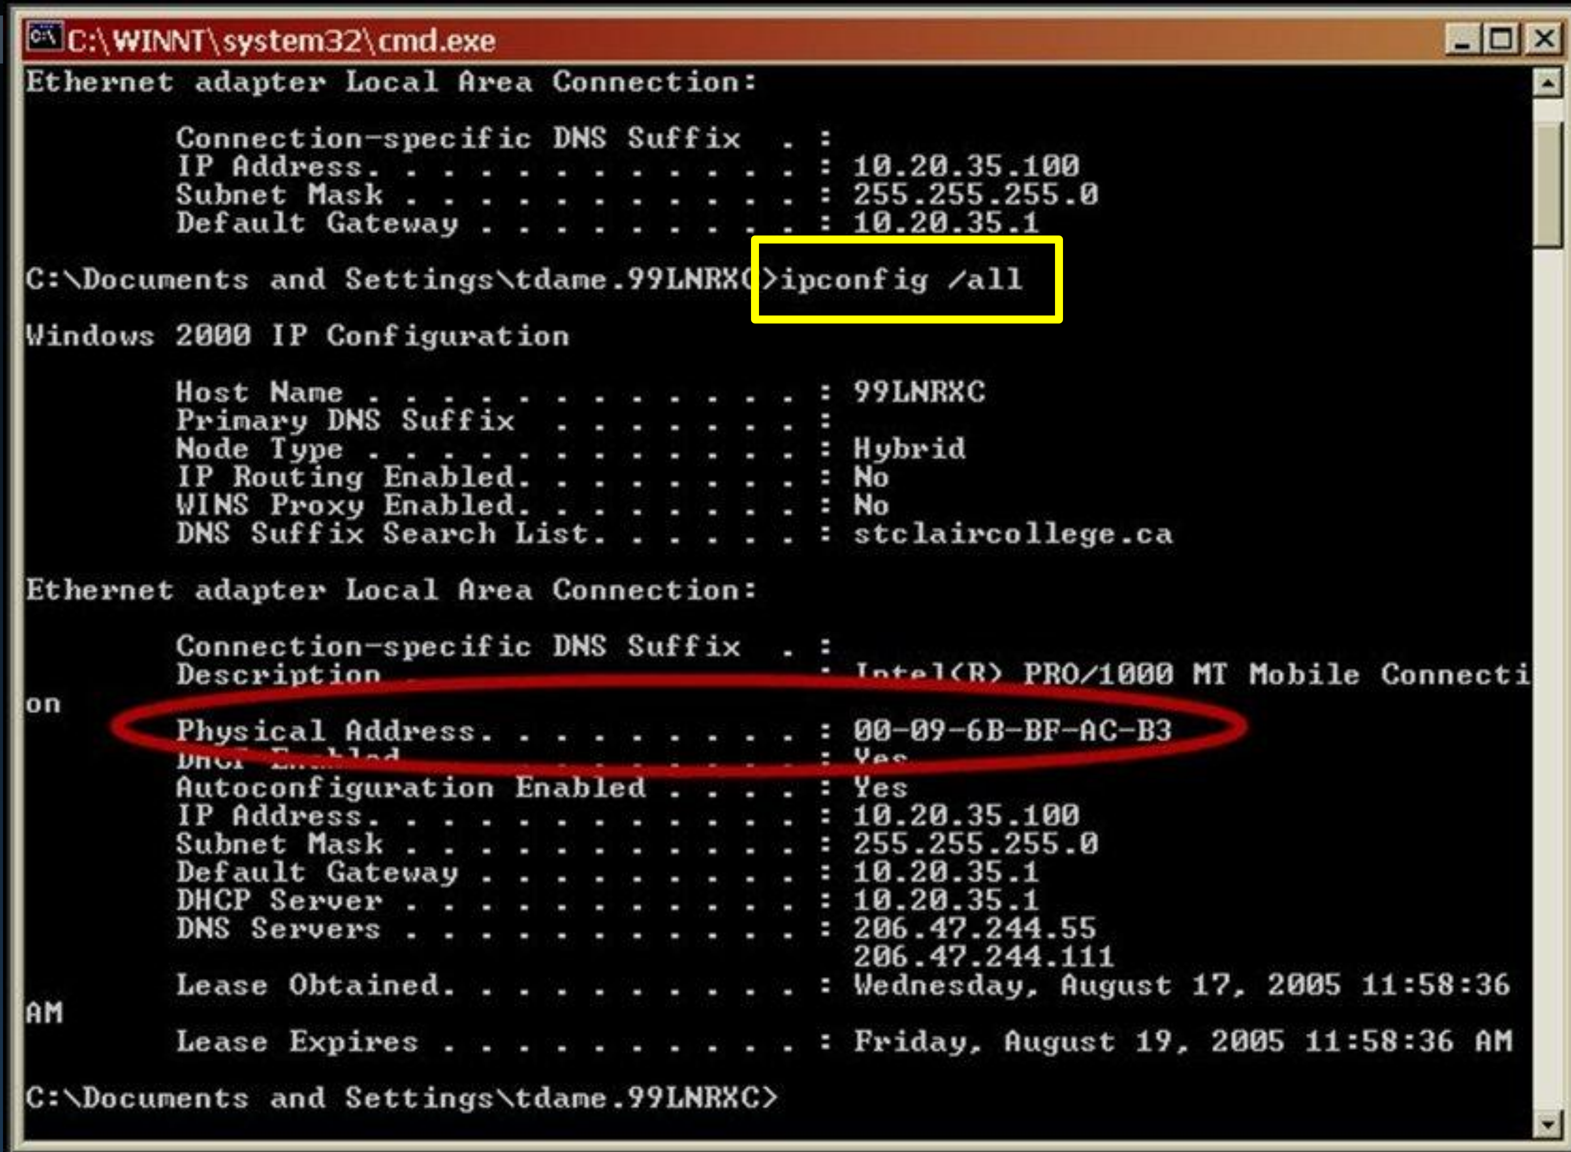

Another Layer of Addressing

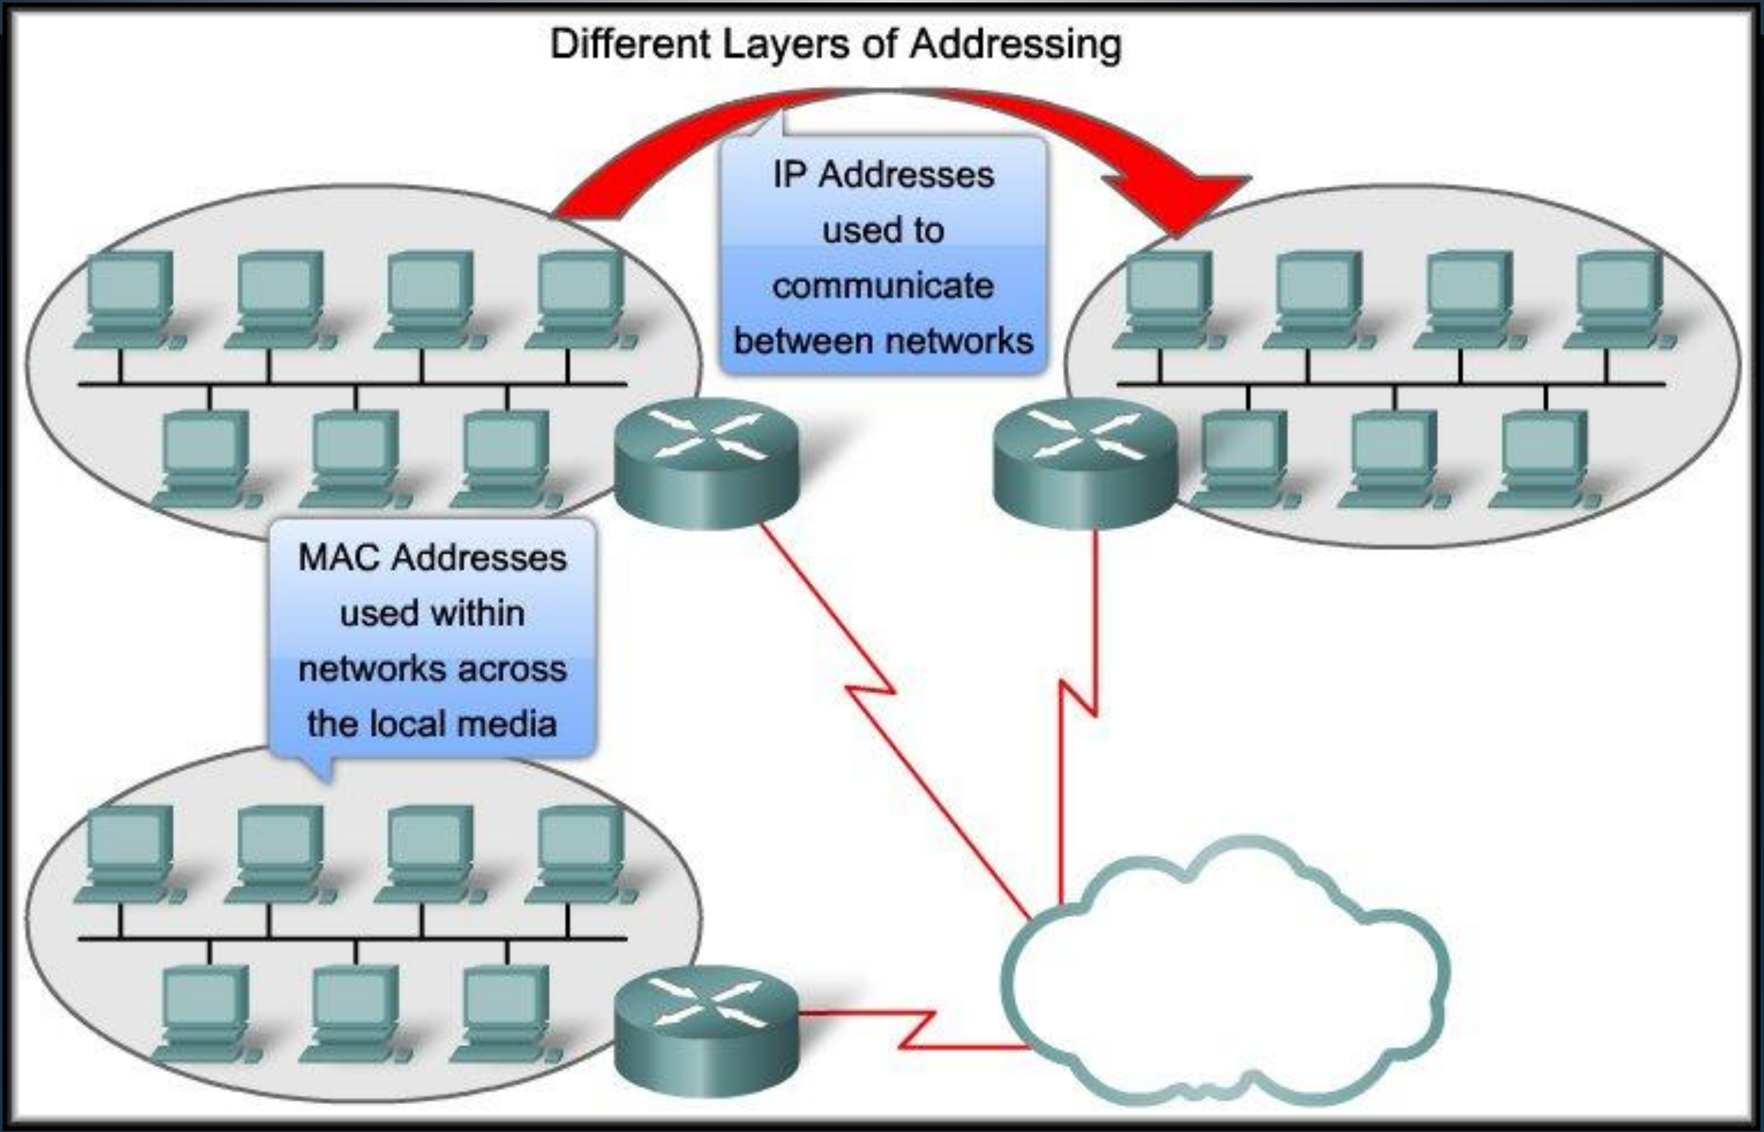

Ethernet Unicast, Multicast and Broadcast

- Different MAC addresses are used to provide different types of communication.
  - Unicast:
    - A unique address identifying a specific host.
  - Multicast:
    - An address recognized by a specific group of hosts.
  - Broadcast:
    - An address used to send information to all hosts.

