

TCP and UDP

Layer 3 of the TCP/IP protocol stack. Transport layer

I. Introduction.

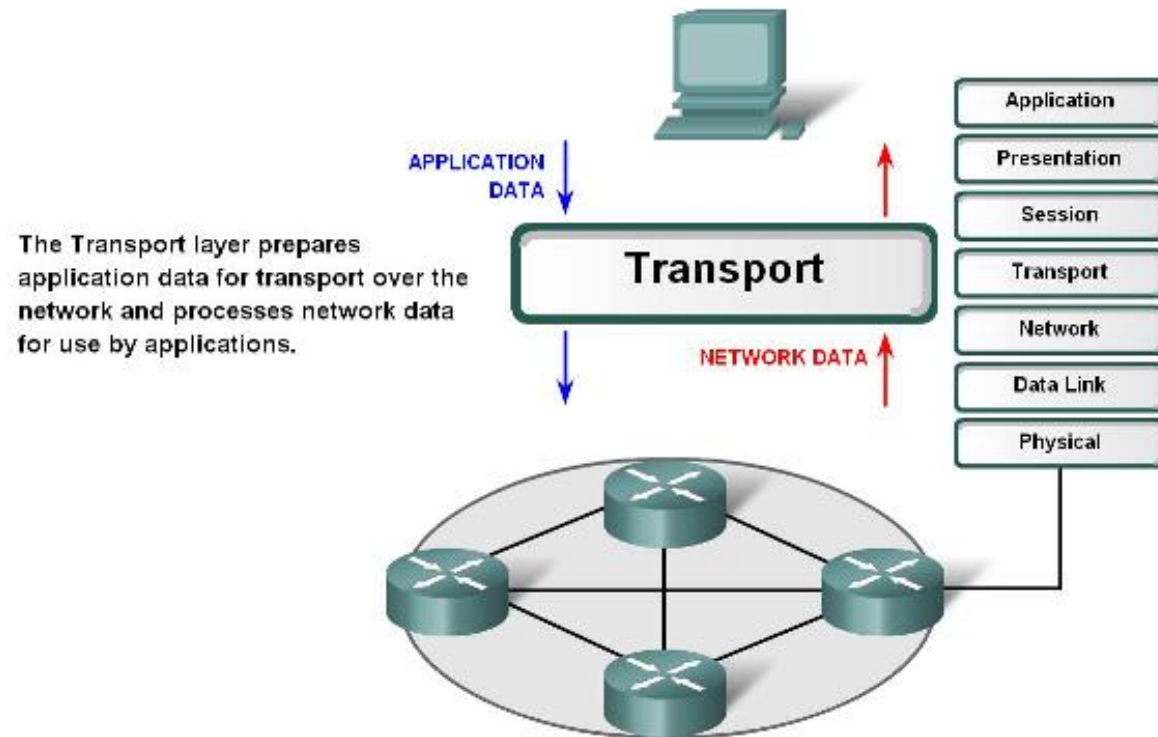
On a single device, people can use multiple services such as e-mail, the web, and instant messaging to send messages or retrieve information. Applications such as e-mail clients, web browsers, and instant messaging clients allow people to use computers and networks to send messages and find information.

Data from each of these applications is packaged, transported, and delivered to the appropriate server daemon or application on the destination device. The processes described in the OSI Transport layer accept data from the Application layer and prepare it for addressing at the Network layer. **The Transport layer is responsible for the overall end-to-end transfer of application data.** The role of the Transport layer is encapsulating application data for use by the Network layer.

The Transport layer also encompasses these functions:

- a). Enables multiple applications to communicate over the network at the same time on a single device*
- b). Ensures that, if required, all the data is received reliably and in order by the correct application*
- c). Employs error handling mechanisms*

