

Role of Transport Layer

- Transport layer
- TCP
- UDP
- TCP/UDP port numbers

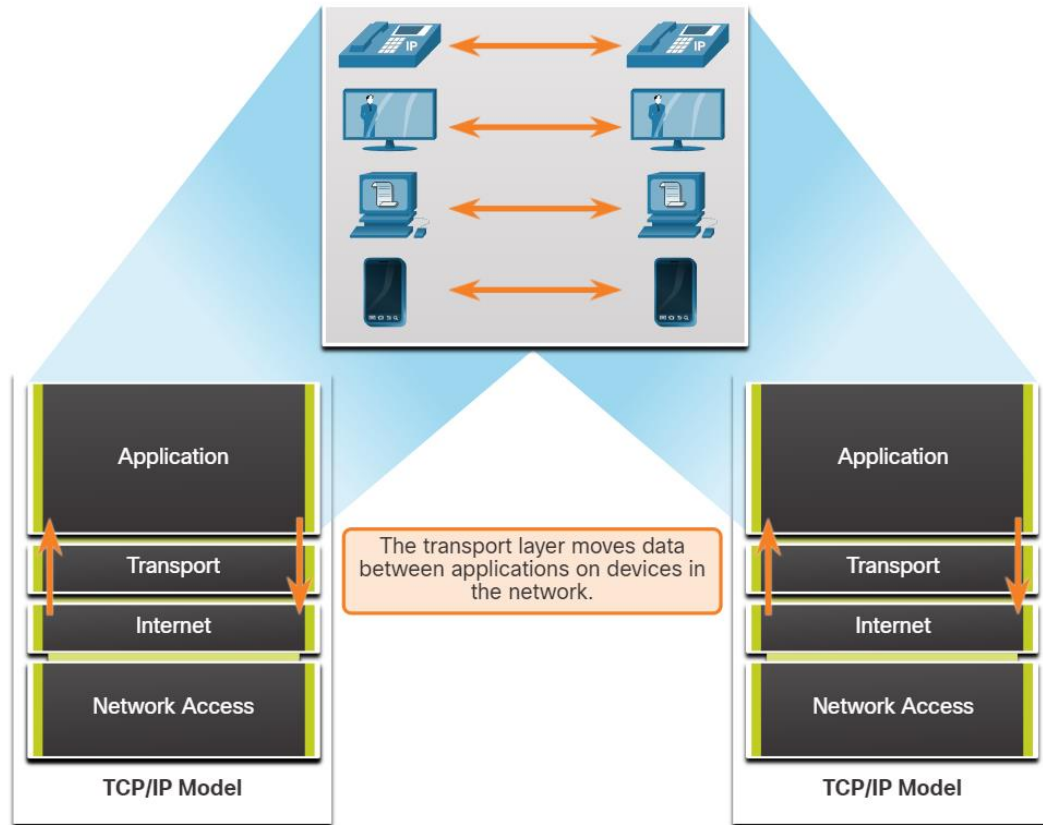
NetAcademy chapter 14.

14

Transport Layer



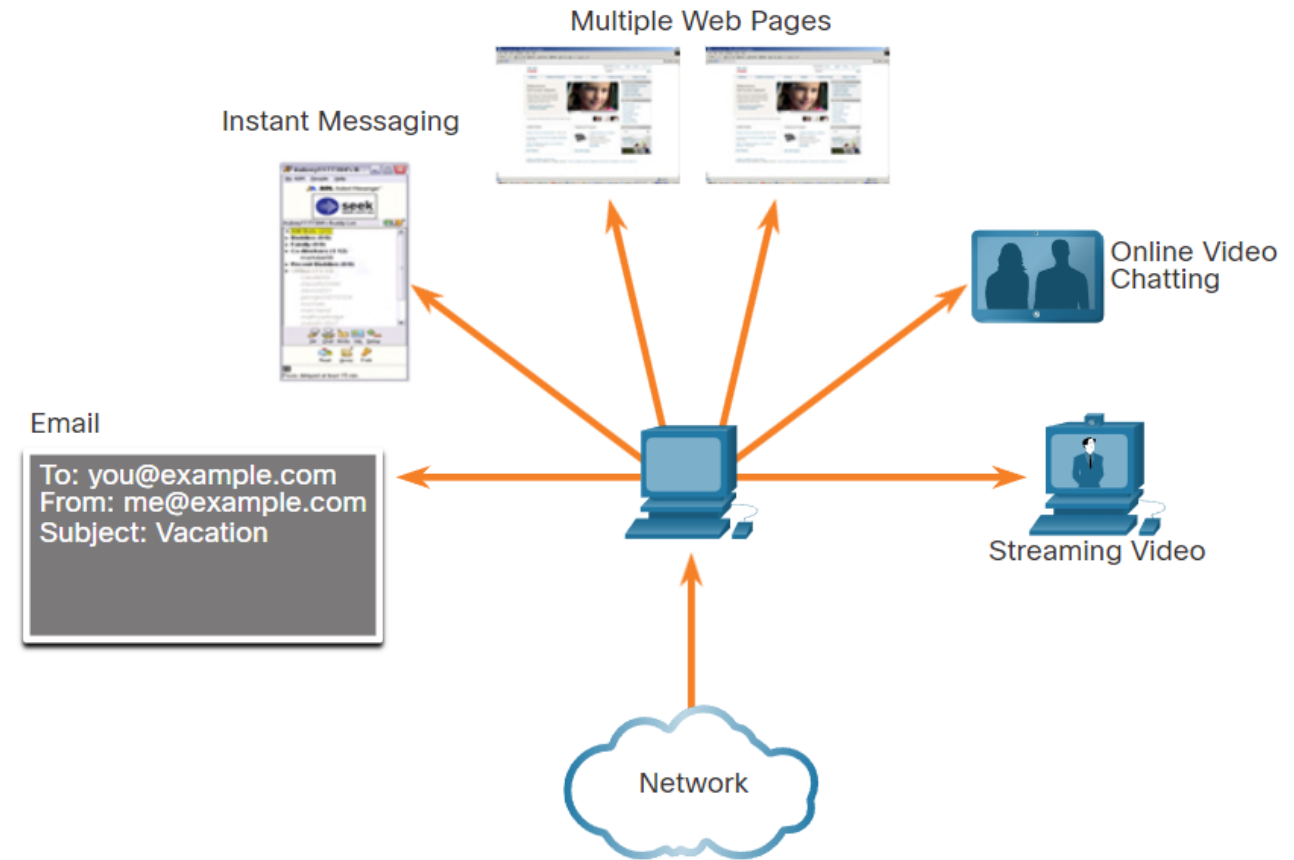
Role of the Transport Layer



- Application layer programs generate data that must be exchanged between source and destination hosts. The transport layer is responsible for logical communications between applications running on different hosts. This may include services such as establishing a temporary session between two hosts and the reliable transmission of information for an application.
- The transport layer has no knowledge of the destination host type, the type of media over 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)**
 - **User Datagram Protocol (UDP)**

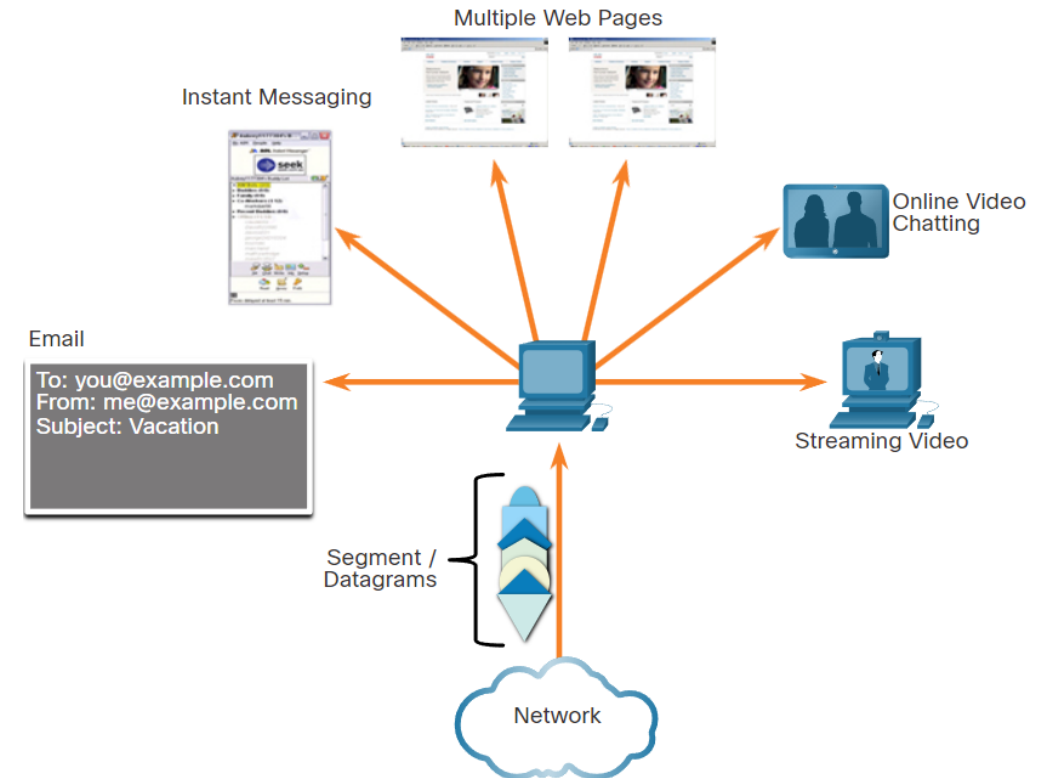
Transport Layer Responsibilities

- At the transport layer, 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**.
- Most networks have a limitation on the amount of data that can be included in a single packet. Therefore, data must be divided into manageable pieces.



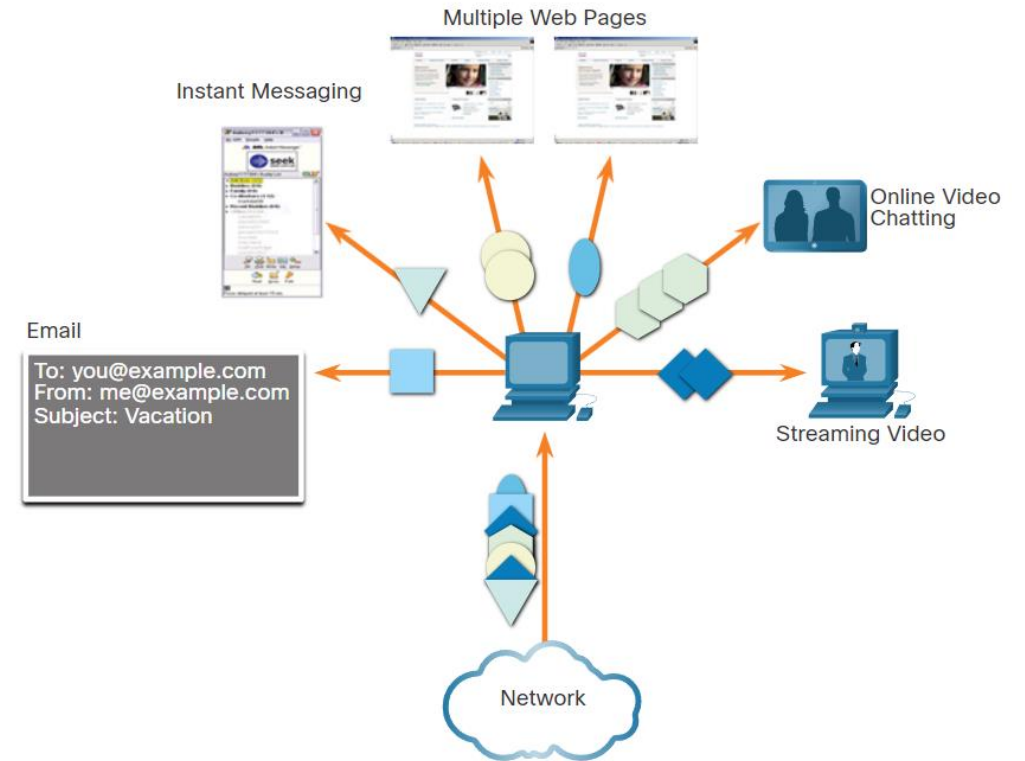
Transport Layer Responsibilities

- **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 illustrates the transport layer using different blocks for each conversation.
- The transport layer **divides the data into smaller blocks** (i.e., segments or datagrams) that are easier to manage and transport.



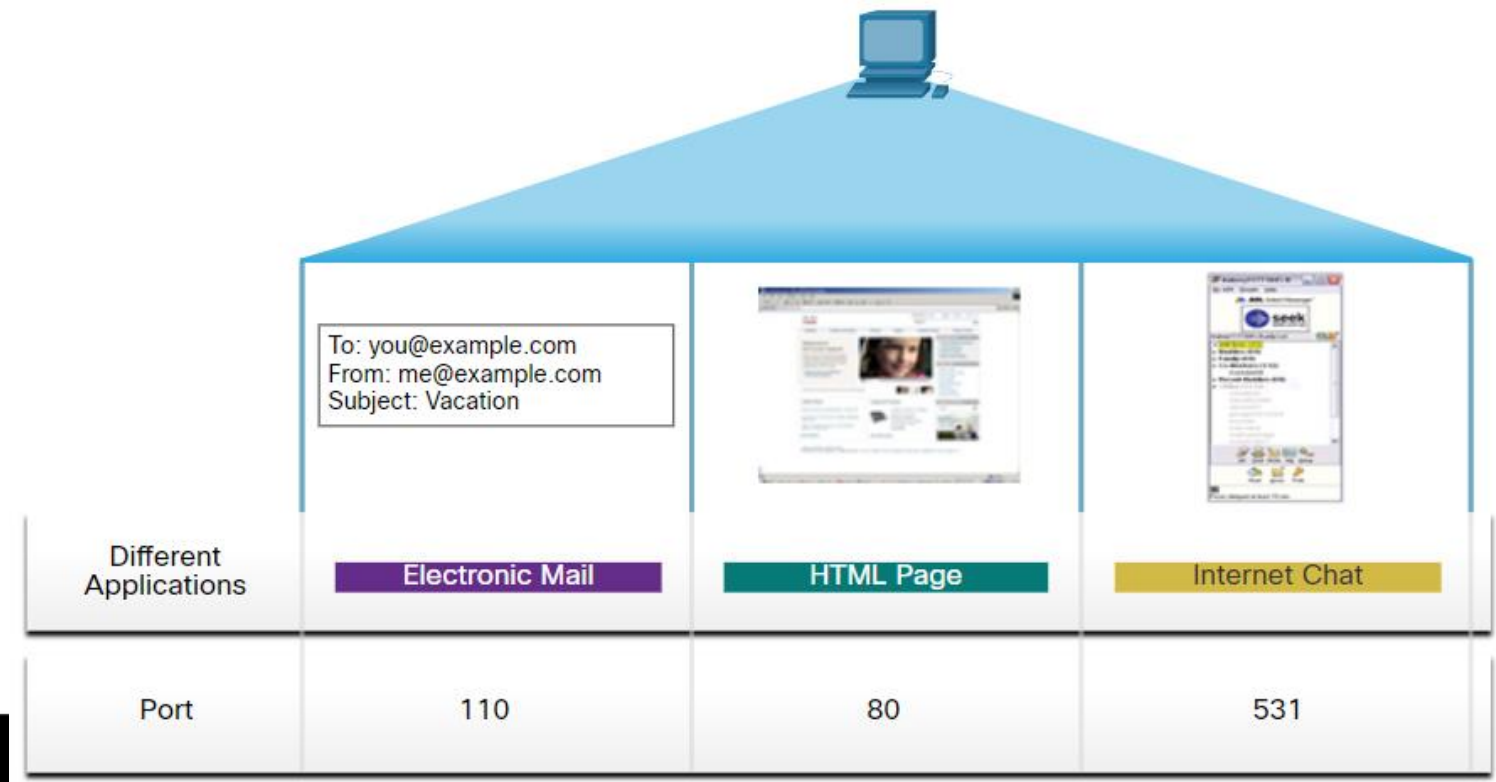
Transport Layer Responsibilities

- **Add Header Information**
- The transport layer protocol also adds **header information containing binary data organized into several fields** to each block of data. It is the values in these fields that enable various transport layer protocols to perform different functions in managing data communication.
- For instance, 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**.



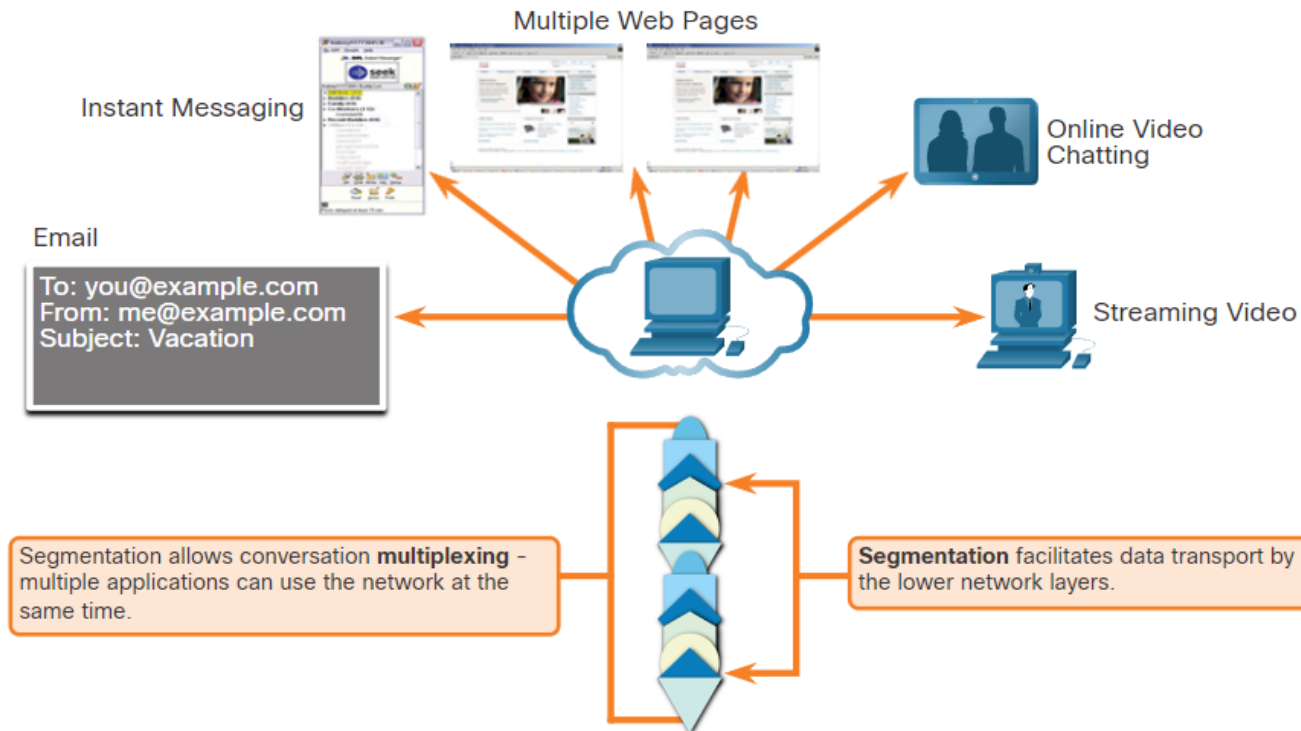
Transport Layer Responsibilities

- **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 illustrated in the figure, each software process that needs to access the network is assigned a port number unique to that host.



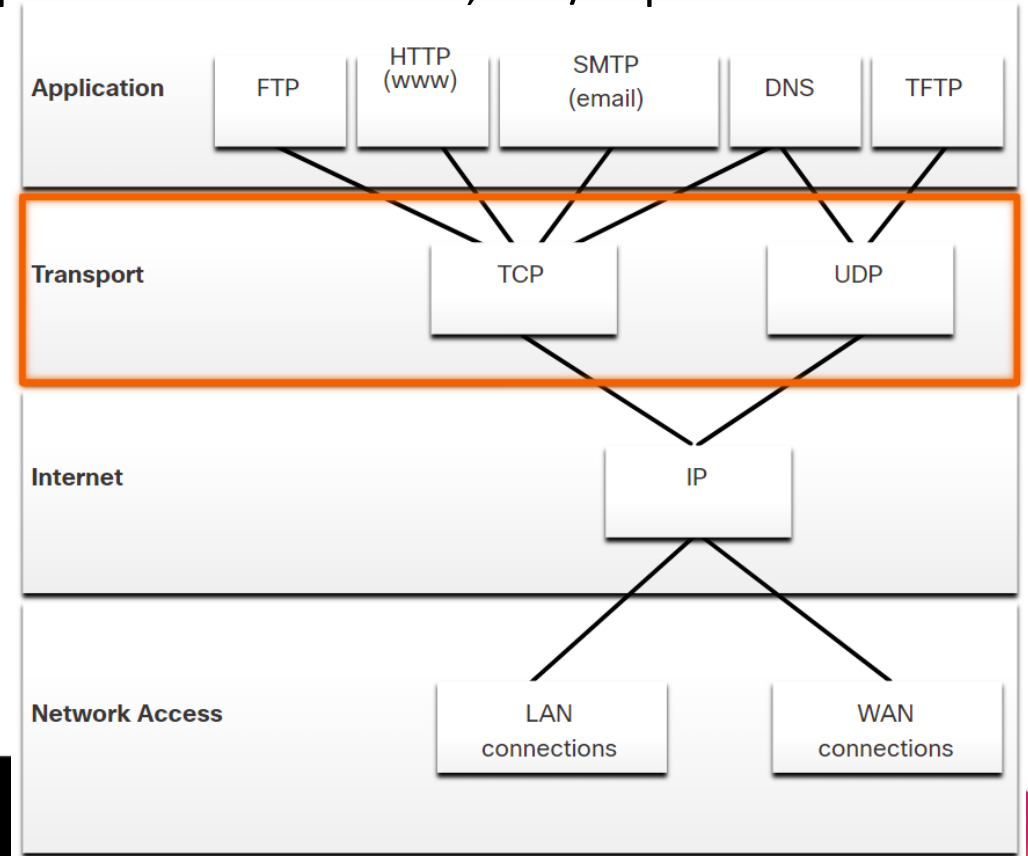
Transport Layer Responsibilities

- **Conversation Multiplexing**
- **Sending** some types of **data** (e.g., a streaming video) across a network, **as one complete communication stream, can consume all the available bandwidth**. **This would prevent other communication conversations from occurring at the same time**. It would also make error recovery and retransmission of damaged data difficult.
- 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 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)

- **IP is not responsible for guaranteeing delivery or determining whether a connection between the sender and receiver needs to be established.**
- 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 divides data into segments.
- TCP transport is analogous to sending packages that are tracked from source to destination. If a shipping order is broken up into several packages, a customer can check online to see the order of the delivery.
- **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
 - 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.

