Practice Exam Questions

JUNIPER NETWORKS

JN0-664

Service Provider Routing and Switching, Professional (JNCIP-SP)

EXAMKILLER

Help Pass Your Exam At First Try

Juniper

Exam JN0-664

Service Provider Professional (JNCIP-SP)

Version: 3.0

[Total Questions: 65]

Question No : 1

Which two statements are correct about the customer interface in an LDP-signaled pseudowire? (Choose two)

A. When the encapsulation is vlan-ccc or extended-vlan-ccc, the configured VLAN tag is not included in the control plane LDP advertisement
B. When the encapsulation is ethernet-ccc, only frames without a VLAN tag are accepted in the data plane
C. When the encapsulation is vLan-ccc or extended-vlan-ccc, the configured VLAN tag is included in the control plane LDP advertisement
D. When the encapsulation is ethemet-ccc, tagged and untagged frames are both accepted in the data plane.

Answer: C,D

Explanation: The customer interface in an LDP-signaled pseudowire is the interface on the PE router that connects to the CE device. An LDP-signaled pseudowire is a type of Layer 2 circuit that uses LDP to establish a point-to-point connection between two PE routers over an MPLS network. The customer interface can have different encapsulation types depending on the type of traffic that is carried over the pseudowire. The encapsulation types are ethernet-ccc, vlan-ccc, extended-vlan-ccc, atm-ccc, frame-relay-ccc, ppp-ccc, cisco-hdlc-ccc, and tcc-ccc. Depending on the encapsulation type, the customer interface can accept or reject tagged or untagged frames in the data plane, and include or exclude VLAN tags in the control plane LDP advertisement. The following table summarizes the behavior of different encapsulation types:

Question No : 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.)

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: Transmission scheduling is a CoS feature that allows you to allocate bandwidth among different queues on an interface. Each queue has a configured bandwidth percentage that determines how much of the available bandwidth it can use. If a

queue exceeds its allocated bandwidth, it is considered to have negative bandwidth credit and its excess traffic will be discarded by default. If a queue does not use all of its allocated bandwidth, it is considered to have positive bandwidth credit and its unused bandwidth can be shared by other queues.

Question No : 3

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

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 No : 4

Exhibit

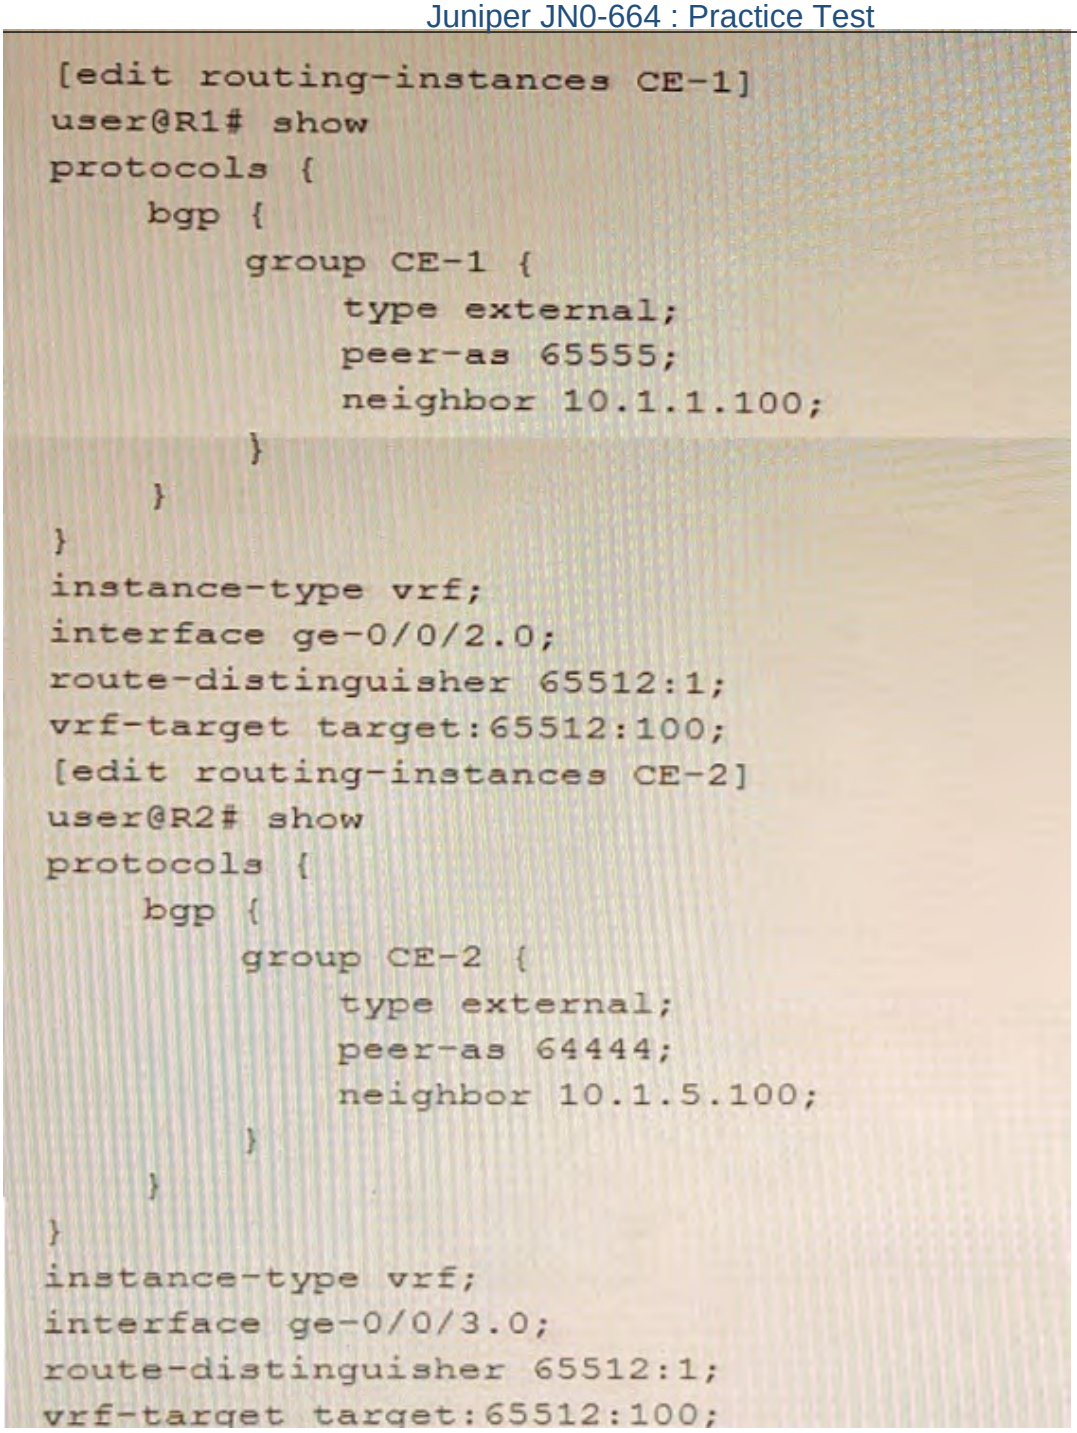

```
[edit routing-instances CE-1]
user@R1# show
protocols {
    bgp {
        group CE-1 {
            type external;
            peer-as 65555;
            neighbor 10.1.1.100;
        }
    }
}
instance-type vrf;
interface ge-0/0/2.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
[edit routing-instances CE-2]
user@R2# show
protocols {
    bgp {
        group CE-2 {
            type external;
            peer-as 64444;
            neighbor 10.1.5.100;
        }
    }
}
instance-type vrf;
interface ge-0/0/3.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
```

