

Juniper

Exam Questions JN0-664

Service Provider - Professional (JNCIP-SP)



NEW QUESTION 1

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 2

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

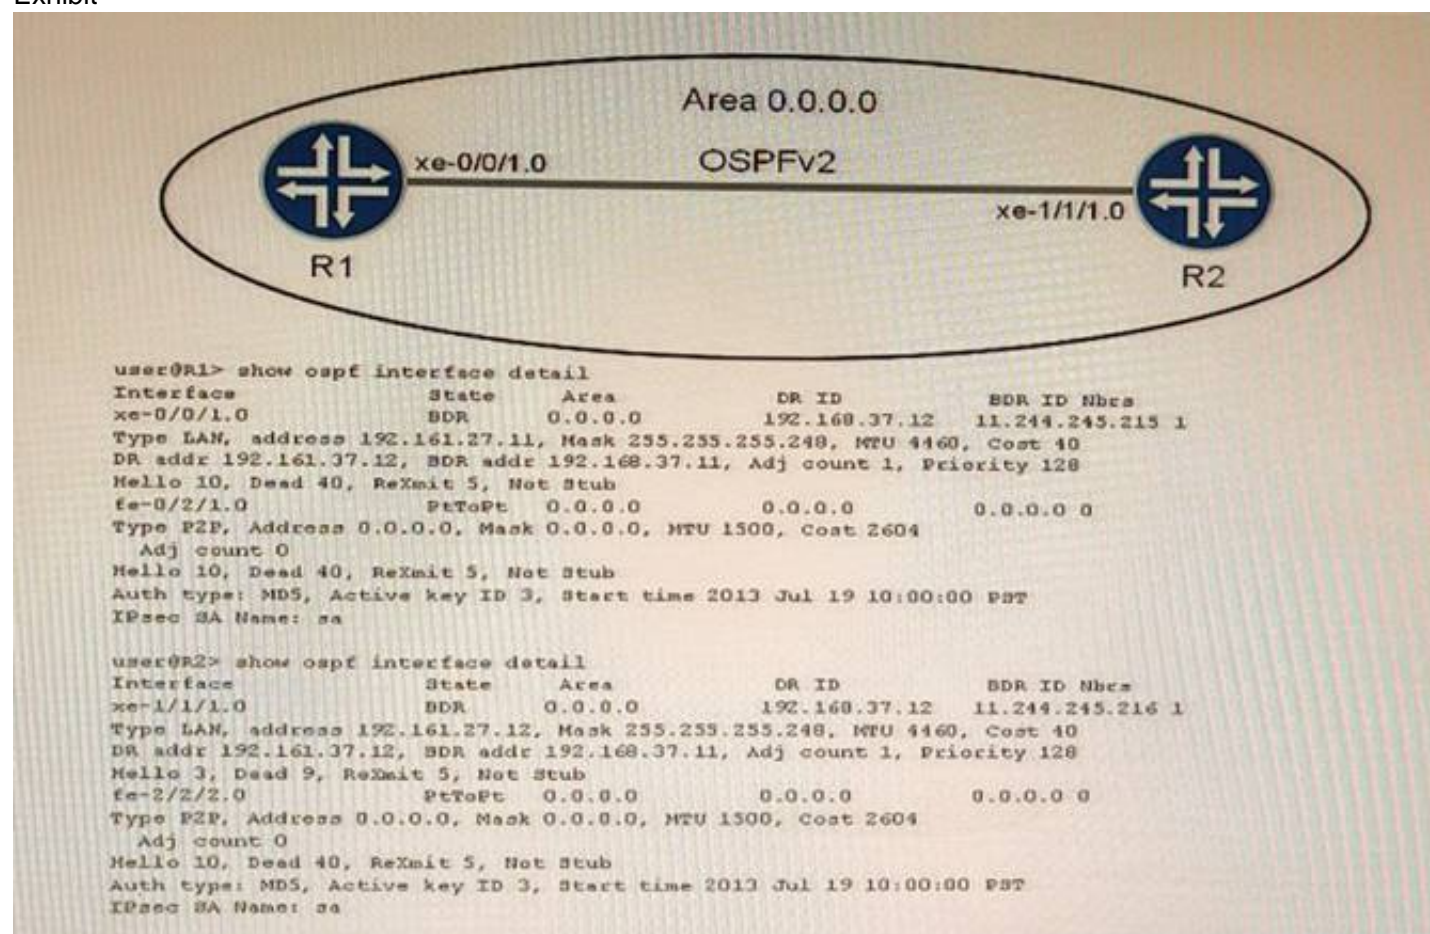
Explanation:

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 3

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 40 seconds.

dead interval of 120 seconds. This causes a mismatch and prevents them from becoming neighbors23.

