Question 1

In IS-IS, which two statements are correct about the designated intermediate system (DIS) on a multi-access network segment? (Choose two)

Options:

A- A router with a priority of 10 wins the DIS election over a router with a priority of 1.

B- A router with a priority of 1 wins the DIS election over a router with a priority of 10.

C- On the multi-access network, each router forms an adjacency to every other router on the segment

D- On the multi-access network, each router only forms an adjacency to the DIS.

Answer:

A, D

Explanation:

In IS-IS, a designated intermediate system (DIS) is a router that is elected on a multi-access network segment (such as Ethernet) to perform some functions on behalf of other routers on the same segment. A DIS is responsible for sending network link-state

advertisements (LSPs), which describe all the routers attached to the network. These LSPs are flooded throughout a single area. A DIS also generates pseudonode LSPs, which represent the multi-access network as a single node in the link-state database. A DIS election is based on the priority value configured on each router's interface connected to the multi-access network. The priority value ranges from 0 to 127, with higher values indicating higher priority. The router with the highest priority becomes the DIS for the area (Level 1, Level 2, or both). If routers have the same priority, then the router with the highest MAC address is elected as the DIS. By default, routers have a priority value of 64. On a multi-access network, each router only forms an adjacency to the DIS, not to every other router on the segment. This reduces the amount of hello packets and LSP

Question 2

A packet is received on an interface configured with transmission scheduling. One of the configured queues In this scenario, which two actions will be taken by default on a Junos device? (Choose two.)

Options:

A- The excess traffic will be discarded

B- The exceeding queue will be considered to have negative bandwidth credit.

C- The excess traffic will use bandwidth available from other queueses

D- The exceeding queue will be considered to have positive bandwidth credit

Answer:

A, B

Explanation:

https://www.juniper.net/documentation/us/en/software/junos/cos-security-devices/topics/concept/cos-transmission-scheduling-security-overview.html

Question 3

Question Type: MultipleChoice

After a recent power outage, your manager asks you to investigate ways to automatically reduce the impact caused by suboptimal routing in your OSPF and OSPFv3 network after devices reboot.

Which three configuration statements accomplish this task? (Choose three.)

Options:

A- set protocols ospf3 realm ipv4-unicast overload timeout 900

B- set protocols ospf overload

C- set protocols ospf overload timeout 900

D- set protocols ospf3 overload

E- set protocols ospf3 overload timeout 900

Answer:

A, C, E

Explanation:

To reduce the impact of suboptimal routing in OSPF and OSPFv3 after devices reboot, you can use the overload feature to prevent a router from being used as a transit router for a specified period of time. This allows the router to stabilize its routing table before forwarding traffic for other routers. To enable the overload feature, you need to do the following:

For OSPF, configure the overload statement under [edit protocols ospf] hierarchy level. You can also specify a timeout value in seconds to indicate how long the router should remain in overload state after it boots up. For example, set protocols ospf overload timeout 900 means that the router will be in overload state for 15 minutes after it boots up.

For OSPFv3, configure the overload statement under [edit protocols ospf3] hierarchy level. You can also specify a realm (ipv4-unicast or ipv6-unicast) and a timeout value in seconds to indicate how long the router should remain in overload state after it boots up for each

realm. For example, set protocols ospf3 realm ipv4-unicast overload timeout 900 means that the router will be in overload state for 15 minutes after it boots up for IPv4 unicast routing.

Question 4

Question Type: MultipleChoice

Exhibit

External BGP

CE-1   PE-1
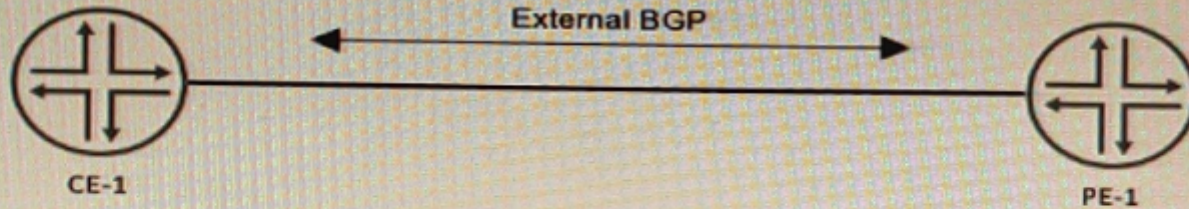
```
user@CE-1# show protocols bgp
group EBGP-to-PE-1 {
    type external;
    local-address 10.10.0.2;
    peer-as 65550;
    local-as 64511;
    neighbor 10.10.0.1 {
        export static-to-bgp;
    }
}


user@CE-1# show policy-options
policy-statement static-to-bgp {
    term export-static {
        from {
            protocol static;
            route-filter 192.168.1.0/24 exact;
            route-filter 192.168.2.0/24 exact;
            route-filter 192.168.3.0/24 exact;
            route-filter 192.168.4.0/24 exact;
            route-filter 192.168.5.0/24 exact;
            route-filter 192.168.6.0/24 exact;
            route-filter 192.168.7.0/24 exact;
            route-filter 192.168.8.0/24 exact;
            route-filter 192.168.9.0/24 exact;
            route-filter 192.168.10.0/24 exact;
        }
        then accept;
    }
    then reject;
}
```

```
user@PE-1# show protocols bgp group EBGP-to-CE-1
type external;
local-address 10.10.0.1;
peer-as 64511;
local-as 65550;
neighbor 10.10.0.2 {
    family inet {
        unicast {
            prefix-limit {
                maximum 5;
                teardown;
            }
        }
    }
}
```

CE-1 must advertise ten subnets to PE-1 using BGP Once CE-1 starts advertising the subnets to PE-1, the BGP peering state changes to Active.

Referring to the CLI output shown in the exhibit, which statement is correct?

Options:

A- CE-1 is advertising its entire routing table.

B- CE-1 is configured with an incorrect peer AS

C- The prefix limit has been reached on PE-1

D- CE-1 is unreachable

Answer:

B

Explanation:

The problem in this scenario is that CE-1 is configured with an incorrect peer AS number for its BGP session with PE-1. The CLI output shows that CE-1 is using AS 65531 as its local AS number and AS 65530 as its peer AS number. However, PE-1 is using AS 65530 as its local AS number and AS 65531 as its peer AS number. This causes a mismatch in the BGP OPEN messages and prevents the BGP session from being established. To solve this problem, CE-1 should configure its peer AS number as 65530 under [edit protocols bgp group external] hierarchy level.

Question 5

Which two statements are correct about VPLS tunnels? (Choose two.)

Options:

A- LDP-signaled VPLS tunnels only support control bit 0.

B- LDP-signaled VPLS tunnels use auto-discovery to provision sites

C- BGP-signaled VPLS tunnels can use either RSVP or LDP between the PE routers.

D- BGP-signaled VPLS tunnels require manual provisioning of sites.

Answer:

A, C

Explanation:

Question 6

Question Type: **MultipleChoice**

Which two statements about IS-IS are correct? (Choose two.)

Options:

A- CSNPs are flooded periodically.

B- PSNPs are flooded periodically.

C- PSNPs contain only descriptions of LSPs.

D- CSNPs contain only descriptions of LSPs.

Answer:

A, C

Explanation:

LSPs contain information about the state and cost of links in the network, and are flooded periodically throughout the network. PSNPs are used to acknowledge receipt of LSPs and request retransmission of missing or corrupted LSPs. PSNPs contain only descriptions of LSPs, such as their sequence numbers and checksums. CSNPs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases.