



**Juniper**

## **Exam Questions JN0-664**

Service Provider - Professional (JNCIP-SP)

### 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.

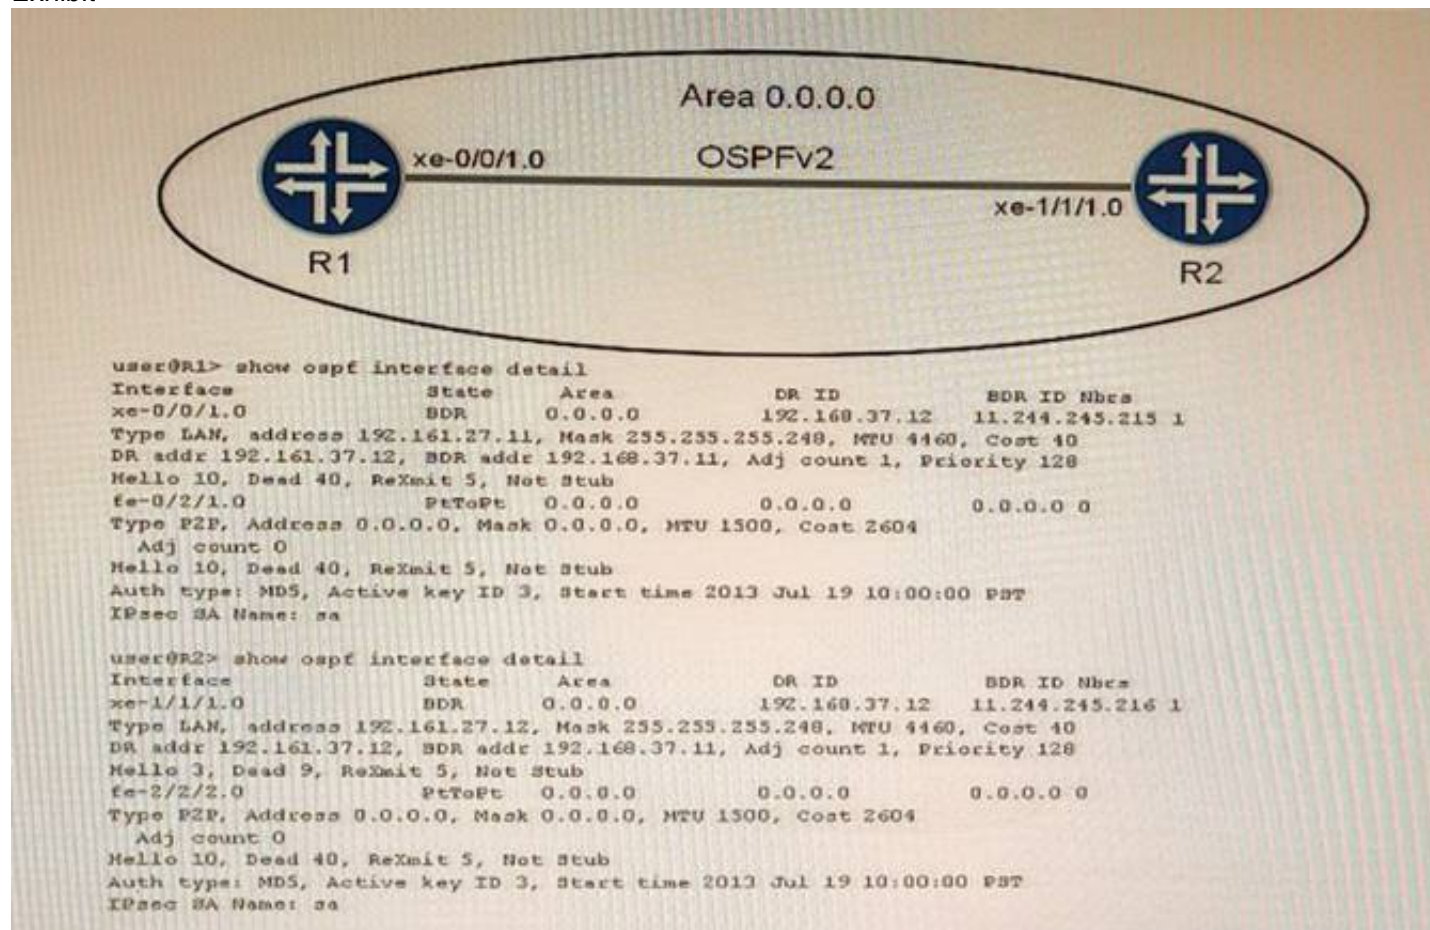
**Answer: C**

#### 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 devices1. 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 VPNs2. 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

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 neighbors23.

