www.atmicnetworks.com

ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

www.atmicnetworks.com

Question: 1

Exhibit.

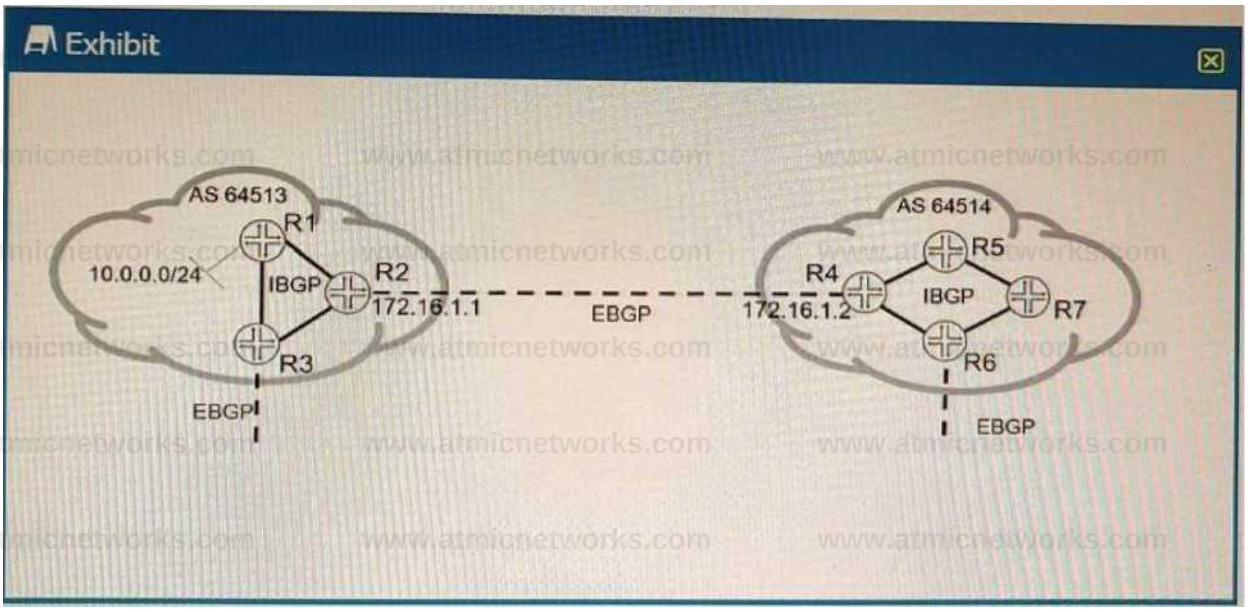

Referring to the exhibit; the 10.0.0.0/24 EBGP route is received on R5; however, the route is being hidden.

What are two solutions that will solve this problem? (Choose two.)

A. On R4, create a policy to change the BGP next hop to 172.16.1.1 and apply it to IBGP as an export policy.

B. On R4, create a policy to change the BGP next hop to itself and apply it to IBGP as an export policy.

C. On R4, add the internal IBGP interface prefixes to the BGP routing tables.

D. On R4, add the external EBGP interface's prefix to the IGP routing tables.

Answer: BD

Explanation:

Question: 2

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.

Question: 3

Exhibit.

```
userBRl# show interfaces
ge-1/2/3 ( unit 0 (
        description tc-R2;
        farr.ily inet {
            address 10.1.1.1/30;

        family iso; )
loO (
    unit 0 { family inet {
            address 192.168.16.1/32; J family iso (
            address 49.0001.1921.6801.6001.00; }


usergRl* show protocols isis (
    interface ge-1/2/3.0 { level 2 disable;

    interface loO.O {
        level 1 disable;
```

www.atmicnetworks.com

```
user@R2s show interfaces
ge-1/2/3 {
    unit 0 ( description to-Rl; family inet ( address 10.1.1.2/30;
        } family iso;

loO { unit 0 ( family inet { address 192.168.16.2/32;

        family iso ( address 49.0001.1921.6801.6002.0 )

u<er9R2* show protocols
isis ( interface ge-1/2/3.0 { level 1 disable;

    interface IcO.O { 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 l disable parameter under the R1 protocols isis interface lo0.0 hierarchy to the level 2 disable parameter.

B. Add IP addresses to the interface ge-1/2/3 unit 0 family iso hierarchy on both R1 and R2.

C. Remove the level 1 disable parameter under the R2 protocols isis interface lo0.0 configuration hierarchy.

D. Change the level 1 disable parameter under the R2 protocols isis interface ge-I/2/3.0 hierarchy to the level 2 disable parameter.

Answer: D

Explanation:

Question: 4

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: B

Explanation:

Question: 5

Exhibit

```
user?router> sho* 12vpn connections
Layer-2 VFK connections: Legend for connection status (St) El -- encapsulation invalid KC EH —
encapsulation mismatch WE vc-On -- virtual circuit down NF CH — control-word mismatch -> CM —
circuit not provisioned <OR — out of   — interface encapsulation not CCC/TCC/VPLS
range                                Up — interface and instance encaps not same
OL -- no outgoing label              On — interface hardware not present
LD — local site signaled down CF RD —— only outbound connection is up
remote site signaled down SC LN — local— only inbound connection is up
site not designated LM RN -- remote site— operational
net   designated   RK   xx   —   unknown-- down
connection status IL KM — MTU mismatch— call admission control failure
                                     MI — local and remote site ID collision
BK — Backup connection               ST — local site ID not minimum designated
PF — Profile parse failure PB RS --  — remote site ID not minimum designated
remote site standby                  SN — no incoming label
LB — Local site not best-site RB VM —— Mesh-Group ID not available
VLAN ID mismatch                     HS — Standby connection
Legend for interface status up --    — Profile busy
operational Dn — down instance: vpn-A— Static Neighbor
Edge protection: Net-Primary         — Remote site not best-site
  Local site: CE1-2 (2) connection-  — Hot-standby Connection
    site Type st          Time
      1              rmt Up     Apr
        Remote PE: 172.17.20.1, Negot.
        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
                                     last up                    • up trans
                                     1 14:35:27 2020                 1
                                     ated control-word: Yes (Null)
```

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: AB

Explanation:

The output is from the show l2vpn connections command on a Juniper router. This command is used to verify the status of Layer 2 VPN (L2VPN) pseudowires between Provider Edge (PE) routers.

Breakdown of Key Information:

Instance: vpn-A

This is the L2VPN instance being monitored.

Connection Status (St)

The connection status is "Up", meaning the pseudowire is operational.

Local Site: CE1-2 (2)

The PE router is attached to a single local site (CE1-2).

Uptime & Connection Flaps

The output shows the last time the connection was up:

Time last up: Apr 11 14:35:27 2020

The "# Up trans" value is 1, meaning this connection has been established once and has not flapped since it was initiated.

VLAN ID Mismatch Check

The legend includes "VM – VLAN ID mismatch", but this status is not present in the connection output.

This means there is NO VLAN ID mismatch.

Flow Labels

The Flow Label Transmit is No, and the Flow Label Receive is No.

This means the PE router does NOT have the capability to pop flow labels.

Question: 6

Exhibit

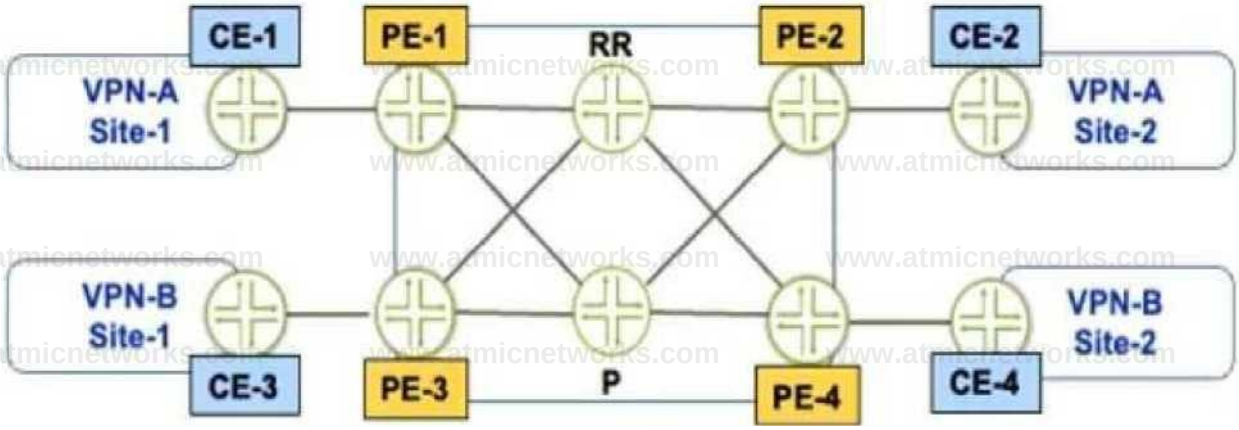

Referring to the exhibit, PE-1 and PE-2 are getting route updates for VPN-B when neither of them service that VPN

Which two actions would optimize this process? (Choose two.)

A. Configure the resolution rib bgp.l3vpn.0 resolution-ribs inet.0 statement on the PEs.

B. Configure the family route-target statement on the RR.

C. Configure the resolution rib bgp.l3vpn.0 resolution-ribs inet.0 statement on the RR.

D. Configure the family route-target statement on the PEs.

Answer: BC

Explanation:

BGP route target filtering can be configured on PE devices or on route reflectors (RRs). Configuring BGP route target filtering on RRs is more efficient and scalable, as it reduces the number of BGP sessions and updates between PE devices. To configure BGP route target filtering on RRs, the following steps are required:

Configure the family route-target statement under the BGP group or neighbor configuration on the RRs. This enables the exchange of the route-target address family between the RRs and their clients (PE devices). Configure the resolution rib bgp.l3vpn.0 resolution-ribs inet.0 statement under the routing-options configuration on the RRs. This enables the RRs to resolve next hops for VPN routes using the inet.0 routing table.

Question: 7

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.

Question: 8

Which statement is correct about IS-IS when it performs the Dijkstra algorithm?

A. The local router moves its own local tuples into the candidate database

B. When a new neighbor ID in the tree database matches a router ID in the LSDB, the neighbor ID is moved to the candidate database

C. Tuples with the lowest cost are moved from the tree database to the LSDB.

D. The algorithm will stop processing once the tree database is empty.

Answer: B

Explanation:

The Dijkstra algorithm in IS-IS operates as follows:

Tree Database Initialization: The local router (root) is added to the tree database with a cost of 0.

Candidate Database Population: Neighbors of the root (from the LSDB) are placed into the candidate database with their associated costs.

Processing Nodes: The node with the lowest cost in the candidate database is moved to the tree database.

Neighbor Evaluation: For each neighbor of the newly added node (from the LSDB), if the neighbor is not already in the tree or candidate database, it is added to the candidate database. If it exists in the candidate with a higher cost, it is updated with the lower cost.

Termination: The algorithm stops when the candidate database is empty, ensuring all shortest paths are

computed.

Analysis of Options:

A . Incorrect. The local router is placed directly into the tree database, not the candidate database.

B . Correct (with context). When a node is added to the tree database, its neighbors (existing in the LSDB) are evaluated. If these neighbors are not already in the tree or candidate database, they are added (not "moved") to the candidate database. The wording "moved" is technically inaccurate, but this option aligns closest with the process of populating the candidate database using LSDB entries during tree database processing.

C . Incorrect. Tuples (nodes) with the lowest cost are moved from the candidate database to the tree database, not from the tree to the LSDB. The LSDB remains static during SPF computation.

D . Incorrect. The algorithm stops when the candidate database is empty, not the tree database. The tree database grows as nodes are processed.

Question: 9

Exhibit

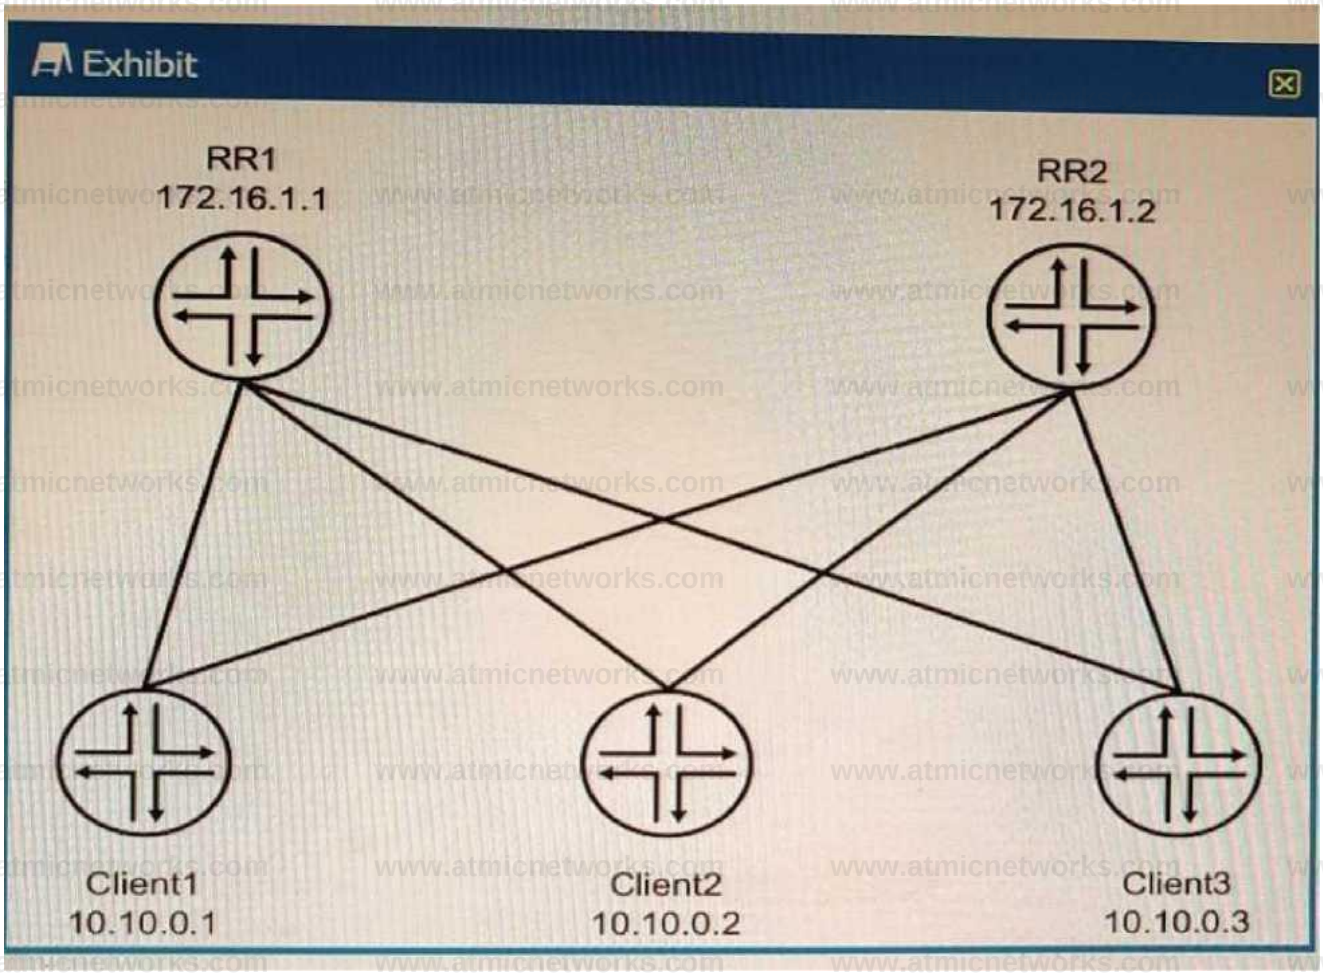

www.atmicnetworks.com

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

A. RR1 is peered to Client2 and RR2

B. RR2 is in an OpenConfirm State until RR1 becomes unreachable.

C. Client1 is peered to Client2 and Client3.

D. Peering is dynamically discovered between all devices.

Answer: A

Explanation:

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

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

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

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

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

Question: 10

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.

Question: 11

Exhibit

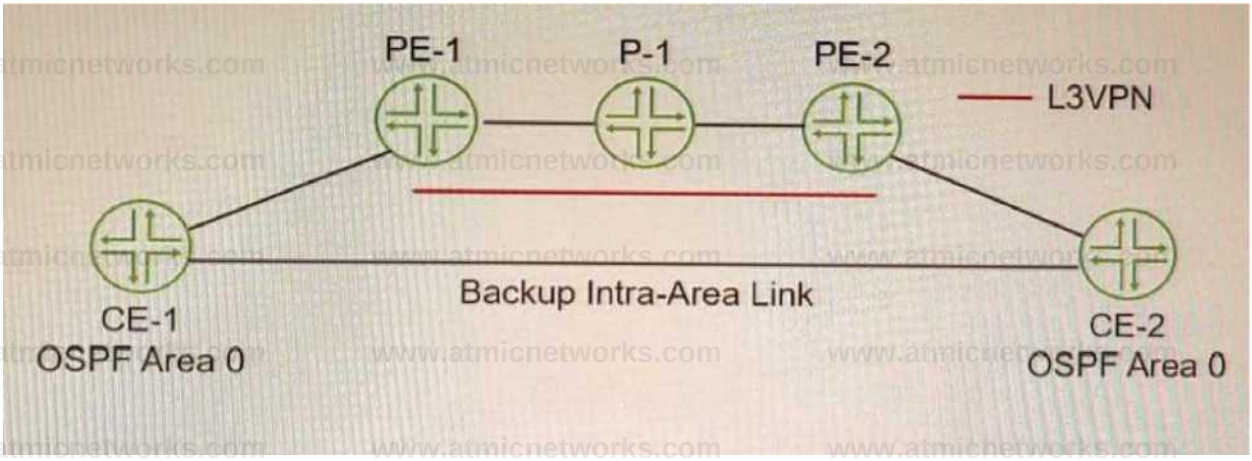

You must ensure that the VPN backbone is preferred over the back door intra-area link as long as the VPN is available. Referring to the exhibit, which action will accomplish this task?

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.

https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/topic-map/configuring- ospfv2-sham-links.html

Question: 12

Exhibit

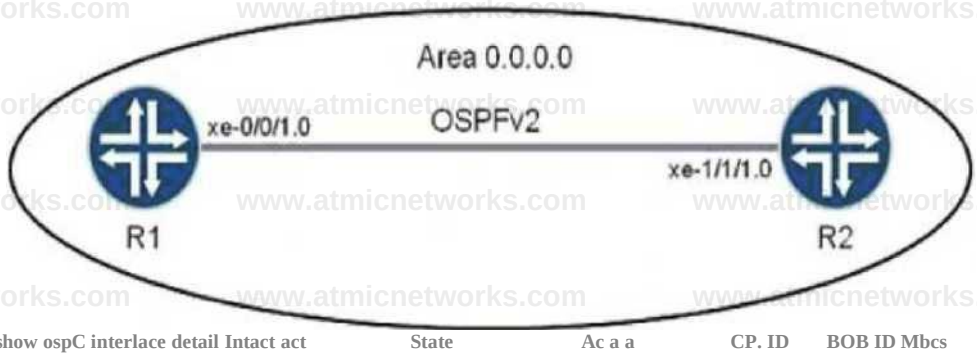


```
user OH > show ospC interlace detail Intact act          State          Ac a a          CP. ID      BOB ID Mbcs
xe-0/0/1.0                BOA 0.0.0.0          190.160.37.12          11.244.215.213 1
Type LAN. address 192.1<1.27.11, Heat 233.233.233.240, WU 4400, Cast 40 on adds 192.101.3 7.12, BOU addc 192.leO.3 7.il, Adj count 1.
Priority 12 8 Mello 10. Dead 40. 8ela.lt 3. Hot Stub te-0/2/1.0                    PtTePt 0.0.0.0      0.0.0.0      0.0.0.0 0
Typo P2P. Address 0.0.0.0. Mask 0.0.0.0. WO 1300, Cost 2604
   Adj count 0
Hello 10, Dead 40, SeXait 5, Hot Stub Auth type: MDS, Active key ID 3, Start time 2013 Jul 19 10:00:00 POT IPaat SA Hama: aa

usorO*2> show ospf inter taco detail
Interlace                 State Area          to ID                  DPS ID Hita
xe-1/1/1.0                BOP 0.0.0.0          192.IM.37.12           11.244.245.2K 1
Type LAN. address 192.KI.27.12, Mask 23S.23S.233.24O, two 4460, Cost 40 DO addc 192.KI.37.12, BM. addc 192.KB.31.il, Adj count 1,
Priority 128 Hello 3. Deed 9. he bait 3, Hot Stub te-2/2/2.0    PtToPt 0.0.0.0      0.0.0.0      0.0.0.0 0
Type P2P. Address 0.0.0.0, Mask 0.0.0.0. MTU 1300, Cost 2604
   Adj count 0
Hello 10. Dead 40, Beirut 3, Hot Stub
Auth type: MDS, Active key ID 3, Start tine 2013 Jul 19 10:00:00 PBT IPsec SA Hane: si
```

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.

Question: 13

Exhibit

user@Rl show configuration interpolated-profile { interpolate {

fill-level [50 75 drop—probability [> }

class-of-service drop-profiles

];

20 60];

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

A. The drop probability jumps immediately from 20% to 60% when the queue level reaches 75% full.

B. The drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full

C. To use this drop profile, you reference it in a scheduler.

D. To use this drop profile, you apply it directly to an interface.

Answer: BC

Explanation:

class-of-service (CoS) is a feature that allows you to prioritize and manage network traffic based on various criteria, such as application type, user group, or packet loss priority. CoS uses different components to classify, mark, queue, schedule, shape, and drop traffic according to the configured policies.

One of the components of CoS is drop profiles, which define how packets are dropped when a queue is congested. Drop profiles use random early detection (RED) algorithm to drop packets randomly before the queue is full, which

helps to avoid global synchronization and improve network performance. Drop profiles can be discrete or interpolated. A discrete drop profile maps a specific fill level of a queue to a specific drop probability. An interpolated drop profile maps a range of fill levels of a queue to a range of drop probabilities and interpolates the values in between.

In the exhibit, we can see that the class-of-service configuration shows an interpolated drop profile with two fill levels (50 and 75) and two drop probabilities (20 and 60). Based on this configuration, we can infer the following statements:

The drop probability jumps immediately from 20% to 60% when the queue level reaches 75% full. This is not correct because the drop profile is interpolated, not discrete. This means that the drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full. The drop probability for any fill level between 50% and 75% can be calculated by using linear interpolation formula.

The drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full. This is correct because the drop profile is interpolated and uses linear interpolation formula to calculate the drop probability for any fill level between 50% and 75%. For example, if the fill level is 60%, the drop probability is 28%, which is calculated by using the formula: (60 - 50) / (75 - 50) * (60 - 20) + 20 = 28.

To use this drop profile, you reference it in a scheduler. This is correct because a scheduler is a component of CoS that determines how packets are dequeued from different queues and transmitted on an interface. A scheduler can reference a drop profile by using the random-detect statement under the [edit class-of-service schedulers] hierarchy level. For example: scheduler test { transmit-rate percent 10; buffer-size percent 10; random-detect test-profile; }

To use this drop profile, you apply it directly to an interface. This is not correct because a drop profile cannot be applied directly to an interface. A drop profile can only be referenced by a scheduler, which can be applied to an interface by using the scheduler-map statement under the [edit class-of- service interfaces] hierarchy level. For example: interfaces ge-0/0/0 { unit 0 { scheduler-map testmap; } }

Question: 14

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

A. If a point-to-point interface is in both L1 and L2, separate hello messages are sent for each level.

B. If a point-to-point interface is in both 11 and L2, one combined hello message is sent for both levels.

C. If a broadcast interface is in both L1 and L2, separate hello messages are sent for each level.

D. If a broadcast interface is in both L1 and L2, one combined hello message is sent for both levels.

Answer: BC

Explanation:

Intermediate System to Intermediate System (IS-IS) is a link-state routing protocol that supports Level 1 (L1), Level

2 (L2), or both (L1/L2) operations. The way IS-IS sends Hello (IIH) packets depends on whether the interface is point-to-point (P2P) or broadcast (LAN).

Evaluating the Answer Choices

Q Option A: "If a point-to-point interface is in both L1 and L2, separate hello messages are sent for each level."

Incorrect!

On point-to-point (P2P) interfaces, only one combined Hello message is sent for both L1 and L2.

IS-IS P2P Hellos include both Level 1 and Level 2 TLVs in the same message.

Reference: Juniper IS-IS documentation confirms that P2P links use a single Hello message with both levels included.

X This statement is incorrect.

Q Option B: "If a point-to-point interface is in both L1 and L2, one combined hello message is sent for both levels."

Correct!

On point-to-point (P2P) links, IS-IS sends a single Hello message that includes TLVs for both L1 and L2.

This reduces overhead and simplifies adjacency formation.

Q This statement is correct.

Q Option C: "If a broadcast interface is in both L1 and L2, separate hello messages are sent for each level."

Correct!

On broadcast (LAN) interfaces, IS-IS sends separate Hello messages for L1 and L2.

This is because L1 and L2 use separate Designated IS (DIS) elections and different multicast addresses:

L1 Hellos: Sent to AllL1IS (01:80:C2:00:00:14)

L2 Hellos: Sent to AllL2IS (01:80:C2:00:00:15)

Reference: Juniper IS-IS Configuration Guide confirms that broadcast interfaces send separate L1 and L2 Hello messages.

Q This statement is correct.

Q Option D: "If a broadcast interface is in both L1 and L2, one combined hello message is sent for both levels."

Incorrect!

As stated above, IS-IS sends separate Hello messages for L1 and L2 on broadcast interfaces because they have independent DIS elections.

X This statement is incorrect.

Final Answer:

Q B. If a point-to-point interface is in both L1 and L2, one combined hello message is sent for both levels.

Q C. If a broadcast interface is in both L1 and L2, separate hello messages are sent for each level.

Verification from Juniper Documentation

Juniper IS-IS Configuration Guide confirms:

Point-to-Point (P2P) interfaces send one combined Hello for both levels.

Broadcast interfaces send separate L1 and L2 Hellos due to separate DIS elections.

RFC 1195 (IS-IS Extensions for IPv4) specifies that broadcast networks require distinct Hellos per level.

Question: 15

Exhibit

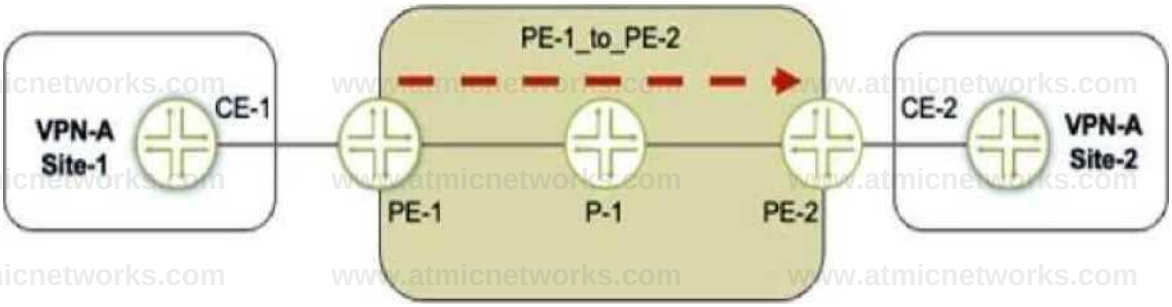

Referring to the exhibit, a working L3VPN exists that connects VPN-A sites CoS is configured correctly to match on the MPLS EXP bits of the LSP, but when traffic is sent from Site-1 to Site-2, PE-2 is not classifying the traffic correctly

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: B

Explanation:

Understanding the Problem in MPLS CoS Classification

How EXP Bits Are Used for CoS in MPLS

Traffic is sent from VPN-A Site-1 → CE-1 → PE-1 → P-1 → PE-2 → CE-2.

The MPLS LSP (Label Switched Path) from PE-1 to PE-2 is expected to carry MPLS EXP bits, which are used for Class of Service (CoS) classification.

PE-2 should classify traffic based on EXP bits received in the MPLS label.

What Happens with PHP (Penultimate Hop Popping)?

By default, the penultimate router (P-1) pops the top MPLS label before sending the packet to PE-2.

Since the EXP bits are in the top MPLS label, they get removed along with the label.

This means that PE-2 no longer sees the correct EXP bits, leading to incorrect traffic classification.

Solution: Configure Explicit-Null on PE-2

Explicit Null (explicit-null) must be configured on PE-2 to ensure that P-1 does NOT remove the MPLS label.

Instead of removing the label, P-1 will send a label of 0 (for IPv4) or 2 (for IPv6) to PE-2.

This preserves the MPLS EXP bits, allowing PE-2 to classify the traffic correctly.

Evaluating the Answer Choices Again

Q B. Configure the explicit-null statement on PE-2.

Correct, because:

PE-2 is the egress LSR, where Ultimate Hop Popping (UHP) must be enabled.

Configuring explicit-null ensures that P-1 does not remove the label, preserving the EXP bits for CoS classification at PE-2.

Configuration on PE-2:

www.atmicnetworks.com

set protocols mpls explicit-null

Juniper Documentation Reference:

"Explicit-null must be configured on the egress LSR to prevent PHP from removing the top MPLS label, thereby preserving the EXP bits."

X A. Configure the explicit-null statement on PE-1.

Incorrect, because:

Explicit-null must be configured on the egress LSR (PE-2), not the ingress LSR (PE-1).

PE-1 only labels the traffic but does not control PHP behavior on P-1.

X C. Configure VPN prefix mapping for the PE-1_to_PE-2 LSP.

Incorrect, because:

VPN prefix mapping is used for mapping VPN routes to LSPs but does not solve the EXP bit issue.

The problem here is label removal (PHP), not route mapping.

X D. Set a static CoS value for the PE-1_to-PE-2 LSP.

Incorrect, because:

This does not preserve the original EXP bits, it only applies a static CoS value.

It's a workaround, not a fix.

Final Answer: Q B. Configure the explicit-null statement on PE-2.

Explanation:

Key Takeaways

Penultimate Hop Popping (PHP) removes the outer MPLS label at P-1, which also removes the EXP bits used for CoS classification.

To keep EXP bits intact, configure explicit-null on the egress PE (PE-2).

This forces P-1 to send a label (0 for IPv4, 2 for IPv6) to PE-2, preserving the EXP bits for CoS classification.

Official Juniper Documentation Reference

Juniper MPLS CoS and PHP Behavior Guide

"To retain CoS EXP bits at the egress LSR, configure explicit-null on the egress PE. This prevents PHP from stripping the MPLS label before reaching the final PE router."

Question: 16

Exhibit

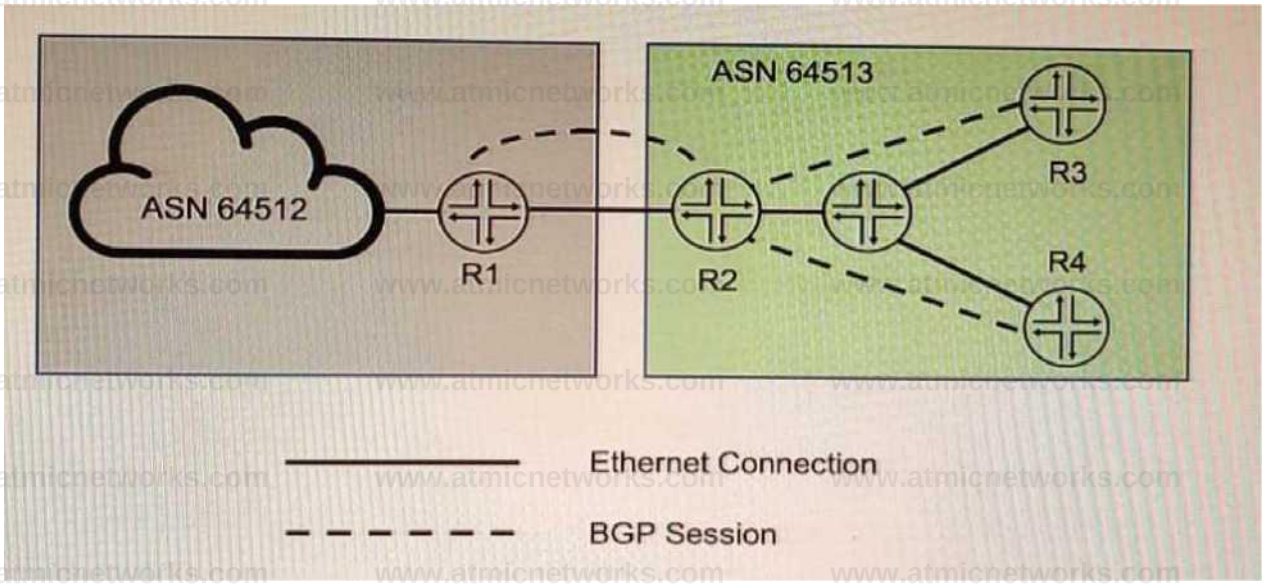

You want to implement the BGP Generalized TTL Security Mechanism (GTSM) on the network

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: B, C, E

Explanation:

https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/ref/statement/multihop-edit-protocols-bgp.html

Question: 17

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:

https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/topic-map/configuring- ospfv2-sham-links.html

Question: 18

Exhibit

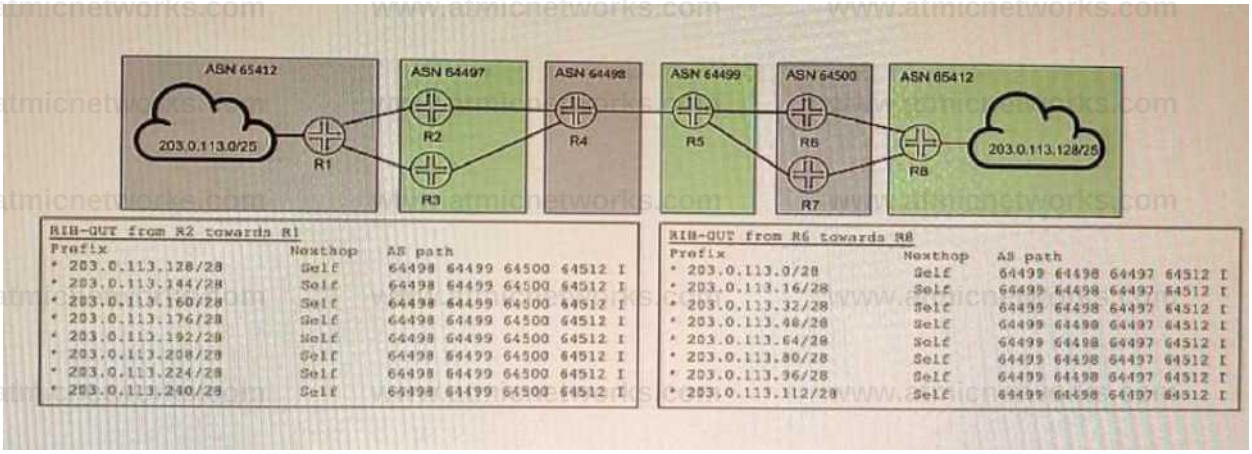

R1 and R8 are not receiving each other's routes

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

A. Configure remove-private on advertisements from AS 64497 toward AS 64498.

B. Configure as-override on advertisement from AS 64500 toward AS 64512.

C. Configure remove-private on advertisements from AS 64500 toward AS 64499.

D. Configure loops and advertise-peer-as on routers in AS 64497 and AS 64450.

E. Configure loops on routers in AS 65412 and advertise-peer-as on routers in AS 64498.

Answer: ABC

Explanation:

www.atmicnetworks.com

Question: 19

Which origin code is preferred by BGP?

A. Internal

B. External

C. Incomplete

D. Null

Answer: A

Explanation:

Prefer the route with the lower origin code. Routes learned from an IGP have a lower origin code than those learned from an exterior gateway protocol (EGP), and both have lower origin codes than incomplete routes (routes whose origin is unknown).

https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/bgp/topics/concept/routing- protocols-address-representation.html

Question: 20

An interface is configured with a behavior aggregate classifier and a multifield classifier How will the packet be processed when received on this interface?

A. The packet will be discarded.

B. The packet will be processed by the BA classifier first, then the MF classifier.

C. The packet will be forwarded with no classification changes.

D. The packet will be processed by the MF classifier first, then the BA classifier.

Answer: B

Explanation:

When a Juniper device receives a packet on an interface with both a Behavior Aggregate (BA) classifier and a Multifield (MF) classifier, Junos OS follows a specific processing order to apply Class of Service (CoS).

Understanding the Classifiers in Junos CoS

1 Qehavior Aggregate (BA) Classifier

Uses packet headers (DSCP, IP precedence, or MPLS EXP bits) to classify traffic into forwarding classes.

Applied at the ingress interface.

www.atmicnetworks.com

Example: A packet with DSCP 46 (Expedited Forwarding) is mapped to a high-priority queue.

2 QMultifield (MF) Classifier

Uses match conditions (like source/destination IP, port numbers, protocol types) to classify traffic.

Typically used for more granular classification beyond what BA can provide.

Junos Processing Order:

When both BA and MF classifiers are configured on an interface, Junos first applies the BA classifier, then the MF classifier.

MF classifier can override the BA classification if necessary.

Evaluating the Answer Choices

8. The packet will be processed by the BA classifier first, then the MF classifier.

Correct, because Junos first applies BA classification based on DSCP/MPLS EXP bits.

After BA classification, the MF classifier is applied, which can refine or override the BA classification.

X A. The packet will be discarded.

Incorrect, because classification does not drop packets unless explicitly configured with a filter or policing action.

X C. The packet will be forwarded with no classification changes.

Incorrect, because both classifiers are applied in a specific order, meaning classification changes will occur.

X D. The packet will be processed by the MF classifier first, then the BA classifier.

Incorrect, because BA classification is always applied first, followed by MF classification.

Final Answer: Q B. The packet will be processed by the BA classifier first, then the MF classifier.

Explanation:

Official Juniper Reference:
"When both BA and MF classifiers are applied on an interface, Junos OS first classifies packets using the BA classifier before applying the MF classifier."

www.atmicnetworks.com

Question: 21

Exhibit

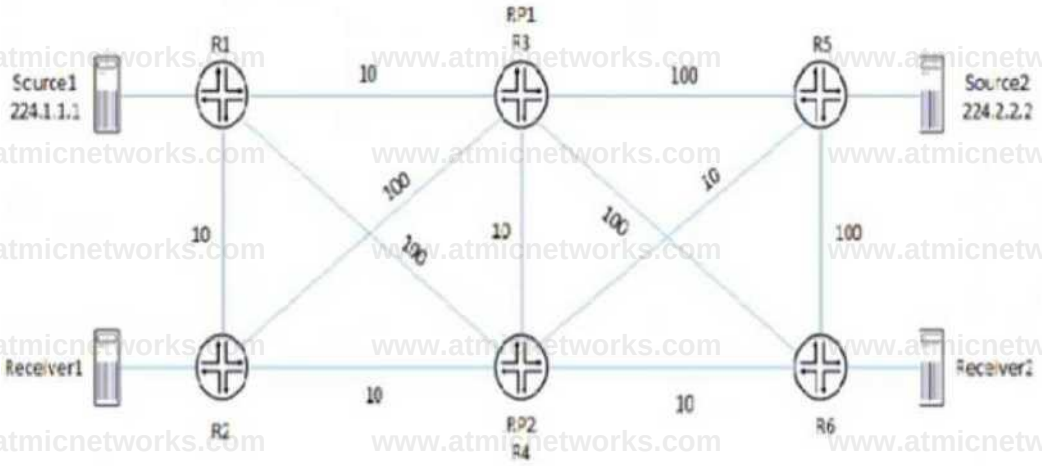

Referring to the exhibit, PIM-SM is configured on all routers, and Anycast-RP with Anycast-PIM is used for the discovery mechanism on RP1 and RP2. The interface metric values are shown for the OSPF area.

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

A. Source2 will use RP2 and Receiver2 will use RP2 forgroup 224.2.2.2.

B. Source2 will use RP1 and Receiver2 will use RP1 forgroup 224.2.2.2.

C. Source1 will use RP1 and Receiver1 will use RP1 forgroup 224.1.1.1.

D. Source1 will use RP1 and Receiver1 will use RP2 for group 224.1 1 1

Answer: AC

Explanation:

A sham link is a logical link between two PE routers that belong to the same OSPF area but are connected through an L3VPN. A sham link makes the PE routers appear as if they are directly connected, and prevents OSPF from preferring an intra-area back door link over the VPN backbone. A sham link creates an OSPF multihop 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 selection1.

www.atmicnetworks.com

Question: 22

Exhibit

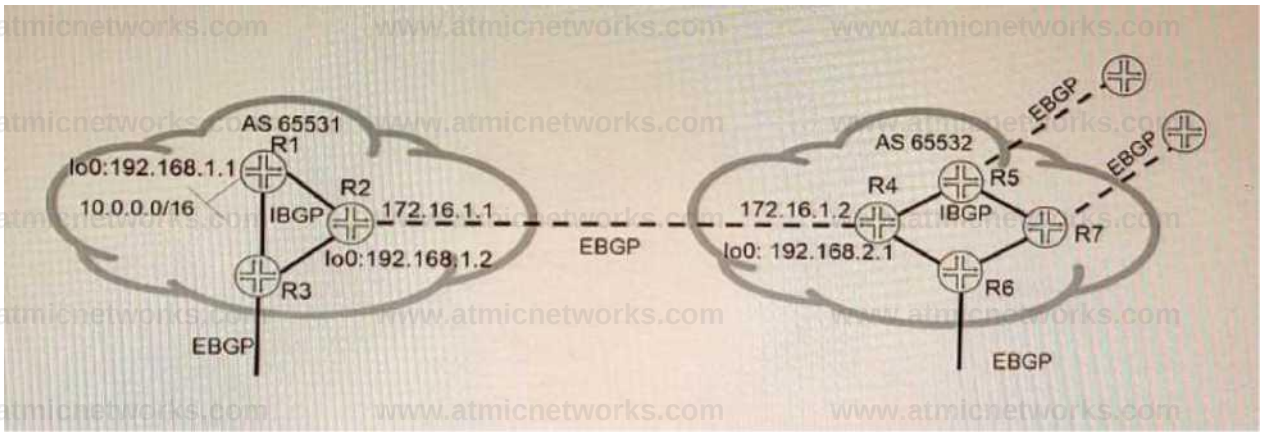

Referring to the exhibit, which three statements are correct about route 10 0 0.0/16 when using the default BGP advertisement rules'? (Choose three.)

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

B. R2 will advertise 10.0.0.0/16 to R3 with 192.168.1.1 as the next hop.

C. R1 will advertise 10.0.0.0/16 to R2 with 192.168.1.1 as the next hop.

D. R1 will prepend AS 65531 when advertising 10.0.0.0/16 to R2.

E. R4 will advertise 10.0.0.0/16 to R6 with 172.16.1.1 as the next hop.

Answer: A, C, E

Explanation:

Question: 23

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.

Question: 24

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: C

Explanation:

Question: 25

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

A. CSNPs are flooded periodically.

B. PSNPs are flooded periodically.

C. PSNPs contain only descriptions of LSPs.

D. CSNPs contain only descriptions of LSPs.

Answer: A, C

Explanation:

Option A (Correct):

Complete Sequence Number PDUs (CSNPs) are periodically flooded by the Designated Intermediate System (DIS) on multi-access networks (e.g., Ethernet).

This ensures all routers on the segment synchronize their Link-State Databases (LSDBs).

Reference: Juniper IS-IS CSNP Overview.

Option C (Correct):

Partial Sequence Number PDUs (PSNPs) contain only the headers (descriptions) of LSPs (e.g., LSP ID, sequence number, checksum).

PSNPs are used to:

Request missing LSPs (when a router detects discrepancies via CSNPs).

Acknowledge LSP receipt (in point-to-point networks).

They do not include the full LSP data.

Reference: Juniper IS-IS PSNP Overview.

Why Other Options Are Incorrect:

Option B: Incorrect. PSNPs are not flooded periodically—they are sent on-demand for specific LSP synchronization.

Option D: Incorrect. While CSNPs do contain LSP descriptions (headers), the term "only" is misleading. CSNPs summarize all LSPs in the LSDB, but they are not limited to "only" descriptions— they serve as a complete database overview.

Key Takeaways:

CSNPs are periodic, broadcast by the DIS, and ensure LSDB consistency.

PSNPs are triggered, contain specific LSP headers, and handle requests/acknowledgments.

IS-IS uses CSNPs and PSNPs to maintain efficient LSDB synchronization without flooding full LSPs unnecessarily.

For further details, refer to Juniper's official IS-IS documentation:

Juniper IS-IS Configuration Guide.

Question: 26

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: AC

Explanation:

https://www.juniper.net/documentation/us/en/software/nce/feature-guide-virtual-private-lan-service/topics/task/vpls-ldp-signaling-solutions.html

https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/topics/concept/vpns-configuring-vpls-routing-instances.html#id-11510150 id-11568648

Question: 27

Exhibit


```
u»rt(£ H Sho* protocols Me grove
tWMo-n-l ( type rittrwl;
      local wo rest It.It.0.2;
      peer-« 45550;
      local at 44511;
      neighbor 10.10.0.1 < export
            static-tii-bgp;
```

```
uieroK-ie shew protocols bgp group fOO to-lE-l type
external;
local-aooress 10.14.0.1;
peer-o* 44511;
locales 45550;
neighbor 10.10.0.2 l
      tally met < unicast ( prefix-lunll ( naxlaua J;
      teardown;
                  )
```

```
iA»rpct-l» snow pol Ky-otUces policy
statement static to top (
      tor* export-static < tree {
            protocol static:

                  route-tilt*'      192.100.1.  0/2* exact;
            route-filter    route-  192.100.2.  0/24 e*act;
            tilt*'      route-filter 192.100.3.  0/24 exact;
            route-filter    route-  192.140.4.  0/24 etact;
            filter      route-filter 192.140.9.  0/24 etact;
            route-filter    route-  192.140.4.  0/24 exact;
            fitter  route-tliter  J 192.140.7.  0/24 exact;
            ther accept;            192.140.4.  0/24 etact;
                                    192.140.9.  0/24 react;
      then reject;                  192.140.10.  0/24 exact;
```

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: C

Explanation:

Analyzing the Exhibit and Understanding the Issue

The exhibit shows BGP configurations on CE-1 and PE-1, which are connected via EBGP.

CE-1 (Customer Edge)

Uses AS 64511 and establishes an EBGP session with PE-1 (AS 65550).

Configured to export 10 static routes (192.168.1.0/24 - 192.168.10.0/24) using the static-to-bgp policy.

www.atmicnetworks.com

PE-1 (Provider Edge)

Uses AS 65550 and is peering with CE-1 (AS 64511).

Configured with a prefix-limit of 5 on received routes from CE-1.

Teardown enabled, meaning if more than 5 prefixes are received, the BGP session is shut down.

Identifying the Problem

CE-1 is correctly configured with peer AS 65550, so Option B ("CE-1 is configured with an incorrect peer AS") is incorrect X.

CE-1 is advertising exactly 10 static routes (as per policy).

PE-1 has a prefix-limit maximum 5 with teardown enabled.

This means that when CE-1 advertises more than 5 prefixes, PE-1 shuts down the BGP session.

BGP moves to the "Active" state, indicating that the session has been disrupted and PE-1 is trying to re-establish the connection.

CE-1 is reachable since the session was initially established before the limit was exceeded, so Option D ("CE-1 is unreachable") is incorrect X.

CE-1 is not advertising its entire routing table, only the static prefixes listed in the policy, so Option A ("CE-1 is advertising its entire routing table") is incorrect X.

Correct Answer

Q C. The prefix limit has been reached on PE-1

Verification from Juniper Documentation

Juniper BGP Prefix Limit Documentation confirms that exceeding the prefix limit with teardown causes the BGP session to go into "Active" state.

Juniper Troubleshooting Guide for BGP Peering Issues states that when a BGP session reaches the prefix limit and has teardown enabled, the session is terminated.

Question: 28

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 ospf3 realm ipv4-unicast overload timeout 900

B. set protocols ospf overload

C. set protocols ospf overload timeout 900

D. set protocols ospf3 overload

E. set protocols ospf3 overload timeout 900

Answer: ACE

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.

Question: 29

A packet is received on an interface configured with transmission scheduling. One of the configured queues In this scenario, which two actions will be taken by default on a Junos device? (Choose two.)

A. The excess traffic will be discarded

B. The exceeding queue will be considered to have negative bandwidth credit.

C. The excess traffic will use bandwidth available from other queues

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

Answer: BC

Explanation:

https://www.juniper.net/documentation/us/en/software/junos/cos-security- devices/topics/concept/cos-transmission-scheduling-security-overview.html

When a Junos device receives a packet on an interface with transmission scheduling, traffic is placed into different queues based on Class of Service (CoS) policies. If a queue exceeds its allocated bandwidth, Junos has default behaviors for handling excess traffic.

Key Junos Behaviors for Transmission Scheduling

Queues Can Borrow Bandwidth from Other Queues Q

If a queue has excess traffic, it can use bandwidth from underutilized queues, as long as bandwidth is available.

Reference from Juniper Documentation:

"By default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues."

Queues Have Credit-Based Tracking Q

A queue that stays within its allocated bandwidth is considered to have positive bandwidth credit.

A queue that exceeds its allocation is considered to have negative bandwidth credit.

Reference from Juniper Documentation:

"A queue receiving traffic in excess of its bandwidth allocation is considered to have negative bandwidth credit."

Evaluating the Answer Choices

Q B. The exceeding queue will be considered to have negative bandwidth credit.

Correct, because when a queue exceeds its allocated bandwidth, Junos assigns it negative bandwidth credit.

This means the queue is in debt and must recover before it can transmit additional packets.

Q C. The excess traffic will use bandwidth available from other queues.

Correct, because Junos allows excess traffic to borrow bandwidth from underutilized queues by default.

If a forwarding class does not use its allocated bandwidth, other queues can borrow the unused bandwidth.

Why the Other Answers Are Incorrect?

X A. The excess traffic will be discarded.

Incorrect, because Junos does not immediately discard excess traffic unless the queue cannot borrow bandwidth.

By default, Junos allows bandwidth sharing, and only if no bandwidth is available does it drop packets.

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

Incorrect, because when a queue exceeds its assigned bandwidth, it gets negative bandwidth credit, not positive credit.

Verified Juniper Official Reference

Junos CoS Transmission Scheduling Overview

"By default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues."

"A queue receiving traffic in excess of its bandwidth allocation is considered to have negative bandwidth credit."

Question: 30

In IS-IS, which two statements are correct about the designated intermediate system (DIS) on a multi-access network segment? (Choose two)

A. A router with a priority of 10 wins the DIS election over a router with a priority of 1.

B. A router with a priority of 1 wins the DIS election over a router with a priority of 10.

C. On the multi-access network, each router forms an adjacency to every other router on the segment

D. On the multi-access network, each router only forms an adjacency to the DIS.

Answer: AC

Explanation:

Option A (Correct):

In IS-IS, the Designated Intermediate System (DIS) is elected based on the highest configured priority (as defined in Junos OS).

If priorities are equal, the router with the highest MAC address becomes the DIS.

A priority value of 10 will always override a lower priority (e.g., 1).

Reference: Juniper IS-IS DIS Election.

Option C (Correct):

On a multi-access network (e.g., Ethernet), all IS-IS routers form adjacencies with every other router on the segment.

Unlike OSPF, IS-IS does not restrict adjacencies to only the DIS.

The DIS is responsible for creating a pseudonode LSP to represent the broadcast network, but full mesh adjacencies are maintained.

Reference: Juniper IS-IS Adjacency Formation.

Why Other Options Are Incorrect:

Option B: Incorrect. Higher priority always wins the DIS election. A priority of 1 cannot override a priority of 10.

Option D: Incorrect. IS-IS routers form adjacencies with all neighbors, not just the DIS.

Key Takeaways:

DIS Election: Prioritizes highest numerical value (e.g., 10 > 1).

Adjacency Behavior: Full mesh adjacencies are maintained, unlike OSPF.

DIS Role: Primarily for generating pseudonode LSPs and optimizing flooding, not adjacency restriction.

For further details, refer to Juniper's official IS-IS documentation:
Juniper IS-IS Configuration Guide.

https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/routing-protocol-is-is-security-designated-router-understanding.html

www.atmicnetworks.com

Question: 31

Exhibit

www.atmicnetworks.com

