



# Juniper

## Exam Questions JN0-664

Service Provider - Professional (JNCIP-SP)

## About ExamBible

*[Your Partner of IT Exam](#)*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

### NEW QUESTION 1

You are configuring a BGP signaled Layer 2 VPN across your MPLS enabled core network. In this scenario, which statement is correct?

- A. You must assign a unique site number to each attached site's configuration.
- B. This type of VPN only supports Ethernet interfaces when connecting to CE devices.
- C. This type of VPN requires the support of the inet-vpn NLRI on all core BGP devices
- D. You must use the same route-distinguisher value on both PE devices.

**Answer: C**

#### Explanation:

BGP signaled Layer 2 VPN is a type of VPN that uses BGP to distribute VPN labels and information for Layer 2 connectivity between sites over an MPLS network. BGP signaled Layer 2 VPN requires the support of the l2vpn NLRI on all core BGP devices<sup>1</sup>. The l2vpn NLRI is a new address family that carries Layer 2 VPN information such as the VPN identifier, the attachment circuit identifier, and the route distinguisher. The l2vpn NLRI is used for both auto-discovery and signaling of Layer 2 VPNs<sup>2</sup>. In this scenario, we are configuring a BGP signaled Layer 2 VPN across an MPLS enabled core network. Therefore, we need to ensure that all core BGP devices support the l2vpn NLRI. References: 1: <https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/topics/concept/vpn-layer-2-overview.html> 2: [https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp\\_l2\\_vpns/configuration/xr-16/mp-l2-vpns-xr-16-book/vpls-bgp-signaling-l2vpn-inter-as-option-a.html](https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l2_vpns/configuration/xr-16/mp-l2-vpns-xr-16-book/vpls-bgp-signaling-l2vpn-inter-as-option-a.html)

### NEW QUESTION 2

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-unicast 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.

### NEW QUESTION 3

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**

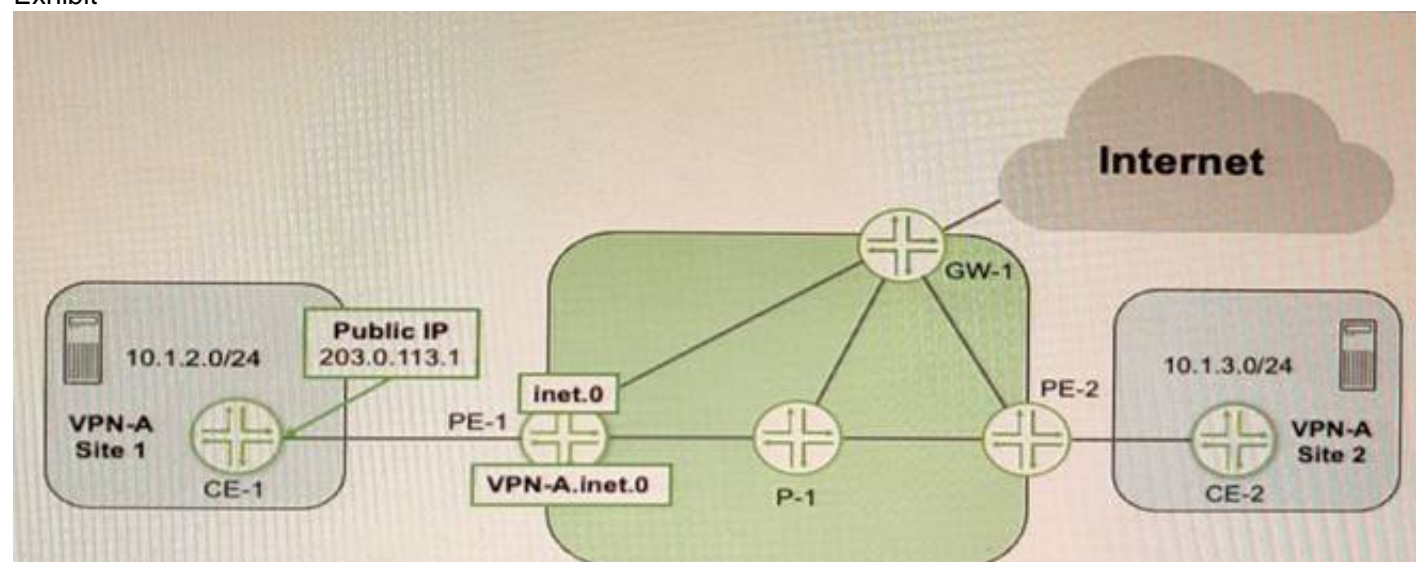
#### Explanation:

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.

### NEW QUESTION 4

Exhibit



Referring to the exhibit, CE-1 is providing NAT services for the hosts at Site 1 and you must provide Internet access for those hosts. Which two statements are correct in this scenario? (Choose two.)

- A. You must configure a static route in the main routing instance for the 10.1.2.0/24 prefix that uses the VPN-A.inet.0 table as the next hop.
- B. You must configure a static route in the main routing instance for the 203.0.113.1/32 prefix that uses the VPN-A.inet.0 table as the next hop.
- C. You must configure a RIB group on PE-1 to leak a default route from the inet.0 table to the VPN-A.inet.0 table.
- D. You must configure a RIB group on PE-1 to leak the 10.1.2.0/24 prefix from the VPN-A.inet.0 table to the inet.0 table.