### 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



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

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

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

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 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

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

- A. LDP-signaled VPLS tunnels only support control bit 0.
- B. LDP-signaled VPLS tunnels use auto-discovery to provision sites
- C. BGP-signaled VPLS tunnels can use either RSVP or LDP between the PE routers.
- D. BGP-signaled VPLS tunnels require manual provisioning of sites.

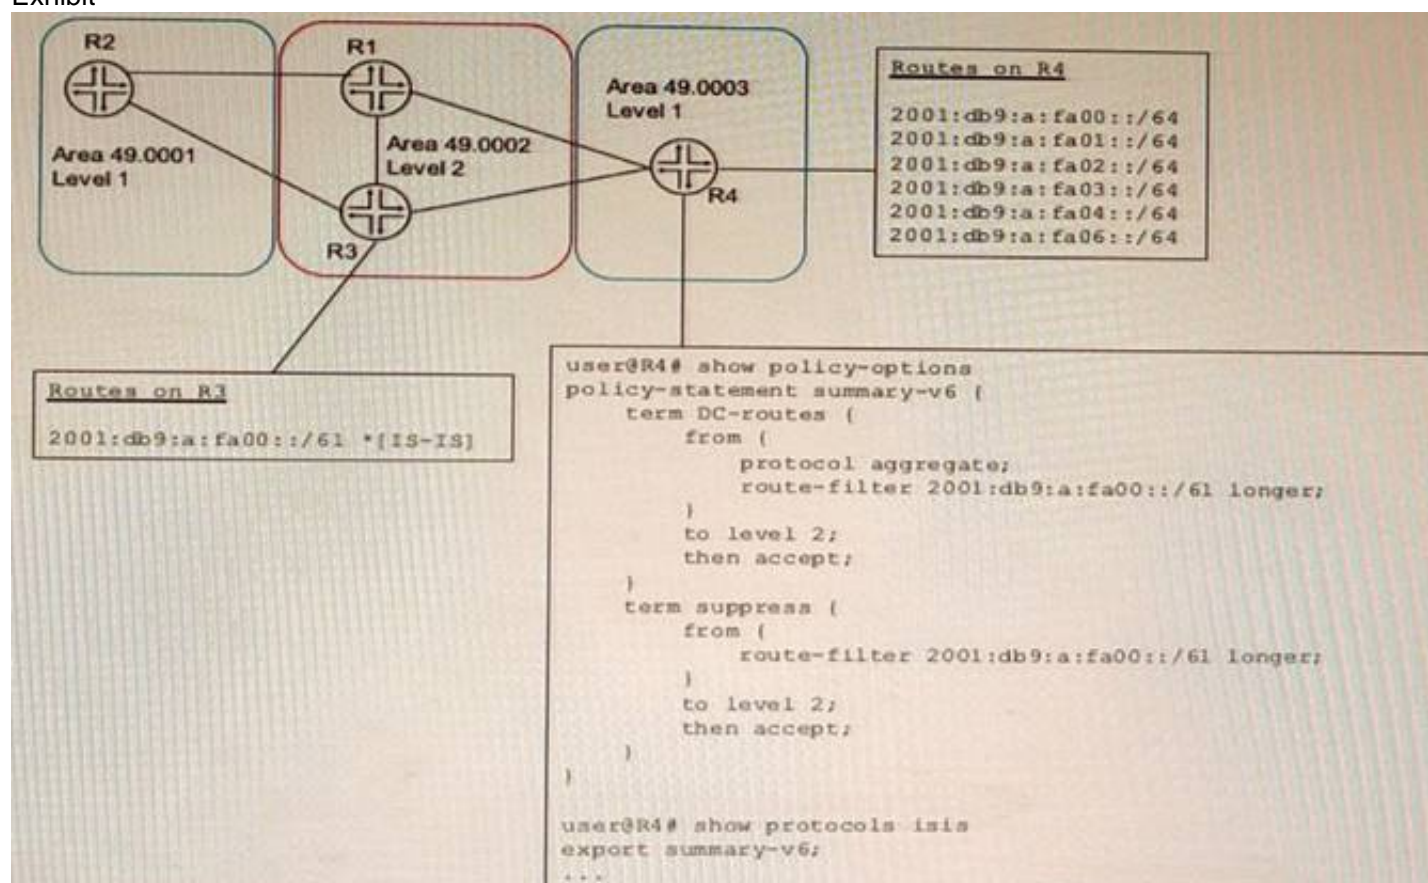
**Answer:** BC

#### Explanation:

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

#### NEW QUESTION 7

Exhibit



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

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

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

**Answer:** BCD

#### Explanation:

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

#### NEW QUESTION 8

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 9**

Exhibit



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

Which action will accomplish this goal?

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

**Answer: D**

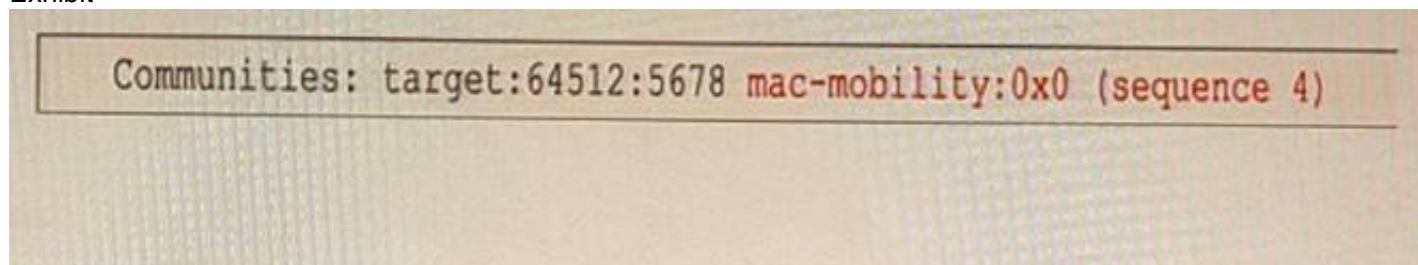
**Explanation:**

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

Reference: 4: <https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/137544-technote-mpls-00.html>

