complete your programming course about resources, doubts and more!

Cisco

(350-501)

Implementing and Operating Cisco Service Provider Network Core Technologies (SPCOR)

Total: **391 Questions**

Link:

Question: 1

DRAG DROP -
Drag and drop the OSs from the left onto the correct descriptions on the right.
Select and Place:

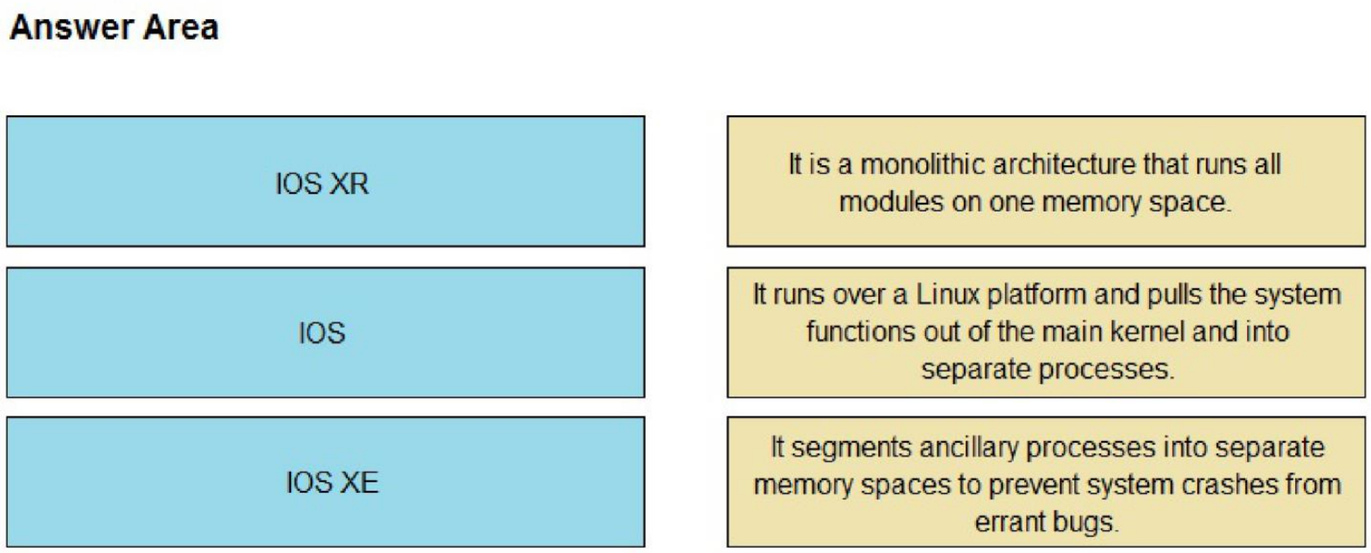

Answer Area

IOS XR	It is a monolithic architecture that runs all modules on one memory space.
IOS	It runs over a Linux platform and pulls the system functions out of the main kernel and into separate processes.
IOS XE	It segments ancillary processes into separate memory spaces to prevent system crashes from errant bugs.

Answer:

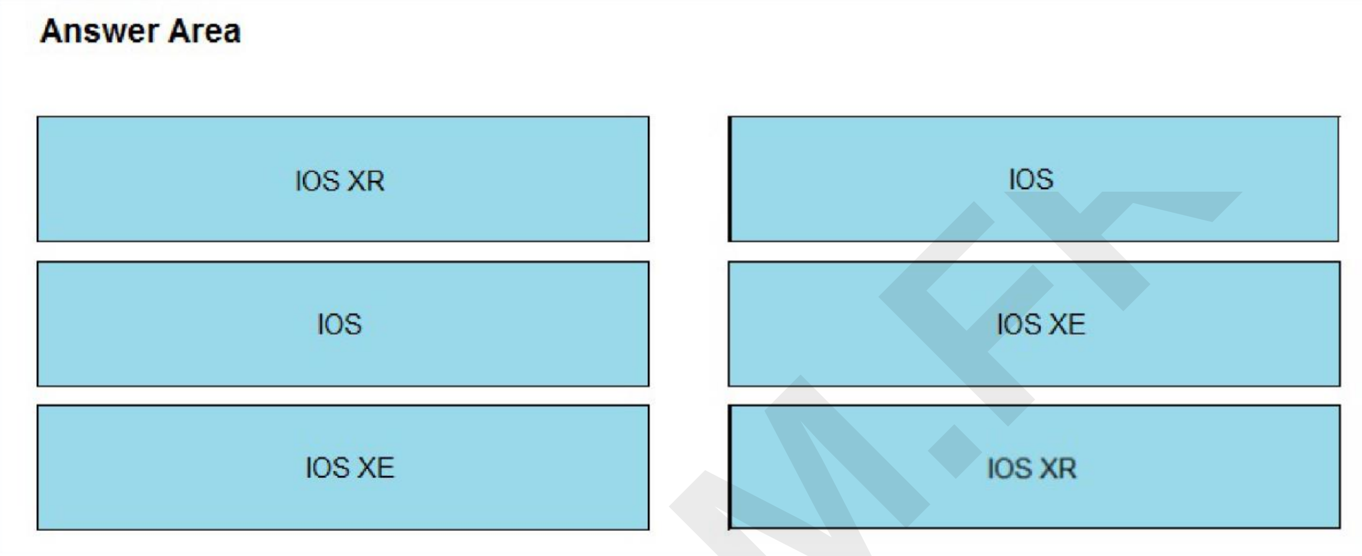

Answer Area

IOS XR	IOS
IOS	IOS XE
IOS XE	IOS XR

Explanation:

Reference:

https://specialties.bayt.com/en/specialties/q/276369/what-is-the-key-difference-between-ios-ios-xe-and-ios-xr-for-cisco-devices/

Question: 2

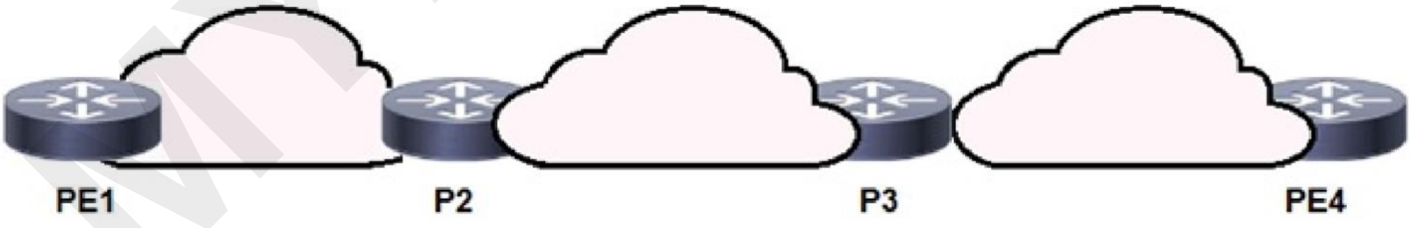

Refer to the exhibit. P3 and PE4 are at the edge of the service provider core and serve as ABR routers.
Aggregation areas are on either side of the core.
Which statement about the architecture is true?

A. To support seamless MPLS, the BGP route reflector feature must be disabled.
B. If each area is running its own IGP, BGP must provide an end-to-end MPLS LSP.
C. If each area is running its own IGP, the ABR routers must redistribute the IGP routing table into BGP.
D. To support seamless MPLS, TDP must be used as the label protocol.

Answer: B

Explanation:

If each area is running its own IGP, BGP must provide an end-to-end MPLS LSP.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9600/software/release/16-12/configuration_guide/mpls/b_1612_mpls_9600_cg/ configuring_seamless_mpls.html

Question: 3

Which component is similar to an EVPN instance?

A. router distinguisher
B. MPLS label
C. IGP router ID
D. VRF

Answer: D

Explanation:

The correct answer is D, VRF (Virtual Routing and Forwarding instance). Here's why:

An EVPN (Ethernet VPN) instance, at its core, provides Layer 2 VPN services over a Layer 3 infrastructure. It's designed to isolate traffic belonging to different customers or service offerings, much like a traditional VLAN, but with enhanced scalability and flexibility. A key requirement for isolating these separate traffic flows is maintaining unique routing and forwarding tables.

A VRF fulfills this exact need. A VRF is a logical separation of routing tables within a router, allowing multiple isolated routing domains to coexist on the same physical device. Each VRF has its own routing table, forwarding table, and set of interfaces, ensuring that traffic within one VRF remains separate from traffic within any other VRF. This provides the necessary isolation that EVPN also offers.

In the context of EVPN, different EVPN instances effectively correspond to different VRFs on the provider edge (PE) routers. When an EVPN route is received for a particular customer or service, it gets associated with the correct VRF, ensuring that traffic is routed accordingly.

Let's look at why the other options are incorrect:

A. Router Distinguisher (RD): An RD is a unique identifier used within an MPLS/BGP VPN to distinguish routes from different VPNs. While essential for EVPN, it does not provide the traffic isolation that an EVPN instance represents. The RD is used within a VRF to make routes specific to that VPN.
B. MPLS Label: MPLS labels are used for forwarding packets across an MPLS network. They are necessary

for EVPN's data plane, but they are transport mechanisms, not analogous to a complete isolation domain like an EVPN instance.

C. IGP Router ID: An IGP (Interior Gateway Protocol) router ID is used to identify a router within a routing protocol domain, and is unrelated to traffic isolation in the context of EVPN.

Therefore, VRF provides the most accurate comparison to an EVPN instance because both function to segment routing and forwarding traffic into separate, isolated logical entities.

Authoritative Resources for Further Research:

1. **Cisco's EVPN documentation:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mpls/configuration/16-9/mp-16-9-book/mp-evpn.html
2. **Juniper's EVPN documentation:**https://www.juniper.net/documentation/en_US/junos/topics/topic-map/evpn-overview.html
3. **VRF Concepts:**https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13709-vrf-lite.html

Question: 4

Why do Cisco MPLS TE tunnels require a link-state routing protocol?

A. The link-state database provides segmentation by area, which improves the path-selection process.

B. The link-state database provides a data repository from which the tunnel endpoints can dynamically select a source ID.

C. Link-state routing protocols use SPF calculations that the tunnel endpoints leverage to implement the tunnel.

D. The tunnel endpoints use the link-state database to evaluate the entire topology and determine the best path.

Answer: D

Explanation:

Here's a detailed justification for why Cisco MPLS TE tunnels rely on link-state routing protocols, focusing on why option D is the correct answer:

Option D correctly states that the tunnel endpoints utilize the link-state database to evaluate the entire network topology and determine the optimal path for traffic engineering. MPLS Traffic Engineering (TE) tunnels are explicitly routed, meaning their path isn't solely based on destination IP like conventional routing.

Instead, they are crafted to meet specific bandwidth, latency, or other quality-of-service requirements. To achieve this, the tunnel endpoints (specifically the ingress router) need a complete picture of the network's topology, including link bandwidth, delay, and administrative metrics. This comprehensive view is provided by the link-state database.

Link-state routing protocols like OSPF and IS-IS maintain a synchronized, consistent view of the network in each router's database. This database is a map of the entire network, which is then used by the traffic engineering calculation algorithm, specifically the Constraint-based Shortest Path First (CSPF). CSPF evaluates various metrics (beyond just shortest path) provided by the link-state database to select an appropriate explicit path for the tunnel that satisfies the specified constraints. Without this detailed topology information, the ingress router cannot dynamically calculate an efficient and compliant path for the TE tunnel. Therefore, option D is the most accurate explanation for the dependency.

Option A is incorrect because while areas in link-state protocols provide segmentation and manage routing domain size, TE tunnels calculate end-to-end paths which are not restricted by areas, utilizing the entire network link-state information. Option B is incorrect because source ID is not dynamically chosen from the link-state database; it is usually a loopback interface IP address, statically assigned to the router. Option C is incorrect because while the link-state protocol's SPF calculation is used by the underlying IGP to forward traffic, it is not directly used by TE tunnel endpoints, which rely on CSPF based on the extended information contained within the link-state advertisements.

Authoritative Links for further research:

Cisco Documentation on MPLS Traffic Engineering:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mpls/configuration/15-mt/mpls-15-mt-book/mp-te-oview.html
RFC 3209 - RSVP-TE: Extensions to RSVP for LSP Tunnels:https://datatracker.ietf.org/doc/html/rfc3209
Understanding Link State Routing Protocols:https://www.geeksforgeeks.org/link-state-routing-protocol/

Question: 5

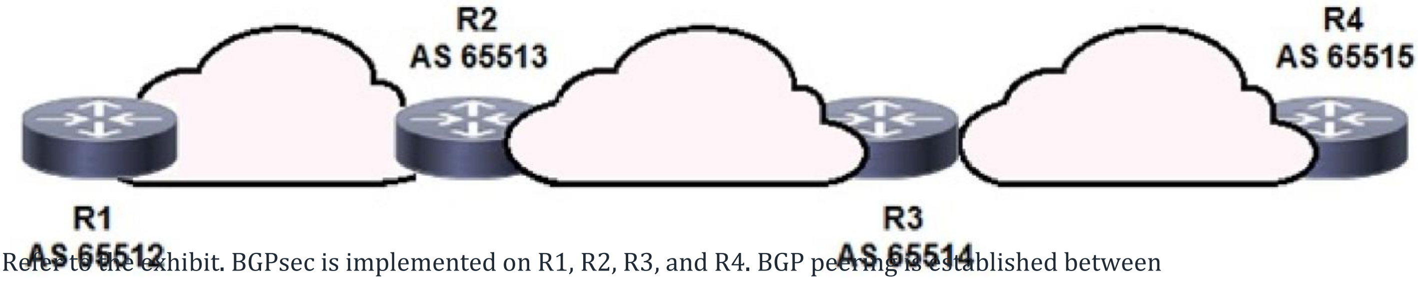

Refer to the exhibit. BGPsec is implemented on R1, R2, R3, and R4. BGP peering is established between neighboring autonomous systems.
Which statement about implementation is true?

A. BGP updates from the iBGP peers are appended with a community of local-as.

B. BGP updates from the all BGP peers are appended with a community of no-export.

C. BGP updates from the eBGP peers are appended with an additional AS path value that is statically set by the domain administrator.

D. BGP updates from the eBGP peers are appended with a BGPsec attribute sequence that includes a public key hash and digital signature.

Answer: D

Explanation:

BGP updates from the eBGP peers are appended with a BGPsec attribute sequence that includes a public key hash and digital signature.

Question: 6

You are configuring MPLS traffic-engineering tunnels in the core. Which two ways exist for the tunnel path across the core? (Choose two.)

A. The dynamic path option is supported only with IS-IS.

B. Tunnels can be configured with dynamic path or explicitly defined path.

C. A zero bandwidth tunnel is not a valid option.

D. The bandwidth statement creates a hard reservation on the link.

E. Tunnel links inherit IGP metrics by default unless overridden.

Answer: BE

Explanation:

Here's the justification for the answer options B and E being correct, and why A, C, and D are incorrect:

Justification for B (Tunnels can be configured with dynamic path or explicitly defined path):

MPLS traffic engineering (MPLS-TE) tunnels offer flexibility in path selection. You can either allow the routing protocol (like OSPF or IS-IS) to dynamically determine the best path based on network conditions, or you can explicitly specify the exact path the tunnel must take, often using strict or loose explicit routing. This provides granular control over traffic flow and allows for optimized resource utilization. The ability to choose between dynamic and explicitly defined paths is a fundamental characteristic of MPLS-TE.

Justification for E (Tunnel links inherit IGP metrics by default unless overridden):

When an MPLS-TE tunnel is established, it initially uses the underlying Interior Gateway Protocol (IGP) metrics of the traversed links to calculate its own path cost. This means that the tunnel's path will likely align with the IGP's shortest path unless specifically modified. However, this behavior can be altered, and you can override the IGP metrics with your own defined metrics. This is useful in situations where traffic needs to be routed in a specific manner that is not the default IGP best path.

Why A is incorrect (The dynamic path option is supported only with IS-IS):

While IS-IS is commonly used with MPLS-TE, the dynamic path option is not exclusively tied to IS-IS. OSPF can also be used for dynamic path computation in MPLS-TE. The choice depends on the network's existing IGP and desired flexibility. This makes option A inaccurate.

Why C is incorrect (A zero bandwidth tunnel is not a valid option):

A zero bandwidth tunnel is indeed a valid configuration. It can be used for various purposes, such as setting up tunnels that are only activated upon certain conditions, for test purposes, or for providing backup paths that should only be used in case of failure of the main tunnel path. It's a useful tool for managing resources and optimizing network utilization.

Why D is incorrect (The bandwidth statement creates a hard reservation on the link):

The bandwidth statement in MPLS-TE tunnels does not create a "hard" reservation. It signals the desired bandwidth requirements to the network, allowing for admission control decisions based on the available bandwidth on links. While it attempts to enforce bandwidth availability, if the link cannot provide that bandwidth, the tunnel may still be established if that’s allowed or may fail setup. This mechanism is known as Constraint Based Routing, or CBR. The bandwidth parameter helps the network to make decisions but does not constitute a guaranteed fixed reservation.

Authoritative Links for Further Research:

Cisco MPLS Traffic Engineering Overview:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te/configuration/15-sy/mp-te-15-sy-book/mp-te-overview.html
Cisco MPLS Configuration Guide:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mpls/configuration/15-mt/mpls-15-mt-book/mp-mpls-cfg.html
Constraint-Based Routing:https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/cbr.html

In conclusion, the ability to define tunnel paths dynamically or explicitly, and the default inheritance of IGP metrics, are core aspects of MPLS-TE. These mechanisms offer flexibility and allow the network administrator to tailor the behavior of traffic engineered tunnels.

Question: 7

Which configuration mode do you use to apply the mpls ldp graceful-restart command in IOS XE Software?

A. MPLS LDP neighbor
B. interface
C. MPLS
D. global

Answer: D

Explanation:

The mpls ldp graceful-restart command configures the graceful restart capability for Label Distribution Protocol (LDP) sessions. This is crucial for maintaining forwarding continuity during control plane outages in MPLS networks, a common technology in service provider environments. Specifically, LDP graceful restart allows an LDP router to maintain its forwarding state and traffic flow while its control plane is restarting or undergoing maintenance. The command, and similar LDP configurations, are applied at a network-wide scope because they affect the global behavior of LDP on a device. Therefore, the command is configured within the global configuration mode of Cisco IOS XE. Option A, MPLS LDP neighbor, configures settings specific to individual LDP adjacencies. Option B, interface, is for configuring interface-specific parameters. Option C, MPLS, is not the correct mode; while it's related to MPLS, it does not contain the command we are looking for, which is applicable to LDP specifically. Thus, option D, global configuration mode, is the appropriate configuration level for the given command.

Authoritative link:

Cisco Command Reference: mpls ldp graceful-restart: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mpls/command/mpls-cr-book/mpls-m2.html#wp1090257
Cisco Documentation: LDP Graceful Restart: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mpls/configuration/15-mt/mpls-15-mt-book/mp-ldp-grst.html

Question: 8

After you analyze your network environment, you decide to implement a full separation model for Internet access and MPLS L3VPN services.
For which reason do you make this decision?

A. It enables EGP and IGP to operate independently.
B. It enables you to choose whether to separate or centralize each individual service.
C. It is easier to manage a system in which services are mixed.
D. It requires only one edge router.

Answer: B

Explanation:

The correct answer is **B. It enables you to choose whether to separate or centralize each individual service.**

A full separation model, as applied to internet access and MPLS L3VPN services, offers flexibility in how you manage various network offerings. It doesn't mandate that all services must be separated; rather, it provides the option to separate or centralize them on a service-by-service basis. This flexibility is key. For instance, you might choose to separate internet access for security or performance reasons, while opting to centralize other services.

Option A, "It enables EGP and IGP to operate independently," is generally true in a separate network design. However, it is not the primary reason for choosing a full separation model. The key factor is the flexibility over service deployment. Option C, "It is easier to manage a system in which services are mixed," is incorrect. A mixed system often creates management complexity and potential conflicts. Option D, "It requires only one edge router," is also incorrect. Full separation usually necessitates multiple edge devices per service to enforce clear boundaries.

In the context of service provider networks, a full separation model allows for better resource allocation, security policy enforcement, and potentially more efficient use of network infrastructure for each individual service. This granular control improves network design and operational efficiency. This approach aligns with the principle of least privilege, as network resources can be more tightly constrained to the services that require them.

For further information, explore resources on network design principles and service provider architecture:

Cisco Design Zone:https://www.cisco.com/c/en/us/solutions/design-zone.html (While not specific to this question, it provides network design principles).

MPLS L3VPN:https://www.cisco.com/c/en/us/solutions/service-provider/mpls/index.html (for a deeper understanding of MPLS L3VPN concepts).

Network segmentation concepts: Numerous articles and whitepapers on general network segmentation are available online.

Question: 9

Which statement about the Cisco MPLS TE forwarding adjacency feature is true?

A. It enables the MPLS core to use EIGRP as the routing protocol.

B. It enables the Cisco MPLS TE tunnel to be advertised into the running IGP.

C. It enables the tailend router to advertise routes to the headend router over the tunnel.

D. It enables the headend and tailend routers to establish a bidirectional tunnel.

Answer: B

Explanation:

The correct answer is **B. It enables the Cisco MPLS TE tunnel to be advertised into the running IGP.**

MPLS Traffic Engineering (TE) tunnels, by default, are not advertised via Interior Gateway Protocols (IGPs) like OSPF or IS-IS. Instead, they are statically configured or established through a signaling protocol like RSVP-TE. The MPLS TE forwarding adjacency feature addresses this limitation. By enabling this feature, the TE tunnel endpoint (the tunnel headend router) advertises the tunnel as a logical point-to-point link into the IGP. This allows other routers within the IGP domain to consider the TE tunnel when calculating their shortest path towards destinations reachable via the tunnel. The IGP considers the TE tunnel as an additional link with its own associated cost metrics. This allows traffic to be forwarded across the TE tunnel via standard IGP routing, enabling traffic engineering based on bandwidth reservation or other TE constraints. It doesn't mean the core can use EIGRP, as stated in option A, nor does it mean the tailend advertises routes to the headend over the tunnel, as per option C, nor does it create a bidirectional tunnel as option D suggests. Bidirectional tunnels are separate but often related concept in MPLS TE. MPLS TE Forwarding Adjacency enhances IGP path calculation by incorporating TE tunnels as usable links.

For further research, refer to Cisco's documentation on MPLS Traffic Engineering and Forwarding Adjacency:Cisco MPLS Traffic Engineering Configuration GuideMPLS Traffic Engineering Overview

Question: 10

While implementing TTL security, you issue the PE(config-router-af)#neighbor 2.2.2.2 ttl-security hops 2 command.
After you issue this command, which BGP packets does the PE accept?

A. to 2.2.2.2, with a TTL of 2 or more

B. from 2.2.2.2, with a TTL of less than 2

C. to 2.2.2.2, with a TTL of less than 253

D. from 2.2.2.2, with a TTL of 253 or more

Answer: D

Explanation:

The command neighbor 2.2.2.2 ttl-security hops 2 configures TTL Security Check (also known as Generalized TTL Security Mechanism or GTSM) for the BGP peer 2.2.2.2. This mechanism is designed to protect against BGP session hijacking by ensuring that BGP packets arrive with an expected TTL value, indicating they haven't traversed excessive hops. The hops 2 parameter dictates that the expected TTL of incoming packets should be 255 minus the configured hop count (2). Therefore, the expected TTL will be 255 - 2 = 253. BGP packets originating from the peer 2.2.2.2, will have a TTL of 255 when they leave the origin router, and this is decremented by each hop they traverse to reach the PE. The PE router will therefore accept packets with a TTL of 253 because it will expect exactly 2 hops between it and the originating peer. Packets with a TTL less than 253 imply more hops, possibly from a malicious actor attempting to spoof a BGP session. Therefore, option D is correct, accepting BGP packets from 2.2.2.2 with a TTL of 253 or more; 253 is the expected value and any higher means the packet has traveled less hops, which is acceptable but anything less means it is an unexpectantly long path. Option A is incorrect because it specifies "to" which is outbound, not inbound. Option B specifies less than 2, which is wrong because it specifies a different meaning and a different range. Option C, indicates an incorrect range of TTL and an incorrect direction.

Further reading on TTL Security Check for BGP can be found at these authoritative links:

Cisco Documentation on BGP GTSM:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/15-mt/irg-15-mt-book/irg-bgp-ttl-security.html
Juniper Networks Knowledge Base on GTSM:https://support.juniper.net/support/s/article/KB14612-Configuring-TTL-Security-Check-on-JunOS
IETF RFC 5082 (The Generalized TTL Security Mechanism (GTSM)):https://www.rfc-editor.org/rfc/rfc5082

Question: 11

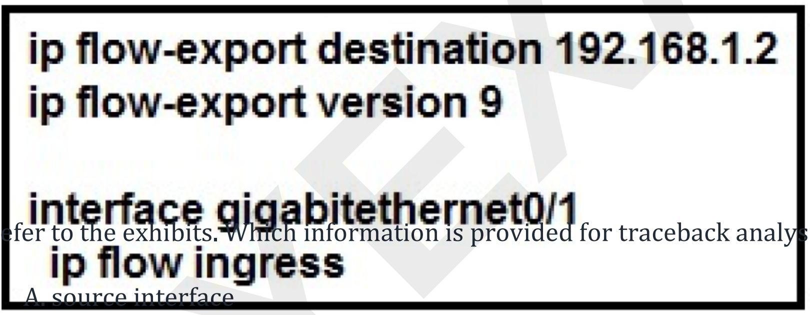

```
ip flow-export destination 192.168.1.2
ip flow-export version 9

interface gigabitethernet0/1
 ip flow ingress
```

Refer to the exhibits. Which information is provided for traceback analysis when this configuration is applied?

A. source interface

B. packet size distribution

C. IP sub flow cache

D. BGP version

Answer: A

Explanation:

Correct answer is A:source interface.

Question: 12

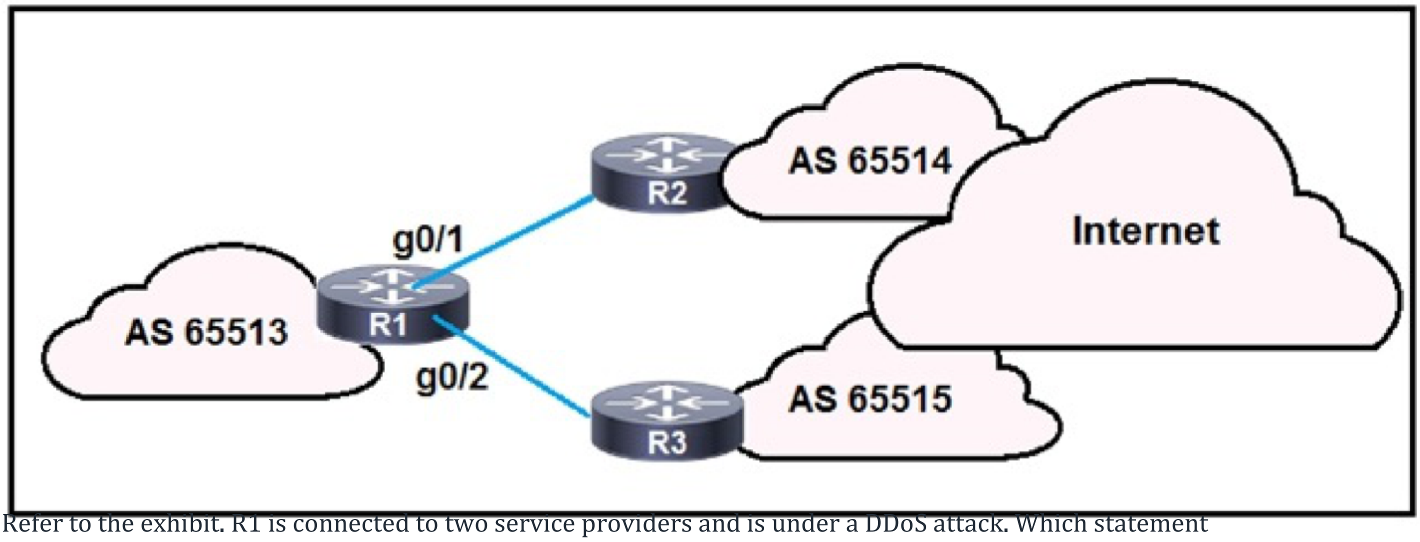

Refer to the exhibit. R1 is connected to two service providers and is under a DDoS attack. Which statement about this design is true if URPF in strict mode is configured on both interfaces?

A. R1 drops all traffic that ingresses either interface that has a FIB entry that exits a different interface. B. R1 drops destination addresses that are routed to a null interface on the router.

C. R1 permits asymmetric routing as long as the AS-PATH attribute entry matches the connected AS. D. R1 accepts source addresses on interface gigabitethernet0/1 that are private addresses.

Answer: A

Explanation:

R1 drops all traffic that ingresses either interface that has a FIB entry that exits a different interface.

Question: 13

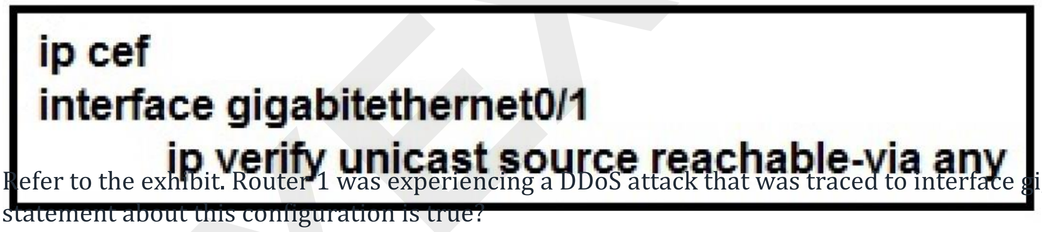

```
ip cef
interface gigabitethernet0/1
        ip verify unicast source reachable-via any
```

Refer to the exhibit. Router 1 was experiencing a DDoS attack that was traced to interface gigabitethernet0/1. Which statement about this configuration is true?

A. Router 1 accepts all traffic that ingresses and egresses interface gigabitethernet0/1.

B. Router 1 drops all traffic that ingresses interface gigabitethernet0/1 that has a FIB entry that exits a different interface.

C. Router 1 accepts source addresses that have a match in the FIB that indicates it is reachable through a real interface.

D. Router 1 accepts source addresses on interface gigabitethernet0/1 that are private addresses.

Answer: C

Explanation:

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_1/nx-os/security/configuration/guide/sec_ nx-os-cfg/sec_urpf.html

Question: 14

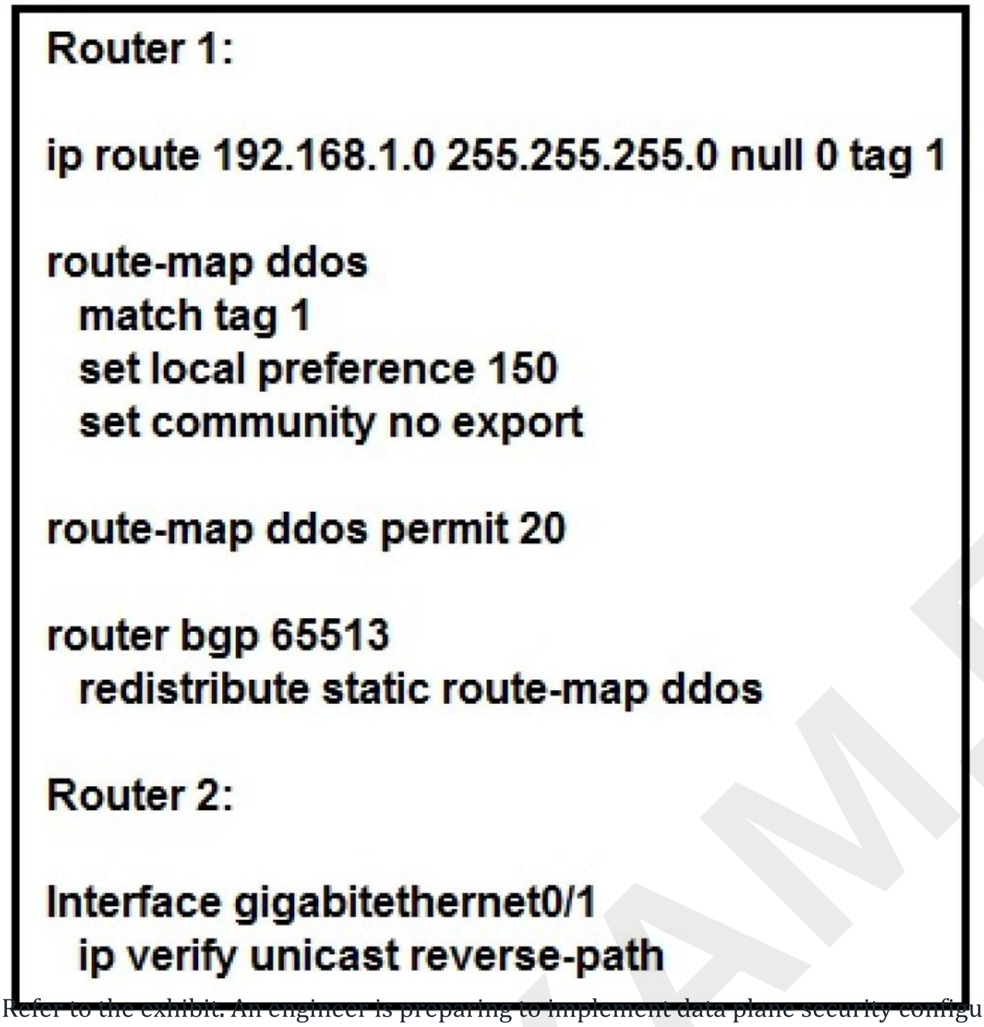

```
Router 1:

ip route 192.168.1.0 255.255.255.0 null 0 tag 1

route-map ddos
  match tag 1
  set local preference 150
  set community no export

route-map ddos permit 20

router bgp 65513
  redistribute static route-map ddos

Router 2:

Interface gigabitethernet0/1
  ip verify unicast reverse-path
```

Refer to the exhibit. An engineer is preparing to implement data plane security configuration. Which statement about this configuration is true?