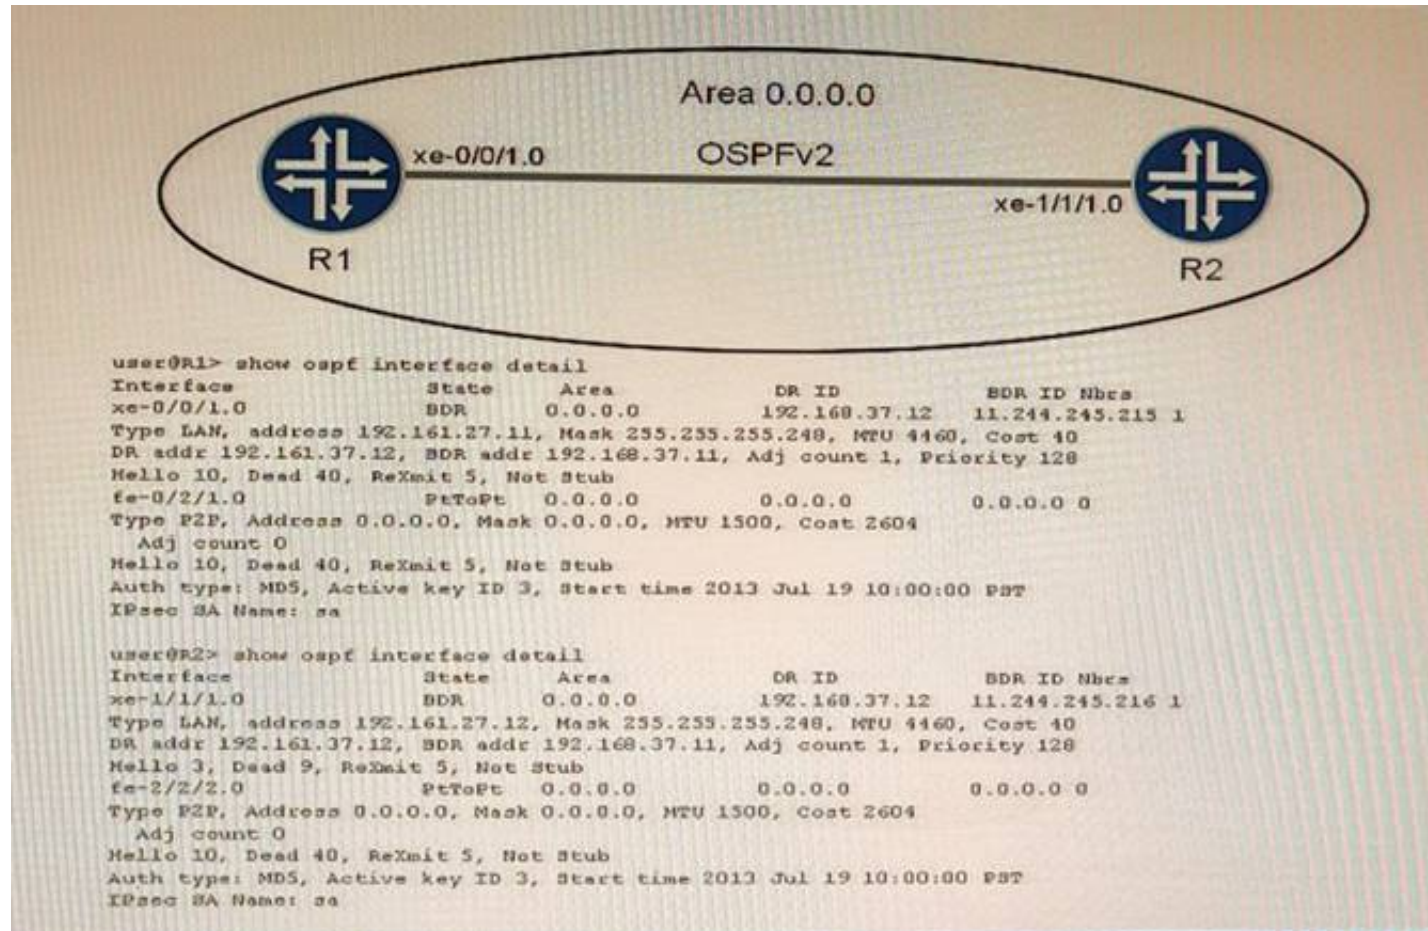
**Answer:** AB

**Explanation:**

To provide Internet access for the hosts at Site 1, you need to configure static routes in the main routing instance on PE-1 that point to the VPN-A.inet.0 table as the next hop. This allows PE-1 to forward traffic from the Internet to CE-1 using MPLS labels and vice versa. You need to configure two static routes: one for the 10.1.2.0/24 prefix that represents the private network of Site 1, and one for the 203.0.113.1/32 prefix that represents the public IP address of CE-1.

**NEW QUESTION 5**

Exhibit



Which two statements are true about the OSPF adjacency displayed in the exhibit? (Choose two.)

- A. There is a mismatch in the hello interval parameter between routers R1 and R2.
- B. There is a mismatch in the dead interval parameter between routers R1 and R2.
- C. There is a mismatch in the OSPF hold timer parameter between routers R1 and R2.
- D. There is a mismatch in the poll interval parameter between routers R1 and R2.

**Answer:** AB

**Explanation:**

The hello interval is the time interval between two consecutive hello packets sent by an OSPF router on an interface. The dead interval is the time interval after which a neighbor is declared down if no hello packets are received from it. These parameters must match between two OSPF routers for them to form an adjacency. In the exhibit, router R1 has a hello interval of 10 seconds and a dead interval of 40 seconds, while router R2 has a hello interval of 30 seconds and a dead interval of 120 seconds. This causes a mismatch and prevents them from becoming neighbors.

**NEW QUESTION 6**

Which two statements describe PIM-SM? (Choose two)

- A. Routers with receivers send join messages to their upstream neighbors.
- B. Routers without receivers must periodically prune themselves from the SPT.
- C. Traffic is initially flooded to all routers and an S,G is maintained for each group.
- D. Traffic is only forwarded to routers that request to join the distribution tree.

**Answer:** AD


**Explanation:**

PIM sparse mode (PIM-SM) is a multicast routing protocol that uses a pull model to deliver multicast traffic. In PIM-SM, routers with receivers send join messages to their upstream neighbors toward a rendezvous point (RP) or a source-specific tree (SPT). The RP or SPT acts as the root of a shared distribution tree for a multicast group. Traffic is only forwarded to routers that request to join the distribution tree by sending join messages. PIM-SM does not flood traffic to all routers or prune routers without receivers, as PIM dense mode does.

**NEW QUESTION 7**

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.

**NEW QUESTION 8**

Exhibit.

**Exhibit**

```

user@R1# show interfaces
ge-1/2/3 {
  unit 0 {
    description to-R2;
    family inet {
      address 10.1.1.1/30;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.16.1/32;
    }
    family iso {
      address 49.0001.1921.6801.6001.00;
    }
  }
}
user@R1# show protocols
isis {
  interface ge-1/2/3.0 {
    level 2 disable;
  }
}
...
user@R2# show interfaces
ge-1/2/3 {
  unit 0 {
    description to-R1;
    family inet {
      address 10.1.1.2/30;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.16.2/32;
    }
    family iso {
      address 49.0001.1921.6801.6002.00;
    }
  }
}
user@R2# show protocols
isis {
  interface ge-1/2/3.0 {
    level 1 disable;
  }
  interface lo0.0 {
    level 1 disable;
  }
}

```

Referring to the exhibit, what must be changed to establish a Level 1 adjacency between routers R1 and R2?

- A. Change the level 1 disable parameter under the R1 protocols isis interface lo0.0 hierarchy to the level 2 disable parameter.
- B. Remove the level 1 disable parameter under the R2 protocols isis interface lo0.0 configuration hierarchy.
- C. Change the level 1 disable parameter under the R2 protocols isis interface ge-1/2/3.0 hierarchy to the level 2 disable parameter.
- D. Add IP addresses to the interface ge-1/2/3 unit 0 family iso hierarchy on both R1 and R2.

**Answer: B**

**Explanation:**

IS-IS routers can form Level 1 or Level 2 adjacencies depending on their configuration and network topology. Level 1 routers are intra-area routers that share the same area address with their neighbors. Level 2 routers are inter-area routers that can connect different areas. Level 1-2 routers are both intra-area and inter-area routers that can form adjacencies with any other router.



In the exhibit, R1 and R2 are in different areas (49.0001 and 49.0002), so they cannot form a Level 1 adjacency. However, they can form a Level 2 adjacency if they are both configured as Level 1-2 routers. R1 is already configured as a Level 1-2 router, but R2 is configured as a Level 1 router only, because of the level 1 disable command under the lo0.0 interface. This command disables Level 2 routing on the loopback interface, which is used as the router ID for IS-IS. Therefore, to establish a Level 1 adjacency between R1 and R2, the level 1 disable command under the R2 protocols isis interface lo0.0 hierarchy must be removed. This will enable Level 2 routing on R2 and allow it to form a Level 2 adjacency with R1.

#### NEW QUESTION 9

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.