Ethernet Unicast, Multicast and Broadcast

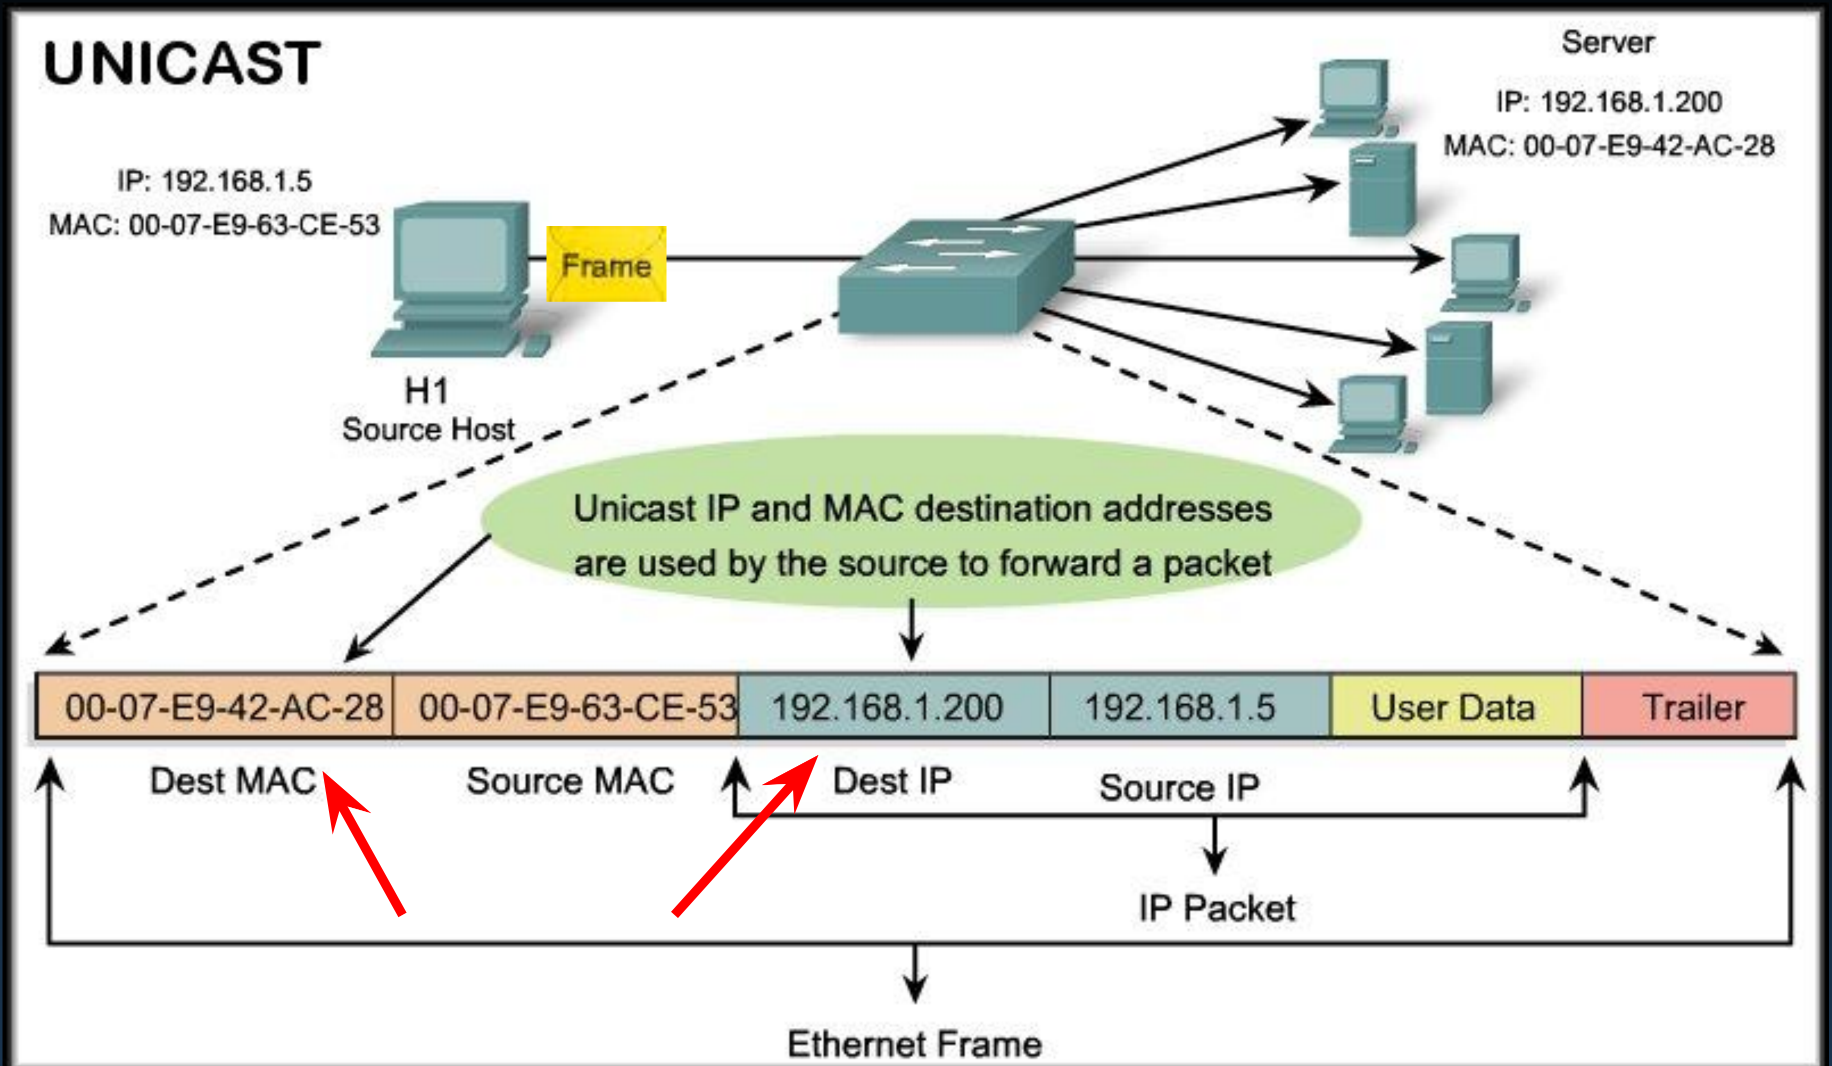

Ethernet Unicast, Multicast and Broadcast

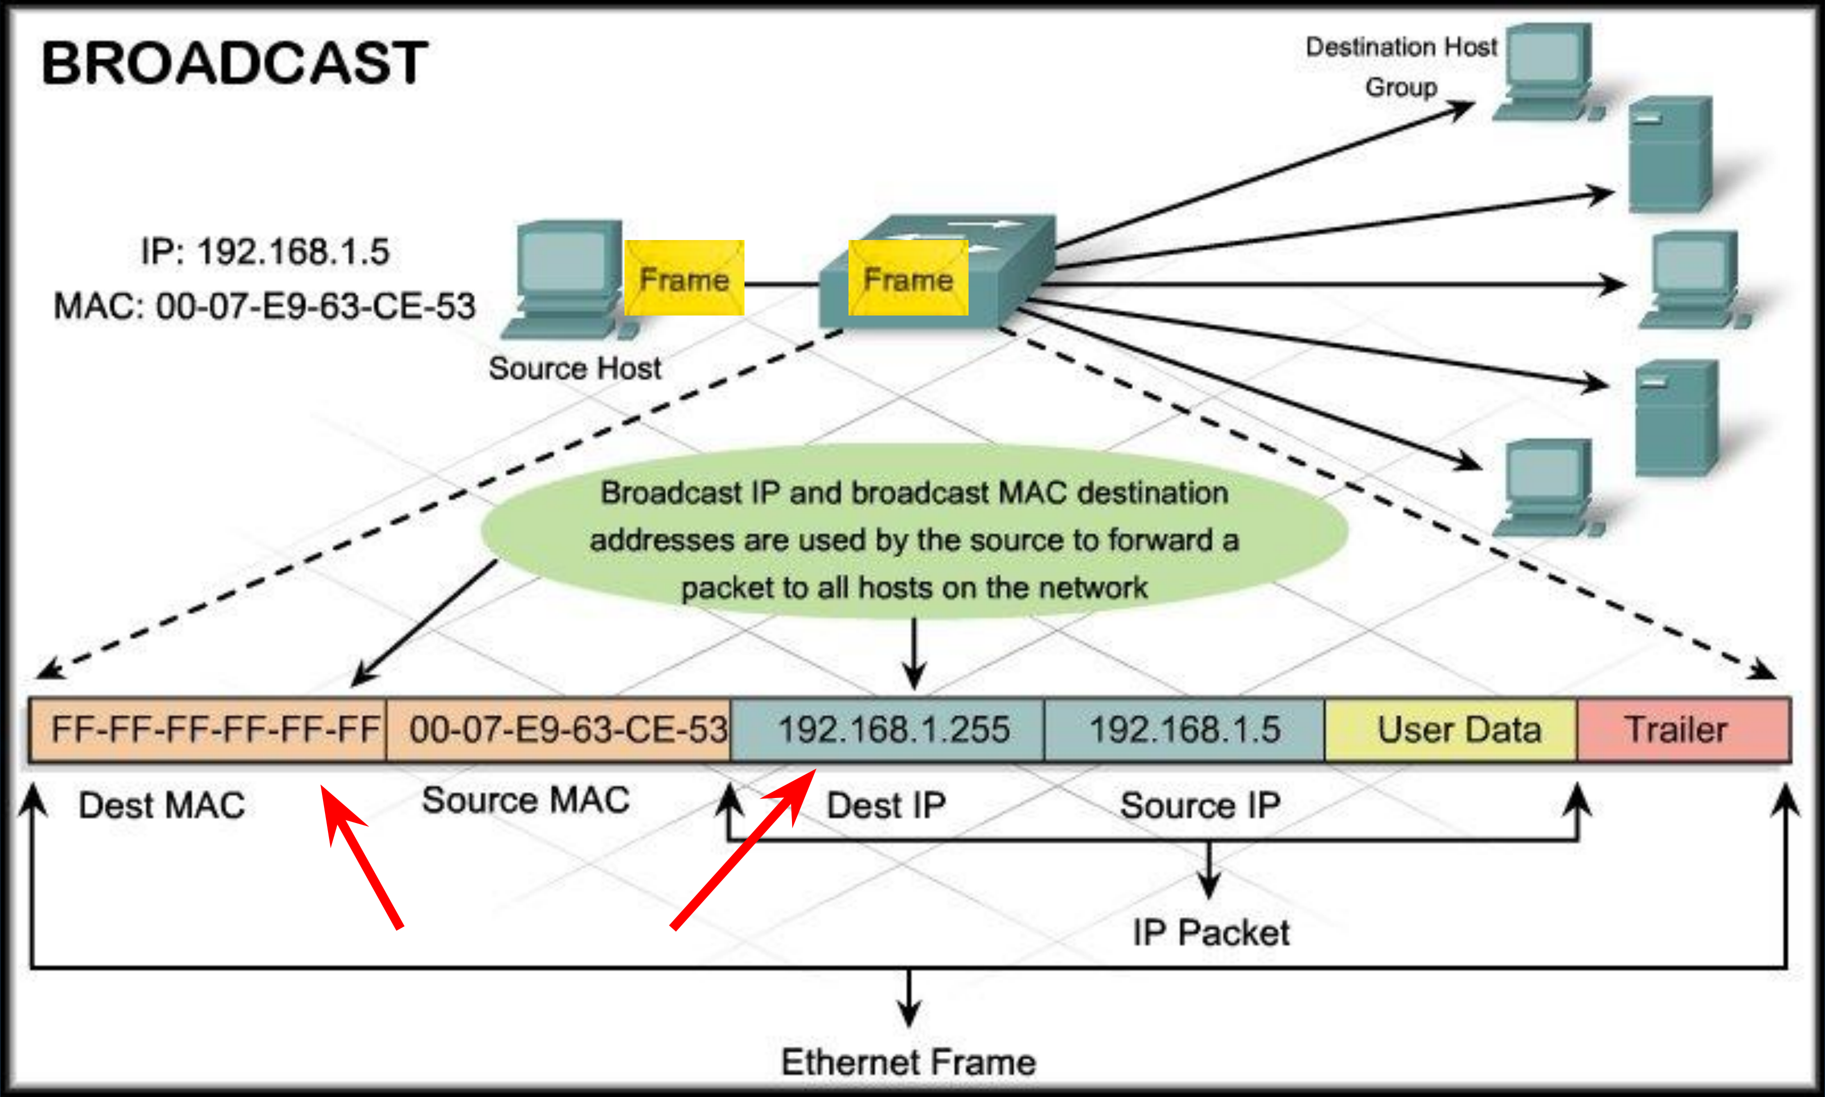

Ethernet Unicast, Multicast and Broadcast

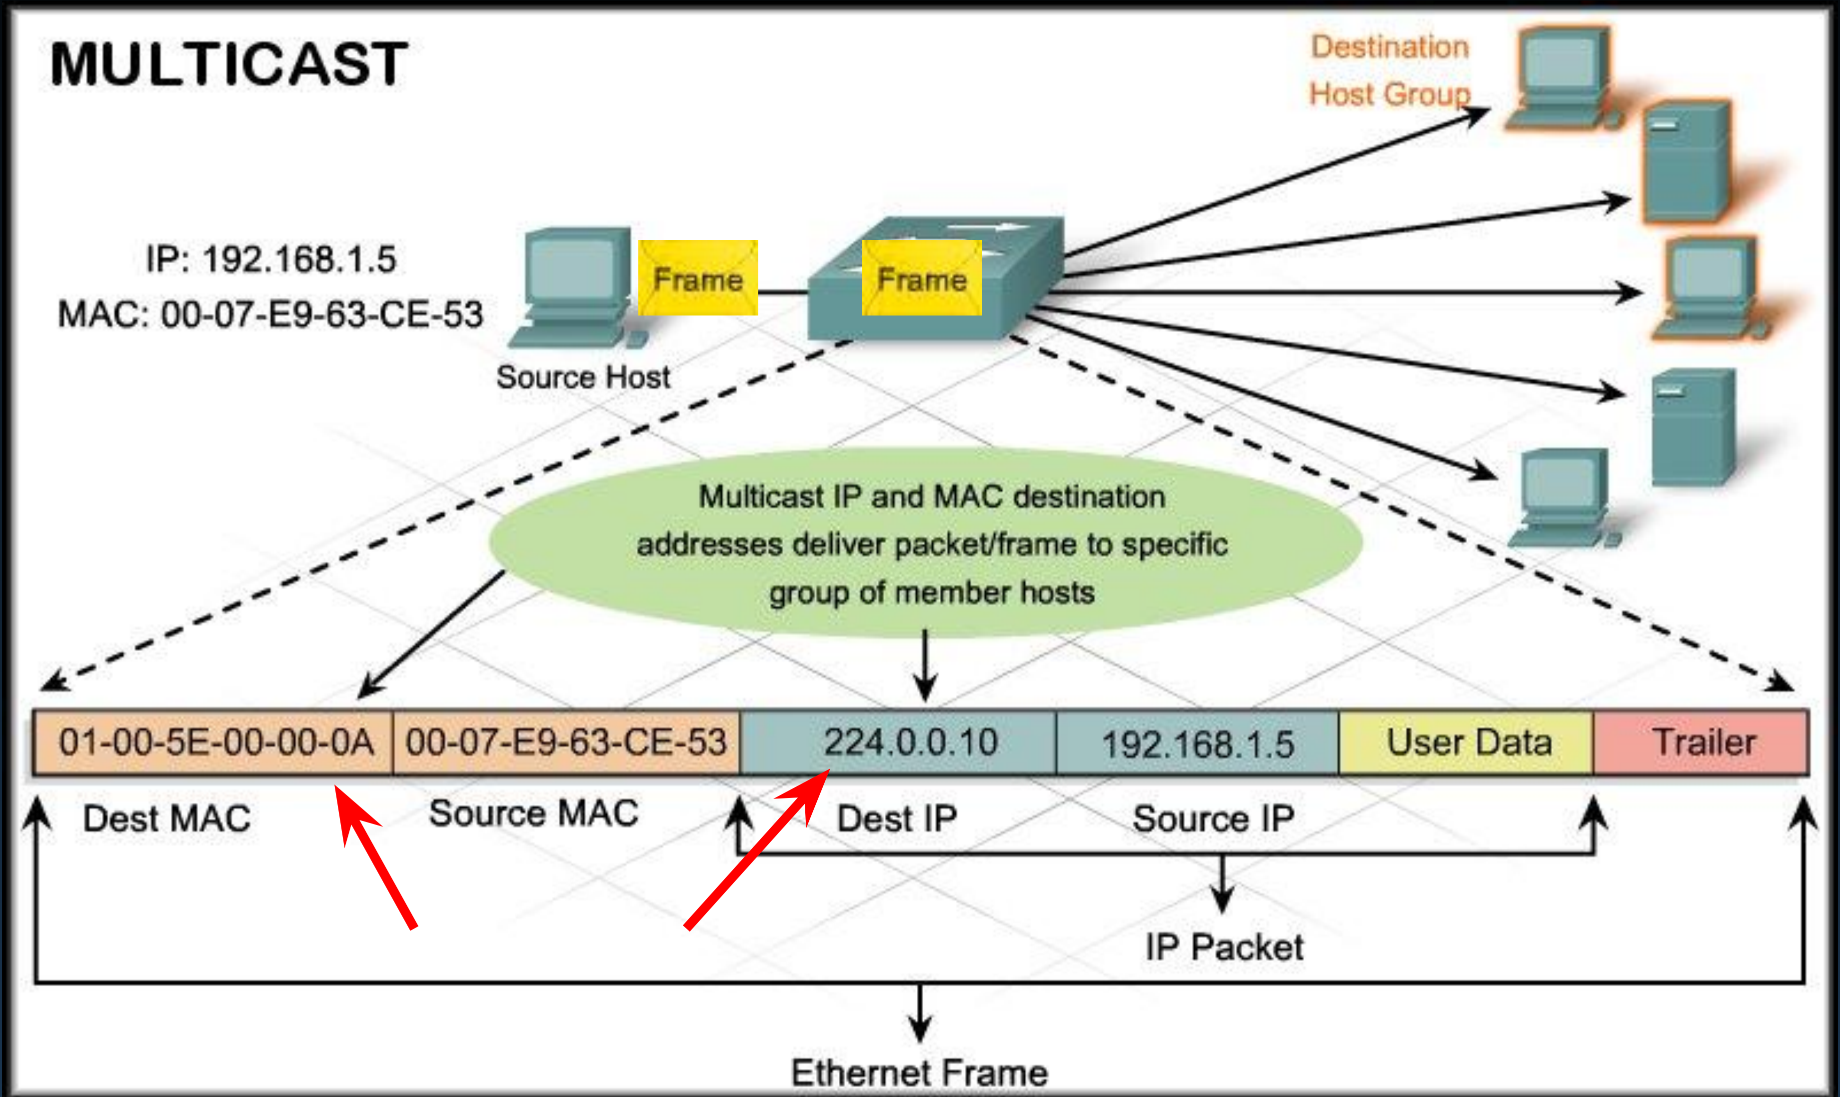

Ethernet MAC
CSMA/CD

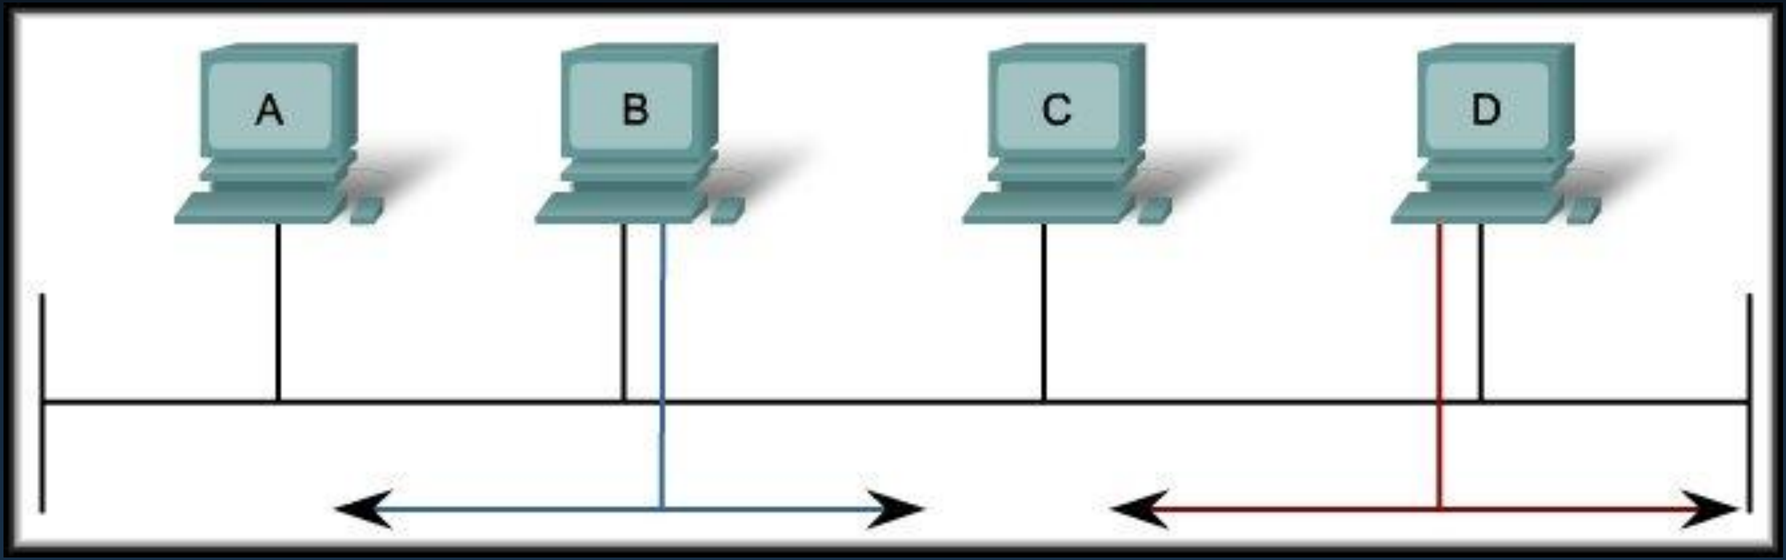

Ethernet MAC method