Referring to the exhibit, which statement is correct?

A. The vrf-target configuration will allow routes to be shared between CE-1 and CE-2.
B. The vrf-target configuration will stop routes from being shared between CE-1 and CE-2.
C. The route-distinguisher configuration will allow overlapping routes to be shared between CE-1 and CE-2.
D. The route-diatinguisher configuration will stop routes from being shared between CE-1 and CE-2.

Answer: C

Explanation: The route distinguisher (RD) is a BGP attribute that is used to create unique VPN IPv4 prefixes for each VPN in an MPLS network. The RD is a 64-bit value that

consists of two parts: an administrator field and an assigned number field. The administrator field can be an AS number or an IP address, and the assigned number field can be any arbitrary value chosen by the administrator. The RD is prepended to the IPv4 prefix to create a VPN IPv4 prefix that can be advertised across the MPLS network without causing any overlap or conflict with other VPNs. In this question, we have two PE routers (PE-1 and PE-2) that are connected to two CE devices (CE-1 and CE-2) respectively. PE-1 and PE-2 are configured with VRFs named Customer-A and Customer-B respectively.

Question No : 5

Exhibit

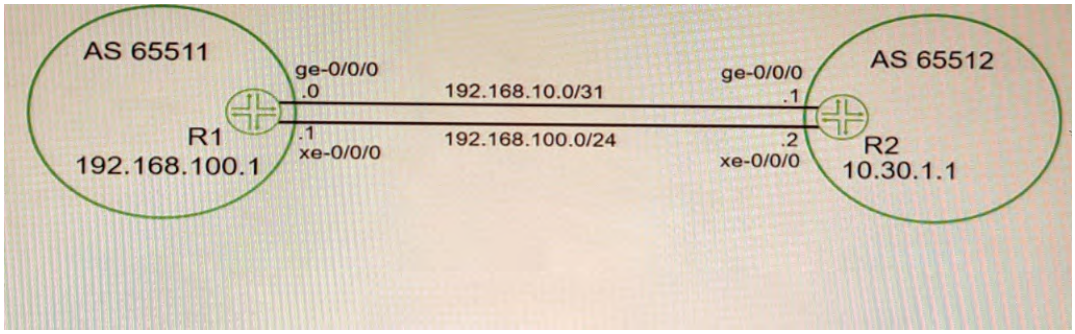

You want to use both links between R1 and R2 Because of the bandwidth difference between the two links, you must ensure that the links are used as much as possible.

Which action will accomplish this goal?

A. Define a policy to tag routes with the appropriate bandwidth community.
B. Disable multipath.
C. Ensure that the metric-out parameter on the Gigabit Ethernet interface is higher than the 10 Gigibit Ethernet interface.
D. Enable per-prefix load balancing.

Answer: D

Explanation: VPLS is a Layer 2 VPN technology that allows multiple sites to connect over a shared IP/MPLS network as if they were on the same LAN. VPLS tunnels can be signaled using either Label Distribution Protocol (LDP) or Border Gateway Protocol (BGP). In this question, we have two links between R1 and R2 with different bandwidths (10 Gbps and 1 Gbps). We want to use both links as much as possible for VPLS traffic. To achieve this, we need to enable per-prefix load balancing on both routers. Per-prefix load balancing is a feature that allows a router to distribute traffic across multiple equal-cost or unequal-cost paths based on the destination prefix of each packet. This improves the utilization of

multiple links and provides better load sharing than per-flow load balancing, which distributes traffic based on a hash of source and destination addresses**4**. Per-prefix load balancing can be enabled globally or per interface using the load-balance per-packet command.
Reference: **4**: https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/137544-technote-mpls-00.html

