

## Exam Questions JN0-664

Service Provider - Professional (JNCIP-SP)

<https://www.2passeasy.com/dumps/JN0-664/>



### 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

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 3

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 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

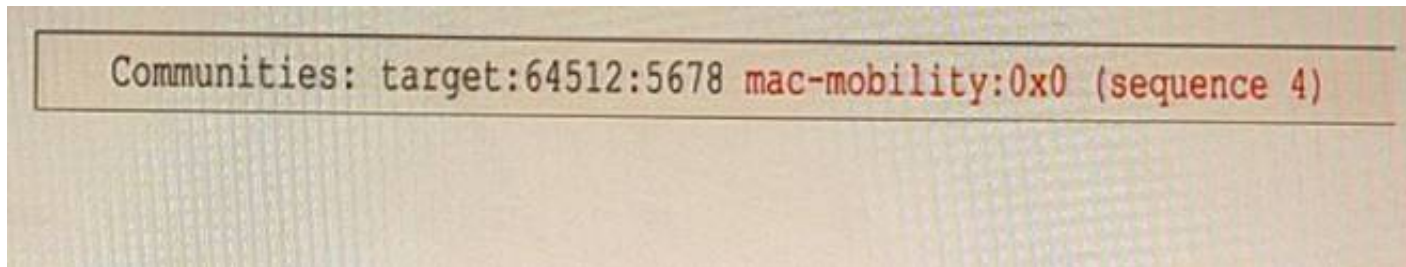
**Explanation:**

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





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 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-xe-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-xe-16-book/bgp-flowspec-route-reflector-support.html)

**NEW QUESTION 7**

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 8**

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 9

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.0 and vpn-a.inet.0 tables.
- B. o and vpn-
- C. ine
- D. o tables.
- E. RIB groups are not needed to leak routes between the inet.0 and vpn-
- 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 10

Which origin code is preferred by BGP?

- A. Internal
- B. External
- C. Incomplete
- D. Null

**Answer:** C

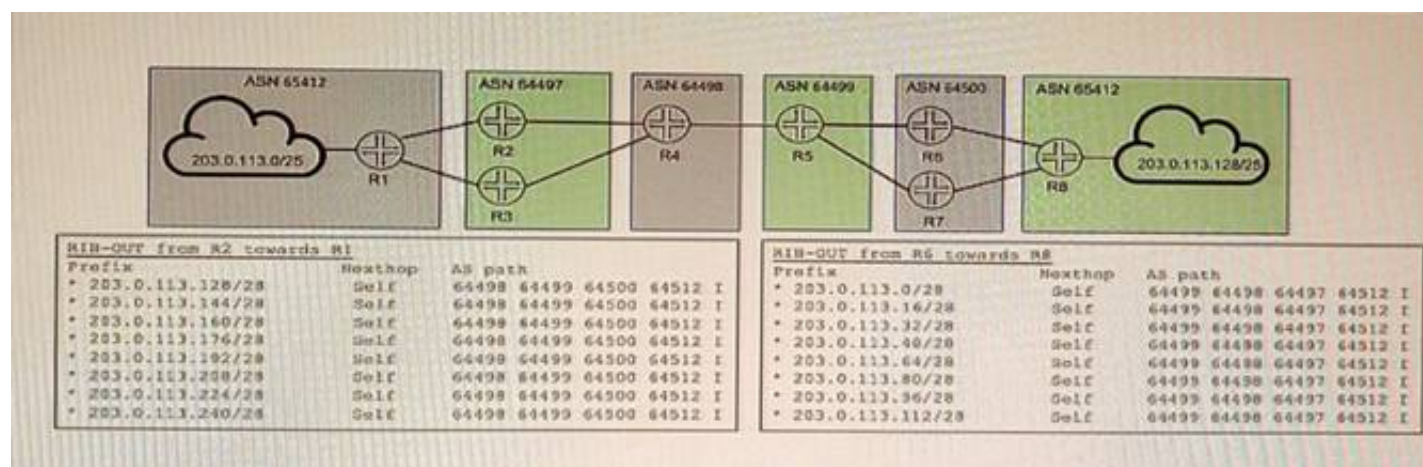
#### Explanation:

BGP uses several attributes to select the best path for a destination prefix. One of these attributes is origin, which indicates how BGP learned about a route. The origin attribute can have one of three values: IGP, EGP, or Incomplete. IGP means that the route was originated by a network or aggregate statement within BGP or by redistribution from an IGP into BGP. EGP means that the route was learned from an external BGP peer (this value is obsolete since BGP version 4). Incomplete means that the route was learned by some other means, such as redistribution from a static route into BGP. BGP prefers routes with lower origin values, so Incomplete is preferred over EGP, which is preferred over IGP.

### NEW QUESTION 10

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:

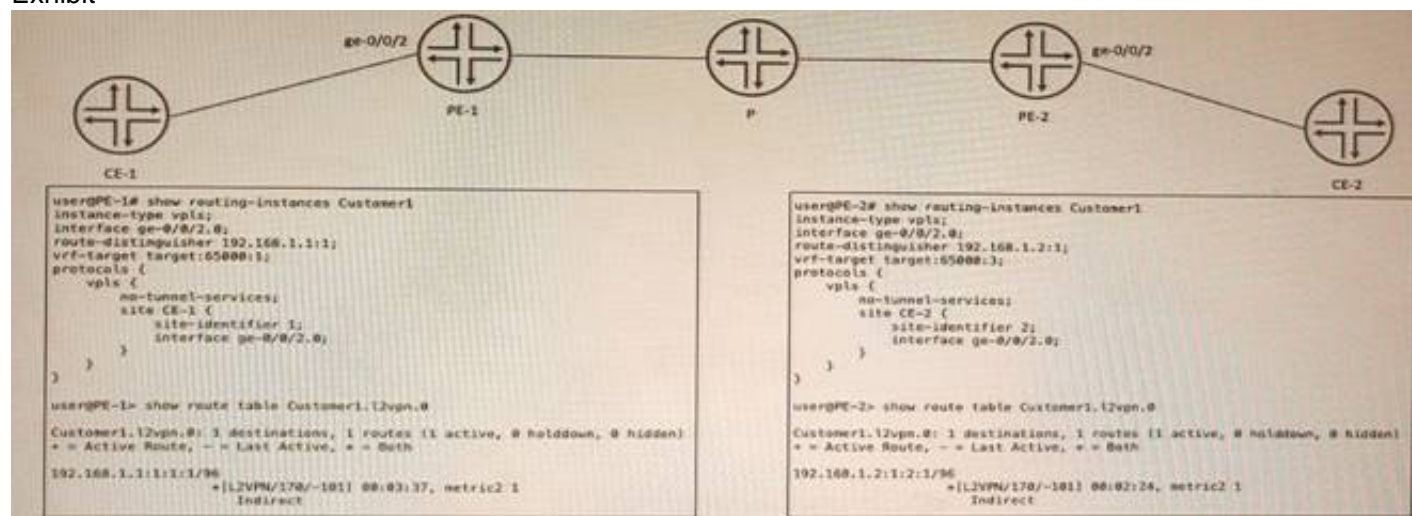
The problem in this scenario is that R1 and R8 are not receiving each other's routes because of private AS numbers in the AS path. Private AS numbers are not globally unique and are not advertised to external BGP peers. To solve this problem, you need to do the following:

? 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 12

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 l2vpn.
- 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.

#### NEW QUESTION 15

By default, which statement is correct about OSPF summary LSAs?

- A. All Type 2 and Type 7 LSAs will be summarized into a single Type 5 LSA
- B. The area-range command must be installed on all routers.
- C. Type 3 LSAs are advertised for routes in Type 1 LSAs.
- D. The metric associated with a summary route will be equal to the lowest metric associated with an individual contributing route

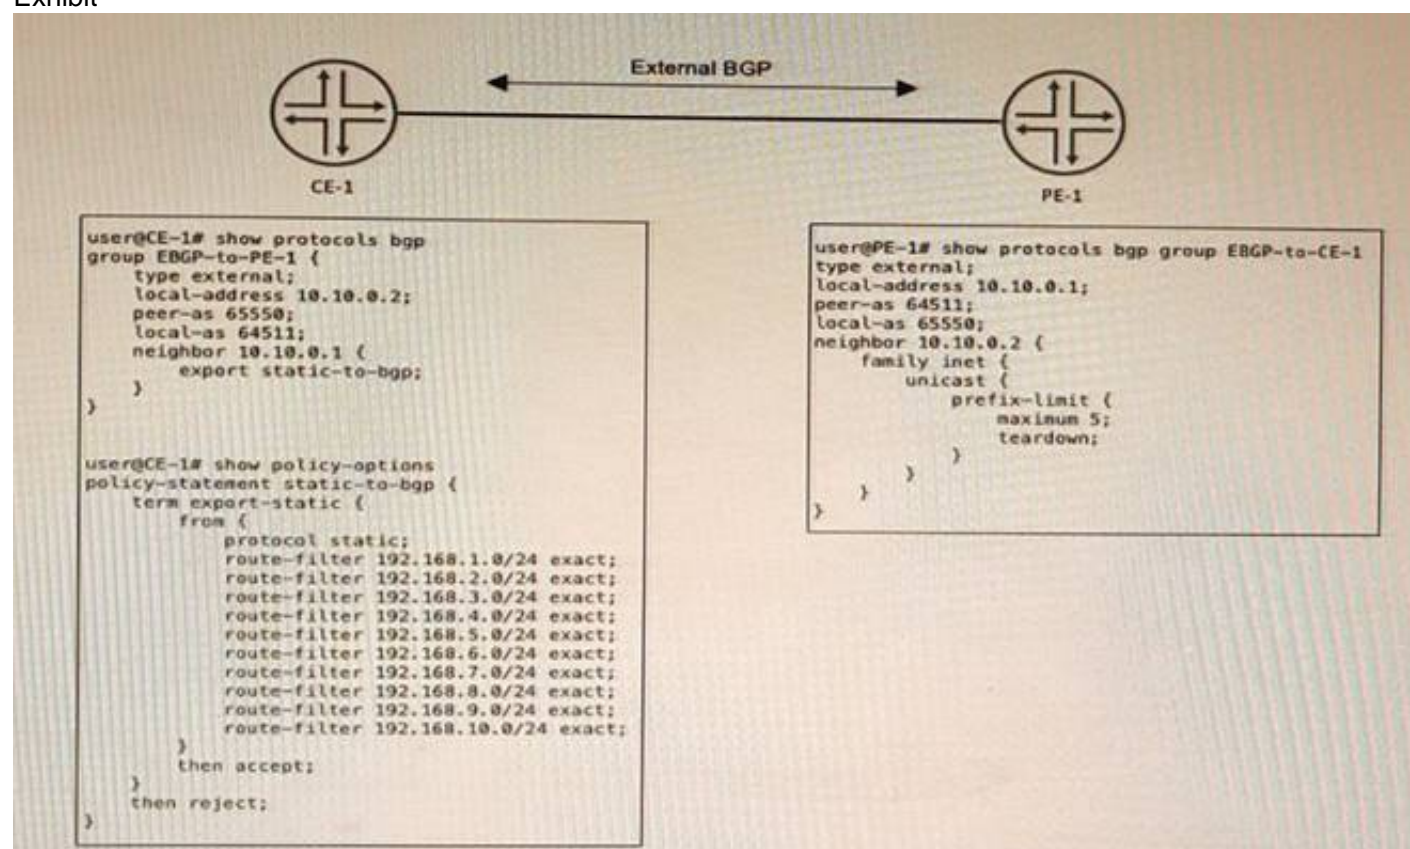
**Answer:** C

#### Explanation:

OSPF uses different types of LSAs to describe different aspects of the network topology. Type 1 LSAs are also known as router LSAs, and they describe the links and interfaces of a router within an area. Type 3 LSAs are also known as summary LSAs, and they describe routes to networks outside an area but within the same autonomous system (AS). By default, OSPF will summarize routes from Type 1 LSAs into Type 3 LSAs when advertising them across area boundaries .

## NEW QUESTION 18

Exhibit



CE-1 must advertise ten subnets to PE-1 using BGP. Once CE-1 starts advertising the subnets to PE-1, the BGP peering state changes to Active. Referring to the CLI output shown in the exhibit, which statement is correct?

- A. CE-1 is advertising its entire routing table.
- B. CE-1 is configured with an incorrect peer AS
- C. The prefix limit has been reached on PE-1
- D. CE-1 is unreachable

**Answer: B**

### Explanation:

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

## NEW QUESTION 22

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 26

Which two statements are correct regarding bootstrap messages that are forwarded within a PIM sparse mode domain? (Choose two.)

- A. Bootstrap messages are forwarded only to routers that explicitly requested the messages within the PIM sparse-mode domain
- B. Bootstrap messages distribute RP information dynamically during an RP election.
- C. Bootstrap messages are used to notify which router is the PIM RP
- D. Bootstrap messages are forwarded to all routers within a PIM sparse-mode domain.

**Answer: BD**

### Explanation:

Bootstrap messages are PIM messages that are used to distribute rendezvous point (RP) information dynamically during an RP election. Bootstrap messages are sent by bootstrap routers (BSRs), which are routers that are elected to perform the RP discovery function for a PIM sparse-mode domain. Bootstrap messages contain information about candidate RPs and their multicast groups, as well as BSR priority and hash mask length. Bootstrap messages are forwarded to all routers within a PIM sparse-mode domain using hop-by-hop flooding.

## NEW QUESTION 28

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

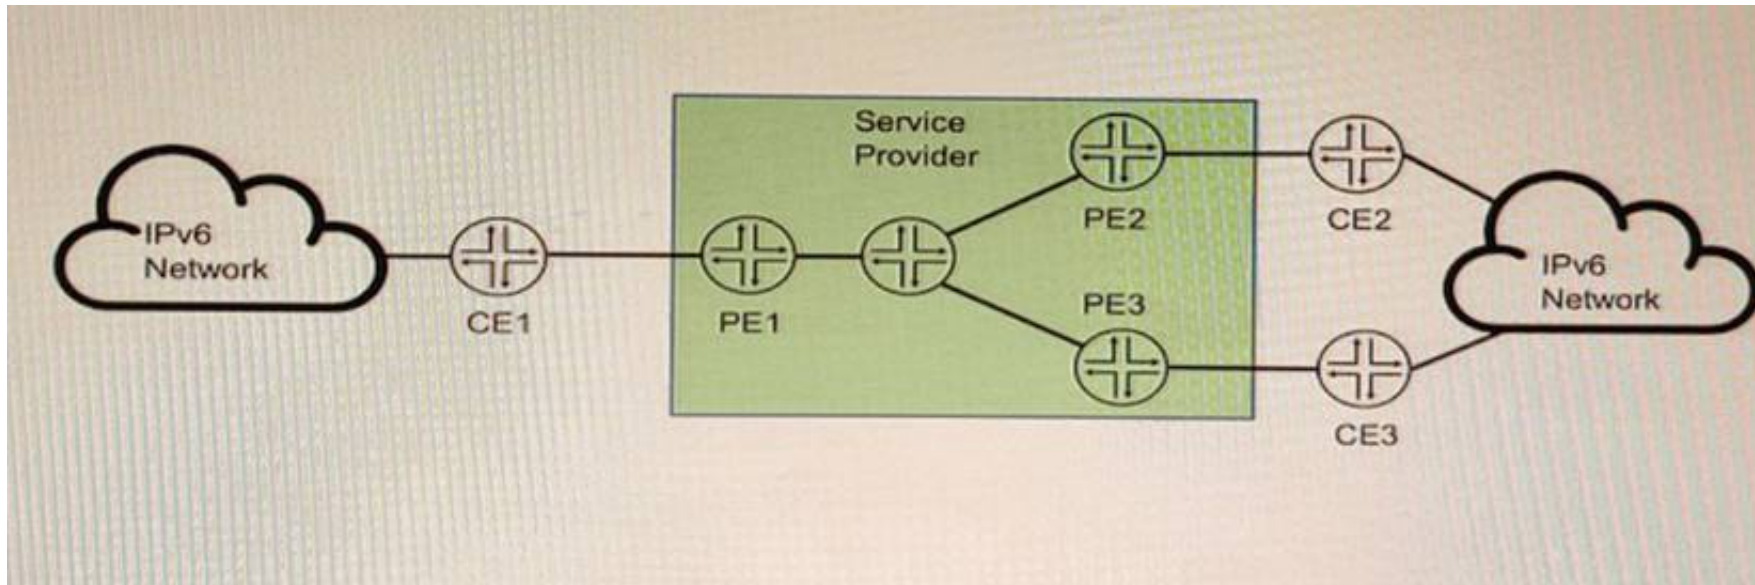
**Answer:** AB

**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.

**NEW QUESTION 33**

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 labeled-unicast between PE routers.
- B. You must configure family inet6 unicast 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

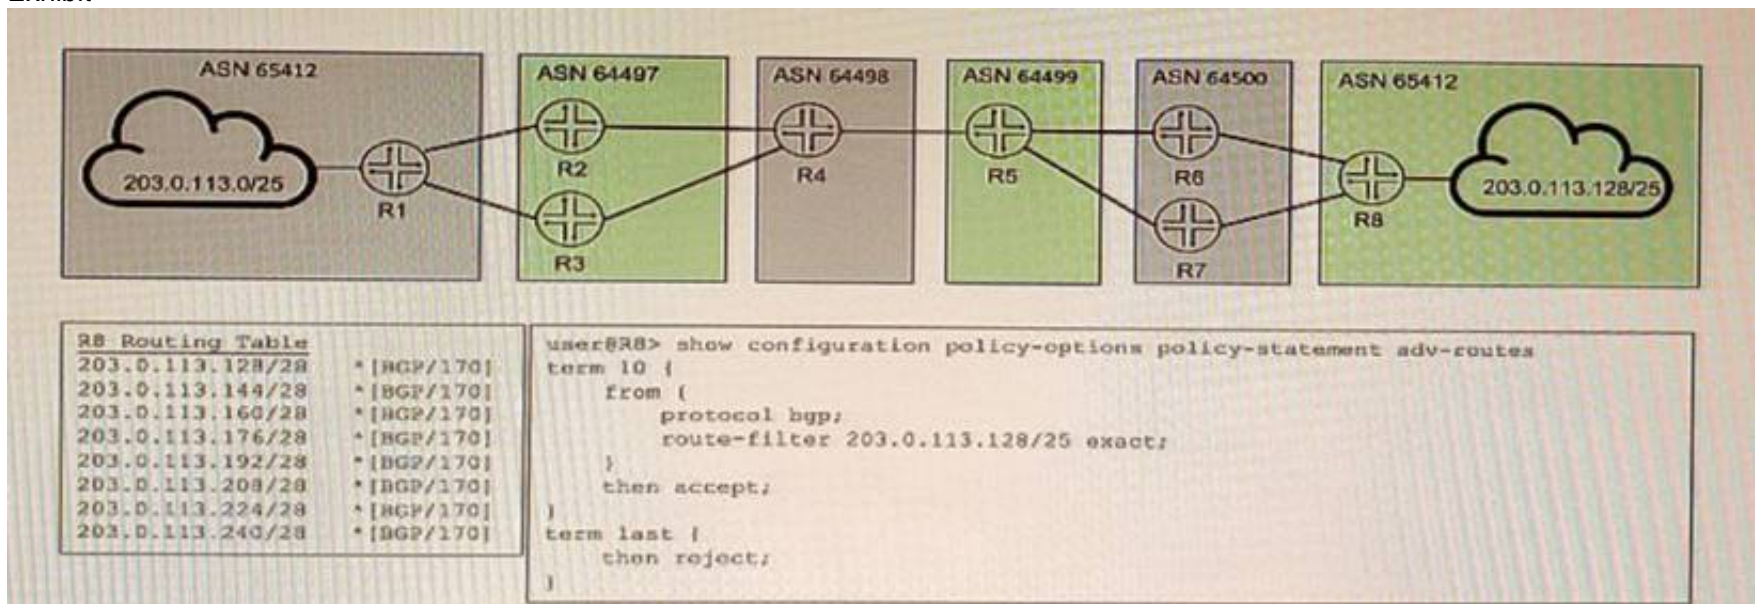
**Answer:** AB

**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.

**NEW QUESTION 35**

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 39

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 44

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.

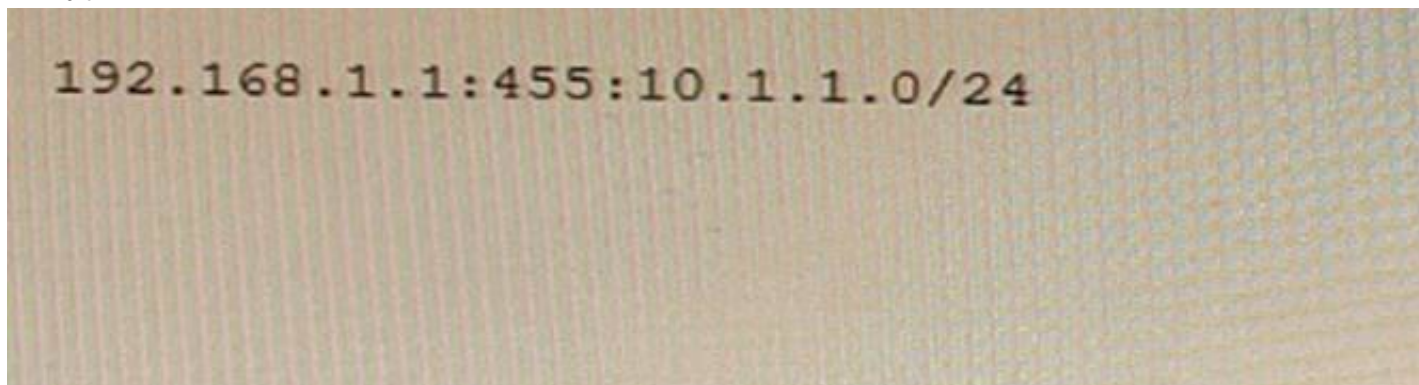
**Answer:** CD

#### 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 49

Exhibit



You are examining an L3VPN route that includes the information shown in the exhibit. Which statement is correct in this scenario?

- A. The information shows a Type 1 route distinguisher.
- B. The information shows a Type 0 route distinguisher.
- C. The information shows a Type 2 route distinguisher.
- D. The information shows a route target.

**Answer:** B

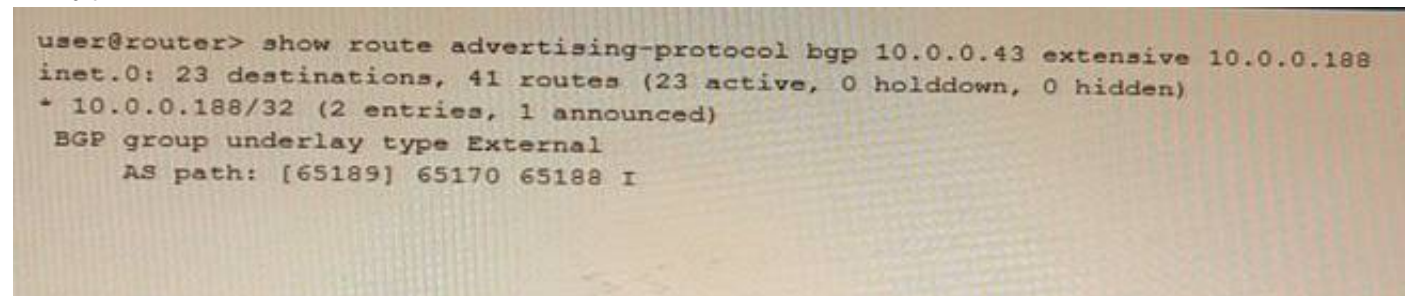
#### Explanation:

The information shows a Type 0 route distinguisher, which is one of the three types of route distinguishers defined by RFC 4364. A route distinguisher is a 64-bit

value that is prepended to an IPv4 address to create a VPN-IPv4 address, which is unique within a VPN routing and forwarding (VRF) table. A Type 0 route distinguisher has two fields: an administrator subfield (2 bytes) and an assigned number subfield (6 bytes). The administrator subfield can be an AS number or an IP address, and the assigned number subfield can be any value assigned by the administrator. In this example, the administrator subfield is 65530 (an AS number) and the assigned number subfield is 1.

## NEW QUESTION 52

Exhibit



```
user@router> show route advertising-protocol bgp 10.0.0.43 extensive 10.0.0.188
inet.0: 23 destinations, 41 routes (23 active, 0 holddown, 0 hidden)
+ 10.0.0.188/32 (2 entries, 1 announced)
  BGP group underlay type External
    AS path: [65189] 65170 65188 I
```

Referring to the exhibit, what do the brackets [ ] in the AS path identify?

- A. They identify the local AS number associated with the AS path if configured on the router, or if AS path prepending is configured
- B. They identify an AS set, which are groups of AS numbers in which the order does not matter
- C. They identify that the autonomous system number is incomplete and awaiting more information from the BGP protocol.
- D. They identify that a BGP confederation is being used to ensure that there are no routing loops.

**Answer:** B

### Explanation:

The brackets [ ] in the AS path identify an AS set, which are groups of AS numbers in which the order does not matter. An AS set is used when BGP aggregates routes from different ASs into a single prefix. For example, if BGP aggregates routes 10.0.0.0/16 and 10.1.0.0/16 from AS 100 and AS 200, respectively, into a single prefix 10.0.0.0/15, then the AS path for this prefix will be [100 200]. An AS set reduces the length of the AS path and prevents routing loops.

## NEW QUESTION 55

.....



## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual JN0-664 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the JN0-664 Product From:

<https://www.2passeasy.com/dumps/JN0-664/>

## Money Back Guarantee

### 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 FirstTry
- \* JN0-664 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year