A. Router 2 is the router receiving the DDoS attack.

B. Router 1 must be configured with uRPF for the RTBH implementation to be effective.

C. Router 1 is the trigger router in a RTBH implementation.

D. Router 2 must configure a route to null 0 for network 192.168.1.0/24 for the RTBH implementation to be complete.

Answer: C

Explanation:

Router 1 is the trigger router in a RTBH implementation.

Question: 15

Which configuration modifies Local Packet Transport Services hardware policies?

A.

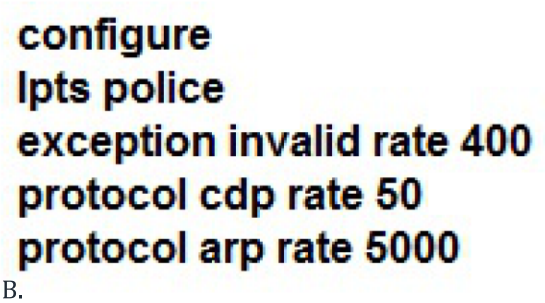

```
configure
lpts police
exception invalid rate 400
protocol cdp rate 50
protocol arp rate 5000
```

B.

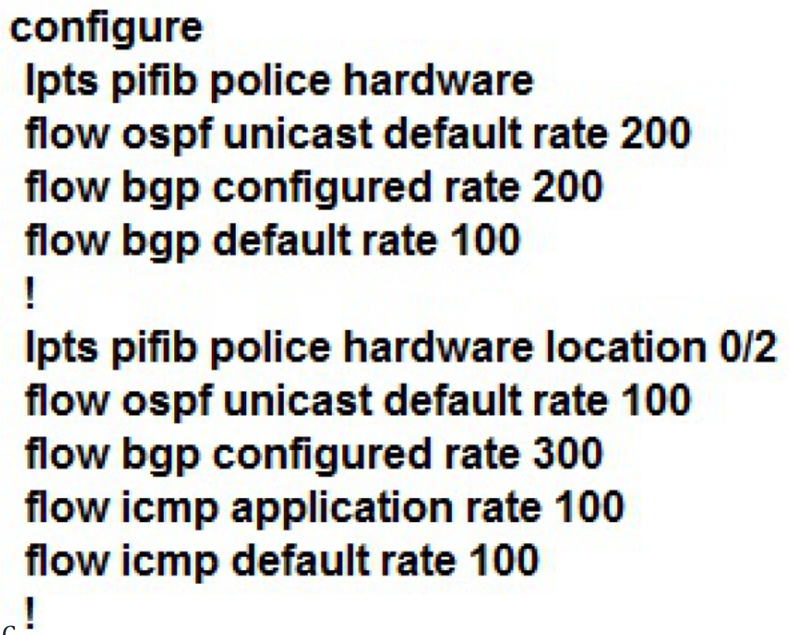

```
configure
 lpts pifib police hardware
 flow ospf unicast default rate 200
 flow bgp configured rate 200
 flow bgp default rate 100
 !
 lpts pifib police hardware location 0/2
 flow ospf unicast default rate 100
 flow bgp configured rate 300
 flow icmp application rate 100
 flow icmp default rate 100
 !
```

C.

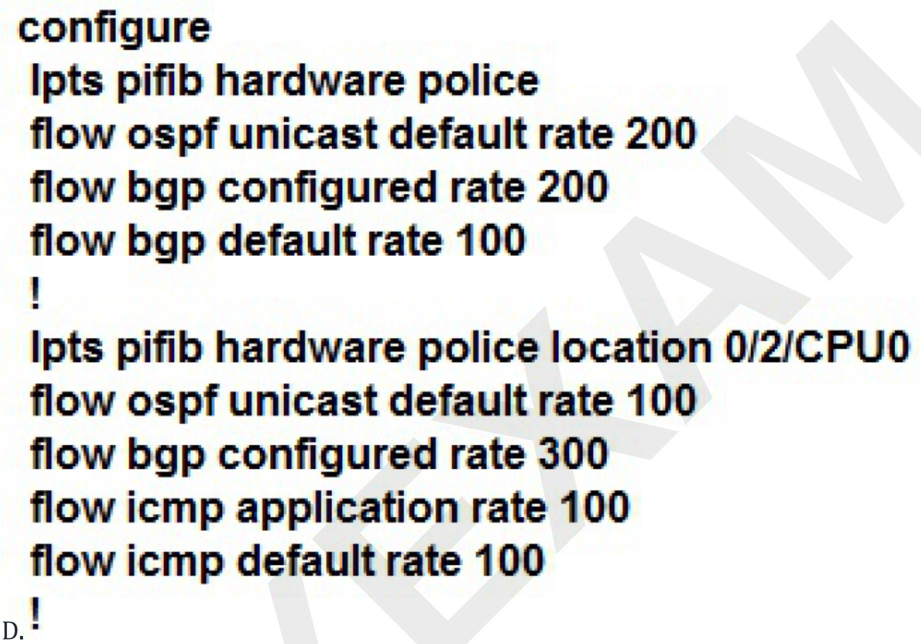

```
configure
 lpts pifib hardware police
 flow ospf unicast default rate 200
 flow bgp configured rate 200
 flow bgp default rate 100
 !
 lpts pifib hardware police location 0/2/CPU0
 flow ospf unicast default rate 100
 flow bgp configured rate 300
 flow icmp application rate 100
 flow icmp default rate 100
 !
```

D.

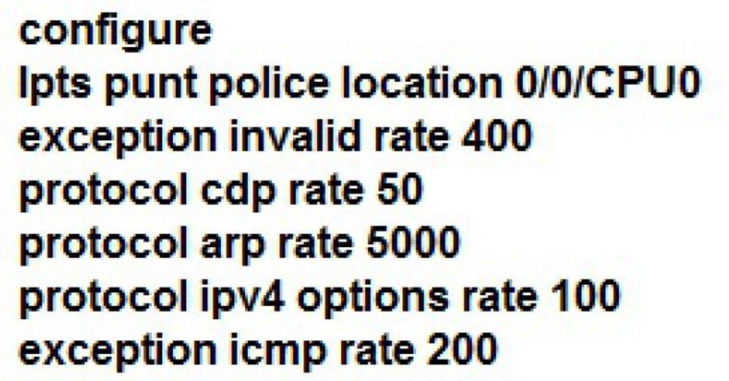

```
configure
lpts punt police location 0/0/CPU0
exception invalid rate 400
protocol cdp rate 50
protocol arp rate 5000
protocol ipv4 options rate 100
exception icmp rate 200
```

Answer: C

Explanation:

Reference:
https://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-1/addr_serv/command/reference/b_ipaddr_cr41crs/ b_ipaddr_cr41crs_chapter_0111.html#wp1754734006

Question: 16

Which additional feature does MPLS DiffServ tunneling support?

A. matching EXP and DSCP values
B. PHB layer management
C. using GRE tunnels to hide markings
D. interaction between MPLS and IGP

Answer: B

Explanation:

The correct answer is **B. PHB layer management.** MPLS DiffServ tunneling's primary advantage lies in its ability to manage per-hop behavior (PHB) at multiple layers. This allows service providers to prioritize traffic based on Quality of Service (QoS) requirements through different segments of their network. While core MPLS labels are used to forward packets, DiffServ principles are applied via the EXP bits within the MPLS header. The EXP bits, in conjunction with DiffServ Code Point (DSCP) markings within the IP header, allow for differentiated treatment of traffic. Options A, C, and D do not accurately represent the core benefits of MPLS DiffServ tunneling. Option A is not a specific additional feature but more of a mechanism of how DiffServ works within MPLS. Option C describes a functionality of GRE tunnels rather than MPLS DiffServ. Option D might have some interaction, but it's not a specific additional feature of the tunnel itself, rather a result of a network setup. PHB layer management is the feature specifically added to MPLS by the DiffServ principle which enables complex QoS management across the network. It creates a powerful framework for consistent end-to-end QoS within the network despite underlying encapsulation changes. This layered QoS approach allows networks to effectively meet the diverse needs of service provider clients.

Here are some relevant resources for further study:

1. **RFC 3270: Multi-Protocol Label Switching (MPLS) Support of Differentiated Services:** https://datatracker.ietf.org/doc/html/rfc3270 - This is the defining RFC for MPLS DiffServ.

2. **Cisco Documentation on QoS in MPLS:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/configuration/15-sy/qos-15-sy-book/qos-mpls.html - Provides practical examples and configuration details for MPLS DiffServ.

3. **Understanding MPLS and Diffserv in Networking** https://www.networkworld.com/article/3268516/understanding-mpls-and-diffserv-in-networking.html A high-level explanation on the interaction of both concepts.

Question: 17

You are creating new Cisco MPLS TE tunnels. Which type of RSVP message does the headend router send to reserve bandwidth on the path to the tailend router?

A. path
B. tear
C. error
D. reservation

Answer: A

Explanation:

The correct answer is **A. path**.

Resource Reservation Protocol (RSVP) is the signaling protocol used in MPLS Traffic Engineering (MPLS TE) to establish and maintain Label Switched Paths (LSPs). The process begins with the headend router, the ingress point for the tunnel, initiating the signaling process. It does this by sending an RSVP **Path** message downstream toward the tailend router, the egress point. This Path message contains vital information about the requested tunnel, such as the destination address (tailend router), traffic characteristics (bandwidth requirements), and the explicit route (if configured). This message effectively probes the network and informs intermediate routers about the tunnel's intended path. It does not reserve any resources itself, but rather sets the stage for the actual resource reservation.

Once the Path message reaches the tailend router, that router then sends an RSVP **Reservation** (Resv) message upstream along the reverse path of the path message. The Resv message is responsible for making bandwidth reservations along the specified path. The 'tear' message is used to remove previously established reservations and paths, and error messages are used to indicate problems that occurred during path or reservation establishment. Therefore, in the context of initiating a new MPLS TE tunnel and reserving resources, the RSVP Path message is the initiating message from the headend router.

Authoritative Link:

Cisco - Understanding RSVP for MPLS: https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/12140-rsvp-mpls.html

Question: 18

Which statement describes the advantage of a Multi-Layer control plane?

A. It provides multivendor configuration capabilities for Layer 3 to Layer 1.
B. It automatically provisions, monitors, and manages traffic across Layer 0 to Layer 3.
C. It supports dynamic wavelength restoration in Layer 0.
D. It minimizes human error configuring converged networks.

Answer: C

Explanation:

The correct answer is C: It supports dynamic wavelength restoration in Layer 0. A multi-layer control plane, especially within service provider networks, allows for coordination and management across different layers of the network stack (Layer 0, 1, 2, and 3). Specifically, Layer 0 refers to the optical layer, which deals with physical transmission of light waves. In a multi-layer control plane, the intelligence of higher layers (like Layer 2 or 3) can be used to manage the physical infrastructure of Layer 0. Dynamic wavelength restoration is a key advantage where the control plane can automatically reroute traffic to an alternative optical path in Layer 0 if a failure occurs. This reduces downtime and improves network resilience.

Option A is incorrect because while a multi-layer control plane can facilitate some level of interoperation between vendors, its primary function isn't to provide multivendor configuration capabilities across layers.

Option B is inaccurate as multi-layer control planes primarily focus on coordinating management and not necessarily auto-provisioning everything across all layers. While Option D suggests minimized human error, it's a general benefit of network automation and not unique to a multi-layer control plane. Furthermore, error minimization does not fully encompass the core advantages of a multi-layer approach.

Dynamic wavelength restoration, as offered by a multi-layer control plane, directly addresses the need for resilient and efficient operation at the lowest network level, which is particularly crucial for service provider networks. This capability is enabled by sharing information between layers, allowing higher layers to instruct lower layers on rerouting traffic during failures.

Authoritative Links for Further Research:

1. **Cisco's Layered Approach in Network Design:** While not specifically about Multi-Layer control plane, this provides context.
https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Campus/HA_Campus_DG/ha_campus_ch2.ht 2. **Juniper's SDN and Multi-layer Control:**https://www.juniper.net/content/dam/www/assets/white-papers/us/en/multilayer-sdn-control.pdf - This focuses on the benefits of a multi-layer approach.
3. **Optical Layer Control Plane:**https://www.ciena.com/insights/what-is/optical-control-plane.html - Focuses on specific benefits of the optical/Layer 0 control plane.

Question: 19

DRAG DROP -
Drag and drop the technologies from the left onto the correct definitions on the right.
Select and Place:

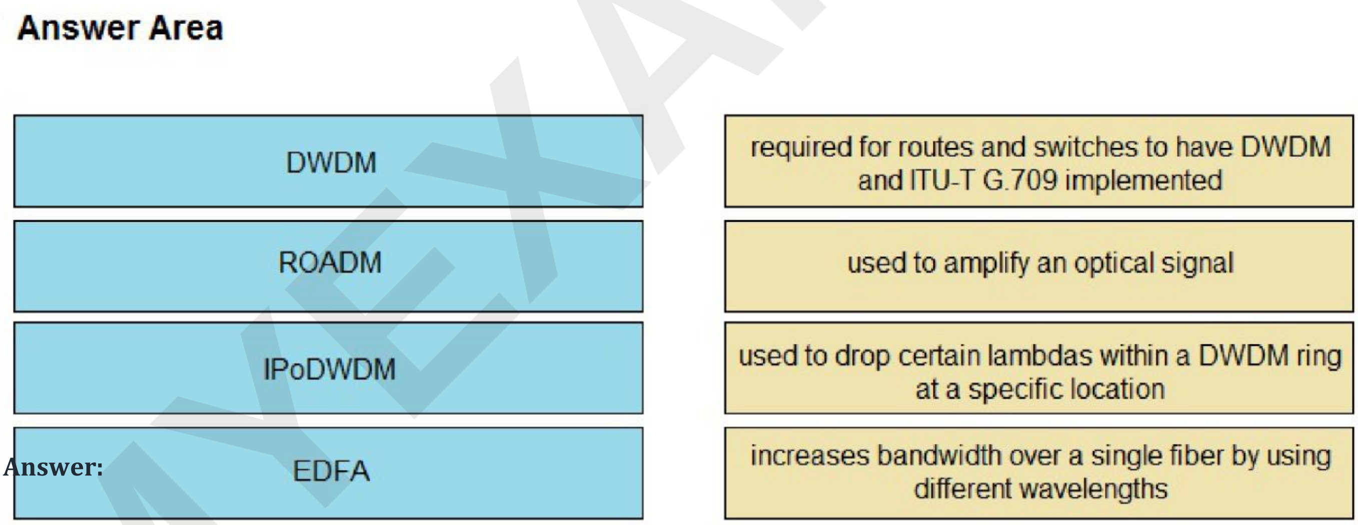

Answer Area

DWDM	required for routes and switches to have DWDM and ITU-T G.709 implemented
ROADM	used to amplify an optical signal
IPoDWDM	used to drop certain lambdas within a DWDM ring at a specific location
EDFA	increases bandwidth over a single fiber by using different wavelengths

Answer:

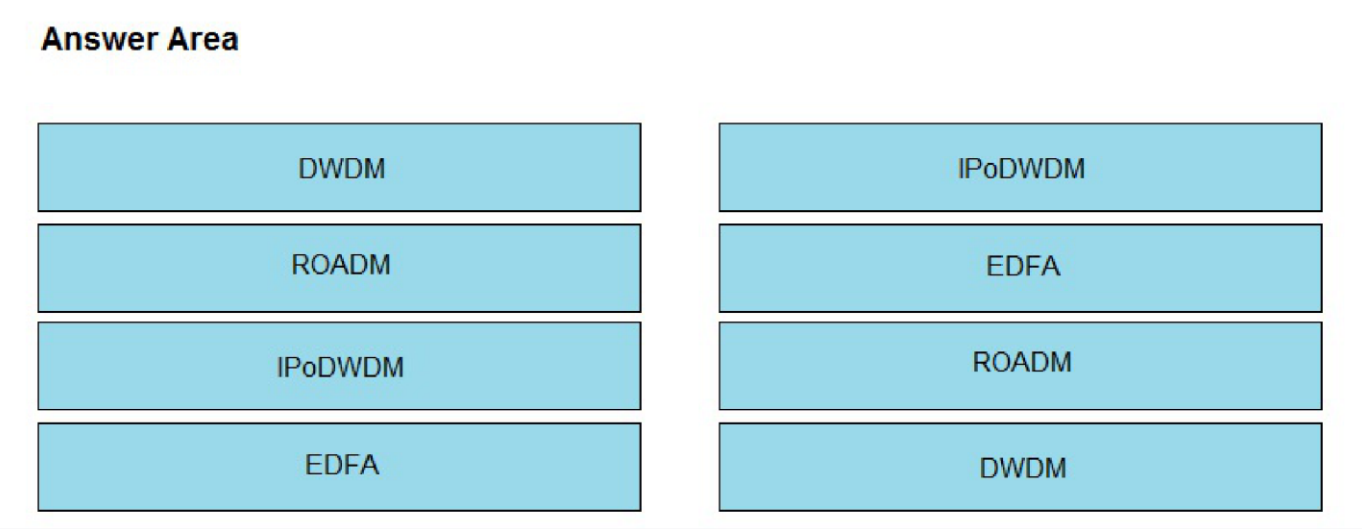

Answer Area

DWDM	IPoDWDM
ROADM	EDFA
IPoDWDM	ROADM
EDFA	DWDM

Explanation:

IPoDWDM = IP over Dense wavelength-division multiplexing.

EDFA = Erbium-Doped Fiber Amplifier.

ROADM = Reconfigurable optical add-drop multiplexer.

DWDM = Dense wavelength-division multiplexing.

Question: 20

An engineer is setting up overlapping VPNs to allow VRF ABC and XYZ to communicate with VRF CENTRAL but wants to make sure that VRF ABC and XYZ cannot communicate.

Which configuration accomplishes these objectives?

A.

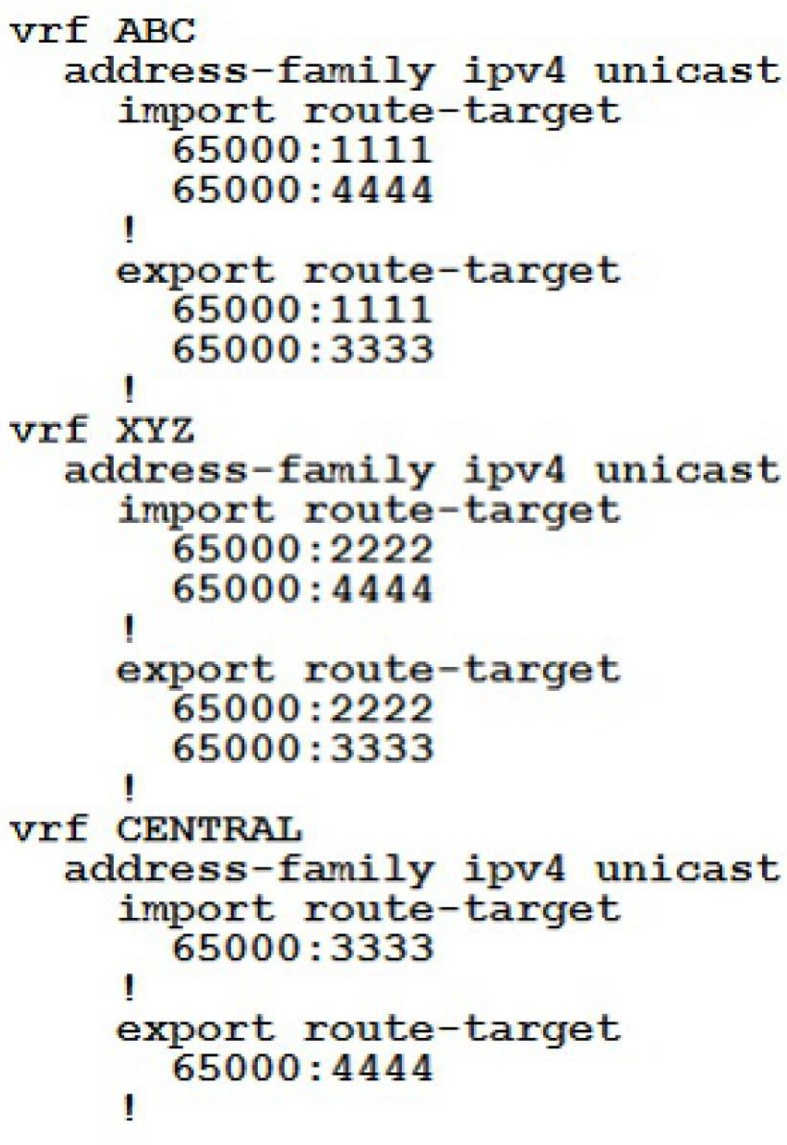

```
vrf ABC
  address-family ipv4 unicast
    import route-target
      65000:1111
      65000:4444
    !
    export route-target
      65000:1111
      65000:3333
    !
vrf XYZ
  address-family ipv4 unicast
    import route-target
      65000:2222
      65000:4444
    !
    export route-target
      65000:2222
      65000:3333
    !
vrf CENTRAL
  address-family ipv4 unicast
    import route-target
      65000:3333
    !
    export route-target
      65000:4444
    !
```

B.

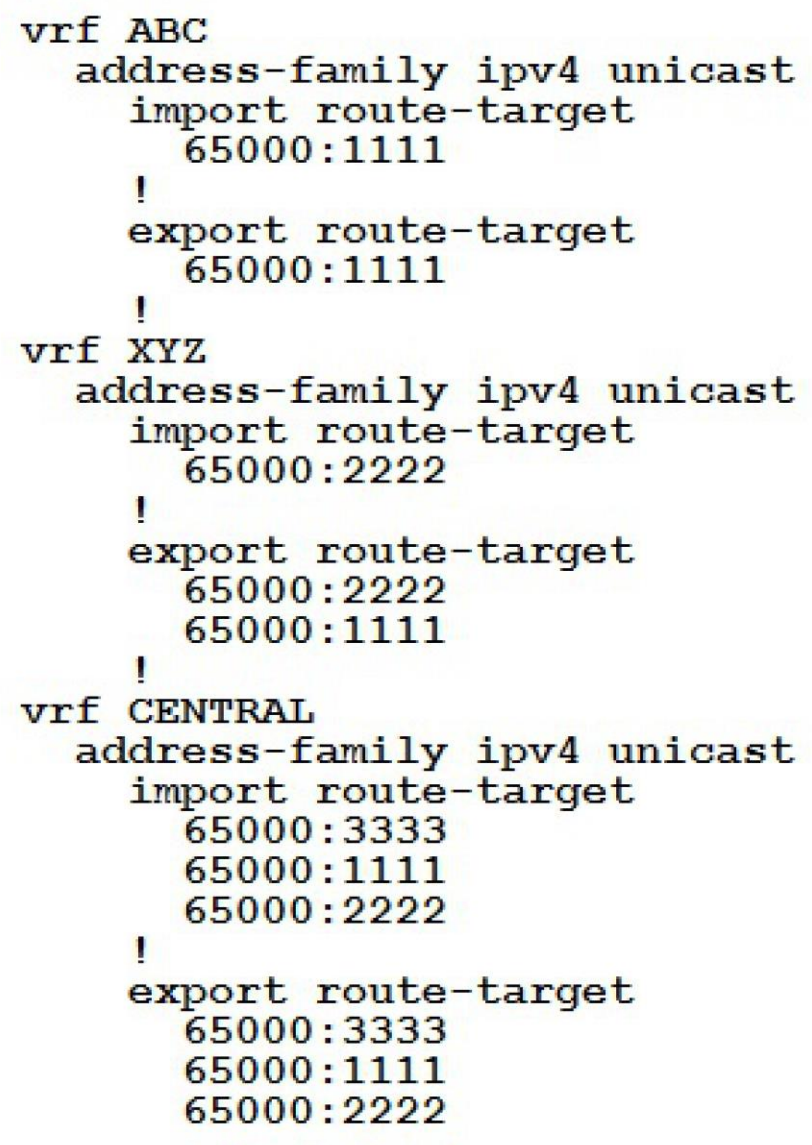

```
vrf ABC
  address-family ipv4 unicast
    import route-target
      65000:1111
    !
    export route-target
      65000:1111
    !
vrf XYZ
  address-family ipv4 unicast
    import route-target
      65000:2222
    !
    export route-target
      65000:2222
      65000:1111
    !
vrf CENTRAL
  address-family ipv4 unicast
    import route-target
      65000:3333
      65000:1111
      65000:2222
    !
    export route-target
      65000:3333
      65000:1111
      65000:2222
```

C.

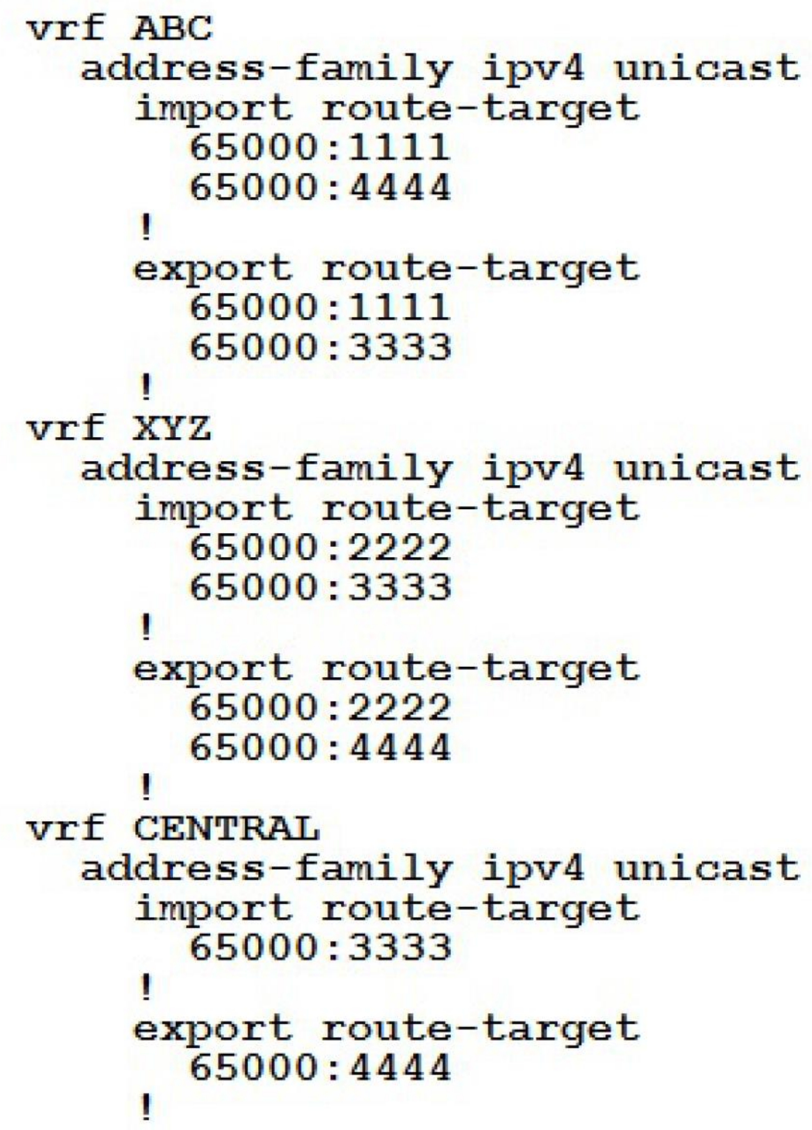

```
vrf ABC
  address-family ipv4 unicast
    import route-target
      65000:1111
      65000:4444
    !
    export route-target
      65000:1111
      65000:3333
    !
vrf XYZ
  address-family ipv4 unicast
    import route-target
      65000:2222
      65000:3333
    !
    export route-target
      65000:2222
      65000:4444
    !
vrf CENTRAL
  address-family ipv4 unicast
    import route-target
      65000:3333
    !
    export route-target
      65000:4444
    !
```

D.

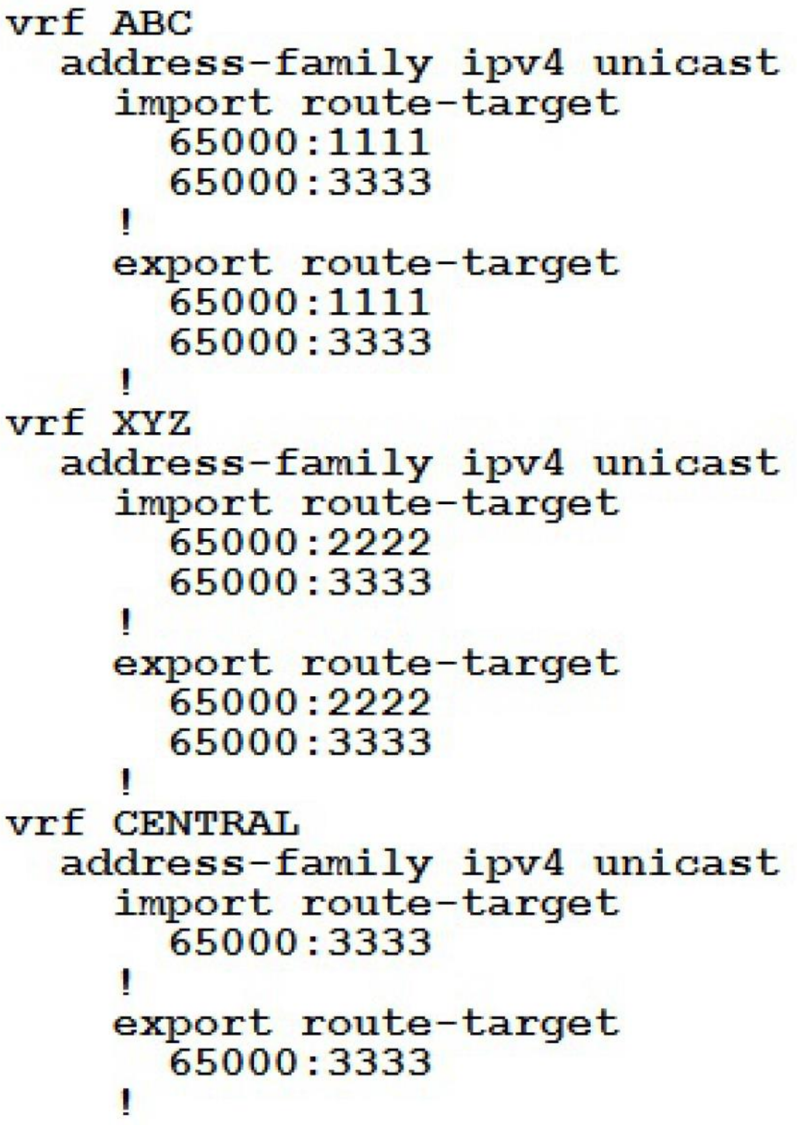

```
vrf ABC
  address-family ipv4 unicast
    import route-target
      65000:1111
      65000:3333
    !
    export route-target
      65000:1111
      65000:3333
    !
vrf XYZ
  address-family ipv4 unicast
    import route-target
      65000:2222
      65000:3333
    !
    export route-target
      65000:2222
      65000:3333
    !
vrf CENTRAL
  address-family ipv4 unicast
    import route-target
      65000:3333
    !
    export route-target
      65000:3333
    !
```

Answer: A
Explanation:

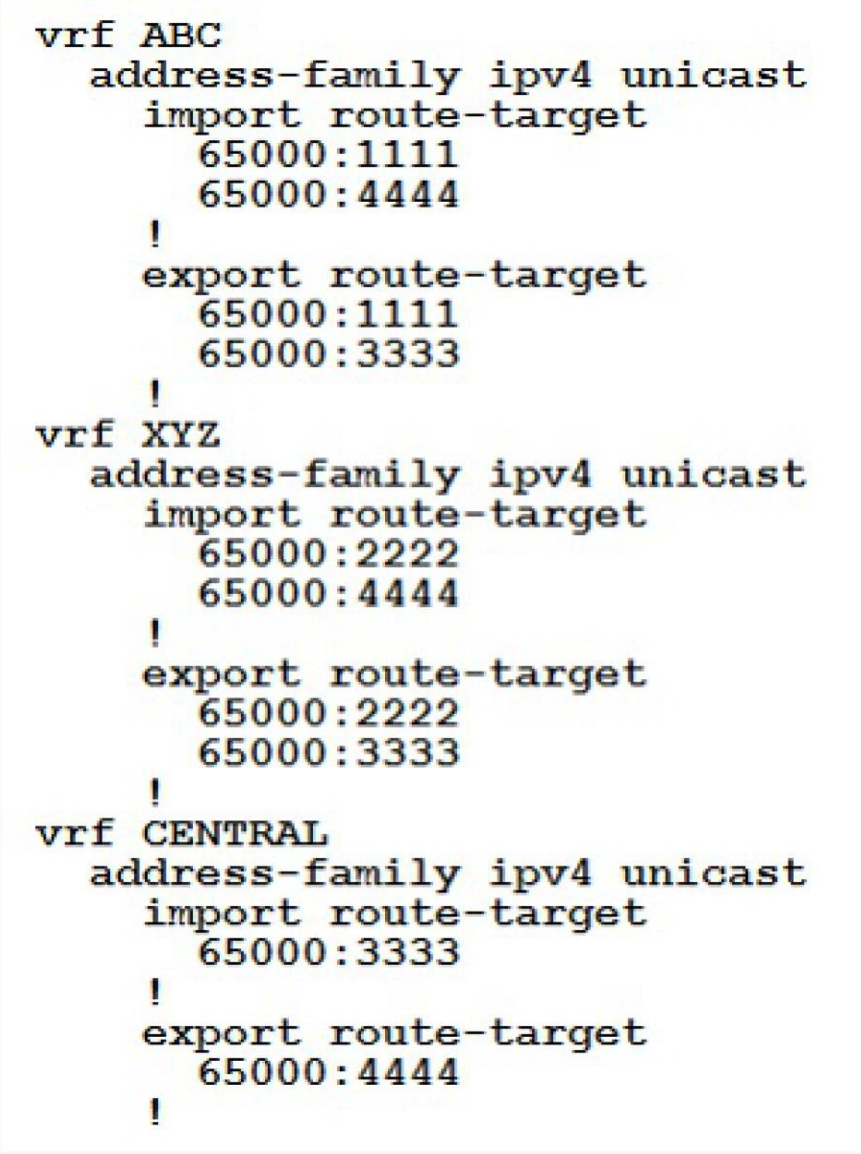

```
vrf ABC
  address-family ipv4 unicast
    import route-target
      65000:1111
      65000:4444
    !
    export route-target
      65000:1111
      65000:3333
    !
vrf XYZ
  address-family ipv4 unicast
    import route-target
      65000:2222
      65000:4444
    !
    export route-target
      65000:2222
      65000:3333
    !
vrf CENTRAL
  address-family ipv4 unicast
    import route-target
      65000:3333
    !
    export route-target
      65000:4444
    !
```