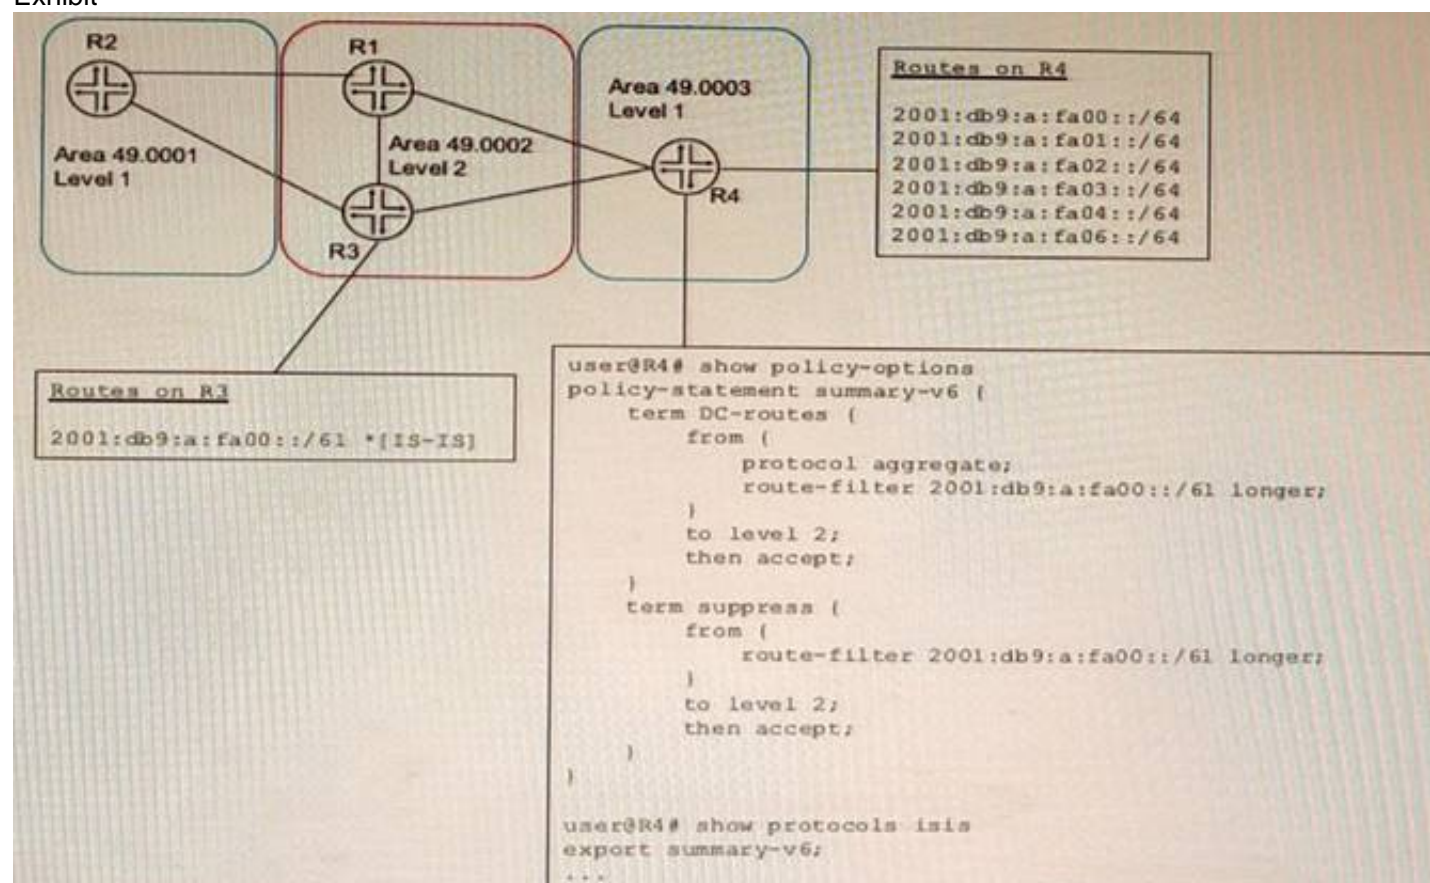
**Answer:** BC

#### 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

#### NEW QUESTION 10

Exhibit



A network designer would like to create a summary route as shown in the exhibit, but the configuration is not working.

Which three configuration changes will create a summary route? (Choose three.)

- A. set policy-options policy-statement leak-v6 term DC-routes then reject
- B. delete policy-options policy-statement leak-v6 term DC-routes from route-filter 2001: db9:a: fa00 : /61 longer
- C. set policy—options policy-statement leak-v€ term DC—routes from route-filter 2001:db9:a:fa00::/61 exact
- D. delete protocols isis export summary-v6
- E. set protocols isis import summary-v6

**Answer:** BCD

#### Explanation:

To create a summary route for IS-IS, you need to configure a policy statement that matches the prefixes to be summarized and sets the next-hop to discard. You also need to configure a summary-address statement under the IS-IS protocol hierarchy that references the policy statement. In this case, the policy statement leak-v6 is trying to match the prefix 2001:db9:a:fa00::/61 exactly, but this prefix is not advertised by any router in the network. Therefore, no summary route is created. To fix this, you need to delete the longer keyword from the route-filter term and change the prefix length to /61 exact. This will match any prefix that falls within the /61 range. You also need to delete the export statement under protocols isis, because this will export all routes that match the policy statement to other IS-IS routers, which is not desired for a summary route.

#### NEW QUESTION 10

Exhibit



```

user@router> show route extensive
...
2:192.168.101.5:65101::22031::02:00:31:06:00:01/304 MAC/IP (2 entries, 1
announced)
TSI:
Page 0 idx 0, (group IBGP-EVPN-Core type Internal) Type 1 val 0xb225964
(adv_entry)
  Advertised metrics:
    Nexthop: 192.168.101.5
    Localpref: 100
    AS path: [65101] I (Originator)
    Cluster list: 2.2.2.2
    Originator ID: 192.168.101.5
    Communities: target:65101:268457487 encapsulation:vxlan(0x8)
    Cluster ID: 3.3.3.3
  Advertise: 00000001
Path 2:192.168.101.5:65101::22031::02:00:31:06:00:01 from 192.168.101.3 Vector
len 4. Val: 0
  *BGP      Preference: 170/-101
            Route Distinguisher: 192.168.101.5:65101
            Next hop type: Indirect, Next hop index: 0
            Address: 0xb2d3490
            Next-hop reference count: 10520
            Source: 192.168.101.3
            Protocol next hop: 192.168.101.5
            Indirect next hop: 0x2 no-forward INH Session ID: 0x0
            State: <Active Int Ext>
            Local AS: 65101 Peer AS: 65101
            Age: 3d 19:56:57      Metric2: 0
            Validation State: unverified
            Task: BGP_65101.192.168.101.3
            Announcement bits (1): 1-BGP_RT_Background
            AS path: I (Originator)
            Cluster list: 2.2.2.2
            Originator ID: 192.168.101.5
            Communities: target:65101:268457487 encapsulation:vxlan(0x8)
            Import Accepted
            Route Label: 22031
            ESI: 05:00:00:fe:4d:00:00:56:0f:00
            Localpref: 100
            Router ID: 192.168.101.3
            Secondary Tables: default-switch.evpn.0
            Indirect next hops: 1
                Protocol next hop: 192.168.101.5
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                Indirect path forwarding next hops: 2
                    Next hop type: Router
                    Next hop: 10.0.2.12 via et-0/0/0.0
                    Session Id: 0x0
                    Next hop: 10.0.2.22 via et-0/0/1.0
                    Session Id: 0x0

192.168.101.5/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 2
Nexthop: 10.0.2.12 via et-0/0/0.0
Session Id: 0
Nexthop: 10.0.2.22 via et-0/0/1.0
Session Id: 0
...

```

Referring to the exhibit, which two statements are true? (Choose two.)

- A. This route is learned through EBGp
- B. This is an EVPN Type-2 route.
- C. The device advertising this route into EVPN is 192.168.101.5.
- D. The devices advertising this route into EVPN are 10.0.2.12 and 10.0.2.22.

**Answer:** BC

**Explanation:**

This is an EVPN Type-2 route, also called a MAC/IP advertisement route, that is used to advertise host IP and MAC address information to other VTEPs in an EVPN network. The route type field in the EVPN NLRI has a value of 2, indicating a Type-2 route. The device advertising this route into EVPN is 192.168.101.5, which is the IP address of the VTEP that learned the host information from the local CE device. This IP address is carried in the MPLS label field of the route as part of the VXLAN encapsulation.