**NEW QUESTION 10**

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**

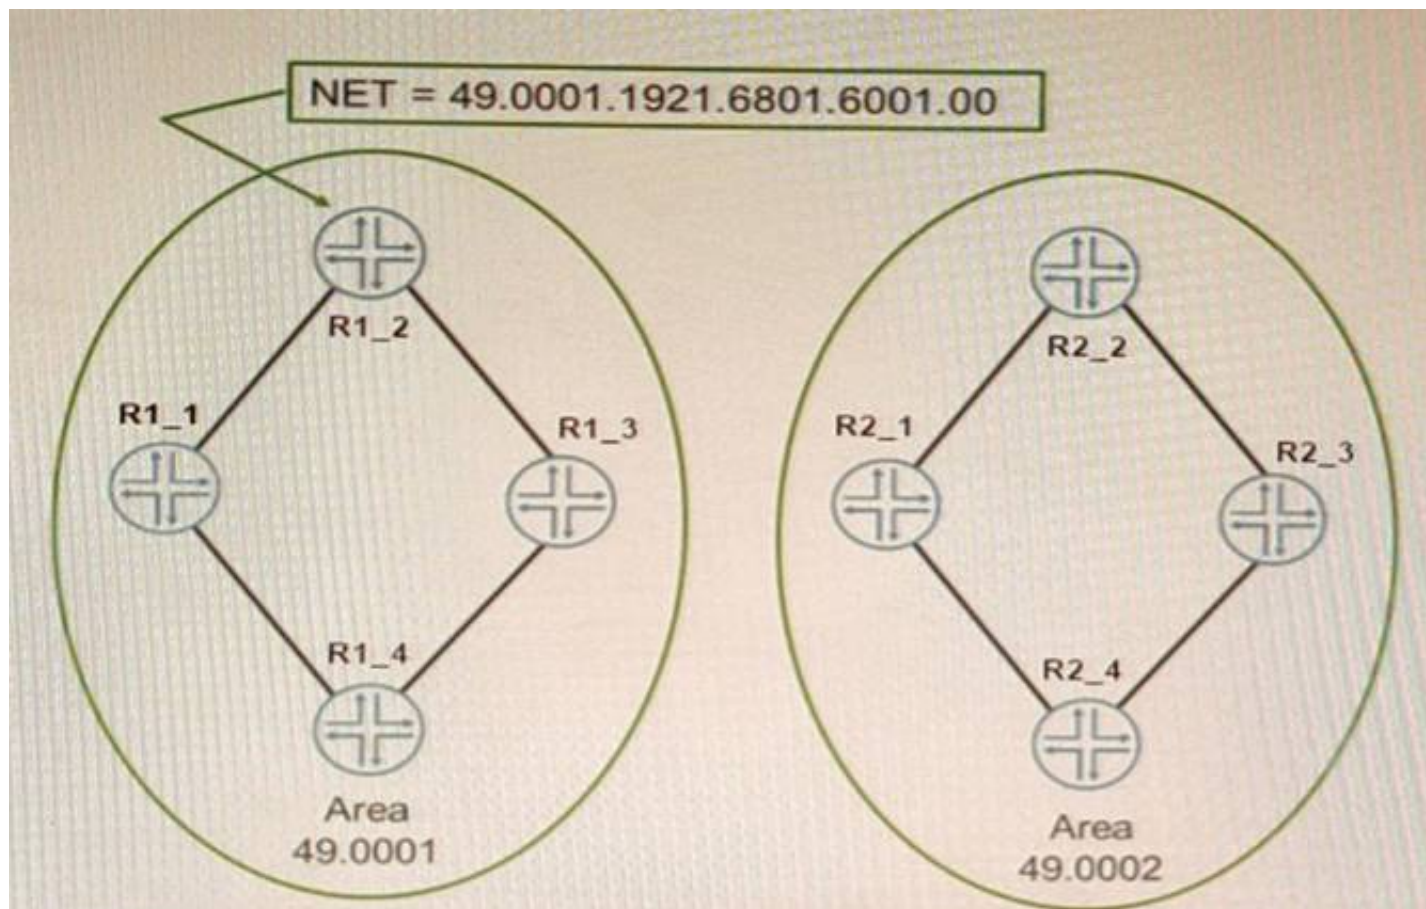
**Explanation:**

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 10**

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/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/isis/topics/concept/isis-routing-overview.html>

**NEW QUESTION 11**

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<sup>3</sup>. 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 14

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 19

You want to ensure that L1 IS-IS routers have only the most specific routes available from L2 IS-IS routers. Which action accomplishes this task?

- A. Configure the ignore-attached-bit parameter on all L2 routers.
- B. Configure all routers to allow wide metrics.
- C. Configure all routers to be L1.
- D. Configure the ignore-attached-bit parameter on all L1 routers

**Answer:** D

#### Explanation:

The attached bit is a flag in an IS-IS LSP that indicates whether a router is connected to another area or level (L2) of the network. By default, L2 routers set this bit when they advertise their LSPs to L1 routers, and L1 routers use this bit to select a default route to reach other areas or levels through L2 routers. However, this may result in suboptimal routing if there are multiple L2 routers with different paths to other areas or levels. To ensure that L1 routers have only the most specific routes available from L2 routers, you can configure the ignore-attached-bit parameter on all L1 routers. This makes L1 routers ignore the attached bit and install all interarea routes learned from L2 routers in their routing tables.

#### NEW QUESTION 21

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 24

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 25**

What is the correct order of packet flow through configurable components in the Junos OS CoS features?