Question: 21

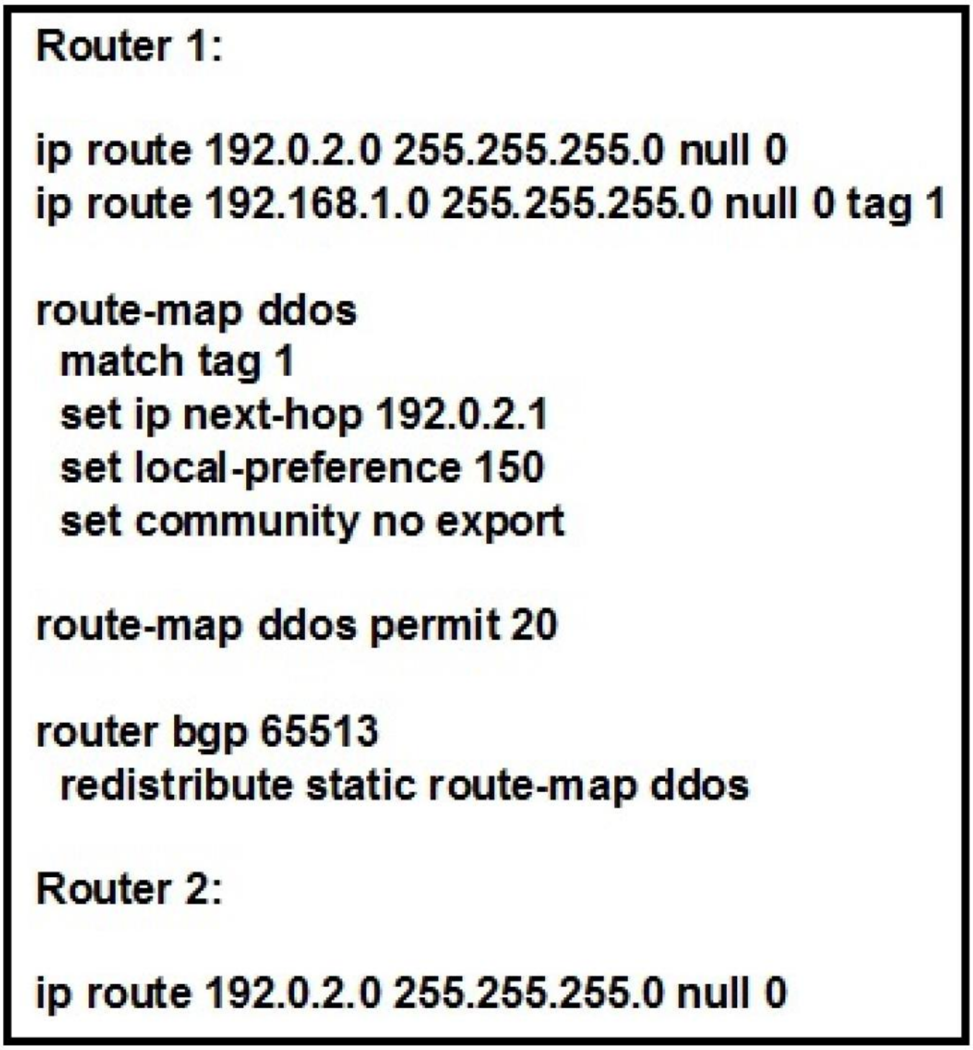

```
Router 1:

ip route 192.0.2.0 255.255.255.0 null 0
ip route 192.168.1.0 255.255.255.0 null 0 tag 1

route-map ddos
  match tag 1
  set ip next-hop 192.0.2.1
  set local-preference 150
  set community no export

route-map ddos permit 20

router bgp 65513
  redistribute static route-map ddos

Router 2:

ip route 192.0.2.0 255.255.255.0 null 0
```

Refer to the exhibit. An engineer is preparing to implement data plane security configuration. Which statement about this configuration is true?

A. Router 1 and Router 2 advertise the route to 192.0.2.0 to all BGP peers. B. All traffic to 192.168.1.0/24 is dropped.

C. All traffic is dropped.

D. Router 1 drops all traffic with a local-preference set to 150.

Answer: B

Explanation:

All traffic to 192.168.1.0/24 is dropped.

Question: 22

Which MPLS design attribute can you use to provide Internet access to a major customer through a separate dedicated VPN?

A. The Internet gateway router is connected as a PE router to the MPLS backbone.

B. The CE router supports VRF-Lite and the full BGP routing table.

C. The Internet gateway inserts the full Internet BGP routing table into the Internet access VPN.

D. The customer that needs the Internet access service is assigned to the same RTs as the Internet gateway.

Answer: D

Explanation:

The correct answer is **D. The customer that needs the Internet access service is assigned to the same RTs as the Internet gateway.** Here's why:

MPLS VPNs rely on Route Target (RT) attributes to control route distribution within the VPN. For a customer to access the internet through a dedicated VPN, the internet gateway's routes must be imported into the customer's VPN routing table. This is achieved by configuring both the internet gateway and the customer's PE router to use matching import RTs. RTs effectively define which VPN routes are shared and thus, the connectivity between networks.

Option A, connecting the internet gateway as a PE router, is standard practice for MPLS internet access but doesn't by itself establish separate VPN access. B, VRF-Lite support on the CE, is useful for segmentation but is irrelevant in the context of proper MPLS VPN internet access. C, the internet gateway inserting the full Internet BGP table into the VPN is also standard, but this table must be distributed using shared RTs.

Therefore, the essential element in achieving a dedicated internet access VPN via MPLS is matching import RTs on the gateway and the customer's PE, allowing the routes to be imported into the customer's VRF. If not, the routes are not shared, and the customer would not be able to reach the internet. Option D is the only one that ensures proper sharing and therefore connectivity.

Authoritative Links:

1. **Cisco's MPLS VPN Configuration Guide:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mpls/configuration/15-mt/mpls-15-mt-book/mp-vpn.html - Specifically, refer to the sections on Route Target usage.
2. **Juniper's Understanding Route Targets:** https://www.juniper.net/documentation/us/en/software/junos/mpls/topics/topic-map/mpls-vpn-route-targets.html - Provides another vendor's explanation of Route Target function.

Question: 23

Which configuration enables BGP FlowSpec client function and installation of policies on all local interfaces?

A. flowspec address-family ipv4 local-install interface-all
B. flowspec address-family ipv4 install interface-all local
C. flowspec address-family ipv4 install interface-all
D. flowspec address-family ipv4 local-install all-interface

Answer: A

Explanation:

Option A, flowspec address-family ipv4 local-install interface-all, is the correct configuration for enabling BGP FlowSpec client functionality and installing policies on all local interfaces. Here's why: BGP FlowSpec allows the propagation of traffic filtering rules via BGP updates. The address-family ipv4 command specifies that the configuration applies to IPv4 address families. The local-install keyword ensures that received FlowSpec routes are installed into the local routing table for policy enforcement. Crucially, interface-all ensures that these FlowSpec policies are applied to all local interfaces. Option B is incorrect because the install keyword should precede the interface-all keyword, and the local keyword placement is incorrect. Option C lacks the critical local-install keyword. Without local-install, the FlowSpec rules might be learned via BGP but won't be activated in the local router's forwarding plane. Option D has incorrect syntax with all-interface. This is not the correct option in cisco ios. FlowSpec's effectiveness relies on this combination: learning the policies through

BGP, installing them locally, and then applying them across all interfaces. Essentially, these steps ensure that the router can both receive traffic manipulation rules and effectively apply those rules to all incoming traffic. This functionality is essential for sophisticated traffic engineering and mitigation in service provider networks, where granular control over traffic flows is vital.

Authoritative links for further research:

Cisco BGP FlowSpec Configuration Guide:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/15-e/irg-15-e-book/irg-flow-spec.html (Specifically look for details on local-install and interface application)
RFC 8955 - Dissemination of Flow Specification Rules:https://datatracker.ietf.org/doc/html/rfc8955 (This RFC provides the foundational information on BGP FlowSpec)

Question: 24

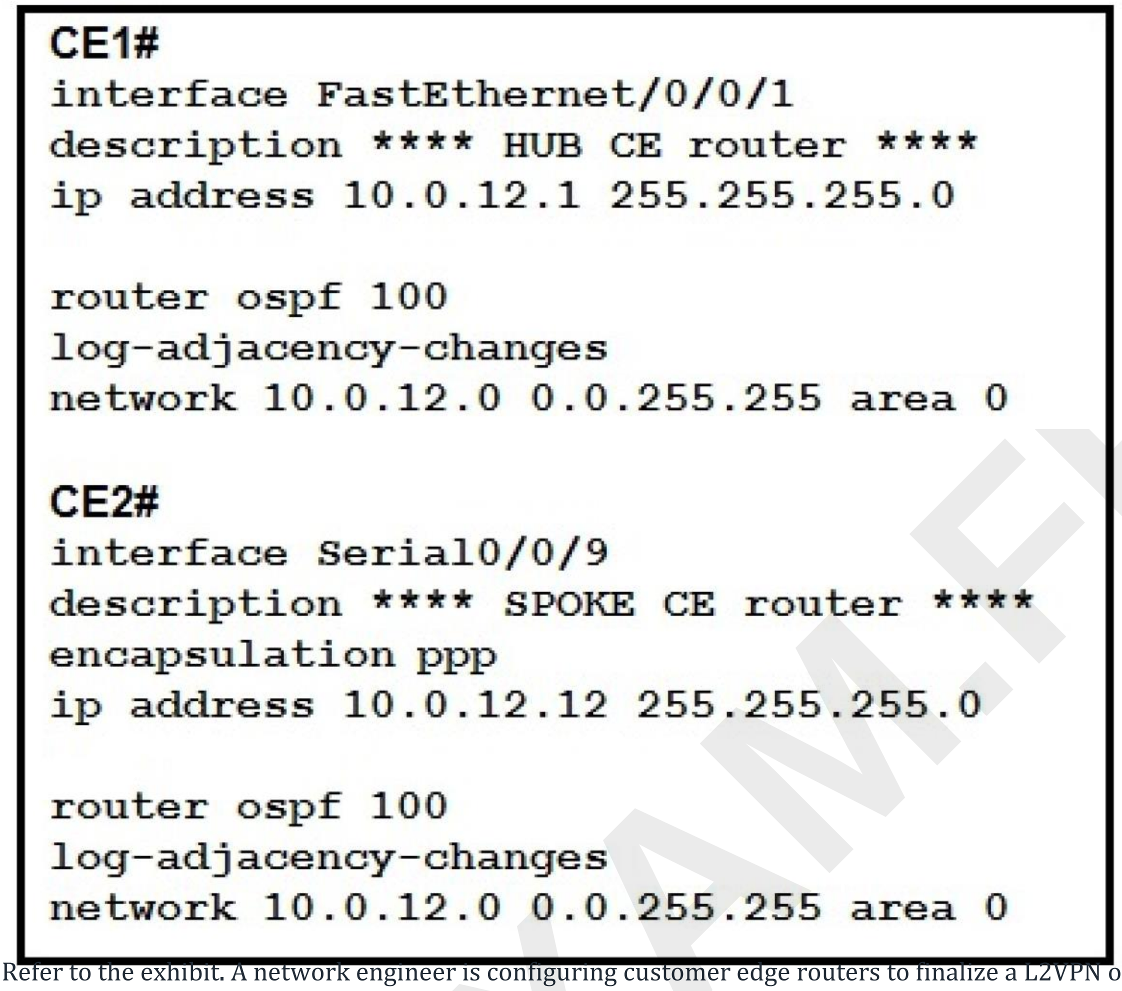

```
CE1#
interface FastEthernet/0/0/1
description **** HUB CE router ****
ip address 10.0.12.1 255.255.255.0

router ospf 100
log-adjacency-changes
network 10.0.12.0 0.0.255.255 area 0

CE2#
interface Serial0/0/9
description **** SPOKE CE router ****
encapsulation ppp
ip address 10.0.12.12 255.255.255.0

router ospf 100
log-adjacency-changes
network 10.0.12.0 0.0.255.255 area 0
```

Refer to the exhibit. A network engineer is configuring customer edge routers to finalize a L2VPN over MPLS deployment. Assume that the AToM L2VPN service that connects the two CEs is configured correctly on the service provider network.
Which action causes the solution to fail?

A. OSPF does not work with L2VPN services.
B. The routing protocol network types are not compatible.
C. A loopback with a /32 IP address has not been used.
D. The xconnect statement has not been defined.

Answer: B

Explanation:

The routing protocol network types are not compatible.

Question: 25

An engineer working for telecommunication company needs to secure the LAN network using a prefix list. Which best practice should the engineer follow when he implements a prefix list?

A. An engineer must identify the prefix list with a number only.

B. The final entry in a prefix list must be /32.

C. An engineer must include only the prefixes for which he needs to log activity.

D. An engineer must use nonsequential sequence numbers in the prefix list so that he can insert additional entries later.

Answer: D

Explanation:

The correct answer is **D.** Using non-sequential sequence numbers in a prefix list is a best practice for maintainability and future modifications. Prefix lists in Cisco devices, used for filtering routes based on IP prefixes, are processed sequentially. Sequence numbers dictate the order in which entries are evaluated. If you use sequential numbers (e.g., 10, 20, 30), adding a new entry in the middle would require renumbering existing lines, which is cumbersome and error-prone. By employing non-sequential numbers (e.g., 10, 50, 100), you can easily insert new prefixes by using an intermediate sequence number (e.g., 60). This flexibility makes the prefix list easier to manage and prevents disruption of existing configurations. Option A is incorrect because prefix lists can be identified by name. Option B is incorrect as not all prefixes need to be /32. Option C is incorrect because logging is not a core function of a prefix list; logging is related to access-lists.

Authoritative links for further research:

Cisco Prefix Lists Configuration:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_flt/configuration/15-sy/ipv6-prefix-list-15-sy.html
IP Routing: Prefix Lists:https://www.networklessons.com/routing/ip-routing-prefix-lists

Question: 26

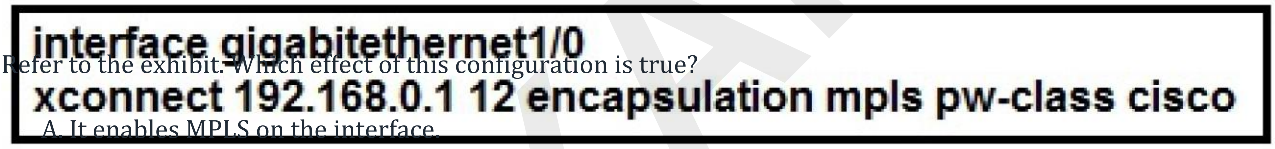

```
interface gigabitethernet1/0
xconnect 192.168.0.1 12 encapsulation mpls pw-class cisco
```

Refer to the exhibit. Which effect of this configuration is true?

A. It enables MPLS on the interface.

B. It creates a pseudowire class named cisco.

C. It enables AToM on interface gigabitethernet1/0.

D. It enables tagging for VLAN 12 on the interface.

Answer: C

Explanation:

Reference:
https://community.cisco.com/t5/service-providers-documents/configuration-example-atom-any-transport-ov

er-mpls-for-like-to/ta-p/3144810

Question: 27

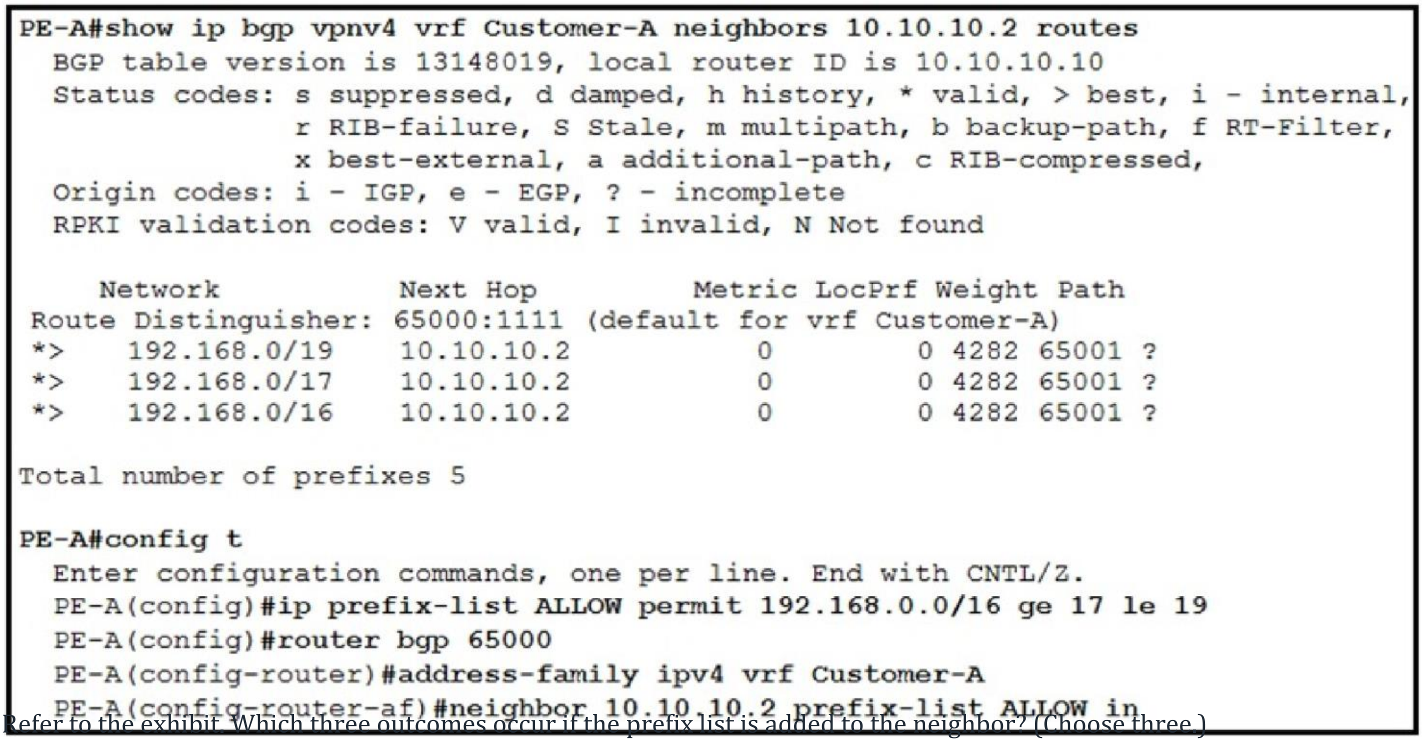

```
PE-A#show ip bgp vpnv4 vrf Customer-A neighbors 10.10.10.2 routes
  BGP table version is 13148019, local router ID is 10.10.10.10
  Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
                r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
                x best-external, a additional-path, c RIB-compressed,
  Origin codes: i - IGP, e - EGP, ? - incomplete
  RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 Route Distinguisher: 65000:1111 (default for vrf Customer-A)
 *>    192.168.0/19    10.10.10.2              0             0 4282 65001 ?
 *>    192.168.0/17    10.10.10.2              0             0 4282 65001 ?
 *>    192.168.0/16    10.10.10.2              0             0 4282 65001 ?

Total number of prefixes 5

PE-A#config t
  Enter configuration commands, one per line. End with CNTL/Z.
  PE-A(config)#ip prefix-list ALLOW permit 192.168.0.0/16 ge 17 le 19
  PE-A(config)#router bgp 65000
  PE-A(config-router)#address-family ipv4 vrf Customer-A
  PE-A(config-router-af)#neighbor 10.10.10.2 prefix-list ALLOW in
```

Refer to the exhibit. Which three outcomes occur if the prefix list is added to the neighbor? (Choose three.)

A. 192.168.0.0/16 is denied.

B. 192.168.0.0/16 is permitted.
C. 192.168.0.0/19 is permitted
D. 192.168.0.0/19 is denied.

E. 192.168.0.0/17 is permitted
F. 192.168.0.0/17 is denied.

Answer: ACE

Explanation:

A. 192.168.0.0/16 is denied.

C. 192.168.0.0/19 is permitted.

E. 192.168.0.0/17 is permitted.

Question: 28

DRAG DROP -
Drag and drop the descriptions from the left onto the corresponding OS types on the right.
Select and Place:

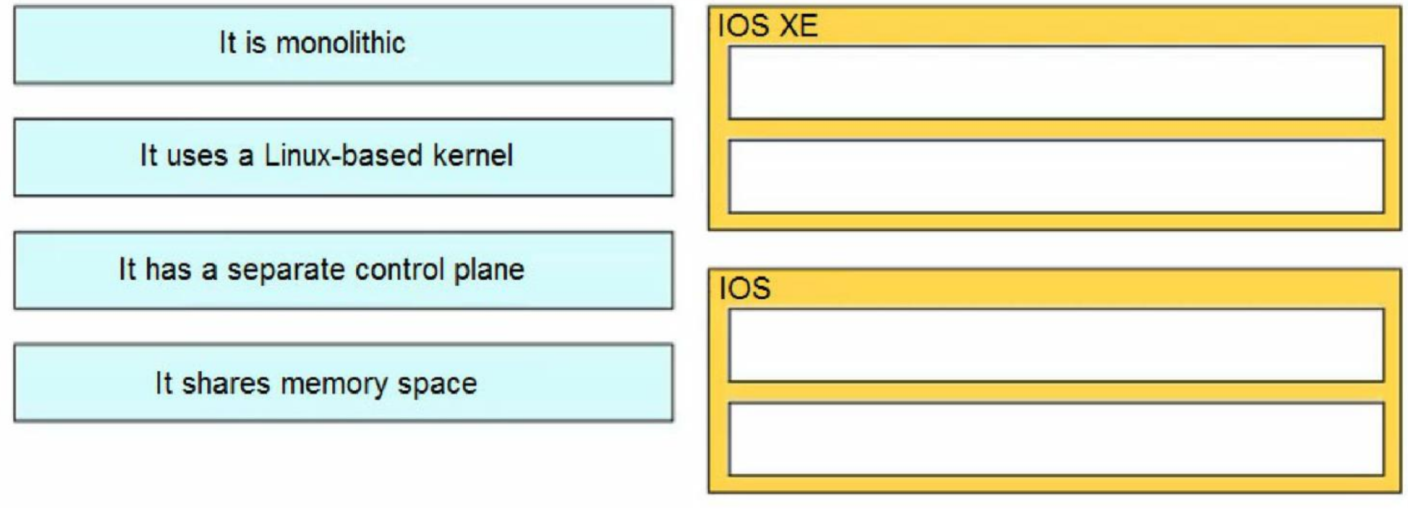

It is monolithic

It uses a Linux-based kernel

It has a separate control plane

It shares memory space

IOS XE

IOS

Answer:

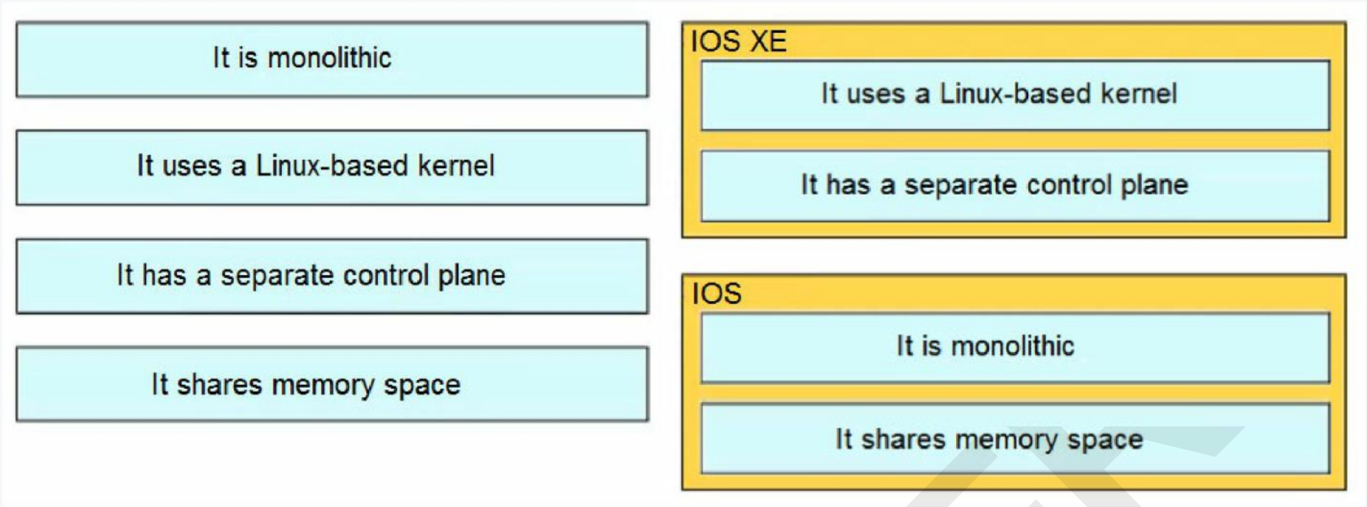

It is monolithic

It uses a Linux-based kernel

It has a separate control plane

It shares memory space

IOS XE

It uses a Linux-based kernel

It has a separate control plane

IOS

It is monolithic

It shares memory space

Question: 29

What does DWDM use to combine multiple optical signals?

A. IP protocols
B. wavelength
C. time slots
D. frequency

Answer: B

Explanation:

Dense Wavelength Division Multiplexing (DWDM) combines multiple optical signals onto a single optical fiber by utilizing different wavelengths of light. Each optical signal is assigned a unique wavelength, effectively creating separate channels on the same fiber. This allows for the simultaneous transmission of numerous data streams, increasing the overall capacity of the fiber optic cable. Think of it like multiple radio stations broadcasting on different frequencies; DWDM does the same but with light wavelengths. This technique is crucial for high-capacity networks and long-distance data transmission. IP protocols, time slots, and frequencies are not the direct mechanisms used by DWDM to combine optical signals. IP protocols operate at a higher network layer, governing how data packets are routed, not how light signals are multiplexed. Time slots, like those used in Time Division Multiplexing (TDM), would divide the transmission time, not the optical spectrum. While frequency is related to wavelength, DWDM specifically uses the wavelength as its defining separation mechanism. The different wavelengths are combined using optical multiplexers and separated using optical demultiplexers at the receiving end. Therefore, the correct answer is wavelength. This allows for an efficient use of the optical fiber and increases bandwidth capacity.

Further reading can be found here:

Cisco - Dense Wavelength-Division Multiplexing (DWDM) Overview
Wikipedia - Wavelength-division multiplexing

Question: 30

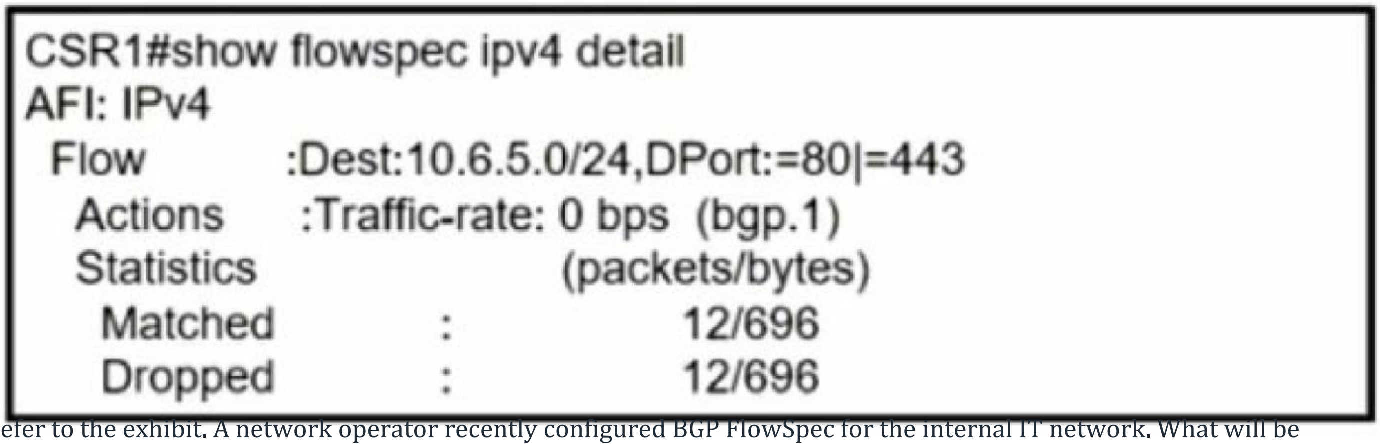

```
CSR1#show flowspec ipv4 detail
AFI: IPv4
  Flow          :Dest:10.6.5.0/24,DPort:=80|=443
    Actions      :Traffic-rate: 0 bps  (bgp.1)
    Statistics                     (packets/bytes)
      Matched             :                12/696
      Dropped             :                12/696
```

Refer to the exhibit. A network operator recently configured BGP FlowSpec for the internal IT network. What will be inferred from the configuration deployed on the network?

A. The policy is configured locally on CSR1 and drops all traffic for TCP ports 80 and 443 B. The
policy is configured locally on CSR1 and currently has no active traffic
C. The policy is learned via BGP FlowSpec and drops all traffic for TCP ports 80 and 443 D. The
policy is learned via BGP FlowSpec and has active traffic

Answer: C

Explanation:

The policy is learned via BGP FlowSpec and drops all traffic for TCP ports 80 and 443.

Question: 31

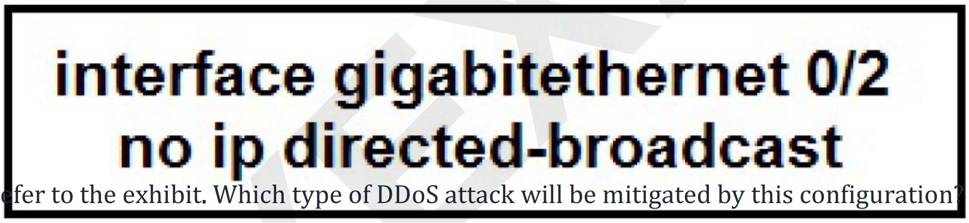

```
interface gigabitethernet 0/2
  no ip directed-broadcast
```

Refer to the exhibit. Which type of DDoS attack will be mitigated by this configuration?

A. teardrop attack
B. smurf attack
C. SYN flood
D. SIP INVITE flood attacks

Answer: B

Explanation:

Correct answer is B:smurf attack.