```
 user?R4> show pim rps Instance: PIM.master address-family INET Rp address Type   Mode Holdtime
Timeout Groups Group prefixes
10.1.255.2    bootstrap   sparse      150     118     0 224.1.1.0/24
10.1.255.3    bootstrap   sparse      150     118     2 224.1.1.0/28
 u*er0R4> show route 10.1.255.2 inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0
hidden) * ■ Active Route, - ■ Last Active, * ■ Both 10.1.255.2/32   '[IS-IS/18] 00:32:27, metric
10
                    > to 10.1.1.2 via ge-0/0/0.0 inet.2: 8 destinations, 8 routes (8 active, 0
holddown, 0 hidden) * = Active Route, - = Last Active, * = Both 0.0.0.0/0  '[Static/5] 00:13:55
                    > to 10.1.1.6 via ge-0/0/1.0 user?R4> show route 10.1.255.3 inet.0: 16
destinations, 16 routes (16 active, 0 holddown, 0 hidden) * = Active Route, - = Last Active, * =
Both 10.1.255.3/32    '[IS-IS/18] 00:32:43, metric 10
                    > to 10.1.1.6 via ge-0/0/1.0 inet.2: 8 destinations, 8 routes (8 active, 0
holddown, 0 hidden) + = Active Route, - = Last Active, ' = Both 0.0.0.0/0  '[static/5] 00:14:25
                    > to 10.1.1.6 via ge-0/0/1.0 [edit] user0R2# show protocols pim rp (
bootstrap ( family inet [ priority 200; ] local ( address 10.1.255.2; group-ranges (
224.1.1.0/24; ) } I interface all; [edit] user@R3# show protocols pim rp ( bootstrap ( family
inet [ priority 210; ) } local ( address 10.1.255.3; group-ranges [ 224.1.1.0/28; } } interface
all;
```

R4 is directly connected to both RPs (R2 and R3) R4 is currently sending all ,o,ns upstream to R3 but you want all joins to go to R2 instead Referring to the exhibit, which configuration change will solve this issue?

A. Change the bootstrap priority on R2 to be higher than R3

B. Change the default route in inet.2 on R4 from R3 as the next hop to R2

C. Change the local address on R2 to be higher than R3.

D. Change the group-range to be more specific on R2 than R3.

Answer: D

Explanation:

The issue arises because R3's group-range (224.1.1.0/28) is more specific than R2's group-range (224.1.1.0/24). In PIM bootstrap (BSR), the RP with the longest prefix (most specific group-range) is preferred, regardless of priority. Even though R3 has a higher bootstrap priority (210 vs. R2's 200), its more specific /28 group-range takes precedence for groups within 224.1.1.0/28.

Why Option D is Correct:

To force R4 to use R2 for all joins, R2's group-range must be more specific than R3's. For example:

If R2's group-range is changed to 224.1.1.0/28 (same as R3) but with a higher priority, R2 would win (priority is compared only when group-ranges are equal).

If R2's group-range is changed to 224.1.1.0/29 (more specific than /28), it will override R3's /28 for groups in the /29 range.

The key is prefix specificity, which overrides priority in BSR elections.

www.atmicnetworks.com

Why Other Options Are Incorrect:

A . Change bootstrap priority on R2 to be higher than R3:

Priority is evaluated only when group-ranges are identical. Since R3's group-range (/28) is more specific than R2's (/24), R3 will still win for groups in 224.1.1.0/28, even if R2's priority is higher.

B . Change the default route in inet.2 on R4:

RPF routes (inet.2) determine how traffic reaches the RP, but they do not influence RP election logic (BSR priority/group-range).

C . Change R2's local address to be higher than R3's:

The RP address is a tiebreaker only if priorities and group-ranges are equal. Since R3's group-range is more specific, this change has no impact.

Key Takeaways:

BSR RP Election Order:

Longest group prefix (most specific).

Highest priority (if prefixes are equal).

Highest RP address (if prefixes and priorities are equal).

To override R3, R2 must advertise a more specific group-range (e.g., /28 or smaller) to ensure it is selected for the desired multicast groups.

Reference:

Juniper PIM Sparse Mode and BSR Configuration.

Question: 32

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:

www.atmicnetworks.com

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.

Reference: 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

Question: 33

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 Gigibit Ethernet interface.
D. Enable per-prefix load balancing.

Answer: A

Explanation:
https://www.juniper.net/documentation/us/en/software/junos/sampling-forwarding-monitoring/bgp/topics/concept/bgp-multipath-unequal-understanding.html

Question: 34

Exhibit

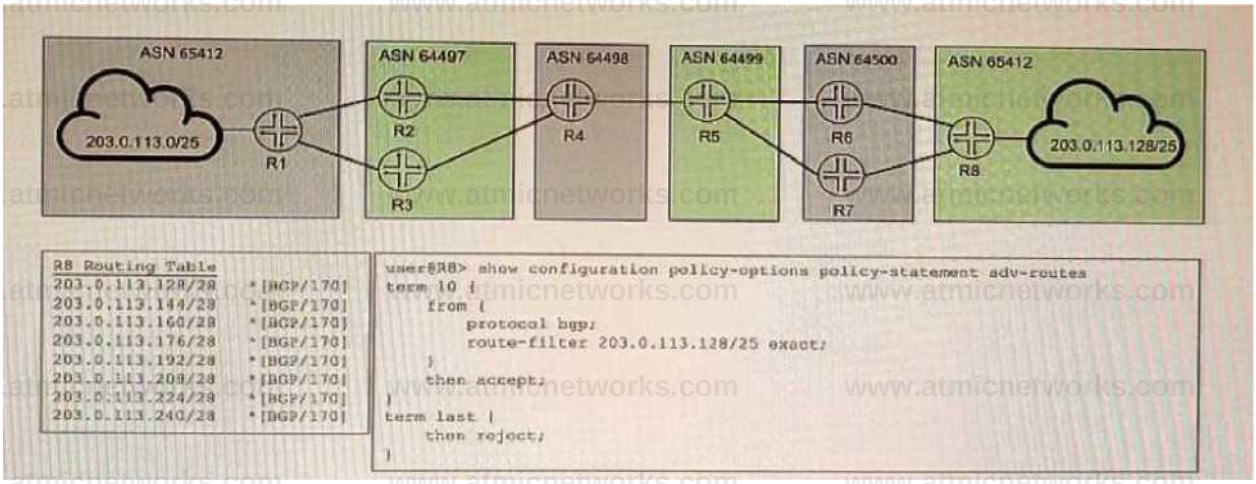

You are attempting to summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500. You implement the export policy shown in the exhibit and all routes from the routing table stop being advertised.

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 tables. -

C. Add the set routing-options static route 203.0.113.123/25 discard command.

D. Replace exact in the export policy with orlonger.

Answer: AC

Explanation:

Question: 35

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

A. Bootstrap messages are forwarded only to routers that explicitly requested the messages within the PIM sparse-mode domain

B. Bootstrap messages distribute RP information dynamically during an RP election.

C. Bootstrap messages are used to notify which router is the PIM RP

D. Bootstrap messages are forwarded to all routers within a PIM sparse-mode domain.

Answer: BD

Explanation:

Bootstrap messages are PIM messages that are used to distribute rendezvous point (RP) information dynamically

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

Question: 36

Exhibit

```
userSPEll shew routing-instances VFN-A ( instance-type vrf;
    interface ge-O/O/l.G;
    vrf-target target:€4512:1234;
    protocols { bgp { group CE { type external; family inet { unicast;

            neighbor 10.0.0.1 peer-as €4512; as-override;
```

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

A. This VPN connects customer sites that use different AS numbers.

B. This VPN connects customer sites that use the same AS number

C. A Layer 2 VPN is configured.

D. A Layer 3 VPN is configured.

Answer: BD

Explanation:

The provided configuration is for a routing instance named VPN-A on a Juniper PE (Provider Edge) router. Let's break it down:

Instance Type: VRF

The instance-type vrf; statement indicates that this is a Layer 3 VPN (L3VPN) using MPLS VPNs (RFC 4364 – BGP/MPLS IP VPNs).

This confirms that Option D (A Layer 3 VPN is configured) is correct Q.

VRF Target and Interface Association

The vrf-target target:64512:1234; defines the route target (RT) for importing and exporting VPN routes.

The interface ge-0/0/1.0; binds this interface to the VRF.

BGP Configuration for CE (Customer Edge) Peering

The group CE section configures external BGP (EBGP) (type external;).

The neighbor 10.0.0.1 is in AS 64512 (peer-as 64512;).

The as-override; statement is used.

Evaluating the Answer Choices

Q Option B: "This VPN connects customer sites that use the same AS number."

The as-override; command allows multiple customer sites that use the same AS number (64512) to communicate over the service provider's MPLS network.

Normally, BGP prevents routes with the same AS in the AS_PATH from being accepted. The as- override feature replaces the customer's AS number with the provider's AS, ensuring proper route advertisement.

Q This statement is correct.

X Option A: "This VPN connects customer sites that use different AS numbers."

If the customer sites had different AS numbers, there would be no need for as-override.

The as-override feature is specifically used when all customer sites share the same AS number, ensuring that BGP routes are accepted.

X This statement is incorrect.

X Option C: "A Layer 2 VPN is configured."

A Layer 2 VPN (L2VPN) configuration would typically use instance-type l2vpn; or EVPN/VPLS-related

parameters (e.g., protocols l2vpn or protocols vpls).

Since this configuration uses instance-type vrf; and BGP with a VRF target, it is clearly a Layer 3 VPN (L3VPN).

X This statement is incorrect.

Q Option D: "A Layer 3 VPN is configured."

The instance-type vrf; confirms this is an MPLS Layer 3 VPN (L3VPN).

VRFs, BGP, and route targets (vrf-target) are specific to Layer 3 VPNs.

Q This statement is correct.

Final Answer:

Q B. This VPN connects customer sites that use the same AS number.

Q D. A Layer 3 VPN is configured.

Verification from Juniper Documentation:

Juniper BGP/MPLS Layer 3 VPNs Guide confirms that instance-type vrf is used for L3VPNs.

Juniper BGP Configuration Guide states that as-override is applied when customer sites use the same AS number.

RFC 4364 (BGP/MPLS IP VPNs) explains how route targets and VRFs are used in L3VPN deployments.

Question: 37

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-distinguiaher value on both PE devices.

Answer: A

Explanation:

Question: 38

Exhibit

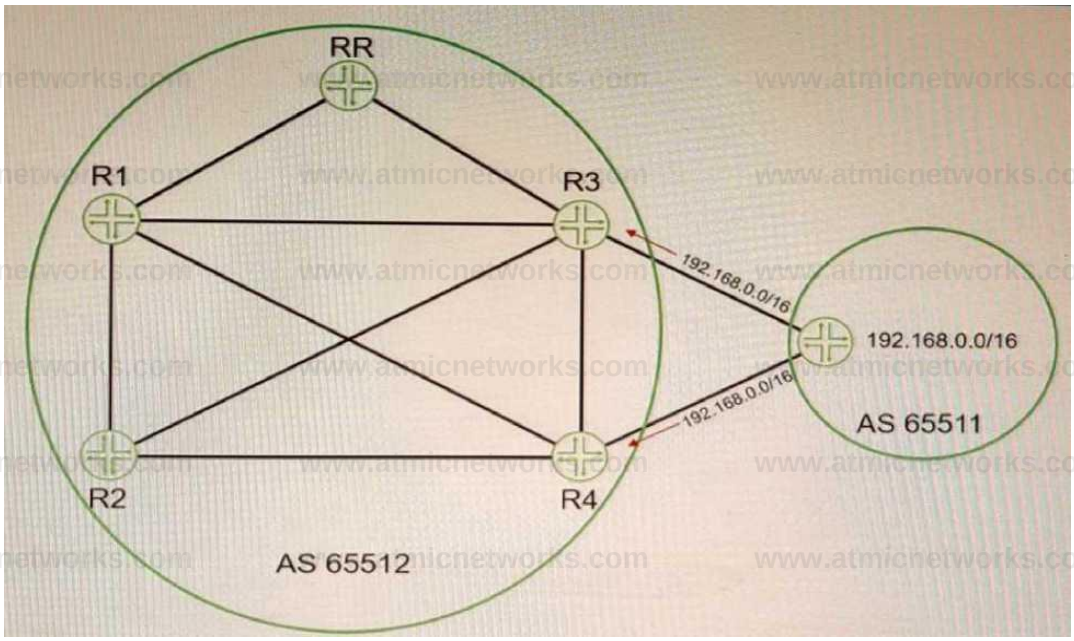

Referring to the exhibit, you are receiving the 192.168 0 0/16 route on both R3 and R4 from your EBGP neighbor You must ensure that R1 and R2 receive both BGP routes from the route reflector

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

A. add-path

B. multihop

C. multipath

D. route-target

Answer: A

Explanation:

BGP add-path is a feature that allows the advertisement of multiple paths through the same peering session for the same prefix without the new paths implicitly replacing any previous paths. This behavior promotes path diversity and reduces multi-exit discriminator (MED) oscillations. BGP addpath is implemented by adding a path identifier to each path in the NLRI. The path identifier can be

considered as something similar to a route distinguisher in VPNs, except that a path ID can apply to any address family. Path IDs are unique to a peering session and are generated for each network3. In this question, we have a route reflector (RR) that receives two routes for the same prefix (192.168.0.0/16) from an EBGP neighbor. By default, the RR will only advertise its best path to its clients (R1 and R2). However, we want R1 and R2 to receive both routes from the RR. To achieve this, we need to configure BGP add-path on the RR and enable it to send multiple paths for the same prefix to its clients.

Reference: 3: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe- 16/irg-xe-16-book/bgp-additional-paths.html

Question: 39

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:

Question: 40

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.

www.atmicnetworks.com

Question: 41

Exhibit

```
user@router> show route advertising-protocol bgp 10.0,0.43 extensive 10.0.0.189
rnet.0: 23 destinations, 41 routes (23 active, 0 holddown, 0 hidden)
* 10.0.0.180/32 (2 entries, 1 announced)
 BSP group underlay type External AS path: [65199] 65170 €5180 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: A

Explanation:

https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/command/show-route-advertising-protocol.html

Question: 42

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

A. OSPFv3 only supports IPv4.

B. OSPFv3 supports both IPv6 and IPv4, but not in the same routing instance.

C. OSPFv3 is not backward compatible with IPv4

D. OSPFv3 supports IPv4 only on interfaces with family inet6 defined

Answer: D

Explanation:

Question: 43

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-unica3t 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.

Question: 44

Exhibit

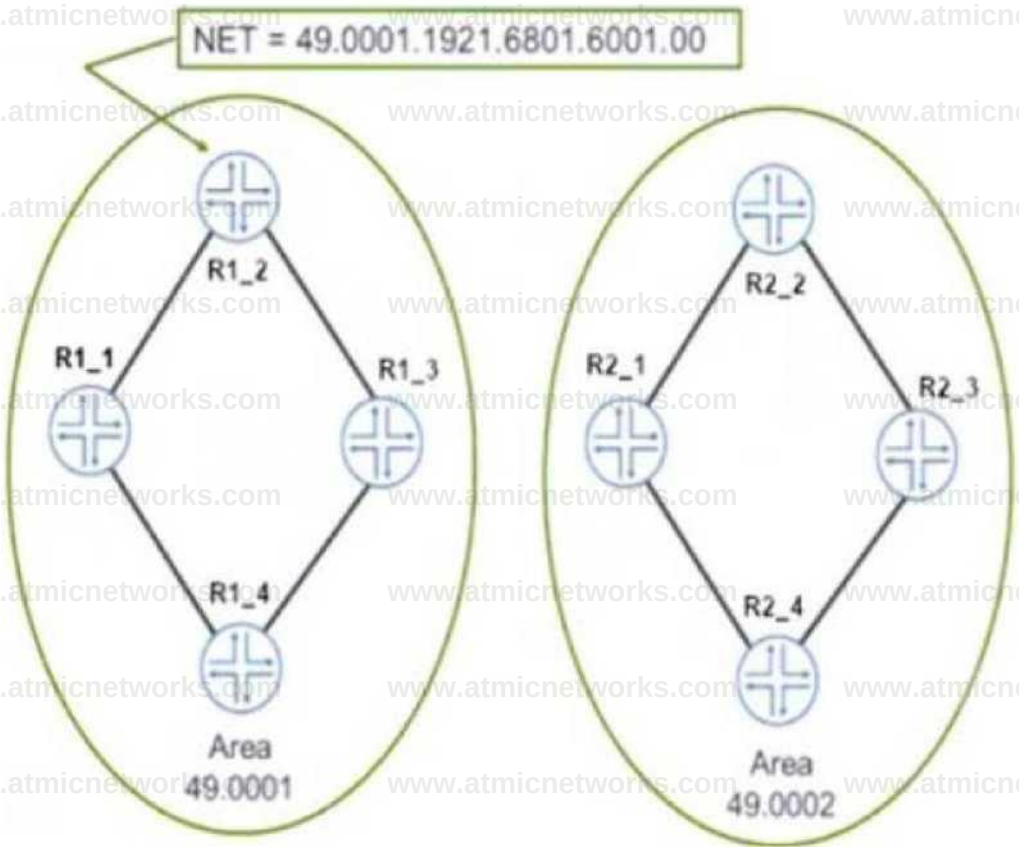

The network shown in the exhibit is based on IS-IS
Which statement is correct in this scenario?

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 IDofR1_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 domain1. 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 layer2. Therefore, the correct statement in this scenario is that the NSEL byte for area 0001 is 01.

Reference: 1: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/configuration/xe- 16/irs-xe-16-book/irs-ovrvw-cf.html 2: https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing- overview.html

Question: 45

Exhibit

(edit policy-options] userJrouter# show policy-statement biock-igap { term 1 (from {
route-filter 224.7.7.7/32 exact; source-address-filter 152.lc3.llC.13732 exac
}
then reject;
1

(edit protocols ignp]
user3routers show
interface ge-0/0/0.C (
group-policy block-igap;
group-limit 25;

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

www.atmicnetworks.com

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: A

Explanation:

This configuration applies to IGMP (Internet Group Management Protocol) and is designed to control multicast group memberships on the interface ge-0/0/0.0.

Breaking Down the Configuration

1 Qolicy-Statement: block-igmp

```
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;
  }
}
```

This policy blocks IGMP joins for group 224.7.7.7 only if the source IP is 192.168.100.10.

If both conditions match, the request is rejected.

2 QGMP Configuration on Interface ge-0/0/0.0

[edit protocols igmp]

user@router# show

www.atmicnetworks.com

```
interface ge-0/0/0.0 {
    group-policy block-igmp;
    group-limit 25;
}
```

group-policy block-igmp applies the policy statement block-igmp, meaning IGMP join requests are evaluated based on this policy.

group-limit 25 means the interface allows up to 25 multicast groups.

Evaluating the Answer Choices

A. Joins for group 224.7.7.7 are rejected if the source address is 192.168.100.10.

Correct, because:

The policy specifically matches group 224.7.7.7 and source IP 192.168.100.10.

If both conditions are met, the join is rejected.

X B. Joins for any group are accepted if the group count value is less than 25.

Incorrect, because:

The group-limit (25) applies to the total number of IGMP groups but does not override explicit policy rules.

Even if there are fewer than 25 groups, a join request can still be rejected by the policy statement.

X C. Joins for group 224.7.7.7 are always rejected, regardless of the group count.

Incorrect, because:

The policy only blocks joins from the specific source 192.168.100.10.

Joins from other sources to 224.7.7.7 are allowed.

X D. Joins for group 224.7.7.7 are accepted if the group count is less than 25.

Incorrect, because:

Joins for 224.7.7.7 from source 192.168.100.10 will always be rejected, even if the group count is below 25.

The group-limit does not override the rejection policy.

"Joins for group 224.7.7.7 are rejected if the source address is 192.168.100.10."

Official Juniper Documentation Reference:
Junos IGMP Policy Configuration Guide

"A group-policy statement allows filtering IGMP joins based on multicast group address and source IP."

Question: 46

Exhibit

[edit routing-instances CE-1] userSRI* show protocols { bgp { group CE-1 { type external; peer-as €5010; neighbor 10.1.1.100;

}

}

instance-type vrf; interface ge-0/0/2.0; route-distinguisher €5512:1; vrf-target target:€5512:100; [edit routing-instances CE-2] user€R2* show protocols { bgp < group CE-2 { type external; peer-as €5020; neighbor 10.1.5.100;

} instance-type vrf; interface ge-0/0/3.0; route-distinguisher €5512:1; vrf-target target:€5512:100;

Referring to the exhibit, which statement is correct?

A. The vrf-target configuration will allow routes to be shared between CE-1 and CE-2.

B. The vrf-target configuration will stop routes from being shared between CE-1 and CE-2.

C. The route-distinguisher configuration will allow overlapping routes to be shared between CE-1 and CE-2.

D. The route-diatinguisher configuration will stop routes from being shared between CE-1 and CE-2.

Answer: A

Explanation:

In the exhibit, we see two VRF (Virtual Routing and Forwarding) instances, CE-1 and CE-2, configured on a Juniper router. Each VRF is associated with a route-distinguisher (RD) and a vrf-target value.

Understanding the Role of vrf-target

The vrf-target is used to define Route Targets (RT), which control the import and export of VPN routes in MPLS Layer 3 VPNs (L3VPNs).

If two VRFs share the same RT, they will import each other's routes, allowing communication between them.

In this case, both VRFs have the same vrf-target:

vrf-target target:65512:100;

Since both CE-1 and CE-2 have the same RT (65512:100), they will import and export each other's routes, enabling route sharing between them.

Understanding route-distinguisher (RD)

The RD (Route Distinguisher) only ensures uniqueness of overlapping IP prefixes within the MPLS network. It does not control route sharing between VRFs.

In the exhibit, both VRFs have the same RD (65512:1), but this does not influence whether they share routes.

