Lecture 10 Reliability Quiz ANS

1. Which statement best captures the transport-layer abstraction difference emphasized in Lecture 10?

A. TCP and UDP both provide a reliable bytestream, but TCP also adds congestion control.
B. TCP provides a reliable bytestream abstraction, while UDP provides a best-effort datagram abstraction.
C. UDP provides a reliable packet abstraction, while TCP provides only ports.
D. TCP and UDP differ only in port numbering, not in application abstraction.
ANS: **B.** TCP is presented as a reliable bytestream service, while UDP gives a best-effort datagram abstraction.

2. What does the transport layer add on top of Layer 3 for both TCP and UDP?

A. De-multiplexing with ports.
B. Encryption of all application data.
C. Guaranteed congestion control.
D. Link-layer retransmission.
ANS: **A.** Both TCP and UDP support de-multiplexing with ports.

3. Why do applications prefer TCP's bytestream abstraction?

A. It lets applications manage individual bits on each link.
B. It hides packet-level best-effort behavior and presents ordered bytes over a connection.
C. It makes routers maintain per-application state.
D. It eliminates the need for IP addresses.
ANS: **B.** The bytestream abstraction lets applications think in terms of ordered bytes and connections rather than individual packets.

4. According to the end-to-end principle discussed in the lecture, where is transport-layer reliability implemented?

A. At intermediate routers.
B. At switches only.
C. At end hosts.
D. At every physical link.
ANS: **C.** The lecture emphasizes that transport-layer features such as reliability are implemented at the end hosts.

5. What is the purpose of an acknowledgment in the single-packet reliability design?

A. To tell the sender the packet arrived successfully.
B. To advertise the next router hop.
C. To set the packet checksum.
D. To increase the TTL.
ANS: **A.** The ack is feedback from the recipient telling the sender the packet was received.

6. How does the sender respond when a packet or its acknowledgment is dropped?

A. It waits forever.
B. It sends a negative acknowledgment.
C. It uses a timer and retransmits when the timer expires.

D. It changes the port number.
ANS: **C.** The sender starts a timer and retransmits if the ack does not arrive before timeout.

7. Why can cumulative acknowledgments be ambiguous when multiple packets are lost?
A. They reveal the exact set of received packets.
B. They tell the sender which packets were corrupted.
C. They do not specify exactly which later packets arrived, only the highest contiguous sequence received.
D. They are larger than individual acks.
ANS: **C.** Cumulative acks tell the sender only the highest contiguous point reached, so multiple-loss cases can be ambiguous.

8. Why are sequence numbers needed when sending multiple packets?
A. To replace IP addresses.
B. To reorder packets and identify duplicates.
C. To choose a routing protocol.
D. To compute RTT directly.
ANS: **B.** Sequence numbers let the receiver identify duplicates and restore packet order.

9. What is the main drawback of stop-and-wait?
A. It is incorrect if packets are reordered.
B. It requires no acks.
C. It achieves only one packet per RTT, making it slow.
D. It cannot use sequence numbers.
ANS: **C.** Stop-and-wait is correct but inefficient because throughput is limited to one packet per RTT.

10. What does a window-based algorithm allow the sender to do?
A. Send multiple packets before receiving all corresponding acknowledgments.
B. Eliminate the need for checksums.
C. Avoid using sequence numbers.
D. Deliver packets out of order to the application.
ANS: **A.** Windowing allows several packets to be in flight at once, which improves throughput.

11. To fill the pipe, Lecture 10 gives the approximate rule:
A. W = packet size / RTT.
B. W = RTT × bandwidth.
C. W = RTT + bandwidth.
D. W = congestion window + advertised window.
ANS: **B.** The lecture uses the bandwidth-delay product idea, expressed as W = RTT × bandwidth.

12. What does flow control protect against?
A. Overloading routers in the core network.
B. Overloading the recipient's buffer space.
C. Corrupting the checksum field.
D. Choosing the wrong Layer 4 protocol.
ANS: **B.** Flow control ensures the sender does not exceed the receiver's available buffering capacity.

13. What value does the receiver use to tell the sender how much more data it can accept?

A. Congestion window.
B. Advertised window.
C. TTL.
D. Fragment offset.
ANS: **B.** The advertised window is sent back by the receiver to show remaining buffer space.

14. What is the congestion window (cwnd)?
A. A receiver-side field that chooses the port number.
B. A sender-side output of the congestion-control algorithm that limits sending rate.
C. A checksum used only on duplicate acks.
D. A fixed constant equal to the MSS.
ANS: **B.** The congestion window is computed by the sender's congestion-control logic to avoid overloading the network.