Question No : 6

Exhibit

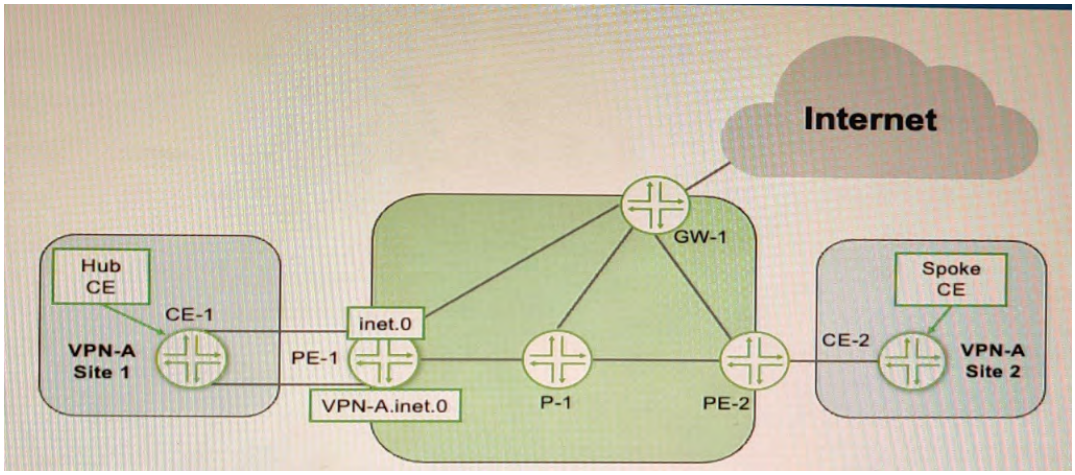

Referring to the exhibit, you must provide Internet access for VPN-A using CE-1 as the hub CE.

Which two statements are correct in this situation? (Choose two.)

A. You must use RIB groups to leak routes between the inet. o and vpn-a. inet. o tables.
B. RIB groups are not needed to leak routes between the inet. 0 and VPN—A. inet. 0 tables,
C. Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> GW-1.
D. Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> CE-1 -> PE-1 -> GW-1.

Answer: A,D

Explanation: To provide Internet access for VPN-A using CE-1 as the hub CE, you need to do the following:

- You must use RIB groups to leak routes between the inet.0 and vpn-a.inet.0 tables on PE-1 and CE-1. RIB groups are routing options that allow you to import routes from one routing table into another routing table based on certain criteria. In this scenario, you need to configure RIB groups on PE-1 and CE-1 to import Internet routes from inet.0 into vpn-a.inet.0 and vice versa.
- Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> CE-1 -> PE-1 -> GW-

1. This is because Site 2 does not have direct Internet access and needs to use CE-1 as its default gateway for Internet traffic. Site 2 sends its Internet traffic to PE-2, which forwards it to PE-1 based on VPN-A routes. PE-1 then sends it to CE-1 based on RIB group import policy. CE-1 then sends it back to PE-1 based on its default route pointing to GW-1. PE-1 then forwards it to GW-1 based on RIB group import policy again.

Question No : 7

Exhibit

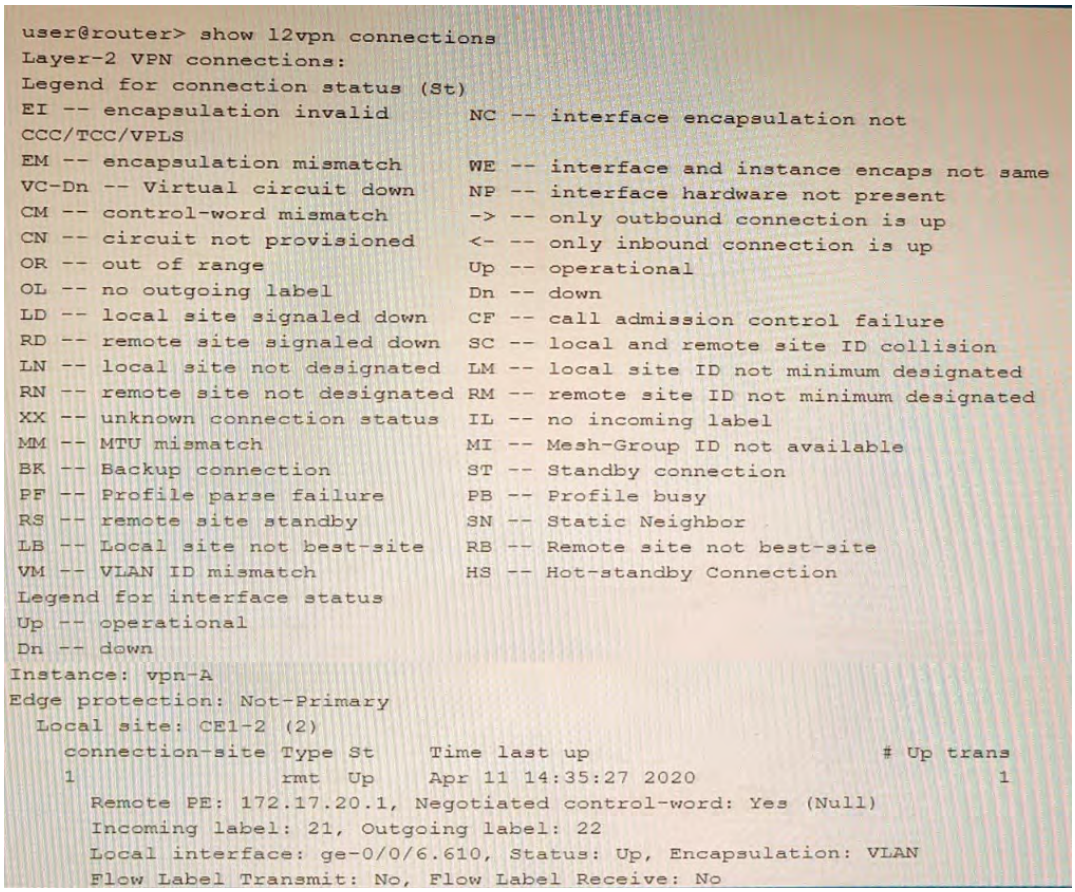

```
user@router> show l2vpn connections
Layer-2 VPN connections:
Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not
CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch           HS -- Hot-standby Connection
Legend for interface status
Up -- operational
Dn -- down
Instance: vpn-A
Edge protection: Not-Primary
  Local site: CE1-2 (2)
    connection-site Type St     Time last up                # Up trans
    1               rmt  Up     Apr 11 14:35:27 2020                1
      Remote PE: 172.17.20.1, Negotiated control-word: Yes (Null)
      Incoming label: 21, Outgoing label: 22
      Local interface: ge-0/0/6.610, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
```