NEW QUESTION 4

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 5

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 6

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.

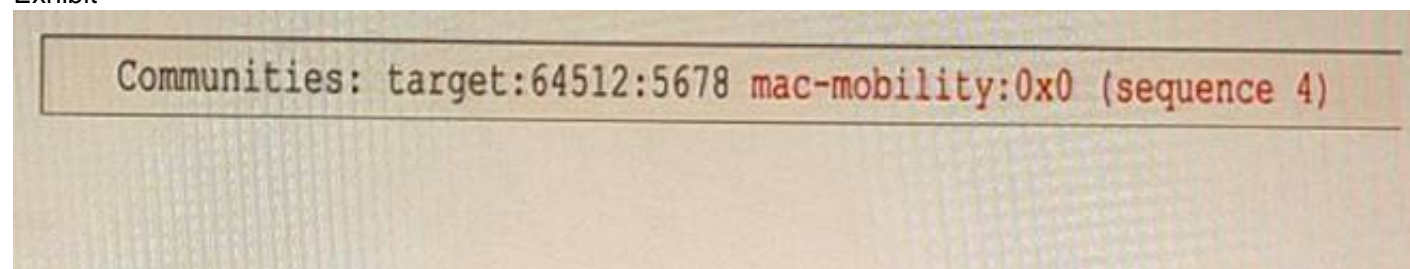
Answer: BC

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 7

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

