


Client_Socket Interfaces (TCP Basic Handling)					
Client Node		Internet	Server Node		EventHelix.com/EventStudio 2.0
Client		Net	Server		
Client App	Client Socket	Network	Server Socket	Server App	04-Jul-03 14:31 (Page 1)

Copyright © 2000-2003 EventHelix.com Inc. All Rights Reserved.

LEG: About TCP

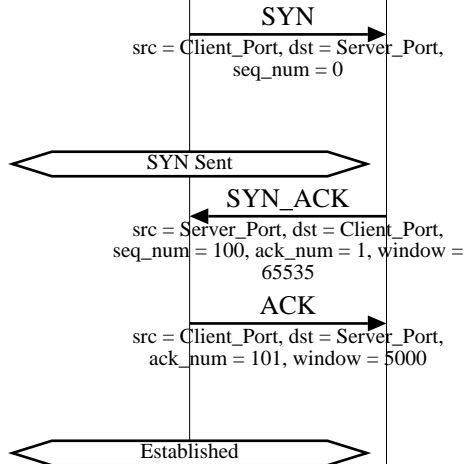
TCP (Transmission Control Protocol) provides a reliable end to end service that delivers packets over the Internet. Packets are delivered in sequence without loss or duplication.



Client Application creates Socket
 The socket is created in the Closed state
 Initial sequence number is set to 0
 Application wishes to communicate with a destination server using a TCP connection. The application opens a socket for the connection in active mode. In this mode, a TCP connection will be attempted with the server.
 Typically, the client will use a well known port number to communicate with the remote Server. For example, HTTP uses port 80.

LEG: Client initiates TCP connection

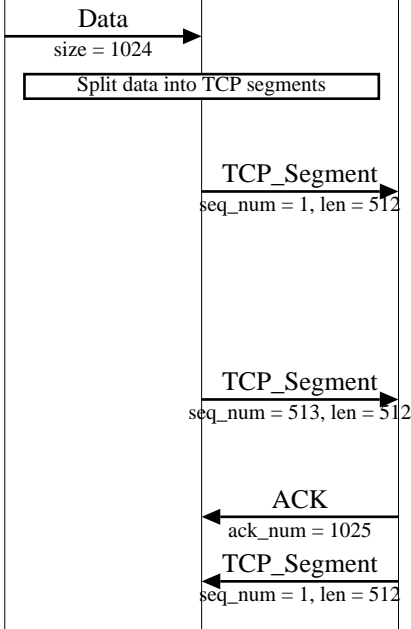
Client initiated three way handshake to establish a TCP connection



Client sets the SYN bit in the TCP header to request a TCP connection. The sequence number field is set to 0. Since the SYN bit is set, this sequence number is used as the initial sequence number
 Socket transitions to the SYN Sent state
 Client receives the SYN_ACK TCP segment
 Client now acknowledges the first segment, thus completing the three way handshake. The receive window is set to 5000. Ack sequence number is set to 101, this means that the next expected sequence number is 101.
 At this point, the client assumes that the TCP connection has been established


LEG: Short data transfer

Data transfer phase: Here a short data transfer takes place, thus TCP slow start has little impact