Question: 32

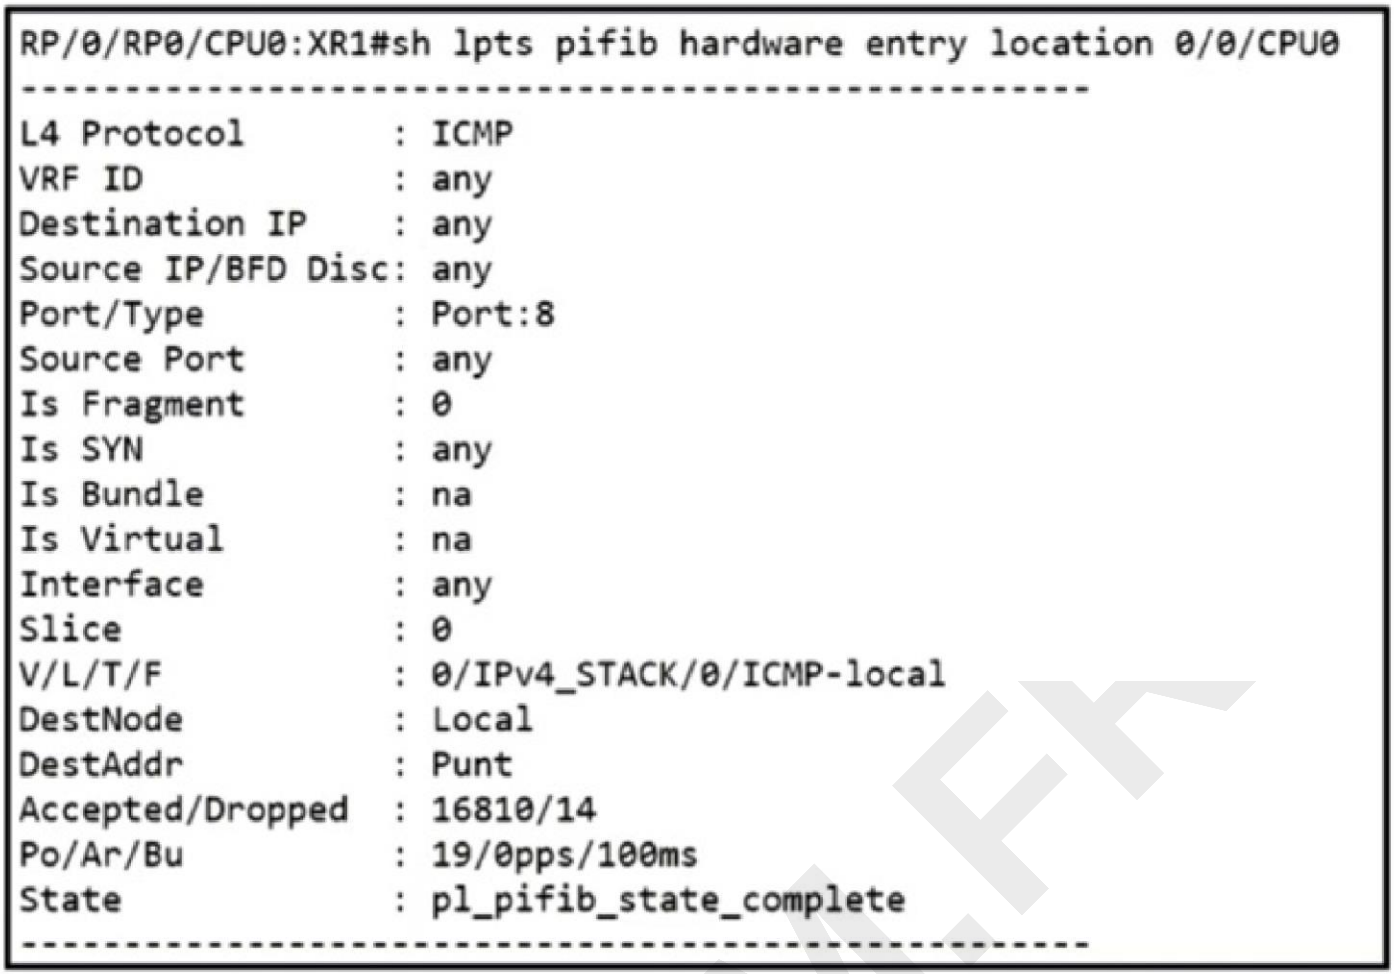

```
RP/0/RP0/CPU0:XR1#sh lpts pifib hardware entry location 0/0/CPU0
----------------------------------------------------
L4 Protocol          : ICMP
VRF ID               : any
Destination IP       : any
Source IP/BFD Disc: any
Port/Type            : Port:8
Source Port          : any
Is Fragment          : 0
Is SYN               : any
Is Bundle            : na
Is Virtual           : na
Interface            : any
Slice                : 0
V/L/T/F              : 0/IPv4_STACK/0/ICMP-local
DestNode             : Local
DestAddr             : Punt
Accepted/Dropped     : 16810/14
Po/Ar/Bu             : 19/0pps/100ms
State                : pl_pifib_state_complete
----------------------------------------------------
```

Refer to the exhibit. While troubleshooting the network, a network operator with an employee id: 1234:55:678 is trying to ping XR1. Which result should the operator expect when trying to ping to an XR1 local address?

A. AII ICMP traffic is dropped
B. AII ICMP traffic responds successfully
C. ICMP traffic works at a policed rate of 19 bytes per second every 100 ms D.
ICMP traffic works at a policed rate of 19 packets every 100 ms

Answer: A

Explanation:

AII ICMP traffic is dropped.

Question: 33

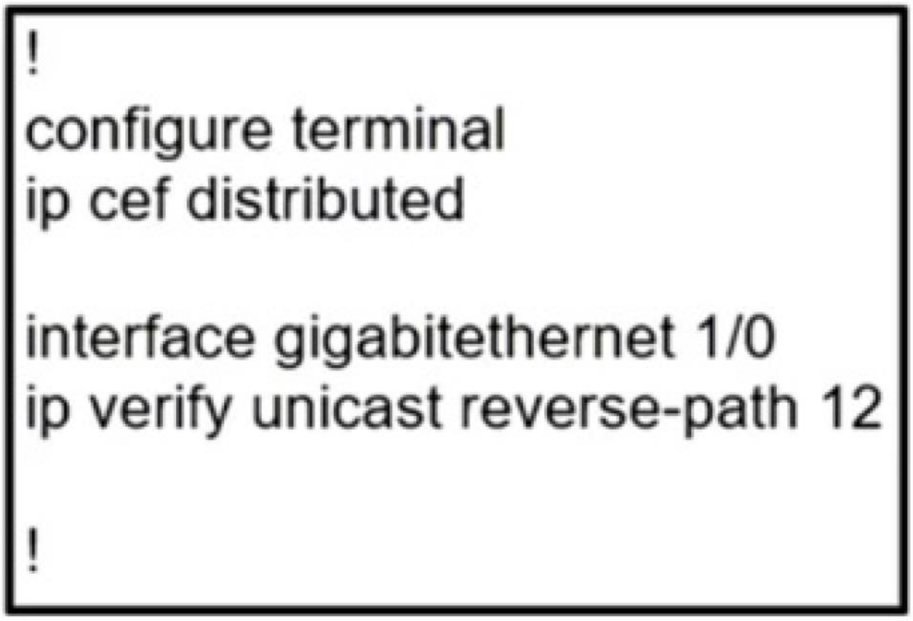

```
!
configure terminal
ip cef distributed

interface gigabitethernet 1/0
ip verify unicast reverse-path 12

!
```

Refer to the exhibit. Which show command should be implemented to display per-interface statistics about uRPF drops and suppressed drops?

A. show cef interface
B. show ip traffic
C. show ip interface
D. show ip interface brief

Answer: C

Explanation:

Reference:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_urpf/configuration/xe-3s/sec-data-urpf-xe-3s-book/cfg-unicast-rpf.html#GUID-3609F224-5D40-422E-8754-B12C60677612

Question: 34

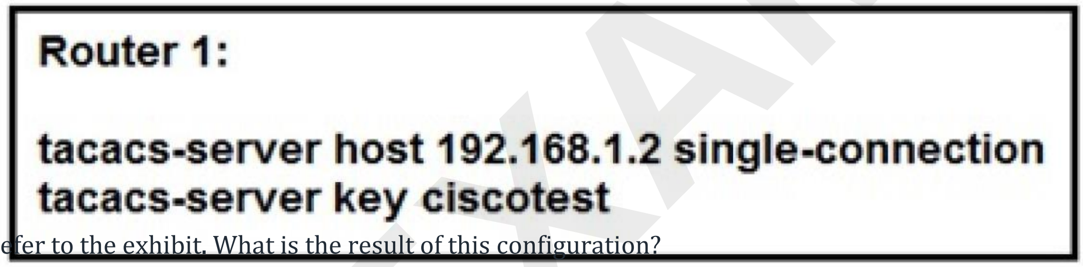

```
Router 1:

tacacs-server host 192.168.1.2 single-connection
tacacs-server key ciscotest
```

Refer to the exhibit. What is the result of this configuration?

A. Router 1 opens and closes a TCP connection to the TACACS+ server every time a user requires authorization B. Router 1 and the TACACS+ server maintain one open connection between them only when network administrator is accessing the router with password ciscotest
C. Router 1 and the TACACS+ server maintain one open connection between them
D. Router 1 opens and closes a TCP connection to the TACACS+ server every time a user requires authentication

Answer: C

Explanation:

Router 1 and the TACACS+ server maintain one open connection between them.

Question: 35

Which QoS model allows hosts to report their QoS needs to the network?

A. IntServ
B. CB-WFQ
C. DiffServ
D. MQC

Answer: A

Explanation:

The correct answer is A. IntServ (Integrated Services) is the QoS model where hosts signal their specific bandwidth and delay requirements to the network. This signaling occurs on a per-flow basis using protocols like RSVP (Resource Reservation Protocol). IntServ aims for strict end-to-end QoS guarantees by explicitly allocating resources for each flow, offering tight control over bandwidth and latency. In contrast, DiffServ (Differentiated Services) classifies traffic into different classes and provides varying levels of service based on these classes. CB-WFQ (Class-Based Weighted Fair Queuing) is a queuing mechanism that prioritizes traffic based on assigned weights but does not inherently involve host signaling. MQC (Modular Quality of Service CLI) is a command-line interface used for configuring QoS policies and is not a QoS model itself.

IntServ's per-flow resource reservation is its defining characteristic, aligning with the concept of hosts actively reporting needs, unlike the other options. It provides a more deterministic QoS compared to other options. Due to this specific mechanism of host signaling, IntServ directly addresses the question's requirement.

Authoritative Links:

Cisco's QoS Technologies - This Cisco documentation provides a comprehensive overview of QoS technologies, including IntServ, DiffServ, and queuing mechanisms.

RFC 1633 - Integrated Services in the Internet Architecture - This RFC provides a detailed explanation of the IntServ architecture and its related protocols.

Question: 36

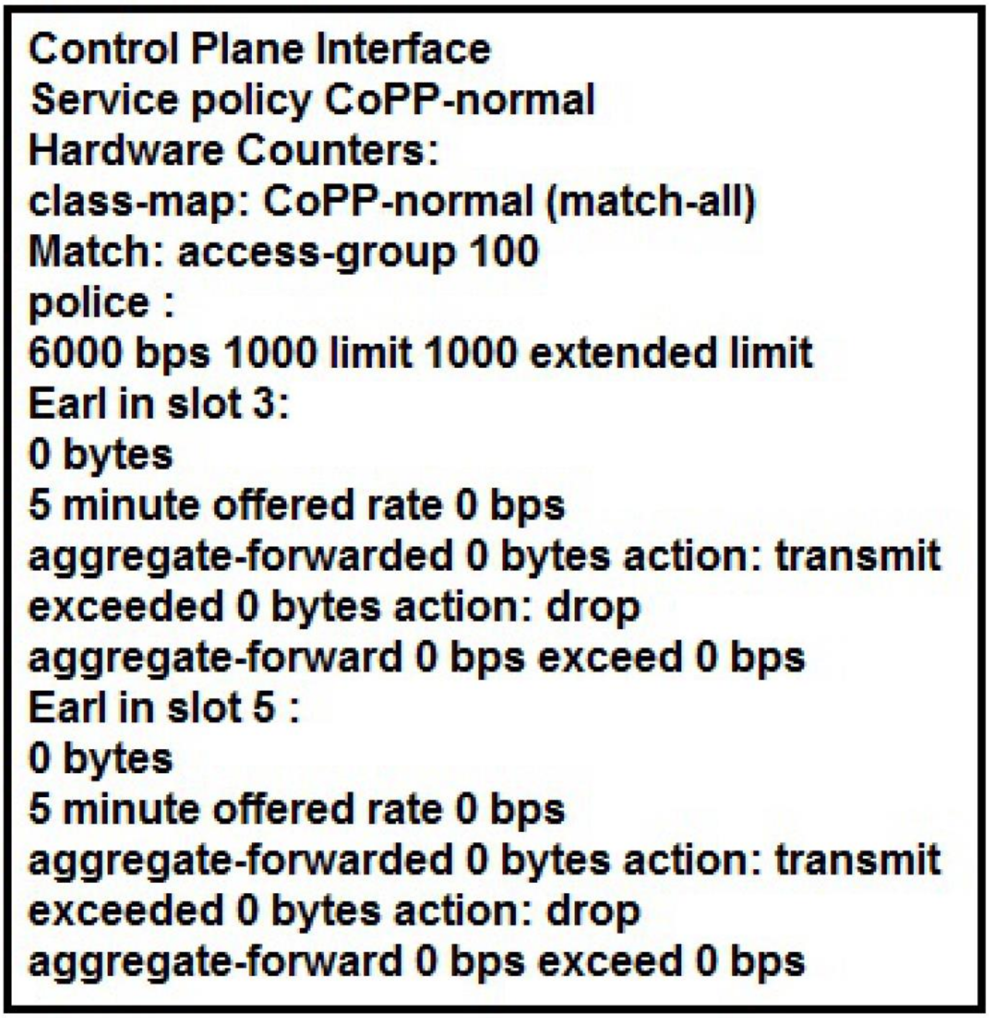

```
Control Plane Interface
Service policy CoPP-normal
Hardware Counters:
class-map: CoPP-normal (match-all)
Match: access-group 100
police :
6000 bps 1000 limit 1000 extended limit
Earl in slot 3:
0 bytes
5 minute offered rate 0 bps
aggregate-forwarded 0 bytes action: transmit
exceeded 0 bytes action: drop
aggregate-forward 0 bps exceed 0 bps
Earl in slot 5 :
0 bytes
5 minute offered rate 0 bps
aggregate-forwarded 0 bytes action: transmit
exceeded 0 bytes action: drop
aggregate-forward 0 bps exceed 0 bps
```

Refer to the exhibit. Which show command shows statistics for the control plane policy and is used to troubleshoot?

A. show control-plane CoPP
B. show policy control-plane
C. show control-plane
D. show policy-map control-plane

Answer: D

Explanation:

show policy-map control-plane.

Question: 37

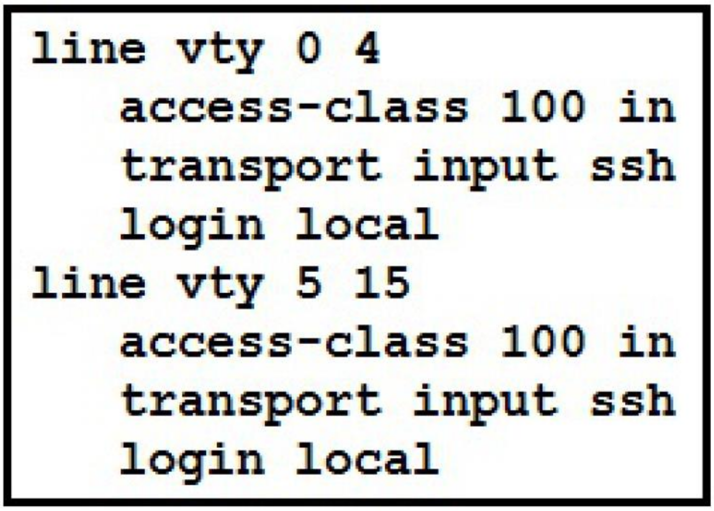

```
line vty 0 4
   access-class 100 in
   transport input ssh
   login local
line vty 5 15
   access-class 100 in
   transport input ssh
   login local
```

Refer to the exhibit. An engineer has started to configure a router for secure remote access as shown. All users who require network access need to be authenticated by the SSH protocol. Which two actions must the engineer implement to complete the SSH configuration? (Choose two.)

A. Configure an IP domain name.

B. Configure ACL 100 to permit access to port 22.

C. Configure a password under the vty lines.

D. Configure crypto keys.

E. Configure service password encryption.

Answer: AD

Explanation:

A. Configure an IP domain name.

D. Configure crypto keys.

Question: 38

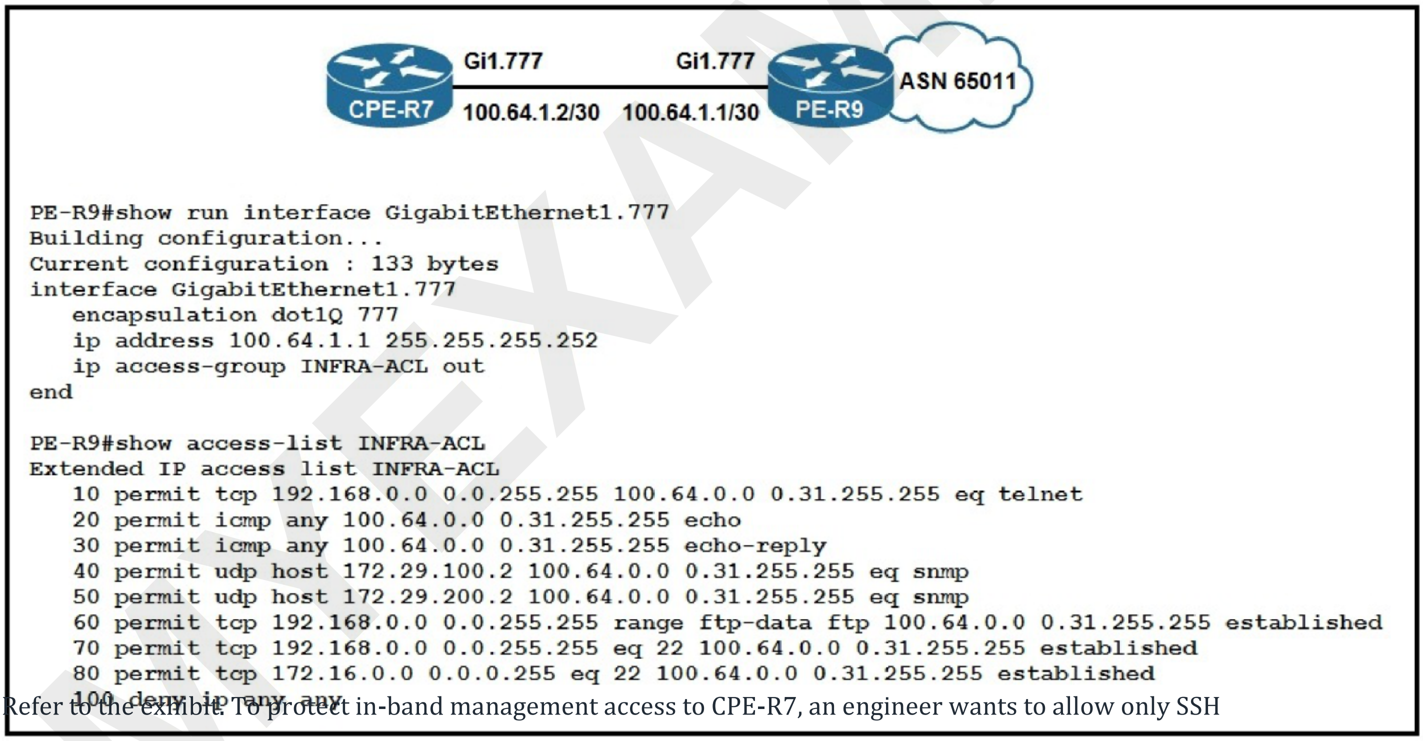

```
PE-R9#show run interface GigabitEthernet1.777
Building configuration...
Current configuration : 133 bytes
interface GigabitEthernet1.777
   encapsulation dot1Q 777
   ip address 100.64.1.1 255.255.255.252
   ip access-group INFRA-ACL out
end

PE-R9#show access-list INFRA-ACL
Extended IP access list INFRA-ACL
   10 permit tcp 192.168.0.0 0.0.255.255 100.64.0.0 0.31.255.255 eq telnet
   20 permit icmp any 100.64.0.0 0.31.255.255 echo
   30 permit icmp any 100.64.0.0 0.31.255.255 echo-reply
   40 permit udp host 172.29.100.2 100.64.0.0 0.31.255.255 eq snmp
   50 permit udp host 172.29.200.2 100.64.0.0 0.31.255.255 eq snmp
   60 permit tcp 192.168.0.0 0.0.255.255 range ftp-data ftp 100.64.0.0 0.31.255.255 established
   70 permit tcp 192.168.0.0 0.0.255.255 eq 22 100.64.0.0 0.31.255.255 established
   80 permit tcp 172.16.0.0 0.0.0.255 eq 22 100.64.0.0 0.31.255.255 established
   100 deny ip any any
```

Refer to the exhibit. To protect in-band management access to CPE-R7, an engineer wants to allow only SSH management and provisioning traffic from management network 192.168.0.0/16. Which infrastructure ACL change must be applied to router PE-R9 to complete this task?

A. ip access-list extended INFRA-ACL 15 permit tcp 192.168.0.0 0.0.255.255 range 49152 65535 100.64.0.0 0.31.255.255 eq 443

B. ip access-list extended INFRA-ACL no 10 15 permit tcp 192.168.0.0 0.0.255.255 range 49152 65535 100.64.0.0 0.31.255.255 eq 22

C. ip access-list extended INFRA-ACL 15 permit tcp 192.168.0.0 0.0.255.255 range 49152 65535 100.64.0.0 0.31.255.255 eq 22

D. ip access-list extended INFRA-ACL no 10 15 permit tcp 192.168.0.0 0.0.255.255 eq 22 100.64.0.0 0.31.255.255 eq 22

Answer: B

Explanation:

ip access-list extended INFRA-ACL no 10 15 permit tcp 192.168.0.0 0.0.255.255 range 49152 65535 100.64.0.0 0.31.255.255 eq 22.

Question: 39

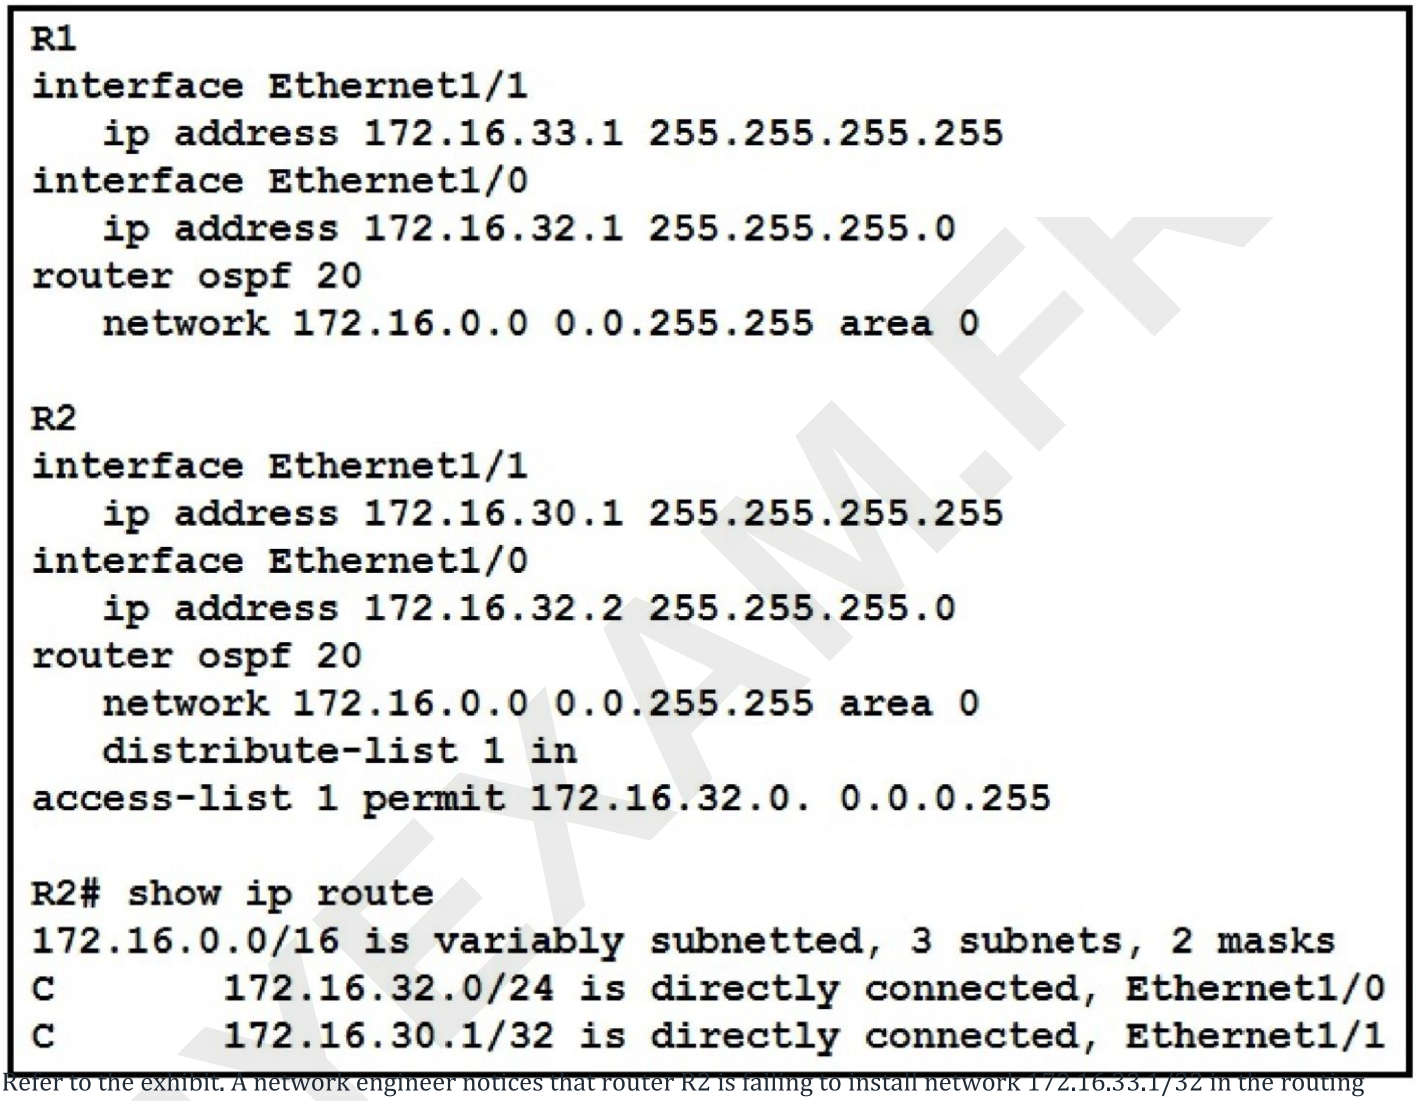

```
R1
interface Ethernet1/1
   ip address 172.16.33.1 255.255.255.255
interface Ethernet1/0
   ip address 172.16.32.1 255.255.255.0
router ospf 20
   network 172.16.0.0 0.0.255.255 area 0

R2
interface Ethernet1/1
   ip address 172.16.30.1 255.255.255.255
interface Ethernet1/0
   ip address 172.16.32.2 255.255.255.0
router ospf 20
   network 172.16.0.0 0.0.255.255 area 0
   distribute-list 1 in
access-list 1 permit 172.16.32.0. 0.0.0.255

R2# show ip route
172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C       172.16.32.0/24 is directly connected, Ethernet1/0
C       172.16.30.1/32 is directly connected, Ethernet1/1
```

Refer to the exhibit. A network engineer notices that router R2 is failing to install network 172.16.33.1/32 in the routing table. Which configuration must the engineer apply to R2 to fix the problem?

A. R2(config)# access-list 1 permit 172.16.33.0 0.0.0.255

B. R2(config)# access-list 1 permit 172.16.33.0 255.255.0.0

C. R2(config)# access-list 1 permit 172.16.33.0 255.255.255.0
D. R2(config)# access-list 1 permit 172.16.33.0 255.0.0.0

Answer: A

Explanation:

R2(config)# access-list 1 permit 172.16.33.0 0.0.0.255.

Question: 40

Which two PHY modes are available to implement an IOS XR 10 Gigabit Ethernet interface? (Choose two.)

A. LAN
B. SONET
C. MAN
D. WAN
E. WDWM

Answer: AD

Explanation:

The correct answer is A. LAN and D. WAN. IOS XR, Cisco's operating system for service provider routers, supports 10 Gigabit Ethernet interfaces operating in two primary physical (PHY) modes: LAN and WAN.

LAN PHY mode, also known as 10GBASE-R, is the standard Ethernet mode commonly used within local area networks. It utilizes a specific encoding scheme and frame format optimized for short-distance, high-bandwidth transmission. This mode is appropriate for connecting routers within a data center or other campus environments where distances are typically limited.

WAN PHY mode, or 10GBASE-W, is designed for longer-distance connections across wide area networks. It uses different encoding and framing mechanisms to make it compatible with SONET/SDH transport technologies, which are standard in traditional telco networks. This mode is generally used when connecting routers across greater distances where lower latency and specialized handling of optical signals are required.

Options B, SONET, and C, MAN, are related concepts but not PHY modes of an interface itself. SONET/SDH is a transport technology that WAN PHY mode is compatible with, not a mode itself. MAN (Metropolitan Area Network) is a network scale, not a PHY mode. Option E, WDWM, likely refers to Wavelength Division Multiplexing which is an optical multiplexing technique, also not a PHY mode itself. Therefore, only LAN and WAN correctly reflect the available PHY modes for 10 Gigabit Ethernet interfaces in Cisco IOS XR. The selection of a PHY mode depends on the network topology and the required transmission distance. LAN for closer proximity and WAN for longer distances requiring compatibility with long-haul telco technologies.

For further research on these topics, refer to Cisco's official documentation regarding interface configurations on IOS XR devices, specifically covering 10 Gigabit Ethernet interfaces and their related physical layer attributes:

Cisco IOS XR Interface and Hardware Component Configuration Guide
Understanding and Troubleshooting WAN PHY and LAN PHY

Question: 41

Which CLI mode must be used to configure the BGP keychain in Cisco IOS XR Software?

A. routing configuration mode
B. BGP neighbor configuration mode
C. global configuration mode
D. BGP address-family configuration mode

Answer: B

Explanation:

The correct answer is **B. BGP neighbor configuration mode**.

BGP keychains, used for MD5 authentication, are configured specifically for a BGP neighbor. This means the configuration isn't a global setting impacting all BGP processes, nor is it tied to a particular routing protocol's overall configuration (like routing configuration mode), or a specific address-family. Instead, authentication parameters, including the keychain, are specified on a per-neighbor basis to secure communication with that particular peer. Cisco IOS XR requires the neighbor <neighbor-ip> password <keychain-name> command under the BGP neighbor configuration to associate a keychain for authentication. This command syntax directly illustrates that authentication is configured within the scope of a neighbor. The keychain itself must also be pre-configured in the global config mode but not associated within it. Thus the specific keychain association to a neighbor is done in BGP neighbor configuration mode. This targeted approach ensures security policies are applied where needed, without unnecessarily broad scope. Configuring the keychain in the global or routing configuration modes would be inefficient, creating unnecessary complexity, and could lead to unintended authentication usage.

Therefore, BGP neighbor configuration mode is the only appropriate CLI context for associating a keychain to authenticate BGP peering sessions.

Relevant links:

Cisco's official documentation on BGP Keychain Configuration: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/16-12/irg-16-12-book/bgp-sec-md5-key.html
Cisco IOS XR BGP Configuration Guide: https://www.cisco.com/c/en/us/td/docs/iosxr/ip-routing/configuration/b-bgp-cg-ios-xr/b-bgp-cg-ios-xr_chapter_010.html

Question: 42

A remote operation center is deploying a set of I-BGP and E-BGP connections for multiple IOS-XR platforms using the same template. The I-BGP sessions exchange prefixes with no apparent issues, but the E-BGP sessions do not exchange routes. What causes this issue?

A. The I-BGP neighbors are mistyped and HELLO packets cannot be exchanged successfully between routers.
B. The E-BGP neighbors are not allowed to exchange Information due to the customer platform's default policy.
C. A PASS ALL policy has not been implemented for the I-BGP neighbors.
D. The next-hop-self command is not implemented on both E-BGP neighbors.

Answer: B

Explanation:

Here's a detailed justification for why option B is the correct answer:

The core issue revolves around the default behavior of network devices, specifically Cisco IOS-XR routers, and how they handle external BGP (E-BGP) peering. By default, many service provider platforms have policies that restrict the exchange of information with external peers, primarily for security reasons. These policies often involve access control lists (ACLs) or route filtering mechanisms that explicitly need to be configured to allow the establishment of E-BGP sessions and the propagation of routes.

Option A is incorrect because mistyped I-BGP neighbors would prevent I-BGP sessions from forming in the first place, which the problem statement explicitly says is not the case. I-BGP sessions are functioning correctly. Option C is also incorrect. A lack of a "pass all" policy would affect I-BGP sessions, which are operating successfully. A pass-all is not required, as long as the required prefixes are permitted. Option D is incorrect because next-hop-self is primarily concerned with how a router advertises routes to its peers, and has no impact on session establishment itself. This command would be needed in certain scenarios, but it would not cause a failure to establish the sessions.

Therefore, option B is the most plausible cause for the described issue. The E-BGP sessions are failing because the customer platform's default security posture is preventing the exchange of routes, a typical configuration to ensure that routers on the network do not inadvertently expose their routing information. The issue isn't about the session forming or the advertisement of the routes themselves but the actual process of receiving and importing the routes. The network administrator needs to explicitly configure policies to enable the E-BGP route exchange with the remote operation center. This could involve configuring an allow-all route policy or filtering prefixes using route-maps and ACLs. Without explicit configuration, the routes will be blocked, and E-BGP will not work.

Authoritative Links:

Cisco IOS XR BGP Configuration Guide: https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k_r6-3/routing/configuration/guide/b_routing_cg_63xasr9k/b_routing_cg_63xasr9k_chapter_011.html (Refer to the section on BGP policy and filtering)
BGP Best Practices:https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13701-best-practices-bgp.html (Provides general recommendations, including the importance of filtering). **Understanding BGP Route Policy and Filtering** : https://networklessons.com/bgp/bgp-route-policy-filtering (Provides understanding on BGP filtering and policy)

Question: 43

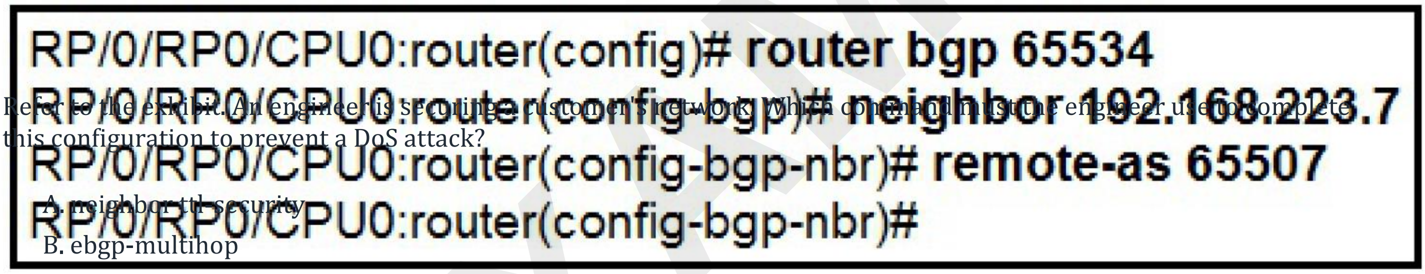

```
RP/0/RP0/CPU0:router(config)# router bgp 65534
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.223.7
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65507
RP/0/RP0/CPU0:router(config-bgp-nbr)#
```

Refer to the exhibit. An engineer is securing a customer's network. Which command must the engineer use to complete this configuration to prevent a DoS attack?

