1. What does the Nagle algorithm state that the receiver should do when the window size becomes zero?
   (Ans: The receiver should wait until either half the receive-buffer is free or 1MSS is free.)
2. The user application gives the following data to TCP soon after the connection is established, with each data coming after 2ms gap: 10B, 200B, 300B, 1000B, 250B. The RTT is 10ms. How many segments are sent out by TCP if MSS = 1460B? What about if RTT = 20ms?
   (Ans: Using Nagle algorithm, after the first segment is sent, we wait until the ack. comes back before sending the data. Since RTT is 10ms and by the time 6ms from transmission of 10B is past, there is more than 1MSS of data, a segment of 1460B is sent. When the ack for 10B comes back, the remaining 290B is sent. If RTT is 20ms also, the same thing will happen.)
3. What are the slow start and congestion avoidance phases of TCP transmission?
   (Ans: Slow start is the first phase of data transmission on a TCP connection. When TCP encounters a retransmission timeout, it enters the congestion avoidance phase by reducing its congestion window to half the previous value.)
4. A receiver of a connection has advertised a window of 64MSS. The sender sends 40 MSS of data at one point of time. It finds that the retransmission timeout occurs at this point of time, what happens at this time? How does the sender proceed with its transmission?
   (Ans: The cwnd is reduced to 32MSS – since the cwnd is exponentially increased, it must have reached 64MSS for 40MSS of data to be transmitted without an ack. It enters congestion avoidance phase and therefore uses additive increase. It increases cwnd to 33MSS when all outstanding acks are received and so on.)
5. If we have congestion in a router, how does TCP congestion algorithm reduce congestion?
   (Ans: TCP uses the multiplicative decrease to reduce congestion immediately and then the additive increase to add data slowly to the network rather than at a shot.)
6. If we have congestion in a router, can it be reduced if the only traffic going through it is UDP traffic?
   (Ans: UDP is an unreliable, connectionless transport layer protocol that does not have any congestion control algorithm. So, unless the routers implement some congestion control at the network layer, there is nothing that can be done about congestion due to UDP traffic.)
7. What are interior and exterior gateway protocols?
   (Ans: Interior gateway protocols are those which are used within an Autonomous System (AS) whereas exterior are those used between ASes)
8. Give examples of interior and exterior gateway protocols.
   (Ans: RIP and OSPF for interior, BGP for exterior)
9. State the three rules of routing table updation in distance vector algorithm.
   (Ans: Add any new n/w, Update iff the newhc < oldhc, Update if the same neighbor advertises the same n/w with a different cost)
10. What is split horizon?
    (Ans: Do not advertise networks learnt on an interface back on the same interface)
11. What is the purpose of using poison reverse?
    (Ans: Help spread bad news fast)
12. Why does the link state protocol use the sequence number?
    (Ans: To prevent duplicate messages from being advt. as well as to ensure that stale information is discarded)
13. Why does a distance vector protocol have slow convergence?
    (Ans: Due to count-to-infinity problem)
14. What is hold-down timeout and how is it useful?
    (Ans: Do not update an entry with cost 16 until at least 2 cycles of periodic update is done. This allows for information about unreachability to reach all the network before

any alternate routes are learnt)

15. Give an example topology and explain the count-to-infinity problem.
16. State two major differences between RIPv1 and RIPv2.
    (Ans: RIPv1 uses bcast whereas V2 uses mcast. RIPv1 is for classful addressing only whereas V2 allows for CIDR)
17. What is the difference in the shortest path found by RIP and OSPF protocols?
    (Ans: RIP finds shortest paths in terms of hop count only whereas OSPF allows for any metric to be used for finding shortest paths as it uses a weighted graph)
18. What is the purpose of age timer in link state algorithm?
    (Ans: If a router reboots, its seq. no. is reset to 1. In such conditions, the routers do not update their network topology information until the seq. no. crosses the value before the router rebooted. This can lead to stale information about the network being maintained. Age timer deletes all entries from routers which are not updated within that time, thus allowing for quicker convergence)
19. The age timer is of duration 120s and a HELLO message in the link state algorithm is sent out every 8s. How many messages are discarded by a neighbor if its sequence number is 18?
    (Ans: After 15 messages (15*8=120s), the age timer deletes the entry and the entry is updated)
20. Why is RIP not currently used often in networks and is limited to enterprises only?
    (Ans: RIP is limited to small networks because 16 is considered infinity and also because of its slow convergence)