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. Because 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 divides data into **datagrams** *that are also referred to as segments*.
- **UDP is a connectionless protocol**. Because **UDP does not provide reliability or flow control, it does not require an established connection**. Because **UDP does not track information sent or received** between the client and server, **UDP is also known as a stateless protocol**.
- **UDP is also known as a best-effort delivery** protocol because **there is no acknowledgment** that the data is received at the destination. With UDP, **there are no transport layer processes that inform the sender of a successful delivery**.
- 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 vs TCP

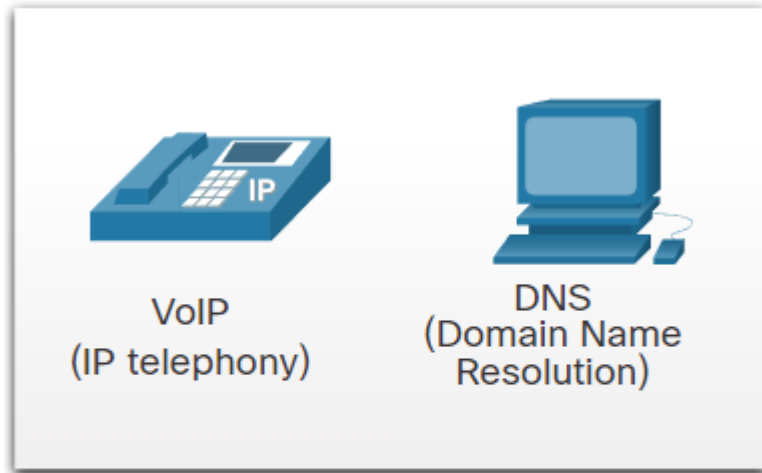
- Some applications can tolerate some data loss during transmission over the network, but delays in transmission are unacceptable. For these applications, UDP is the better choice because it requires less network overhead. UDP is preferable for applications such as Voice over IP (VoIP). Acknowledgments and retransmission would slow down delivery and make the voice conversation unacceptable.
- UDP is also used by request-and-reply applications where the data is minimal, and retransmission can be done quickly. For example, Domain Name System (DNS) uses UDP for this type of transaction. The client requests IPv4 and IPv6 addresses for a known domain name from a DNS server. If the client does not receive a response in a predetermined amount of time, it simply sends the request again.
- For example, if one or two segments of a live video stream fail to arrive, it creates a momentary disruption in the stream. This may appear as distortion in the image or sound, but may not be noticeable to the user. If the destination device had to account for lost data, the stream could be delayed while waiting for retransmissions, therefore causing the image or sound to be greatly degraded. In this case, it is better to render the best media possible with the segments received, and forego reliability.

UDP vs TCP

- For other applications it is important that all the data arrives and that it can be processed in its proper sequence. For these types of applications, TCP is used as the transport protocol. For example, applications such as databases, web browsers, and email clients, require that all data that is sent arrives at the destination in its original condition. Any missing data could corrupt a communication, making it either incomplete or unreadable. For example, it is important when accessing banking information over the web to make sure all the information is sent and received correctly.
- Application developers must choose which transport protocol type is appropriate based on the requirements of the applications. Video may be sent over TCP or UDP. Applications that stream stored audio and video typically use TCP. The application uses TCP to perform buffering, bandwidth probing, and congestion control, in order to better control the user experience.
- Real-time video and voice usually use UDP, but may also use TCP, or both UDP and TCP. A video conferencing application may use UDP by default, but because many firewalls block UDP, the application can also be sent over TCP.
- Applications that stream stored audio and video use TCP. For example, if your network suddenly cannot support the bandwidth needed to watch an on-demand movie, the application pauses the playback. During the pause, you might see a “buffering...” message while TCP works to re-establish the stream. When all the segments are in order and a minimum level of bandwidth is restored, your TCP session resumes, and the movie resumes playing.

UDP vs TCP

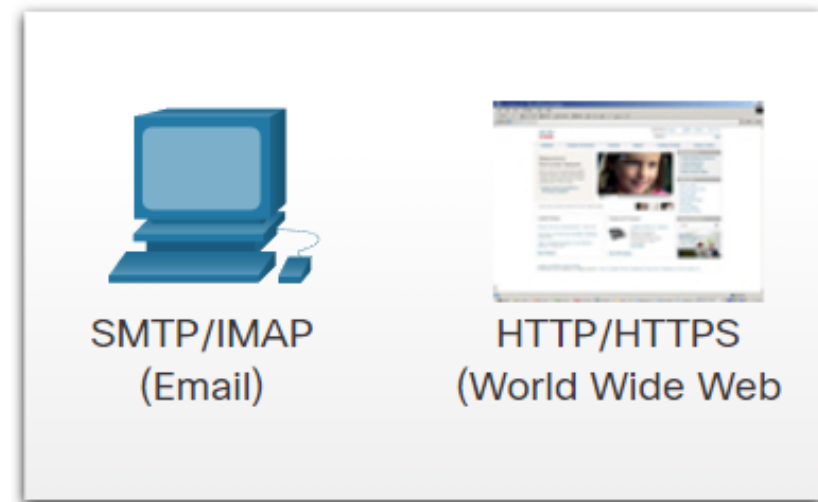
UDP



Required protocol properties:

- Fast
- Low overhead
- Does not require acknowledgements
- Does not resend lost data
- Delivers data as it arrives

TCP



Required protocol properties:

- Reliable
- Acknowledges data
- Resends lost data
- Delivers data in sequenced order

TCP Features

In addition to supporting the basic functions of data segmentation and reassembly, TCP also provides the following services:

Establishes a Session - TCP is a connection-oriented protocol that negotiates and establishes a permanent connection (or session) between source and destination devices prior to forwarding any traffic. Through session establishment, the devices negotiate the amount of traffic that can be forwarded at a given time, and the communication data between the two can be closely managed.

Ensures Reliable Delivery - For many reasons, it is possible for a segment to become corrupted or lost completely, as it is transmitted over the network. TCP ensures that each segment that is sent by the source arrives at the destination.

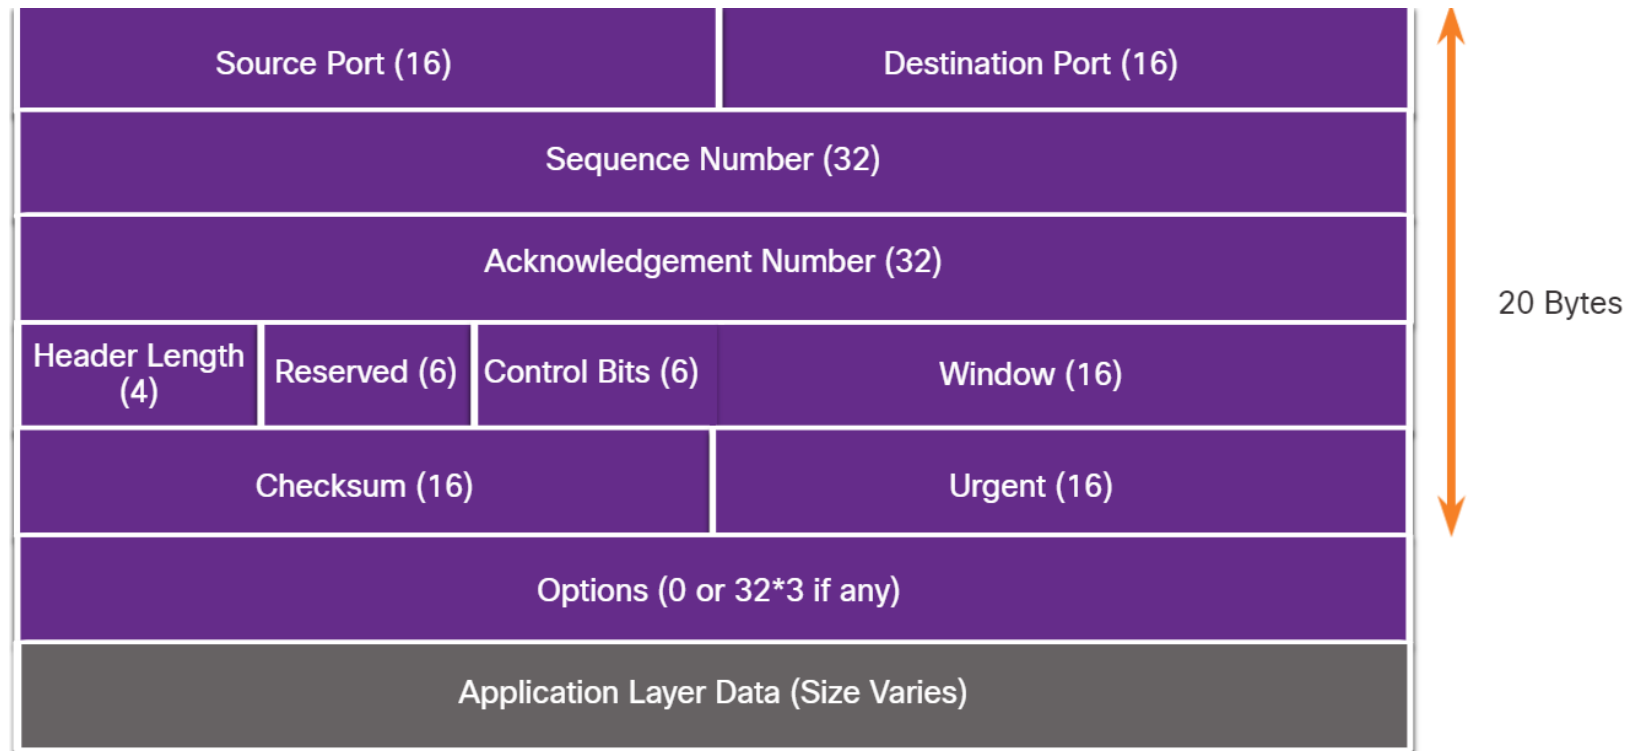
Provides Same-Order Delivery - Because networks may provide multiple routes that can have different transmission rates, data can arrive in the wrong order. By numbering and sequencing the segments, TCP ensures segments are reassembled into the proper order.

Supports Flow Control - Network hosts have limited resources (i.e., memory and processing power). When TCP is aware that these resources are overtaxed, it can request that the sending application reduce the rate of data flow. This is done by TCP regulating the amount of data the source transmits. Flow control can prevent the need for retransmission of the data when the resources of the receiving host are overwhelmed.

TCP Header

TCP is a stateful protocol which means 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 (i.e., 160 bits) of overhead when encapsulating the application layer data.

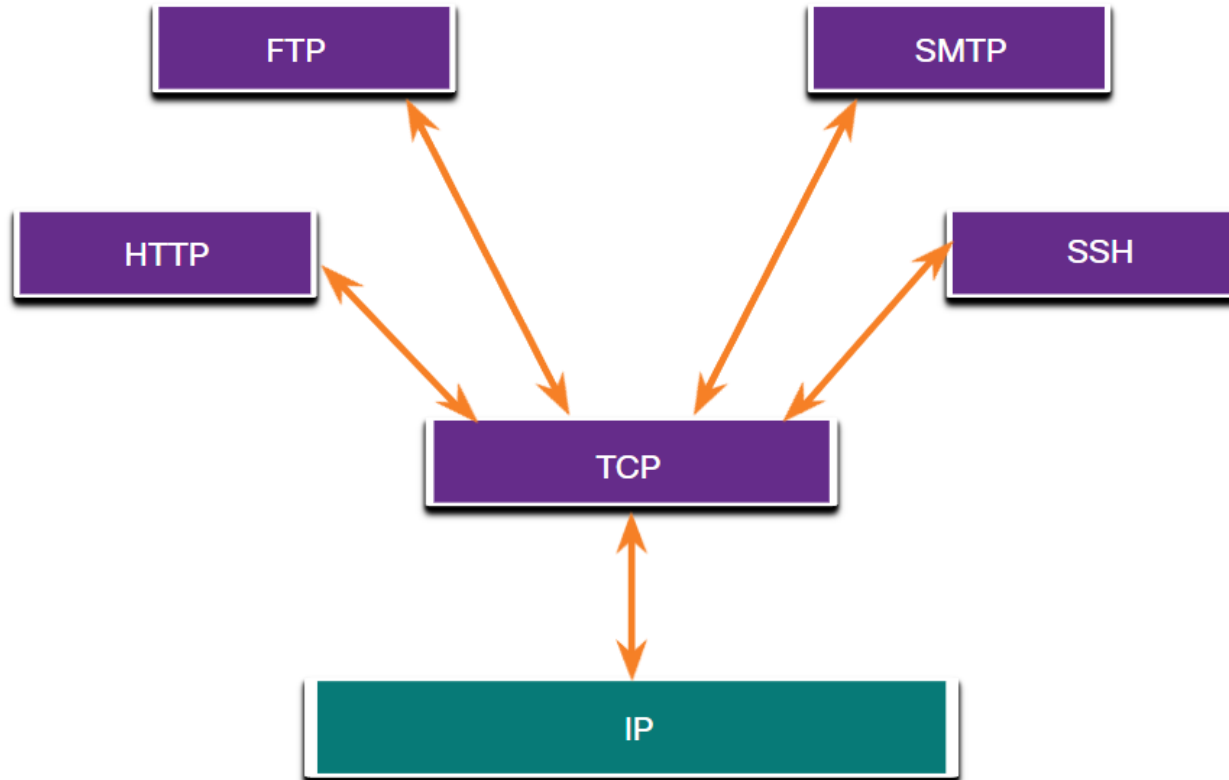


TCP Header Fields

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.

