Module 9: The Transport Layer

Instructor Materials

CyberOps Associate v1.0

Module Objectives

- Module Title: The Transport Layer
- Module Objective: Explain how transport layer protocols support network functionality.

Topic Title	Topic Objective
Transport Layer Characteristics	Explain how transport layer protocols support network communication.
Transport Layer Session Establishment	Explain how the transport layer establishes communication sessions.
Transport Layer Reliability	Explain how the transport layer establishes reliable communications.

9.1 Transport Layer Characteristics

Role of the Transport Layer

- The transport layer is responsible for logical communications between applications running on different hosts.
- As shown in the figure, the transport layer is the link between the application layer and the lower layers that are responsible for network transmission.
- The transport layer has no knowledge of the destination host type, the type of media for which the data must travel, the path taken by the data, the congestion on a link, or the size of the network.
- The transport layer includes two protocols, Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).

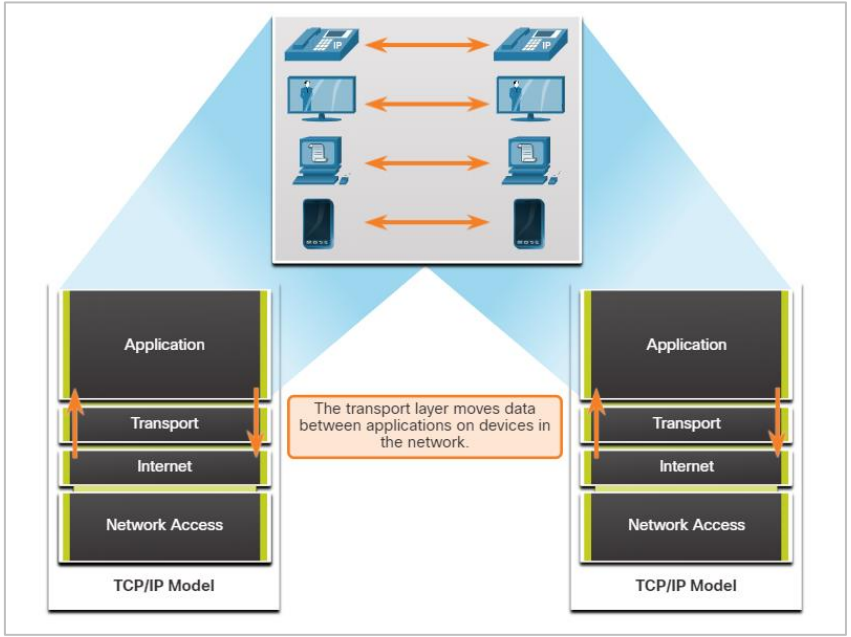

Transport Layer Responsibilities

The transport layer has many responsibilities.

Tracking Individual Conversations

- Each set of data flowing between a source application and a destination application is known as a conversation and is tracked separately.
- It is the responsibility of the transport layer to maintain and track these multiple conversations.
- As shown in the figure, a host may have multiple applications that are communicating across the network simultaneously.
- Most networks have a limitation on the amount of data that can be included in a single packet. Data must be divided into manageable pieces.

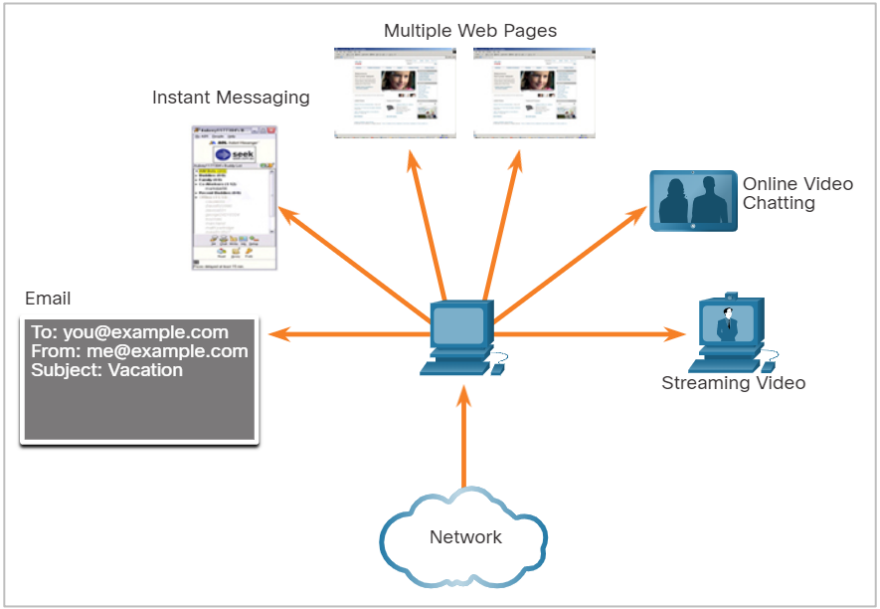

Transport Layer Responsibilities (Contd.)

Segmenting Data and Reassembling Segments

- It is the transport layer responsibility to divide the application data into appropriately sized blocks.
- Depending on the transport layer protocol used, the transport layer blocks are called either segments or datagrams.
- The figure shows the transport layer using different blocks for each conversation.
- The transport layer divides the data into smaller blocks (segments or datagrams) that are easier to manage and transport.

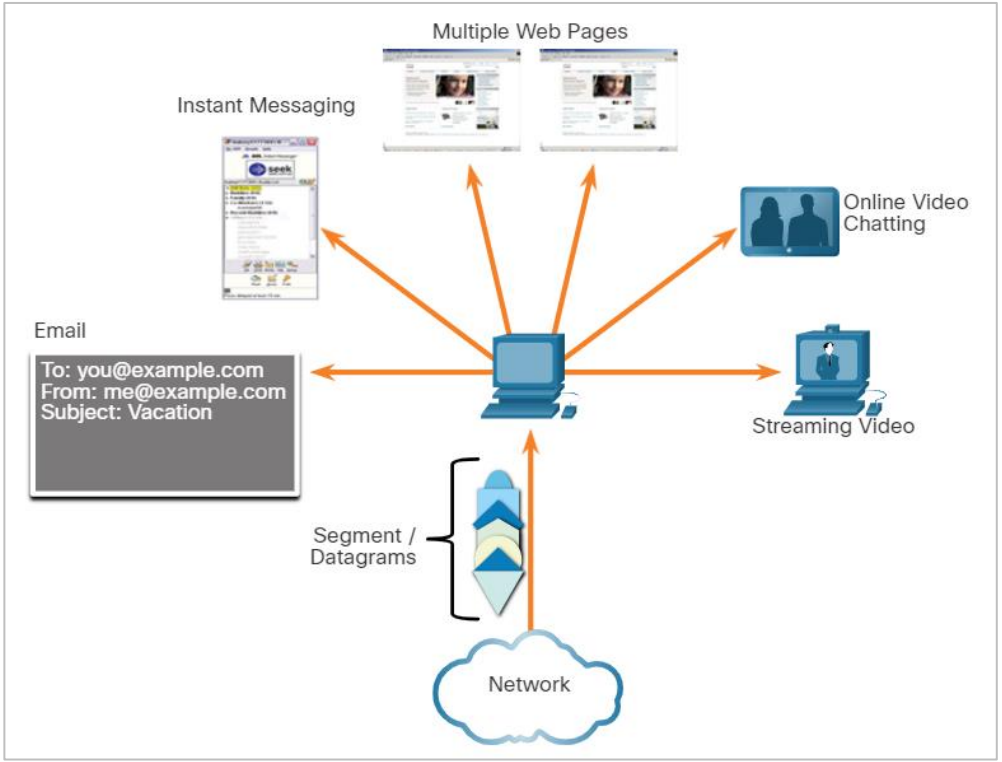

Transport Layer Responsibilities (Contd.)

Add Header Information

- The transport layer protocol also adds header information containing binary data organized into several fields to each block of data.
- The values in these fields enable various transport layer protocols to perform different functions in managing data communication.
- The header information is used by the receiving host to reassemble the blocks of data into a complete data stream for the receiving application layer program.
- The transport layer ensures that even with multiple application running on a device, all applications receive the correct data.

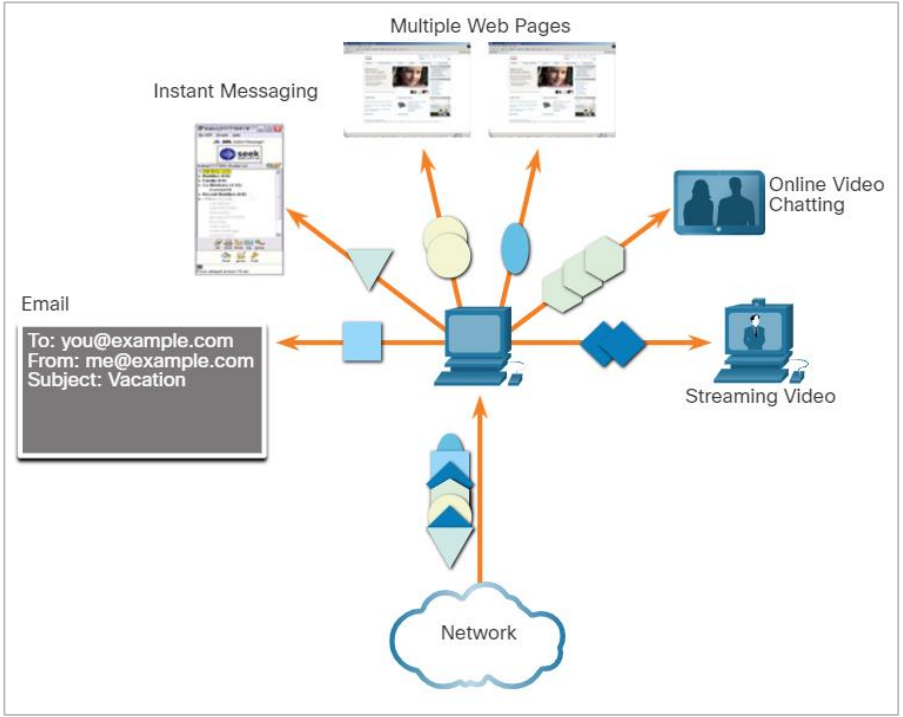

Transport Layer Responsibilities (Contd.)

Identifying the Applications

- The transport layer must be able to separate and manage multiple communications with different transport requirement needs.
- To pass data streams to the proper applications, the transport layer identifies the target application using an identifier called a port number.
- As shown in the figure, each software process that needs to access the network is assigned a port number unique to that host.

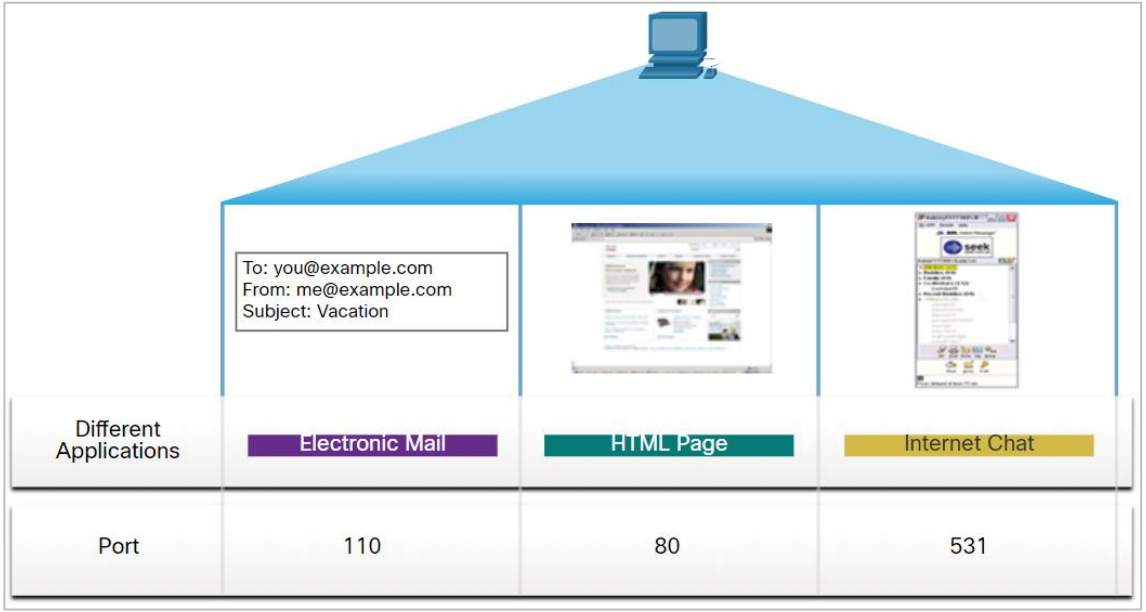

Transport Layer Responsibilities (Contd.)

Conversation Multiplexing

- Sending some types of data across a network, as one complete communication stream, can consume all the available bandwidth.
- This prevents other communication conversations from occurring at the same time and also make error recovery and retransmission of damaged data difficult.
- As shown in the figure, the transport layer uses segmentation and multiplexing to enable different communication conversations to be interleaved on the same network.
- Error checking can be performed on the data in the segment, to determine if the segment was altered during transmission.

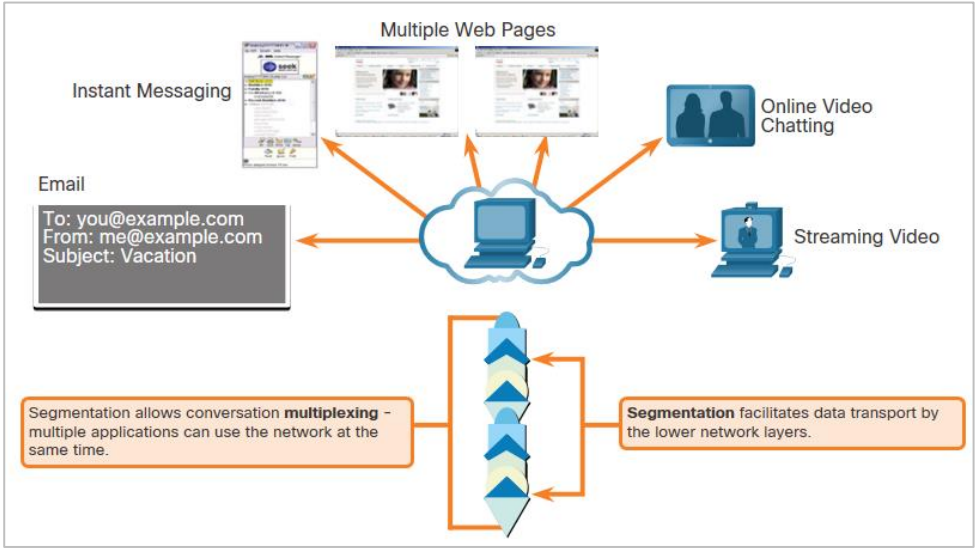

Transport Layer Protocols

- IP is concerned only with the structure, addressing, and routing of packets.
- IP does not specify how the delivery or transportation of the packets takes place.
- Transport layer protocols (TCP and UDP) specify how to transfer messages between hosts, and are responsible for managing reliability requirements of a conversation.
- The transport layer includes the TCP and UDP protocols.
- Different applications have different transport reliability requirements. Therefore, TCP/IP provides two transport layer protocols, as shown in the figure.

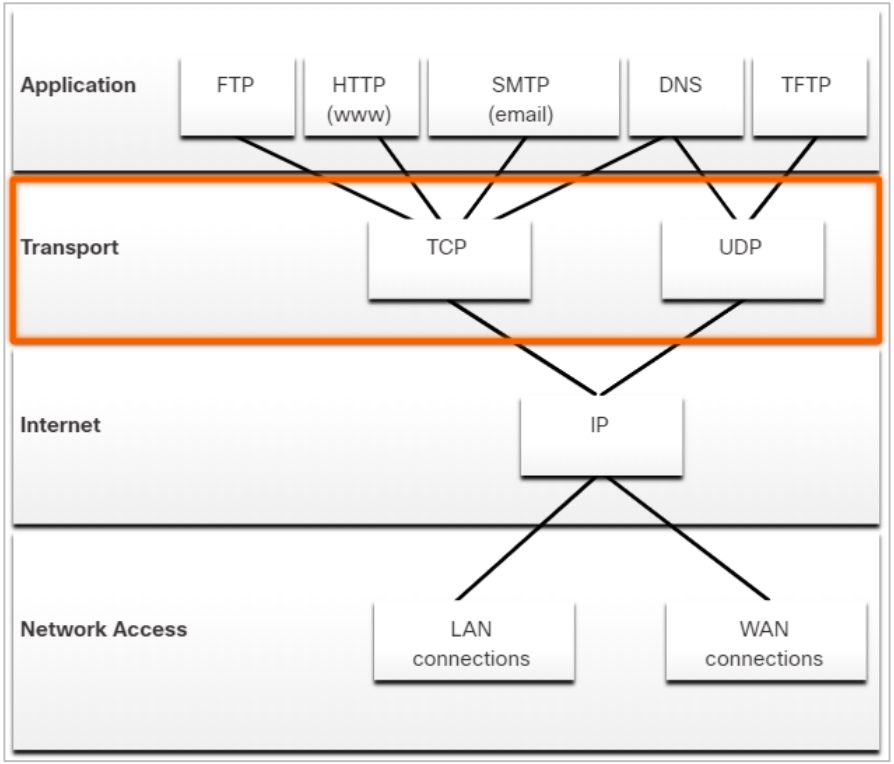

Transmission Control Protocol (TCP)

- TCP is considered a reliable, full-featured transport layer protocol, which ensures that all of the data arrives at the destination.
- TCP includes fields which ensure the delivery of the application data. These fields require additional processing by the sending and receiving hosts.
- TCP transport is analogous to sending packages that are tracked from source to destination.
- TCP provides reliability and flow control using these basic operations:
  - Number and track data segments transmitted to a specific host from a specific application
  - Acknowledge received data
  - Retransmit any unacknowledged data after a certain amount of time
  - Sequence data that might arrive in wrong order
  - Send data at an efficient rate that is acceptable by the receiver
- ***Note**: TCP divides data into segments.*

Transmission Control Protocol (TCP) (Contd.)

In order to maintain the state of a conversation and track the information, TCP must first establish a connection between the sender and the receiver. This is why TCP is known as a connection-oriented protocol.

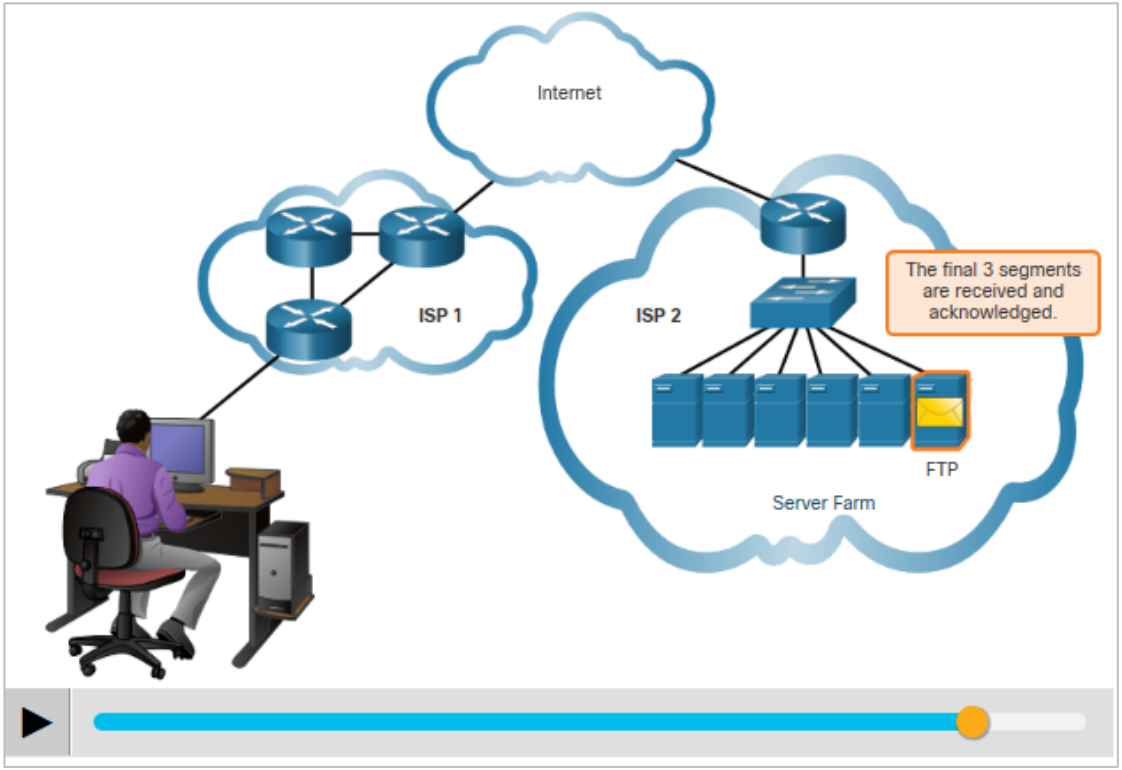

TCP Header

- TCP is a stateful protocol as it keeps track of the state of the communication session.
- To track the state of a session, TCP records which information it has sent and which information has been acknowledged.
- The stateful session begins with the session establishment and ends with the session termination.
- A TCP segment adds 20 bytes (160 bits) of overhead when encapsulating the application layer data. The figure shows the fields in a TCP header.

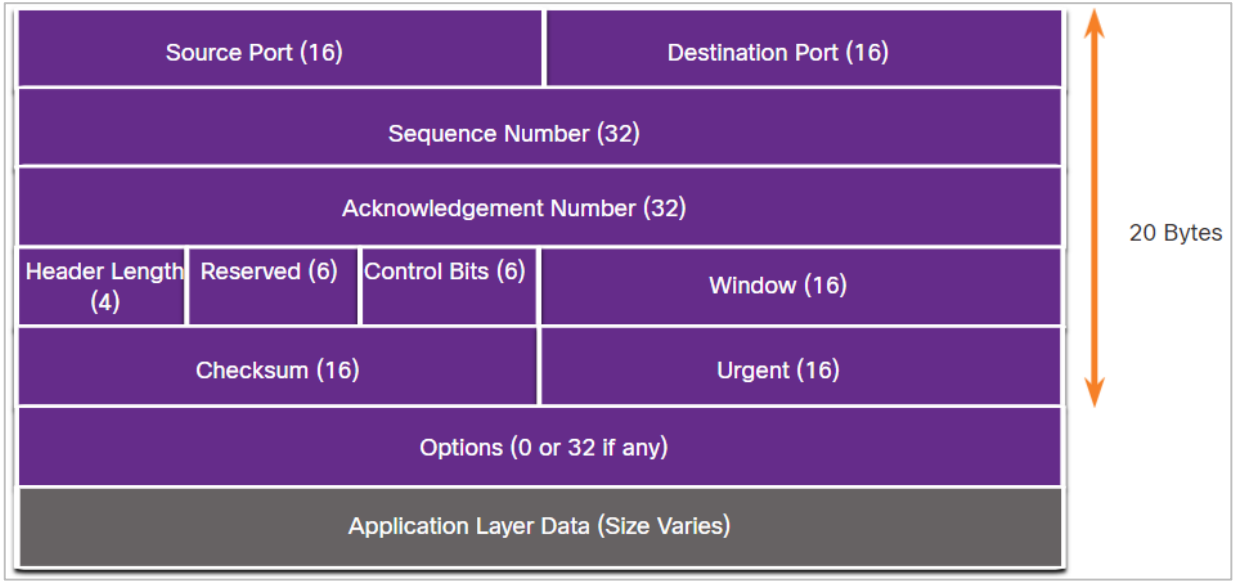

TCP Header Fields

TCP Header Field	Description
Source Port	A 16-bit field used to identify the source application by port number.
Destination Port	A 16-bit field used to identify the destination application by port number.
Sequence Number	A 32-bit field used for data reassembly purposes.
Acknowledgment Number	A 32-bit field used to indicate that data has been received and the next byte expected from the source.
Header Length	A 4-bit field known as "data offset" that indicates the length of the TCP segment header.
Reserved	A 6-bit field that is reserved for future use.
Control bits	A 6-bit field that includes bit codes, or flags, which indicate the purpose and function of the TCP segment.
Window size	A 16-bit field used to indicate the number of bytes that can be accepted at one time.
Checksum	A 16-bit field used for error checking of the segment header and data.
Urgent	A 16-bit field used to indicate if the contained data is urgent.

User Datagram Protocol (UDP)

- UDP is a simpler transport layer protocol than TCP.
- It does not provide reliability and flow control, which means it requires fewer header fields.
- The sender and the receiver UDP processes do not have to manage reliability and flow control, this means UDP datagrams can be processed faster than TCP segments.
- UDP provides the basic functions for delivering datagrams between the appropriate applications, with very little overhead and data checking.
- UDP is a connectionless protocol. Because UDP does not provide reliability or flow control, it does not require an established connection.
- UDP is also known as a stateless protocol. Because UDP does not track information sent or received between the client and server.
- ***Note****: UDP divides data into datagrams that are also referred to as segments.*

User Datagram Protocol (UDP) (Contd.)

- UDP is also known as a best-effort delivery protocol because there is no acknowledgment that the data is received at the destination.
- UDP is like placing a regular, nonregistered, letter in the mail. The sender of the letter is not aware of the availability of the receiver to receive the letter. Nor is the post office responsible for tracking the letter or informing the sender if the letter does not arrive at the final destination.

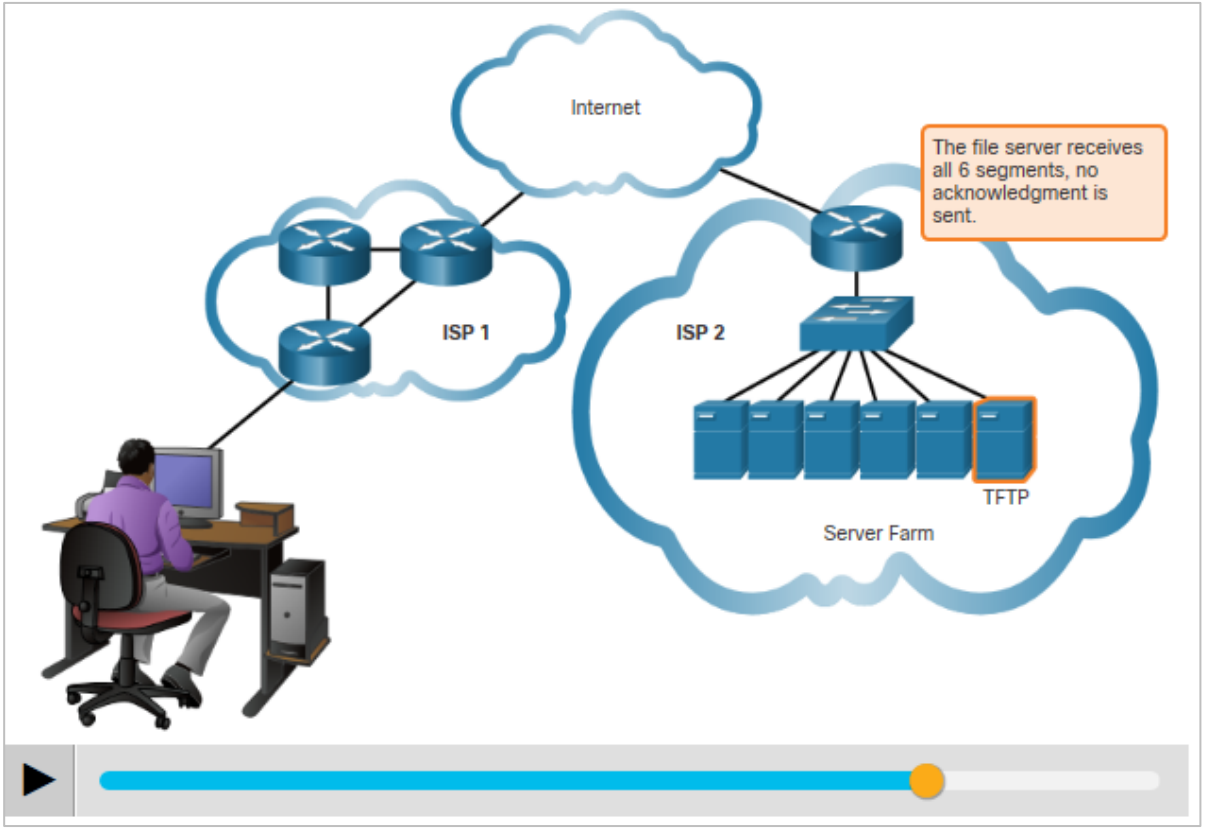

UDP Header

- UDP is a stateless protocol meaning neither the client, nor the server, tracks the state of the communication session. If reliability is required when using UDP as the transport protocol, it must be handled by the application.
- The requirements for delivering live video and voice over the network is the data continues to flow quickly. Live video and voice applications can tolerate some data loss and are perfectly suited to UDP.
- The blocks of communication in UDP are called datagrams, or segments. These datagrams are sent as best effort by the transport layer protocol.
- The UDP header only has four fields and requires 8 bytes (64 bits). The figure shows the fields in a UDP header.

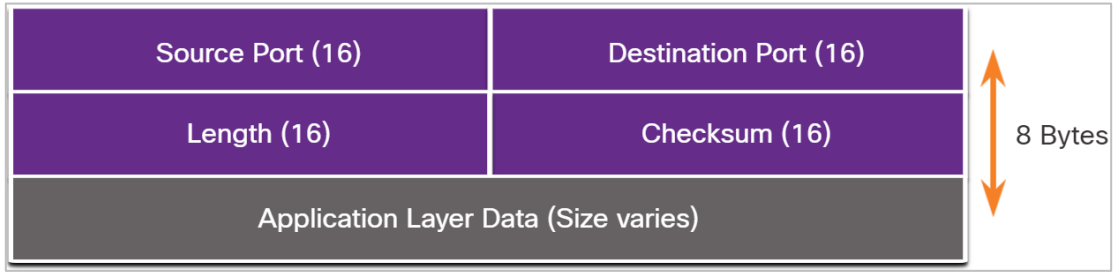

UDP Header Fields

UDP Header Field	Description
Source Port	A 16-bit field used to identify the source application by port number.
Destination Port	A 16-bit field used to identify the destination application by port number.
Sequence Number	A 32-bit field used for data reassembly purposes.
Length	A 16-bit field that indicates the length of the UDP datagram header.
Checksum	A 16-bit field used for error checking of the datagram header and data.

Socket Pairs

- The source and destination ports are placed within the segment. The segments are then encapsulated within an IP packet.
- The IP packet contains the IP address of the source and destination. The combination of the source IP address and source port number, or the destination IP address and destination port number is known as a socket.
- Sockets enable multiple processes, running on a client, to distinguish themselves from each other, and multiple connections to a server process to be distinguished from each other.
- The source port number acts as a return address for the requesting application.
- The transport layer keeps track of this port and the application that initiated the request so that when a response is returned, it can be forwarded to the correct application.

Socket Pairs (Contd.)

- In the figure, the PC is simultaneously requesting FTP and web services from the destination server.
- The FTP request generated by the PC includes the Layer 2 MAC addresses and the Layer 3 IP addresses. The request also identifies the source port number 1305 and destination port, identifying the FTP services on port 21.
- The host also has requested a web page from the server using the same Layer 2 and Layer 3 addresses.

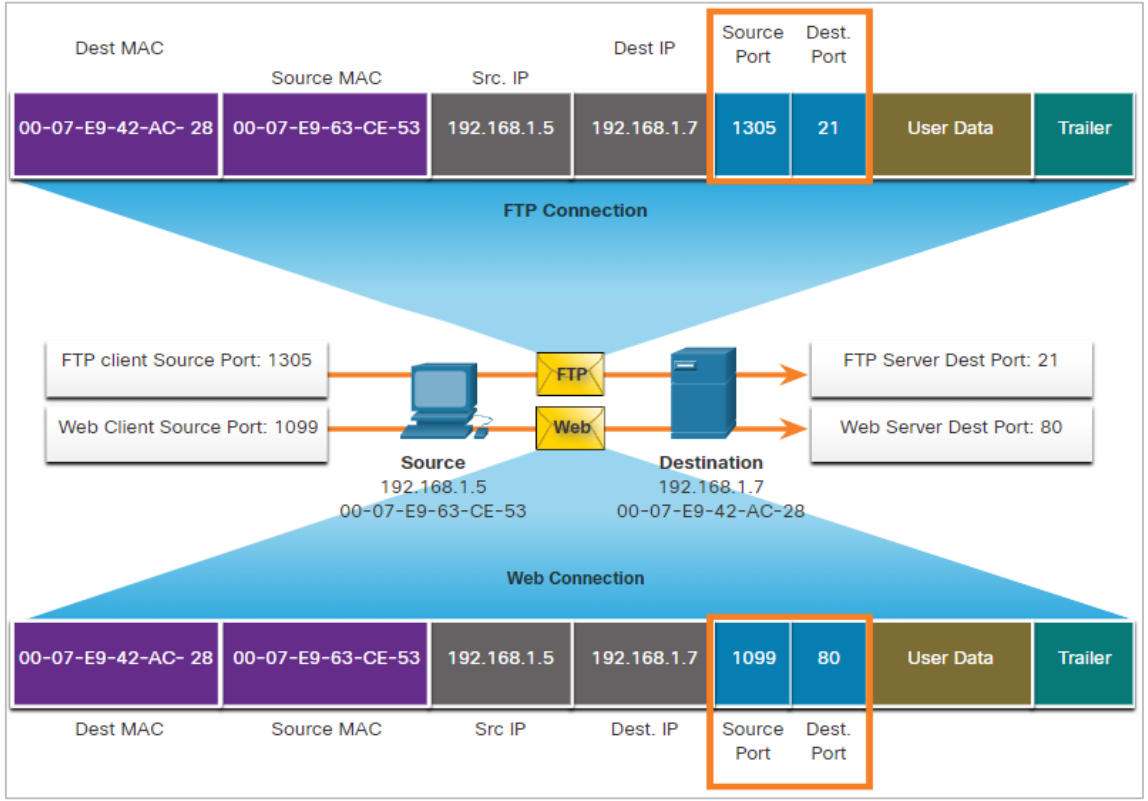

Socket Pairs (Contd.)

- It is using the source port number 1099 and destination port identifying the web service on port 80.
- The socket is used to identify the server and service being requested by the client.
- A client socket with 1099 representing the source port number might be 192.168.1.5:1099. The socket on a web server might be 192.168.1.7:80. Together, these two sockets combine to form a *socket pair*: 192.168.1.5:1099, 192.168.1.7:80

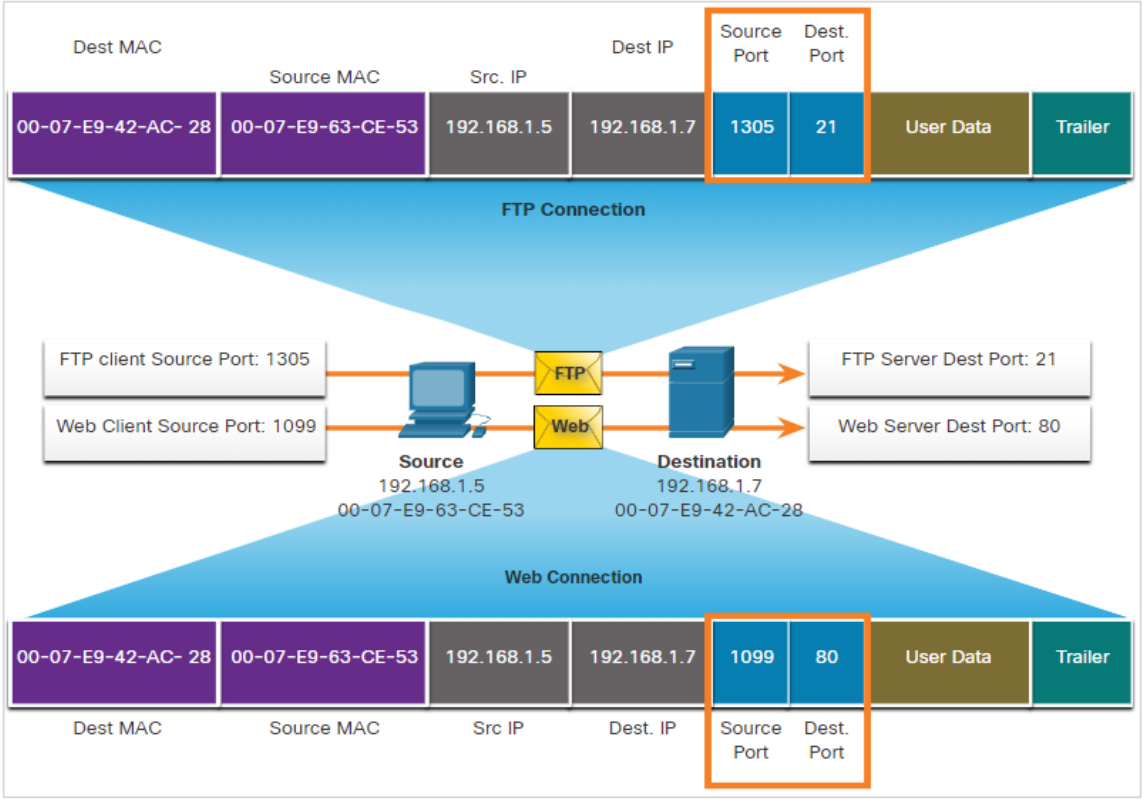

9.2 Transport Layer Session Establishment

TCP Server Processes

- Each application process running on a server is configured to use a port number. The port number is either automatically assigned or configured manually by a system administrator.
- An individual server cannot have two services assigned to the same port number within the same transport layer services.
- A host running a web server application and a file transfer application cannot have both configured to use the same port, such as TCP port 80.
- An active server application assigned to a specific port is considered open, which means that the transport layer accepts, and processes segments addressed to that port.
- Any incoming client request addressed to the correct socket is accepted, and the data is passed to the server application.
- There can be many ports open simultaneously on a server, one for each active server application.

TCP Server Processes (Contd.)

Clients Sending TCP Requests

Client 1 is requesting web services and Client 2 is requesting email service of the same sever.

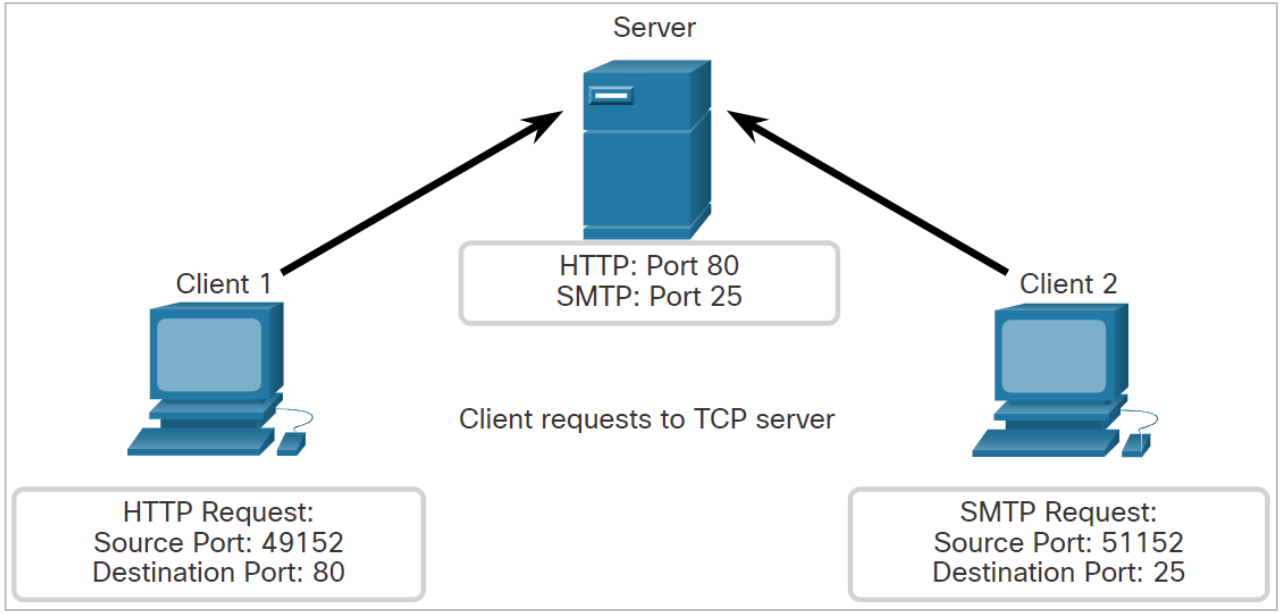

TCP Server Processes (Contd.)

Request Destination Ports
Client 1 is requesting web services using well-known destination port 80 (HTTP) and Client 2 is requesting email service using well-known port 25 (SMTP).

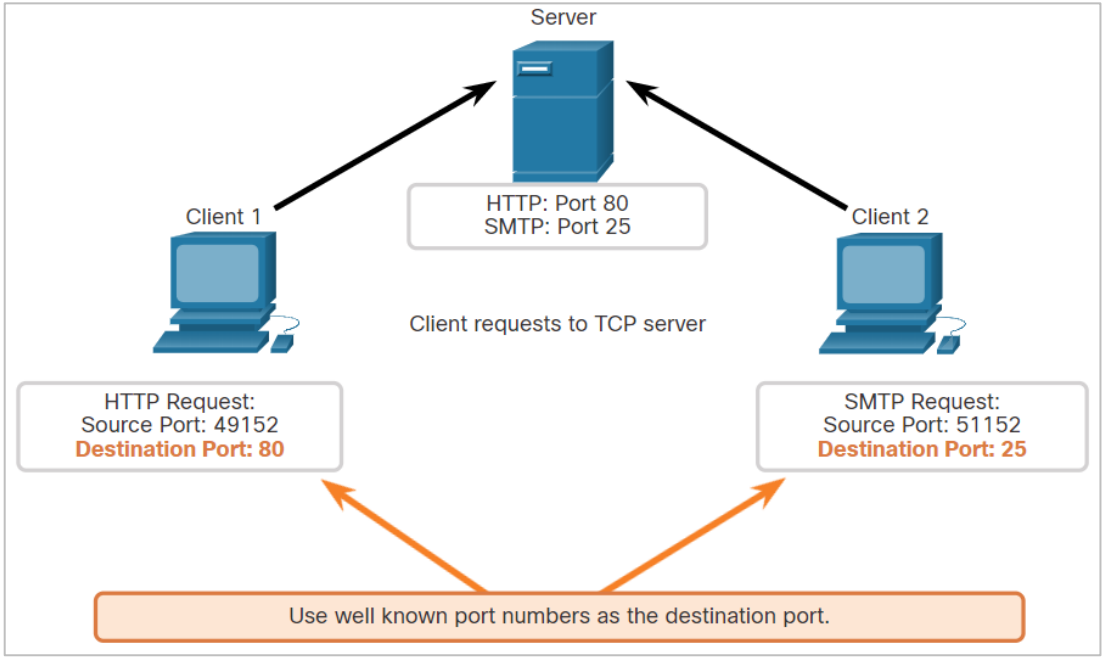

TCP Server Processes (Contd.)

Request Source Ports

Client requests dynamically generate a source port number. In this case, Client 1 is using source port 49152 and Client 2 is using source port 51152.

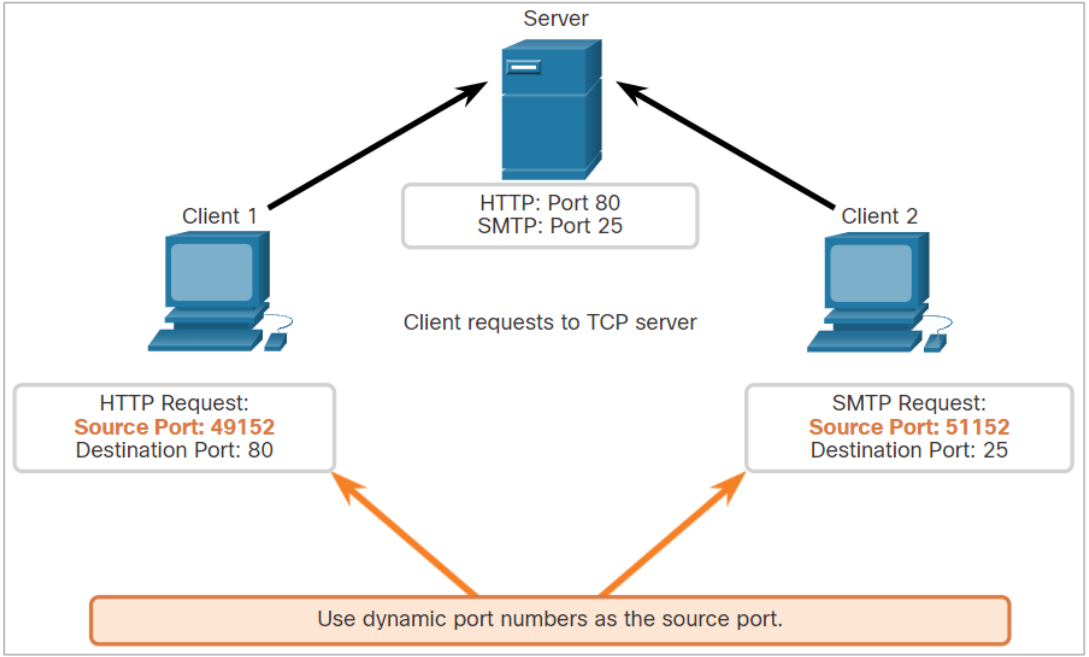

TCP Server Processes (Contd.)

Response Destination Ports

When the server responds to the client requests, it reverses the destination and source ports of the initial request. Notice that the Server response to the web request now has destination port 49152 and the email response now has destination port 51152.

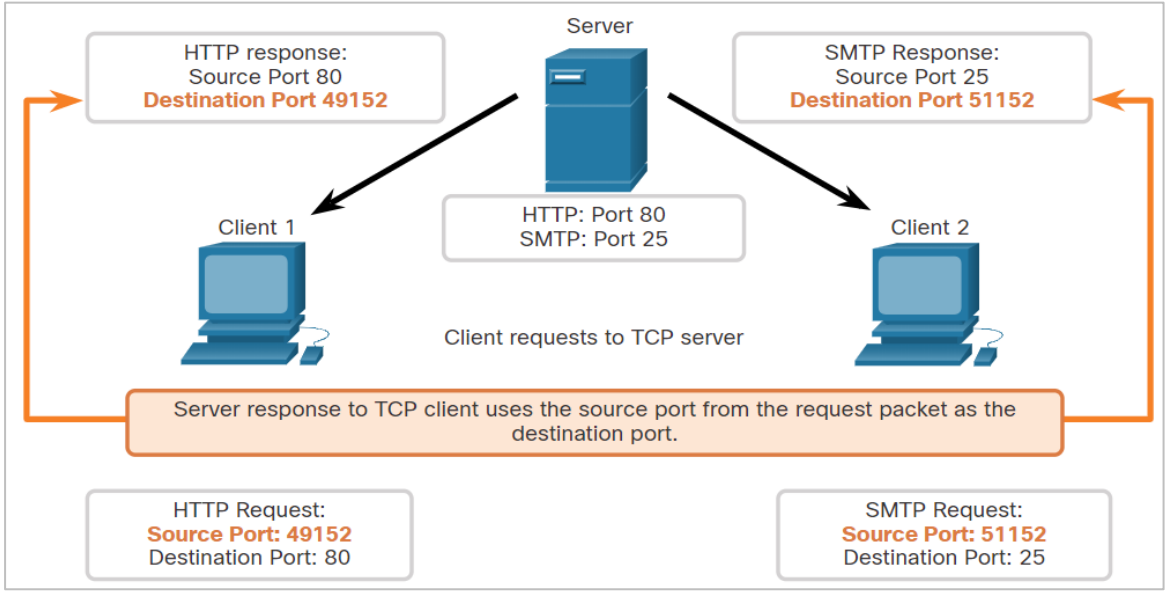

TCP Server Processes (Contd.)

Response Source Ports
The source port in the server response is the original destination port in the initial requests.

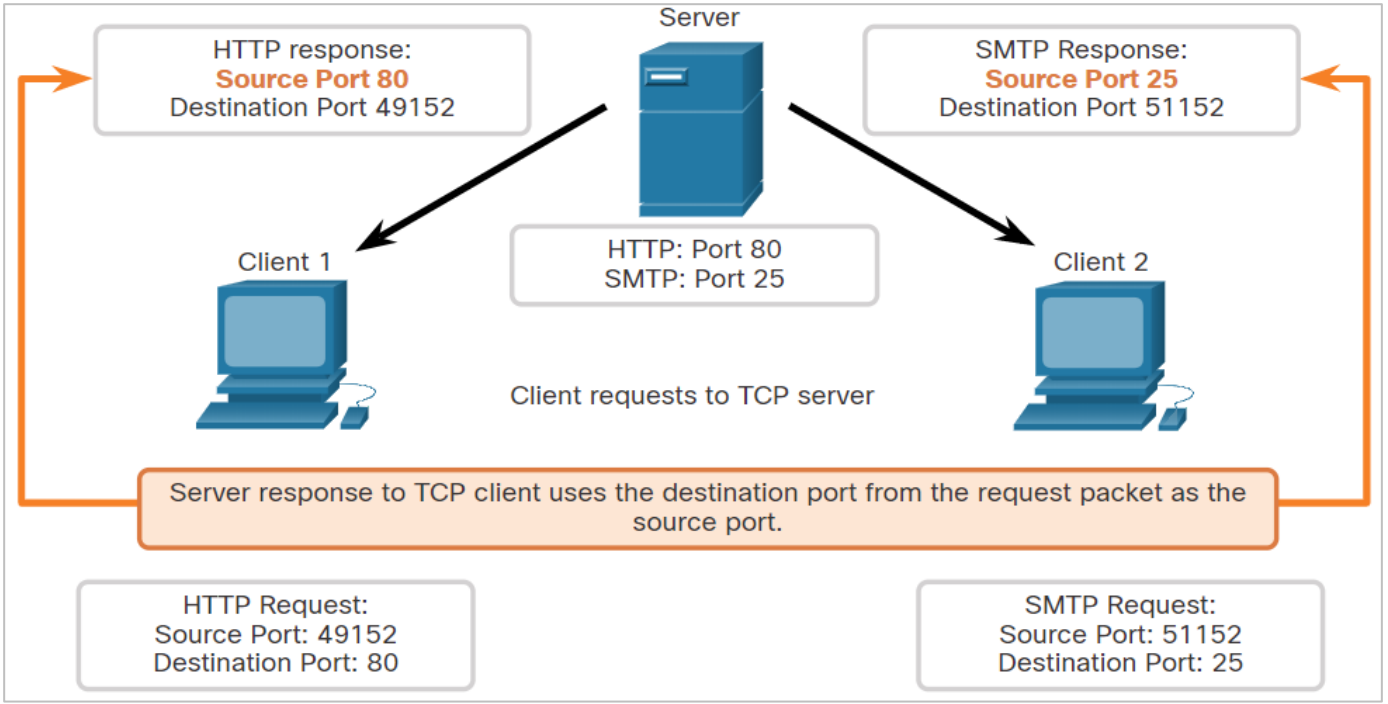

TCP Connection Establishment

- In TCP connections, the host client establishes the connection with the server using the three-way handshake process.
- The three-way handshake validates that the destination host is available to communicate.
- The TCP connection establishment steps are:
  - **Step 1. SYN:** The initiating client requests a client-to-server communication session with the server.

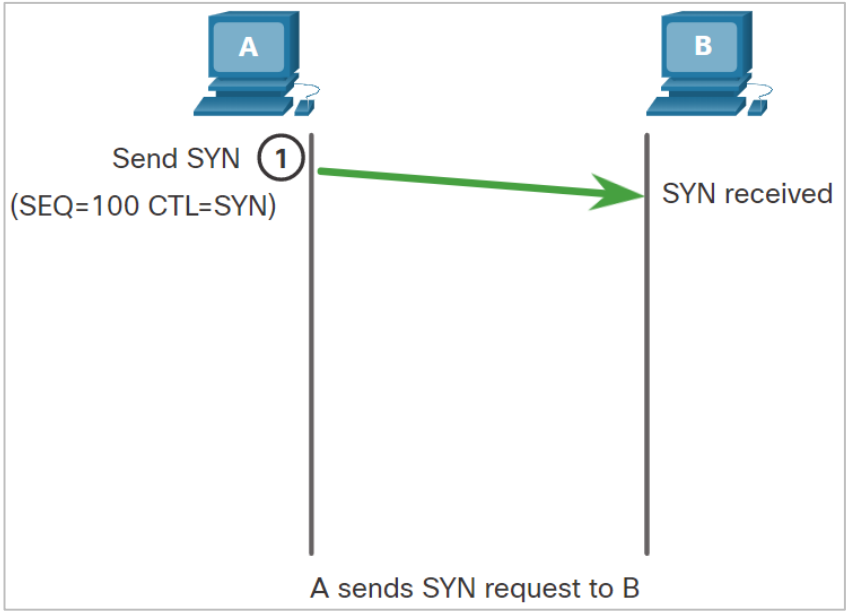

TCP Connection Establishment (Contd.)

Step 2. ACK and SYN: The server acknowledges the client-to-server communication session and requests a server-to-client communication session.

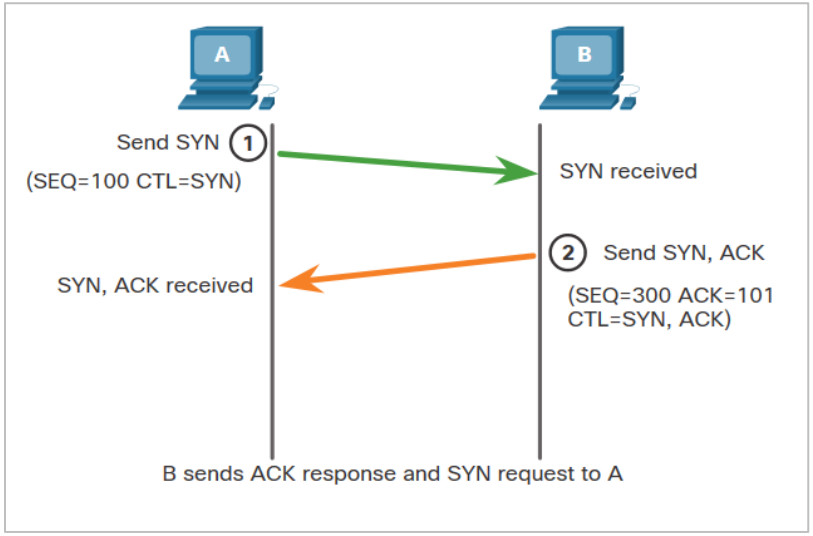

Step 3. ACK: The initiating client acknowledges the server-to-client communication session.

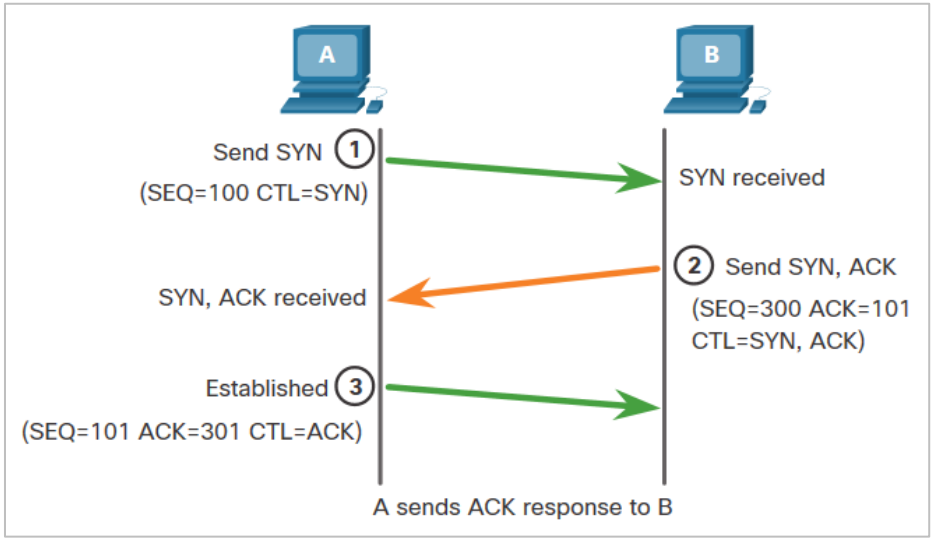

Session Termination

- To close a connection, the Finish (FIN) control flag must be set in the segment header.
- To end each one-way TCP session, a two-way handshake, consisting of a FIN segment and an Acknowledgment (ACK) segment, is used.
- Therefore, to terminate a single conversation supported by TCP, four exchanges are needed to end both sessions. Either the client or the server can initiate the termination.
- The terms client and server are used as a reference for simplicity, but any two hosts that have an open session can initiate the termination process.
- When all segments have been acknowledged, the session is closed.

Session Termination (Contd.)

The session termination steps are:

Step 1. FIN: When the client has no more data to send in the stream, it sends a segment with the FIN flag set.

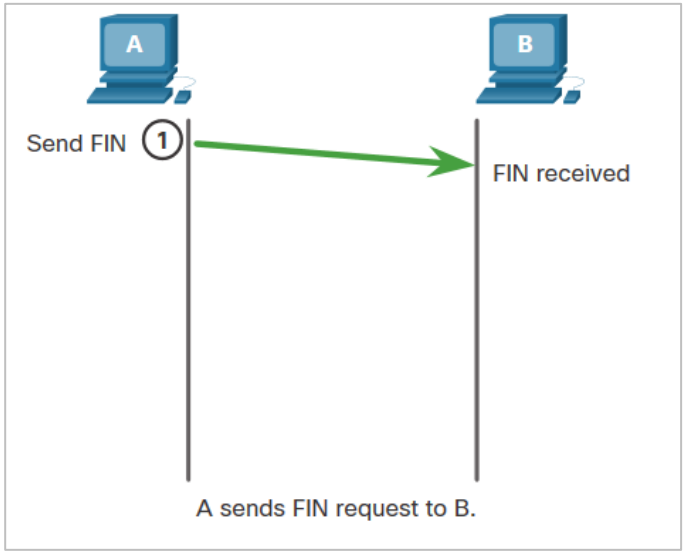

- **Step 2. ACK:** The server sends an ACK to acknowledge the receipt of the FIN to terminate the session from client to server.

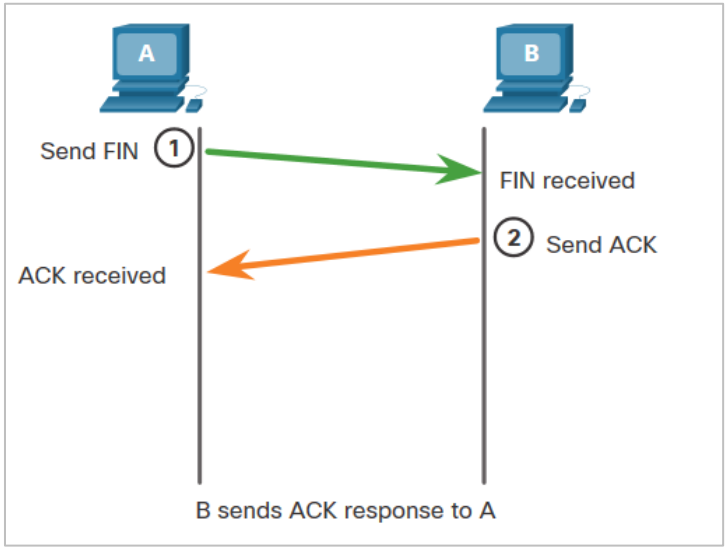

Session Termination (Contd.)

Step 3. FIN: The server sends a FIN to the client to terminate the server-to-client session.

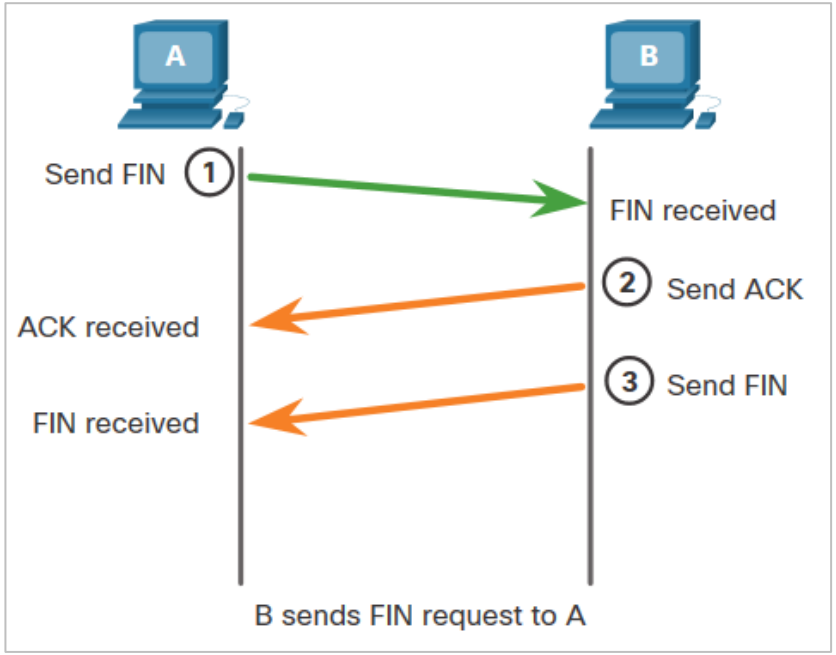

Step 4. ACK: The client responds with an ACK to acknowledge the FIN from the server.

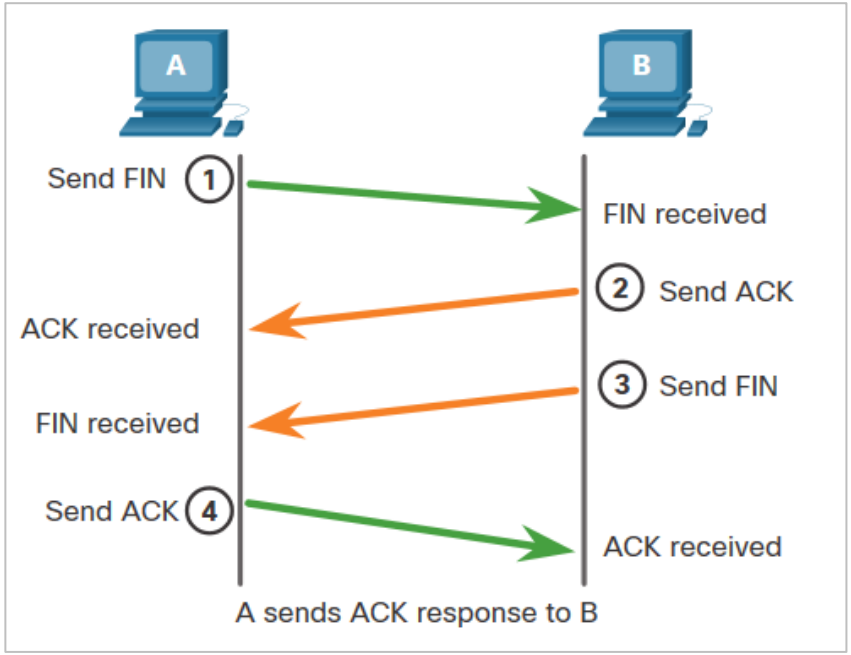

TCP Three-way Handshake Analysis

- Hosts maintain state, track each data segment within a session, and exchange information about the data is received using the information in the TCP header.
- TCP is a full-duplex protocol, where each connection represents two one-way communication sessions. To establish the connection, the hosts perform a three-way handshake. As shown in the figure, control bits in the TCP header indicate the progress and status of the connection.
- The functions of the three-way handshake are:
  - It establishes that the destination device is present on the network.
  - It verifies that the destination device has an active service and is accepting requests on the destination port number that the initiating client intends to use.
  - It informs the destination device that the source client intends to establish a communication session on that port number.
- After the communication is completed the sessions are closed, and the connection is terminated. The connection and session mechanisms enable TCP reliability function.

TCP Three-way Handshake Analysis (Contd.)

- The six bits in the Control Bits field of the TCP segment header are also known as flags. A flag is a bit that is set to either on or off. The six control bits flags are as follows:
- **URG** - Urgent pointer field significant
- **ACK** - Acknowledgment flag used in connection establishment and session termination
- **PSH** - Push function
- **RST** - Reset the connection when an error or timeout occurs
- **SYN** - Synchronize sequence numbers used in connection establishment
- **FIN** - No more data from sender and used in session termination

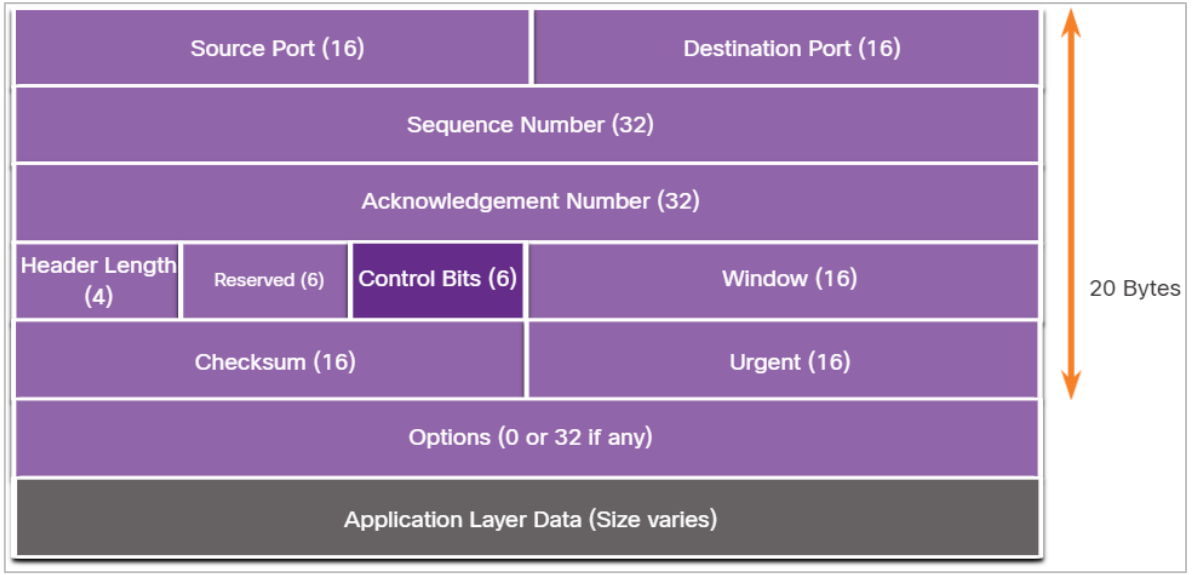

Lab – Using Wireshark to Observe the TCP 3-Way Handshake

In this lab, you will complete the following objectives:

- **Part 1:** Prepare the Hosts to Capture the Traffic
- **Part 2:** Analyze the Packets using Wireshark
- **Part 3:** View the Packets using tcpdump

9.3 Transport Layer Reliability

Guaranteed and Ordered Delivery

- There may be times when either TCP segments do not arrive at their destination or arrive out of order.
- For the original message to be understood by the recipient, all the data must be received and the data in these segments must be reassembled into the original order.
- Sequence numbers are assigned in the header for each packet to achieve this goal. The sequence number represents the first data byte of the TCP segment.
- During session setup, an initial sequence number (ISN) is set, which represents the starting value of the bytes that are transmitted to the receiving application.
- As data is transmitted during the session, the sequence number is incremented by the number of bytes that have been transmitted.
- This data byte tracking enables each segment to be uniquely identified and acknowledged. Missing segments can then be identified.
- The ISN is effectively a random number which prevents certain types of malicious attacks.

Guaranteed and Ordered Delivery (Contd.)

- Segment sequence numbers indicate how to reassemble and reorder received segments, as shown in the figure.
- The receiving TCP process places the data from a segment into a receiving buffer.
- Segments are then placed in the proper sequence order and passed to the application layer when reassembled.
- Any segments that arrive with sequence numbers that are out of order are held for later processing.
- Then, when the segments with the missing bytes arrives, these segments are processed in order.

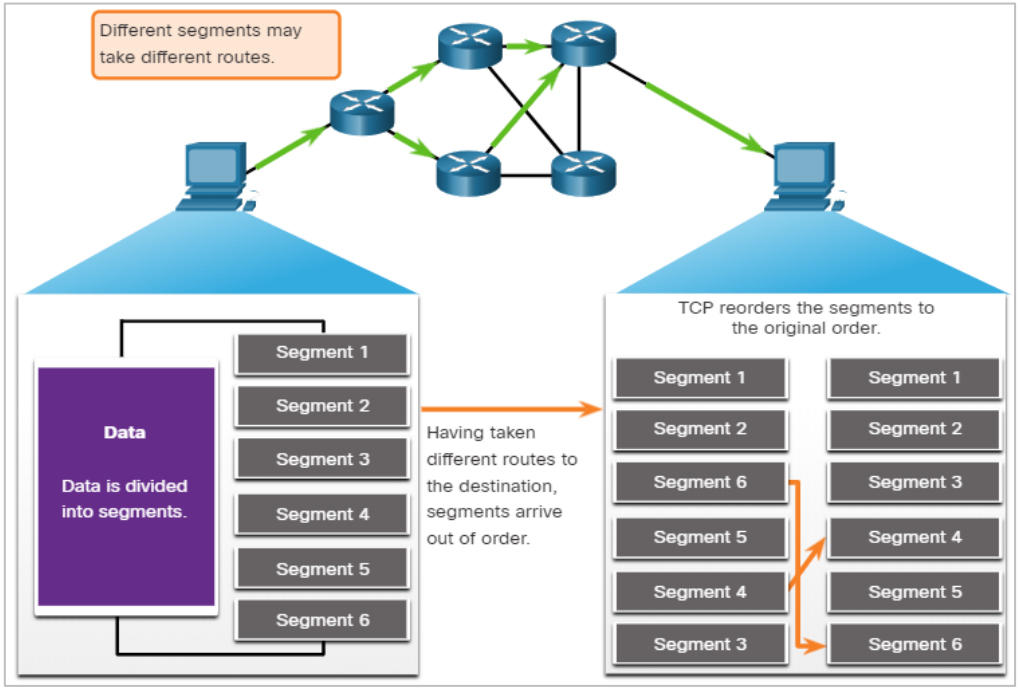

Data Loss and Retransmission

- TCP provides methods of managing the segment losses by retransmitting the segments for unacknowledged data.
- The sequence (SEQ) number and acknowledgement (ACK) number are used together to confirm receipt of the bytes of data contained in the transmitted segments.
- The SEQ number identifies the first byte of data in the segment being transmitted.
- TCP uses the ACK number sent back to the source to indicate the next byte that the receiver expects to receive. This is called expectational acknowledgement.
- Prior to later enhancements, TCP could only acknowledge the next byte expected.

Data Loss and Retransmission (Contd.)

- In the figure, Host A sends segments 1 through 10 to host B. If all the segments arrive except segments 3 and 4, host B would reply with acknowledgment specifying that the next segment expected is segment 3.
- Host A has no idea if any other segments arrived or not. It would resend segments 3 through 10.
- If all the resent segments arrived successfully, segments 5 through 10 would be duplicates. This can lead to delays, congestion, and inefficiencies.

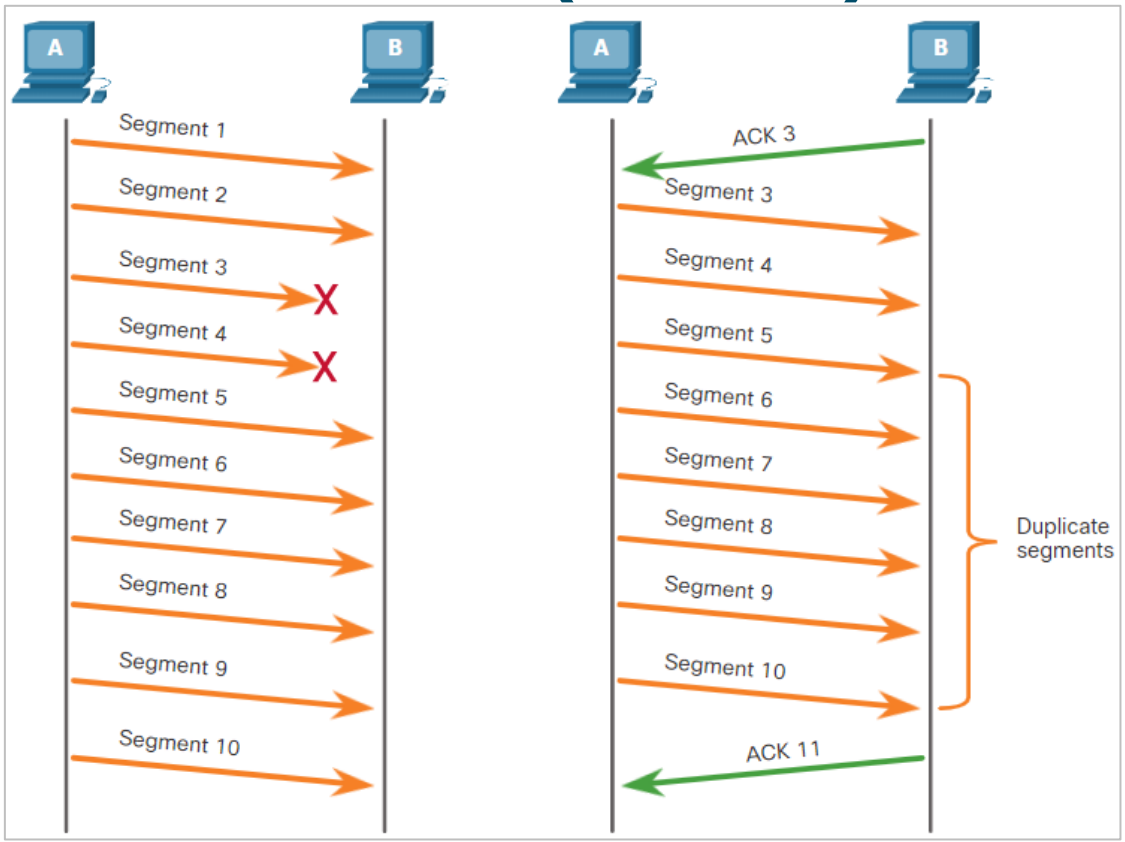

Data Loss and Retransmission (Contd.)

- Host operating systems employ an optional TCP feature called selective acknowledgment (SACK), negotiated during the three-way handshake.
- If both hosts support SACK, the receiver can acknowledge which segments (bytes) were received including any discontinuous segments.
- The sending host would only need to retransmit the missing data.
- In the figure, host A sends segments 1 through 10 to host B.
- If all the segments arrive except for segments 3 and 4, host B can acknowledge that it has received segments 1 and 2 (ACK 3), and selectively acknowledge segments 5 through 10 (SACK 5-10). Host A would only need to resend segments 3 and 4.

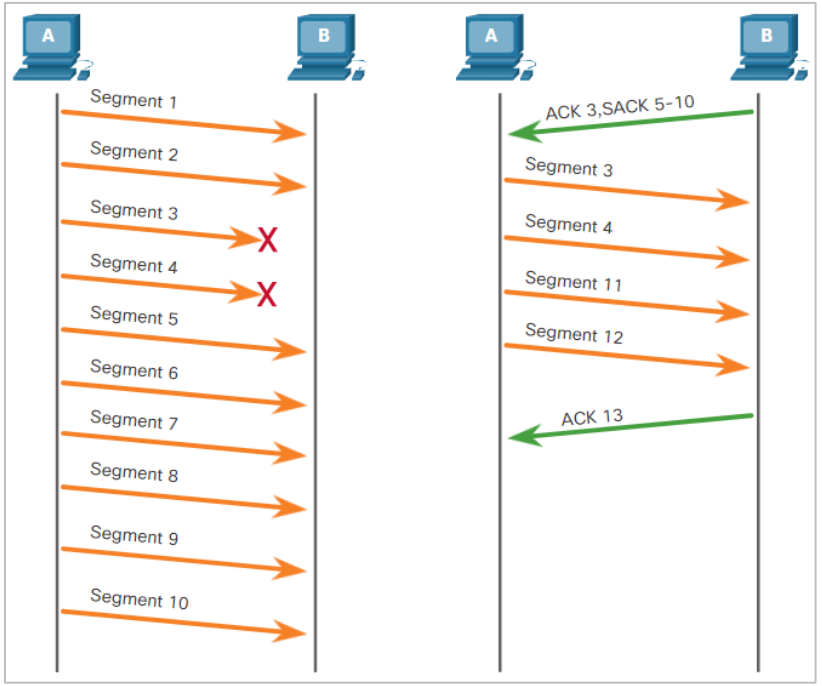

Flow Control - Window and ACKs

- TCP also provides mechanisms for flow control. Flow control is the amount of data that the destination can receive and process reliably.
- Flow control helps maintain the reliability of TCP transmission by adjusting the rate of data flow between source and destination for a given session.
- To accomplish this, the TCP header includes a 16-bit field called the window size.
- The window size that determines the number of bytes that can be sent before expecting an acknowledgment.
- The acknowledgment number is the number of the next expected byte.
- The window size is the number of bytes that the destination device of a TCP session can accept and process at one time.

Flow Control - Window and ACKs (Contd)

- The figure shows an example of window size and acknowledgments.
- The window size is included in every TCP segment so the destination can modify the window size at any time depending on buffer availability.
- The initial window size is agreed upon when the TCP session is established during the three-way handshake.
- The source device must limit the number of bytes sent to the destination device based on the window size of the destination. Only after the source receives an acknowledgment, it can continue sending more data for the session.

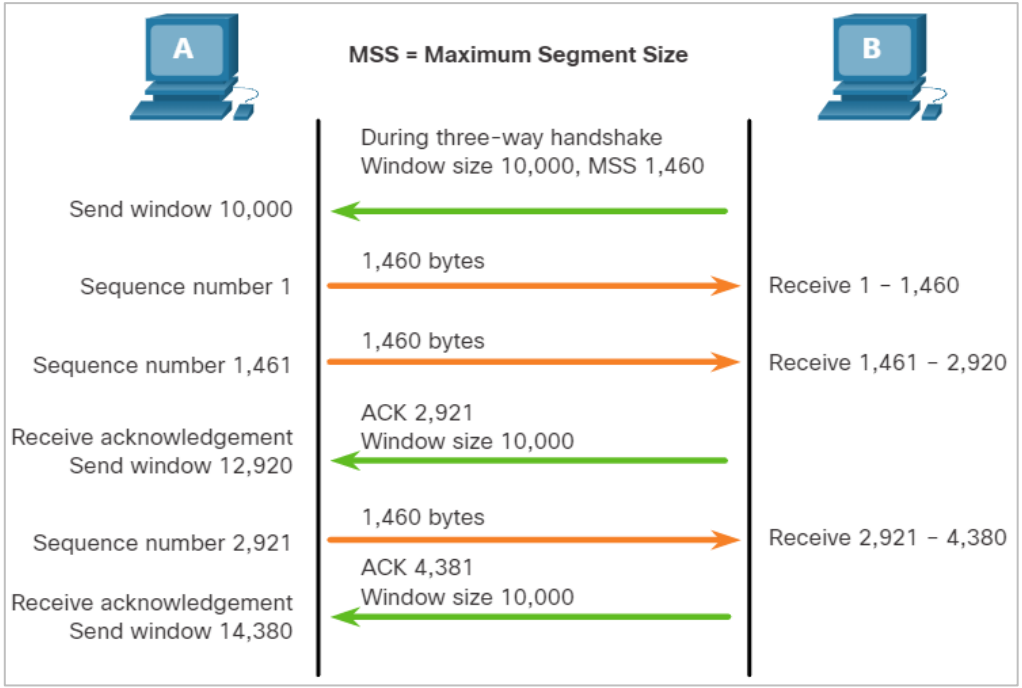

Flow Control - Window and ACKs (Contd)

- The destination will not wait for all the bytes for its window size to be received before replying with an acknowledgment.
- As the bytes are received and processed, the destination will send acknowledgments to inform the source that it can continue to send additional bytes.
- A destination sending acknowledgments as it processes bytes received, and the continual adjustment of the source send window, is known as sliding windows.
- If the availability of the destination's buffer space decreases, it may reduce its window size to inform the source to reduce the number of bytes it should send without receiving an acknowledgment.

***Note**: Devices today use the sliding windows protocol. The receiver sends an acknowledgment after every two segments it receives. The advantage of sliding windows is that it allows the sender to continuously transmit segments, as long as the receiver is acknowledging previous segments.*

TCP Flow Control - Maximum Segment Size

- In the figure, the source is transmitting 1,460 bytes of data within each TCP segment. This is the Maximum Segment Size (MSS) that the destination device can receive.
- The MSS is part of the options field in the TCP header that specifies the largest amount of data, in bytes, that a device can receive in a single TCP segment.
- The MSS size does not include the TCP header.
- The MSS is included during the three-way handshake.

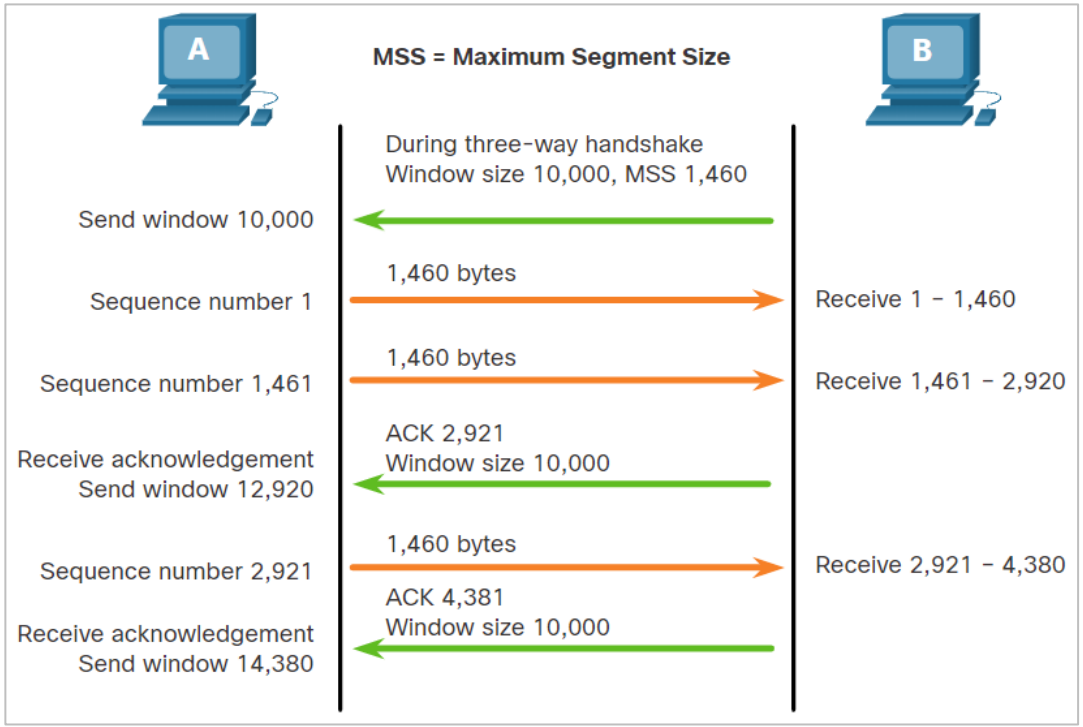

Flow Control – MSS (Contd.)

- A common MSS is 1,460 bytes when using IPv4. A host determines the value of its MSS field by subtracting the IP and TCP headers from the Ethernet maximum transmission unit (MTU).
- On an Ethernet interface, the default MTU is 1500 bytes. Subtracting the IPv4 header of 20 bytes and the TCP header of 20 bytes, the default MSS size will be 1460 bytes, as shown in the figure.

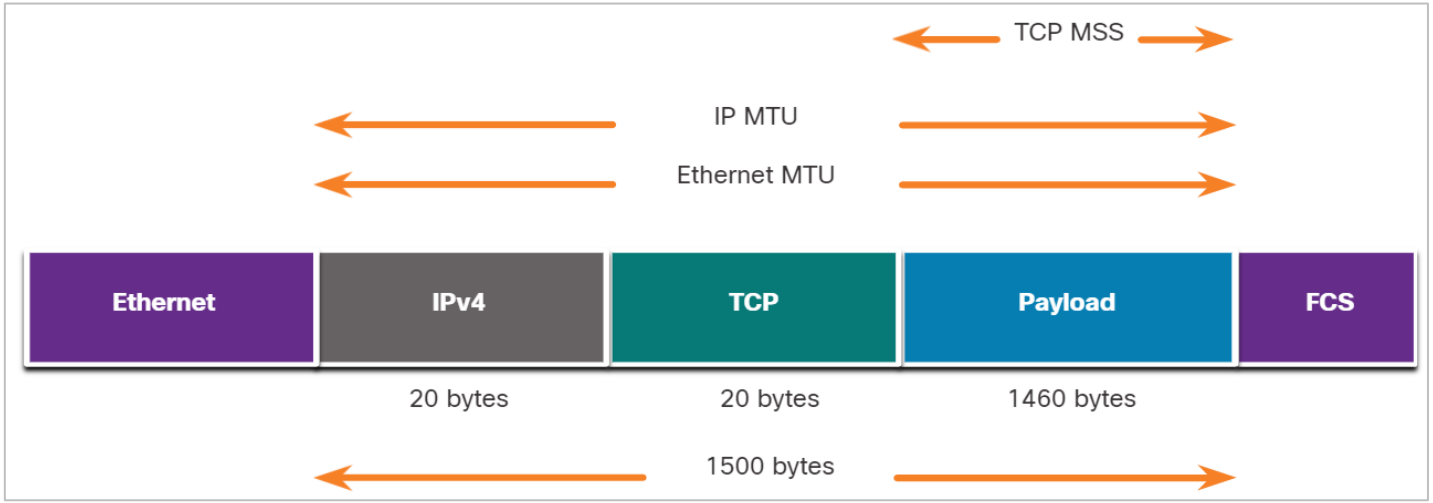

TCP Flow Control - Congestion Avoidance

- When congestion occurs on a network, it results in packets being discarded by the overloaded router.
- When packets containing TCP segments do not reach their destination, they are left unacknowledged.
- By determining the rate at which TCP segments are sent but not acknowledged, the source can assume a certain level of network congestion.
- Whenever there is congestion, retransmission of lost TCP segments from the source will occur.
- If the retransmission is not properly controlled, the additional retransmission of the TCP segments can make the congestion even worse.
- Not only are new packets with TCP segments introduced into the network, but the feedback effect of the retransmitted TCP segments that were lost will also add to the congestion.
- To avoid and control congestion, TCP employs several congestion handling mechanisms, timers, and algorithms.

Flow Control – CA (Contd.)

- If the source determines that the TCP segments are either not being acknowledged or not acknowledged in a timely manner, then it can reduce the number of bytes it sends before receiving an acknowledgment.
- As shown in the figure, PC A senses there is congestion and therefore, reduces the number of bytes it sends before receiving an acknowledgment from PC B.
- Acknowledgment numbers are for the next expected byte and not for a segment. The segment numbers used are simplified for illustration purposes.

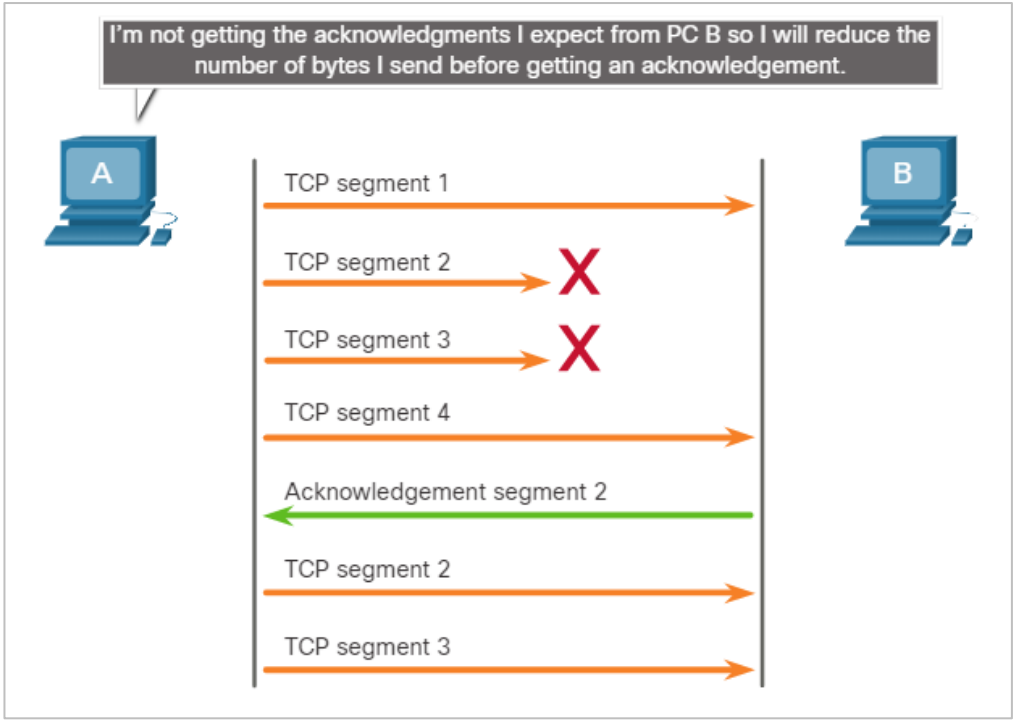

Lab – Exploring Nmap

- Port scanning is usually part of a reconnaissance attack.
- There are a variety of port scanning methods that can be used.
- We will explore how to use the Nmap utility. Nmap is a powerful network utility that is used for network discovery and security auditing.

Thank you! Questions?

Vladimír Veselý

updated: 2024-02-24

https://www.fit.vutbr.cz/research/groups/nes@fit

Module 10: Network Services

Instructor Materials

CyberOps Associate v1.0

Module Objectives

Module Title: Network Services

Module Objective: Explain how network services enable network functionality.

Topic Title	Topic Objective
DHCP	Explain how DHCP services enable network functionality.
DNS	Explain how DNS services enable network functionality.
NAT	Explain how NAT services enable network functionality.
File Transfer and Sharing Services	Explain how file transfer services enable network functionality.
Email	Explain how email services enable network functionality.
HTTP	Explain how HTTP services enable network functionality.

10.1 DHCP

Dynamic Host Configuration Protocol

- Two types of addressing:
  - **Dynamic** – Dynamic Host Configuration Protocol (DHCP) for IPv4 service automates the assignment of IPv4 addresses, subnet masks, gateways, and other IPv4 networking parameters.
  - **Static** – The network administrator manually enters IP address information on hosts.
- When a host connects to the network, the DHCP server chooses an address from a configured range of addresses called a pool and assigns it to the host.
- DHCP can allocate IP addresses for a configurable period of time, called a lease period.

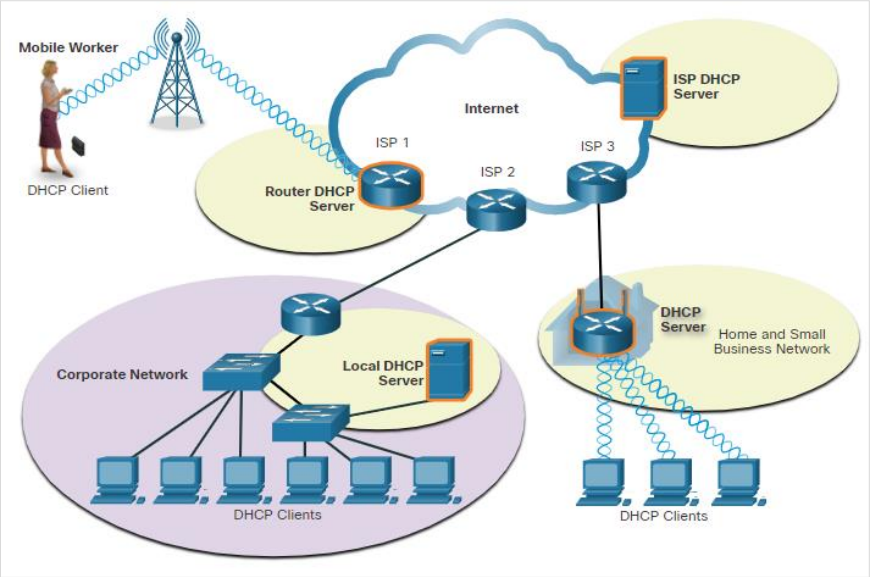

Medium-to-large networks – DHCP server is a local PC-based server
Home network – DHCP server is on the local router connecting the home network to the ISP.

DHCP Operation

- DHCP operation includes: DHCPDISCOVER, DHCPOFFER, DHCPREQUEST, DHCPACK, and DHCPNAK.
- When DHCP-configured device connects to the network, the client broadcasts a **DHCPDISCOVER** message to identify any available DHCP servers on the network.
- A DHCP server replies with a **DHCPOFFER** message, which offers a lease to the client.
- The client sends a **DHCPREQUEST** message that identifies the explicit server and lease offer that the client is accepting.

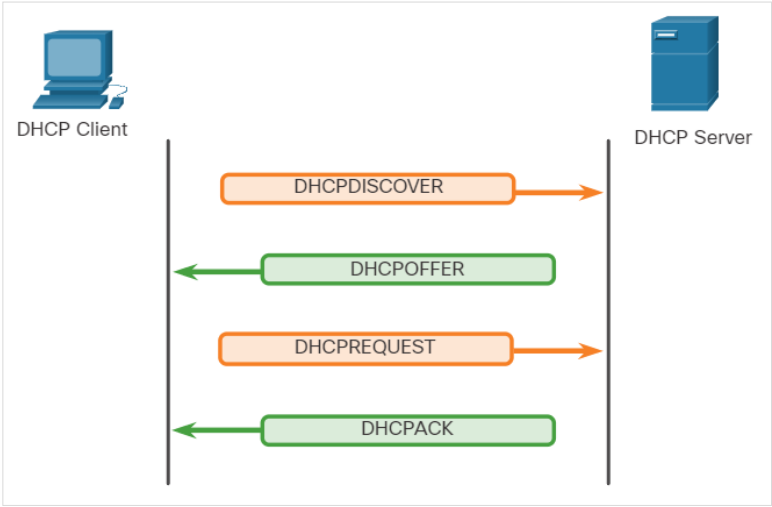

- If the IPv4 address requested by the client, or offered by the server, is still available, the server returns the **DHCPACK** message. If the offer is no longer valid, then the selected server responds with a **DHCPNAK** message. If a **DHCPNAK** message is returned, then the selection process begins again with a new **DHCPDISCOVER** message being transmitted.

DHCP Message Format

- The DHCPv4 message format is used for all DHCPv4 transactions.
- The DHCPv4 messages are encapsulated within the UDP transport protocol.
- The below table lists the fields covered in the structure of the DHCPv4 message.

Fields in the structure of DHCPv4 Message		
Operation (OP) Code	Seconds	Gateway IP Address
Hardware Type	Flags	Client Hardware Address .
Hardware Address Length	Client IP Address	Server Name
Hops	Your IP Address	Boot Filename
Transaction Identifier	Server IP Address	DHCP Options

10.2 DNS

DNS Overview

- Domain Name System (DNS) provides domain names and their associated IP addresses.
- The DNS system consists of a global hierarchy of distributed servers that contain databases of name to IP address mappings.
- The client computer in the figure will send a request to the DNS server to get the IP address for www.cisco.com so that it can address packets to that server.
- Malicious DNS traffic can be detected through protocol analysis and the inspection of DNS monitoring information.

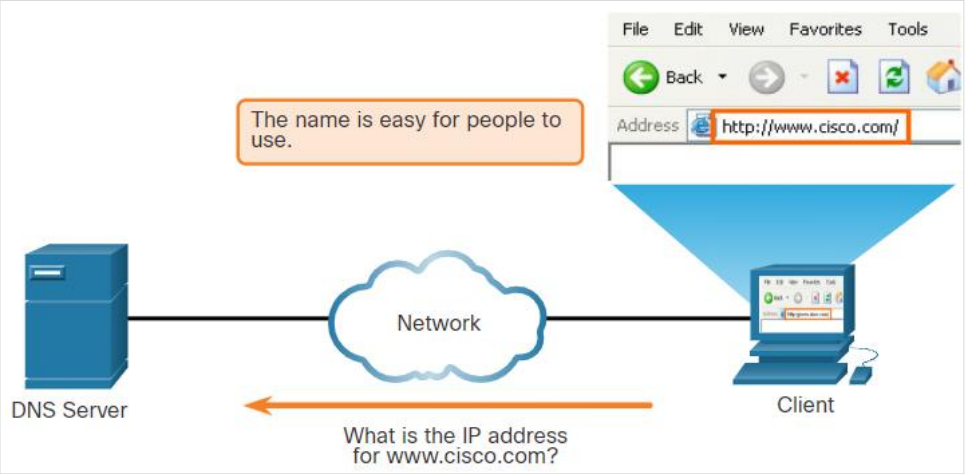

DNS Resolves Names to IP Addresses

The DNS Domain Hierarchy

- DNS consists of a hierarchy of generic top-level domains and numerous country-level domains.
- The second-level domains are represented by a domain name that is followed by a top-level domain.
- Subdomains are found at the next level of the DNS hierarchy and represent some division of the second-level domain.
- Fourth level domain can represent a host in a subdomain.
- Top-level domains represent either the type of organization or country of origin. Examples: (**.org)** - a non-profit organization, **(.au)** – Australia.

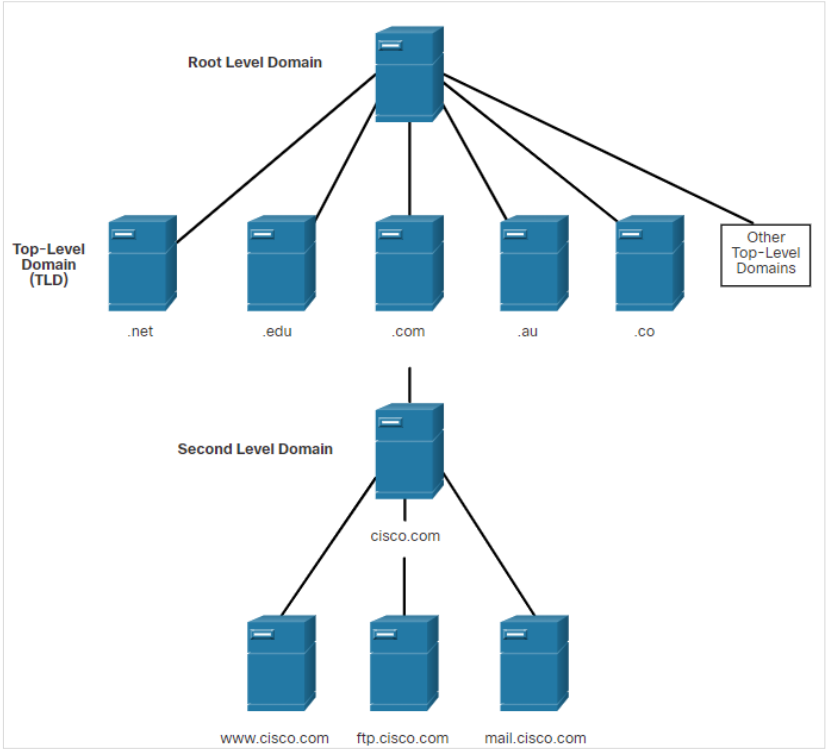

DNS Hierarchy

The DNS Lookup Process

- To resolve a name to an IP address, the resolver, will first check its local DNS cache. If the mapping is not found, a query will be issued to the DNS server .
- If the mapping is not found there, the DNS server will query other higher-level DNS servers that are authoritative for the top-level domain in order to find the mapping. These are known as **recursive queries**.
- The caching DNS servers can resolve recursive queries without forwarding the queries to higher level servers.

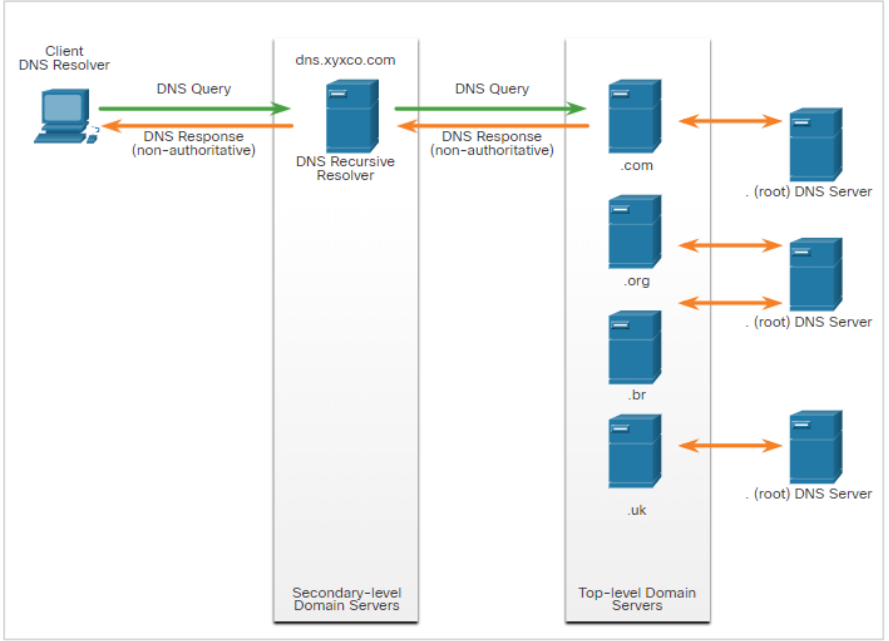

- If a server requires data for a zone, it will request a transfer of that data from an authoritative server for that zone. The process of transferring DNS data between servers is known as zone transfer.

The DNS Lookup Process(Contd.)

Steps involved in DNS resolution:

Step 1 - The user types an FQDN into a browser application Address field.

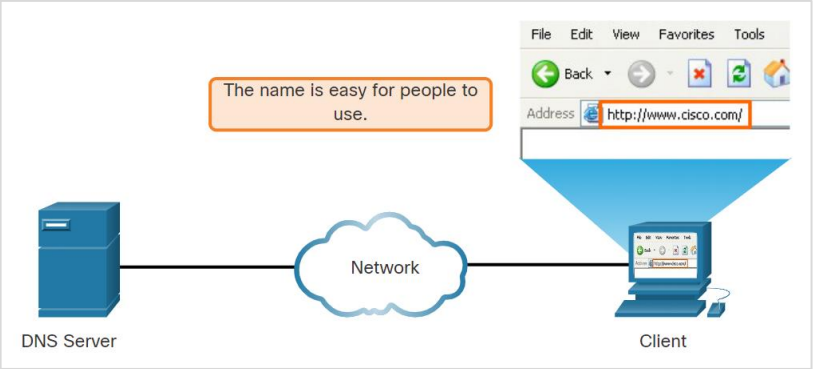

Step 2 - A DNS query is sent to the designated DNS server for the client computer.

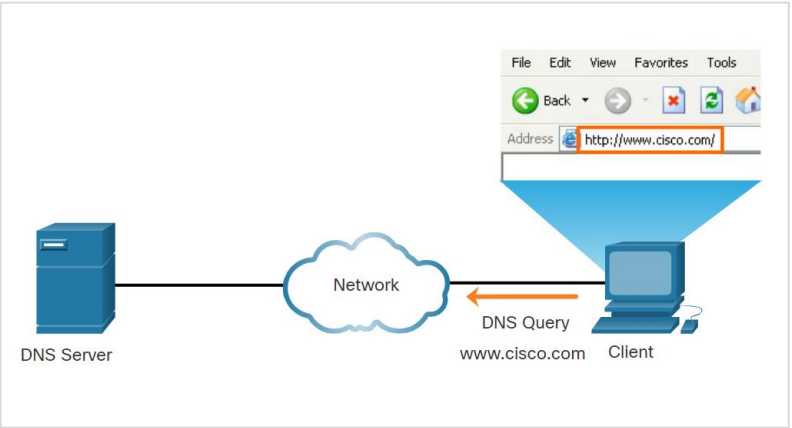

The DNS Lookup Process(Contd.)

Steps involved in DNS resolution:

Step 3 - The DNS server matches the FQDN with its IP address.

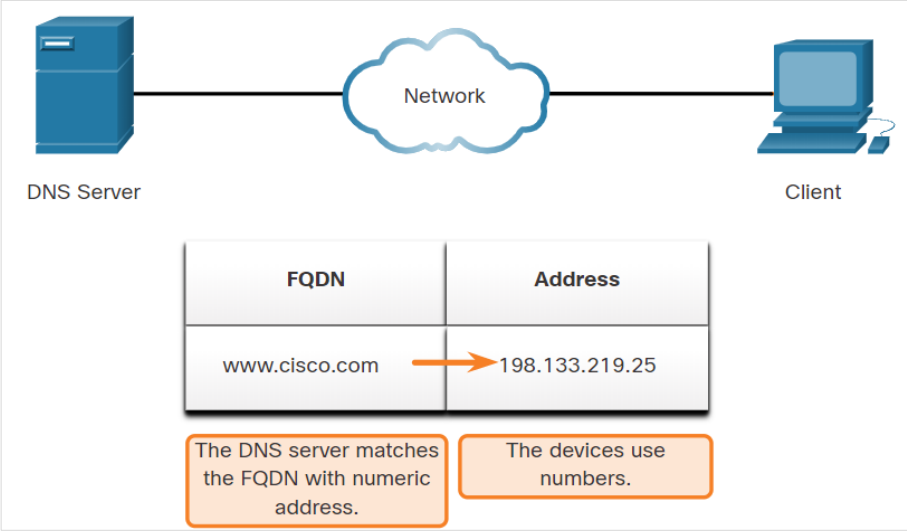

Step 4 - The DNS query response is sent back to the client with the IP address for the FQDN.

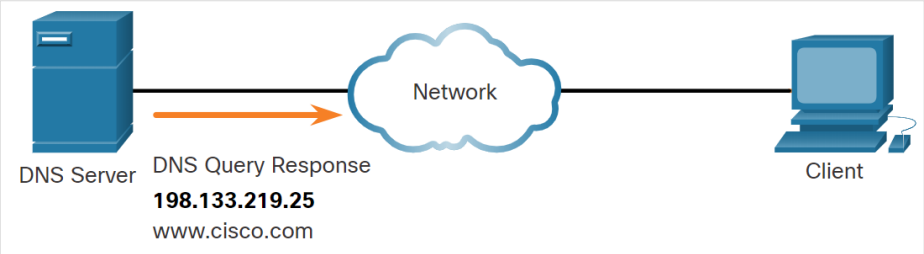

The DNS Lookup Process(Contd.)

Steps involved in DNS resolution:

Step 5 - The DNS server matches the FQDN with its IP address.

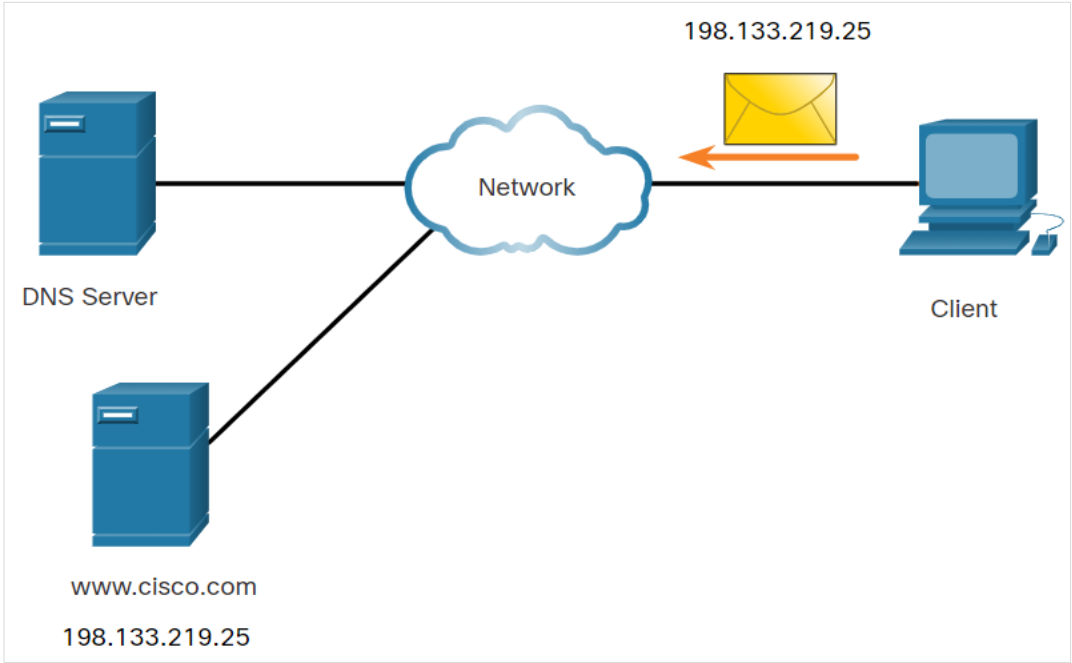

DNS Message Format

- DNS uses UDP port 53 for DNS queries and responses.
- If a DNS response exceeds 512 bytes, Dynamic DNS (DDNS) is used.
- The DNS protocol communications use a single format called a **message**.
- DNS uses the same message format for all types of client queries and server responses, error messages, and transfer of resource record information.

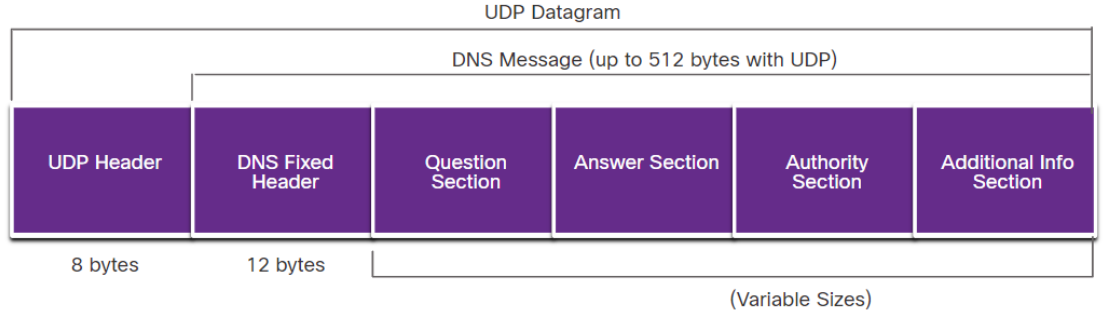

DNS Message Format (Contd.)

Sections of DNS message format :

DNS message section	Description
Question	The question for the server. It contains the domain name to be resolved, the class of domain, and the query type.
Answer	The DNS resource record, or RR, for the query including the resolved IP address depending on the RR type.
Authority	Contains the RRs for the domain authority.
Additional	Relevant to query responses only. Consists of RRs that hold additional information that will make query resolution more efficient

Dynamic DNS

- Dynamic DNS (DDNS) allows a user or organization to register an IP address with a domain name as in DNS.
- The subdomain is mapped to the IP address of the user's server, or home router connection to the internet.
- When a change is detected, the DDNS provider is immediately informed of the change and the mapping between the user's subdomain and the internet IP address is immediately updated.

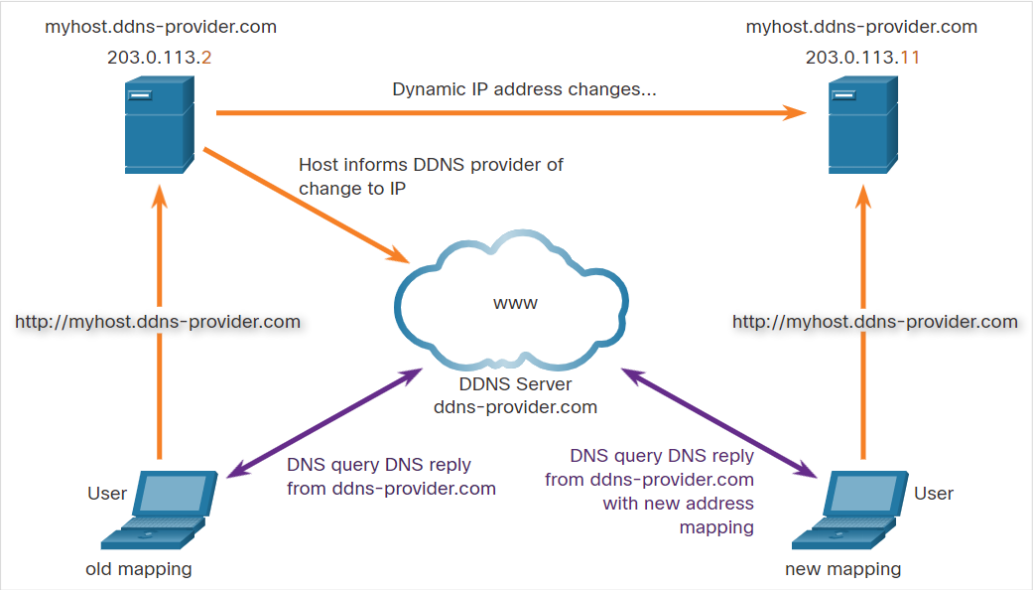

- The DDNS provider service supplies that IP address to the resolver's second level DNS server. This DNS server, either at the organization or ISP, provides the DDNS IP address to the resolver.

The WHOIS Protocol

- WHOIS is a TCP-based protocol that is used to identify the owners of internet domains through the DNS system.
- The WHOIS application uses a query, in the form of a FQDN.
- WHOIS is a starting point for identifying potentially dangerous internet locations that may have been reached through the network.
- ICANN Lookup, an internet-based WHOIS tool, is used to obtain the registration record a URL.

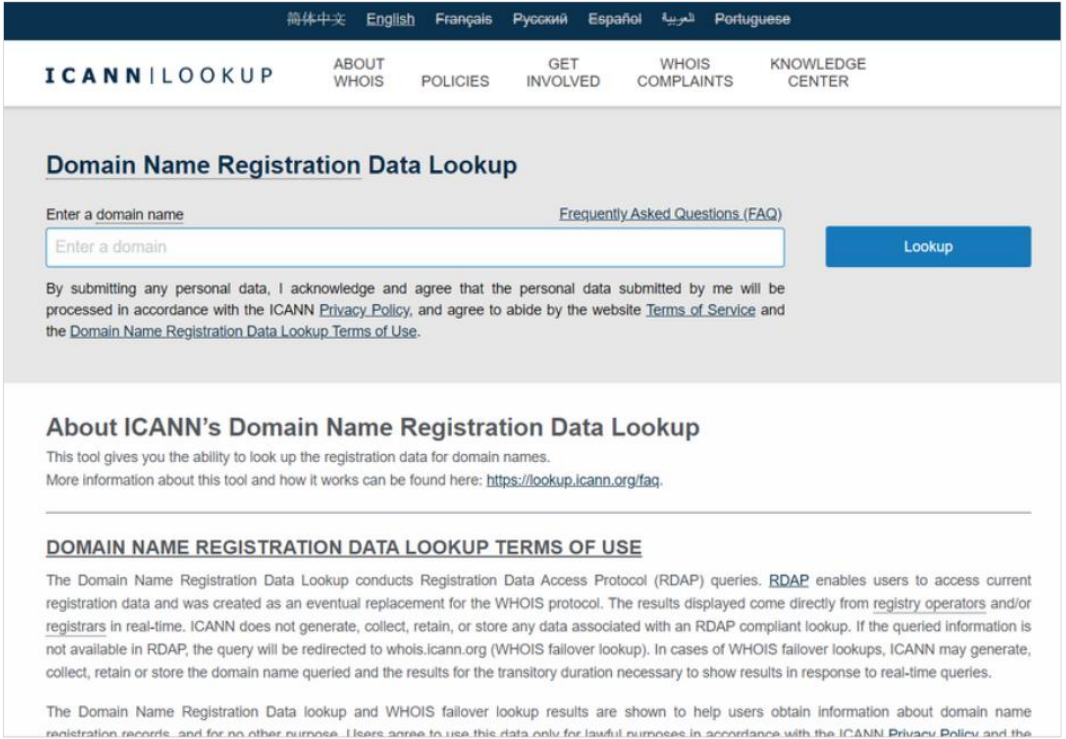

DNS Security

- DNS is now used by many types of malware. Some varieties of malware use DNS to communicate with command-and-control (CnC) servers and to exfiltrate data in traffic disguised as normal DNS queries.
- Malware could encode stolen data as the subdomain portion of a DNS lookup for a domain where the nameserver is under control of an attacker.
- A DNS lookup for ‘long-string-of-exfiltrated-data.example.com’ would be forwarded to the nameserver of example.com, which would record ‘long-string-of-exfiltrated-data’ and reply back to the malware with a coded response. This use of the DNS subdomain is shown in the figure. The exfiltrated data is the encoded text shown in the box. The threat actor collects the encoded data, decodes and combines it, and now has access to an entire data file.