- In a shared media environment, all devices have guaranteed access to the medium but they have no prioritized claim on it.
- If more than one device transmits simultaneously
  - The physical signals collide.
  - The network must recover.

- *Collisions are the cost that Ethernet pays to get the low overhead associated with each transmission.*

- Ethernet uses Carrier Sense Multiple Access with Collision Detection (CSMA/CD) to detect and handle collisions and manage the resumption of communications.

CSMA/CD: The Process

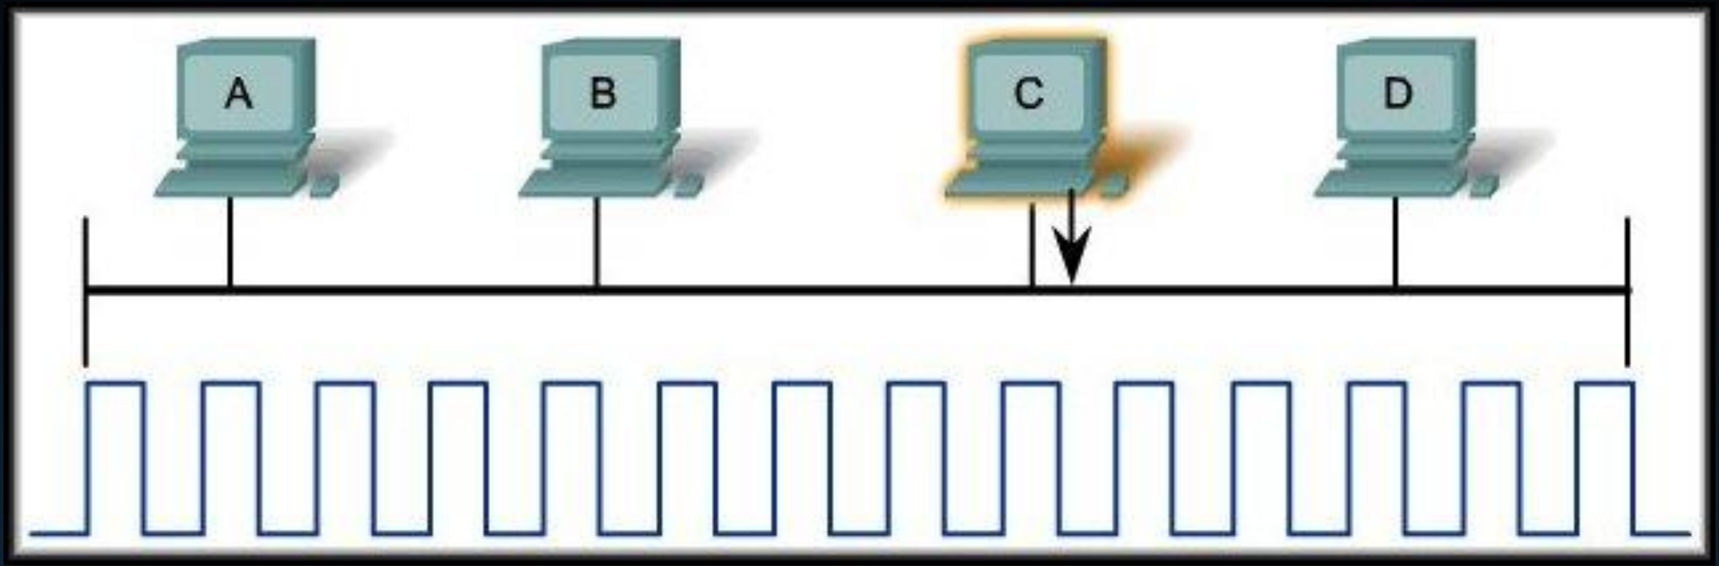

- To transmit, each host will listen on the media.