Correct Answer Selection

A (Correct): The vrf-target configuration enables route sharing between CE-1 and CE-2 since they have the same RT (65512:100).

B (Incorrect): The vrf-target does the opposite—it allows sharing, not blocking.

C (Incorrect): The route-distinguisher only provides unique route identification, but does not affect route sharing.

D (Incorrect): Again, route-distinguisher has no impact on route sharing.

Reference from Juniper Official Documentation

Q Juniper Documentation - Junos MPLS VPNs Configuration Guide:
"Route targets (vrf-target) are used to control the import and export of VPN routes between different VRFs. VRFs with the same route target can import and export routes to each other, enabling inter- VRF communication."

Thus, the correct answer is:

Q A. The vrf-target configuration will allow routes to be shared between CE-1 and CE-2.

Question: 47

Exhibit

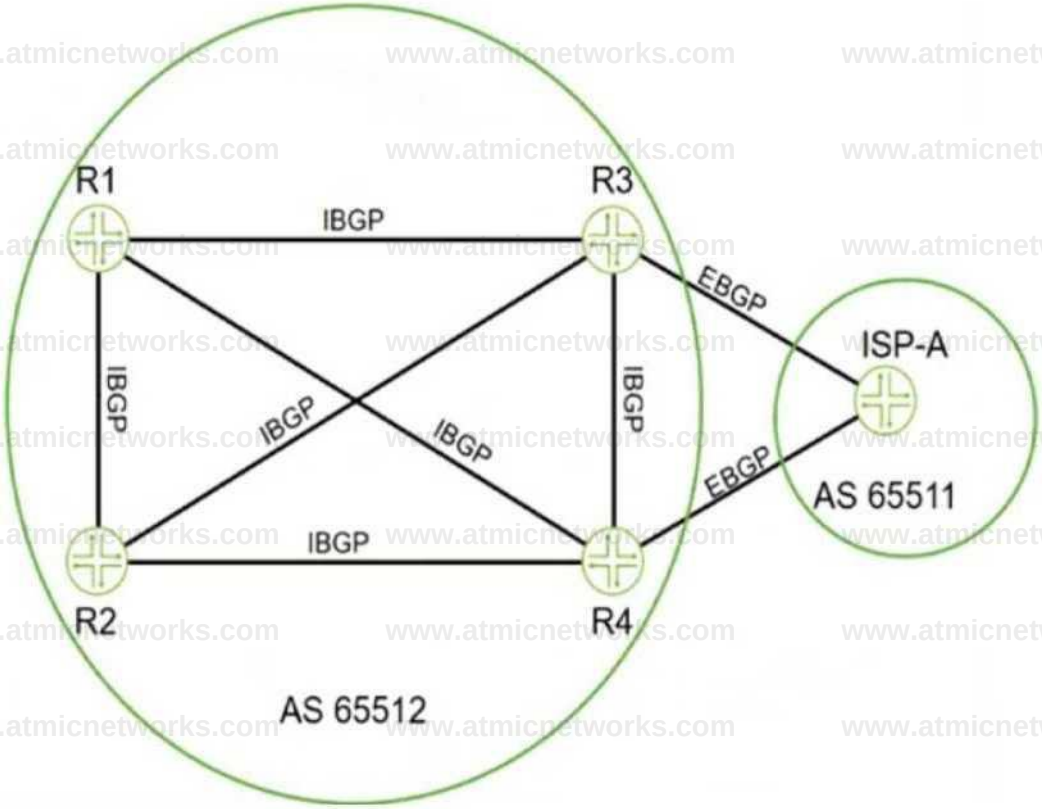

Click the Exhibit button-Referring to the exhibit, which two statements are correct about BGP routes on R3 that are learned from the ISP-A neighbor? (Choose two.)

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: BD

Explanation:

Analyzing the Exhibit

The diagram represents BGP peering between:

www.atmicnetworks.com

AS 65512 (Enterprise Network)

AS 65511 (ISP-A)

R3 and R4 are peering with ISP-A using EBGP.

R1, R2, R3, and R4 are peering within AS 65512 using IBGP.

Understanding BGP Route Behavior

Option A: "By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3." X

Incorrect!

EBGP behavior: When a BGP route is advertised via EBGP, the next-hop IP is changed to the router's own IP by default.

Since ISP-A is advertising routes via EBGP to R3, the next-hop is changed to ISP-A's IP.

Thus, this statement is incorrect.

Option B: "The BGP local-preference value that is used by ISP-A is not advertised to R3." Q

Correct!

BGP Local Preference (LOCAL_PREF) is an IBGP-only attribute.

Local Preference is NOT shared over EBGP because it is used within an AS to influence route selection.

ISP-A will not send LOCAL_PREF to R3, as R3 is in a different AS.

Thus, this statement is correct.

Option C: "All BGP attribute values must be removed before receiving the routes." X

Incorrect!

BGP does not remove all attributes when advertising routes. Some attributes are modified (e.g., next-hop, AS-PATH), but others (like MED, community) may be preserved.

Thus, this statement is incorrect.

Option D: "The next-hop value for these routes is changed by ISP-A before being sent to R3." Q

Correct!

As per default EBGP behavior, the next-hop is changed when a route is advertised to an EBGP peer.

This means ISP-A changes the next-hop to its own IP before sending it to R3.

Thus, this statement is correct.

Final Answer:

Q B. The BGP local-preference value that is used by ISP-A is not advertised to R3.

Q D. The next-hop value for these routes is changed by ISP-A before being sent to R3.

Verification from Juniper Documentation:

Juniper BGP Configuration Guide confirms that LOCAL_PREF is not advertised over EBGP.

RFC 4271 (BGP-4) specifies that next-hop is changed by default when advertising routes via EBGP.

Question: 48

Exhibit

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

You have MAC addresses moving in your EVPN environment

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

A. It identifies MAC addresses that should be discarded.

B. It resolves conflicting MAC address ownership claims.

C. It helps the local PE to identify the latest advertisement.

D. It is advertised using a Type 2 message

Answer: CD

Explanation:

In an EVPN (Ethernet Virtual Private Network) environment, MAC address mobility is a critical feature that allows devices to move across different locations while ensuring the network consistently tracks their MAC addresses.

Let's break down the components in the exhibit and analyze the correct statements.

Understanding MAC Mobility and Sequence Numbers in EVPN

In EVPN, MAC mobility is managed through sequence numbers that are included in Type 2 MAC/IP advertisements.

The sequence number tracks MAC movement events and is used to determine the most recent update when a MAC address appears on different PEs (Provider Edge devices).

When a MAC address moves between locations, the EVPN PEs increment the sequence number and advertise it to resolve conflicts and determine which PE has the most up-to-date information.

Now, Let's Review the Options:

Q C. It helps the local PE to identify the latest advertisement.

Correct:

The sequence number plays a key role in resolving MAC address conflicts. If multiple PEs advertise the same MAC address, the PE compares the sequence numbers to determine which update is the latest.

A higher sequence number indicates a more recent MAC update.

Q D. It is advertised using a Type 2 message.

Correct:

EVPN MAC/IP advertisements use BGP EVPN Type 2 messages to carry MAC addresses, IP addresses (optional), and their associated sequence numbers.

Type 2 advertisements are used to track MAC mobility and IP reachability information in the EVPN.

Why the Other Options Are Incorrect:

X A. It identifies MAC addresses that should be discarded.

Incorrect:

The sequence number doesn't identify MAC addresses that need to be discarded.

Instead, it resolves conflicts by determining the most recent MAC address advertisement based on the highest sequence number.

X B. It resolves conflicting MAC address ownership claims.

Partially true, but misleading:

While it's true that sequence numbers are used in conflict resolution, the sequence number itself doesn't directly resolve ownership claims. It only helps determine which advertisement is more recent. The actual conflict resolution happens through the comparison of the advertisements and sequence numbers.

Final Answer:

Q C. It helps the local PE to identify the latest advertisement.

Q D. It is advertised using a Type 2 message.

Reference from Juniper Documentation:

Juniper EVPN Configuration Guide:

"In EVPN MAC/IP advertisements, sequence numbers track the mobility of MAC addresses and are used to resolve conflicts when the same MAC address is advertised by multiple PEs. The PE with the higher sequence number has the most recent information."
Juniper BGP EVPN Mobility Documentation

Question: 49

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.

Question: 50

Exhibit

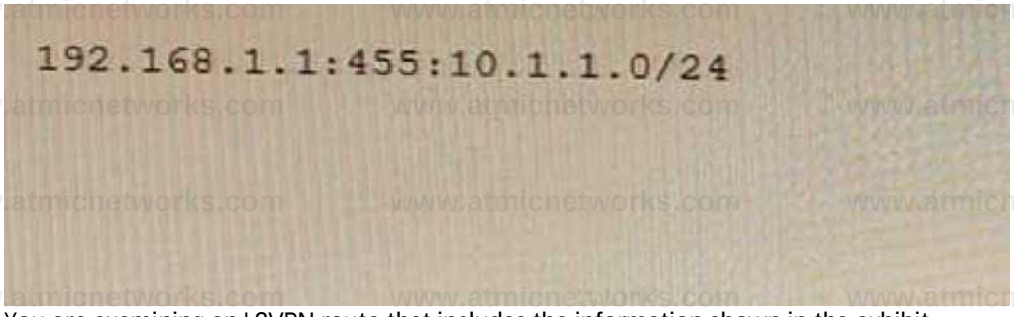

You are examining an L3VPN route that includes the information shown in the exhibit

Which statement is correct in this scenario?

A. The information shows a Type 1 route distinguisher.

B. The information shows a Type 0 route distinguisher

C. The information shows a Type 2 route distinguisher.

D. The information shows a route target

Answer: A

Explanation:

Type 1: When Type value is 1, the Administrator field is 4-bytes and Assigned Number field is 2- bytes. The Administrator field should be set to the IP address (public IP addresses should be used). The Assigned Number field contains a number from a numbering space that is administered by the enterprise to which the IP address has been assigned by the appropriate authority.

Question: 51

Exhibit

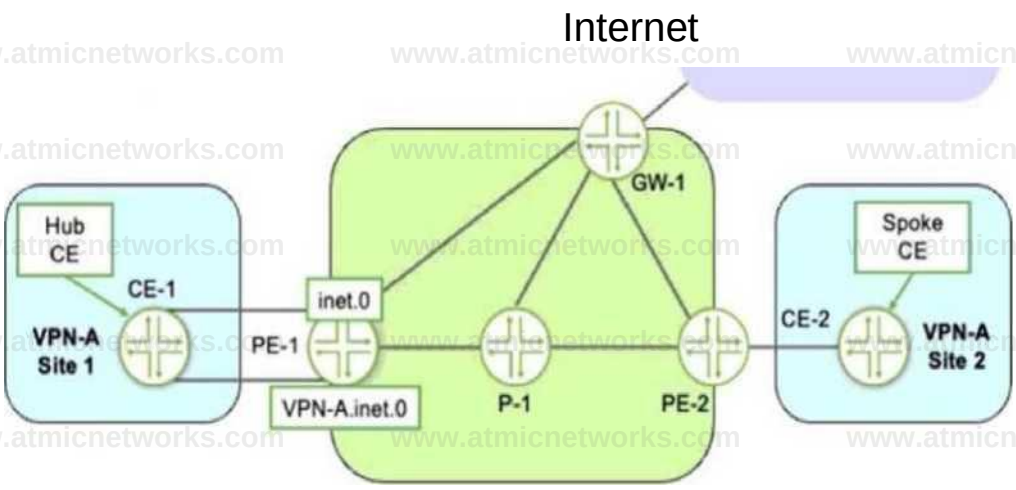

Referring to the exhibit, you must provide Internet access for VPN-A using CE-1 as the hub CE.

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

A. You must use RIB groups to leak routes between the inet. o and vpn-a. inet. o tables.

B. RIB groups are not needed to leak routes between the inet. 0 and VPN—A. inet. 0 tables,

C. Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> GW-1.

D. Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> CE-1 -> PE-1 -> GW-1.

Answer: BD

Explanation:

Question: 52

Exhibit

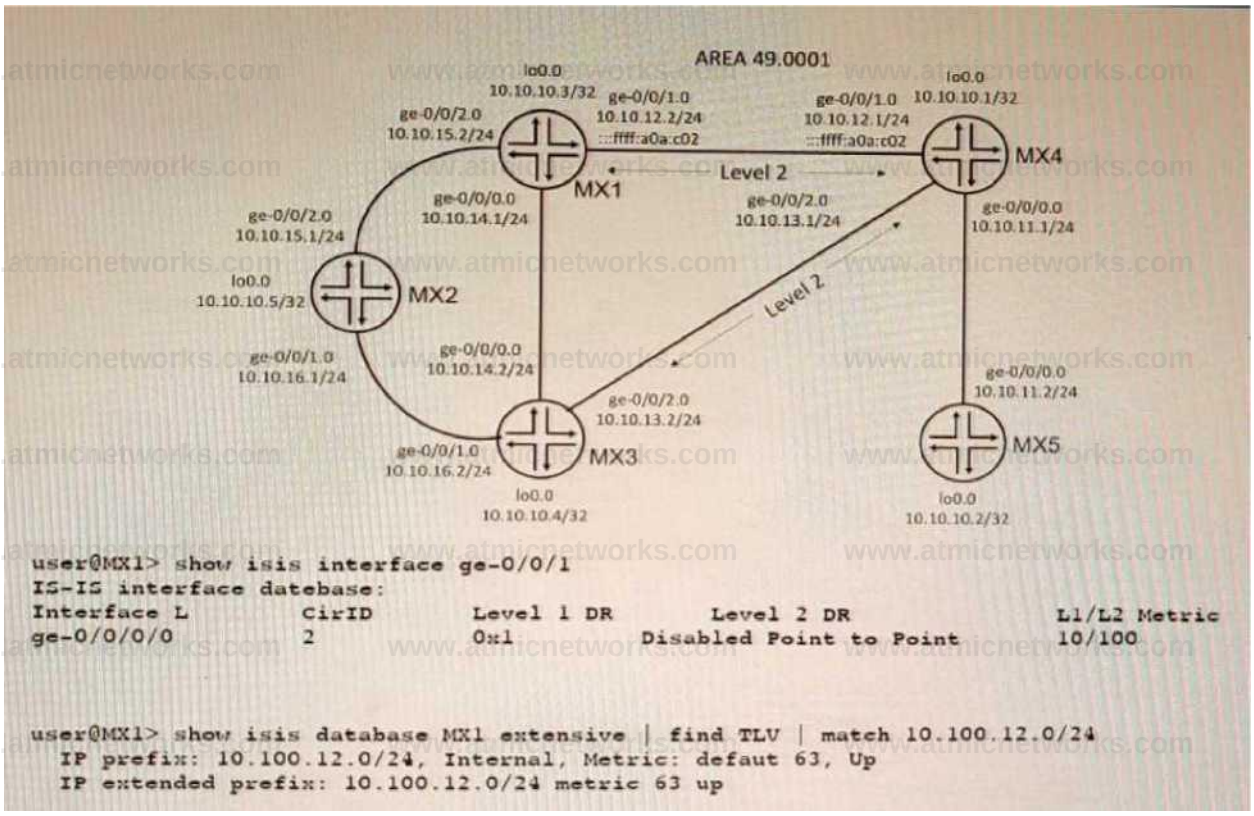

A network is using IS-IS for routing.

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

A. There are both narrow and wide metric devices in the topology

B. The interface specified a metric of 100 for L2.

C. Wide metrics have specifically been requested

D. Both IPv4 and IPv6 are being used in the topology

Answer: A

Explanation:

TLVs are tuples of (Type, Length, Value) that can be advertised in IS-IS packets. TLVs can carry different kinds of information in the Link State Packets (LSPs). IS-IS supports both narrow and wide metrics for link costs. Narrow metrics use a single octet to encode the link cost, while wide metrics use three octets. Narrow metrics have a maximum value of 63, while wide metrics have a maximum value of 16777215. If there are both narrow and wide metric devices in the topology, IS-IS will advertise two TLVs for each link: one with the narrow metric and one with the wide metric. This allows backward compatibility with older devices that only support narrow metrics12.

Question: 53

Exhibit

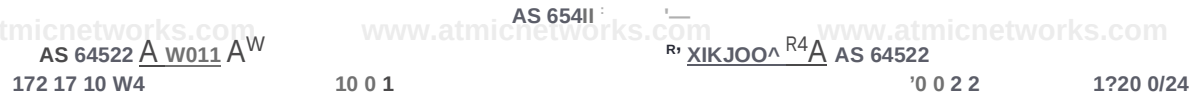

You are asked to exchange routes between R1 and R4 as shown in the exhibit. These two routers use the same AS number Which two steps will accomplish this task? (Choose two.)

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: B, C

Explanation:

https://www.juniper.net/documentation/us/en/software/junos/routing-policy/bgp/topics/example/bgp-advertise-peer-as.html

Question: 54

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 .

Question: 55

Exhibit

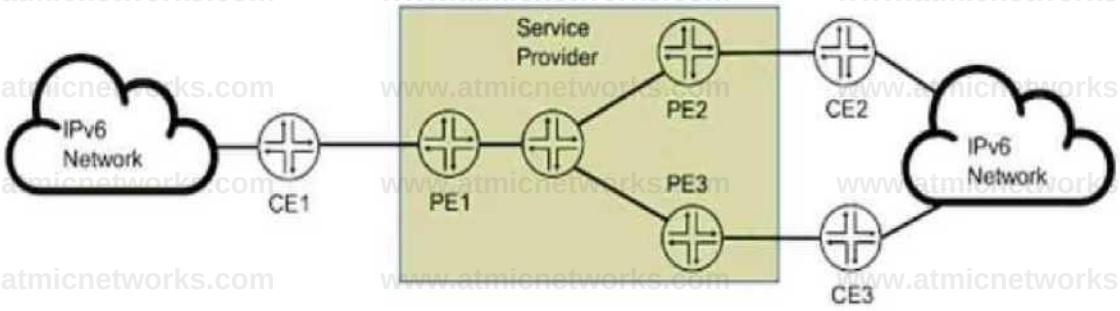

You are running a service provider network and must transport a customer's IPv6 traffic across your IPv4-based MPLS network using BGP You have already configured mpis ipv6-tunneling on your PE routers.

Which two statements are correct about the BGP configuration in this scenario? (Choose two.)

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

B. You must configure family inet6 add-path between PE and CE routers.

C. You must configure family inet6 unicast between PE and CE routers.

D. You must configure family inet6 unicast between PE routers.

Answer: AC

Explanation:

To transport IPv6 traffic over an IPv4-based MPLS network using BGP, you need to configure two address families: family inet6 labeled-unicast and family inet6 unicast. The former is used to exchange IPv6 routes with MPLS labels between PE routers, and the latter is used to exchange IPv6 routes without labels between PE and CE routers. The mpis ipv6-tunneling command enables the PE routers to encapsulate the IPv6 packets with an MPLS label stack and an IPv4 header before sending them over the MPLS network.

Question: 56

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

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

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

B. VXLAN must be configured on all PE routers.

C. EVPN uses Type 2 routes to advertise MAC address and IP address pairs learned using ARP snooping

D. EVPN uses Type 3 routes to join a multicast tree to flood traffic.

Answer: CD

Explanation:

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

Question: 57

Exhibit

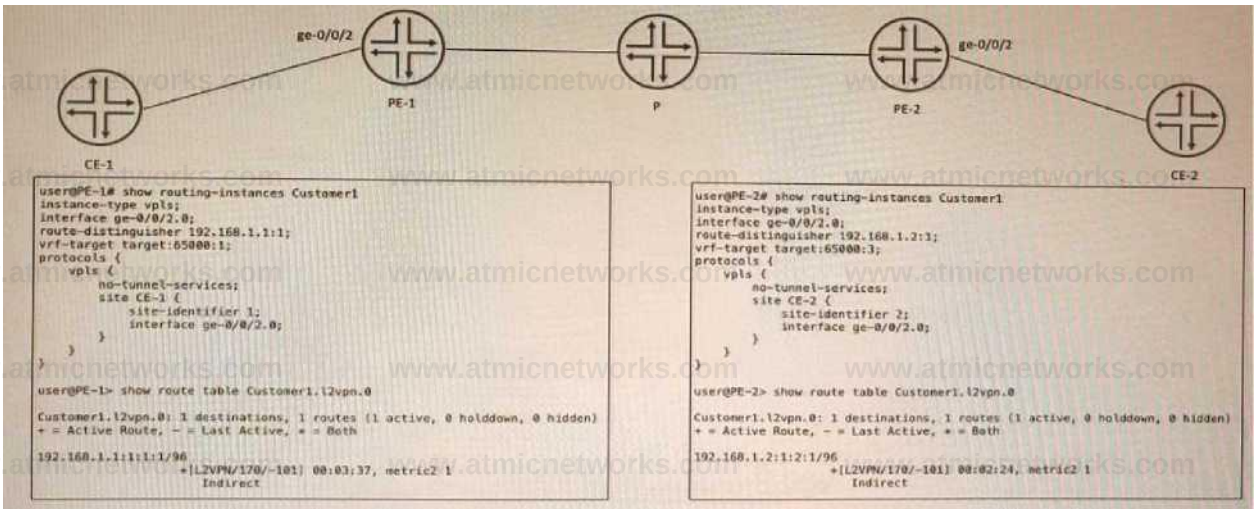

CE-1 and CE-2 are part of a VPLS called Customer1 No connectivity exists between CE-1 and CE-2. In the process of troubleshooting, you notice PE-1 is not learning any routes for this VPLS from PE-2, and PE-2 is not learning any routes for this VPLS from PE-1.

A. The route target must match on PE-1 and PE-2.

B. The route distinguisher must match on PE-1 and PE-2.

C. The instance type should be changed to l2vpn.

D. The no-tunnel-services statement should be deleted on both PEs.

Answer: A

Explanation:

VPLS is a technology that provides Layer 2 VPN services over an MPLS network. VPLS uses BGP as its control

protocol to exchange VPN membership information between PE routers. The route target is a BGP extended community attribute that identifies which VPN a route belongs to. The route target must match on PE routers that participate in the same VPLS instance, otherwise they will not accept or advertise routes for that VPLS.

Question: 58

Exhibit

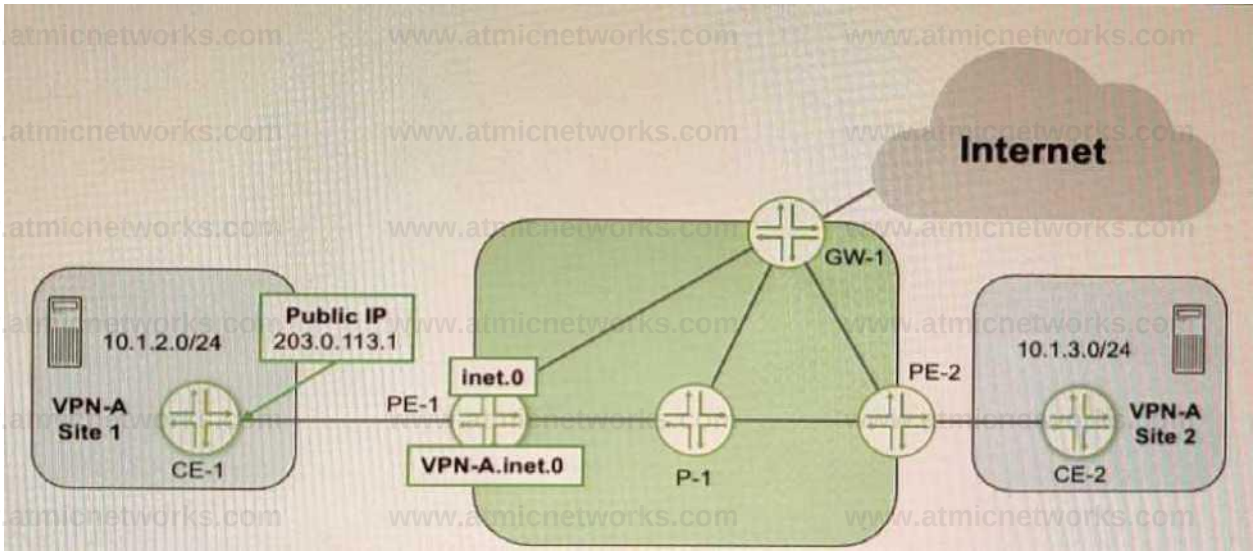

Referring to the exhibit, CE-1 is providing NAT services for the hosts at Site 1 and you must provide Internet access for those hosts

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

A. You must configure a static route in the main routing instance for the 10 1 2.0/24 prefix that uses the VPN-A.inet.0 table as the next hop

B. You must configure a static route in the main routing instance for the 203.0.113.1/32 prefix that uses the VPN-A.inet.0 table as the next hop.