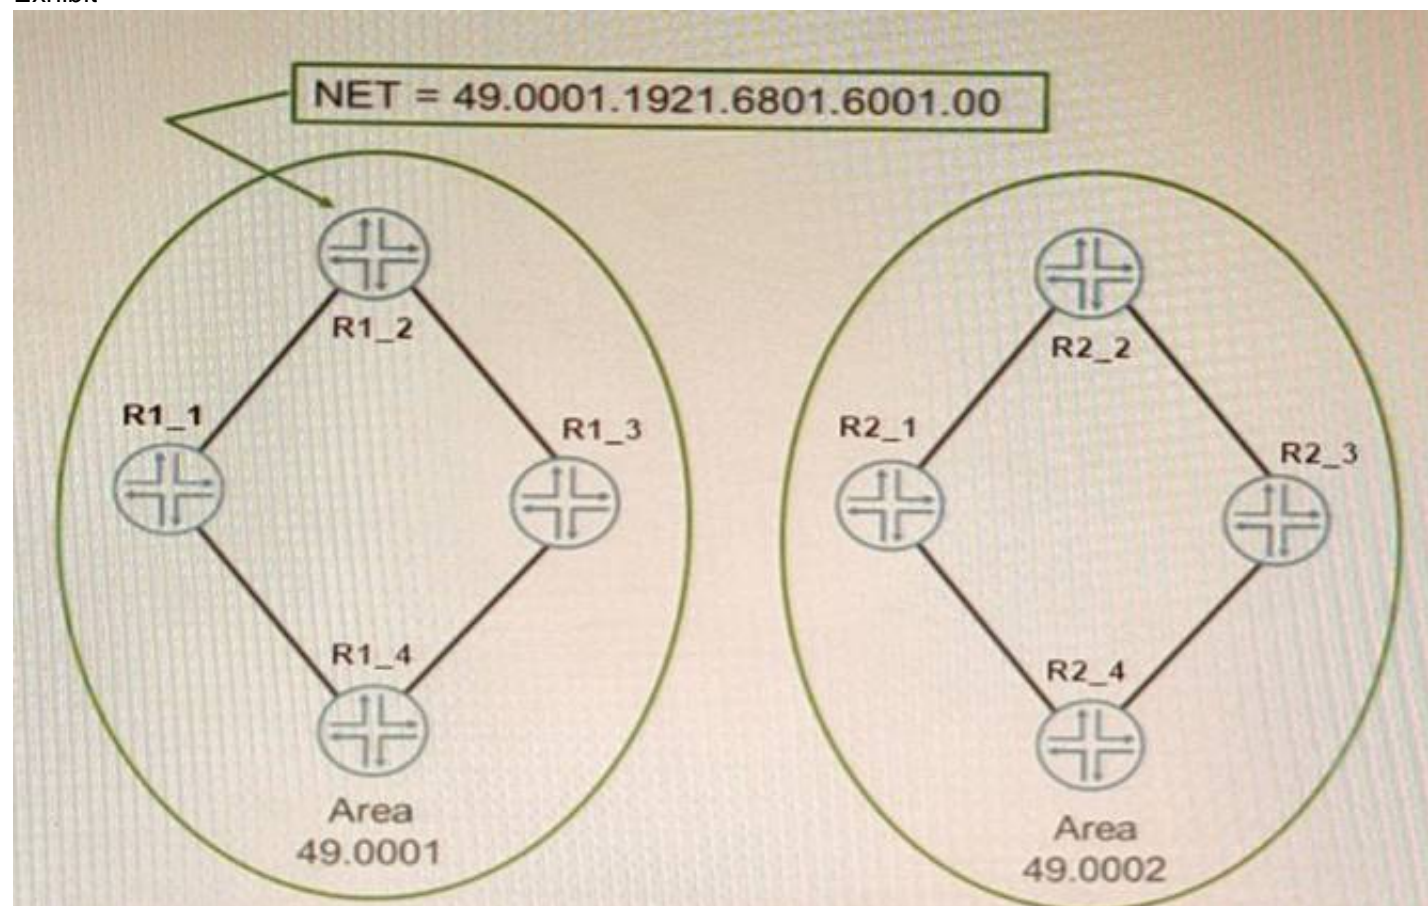
Answer: BC

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 8

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-ovrvw-cf.html 2: <https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing-overview.html>

NEW QUESTION 9

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

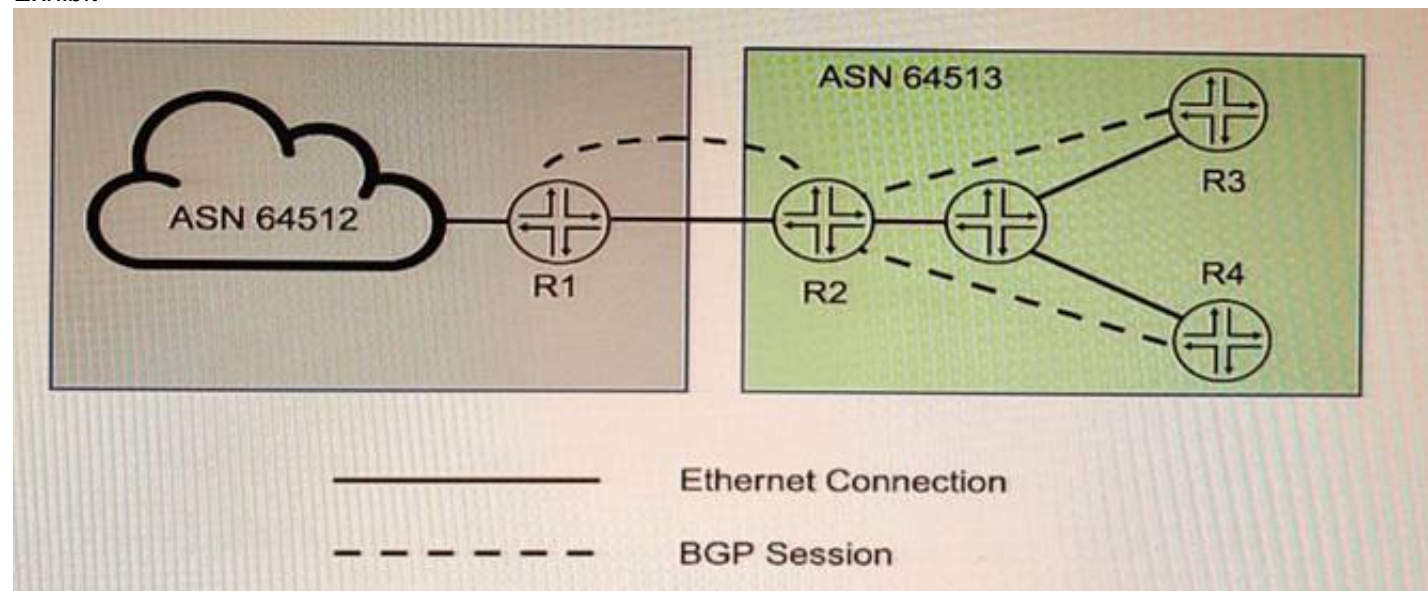
Explanation:

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 actions¹. 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 analyzer². BGP FlowSpec rules are internally converted to equivalent Cisco Common Classification Policy Language (C3PL) representing corresponding match and action parameters². 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

NEW QUESTION 10

Exhibit



You want to implement the BGP Generalized TTL Security Mechanism (GTSM) on the network. Which three statements are correct in this scenario? (Choose three)

- A. You can implement BGP GTSM between R2, R3, and R4
- B. BGP GTSM requires a firewall filter to discard packets with incorrect TTL.
- C. You can implement BGP GTSM between R2 and R1.
- D. BGP GTSM requires a TTL of 1 to be configured between neighbors.
- E. BGP GTSM requires a TTL of 255 to be configured between neighbors.

Answer: ADE

Explanation:

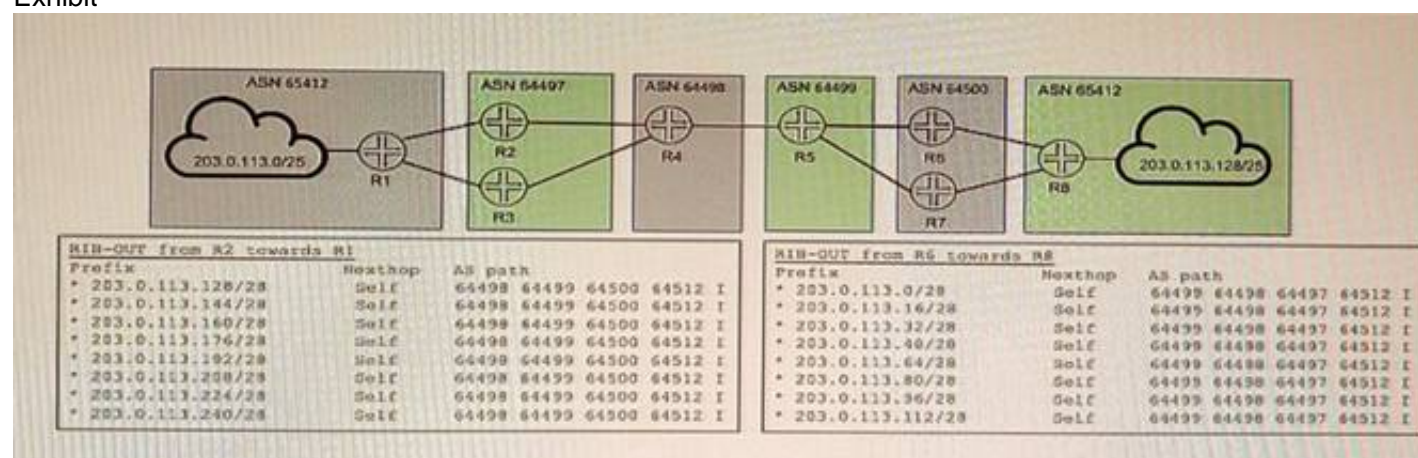
BGP GTSM is a technique that protects a BGP session by comparing the TTL value in the IP header of incoming BGP packets against a valid TTL range. If the TTL value is within the valid TTL range, the packet is accepted. If not, the packet is discarded. The valid TTL range is from 255 – the configured hop count + 1 to 255. When GTSM is configured, the BGP packets sent by the device have a TTL of 255. GTSM provides best protection for directly connected EBGP sessions, but not for multihop EBGP or IBGP sessions because the TTL of packets might be modified by intermediate devices.