Which two statements about the output shown in the exhibit are correct? (Choose two.)

A. The PE is attached to a single local site.
B. The connection has not flapped since it was initiated.
C. There has been a VLAN ID mismatch.
D. The PE router has the capability to pop flow labels

Answer: A,D

Explanation:

According to **1** and **2**, BGP Layer 2 VPNs use BGP to distribute endpoint provisioning information and set up pseudowires between PE devices. BGP uses the Layer 2 VPN (L2VPN) Routing Information Base (RIB) to store endpoint provisioning information, which is updated each time any Layer 2 virtual forwarding instance (VFI) is configured. The prefix and path information is stored in the L2VPN database, which allows BGP to make decisions about the best path.

In the output shown in the exhibit, we can see some information about the L2VPN RIB and the pseudowire state. Based on this information, we can infer the following statements:

- The PE is attached to a single local site. This is correct because the output shows only one local site ID (1) under the L2VPN RIB section. A local site ID is a unique identifier for a site within a VPLS domain. If there were multiple local sites attached to the PE, we would see multiple local site IDs with different prefixes.
- The connection has not flapped since it was initiated. This is correct because the output shows that the uptime of the pseudowire is equal to its total uptime (1w6d). This means that the pseudowire has been up for one week and six days without any interruption or flap.
- There has been a VLAN ID mismatch. This is not correct because the output shows that the remote and local VLAN IDs are both 0 under the pseudowire state section. A VLAN ID mismatch occurs when the remote and local VLAN IDs are different, which can cause traffic loss or misdelivery. If there was a VLAN ID mismatch, we would see different values for the remote and local VLAN IDs.
- The PE router has the capability to pop flow labels. This is correct because the output shows that the flow label pop bit is set under the pseudowire state section. The flow label pop bit indicates that the PE router can pop (remove) the MPLS flow label from the packet before forwarding it to the CE device. The flow label is an optional MPLS label that can be used for load balancing or traffic engineering purposes.

Question No : 8

Exhibit

```
user@R4> show pim rps
Instance: PIM.master
address-family INET
RP address       Type          Mode   Holdtime Timeout Groups Group prefixes
10.1.255.2       bootstrap     sparse      150     118      0 224.1.1.0/24
10.1.255.3       bootstrap     sparse      150     118      2 224.1.1.0/28
user@R4> show route 10.1.255.2
inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.1.255.2/32      *[IS-IS/18] 00:32:27, metric 10
                     >  to 10.1.1.2 via ge-0/0/0.0
inet.2: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0           *[Static/5] 00:13:55
                     >  to 10.1.1.6 via ge-0/0/1.0
user@R4> show route 10.1.255.3

inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.1.255.3/32      *[IS-IS/18] 00:32:43, metric 10
                     >  to 10.1.1.6 via ge-0/0/1.0
inet.2: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0           *[Static/5] 00:14:25
                     >  to 10.1.1.6 via ge-0/0/1.0
[edit]
user@R2# show protocols pim
rp {
    bootstrap {
        family inet {
            priority 200;
        }
    }
    local {
        address 10.1.255.2;
        group-ranges {
            224.1.1.0/24;
        }
    }
}
interface all;
[edit]
user@R3# show protocols pim
rp {
    bootstrap {
        family inet {
            priority 210;
        }
    }
    local {
        address 10.1.255.3;
        group-ranges {
            224.1.1.0/28;
        }
    }
}
interface all;
```

R4 is directly connected to both RPs (R2 and R3) R4 is currently sending all ,o,ns upstream to R3 but you want all joins to go to R2 instead Referring to the exhibit, which configuration change will solve this issue?

A. Change the bootstrap priority on R2 to be higher than R3
B. Change the default route in inet.2 on R4 from R3 as the next hop to R2
C. Change the local address on R2 to be higher than R3.
D. Change the group-range to be more specific on R2 than R3.

Answer: A

Explanation: PIM Bootstrap Router (BSR) is a mechanism that allows PIM routers to discover and announce rendezvous point (RP) information for multicast groups. BSR uses two roles: candidate BSR and candidate RP. Candidate BSR is the router that collects information from all available RPs in the network and advertises it throughout the network. Candidate RP is the router that wants to become the RP and registers itself with the BSR. There can be only one active BSR in the network, which is elected based on the highest priority or highest IP address if the priority is the same. The BSR priority can be configured manually or assigned automatically. The default priority is 0 and the highest priority is 255**1** . In this question, R4 is directly connected to both RPs (R2 and R3) and is currently sending all joins upstream to R3 but we want all joins to go to R2 instead. To achieve this, we need to change the BSR priority on R2 to be higher than R3 so that R2 becomes the active BSR and advertises its RP information to R4.
Reference: **1**: https://study-ccnp.com/multicast-rendezvous-points-explained/