  - If a signal from another device is present, it will wait for a specific amount of time and listen again.
  - If no signal is present, it will transmit.

CSMA/CD: The Process

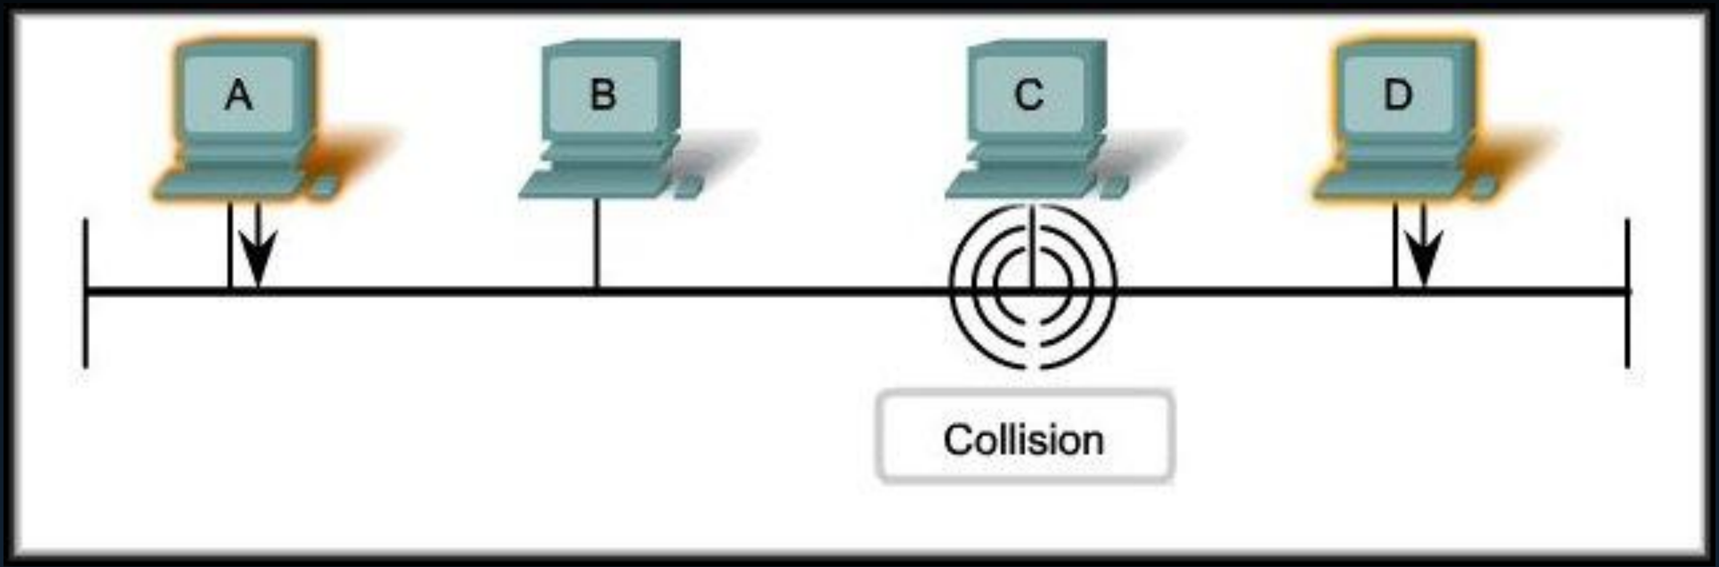

- It can happen that two devices will determine that it is safe to transmit at exactly the same time.
  - In that case, both will transmit their frame.

Collision!

CSMA/CD: The Process

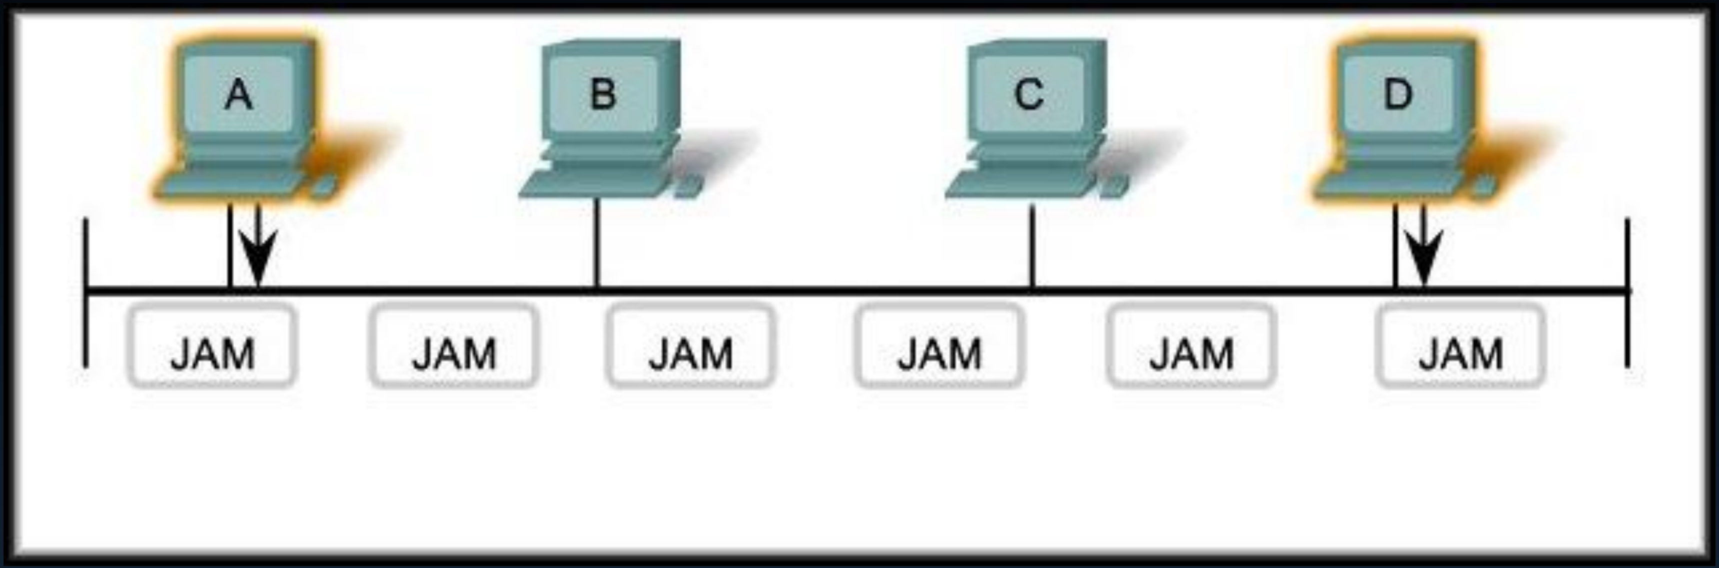

- Both devices detect the collision and send out a jamming signal.
  - The jamming signal is detected by all devices and all devices now know that a collision has occurred on the network.

CSMA/CD: The Process

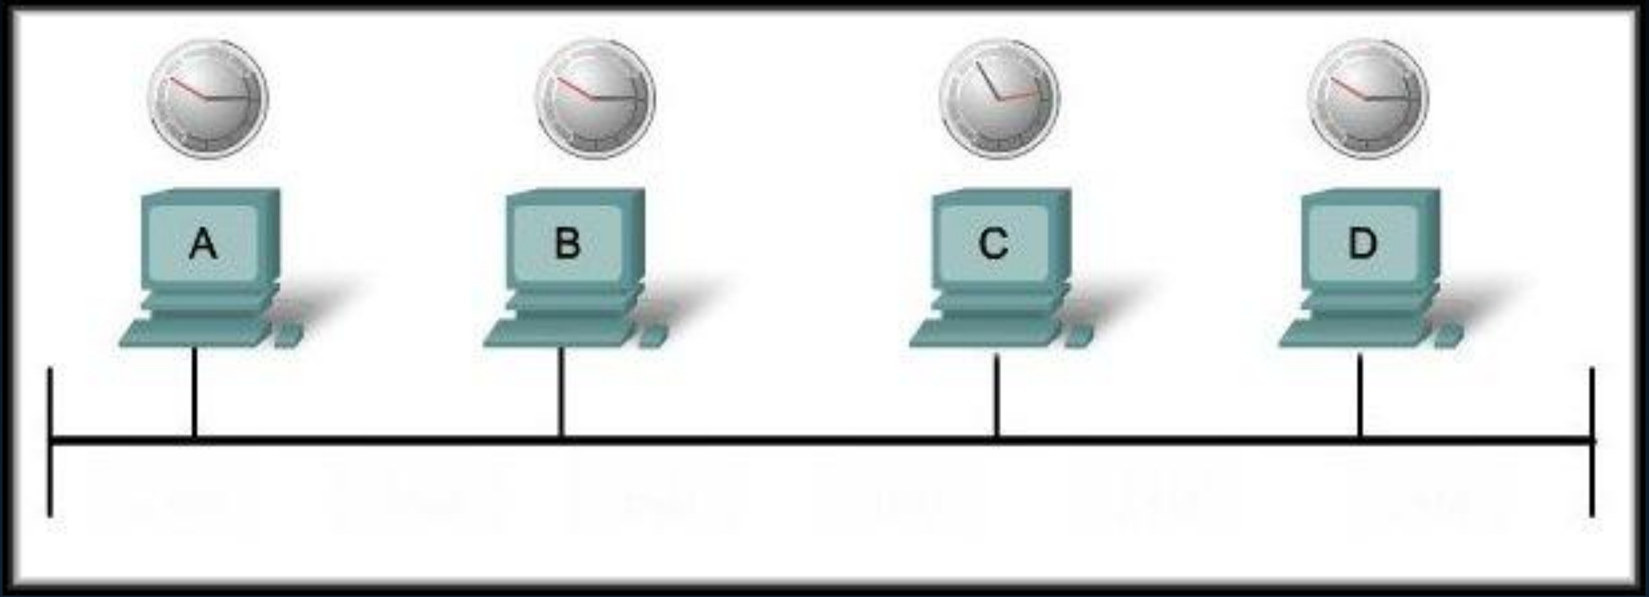

- The jamming signal causes each device to invoke a backoff algorithm.
  - Devices wait a random amount of time before returning to listening mode.
  - The random time ensures that the original devices that caused the collision won't repeat it.