In the exhibit, we can see that R2, R3, and R4 are in the same AS (AS 20) and R1 is in a different AS (AS 10). Based on this information, we can infer the following statements:

- ? You can implement BGP GTSM between R2, R3, and R4. This is not correct because R2, R3, and R4 are IBGP peers and GTSM does not provide effective protection for IBGP sessions. The TTL of packets between IBGP peers might be changed by intermediate devices or routing protocols.
- ? BGP GTSM requires a firewall filter to discard packets with incorrect TTL. This is not correct because BGP GTSM does not require a firewall filter to discard packets with incorrect TTL. BGP GTSM uses TCP option 19 to negotiate GTSM capability between peers and uses TCP option 20 to carry the expected TTL value in each packet. The receiver checks the expected TTL value against the actual TTL value and discards packets with incorrect TTL values.
- ? You can implement BGP GTSM between R2 and R1. This is correct because R2 and R1 are EBGP peers and GTSM provides effective protection for directly connected EBGP sessions. The TTL of packets between directly connected EBGP peers is not changed by intermediate devices or routing protocols.
- ? BGP GTSM requires a TTL of 1 to be configured between neighbors. This is not correct because BGP GTSM requires a TTL of 255 to be configured between neighbors. The sender sets the TTL of packets to 255 and the receiver expects the TTL of packets to be 255 minus the configured hop count.
- ? BGP GTSM requires a TTL of 255 to be configured between neighbors. This is correct because BGP GTSM requires a TTL of 255 to be configured between neighbors. The sender sets the TTL of packets to 255 and the receiver expects the TTL of packets to be 255 minus the configured hop count.

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 11

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 15

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

- A. All Type 2 and Type 7 LSAs will be summanzed 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