**NEW QUESTION 12**

Exhibit



Communities: target:64512:5678 mac-mobility:0x0 (sequence 4)

You have MAC addresses moving in your EVPN environment

Referring to the exhibit, which two statements are correct about the sequence number? (Choose two)

- A. It identifies MAC addresses that should be discarded.
- B. It resolves conflicting MAC address ownership claims.
- C. It helps the local PE to identify the latest advertisement.
- D. It is advertised using a Type 2 message

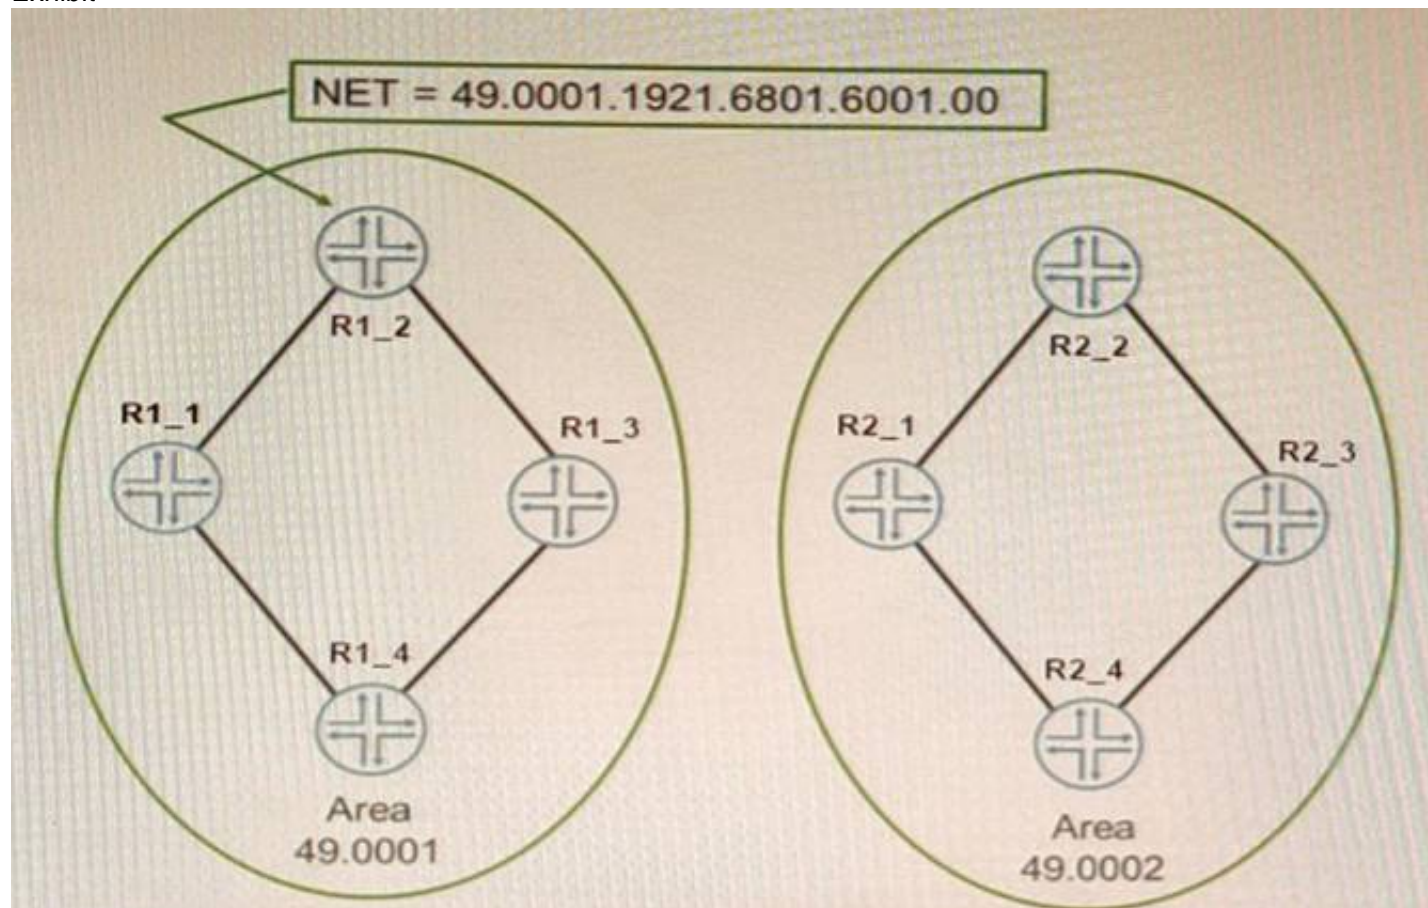
**Answer:** BC

**Explanation:**

The sequence number is a field in the MAC mobility extended community that is used to resolve conflicting MAC address ownership claims and to help the local PE to identify the latest advertisement. The sequence number is incremented by one for every MAC address mobility event, such as when a host moves from one Ethernet segment to another segment in the EVPN network. The PE device that receives multiple MAC advertisements for the same MAC address chooses the one with the highest sequence number as the most recent and valid advertisement.

**NEW QUESTION 15**

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 ID of R1\_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. 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. 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/xr-16/ios-xr-16-book/ios-ovr-vw-cf.html](https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/configuration/xr-16/ios-xr-16-book/ios-ovr-vw-cf.html) 2: <https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing-overview.html>

**NEW QUESTION 18**

Exhibit



```
[edit policy-options]
user@router# show
policy-statement block-igmp {
  term 1 {
    from {
      route-filter 224.7.7.7/32 exact;
      source-address-filter 192.168.100.10/32 exact;
    }
    then reject;
  }
}
[edit protocols igmp]
user@router# show
interface ge-0/0/0.0 {
  group-policy block-igmp;
  group-limit 25;
}
```

Based on the configuration contents shown in the exhibit, which statement is true?

- A. Joins for group 224.7.7.7 are rejected if the source address is 192.168.100.10
- B. Joins for any group are accepted if the group count value is less than 25.
- C. Joins for group 224.7.7.7 are always rejected, regardless of the group count.
- D. Joins for group 224.7.7.7 are accepted if the group count is less than 25

