

# Juniper

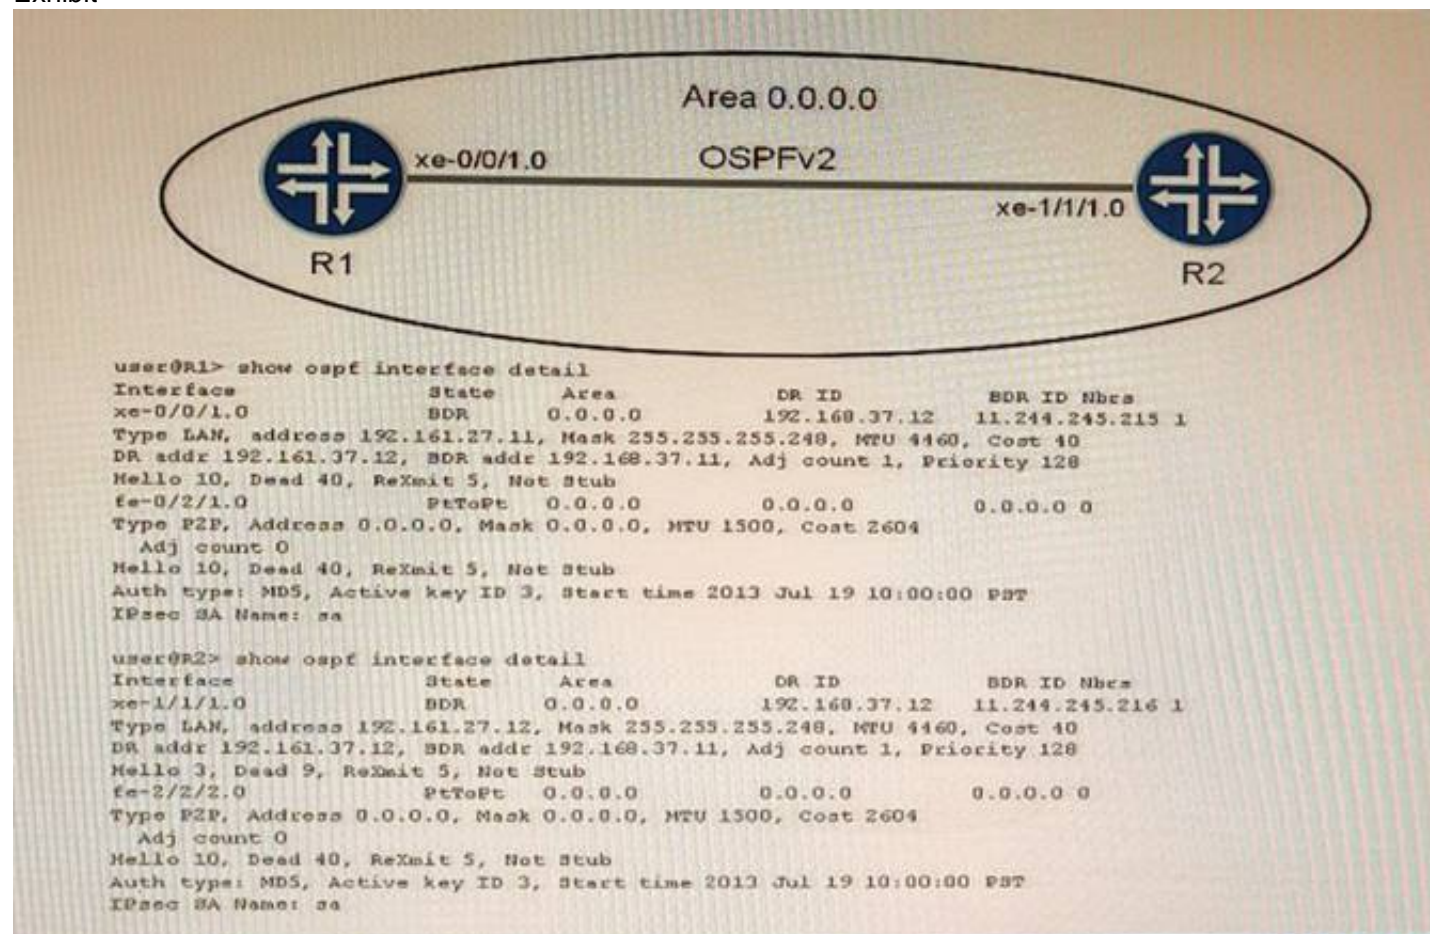
## Exam Questions JN0-664

Service Provider - Professional (JNCIP-SP)



## NEW QUESTION 1

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 2

Which three mechanisms are used by Junos platforms to evaluate incoming traffic for CoS purposes? (Choose three )

- A. rewrite rules
- B. behavior aggregate classifiers
- C. traffic shapers
- D. fixed classifiers
- E. multifield classifiers

**Answer:** BDE

### Explanation:

Junos platforms use different mechanisms to evaluate incoming traffic for CoS purposes, such as:

? Behavior aggregate classifiers: These classifiers use a single field in a packet header to classify traffic into different forwarding classes and loss priorities based on predefined or user-defined values.