- A. Multifield Classifier -> Behavior Aggregate Classifier -> Input Policer -> Forwarding Policy Options -> Fabric Scheduler -> Output Policer -> Rewrite Marker -> Scheduler/Shaper/RED
- B. Behavior Aggregate Classifier -> Multifield Classifier -> Input Policer -> Forwarding Policy Options -> Fabric Scheduler -> Output Policer -> Scheduler/Shaper/RED -> Rewrite Marker
- C. Behavior Aggregate Classifier -> Input Policer -> Multifield Classifier -> Forwarding Policy Options -> Fabric Scheduler -> Output Policer -> Scheduler/Shaper/RED -> Rewrite Marker
- D. Behavior Aggregate Classifier -> Multifield Classifier -> Input Policer -> Forwarding Policy Options -> Fabric Scheduler -> Scheduler/Shaper/RED -> Output Policer -> Rewrite Marker

**Answer:** C

**Explanation:**

The correct order of packet flow through configurable components in the Junos OS CoS features is as follows:

? Behavior Aggregate Classifier: This component uses a single field in a packet header to classify traffic into different forwarding classes and loss priorities based on predefined or user-defined values.

? Input Policer: This component applies rate-limiting and marking actions to incoming traffic based on the forwarding class and loss priority assigned by the classifier.

? Multifield Classifier: This component uses multiple fields in a packet header to classify traffic into different forwarding classes and loss priorities based on user-defined values and filters.

? Forwarding Policy Options: This component applies actions such as load balancing, filtering, or routing to traffic based on the forwarding class and loss priority assigned by the classifier.

? Fabric Scheduler: This component schedules traffic across the switch fabric based on the forwarding class and loss priority assigned by the classifier.

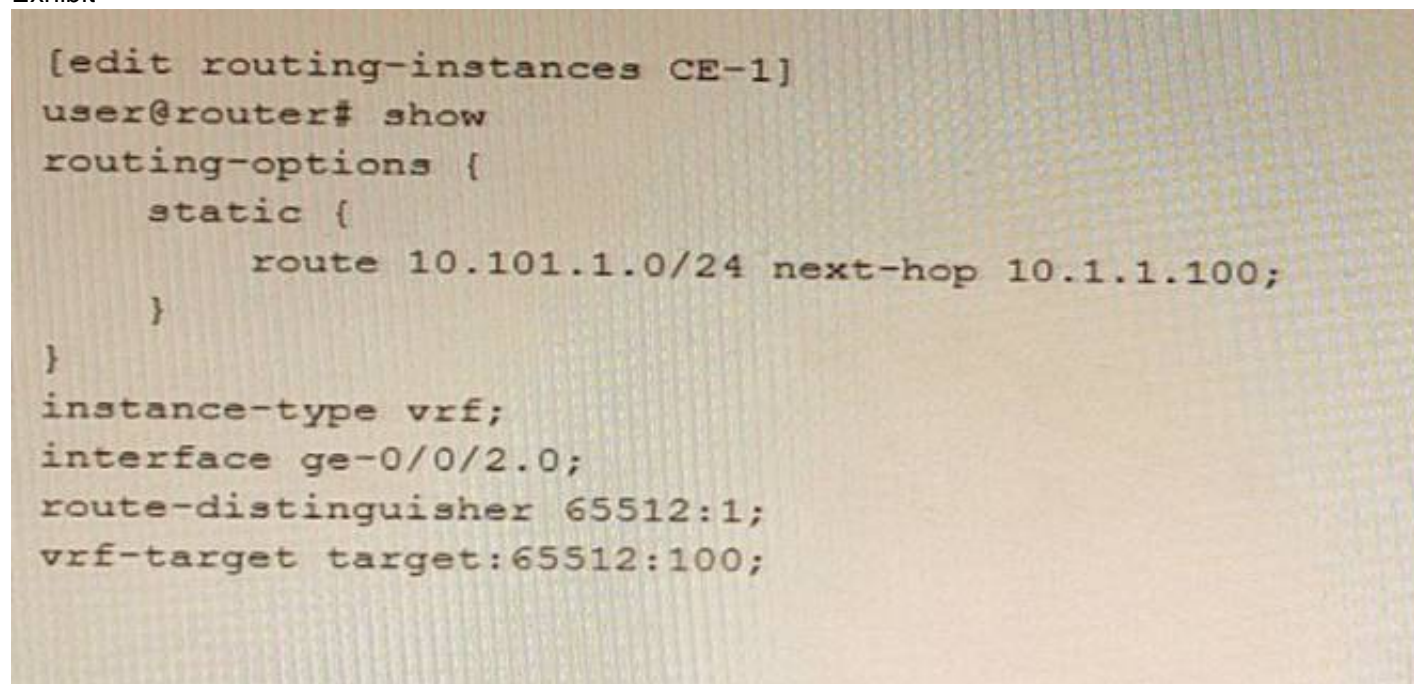
? Output Policer: This component applies rate-limiting and marking actions to outgoing traffic based on the forwarding class and loss priority assigned by the classifier.