DNS (Contd.)

- It is likely that the subdomain part of such requests would be much longer than usual requests. Cyber analysts can use the distribution of the lengths of subdomains within DNS requests to construct a mathematical model that describes normality.
  - They can then use this to compare their observations and identify an abuse of the DNS query process. For example, it would not be normal to see a host on the network sending a query to aW4gcGxhY2UgdG8gcHJvdGVjdC.example.com.
  - DNS queries for randomly generated domain names, or extremely long random-appearing subdomains, should be considered suspicious, especially if their occurrence spikes dramatically on the network.
- DNS proxy logs can be analyzed to detect these conditions.
- Alternatively, services such as the Cisco Umbrella passive DNS service can be used to block requests to suspected CnC and exploit domains

Lab - Using Wireshark to Examine a UDP DNS Capture

- In this lab, you will complete the following objectives:
  - Communicate with a DNS server by sending a DNS query using the UDP transport protocol.
  - Use Wireshark to examine the DNS query and response exchanges with the same server.

10.4 File Transfer and Sharing Services

FTP and TFTP

- FTP allows data transfers between a client and a server.
- An FTP client runs on a computer and is used to push and pull data from an FTP server.
- FTP connections between the client and server:
  - **Control Connection**: The client opens the first connection to the server for control traffic.
  - **Data Connection**: The client opens the second connection to the server for data traffic.
- Trivial File Transfer Protocol (TFTP) is a simplified file transfer protocol that uses well-known UDP port number 69.

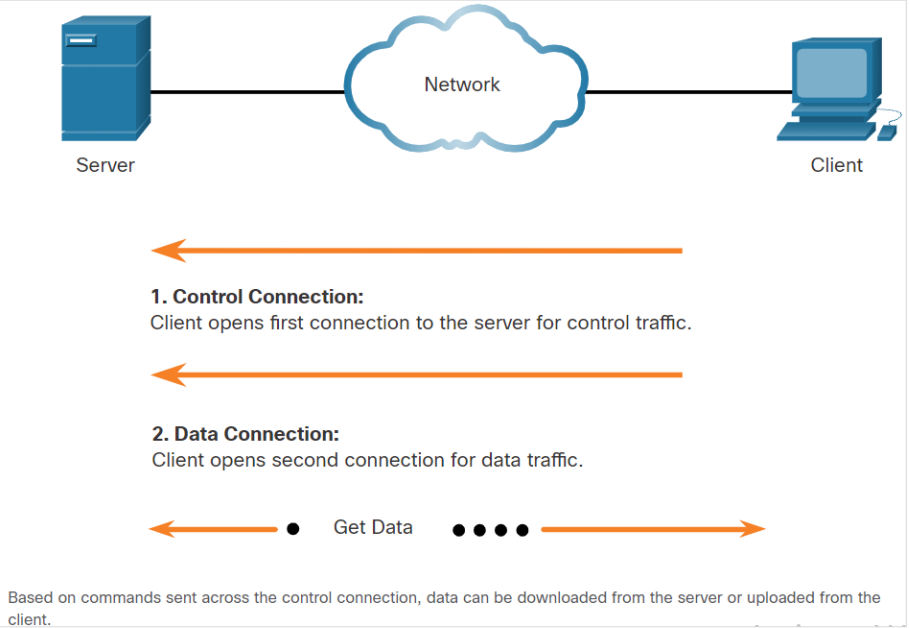

SMB

- The Server Message Block (SMB) is a client/server file sharing protocol that describes the structure of shared network resources.
- SMB is a client/server, request-response protocol.
- Servers can make their own resources available to clients on the network.

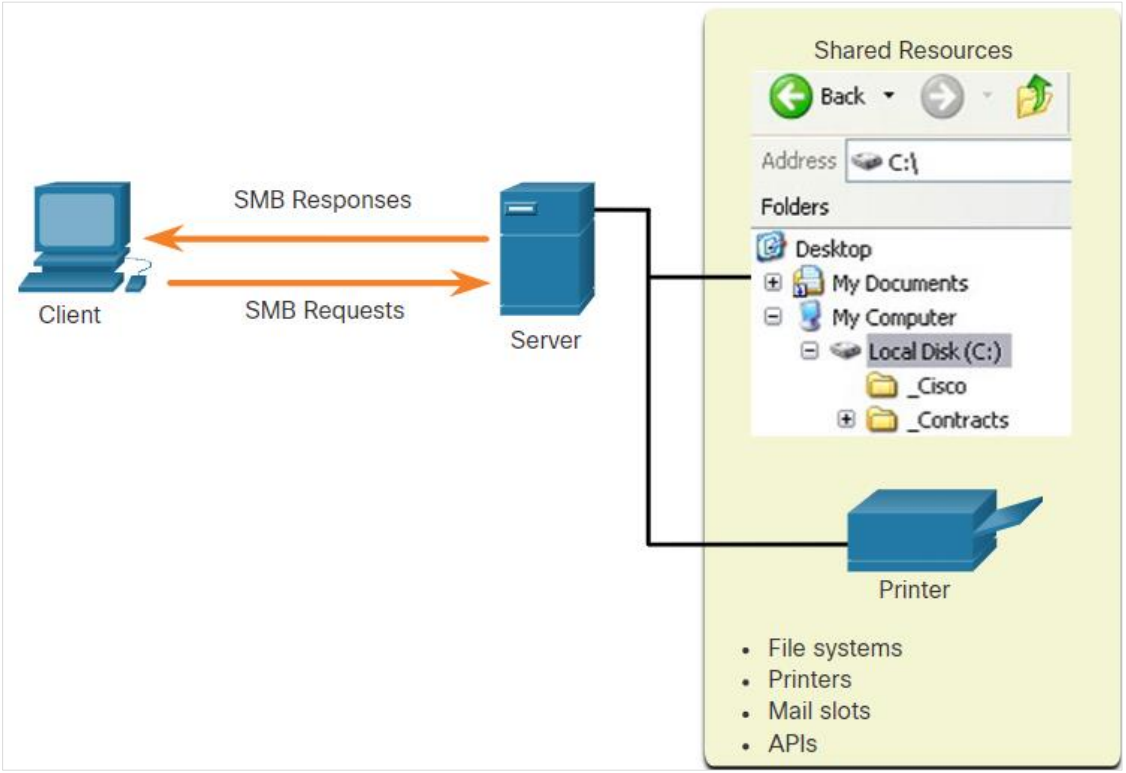

SMB (Contd.)

- SMB messages can start, authenticate, and terminate sessions, control file and printer access, and allow an application to send or receive messages to or from another device.
- SMB file sharing and print services have become the mainstay of Microsoft networking.
- A file may be copied from PC to PC with Windows Explorer using the SMB protocol.

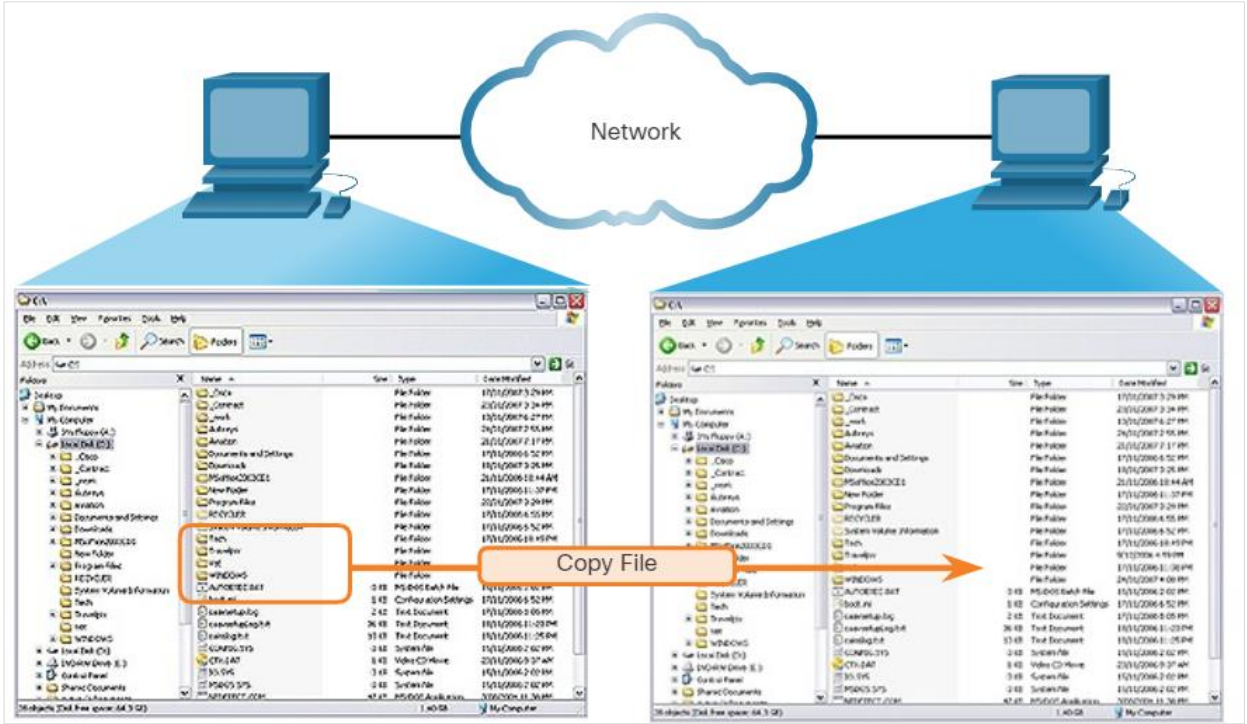

Lab - Using Wireshark to Examine TCP and UDP Captures

- In this lab, you will complete the following objectives:
  - Identify TCP Header Fields and Operation Using a Wireshark FTP Session Capture
  - Identify UDP Header Fields and Operation Using a Wireshark TFTP Session Capture

10.5 Email

Email protocols

- Email is a store-and-forward method of sending, storing, and retrieving electronic messages across a network.
- Email clients communicate with mail servers to send and receive email.
- Email supports three separate protocols for operation: Simple Mail Transfer Protocol (SMTP), Post Office Protocol (POP), and IMAP.
- A client retrieves email using one of the two application layer protocols: POP or IMAP.

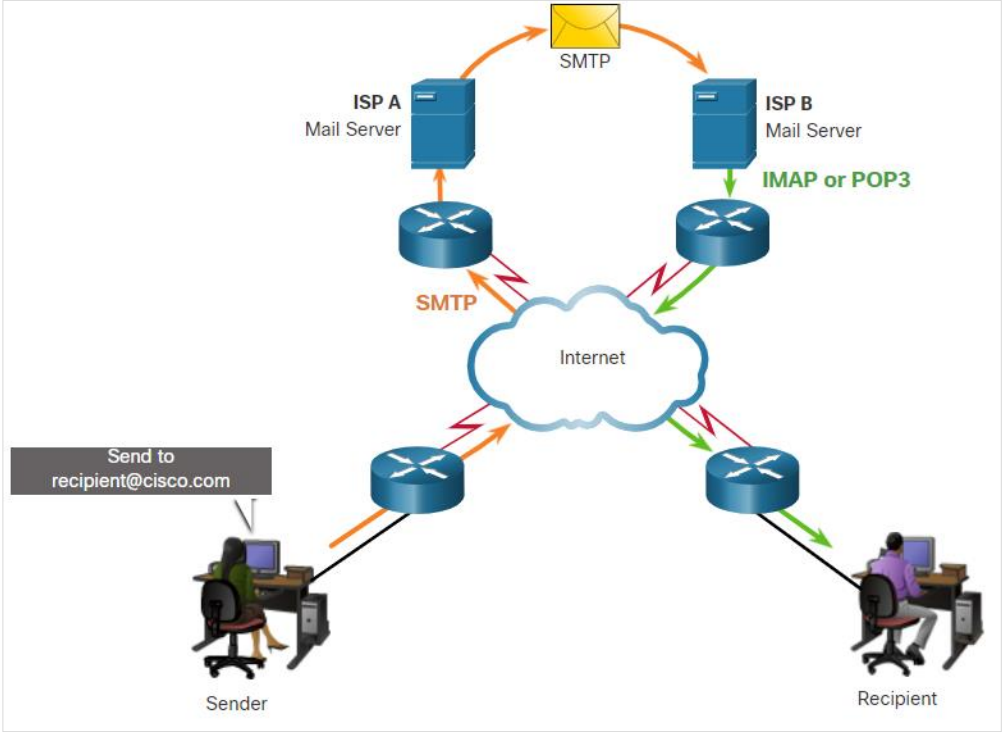

SMTP

- SMTP message formats require a message header and a message body.
- When a client sends an email, the client SMTP process connects with a server SMTP process on a well-known port 25.
- When the server receives the message, it either places the message in a local account, if the recipient is local, or forwards the message to another mail server for delivery.
- Periodically, the server checks the queue for messages and attempts to send them again. If the message is still not delivered after a predetermined expiration time, it is returned to the sender as undeliverable.

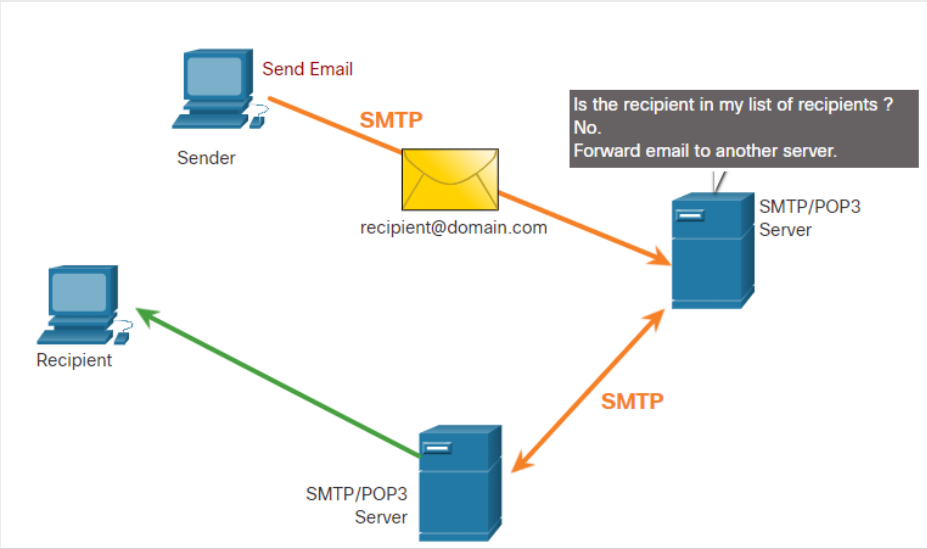

POP3

- POP3 is used by an application to retrieve a mail from a mail server.
- With POP3, email messages are downloaded to the client and removed from the server.
- The server starts the POP3 service by passively listening on TCP port 110 for client connection requests.
- The client sends a request to establish a TCP connection with the server.
- Once the connection is established, the POP3 server sends a greeting. The client and POP3 server then exchange commands and responses until the connection is closed or aborted.

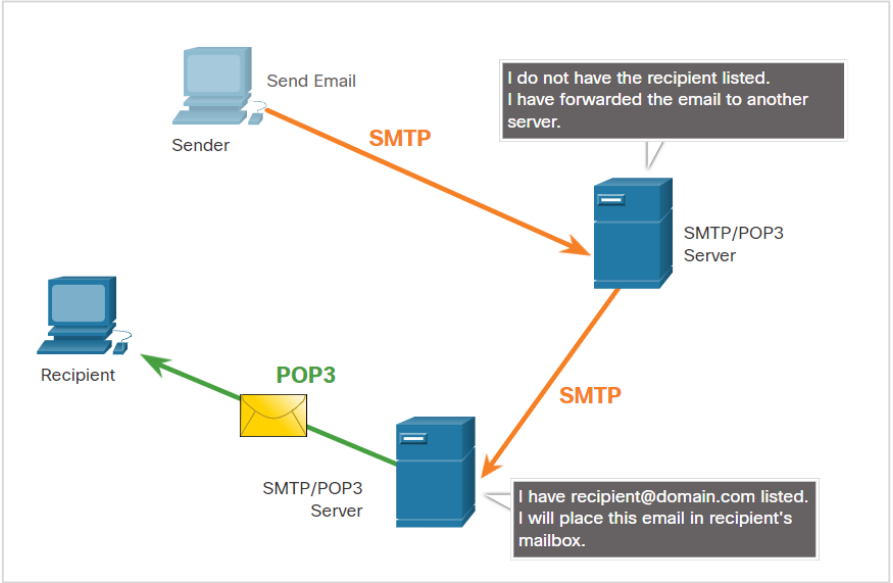

IMAP

- IMAP is the protocol that describes a method to retrieve email messages.
- When the user connects to an IMAP-capable server, copies of the messages are downloaded to the client application. The original messages are kept on the server until manually deleted.
- Users view copies of the messages in their email client software.
- Users can create a file hierarchy on the server to organize and store mail.
- When a user decides to delete a message, the server synchronizes that action and deletes the message from the server.

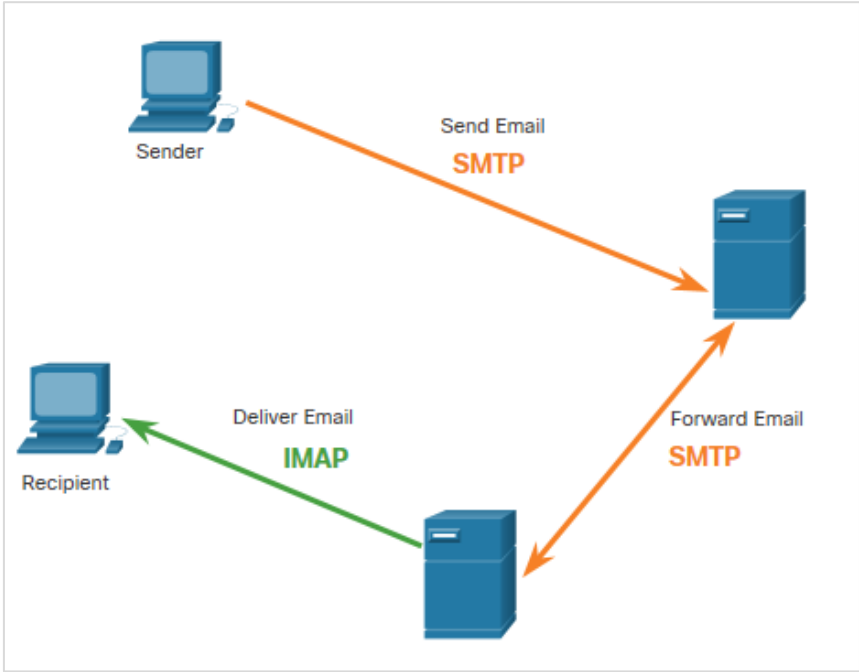

Email Protocols Security

- Email protocols such as SMTP, POP3, and IMAP can be used by threat actors to spread malware, exfiltrate data, or provide channels to malware CnC servers, as shown in the figure.
- SMTP sends data from a host to a mail server and between mail servers.
- IMAP and POP3 are used to download email messages from a mail server to the host computer. They are the application protocols that are responsible for bringing malware to the host.
- Security monitoring can identify when a malware attachment entered the network and which host it first infected.

Email Protocol Threats

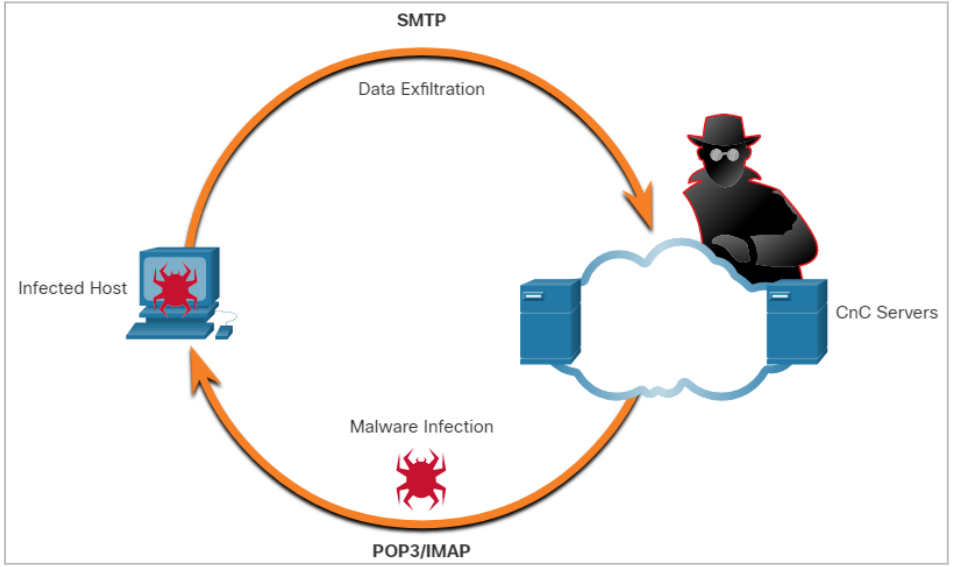

10.6 HTTP

HTTP and HTML

- When a web address or Uniform Resource Locator (URL) is typed into a web browser, the web browser establishes a connection with the web service that is using the HTTP protocol.
- Lets take a look on how a web page is opened in a browser.
  Example: http://www.cisco.com/index.html

Step 1: The browser interprets the three parts of the URL:

- http (the protocol or scheme)
- www.cisco.com (the server name)
- index.html (the specific filename requested)

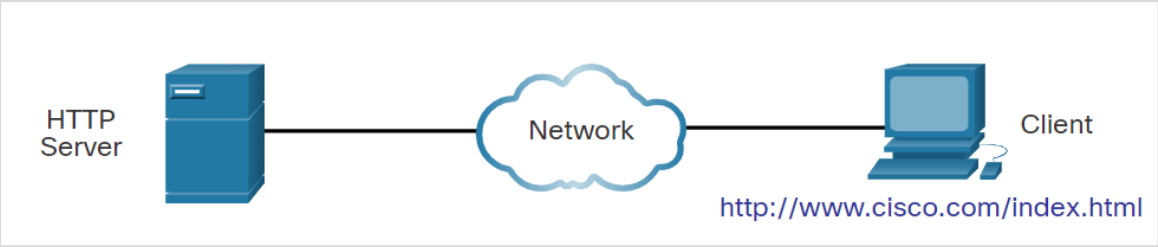

HTTP and HTML (Contd.)

- **Step 2:** The client initiates an HTTP request to a server by sending a GET request to the server and asks for the index.html file.

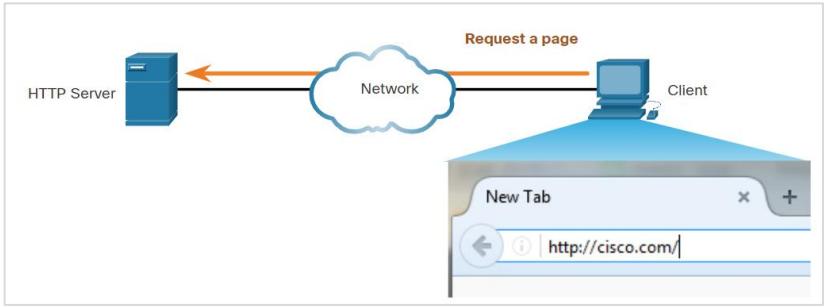

- **Step 3:** In response to the request, the server sends the HTML code for this web page to the browser.

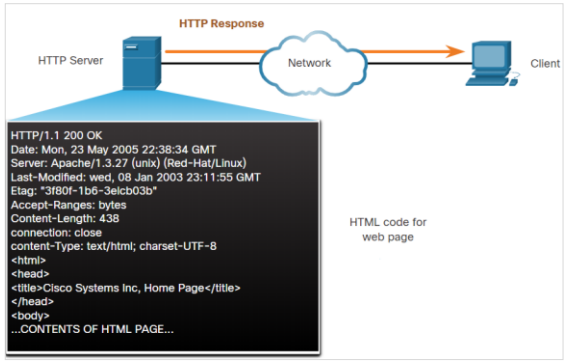

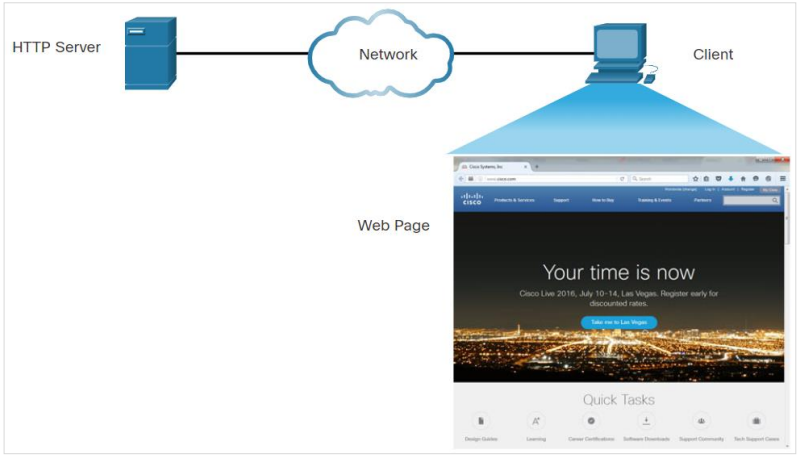

- **Step 4:** The browser deciphers the HTML code and formats the page for the browser window.

The HTTP URL

- HTTP URLs can specify the port on the server that should handle the HTTP methods.
- It can specify a query string and fragment.
- Query strings are preceded by a “?” character and typically consist of a series of name and value pairs.
- A fragment is preceded by a “#” character. It refers to a subordinate part of the resource that is requested in the URL.
- The parts of an HTTP URL are shown in the below figure:

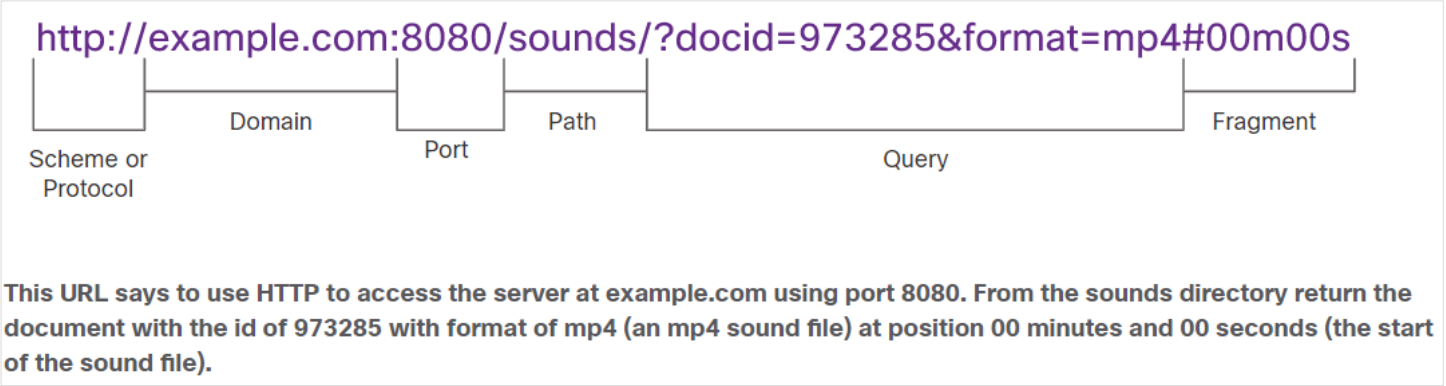

This URL says to use HTTP to access the server at example.com using port 8080. From the sounds directory return the document with the id of 973285 with format of mp4 (an mp4 sound file) at position 00 minutes and 00 seconds (the start of the sound file).

HTTP Operation

- HTTP is a request/response protocol that uses TCP port 80.It is flexible but not a secure protocol.
- When a client sends a request to a web server, it will use one of the six methods specified by HTTP:
  - GET
  - POST
  - PUT
  - DELETE
  - OPTIONS
  - CONNECT

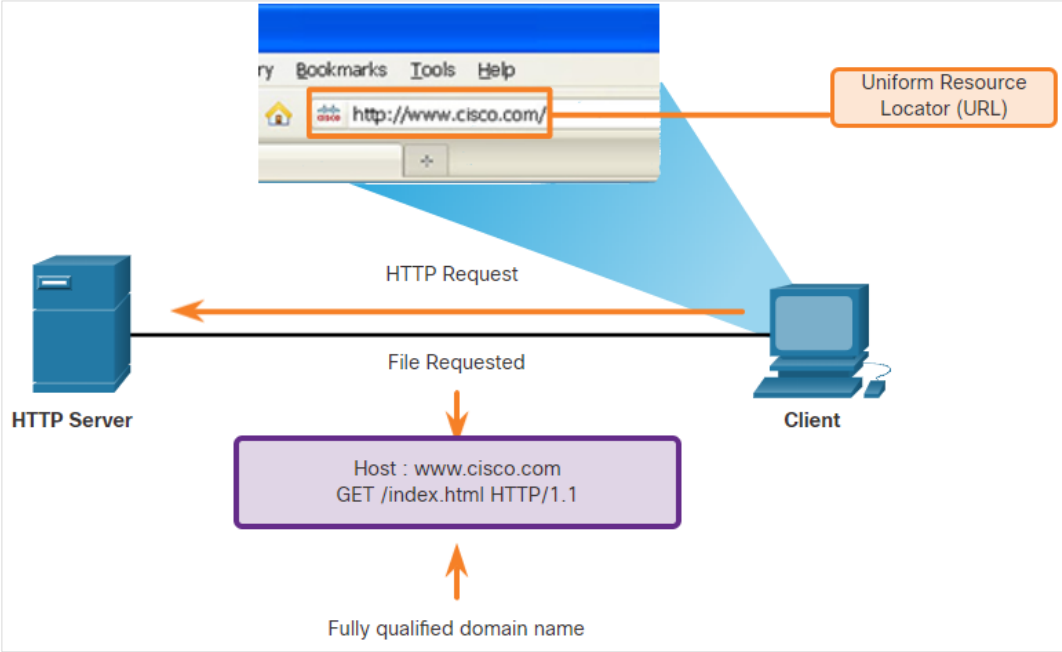

HTTP Status Codes

- The HTTP Status codes are numeric, with the first number in the code indicating the type of message.
- The five status code groups are **1xx** - Informational, **2xx** - Success, **3xx** - Redirection , **4xx** - Client Error and **5xx** - Server Error
- The below table explains some common status codes:

Code	Status	Meaning
1xx - Informational		
100	Continue	The client should continue with the request. The Server has verified that the request can be fulfilled.
2xx - Success		
200	OK	The request completed successfully.
202	Accepted	The request has been accepted for processing, but processing is not completed.

HTTP Status Codes (Contd.)

Code	Status	Meaning
4xx – Client Error		
403	Forbidden	The request is understood by the server, but the resource will not be fulfilled. This is possibly because the requester is not authorized to view the resource.
404	Not Found	The server could not find the requested resource. This can be caused by an out-of-date or incorrect URL.

HTTP/2

- The purpose of HTTP/2 is to improve HTTP performance by addressing latency issues that existed in the HTTP 1.1 version of the protocol.
- HTTP/2 uses the same header format as HTTP 1.1 and uses the same status codes.
- Few important features of HTTP/2 that a cybersecurity analyst must be aware of:
  - Multiplexing
  - Server PUSH
  - A binary protocol
  - Header compression

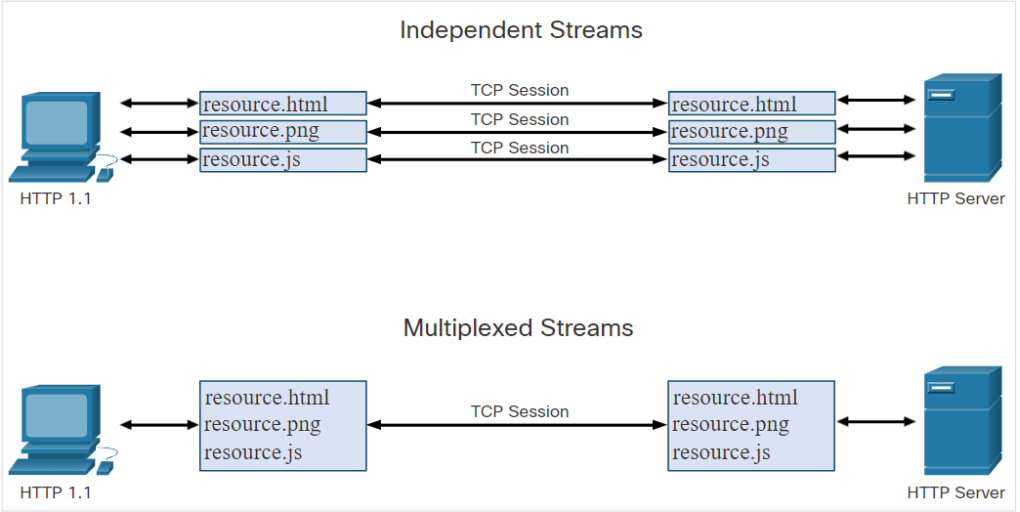

HTTP and HTTPS

- **Hypertext Transfer Protocol (HTTP)** is the backbone protocol of the World Wide Web.
- All information carried in HTTP/1.1 is transmitted in plaintext from the source computer to the destination on the internet.
- HTTP does not protect data from alteration or interception by malicious parties, which is a serious threat to privacy, identity, and information security.
- All browsing activity should be considered to be at risk.

Securing HTTP – HTTPS

- For secure communication across the internet, the HTTP Secure (HTTPS) protocol is used.
- HTTPS uses authentication and encryption to secure data as it travels between the client and the server.
- HTTPS uses the same client request-server response process as HTTP, but the data stream is encrypted with **Secure Socket Layer (SSL)**, or **Transport Layer Security (TLS)**, before being transported across the network.
  - This makes the HTTP data unreadable as it leaves the source computer until it reaches the server.
  - HTTPS is not a mechanism for web server security. It only secures HTTP protocol traffic while it is in transit.
- HTTPS/2 is specified to use HTTPS over TLS with the Application-Layer Protocol Negotiation (ALPN) extension for TLS 1.2 or newer.

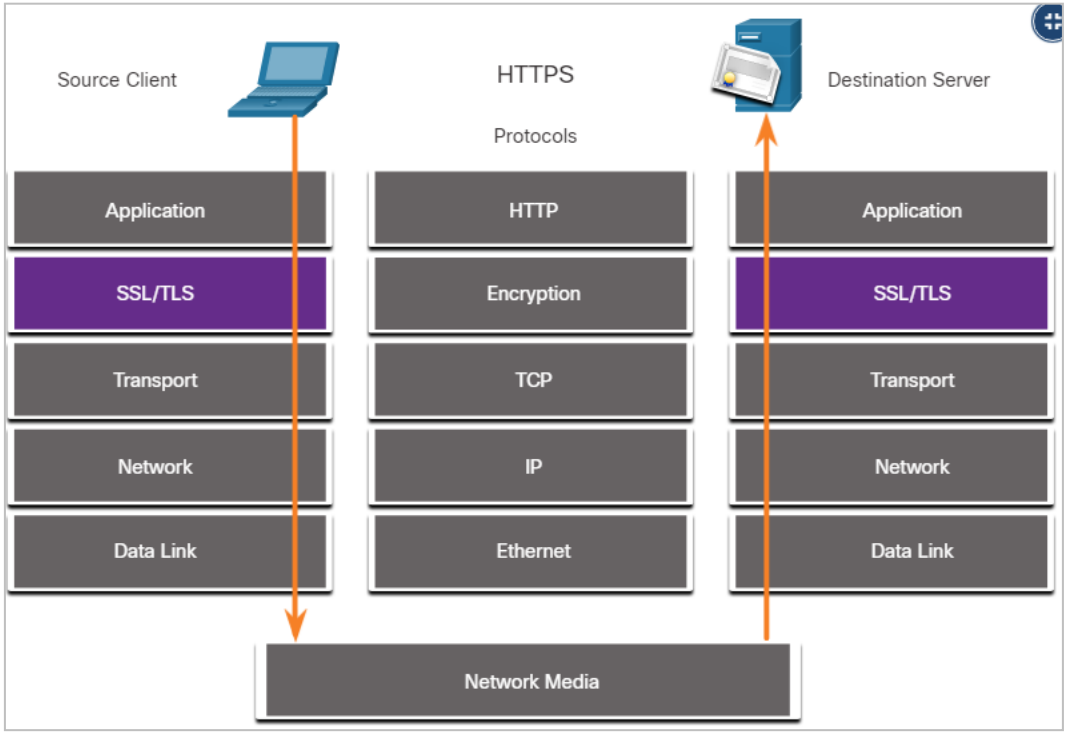

Securing HTTP – HTTPS(Contd.)

- Unfortunately, the encrypted HTTPS traffic complicates network security monitoring.
- Some security devices include SSL decryption and inspection; however, this can present processing and privacy issues.
- HTTPS adds complexity to packet captures due to the additional messaging involved in establishing the encrypted connection.
- This process is summarized in the figure and represents additional overhead on top of HTTP.