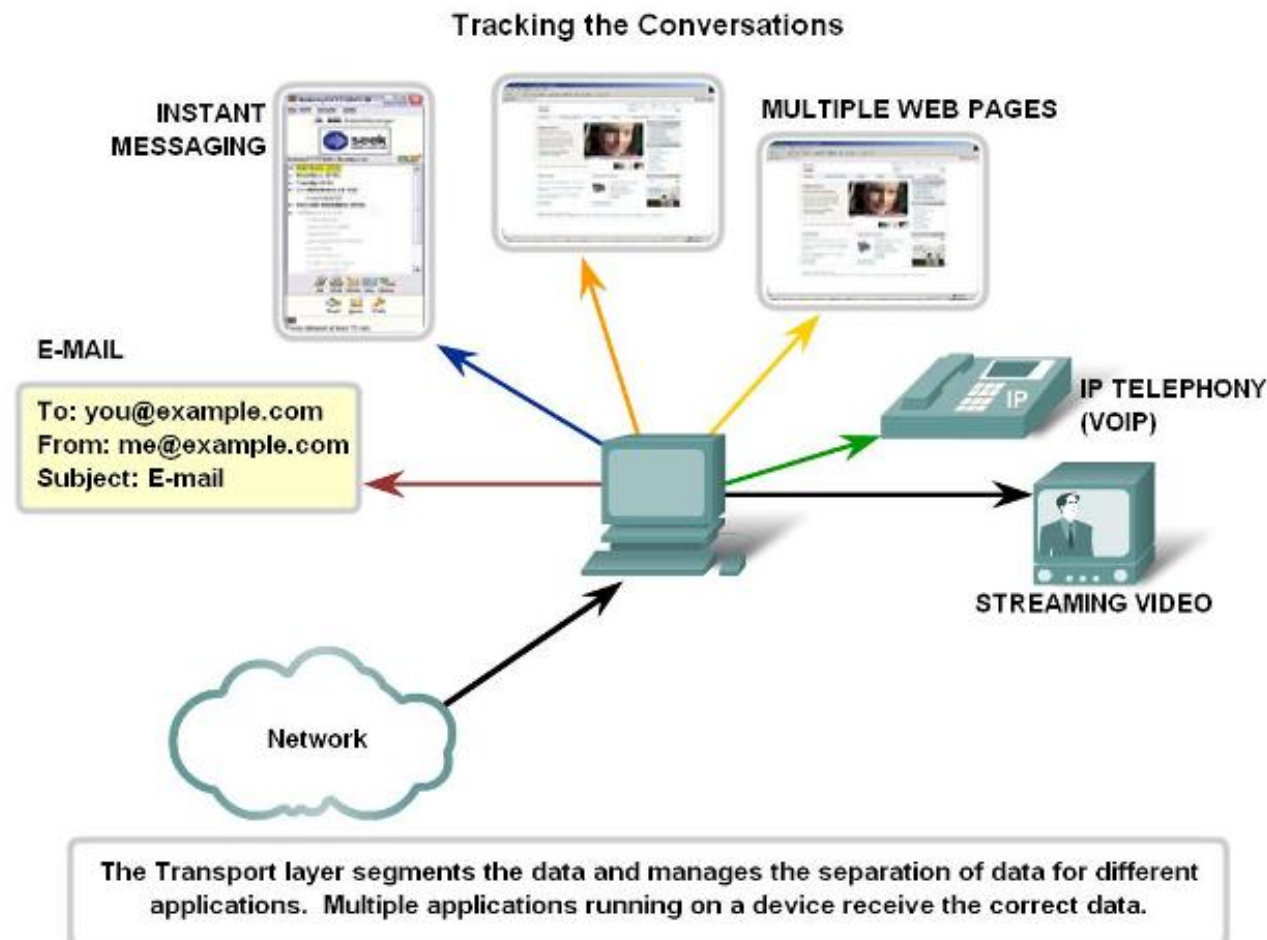
OSI – Transport Layer



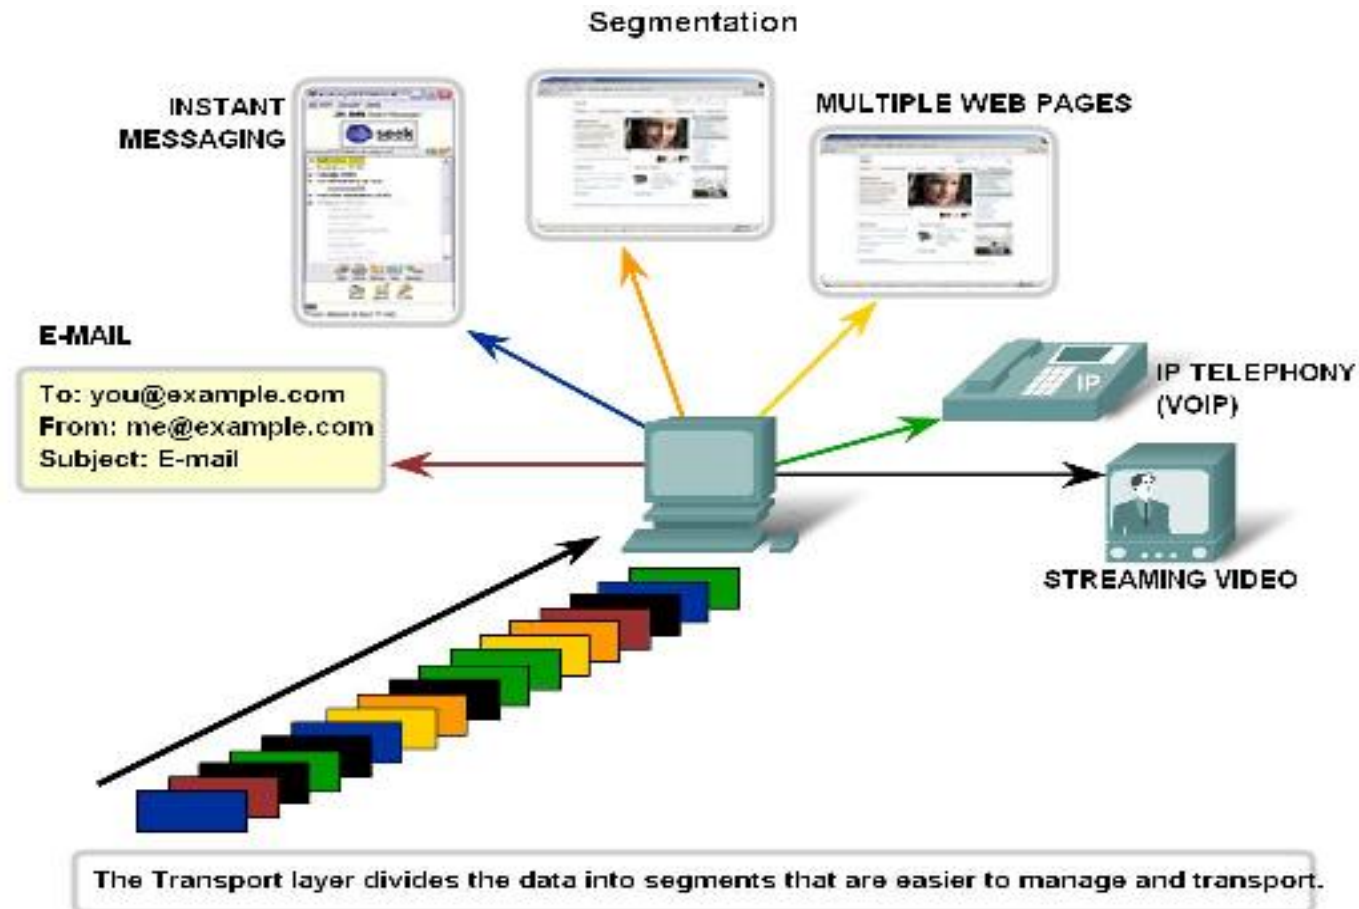
II. Purpose of Transport Layer

1. Tracking the individual communication between applications on the source and destination hosts:

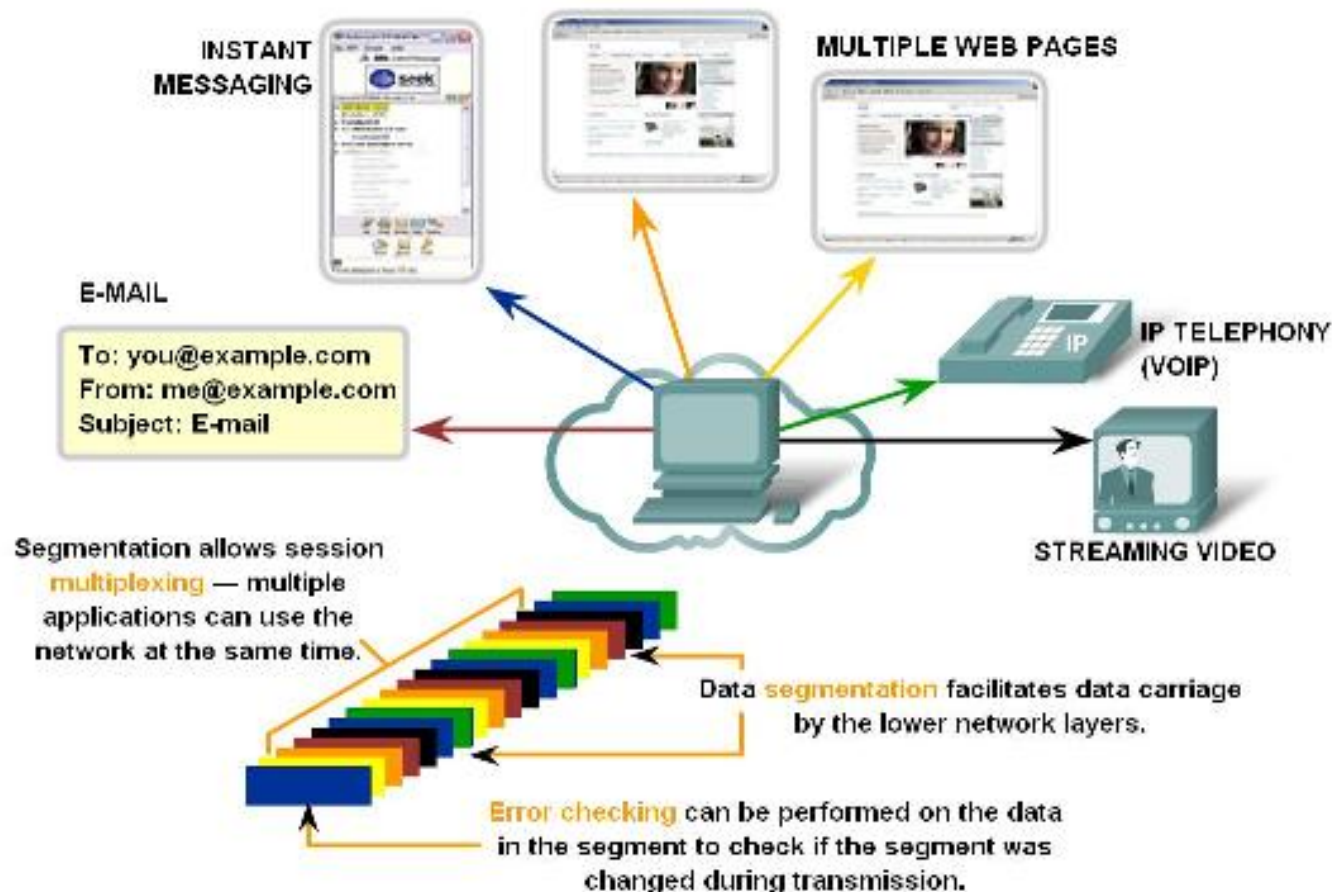
Any host may have multiple applications that are communicating across the network. Each of these applications will be communicating with one or more applications on remote hosts. **It is the responsibility of the Transport layer to maintain the multiple communication streams between these applications.**



2) Segmenting data and managing each piece: As each application creates a stream data to be sent to a remote application, this data must be prepared to be sent across the media in manageable pieces. The Transport layer protocols describe services that segment this data from the Application layer. This includes the encapsulation required on each piece of data. Each piece of application data requires headers to be added at the Transport layer to indicate to which communication it is associated.



3) Reassembling the segments into streams of application data: At the receiving host, each piece of data may be directed to the appropriate application. Additionally, these individual pieces of data must also be reconstructed into a complete data stream that is useful to the Application layer. The protocols at the Transport layer describe the how the Transport layer header information is used to reassemble the data pieces into streams to be passed to the Application layer.