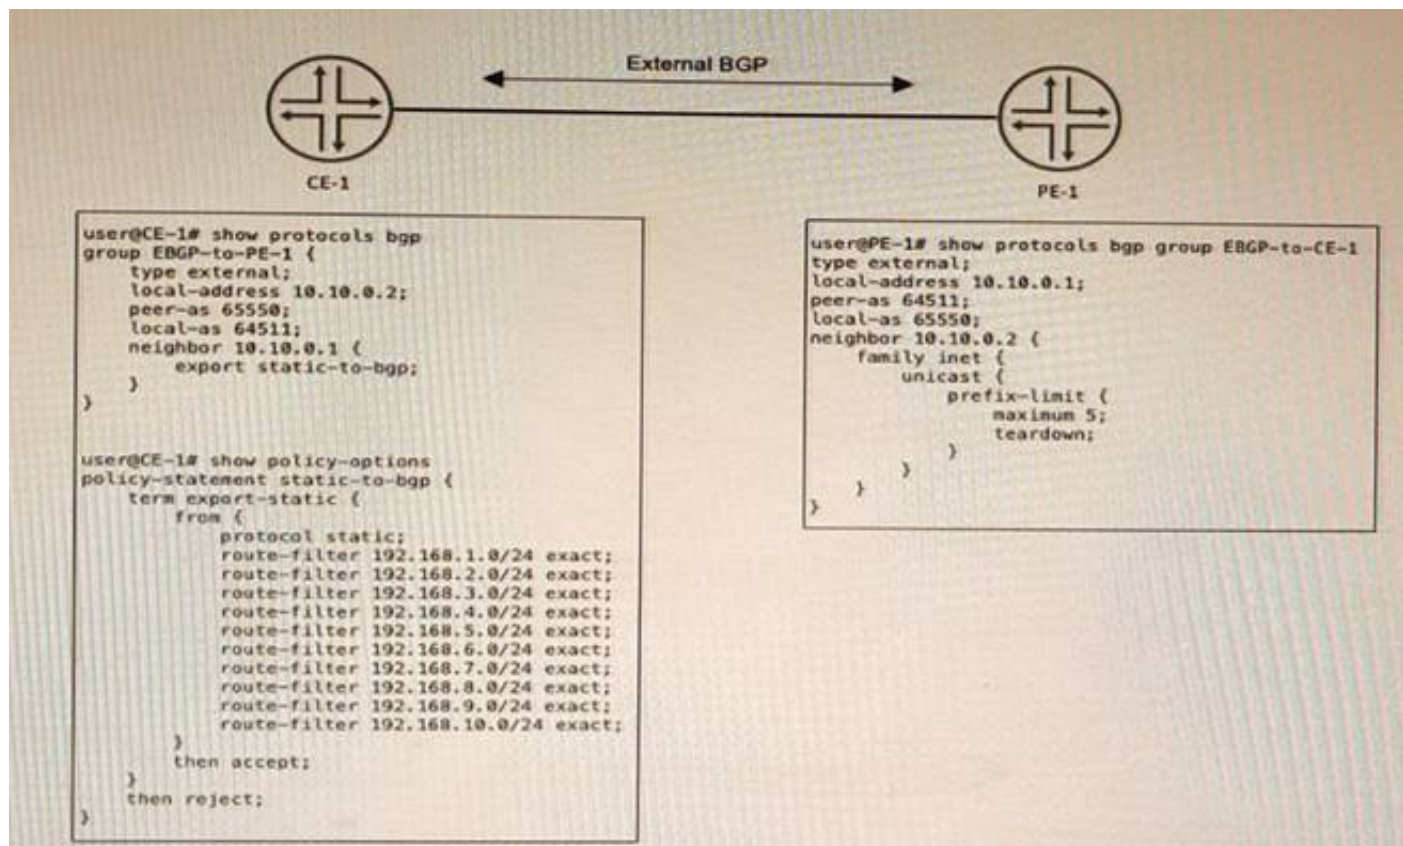
Answer: C

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

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

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

Answer: CD

Explanation:

OSPF is an interior gateway protocol that uses link-state routing to exchange routing information among routers within a single autonomous system. OSPF prevents routing loops in multi-area networks by using two methods: area hierarchy and SPF algorithm. Area hierarchy is the concept of dividing a large OSPF network into smaller areas that are connected to a backbone area (area 0). This reduces the amount of routing information that each router has to store and process, and also limits the scope of link-state updates within each area. All areas are required to connect to area 0 either directly or through virtual links². SPF algorithm is the method that OSPF uses to calculate the shortest path to each destination in the network based on link-state information. The SPF algorithm runs on each router and builds a shortest-path tree that represents the topology of the network from the router's perspective. The SPF algorithm prunes looped paths within an area by choosing only one best path for each destination³.

References: 2: <https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/concept/ospf-area-overview.html> 3:

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

NEW QUESTION 24

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 27

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:

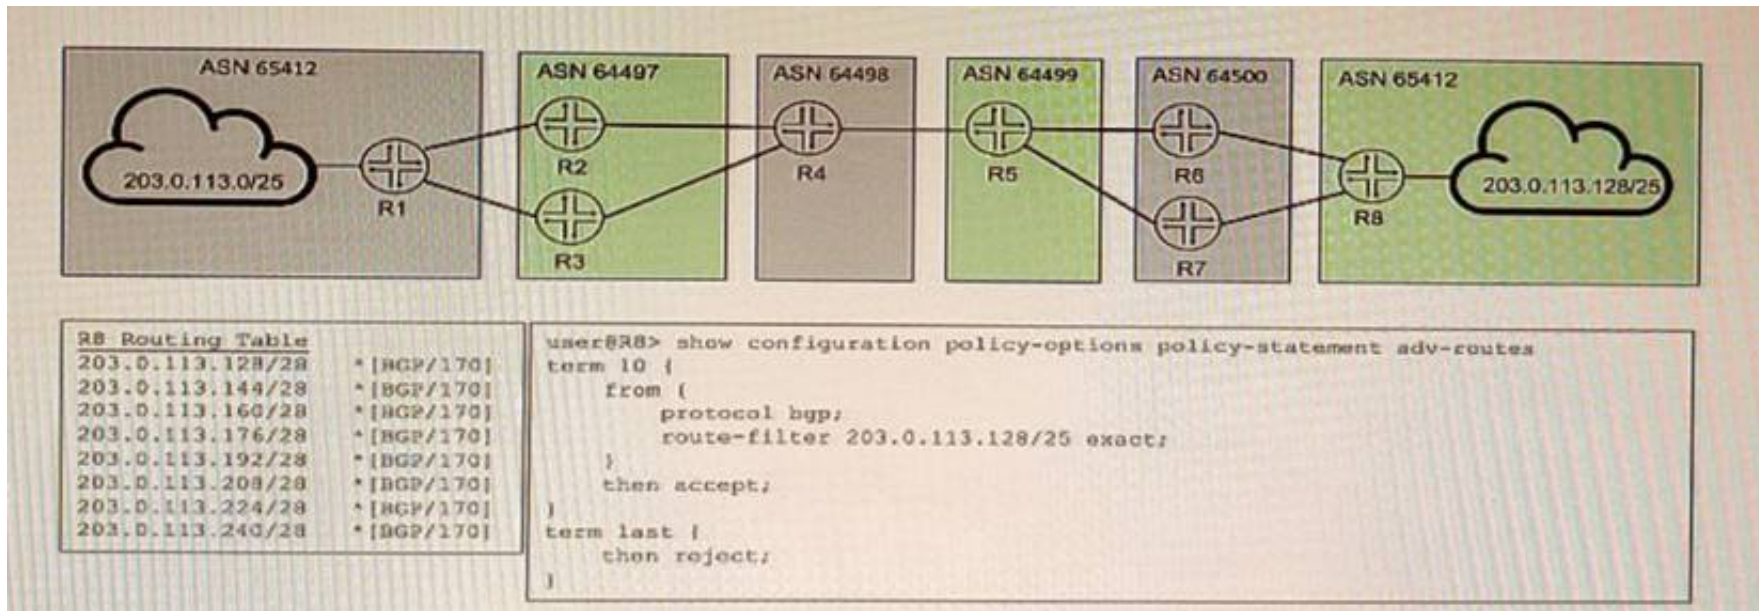
? By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3. This is because the default behavior of EBGP is to preserve the next-hop attribute of the routes received from another EBGP neighbor. The next- hop attribute indicates the IP address of the router that 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 32