C. You must configure a RIB group on PE-1 to leak a default route from the inet.0 table to the VPN- A.inet.0 table.

D. You must configure a RIB group on PE-1 to leak the 10 1 2.0/24 prefix from the VPN-A.inet.0 table to the inet.0 table.

Answer: B, C

Explanation:

Question: 59

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.

Question: 60

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.

Question: 61

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: A

Explanation:

The vrf-target statement is used in routing-instances to define route-target communities for VPN route import and export policies.

vrf-target export → Controls which route targets (RTs) are added to advertised routes (used when sending routes to remote PEs).

vrf-target import → Controls which VPN routes are accepted into the VRF (used when receiving routes from remote PEs).

Question: 62

Exhibit

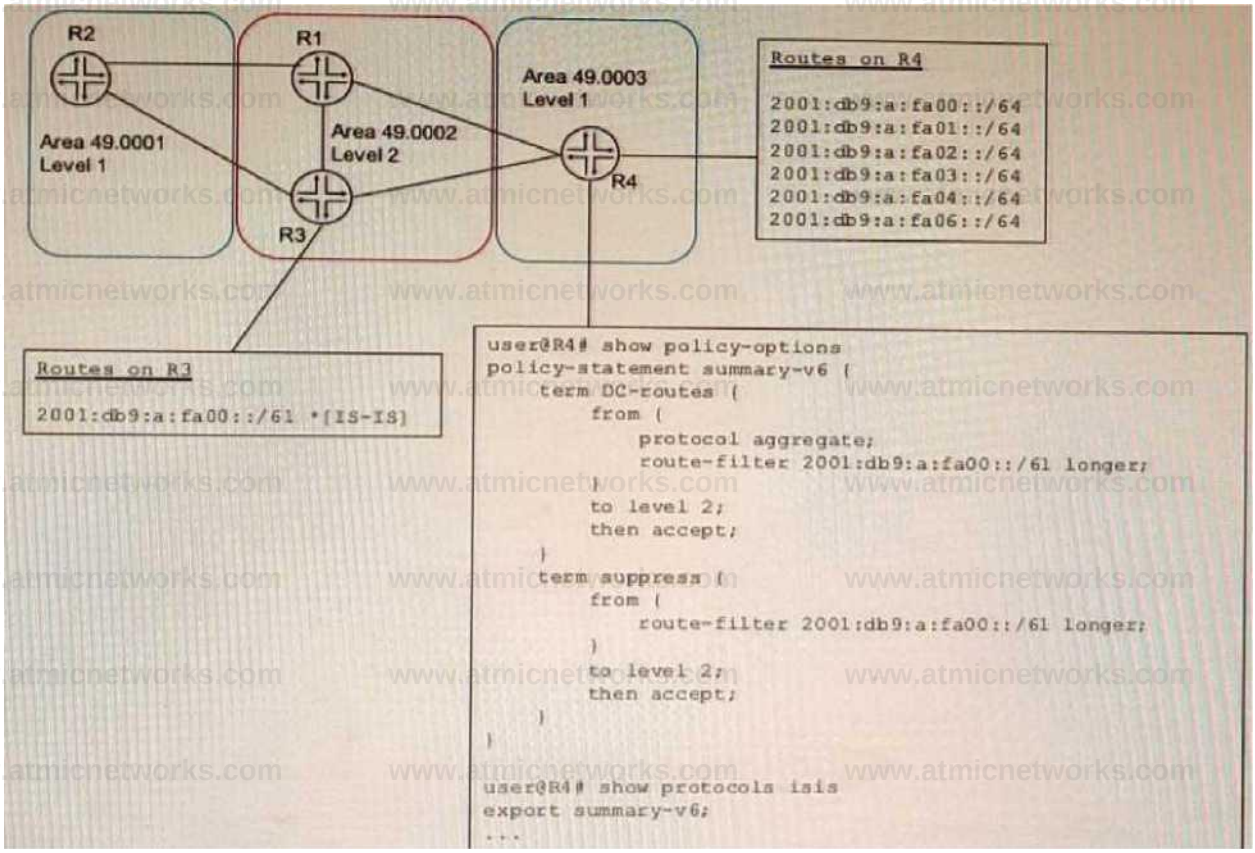

A network designer would like to create a summary route as shown in the exhibit, but the configuration is not working.

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 : :/6l longer

C. set policy—options policy-statement leak-v€ term DC—routes from route-filter 2001:db9:a:faOO::/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 summaryaddress 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

www.atmicnetworks.com

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.

Question: 63

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): l-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.O
   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.

Question: 64

Exhibit

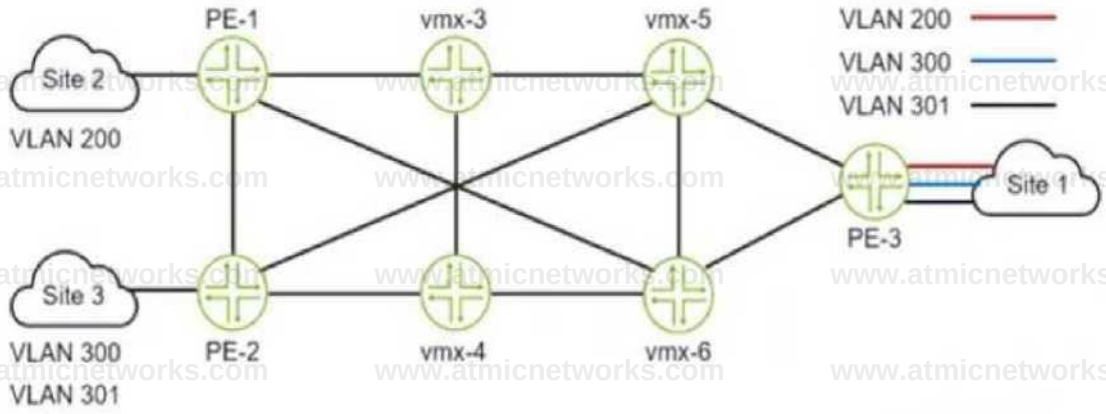

You want Site 1 to access three VLANs that are located in Site 2 and Site 3 The customer-facing interface on the PE-1 router is configured for Ethernet-VLAN encapsulation.

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.

Question: 65

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: B

Explanation:

https://www.juniper.net/documentation/us/en/software/junos/cos/topics/concept/packet-flow-cos- process-cos-config-guide.html

Question: 66

Exhibit.

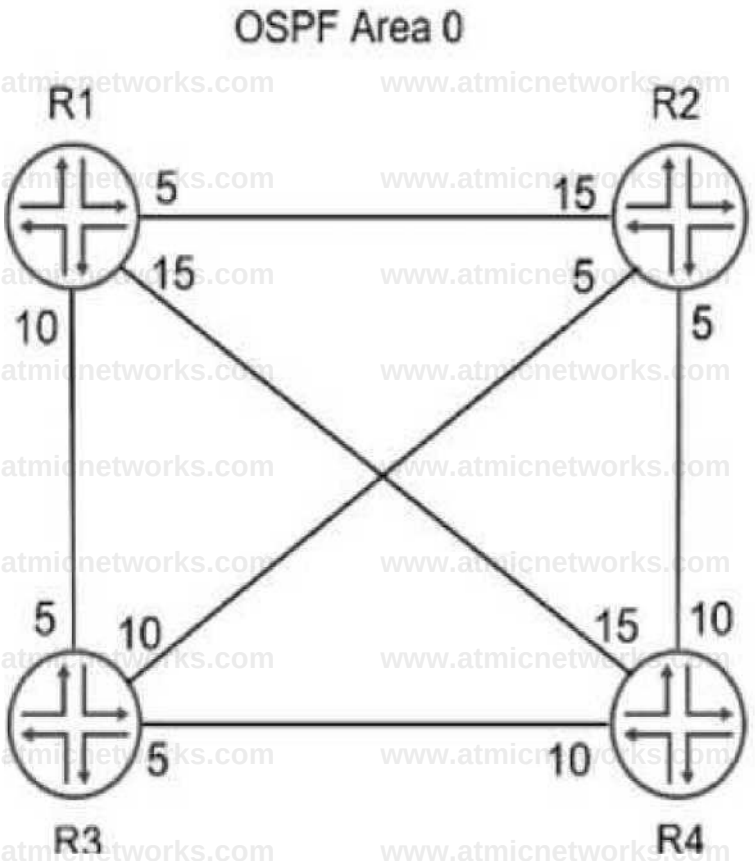

Referring to the exhibit, which path would traffic passing through R1 take to get to R4?

A. R1 -> R3 -> R4

B. R1 -> R2 -> R3 -> R4

C. R1 -> R2 -> R4

D. R1 -> R4

Answer: C

Explanation:

The OSPF cost is carried in the LSAs that are exchanged within an OSPF area. When a router calculates the cost to a destination it uses the cost of the exit interface of each router in the path to the destination.

Question: 67

Exhibit

(6500DR1 R2-- R3 (65003) [edit] user@R2# run show route 11.11.11.0/24 inet.O: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden) ♦ • Active Route, - • Last Active, * ■ Both

11.11.11.0/24 '[BGP/170] 00:00:50, localpref 100

www.atmicnetworks.com

```
AS path: 65001 I, validation-state: unverified > to 172.16.1.1 via ge-
0/0/0.0 (BGP/170] 00:00:50, localpref 100
AS path: 65003 I, validation-state: unverified (edit] user@R2# show
protocols bgp group RI { neighbor 172.16.1.1 { peer-as 65001;

group R3 {
    neighbor 172.16.2.1 { peer-as 65003;

 local-as 65002; [edit]
user8R2# show policy-options policy-statement lb { then (
        load-balance per-packet;

policy-statement prepend ( term 1 (
        then as-path-prepend 65001;

[edit]
ua«r@R2# show routing-options forwarding-table ( export lb;
}
```

R2 is receiving the same route from R1 and R3. You must ensure that you can load balance traffic for that route. Referring to the exhibit, which configuration change will allow load balancing?

A. Configure the multipath parameter under the global BGP configuration.

B. Apply the prepend policy as an import policy under group R1.

C. Configure the multipath multiple-as parameter under the global BGP configuration.

D. Apply the prepend policy as an import policy under group R3.

Answer: C

Explanation:

Question: 68

Exhibit

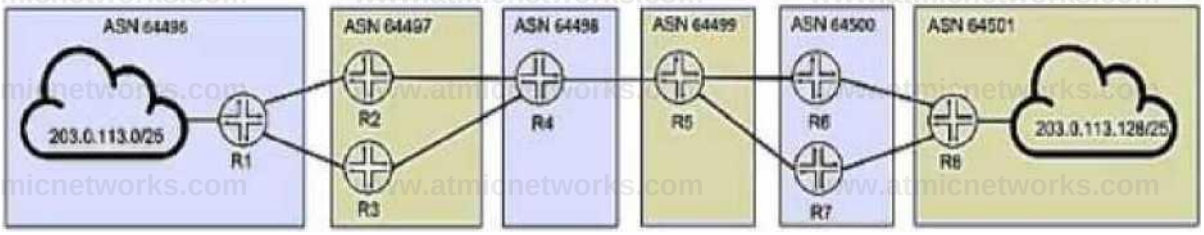

usen;Rl» thov route forwardvig-taBU Batching 243.6.113.124/24
Routing table: default.inet
Internet:

Oeitinotion	Type	RtRef	Next Teo	Type	Index	NMW	Netif
203.0.113.174/28	liter	0	10.1.1.1	uctt	176	11	ge-0/0/4.0
203.0.113.144/28	user	8	10.1.1.1	uc»t	576	11	ge-0/0/4.0
203.0.113.160/28	user	0	10.1.1.1	uc»t	576	11	ge-0/0/4.0
203.0.113.176/28	user	0	10.1.1.1	util	576	11	ge-0/0/4.0
203.0.113.102/28	user	0	10.1.1.1	uc»t	576	11	ge-0/0/4.0
203.0.113.208/28	user	8	10.1.1.1	vest	576	11	ge-0/0/4.0
203.0.113.224/28	user	0	10.1.1.1	uett	576	11	ge-0/0/4.0
203.O.1L3.24O/28	user	0	10.1.1.1	utti	576	11	ge-0/0/4.0

You are troubleshooting the connection between AS 64496 and AS 64497 and notice that only one of the paths is being used for traffic forwarding.

Referring to the exhibit, which three actions will ensure that R1 is configured properly for load balancing BGP routes? (Choose three.)

A. Verify that the routing table on R1 has BGP routes for 203.0.113.128/25 with multiple next hops.

B. Verify that the multipath option is configured under protocols bgp on both R2 and R3.

C. Verify that there is a load balancing export policy under routing-options for the received BGP routes on R1.

D. Verify that the multipath option is configured under protocols bgp on R1.

E. Verify that an import load balancing policy exists under protocols bgp for the received BGP routes on R1.

Answer: A, C, D

Explanation:

https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/load- balancing-bgp-session.html

Question: 69

You are configuring anycast RP for load balancing and redundancy in your PIM-SM domain. You want to share active sources between RPs.

In this scenario, what are two solutions that will accomplish this task? (Choose two.)

A. Configure MSDP on each RP router.

B. Configure anycast PIM with the rp-set statement on each RP router.

C. Configure anycast PIM with the rp-set statement on each source DR router.

D. Configure MSDP on each source DR router.

Answer: A, B

Explanation:

In a PIM Sparse Mode (PIM-SM) domain, Anycast RP is used for load balancing and redundancy by configuring multiple RPs with the same IP address. However, for active multicast sources to be shared between RPs, an additional mechanism is needed since PIM-SM does not automatically synchronize sources between RPs.

Evaluating the Answer Choices

Option A: "Configure MSDP on each RP router." Q

Multicast Source Discovery Protocol (MSDP) is required in an Anycast RP setup to share active source information between RPs.

MSDP allows RPs to exchange source-active (SA) messages, ensuring that multicast receivers in different regions can still receive traffic from sources registered with different RPs.

Juniper Documentation confirms that MSDP is used to synchronize active sources across multiple RPs in an Anycast RP deployment.

This is a correct answer.

Option B: "Configure anycast PIM with the rp-set statement on each RP router." Q

Anycast PIM allows multiple RPs to share the same IP address, and the rp-set statement is used to define the set of Anycast RPs.

This enables receivers and sources to register with the closest RP.

However, Anycast PIM alone does not share active source information between RPs; MSDP is still needed for that.

The combination of Anycast PIM (rp-set) and MSDP is the correct approach.

This is a correct answer.

Question: 70

Exhibit

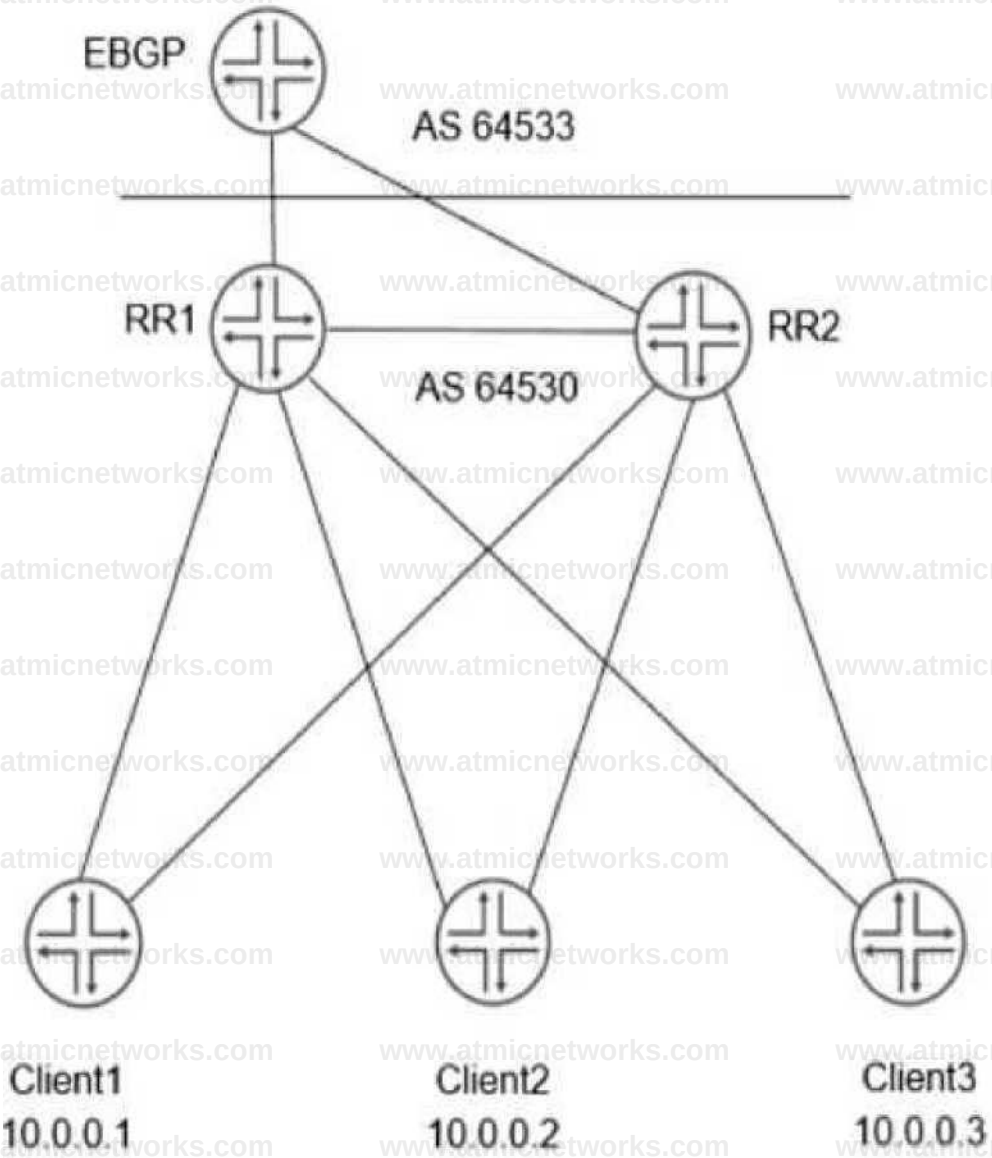

Referring to the exhibit, which two statements are correct about the dual route reflectors within a cluster? (Choose two.)

A. RR1 and RR2 must have the same duster ID to exchange routes learned from the client.

B. RR1 and RR2 append the duster ID when advertising routes from dient to dient.

C. RR1 and RR2 advertise routes learned from the clients to EBGP peers, using itself as the next hop.

D. RR1 advertises routes from the client to RR2. using itself as the next hop.

Answer: B, C

Explanation:

Question: 71

You want to ensure that a single-area OSPF network will be loop free.

In this scenario, what are two requirements that satisfy this requirement? (Choose two.)

www.atmicnetworks.com

A. The DR/BDR ensures that each node within an area has the same information in their LSDBs.

B. The Shortest Path First algorithm must prune looped paths.

C. Nodes within an area must connect in a full mesh.

D. All nodes within an area must have the same information in their LSDBs.

Answer: B, D

Explanation:

Question: 72

Your network is receiving the 203.0.113.0/24 network using EBGP from AS 64500 and AS 64501. Both of these advertisements have identical local-preference values, AS-path lengths, and BGP origin codes. You want to influence the way your AS sends traffic to the 203.0.113.0/24 network.

In this scenario, which attribute would you consider next when selecting the best path?

A. router ID

B. MED value

C. peer IP address

D. IGP metric

Answer: B

Explanation:

To determine the correct answer, let's analyze the BGP path selection process and identify which attribute would be considered next in this scenario.

Background on BGP Path Selection

When multiple paths to the same destination are received via BGP, the router uses a step-by-step process to select the best path. The order of attributes considered is as follows (simplified for this scenario):

Highest Local Preference : The path with the highest local preference is preferred.

Shortest AS Path : The path with the shortest AS path length is preferred.

Lowest Origin Code : Paths with an origin code of IGP are preferred over EGP, and EGP is preferred over

www.atmicnetworks.com

Incomplete.

Lowest MED (Multi-Exit Discriminator) : If the first three attributes are identical, the path with the lowest MED value is preferred.

eBGP over iBGP : eBGP paths are preferred over iBGP paths.

IGP Metric to Next Hop : The path with the lowest IGP metric to the next-hop router is preferred.

Router ID : If all else is equal, the path from the router with the lowest Router ID is preferred.

Peer IP Address : As a last tiebreaker, the path from the peer with the lowest IP address is preferred.

Scenario Analysis

In this scenario:

You are receiving the 203.0.113.0/24 network via EBGP from two different autonomous systems (AS 64500 and AS 64501).

Both advertisements have identical local-preference values , AS-path lengths , and BGP origin codes .

Given that the first three attributes in the BGP path selection process are identical, the next attribute to consider is the MED (Multi-Exit Discriminator) value.

Analysis of the Options

Option A: Router ID

Incorrect : The Router ID is considered much later in the BGP path selection process, only after other attributes like MED and IGP metric have been evaluated. Since MED is still relevant here, Router ID is not the next attribute to consider.

Option B: MED value

Correct : The MED value is used to influence inbound traffic from neighboring ASes. When local preference, AS path length, and origin code are identical, the path with the lowest MED value is preferred. This makes MED the next attribute to consider in this scenario.

Option C: Peer IP Address

Incorrect : The peer IP address is a tiebreaker used only at the very end of the BGP path selection process, after all other attributes have been evaluated. It is not relevant here because MED has not yet been considered.

Option D: IGP Metric

Incorrect : The IGP metric to the next-hop router is considered after MED. Since MED is still relevant in this scenario, IGP metric is not the next attribute to evaluate.

Final Answer

The correct answer is:

B . MED value

Summary

When local preference, AS path length, and origin code are identical, the MED value is the next attribute considered in the BGP path selection process.

MED is used to influence how traffic enters your AS from neighboring ASes.

Question: 73

Refer to the exhibit.

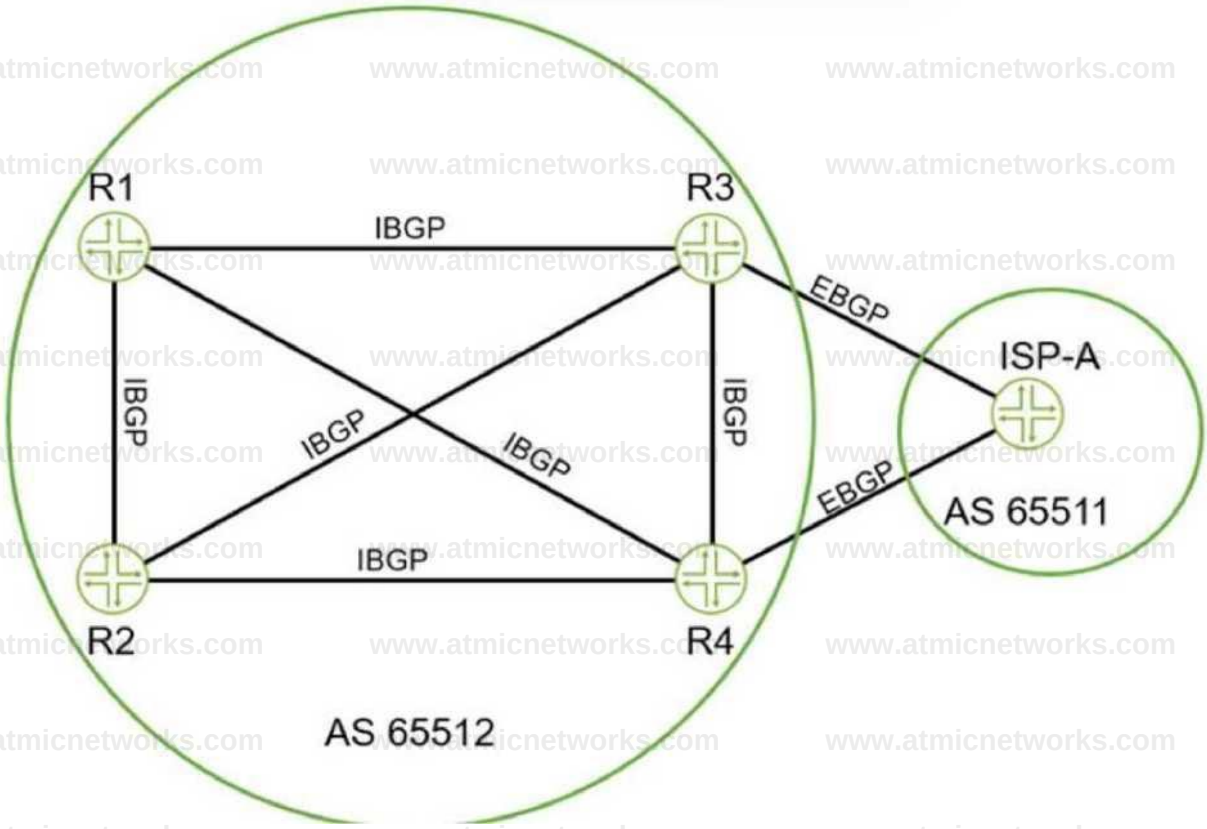

www.atmicnetworks.com

Click the Exhibit button.