**Answer: D**

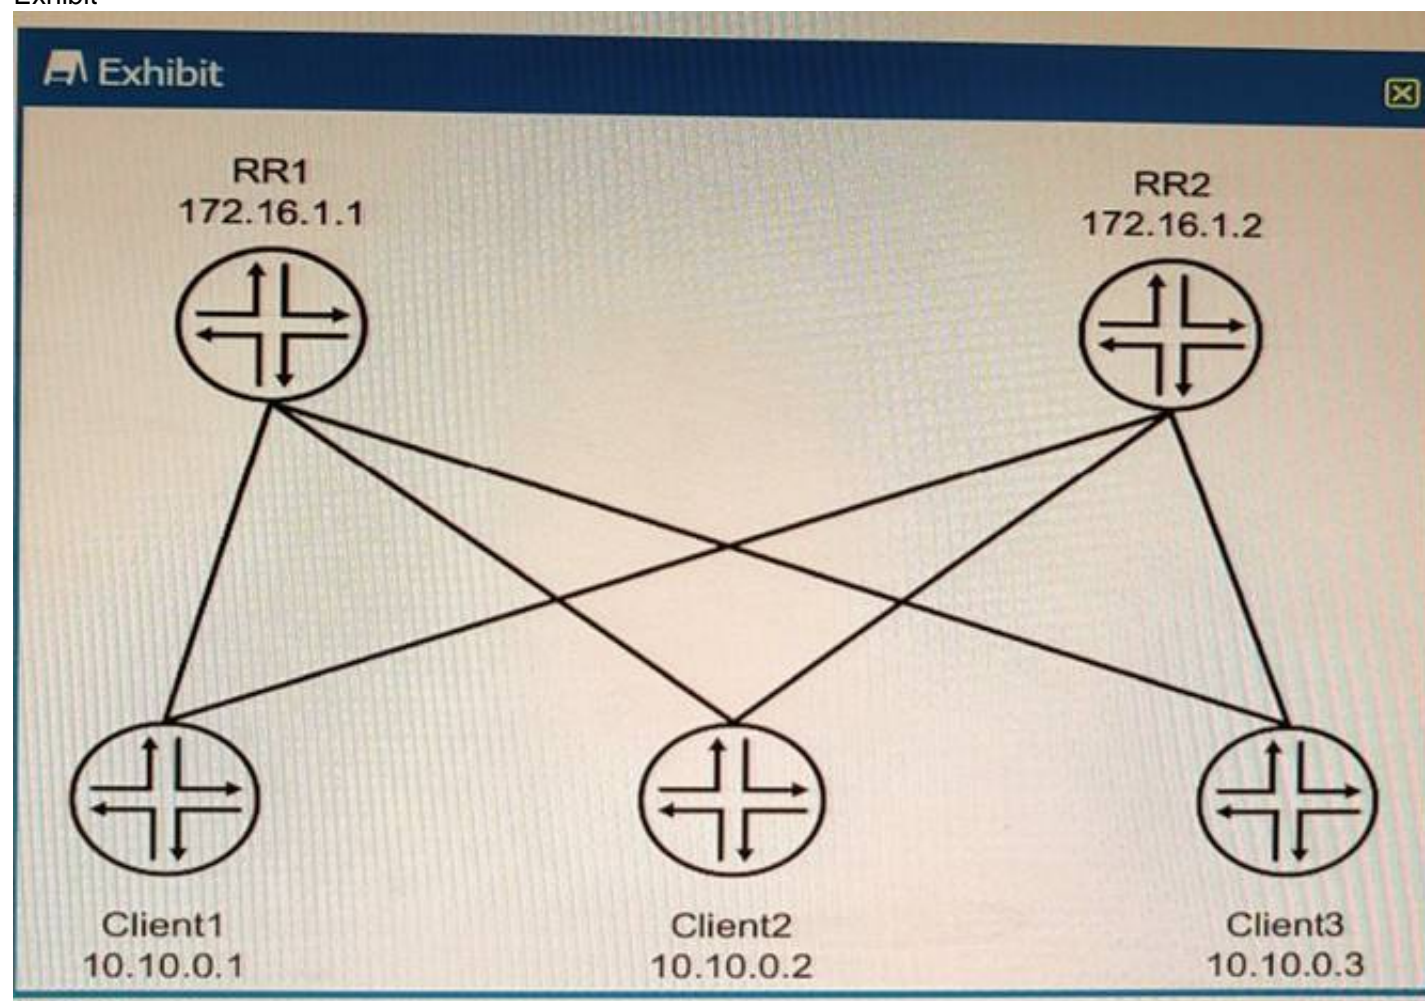
**Explanation:**

BGP policy framework is a set of tools that allows you to control the flow of routing information and apply routing policies based on various criteria. BGP policy framework consists of several components, such as route maps, prefix lists, community lists, AS path lists, and route filters. Route maps are used to define routing policies by matching certain conditions and applying certain actions. Prefix lists are used to filter routes based on their prefixes. Community lists are used to filter routes based on their community attributes. AS path lists are used to filter routes based on their AS path attributes. Route filters are used to filter routes based on their prefix length or range. In this question, we have a route map named ISP-A that has two clauses: clause 10 and clause 20. Clause 10 matches any route with a prefix length between 8 and 24 bits and sets the local preference to 200. Clause 20 matches any route with a prefix of 224.7.7.7/32 and rejects it. The route map is applied inbound on the BGP neighborship with ISP-A. Based on this configuration, the correct statement is that joins for group 224.7.7.7 are always rejected, regardless of the group count. This is because clause 20 explicitly denies any route with a prefix of 224.7.7.7/32, which corresponds to the multicast group 224.7.7.7.

Reference: 3: [https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute\\_bgp/configuration/xr-16/irg-xr-16-book/bgp-policy-framework.html](https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xr-16/irg-xr-16-book/bgp-policy-framework.html)

**NEW QUESTION 21**

Exhibit



The environment is using BGP All devices are in the same AS with reachability redundancy Referring to the exhibit, which statement is correct?

- A. RR1 is peered to Client2 and RR2
- B. RR2 is in an OpenConfirm State until RR1 becomes unreachable.
- C. Client1 is peered to Client2 and Client3.
- D. Peering is dynamically discovered between all devices.

**Answer:** A

**Explanation:**

BGP route reflectors are BGP routers that are allowed to ignore the IBGP loop avoidance rule and advertise IBGP learned routes to other IBGP peers under specific conditions. BGP route reflectors can reduce the number of IBGP sessions and updates in a network by eliminating the need for a full mesh of IBGP peers. BGP route reflectors can have three types of peerings:

? EBGp neighbor: A BGP router that belongs to a different autonomous system (AS) than the route reflector.

? IBGP client neighbor: An IBGP router that receives reflected routes from the route reflector. A client does not need to peer with other clients or non-clients.

? IBGP non-client neighbor: An IBGP router that does not receive reflected routes from the route reflector. A non-client needs to peer with other non-clients and the route reflector.

In the exhibit, we can see that RR1 and RR2 are route reflectors in the same AS with reachability redundancy. They have two types of peerings: EBGp neighbors (R1 and R4) and IBGP client neighbors (Client1, Client2, and Client3). RR1 and RR2 are also peered with each other as IBGP non-client neighbors.

**NEW QUESTION 25**

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:** AD

**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

**NEW QUESTION 29**

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
- B. o and vpn-
- C. ine
- D. o tables.
- E. RIB groups are not needed to leak routes between the ine
- F. 0 and VPN—
- G. ine
- H. 0 tables,
- I. Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> GW-1.
- J. Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> CE-1 -> PE-1 -> GW-1.

**Answer:** AD

**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.

**NEW QUESTION 30**

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 joins 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:

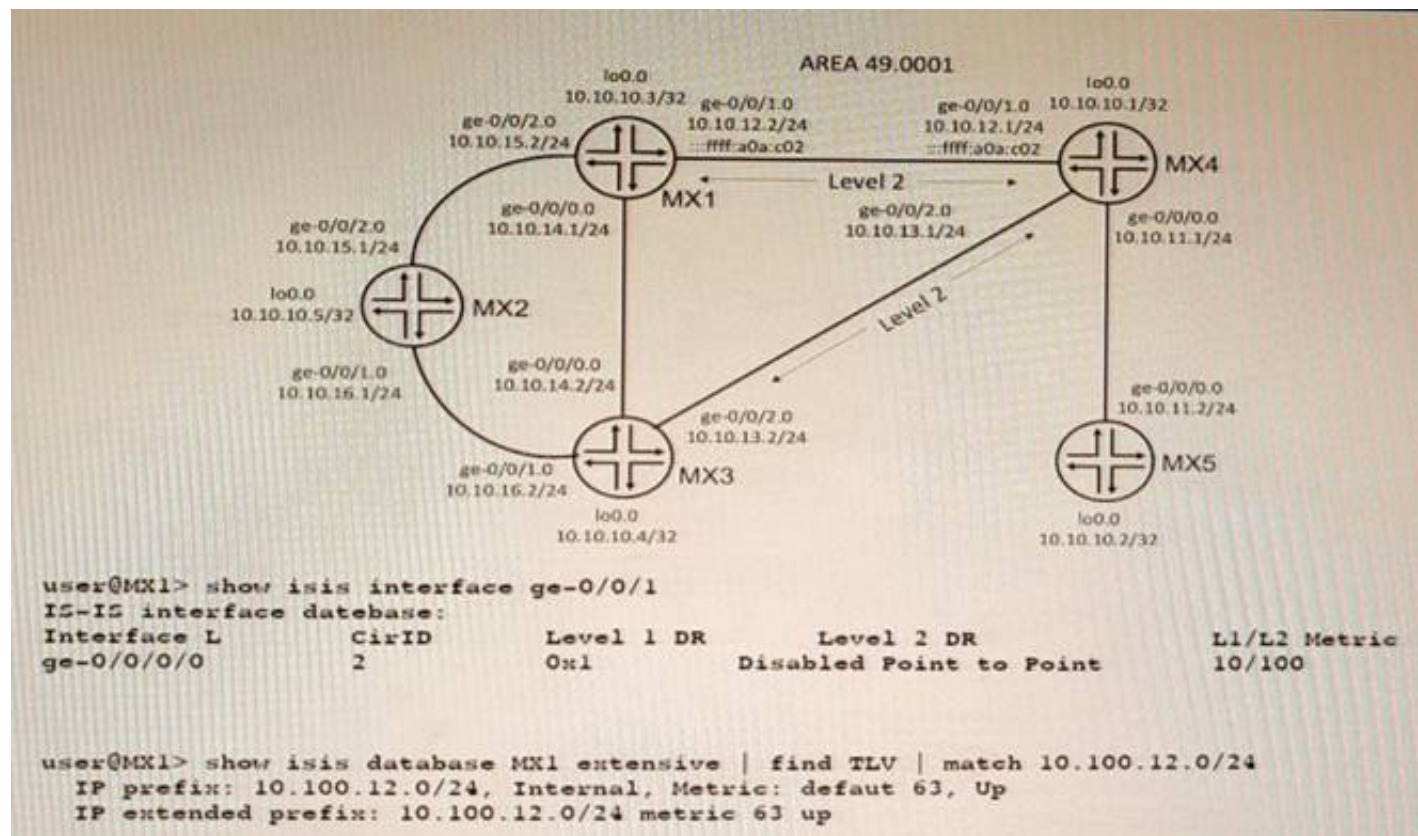
PIM Bootstrap Router (BSR) is a mechanism that allows PIM routers to discover and announce rendezvous point (RP) information for multicast groups. BSR uses two roles: candidate BSR and candidate RP. Candidate BSR is the router that collects information from all available RPs in the network and advertises 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 2515. 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/>

#### NEW QUESTION 34

Exhibit





A network is using IS-IS for routing.