Question No : 9

Exhibit

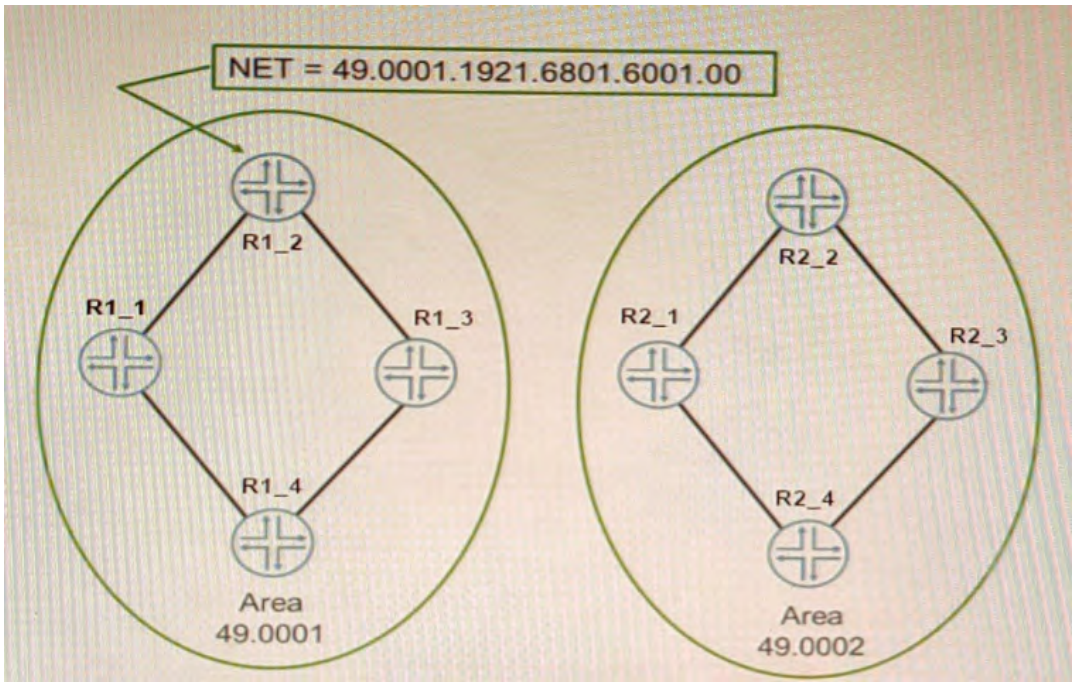

The network shown in the exhibit is based on IS-IS

Which statement is correct in this scenario?

A. The NSEL byte for Area 0001 is 00.
B. The area address is two bytes.

C. The routers are using unnumbered interfaces
D. The system IDofR1_2 is 192.168.16.1

Answer: A

Explanation: IS-IS is an interior gateway protocol that uses link-state routing to exchange routing information among routers within a single autonomous system. IS-IS uses two types of addresses to identify routers and areas: system ID and area address. The system ID is a unique identifier for each router in an IS-IS domain. The system ID is 6 octets long and can be derived from the MAC address or manually configured. The area address is a variable-length identifier for each area in an IS-IS domain. The area address can be 1 to 13 octets long and is composed of high-order octets of the address. An IS-IS instance may be assigned multiple area addresses, which are considered synonymous. Multiple synonymous area addresses are useful when merging or splitting areas in the domain**1**. In this question, we have a network based on IS-IS with four routers (R1_1, R1_2, R2_1, and R2_2) belonging to area 0001. The area address for area 0001 is 49.0001. The NSEL byte for area 0001 is the last octet of the address, which is 01. The NSEL byte stands for Network Service Access Point Selector (NSAP Selector) and indicates the type of service requested from the network layer**2**. Therefore, the correct statement in this scenario is that the NSEL byte for area 0001 is 01.
References: **1**: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/configuration/xe-16/irs-xe-16-book/irs-ovrvw-cf.html **2**: https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing-overview.html

Question No : 10

Exhibit

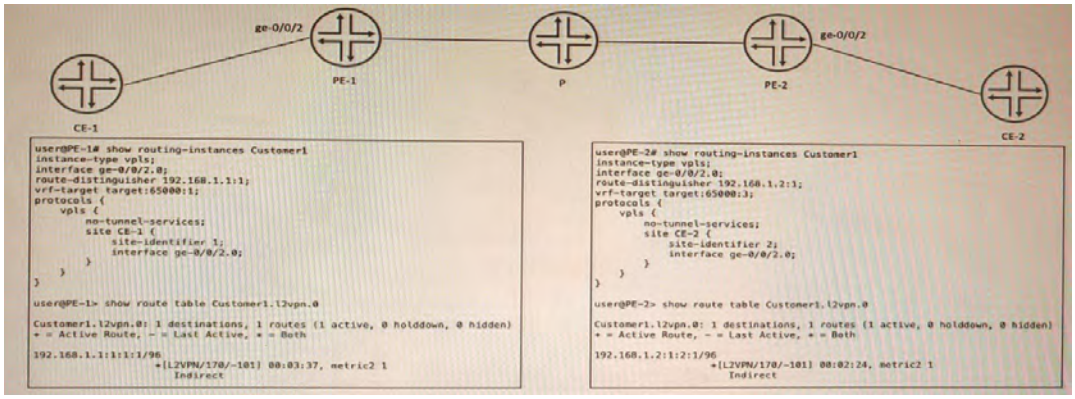

CE-1 and CE-2 are part of a VPLS called Customer1 No connectivity exists between CE-1 and CE-2. In the process of troubleshooting, you notice PE-1 is not learning any routes for

this VPLS from PE-2, and PE-2 is not learning any routes for this VPLS from PE-1.