? Fixed classifiers: These classifiers use a fixed field in a packet header to classify traffic into different forwarding classes and loss priorities based on predefined values.

? Multifield classifiers: These classifiers use multiple fields in a packet header to classify traffic into different forwarding classes and loss priorities based on user-defined values and filters.

Rewrite rules and traffic shapers are not used to evaluate incoming traffic for CoS purposes, but rather to modify or shape outgoing traffic based on CoS policies.

## NEW QUESTION 3

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 4

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

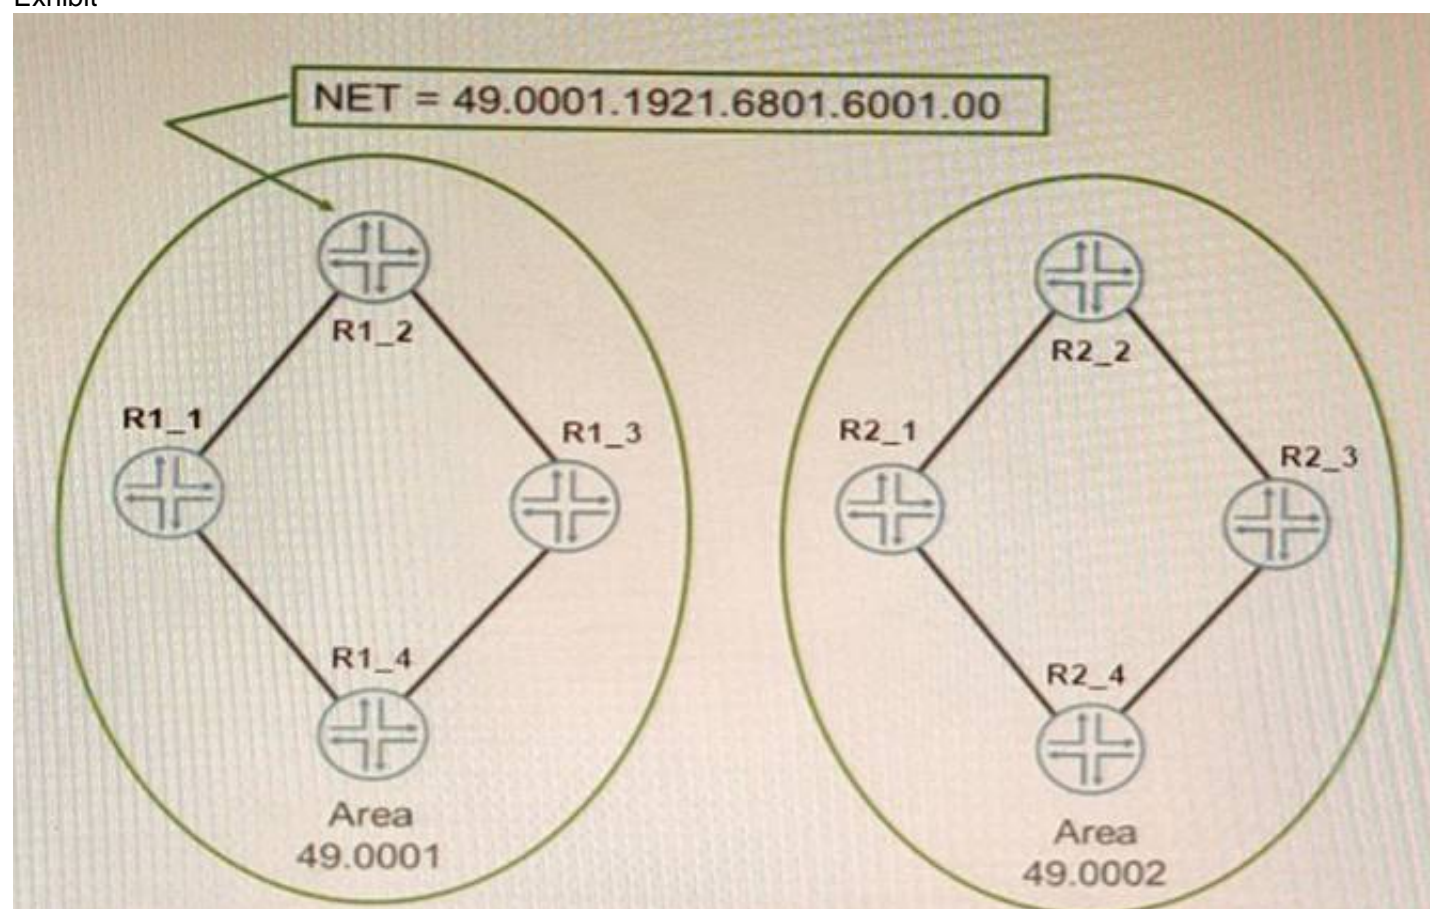
#### Explanation:

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 5

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<sup>1</sup>. 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<sup>2</sup>. 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/irs-xr-16-book/irs-ovrvw-cf.html](https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/configuration/xr-16/irs-xr-16-book/irs-ovrvw-cf.html) 2:

<https://www.juniper.net/documentation/us/en/software/junos/is-isis/topics/concept/is-isis-routing-overview.html>

## NEW QUESTION 6

Which statement is true regarding BGP FlowSpec?

- A. It uses a remote triggered black hole to protect a network from a denial-of-service attack.
- B. It uses dynamically created routing policies to protect a network from denial-of-service attacks
- C. It is used to protect a network from denial-of-service attacks dynamically
- D. It verifies that the source IP of the incoming packet has a resolvable route in the routing table

**Answer: B**

### Explanation:

BGP FlowSpec is a feature that extends the Border Gateway Protocol (BGP) to enable routers to exchange traffic flow specifications, allowing for more precise control of network traffic. The BGP FlowSpec feature enables routers to advertise and receive information about specific flows in the network, such as those originating from a particular source or destined for a particular destination. Routers can then use this information to construct traffic filters that allow or deny packets of a certain type, rate limit flows, or perform other actions<sup>1</sup>. BGP FlowSpec can also help in filtering traffic and taking action against distributed denial of service (DDoS) attacks by dropping the DDoS traffic or diverting it to an analyzer<sup>2</sup>. BGP FlowSpec rules are internally converted to equivalent Cisco Common Classification Policy Language (C3PL) representing corresponding match and action parameters<sup>2</sup>. Therefore, BGP FlowSpec uses dynamically created routing policies to protect a network from denial-of-service attacks.

References: 1: <https://www.networkingsignal.com/what-is-bgp-flowspec/> 2: [https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute\\_bgp/configuration/xr-16/irg-xr-16-book/bgp-flowspec-route-reflector-support.html](https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xr-16/irg-xr-16-book/bgp-flowspec-route-reflector-support.html)