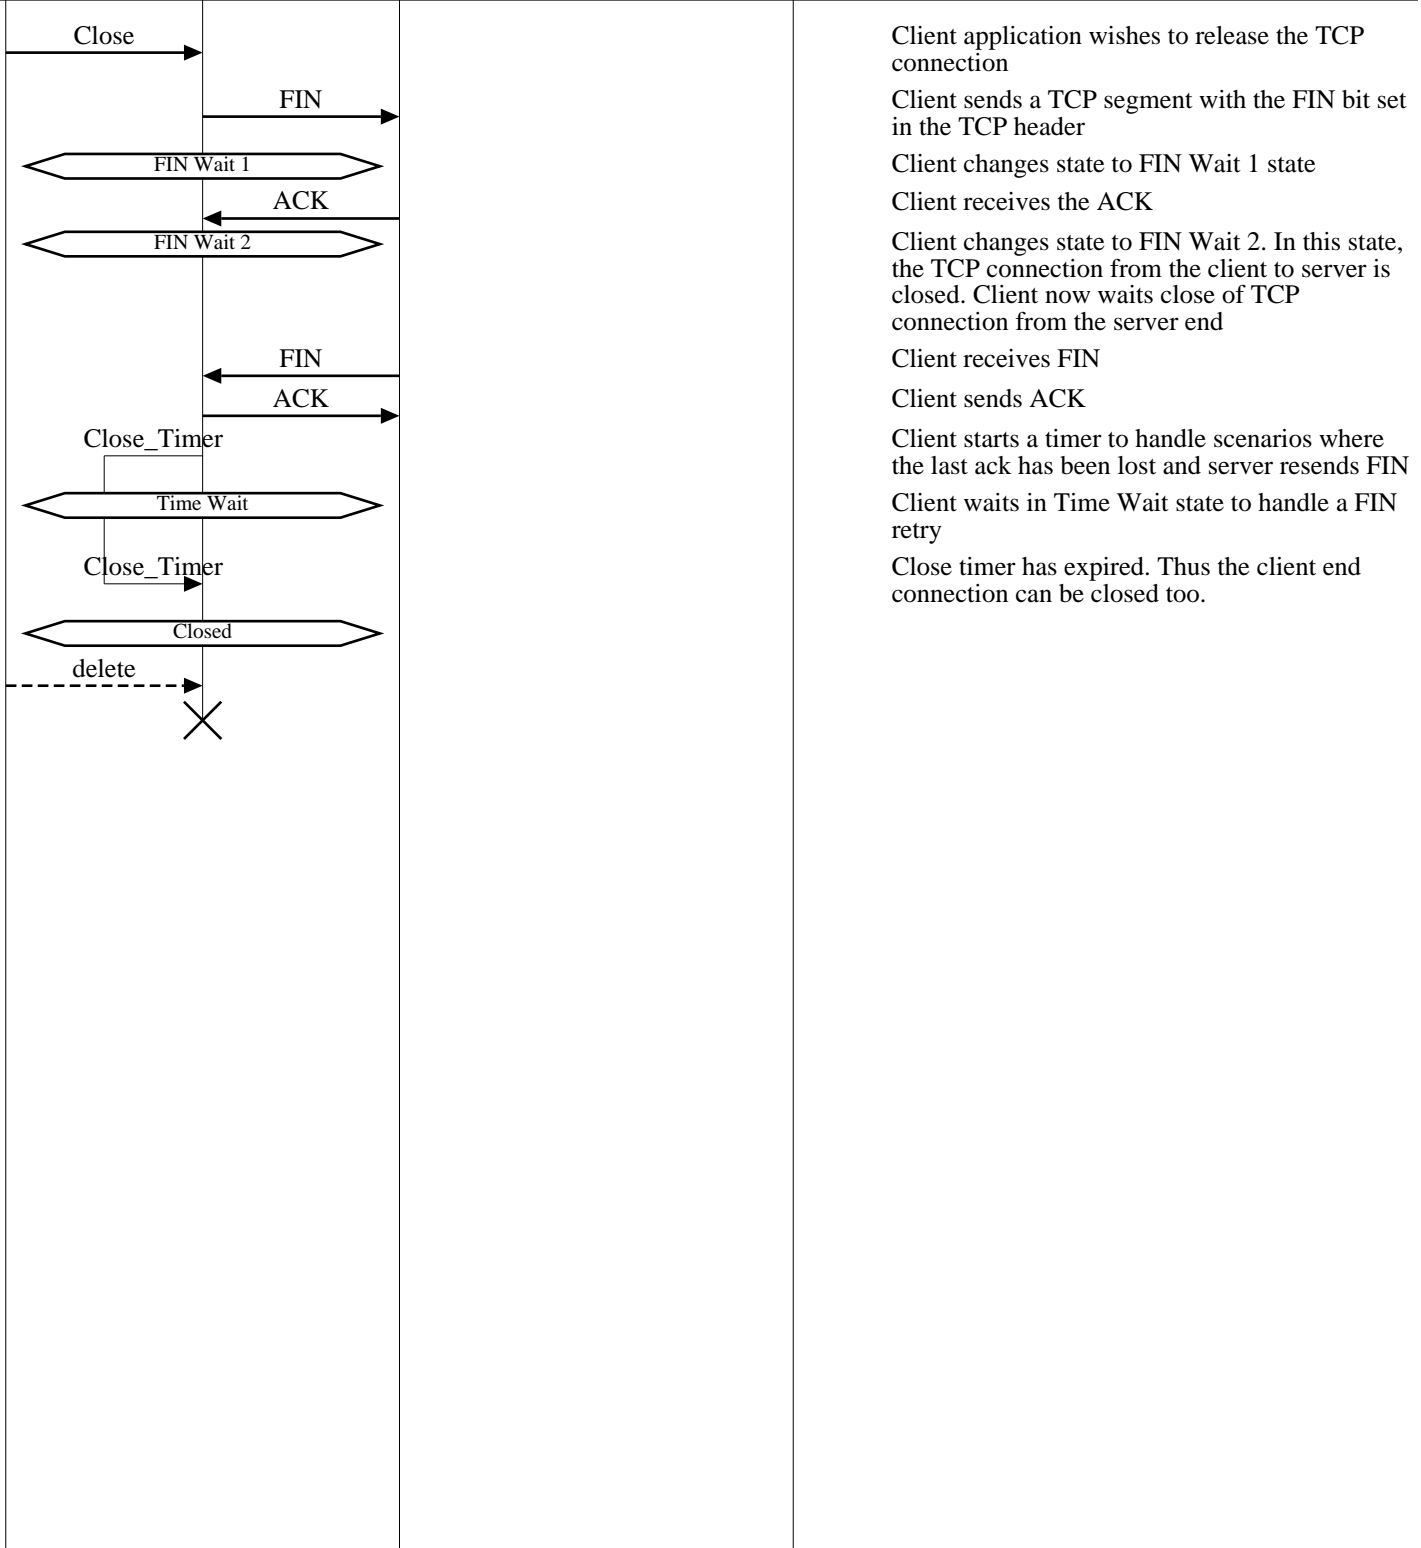


Client application sends 1024 bytes of data to the socket
 This TCP connection limits TCP segments to 512 bytes, thus the received data is split into 2 TCP segments
 The first TCP segment is sent with a sequence number of 1. This is the sequence number for the first byte in the segment.
 (Note that unlike other protocols, TCP maintains sequence numbers at byte level. The sequence number field in the TCP header corresponds to the first byte in the segment.)
 Bytes in the first TCP segment correspond to 1 to 512 sequence numbers. Thus, the second TCP segment contains data starting with 513 sequence number

Client_Socket Interfaces (TCP Basic Handling)					
Client Node		Internet	Server Node		EventHelix.com/EventStudio 2.0
Client		Net	Server		
Client App	Client Socket	Network	Server Socket	Server App	04-Jul-03 14:31 (Page 2)



Client initiates TCP connection close



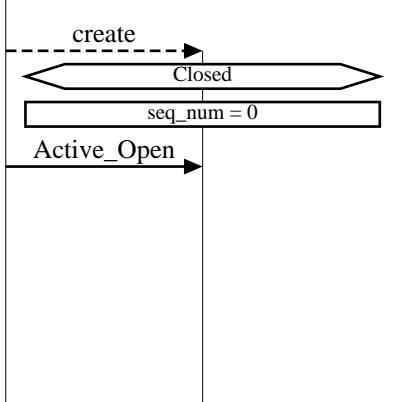
Client_Socket Interfaces (TCP Slow Start)					
Client Node		Internet	Server Node		EventHelix.com/EventStudio 2.0
Client		Net	Server		
Client App	Client Socket	Network	Server Socket	Server App	04-Jul-03 14:31 (Page 3)

Copyright © 2000-2003 EventHelix.com Inc. All Rights Reserved.