HTTPS Transactions

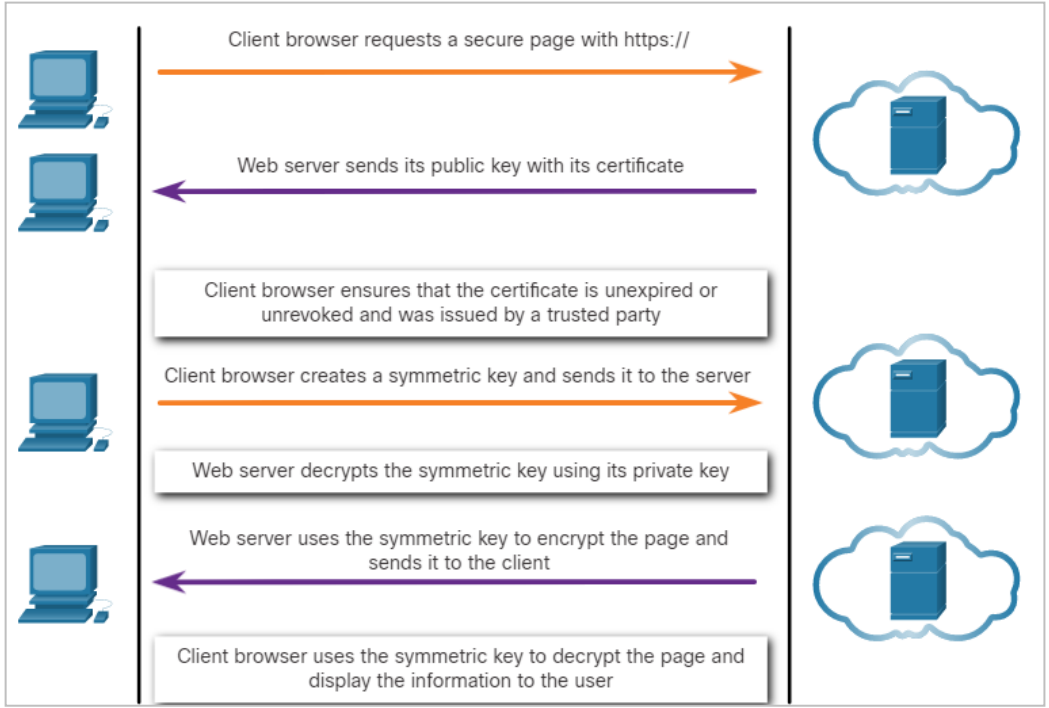

iFrame Injection

- A common exploit of HTTP is called **iFrame (inline frame) injection**.
  - In iFrame injection, a threat actor compromises a webserver and plants malicious code which creates an invisible iFrame on a commonly visited webpage.
  - When the iFrame loads, malware is downloaded, frequently from a different URL than the webpage that contains the iFrame code.
- Network security services, such as Cisco Web Reputation filtering, can detect when a website attempts to send content from an untrusted website to the host, even when sent from an iFrame.

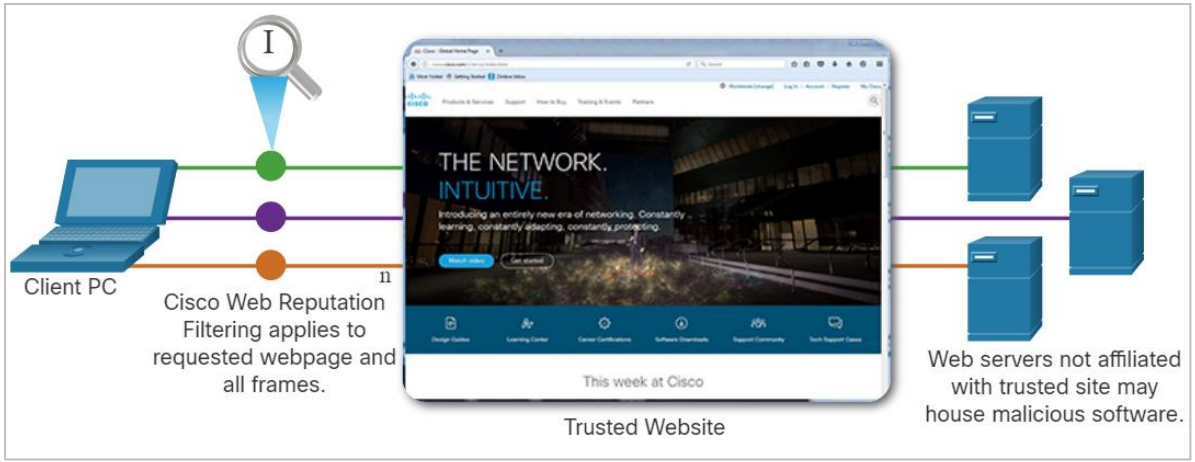

Lab - Using Wireshark to Examine HTTP and HTTPS Traffic

- In this lab, you will complete the following objectives:
  - Capture and view HTTP traffic
  - Capture and view HTTPS traffic

Thank you! Questions?

Vladimír Veselý

updated: 2024-02-24

https://www.fit.vutbr.cz/research/groups/nes@fit

Module 11: Network Communication Devices

Instructor Materials

CyberOps Associate v1.0

Module Objectives

Module Title: Network Communication Devices

Module Objective: Explain how network devices enable wired and wireless network communication.

Topic Title	Topic Objective
Network Devices	Explain how network devices enable network communication.
Wireless Communications	Explain how wireless devices enable network communication.

11.1 Network Devices

End Devices

- The most familiar network devices are end devices. An end device is either the source or destination of a message transmitted over the network.
- To distinguish one end device from another, each interface of end device on a network has an address.
- When an end device initiates communication, it uses the address of the destination to specify where to deliver the message.

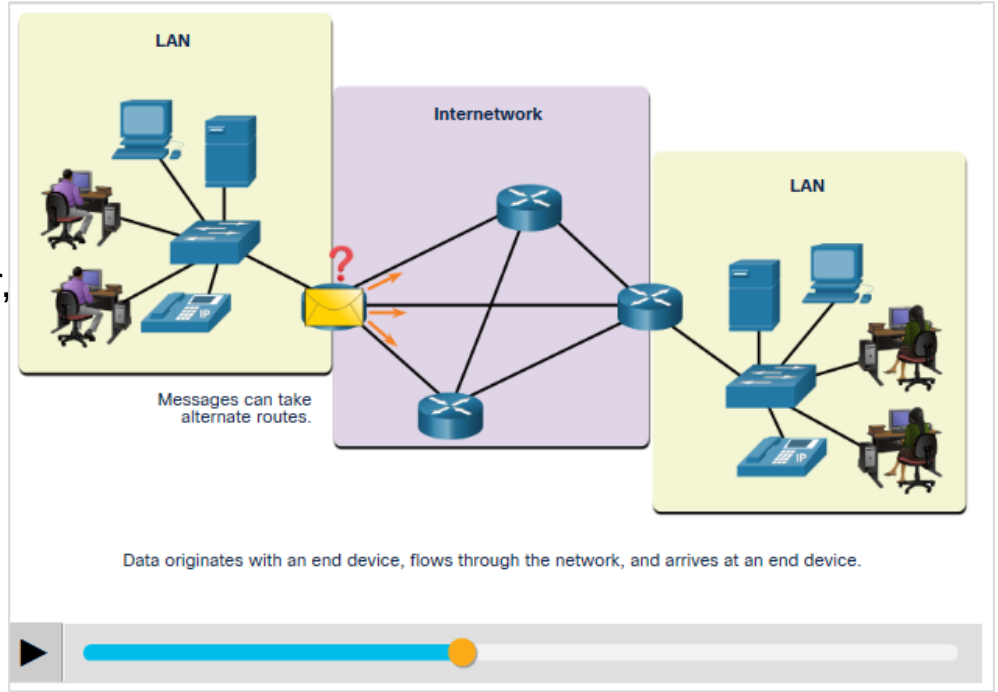

Routers

- Routers are devices that operate at the OSI network layer (Layer 3).
- Routers are used to interconnect remote sites.
- The **routing process** uses network routing tables, protocols, and algorithms to determine the most efficient path for forwarding an IP packet.

The Router Connection

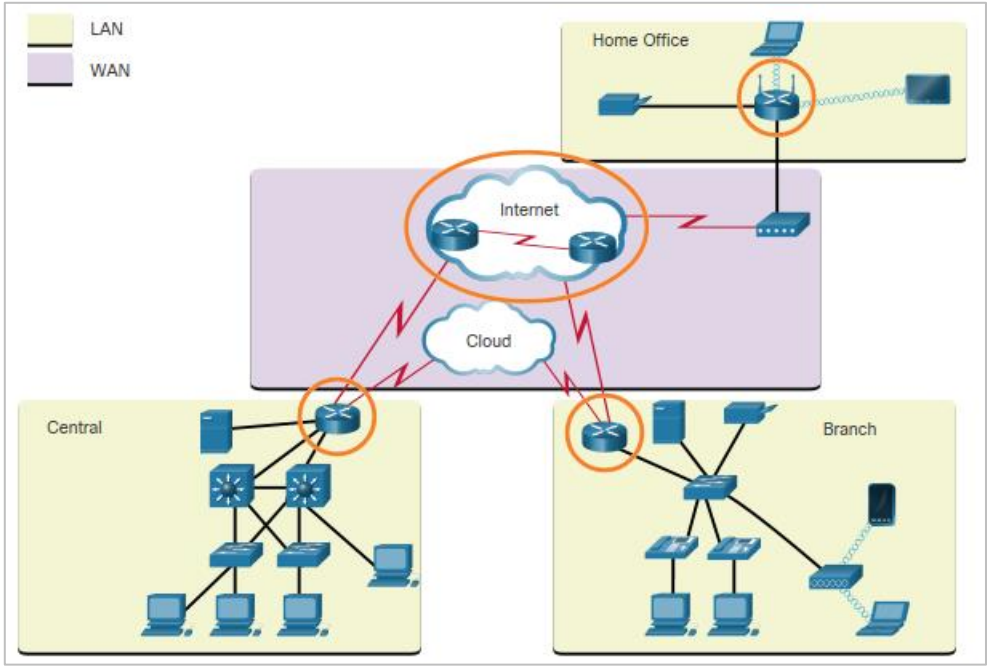

Routers (Contd.)

- Routers have two primary functions: path determination and packet forwarding
- To perform path determination, each router builds and maintains a **routing table** which is a database of known networks and how to reach them.
  - The routing table can be built manually and contain static routes or can be built using a dynamic routing protocol.
- Packet forwarding is accomplished by using a switching function.
  - Switching is the process used by a router to accept a packet on one interface and forward it out of another interface.
  - A primary aim of the switching function is to encapsulate packets in the appropriate data link frame type for the outgoing data link.

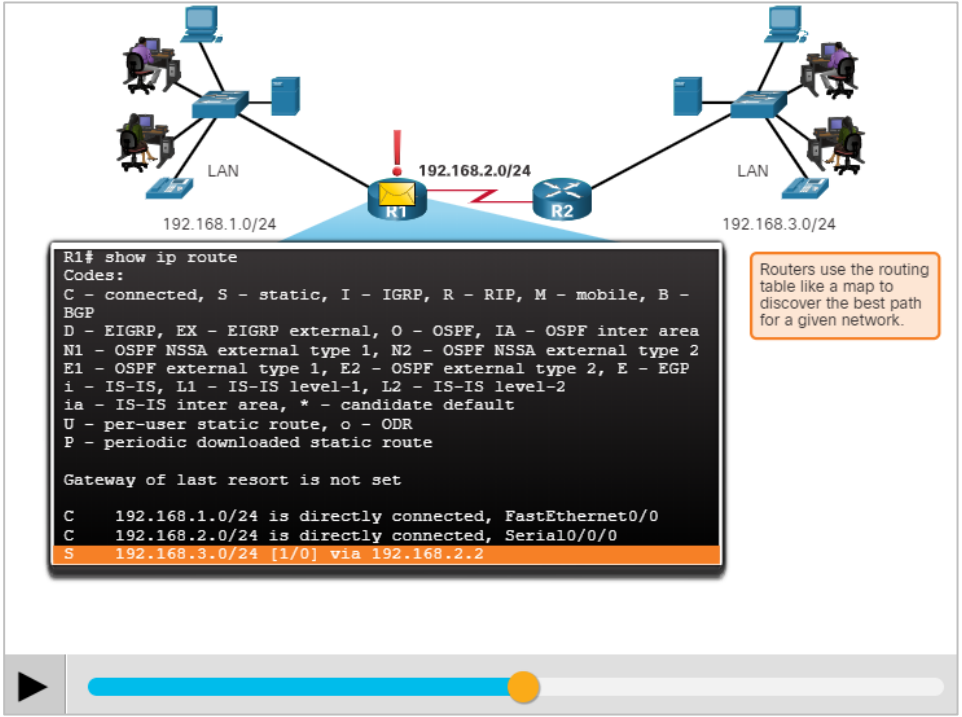

Routers (Contd.)

- After the router has determined the exit interface using path determination, the router must encapsulate the packet into the data link frame of the outgoing interface.
- When a packet received from one network and destined for another network, the router performs the following three major steps:
  1) It de-encapsulates the Layer 2 frame header and trailer to expose the Layer 3 packet.
  2) It examines the destination IP address of the IP packet to find the best path in the routing table.
  3) If the router finds a path to the destination, it encapsulates the Layer 3 packet into a new Layer 2 frame and forwards that frame out the exit interface.
- As a packet travels from the source device to the final destination device, the Layer 3 IP addresses do not change as the Layer 3 PDU does not change. The Layer 2 data link addresses change at every router on the path to the destination, as the packet is de-encapsulated and re-encapsulated in a new Layer 2 frame.

Encapsulating and De-Encapsulating Packets

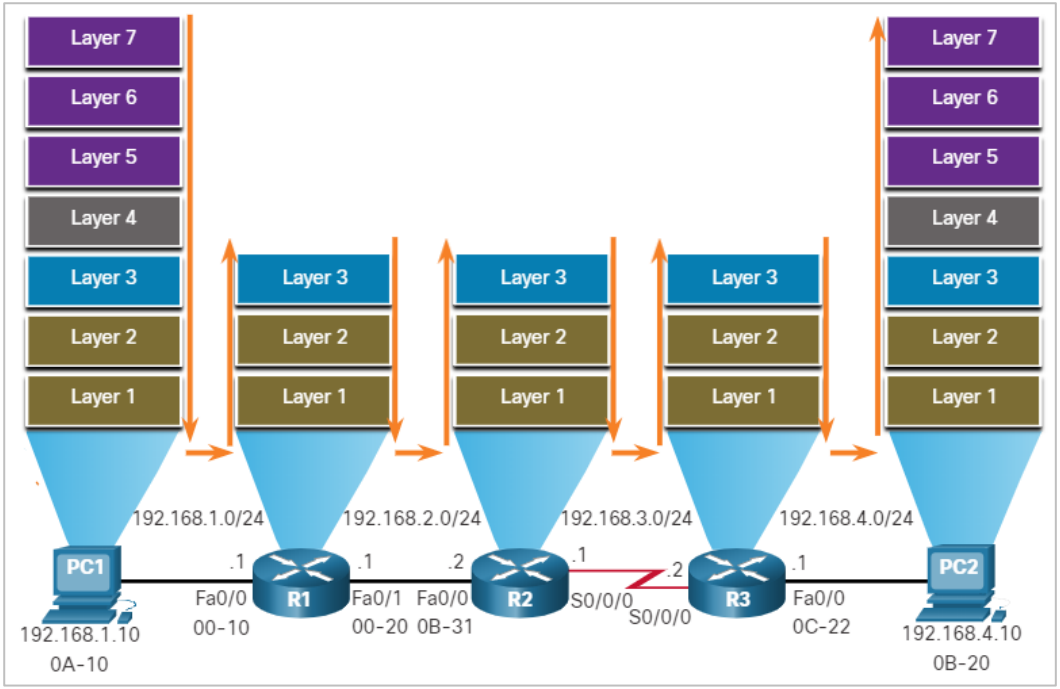

Packet Forwarding Decision Process

- Now, router must determine how to encapsulate the packet and forward it out the correct egress interface.
- The following steps describe the packet forwarding process shown in the figure:
  1) The data link frame with an encapsulated IP packet arrives on the ingress interface.
  2) The router examines the destination IP address in the packet header and consults its IP routing table.
  3) The router finds the longest matching prefix in the routing table.
  4) The router encapsulates the packet in a data link frame and forwards it out the egress interface.
  5) The destination could be a device connected to the network or a next-hop router.
  6) If there is no matching route entry the packet is dropped.

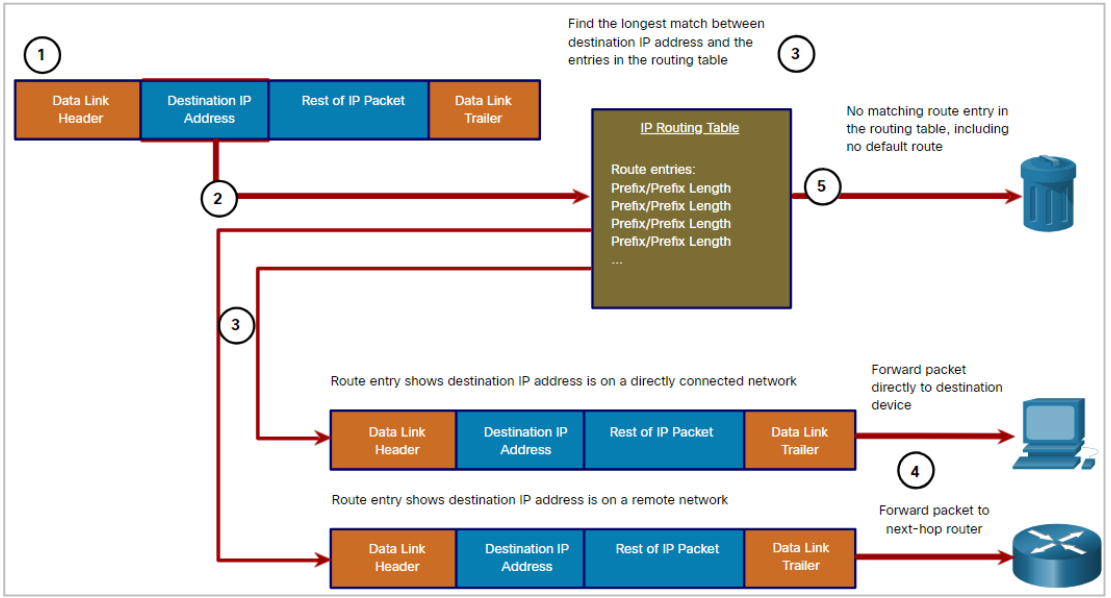

Packet Forwarding Decision Process (Contd.)

The three actions a router can perform with a packet, after the best path is determined:

1) Forwards the Packet to a Device on a Directly Connected Network
2) Forwards the Packet to a Next-Hop Router
3) Drops the Packet - No Match in Routing Table

1) Forwards the Packet to a Device on a Directly Connected Network

- If the route entry indicates that the egress interface is a directly connected network, this means that the destination IP address of the packet belongs to a device on the directly connected network.
- The packet can be forwarded directly to the destination device, an end device on an Ethernet LAN, which means the packet must be encapsulated in an Ethernet frame.
- To encapsulate the packet in the Ethernet frame, the router needs to determine the destination MAC address associated with the destination IP address of the packet.

Packet Forwarding Decision Process (Contd.)

- The process varies based on whether the packet is an IPv4 or IPv6 packet.

- **IPv4 packet:**
  - The router checks its ARP table for the destination IPv4 address and an associated Ethernet MAC address.
  - If there is no match, the router sends an ARP Request and the destination device will return an ARP Reply with its MAC address.
  - The router can now forward the IPv4 packet in an Ethernet frame with the proper destination MAC address.
- **IPv6 packet:**
  - The router checks its neighbor cache for the destination IPv6 address and an associated Ethernet MAC address.
  - If there is no match, the router sends an ICMPv6 Neighbor Solicitation (NS) message and the destination device will return an ICMPv6 Neighbor Advertisement (NA) message with its MAC address.
  - The router can now forward the IPv6 packet in an Ethernet frame with the proper destination MAC addresss

Packet Forwarding Decision Process (Contd.)

2) Forwards the Packet to a Next-Hop Router

- If the route entry indicates that the destination IP address is on a remote network, this means the destination IP address of the packet belongs to a device on network that is not directly connected.
- Therefore, the packet must be forwarded to another router, specifically a next-hop router. The next-hop address is indicated in the route entry.
- If the forwarding router and the next-hop router are on an Ethernet network, a similar process (ARP and ICMPv6 Neighbor Discovery) will occur for determining the destination MAC address of the packet. The difference is that the router will search for the IP address of the next-hop router in its ARP table or neighbor cache, instead of the destination IP address.

3) Drops the Packet - No Match in Routing Table

- If there is no match between the destination IP address and a prefix in the routing table, and if there is no default route, the packet will be dropped.

Routing Information

- The routing table stores the following information:
  - **Directly connected routes -** These routes come from the active router interfaces. Routers add a directly connected route when an interface is configured with an IP address and is activated.
  - **Remote routes -** These are remote networks connected to other routers.
- A routing table is a data file in RAM that is used to store route information about directly connected and remote networks.
- The routing table contains network or next hop associations. These associations tell a router that a particular destination can be optimally reached by sending the packet to a specific router that represents the next hop on the way to the final destination.
- The next hop association can also be the outgoing or exit interface to the next destination.

Routing Information (Contd.)

Directly Connected and Remote Network Routes

The figure identifies the directly connected networks and remote networks of router R1.

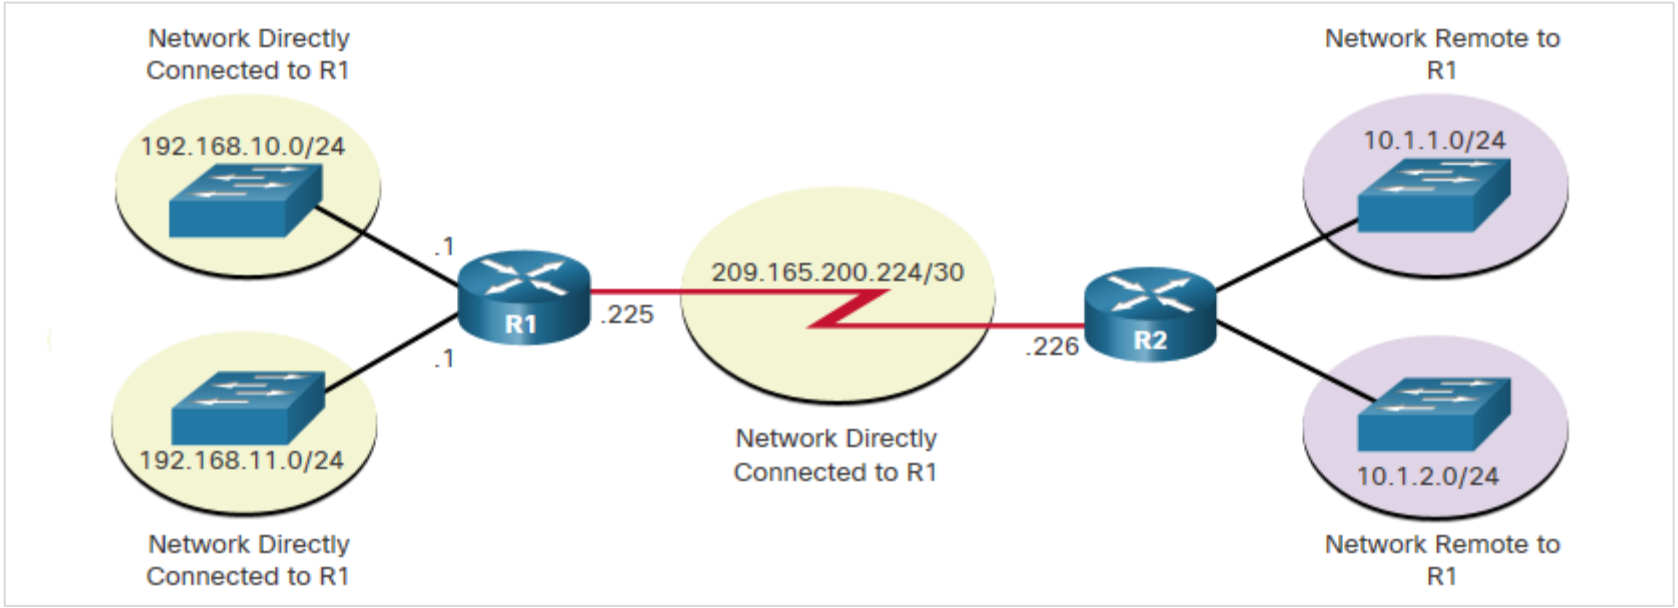

Routing Information (Contd.)

- The destination network entries in the routing table can be added in several ways:
  - **Local Route interfaces –** These are 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 –** These are added to the routing table when an interface is configured and active.
  - **Static routes –** These are added when a route is manually configured and the exit interface is active.
  - **Dynamic routing protocol –** This is added when routing protocols that dynamically learn about the network, such as EIGRP or OSPF, are implemented and networks are identified.
- Dynamic routing protocols exchange network reachability information between routers and dynamically adapt to network changes.

Routing Information (Contd.)

- One of the first dynamic routing protocols was RIP. RIPv1 was released in 1988.
- To address the needs of larger networks, two advanced routing protocols Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) were developed.
- Cisco developed the Interior Gateway Routing Protocol (IGRP) and Enhanced IGRP (EIGRP) which also scales well in larger network implementations.
- The Border Gateway Protocol (BGP) is now used between Internet Service Providers (ISPs) and their larger private clients to exchange routing information.
- The following table classifies the protocols:

Protocol	Interior Gateway Protocols				Exterior Gateway Protocols
	Distance Vector		Link State		Path Vector
IPv4	RIPv2	EIGRP	OSPFv2	IS-IS	BGP-4
IPv6	RIPng	EIGRP for IPv6	OSPFv3	IS-IS for IPv6	BGP-MP

End-to-End Packet Forwarding

- The primary responsibility of the packet forwarding function is to encapsulate packets in the appropriate data link frame type for the outgoing interface.
- The process of packet forwarding is described through the following example.
- **PC1 Sends Packet to PC2**
  - In the first animation, PC1 sends a packet to PC2.
  - Note that if an ARP entry does not exist in the ARP table for the default gateway of 192.168.1.1, PC1 sends an ARP request.
  - Router R1 then return an ARP reply.

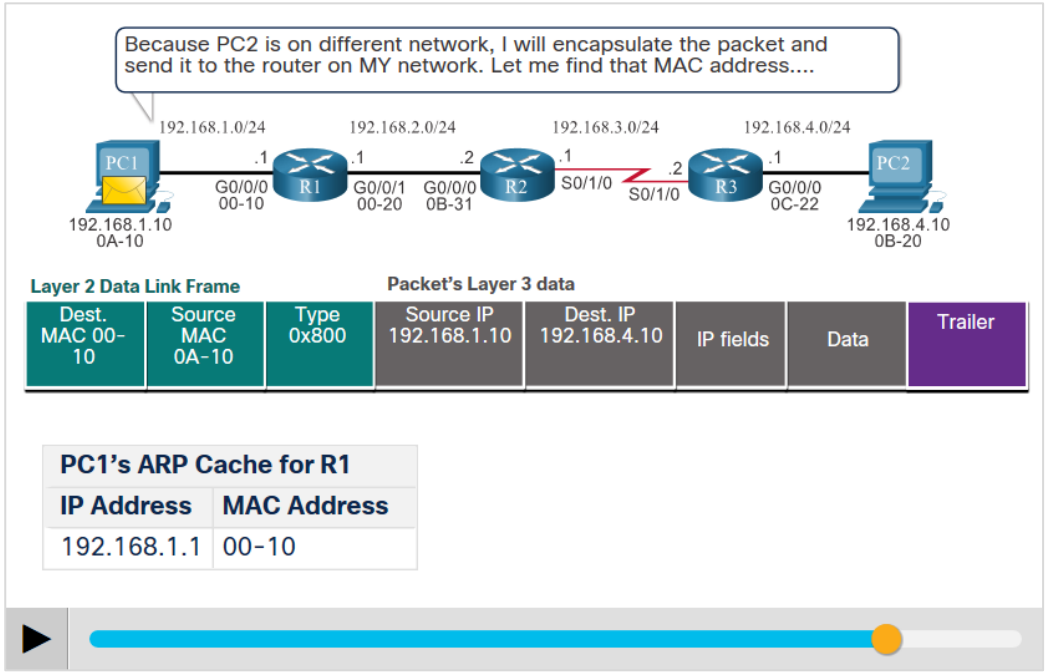

End-to-End Packet Forwarding (Contd.)

- **R1 Forwards the Packet to PC2**
  - Because the exit interface is on an Ethernet network, R1 must resolve the next-hop IPv4 address with a destination MAC address using its ARP table.
  - If an ARP entry does not exist in the ARP table for the next-hop interface of 192.168.2.2, R1 sends an ARP request.
  - R2 would then return an ARP Reply.

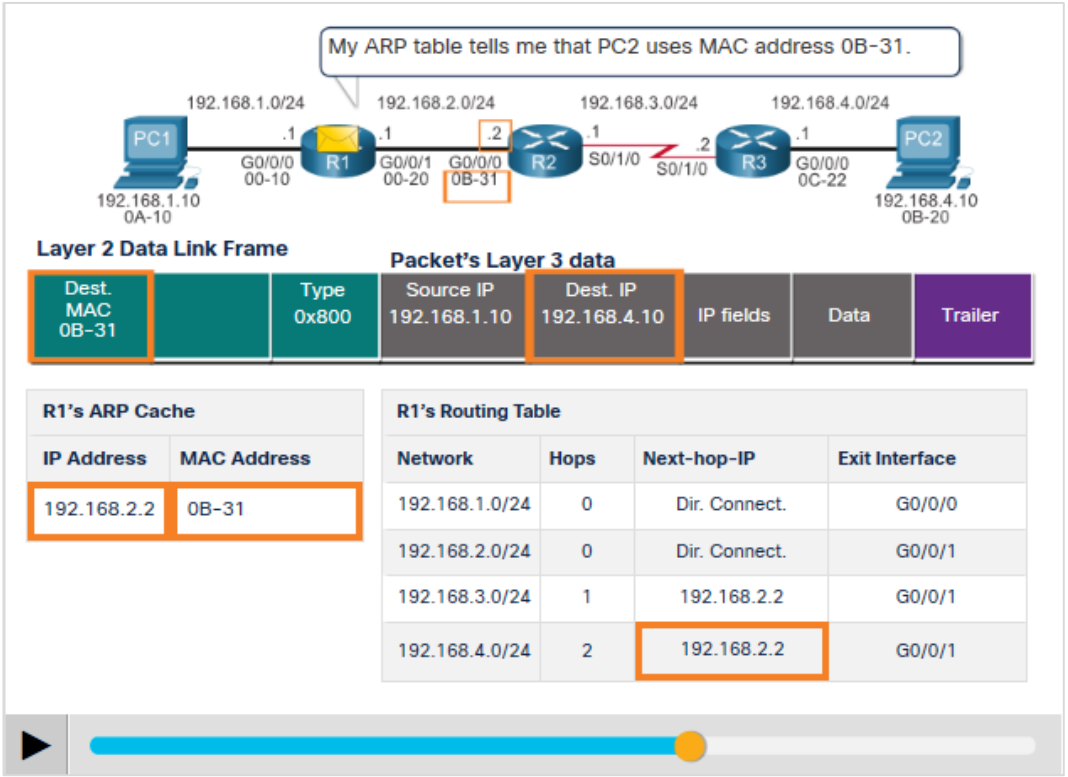

End-to-End Packet Forwarding (Contd.)

R2 Forwards the Packet to R3

- R2 now forwards the packet to R3.
- As the exit interface is not an Ethernet network, R2 does not have to resolve the next-hop IPv4 address with a destination MAC address.
- When the interface is a point-to-point (P2P) serial connection, the router encapsulates the IPv4 packet into the proper data link frame format used by the exit interface.
- As there are no MAC addresses on serial interfaces, R2 sets the data link destination address to an equivalent of a broadcast.

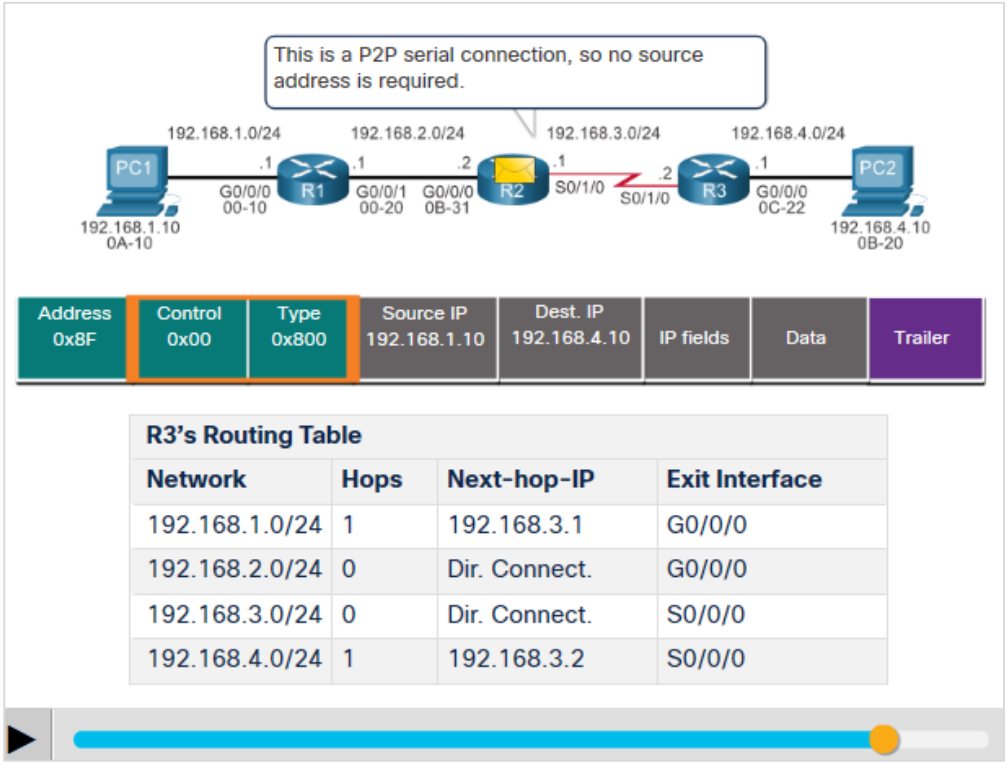

R3's Routing Table

Network	Hops	Next-hop-IP	Exit Interface
192.168.1.0/24	1	192.168.3.1	G0/0/0
192.168.2.0/24	0	Dir. Connect.	G0/0/0
192.168.3.0/24	0	Dir. Connect.	S0/0/0
192.168.4.0/24	1	192.168.3.2	S0/0/0

End-to-End Packet Forwarding (Contd.)

- **R3 Forwards the Packet to PC2**
  - As the destination IPv4 address is on a directly connected Ethernet network, R3 must resolve the destination IPv4 address of the packet with its associated MAC address.
  - If the entry is not in the ARP table, R3 sends an ARP request out of its FastEthernet 0/0 interface.
  - PC2 would then return an ARP reply with its MAC address.

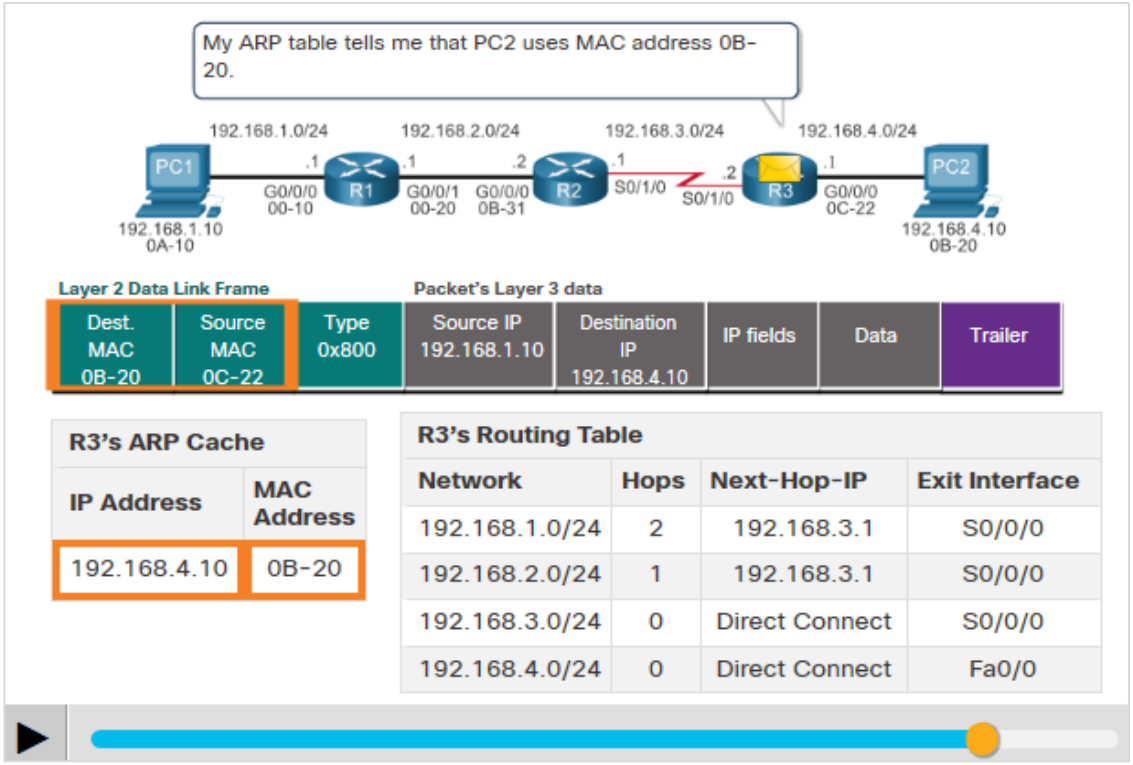

Hubs, Bridges, LAN Switches

- An Ethernet **hub** acts as a multiport repeater that receives an incoming electrical signal (data) on a port. It then immediately forwards a regenerated signal out all other ports. Hubs use physical layer processing to forward data.
- **Bridges** have two interfaces and are connected between hubs to divide the network into multiple collision domains. Each collision domain can have only one sender at a time.
- LAN **switches** are multiport bridges that connect devices into a star topology. Switches also segment a LAN into separate collision domains, one for each switch port. A switch makes forwarding decisions based on Ethernet MAC addresses.

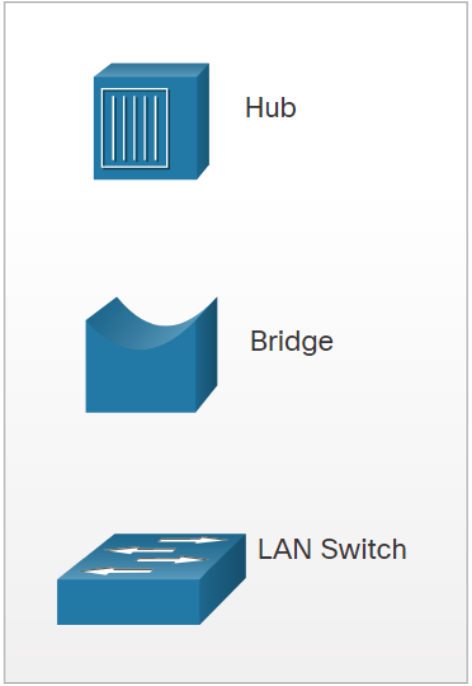

Switching Operation (Contd.)