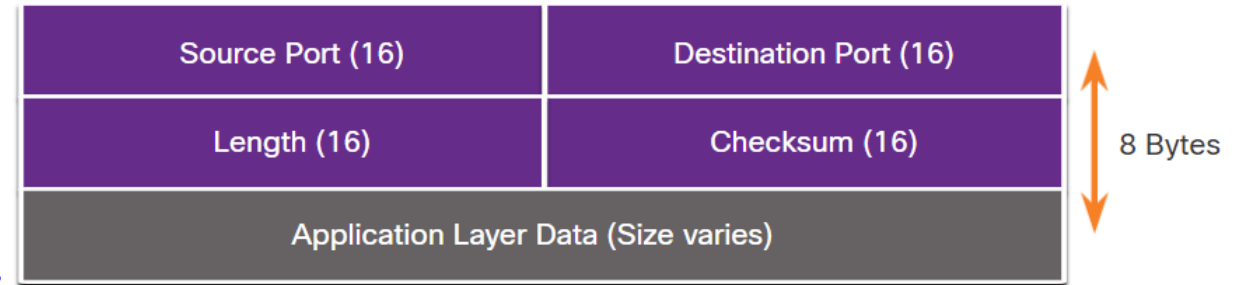
Applications that use TCP

- TCP is a good example of how the different layers of the TCP/IP protocol suite have specific roles. TCP handles all tasks associated with dividing the data stream into segments, providing reliability, controlling data flow, and reordering segments.
- TCP frees the application from having to manage any of these tasks. Applications, like those shown in the figure, can simply send the data stream to the transport layer and use the services of TCP.



UDP Features

- UDP features include the following:
 - Data is reconstructed in the order that it is received.
 - Any segments that are lost are not resent.
 - There is no session establishment.
 - The sending is not informed about resource availability.

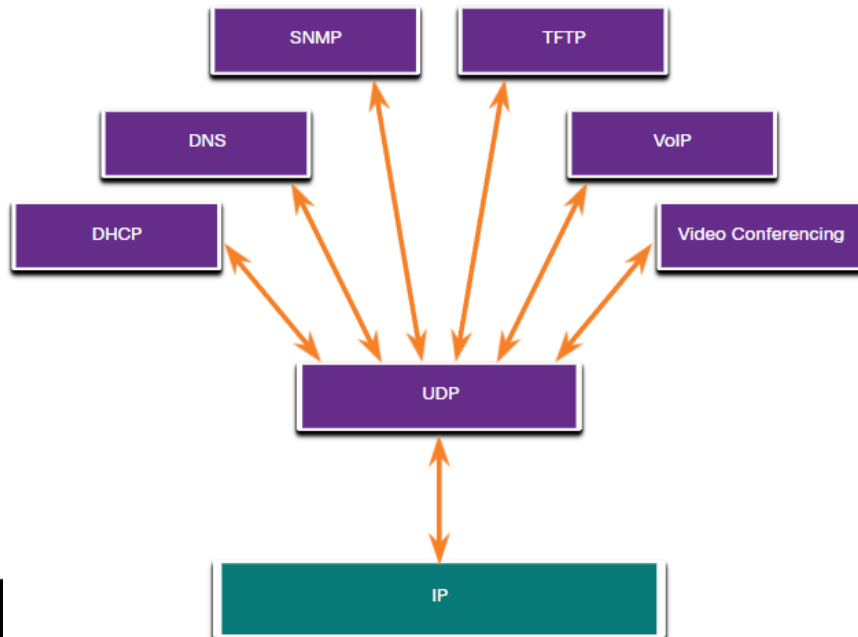


- **UDP is a stateless protocol**, meaning neither the client, nor the server, tracks the state of the communication session. If reliability is required when using UDP as the transport protocol, it must be handled by the application.
- One of the most important requirements for delivering live video and voice over the network is that the data continues to flow quickly. Live video and voice applications can tolerate some data loss with minimal or no noticeable effect, and are perfectly suited to UDP.

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.
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.

Applications that use UDP

- There are three types of applications that are best suited for UDP:
 - **Live video and multimedia applications** - These applications can tolerate some data loss, but require little or no delay. Examples include **VoIP** and **live streaming video**.
 - **Simple request and reply applications** - Applications with simple transactions where a host sends a request and may or may not receive a reply. Examples include **DNS (Domain Name System/Service)** and **DHCP (Dynamic Host configuration Protocol)**.
 - **Applications that handle reliability themselves** - Unidirectional communications where flow control, error detection, acknowledgments, and error recovery is not required, or can be handled by the application. Examples include **SNMP (Simple Network Management Protocol)** and **TFTP (Trivial File Transfer Protocol)**.



Although DNS and SNMP use UDP by default, both can also use TCP. DNS will use TCP if the DNS request or DNS response is more than 512 bytes, such as when a DNS response includes many name resolutions. Similarly, under some situations the network administrator may want to configure SNMP to use TCP.

Port Numbers-Multiple Separate Communications

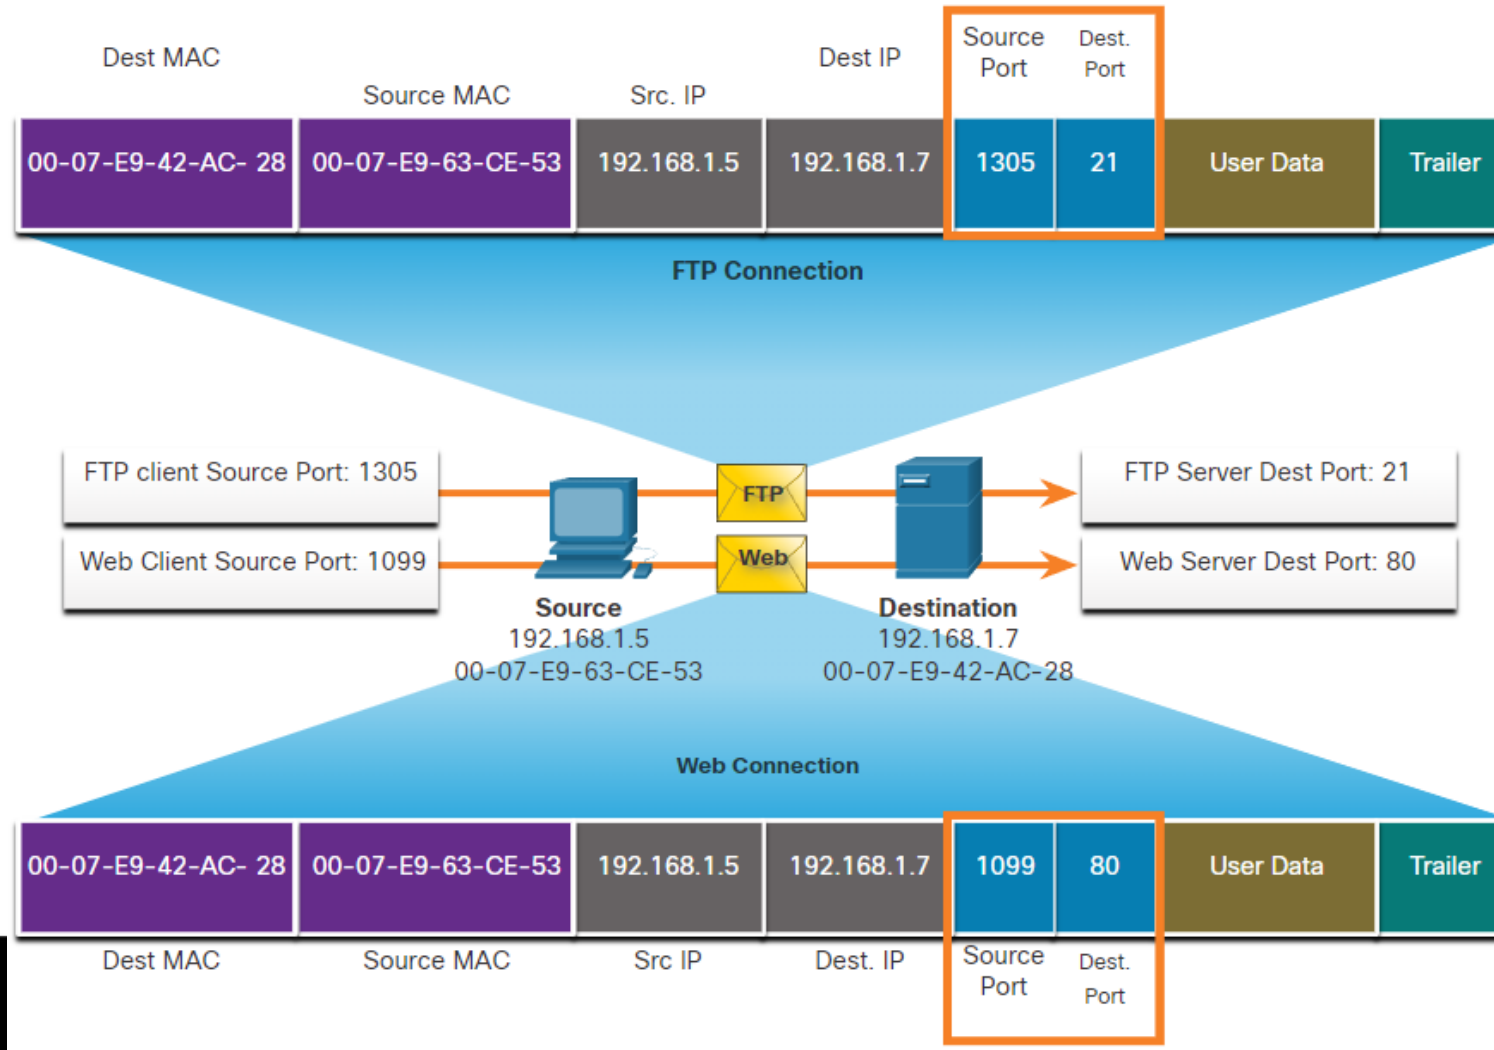
- No matter what type of data is being transported, both TCP and UDP use port numbers.
- The TCP and UDP transport layer protocols use port numbers to manage multiple, simultaneous conversations.
- The source port and destination port header fields are 2 bytes each



- The source port number is associated with the originating application on the local host whereas the destination port number is associated with the destination application on the remote host.
- For instance, assume a host is initiating a web page request from a web server. When the host initiates the web page request, the source port number is dynamically generated by the host to uniquely identify the conversation. Each request generated by a host will use a different dynamically created source port number. This process allows multiple conversations to occur simultaneously.
- In the request, the destination port number is what identifies the type of service being requested of the destination web server. For example, when a client specifies port 80 in the destination port, the server that receives the message knows that web services are being requested.
- A server can offer more than one service simultaneously such as web services on port 80 while it offers File Transfer Protocol (FTP) connection establishment on port 21.

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.



Socket Pairs

- In the example on the picture on the previous slide, 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 (i.e., dynamically generated by the host) 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. However, it is using the source port number 1099 (i.e., dynamically generated by the host) 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 might look like this, with 1099 representing the source port number: **192.168.1.5:1099**, and 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**
- 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.

Port Number Groups

- The Internet Assigned Numbers Authority (**IANA**) is the standards organization responsible for assigning various addressing standards, including the 16-bit port numbers. The 16 bits used to identify the source and destination port numbers provides a range of ports from 0 through 65535.

Port Group	Number Range	Description
Well-known Ports	0 to 1,023	<ul style="list-style-type: none">•These port numbers are reserved for common or popular services and applications such as web browsers, email clients, and remote access clients.•Defined well-known ports for common server applications enables clients to easily identify the associated service required.
Registered Ports	1,024 to 49,151	<ul style="list-style-type: none">•These port numbers are assigned by IANA to a requesting entity to use with specific processes or applications.•These processes are primarily individual applications that a user has chosen to install, rather than common applications that would receive a well-known port number.•For example, Cisco has registered port 1812 for its RADIUS server authentication process.
Private and/or Dynamic Ports	49,152 to 65,535	<ul style="list-style-type: none">•These ports are also known as <i>ephemeral ports</i>.•The client's OS usually assign port numbers dynamically when a connection to a service is initiated.•The dynamic port is then used to identify the client application during communication.