LEG: About TCP Slow Start

TCP is an end to end protocol which operates over the heterogeneous Internet. TCP has no advance knowledge of the network characteristics, thus it has to adjust its behavior according to the current state of the network. TCP has built in support for congestion control. Congestion control ensures that TCP does not pump data at a rate higher than what the network can handle.

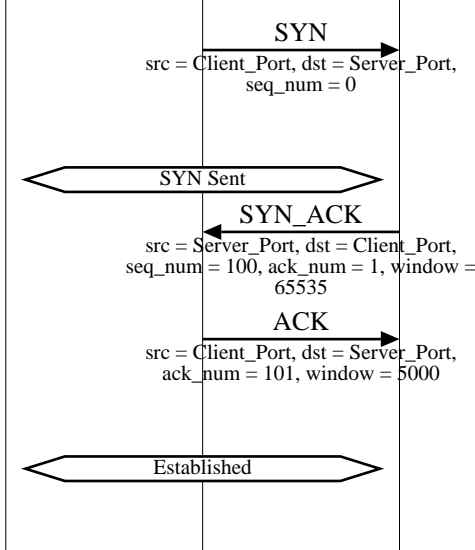
In this sequence diagram we will analyse "Slow start", an important part of the congestion control mechanisms built right into TCP. As the name suggests, "Slow Start" starts slowly, increasing its window size as it gains confidence about the networks throughput.



Client Application creates Socket
 The socket is created in the Closed state
 Initial sequence number is set to 0
 Application wishes to communicate with a destination server using a TCP connection. The application opens a socket for the connection in active mode. In this mode, a TCP connection will be attempted with the server.
 Typically, the client will use a well known port number to communicate with the remote Server. For example, HTTP uses port 80.

LEG: Client initiates TCP connection

Client initiated three way handshake to establish a TCP connection



Client sets the SYN bit in the TCP header to request a TCP connection. The sequence number field is set to 0. Since the SYN bit is set, this sequence number is used as the initial sequence number

Socket transitions to the SYN Sent state
 Client receives the SYN_ACK TCP segment


Client now acknowledges the first segment, thus completing the three way handshake. The receive window is set to 5000. Ack sequence number is set to 101, this means that the next expected sequence number is 101.

At this point, the client assumes that the TCP connection has been established

LEG: TCP Slow Start

A TCP connection starts in the "Slow Start" state. In this state, TCP adjusts its transmission rate based on the rate at which the acknowledgements are received from the other end.

TCP Slow start is implemented using two variables, viz cwnd (Congestion Window) and ssthresh (Slow Start Threshold). cwnd is a self imposed transmit window restriction at the sender end. cwnd will increase as TCP gains more confidence on the networks ability to handle traffic. ssthresh is the threshold for determining the point at which TCP exits slow start. If cwnd increases beyond ssthresh, the TCP session in that direction is considered to be out of slow start phase




Client maintains a congestion window (cwnd). Initially the window is set to lower of the maximum TCP segment size and receiver's allowed window size. In most cases the segment size is smaller than receiver window, thus cwnd is set to the maximum TCP segment size (512 in this example)

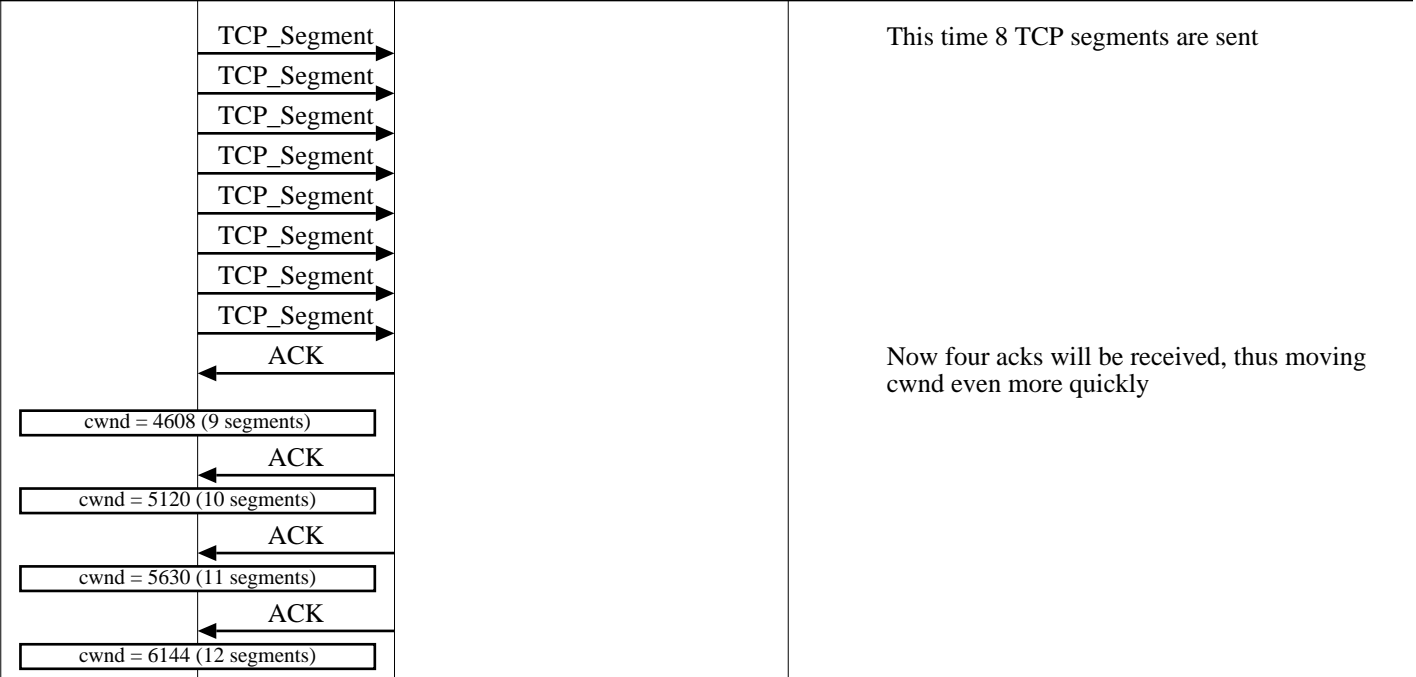
Note here that cwnd implements a transmitter end flow control. The receiver advertised window implements a receiver enforced flow control.