4) Identifying the different applications: In order to pass data streams to the proper applications, the Transport layer must identify the target application. To accomplish this, the Transport layer assigns an application an identifier. The TCP/IP protocols call this identifier a port number. Each software process that needs to access the network is assigned a port number unique in that host. This port number is used in the transport layer header to indicate to which application that piece of data is associated.

III. Reliable Communication

Different applications have different requirements for their data, and therefore different Transport protocols have been developed to meet these requirements. A Transport layer protocol can implement is a method to ensure reliable delivery of the data. In networking terms, **reliability means ensuring that each piece of data that the source sends arrives at the destination**. At the Transport layer the three basic operations of reliability are:

- a) tracking transmitted data
- b) acknowledging received data
- c) retransmitting any unacknowledged data

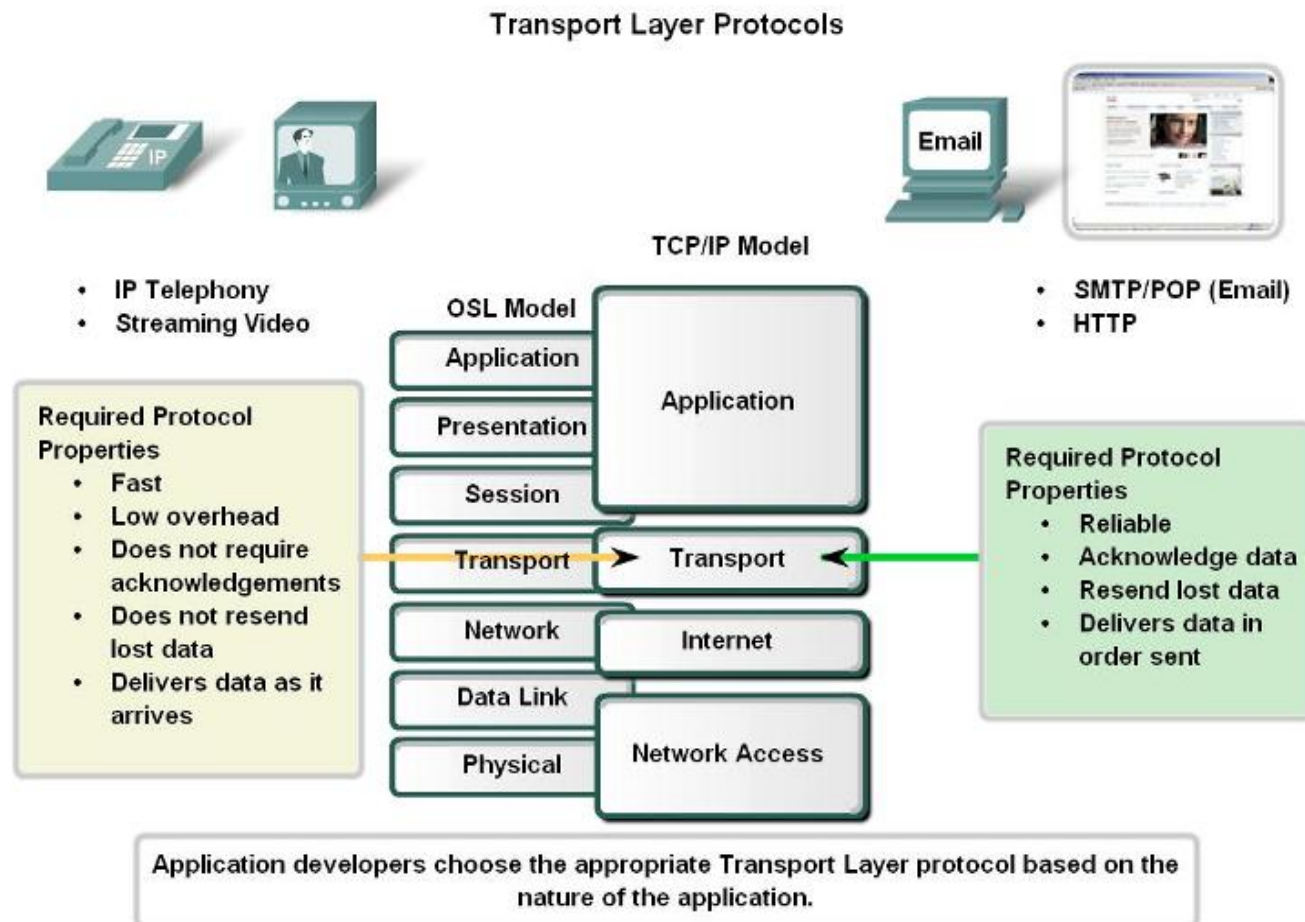
This requires the processes of Transport layer of the source to keep track of all the data pieces of each conversation and the retransmit any of data that did were not acknowledged by the destination. The Transport layer of the receiving host must also track the data as it is received and acknowledge the receipt of the data.

These reliability processes place additional overhead on the network resources due to the acknowledgement, tracking, and retransmission. To support these reliability operations, more control data is exchanged between the sending and receiving hosts. This control information is contained in the Transport Layer header.

2. Determining the Need for Reliability

Applications, such as databases, web pages, and e-mail, require that all of the sent data arrive at the destination in its original condition, in order for the data to be useful. Any missing data could cause a corrupt communication that is either incomplete or unreadable. Therefore, these applications are designed to use a Transport layer protocol that implements reliability. The additional network overhead is considered to be required for these applications.

Other applications are more tolerant of the loss of small amounts of data. For example, if one or two segments of a video stream fail to arrive, it would only create a momentary disruption in the stream. This may appear as distortion in the image but may not even be noticeable to the user.



IV. TCP and UDP Protocols

The two most common Transport layer protocols of TCP/IP protocol suite are Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). Both protocols manage the communication of multiple applications. The differences between the two are the specific functions that each protocol implements.

User Datagram Protocol (UDP)

UDP is a simple, connectionless protocol, described in RFC 768. It has the advantage of providing for low overhead data delivery. The pieces of communication in UDP are called datagrams. These datagrams are sent as "best effort" by this Transport layer protocol.

Applications that use UDP include:

- Domain Name System (DNS)
- Video Streaming
- Voice over IP (VoIP)

Transmission Control Protocol (TCP)

TCP is a connection-oriented protocol, described in RFC 793. TCP incurs additional overhead to gain functions. Additional functions specified by TCP are the same order delivery, reliable delivery, and flow control. Each TCP segment has 20 bytes of overhead in the header encapsulating the Application layer data, whereas each UDP segment only has 8 bytes of overhead.

Applications that use TCP are:

- Web Browsers
- E-mail
- File Transfers

V. Transmission Control Protocol (TCP)

The reliability of TCP communication is performed using connection-oriented sessions. Before a host using TCP sends data to another host, the Transport layer initiates a process to create a connection with the destination.

- * Establishes a session between source host and source destination (this ensures that each host is prepared and aware for the connection).
- * The destination host sends acknowledgements to the source for the segments that it receives.
- * As the source receives an acknowledgement, it knows that the data has been successfully delivered and can quit tracking that data.
- * If the source does not receive an acknowledgement within a predetermined amount of time, it retransmits that data to the destination.
- * The establishment of the sessions creates overhead in the form of additional segments being exchanged.
- * There is also additional overhead on the individual hosts created by the necessity to keep track of which segments are awaiting acknowledgement and by the retransmission process.