Referring to the exhibit, which two statements are correct about BGP routes on R3 that are advertised to R1? (Choose two.)

A. By default, the next-hop value for these routes is changed by R3 before being sent to R1.

B. By default, all BGP attributes values must be removed before advertising the routes to R1.

C. By default, the BGP local-preference value that is assigned on R3 is advertised to R1.

D. By default, the next-hop value for these routes is not changed by R3 before being sent to R1.

Answer: CD

Explanation:

In the exhibit, we see an internal BGP (iBGP) setup within AS 65512, and an external BGP (eBGP) connection between R3 and ISP-A (AS 65511). The questions focus on the behavior of BGP routes advertised from R3 to R1 within the same AS.

1. **BGP Next-Hop Attribute (Option A and D)**:

- In iBGP, the next-hop attribute is **not** changed when a route is advertised to another iBGP peer. This means that when R3 advertises a route to R1, it retains the original next-hop value as learned from the eBGP peer (ISP-A).

- Therefore, Option D is correct: "By default, the next-hop value for these routes is not changed by R3 before being sent to R1."

2. **BGP Attributes (Option B and C)**:

- BGP attributes such as local preference, AS-path, and others are crucial for BGP route selection. The local preference attribute is used within an AS to indicate the preferred path for outbound traffic.

- When R3 advertises BGP routes to R1, it includes the local preference value assigned to those routes. This value is not removed and is propagated within the iBGP mesh.

- Therefore, Option C is correct: "By default, the BGP local-preference value that is assigned on R3 is advertised to R1."

Reference:

- Juniper Networks documentation on BGP behavior provides detailed insights into the propagation of BGP attributes within iBGP and eBGP contexts. Specifically, the Junos OS documentation covers the default behavior of next-hop and local preference attributes in BGP configurations.

- Junos OS BGP Configuration Guide: [Junos OS BGP Configuration Guide](https://www.juniper.net/documentation/en_US/junos/topics/concept/bgp-overview.html)

- For a deeper understanding of BGP attributes and their default behaviors, the "Juniper Networks Day One: Exploring BGP" book is an excellent resource.

Question: 74

Refer to the exhibit.

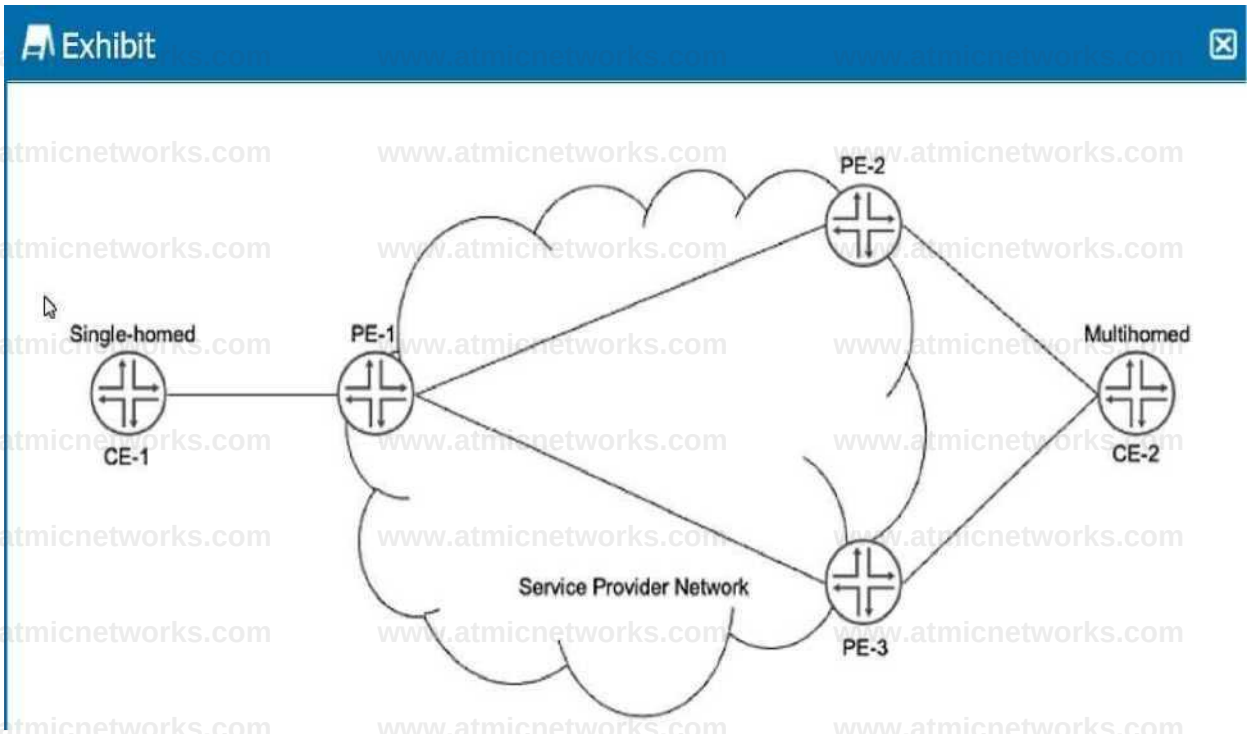

www.atmicnetworks.com

Click the Exhibit button.

You have an EVI implemented between PE-1, PE-2, and PE-3 to allow communication between CE-1 and CE-2. CE-2 receives unicast traffic from CE-1 on both links to PE-2 and

PE-3. When CE-1 sends broadcast traffic, CE-2 receives it on only one of the multihomed links.

Referring to the exhibit, which EVPN route type enables this behavior?

A. Type 4

B. Type 3

C. Type 1

D. Type 2

Answer: B

Explanation:

In the context of Ethernet VPN (EVPN) and the behavior described in the exhibit, it's essential to understand the different EVPN route types and their specific functionalities. Here, CE-2 is receiving unicast traffic on both of its multihomed links to PE-2 and PE-3, but broadcast traffic is received only on one of these links.

Explanation of EVPN Route Types:

1. **Type 1 (Ethernet Auto-Discovery Routes)**:
   - These routes are used for auto-discovery of Ethernet segments and for advertising VLAN membership.
   - They do not directly influence the behavior described in the question.

www.atmicnetworks.com

2. **Type 2 (MAC/IP Advertisement Routes)**:

   - These routes are used to advertise MAC addresses and IP-to-MAC bindings within the EVPN.

   - They handle unicast traffic forwarding and are crucial for populating the MAC address tables on the PE devices.

   - While important, they do not explain the selective broadcast behavior.

3. **Type 3 (Inclusive Multicast Ethernet Tag Routes)**:

   - These routes are used to build multicast distribution trees for delivering broadcast, unknown unicast, and multicast (BUM) traffic.

   - They ensure that BUM traffic is sent only once per Ethernet segment, preventing duplicate frames from being sent to multihomed CEs.

   - This aligns with the behavior described where CE-2 receives broadcast traffic on only one link to prevent duplication.

4. **Type 4 (Ethernet Segment Routes)**:

   - These routes are used to advertise the presence of an Ethernet segment and are crucial for Designated Forwarder (DF) election processes in multihoming scenarios.

   - While relevant to multihoming, they are not directly responsible for the selective broadcast behavior.

Conclusion:

The behavior described, where CE-2 receives broadcast traffic on only one of its multihomed links, is controlled by Type 3 routes. These routes are specifically designed to handle inclusive multicast and broadcast traffic efficiently in EVPN environments, ensuring that such traffic is not duplicated across multiple links to the same CE.

Reference:

- Juniper Networks EVPN Documentation: [EVPN Overview](https://www.juniper.net/documentation/en_US/junos/topics/concept/evpn- overview.html)

- RFC 7432, BGP MPLS-Based Ethernet VPN: [RFC 7432](https://tools.ietf.org/html/rfc7432) provides detailed descriptions of EVPN route types and their functions.

- Junos OS EVPN Configuration Guide: [Junos OS EVPN Configuration Guide](https://www.juniper.net/documentation/en_US/junos/topics/topic-map/evpn.html)

Question: 75

Refer to the exhibit.

```
user@router> show route receive-protocol bgp 10.16.40.1 extensive inet.O: 233 destinations, 233
routes (233 active, 0 holddown, 0 hidden) * 10.10.0.0/24 (1 entry, 1 announced)
     Accepted
     Nexthop: 10.16.40.1
     AS path: 65000 {65137 65224) I
     Aggregator: 65.0.00 10.11.11.11 user@router> show route 10.10'. 0.0/24 inet.O: 233 destinations,
233 routes (233 active, 0 holddown, 0 hidden) + - Active Route, - - Last Active, ■* - Both
10.10.0.0/24        *[BGP/170] 00:12:17, localpref 100
                      AS path: 65000 (65137 65224} I, validation-state: unverified > to 10.16.40.1
                    via ge-0/0/2.0
```

Click the Exhibit button.You are troubleshooting an issue for a customer site that uses 10.10.0.0/24 in AS 65224, but you see another AS in the AS path.

Referring to the exhibit, what is the cause of the problem?

A. AS 65000 is pre-pending AS 65137 to route advertisements.

B. The local AS is receiving two equal cost routes to 10.10.0.0/24.

C. The local AS is in the process of withdrawing the route from AS 65137.

D. D. AS 65137 is advertising the 10.10.0.0/24 prefix.

Answer: D

Explanation:

Question: 76

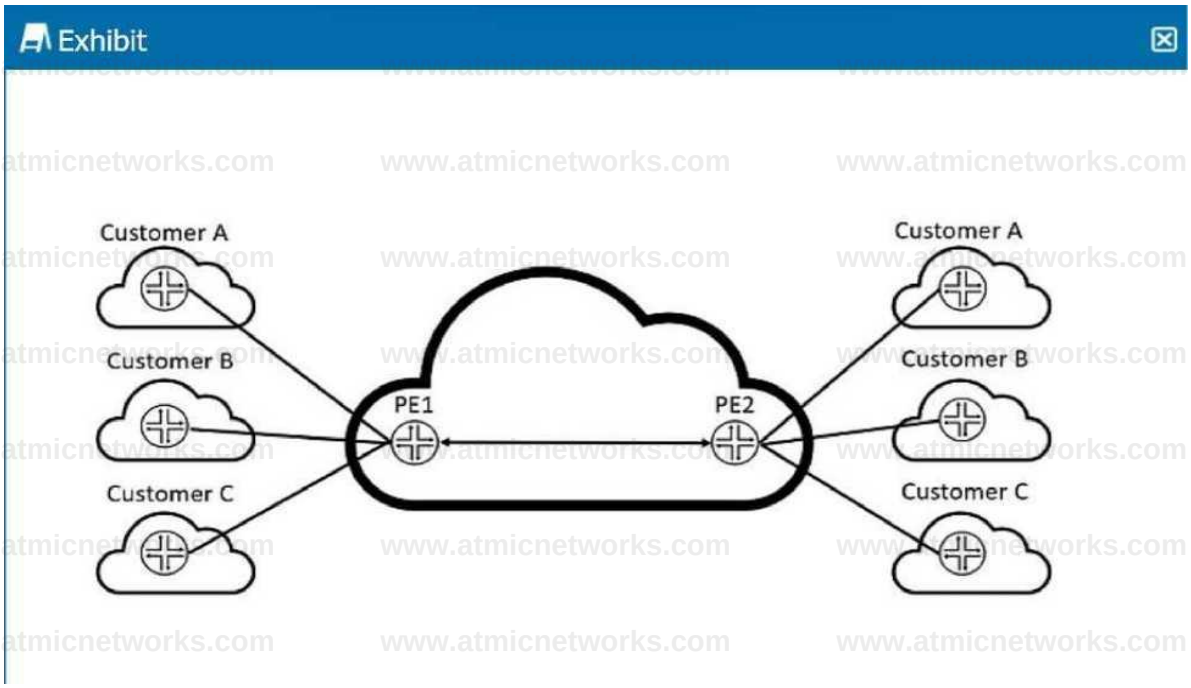

Refer to the exhibit.

www.atmicnetworks.com

Click the Exhibit button.

After adding Customer C to your Layer 3 VPN, you must ensure that PE2 is receiving VPN routes for all customers attached to PE1, as shown in the exhibit.

Which operational command displays this information?

A. show route table inet.0

B. show route table bgp.l3vpn.0

C. show route table customer-c.inet.0

D. show route summary

Answer: B

Explanation:

Understanding the Exhibit and the Problem Statement

The diagram shows a Layer 3 VPN (L3VPN) setup where multiple customers (Customer A, B, and C) are connected across a service provider network using PE (Provider Edge) routers.

PE1 and PE2 exchange VPN routes for all customers using BGP/MPLS Layer 3 VPN (L3VPN) routing.

The question asks how to verify if PE2 is receiving VPN routes for Customer C.

Evaluating the Answer Choices

Q B. show route table bgp.l3vpn.0 (Correct Answer)

Why?

This command shows all VPN routes stored in the BGP Layer 3 VPN table (bgp.l3vpn.0).

Since PE routers exchange VPN routes using MP-BGP, this table contains the VPN-IPv4 or VPN-IPv6 routes for all customers.

If PE2 is receiving routes from PE1 for Customer C, they will appear in bgp.l3vpn.0.

X A. show route table inet.0 (Incorrect)

Why?

The inet.0 table contains global unicast routes for the service provider's network.

VPN routes do not appear here because they are stored in VRF-specific tables.

This command won't help verify VPN route exchange between PE1 and PE2.

X C. show route table customer-c.inet.0 (Incorrect)

Why?

The table customer-c.inet.0 represents the VRF routing table for Customer C on the local PE.

It only shows locally installed routes for Customer C but does not confirm if PE2 is receiving routes from PE1.

This command is useful to check local VPN routing but not BGP route propagation.

X D. show route summary (Incorrect)

Why?

This command only provides a summary of route counts per protocol (BGP, OSPF, etc.).

It does not display specific VPN routes.

It is useful for general troubleshooting but doesn't confirm VPN route receipt.

Final Answer: Q show route table bgp.l3vpn.0 (Option B)

Explanation:

Official Juniper Documentation Reference

Junos MPLS VPNs Configuration Guide

Juniper Documentation

"The show route table bgp.l3vpn.0 command displays all VPN-IPv4 routes learned via MP-BGP for Layer 3 VPNs."

Question: 77

Refer to the exhibit.

B Exhibit 0

```
user@Rl>show .pi® join extensive 232.1.1.1 instance: PIM.master Family: INET
R = Rendezvous Point Tree, S J Sparse, W = Wildcard Group: 232.. 1.1.1   '
    Source: *
    RP: 10.1.255.112
    Flags: sparse,.rptr.ee, wildcard Upstream interface: ge-0/0/0.0. Upstream neighbor: 10.1.11.1 Upstream state: Join to RP Uptime:
    00:04:10 Downstream neighbors:
        Interface: Local Interface: ge-0/0/2.0 10.1.1.1 State: Join Flags: SRW Uptime: 00:04:10 Time since last
    Number of downstream interfaces: 2 NtBiber .of downstream neighbors: 1
Group: 232.1.1.1                                    Timeout: Infinity
                                                    Join: 00:04:10
    Source: 172.16.1.2
    Flags: sparse, spt
    Upstream interface: ge-O/C/l.O
    Upstream neighbor: 10.1.21,1
    Upstream state: Join to Source, Prune to RP
    Keepalive timeout: 317
    Uptime: 00:01.: 39                                                    &
    Downstream neighbors:
        Interface: Local
        Interface: ge-0/0/2.0
            10.1.1.1 state: Join Flags: ;S Timeout: Infinity
            Uptime: 00:01:39 Time since last Join: 0.0:01:39
    Number .of downstream interfaces:: 2
    Number ff doWlistteam. neighbors: 1.
```

Click the Exhibit button.

Referring to the exhibit, which two statements are correct regarding the output shown in the exhibit? (Choose two.)

A. The multicast group is an ASM group.

B. The multicast traffic is using the SPT.

C. The multicast group is an SSM group.

D. The multicast traffic is using the RPT.

Answer: AB

Explanation:

www.atmicnetworks.com

In the provided exhibit, the output of the `show pim join extensive 232.1.1.1` command is shown. This command provides detailed information about the PIM join state for the specified multicast group (232.1.1.1) on the router R1. To determine the correct statements regarding the multicast traffic, let's analyze the output and the terms involved:

1. **ASM vs. SSM**:

- **ASM (Any-Source Multicast)**: In ASM, receivers are interested in receiving multicast traffic from any source sending to a particular multicast group.
- **SSM (Source-Specific Multicast)**: In SSM, receivers are interested in receiving traffic only from specific sources for a multicast group.

- **Group Address Range**:
  - ASM uses the range 224.0.0.0 to 239.255.255.255.
  - SSM uses the range 232.0.0.0 to 232.255.255.255.

Since the group address 232.1.1.1 falls within the SSM range (232.0.0.0/8), there might be confusion. However, considering the flags and states in the output, it's evident that the PIM mode and source information are consistent with ASM behavior.

www.atmicnetworks.com

2. †Multicast Trees**:

- **RPT (Rendezvous Point Tree)**: Multicast traffic initially uses the RPT, where the Rendezvous Point (RP) acts as an intermediate point.
- **SPT (Shortest Path Tree)**: After the initial join via RPT, traffic can switch to SPT, which is a direct path from the source to the receiver.

3. **Output Analysis**:

Flags:

- The flags `sparse, rp-tree, wildcard` indicate that the group 232.1.1.1 is currently using RPT. This is typical for ASM, where traffic initially goes through the RP.
- The flags `sparse, spt` indicate that for the source 172.16.1.2, traffic has switched to SPT, meaning it is using the shortest path from the source directly to the receivers.

Conclusion:

Based on the analysis:

- **A. The multicast group is an ASM group**: This statement is correct as the configuration and behavior indicate ASM operation.
- **B. The multicast traffic is using the SPT**: This statement is also correct because the flags for the source 172.16.1.2 indicate that the traffic is using the SPT.

Thus, the correct answers are:

* *A. The multicast group is an ASM group.**
* *B. The multicast traffic is using the SPT.**

* Juniper Networks PIM Documentation: [PIM Overview](https://www.juniper.net/documentation/en_US/junos/topics/concept/pim- overview.html)
* Junos OS Multicast Routing Configuration Guide: [Multicast Routing Configuration

†Reference**:

Guide](https://www.juniper.net/documentation/en_US/junos/topics/topic-map/multicast- routing.html)

Question: 78

Refer to the exhibit.

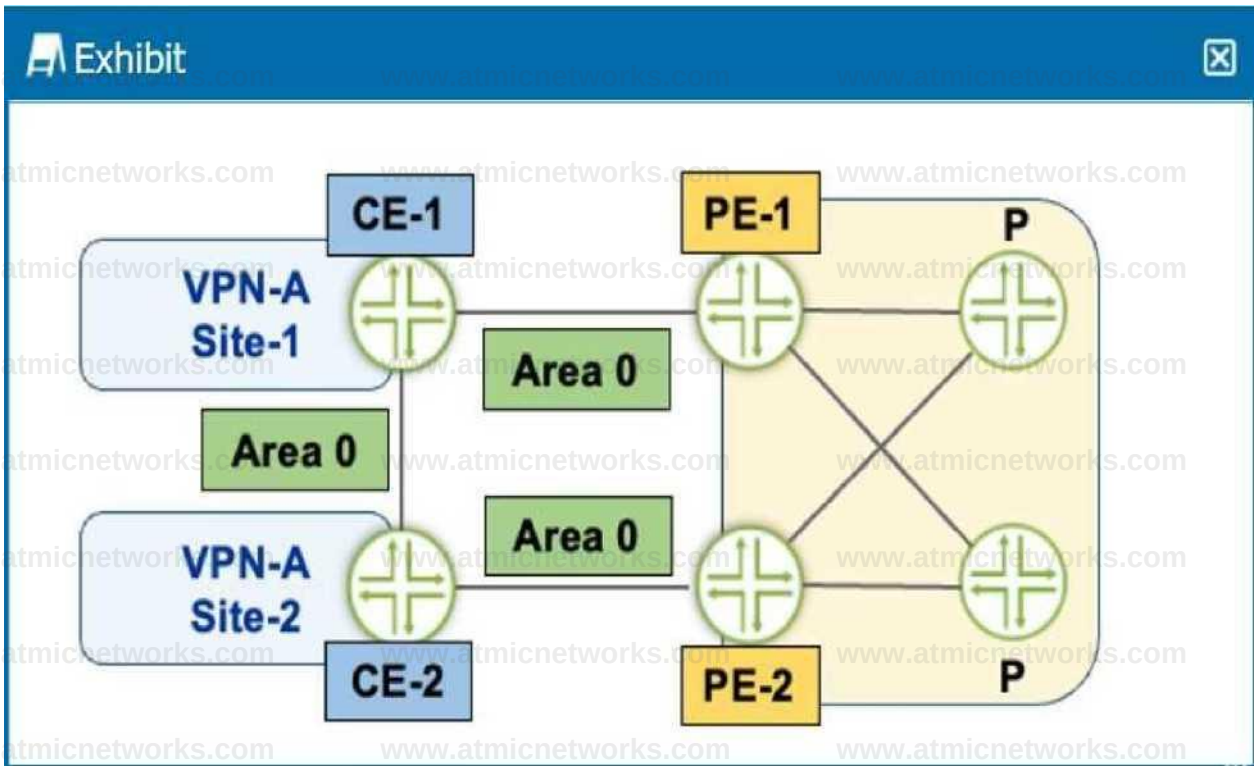

Click the Exhibit button.

Referring to the exhibit, the PE-to-CE protocol being used is OSPF for the L3VPN. Also, there is an OSPF neighborship between CE-1 and CE-2.

Which statement is correct in this situation?

A. You must set a high metric on the CE-1 to CE-2 link for hosts at Site-1 to use the L3VPN to reach hosts at Site-2.

B. Hosts at Site-1 will reach hosts at Site-2 through the CE-1 and CE-2 link by default.

C. Hosts at Site-1 will reach hosts at Site-2 through the L3VPN by default.

www.atmicnetworks.com

D. You must set a high metric on the CE-1 to PE-1 link for hosts at Site-1 to use the CE-1 to CE-2 link to reach hosts at Site-2.

Answer: B

Explanation:

In the exhibit, the PE-to-CE protocol used is OSPF, and there is an OSPF neighborship between CE-1 and CE-2 within the same Area 0. Let's analyze the default OSPF routing behavior in this setup to determine the correct statement.

1. **OSPF Neighborship**:
   - CE-1 and CE-2 have an OSPF neighborship directly within Area 0.
   - OSPF prefers intra-area routes over inter-area and external routes.

2. **Default Routing Behavior**:
   - Since CE-1 and CE-2 are directly connected through an OSPF link within the same area, OSPF will prefer this direct intra-area path over any other paths learned via the PE routers and the L3VPN.
   - This is because intra-area routes have a lower metric compared to inter-area or external routes.

3. **Metric Considerations**:
   - By default, OSPF will route traffic between Site-1 and Site-2 through the direct link between CE-1 and CE-2, unless the link's metric is artificially increased to make it less preferable.
   - There is no need to adjust metrics for the CE-1 to PE-1 link to prefer the CE-1 to CE-2 path, as OSPF already prefers direct intra-area paths.

Conclusion:

Given the default behavior of OSPF and the topology shown in the exhibit, the correct statement is:

- *B. Hosts at Site-1 will reach hosts at Site-2 through the CE-1 and CE-2 link by default.**

Reference:

- OSPF Design Guide: [Juniper Networks OSPF Design Guide](https://www.juniper.net/documentation/en_US/junos/topics/concept/ospf-design- overview.html)

- Juniper Networks Technical Documentation on OSPF: [Junos OS OSPF Configuration Guide](https://www.juniper.net/documentation/en_US/junos/topics/concept/ospf-routing- overview.html)

Question: 79

Refer to the exhibit.

```
user@Rl> show route protocol bgp
inet.O: 8 destinations, 12 routes (8 active, 0 holddown, 0 hidden) + = Active
Route, - = Last Active, * = Both
172.16.20.4/30     *[BGP/170] 00:49:55, localpref 100
                       AS path: 2 I, validation-state: unverified
                        > to 10.0.18.2 via ge-1/0/4.0
                          to 10.0.19.2 via ge-1/0/5.0
                       [BGP/170] 00:49:55, localpref 100
                        AS path: 2 I, validation-state: unverified
                        > to 10.0.19.2 via ge-1/0/5.0
```

Click the Exhibit button.

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

A. The multihop configuration is used for load balancing.

B. This route is learned from two different AS numbers.

C. This route is learned from the same AS number.

www.atmicnetworks.com

D. The multipath configuration is used for load balancing.

Answer: CD

Explanation:

In the exhibit, the output of the `show route protocol bgp` command is shown for the prefix `172.16.20.4/30`. Let's analyze the provided BGP routing table to determine which statements are correct.