Client_Socket Interfaces (TCP Slow Start)					
Client Node		Internet	Server Node		EventHelix.com/EventStudio 2.0
Client		Net	Server		
Client App	Client Socket	Network	Server Socket	Server App	04-Jul-03 14:31 (Page 4)
				<p>TCP connections start with ssthresh set to 64K. This variable will be used to determine the point at which TCP exits slow start</p> <p>Client end TCP connection moves to slow start state</p> <p>Client application sends 5120 bytes of data to the socket</p>	
Roundtrip #1 of data transmission					
				<p>The first TCP segment is sent with a sequence number of 1. This is the sequence number for the first byte in the segment.</p> <p>Client receives the acknowledgement for the first TCP data segment</p> <p>As the TCP session is in slow start, receipt of an acknowledgement increments the congestion window by one 1 segment.</p>	
Roundtrip #2 of data transmission					
				<p>Since the congestion window has increased to 2, TCP can now send two segments without waiting for an ack</p> <p>Receipt for ack again moves the congestion window</p>	
Roundtrip #3 of data transmission					
				<p>Now three segments can be sent without waiting for an ack</p> <p>The TCP acknowledgements again increment cwnd. This time two acks are received, so cwnd will get incremented by 2</p> <p>Since cwnd has reached 5 segments, TCP is allowed to send 5 segments without waiting for the ack</p>	
Roundtrip #4 of data transmission					
				<p>Three acknowledgements will be received for the 5 TCP segments. Now the cwnd has almost started increasing geometrically for every round</p>	

Client_Socket Interfaces (TCP Slow Start)					
Client Node		Internet	Server Node		EventHelix.com/EventStudio 2.0
Client		Net	Server		
Client App	Client Socket	Network	Server Socket	Server App	04-Jul-03 14:31 (Page 5)



Roundtrip #5 of data transmission

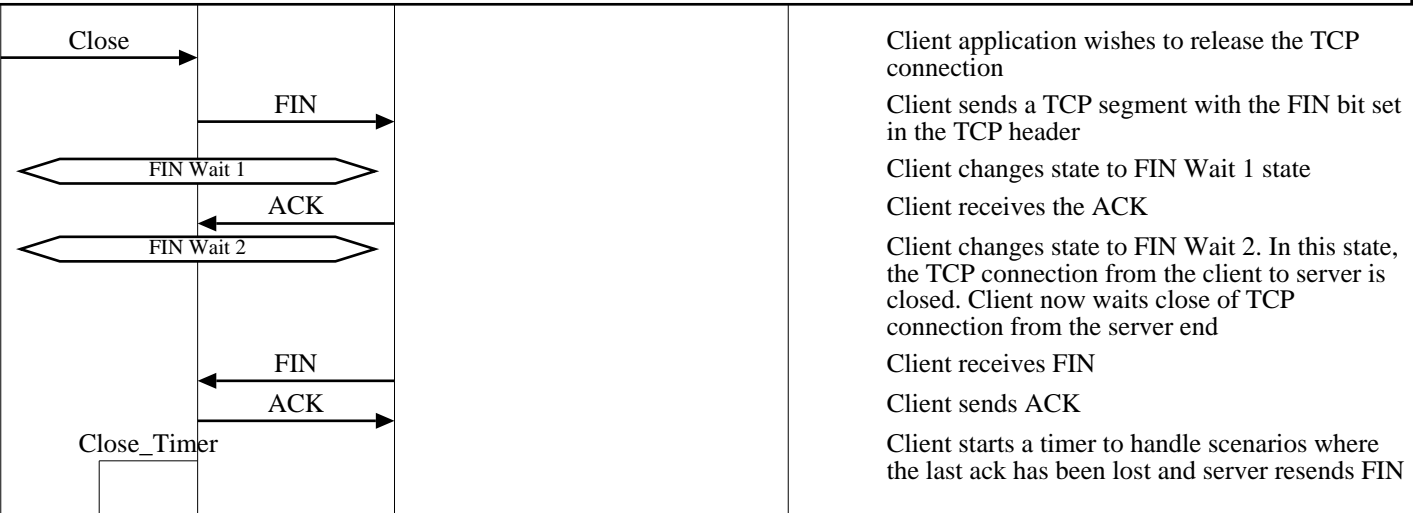


Within a few more roundtrip interactions cwnd will exceed ssthresh. At this point the session will be considered out of slow start. Note that the TCP connection from the client side is out of slow start but the server end is still in slow start as it has not sent any data to the client.


Exiting slow start signifies that the TCP connection has reached an equilibrium state where the congestion window closely matches the networks capacity. From this point on, the congestion window will not move geometrically. cwnd will move linearly once the connection is out of slow start.



Client initiates TCP connection close



Client_Socket Interfaces (TCP Slow Start)					
Client Node		Internet	Server Node		EventHelix.com/EventStudio 2.0
Client		Net	Server		
Client App	Client Socket	Network	Server Socket	Server App	04-Jul-03 14:31 (Page 6)




Client waits in Time Wait state to handle a FIN retry
 Close timer has expired. Thus the client end connection can be closed too.

Client_Socket Interfaces (TCP Congestion Avoidance)					
Client Node		Internet	Server Node		EventHelix.com/EventStudio 2.0
Client		Net	Server		
Client App	Client Socket	Network	Server Socket	Server App	04-Jul-03 14:31 (Page 7)

Copyright © 2000-2003 EventHelix.com Inc. All Rights Reserved.