In this scenario, why are there two TLVs shown in the exhibit?

- A. There are both narrow and wide metric devices in the topology
- B. The interface specified a metric of 100 for L2.
- C. Wide metrics have specifically been requested
- D. Both IPv4 and IPv6 are being used in the topology

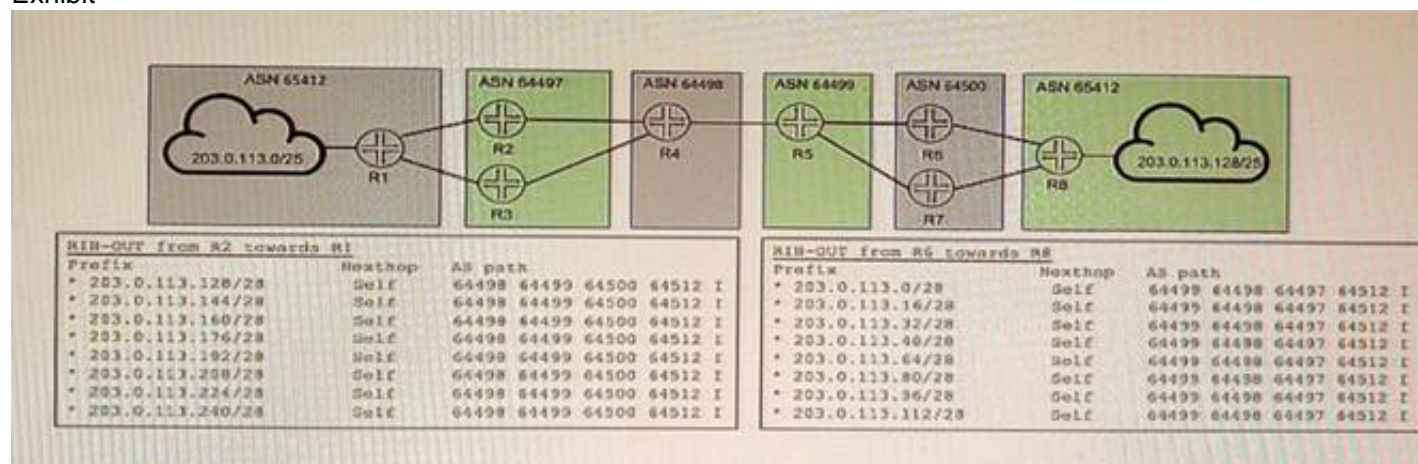
**Answer: A**

#### Explanation:

TLVs are tuples of (Type, Length, Value) that can be advertised in IS-IS packets. TLVs can carry different kinds of information in the Link State Packets (LSPs). IS-IS supports both narrow and wide metrics for link costs. Narrow metrics use a single octet to encode the link cost, while wide metrics use three octets. Narrow metrics have a maximum value of 63, while wide metrics have a maximum value of 16777215. If there are both narrow and wide metric devices in the topology, IS-IS will advertise two TLVs for each link: one with the narrow metric and one with the wide metric. This allows backward compatibility with older devices that only support narrow metrics.

#### NEW QUESTION 35

Exhibit



R1 and R8 are not receiving each other's routes

Referring to the exhibit, what are three configuration commands that would solve this problem? (Choose three.)

- A. Configure loops and advertise-peer-as on routers in AS 64497 and AS 64450.
- B. Configure loops on routers in AS 65412 and advertise-peer-as on routers in AS 64498.
- C. Configure as-override on advertisement from AS 64500 toward AS 64512.
- D. Configure remove-private on advertisements from AS 64497 toward AS 64498
- E. Configure remove-private on advertisements from AS 64500 toward AS 64499

**Answer: BDE**

#### Explanation:

The problem in this scenario is that R1 and R8 are not receiving each other's routes because of private AS numbers in the AS path. Private AS numbers are not globally unique and are not advertised to external BGP peers. To solve this problem, you need to do the following:

? Configure loops on routers in AS 65412 and advertise-peer-as on routers in AS 64498. This allows R5 and R6 to advertise their own AS number (65412) instead of their peer's AS number (64498) when sending updates to R7 and R8. This prevents a loop detection issue that would cause R7 and R8 to reject the routes from R5 and R62.

? Configure remove-private on advertisements from AS 64497 toward AS 64498 and from AS 64500 toward AS 64499. This removes any private AS numbers from the AS path before sending updates to external BGP peers. This allows R2 and R3 to receive the routes from R1 and R4, respectively3.

#### NEW QUESTION 39

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.

Option D is correct because 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. Option A is not correct because by default on PE-2, the site's local ID is automatically assigned a value of 0 and does not need to be configured to match the total number of attached sites. Option B is not correct because you do not need to create a unique Layer 2 VPN routing instance for each site on the PE-2 device. You can create one routing instance for all sites within a VFI. Option C is not correct because you do not need to use separate physical interfaces to connect PE-2 to each site. You can use subinterfaces or service instances on a single physical interface.

**NEW QUESTION 44**

When using OSPFv3 for an IPv4 environment, which statement is correct?

- A. OSPFv3 only supports IPv4.
- B. OSPFv3 supports both IPv6 and IPv4, but not in the same routing instance.
- C. OSPFv3 is not backward compatible with IPv4
- D. OSPFv3 supports IPv4 only on interfaces with family inet6 defined

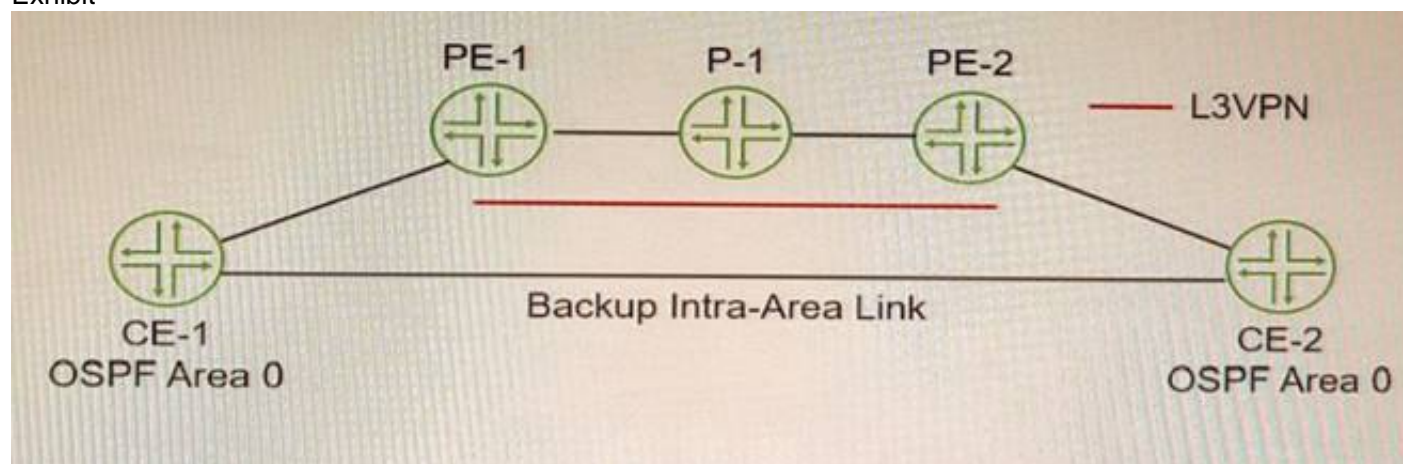
**Answer: C**

**Explanation:**

OSPFv3 is an extension of OSPFv2 that supports IPv6 routing and addressing. OSPFv3 is not backward compatible with IPv4 because it uses a different packet format and a different link-state advertisement (LSA) structure than OSPFv2. OSPFv3 also uses IPv6 link-local addresses as router IDs and neighbor addresses, instead of IPv4 addresses. To use OSPFv3 for an IPv4 environment, you need to enable the IPv4 unicast address family under [edit protocols ospf3] hierarchy level and configure IPv4 addresses on the interfaces.