Well-Known Port Numbers

Port Number	Protocol	Application
20	TCP	File Transfer Protocol (FTP) - Data
21	TCP	File Transfer Protocol (FTP) - Control
22	TCP	Secure Shell (SSH)
23	TCP	Telnet
25	TCP	Simple Mail Transfer Protocol (SMTP)
53	UDP, TCP	Domain Name System (DNS)
67	UDP	Dynamic Host Configuration Protocol (DHCP) - Server
68	UDP	Dynamic Host Configuration Protocol - Client
69	UDP	Trivial File Transfer Protocol (TFTP)
80	TCP	Hypertext Transfer Protocol (HTTP)
110	TCP	Post Office Protocol version 3 (POP3)
143	TCP	Internet Message Access Protocol (IMAP)
161	UDP	Simple Network Management Protocol (SNMP)
443	TCP	Hypertext Transfer Protocol Secure (HTTPS)

Some applications may use both TCP and UDP. For example, **DNS uses UDP when clients send requests to a DNS server**. However, **communication between two DNS servers always uses TCP**.

Well-Known Port Numbers

Unexplained TCP connections can pose a major security threat. They can indicate that something or someone is connected to the local host. Sometimes it is necessary to know which active TCP connections are open and running on a networked host. **Netstat** is an important network utility that can be used to verify those connections.

By default, the **netstat** command will attempt to resolve IP addresses to domain names and port numbers to well-known applications. The **-n** option can be used to display IP addresses and port numbers in their numerical form.

```
C:\> netstat
```

```
Active Connections
```

Proto	Local Address	Foreign Address	State
TCP	192.168.1.124:3126	192.168.0.2:netbios-ssn	ESTABLISHED
TCP	192.168.1.124:3158	207.138.126.152:http	ESTABLISHED
TCP	192.168.1.124:3159	207.138.126.169:http	ESTABLISHED
TCP	192.168.1.124:3160	207.138.126.169:http	ESTABLISHED
TCP	192.168.1.124:3161	sc.msn.com:http	ESTABLISHED
TCP	192.168.1.124:3166	www.cisco.com:http	ESTABLISHED

```
(output omitted)
```

```
C:\>
```

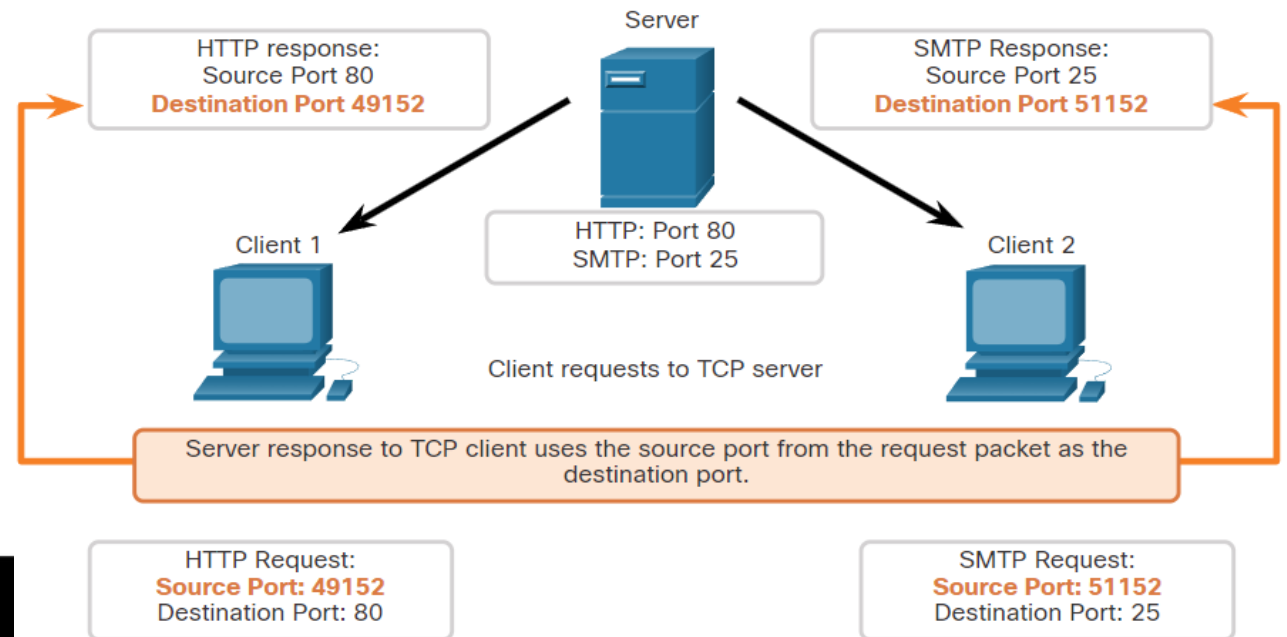
- netstat
- netstat -n

TCP Communication Process

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.

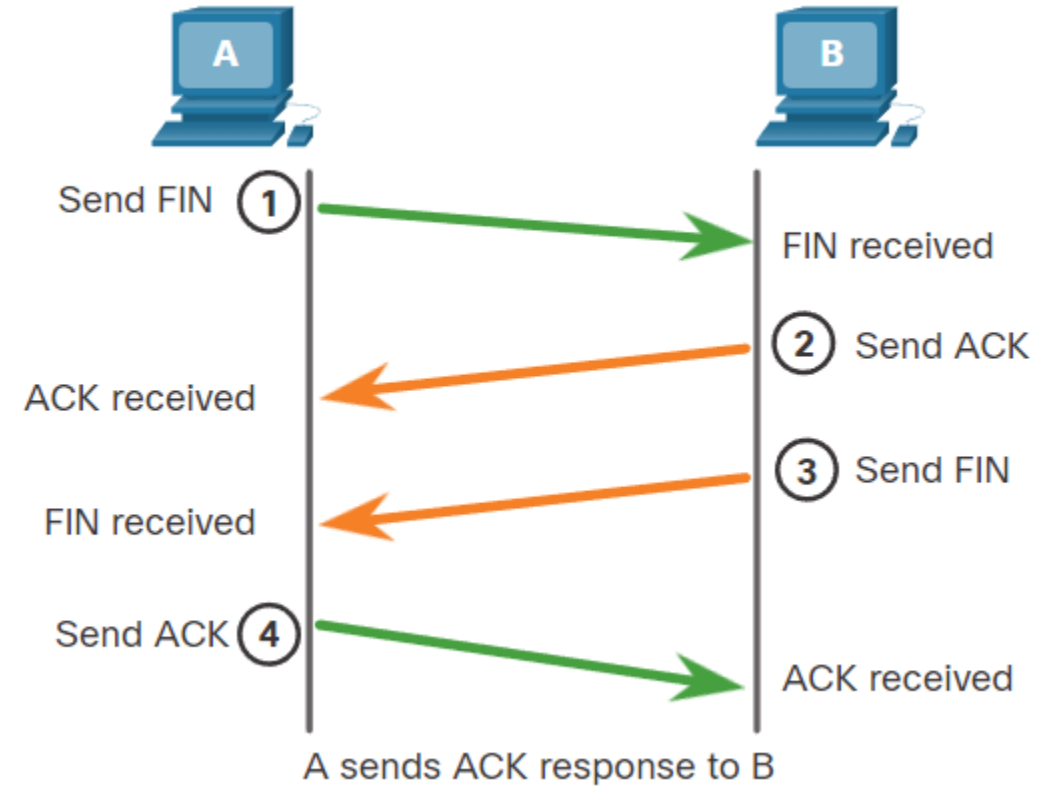
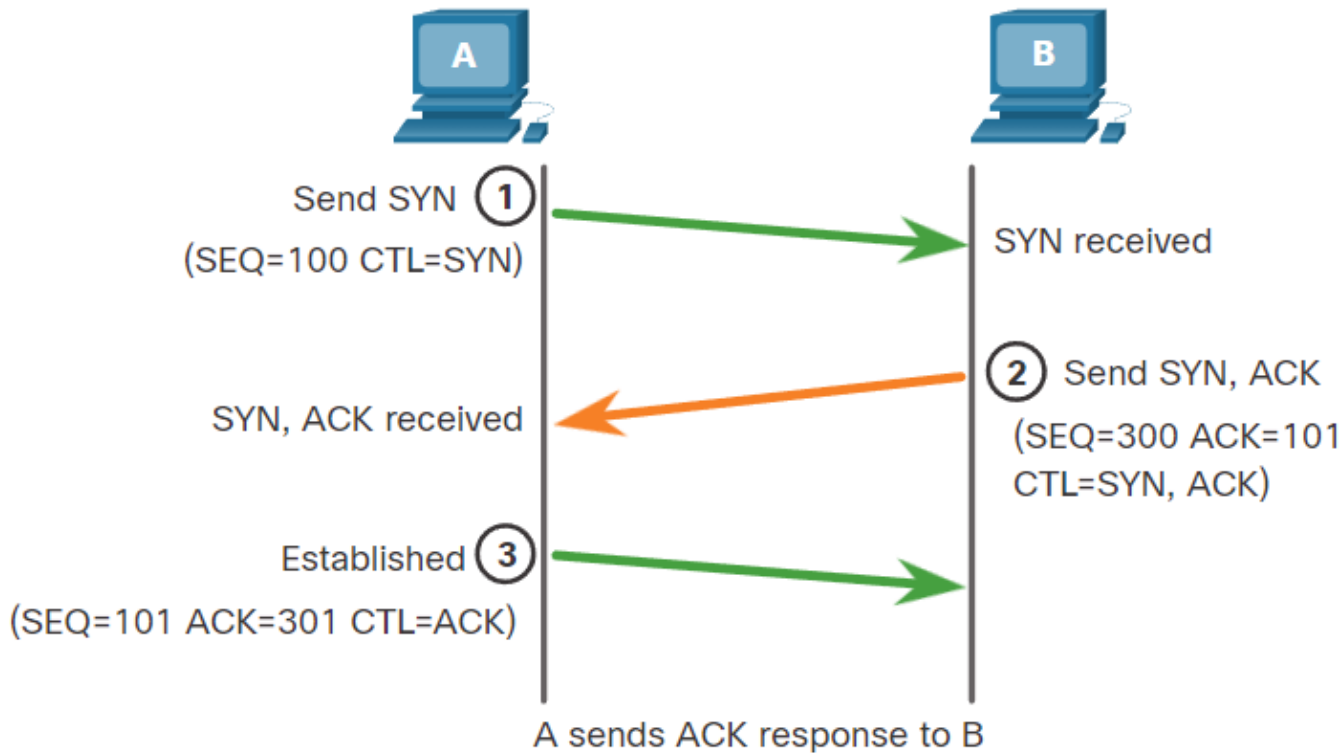
An individual **server cannot have two services assigned to the same port number** within the same transport layer services. For example, a host running a web server application and a file transfer application cannot have both configured to use the same port, such as TCP port 80.

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.



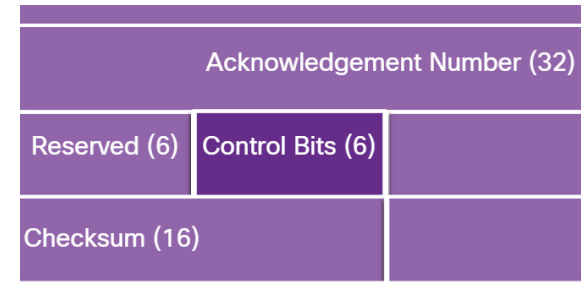
TCP Connection Establishment and termination

three-way handshake process



TCP Connection Establishment and termination

Hosts maintain state, track each data segment within a session, and exchange information about what 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.** Control bits in the TCP header indicate the progress and status of the connection.