LEG: About TCP Congestion Avoidance

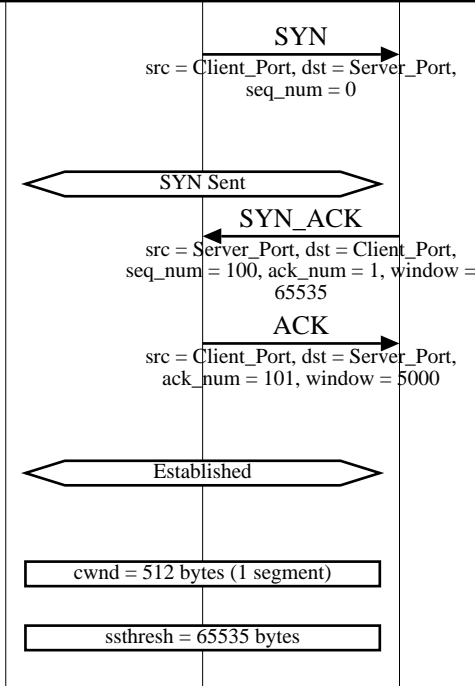
We have already seen that TCP connection starts up in slow start mode, geometrically increasing the congestion window (cwnd) until it crosses the slow start threshold (ssthresh). Once cwnd is greater than ssthresh, TCP enters the congestion avoidance mode of operation. In this mode, the primary objective is to maintain high throughput without causing congestion. If TCP detects segment loss, it assumes that congestion has been detected over the internet. As a corrective action, TCP reduces its data flow rate by reducing cwnd. After reducing cwnd, TCP goes back to slow start.



Client Application creates Socket
 The socket is created in the Closed state
 Initial sequence number is set to 0
 Application wishes to communicate with a destination server using a TCP connection. The application opens a socket for the connection in active mode. In this mode, a TCP connection will be attempted with the server.
 Typically, the client will use a well known port number to communicate with the remote Server. For example, HTTP uses port 80.

LEG: Client initiates TCP connection

Client initiated three way handshake to establish a TCP connection




Client sets the SYN bit in the TCP header to request a TCP connection. The sequence number field is set to 0. Since the SYN bit is set, this sequence number is used as the initial sequence number
 Socket transitions to the SYN Sent state
 Client receives the SYN_ACK TCP segment

Client now acknowledges the first segment, thus completing the three way handshake. The receive window is set to 5000. Ack sequence number is set to 101, this means that the next expected sequence number is 101.
 At this point, the client assumes that the TCP connection has been established

LEG: TCP Congestion Avoidance Operation

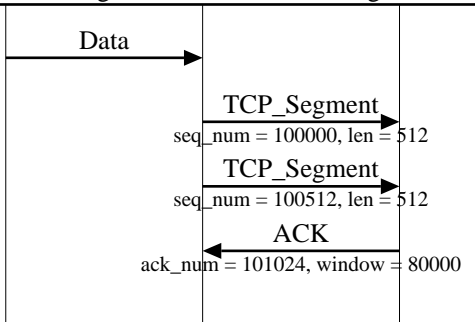
TCP connection begins with a congestion window size of 1 segment
 The slow start threshold starts with 64 Kbytes as the threshold value.

TCP session begins with "Slow Start". See the sequence diagram on slow start for details



Since cwnd < ssthresh, TCP state is slow start

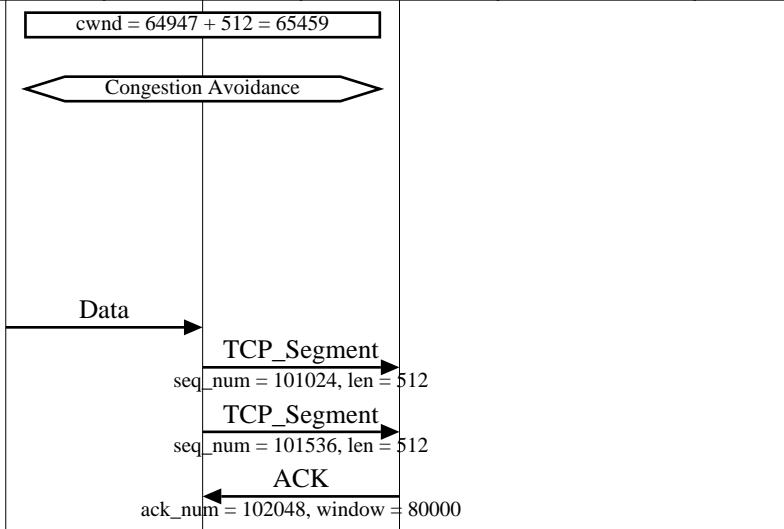
TCP congestion window grows from 512 bytes (1 segment) to 64947 (assuming no segment losses are detected during slow start). During slow start the congestion window was being incremented by 1 segment for every TCP Ack from the other end.



Client Application sends data for transmission over the TCP Socket
 Data is split into TCP Segments. The segments are sent over the Internet


Client_Socket Interfaces (TCP Congestion Avoidance)

Client Node		Internet	Server Node		EventHelix.com/EventStudio 2.0
Client		Net	Server		
Client App	Client Socket	Network	Server Socket	Server App	04-Jul-03 14:31 (Page 8)



Since TCP is in slow start, every ack leads to the window growing by one segment.
 At this point cwnd (=65459) > ssthresh (=65535) thus TCP changes state to congestion avoidance. Now TCP window growth will be much more conservative. If no segment or ack losses are detected, the congestion window will grow no more than one segment per roundtrip. (Compare this with geometric growth of 1 segment per TCP ack in slow start)
 More data is received from the client application
 Client data is split into TCP segments

cwnd is incremented using the formula: $cwnd = cwnd + (segment_size \cdot segment_size) / cwnd$




Now TCP is in congestion avoidance mode, so the TCP window advances very slowly. Here the window increased by only 4 bytes.
 Data to be sent to server
 TCP session sends out the data as a single segment
 TCP session starts a ack timer, awaiting the TCP ack for this segment.
 Note: The above timer is started for every segment. The timer is not shown at other places as it played role in our analysis
 TCP times out for a TCP ACK from the other end. This will be treated as a sign of congestion by TCP

When TCP detects congestion, it stores half of the current congestion window in ssthresh variable. In this case, ssthresh has been reduced from 65535 to 32731. This signifies that TCP now has less confidence on the network's ability to support big window sizes. Thus if the window size falls due to congestion, rapid window size increases will be carried out only until the window reaches 32731. Once this lowered ssthresh value is reached, window growth will be much slower.