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 35

Which two statements are correct about a sham link? (Choose two.)

- A. It creates an OSPF multihop neighborhood between two PE routers.
- B. It creates a BGP multihop neighborhood between two PE routers.
- C. The PEs exchange Type 1 OSPF LSAs instead of Type 3 OSPF LSAs for the L3VPN routes
- D. The PEs exchange Type 3 OSPF LSAs instead of Type 1 OSPF LSAs for the L3VPN routes.

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 selection.

NEW QUESTION 37

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 38

Exhibit


```

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

```

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

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

Answer: AD


Explanation:

According to 1 and 2, BGP Layer 2 VPNs use BGP to distribute endpoint provisioning information and set up pseudowires between PE devices. BGP uses the Layer 2 VPN (L2VPN) Routing Information Base (RIB) to store endpoint provisioning information, which is updated each time any Layer 2 virtual forwarding instance (VFI) is configured. The prefix and path information is stored in the L2VPN database, which allows BGP to make decisions about the best path. In the output shown in the exhibit, we can see some information about the L2VPN RIB and the pseudowire state. Based on this information, we can infer the following statements:

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

NEW QUESTION 41

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

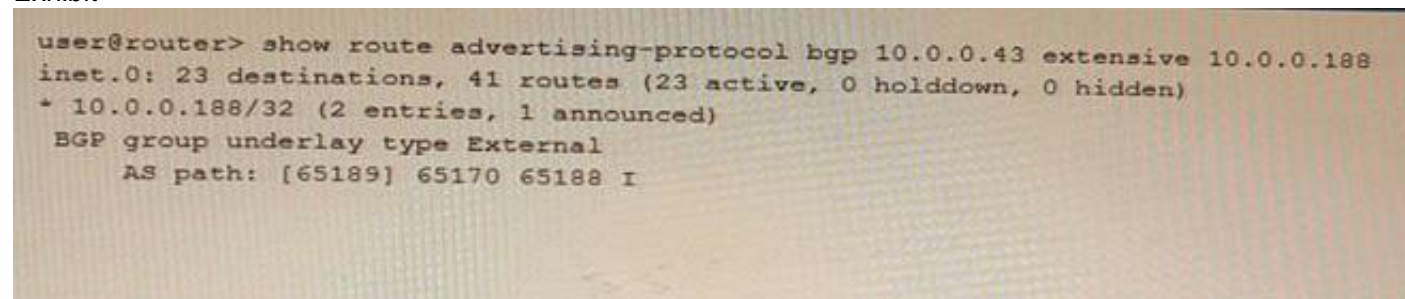
Answer: A

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 policy². 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 42

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 46

.....

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](#)