A. neighbor ttl-security

B. ebgp-multihop

C. neighbor ebgp-multihop

D. ttl-security

Answer: D

Explanation:

Correct answer is D: ttl-security.

Question: 44

What is the function of the FEC field within the OTN signal structure?

A. It allows the sending devices to apply QoS within the OTN forwarding structure.

B. It allows deep inspection of data payload fields.

C. It allows receivers to correct errors upon data arrival

D. It allows source nodes to discard payload errors before transmitting data on the network.

Answer: C

Explanation:

The correct answer is C, stating that the FEC (Forward Error Correction) field within the Optical Transport Network (OTN) signal structure enables receivers to correct errors upon data arrival. This is fundamentally the primary purpose of FEC. In the context of high-speed optical networks, data transmitted over long distances or through noisy channels is prone to bit errors. The FEC field contains redundant information calculated based on the transmitted data. When a receiver encounters corrupted data, it uses this redundant information to identify and rectify errors without needing retransmission, enhancing the reliability and integrity of the transmission. The FEC calculation happens at the transmitter, while the error correction process occurs at the receiver. This capability directly combats issues like signal degradation and interference, ensuring successful delivery of data even in less-than-ideal conditions. Other options are incorrect because FEC does not manage QoS (Option A), perform deep packet inspection (Option B), or discard errors before transmission (Option D).

Instead, its role is specifically in error correction at the receiving end. The core concept behind FEC aligns with reliability principles in cloud networking where error-free transmission is crucial for applications and services. Cloud providers rely heavily on robust transport mechanisms like OTN with FEC to guarantee network performance and user experience.

For further research, consider exploring the following links:

1. **ITU-T G.709 - Interfaces for the Optical Transport Network (OTN):**https://www.itu.int/rec/T-REC-G.709/en This is the foundational standard for OTN and provides in-depth details on the signal structure and FEC mechanisms.
2. **Cisco - Understanding OTN with FEC:** https://www.cisco.com/c/en/us/td/docs/optical/transceiver_modules/100G/100g_otn.html This link provides a good explanation of the use of FEC in the context of Cisco's optical transport solutions. 3. **Wikipedia - Forward Error Correction:**https://en.wikipedia.org/wiki/Forward_error_correction This is a general resource providing a good overview of different types of FEC and their uses.

Question: 45

A customer of an ISP requests support to setup a BGP routing policy. Which BGP attribute should be configured to choose specific BGP speakers as preferred points for the customer AS?

A. lowest multi-exit discriminator

B. highest local preference outbound

C. lowest local preference inbound

D. highest local preference inbound

Answer: D

Explanation:

The correct answer is **D. highest local preference inbound.** Here's why:

Local Preference (LOCAL_PREF) is a BGP attribute used within an Autonomous System (AS) to influence the path selection process for outbound traffic. It indicates the degree of preference given to a particular route. A higher local preference value signifies a more preferred route. When a BGP speaker receives multiple paths to the same destination, it prioritizes the path with the highest local preference. This attribute is not advertised to external ASes, so it's relevant for path selection within the AS.

In the scenario, the customer desires specific BGP speakers as their preferred entry points into their AS. They can signal this preference to their upstream ISP by setting the local preference attribute on received routes from their preferred speakers, thereby ensuring outbound traffic from the customer AS will use that path. Higher values of LOCAL_PREF will be preferred.

Multi-Exit Discriminator (MED) is used between ASes to indicate preference of paths, but is not relevant for the customer preference, which is internal to its AS. Local preference is specific to the router receiving and learning the routes, thus it would need to be configured inbound. Lower values of LOCAL_PREF are not preferred and would not influence the routing in the direction the customer is requesting.

Therefore, configuring the highest local preference inbound on the customer's AS will achieve the desired outcome of selecting preferred BGP speakers.

Authoritative Links:

Cisco BGP Attributes:https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13704-21.html
Understanding BGP Path Selection:https://www.juniper.net/documentation/en_US/junos/topics/topic-map/bgp-path-selection.html
RFC 4271 - A Border Gateway Protocol 4 (BGP-4):https://datatracker.ietf.org/doc/html/rfc4271

Question: 46

Which three OSPF parameters must match before two devices can establish an OSPF adjacency? (Choose three.)

A. IP address
B. subnet mask
C. interface cost
D. process ID
E. area number
F. hello timer setting

Answer: BEF

Explanation:

Here's a breakdown of why options B (subnet mask), E (area number), and F (hello timer setting) are the correct parameters that must match for OSPF adjacency formation, along with why the others are incorrect:

OSPF (Open Shortest Path First) relies on precise neighbor discovery mechanisms. To form an adjacency, routers must agree on certain parameters, signaling they belong to the same OSPF network segment. Subnet masks (B) are critical because they define the network to which an interface belongs. If the masks mismatch, routers cannot understand they are on the same network. Area number (E) dictates the OSPF area a router belongs to. Routers within the same area form adjacencies; if area numbers don't match, adjacent formation fails. Hello timers (F) govern the frequency OSPF hello packets are exchanged; mismatches prevent routers from detecting each other.

IP addresses (A) are unique identifiers and don't directly affect adjacency formation; OSPF instead focuses on network relationships. The interface cost (C) parameter affects path calculations post-adjacency, but does not prevent initial adjacencies from forming. OSPF Process IDs (D) are locally significant and don't need to match for adjacency formation; they are used to differentiate multiple OSPF processes on a single router.

Therefore, consistent subnet masks, area numbers, and hello timer settings are fundamental for two routers to discover and establish an OSPF adjacency, enabling the reliable exchange of routing information. These parameters ensure seamless routing within the OSPF network.

Authoritative Links for further research:

Cisco's OSPF documentation: https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html
RFC 2328, OSPF Version 2: https://www.rfc-editor.org/rfc/rfc2328

Question: 47

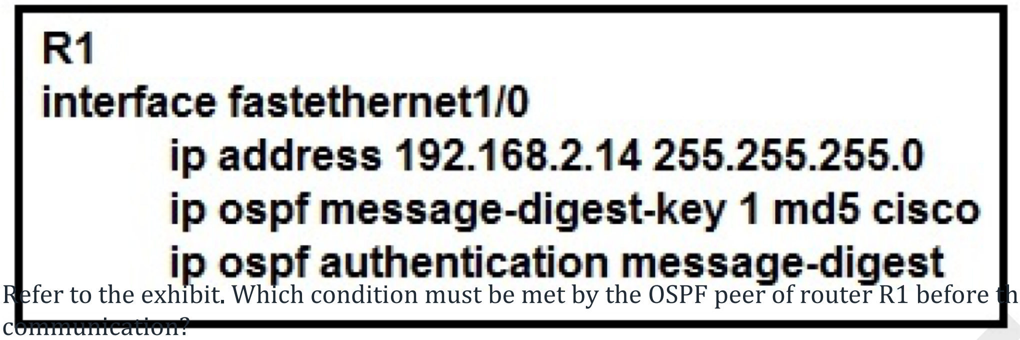

```
R1
interface fastethernet1/0
        ip address 192.168.2.14 255.255.255.0
        ip ospf message-digest-key 1 md5 cisco
        ip ospf authentication message-digest
```

Refer to the exhibit. Which condition must be met by the OSPF peer of router R1 before the two devices can establish communication?

A. The OSPF peer must use clear-text authentication.

B. The OSPF peer must be configured as an OSPF stub router.

C. The interface on the OSPF peer may have a different key ID, but it must use the same key value as the configured interface.

D. The interface on the OSPF peer must use the same key ID and key value as the configured interface.

Answer: D

Explanation:

Reference:
https://networklessons.com/ospf/how-to-configure-ospf-md5-authentication

Question: 48

DRAG DROP -
Drag and drop the OSPF area types from the left onto the correct statements on the right.
Select and Place:

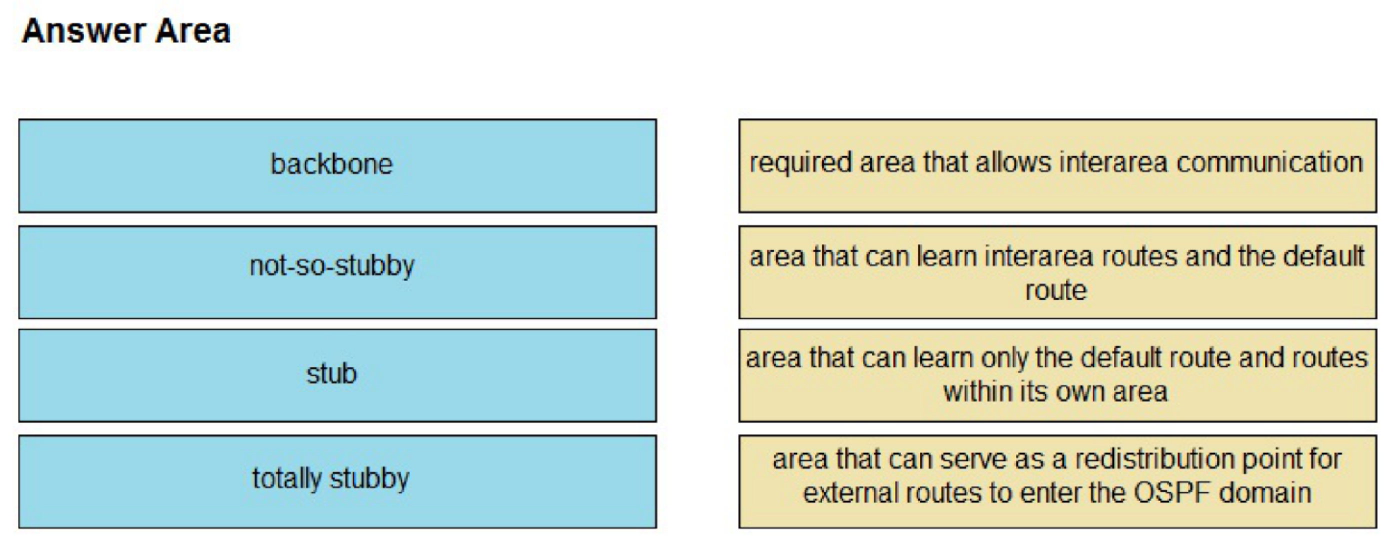

Answer Area

backbone	required area that allows interarea communication
not-so-stubby	area that can learn interarea routes and the default route
stub	area that can learn only the default route and routes within its own area
totally stubby	area that can serve as a redistribution point for external routes to enter the OSPF domain

Answer:

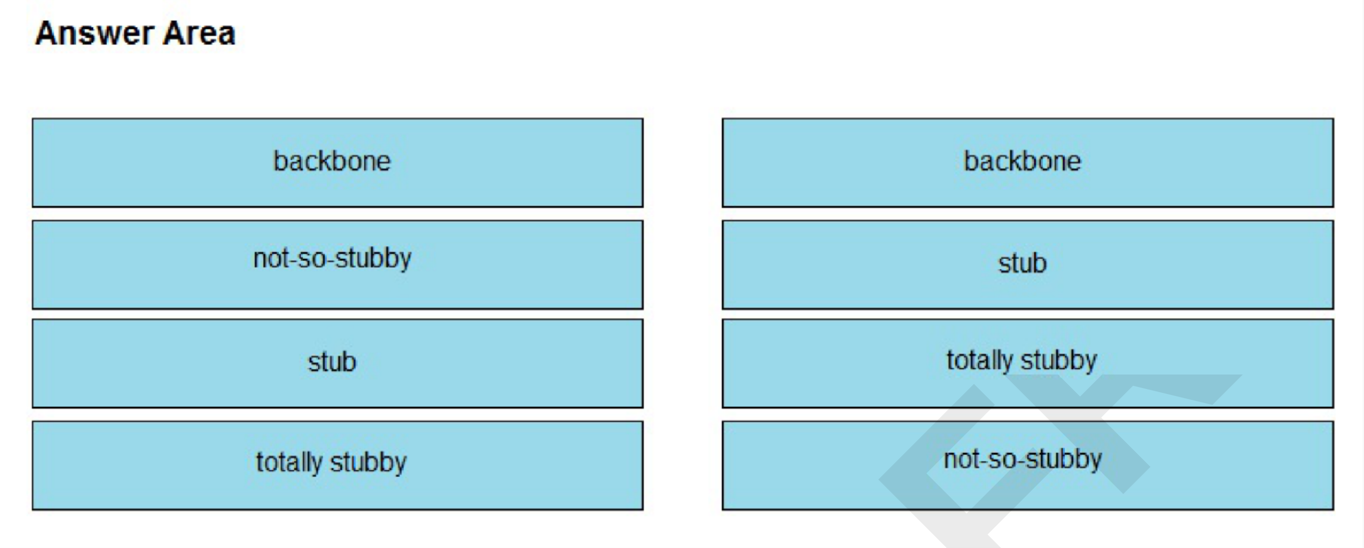

Answer Area

backbone	backbone
not-so-stubby	stub
stub	totally stubby
totally stubby	not-so-stubby

Explanation:

Backbone.

Stub.

Totally stubby.

not-so Stubby.

Question: 49

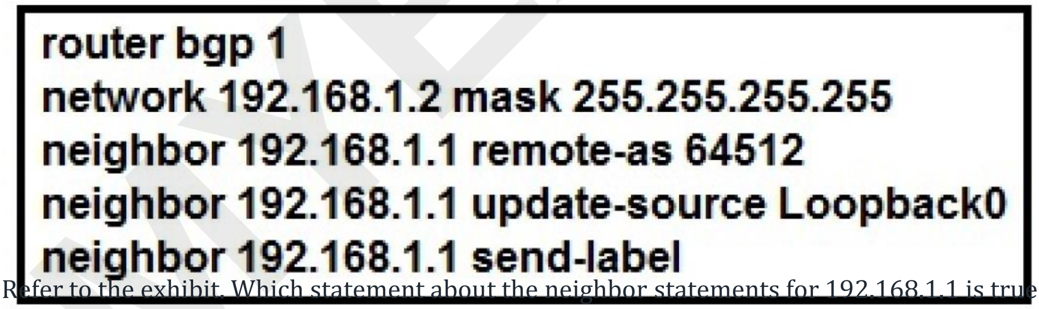

```
router bgp 1
network 192.168.1.2 mask 255.255.255.255
neighbor 192.168.1.1 remote-as 64512
neighbor 192.168.1.1 update-source Loopback0
neighbor 192.168.1.1 send-label
```

Refer to the exhibit. Which statement about the neighbor statements for 192.168.1.1 is true?

A. The router sends BGP labels for its prefixes to this peer.

B. The router must have TDP configured for the send-label command to operate. C. The neighbor router receives at least four labels from this router.

D. The router sends only a label for the prefix for Loopback0.

Answer: A

Explanation:

The router sends BGP labels for its prefixes to this peer.

Question: 50

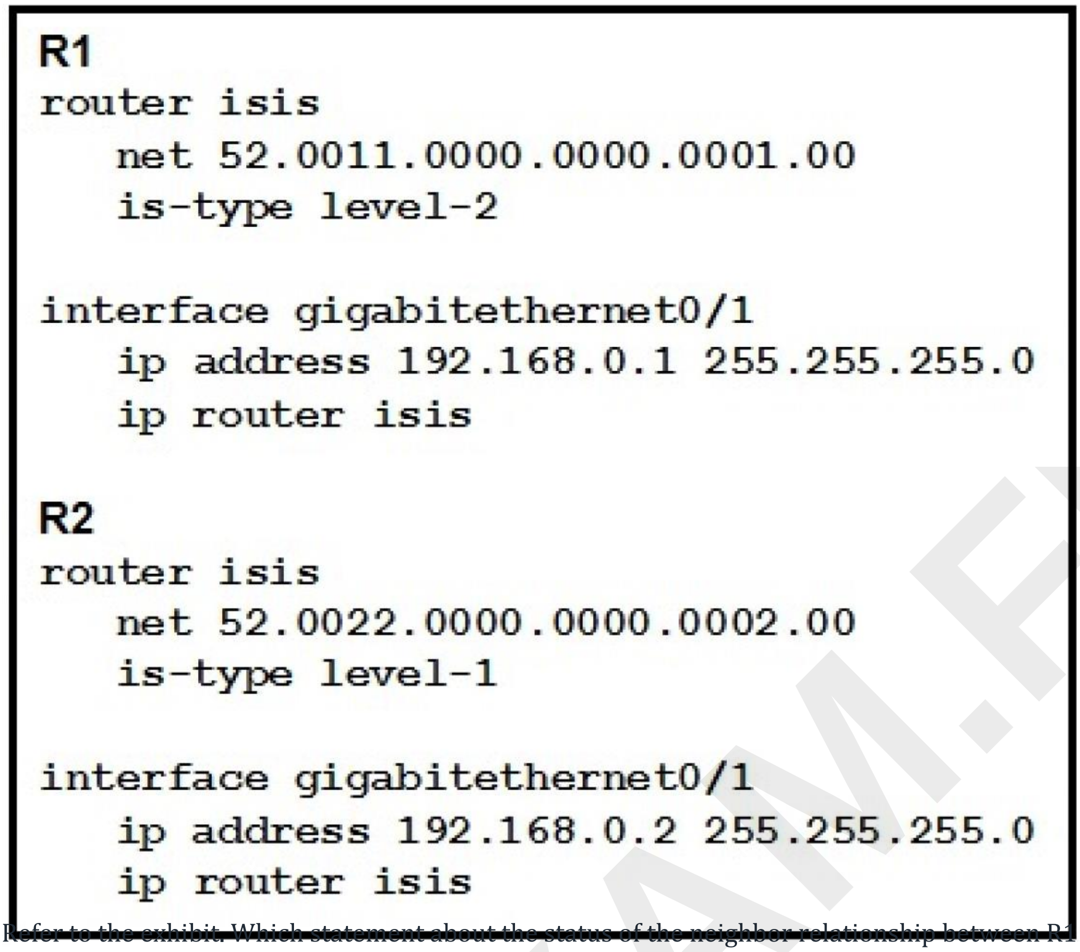

```
R1
router isis
   net 52.0011.0000.0000.0001.00
   is-type level-2

interface gigabitethernet0/1
   ip address 192.168.0.1 255.255.255.0
   ip router isis

R2
router isis
   net 52.0022.0000.0000.0002.00
   is-type level-1

interface gigabitethernet0/1
   ip address 192.168.0.2 255.255.255.0
   ip router isis
```

Refer to the exhibit. Which statement about the status of the neighbor relationship between R1 and R2 is true?

A. The neighbor relationship is down because the two routers are configured with different area types. B. The neighbor relationship is down because the two routers are in the same subnet.

C. The neighbor relationship is up because R2 is level 1 and Ievel 2 router.

D. The neighbor relationship is down because R2 is operating as a Level 1 router and the two routers are in different areas.

Answer: D

Explanation:

The neighbor relationship is down because R2 is operating as a Level 1 router and the two routers are in different areas.

Question: 51

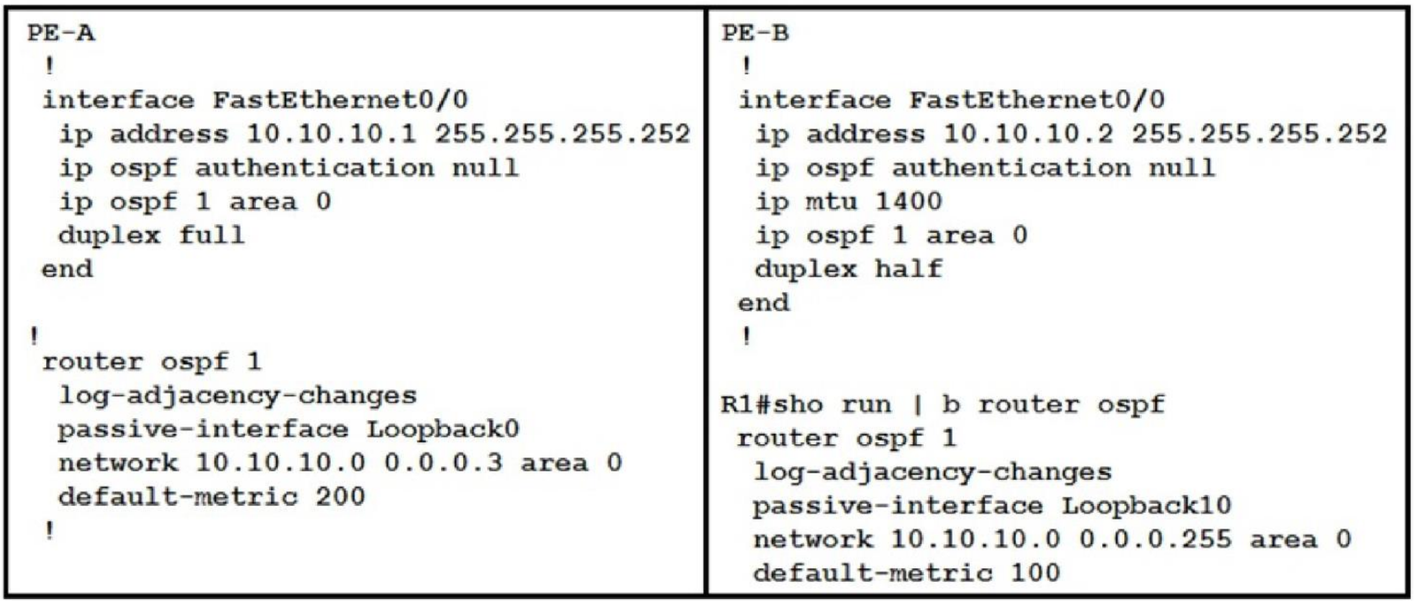

```
PE-A
 !
 interface FastEthernet0/0
  ip address 10.10.10.1 255.255.255.252
  ip ospf authentication null
  ip ospf 1 area 0
  duplex full
 end

!
 router ospf 1
  log-adjacency-changes
  passive-interface Loopback0
  network 10.10.10.0 0.0.0.3 area 0
  default-metric 200
 !
```

```
PE-B
 !
 interface FastEthernet0/0
  ip address 10.10.10.2 255.255.255.252
  ip ospf authentication null
  ip mtu 1400
  ip ospf 1 area 0
  duplex half
 end
 !

R1#sho run | b router ospf
 router ospf 1
  log-adjacency-changes
  passive-interface Loopback10
  network 10.10.10.0 0.0.0.255 area 0
  default-metric 100
```

Refer to the exhibit. Which configuration prevents the OSPF neighbor from establishing?

A. default-metric
B. duplex
C. network statement
D. mtu

Answer: D

Explanation:

Correct answer is D:mtu.

Question: 52

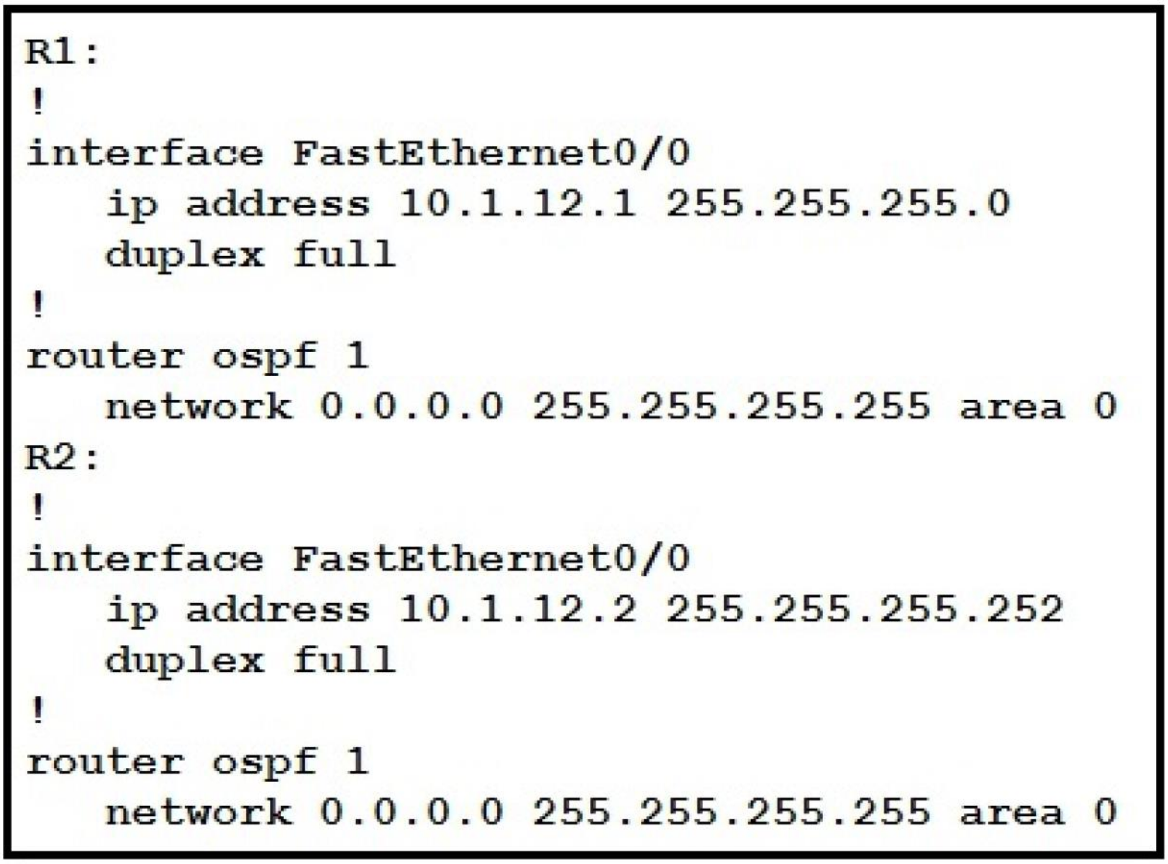

```
R1:
!
interface FastEthernet0/0
   ip address 10.1.12.1 255.255.255.0
   duplex full
!
router ospf 1
   network 0.0.0.0 255.255.255.255 area 0
R2:
!
interface FastEthernet0/0
   ip address 10.1.12.2 255.255.255.252
   duplex full
!
router ospf 1
   network 0.0.0.0 255.255.255.255 area 0
```

Refer to the exhibit. R1 and R2 are directly connected with Fast Ethernet interfaces and have the above configuration applied OSPF adjacency is not formed.
When the debug ip ospf hello command is issued on R1, these log messages are seen:

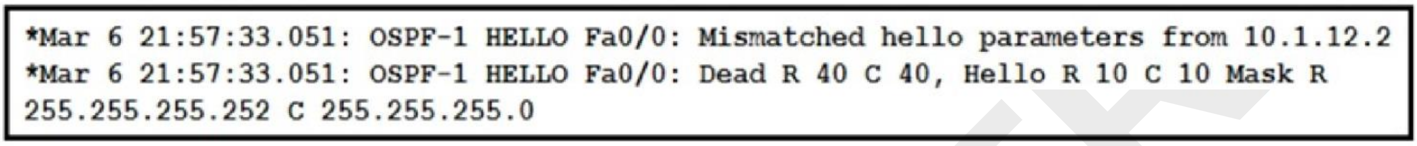

```
*Mar 6 21:57:33.051: OSPF-1 HELLO Fa0/0: Mismatched hello parameters from 10.1.12.2
*Mar 6 21:57:33.051: OSPF-1 HELLO Fa0/0: Dead R 40 C 40, Hello R 10 C 10 Mask R
255.255.255.252 C 255.255.255.0
```

Which command can be configured on routers R1 and R2 on f0/0 interfaces to form OSPF adjacency?

A. ip ospf network point-to-multipoint non-broadcast
B. ip ospf network non-broadcast
C. ip ospf network broadcast
D. ip ospf network point-to-point

Answer: D

Explanation:

Reference:
https://community.cisco.com/t5/routing/ospf-point-to-point-links/td-p/1913398

Question: 53

Which two tasks must you perform when you implement LDP NSF on your network? (Choose two.)

A. Enable NSF for BGP.
B. Implement direct connections for LDP peers.
C. Enable NSF for EIGRP.
D. Disable Cisco Express Forwarding.
E. Enable NSF for the link-state routing protocol that is in use on the network.

Answer: AE

Explanation:

The correct answer is A and E. LDP Non-Stop Forwarding (NSF) relies on the underlying routing protocol's NSF capabilities for a graceful restart during a route processor switchover. Option A, "Enable NSF for BGP", is necessary because BGP often carries the VPNv4 routes used in MPLS networks that LDP might be using to establish adjacencies and label distribution. If BGP does not support NSF, then an LDP switchover might disrupt the VPN service. Option E, "Enable NSF for the link-state routing protocol that is in use on the network," is also crucial because LDP uses the underlying link-state routing protocol (such as OSPF or IS-IS) to discover peers. If the link-state routing protocol is not NSF-aware, the LDP session will be lost. Option B, "Implement direct connections for LDP peers," is not strictly required for LDP NSF, as LDP sessions can be established over routed links. Option C, "Enable NSF for EIGRP," is incorrect because EIGRP is not a common routing protocol used for service provider core networks where LDP is deployed. Option D, "Disable Cisco Express Forwarding (CEF)," is incorrect as CEF is required for the operation of LDP and therefore also for LDP NSF. LDP NSF depends on the ability to maintain forwarding table information during a route processor failure, and CEF plays a crucial role in providing this forwarding capability. Therefore, both the BGP and link-state routing protocol must be NSF-capable for LDP NSF to work effectively.

Further research:

1. Cisco Documentation on LDP NSF: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mpls/configuration/15-mt/mp-15-mt-book/mp-ldp-nsf.html
2. Understanding LDP NSF: https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/label-distribution-protocol-ldp/117482-technote-ldp-00.html

Question: 54

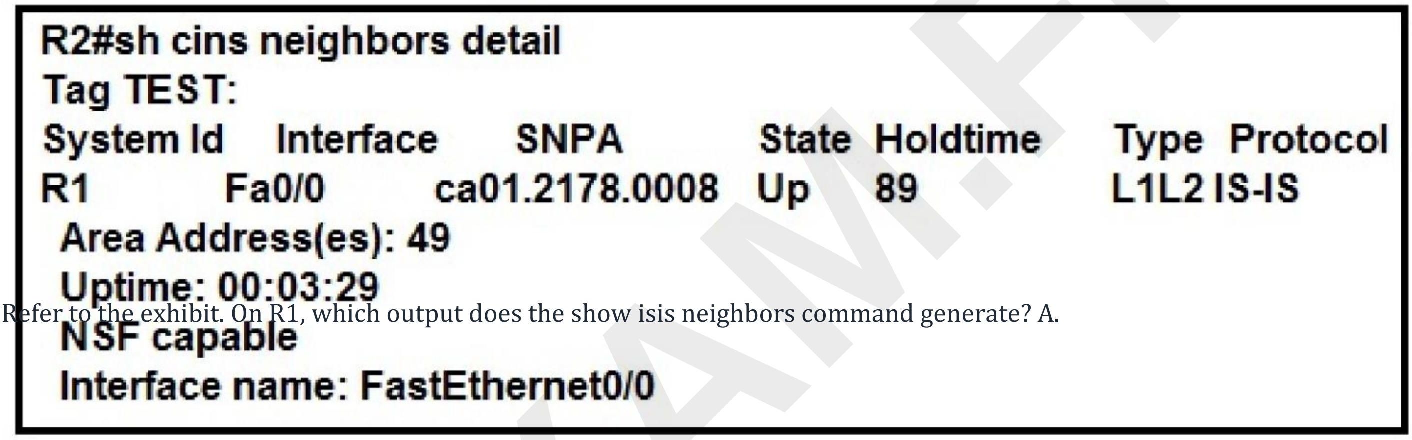

```
R2#sh cins neighbors detail
Tag TEST:
System Id    Interface     SNPA            State  Holdtime    Type  Protocol
R1           Fa0/0         ca01.2178.0008  Up     89          L1L2  IS-IS
  Area Address(es): 49
  Uptime: 00:03:29
  NSF capable
  Interface name: FastEthernet0/0
```

Refer to the exhibit. On R1, which output does the show isis neighbors command generate? A.

B.

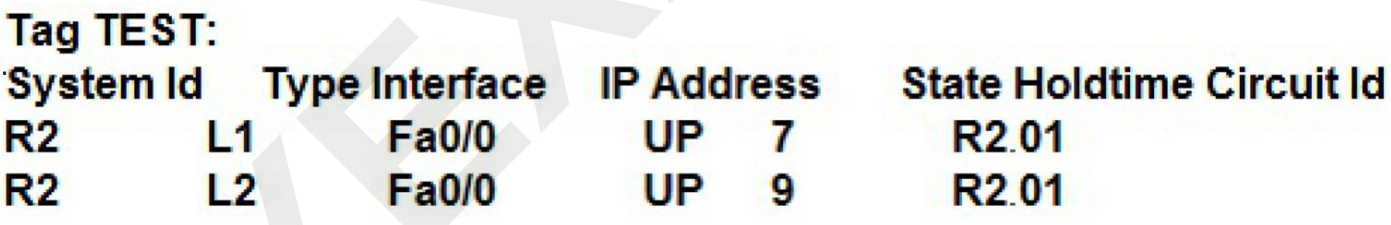

```
Tag TEST:
System Id    Type Interface    IP Address      State Holdtime Circuit Id
R2           L1        Fa0/0         UP      7        R2.01
R2           L2        Fa0/0         UP      9        R2.01
```

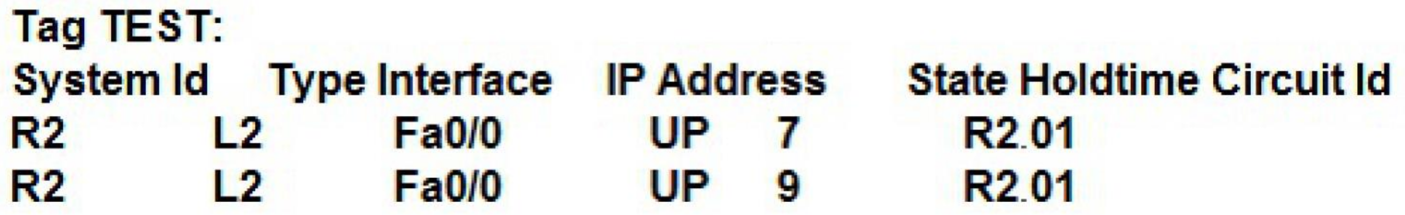

```
Tag TEST:
System Id    Type Interface    IP Address     State Holdtime Circuit Id
R2           L2        Fa0/0       UP     7        R2.01
R2           L2        Fa0/0       UP     9        R2.01
```

C.