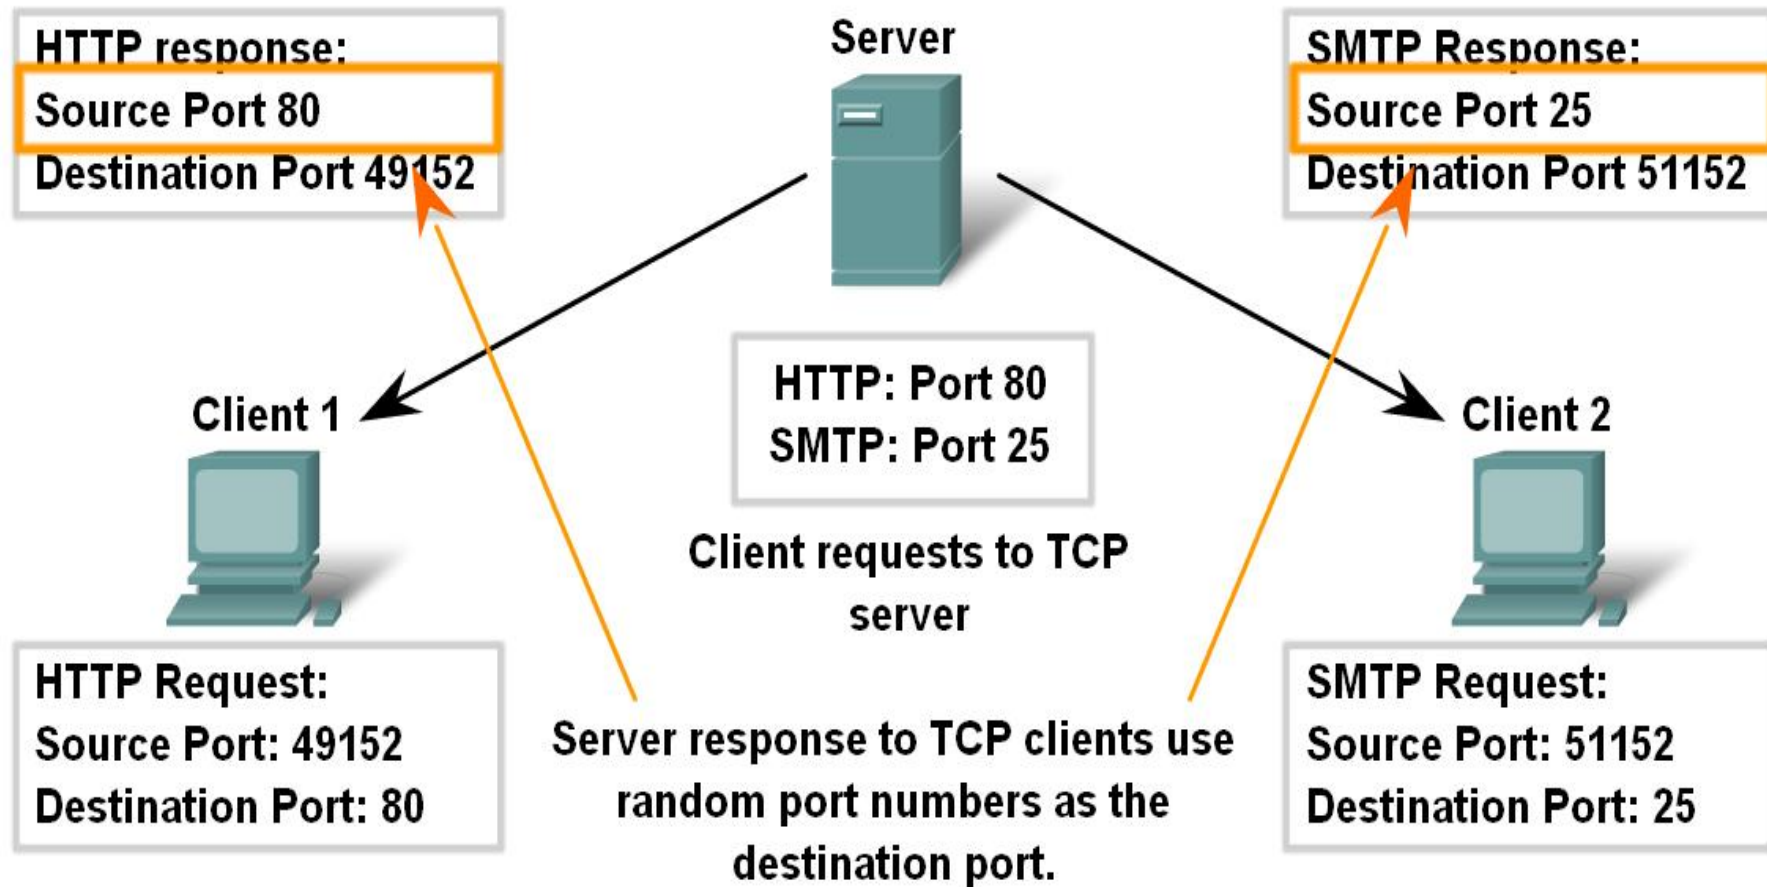
TCP Segment Structure

TCP Header

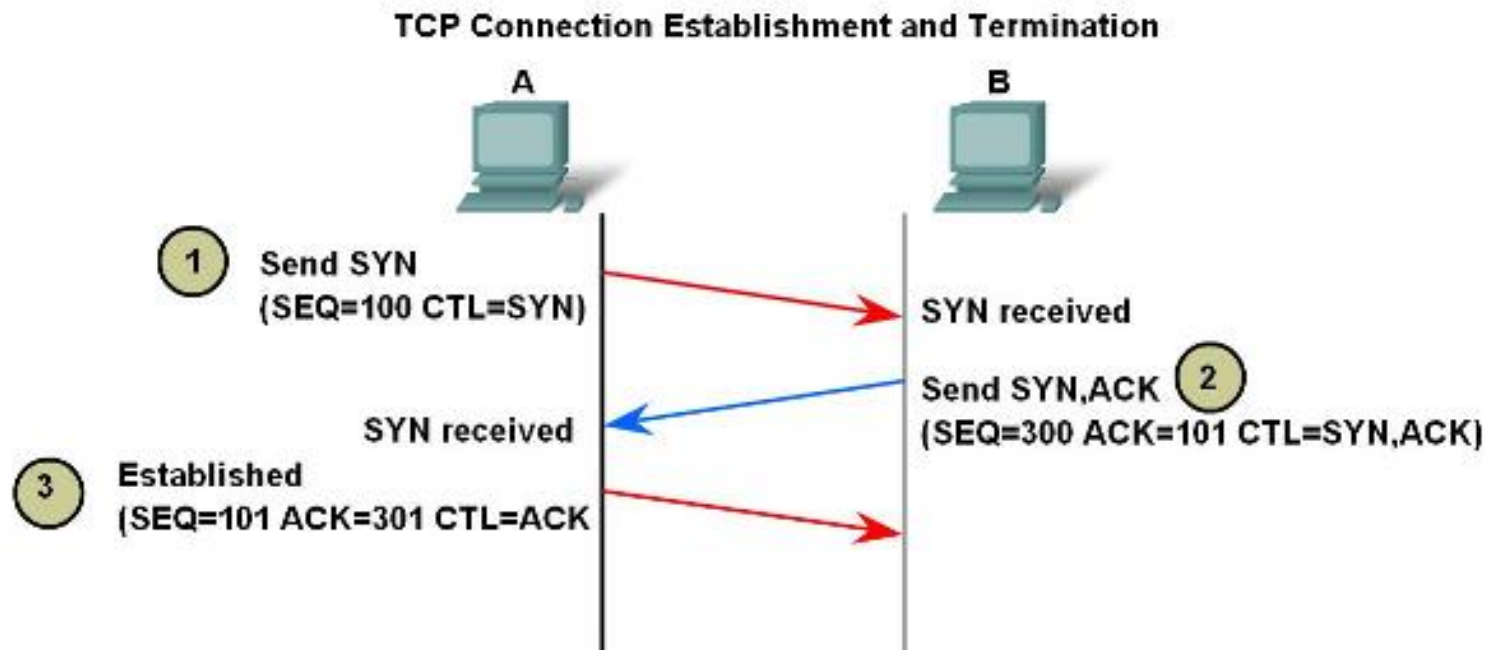
Bit offset	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	Source port																Destination port															
32	Sequence number																															
64	Acknowledgment number																															
96	Data offset				Reserved				C W R	E C E	U R G	A C K	P S H	R S T	S Y N	F I N	Window Size															
128	Checksum																Urgent pointer															
160	Options (if Data Offset > 5)																															
...	...																															