## NEW QUESTION 7

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**

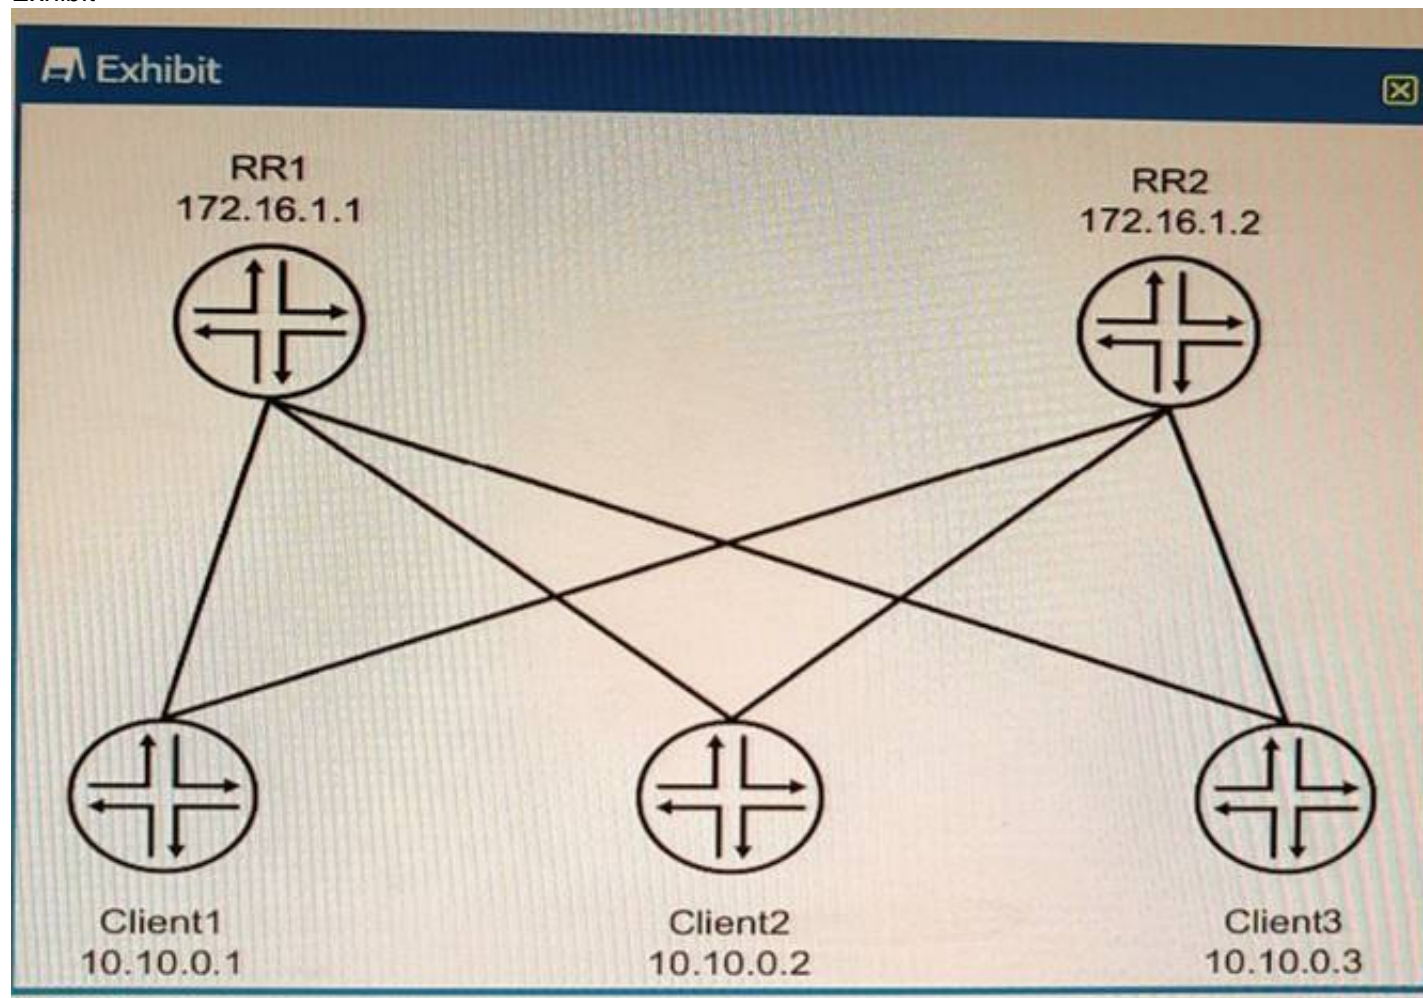
**Explanation:**

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 8**

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 9**

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

- A. PSNPs are flooded periodically.
- B. PSNPs contain only descriptions of LSPs.
- C. CSNPs are flooded periodically
- D. CSNPs contain only descriptions of LSPs.

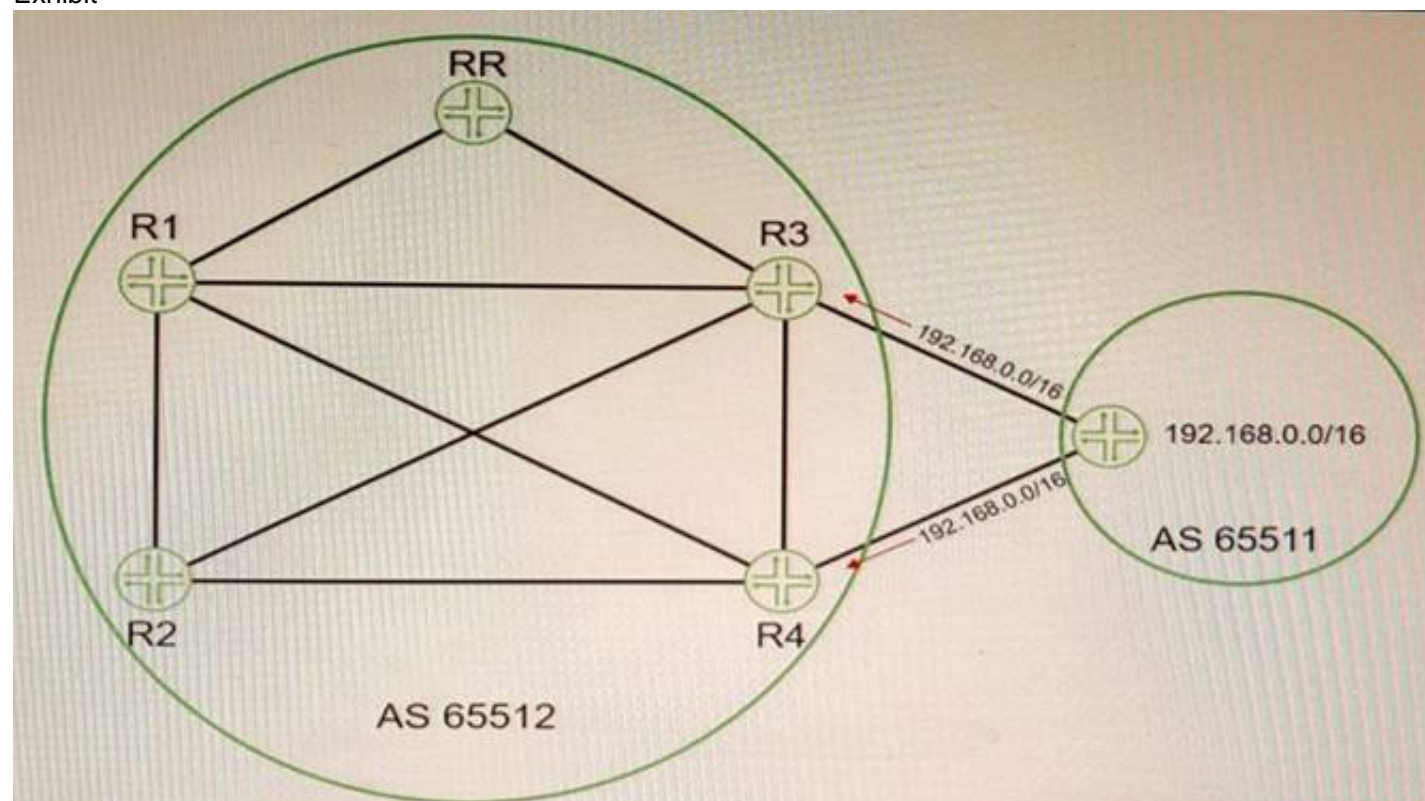
**Answer:** BC

**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 packets to synchronize link-state databases among routers: Link State Packets (LSPs) and Partial Sequence Number Packets (PSNPs). 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<sup>3</sup>. IS-IS also uses another type of packet called Complete Sequence Number Packets (CSNPs), which are used to summarize the entire link-state database at regular intervals or when a new adjacency is formed. CSNPs are flooded periodically throughout the network and contain only descriptions of LSPs<sup>4</sup>. Therefore, PSNPs contain only descriptions of LSPs and CSNPs are flooded periodically. References: 3: <https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/routing-policy-is-is-partial-sequence-number-packet-psnp.html> 4: <https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/routing-policy-is-is-complete-sequence-number-packet-csnp.html>

**NEW QUESTION 10**

Exhibit



Referring to the exhibit, you are receiving the 192.168.0.0/16 route on both R3 and R4 from your EBGp neighbor. You must ensure that R1 and R2 receive both BGP routes from the route reflector.

In this scenario, which BGP feature should you configure to accomplish this behavior?

- A. add-path
- B. multihop
- C. multipath
- D. route-target

**Answer:** A

**Explanation:**

BGP add-path is a feature that allows the advertisement of multiple paths through the same peering session for the same prefix without the new paths implicitly replacing any previous paths. This behavior promotes path diversity and reduces multi-exit discriminator (MED) oscillations. BGP add-path is implemented by adding a path identifier to each path in the NLRI. The path identifier can be considered as something similar to a route distinguisher in VPNs, except that a path ID can apply to any address family. Path IDs are unique to a peering session and are generated for each network<sup>3</sup>. In this question, we have a route reflector (RR) that receives two routes for the same prefix (192.168.0.0/16) from an EBGp neighbor. By default, the RR will only advertise its best path to its clients (R1 and R2). However, we want R1 and R2 to receive both routes from the RR. To achieve this, we need to configure BGP add-path on the RR and enable it to send multiple paths for the same prefix to its clients.

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

#### NEW QUESTION 10

Which two EVPN route types are used to advertise a multihomed Ethernet segment? (Choose two )

- A. Type 1
- B. Type 3
- C. Type 4
- D. Type 2

**Answer:** AC

#### Explanation:

EVPN is a solution that provides Ethernet multipoint services over MPLS networks. EVPN uses BGP to distribute endpoint provisioning information and set up pseudowires between PE devices. EVPN uses different route types to convey different information in the control plane. The following are the main EVPN route types:

? Type 1 - Ethernet Auto-Discovery Route: This route type is used for network-wide messaging and discovery of other PE devices that are part of the same EVPN instance. It also carries information about the redundancy mode and load balancing algorithm of the PE devices.

? Type 2 - MAC/IP Advertisement Route: This route type is used for MAC and IP address learning and advertisement between PE devices. It also carries information about the Ethernet segment identifier (ESI) and the label for forwarding traffic to the MAC or IP address.

? Type 3 - Inclusive Multicast Ethernet Tag Route: This route type is used for broadcast, unknown unicast, and multicast (BUM) traffic forwarding. It also carries information about the multicast group and the label for forwarding BUM traffic.

? Type 4 - Ethernet Segment Route: This route type is used for multihoming scenarios, where a CE device is connected to more than one PE device. It also carries information about the ESI and the designated forwarder (DF) election process.

#### NEW QUESTION 13

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.

#### NEW QUESTION 17

Exhibit



```
user@PE1# show routing-instances
VPN-A {
    instance-type vrf;
    interface ge-0/0/1.0;
    vrf-target target:64512:1234;
    protocols {
        bgp {
            group CE {
                type external;
                family inet {
                    unicast;
                }
                neighbor 10.0.0.1 {
                    peer-as 64512;
                    as-override;
                }
            }
        }
    }
}
```

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

- A. This VPN connects customer sites that use different AS numbers.
- B. This VPN connects customer sites that use the same AS number
- C. A Layer 2 VPN is configured.
- D. A Layer 3 VPN is configured.

**Answer:** AD

**Explanation:**

The configuration shown in the exhibit is for a Layer 3 VPN that connects customer sites that use different AS numbers. A Layer 3 VPN is a type of VPN that uses MPLS labels to forward packets across a provider network and BGP to exchange routing information between PE routers and CE routers. A Layer 3 VPN allows customers to use different routing protocols and AS numbers at their sites, as long as they can peer with BGP at the PE-CE interface. In this example, CE-1 is using AS 65530 and CE-2 is using AS 65531, but they can still communicate through the VPN because they have BGP sessions with PE-1 and PE-2, respectively.

**NEW QUESTION 18**

Your organization manages a Layer 3 VPN for multiple customers To support advanced route than one BGP community on advertised VPN routes to remote PE routers.

Which routing-instance configuration parameter would support this requirement?

- A. vrf-export
- B. vrf-import
- C. vrf-target export
- D. vrf-target import

**Answer:** C

**Explanation:**

The vrf-target export parameter is used to specify one or more BGP extended community attributes that are attached to VPN routes when they are exported from a VRF routing instance to remote PE routers. This parameter allows you to control which VPN routes are accepted by remote PE routers based on their import policies. You can specify more than one vrf-target export value for a VRF routing instance to support advanced route filtering or route leaking scenarios.

**NEW QUESTION 21**

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.)

- A. set protocols ospf overload timeout 900

- B. set protocols ospf3 realm ipv4-unicast overload timeout 900
- C. set protocols ospf overload
- D. set protocols oapf3 overload timeout 900
- E. set protocols ospf3 overload

**Answer:** AE

**Explanation:**

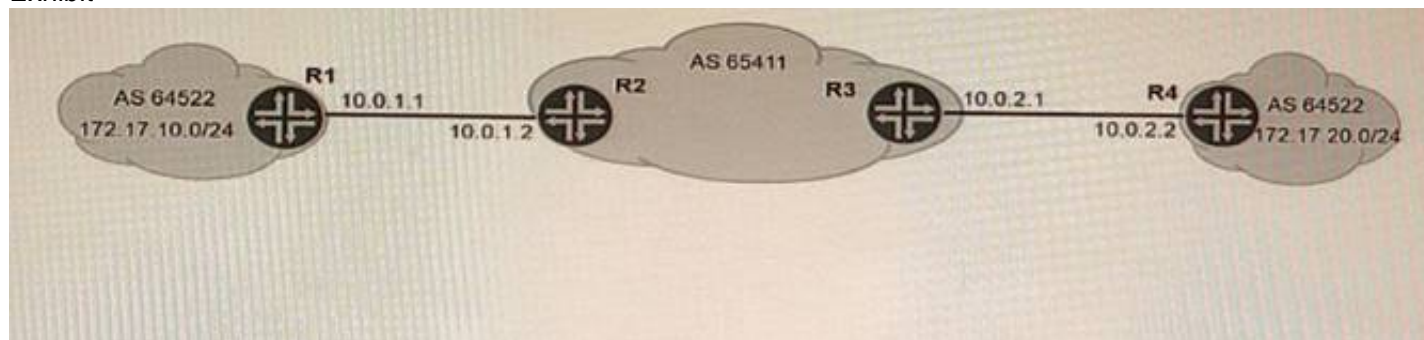
To reduce the impact of suboptimal routing in OSPF and OSPFv3 after devices reboot, you can use the overload feature to prevent a router from being used as a transit router for a specified period of time. This allows the router to stabilize its routing table before forwarding traffic for other routers.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

**NEW QUESTION 24**

Exhibit



You are asked to exchange routes between R1 and R4 as shown in the exhibit. These two routers use the same AS number Which two steps will accomplish this task? (Choose two.)

- A. Configure the BGP group with the advertise-peer-as parameter on R1 and R4.
- B. Configure the BGP group with the as-override parameter on R2 and R3
- C. Configure the BGP group with the advertise-peer-as parameter on R2 and R3.
- D. Configure the BGP group with the as-override parameter on R1 and R4

**Answer:** AB

**Explanation:**

The advertise-peer-as parameter allows a router to advertise its peer's AS number as part of the AS path attribute when sending BGP updates to other peers. This parameter is useful when two routers in the same AS need to exchange routes through another AS, such as in the case of R1 and R4. By configuring this parameter on R1 and R4, they can advertise each other's AS number to R2 and R3, respectively.

The as-override parameter allows a router to replace the AS number of its peer with its own AS number when receiving BGP updates from that peer. This parameter is useful when two routers in different ASes need to exchange routes through another AS that has the same AS number as one of them, such as in the case of R2 and R3. By configuring this parameter on R2 and R3, they can override the AS number of R1 and R4 with their own AS number when sending BGP updates to each other.

**NEW QUESTION 29**

You are responding to an RFP for a new MPLS VPN implementation. The solution must use LDP for signaling and support Layer 2 connectivity without using BGP The solution must be scalable and support multiple VPN connections over a single MPLS LSP The customer wants to maintain all routing for their Private network In this scenario, which solution do you propose?

- A. circuit cross-connect
- B. BGP Layer 2 VPN
- C. LDP Layer 2 circuit
- D. translational cross-connect

**Answer:** C

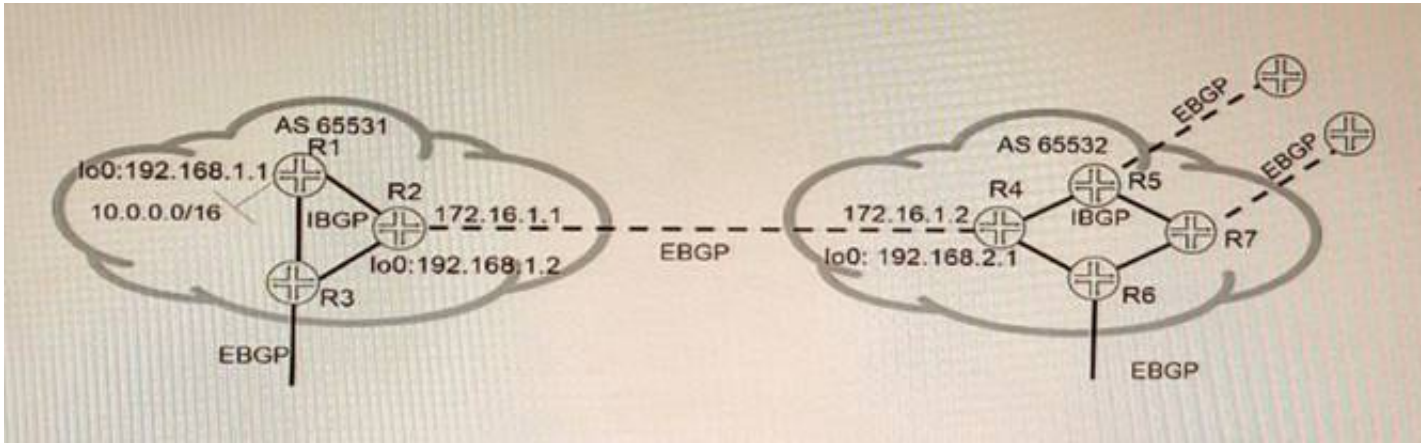
**Explanation:**

AToM (Any Transport over MPLS) is a framework that supports various Layer 2 transport types over an MPLS network core. One of the transport types supported by AToM is LDP Layer 2 circuit, which is a point-to-point Layer 2 connection that uses LDP for signaling and MPLS for forwarding. LDP Layer 2 circuit can support Layer 2 connectivity without using BGP and can be scalable and efficient by using a single MPLS LSP for multiple VPN connections. The customer can maintain all routing for their private network by using their own CE switches.