Since current congestion has been detected by timeout, TCP takes the drastic action of reducing the congestion window to 1. As you can see, this will have a big impact on the throughput.
 cwnd (=1) is now lower than ssthresh (=32731) so TCP goes back to slow start.

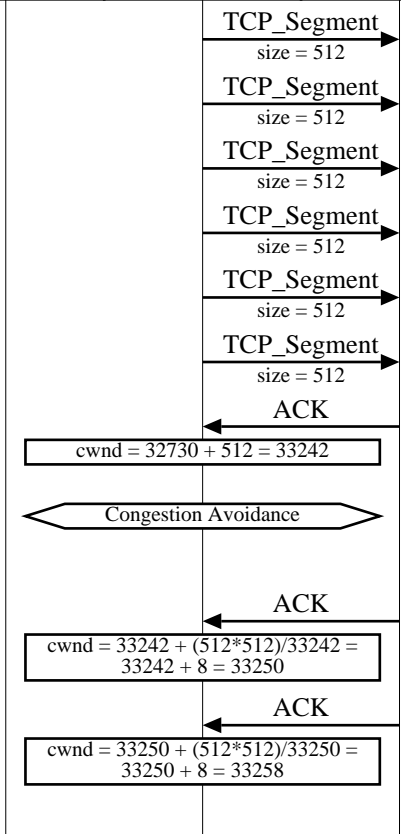
Since TCP is in slow start, a TCP acknowledgement results in the window growing by one segment

TCP window continues to grow exponentially until it reaches the ssthresh (=32731) value.



Client_Socket Interfaces (TCP Congestion Avoidance)

Client Node		Internet	Server Node		EventHelix.com/EventStudio 2.0
Client		Net	Server		
Client App	Client Socket	Network	Server Socket	Server App	04-Jul-03 14:31 (Page 9)



Six TCP segments are transmitted in the slow start mode

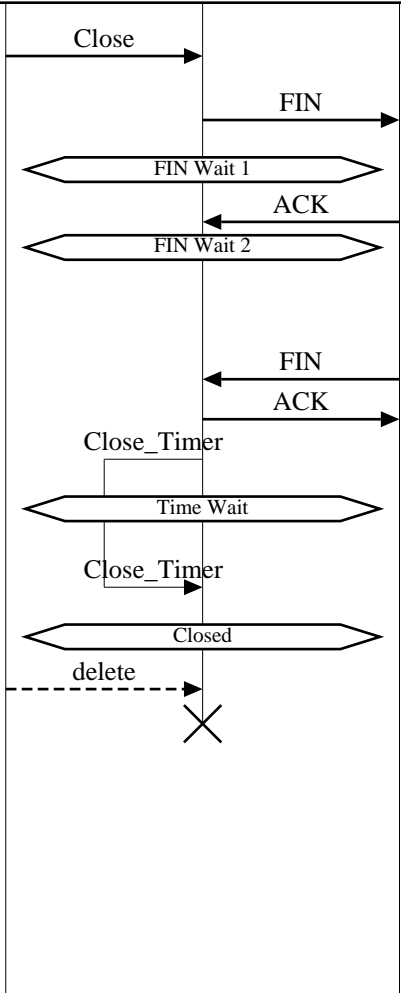
Ack for the first two segments is received
 TCP is in slow start so the congestion window is increased by one segment
 Now $cwnd (=33242) > ssthresh (=32731)$, thus the TCP session moves into congestion avoidance

Ack for the next two segments is received
 Now the TCP window is growing very slowly by approximately 8 bytes per ack

Ack for the last two segments is received
 Congestion window continues to advance at a slow rate

LEG: Client initiates TCP connection close

Client initiates TCP connection close



Client application wishes to release the TCP connection

Client sends a TCP segment with the FIN bit set in the TCP header

Client changes state to FIN Wait 1 state

Client receives the ACK

Client changes state to FIN Wait 2. In this state, the TCP connection from the client to server is closed. Client now waits close of TCP connection from the server end

Client receives FIN

Client sends ACK

Client starts a timer to handle scenarios where the last ack has been lost and server resends FIN

Client waits in Time Wait state to handle a FIN retry

Close timer has expired. Thus the client end connection can be closed too.


Client_Socket Interfaces (TCP Fast Transmit and Recovery)					
Client Node		Internet	Server Node		EventHelix.com/EventStudio 2.0
Client		Net	Server		
Client App	Client Socket	Network	Server Socket	Server App	04-Jul-03 14:31 (Page 10)

Copyright © 2000-2003 EventHelix.com Inc. All Rights Reserved.

LEG: About Fast Retransmit and Fast Recovery

TCP Slow Start and Congestion Avoidance lower the data throughput drastically when segment loss is detected. Fast Retransmit and Fast Recovery have been designed to speed up the recovery of the connection, without compromising its congestion avoidance characteristics.