? Scheduler/Shaper/RED: This component schedules, shapes, and drops traffic at the egress interface based on the forwarding class and loss priority assigned by the classifier.

? Rewrite Marker: This component rewrites the code-point bits of packets leaving an interface based on the forwarding class and loss priority assigned by the classifier.

**NEW QUESTION 27**

Exhibit



```
[edit routing-instances CE-1]
user@router# show
routing-options {
    static {
        route 10.101.1.0/24 next-hop 10.1.1.100;
    }
}
instance-type vrf;
interface ge-0/0/2.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
```

Referring to the exhibit, which statement is true?

- A. The 10.101.1.0/24 route will be shared if the vrf-table-label parameter is configured.
- B. The 10.101.1.0/24 route will only be shared if BGP is configured in the routing instance
- C. The 10.101.1.0/24 route will be shared if there are other VRFs that use the same route target community
- D. The 10.101.1.0/24 route will be shared if the auto-export parameter is configured

**Answer:** D

**Explanation:**

The auto-export parameter is a routing option that allows a routing instance to share routes with other routing instances or the master routing table. The auto-export parameter automatically exports routes from one routing instance to another based on the route target communities attached to the routes. In this scenario, the 10.101.1.0/24 route will be shared if the auto-export parameter is configured under [edit routing-options] hierarchy level.