1. **AS Path Analysis**:
   - The AS path for the route `172.16.20.4/30` is shown as `2 I`.
   - This indicates that the route was learned from AS 2 and it is an internal (iBGP) route within the same AS.

2. **Multiple Paths**:
   - The route has two next-hop IP addresses: `10.0.18.2` via interface `ge-1/0/4.0` and `10.0.19.2` via interface `ge-1/0/5.0`.
   - This indicates that BGP multipath is configured, which allows multiple equal-cost paths to be used for load balancing.
   - BGP multipath must be explicitly configured to use multiple paths for the same prefix.

3. **Multihop vs. Multipath**:
   - **Multihop Configuration**: This is typically used for establishing BGP sessions with peers that are not directly connected. It is not related to load balancing.
   - **Multipath Configuration**: This is used to enable load balancing across multiple paths for the same prefix, which is the case here.

Conclusion:

Given the above analysis:

- **C. This route is learned from the same AS number**: Correct. The AS path `2 I` indicates the route was

www.atmicnetworks.com

learned from the same AS number (AS 2).

- **D. The multipath configuration is used for load balancing**: Correct. The presence of multiple next-hops indicates that BGP multipath is configured for load balancing.

Thus, the correct answers are:

* *C. This route is learned from the same AS number.**

* *D. The multipath configuration is used for load balancing.**

Reference:

* Junos OS BGP Multipath Documentation: [Junos OS BGP Multipath](https://www.juniper.net/documentation/en_US/junos/topics/topic-map/bgp- multipath.html)

* Junos OS BGP Configuration Guide: [Junos OS BGP Configuration](https://www.juniper.net/documentation/en_US/junos/topics/concept/bgp-routing-overview.html)

Question: 80

Which three statements about IS-IS in a multi-area network are correct? (Choose three.)

A. Internal L1 PDUs are flooded to the local area's L2 routers.

B. External L2 PDUs are flooded to all L2 routers in other areas.

C. Internal L1 PDUs are flooded to all L1 routers in other areas.

D. Internal L1 PDUs are only flooded to the local area's L1 routers.

E. External L2 PDUs are only flooded to the local area's L2 routers.

www.atmicnetworks.com

Answer: ABD

Explanation:

Intermediate System to Intermediate System (IS-IS) is a link-state routing protocol designed to move information efficiently within a computer network, a group of physically connected computers or similar devices. It operates in two levels, Level 1 (L1) and Level 2 (L2), and supports hierarchical routing within a multi-area network.

Let's analyze each statement to determine its correctness in the context of IS-IS multi-area networks.

1. **Statement A: Internal L1 PDUs are flooded to the local area's L2 routers.**

- This statement is correct. L1 PDUs (Protocol Data Units) are flooded within the L1 area and also to the L2 routers that are present in the same area. These L2 routers act as the boundary routers that connect the local L1 area to other L1 areas via L2.

2. **Statement B: External L2 PDUs are flooded to all L2 routers in other areas.**

- This statement is correct. L2 PDUs are flooded throughout the entire L2 backbone, which includes all L2 routers in different areas. This ensures that inter-area routing information is shared across the network.

3. **Statement C: Internal L1 PDUs are flooded to all L1 routers in other areas.**

- This statement is incorrect. Internal L1 PDUs are only flooded within the local L1 area. They do not cross L1 area boundaries; inter-area communication is handled by L2 routers.

4. **Statement D: Internal L1 PDUs are only flooded to the local area's L1 routers.**

- This statement is correct. Internal L1 PDUs are indeed only flooded within their local L1 area, and do not go beyond it.

5. **Statement E: External L2 PDUs are only flooded to the local area's L2 routers.**

- This statement is incorrect. External L2 PDUs are flooded to all L2 routers throughout the IS-IS network, not just to those in the local area. This allows L2 routers to maintain a complete map of the network's topology.

Conclusion:

Given the analysis, the correct answers are:

www.atmicnetworks.com

* *A. Internal L1 PDUs are flooded to the local area's L2 routers.**

* *B. External L2 PDUs are flooded to all L2 routers in other areas.**

* *D. Internal L1 PDUs are only flooded to the local area's L1 routers.**

Reference:

* Juniper Networks Documentation on IS-IS: [IS-IS Overview](https://www.juniper.net/documentation/en_US/junos/topics/concept/is-is-routing- overview.html)

* RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments: [RFC 1195](https://tools.ietf.org/html/rfc1195) which details the operation of IS-IS in multi-area networks.

Question: 81

Refer to the exhibit.

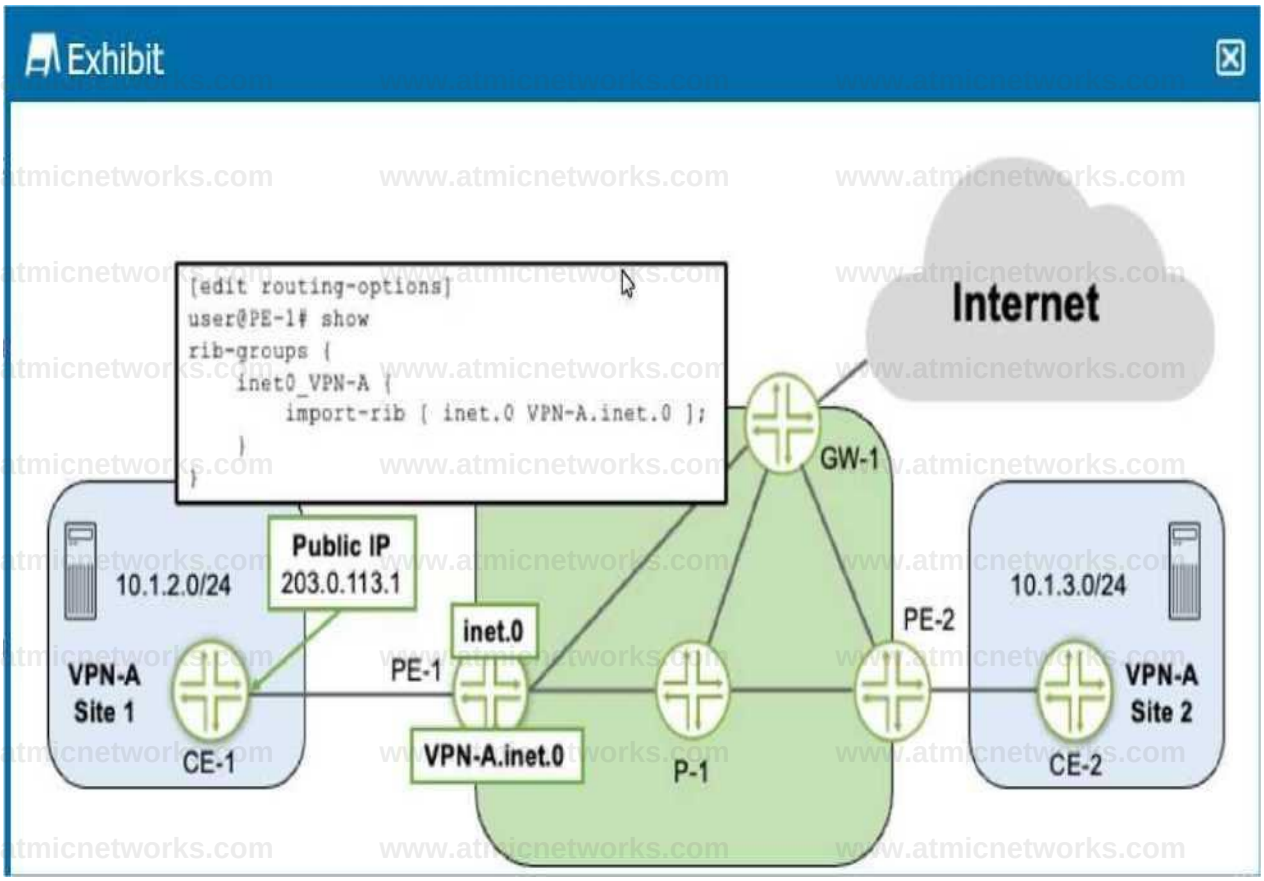

www.atmicnetworks.com

Click the Exhibit button.

Referring to the exhibit, you must provide VRF Internet access over a single connection for VPN-A Site 1, which connects to PE-1.

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

A. You must use the RIB group to move a default route, which is learned through BGP, from the inet. o table to the VPN-A. inet. 0 table.

B. You do not need to use the RIB group to move interface routes from the inet. o table to the VPN-A. inet. 0 table.

C. You do not need to use the RIB group default route, which is learned through BGP, from the inet. o table to the VPN-A. inet. 0 table.

D. You must use the RIB group to move interface routes from the inet . 0 table to the VPN-A. inet. 0 table.

Answer: AB

Explanation:

In the provided exhibit, the configuration involves using a RIB (Routing Information Base) group to facilitate internet access for VPN-A Site 1 through PE-1. The goal is to provide VRF Internet access over a single connection.

1. **Understanding RIB Groups**:
   - RIB groups allow for the import and export of routes between different routing tables.
   - In this scenario, we have two RIBs: `inet.0` (the main routing table) and `VPN-A.inet.0` (the VRF- specific routing table).

www.atmicnetworks.com

2. **Statement Analysis**:

- **A. You must use the RIB group to move a default route, which is learned through BGP, from the inet.0 table to the VPN-A.inet.0 table.**
  - Correct. To provide Internet access to VPN-A, the default route (0.0.0.0/0) learned via BGP in the `inet.0` table must be made available in the `VPN-A.inet.0` table. This is done using the RIB group to import the default route.

- **B. You do not need to use the RIB group to move interface routes from the inet.0 table to the VPN-A.inet.0 table.**
  - Correct. Interface routes (connected routes) are typically directly added to both the global and the VRF routing tables without needing a RIB group. These routes are known to the VRF because the interfaces are part of the VRF configuration.

- **C. You do not need to use the RIB group default route, which is learned through BGP, from the inet.0 table to the VPN-A.inet.0 table.**
  - Incorrect. As discussed, the default route needs to be imported into the VRF's routing table using a RIB group to enable Internet access for the VRF.

- **D. You must use the RIB group to move interface routes from the inet.0 table to the VPN-A.inet.0 table.**
  - Incorrect. Interface routes are directly associated with the VRF interfaces and are automatically known to the VRF routing table. There is no need to use a RIB group for these routes.

Conclusion:

The correct answers are:

* *A. You must use the RIB group to move a default route, which is learned through BGP, from the inet.0 table to the VPN-A.inet.0 table.**
* *B. You do not need to use the RIB group to move interface routes from the inet.0 table to the VPN-A.inet.0

table.**

Reference:

* Juniper Networks Documentation on RIB Groups: [RIB Groups Overview](https://www.juniper.net/documentation/en_US/junos/topics/concept/rib-groups- overview.html)

* Junos OS VPNs Configuration Guide: [Junos VPNs Configuration](https://www.juniper.net/documentation/en_US/junos/topics/concept/vpns- overview.html)

Question: 82

Which two statements are correct regarding the PIM DR in a PIM-SM domain? (Choose two.)

A. The source DR sends PIM register messages from the source network to the RP.
B. If the DR priorities match, the router with the lowest IP address is selected as the DR.

C. The receiver DR sends PIM join and PIM prune messages from the receiver network toward the RP.

D. By default, PIM DR election is performed on point-to-point links.

Answer: AC

Explanation:

In PIM-SM (Protocol Independent Multicast - Sparse Mode), the Designated Router (DR) plays a crucial role in multicast forwarding. The DR is responsible for various tasks depending on whether it is connected to the source or the receiver. Let's analyze each statement regarding the PIM DR in a PIM-SM domain.

1. **Statement A: The source DR sends PIM register messages from the source network to the RP.**

- Correct. In PIM-SM, the DR on the source's local network is responsible for encapsulating multicast packets in PIM Register messages and sending them to the Rendezvous Point (RP). This process ensures that the RP is

www.atmicnetworks.com

aware of active sources.

2. **Statement B: If the DR priorities match, the router with the lowest IP address is selected as the DR.**

- Incorrect. The correct rule is that if the DR priorities match, the router with the **highest** IP address is selected as the DR. The election process first compares priorities; if priorities are equal, the IP addresses are compared to select the DR.

3. **Statement C: The receiver DR sends PIM join and PIM prune messages from the receiver network toward the RP.**

- Correct. In PIM-SM, the DR on the receiver's local network sends PIM Join messages toward the RP to join the multicast distribution tree. Similarly, it sends PIM Prune messages to leave the tree when there are no interested receivers.

4. **Statement D: By default, PIM DR election is performed on point-to-point links.**

- Incorrect. By default, PIM DR election is performed on multi-access networks (e.g., Ethernet). On point-to-point links, there is no need for a DR election as there are only two routers involved.

Conclusion:

The correct statements regarding the PIM DR in a PIM-SM domain are:

A. The source DR sends PIM register messages from the source network to the RP.

C. The receiver DR sends PIM join and PIM prune messages from the receiver network toward the RP.

Reference:

- Juniper Networks Documentation on PIM-SM: [PIM-SM Overview](https://www.juniper.net/documentation/en_US/junos/topics/concept/pim-sparse-mode-overview.html)
- RFC 7761, Protocol Independent Multicast - Sparse Mode (PIM-SM): [RFC 7761](https://tools.ietf.org/html/rfc7761) which details the PIM-SM protocol, including DR roles and election

www.atmicnetworks.com

procedures.

Question: 83

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

A. PSNPs are used to acknowledge a received LSP.

B. CSNPs are used to acknowledge a received LSP.

C. CSNPs are used to request a missing LSP.

D. PSNPs are used to request a missing LSP.

Answer: AD

Explanation:

Intermediate System to Intermediate System (IS-IS) is a link-state routing protocol used to move information efficiently within a computer network. It uses a series of Protocol Data Units (PDUs) to manage the network's topology and ensure consistency across all routers in the network. Specifically, Link State PDUs (LSPs), Complete Sequence Number PDUs (CSNPs), and Partial Sequence Number PDUs (PSNPs) play crucial roles in this process.

1. **PSNPs (Partial Sequence Number PDUs)**:
   - **Acknowledge a received LSP**: PSNPs are used to acknowledge the receipt of LSPs. When a router receives an LSP, it sends a PSNP back to the sender to confirm that the LSP has been received.
   - **Request a missing LSP**: PSNPs are also used to request missing LSPs. If a router identifies a missing LSP based on sequence numbers, it can send a PSNP to request the specific LSP from its neighbors.

2. **CSNPs (Complete Sequence Number PDUs)**:

www.atmicnetworks.com

- **Summarize LSPs**: CSNPs are used to summarize all the LSPs known to a router. They are typically sent at regular intervals to provide a complete list of LSPs in a database. They are not used to acknowledge or request specific LSPs but provide an overview of all LSPs for database synchronization.

Based on this understanding, let's evaluate the statements:

- **A. PSNPs are used to acknowledge a received LSP.**

- Correct. PSNPs serve the purpose of acknowledging LSPs received from other routers.

- **B. CSNPs are used to acknowledge a received LSP.**

- Incorrect. CSNPs are not used for acknowledging LSPs; they are used to provide a summary of all LSPs.

- **C. CSNPs are used to request a missing LSP.**

- Incorrect. CSNPs are not used to request missing LSPs; this is the role of PSNPs.

- **D. PSNPs are used to request a missing LSP.**

- Correct. PSNPs are used to request specific missing LSPs when a router detects that it is missing information.

Conclusion:

The correct statements about IS-IS are:

* *A. PSNPs are used to acknowledge a received LSP.**

* *D. PSNPs are used to request a missing LSP.**

Reference:

* Juniper Networks Documentation on IS-IS: [IS-IS Overview](https://www.juniper.net/documentation/en_US/junos/topics/concept/is-is-routing- overview.html)

* RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments: [RFC 1195](https://tools.ietf.org/html/rfc1195) which details the operation and use of IS-IS, including the roles of PSNPs and CSNPs.

Question: 84

Refer to the exhibit.

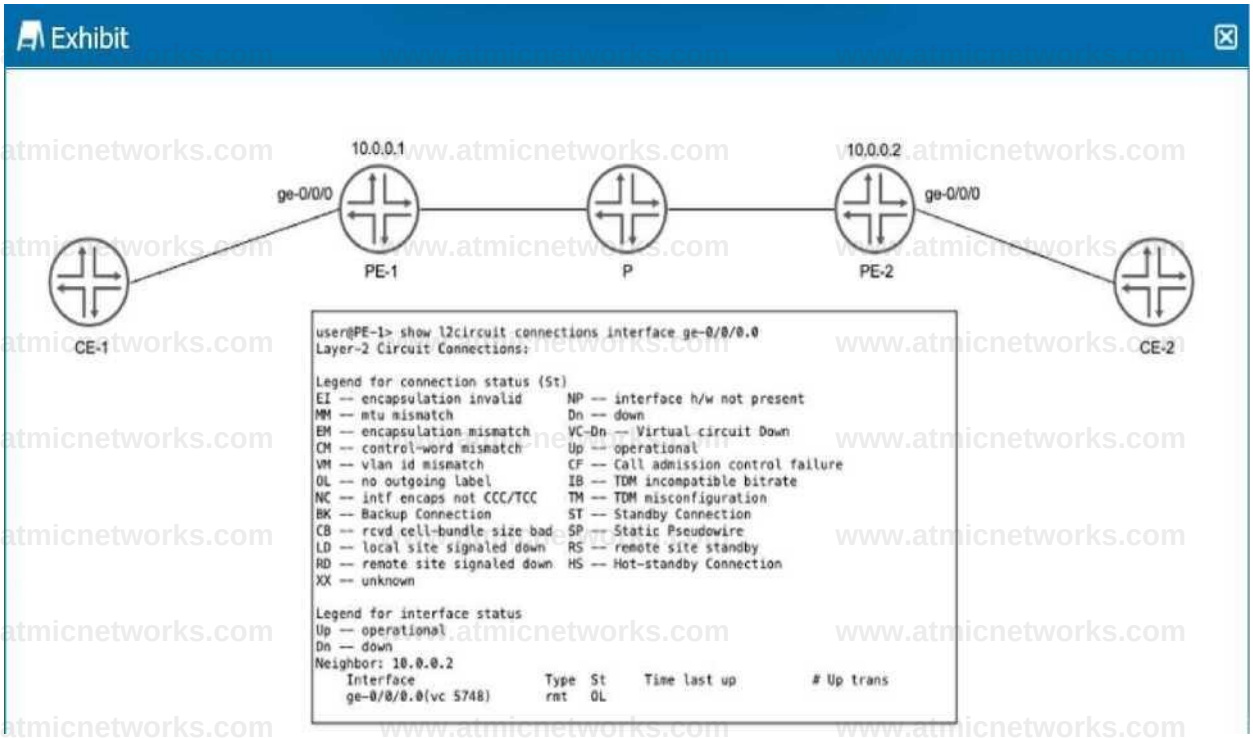

Click the Exhibit button.

PE-1 and PE-2 are configured with LDP-signaled pseudowires to provide connectivity between CE-1 and CE-2. You notice no connectivity exists between CE-1 and CE-2.

Referring to the exhibit, which two statements describe potential causes for this fault? (Choose two.)

A. The VC IDs are mismatched.

B. There is no LSP configured from PE-1 to PE-2.

C. Interface ge-0/0/0 on PE-1 is down.

D. There is no LSP configured from PE-2 to PE-1.

Answer: AD

Explanation:

Question: 85

You have an L2VPN connecting two CEs across a provider network that runs OSPF. You have OSPF configured on both CEs.

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

A. OSPF neighborship is formed between the CEs and PEs.

B. The CE and PE OSPF areas can be different.

C. The CE and PE OSPF areas must match.

D. OSPF neighborship is formed between the two CEs.

Answer: BD

Explanation:

In an L2VPN scenario, the provider network connects two customer edge (CE) devices across a Layer 2 virtual private network. Let's analyze how OSPF operates in this setup.

1. ‡OSPF Neighborship in L2VPN**:

- An L2VPN provides a Layer 2 connection between two sites, making it transparent to Layer 3 protocols like

‡Conclusion**:

www.atmicnetworks.com

OSPF. This means the CEs can form OSPF adjacencies directly with each other as if they were on the same local network.

2. **OSPF Configuration on CEs and PEs**:

- **Statement A: OSPF neighborship is formed between the CEs and PEs**:

- Incorrect. In an L2VPN, the provider's network is transparent to the OSPF running on the CEs. OSPF neighborship forms directly between the CEs, not between the CEs and PEs.

- **Statement B: The CE and PE OSPF areas can be different**:

- Correct. Since OSPF adjacencies form directly between the CEs and not between CEs and PEs, the OSPF areas on the CEs and PEs can be different. The provider network acts as a transparent bridge, and OSPF doesn't see the PEs.

- **Statement C: The CE and PE OSPF areas must match**:

- Incorrect. As noted above, because the OSPF neighborship forms directly between the CEs, the OSPF areas on the CEs and PEs do not need to match.

- **Statement D: OSPF neighborship is formed between the two CEs**:

- Correct. The L2VPN makes the connection between the two CEs appear as a direct Layer 2 link, allowing them to form an OSPF adjacency directly.

Given the above analysis, the correct statements are:

* *B. The CE and PE OSPF areas can be different.**

* *D. OSPF neighborship is formed between the two CEs.**

Reference:

* Juniper Networks Documentation on L2VPNs: [Configuring Layer 2 VPNs](https://www.juniper.net/documentation/en_US/junos/topics/task/configuration/layer-2-vpns-configuring.html)

www.atmicnetworks.com

* OSPF Configuration Guide: [Junos OS OSPF Configuration](https://www.juniper.net/documentation/en_US/junos/topics/concept/ospf-routing-overview.html)

Question: 86

* ou have an L2VPN connecting two CEs across a provider network. The CEs and provider network are configured with the default MTU setting. You use the ping command from one CE to the other CE with a size of 1500 bytes.

In this scenario, which statement is correct when using the ping command?

A. You expect the ping results to be fragmented.

B. You expect a silent discard.

C. You expect an echo reply.

D. You expect an ICMP message too long error.

Answer: B

Explanation:

Layer 2 VPNs don't support fragmentation in the provider network. It is critical that the provider network supports the largest frame that the CE devices can generate after the MPLS and virtual routing and forwarding (VRF) labels are added by the PE devices. This example leaves the CE devices at the default 1500-byte maximum transmission unit (MTU) while configuring the provider core to support a 4000 byte MTU. This configuration avoids discards by ensuring the CE devices cannot exceed the MTU in the provider's network.

www.atmicnetworks.com

Question: 87

Refer to the exhibit.

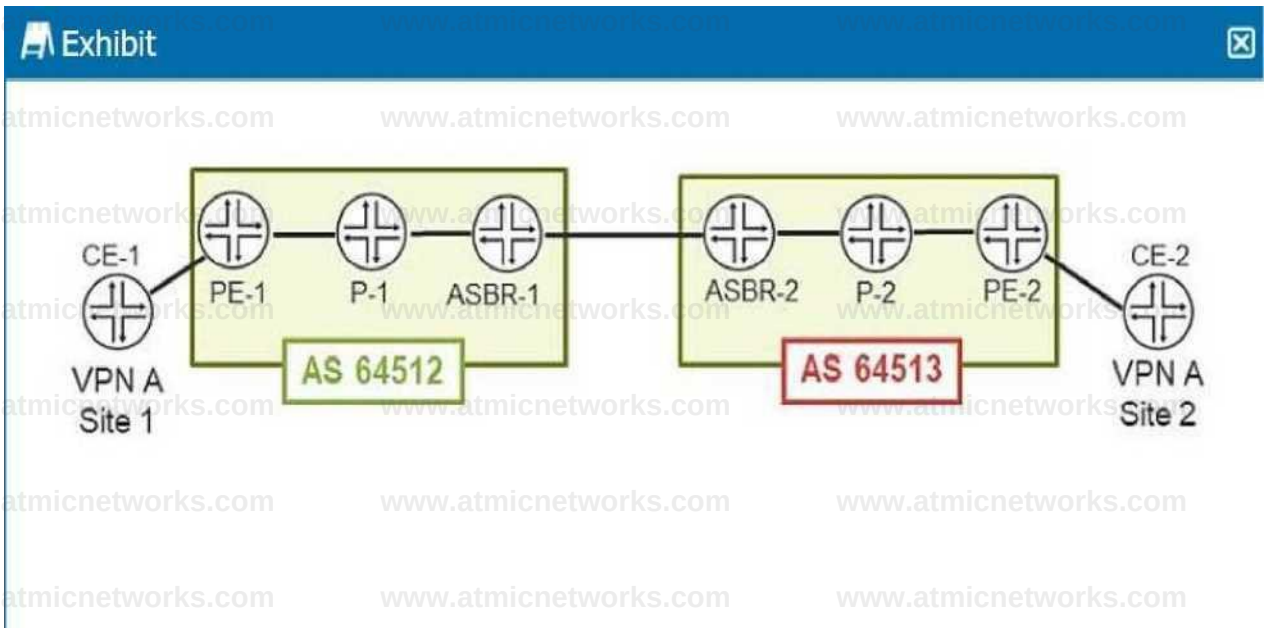

Click the Exhibit hutton.