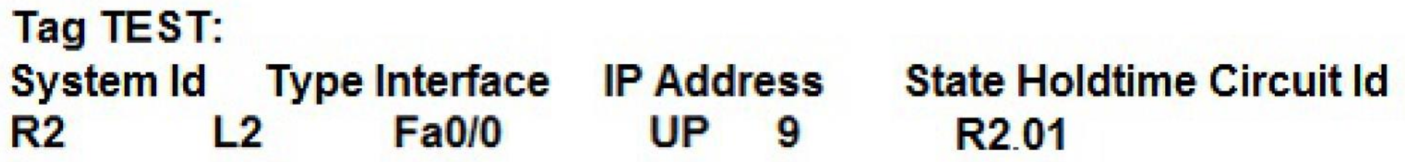

```
Tag TEST:
System Id    Type Interface    IP Address     State Holdtime Circuit Id
R2           L2        Fa0/0       UP     9       R2.01
```

D.

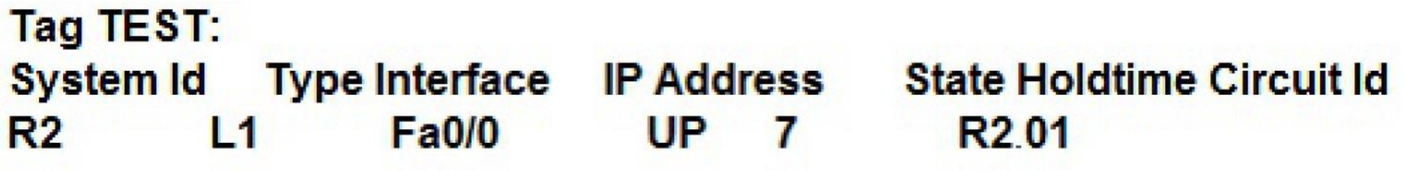

```
Tag TEST:
System Id    Type Interface    IP Address     State Holdtime Circuit Id
R2           L1        Fa0/0       UP     7       R2.01
```

Answer: A

Explanation:

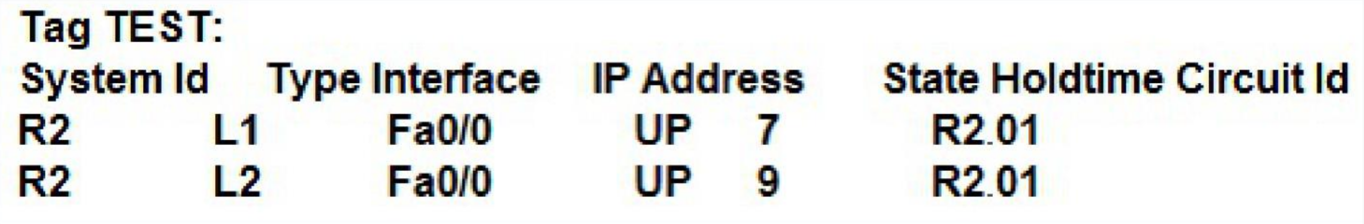

```
Tag TEST:
System Id    Type Interface    IP Address     State Holdtime Circuit Id
R2           L1        Fa0/0       UP     7       R2.01
R2           L2        Fa0/0       UP     9       R2.01
```

Question: 55

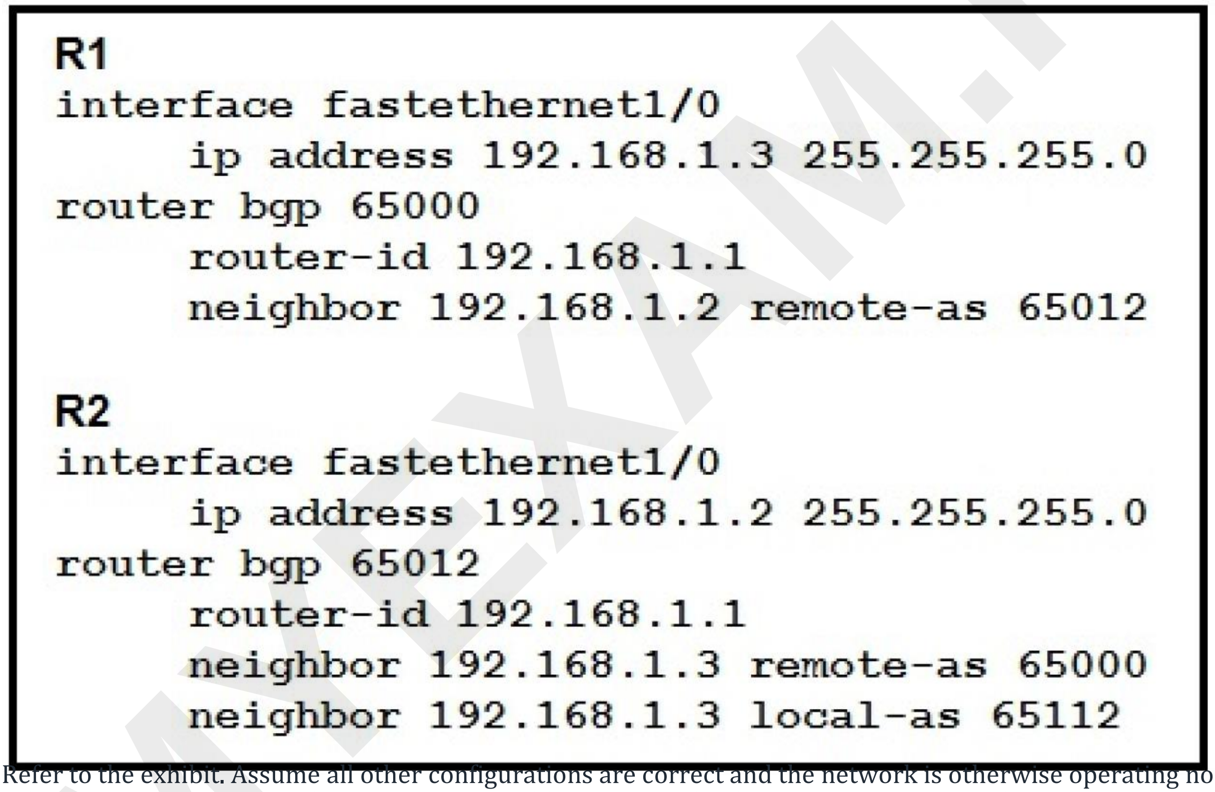

```
R1
interface fastethernet1/0
     ip address 192.168.1.3 255.255.255.0
router bgp 65000
     router-id 192.168.1.1
     neighbor 192.168.1.2 remote-as 65012

R2
interface fastethernet1/0
     ip address 192.168.1.2 255.255.255.0
router bgp 65012
     router-id 192.168.1.1
     neighbor 192.168.1.3 remote-as 65000
     neighbor 192.168.1.3 local-as 65112
```

Refer to the exhibit. Assume all other configurations are correct and the network is otherwise operating normally.
Which conclusion can you draw about the neighbor relationship between routers R1 and R2?

A. The neighbor relationship is up.

B. The neighbor relationship will be up only if the two devices have activated the correct neighbor relationships under the IPv4 address family.

C. The neighbor is down because the local-as value for R2 is missing in the R1 neighbor statement.

D. The neighbor relationship is down because R1 believes R2 is in AS 65012.

Answer: D

Explanation:

The neighbor relationship is down because R1 believes R2 is in AS 65012.

Question: 56

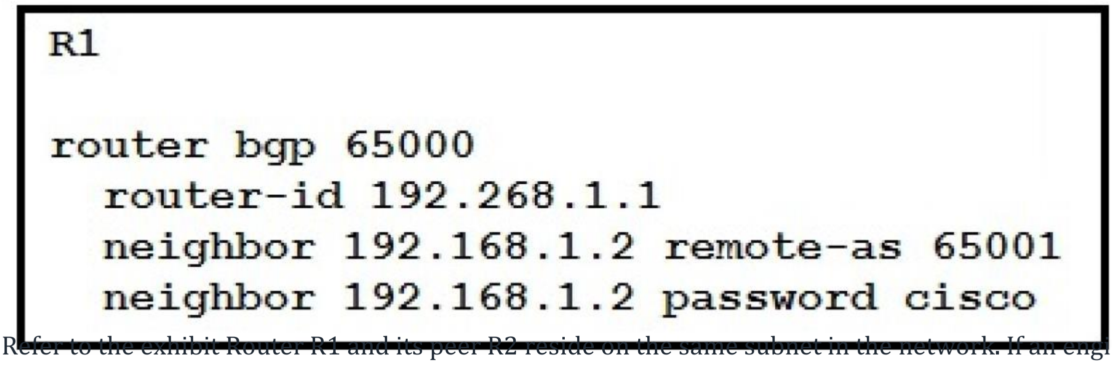

```
R1

router bgp 65000
  router-id 192.268.1.1
  neighbor 192.168.1.2 remote-as 65001
  neighbor 192.168.1.2 password cisco
```

Refer to the exhibit Router R1 and its peer R2 reside on the same subnet in the network. If an engineer implements this configuration to R1, how does it make connections to R2?

A. R1 establishes TCP connections that are authenticated with a clear-text password.

B. R1 establishes UDP connections that are authenticated with an MD5 password. C. R1 establishes UDP connections that are authenticated with a clear-text password.

D. R1 establishes TCP connections that are authenticated with an MD5 password.

Answer: D

Explanation:

R1 establishes TCP connections that are authenticated with an MD5 password.

Question: 57

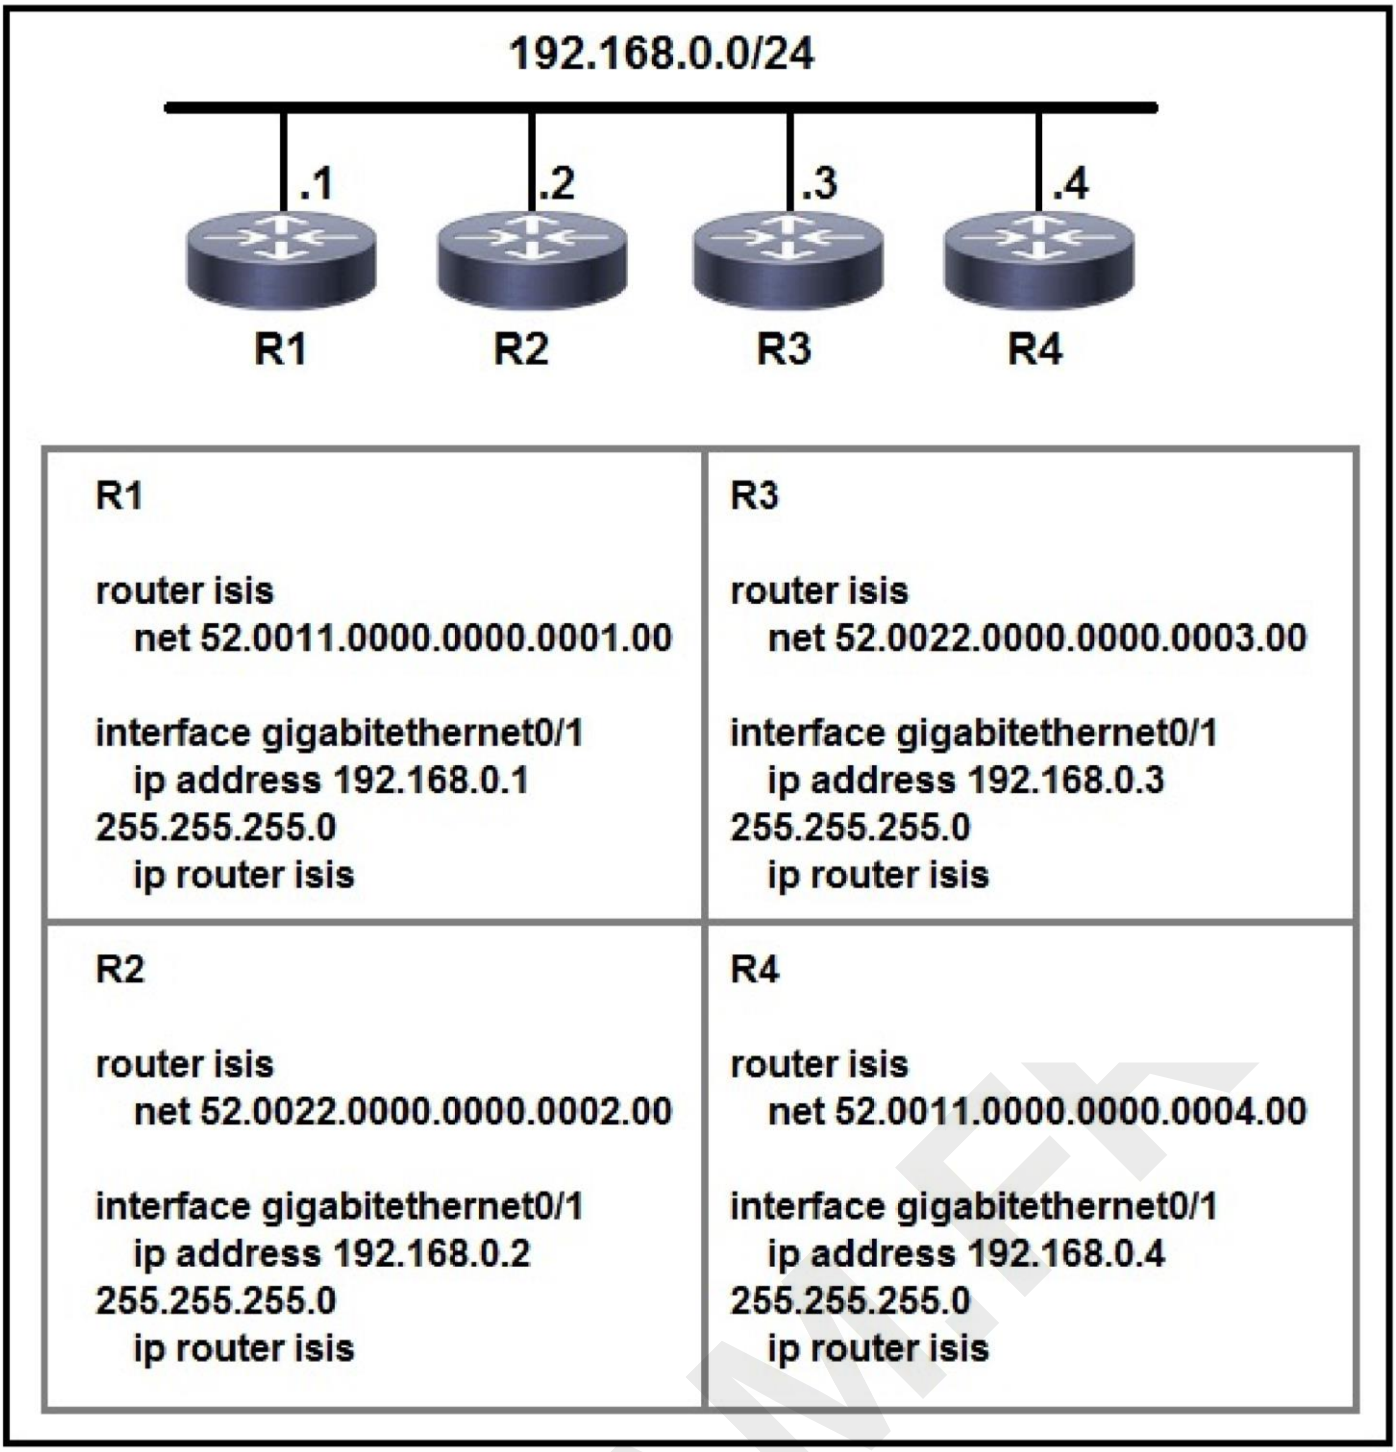

```
R1

router isis
   net 52.0011.0000.0000.0001.00

interface gigabitethernet0/1
   ip address 192.168.0.1
255.255.255.0
   ip router isis
```

```
R3

router isis
   net 52.0022.0000.0000.0003.00

interface gigabitethernet0/1
   ip address 192.168.0.3
255.255.255.0
   ip router isis
```

```
R2

router isis
   net 52.0022.0000.0000.0002.00

interface gigabitethernet0/1
   ip address 192.168.0.2
255.255.255.0
   ip router isis
```

```
R4

router isis
   net 52.0011.0000.0000.0004.00

interface gigabitethernet0/1
   ip address 192.168.0.4
255.255.255.0
   ip router isis
```

Refer to the exhibit. Which two statements about the IS-IS topology are true? (Choose two.)

A. R1 and R4 are Level 2 neighbors.

B. All four routers are operating as Level 1-2 routers.
C. All four routers are operating as Level 2 routers only.
D. All four routers are operating as Level 1 routers only.
E. R1 and R2 are Level 2 neighbors.

Answer: BE

Explanation:

B. All four routers are operating as Level 1-2 routers.

E. R1 and R2 are Level 2 neighbors.

Question: 58

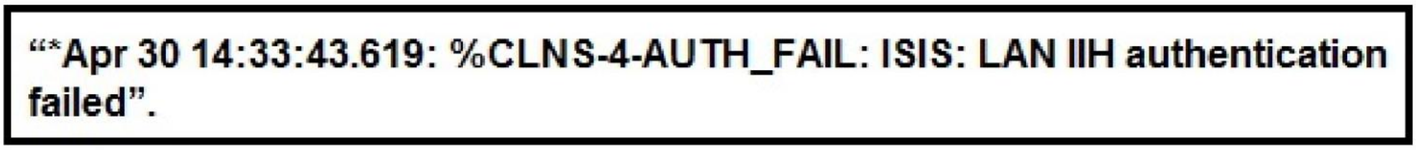

```
"**Apr 30 14:33:43.619: %CLNS-4-AUTH_FAIL: ISIS: LAN IIH authentication
failed".
```

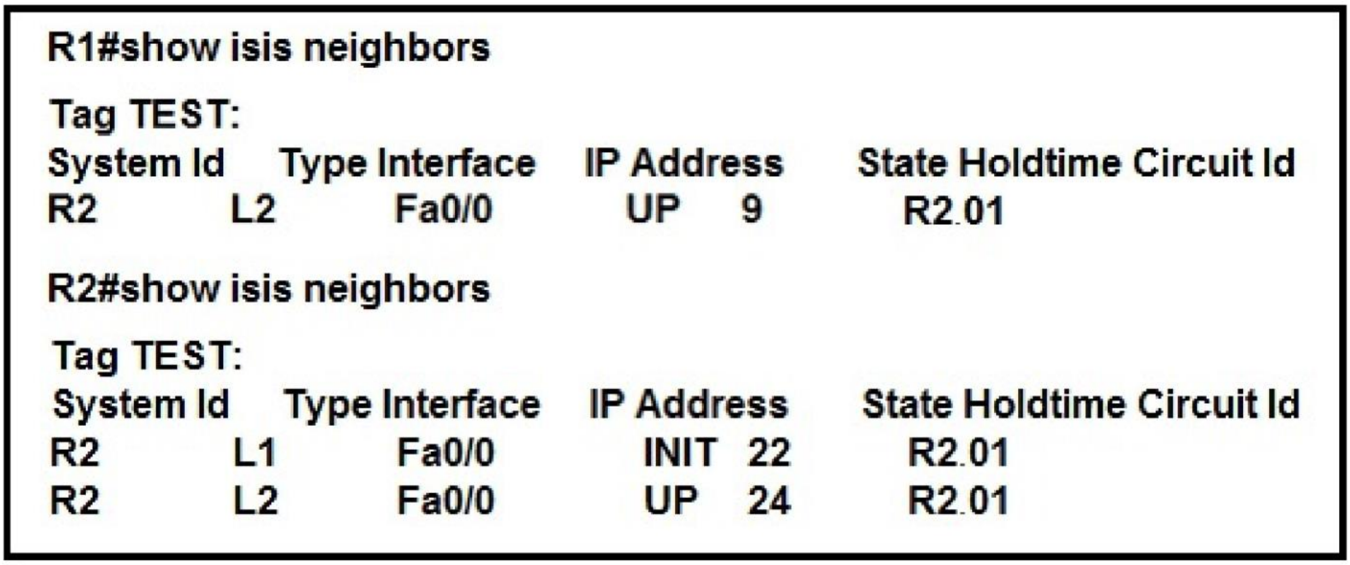

```
R1#show isis neighbors

Tag TEST:
System Id    Type Interface    IP Address     State Holdtime Circuit Id
R2           L2         Fa0/0       UP    9        R2.01

R2#show isis neighbors

Tag TEST:
System Id    Type Interface    IP Address     State Holdtime Circuit Id
R2           L1         Fa0/0       INIT  22       R2.01
R2           L2         Fa0/0       UP    24       R2.01
```

Refer to the exhibits. R1 and R2 are directly connected and IS-IS routing has been enabled between R1 and R2. R1 generates the above log message periodically.
Based on this output, which statement is true?

A. IS-IS neighbor authentication is failing for Level 2 PDUs only.

B. IS-IS neighbor authentication is failing for Level 2 first and then for Level 1 PDUs. C. IS-IS neighbor authentication is failing for Level 1 and Level 2 PDUs.

D. IS-IS neighbor authentication is failing for Level 1 PDUs only.

Answer: D

Explanation:

IS-IS neighbor authentication is failing for Level 1 PDUs only.

Question: 59

Which BGP attribute is used first when determining the best path?

A. origin
B. AS path
C. local preference
D. weight

Answer: D

Explanation:

The BGP (Border Gateway Protocol) path selection process is a multi-step algorithm to determine the best route for a destination. Weight is the first attribute evaluated in this process. Cisco devices prioritize routes with a higher weight value. Weight is a Cisco-specific attribute, meaning it is locally significant to the router where it's configured and not propagated to other routers. If multiple paths exist to the same destination, the path with the highest weight is chosen, regardless of other attributes. If weights are equal, the algorithm proceeds to subsequent attributes like local preference, shortest AS path, origin, MED (Multi-Exit Discriminator) and finally, the IGP cost to the next-hop. Local preference, while influential, is not the first selection criterion. Origin refers to how the route was injected into BGP and AS path denotes the sequence of AS numbers traversed. These are lower in priority than weight in Cisco's BGP best-path selection algorithm.

Therefore, the highest weight is the deciding factor first when choosing between multiple BGP paths.

For further information, consult these resources:

Cisco Documentation on BGP Path Selection: https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html
BGP Best Path Selection Algorithm: https://networklessons.com/bgp/bgp-path-selection-algorithm

Question: 60

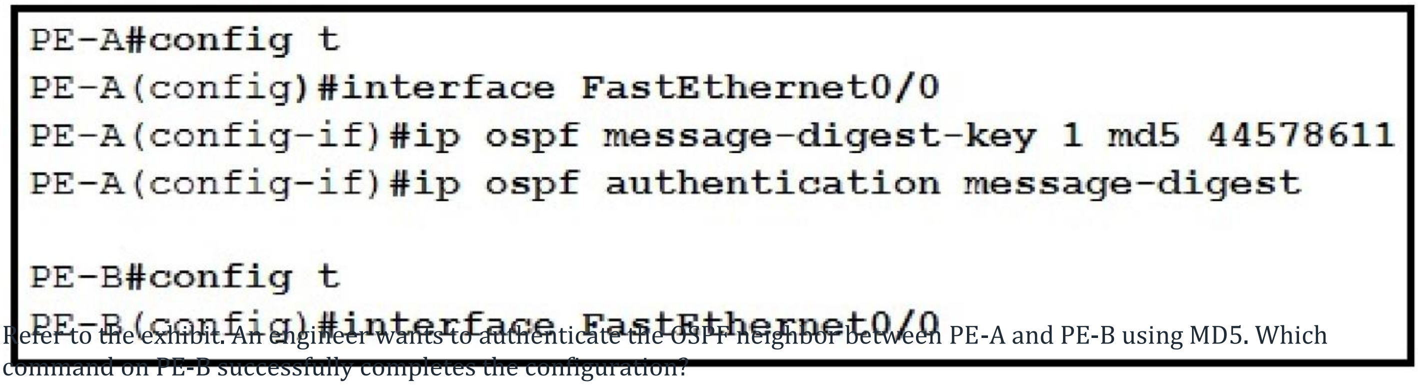

```
PE-A#config t
PE-A(config)#interface FastEthernet0/0
PE-A(config-if)#ip ospf message-digest-key 1 md5 44578611
PE-A(config-if)#ip ospf authentication message-digest

PE-B#config t
PE-B(config)#interface FastEthernet0/0
```

Refer to the exhibit. An engineer wants to authenticate the OSPF neighbor between PE-A and PE-B using MD5. Which command on PE-B successfully completes the configuration?

A. PE-B(config-if)#ip ospf message-digest-key 1 md5 44578611 PE-B(config-if)#ip ospf authentication null

B. PE-B(config-if)#ip ospf message-digest-key 1 md5 44578611 PE-B(config-if)#ip ospf authentication key-chain 44578611

C. PE-B(config-if)#ip ospf message-digest-key 1 md5 44568611 PE-B(config-if)#ip ospf authentication null

D. PE-B(config-if)#ip ospf message-digest-key 1 md5 44578611 PE-B(config-if)#ip ospf authentication message-digest

Answer: D

Explanation:

PE-B(config-if)#ip ospf message-digest-key 1 md5 44578611 PE-B(config-if)#ip ospf authentication message-digest.

Question: 61

DRAG DROP -

Drag and drop each NAT64 description from the left onto the correct NAT64 type on the right.
Select and Place:

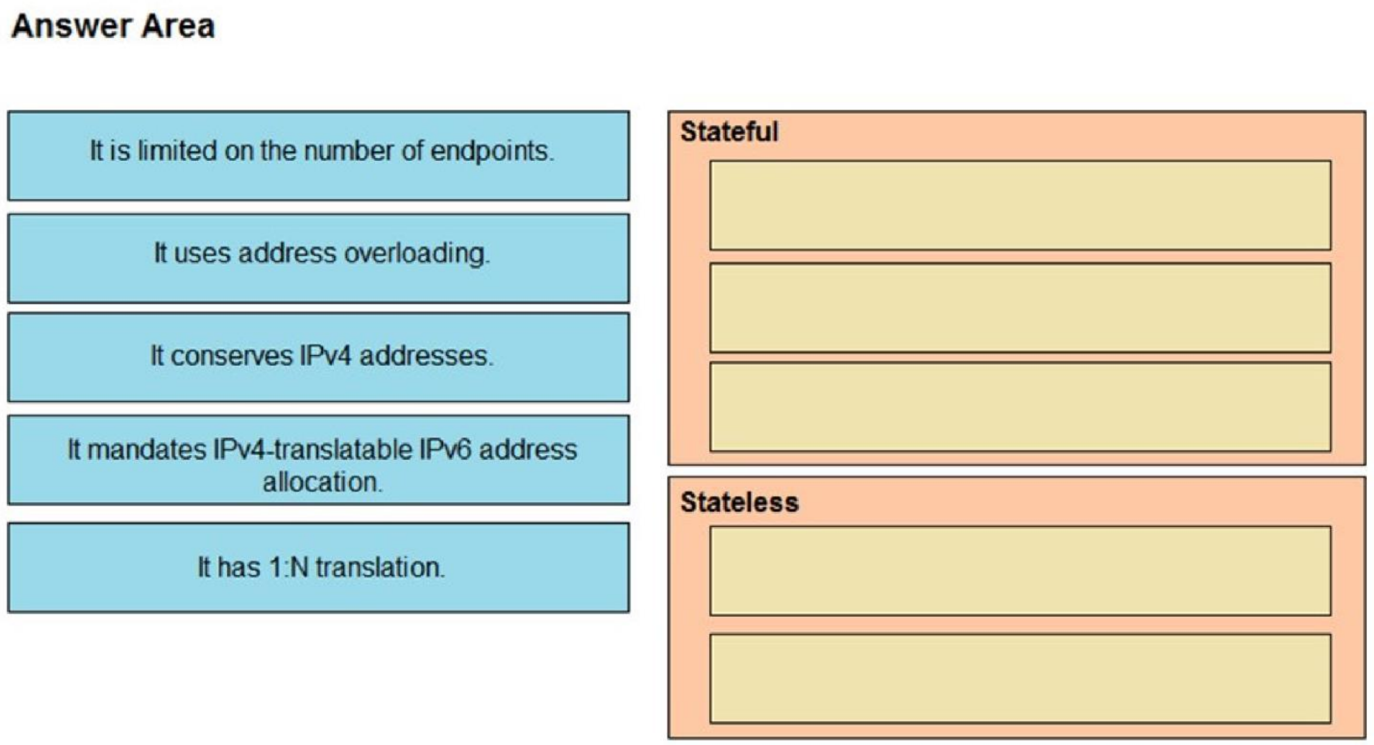

Answer Area

It is limited on the number of endpoints.

It uses address overloading.

It conserves IPv4 addresses.

It mandates IPv4-translatable IPv6 address allocation.

It has 1:N translation.

Stateful

Stateless

Answer:

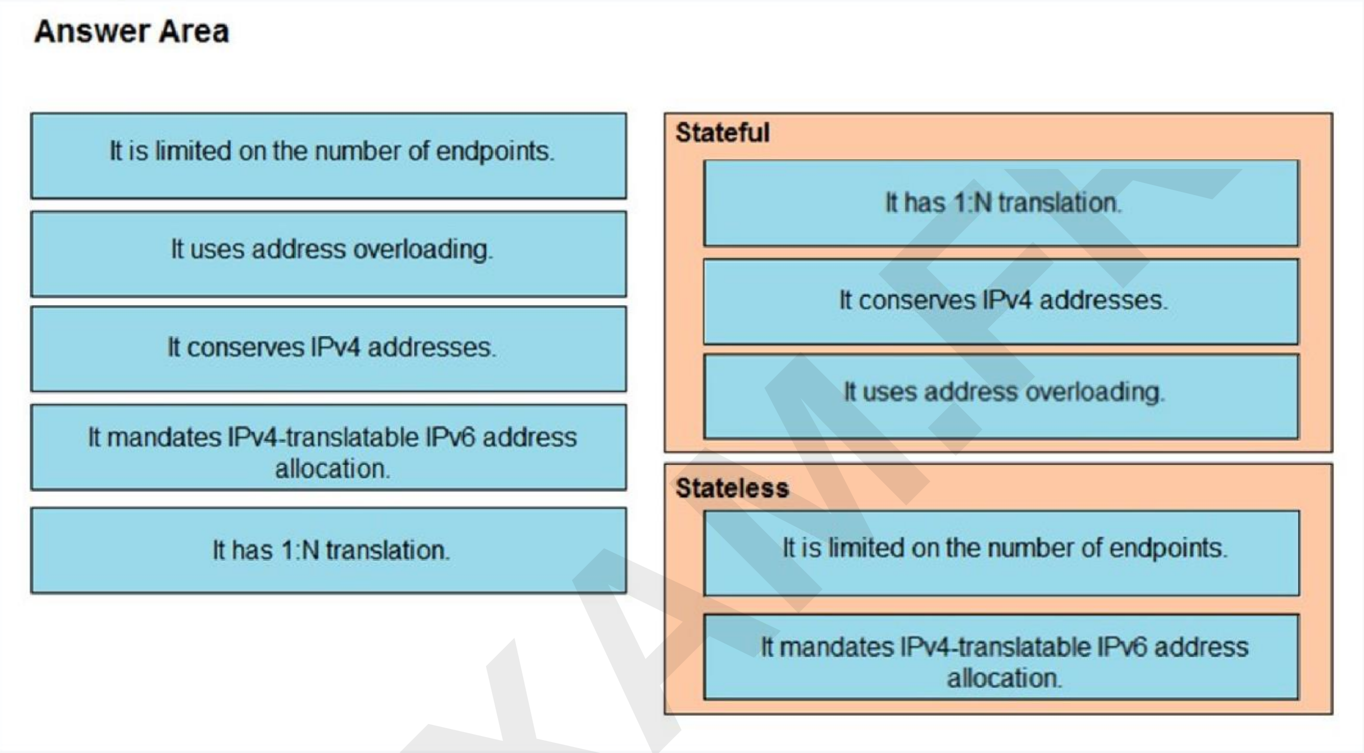

Answer Area

It is limited on the number of endpoints.

It uses address overloading.

It conserves IPv4 addresses.

It mandates IPv4-translatable IPv6 address allocation.

It has 1:N translation.

Stateful

It has 1:N translation.

It conserves IPv4 addresses.

It uses address overloading.

Stateless

It is limited on the number of endpoints.

It mandates IPv4-translatable IPv6 address allocation.