- The following two-step process is performed on every Ethernet frame that enters a switch.
- **Learn – Examining the Source MAC Address**
  - Every frame that enters a switch is checked for new MAC address information by examining the frame's source MAC address and the port number where the frame entered the switch.
  - If the source MAC address is not in the table, it is added to the MAC address table along with the incoming port number.
  - If the source MAC address does exist in the table, the switch updates the refresh timer for that entry. By default, most Ethernet switches keep an entry in the table for five minutes.
- *Note: If the source MAC address does exist in the table but on a different port, the switch treats this as a new entry. The entry is replaced using the same MAC address, but with the more current port number.*

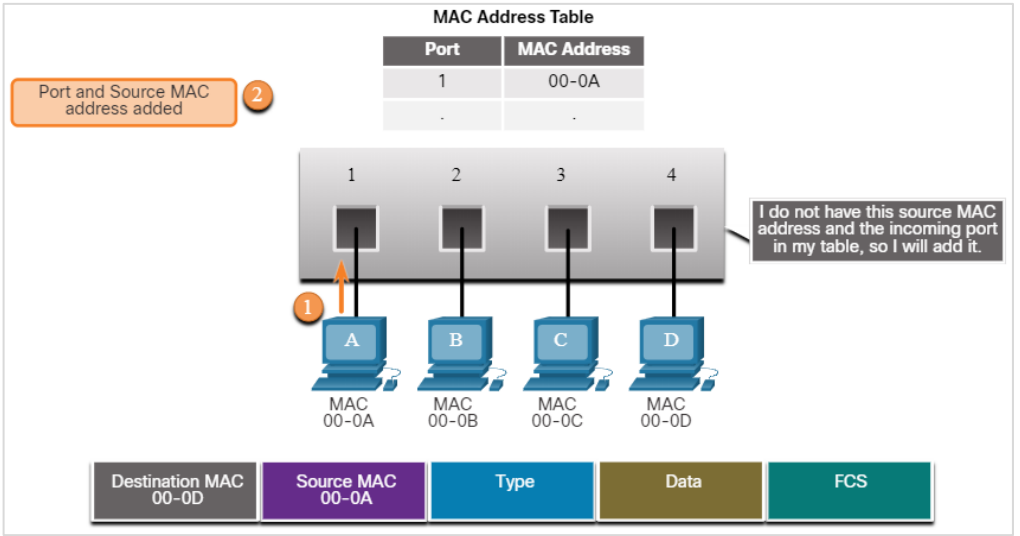

Switching Operation (Contd.)

- **Forward – Examining the Destination MAC Address**
  - If the destination MAC address is a unicast address, the switch will look for a match between the destination MAC address of the frame and an entry in its MAC address table.
  - If the destination MAC address is in the table, it will forward the frame out the specified port.
  - If the destination MAC address is not in the table, the switch will forward the frame out all ports except the incoming port. This is called an unknown unicast.
- *Note: If the destination MAC address is a broadcast or a multicast, the frame is also flooded out all ports except the incoming port.*

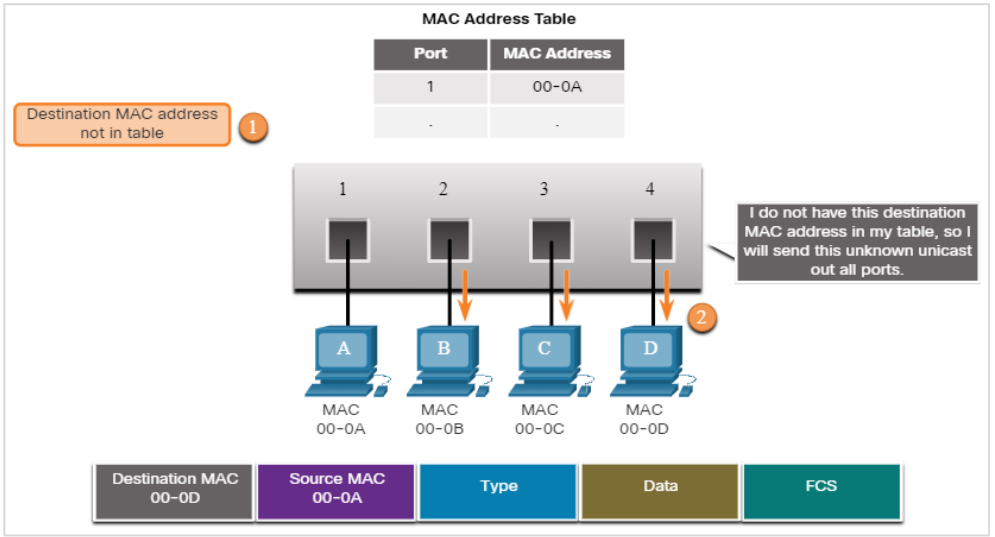

VLANs

- VLANs provide a way to group devices within a LAN.
- **VLAN** provides segmentation and organizational flexibility within a switched internetwork.
- It allows an administrator to segment networks based on factors such as function, project team, or application, without regard for the physical location of the user or device.
- It creates a logical broadcast domain that can span multiple physical LAN segments.
- It prevent users on different VLANs from snooping on each other's traffic.

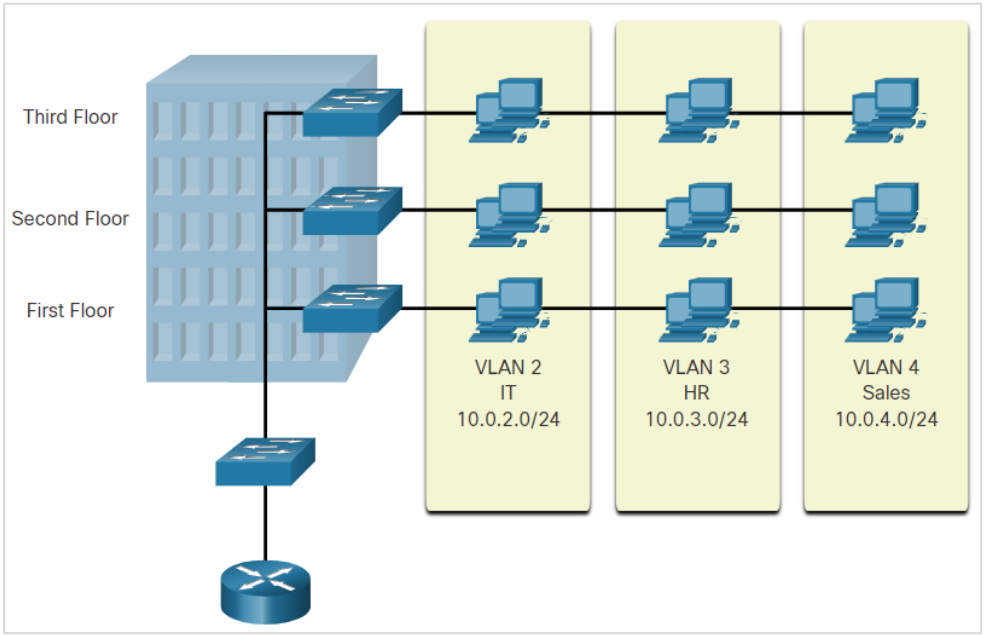

STP

- The **Spanning Tree Protocol** is used to maintain one loop-free path in the Layer 2 network, at any time.
- Loops and duplicate frames have severe consequences for a switched network. STP was developed to address these issues.
- It ensures that there is one logical path between all destinations on the network by blocking redundant paths.
- A port is considered blocked when user data is prevented from entering or leaving that port. This does not include bridge protocol data unit (BPDU) frames that are used by STP to prevent loops.
- If the path is ever needed to compensate for a network cable or switch failure, STP recalculates the paths and unblocks the necessary ports to allow the redundant path to become active.

Multilayer Switching

- Multilayer switches (**Layer 3 switches**) perform Layer 2 switching and also forward frames based on Layer 3 and 4 information.
- All Cisco Catalyst multilayer switches support the following types of Layer 3 interfaces:
  - **Routed port -** A pure Layer 3 interface similar to a physical interface on a Cisco IOS router.
  - **Switch virtual interface (SVI) -** A virtual VLAN interface for inter-VLAN routing. In other words, SVIs are the virtual-routed VLAN interfaces.

Multilayer Switching (Contd.)

- **Routed Ports**
  - A routed port is a physical port that acts similarly to an interface on a router.
  - Unlike an access port, a routed port is not associated with a particular VLAN. It behaves like a regular router interface.
  - Also, as Layer 2 functionality has been removed, Layer 2 protocols, such as STP, do not function on a routed interface.
  - Some protocols, such as LACP and EtherChannel, do function at Layer 3. Unlike Cisco IOS routers, routed ports on a Cisco IOS switch do not support sub interfaces.

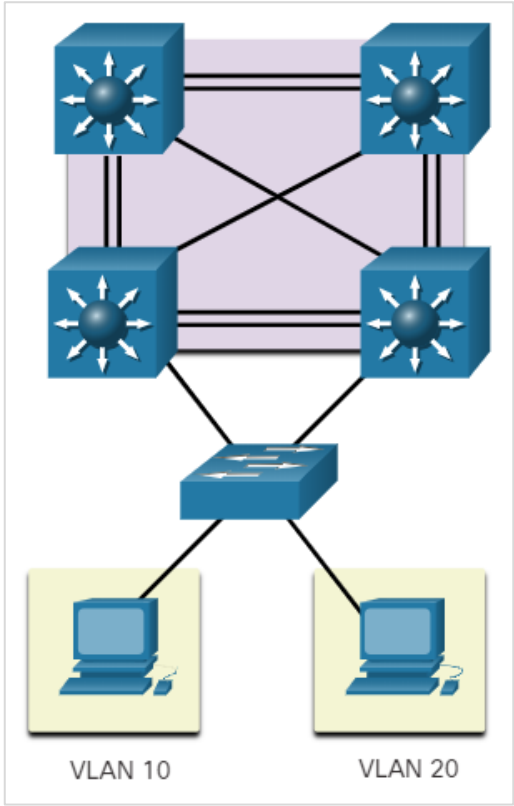

Multilayer Switching (Contd.)

- **Switch Virtual Interface**
  - An SVI is a virtual interface that is configured within a multilayer switch. Unlike the basic Layer 2 switches, a multilayer switch can have multiple SVIs. An SVI can be created for any VLAN that exists on the switch.
  - An SVI is considered to be virtual as there is no physical port dedicated to the interface. It can perform the same functions for the VLAN as a router interface and can be configured in much the same way as a router interface.
  - The SVI for the VLAN provides Layer 3 processing for packets to or from all switch ports associated with that VLAN.

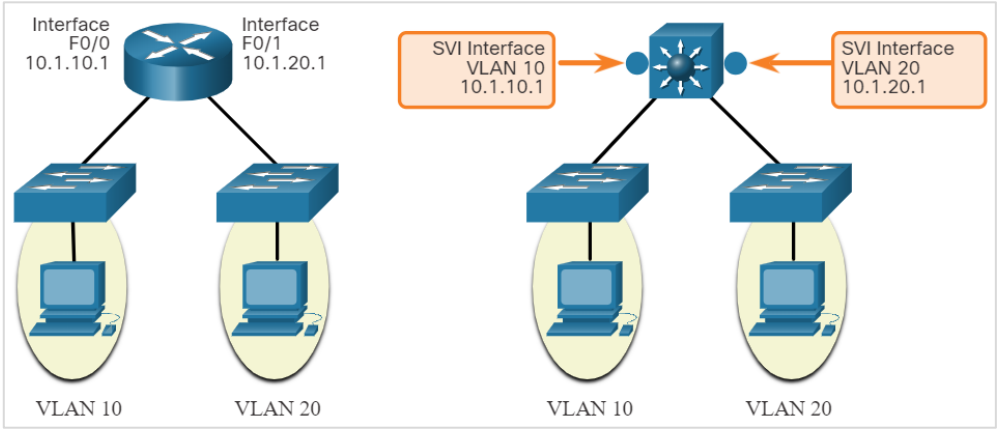

11.2 Wireless Communications

Wireless versus Wired LANs

- **WLANs** use Radio Frequencies (RF) instead of cables at the physical layer and MAC sublayer of the data link layer.
- The IEEE has adopted the 802 LAN/MAN portfolio of computer network architecture standards which includes two dominant working groups **802.3 Ethernet**, which defined Ethernet for wired LANs and **802.11 which defined Ethernet for WLANs**.
- WLANs also differ from wired LANs as follows:
  - WLANs connect clients to the network through a wireless access point (AP) or wireless router, instead of an Ethernet switch.
  - WLANs connect mobile devices that are often battery powered, as opposed to plugged-in LAN devices. Wireless NICs tend to reduce the battery life of a mobile device.
  - WLANs use a different frame format than wired Ethernet LANs. WLANs require additional information in the Layer 2 header of the frame.
  - WLANs raise more privacy issues because radio frequencies can reach outside the facility.

Wireless versus Wired LANs (Contd.)

The difference between WLAN and Wired LAN is summarized in the following table.

Characteristic	802.11 Wireless LAN	802.3 Wired Ethernet LANs
Physical Layer	Radio frequency (RF)	Physical cables
Media Access	Collision avoidance	Collision detection
Availability	Anyone with a wireless NIC in range of an access point	Physical cable connection required
Signal Interference	Yes	Minimal
Regulation	Different regulations by country	IEEE standard dictates

802.11 Frame Structure

- All 802.11 wireless frames contain the following fields:
  - **Frame Control** – This identifies the type of wireless frame and contains subfields for protocol version, frame type, address type, power management, and security settings.
  - **Duration** – This is used to indicate the remaining duration needed to receive the next frame transmission.
  - **Address1** – This contains the MAC address of the receiving wireless device or AP.
  - **Address2** – This contains the MAC address of the transmitting wireless device or AP.

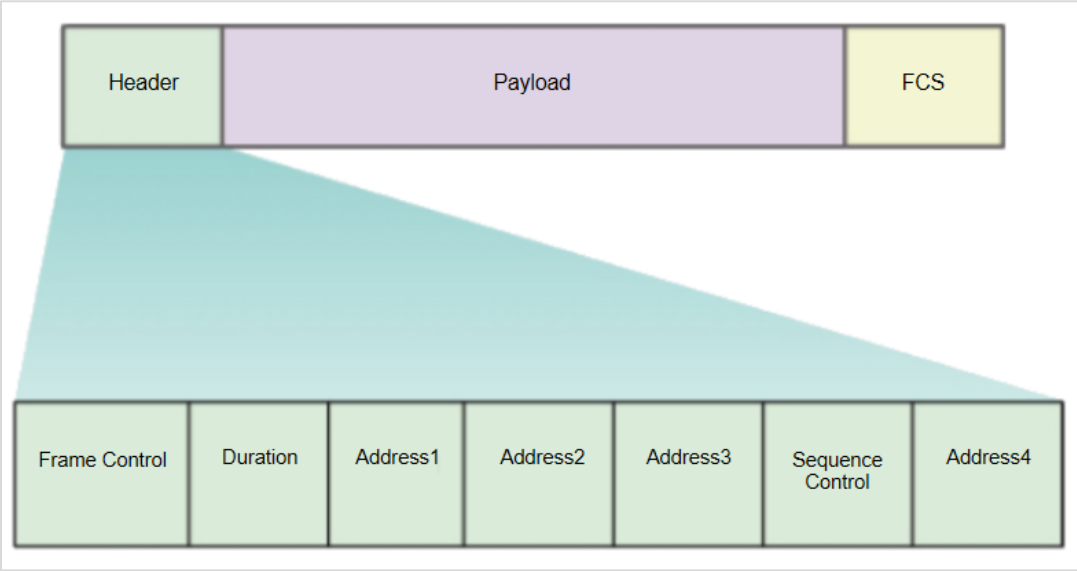

802.11 Frame Structure (Contd.)

- **Address3** - This contains the MAC address of the destination, such as the router interface with AP attached.
- **Sequence Control** – This contains information to control sequencing and fragmented frames.
- **Address4** - This is usually missing as it is used only in ad hoc mode.
- **Payload** – This contains the data for transmission.
- **FCS** – This is used for Layer 2 error control.

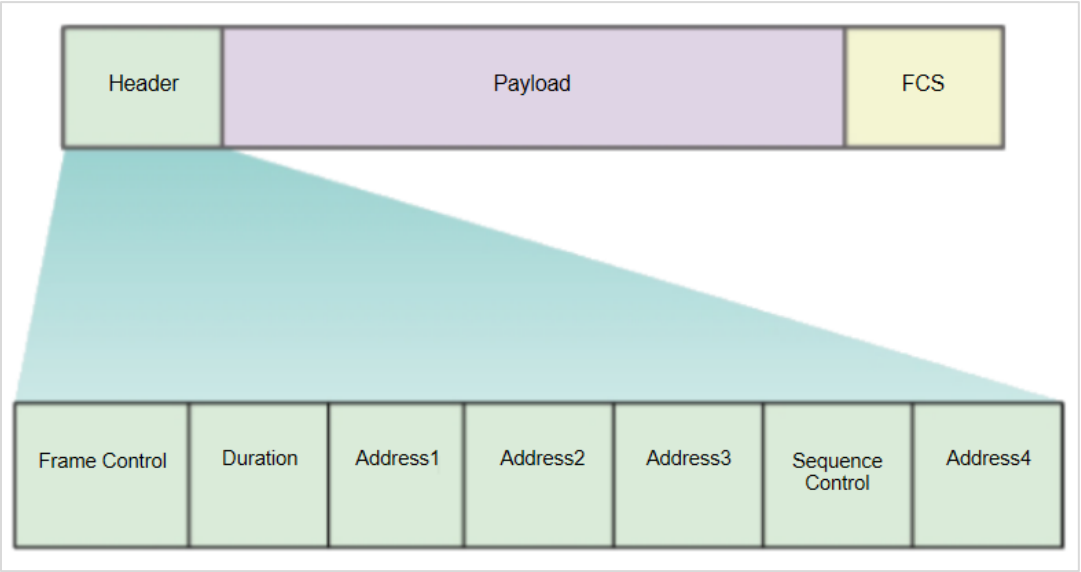

CSMA/CA

- WLANs are half-duplex, shared media configurations.
  - Half-duplex means that only one client can transmit or receive at any given moment.
  - Shared media means that wireless clients can all transmit and receive on the same radio channel.
- This creates a problem because a wireless client cannot hear while it is sending, which makes it impossible to detect a collision.
- To resolve this problem, WLANs use **Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA)** method to determine how and when to send data on the network.
- Wireless client does the following:

1) Listens to the channel to see if it is idle. The channel is also called the carrier.
2) Sends a Ready To Send (RTS) message to the AP to request dedicated access to the network.
3) Receives a Clear To Send (CTS) message from the AP granting access to send.
4) If the wireless client does not receive a CTS message, it waits a random amount of time before restarting the process.
5) After it receives the CTS, it transmits the data.
6) All transmissions are acknowledged.

CSMA/CA (Contd.)

A wireless client does the following:

- Listens to the channel to see if it is idle. The channel is also called the carrier.
- Sends a Ready To Send (RTS) message to the AP to request dedicated access to the network.
- Receives a Clear To Send (CTS) message from the AP granting access to send.
- If the wireless client does not receive a CTS message, it waits a random amount of time before restarting the process.
- After it receives the CTS, it transmits the data.
- All transmissions are acknowledged.

Wireless Client and AP Association

- For wireless devices to communicate over a network, they must first associate with an AP or wireless router.
- An important part of the 802.11 process is discovering a WLAN and subsequently connecting to it.
- Wireless devices complete the following three stage process, as shown in the figure:
  - Discover a wireless AP
  - Authenticate with AP
  - Associate with AP

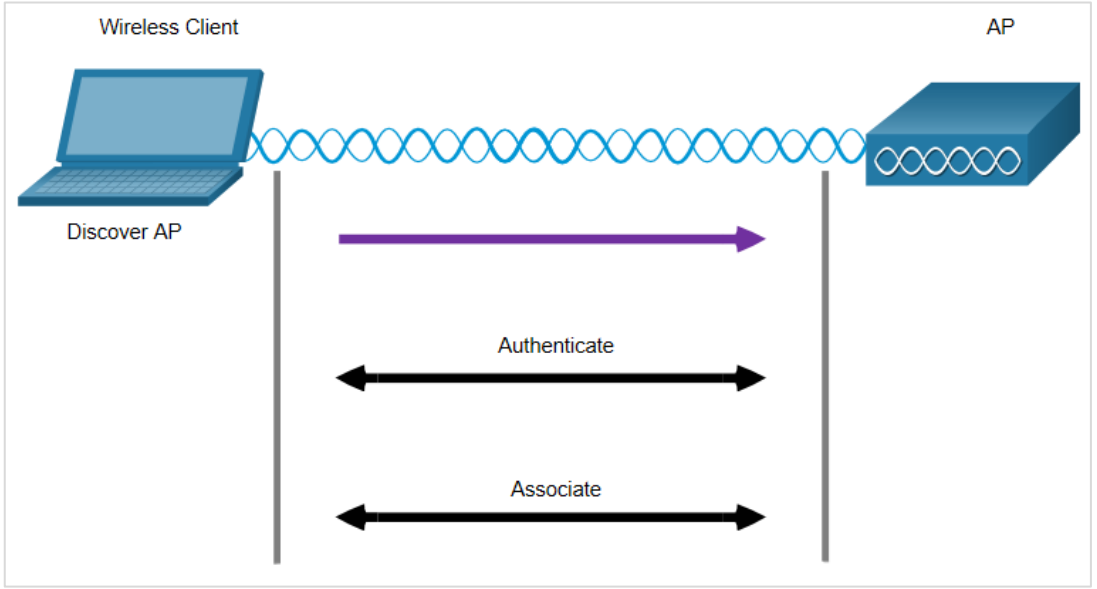

Wireless Client and AP Association (Contd.)

- In order to have a successful association, a wireless client and an AP must agree on specific parameters.
- Parameters must then be configured on the AP and subsequently on the client. The configurable wireless parameters include:
  - **SSID** -The SSID name appears in the list of available wireless networks on a client.
  - **Password –** This is required from the wireless client to authenticate to the AP.
  - **Network mode** - This refers to the 802.11a/b/g/n/ac/ad WLAN standards.
  - **Security mode** - This refers to the security parameter settings, such as WEP, WPA, WPA2, WPA3. Always enable the highest security level supported.
  - **Channel settings** - This refers to the frequency bands used to transmit wireless data.

Passive and Active Discover Mode

- Wireless devices must discover and connect to an AP or wireless router.
- Wireless clients connect to the AP using a scanning (probing) process such as passive and active.
- **Passive Mode**
  - In this mode, the AP openly advertises its service by periodically sending broadcast **beacon frames** containing the SSID, supported standards, and security settings.
  - The primary purpose of the beacon is to allow wireless clients to learn which networks and APs are available in a given area.
  - This allows the wireless clients to choose which network and AP to use.

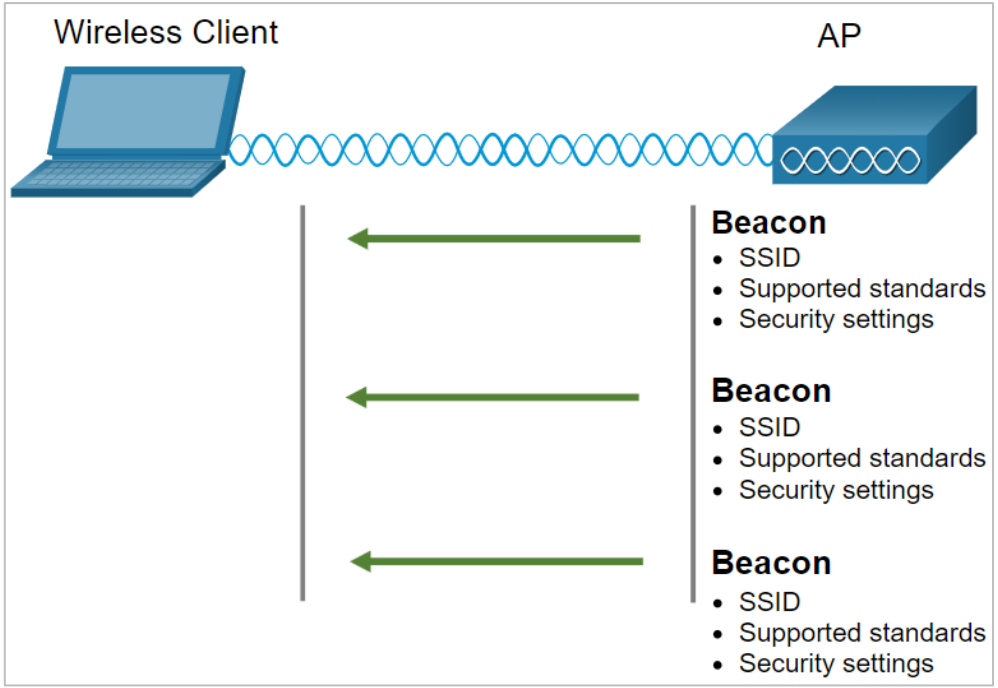

Passive and Active Discover Mode (Contd.)

- **Active Mode**
  - In this mode, wireless clients must know the name of the SSID.
  - The wireless client initiates the process by broadcasting a **probe request** frame on multiple channels includes the SSID name and standards supported.
  - APs configured with the SSID will send a **probe response** that includes the SSID, supported standards, and security settings.
  - Active mode may be required if an AP is configured to not broadcast beacon frames.
  - A wireless client could also send a probe request without a SSID name to discover nearby WLAN networks. APs configured to broadcast beacon frames would respond to the wireless client with a probe response and provide the SSID name.

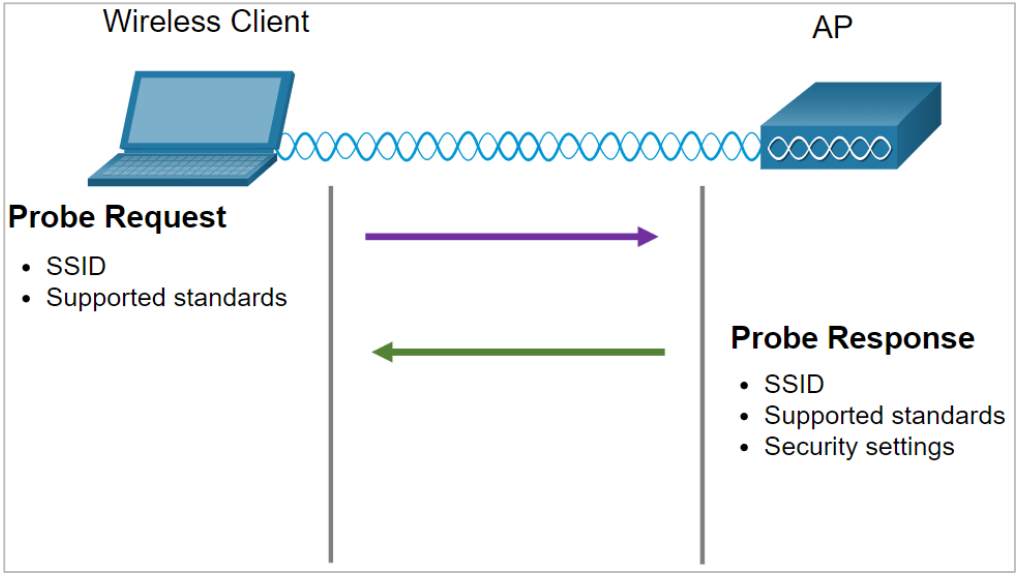

AP, LWAP, and WLC

- All of the control and management functions of the APs on a network can be centralized into a **Wireless LAN Controller (WLC)**
  - All management functions, such as defining SSIDs and authentication are conducted on the centralized WLC rather than on each individual AP.
  - A major benefit of centralizing the AP management functions in the WLC is simplified configuration and monitoring of numerous access points, among many other benefits.
- When using a WLC, the APs no longer act autonomously, but instead act as **lightweight APs (LWAPs)**
  - LWAPs only forward data between the wireless LAN and the WLC.

Thank you! Questions?

Vladimír Veselý

updated: 2024-02-24

https://www.fit.vutbr.cz/research/groups/nes@fit

Module 12: Network Security Infrastructure

Instructor Materials

CyberOps Associate v1.0

Module Objectives

Module Title: Network Security Infrastructure

Module Objective: Explain how devices and services are used to enhance network security.

Topic Title	Topic Objective
Network Topologies	Explain how network designs influence the flow of traffic through the network.
Security Devices	Explain how specialized devices are used to enhance network security.
Security Services	Explain how network services enhance network security.

12.1 Network Topologies

Network Representations

- Network diagrams, often called topology diagrams, use symbols to represent different devices and connections within the network.
- The important terminologies to be known include:
  - **Network Interface Card (NIC)**
  - **Physical Port**
  - **Interface**

***Note**: The terms port and interface are often used interchangeably.*

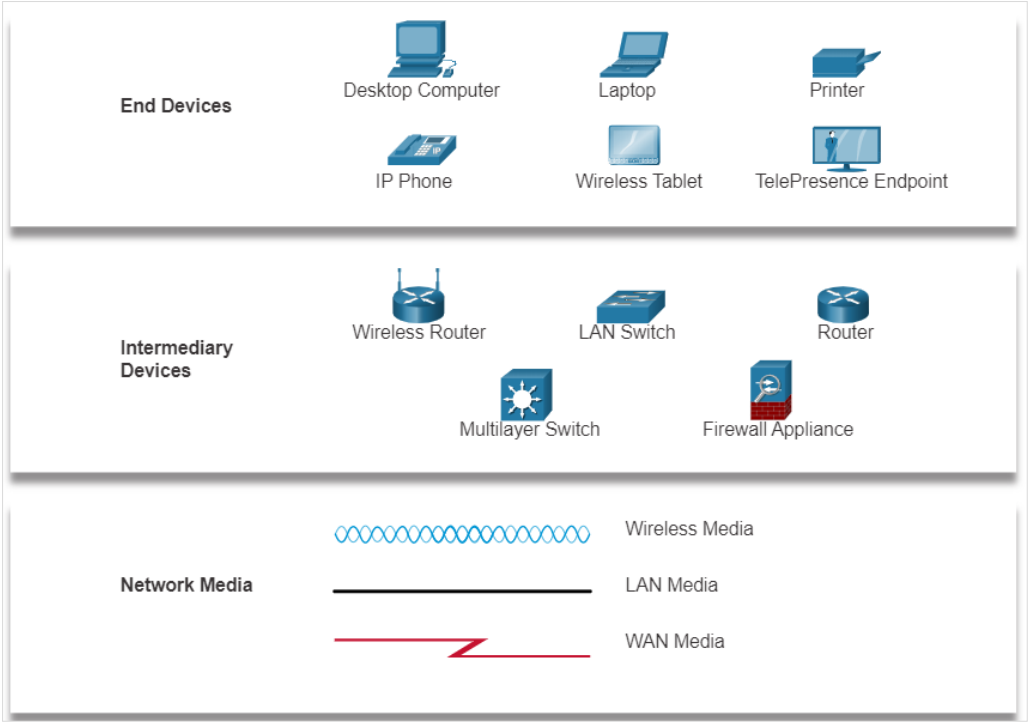

Topology Diagrams

Physical topology diagrams illustrate the physical location of intermediary devices and cable installation.

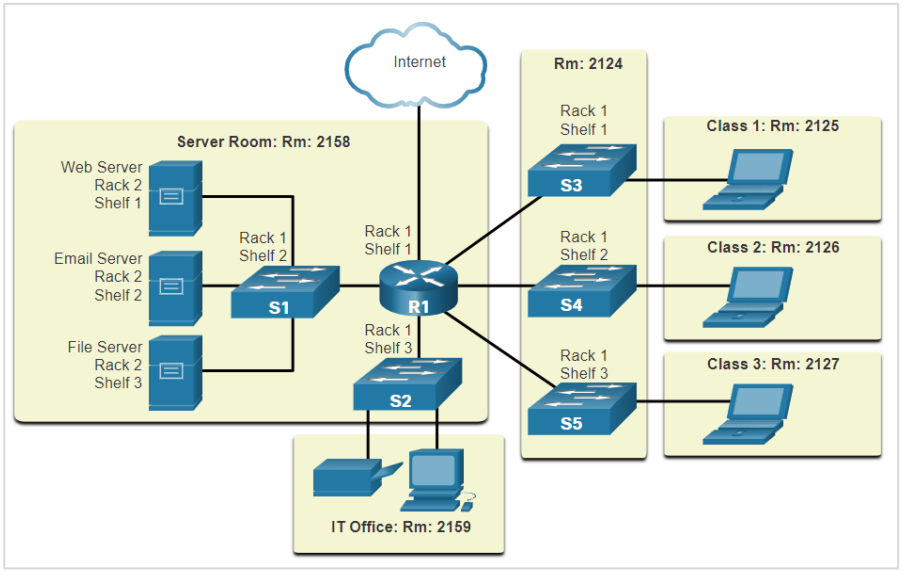

Logical topology diagrams illustrate devices, ports, and the addressing scheme of the network.

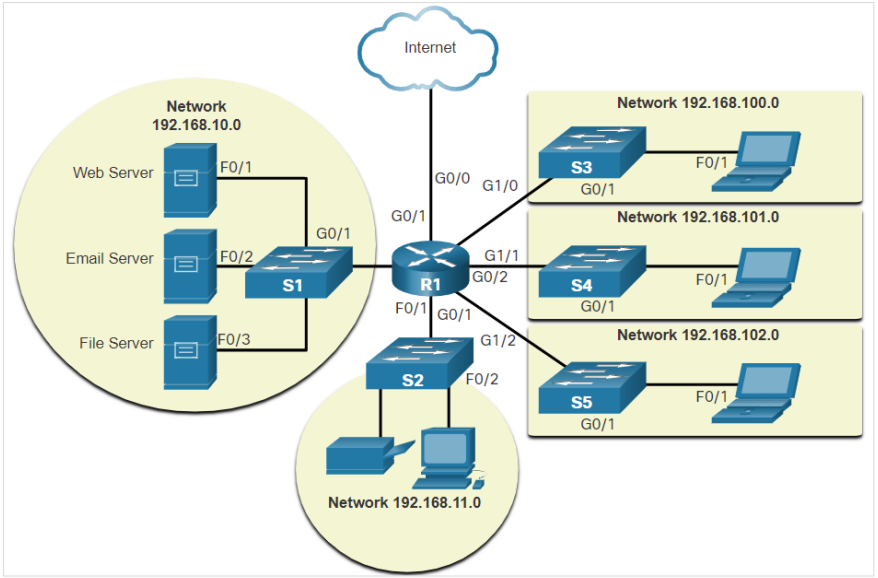

LANs and WANs

- Network infrastructures vary greatly in terms of:
  - Size of the area covered
  - Number of users connected
  - Number and types of services available
  - Area of responsibility
- The two most common types of network infrastructures are
  - **Local Area Networks (LANs)**
  - **Wide Area Networks (WANs)**

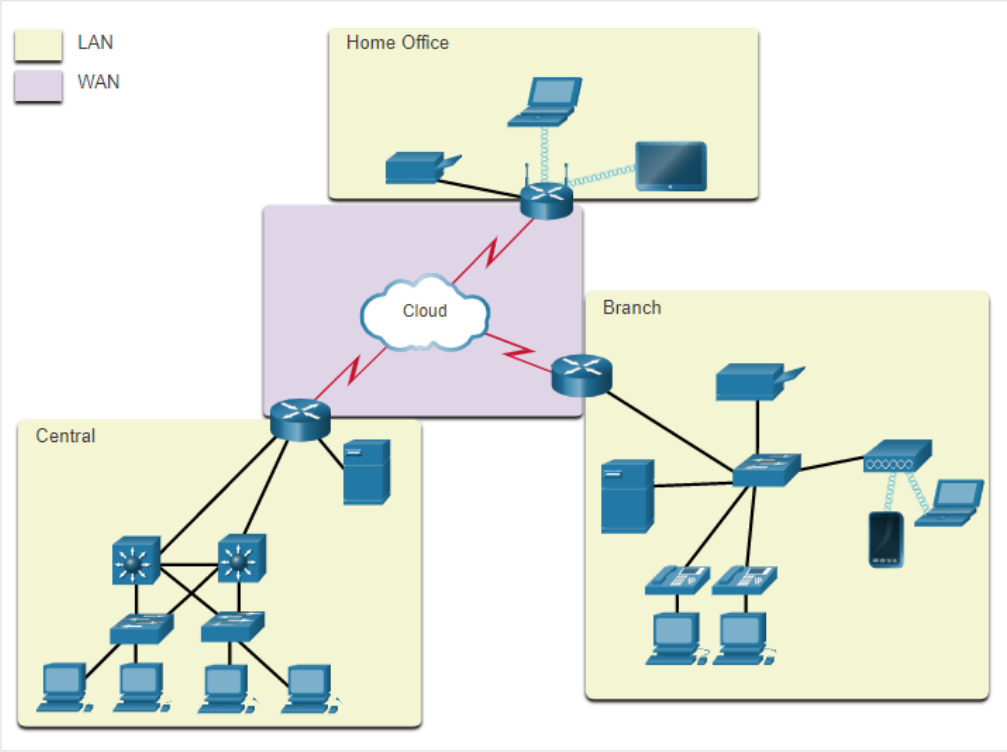

LANs connected to a WAN

LANs and WANs (Contd.)

A **LAN** is a network infrastructure that spans a small geographical area.

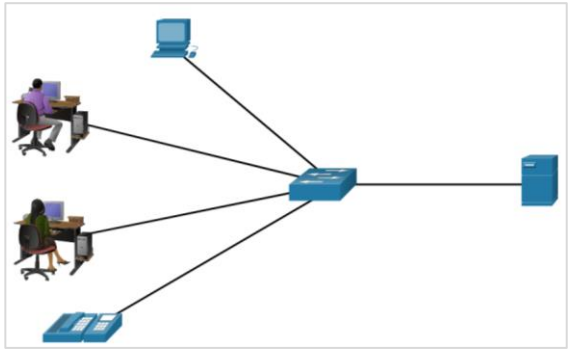

A **WAN** is a network infrastructure that spans a wide geographical area.

