RAT Test Run

Question 1: A logically centralised control plane means:

A that a logically central routing controller computes the forwarding table within a router.

B that a logically central control plane is implemented in each router.

C that a logically central forwarding table is computed in each router.

D that a logically central routing controller computes the forwarding table of a router and distributes it to the router.

Question 2: A global routing algorithm computes the least-cost path between a source and destination:

A by using complete, global knowledge about the network in each router or in a centralised control plane.

B in an iterative manner where no node has the complete information about the costs of all network links.

C by using complete, global knowledge about the network in a centralised data plane.

D in an iterative manner where a central control plane has the complete information about the costs of all network links.

Question 3: A distributed routing algorithm computes the least-cost path between a source and destination:

A by using complete, global knowledge about the network in each router or in a centralised control plane.

B by using complete, global knowledge about the network in a distributed data plane.

C in an iterative manner where a distributed control plane has the complete information about the costs of all network links.

D in an iterative manner where no node has the complete information about the costs of all network links.

Question 4: A dynamic routing algorithm changes the routing paths:

A when link cost or topology changes.

B when routing loops occur.

C when oscillation in routes is detected.

D for each newly created flow (e.g. a TCP flow).

Question 5: The count-to-infinity problem:

A can occur in both distance-vector and link-state routing.

B refers to a problem in distance-vector routing.

C refers to a problem in link-state routing.

D results from using the *poisoned reverse* technique.

Question 6: Comparing link-state (LS) with distance-vector (DV) algorithms we can say that:

A LS typically has lower message complexity.

B LS typically has a higher message complexity.

C DV converges faster.

D DV is more robust to erroneous routing information.

Question 7: Assume a network with three nodes: k, l and m. Suppose the initial link costs are $c(k,l) = 5$, $c(k,m) = 20$ and $c(l,m) = 3$ and that a distance-vector algorithm computes all the least-cost paths. If suddenly the cost of link (k,l) increases from 5 to 40:

A we will have a transient routing loop between m and l.

B we will have a transient routing loop between k and l.

C a new route from l through m to k will be immediately calculated.

D a new route from m to k via the direct link will be immediately calculated.

Question 8: Routing loops can result in:

A increased packet loss if the TTL of an IP datagram is reached.

B increased delay until the loop is removed but not in losses.

C increased routing information being sent by LS routing protocols.

D increased routing information being sent by DV routing protocols.

Question 9: The shortest hop-count path is the same as the least-cost path:

A always.

B for all paths with the same destination.

C for all paths with the same source.

D when all the edges of the graph have the same cost.

Question 10: Traceroute works as follows:

A it sends probe packets (often UDP to high, unused ports) with increasing TTL values; each router that discards a probe whose TTL reaches zero returns an ICMP Time Exceeded message; when the destination host is reached, it typically returns an ICMP Port Unreachable.

B it establishes a TCP connection to the destination and counts SYN/ACK responses from each hop.

C routers along the path reply with ICMP Echo Reply messages to each probe.

D routers forward probes even when the TTL is zero; only the destination returns ICMP Time Exceeded.