A. The route target must match on PE-1 and PE-2.
B. The route distinguisher must match on PE-1 and PE-2.
C. The instance type should be changed to I2vpn.
D. The no-tunnel-services statement should be deleted on both PEs.

Answer: A

Explanation: VPLS is a technology that provides Layer 2 VPN services over an MPLS network. VPLS uses BGP as its control protocol to exchange VPN membership information between PE routers. The route target is a BGP extended community attribute that identifies which VPN a route belongs to. The route target must match on PE routers that participate in the same VPLS instance, otherwise they will not accept or advertise routes for that VPLS.

Question No : 11

In which two ways does OSPF prevent routing loops in multi-area networks? (Choose two.)

A. All areas are required to connect as a full mesh.
B. The LFA algorithm prunes all looped paths within an area.
C. All areas are required to connect to area 0.
D. The SPF algorithm prunes looped paths within an area.

Answer: C,D

Explanation: OSPF is an interior gateway protocol that uses link-state routing to exchange routing information among routers within a single autonomous system. OSPF prevents routing loops in multi-area networks by using two methods: area hierarchy and SPF algorithm. Area hierarchy is the concept of dividing a large OSPF network into smaller areas that are connected to a backbone area (area 0). This reduces the amount of routing information that each router has to store and process, and also limits the scope of link-state updates within each area. All areas are required to connect to area 0 either directly or through virtual links**2**. SPF algorithm is the method that OSPF uses to calculate the shortest path to each destination in the network based on link-state information. The SPF algorithm runs on each router and builds a shortest-path tree that represents the topology of the network from the router's perspective. The SPF algorithm prunes looped paths within an area by choosing only one best path for each destination**3**.
References: **2**:
https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/concept/ospf-area-overview.html **3**:

https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/concept/ospf-spf-algorithm-overview.html

Question No : 12

Exhibit

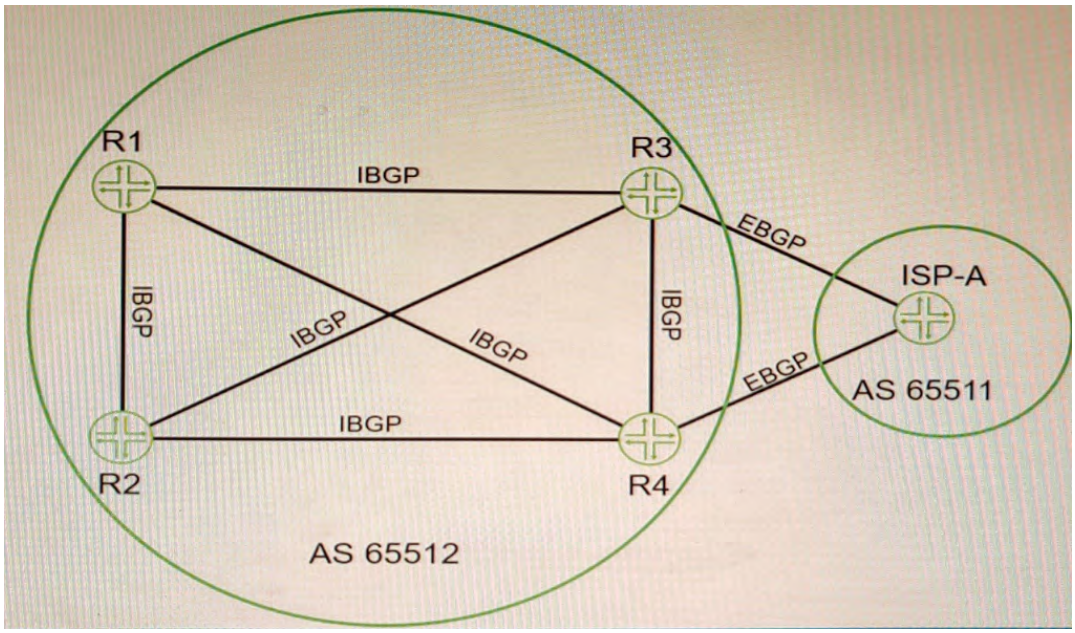

Click the Exhibit button-Referring to the exhibit, which two statements are correct about BGP routes on R3 that are learned from the ISP-A neighbor? (Choose two.)

A. By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3.
B. The BGP local-preference value that is used by ISP-A is not advertised to R3.
C. All BGP attribute values must be removed before receiving the routes.
D. The next-hop value for these routes is changed by ISP-A before being sent to R3.

Answer: A,B

Explanation:

BGP is an exterior gateway protocol that uses path vector routing to exchange routing information among autonomous systems. BGP uses various attributes to select the best path to each destination and to propagate routing policies. Some of the common BGP attributes are AS path, next hop, local preference, MED, origin, weight, and community. BGP attributes can be classified into four categories: well-known mandatory, well-known discretionary, optional transitive, and optional nontransitive. Well-known mandatory attributes are attributes that must be present in every BGP update message and must be recognized by every BGP speaker. Well-known discretionary attributes are attributes that

may or may not be present in a BGP update message but must be recognized by every BGP speaker. Optional transitive attributes are attributes that may or may not be present in a BGP update message and may or may not be recognized by a BGP speaker. If an optional transitive attribute is not recognized by a BGP speaker, it is passed along to the next BGP speaker. Optional nontransitive attributes are attributes that may or may not be present in a BGP update message and may or may not be recognized by a BGP speaker. If an optional nontransitive attribute is not recognized by a BGP speaker, it is not passed along to the next BGP speaker. In this question, we have four routers (R1, R2, R3, and R4) that are connected in a full mesh topology and running IBGP. R3 receives the 192.168.0.0/16 route from its EBGP neighbor and advertises it to R1 and R4 with different BGP attribute values. We are asked which statements are correct about the BGP routes on R3 that are learned from the ISP-A neighbor. Based on the information given, we can infer that the correct statements are:

- By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3. This is because the default behavior of EBGP is to preserve the next-hop attribute of the routes received from another EBGP neighbor. The next-hop attribute indicates the IP address of the router that should be used as the next hop to reach the destination network.
- The BGP local-preference value that is used by ISP-A is not advertised to R3. This is because the local-preference attribute is a well-known discretionary attribute that is used to influence the outbound traffic from an autonomous system. The local-preference attribute is only propagated within an autonomous system and is not advertised to external neighbors.