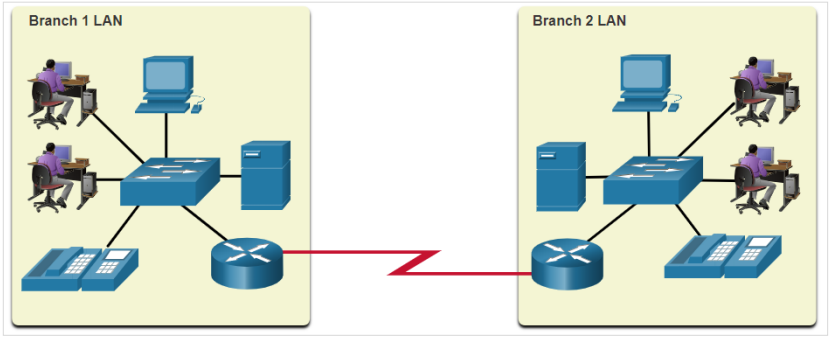

LAN	WAN
Interconnect end devices in a limited area.	Interconnect LANs over wide geographical areas.
Administered by a single organization or individual.	Typically administered by multiple service providers.
Provide high-speed bandwidth to internal end devices and intermediary devices.	Typically provide slower speed links between LANs.

The Three-Layer Network Design Model

- The campus wired LAN uses a hierarchical design model to separate the network topology into modular groups or layers.
- The hierarchical LAN design includes three layers:
  - **Access** - Provides endpoints and users direct access to the network.
  - **Distribution** - Aggregates access layers and provides connectivity to services.
  - **Core** - Provides connectivity between distribution layers for large LAN environments.

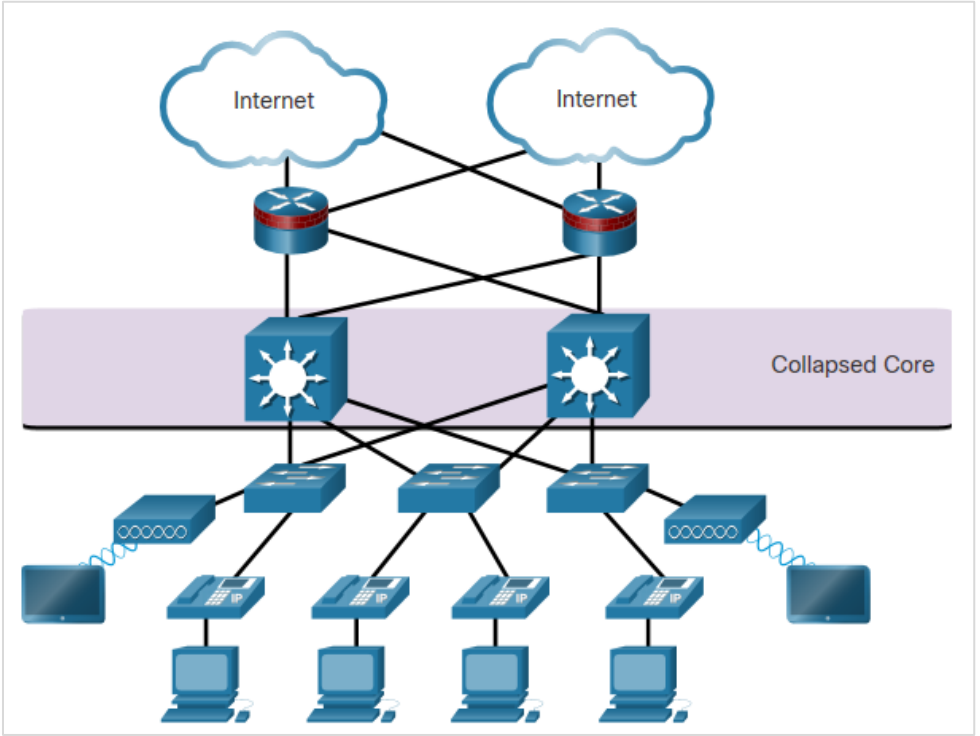

Hierarchical Design Model

Common Security Architectures

- Firewall design is primarily about device interfaces permitting or denying traffic based on the source, the destination, and the type of traffic. The three firewall designs are:
- **Public and Private**
  - The public network (or outside network) is untrusted, and the private network (or inside network) is trusted.

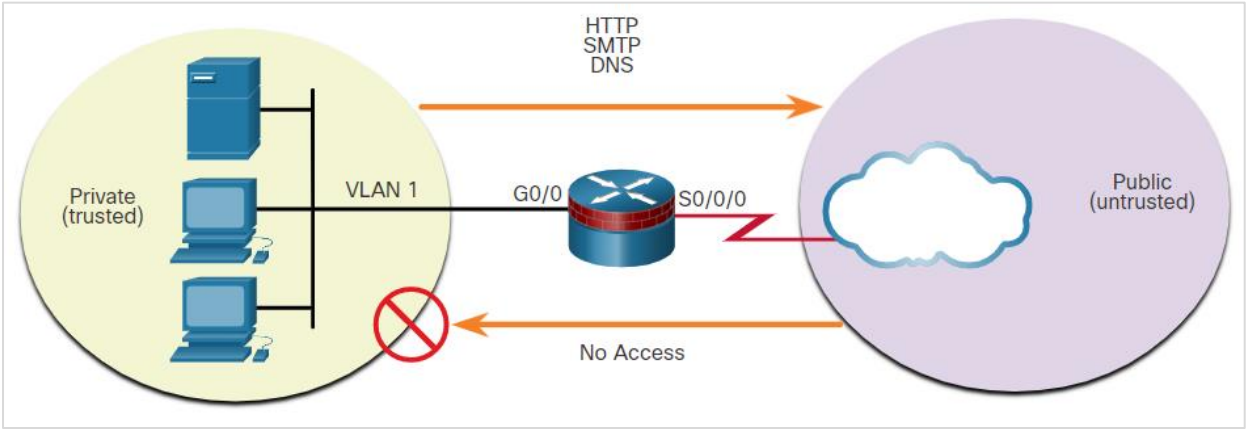

Common Security Architectures (Contd.)

- **Demilitarized Zone (DMZ)**
  - Inside interface connected to the private network
  - Outside interface connected to the public network
  - DMZ interface

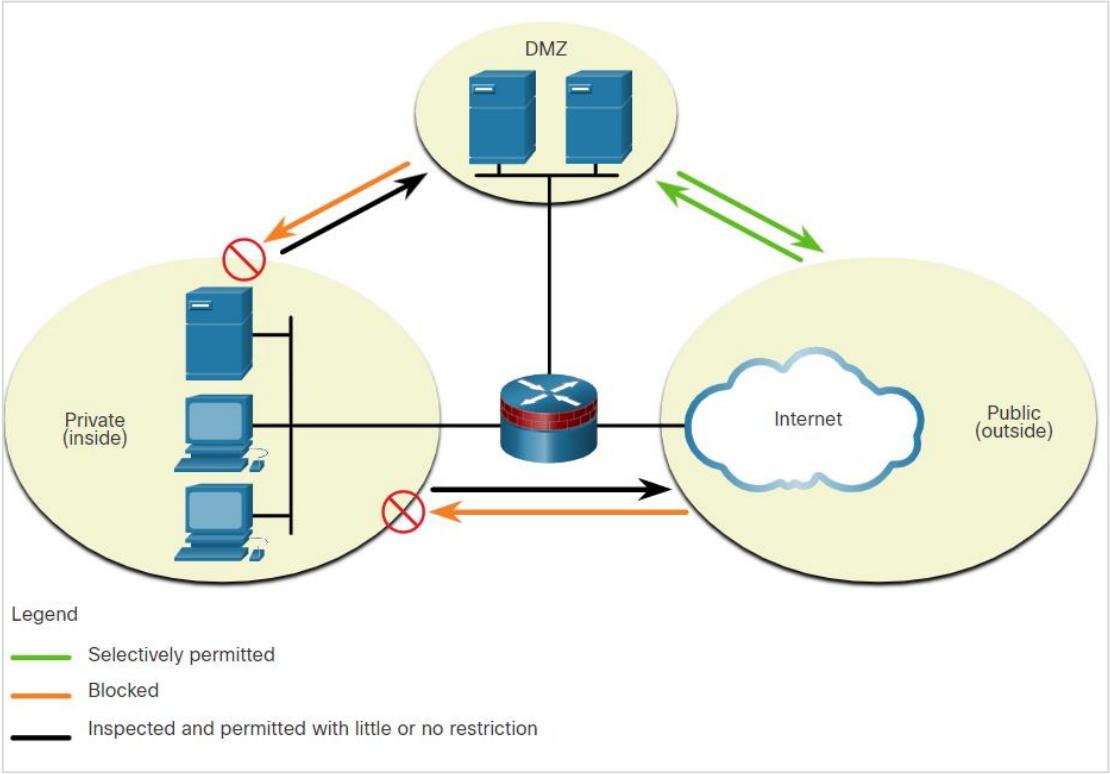

Common Security Architectures (Contd.)

- **Zone-based Policy Firewalls (ZPFs)**
  - ZPFs use the concept of zones to provide additional flexibility.
  - A zone is a group of one or more interfaces that have similar functions or features.
  - Zones help to specify where a Cisco IOS firewall rule or policy should be applied.

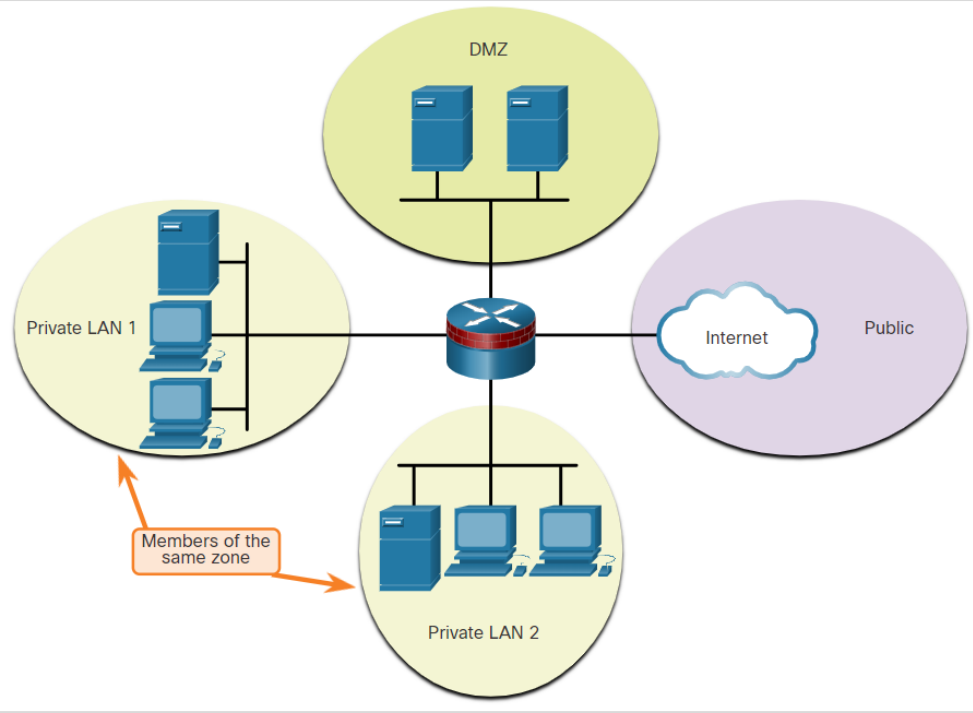

12.2 Security Devices

Firewalls

- A **firewall** is a system, or group of systems, that enforces an access control policy between networks
- Common firewall properties:
  - Resistant to network attacks
  - The only transit point between internal corporate networks and external networks because all traffic flows through the firewall
  - Enforce the access control policy

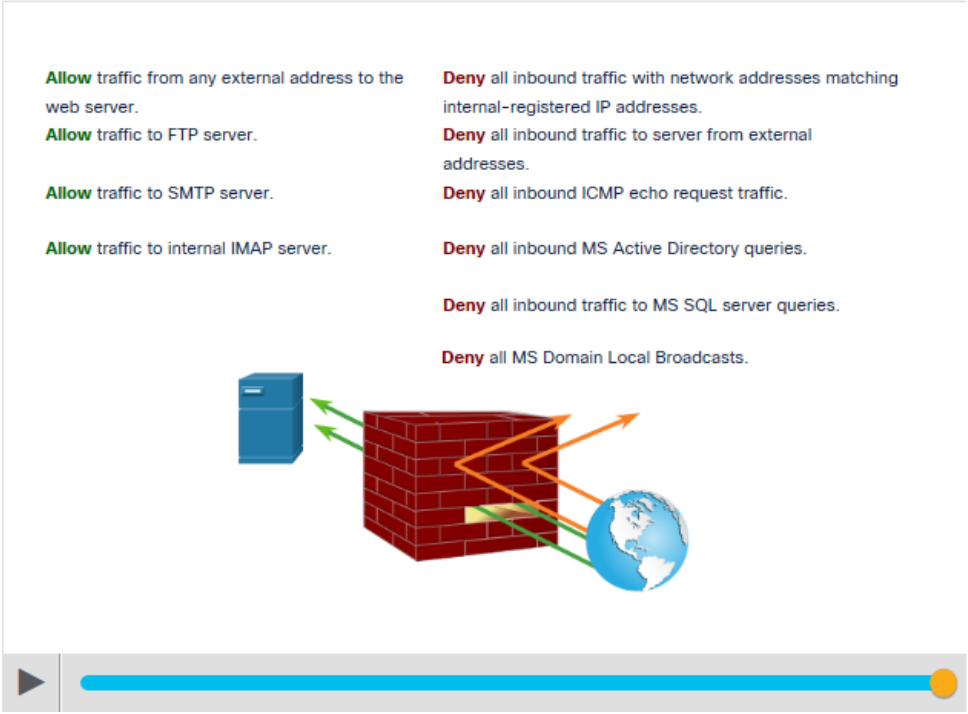

Firewalls (Contd.)

Following are the benefits and limitations of firewalls:

Firewall Benefits	Firewall Limitations
Prevent the exposure of sensitive hosts, resources, and applications to untrusted users.	A misconfigured firewall can have serious consequences for the network, such as becoming a single point of failure.
Sanitize protocol flow, which prevents the exploitation of protocol flaws.	The data from many applications cannot be passed over firewalls securely.
Block malicious data from servers and clients.	Users might proactively search for ways around the firewall to receive blocked material, which exposes the network to potential attack.
Reduce security management complexity.	Network performance can slow down.
	Unauthorized traffic can be tunnelled or hidden as legitimate traffic through the firewall.

Firewall Type Descriptions

- **Packet Filtering (Stateless) Firewall**
  - Packet Filtering firewalls are part of a router firewall, which permits or denies traffic based on Layer 3 and Layer 4 information.
  - They are stateless firewalls that use a simple policy table look-up that filters traffic based on specific criteria.

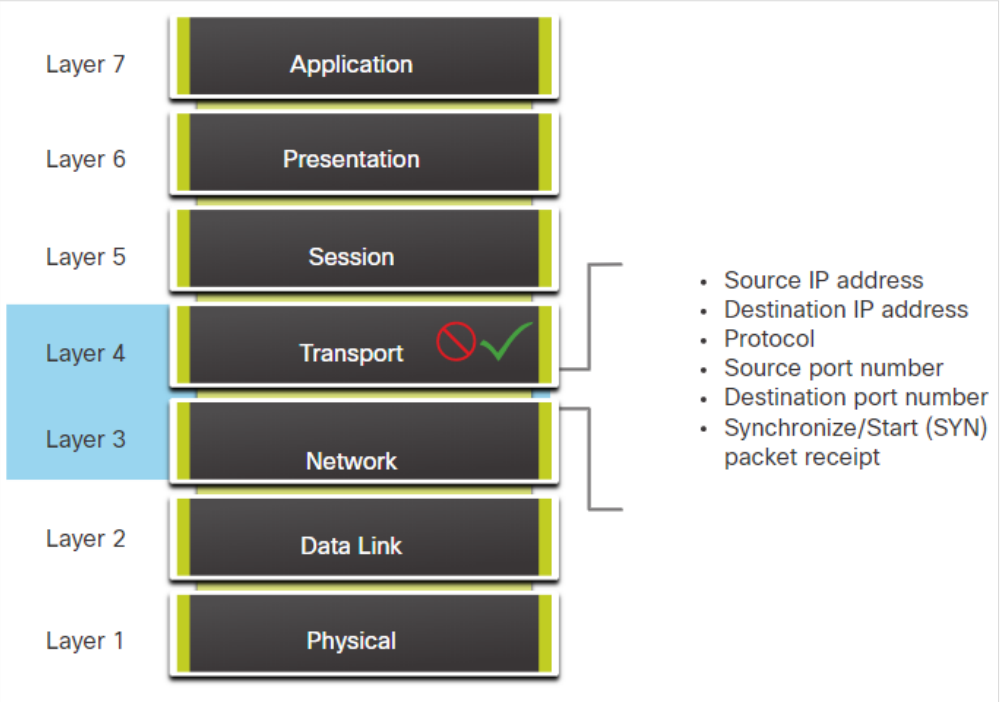

Firewall Type Descriptions (Contd.)

- **Stateful Firewalls**
  - Stateful firewalls are the most versatile and the most common firewall technologies in use.
  - These firewalls provide stateful packet filtering by using connection information maintained in a state table.

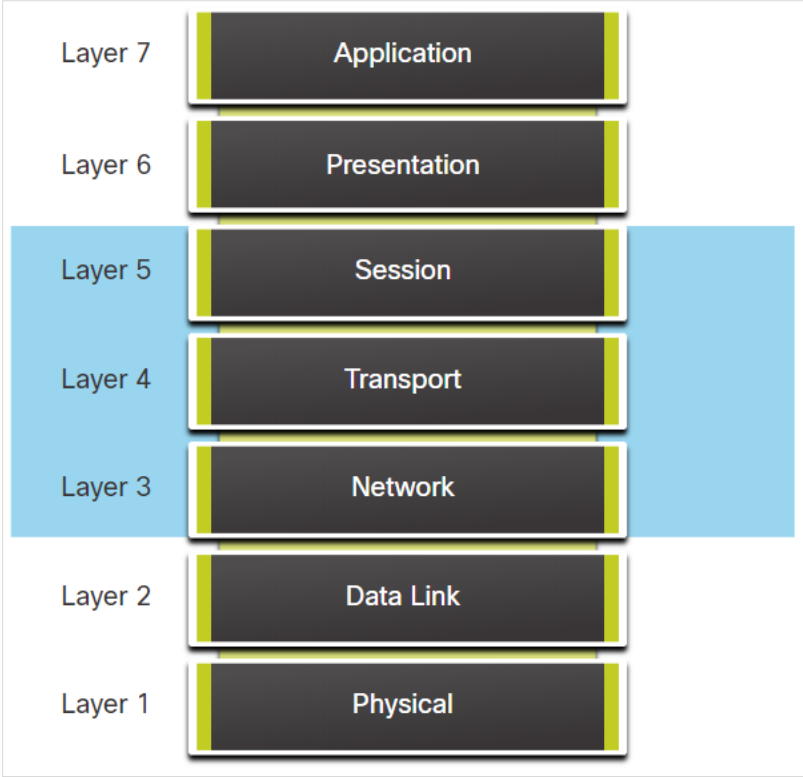

Firewall Type Descriptions (Contd.)

- **Application gateway firewall (proxy firewall)**
  - Application gateway firewall filters information at Layers 3, 4, 5, and 7 of the OSI reference model.
  - Most of the firewall control and filtering is done in the software.

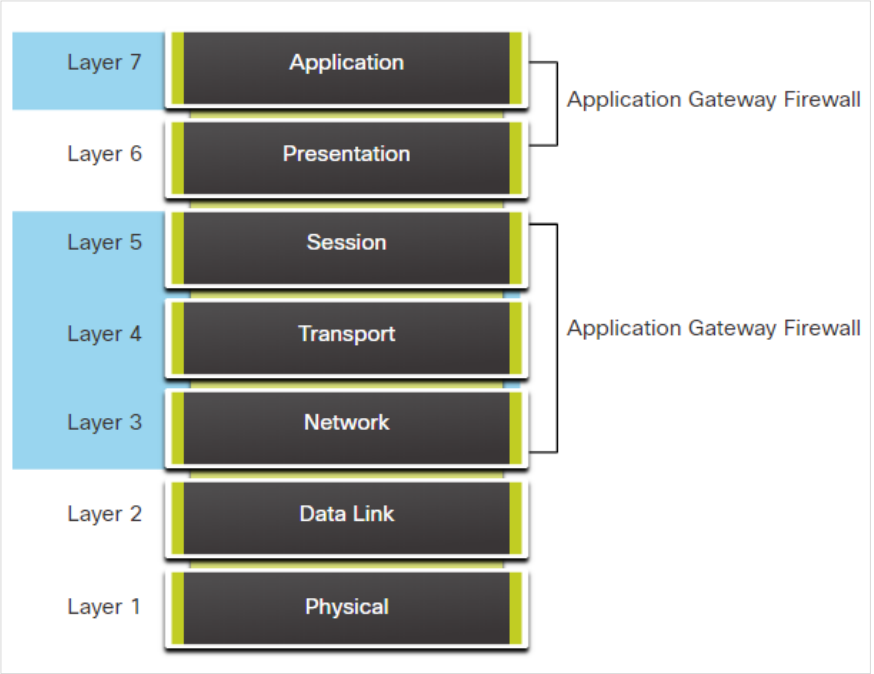

Firewall Type Descriptions (Contd.)

- **Next-generation firewalls (NGFW)** go beyond stateful firewalls by providing:
  - Integrated intrusion prevention
  - Application awareness and control to see and block risky apps
  - Upgrade paths to include future information feeds
  - Techniques to address evolving security threats

Intrusion Prevention and Detection Devices

- A networking architecture paradigm shift is required to defend against fast-moving and evolving attacks. This must include cost effective and prevention systems such as:
  - **Intrusion Detection Systems (IDS)**
  - **Intrusion Prevention Systems (IPS)**
- The network architecture integrates these solutions into the entry and exit points of the network.
- The figure shows how an IPS device handles malicious traffic.

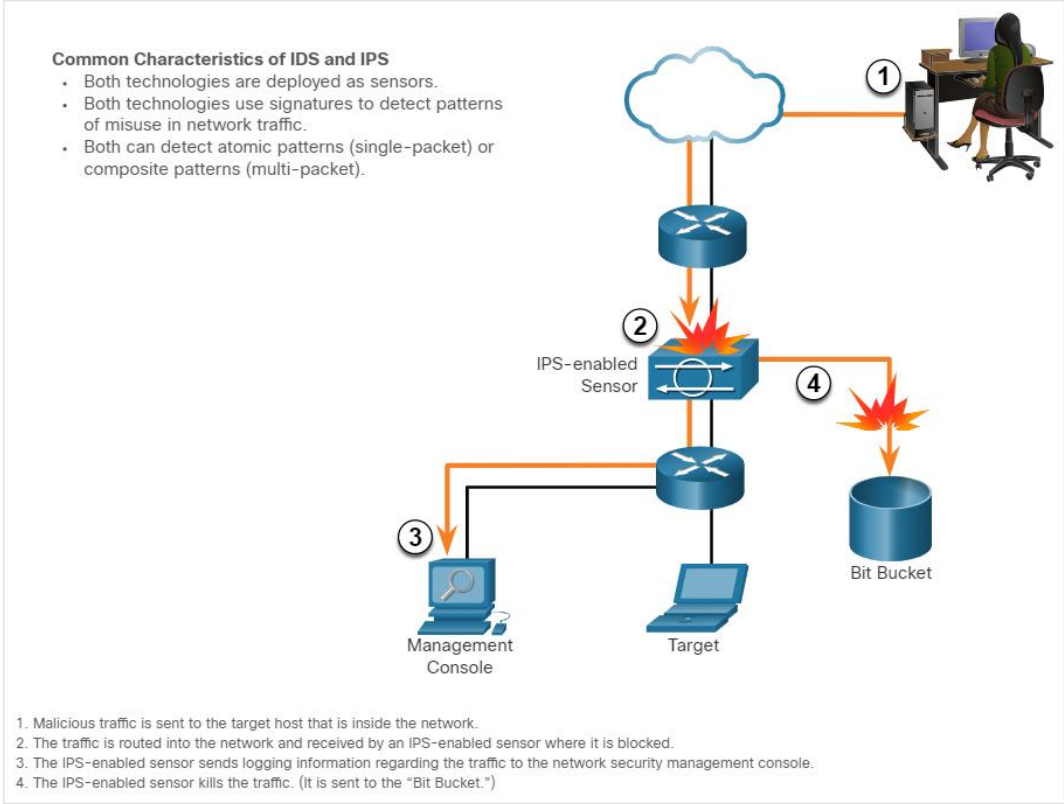

(Dis)Advantages of IDS and IPS

The table lists the advantages and disadvantages of IDS and IPS:

Solution	Advantages	Disadvantages
IDS	• No Impact on network (latency, jitter) • No Network impact if there is a sensor failure • No network impact if there is sensor overload	• Response action cannot stop trigger packets • Correct tuning required for response actions • More vulnerable to network security evasion techniques
IPS	• Stops trigger packets • Can use stream normalization techniques	• Sensor issues might affect network traffic • Sensor overloading impacts the network • Some impact on network (latency, jitter)

Deployment Consideration:

- IPS and IDS technologies can complement each other.
- Deciding which implementation to use is based on the security goals of the organization as stated in their network security policy.

Types of IPS

- There are two primary kinds of IPS :
  - Host-based IPS
  - Network-based IPS
- **Host-based IPS (HIPS)**
  - HIPS is a software installed on a host to monitor and analyze suspicious activity.

Advantages	Disadvantages
• Provides protection specific to a host operating system • Provides operating system and application level protection • Protects the host after the message is decrypted	• Operating system dependent • Must be installed on all hosts

Types of IPS (Contd.)

- **Network-based IPS**
  - Network-based IPS are implemented using a dedicated or non-dedicated IPS device.
  - Sensors detect malicious and unauthorized activity in real time and can take action when required.

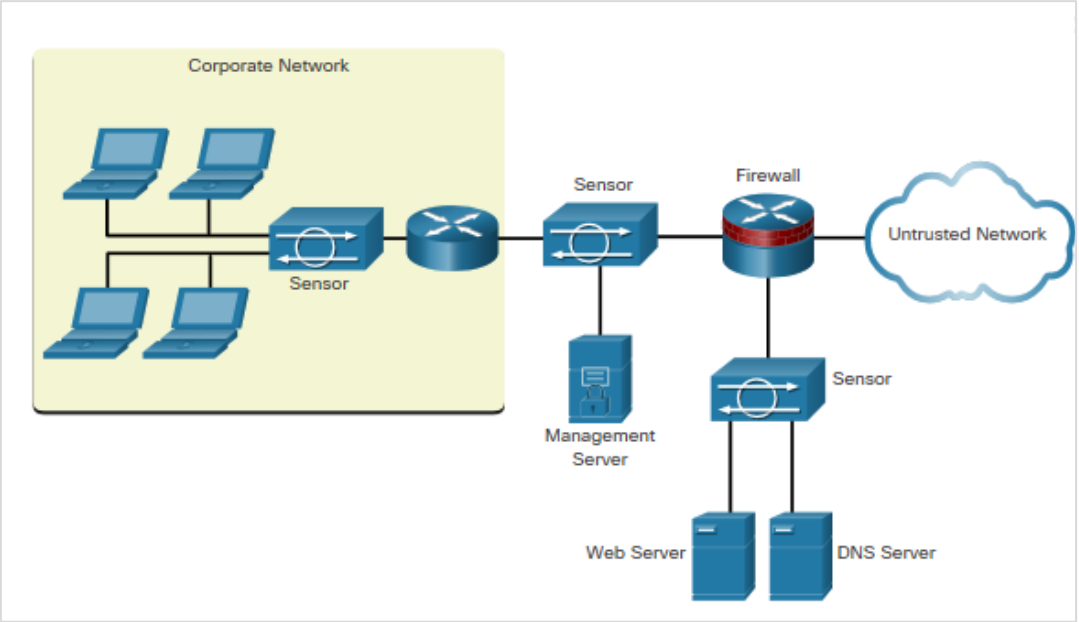

Specialized Security Appliances

Few examples of specialized security appliances.

Cisco Advanced Malware Protection (AMP)	Cisco Web Security Appliance (WSA)	Cisco Email Security Appliance (ESA)
An enterprise-class advanced malware analysis and protection solution	A secure web gateway that combines leading protections to help organizations address the growing challenges of securing and controlling web traffic	ESA/Cisco Cloud Email Security helps to mitigate email-based threats and the ESA defends mission-critical email systems
It provides comprehensive malware protection for organizations before, during, and after an attack	Protects the network by automatically blocking risky sites and testing unknown sites before allowing users to access them	Constantly updated by real-time feeds from Cisco Talos, which detects and correlates threats using a worldwide database monitoring system
		Features: Global threat intelligence, Spam blocking, Advanced Malware Protection, Outbound Message Control

12.3 Security Services

SNMP

- **Simple Network Management Protocol (SNMP)** is an application layer protocol that provides a message format for communication between managers and agents.
- It allows network administrators to perform the following:
  - Manage end devices such as servers, workstations, routers, switches, and security appliances, on an IP network.
  - Monitor and manage network performance.
  - Find and solve network problems.
  - Plan for network growth.
- The SNMP system consists of two elements:
  - **SNMP manager**: Runs SNMP management software.
  - **SNMP agents**: Nodes being monitored and managed.

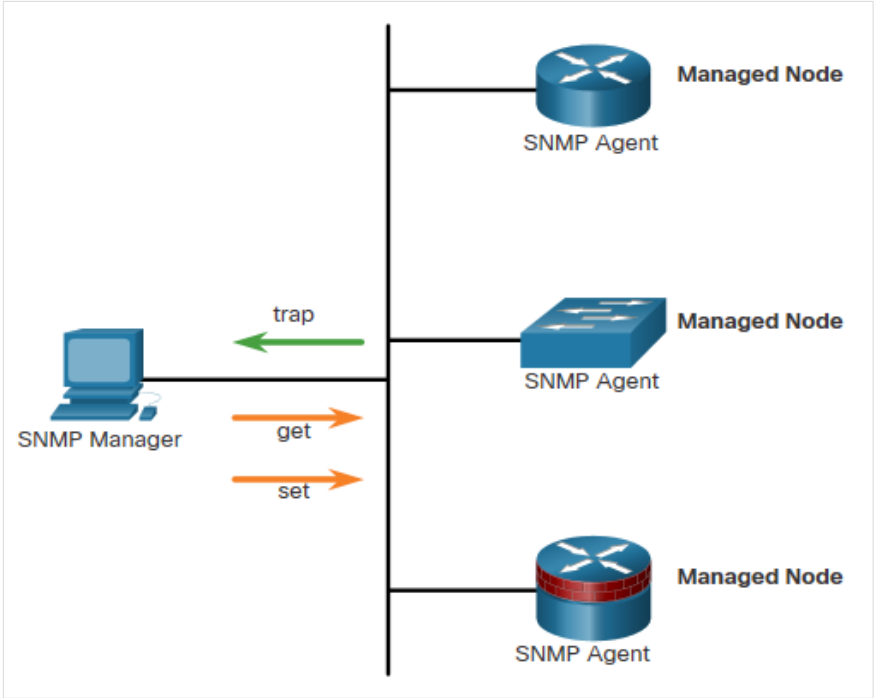

NetFlow

- **NetFlow** is technology that provides statistics on packets flowing through a router or multilayer switch.
- NetFlow provides data to enable:
  - network and security monitoring,
  - network planning
  - traffic analysis to include identification of network bottlenecks
  - IP accounting for billing purposes.
- NetFlow can monitor application connection, tracking byte and packet counts for that individual application flow.
- It then pushes the statistics over to an external server called a **NetFlow collector**.

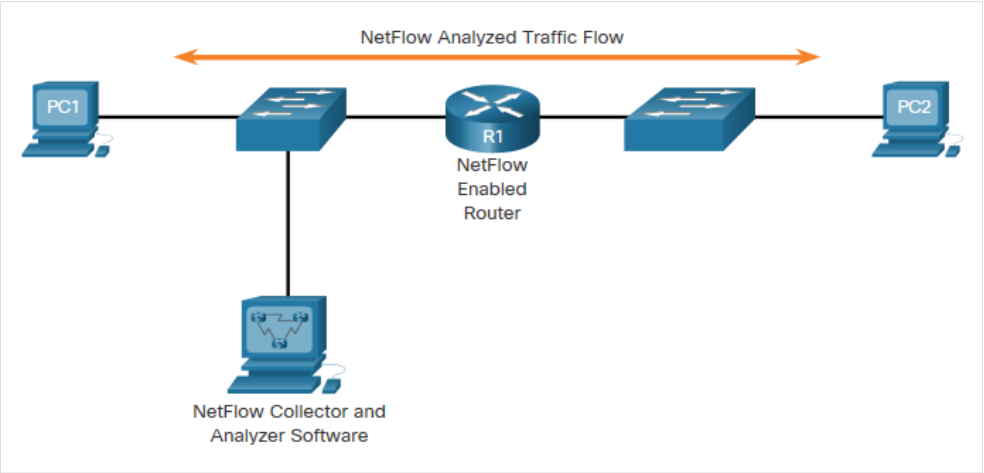

PC 1 connects to PC 2 using HTTPS

Port Mirroring

- **Port mirroring** is a feature that allows a switch to make duplicate copies of traffic passing through a switch, and then sending it out a port with a network monitor attached.

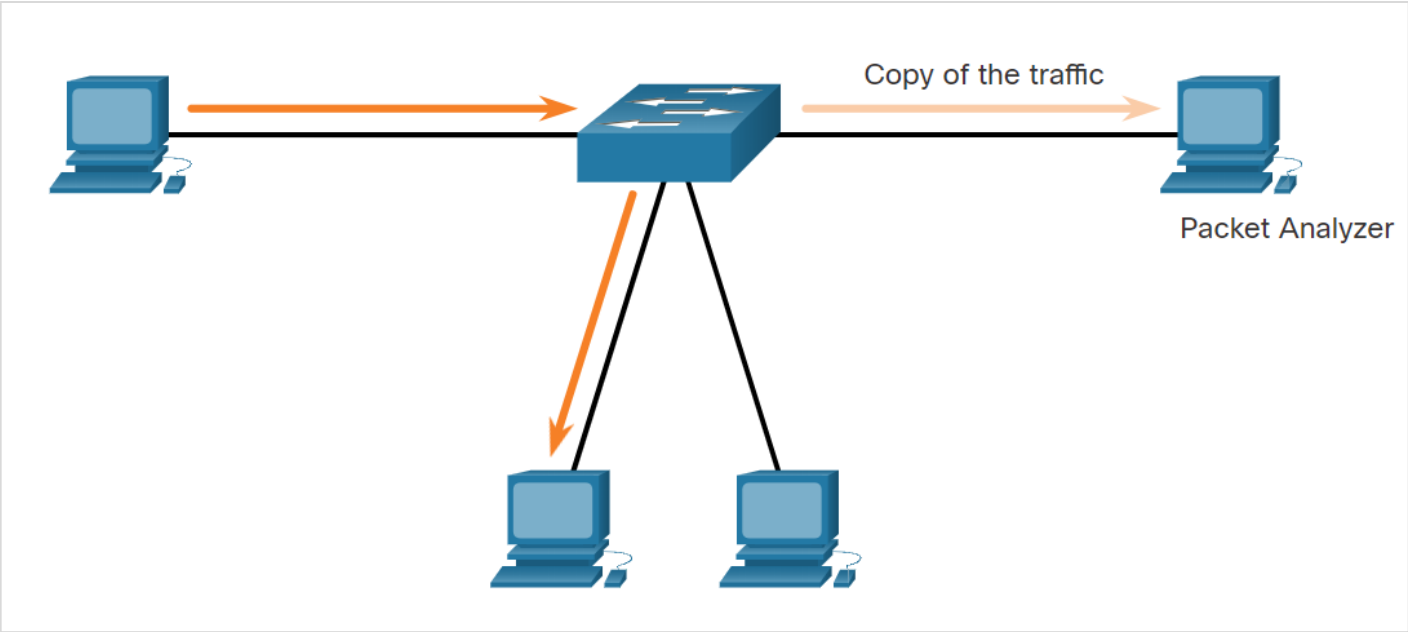

AAA Servers

The below table lists the three independent security functions provided by the AAA architectural framework.

Functions	Description
Authentication	• Users and administrators must prove that they are who they say they are. • Authentication can be established using username and password combinations, challenge and response questions, token cards, and other methods. • AAA authentication provides a centralized way to control access to the network.
Authorization	• After the user is authenticated, authorization services determine which resources the user can access and which operations the user is allowed to perform. • An example is “User ‘student’ can access host serverXYZ using SSH only.”
Accounting	• Accounting records what the user does, including what is accessed, the amount of time the resource is accessed, and any changes that were made. • Accounting keeps track of how network resources are used. • An example is "User ‘student’ accessed host serverXYZ using SSH for 15 minutes."

AAA Servers (Contd.)

The below table lists the difference between Terminal Access Controller Access-Control System Plus (TACACS+) and Remote Authentication Dial-In User Service (RADIUS) protocols.

	TACACS+	RADIUS
Functionality	Separates AAA according to the AAA architecture,	Combines authentication and authorization but separates accounting,
Standard	Mostly Cisco supported	Open/RFC standard
Transport	TCP	UDP
Protocol CHAP	Bidirectional challenge and response as used in Challenge Handshake Authentication Protocol (CHAP)	Unidirectional challenge and response from the RADIUS security server to the RADIUS client
Confidentiality	Entire packet encrypted	Password encrypted
Customization	Provides authorization of router commands on per-user or per-group basis	No option to authorize router commands on a per-user or per-group basis
Accounting	Limited	Extensive

Thank you! Questions?

Vladimír Veselý

updated: 2024-02-24

https://www.fit.vutbr.cz/research/groups/nes@fit

Module 24: Technologies and Protocols

Instructor Materials

CyberOps Associate v1.0

Module Objectives

Module Title: Technologies and Protocols

Module Objective: Explain how security technologies affect security monitoring.

Topic	Topic Objective
Monitoring Common Protocols	Explain the behavior of common network protocols in the context of security monitoring.
Security Technologies	Explain how security technologies affect the ability to monitor common network protocols.