**NEW QUESTION 45**

Exhibit



You must ensure that the VPN backbone is preferred over the back door intra-area link as long as the VPN is available. Referring to the exhibit, which action will accomplish this task?

- A. Configure an import routing policy on the CE routers that rejects OSPF routes learned on the backup intra-area link.
- B. Enable OSPF traffic-engineering.
- C. Configure the OSPF metric on the backup intra-area link that is higher than the L3VPN link.
- D. Create an OSPF sham link between the PE routers.

**Answer: D**

**Explanation:**

A sham link is a logical link between two PE routers that belong to the same OSPF area but are connected through an L3VPN. A sham link makes the PE routers appear as if they are directly connected, and prevents OSPF from preferring an intra-area back door link over the VPN backbone. To create a sham link, you need to configure the local and remote addresses of the PE routers under the [edit protocols ospf area area-id] hierarchy level1.

**NEW QUESTION 46**

Which two statements are correct about reflecting inet-vpn unicast prefixes in BGP route reflection? (Choose two.)

- A. Route reflectors do not change any existing BGP attributes by default when advertising routes.
- B. A BGP peer does not require any configuration changes to become a route reflector client.
- C. Clients add their originator ID when advertising routes to their route reflector
- D. Route reflectors add their cluster ID to the AS path when readvertising client routes.

**Answer: AB**

**Explanation:**

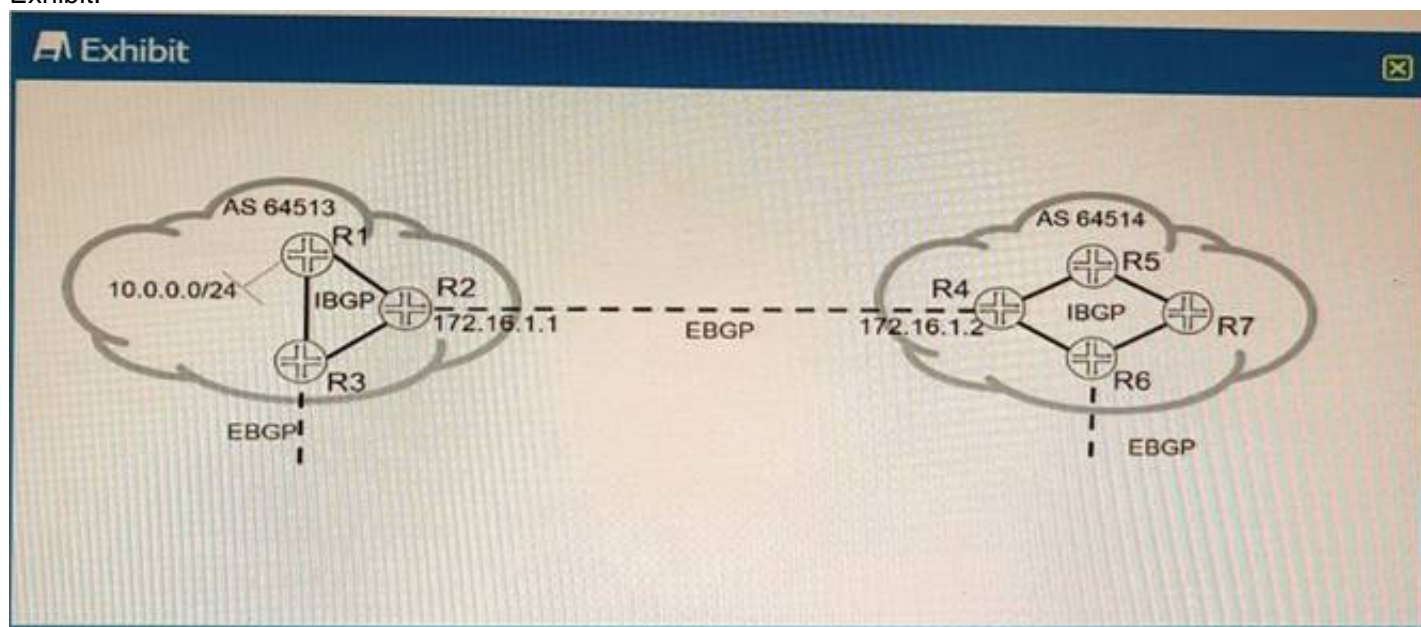
Route reflection is a BGP feature that allows a router to reflect routes learned from one IBGP peer to another IBGP peer, without requiring a full-mesh IBGP topology. Route reflectors do not change any existing BGP attributes by default when advertising routes, unless explicitly configured to do so. A BGP peer does



not require any configuration changes to become a route reflector client, only the route reflector needs to be configured with the client parameter under [edit protocols bgp group group-name neighbor neighbor- address] hierarchy level.

#### NEW QUESTION 51

Exhibit.



Referring to the exhibit; the 10.0.0.0/24 EBGP route is received on R5; however, the route is being hidden. What are two solutions that will solve this problem? (Choose two.)

- A. On R4, create a policy to change the BGP next hop to itself and apply it to IBGP as an export policy
- B. Add the external interface prefix to the IGP routing tables
- C. Add the internal interface prefix to the BGP routing tables.
- D. On R4, create a policy to change the BGP next hop to 172.16.1.1 and apply it to IBGP as an export policy

**Answer: AB**


#### Explanation:

the default behavior for iBGP is to propagate EBGP-learned prefixes without changing the next-hop. This can cause issues if the next-hop is not reachable via the IGP. One solution is to use the next-hop self command on R4, which will change the next-hop attribute to its own loopback address. This way, R5 can reach the next-hop via the IGP and install the route in its routing table.

Another solution is to add the external interface prefix (120.0.4.16/30) to the IGP routing tables of R4 and R5. This will also make the next-hop reachable via the IGP and allow R5 to use the route. According to2, this is a possible workaround for a pure IP network, but it may not work well for an MPLS network.