References: : https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html : https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13762-40.html : https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13759-37.html

Question No : 13

Exhibit

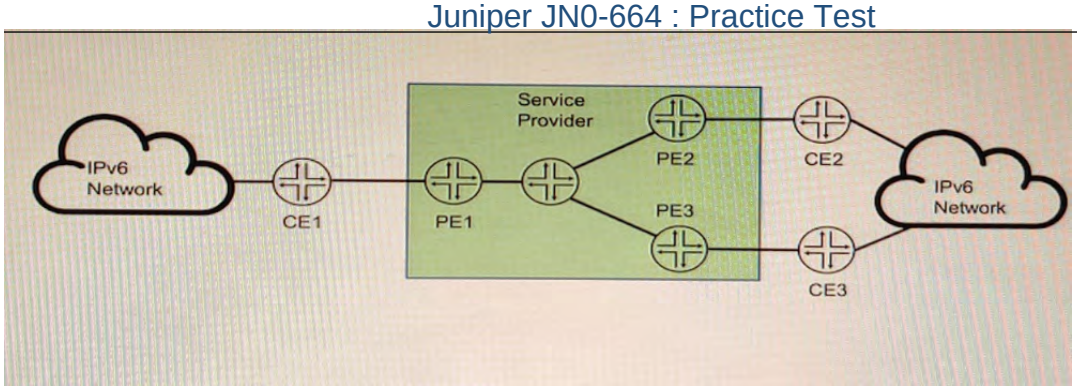

You are running a service provider network and must transport a customer's IPv6 traffic across your IPv4-based MPLS network using BGP You have already configured mpis ipv6-tunneling on your PE routers.

Which two statements are correct about the BGP configuration in this scenario? (Choose two.)

A. You must configure family inet6 labcled-unicast between PE routers.
B. You must configure family inet6 unicaat between PE and CE routers.
C. You must configure family inet6 add-path between PE and CE routers.
D. You must configure family inet6 unicast between PE routers

Answer: A,B

Explanation: To transport IPv6 traffic over an IPv4-based MPLS network using BGP, you need to configure two address families: family inet6 labeled-unicast and family inet6 unicast. The former is used to exchange IPv6 routes with MPLS labels between PE routers, and the latter is used to exchange IPv6 routes without labels between PE and CE routers. The mpis ipv6-tunneling command enables the PE routers to encapsulate the IPv6 packets with an MPLS label stack and an IPv4 header before sending them over the MPLS network.

Question No : 14

You are asked to protect your company's customers from amplification attacks. In this scenario, what is Juniper's recommended protection method?

A. ASN prepending
B. BGP FlowSpec
C. destination-based Remote Triggered Black Hole
D. unicast Reverse Path Forwarding

Answer: C

Explanation: amplification attacks are a type of distributed denial-of-service (DDoS) attack

that exploit the characteristics of certain protocols to amplify the traffic sent to a victim. For example, an attacker can send a small DNS query with a spoofed source IP address to a DNS server, which will reply with a much larger response to the victim. This way, the attacker can generate a large amount of traffic with minimal resources.
One of the methods to protect against amplification attacks is destination-based Remote Triggered Black Hole (RTBH) filtering. This technique allows a network operator to drop traffic destined to a specific IP address or prefix at the edge of the network, thus preventing it from reaching the victim and consuming bandwidth and resources. RTBH filtering can be implemented using BGP to propagate a special route with a next hop of 192.0.2.1 (a reserved address) to the edge routers. Any traffic matching this route will be discarded by the edge routers.

Question No : 15

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

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: B,C

Explanation: VPLS is a Layer 2 VPN technology that allows multiple sites to connect over a shared IP/MPLS network as if they were on the same LAN. VPLS tunnels can be signaled using either Label Distribution Protocol (LDP) or Border Gateway Protocol (BGP). LDP-signaled VPLS tunnels use auto-discovery to provision sites, meaning that PE routers can automatically discover other PE routers that belong to the same VPLS instance

Question No : 16

Exhibit

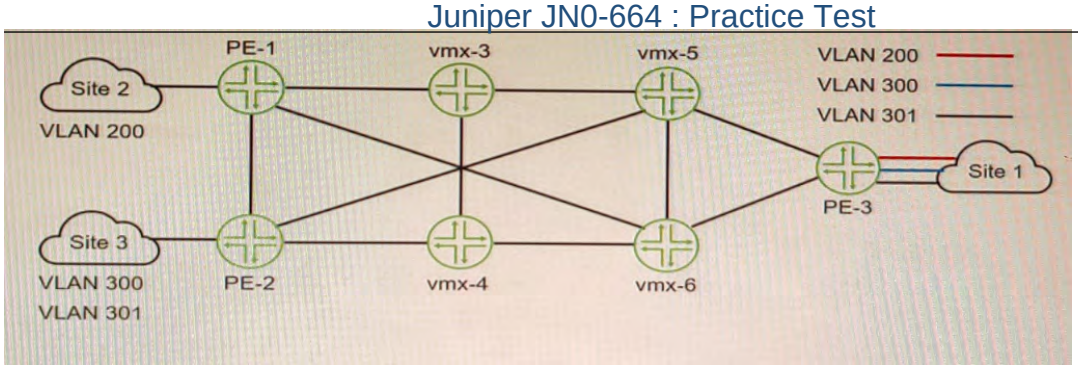

You want Site 1 to access three VLANs that are located in Site 2 and Site 3 The customer-facing interface on the PE-1 router is configured for Ethernet-VLAN encapsulation.