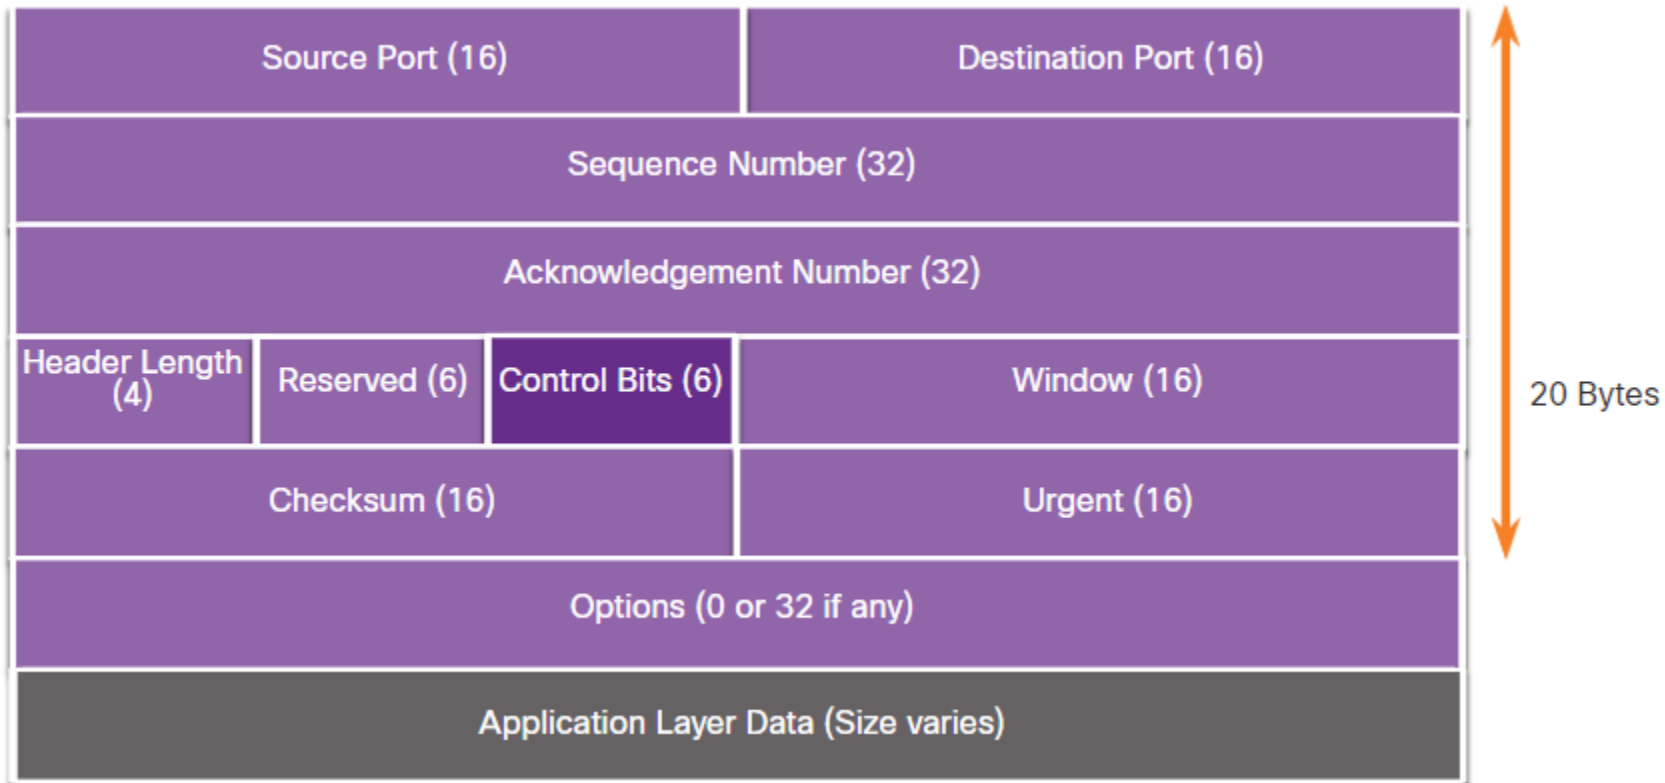


These are the functions of the three-way handshake:

- 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 Connection Establishment and termination



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

TCP Connection Establishment and termination

three-way handshake process

192.168.43.223	104.19.141.57	TCP	66	49441+80	[SYN] Seq=0	win=8192 Len
104.19.141.57	192.168.43.223	TCP	66	80+49441	[SYN, ACK] Seq=0 Ack=1 w	
192.168.43.223	104.19.141.57	TCP	54	49441+80	[ACK] Seq=1	Ack=1 win=66

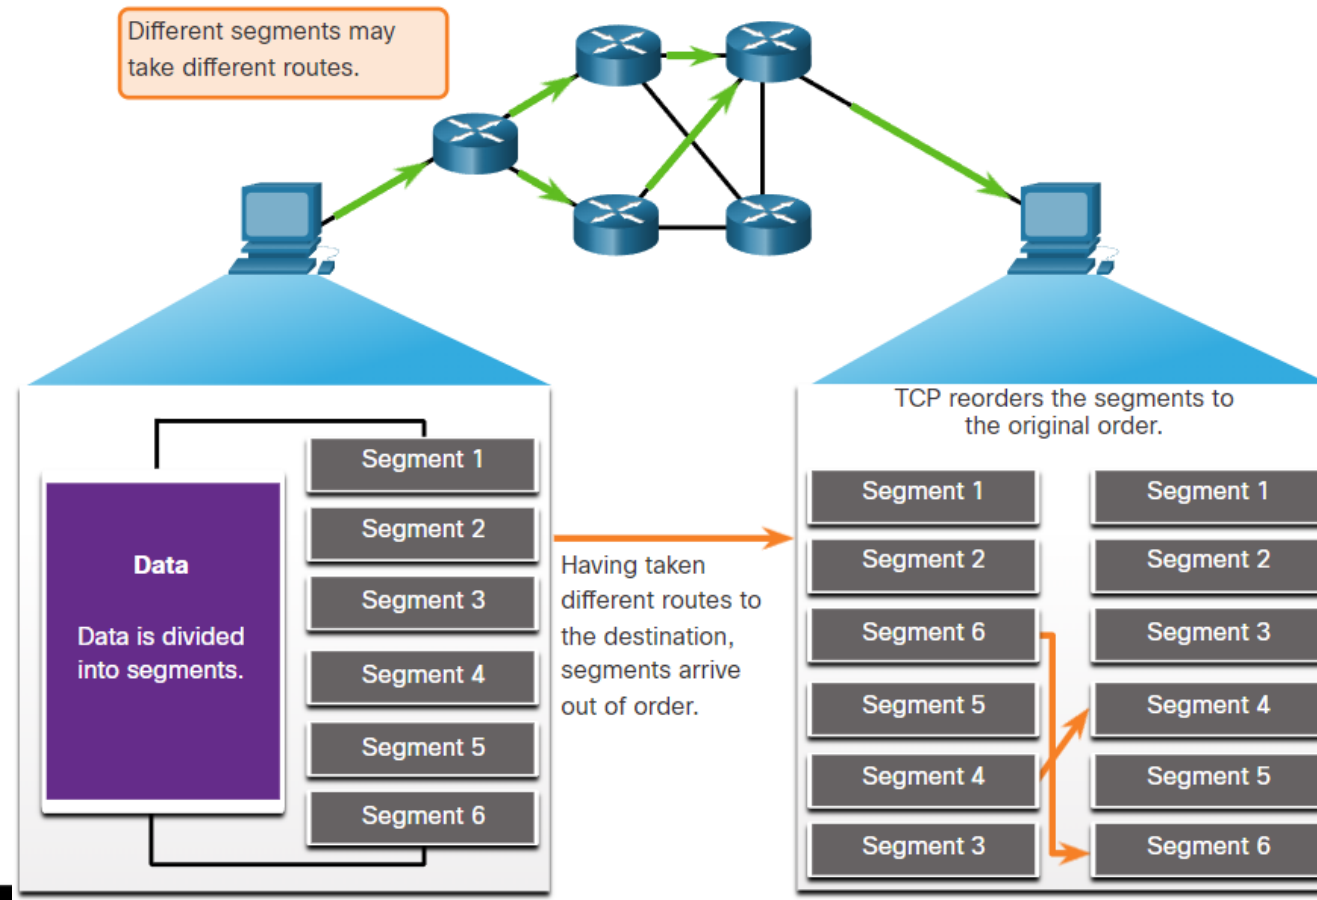
Session termination

This ACK is for some of the previous segments

192.168.43.223	104.19.141.57	TCP	54	49607+80	[FIN, ACK]	Seq
104.19.141.57	192.168.43.223	TCP	54	80+49607	[FIN, ACK]	Seq
192.168.43.223	104.19.141.57	TCP	54	49607+80	[ACK] Seq=2	A

TCP Reliability - Guaranteed and Ordered Delivery

During session setup, an initial sequence number (ISN) is set. This ISN 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.



TCP Reliability - Data Loss and Retransmission

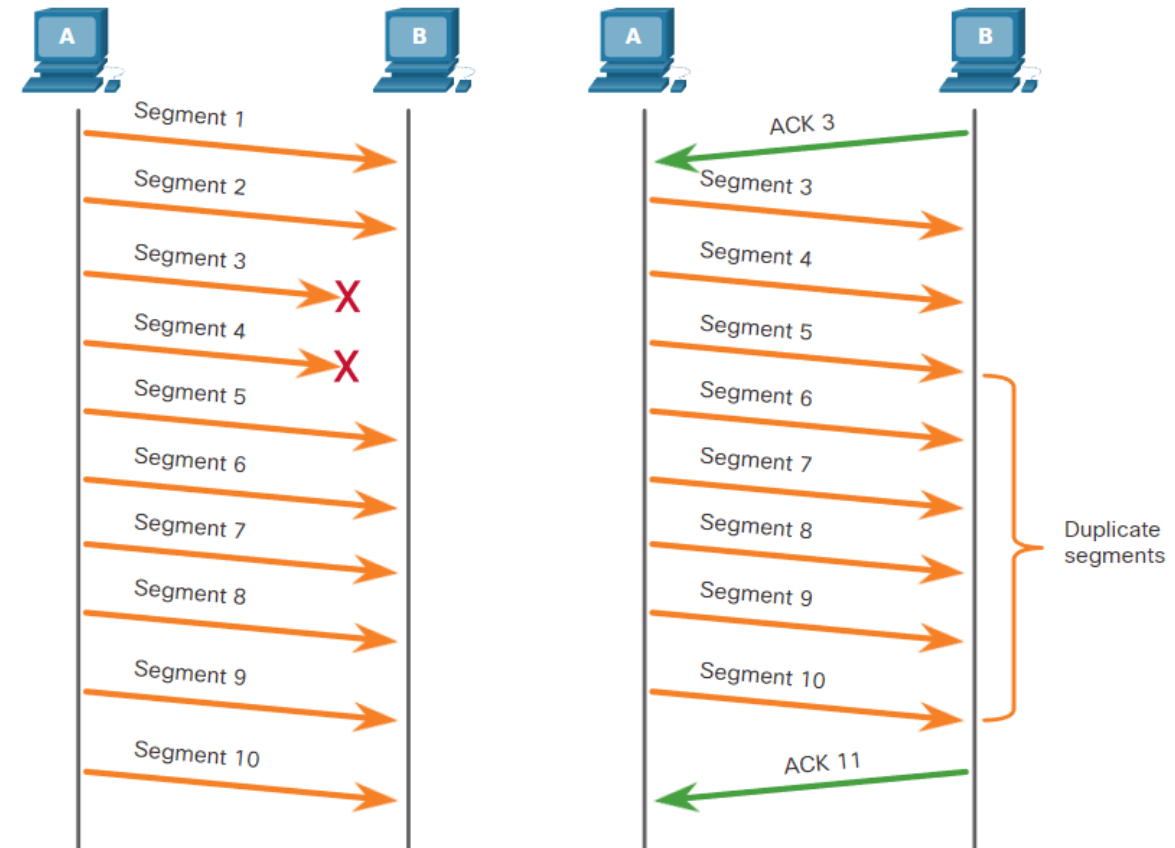
No matter how well designed a network is, data loss occasionally occurs. TCP provides methods of managing these segment losses. Among these is a mechanism to retransmit segments for unacknowledged data.

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**.