- * Source port (16 bits) – identifies the sending port
- * Destination port (16 bits) – identifies the receiving port
- * Sequence number (32 bits) – has a dual role
 - * If the SYN flag is set, then this is the initial sequence number. The sequence number of the actual first data byte (and the acknowledged number in the corresponding ACK) will then be this sequence number plus 1.
 - * If the SYN flag is clear, then this is the sequence number of the first data byte
- * Acknowledgment number (32 bits)
- * Data offset (4 bits) – specifies the size of the TCP header in 32-bit words
- * Reserved (4 bits) – for future use and should be set to zero
- * Flags (8 bits) (aka Control bits) – contains 8 1-bit flags
- * Window (16 bits) – the size of the receive window, which specifies the number of bytes that the receiver is currently willing to receive.
- * Checksum (16 bits) – The 16-bit checksum field is used for error-checking of the header and data
- * Urgent pointer (16 bits) – if the URG flag is set, then this 16-bit field is an offset from the sequence number indicating the last urgent data byte.

Clients Sending TCP Requests



TCP Connection Establishment

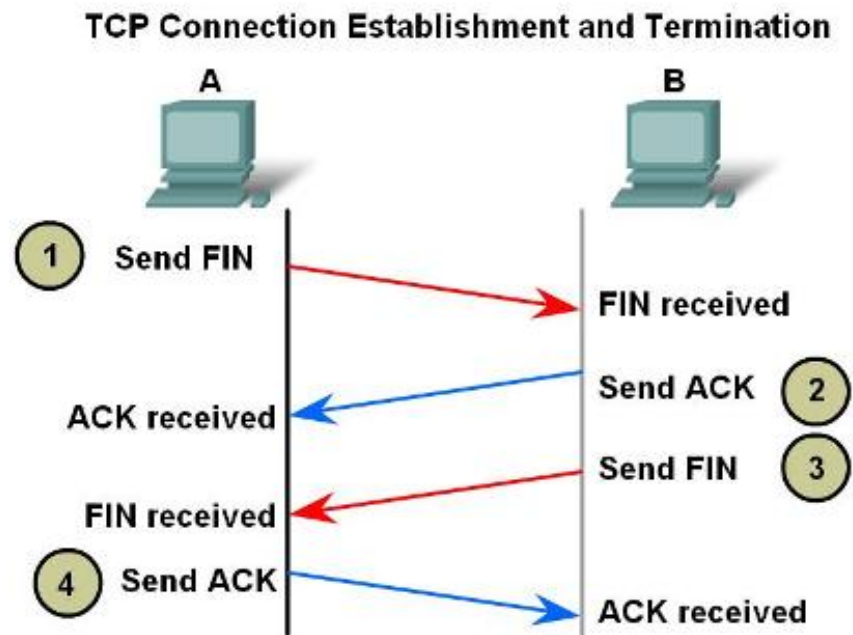


A sends SYN request to B

B sends ACK response and SYN request to A

A sends ACK response to B

TCP Connection Termination



A sends FIN request to B

B sends ACK response to A

B sends FIN request to A

A sends ACK response to B

TCP 3-way Handshake (SYN)

13	6.201109	192.168.254.254	10.1.1.1	DNS	Standard query r
14	6.202100	10.1.1.1	192.168.254.254	TCP	1069 > http [SYN
15	6.202513	192.168.254.254	10.1.1.1	TCP	http > 1069 [SYN
16	6.202543	10.1.1.1	192.168.254.254	TCP	1069 > http [ACK
17	6.202651	10.1.1.1	192.168.254.254	HTTP	GET / HTTP/1.1

⊕	Frame 14 (62 bytes on wire, 62 bytes captured)
⊕	Ethernet II, Src: QuantaCo_bd:0c:7c (00:c0:9f:bd:0c:7c), Dst: Cisco_cf:66:40
⊕	Internet Protocol, Src: 10.1.1.1 (10.1.1.1), Dst: 192.168.254.254 (192.168.2
⊖	Transmission Control Protocol, Src Port: 1069 (1069), Dst Port: http (80), s
	Source port: 1069 (1069)
	Destination port: http (80)
	Sequence number: 0 (relative sequence number)
	Header length: 28 bytes
⊖	Flags: 0x02 (SYN)
	0... = Congestion window Reduced (CWR): Not set
	.0.. = ECN-Echo: Not set
	..0. = Urgent: Not set
	0... = Acknowledgment: Not set

Protocol Analyzer shows initial client request for session in frame 14

TCP segment in this frame shows:

- SYN flag set to validate an initial Sequence number
- Randomized sequence number valid (relative value is 0)
- Random source port 1069
- Well known destination port is 80 (HTTP port) indicates web server (httpd)

TCP 3-way Handshake (SYN, ACK)

13	6.201109	192.168.254.254	10.1.1.1	DNS	Standard query
14	6.202100	10.1.1.1	192.168.254.254	TCP	1069 > http [SYN]
15	6.202513	192.168.254.254	10.1.1.1	TCP	http > 1069 [ACK]
16	6.202543	10.1.1.1	192.168.254.254	TCP	1069 > http [ACK]
17	6.202651	10.1.1.1	192.168.254.254	HTTP	GET / HTTP/1.1

⊕ Frame 15 (62 bytes on wire, 62 bytes captured)

⊕ Ethernet II, Src: Cisco_cf:66:40 (00:0c:85:cf:66:40), Dst: QuantaCo_bd:0c:

⊕ Internet Protocol, Src: 192.168.254.254 (192.168.254.254), Dst: 10.1.1.1 (10.1.1.1)

⊖ Transmission Control Protocol, Src Port: http (80), Dst Port: 1069 (1069),

Source port: http (80)

Destination port: 1069 (1069)

Sequence number: 0 (relative sequence number)

Acknowledgement number: 1 (relative ack number)

Header length: 28 bytes

⊖ Flags: 0x12 (SYN, ACK)

0... = Congestion window Reduced (CWR): Not set

.0.. = ECN-Echo: Not set

^ Urgent: Not set

A protocol analyzer shows server response in frame 15

- ACK flag set to indicate a valid Acknowledgement number
- Acknowledgement number response to initial sequence number as relative value of 1
- SYN flag set to indicate the Initial sequence number for the server to client session
- Destination port number of 1069 to corresponding to the clients source port
- Source port number of 80 (HTTP) indicating the web server service (httpd)

TCP 3-way Handshake (ACK)

13	6.201109	192.168.254.254	10.1.1.1	DNS	Standard query re
14	6.202100	10.1.1.1	192.168.254.254	TCP	1069 > http [SYN]
15	6.202513	192.168.254.254	10.1.1.1	TCP	http > 1069 [SYN,
16	6.202543	10.1.1.1	192.168.254.254	TCP	1069 > http [ACK]
17	6.202651	10.1.1.1	192.168.254.254	HTTP	GET / HTTP/1.1

+	Frame 16 (54 bytes on wire, 54 bytes captured)
+	Ethernet II, Src: QuantaCo_bd:0c:7c (00:c0:9f:bd:0c:7c), Dst: Cisco_cf:66:40
+	Internet Protocol, Src: 10.1.1.1 (10.1.1.1), Dst: 192.168.254.254 (192.168.254.254)
-	Transmission Control Protocol, Src Port: 1069 (1069), Dst Port: http (80), Seq: 1069, Win: 0, Len: 0
	Source port: 1069 (1069)
	Destination port: http (80)
	Sequence number: 1 (relative sequence number)
	Acknowledgement number: 1 (relative ack number)
	Header length: 20 bytes
-	Flags: 0x10 (ACK)
	0... = Congestion window Reduced (CWR): Not set
	.0.. = ECN-Echo: Not set
	..0. = Urgent: Not set

Protocol Analyzer shows client response to session in frame 16

The TCP segment in this frame shows:

- ACK flag set to indicate a valid Acknowledgement number
- Acknowledgement number response to initial sequence number as relative value of 1
- Source port number of 1069 to corresponding
- Destination port number of 80 (HTTP) indicating the web server service (httpd)

TCP Session Termination (FIN)

The image shows a Wireshark packet capture interface. The top pane displays a list of packets. Packet 20 is highlighted, showing a TCP FIN request from 192.168.254.254 to 10.1.1.1 on port 1069. The bottom pane shows the detailed view of packet 20, including Ethernet II, Internet Protocol, and Transmission Control Protocol (TCP) details. The TCP details show the source port as http (80), destination port as 1069 (1069), sequence number as 440 (relative sequence number), and acknowledgement number as 414 (relative ack number). The flags are 0x11 (FIN, ACK). The congestion window is reduced (CWR): Not set.

No.	Time	Source	Destination	Protocol	Length	Info
19	6.203857	192.168.254.254	10.1.1.1	HTTP	200	OK (200)
20	6.203876	192.168.254.254	10.1.1.1	TCP	60	http > 1069 [FIN, Seq=440]
21	6.203899	10.1.1.1	192.168.254.254	TCP	60	1069 > http [ACK, Seq=414]
22	6.204139	10.1.1.1	192.168.254.254	TCP	60	1069 > http [FIN, Seq=414]
23	6.204416	192.168.254.254	10.1.1.1	TCP	60	http > 1069 [ACK, Seq=415]
24	6.204668	10.1.1.1	192.168.254.254	DNS	Standard query response	

Frame 20 (60 bytes on wire, 60 bytes captured)

Ethernet II, Src: Cisco_cf:66:40 (00:0c:85:cf:66:40), Dst: QuantaCo_bd:0c:7c (08:00:0c:2c:7c:00)

Internet Protocol, src: 192.168.254.254 (192.168.254.254), dst: 10.1.1.1 (10.1.1.1)

Transmission Control Protocol, Src Port: http (80), Dst Port: 1069 (1069), Seq: 440, Win: 0, Len: 0

Source port: http (80)

Destination port: 1069 (1069)

Sequence number: 440 (relative sequence number)

Acknowledgement number: 414 (relative ack number)

Header length: 20 bytes

Flags: 0x11 (FIN, ACK)

0... = Congestion window reduced (CWR): Not set

- A protocol analyzer shows details of frame 20, TCP FIN request.

- The destination and source ports
- The header field contents and values

TCP Session Termination (ACK)

19	6.203857	192.168.254.254	10.1.1.1	HTTP	HTTP/1.1 200 OK (
20	6.203876	192.168.254.254	10.1.1.1	TCP	http > 1069 [FIN,
21	6.203899	10.1.1.1	192.168.254.254	TCP	1069 > http [ACK]
22	6.204139	10.1.1.1	192.168.254.254	TCP	1069 > http [FIN,
23	6.204416	192.168.254.254	10.1.1.1	TCP	http > 1069 [ACK]
24	6.602668	10.1.1.1	192.168.254.254	DNS	Standard query A

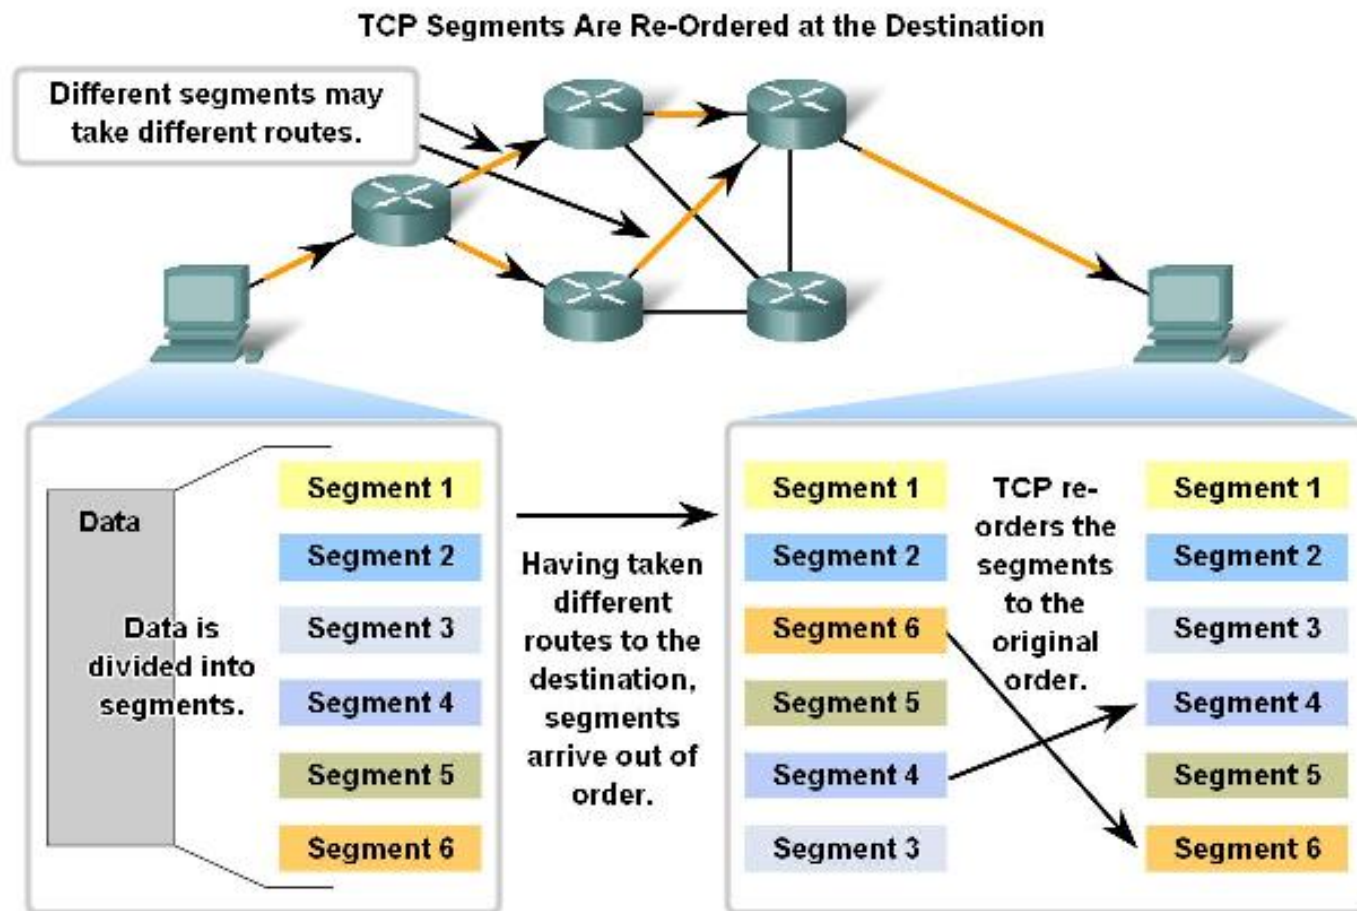
⊕	Frame 21 (54 bytes on wire, 54 bytes captured)
⊕	Ethernet II, Src: QuantaCo_bd:0c:7c (00:c0:9f:bd:0c:7c), Dst: Cisco_cf:66:40
⊕	Internet Protocol, Src: 10.1.1.1 (10.1.1.1), Dst: 192.168.254.254 (192.168.254.254)
⊖	Transmission Control Protocol, Src Port: 1069 (1069), Dst Port: http (80), Seq
	Source port: 1069 (1069)
	Destination port: http (80)
	Sequence number: 414 (relative sequence number)
	Acknowledgement number: 441 (relative ack number)
	Header length: 20 bytes
⊖	Flags: 0x10 (ACK)
	0 - Congestion window reduced (CWR): Not set

- A protocol analyzer shows details of frame 21, TCP ACK response.

- The destination and source ports
- The header field contents and values

TCP Segment Reassembly

When services send data using TCP, segments may arrive at their destination out of order. For the original message to be understood by the recipient, the data in these segments is reassembled into the original order. Sequence numbers are assigned in the header of each packet to achieve this goal.