What is the minimum number of L2VPN routing instances to be configured to accomplish this task?

A. 1
B. 3
C. 2
D. 6

Answer: B

Explanation: To allow Site 1 to access three VLANs that are located in Site 2 and Site 3, you need to configure three L2VPN routing instances on PE-1, one for each VLAN. Each L2VPN routing instance will have a different VLAN ID and a different VNI for VXLAN encapsulation. Each L2VPN routing instance will also have a different vrf-target export value to identify which VPN routes belong to which VLAN. This way, PE-1 can forward traffic from Site 1 to Site 2 and Site 3 based on the VLAN tags and VNIs.

Question No : 17

An interface is configured with a behavior aggregate classifier and a multifield classifier How will the packet be processed when received on this interface?

A. The packet will be discarded.
B. The packet will be processed by the BA classifier first, then the MF classifier.
C. The packet will be forwarded with no classification changes.
D. The packet will be processed by the MF classifier first, then the BA classifier.

Answer: C

Explanation: behavior aggregate (BA) classifiers and multifield (MF) classifiers are two types of classifiers that are used to assign packets to a forwarding class and a loss priority based on different criteria. The forwarding class determines the output queue for a packet.

The loss priority is used by a scheduler to control packet discard during periods of congestion.

A BA classifier maps packets to a forwarding class and a loss priority based on a fixed-length field in the packet header, such as DSCP, IP precedence, MPLS EXP, or IEEE 802.1p CoS bits. A BA classifier is computationally efficient and suitable for core devices that handle high traffic volumes. A BA classifier is useful if the traffic comes from a trusted source and the CoS value in the packet header is trusted.

An MF classifier maps packets to a forwarding class and a loss priority based on multiple fields in the packet header, such as source address, destination address, protocol type, port number, or VLAN ID. An MF classifier is more flexible and granular than a BA classifier and can match packets based on complex filter rules. An MF classifier is suitable for edge devices that need to classify traffic from untrusted sources or rewrite packet headers.

You can configure both a BA classifier and an MF classifier on an interface. If you do this, the BA classification is performed first and then the MF classification. If the two classification results conflict, the MF classification result overrides the BA classification result.

Based on this information, we can infer the following statements:

- The packet will be discarded. This is not correct because the packet will not be discarded by the classifiers unless it matches a filter rule that specifies discard as an action. The classifiers only assign packets to a forwarding class and a loss priority based on their match criteria.
- The packet will be processed by the BA classifier first, then the MF classifier. This is correct because if both a BA classifier and an MF classifier are configured on an interface, the BA classification is performed first and then the MF classification. If they conflict, the MF classification result overrides the BA classification result.
- The packet will be forwarded with no classification changes. This is not correct because the packet will be classified by both the BA classifier and the MF classifier if they are configured on an interface. The final classification result will determine which output queue and which discard policy will be applied to the packet.
- The packet will be processed by the MF classifier first, then the BA classifier. This is not correct because if both a BA classifier and an MF classifier are configured on an interface, the BA classification is performed first and then the MF classification. If they conflict, the MF classification result overrides the BA classification result.

Question No : 18

When building an interprovider VPN, you notice on the PE router that you have hidden routes which are received from your BGP peer with family inet labeled-unica3t configured.

Which parameter must you configure to solve this problem?

A. Under the family inet labeled-unicast hierarchy, add the explicit null parameter.
B. Under the protocols ospf hierarchy, add the traffic-engineering parameter.
C. Under the family inet labeled-unicast hierarchy, add the resolve-vpn parameter.
D. Under the protocols mpls hierarchy, add the traffic-engineering parameter

Answer: C

Explanation: The resolve-vpn parameter is a BGP option that allows a router to resolve labeled VPN-IPv4 routes using unlabeled IPv4 routes received from another BGP peer with family inet labeled-unicast configured. This option enables interprovider VPNs without requiring MPLS labels between ASBRs or using VRF tables on ASBRs. In this scenario, you need to configure the resolve-vpn parameter under [edit protocols bgp group external family inet labeled-unicast] hierarchy level on both ASBRs.

Question No : 19

You are configuring a BGP signaled Layer 2 VPN across your MPLS enabled core network. Your PE-2 device connects to two sites within the s VPN

In this scenario, which statement is correct?

A. By default on PE-2, the site's local ID is automatically assigned a value of 0 and must be configured to match the total number of attached sites.
B. You must create a unique Layer 2 VPN routing instance for each site on the PE-2 device.
C. You must use separate physical interfaces to connect PE-2 to each site.
D. By default on PE-2, the remote site IDs are automatically assigned based on the order that you add the interfaces to the site configuration.

Answer: D

Explanation: BGP Layer 2 VPNs use BGP to distribute endpoint provisioning information and set up pseudowires between PE devices. BGP uses the Layer 2 VPN (L2VPN) Routing Information Base (RIB) to store endpoint provisioning information, which is updated each time any Layer 2 virtual forwarding instance (VFI) is configured. The prefix and path information is stored in the L2VPN database, which allows BGP to make decisions about the best path.
In BGP Layer 2 VPNs, each site has a unique site ID that identifies it within a VFI. The site ID can be manually configured or automatically assigned by the PE device. By default, the site ID is automatically assigned based on the order that you add the interfaces to the site configuration. The first interface added to a site configuration has a site ID of 1, the second interface added has a site ID of 2, and so on.