**NEW QUESTION 32**

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

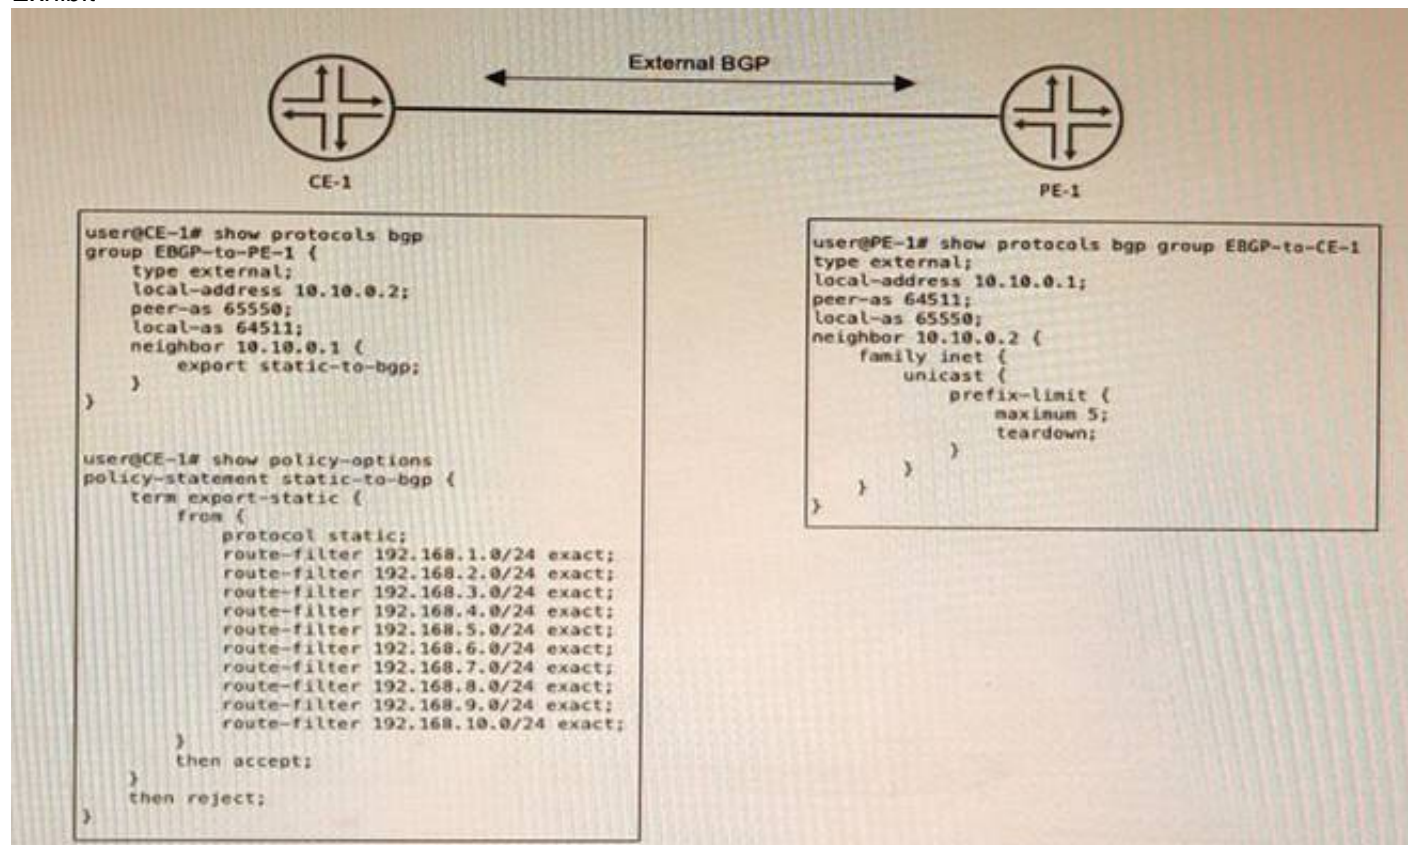
**Answer: C**

**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 36**

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 38**

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

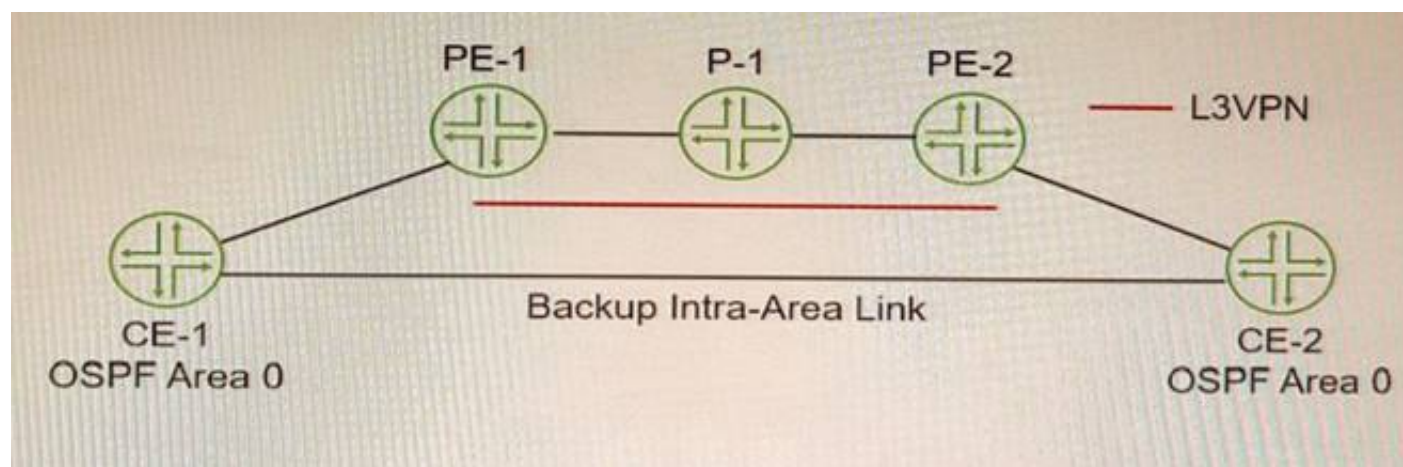
**Answer: C**

**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 40**

Exhibit



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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 43**

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

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

**Answer:** AB

**Explanation:**

Route reflection is a BGP feature that allows a router to reflect routes learned from one IBGP peer to another IBGP peer, without requiring a full-mesh IBGP topology. Route reflectors do not change any existing BGP attributes by default when advertising routes, unless explicitly configured to do so. A BGP peer does not require any configuration changes to become a route reflector client, only the route reflector needs to be configured with the client parameter under [edit protocols bgp group group-name neighbor neighbor- address] hierarchy level.