24.1 Monitoring Common Protocols

Syslog Servers

- The most common method of accessing system messages is to use a protocol called syslog.
- The **Syslog** protocol allows networking devices to send their system messages across the network to syslog servers.
- It provides three primary functions:
  - The ability to gather logging information for monitoring and troubleshooting
  - The ability to select the type of logging information that is captured
  - The ability to specify the destination of captured syslog messages

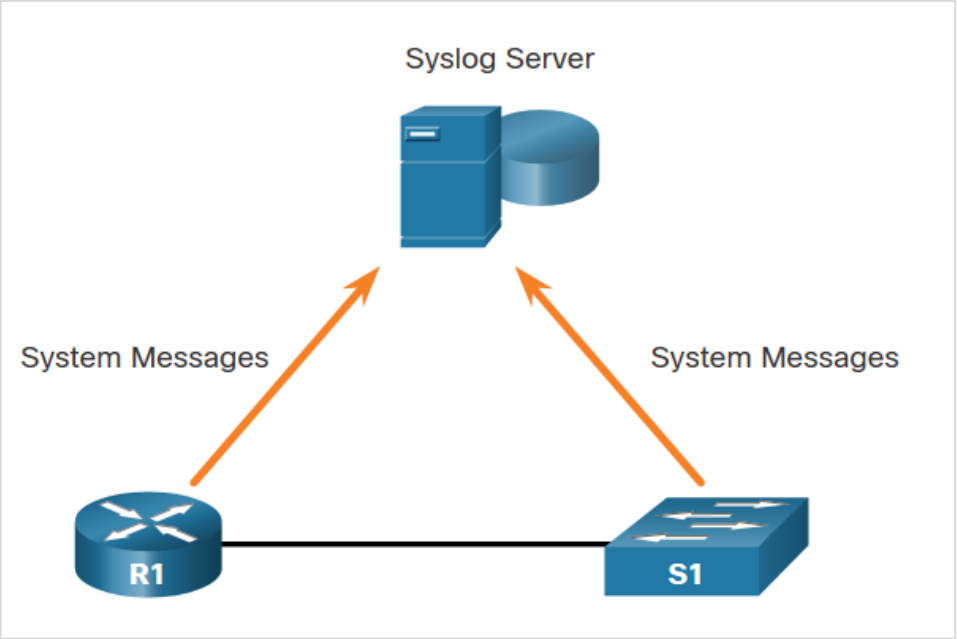

Syslog

Syslog

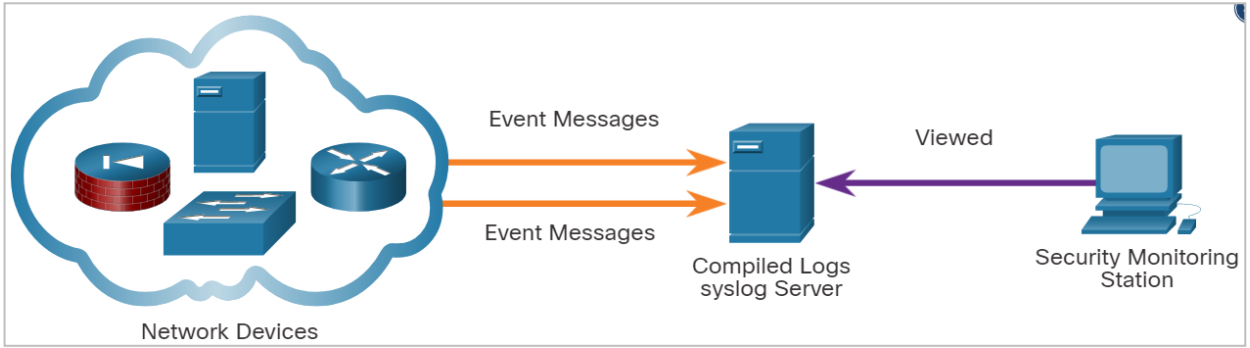

- Syslog offers system-neutral means of transmitting, storing, and analyzing messages.
- Servers that run syslog typically listen on UDP port 514.
- Syslog servers may be a target for threat actors.
  - Some exploits, such as those involving data exfiltration, can take a long time to complete due to the very slow ways in which data is secretly stolen from the network. Some attackers may try to hide the fact that exfiltration is occurring. They attack the syslog servers that contain the information that could lead to detection of the exploit.
  - Hackers may attempt to block the transfer of data from syslog clients to servers, tamper with or destroy log data, or tamper with the software that creates and transmits log messages.
  - The next generation (ng) syslog implementation, known as syslog-ng, offers enhancements that can help prevent some of the exploits that target syslog.

NTP

- It is important to synchronize the time across all devices on the network (for instance for Syslog messages).
- The date and time settings on a network device can be set using one of two methods:
  - Manual configuration of the date and time
  - Configuring the **Network Time Protocol (NTP)**
- NTP networks use a hierarchical system of time sources for synchronization, where each level in this system is called a stratum. NTP servers are arranged in three levels known as strata:
  - **Stratum 0**: An NTP network gets the time from authoritative time sources.
  - **Stratum 1**: Devices are directly connected to the authoritative time sources.
  - **Stratum 2** and lower strata: Stratum 2 devices, such as NTP clients, synchronize their time using the NTP packets from stratum 1 servers.

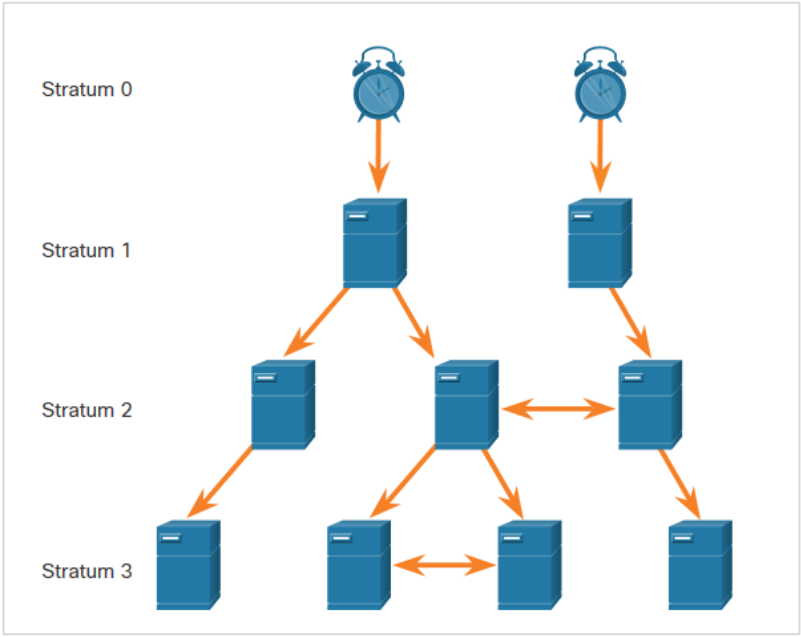

NTP Stratum Levels

NTP

- NTP operates on UDP port 123.
- Threat actors may attempt to attack the NTP infrastructure in order to corrupt time information used to correlate logged network events.
- Threat actors have been known to use NTP systems to direct DDoS attacks through vulnerabilities in client or server software. These attacks can disrupt network availability.

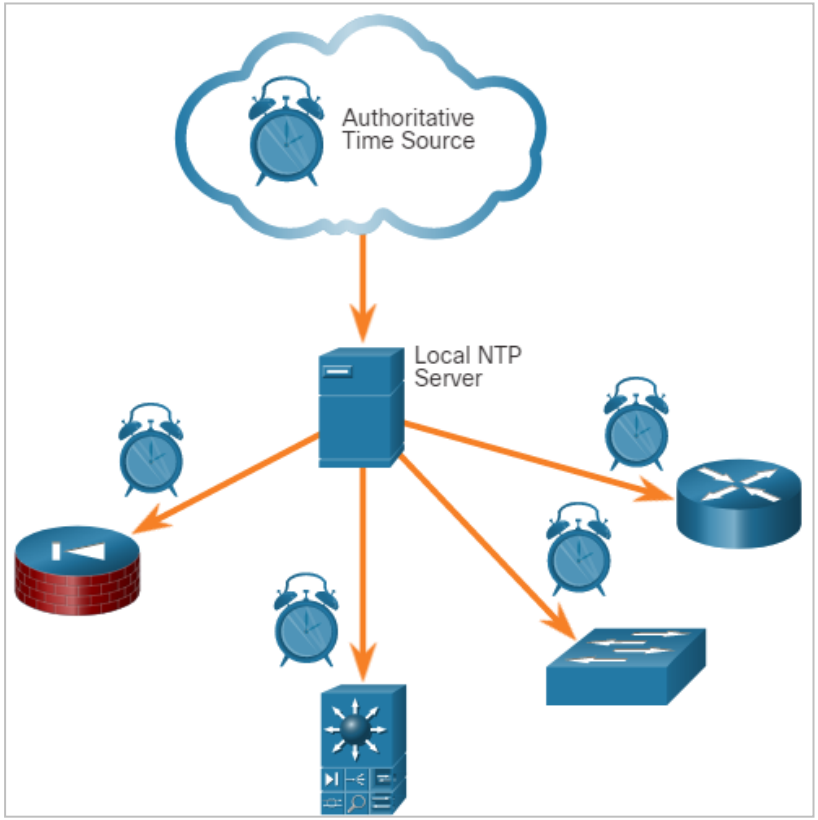

ICMP

- ICMP can be used to identify hosts on a network, the structure of a network, and determine the operating systems at use on the network. It can also be used as a vehicle for various types of DoS attacks.
- ICMP can also be used for data exfiltration.
- Because of the concern that ICMP can be used to surveil or deny service from outside of the network, ICMP traffic from inside the network is sometimes overlooked.
- Some varieties of malware use crafted ICMP packets to transfer files from infected hosts to threat actors using this method, which is known as ICMP tunneling.

10.3 NAT

IPv4 Private Address Space

- To allow a device with a private IPv4 address to access devices and resources outside the local network, the private address must be translated to a public address.
- NAT provides the translation of private addresses to public addresses.
- A single, public IPv4 address can be shared by thousands of devices, each configured with a unique private IPv4 address.
- The solution to the exhaustion of IPv4 address space and the limitations of NAT is the eventual transition to IPv6.

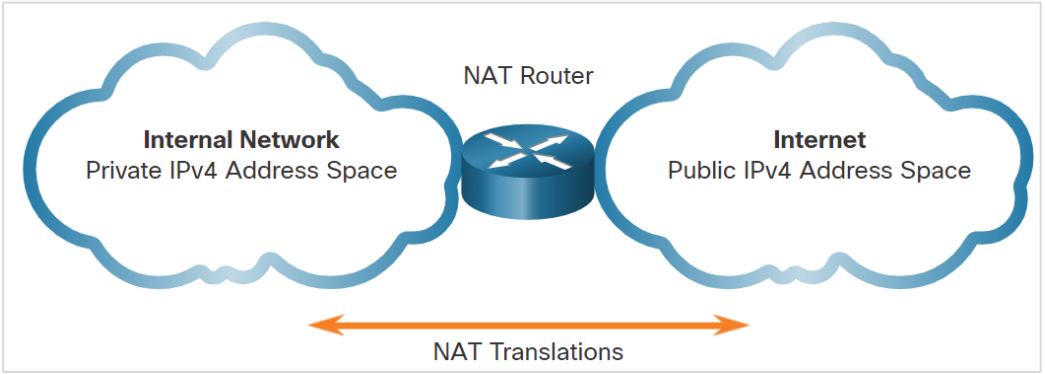

What is NAT?

- NAT is used to conserve public IPv4 addresses.
- NAT-enabled routers can be configured with one or more valid public IPv4 addresses which are known as the **NAT pool**.
- A NAT router typically operates at the border of a stub network.
- When a device inside the stub network wants to communicate with a device outside its network, the packet is forwarded to the border router and the router performs the NAT process.

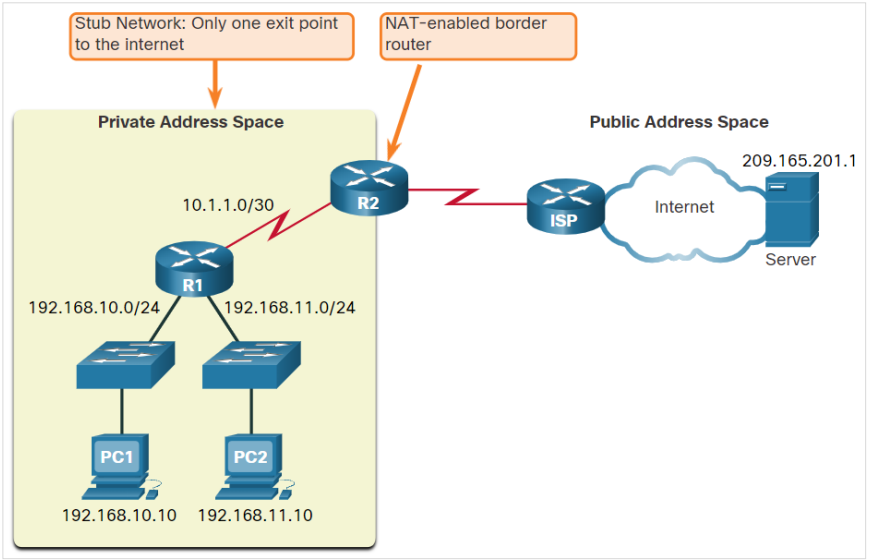

***Note**: The connection to the ISP may use a private address or a public address that is shared among customers. For the purposes of this module, a public address is shown.*

How NAT works?

- In this example, PC1 with private address 192.168.10.10 wants to communicate with an outside web server with public address 209.165.201.1.
- Click the Play button in the figure to view the animation.

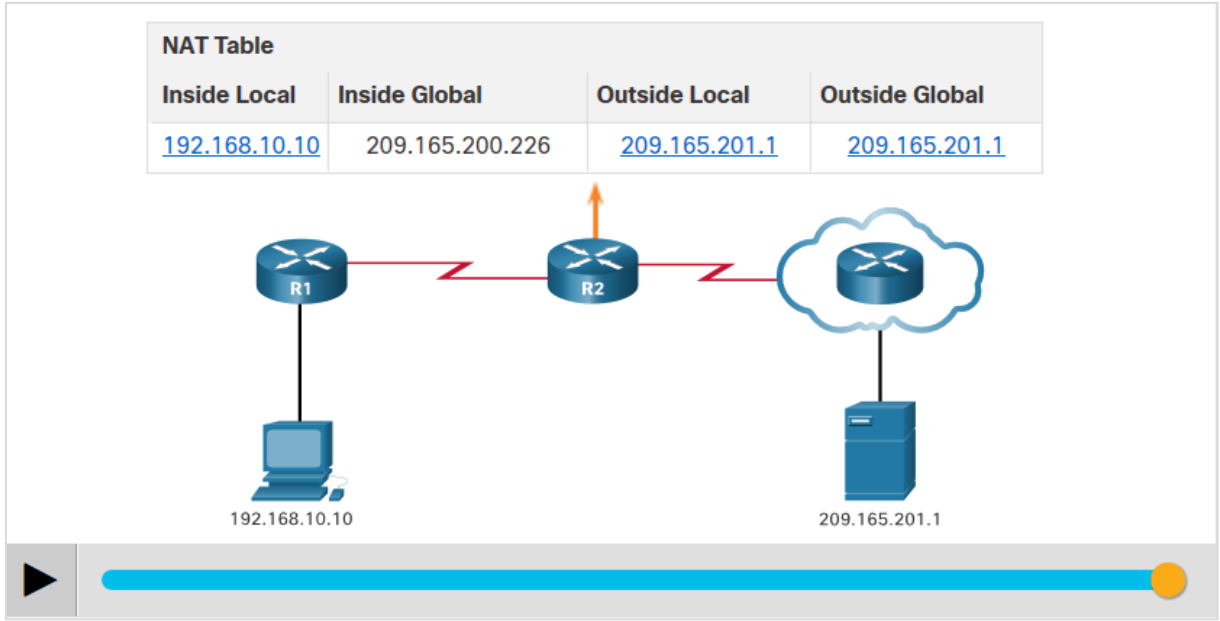

Port Address Translation

- Port Address Translation (PAT), also known as NAT overload, maps multiple private IPv4 addresses to a single public IPv4 address or a few addresses.
- When a device initiates a TCP/IP session, it generates a TCP or UDP source port value, or a specially assigned query ID for ICMP, to uniquely identify the session.
- PAT ensures that devices use a different TCP port number for each session with a server on the internet.
- PAT adds unique source port numbers to the inside global address to distinguish between translations.

Click Play in the figure to view an animation of the PAT process.

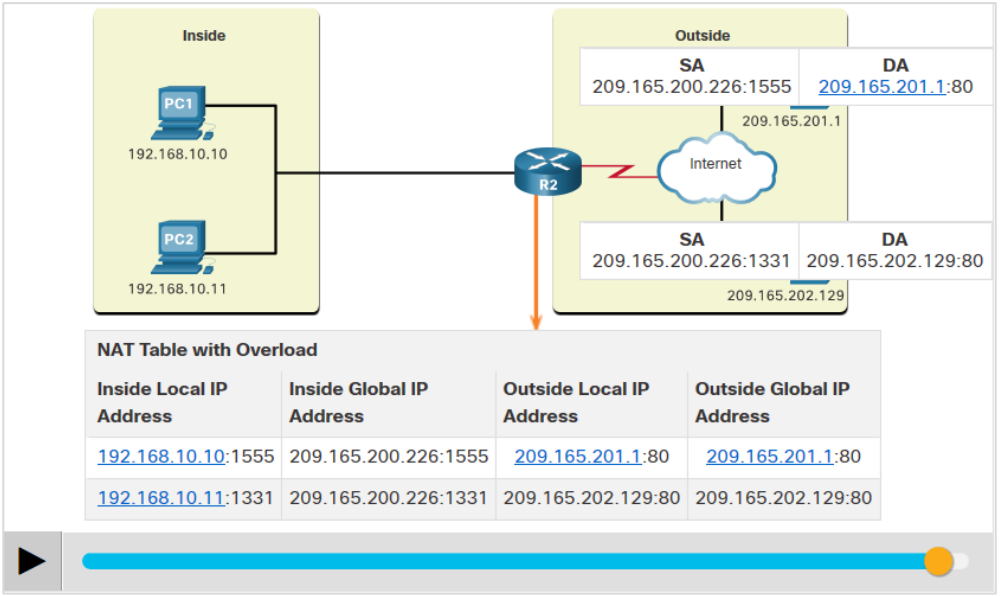

NAT and PAT

- Network Address Translation (NAT) and Port Address Translation (PAT) can complicate security monitoring.
- The figure shows the relationship between internal and external addresses that are used as Source Addresses (SA) and Destination Addresses (DA).
- If PAT is in effect, it could be difficult to log the specific inside device that is requesting and receiving the traffic when it enters the network.
- This problem can be relevant with NetFlow data. NetFlow flows are unidirectional and are defined by the addresses and ports that they share.

Network Address Translation

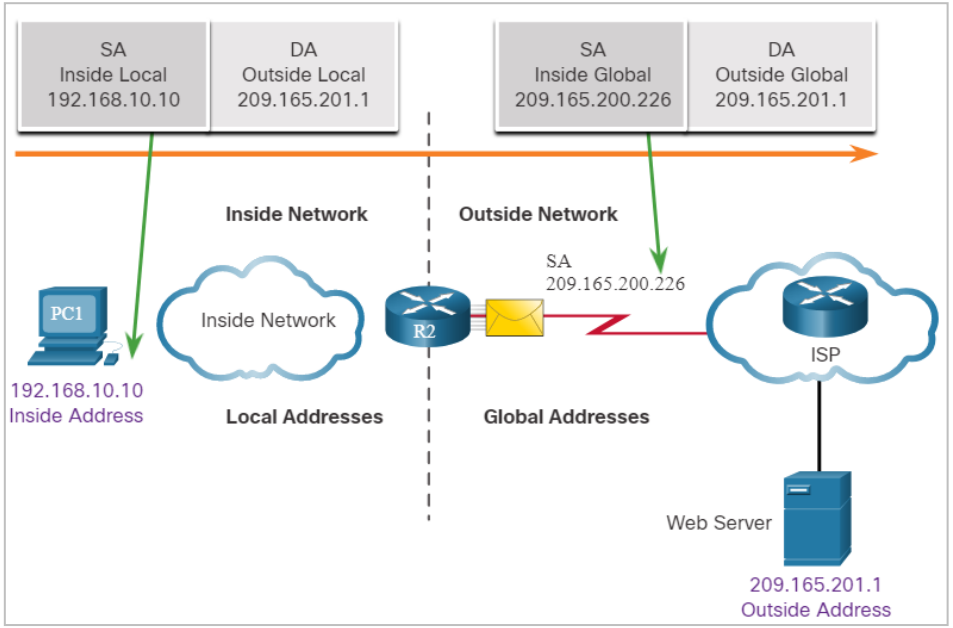

24.2 Security Technologies

Traffic Control with ACLs

- An **Access Control List (ACL)** is a series of commands that control whether a device forwards or drops packets based on information found in the packet header.
- When configured, ACLs perform the following tasks:
  - Limit network traffic to increase network performance.
  - Provide traffic flow control.
  - Provide basic level of security for network access.
  - Filter traffic based on traffic type.
  - Screen hosts to permit or deny access to network services.

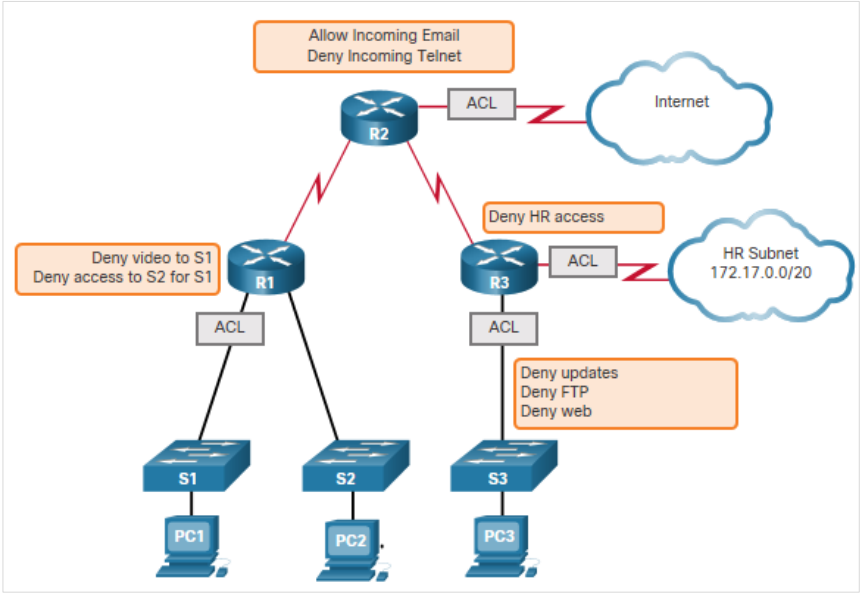

Sample Topology with ACLs applied to routers R1, R2, and R3.

ACLs: Important Features

- **Standard ACL** - Used to permit or deny traffic only from source IPv4 addresses.
- **Extended ACL** - Filters IPv4 packets based on several attributes that include:
  - Protocol type
  - Source IPv4 address
  - Destination IPv4 address
  - Source TCP or UDP ports
  - Destination TCP or UDP ports
  - Optional protocol type information for finer control
- Standard and extended ACLs can be created using either a number or a name to identify the ACL and its list of statements.

ACLs

- Access Control Lists (ACLs) and packet filtering are technologies that contribute to an evolving set of network security protections.
- The figure shows the use of ACLs to permit only specific types of Internet Control Message Protocol (ICMP) traffic. The server at 192.168.1.10 is part of the inside network and is allowed to send ping requests to the outside host at 209.165.201.3.
- The outside host's return ICMP traffic is allowed if it is an ICMP reply or any ICMP unreachable message. All other ICMP traffic types are denied.

Mitigating ICMP Abuse

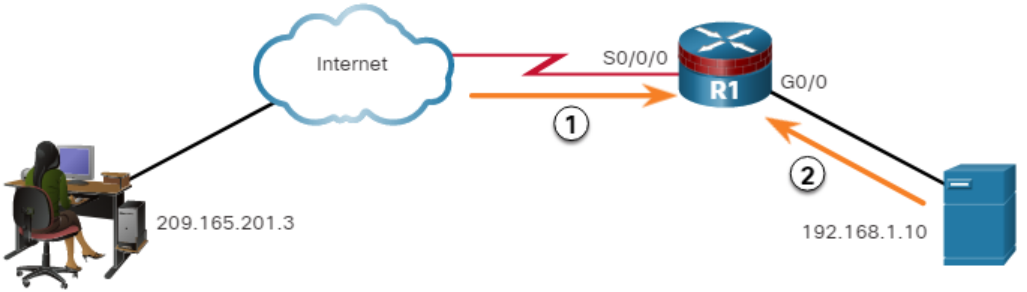

1. Rules on R1 for ICMP traffic from the Internet

```
access-list 112 permit icmp any any echo-reply
access-list 112 permit icmp any any source-quench
access-list 112 permit icmp any any unreachable
access-list 112 deny icmp any any
access-list 112 permit ip any any
```

2. Rules on R1 for ICMP traffic from inside the network

```
access-list 114 permit icmp 192.168.1.0 0.0.0.255 any echo
access-list 114 permit icmp 192.168.1.0 0.0.0.255 any parameter-problem
access-list 114 permit icmp 192.168.1.0 0.0.0.255 any packet-too-big
access-list 114 permit icmp 192.168.1.0 0.0.0.255 any source-quench
access-list 114 deny icmp any any
access-list 114 permit ip any any
```

ACLs (Contd.)

- Attackers can determine which IP addresses, protocols, and ports are allowed by ACLs. This can be done either by port scanning or penetration testing, or through other forms of reconnaissance.
- Attackers can craft packets that use spoofed source IP addresses.
- Applications can establish connections on arbitrary ports. Other features of protocol traffic can also be manipulated, such as the established flag in TCP segments. Rules cannot be anticipated and configured for all emerging packet manipulation techniques.
- In order to detect and react to packet manipulation, more sophisticated behavior and context-based measures need to be taken.
- Cisco Next Generation firewalls, Advanced Malware Protection (AMP), and email and web content appliances are able to address the shortcomings of rule-based security measures.

VPN

- Virtual Private Network is created over a public network (usually the internet).
- A VPN uses virtual connections routed through the Internet from the organization to the remote site.
- A VPN is a communications environment in which access is strictly controlled to permit peer connections within a defined community of interest.
- Confidentiality is achieved by encrypting the traffic within the VPN.
- In short, **VPN** connects two endpoints over a public network, to form a logical connection which can be made at Layer 2 or Layer 3.

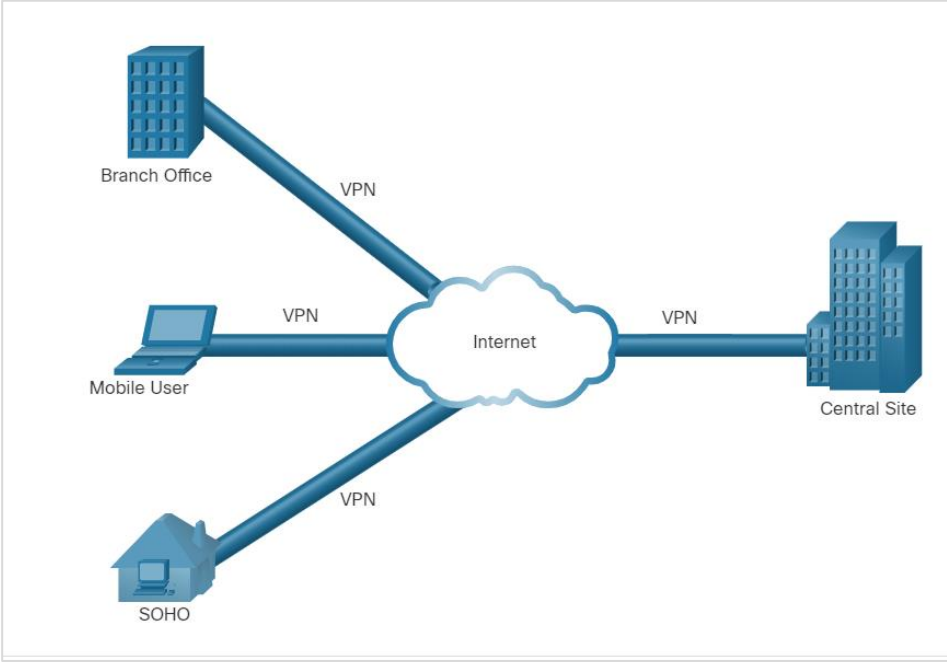

Virtual Private Network

Encryption, Encapsulation, and Tunneling

- Encryption can present challenges to security monitoring by making packet details unreadable.
- Encryption is part of VPN technologies. In VPNs, IP is used to carry encrypted traffic.
- The encrypted traffic essentially establishes a virtual point-to-point connection between networks over public facilities.
- Encryption makes the traffic unreadable to any other devices but the VPN endpoints.
- A similar technology can be used to create a virtual point-to-point connection between an internal host and threat actor devices.
- Malware can establish an encrypted tunnel that rides on a common and trusted protocol, and use it to exfiltrate data from the network.

Peer-to-Peer Networking and Tor

- In peer-to-peer (P2P) networking hosts can operate in both client and server roles.
- The three types of P2P applications are **file sharing, processor sharing,** and **instant messaging**.

P2P

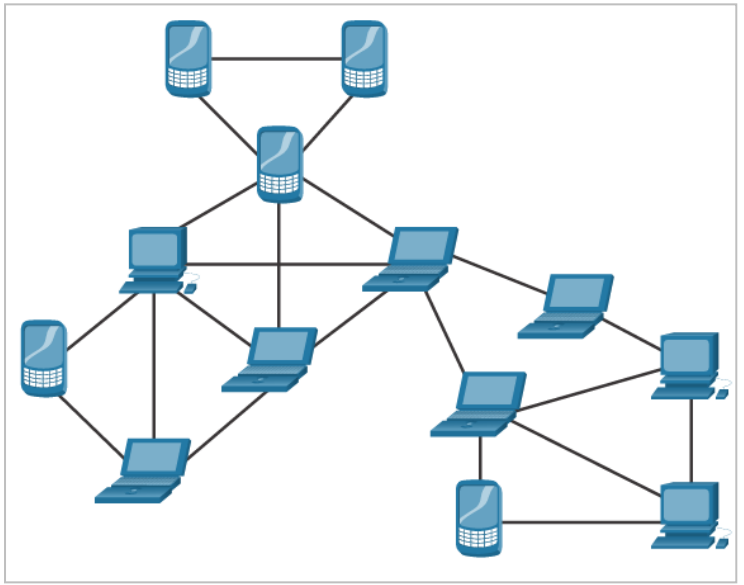

- **File sharing P2P**
  - In file sharing P2P, files on a participating machine are shared with members of the P2P network.
  - File-sharing P2P applications should not be allowed on corporate networks. P2P network activity can avoid firewall protections and is a common vector for the spread of malware.
  - BitTorrent is a P2P file sharing network.

Peer-to-Peer Networking and Tor (Contd.)

- P2P is inherently dynamic. It can operate by connecting to numerous destination IP addresses, and it can also use dynamic port numbering.
- **Processor sharing P2P networks** donate processor cycles to distributed computational tasks.
  - Cancer research, searching for extraterrestrials, and scientific research use donated processor cycles to distribute computational tasks.
- **Instant messaging (IM)** is also considered to be a P2P application.
  - IM has legitimate value within organizations that have geographically distributed project teams.
  - In this case, specialized IM applications are available, such as the Webex Teams platform, which are more secure than IM that uses public servers.

Peer-to-Peer Networking and Tor (Contd.)

- Tor is a software platform and network of P2P hosts that function as internet routers on the Tor network.
- The **Tor** network allows users to browse the internet anonymously. Users access the Tor network by using a special browser.
- When browsing begins, the browser constructs a layered end-to-end path across the Tor server network that is encrypted, as shown in the figure.

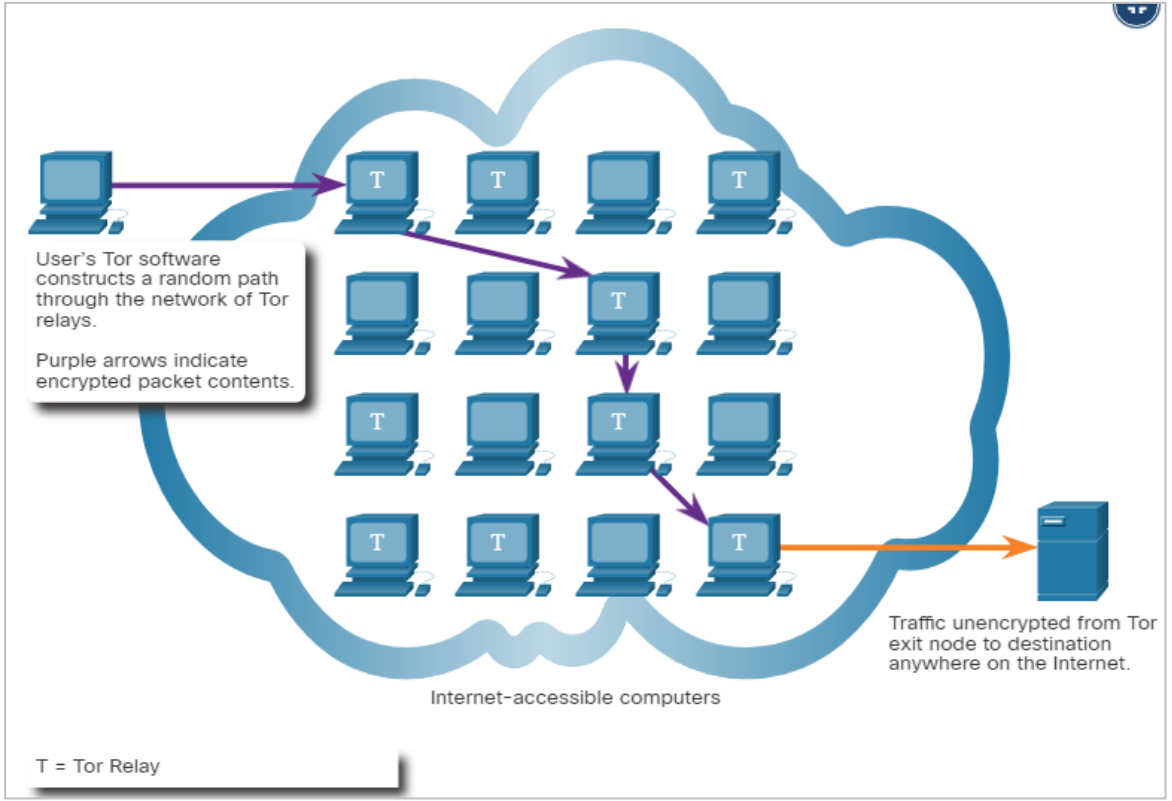

Peer-to-Peer Networking and Tor (Contd.)

- Each encrypted layer is "peeled away" like the layers of an onion as the traffic traverses a Tor relay. The layers contain encrypted next-hop information that can only be read by the router that needs to read the information.
- In this way, no single device knows the entire path to the destination, and routing information is readable only by the device that requires it.
- Finally, at the end of the Tor path, the traffic reaches its internet destination.
- When traffic is returned to the source, an encrypted layered path is again constructed.
- Tor presents a number of challenges to cybersecurity analysts.
  - First, Tor is widely used by criminal organizations on the "dark net."
  - Also, Tor has been used as a communications channel for malware CnC.
- As the destination IP address of Tor traffic is confused by encryption, with only the next-hop Tor node known, Tor traffic avoids blacklists that have been configured on security devices.

Load Balancing

- **Load balancing** involves the distribution of traffic between devices or network paths to prevent overwhelming network resources with too much traffic.
- If redundant resources exist, a load balancing algorithm or device will work to distribute traffic between those resources, as shown in the figure.
- One way this is done is through techniques that use DNS to send traffic to resources that have the same domain name but multiple IP addresses.

Load Balancing with DNS Delegation

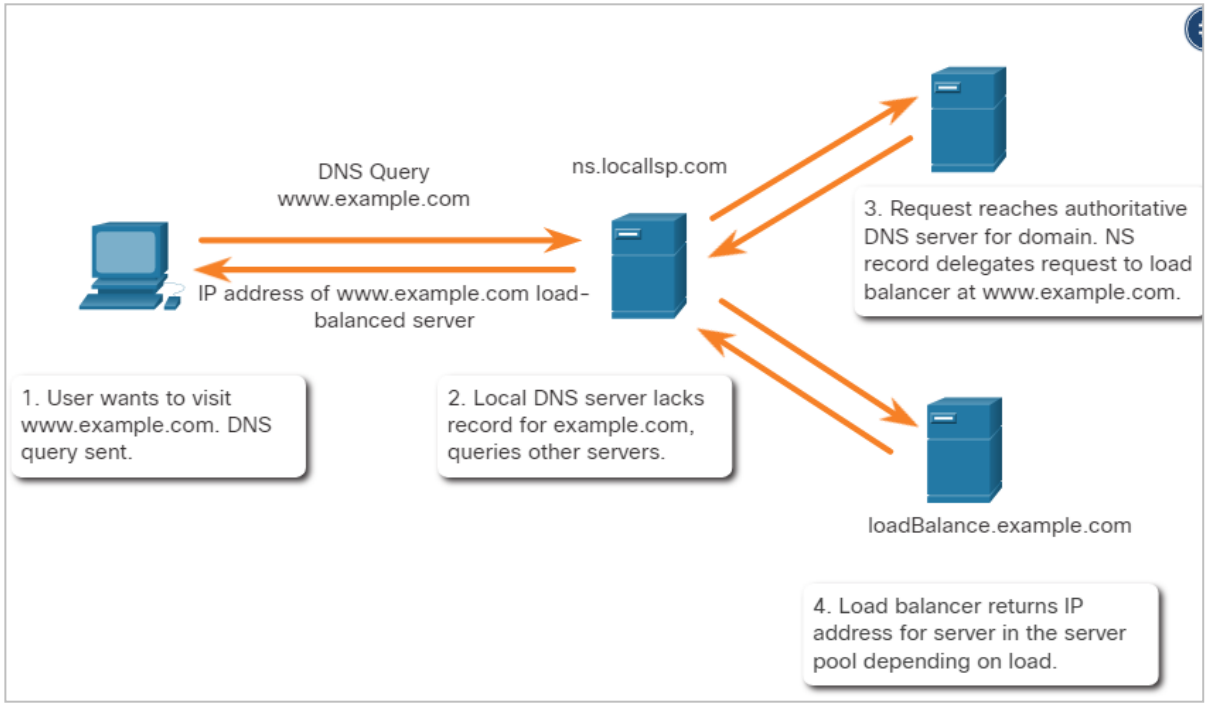

Load Balancing (Contd.)

- In some cases, the distribution may be to servers that are distributed geographically.
  - This results in single internet transaction which is represented by multiple IP addresses on the incoming packets. This may cause suspicious features to appear in packet captures.
- Also, some load balancing manager (LBM) devices use probes to test for the performance of different paths and the health of different devices.
  - An LBM may send probes to the different servers that it is load balancing traffic to in order to detect that the servers are operating.
  - This is done to avoid sending traffic to a resource that is not available.
  - These probes can appear to be suspicious traffic if the cybersecurity analyst is not aware that this traffic is part of the operation of the LBM.

Thank you! Questions?

Vladimír Veselý

updated: 2024-02-24

https://www.fit.vutbr.cz/research/groups/nes@fit