Explanation:

Reference:
https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/enterprise-ipv6-solution/white_paper_ c11-676277.html

Question: 62

DRAG DROP -
Drag and drop the functionalities from the left onto the correct target fields on the right.
Select and Place:

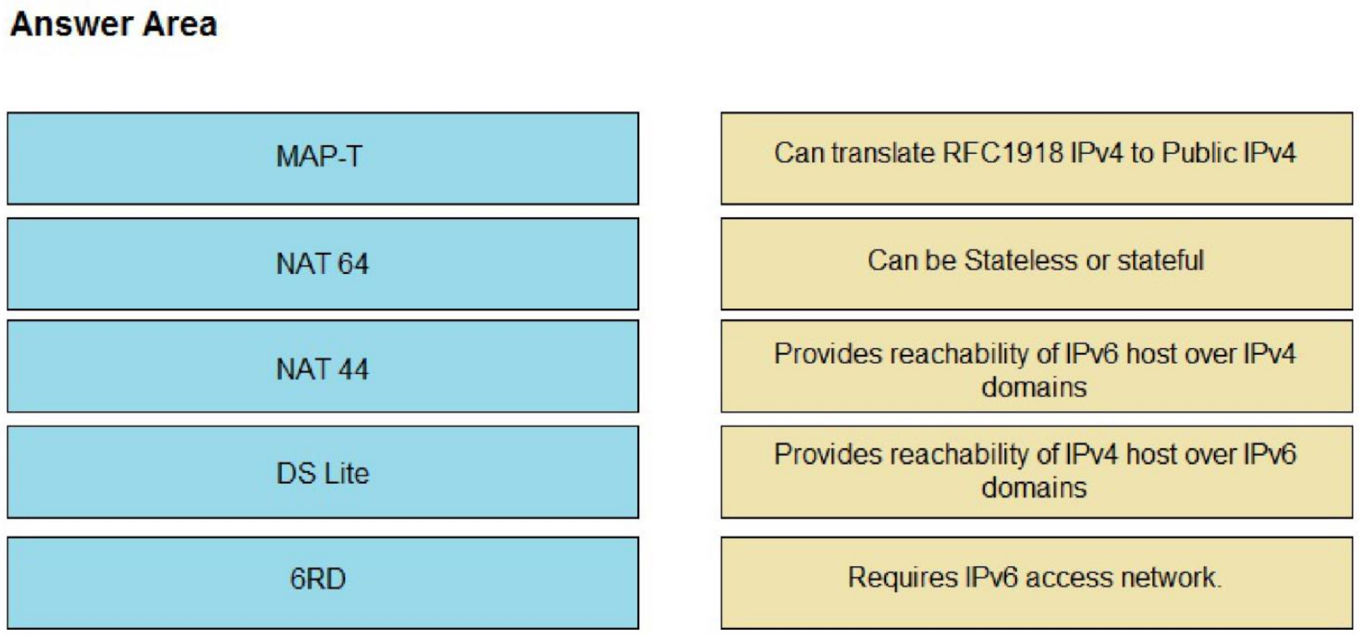

Answer Area

MAP-T	Can translate RFC1918 IPv4 to Public IPv4
NAT 64	Can be Stateless or stateful
NAT 44	Provides reachability of IPv6 host over IPv4 domains
DS Lite	Provides reachability of IPv4 host over IPv6 domains
6RD	Requires IPv6 access network.

Answer:

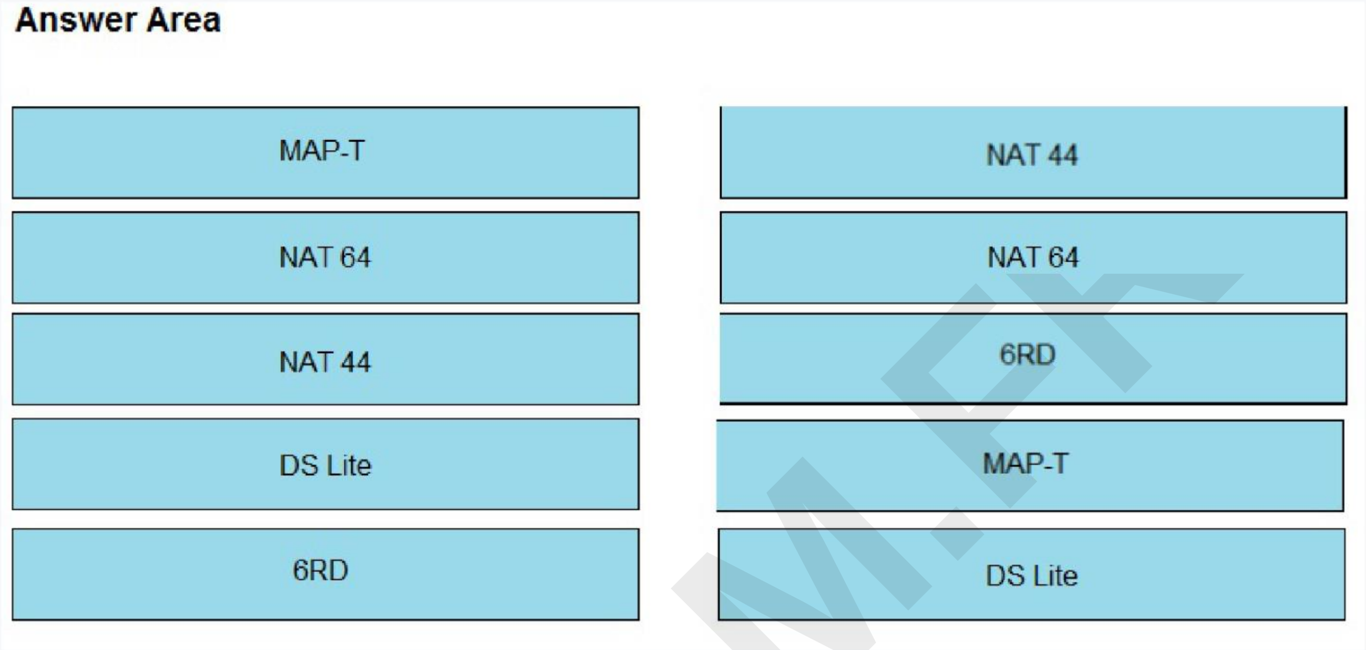

Answer Area

MAP-T	NAT 44
NAT 64	NAT 64
NAT 44	6RD
DS Lite	MAP-T
6RD	DS Lite

Explanation:

NAT 44

NAT 64

6RD

MAP-T

DS Lite

Question: 63 How much must the MTU be increased when configuring the 802.1q VLAN tag?

A. 2 bytes
B. 4 bytes
C. 8 bytes
D. 12 bytes

Answer: B

Explanation:

The correct answer is **B. 4 bytes**. When an Ethernet frame is tagged with an 802.1q VLAN tag, a 4-byte header is inserted into the frame. This tag includes the Tag Protocol Identifier (TPID) field (2 bytes, typically 0x8100) and the Tag Control Information (TCI) field (2 bytes). The TPID identifies the frame as 802.1q tagged, allowing switches to recognize and process VLAN information. The TCI field contains the Priority Code Point (PCP) for quality of service, Drop Eligible Indicator (DEI), and the VLAN Identifier (VID), which dictates which VLAN the frame belongs to. Because this 4-byte tag is added, the Maximum Transmission Unit (MTU) of the underlying Ethernet frame effectively increases by 4 bytes. To avoid fragmentation, especially at the IP layer, you might need to increase the configured MTU on interfaces where 802.1q tagging is used. While the actual physical layer frame does not grow in size, the logical increase in frame size with the added tag requires considering the overall path MTU. Therefore, the standard 1500 byte Ethernet MTU can result in fragmented IP packets. Therefore, the effective increase in MTU due to the VLAN tag is 4 bytes.

For further reading:

IEEE 802.1Q: https://standards.ieee.org/ieee/802/1Q/
Cisco documentation on VLAN tagging: (search "Cisco 802.1q") https://www.cisco.com/c/en/us/index.html
Network MTU considerations (general): https://www.cloudflare.com/learning/network-layer/what-is-mtu/

Question: 64

Egress PE NAT is being used via a single centralized router to provide Internet access to L3VPN customers. Which description of the NAT operation is true?

A. The NAT table contains a field to identify the inside VRF of a translation.
B. Multiple address pools are needed for the same L3VPN because each site has a separate NAT.
C. The different L3VPNs using the Internet access must not have IP overlaps internally.
D. Users in different VRFs cannot share the same outside global IP address.

Answer: A

Explanation:

Okay, let's break down why option A is the correct answer regarding egress PE NAT in a service provider network for L3VPNs, and why the other options are not.

Egress PE NAT, in this context, refers to Network Address Translation being performed on the egress Point of Presence (PE) router before traffic from L3VPN customers is routed to the internet. The core challenge here is maintaining separation and unique addressing for customers who might have overlapping internal IP address spaces.

Why Option A is Correct: The key is understanding how the NAT router keeps track of which L3VPN or VRF the traffic originates from. The NAT table must include a field that identifies the inside VRF of each translated session. This is vital because it allows the router to accurately translate return traffic back to the correct customer's VRF. Without this VRF awareness, return packets would be misrouted, breaking the communication. This separation is crucial for maintaining the logical isolation promised by VRFs. It ensures that even if multiple customers use the same private IP address range, their traffic is kept separate and directed to their specific endpoint behind the NAT device.

Why Other Options are Incorrect:

B. Multiple address pools are needed for the same L3VPN because each site has a separate NAT: This is incorrect. In a centralized model, there's typically only one NAT router serving all L3VPNs for internet access.

Multiple address pools might be used but not for different sites within the same L3VPN. Different address pools are used to avoid clashes for different VPNs.

C. The different L3VPNs using the Internet access must not have IP overlaps internally: This is incorrect. L3VPNs are designed to allow internal IP overlaps between different customers. NAT is precisely used to resolve IP overlap issues that may happen. The NAT process translates the private, potentially overlapping, internal addresses to globally unique public addresses on the Internet, resolving the problem of duplicated addresses.

D. Users in different VRFs cannot share the same outside global IP address: This statement is also incorrect. If the NAT implementation supports it, Multiple VRFs can use the same outside global IP address through PAT (Port Address Translation), by differentiating on source ports. Also with NAT overload, several private addresses can share the same public address. In most real-world scenarios, providers do overload private IPs to a pool of public IPs, allowing many customers to share the same outside global IP, which is cost-effective and doesn't require a 1:1 public IP mapping for every private customer IP.

Authoritative Links:

Cisco: Understanding VRF-Aware NAT:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_nat/configuration/15-sy/nat-15-sy-book/nat-vrf-aware.html This Cisco documentation explains VRF-aware NAT and how it keeps track of the routing context.

Juniper: VRF-Aware NAT Overview:
https://www.juniper.net/documentation/us/en/software/junos/nat/topics/topic-map/nat-vrf-aware.html
Juniper provides similar concepts about VRF NAT and the advantages.

In summary, the ability to identify the inside VRF within the NAT table (option A) is critical for correctly performing NAT in a multi-VRF environment and is therefore the correct choice. The other options contain inaccuracies that do not represent how NAT works within a L3VPN service provider context.

Question: 65

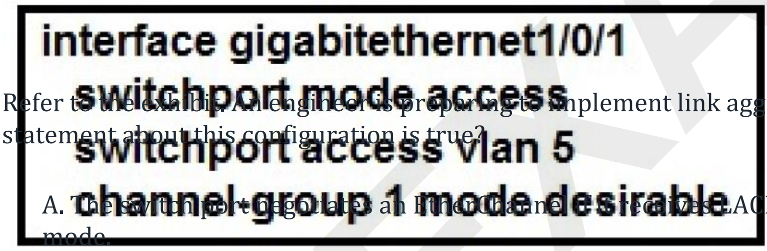

```
interface gigabitethernet1/0/1
 switchport mode access
 switchport access vlan 5
 channel-group 1 mode desirable
```

Refer to the exhibit. An engineer is preparing to implement link aggregation configuration. Which statement about this configuration is true?

A. The switchport negotiates an EtherChannel if it receives LACP packets from a connected peer running passive mode.

B. The switch port actively sends packets to negotiate an EtherChannel using PAgP.

C. The switch port passively negotiates an EtherChannel if it receives PAgP packets from a connected peer.

D. The switch port accepts LACP and PAgP packets from a connected peer and negotiate an EtherChannel using the common EtherChannel mode.

Answer: B

Explanation:

Reference:
https://www.cisco.com/en/US/docs/switches/metro/me3600x_3800x/trash/swethchl.html

Question: 66

You are writing an RPL script to accept routes only from certain autonomous systems. Consider this code:

```
RP/0/RP0/CPU0:router(config-rpl)# if as-path in (ios-regex '.*77 )
RP/0/RP0/CPU0:router(config-rpl-if)# pass
RP/0/RP0/CPU0:router(config-rpl-if)# endif
```

If you apply this code to BGP filters, which effect does the code have on your router?

A. denies routes from AS 7070
B. allows routes from AS 7077
C. denies routes from AS 7007
D. allows routes from AS 770

Answer: B

Explanation:

The provided RPL (Routing Policy Language) script uses an if statement with an AS-path check employing a regular expression. The core logic lies in this line: if as-path in (ios-regex '.*77*

#39;). Let's break it down. *as-path* refers to the sequence of Autonomous System numbers a route has traversed. *ios-regex '.77*
#39; is the regular expression being applied to the AS-path. The . matches any single character, * means zero or more occurrences of the preceding character, 77 matches the literal string "77", and $ matches the end of the string. Therefore, the regex essentially checks if the AS-path ends with the numbers "77".

The pass command inside the if block dictates that if the condition (the AS-path ending with "77") is true, then the route will be accepted. Conversely, routes that do not end with "77" will be implicitly denied because no other explicit pass or reject action is specified.

Therefore, the effect of the script is to specifically allow routes where the last Autonomous System in the AS-path is '77', regardless of the preceding ASNs. This is why Option B, "allows routes from AS 7077," is the correct answer. The other options are incorrect because: A mentions AS 7070 which does not match the regex. C mentions AS 7007 which does not match the regex. D mentions AS 770 which does not match the regex because it does not end in '77'.

Further research on RPL and AS-Path filtering can be done through Cisco's official documentation. Here is a relevant link to Cisco's documentation on RPL:

https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/routing/65x/b-routing-cg-ncs5500-65x/b-routing-cg-ncs5500-65x_chapter_011.html

Question: 67

Which regular expression query modifier function indicates the start of a string?

A. +
B. ^
C. $
D. [^]

Answer: B

Explanation:

The correct answer is **B. ^**. In regular expressions (regex), the caret symbol ^ is an anchor that asserts the position at the beginning of the string or a line within a multiline string. It doesn't match any characters themselves, but rather specifies a location where the subsequent pattern must start. Conversely, $ (option C) indicates the end of the string or line. The plus sign + (option A) is a quantifier, matching one or more occurrences of the preceding element. Finally, [^] (option D) denotes a negated character set, meaning any character not within the brackets, which has nothing to do with string start position. The ^ is used in regex syntax in various programming languages and tools, including those used for network configuration and scripting, and is fundamental in pattern matching for tasks like data validation or log analysis. In the context of Cisco networking, understanding regular expressions is essential for tasks like filtering logs, identifying specific patterns in configurations, and manipulating network data. Therefore, ^ consistently designates the beginning position in regex patterns and is the correct modifier to use in such contexts.

Authoritative Links:

1. **Regular-Expressions.info: Anchors:**https://www.regular-expressions.info/anchors.html
2. **MDN Web Docs: Regular Expressions:**https://developer.mozilla.org/en-US/docs/Web/JavaScript/Guide/Regular_expressions (This explains regex concepts generally, applicable across many programming contexts.)
3. **Cisco Documentation: Using Regular Expressions for Configuration:** (Specific Cisco documentation may vary depending on the device, but often includes regex usage in config management). Refer to Cisco's documentation portal for relevant articles by searching for "Cisco regular expressions configuration".

Question: 68

A network engineer is configuring a BGP route policy for the SUBNET prefix set. Matching traffic must be dropped, and other traffic must have its MED value set to
400 and community 4:400 added to the route. Which configuration must an engineer apply?
A.

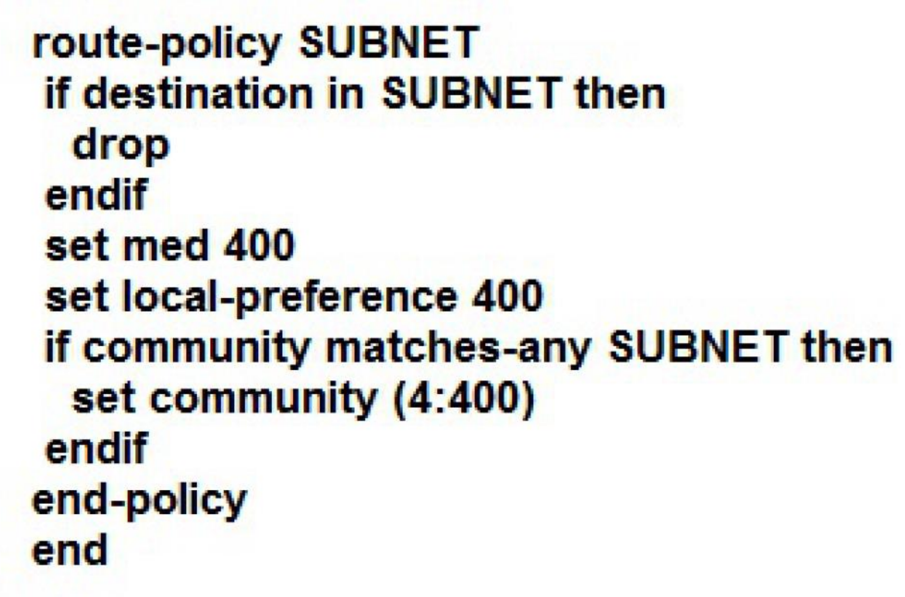

```
route-policy SUBNET
 if destination in SUBNET then
   drop
 endif
 set med 400
 set local-preference 400
 if community matches-any SUBNET then
   set community (4:400)
 endif
end-policy
end
```

B.

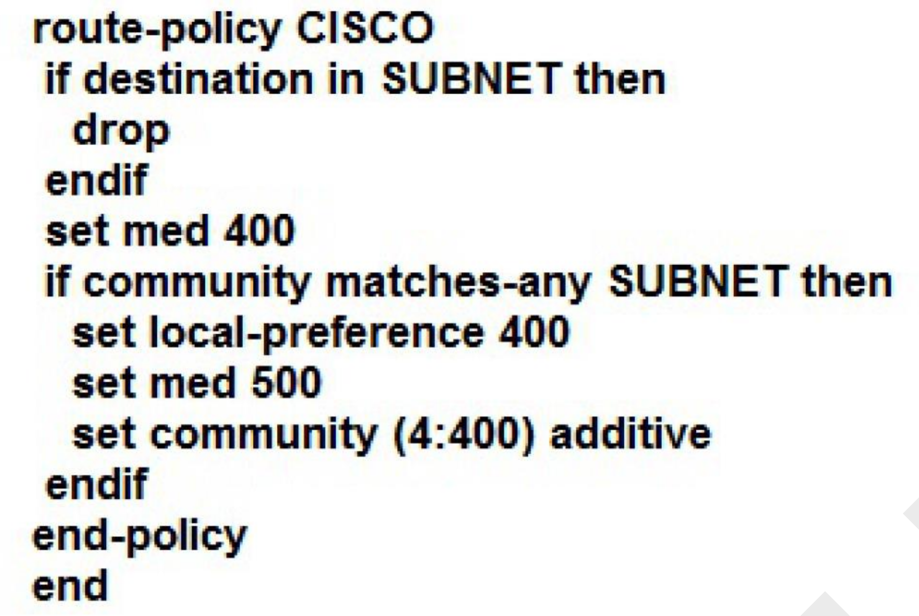

```
route-policy CISCO
 if destination in SUBNET then
   drop
 endif
 set med 400
 if community matches-any SUBNET then
   set local-preference 400
   set med 500
   set community (4:400) additive
 endif
end-policy
end
```

C.

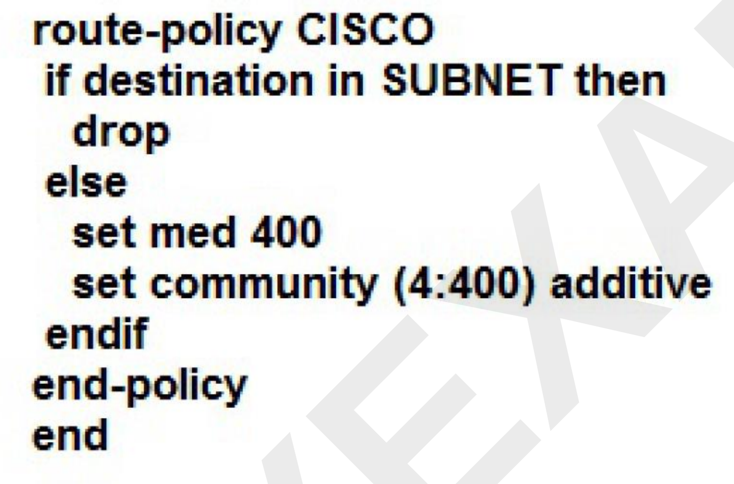

```
route-policy CISCO
 if destination in SUBNET then
   drop
 else
   set med 400
   set community (4:400) additive
 endif
end-policy
end
```

D.

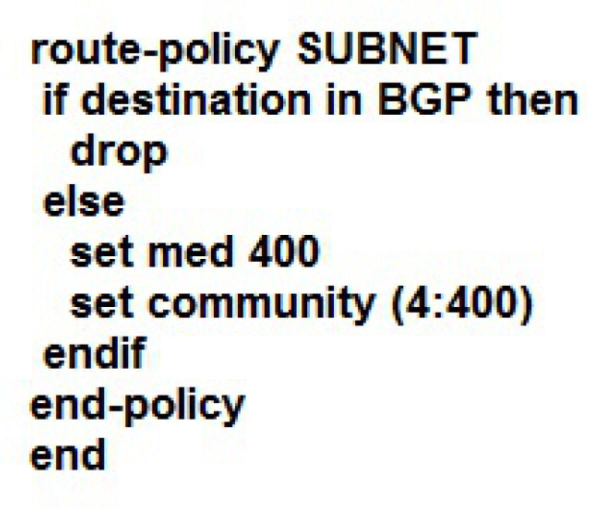

```
route-policy SUBNET
 if destination in BGP then
   drop
 else
   set med 400
   set community (4:400)
 endif
end-policy
end
```

Answer: C
Explanation:

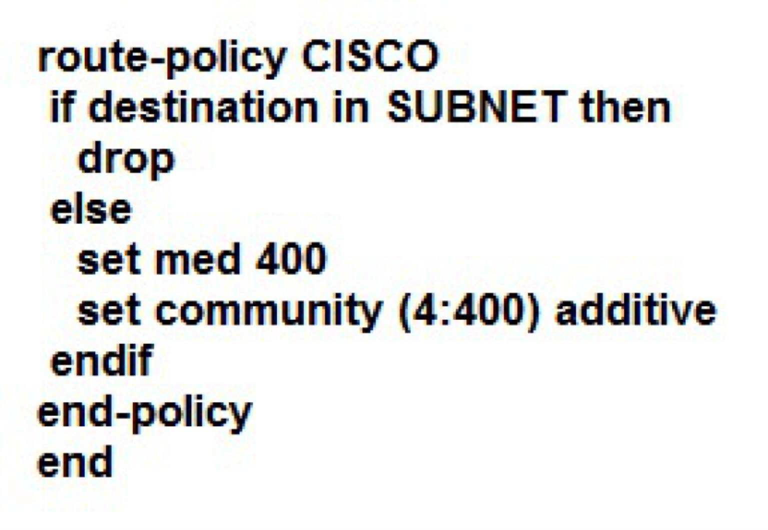

```
route-policy CISCO
 if destination in SUBNET then
   drop
 else
   set med 400
   set community (4:400) additive
 endif
end-policy
end
```

Question: 69

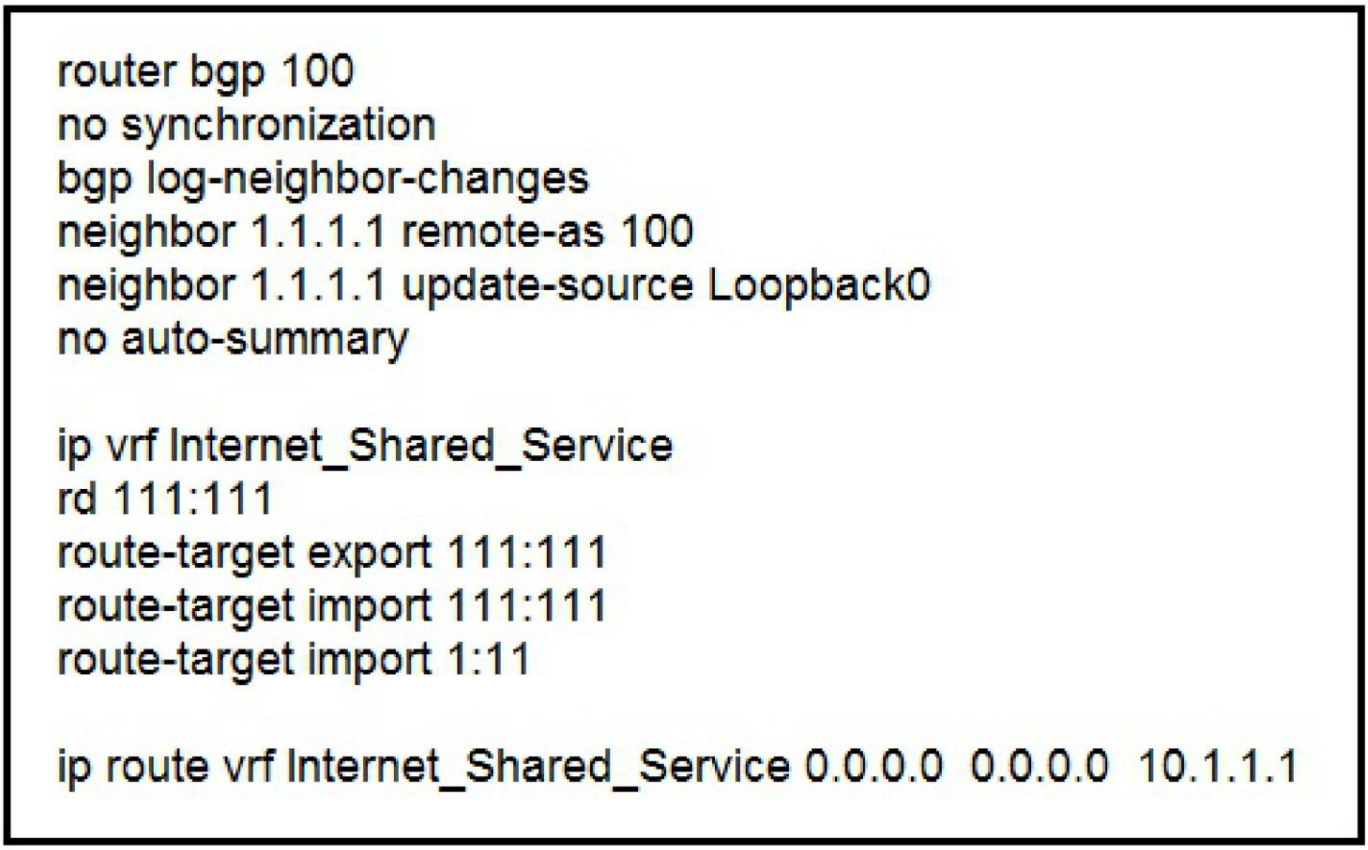

```
router bgp 100
no synchronization
bgp log-neighbor-changes
neighbor 1.1.1.1 remote-as 100
neighbor 1.1.1.1 update-source Loopback0
no auto-summary

ip vrf Internet_Shared_Service
rd 111:111
route-target export 111:111
route-target import 111:111
route-target import 1:11

ip route vrf Internet_Shared_Service 0.0.0.0  0.0.0.0  10.1.1.1
```

Refer to the exhibit. Which additional configuration must an engineer apply to the edge router to inject a default route into the MP-BGP address family for the
Internet_Shared_Service dedicated VRF?

A.

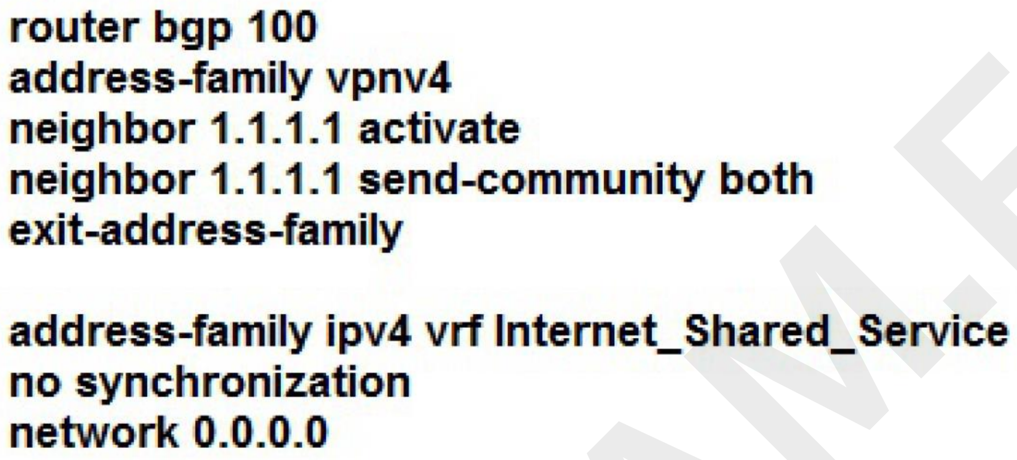

```
router bgp 100
address-family vpnv4
neighbor 1.1.1.1 activate
neighbor 1.1.1.1 send-community both
exit-address-family

address-family ipv4 vrf Internet_Shared_Service
no synchronization
network 0.0.0.0
```

B.

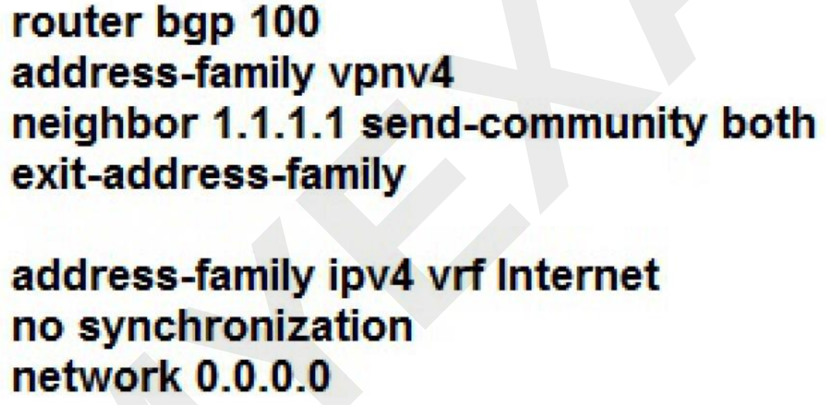

```
router bgp 100
address-family vpnv4
neighbor 1.1.1.1 send-community both
exit-address-family

address-family ipv4 vrf Internet
no synchronization
network 0.0.0.0
```

C.

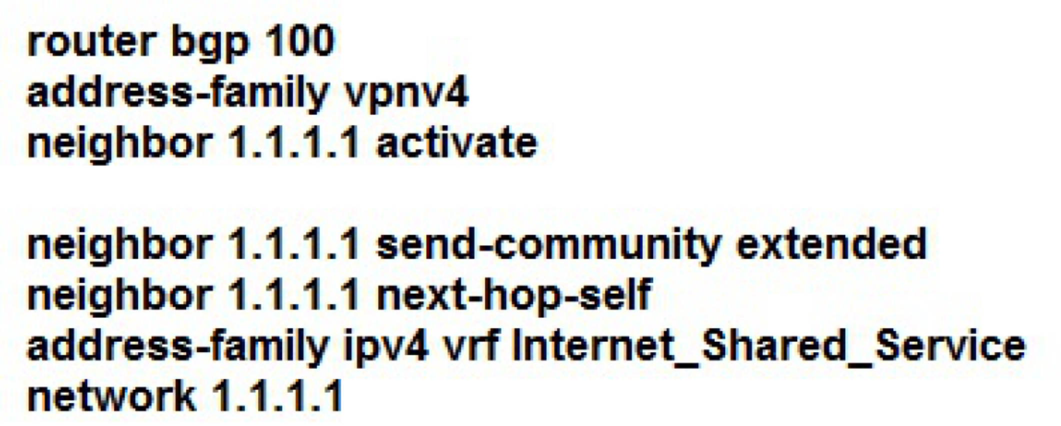

```
router bgp 100
address-family vpnv4
neighbor 1.1.1.1 activate

neighbor 1.1.1.1 send-community extended
neighbor 1.1.1.1 next-hop-self
address-family ipv4 vrf Internet_Shared_Service
network 1.1.1.1
```

D.