**NEW QUESTION 30**

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:**

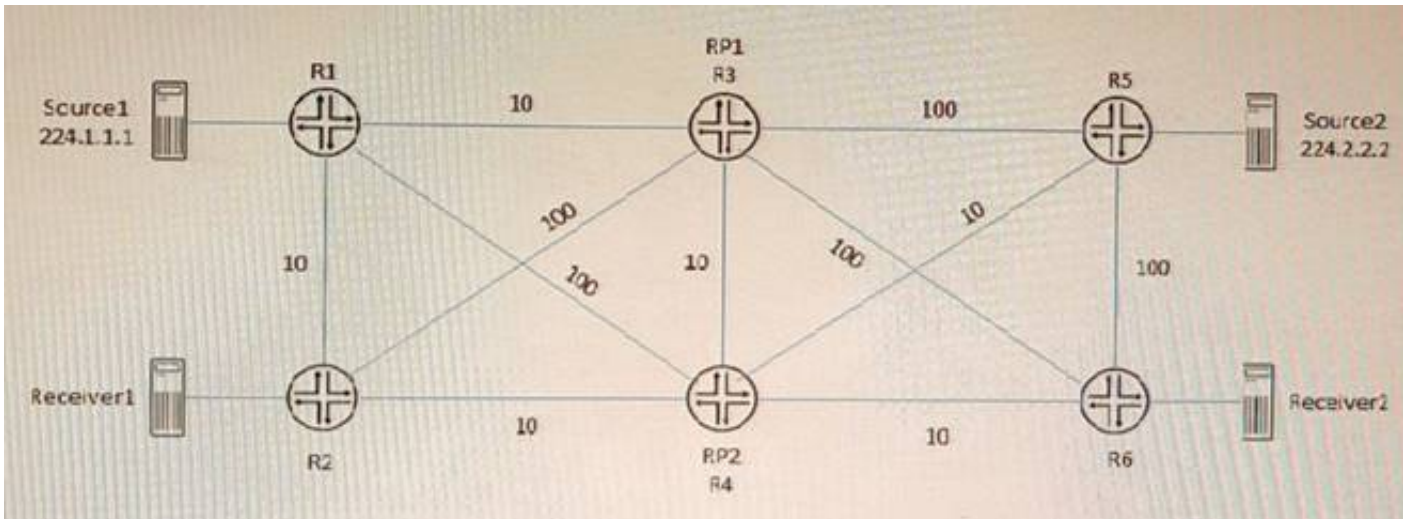
The problem in this scenario is that R1 and R8 are not receiving each other's routes because of private AS numbers in the AS path. Private AS numbers are not globally unique and are not advertised to external BGP peers. To solve this problem, you need to do the following:

? 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 34**

Exhibit



Referring to the exhibit, PIM-SM is configured on all routers, and Anycast-RP with Anycast- PIM is used for the discovery mechanism on RP1 and RP2. The interface metric values are shown for the OSPF area.

In this scenario, which two statements are correct about which RP is used? (Choose two.)

- A. Source2 will use RP2 and Receiver1 will use RP2 for group 224.2.2.2.
- B. Source2 will use RP1 and Receiver2 will use RP1 for group 224.2.2.2.
- C. Source1 will use RP1 and Receiver1 will use RP1 for group 224.1.1.1.
- D. Source1 will use RP1 and Receiver1 will use RP2 for group 224.1 1 1

**Answer: AC**

**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. A sham link creates an OSPF multihop neighborhood between the PE routers using TCP port 646. The PEs exchange Type 1 OSPF LSAs instead of Type 3 OSPF LSAs for the L3VPN routes, which allows OSPF to use the correct metric for route selection1.

**NEW QUESTION 36**

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### JN0-664 Practice Exam Features:

- \* JN0-664 Questions and Answers Updated Frequently
- \* JN0-664 Practice Questions Verified by Expert Senior Certified Staff
- \* JN0-664 Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* JN0-664 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The JN0-664 Practice Test Here](#)**