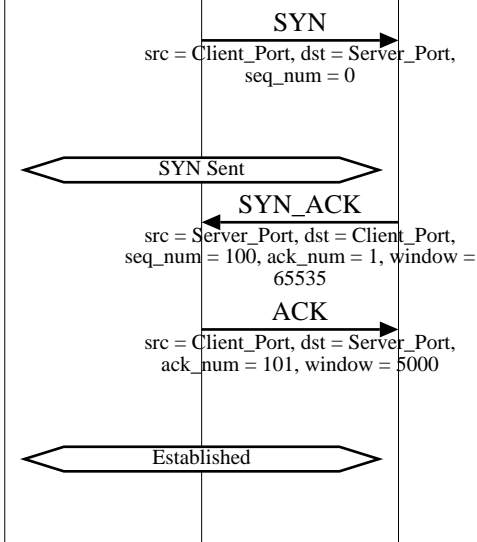
Fast Retransmit and Recovery detect a segment loss via duplicate acknowledgements. When a segment is lost, TCP at the receiver will keep sending ack segments indicating the next expected sequence number. This sequence number would correspond to the lost segment. If only one segment is lost, TCP will keep generating acks for the following segments. This will result in the transmitter getting duplicate acks (i.e. acks with the same ack sequence number)



Client Application creates Socket
 The socket is created in the Closed state
 Initial sequence number is set to 0
 Application wishes to communicate with a destination server using a TCP connection. The application opens a socket for the connection in active mode. In this mode, a TCP connection will be attempted with the server.
 Typically, the client will use a well known port number to communicate with the remote Server. For example, HTTP uses port 80.

LEG: Client initiates TCP connection

Client initiated three way handshake to establish a TCP connection



Client sets the SYN bit in the TCP header to request a TCP connection. The sequence number field is set to 0. Since the SYN bit is set, this sequence number is used as the initial sequence number

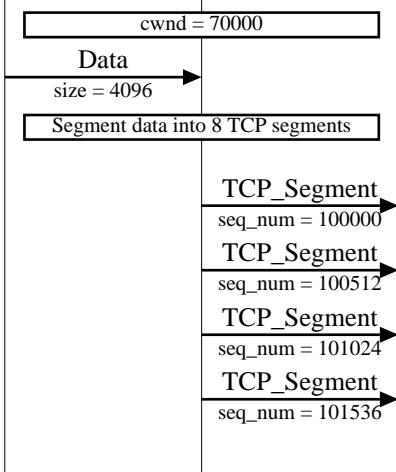
Socket transitions to the SYN Sent state
 Client receives the SYN_ACK TCP segment

Client now acknowledges the first segment, thus completing the three way handshake. The receive window is set to 5000. Ack sequence number is set to 101, this means that the next expected sequence number is 101.

At this point, the client assumes that the TCP connection has been established

LEG: Fast Retransmit and Recovery

TCP Connection begins with slow start. The congestion window grows from an initial 512 bytes to 70000 bytes



Congestion window has reached 70000 bytes
 Client App transmits 4Kbytes of data

TCP segments data to 8 TCP segments (each segment is 512 bytes)

TCP segment (start sequence number = 100000) is transmitted


TCP segment (start sequence number = 100512) is transmitted

TCP segment (start sequence number = 101024) is transmitted

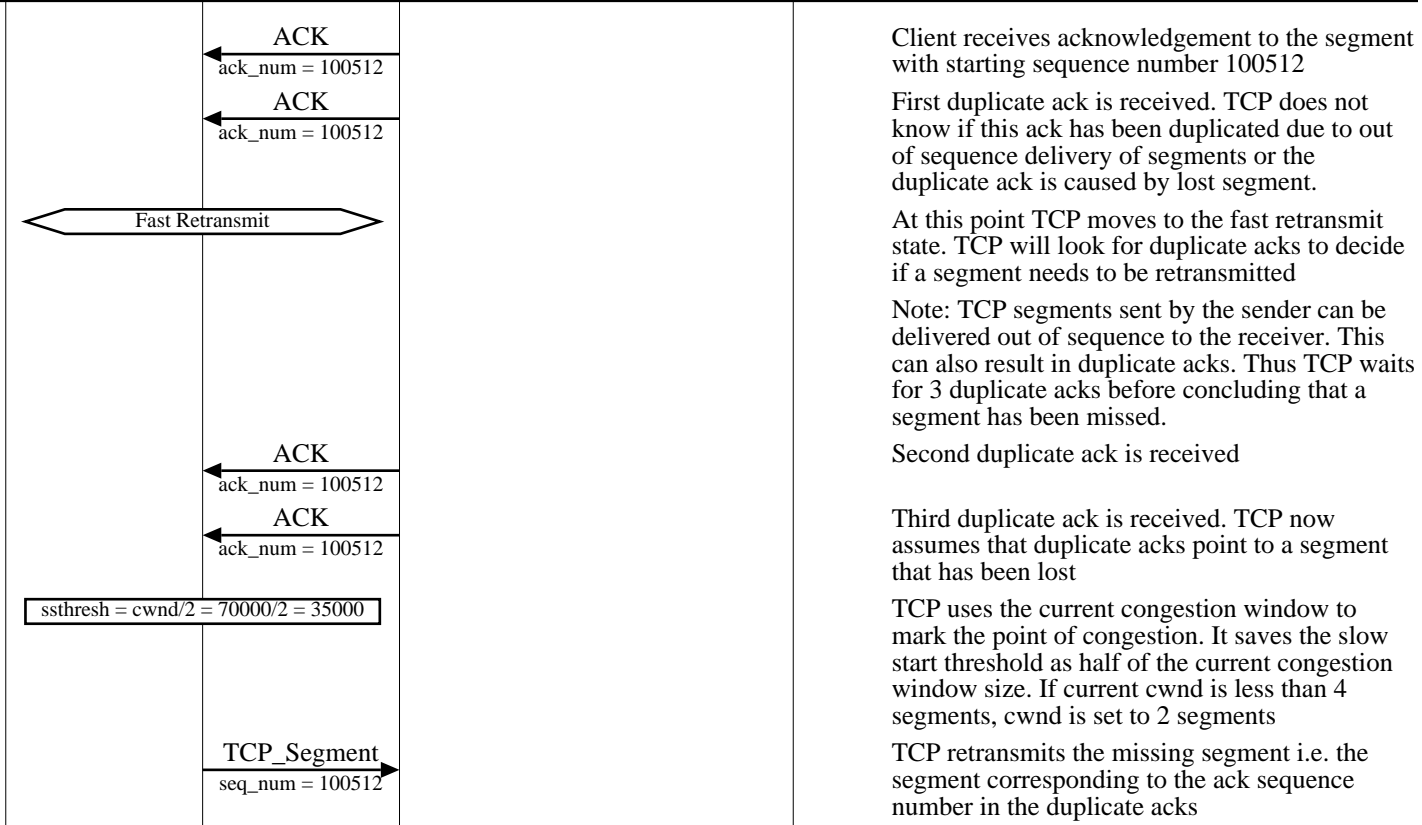
TCP segment (start sequence number = 101536) is transmitted