You are configuring an interprovider Option C Layer 3 VPN to connect two customer sites.

Referring to the exhibit, which three statements are correct? (Choose three.)

A. ASBR routers maintain the internal routes from its own AS and the loopback addresses from the other AS PEs.

B. PE routers maintain the internal routes from its own AS, the loopback address from the other AS PEs, and the L3VPN routes.

C. P routers only maintain the internal routes from their own AS.

D. P routers maintain the internal routes from its own AS and the loopback address from the other AS PEs.

E. ASBR routers maintain the internal routes from its own AS, the loopback address from the other AS PEs, and

www.atmicnetworks.com

the L3VPN routes.

Answer: ABC

Explanation:

Interprovider Option C for Layer 3 VPNs involves the use of Autonomous System Boundary Routers (ASBRs) to exchange labeled VPN-IPv4 routes between different Autonomous Systems (AS). This option requires BGP sessions between ASBRs, and the VPN routes are carried end-to-end using MPLS labels. Here's a detailed analysis of the roles of different routers in this scenario:

1. **ASBR Routers**:

- ASBRs are responsible for exchanging VPN-IPv4 routes between different ASes.

- **A. ASBR routers maintain the internal routes from its own AS and the loopback addresses from the other AS PEs.**

- Correct. ASBRs maintain routes to internal destinations within their own AS, and they also need to know the loopback addresses of PEs in the other AS to set up the BGP sessions and MPLS tunnels.

2. §PE Routers**:

- PE routers are responsible for maintaining VPN routes and label information to forward VPN traffic correctly.

- **B. PE routers maintain the internal routes from its own AS, the loopback address from the other AS PEs, and the L3VPN routes.**

- Correct. PE routers need to maintain:

- Internal routes within their AS for routing.

- Loopback addresses of other AS PEs for establishing MPLS LSPs.

- L3VPN routes to provide end-to-end VPN connectivity.

3. **P Routers**:

- P routers are the core routers that do not participate in BGP VPN routing but forward labeled packets based on MPLS labels.

§Conclusion**:

www.atmicnetworks.com

- **C. P routers only maintain the internal routes from their own AS.**

- Correct. P routers maintain the internal routing information to forward packets within the AS and use MPLS labels for forwarding VPN packets. They do not maintain VPN routes or routes from other ASes.

4. **Incorrect Statements**:

- **D. P routers maintain the internal routes from its own AS and the loopback address from the other AS PEs.**

- Incorrect. P routers do not need to maintain the loopback addresses of other AS PEs. They only maintain internal routing and MPLS label information.

- **E. ASBR routers maintain the internal routes from its own AS, the loopback address from the other AS PEs, and the L3VPN routes.**

- Incorrect. ASBR routers do not maintain L3VPN routes. They exchange labeled VPN-IPv4 routes with other ASBRs and forward them to PE routers.

The correct answers are:

* *A. ASBR routers maintain the internal routes from its own AS and the loopback addresses from the other AS PEs.**

* *B. PE routers maintain the internal routes from its own AS, the loopback address from the other AS PEs, and the L3VPN routes.**

* *C. P routers only maintain the internal routes from their own AS.**

Reference:

* Juniper Networks Documentation on Interprovider VPNs: [Interprovider VPN Configuration](https://www.juniper.net/documentation/en_US/junos/topics/topic-map/mpls-vpn-interprovider.html)

* MPLS and VPN Architectures, CCIP Edition by Ivan Pepelnjak and Jim Guichard

Question: 88

* ou are configuring a Layer 3 VPN between two sites. You are configuring the vrf-target target : 65100:100

www.atmicnetworks.com

statement in your routing instance.

In this scenario, which two statements describe the vrf-target configuration? (Choose two.)

A. This value is used to identify BGP routes learned from the local CE device.

B. This value is used to identify BGP routes learned from the remote PE device.

C. This value is used to add a target community to BGP routes advertised to the local CE device.

D. This value is used to add a target community to BGP routes advertised to the remote PE device.

Answer: BD

Explanation:

The `vrf-target` statement in a Layer 3 VPN configuration is used to control the import and export of VPN routes by attaching a target community to the routes. This helps in defining which VPN routes should be imported into or exported from a particular VRF (Virtual Routing and Forwarding) instance.

1. **Understanding VRF Target**:

- The `vrf-target` statement specifies the extended community attributes (route targets) that are used to control the import and export of routes in a VRF.

- These attributes help in identifying which routes should be shared between different VRFs, particularly across different PE (Provider Edge) devices.

2. **Statements Analysis**:

- **A. This value is used to identify BGP routes learned from the local CE device.**

- Incorrect. The `vrf-target` attribute is not used to identify routes learned from the local CE device. It is used to manage routes between PE devices and within the provider's MPLS network.

- **B. This value is used to identify BGP routes learned from the remote PE device.**

- Correct. The `vrf-target` value helps in identifying which routes from remote PE devices should be imported into the local VRF. It essentially acts as a filter for importing BGP routes with matching target communities.

- **C. This value is used to add a target community to BGP routes advertised to the local CE device.**

- Incorrect. Routes advertised to the local CE device do not use the `vrf-target` attribute. Instead, these routes are typically managed within the local VRF routing table.

- **D. This value is used to add a target community to BGP routes advertised to the remote PE device.**

- Correct. When advertising routes from the local PE to remote PE devices, the `vrf-target` value is added to these routes. This target community ensures that the correct routes are shared across the VPN.

Conclusion:

The correct statements about the `vrf-target` configuration in a Layer 3 VPN scenario are:

* *B. This value is used to identify BGP routes learned from the remote PE device.**

* *D. This value is used to add a target community to BGP routes advertised to the remote PE device.**

Reference:

* Juniper Networks Documentation on VRF Target: [VRF Target Configuration](https://www.juniper.net/documentation/en_US/junos/topics/topic-map/layer-3-vpns.html)
* MPLS and VPN Architectures by Ivan Pepelnjak and Jim Guichard

Question: 89

You must alter class-of-service values in packets on the outbound interface of an edge router.

In this scenario, which CoS component allows you to accomplish this task?

www.atmicnetworks.com

A. output policer

B. scheduler

C. rewrite rules

D. forwarding classes

Answer: C

Explanation:

Class of Service (CoS) in networking is used to manage traffic by classifying, scheduling, and sometimes modifying packets to ensure network performance and Quality of Service (QoS). Different CoS components are used to achieve these goals. Let's analyze each option to determine which CoS component allows you to alter class-of-service values on the outbound interface of an edge router.

1. **Output Policer**:

- Policing is used to control the rate of traffic sent to or from a network interface. It can drop or remark traffic that exceeds a certain rate.

- Policing is not typically used to alter CoS values but to enforce traffic limits.

2. **Scheduler**:

- A scheduler is responsible for managing the order in which packets are transmitted out of an interface based on their CoS markings. It can allocate bandwidth and prioritize traffic.

- The scheduler manages how packets are queued and sent but does not alter the CoS values of packets.

3. **Rewrite Rules**:

- Rewrite rules are used to modify the CoS values of packets, such as DSCP (Differentiated Services Code Point) or 802.1p bits, as they exit an interface.

www.atmicnetworks.com

- Rewrite rules can alter the class-of-service values in the packet headers to match the desired policies of the outbound interface.
- Therefore, rewrite rules are the correct component for altering CoS values on an outbound interface.

4. **Forwarding Classes**:
   - Forwarding classes are used to categorize packets into different traffic classes within a router for QoS handling.
   - They help in defining how packets should be treated by the scheduler but do not directly modify the CoS values.

Conclusion:

To alter class-of-service values in packets on the outbound interface of an edge router, the correct CoS component to use is:

C. rewrite rules

Reference:

- Juniper Networks Documentation on CoS: [Class of Service Overview](https://www.juniper.net/documentation/en_US/junos/topics/concept/class-of-service-overview.html)
- Junos OS CoS Configuration Guide: [Rewrite Rules](https://www.juniper.net/documentation/en_US/junos/topics/topic-map/class-of-service- rewrite-rules.html)

Question: 90

Refer to the exhibit.

```
root@RI> show ospf interface extensive
interface            State   Area            DR ID           BDR ID          Nbrs
```

www.atmicnetworks.com

```
et-0/0/33.0          DR      0.0.0*0          192.168.252.0    0.0.0.0          0
  Type: LAN, Address: 192.168.254*0, Mask: 255.255.255.254, MTU: 9202, Cost: 1
  DR addr: 192.168.254.0, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
root@R4> show ospf interface extensive
Interface            State Area             DR ID            BDR ID           Nbrs
et-0/0/48.0          Waiting 0.0.0.0        0.0.0.0          0.0.0.0          0
  Type: LAN, Address: 192.168.254.1, Mask: 255.255.255.254, MTU: 9202, Cost: 1
  Priority: 128
    Adj count: 0
  Hello: 5, Dead: 20, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost; 1
root@R2> show ospf interface et-0/0/33.0 extensive
Interface            State    Area          DR ID            BDR ID           Nbrs
et-0/0/33.0          BDR      0.0^0.0       192.168.253.0    192.168.252.1    1
    Type: LAN, Address: 192.168,254.8, Mask: 255.255.255.254, MTU: 9202, Cost: 1
    DR addr: 192.168.254.9, BDR addr: 192.168.254.8, Priority: 128
    Adj count: 1
    Hello: 10, Dead: 40, ReXmit: 5, Not Stub
    Auth type: None
    Protection type: None
    Topology default (ID 0) -> Cost: 1
```

Click the Exhibit button.

You have an OSPF environment. You have recently added a router called R4 that is directly connected to R1 and R2. You discover that R4 is only peering with R2.

Referring to the exhibit, how would you correct the peering?

A. Adjust the Priority on R1 to be lower than the Priority on R4.

B. Change the MTU size on R1 and R2 to be 22 bytes higher than R4's MTU size.

C. Adjust the Dead Interval on R4 to match the Dead Interval on R1 and R2.

D. Adjust the Hello Interval on R1 and R2 to match the Hello Interval on R4.

Answer: D

Explanation:

www.atmicnetworks.com

Question: 91

A router running IS-IS is configured with an ISO address of 49.0001.00a0.c96b.c490.00.

Which part of this address is the system ID?

A. 00a0.c96b.c490 is the system identifier.

B. 0001.00a0.c96b.c490 is the system identifier.

C. c96b.c490 is the system identifier.

D. c490 is the system identifier.

Answer: A

Explanation:

In IS-IS (Intermediate System to Intermediate System) routing, each router is identified by a unique ISO (International Organization for Standardization) address, also known as a Network Entity Title (NET). The NET consists of three parts:

1. **Area Identifier**: Indicates the area to which the router belongs.
2. **System Identifier**: Uniquely identifies the router within the area.
3. **NSAP Selector (NSEL)**: Typically set to 00 for a router, indicating the Network Service Access Point.

The format of the ISO address is `49.XXXX.YYYY.YYYY.ZZZZ.ZZZZ.00`, where:

www.atmicnetworks.com

- `49` is the AFI (Authority and Format Identifier) indicating a private address.
- `XXXX` is the Area Identifier.
- `YYYY.YYYY.YYYY` is the System Identifier.
- `ZZZZ.ZZZZ` is the NSAP Selector.

Given the address `49.0001.00a0.c96b.c490.00`:

- **Area Identifier**: `49.0001`
- **System Identifier**: `00a0.c96b.c490`
- **NSAP Selector**: `00`

Explanation:

- **A. 00a0.c96b.c490 is the system identifier**:
- Correct. The System Identifier in an ISO address is a 48-bit (6-byte) field used to uniquely identify the router. In this address, `00a0.c96b.c490` is the correct 6-byte System Identifier.

- **B. 0001.00a0.c96b.c490 is the system identifier**:
- Incorrect. This includes the Area Identifier as part of the System Identifier, which is not correct.

- **C. c96b.c490 is the system identifier**:
- Incorrect. This is only part of the System Identifier. The full System Identifier must be 6 bytes long.
- **D. c490 is the system identifier**:
- Incorrect. This is an incomplete and incorrect part of the System Identifier.

Conclusion:

The correct part of the address that represents the System Identifier is:

- *A. 00a0.c96b.c490 is the system identifier.**

www.atmicnetworks.com

Reference:

- Juniper Networks Documentation on IS-IS: [IS-IS Configuration](https://www.juniper.net/documentation/en_US/junos/topics/task/configuration/isis-configuring.html)
- ISO/IEC 10589, the IS-IS routing protocol standard.

Question: 92

- ou are using a Layer 3 VPN to connect two customer sites. The VPN routes for the customer networks appear as hidden in the bgp. 13vpn. o routing table on the PE routers.

What is causing this problem?

A. The routes use overlapping IP addresses.

B. There is not an established MPLS LSP between the two PE routers.

C. There is a routing loop in the service provider backbone.

D. Route targets are not configured.

Answer: B

Explanation:

For a Layer 3 VPN to function correctly, an MPLS Label Switched Path (LSP) must be established between the Provider Edge (PE) routers. The MPLS LSP is necessary for the transport of VPN traffic across the service provider's backbone network. If the MPLS LSP is not established, the PE routers cannot forward the VPN traffic properly, causing the routes to be hidden in the BGP routing table.

Here's a breakdown of why the other options are less likely:

A . The routes use overlapping IP addresses.

Overlapping IP addresses might cause issues with route advertisement and selection, but they would not typically cause routes to be hidden in the bgp.l3vpn.0 table.

C . There is a routing loop in the service provider backbone.

While routing loops are problematic, they would not specifically cause the routes to be hidden in the bgp.l3vpn.0 table. Routing loops would more likely result in dropped packets or increased latency.

D . Route targets are not configured.

Incorrect or missing route target configuration would prevent routes from being imported into the correct VRF, but it would not usually result in the routes being hidden. Instead, they would simply not appear in the relevant VRF.

Thus, the absence of an established MPLS LSP is the most plausible cause for the routes being hidden.

Question: 93

Refer to the exhibit.

```
user0Rl> show configuration class-of-service
interfaces {
    ge-0/0/0 { scheduler-map customer-traffic;
    1
]
scheduler-maps {
    custemer traffic (
        forwarding-class best-effort scheduler best-effort-scheduler;
        forwarding-class expedited-forwarding scheduler priority-scheduler |
}
schedulers {
    best-effort-scheduler { transmit-rate percent 40; priority low;
    priority-scheduler { transmit-rate percent 30.; priority high;
}
```

Click the Exhibit button.

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

A. Incoming traffic will not be classified because no classifier exists in the configuration.

B. The best-effort queue can transmit more than 40% of the total bandwidth on the ge-0/0/0 interface, if no other queue is using that bandwidth.

C. Incoming traffic will be classified using the default classifier.

D. The best-effort queue can never transmit more than 40% of the total bandwidth on the ge-0/0/0 interface, even if that bandwidth is available.

Answer: BC

Explanation:

https://www.juniper.net/documentation/us/en/software/junos/traffic-mgmt-qfx/topics/concept/cos-qfx-series-default-scheduling-classification-understanding.html

Question: 94

Refer to the Exhibit:

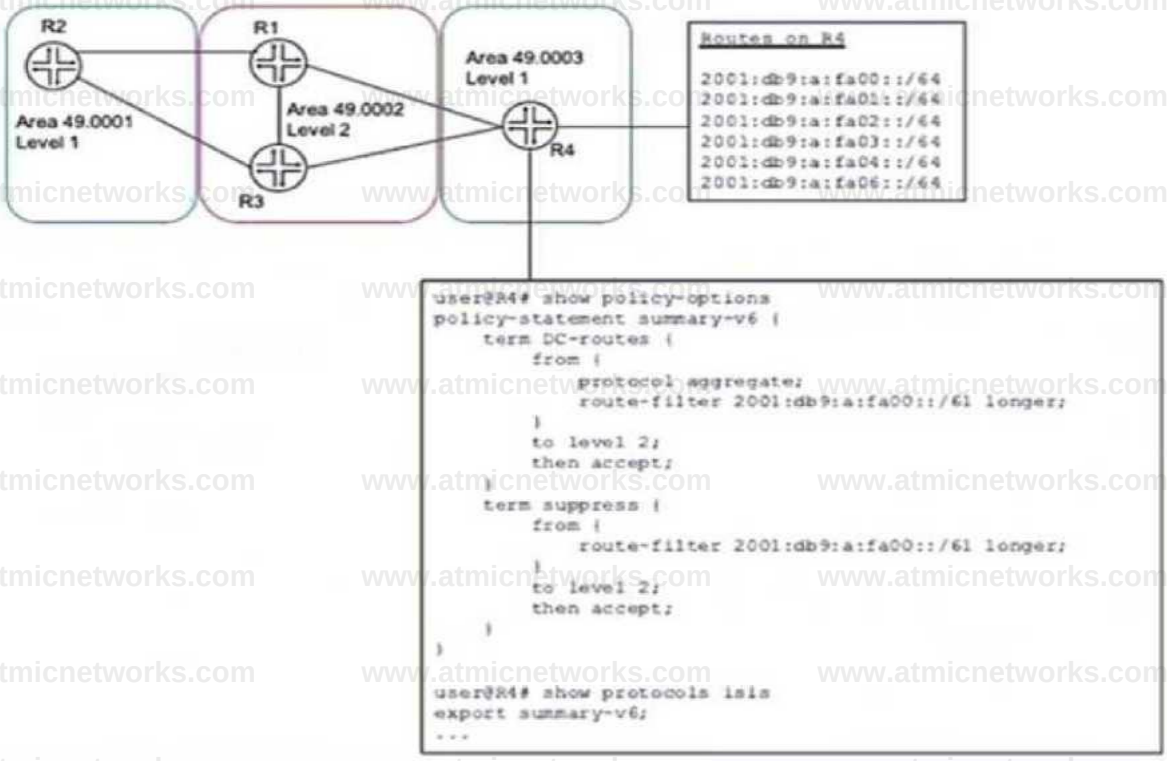

A network designer would like to advertise a single summary route from R4 to IS-IS level 2 neighbors as shown in the exhibit, but the configuration is not working.
Which three configuration changes will accomplish this task? (Choose three.)

A. delete protocols isis export summary-v6

B. set protocols isis import summary-v6

C. delete policy-options policy-statement summary-v6 term DC-routes from route-filter 2001:db5:a:fa00::/61 longer

D. set policy-options policy-statement summary-v6 term DC-routes from route-filter 2001:dbS:a:fa00::/6l exact

E. set policy-options policy-statement summary-v6 term suppress then reject

Answer: CDE

Explanation:

Question: 95

You are configuring schedulers to define the class-of-service properties of output queues. You want to control packet drops during periods of congestion.

In this scenario, which CoS configuration parameter would be used to accomplish this task?

A. buffer size

B. priority

C. shaping rate

D. drop profile

Answer: D

Explanation:

When configuring Class of Service (CoS) properties for output queues, we need to manage packet

drops during periods of congestion. Juniper's CoS framework provides several tools to manage congestion, including drop profiles, buffer sizes, and scheduling mechanisms. Let's break down each option and identify the correct one.

Evaluating the Answer Choices

Q D. drop profile (Correct Answer)

Why?

A drop profile defines when packets should be dropped based on the queue fill level.

Random Early Detection (RED) or Tail Drop can be used to manage congestion by discarding lower- priority packets first.

Drop profiles are configured under the scheduler to determine how aggressive packet dropping should be during congestion.

Example Juniper Configuration:

```
schedulers {
    best-effort {
        drop-profile low-drop;
    }
}

drop-profiles {
    low-drop {
        fill-level 80 drop-probability 50;
    }
}
```

fill-level 80 → When the queue reaches 80% full, packet drops begin.

drop-probability 50 → There is a 50% chance of dropping packets once the threshold is reached.

Official Juniper Documentation Reference:
Junos Class of Service Configuration Guide

www.atmicnetworks.com

"A drop profile determines how packets are discarded based on the queue fill level, allowing control over congestion behavior."

Why the Other Options Are Incorrect?

X A. buffer size (Incorrect)

Why?

The buffer size determines how many packets the queue can store before congestion occurs.

A larger buffer can delay drops, but it does not actively control dropping behavior.

It affects latency rather than controlling packet drops.

X B. priority (Incorrect)

Why?

Priority controls which queue gets serviced first, not how drops are handled.

Higher priority queues are serviced before lower-priority queues, but this does not prevent congestion-related drops.

X C. shaping rate (Incorrect)

Why?

Shaping limits the maximum transmission rate of the queue.

While shaping helps reduce congestion, it does not control which packets get dropped during congestion.

Shaping is useful for traffic smoothing, but it does not actively drop packets based on queue fill levels.

Final Answer: Q D. drop profile

Explanation:

Controls packet drops based on queue congestion.

Defines RED (Random Early Detection) or Tail Drop mechanisms.

Directly influences drop probability as the queue fills up.

Official Juniper Reference:

"Drop profiles are used to manage congestion by determining when and how aggressively packets are dropped based on queue fill level."

Question: 96

Click the Exhibit button.

```
[edit routing-instances VPS-A] user?PEC show instance-type vrf; routing-options ( static f
        route 10.1.0.0/16 next-hop 10.1.0.1:
    I
J
Interface ge-0/0/2.0: route-distinguisher 172.1?.20.1:1; vtf-export vpn-a-export;
vrf-target target:<S512:-: [edit policy-options policy-statement vpn-a-export J userJFEs show term add-RTs I then (
conssunity add vpn-a-target: conounity add vpn-n-target; accept;
} [edit policy-options] userJPE: show natch com

cotoDunity vpn-a-target members target:65512:1; coewunity vpn-m-target members target:65512:2;
```

Referring to the exhibit, which statement is correct?

A. VPN routes are exported with the target:65512:1 and target:65512:2 route targets.

B. You cannot use the vrf-target and vrf-export statements in the same VRF.

C. VPN routes with the target:65512:1 and target:65512:2 route targets are imported.

D. VPN routes are exported with only the target:65512:1 route target

Answer: A

Explanation:

The exhibit shows the configuration of a VRF (Virtual Routing and Forwarding) instance on a Juniper PE router.

Let's break down the key components:

VRF Configuration (VPN-A)

The instance type is VRF, meaning this is an L3VPN (Layer 3 VPN).

The routing instance contains a static route (10.1.0.0/16 next-hop 10.1.0.1).

The interface ge-0/0/2.0 is assigned to the VRF.

Route Distinguisher (RD): 172.17.20.1:1

VRF-Export Policy: vpn-a-export

www.atmicnetworks.com

VRF-Target: target:65512:1 (This defines which routes will be imported into the VRF).

VRF Export Policy (vpn-a-export)

The vpn-a-export policy adds two BGP communities (route targets) to exported VPN routes:

community add vpn-a-target;

community add vpn-m-target;

accept;

The vpn-a-target community corresponds to target:65512:1.

The vpn-m-target community corresponds to target:65512:2.

Policy-Options (Community Definitions)

community vpn-a-target members target:65512:1;

community vpn-m-target members target:65512:2;

This confirms that routes exported from this VRF will have BOTH target:65512:1 and target:65512:2.

Evaluating the Answer Choices

Q Option A: "VPN routes are exported with the target:65512:1 and target:65512:2 route targets."

The vpn-a-export policy explicitly adds both vpn-a-target (65512:1) and vpn-m-target (65512:2) to exported routes.

This is correct. Q

X Option B: "You cannot use the vrf-target and vrf-export statements in the same VRF."

This is incorrect.

Juniper allows the use of both vrf-target and vrf-export in the same VRF:

vrf-target is used for importing routes.

vrf-export defines export policies (which can add additional route targets).

This is incorrect. X

X Option C: "VPN routes with the target:65512:1 and target:65512:2 route targets are imported."

The vrf-target target:65512:1; statement only controls importing routes.

The import policy does not include target:65512:2, so routes tagged with target:65512:2 alone would not be imported into this VRF.

This is incorrect. X

X Option D: "VPN routes are exported with only the target:65512:1 route target."

The export policy (vpn-a-export) clearly adds both 65512:1 and 65512:2.

This is incorrect. X

Final Answer:

Q A. VPN routes are exported with the target:65512:1 and target:65512:2 route targets.

Verification from Juniper Documentation

Juniper MPLS L3VPN Configuration Guide confirms that vrf-target is used for importing, while vrf- export can be used for exporting multiple route targets.

Juniper Routing Policy Documentation states that export policies can add multiple BGP communities

(route targets).

RFC 4364 (BGP/MPLS IP VPNs) defines the use of route targets for VPN route control.