#### NEW QUESTION 54

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: AB**

#### 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:

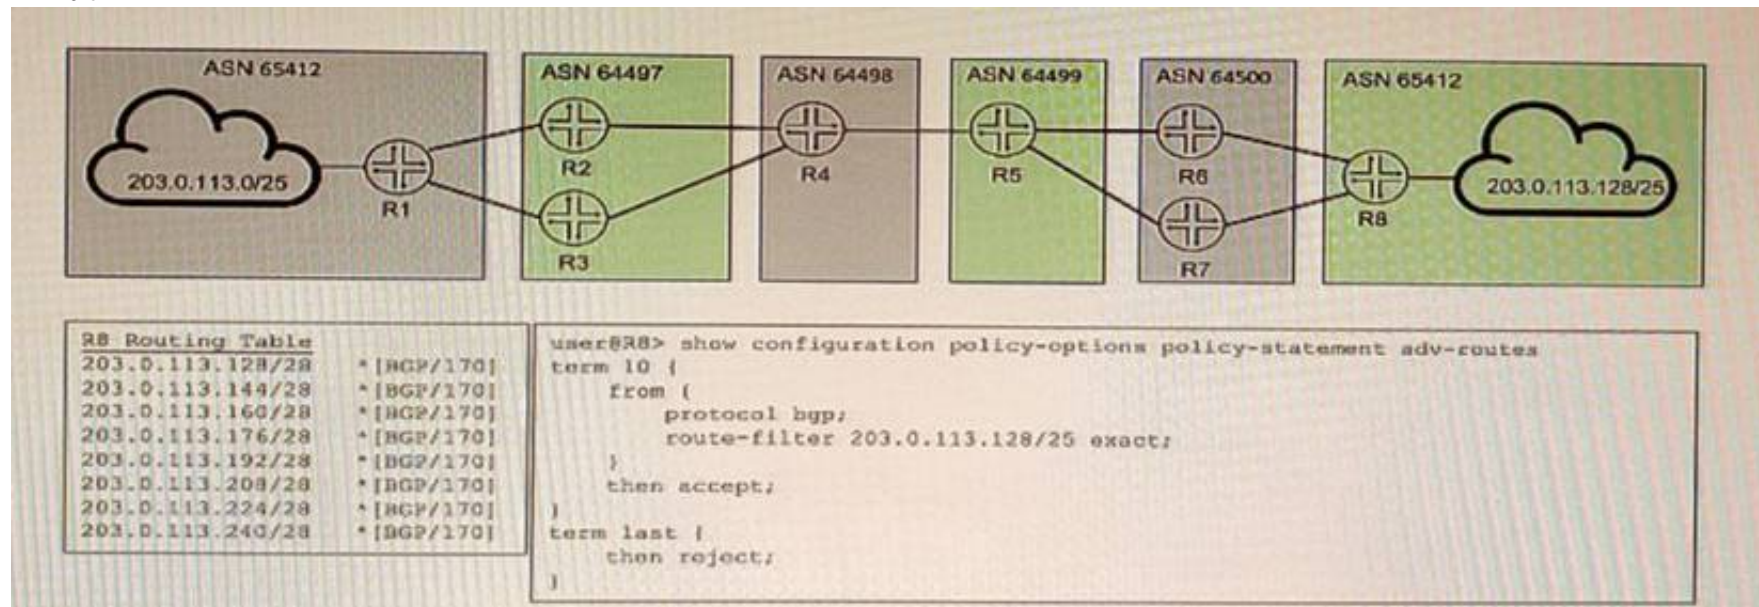
? By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3. This is because the default behavior of EBGP is to preserve the next-hop attribute of the routes received from another EBGP neighbor. The next-hop attribute indicates the IP address of the router that should 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>

## NEW QUESTION 55

Exhibit



You are attempting to summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500. You implement the export policy shown in the exhibit and all routes from the routing table stop being advertised.

In this scenario, which two steps would you take to summarize the route in BGP? (Choose two.)

- A. Remove the from protocol bgp command from the export policy.
- B. Add the set protocols bgp family inet unicast add-path command to allow additional routes to the RIB table
- C. -
- D. Add the set routing-options static route 203.0.113.123/25 discard command.
- E. Replace exact in the export policy with orlonger.

**Answer:** CD

### Explanation:

To summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500, you need to do the following:

? Add the set routing-options static route 203.0.113.128/25 discard command. This creates a static route for the summary prefix and discards any traffic destined to it. This is necessary because BGP can only advertise routes that are present in the routing table.

? Replace exact in the export policy with orlonger. This allows R8 to match and advertise any route that is equal or more specific than the summary prefix. The exact term only matches routes that are exactly equal to the summary prefix, which is not present in the routing table.

## NEW QUESTION 57

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:** CD

### 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:

## NEW QUESTION 58

You are a network architect for a service provider and want to offer Layer 2 services to your customers You want to use EVPN for Layer 2 services in your existing MPLS network.

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

- A. Segment routing must be configured on all PE routers.

- B. VXLAN must be configured on all PE routers.
- C. EVPN uses Type 2 routes to advertise MAC address and IP address pairs learned using ARP snooping
- D. EVPN uses Type 3 routes to join a multicast tree to flood traffic.

**Answer:** CD

**Explanation:**

EVPN is a technology that connects L2 network segments separated by an L3 network using a virtual Layer 2 network overlay over the Layer 3 network. EVPN uses BGP as its control protocol to exchange different types of routes for different purposes. Type 2 routes are used to advertise MAC address and IP address pairs learned using ARP snooping from the local CE devices. Type 3 routes are used to join a multicast tree to flood traffic such as broadcast, unknown unicast, and multicast (BUM) traffic.

**NEW QUESTION 63**

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

- A. If a broadcast interface is in both L1 and L2, one combined hello message is sent for both levels.
- B. If a point-to-point interface is in both L1 and L2, separate hello messages are sent for each level.
- C. If a point-to-point interface is in both L1 and L2, one combined hello message is sent for both levels.
- D. If a broadcast interface is in both L1 and L2, separate hello messages are sent for each level


**Answer:** BD

**Explanation:**

IS-IS supports two levels of routing: Level 1 (intra-area) and Level 2 (interarea). An IS-IS router can be either Level 1 only, Level 2 only, or both Level 1 and Level 2. A router that is both Level 1 and Level 2 is called a Level 1-2 router. A Level 1-2 router sends separate hello messages for each level on both point-to-point and broadcast interfaces. A point-to-point interface provides a connection between a single source and a single destination. A broadcast interface behaves as if the router is connected to a LAN.

**NEW QUESTION 65**

Exhibit



Referring to the exhibit, a working L3VPN exists that connects VPN-A sites CoS is configured correctly to match on the MPLS EXP bits of the LSP, but when traffic is sent from Site-1 to Site-2, PE-2 is not classifying the traffic correctly. What should you do to solve the problem?

- A. Configure the explicit-null statement on PE-1.
- B. Configure the explicit-null statement on PE-2
- C. Configure VPN prefix mapping for the PE-1\_to\_PE-2 LSP
- D. Set a static CoS value for the PE-1\_to\_PE-2 LSP


**Answer:** A

**Explanation:**

The explicit-null statement enables the PE router to send an MPLS label with a value of 0 (explicit null) instead of an IP header for packets destined to the VPN customer sites. This allows the penultimate hop router (the router before the egress PE router) to preserve the EXP bits of the MPLS label and pass them to the egress PE router. The egress PE router can then use these EXP bits to classify the traffic according to the CoS policy. In this example, PE-1 should configure the explicit-null statement under [edit protocols mpls label-switched-path PE-1\_to\_PE-2] hierarchy level.

**NEW QUESTION 69**

Exhibit



Referring to the exhibit, which three statements are correct about route 10 0 0.0/16 when using the default BGP advertisement rules'? (Choose three.)

- A. R1 will prepend AS 65531 when advertising 10 0 0.0/16 to R2.
- B. R1 will advertise 10.0.0.0/16 to R2 with 192 168 1 1 as the next hop.
- C. R2 will advertise 10.0.0.0/16 to R3 with 192.168.1 1 as the next hop
- D. R4 will advertise 10 0 0.0/16 to R6 with 172.16 1 1 as the next hop



E. R2 will advertise 10.0.0.0/16 to R4 with 172.16.1.1 as the next hop

**Answer:** BDE

**Explanation:**

The problem in this scenario is that R1 and R8 are not receiving each other's routes because of private AS numbers in the AS path. Private AS numbers are not globally unique and are not advertised to external BGP peers. To solve this problem, you need to do the following:

? Configure loops on routers in AS 65412 and advertise-peer-as on routers in AS 64498. This allows R5 and R6 to advertise their own AS number (65412) instead of their peer's AS number (64498) when sending updates to R7 and R8. This prevents a loop detection issue that would cause R7 and R8 to reject the routes from R5 and R62.

? Configure remove-private on advertisements from AS 64497 toward AS 64498 and from AS 64500 toward AS 64499. This removes any private AS numbers from the AS path before sending updates to external BGP peers. This allows R2 and R3 to receive the routes from R1 and R4, respectively3.

**NEW QUESTION 73**

.....

## Relate Links

**100% Pass Your JN0-664 Exam with ExamBible Prep Materials**

<https://www.exambible.com/JN0-664-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>