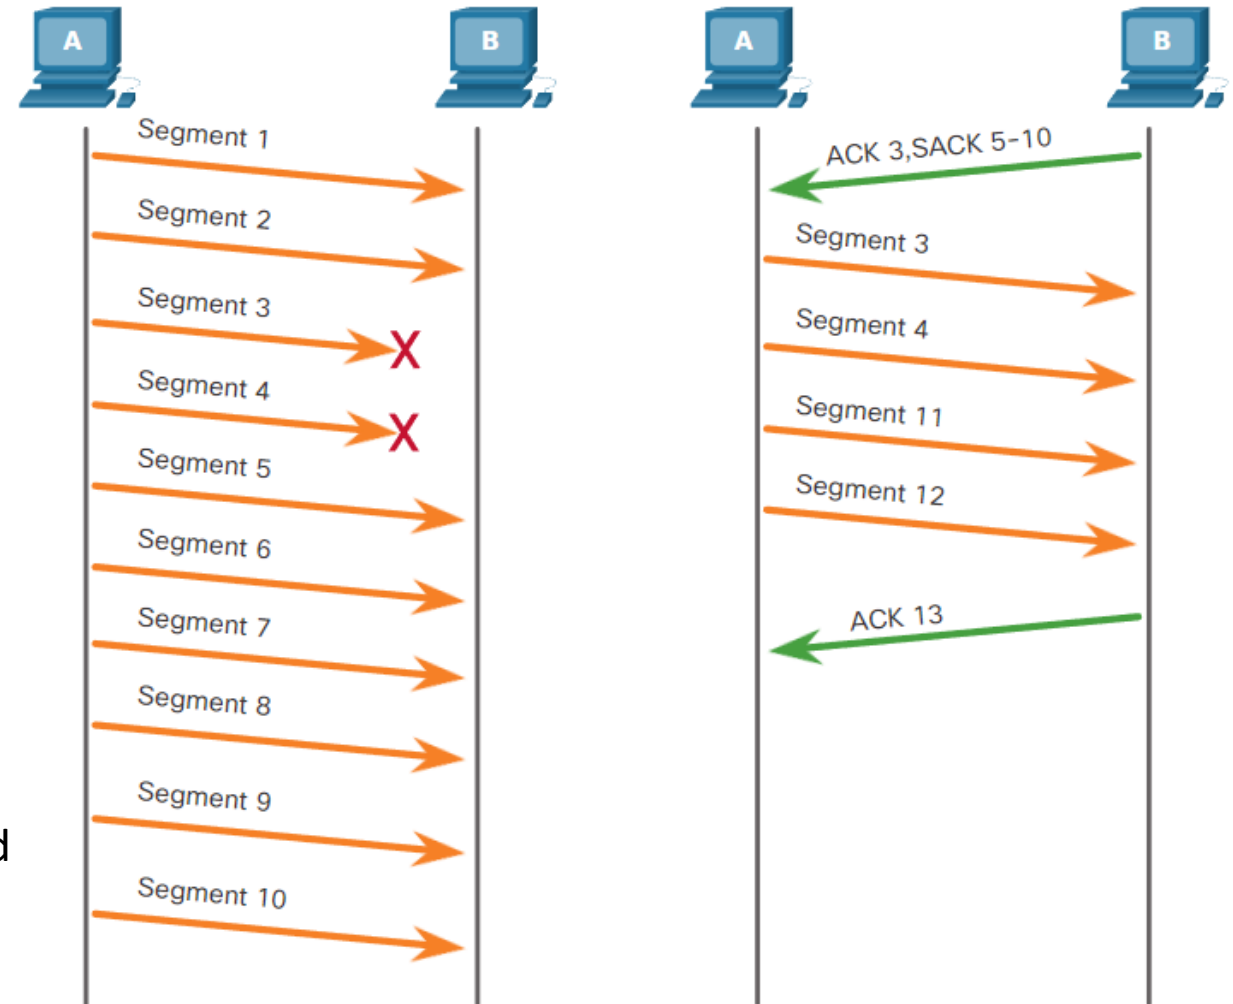
TCP Reliability - Data Loss and Retransmission

- Prior to later enhancements, TCP could only acknowledge the next byte expected.
- For example, host A sends segments 1 through 10 to host B. If all the segments arrive except for 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. Host A would, therefore, 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.



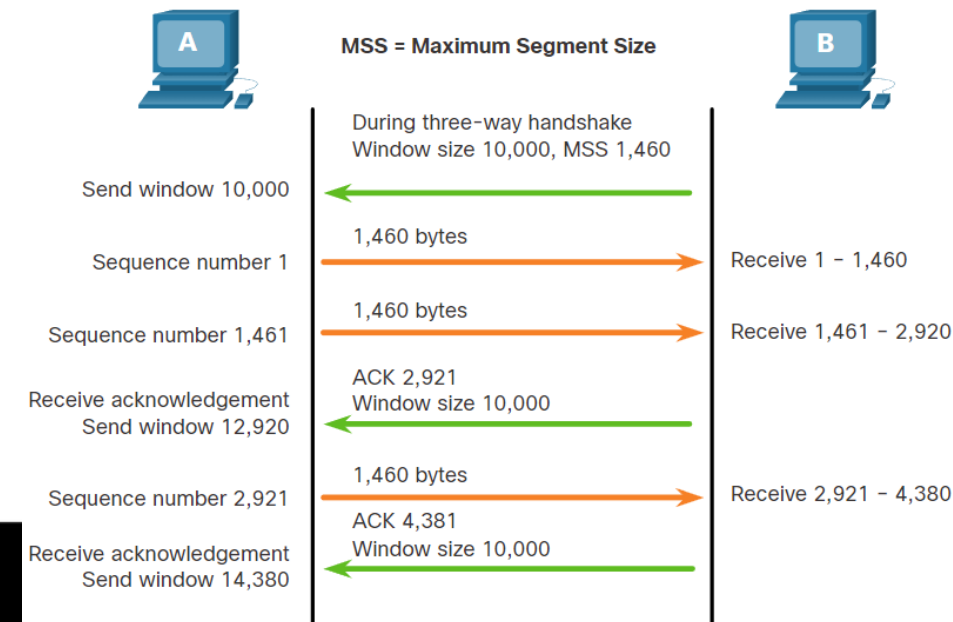
TCP Reliability - Data Loss and Retransmission

- Host operating systems today typically employ an optional TCP feature called selective acknowledgment (SACK), negotiated during the three-way handshake.
- If both hosts support SACK, the receiver can explicitly acknowledge which segments (bytes) were received including any discontinuous segments.
- The sending host would therefore only need to retransmit the missing data.
- For example, in the next figure, again using segment numbers for simplicity, 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.
- TCP uses timers to know how long to wait before resending a segment



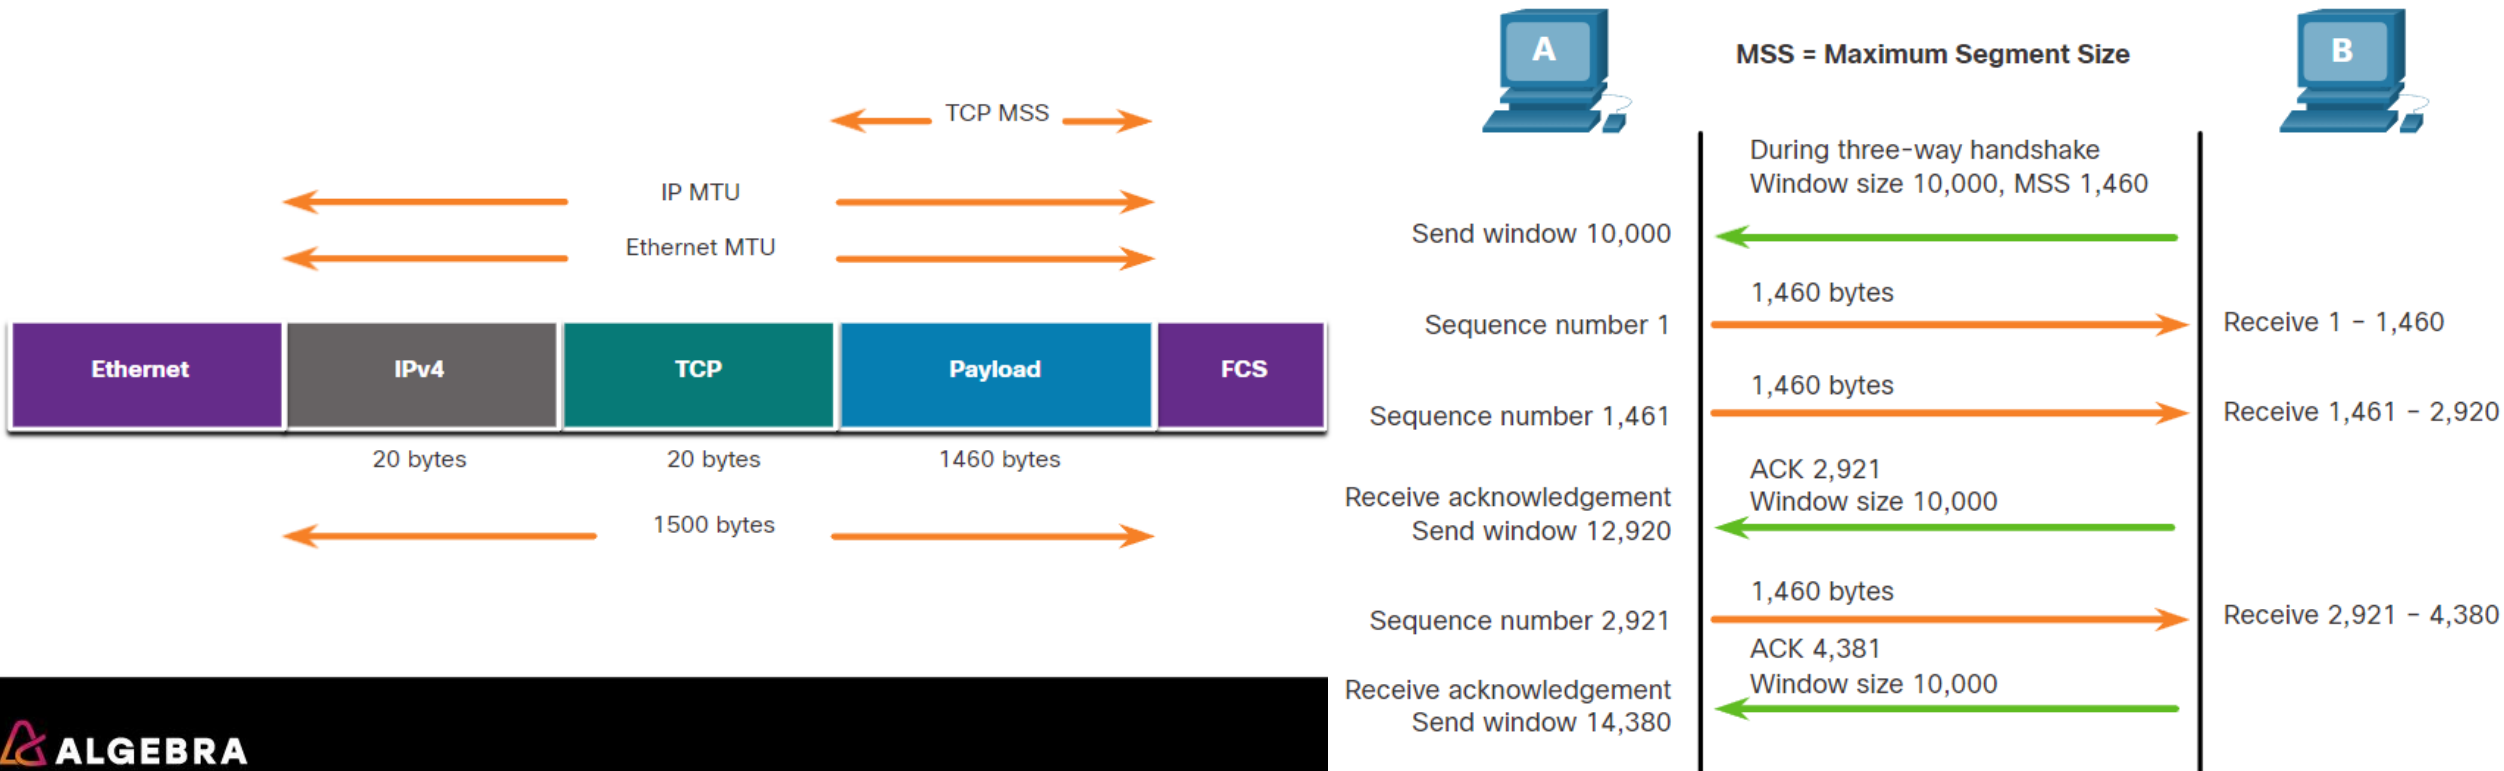
TCP Flow Control - Window Size and Acknowledgments

- TCP also provides mechanisms for **flow control**. 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 determines the number of bytes that can be sent before expecting an acknowledgment.
- The window size is the number of bytes that the destination device of a TCP session can accept and process at one time.
- In this example, the PC B initial window size for the TCP session is 10,000 bytes. Starting with the first byte, byte number 1, the last byte PC A can send without receiving an acknowledgment is byte 10,000. This is known as the **send window of PC A**. 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 device receives an acknowledgment that the bytes have been received, can it continue sending more data for the session.
- A destination sending acknowledgments as it processes bytes received, and the **continual adjustment of the source send window**, is known as **sliding windows**.