**NEW QUESTION 47**

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 ospf3 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 51**

Exhibit





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

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

**Answer: AB**

**Explanation:**

BGP is an exterior gateway protocol that uses path vector routing to exchange routing information among autonomous systems. BGP uses various attributes to select the best path to each destination and to propagate routing policies. Some of the common BGP attributes are AS path, next hop, local preference, MED, origin, weight, and community. BGP attributes can be classified into four categories: well-known mandatory, well-known discretionary, optional transitive, and optional nontransitive. Well-known mandatory attributes are attributes that must be present in every BGP update message and must be recognized by every BGP speaker. Well-known discretionary attributes are attributes that may or may not be present in a BGP update message but must be recognized by every BGP speaker. Optional transitive attributes are attributes that may or may not be present in a BGP update message and may or may not be recognized by a BGP speaker. If an optional transitive attribute is not recognized by a BGP speaker, it is passed along to the next BGP speaker. Optional nontransitive attributes are attributes that may or may not be present in a BGP update message and may or may not be recognized by a BGP speaker. If an optional nontransitive attribute is not recognized by a BGP speaker, it is not passed along to the next BGP speaker. In this question, we have four routers (R1, R2, R3, and R4) that are connected in a full mesh topology and running IBGP. R3 receives the 192.168.0.0/16 route from its EBGP neighbor and advertises it to R1 and R4 with different BGP attribute values. We are asked which statements are correct about the BGP routes on R3 that are learned from the ISP-A neighbor. Based on the information given, we can infer that the correct statements are:

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

? The BGP local-preference value that is used by ISP-A is not advertised to R3. This is because the local-preference attribute is a well-known discretionary attribute that is used to influence the outbound traffic from an autonomous system. The local-preference attribute is only propagated within an autonomous system and is not advertised to external neighbors.

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

**NEW QUESTION 55**

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

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


**Answer: CD**

**Explanation:**

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

**NEW QUESTION 58**

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 60**

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

- A. It creates an OSPF multihop neighborship between two PE routers.
- B. It creates a BGP multihop neighborship 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 neighborship 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<sup>1</sup>.

**NEW QUESTION 61**

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 65**

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 67

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:**

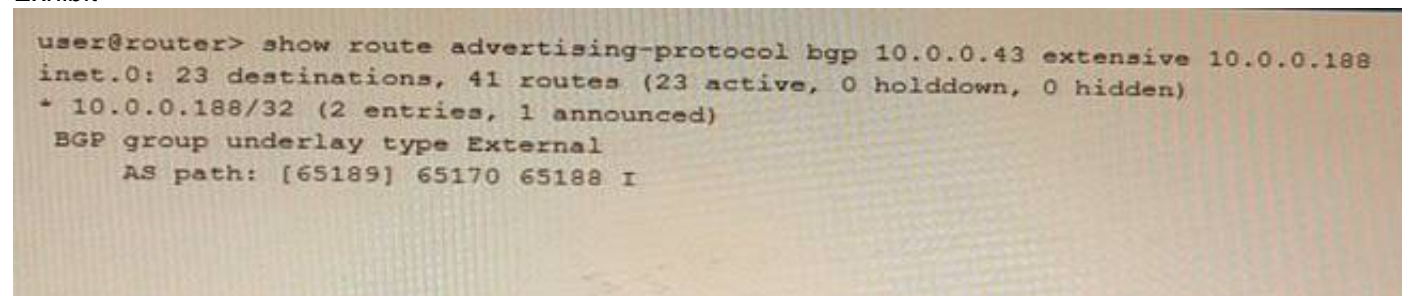
The problem in this scenario is that R1 and R8 are not receiving each other's routes because of private AS numbers in the AS path. Private AS numbers are not globally unique and are not advertised to external BGP peers. To solve this problem, you need to do the following:

? 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 68**

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 73**

.....



## 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](#)**