Given that the routing table entries in a router R1 are as follows:

Dest. Network	Gateway	Interface	Hop Cost
N0	Direct	Eth0	1
N1	Direct	Eth1	1

21. If router R1 receives an advertisement from a router with ID N1.R2 at time 0 – ((N2, 1), (N3, 2)) in DV algorithm, what is the topology of the network?
22. What are the new routing table entries after the above advertisement is received by R1?
    (Ans: (N2, N1.R2, eth1, 2), (N3, N1.R2, eth1, 3))
23. At 20s, R1 receives a new advertisement from N1.R2 with the following: (N2, 16). Which of the optimizations to deal with count-to-infinity problem are definitely enabled based on the two advertisements received?
    (Ans: Split horizon, poison reverse, triggered updates)
24. R1 receives an advertisement of ((N3, 1), (N2, 2)) from N0.R3 at time 50s after the two above advertisements. If hold-down timeout is enabled in R1, what will be the routing table entries in R1 after receiving this advertisement?
    (Ans: N3, N0.R3, eth0, 2) – N2 is not updated because of hold-down timeout being enabled)
25. If triggered updates are enabled in R1, when will the triggered update be sent from R1?
    (Ans: As soon as any entry is changed – so at 50s)
26. What is the network topology based on all the advertisements received so far?
27. A router Ri has received the following messages in the link-state algorithm: (R1, 1, (R2, R5)), (R2, 10, (R1, R3, R4)), (R3, 5, (R5, R2)), (R4, 2, (R2)), (R5, 1, (R1, R3)). What is the network topology?
    (Ans: Use the neighbor info. to construct the topology)
28. Given the topology below, what is the path for packets destined from host H1 to H2 given that link-state algorithm is used?

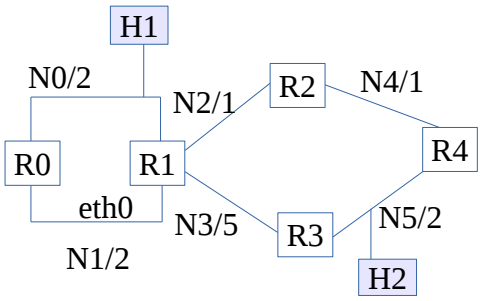

(Ans: H1->R1->R2->R4->H2)

29. In the topology below, will split horizon be sufficient to prevent count-to-infinity problem?

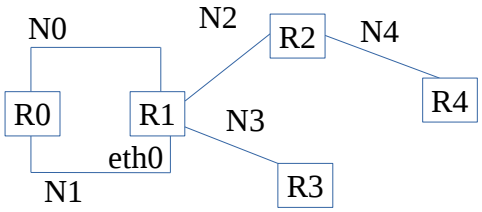

(Ans: Yes. All routers have information about N1 received only from R1 even though there is a cycle in the topology. And, there is no cycle back to R1 from the other routers which learn about N1 from R1)

30. If in the above topology only triggered updates are enabled, what will be the effect of eth0 going down? Will it lead to count-to-infinity problem or not?
(Ans: It can still lead to it as while R1 advt. N1 unreachability to R2, R3 may advt. To R1)
31. What is the problem if the entry is changed when the hop cost is equal to the cost in RT?
(Ans: Network oscillations leading to more bandwidth being consumed by control plane messages and a lot of churn in the paths followed by datagrams)
32. If in the below topology, R2's advt. reaches R1 first and then R3's advt., which one is used for the routing table entry to N2?
(Ans: Since the cost is the same, the first advt. is the one used)