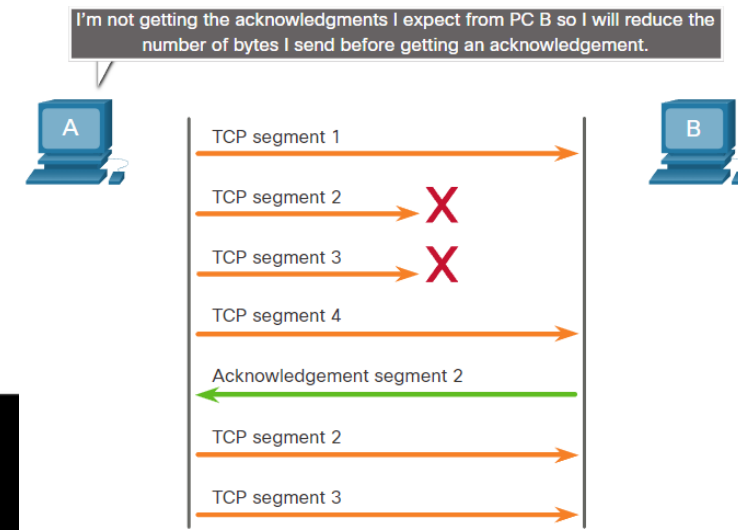
TCP Flow Control - Maximum Segment Size (MSS)

- In the figure, the source is transmitting **1,460 bytes of data within each TCP segment**. This is typically 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 typically included during the three-way handshake.
- 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.
- 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 illustrated 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.

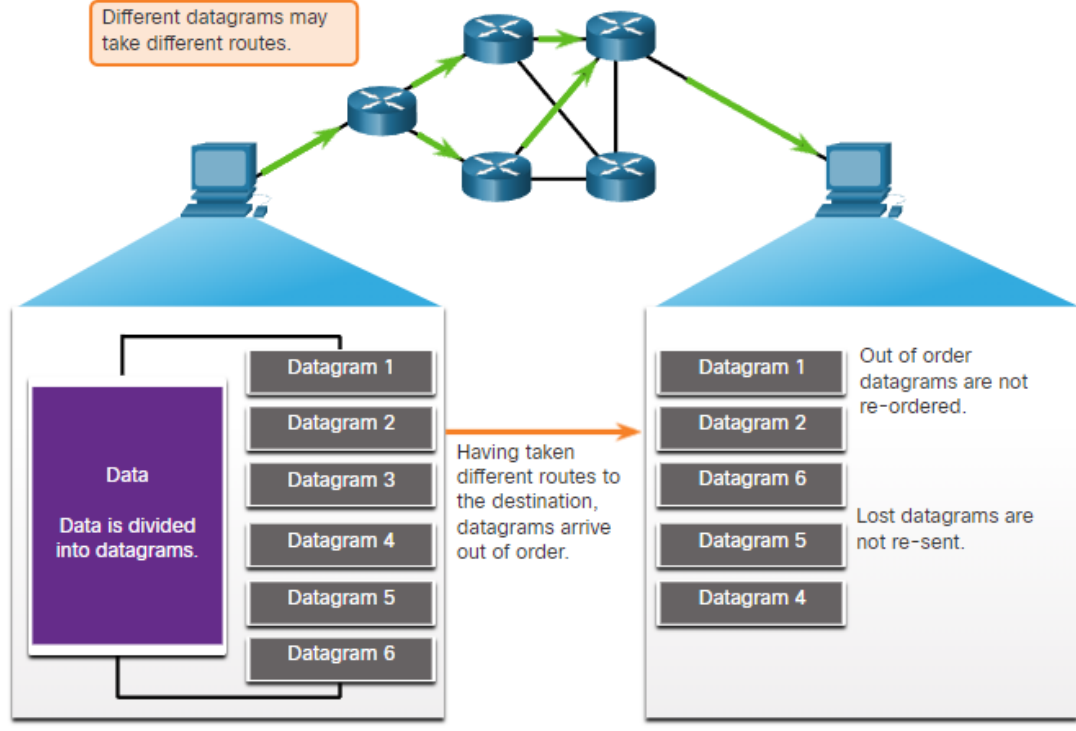


UDP Communication



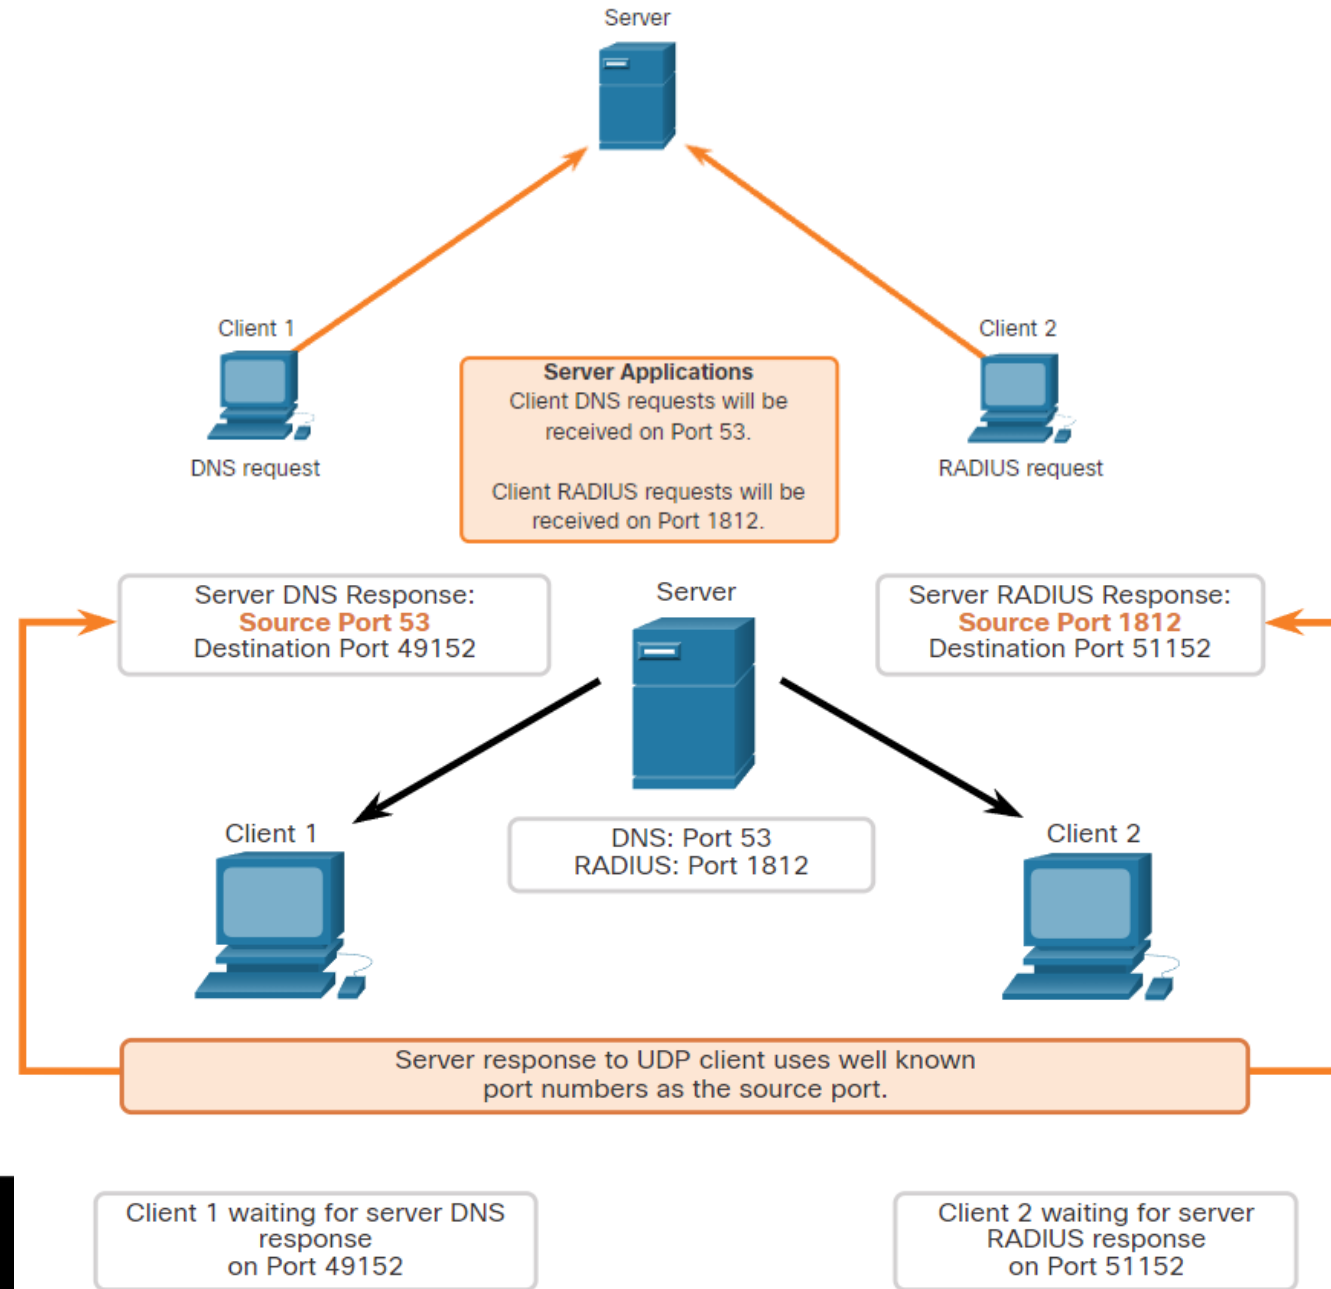
UDP does not establish a connection before sending data.

Different datagrams may take different routes.



- UDP does not establish a connection. UDP provides low overhead data transport because it has a small datagram header and no network management traffic
- Like segments with TCP, when UDP datagrams are sent to a destination, they often take different paths and arrive in the wrong order. UDP does not track sequence numbers the way TCP does. UDP has no way to reorder the datagrams into their transmission order
- Therefore, UDP simply reassembles the data in the order that it was received and forwards it to the application. If the data sequence is important to the application, the application must identify the proper sequence and determine how the data should be processed.

UDP Server and Client Processes and Requests



- Like TCP-based applications, UDP-based server applications are assigned **well-known or registered port numbers**, as shown in the figure. When these applications or processes are running on a server, they accept the data matched with the assigned port number. When UDP receives a datagram destined for one of these ports, **it forwards the application data to the appropriate application based on its port number.**
 - As with TCP, **client-server communication is initiated by a client application that requests data from a server process.** The UDP client process dynamically selects a port number from the range of port numbers and uses this as the source port for the conversation. The **destination port is usually the well-known or registered port number assigned to the server process.**
- After a client has selected the source and destination ports, the same pair of ports are used in the header of all datagrams in the transaction. **For the data returning to the client from the server, the source and destination port numbers in the datagram header are reversed.**