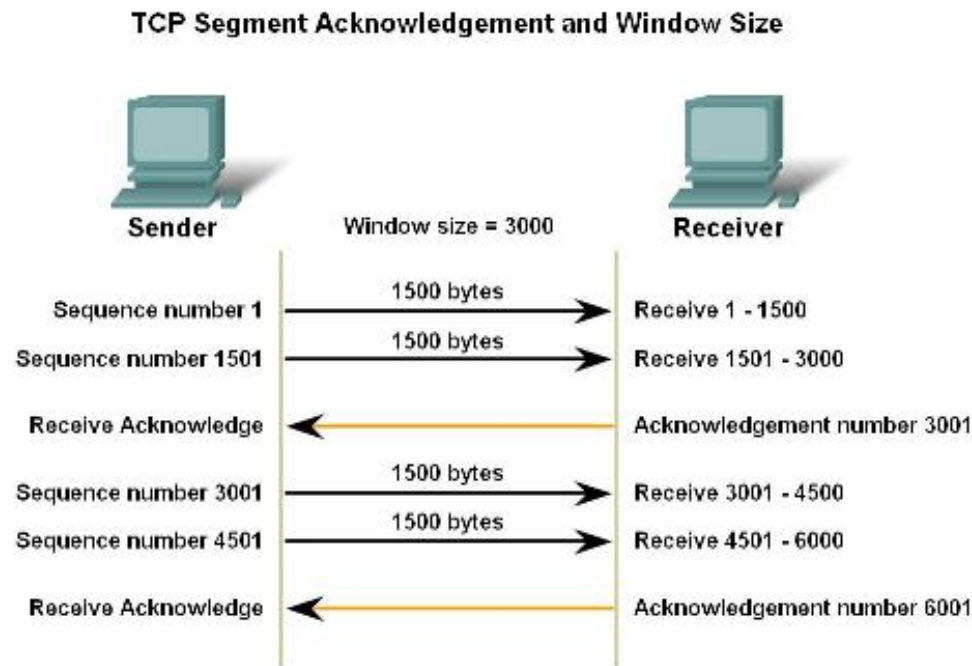


TCP Flow Control

* TCP also provides mechanisms for flow control. Flow control assists the reliability of TCP transmission by adjusting the effective rate of data flow between the two services in the session. When the source is informed that the specified amount of data in the segments is received, it can continue sending more data for this session.

* This Window Size field in the TCP header specifies the amount of data that can be transmitted before an acknowledgement must be received. The initial window size is determined during the session startup via the three-way handshake.

* TCP feedback mechanism adjusts the effective rate of data transmission to the maximum flow that the network and destination device can support without loss. TCP attempts to manage the rate of transmission so that all data will be received and retransmissions will be minimized.



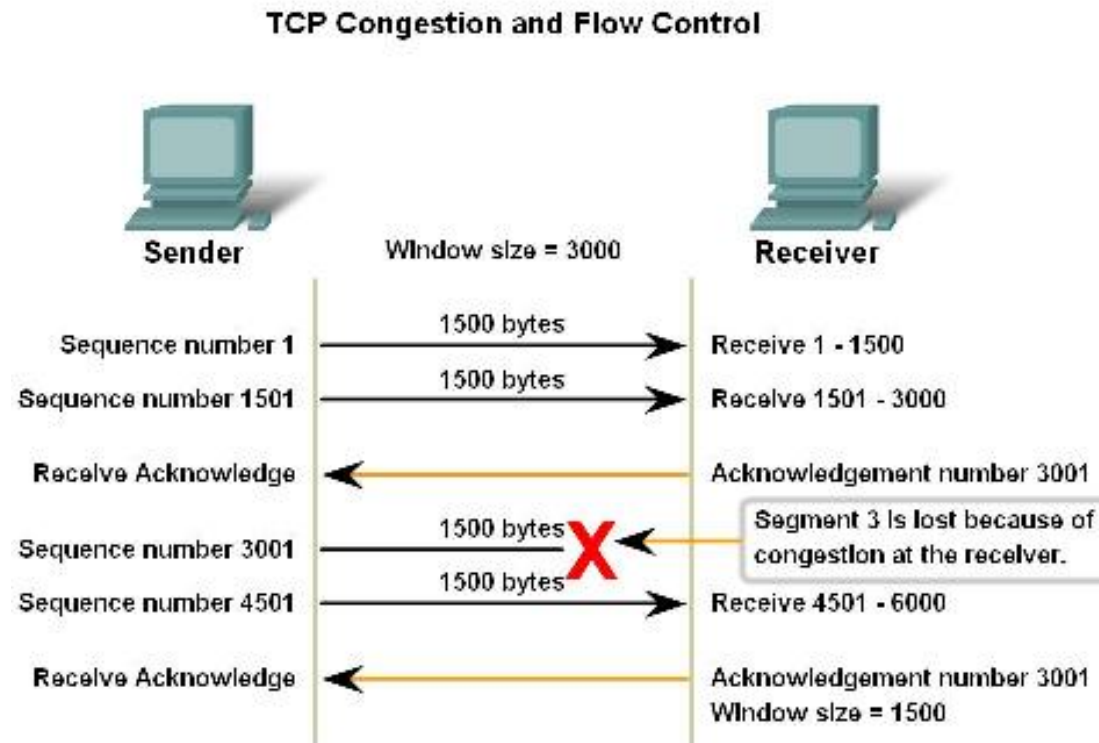
The **window size** determines the number of bytes sent before an acknowledgment is expected.

The **acknowledgement** number is the number of the next expected byte.

Reducing Window Size

* Another way to control the data flow is to use dynamic window sizes. When network resources are constrained, TCP can reduce the window size to require that received segments be acknowledged more frequently. This effectively slows down the rate of transmission because the source waits for data to be acknowledged more frequently.

* The TCP receiving host sends the window size value to the sending TCP to indicate the number of bytes that it is prepared to receive as a part of this session. If the destination needs to slow down the rate of communication because of limited buffer memory, it can send a smaller window size value to the source as part of an acknowledgement.



If segments are lost because of congestion, the Receiver will acknowledge the last received sequential segment and reply with a reduced window size.

VI. UDP Protocol

UDP is a simple protocol that provides the basic Transport layer functions. It has much lower overhead than TCP, since it is not connection-oriented and does not provide the sophisticated retransmission, sequencing, and flow control mechanisms.

This does not mean that applications that use UDP are always unreliable. It simply means that these functions are not provided by the Transport layer protocol and must be implemented elsewhere if required.

Although the total amount of UDP traffic found on a typical network is often relatively low, key Application layer protocols **that use UDP include:**

- Domain Name System (DNS)
- Simple Network Management Protocol (SNMP)
- Dynamic Host Configuration Protocol (DHCP)
- Routing Information Protocol (RIP)
- Trivial File Transfer Protocol (TFTP)
- Online games

Some applications, such as online games or VoIP, can tolerate some loss of some data. If these applications used TCP, they may experience large delays while TCP detects data loss and retransmits data. These delays would be more detrimental to the application than small data losses. Some applications, such as DNS, will simply retry the request if they do not receive a response, and therefore they do not need TCP to guarantee the message delivery. The low overhead of UDP makes it very desirable for such applications.