Client_Socket Interfaces (TCP Fast Transmit and Recovery)

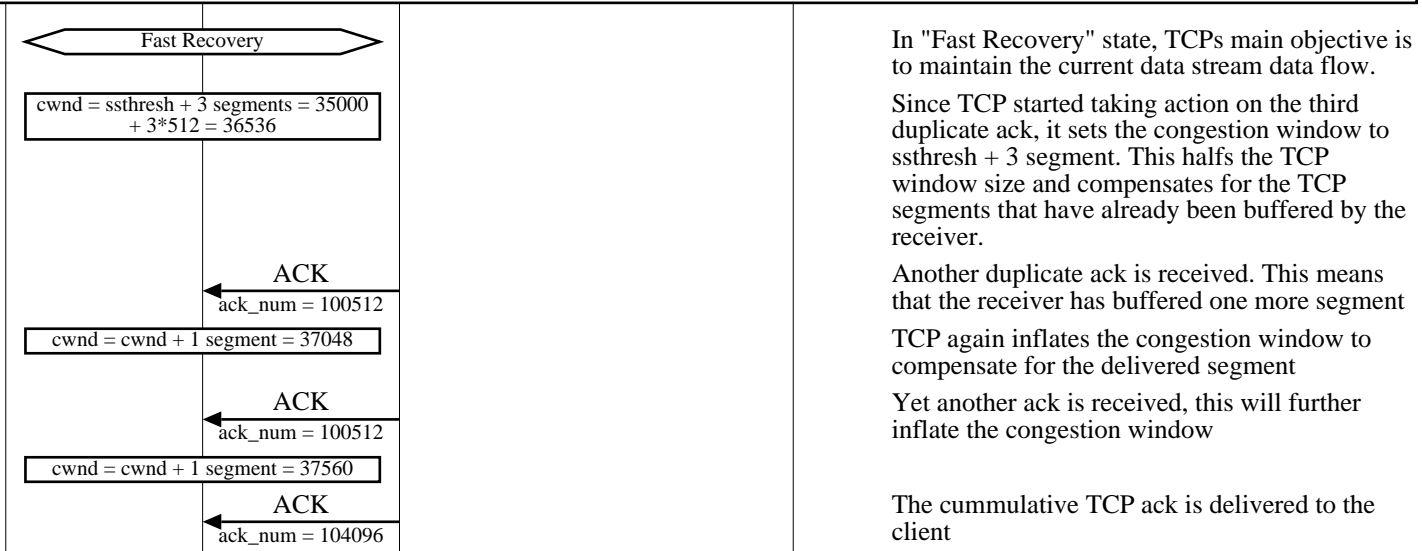
Client Node		Internet	Server Node		EventHelix.com/EventStudio 2.0
Client		Net	Server		
Client App	Client Socket	Network	Server Socket	Server App	04-Jul-03 14:31 (Page 11)



Fast Retransmit: TCP receives duplicate acks and it decides to retransmit the segment, without waiting for the segment timer to expire. This speeds up recovery of the lost segment




Fast Recovery: Once the lost segment has been transmitted, TCP tries to maintain the current data flow by not going back to slow start. TCP also adjusts the window for all segments that have been buffered by the receiver.



Congestion Avoidance

Client_Socket Interfaces (TCP Fast Transmit and Recovery)

Client Node		Internet	Server Node		EventHelix.com/EventStudio 2.0
Client		Net	Server		
Client App	Client Socket	Network	Server Socket	Server App	04-Jul-03 14:31 (Page 12)

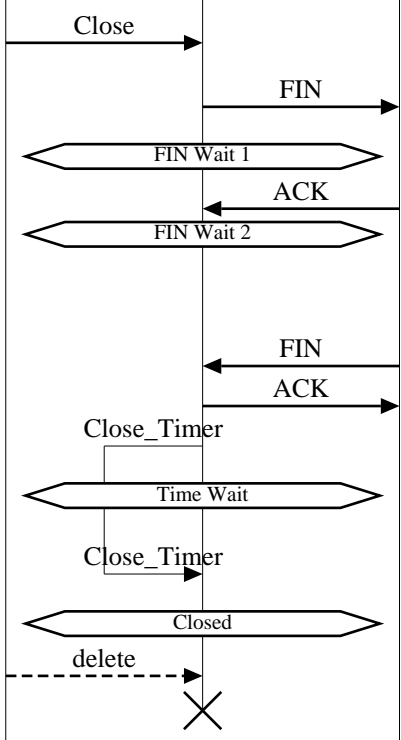


The connection has moved back to the congestion avoidance state.

TCP takes a congestion avoidance action and sets the segment size back to the slow start threshold. The TCP window will now increase by a maximum of one segment per round trip

LEG: Client initiates TCP connection close

Client initiates TCP connection close



Client application wishes to release the TCP connection

Client sends a TCP segment with the FIN bit set in the TCP header

Client changes state to FIN Wait 1 state

Client receives the ACK

Client changes state to FIN Wait 2. In this state, the TCP connection from the client to server is closed. Client now waits close of TCP connection from the server end

Client receives FIN

Client sends ACK

Client starts a timer to handle scenarios where the last ack has been lost and server resends FIN

Client waits in Time Wait state to handle a FIN retry

Close timer has expired. Thus the client end connection can be closed too.