CSMA/CD: The Process

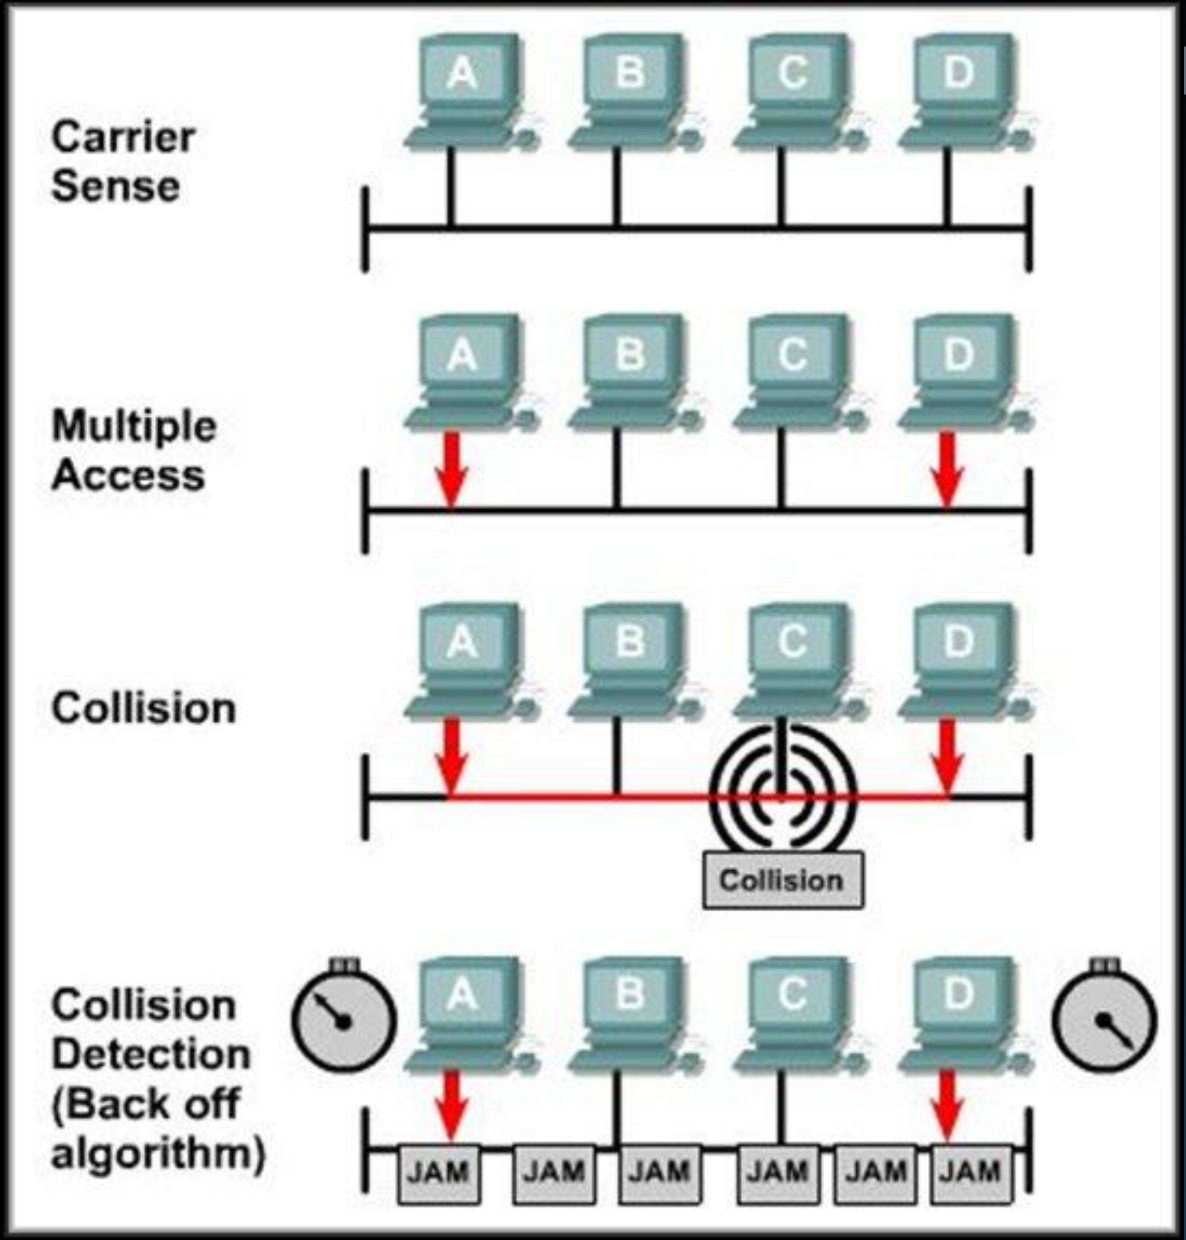

Ethernet Timing

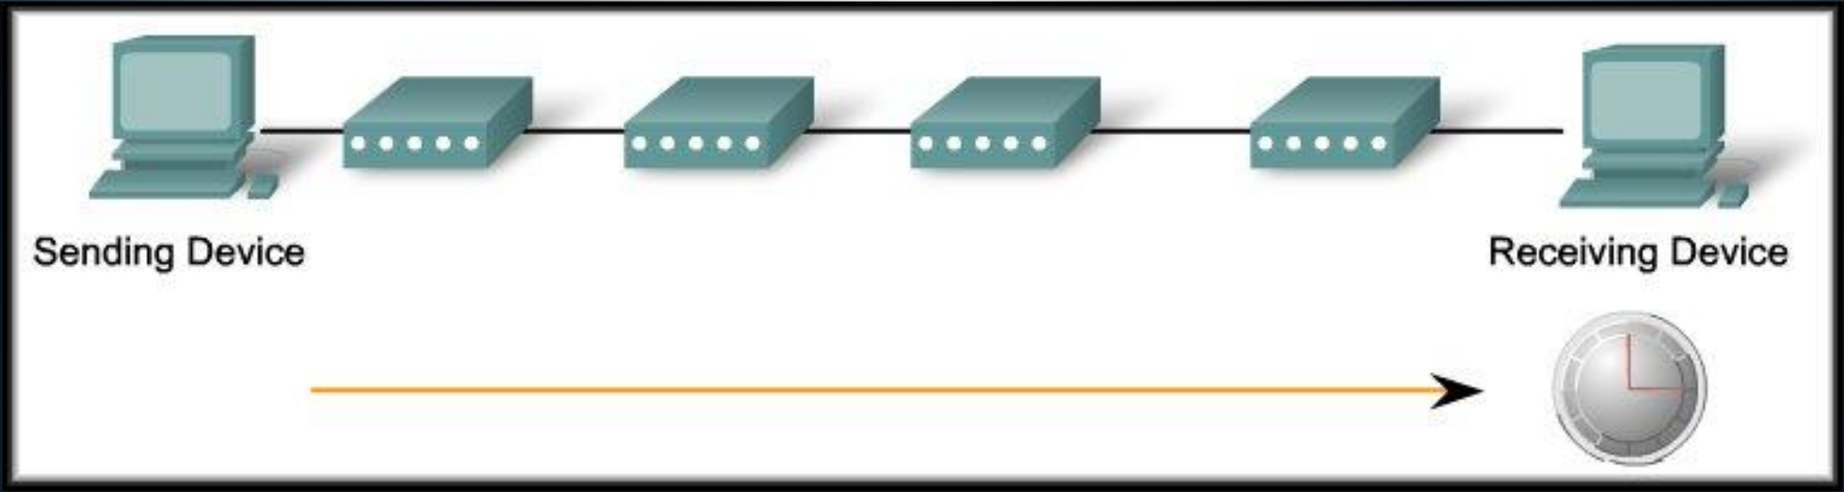

- Latency:
  - Each transmission encounters a certain amount of delay before reaching the destination.
  - Every network device encountered in the path adds to the delay or increases the latency of the transmission.
  - Increases the chance of collisions.