UDP Datagram Structure

bits	0 - 15	16 - 31
0	Source Port	Destination Port
32	Length	Checksum
64	Data	

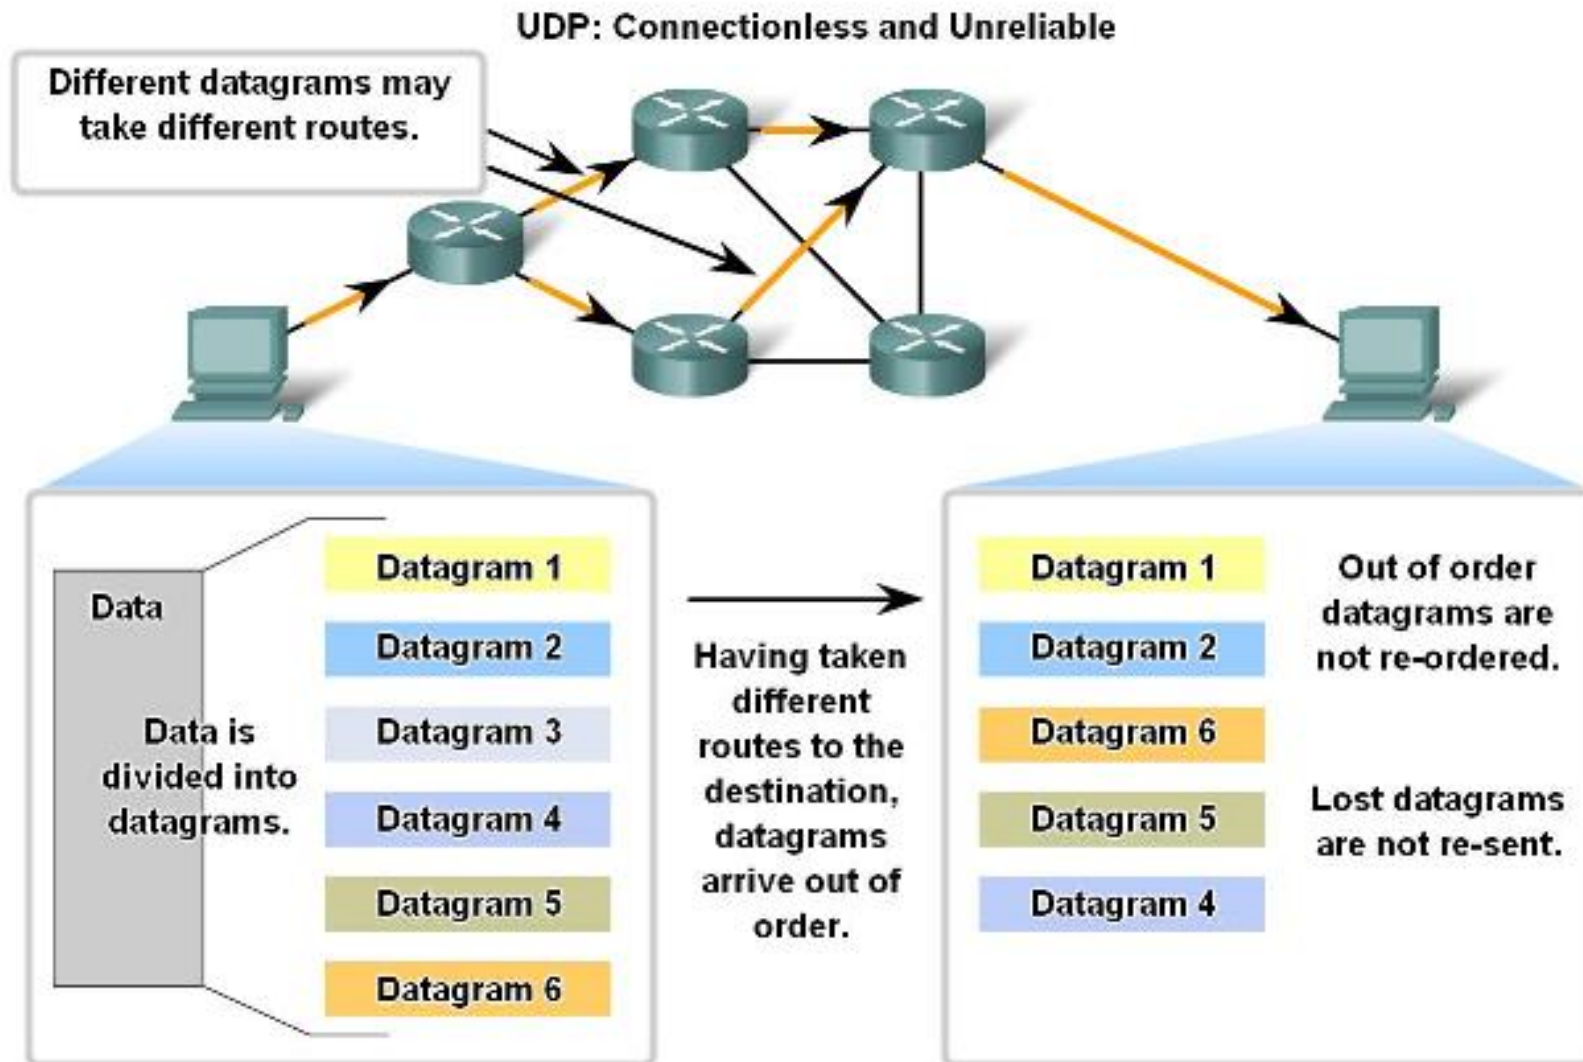
* Source port: This field identifies the sending port when meaningful and should be assumed to be the port to reply to if needed. If not used, then it should be zero.

* Destination port: This field identifies the destination port and is required.

* Length: A 16-bit field that specifies the length in bytes of the entire datagram: header and data. The minimum length is 8 bytes since that's the length of the header. The field size sets a theoretical limit of 65,535 bytes (8 byte header + 65527 bytes of data) for a UDP datagram. The practical limit for the data length which is imposed by the underlying IPv4 protocol is 65,507 bytes.

* Checksum: The 16-bit checksum field is used for error-checking of the header and data. The algorithm for computing the checksum is different for transport over IPv4 and IPv6. If the checksum is omitted in IPv4, the field uses the value all-zeros. This field is not optional for IPv6.

UDP Datagram Reassembly



VII. Ports

* The TCP and UDP based services keep track of the various applications that are communicating. To differentiate the segments and datagrams for each application, both TCP and UDP have header fields that can uniquely identify these applications. These unique identifiers are the port numbers.

* In the header of each segment or datagram, there is a source and destination port. The source port number is the number for this communication associated with the originating application on the local host. The destination port number is the number for this communication associated with the destination application on the remote host.

* Port numbers are assigned in various ways, depending on whether the message is a request or a response. While server processes have static port numbers assigned to them, clients dynamically chooses a port number for each conversation.

* The combination between IP address and port number is called socket and it's unique connection.

```
C:\>netstat
```

```
Active Connections
```

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

```
C:\>
```

Port Numbers

Well Known Ports (Numbers 0 to 1023) - These numbers are reserved for services and applications. They are commonly used for applications such as HTTP (web server) POP3/SMTP (e-mail server) and Telnet. By defining these well-known ports for server applications, client applications can be programmed to request a connection to that specific port and its associated service.

Registered Ports (Numbers 1024 to 49151) - These port numbers are assigned to user 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. When not used for a server resource, these ports may also be used dynamically selected by a client as its source port.

Dynamic or Private Ports (Numbers 49152 to 65535) - Also known as Ephemeral Ports, these are usually assigned dynamically to client applications when initiating a connection. It is not very common for a client to connect to a service using a Dynamic or Private Port (although some peer-to-peer file sharing programs do).

Using both TCP and UDP

Some applications may use both TCP and UDP. For example, the low overhead of UDP enables DNS to serve many client requests very quickly. Sometimes, however, sending the requested information may require the reliability of TCP. In this case, the well known port number of 53 is used by both protocols with this service.

TCP Ports

Port Number Range	Port Group
0 to 1023	Well Known (Contact) Ports
1024 to 49151	Registered Ports
49152 to 65535	Private and/or Dynamic Ports

Registered TCP Ports:
 1863 MSN Messenger
 8008 Alternate HTTP
 8080 Alternate HTTP

Well Known TCP Ports
 21 FTP
 23 Telnet
 25 SMTP
 80 HTTP
 110 POP3
 194 Internet Relay Chat (IRC)
 443 Secure HTTP (HTTPS)

UDP Ports

Port Number Range	Port Group
0 to 1023	Well Known (Contact) Ports
1024 to 49151	Registered Ports
49152 to 65535	Private and/or Dynamic Ports

Registered UDP Ports:
 1812 RADIUS Authentication Protocol
 2000 Cisco SCCP (VoIP)
 5004 RTP (Voice and Video Transport Protocol)
 5060 SIP (VoIP)

Well Known UDP Ports:
 69 TFTP
 520 RIP

TCP and UDP Ports

Port Number Range	Port Group
0 to 1023	Well Known (Contact) Ports
1024 to 49151	Registered Ports
49152 to 65535	Private and/or Dynamic Ports

Registered TCP/UDP Common Ports:
 1433 MS SQL
 2948 WAP (MMS)

Well Known TCP/UDP Common Ports:
 53 DNS
 161 SNMP
 531 AOL Instant Messenger, IRC