Ethernet Timing

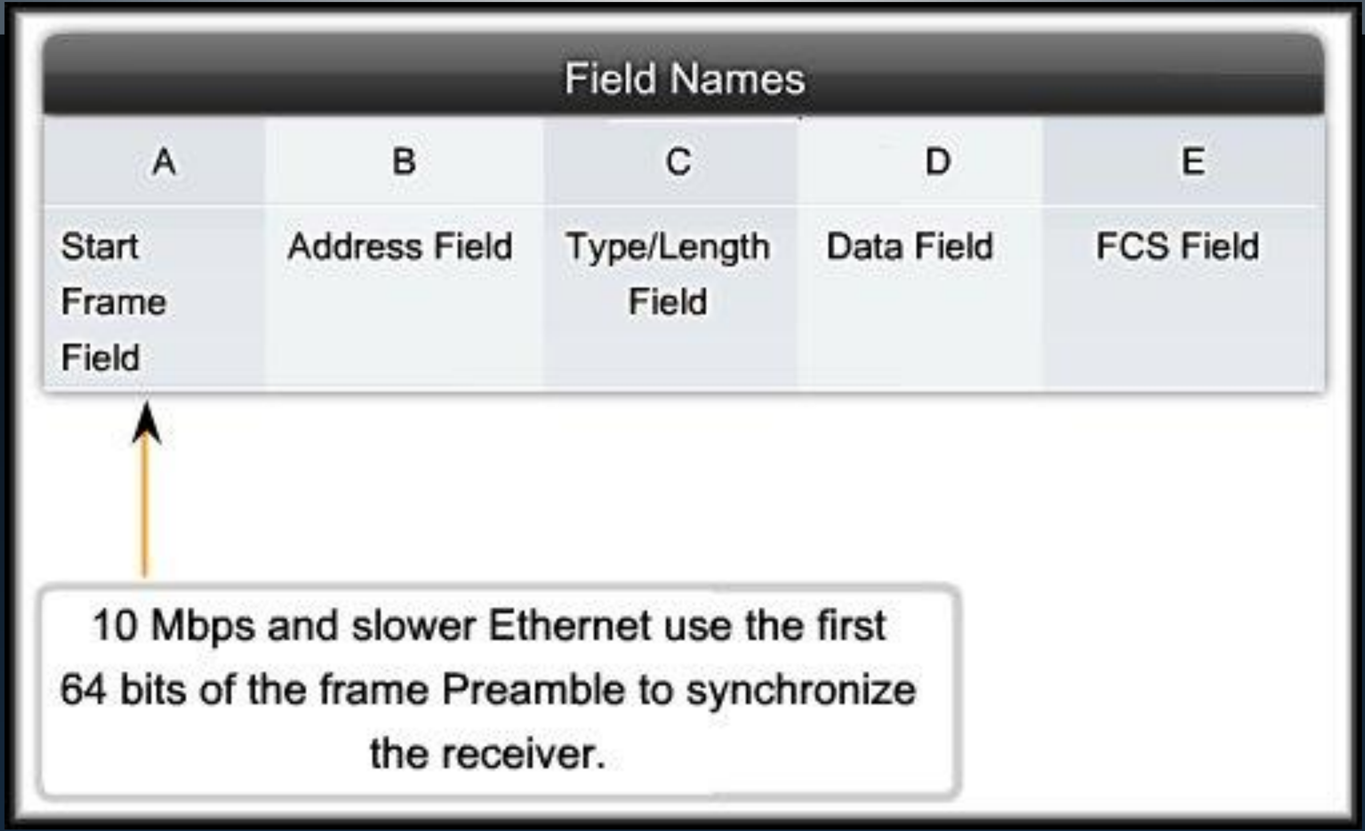

- Timing and Synchronization:
  - The 8 byte (64 bit) preamble is transmitted at the start of the frame.

Ethernet Timing

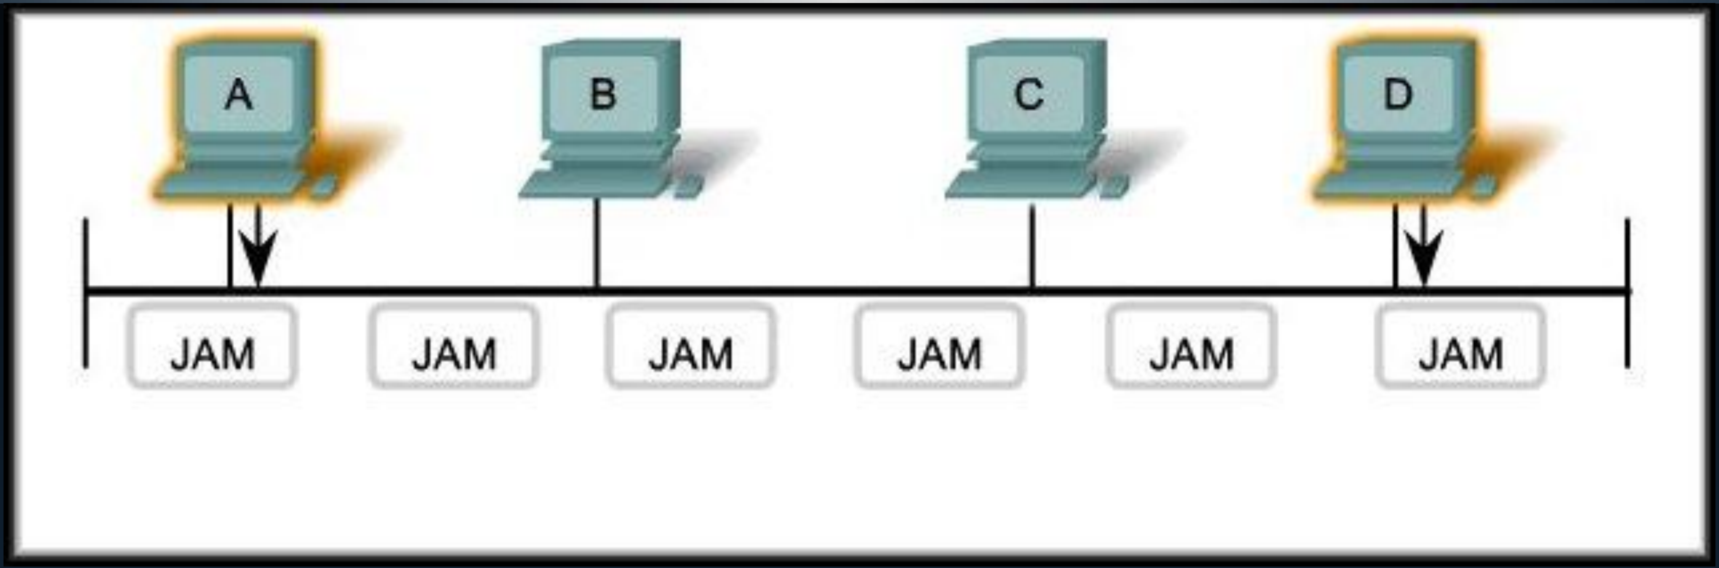

- Jam Signal:
  - As soon as a collision is detected, the sending devices transmit a 32-bit "jam" signal - simply a repeating 1, 0, 1, 0 pattern.
  - Less than 64 bytes (runt).
  - Avoids detection of the jam signal as a frame.

Ethernet Timing

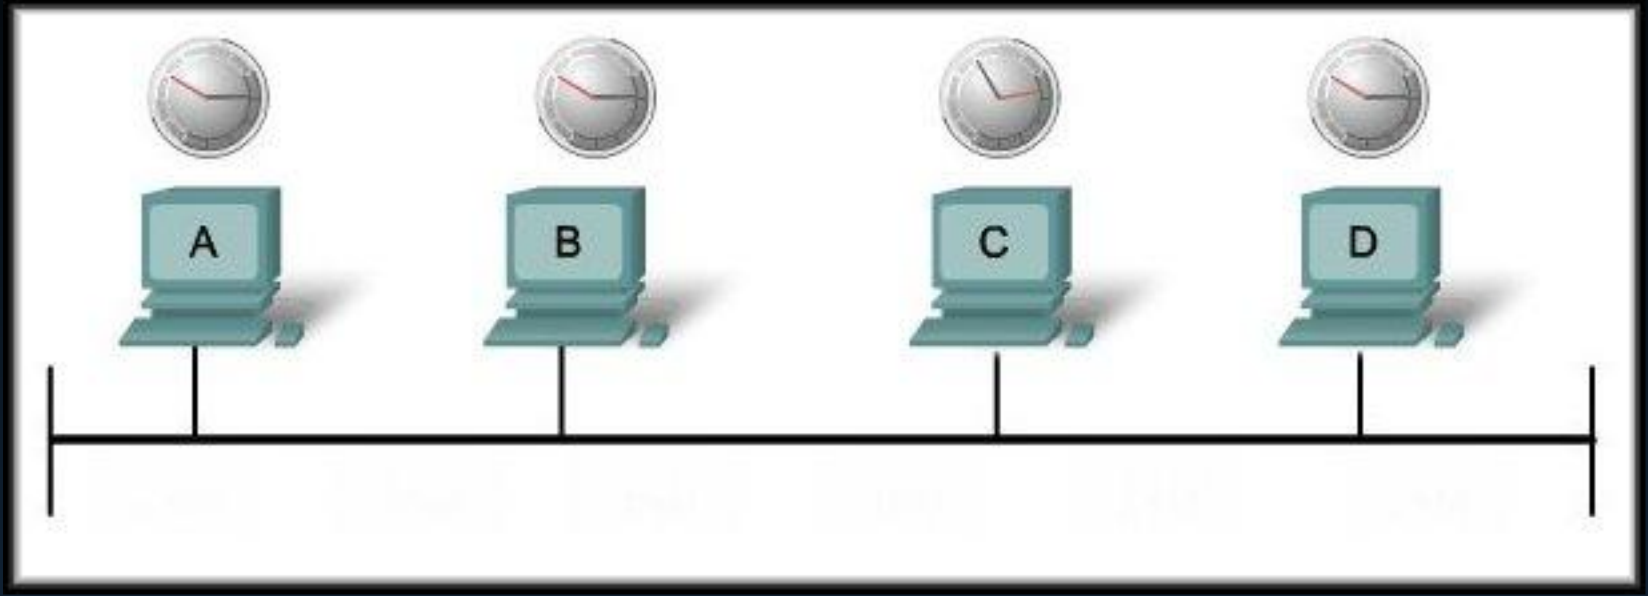

- Backoff Timing:
  - If a collision occurs, all devices wait a random amount of time before listening again.
  - If media congestion results in the MAC layer unable to send the frame after 16 attempts, it gives up and generates an error to the Network layer.