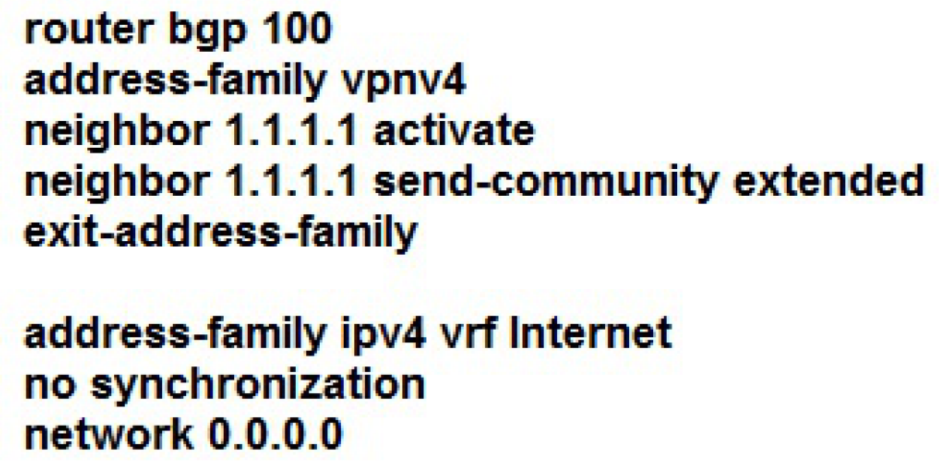

```
router bgp 100
address-family vpnv4
neighbor 1.1.1.1 activate
neighbor 1.1.1.1 send-community extended
exit-address-family

address-family ipv4 vrf Internet
no synchronization
network 0.0.0.0
```

Answer: A

Explanation:

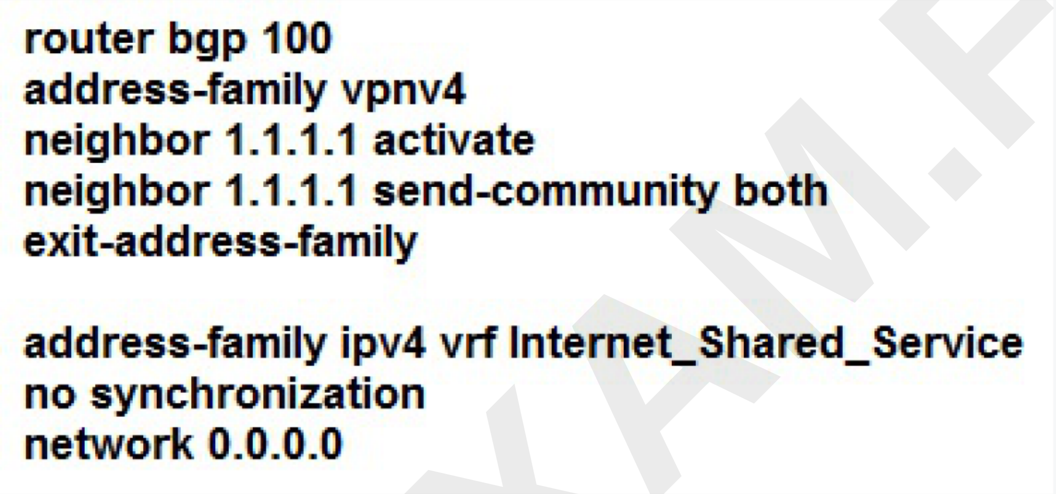

```
router bgp 100
address-family vpnv4
neighbor 1.1.1.1 activate
neighbor 1.1.1.1 send-community both
exit-address-family

address-family ipv4 vrf Internet_Shared_Service
no synchronization
network 0.0.0.0
```

Question: 70

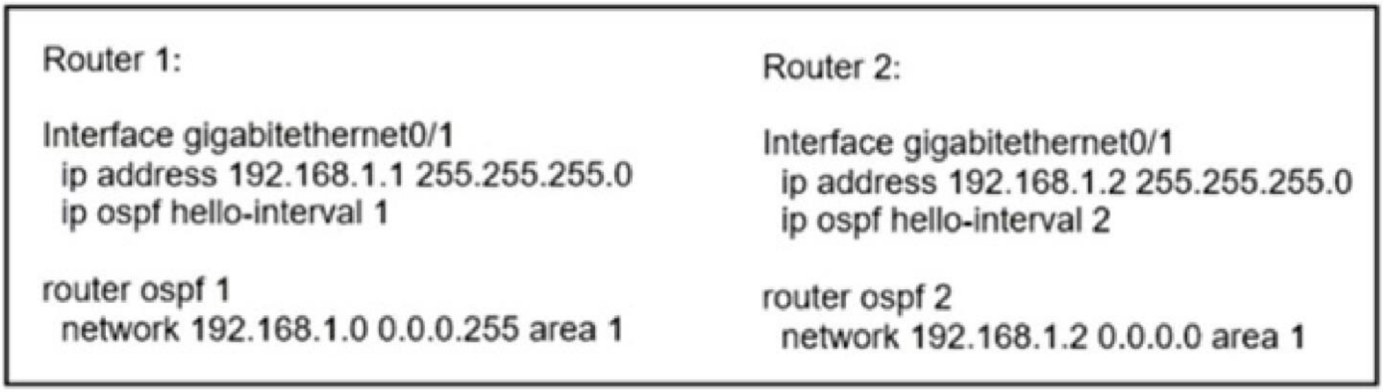

```
Router 1:

Interface gigabitethernet0/1
 ip address 192.168.1.1 255.255.255.0
 ip ospf hello-interval 1

router ospf 1
 network 192.168.1.0 0.0.0.255 area 1
```

```
Router 2:

Interface gigabitethernet0/1
 ip address 192.168.1.2 255.255.255.0
 ip ospf hello-interval 2

router ospf 2
 network 192.168.1.2 0.0.0.0 area 1
```

Refer to the exhibit. What reestablishes the OSPF neighbor relationship between Router 1 and Router 2?

A. OSPF process IDs match
B. authentication is added to the configuration
C. correct wildcard mask is used on Router 2
D. hello intervals match

Answer: D

Explanation:

Correct answer is D: hello intervals match.

Question: 71

A network engineer is deploying VRF on ASBR router R1. The interface must have connectivity over an MPLS VPN Inter-AS Option AB network. Which configuration must the engineer apply on the router to accomplish this task? A.

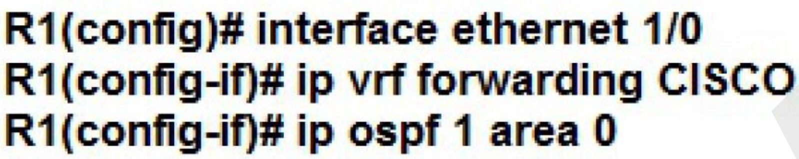

```
R1(config)# interface ethernet 1/0
R1(config-if)# ip vrf forwarding CISCO
R1(config-if)# ip ospf 1 area 0
```

B.

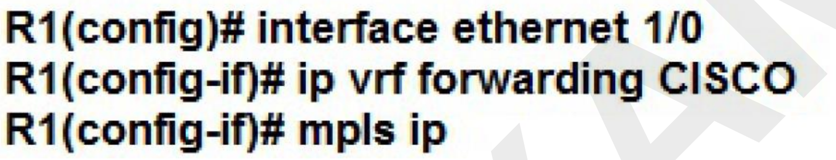

```
R1(config)# interface ethernet 1/0
R1(config-if)# ip vrf forwarding CISCO
R1(config-if)# mpls ip
```

C.

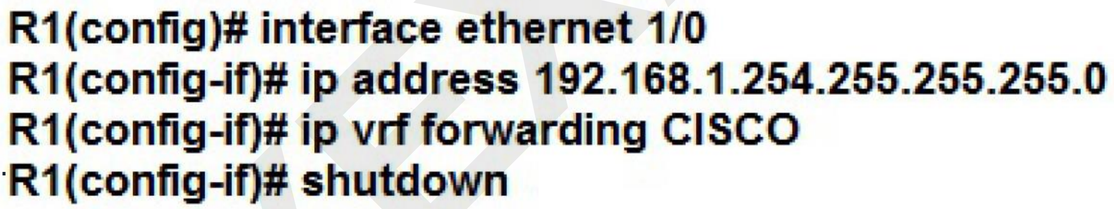

```
R1(config)# interface ethernet 1/0
R1(config-if)# ip address 192.168.1.254.255.255.255.0
R1(config-if)# ip vrf forwarding CISCO
R1(config-if)# shutdown
```

D.

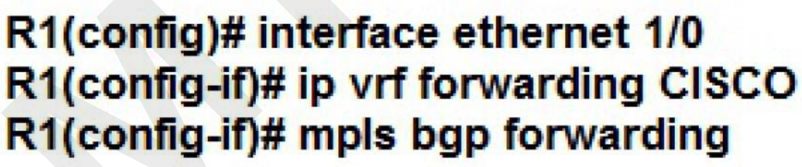

```
R1(config)# interface ethernet 1/0
R1(config-if)# ip vrf forwarding CISCO
R1(config-if)# mpls bgp forwarding
```

Answer: D

Explanation:

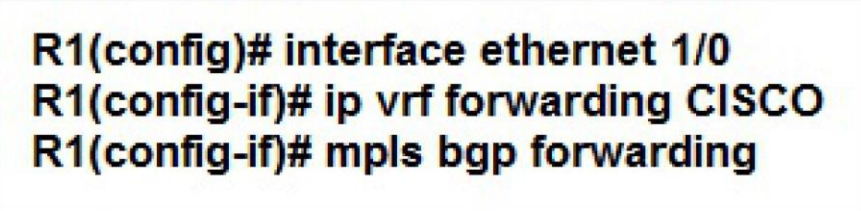

```
R1(config)# interface ethernet 1/0
R1(config-if)# ip vrf forwarding CISCO
R1(config-if)# mpls bgp forwarding
```

Question: 72

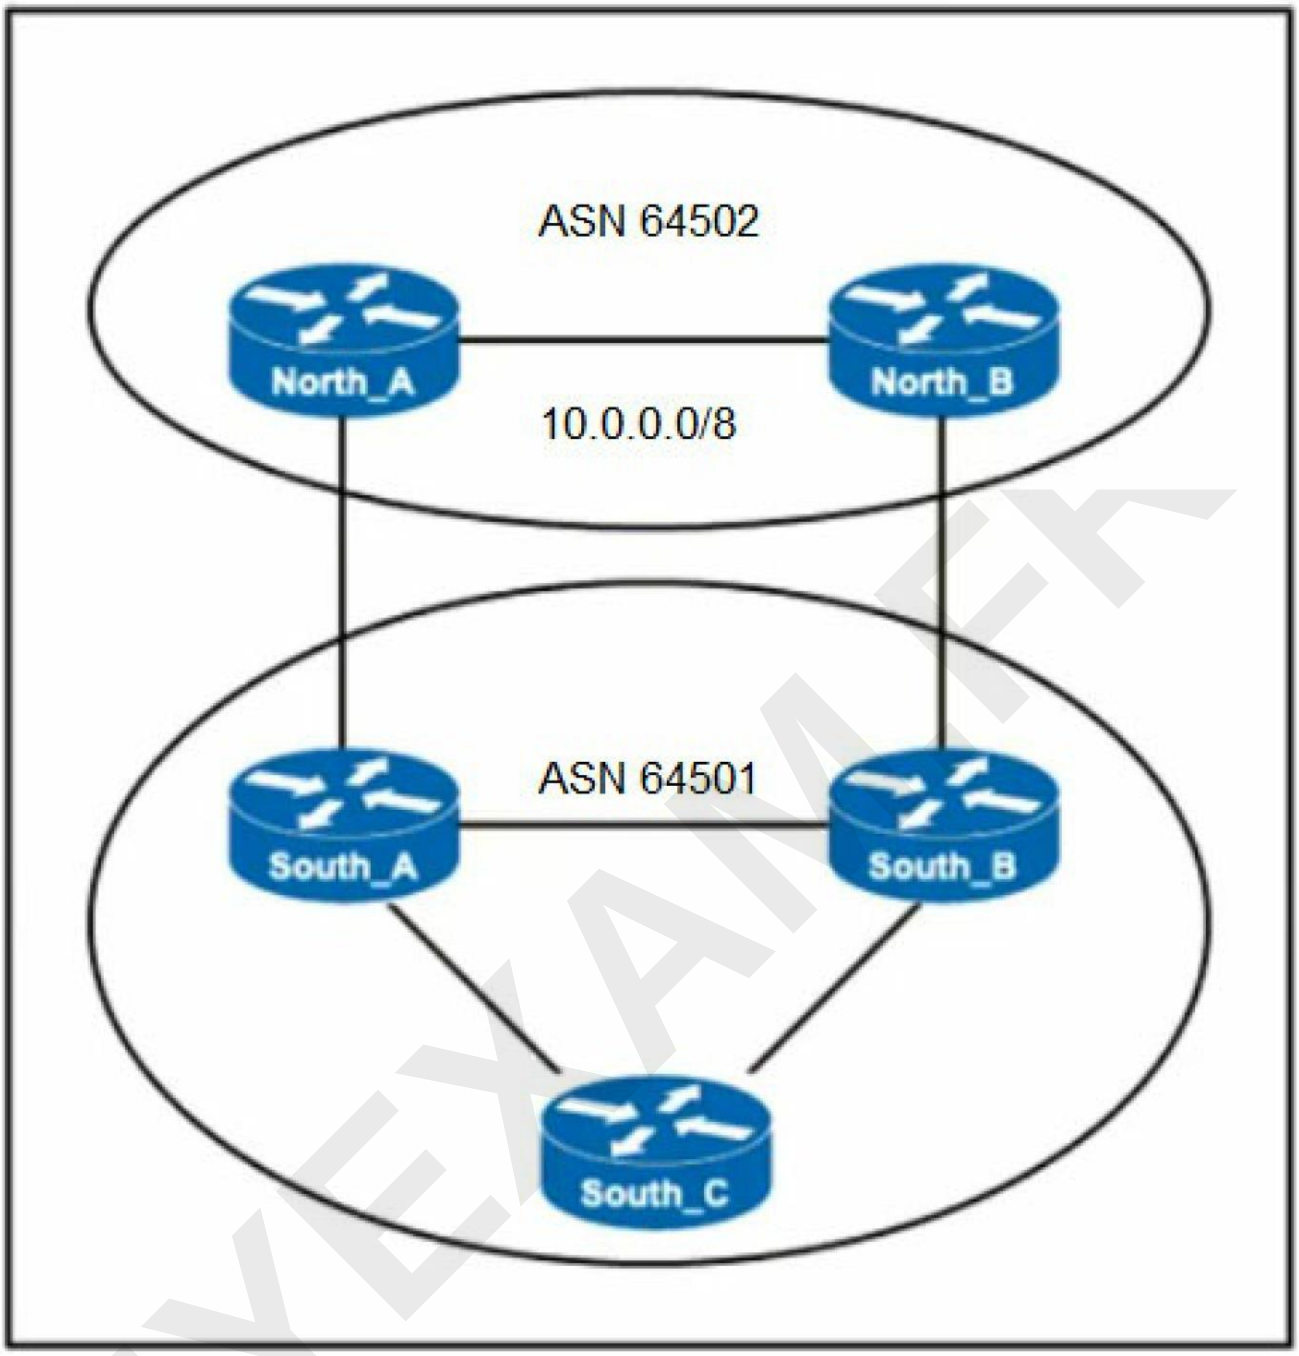

Refer to the exhibit. ASN 64501 currently reaches the networks under the 10.0.0.0/8 prefix via the North_B router, which is a slow backup link. The administrator of ASN 64502 wants traffic from ASN 64501 to 10.0.0.0/8 to travel via the primary link, North_A. Which change to the network configuration accomplishes this task?

A. Set a higher local preference between North_A and South-A
B. Set a lower MED between North_B and South_B
C. Advertise the 10.0.0.0/8 prefix through North_B and specific subnets through North_A D. Set a lower Weight value for incoming traffic on North_A

Answer: C

Explanation:

Advertise the 10.0.0.0/8 prefix through North_B and specific subnets through North_A.

Question: 73

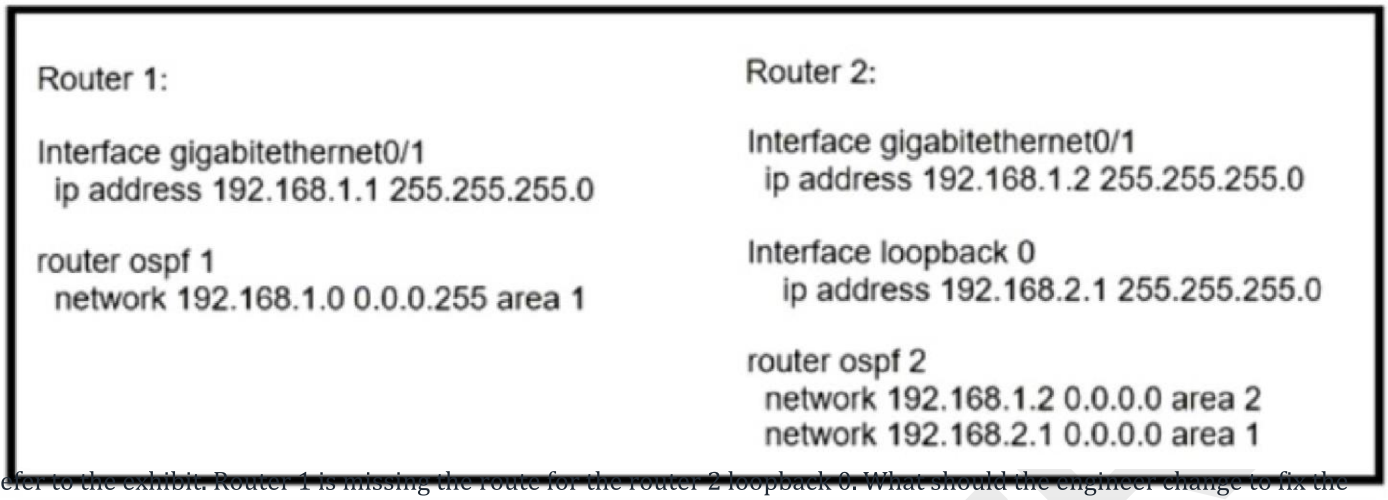

```
Router 1:

Interface gigabitethernet0/1
 ip address 192.168.1.1 255.255.255.0

router ospf 1
 network 192.168.1.0 0.0.0.255 area 1
```

```
Router 2:

Interface gigabitethernet0/1
 ip address 192.168.1.2 255.255.255.0

Interface loopback 0
   ip address 192.168.2.1 255.255.255.0

router ospf 2
 network 192.168.1.2 0.0.0.0 area 2
 network 192.168.2.1 0.0.0.0 area 1
```

Refer to the exhibit. Router 1 is missing the route for the router 2 loopback 0. What should the engineer change to fix the problem?

A. Router 1 to be an ABR
B. the wildcard mask network statement in OSPF of Router 2 C. the hello timers on Router 1 and Router 2 to be different D. the area numbers on Router 1 and Router 2 to be similar

Answer: D

Explanation:

the area numbers on Router 1 and Router 2 to be similar.

Question: 74

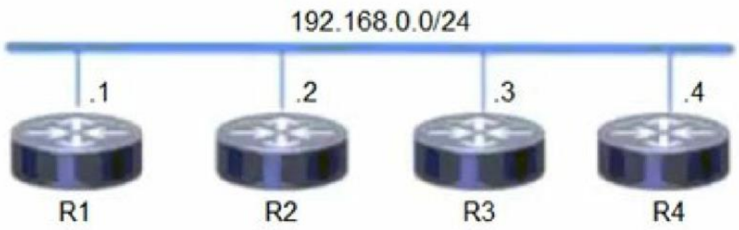

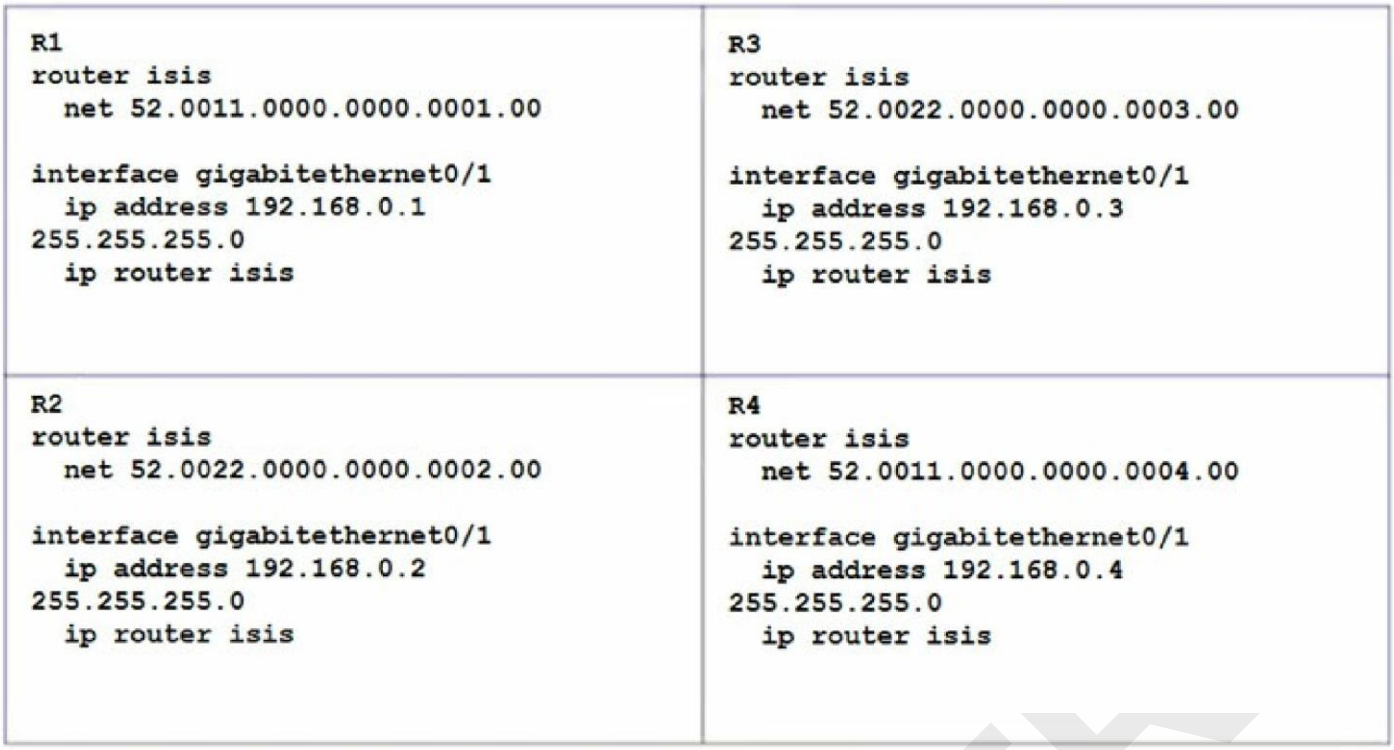

```
R1
router isis
  net 52.0011.0000.0000.0001.00

interface gigabitethernet0/1
  ip address 192.168.0.1
255.255.255.0
  ip router isis
```

```
R3
router isis
  net 52.0022.0000.0000.0003.00

interface gigabitethernet0/1
  ip address 192.168.0.3
255.255.255.0
  ip router isis
```

```
R2
router isis
  net 52.0022.0000.0000.0002.00

interface gigabitethernet0/1
  ip address 192.168.0.2
255.255.255.0
  ip router isis
```

```
R4
router isis
  net 52.0011.0000.0000.0004.00

interface gigabitethernet0/1
  ip address 192.168.0.4
255.255.255.0
  ip router isis
```

Refer to the exhibit. Which two topology changes happen to the IS-IS routers? (Choose two.)

A. R1 and R4 are Level 2 neighbors
B. All four routers are operating as Level 1-2 routers
C. All four routers are operating as Level 2 routers only
D. R1 and R2 are Level 2 neighbors
E. All four routers are operating as Level 1 routers only

Answer: BD

Explanation:

B. All four routers are operating as Level 1-2 routers.

D. R1 and R2 are Level 2 neighbors.

Question: 75

An engineer is trying to implement BGP in a multihomed architecture. What must the engineer configure to influence inbound path selection?

A. A route map with AS-PATH attribute to control the inbound traffic
B. An offset list to set the metric for routes received from neighboring autonomous systems C. An
access list to identify traffic and enable it on both of the provider-facing interfaces

D. A route map with WEIGHT attribute to control the inbound traffic

Answer: A

Explanation:

The correct answer is **A. A route map with AS-PATH attribute to control the inbound traffic**. BGP's path selection process is complex, but for influencing inbound traffic, the receiving router (in this case, the customer edge device) has limited direct control. It advertises routes to its peers, and the peer (the upstream ISP) then chooses its best path, impacting how it sends traffic to the customer. The AS_PATH attribute manipulation is a common technique used by the advertising router (the ISP) to influence how its peers select a route. By prepending its own AS number multiple times to the AS_PATH, the ISP effectively makes that path appear less desirable to other ISPs, forcing inbound traffic destined for the customer to take another preferred path.

Option B is incorrect because an offset list adjusts metrics for outgoing traffic, and it does not directly influence the path selection on another AS's side. Option C is unsuitable because access lists filter traffic and don't influence path selection. Option D is wrong because the WEIGHT attribute is local to a single router and is not shared in BGP updates, meaning it does not influence other peer's route selection. Instead, attributes such as AS-PATH (for path length), LOCAL_PREF (for preference within an AS), and MED (for preferences between AS's) are key factors for influencing path selection. Therefore, the AS-PATH attribute manipulation via a route-map applied on the advertising ISP is the tool available to impact inbound traffic selection by its peers.

Authoritative Links:

1. **Cisco BGP Path Selection:**https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html
2. **BGP Attributes - AS_PATH:** https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/bgp-as-path-attributes.html
3. **BGP Route Maps:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/15-sy/irg-15-sy-book/irg-route-map.html

Question: 76

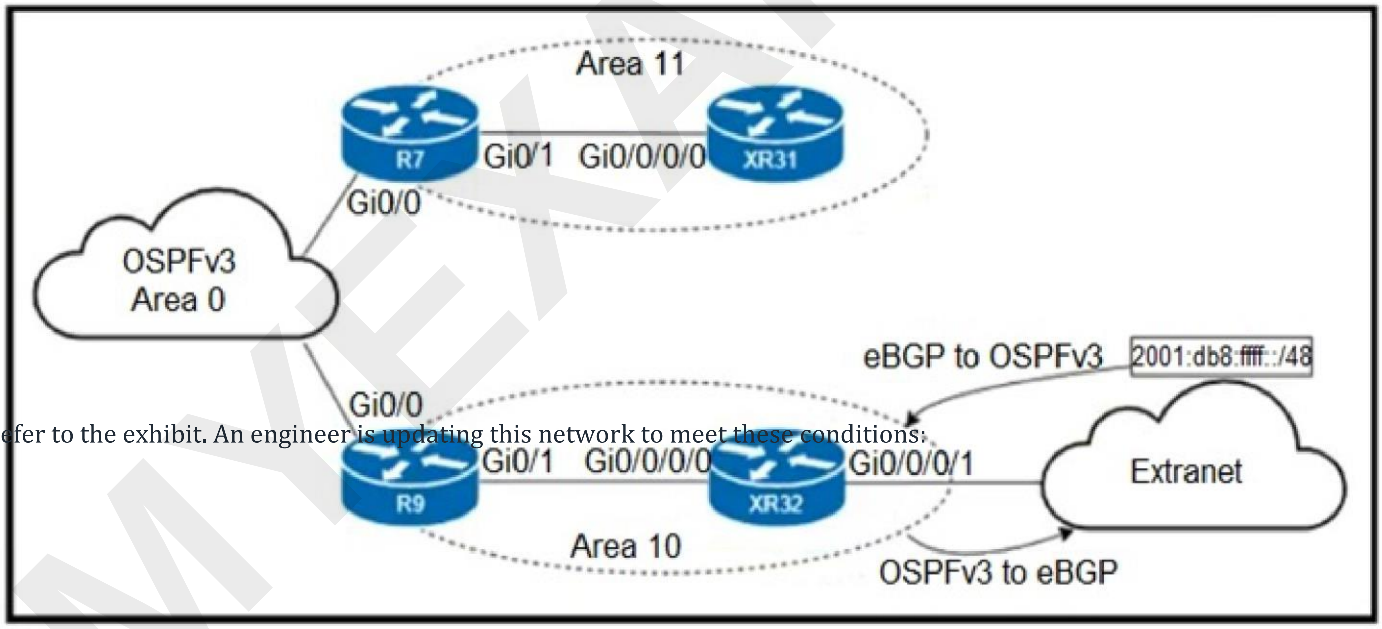

Refer to the exhibit. An engineer is updating this network to meet these conditions:

✑ Area 10 will receive inter-area routes and support mutual redistribution of external routes with the extranet.✑ The ::/0 route is prohibited in Area 10.
✑ Area 11 will receive only the ::/0 route from the ABR.
✑ External route redistribution is not supported in Area 11.
✑ The ABR in Area 11 will advertise no interarea routes.
Which two configurations must be performed to meet the requirements? (Choose two.)

A. Configure area 10 as nssa on R9 and XR32

B. Configure area 11 as stub no-summary on R7 and as stub on XR31

C. Configure area 11 as nssa no-summary on R7 and as nssa on XR31

D. Configure area 11 as nssa default-information-originate on R7 and as nssa on XR31

E. Configure area 10 as stub on R9 and XR32

Answer: AB

Explanation:

A. Configure area 10 as nssa on R9 and XR32.

B. Configure area 11 as stub no-summary on R7 and as stub on XR31.

Question: 77

A network team has failed to implement IS-IS multitopology. What is the reason for it?

A. The routing process supported Level 1 only

B. The router did not support VRFs

C. The routing process did not support extended metrics

D. The router did not have Cisco Discovery Protocol and Cisco Express Forwarding disabled

Answer: C

Explanation:

The correct answer is **C. The routing process did not support extended metrics.**

IS-IS multitopology requires the use of extended metrics to differentiate between routes belonging to different topologies. Standard IS-IS metrics are not sufficient to support multiple distinct topologies. When extended metrics are not enabled or supported on the routing process, IS-IS will not be able to segregate routes based on topology, thereby preventing the implementation of multitopology.

Option A is incorrect because IS-IS supports Level 1 and Level 2 routing, and the issue is not related to the level of routing. Option B is incorrect because VRFs (Virtual Routing and Forwarding instances) are not directly related to implementing IS-IS multitopology. While VRFs can be used in conjunction with multitopology in some deployments, the core issue preventing its implementation here is not the lack of VRFs.

Option D is incorrect because Cisco Discovery Protocol (CDP) and Cisco Express Forwarding (CEF) are not necessary for the correct function of IS-IS multitopology. CDP is a device discovery protocol, and CEF is Cisco's packet forwarding mechanism, neither of which directly impacts the implementation of routing protocol features such as multitopology.

In summary, the core requirement for IS-IS multitopology is the ability to differentiate between routes of different topologies through extended metrics. Lack of support for this crucial feature is the primary reason for its failure.

Authoritative Link:

Cisco Official Documentation on IS-IS Multitopology:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/configuration/15-mt/irs-15-mt-book/irs-isis-multi-topo.html (Refer to the section on Extended Metrics)

Question: 78

ASN 65001 is peering with ASN 65002 to exchange IPv6 BGP routes. All routes that originate in ASN 65001 have a standard community value of 65001:100, and
ASN 65002 is allowed to advertise only 2001:db8:aaaa::/48. An engineer needs to update the ASN 65001 route-filtering configuration to meet these conditions:

➾ Looped routes into ASN 65001 and routes that have traversed 10 or more ASNs must be denied.

➾ Routes accepted into ASN 65001 must be assigned a community value of 65001:200.

Which configuration must be engineer apply to the ASN 65001 border router?

A.

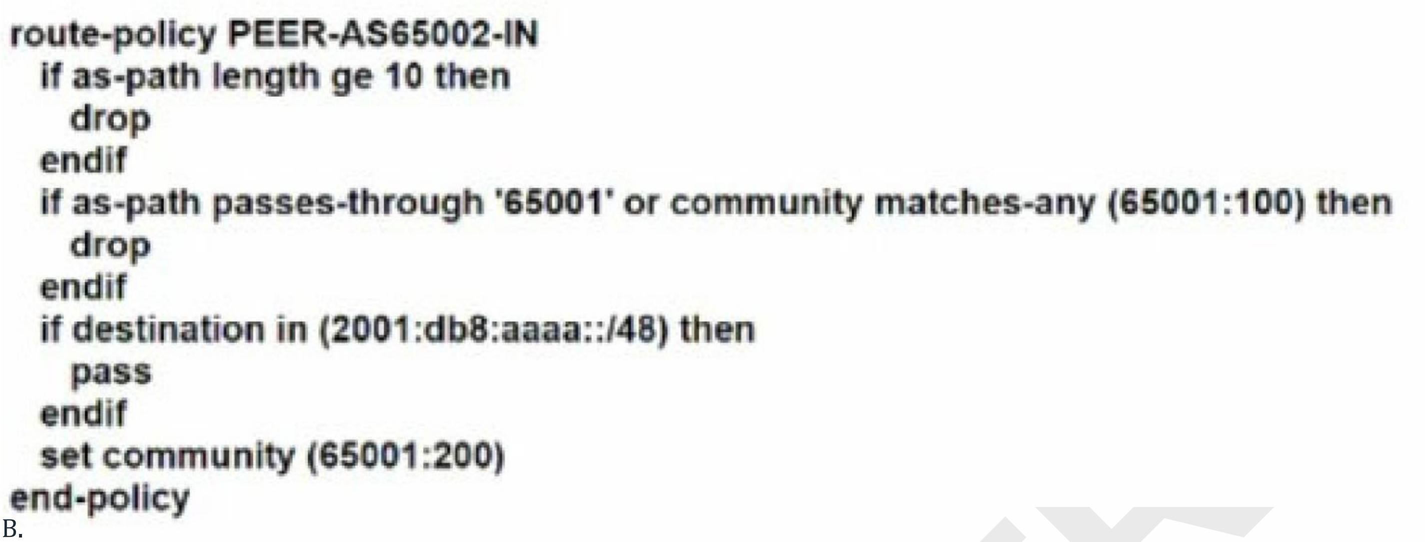

```
route-policy PEER-AS65002-IN
  if as-path length ge 10 then
    drop
  endif
  if as-path passes-through '65001' or community matches-any (65001:100) then
    drop
  endif
  if destination in (2001:db8:aaaa::/48) then
    pass
  endif
  set community (65001:200)
end-policy
```

B.

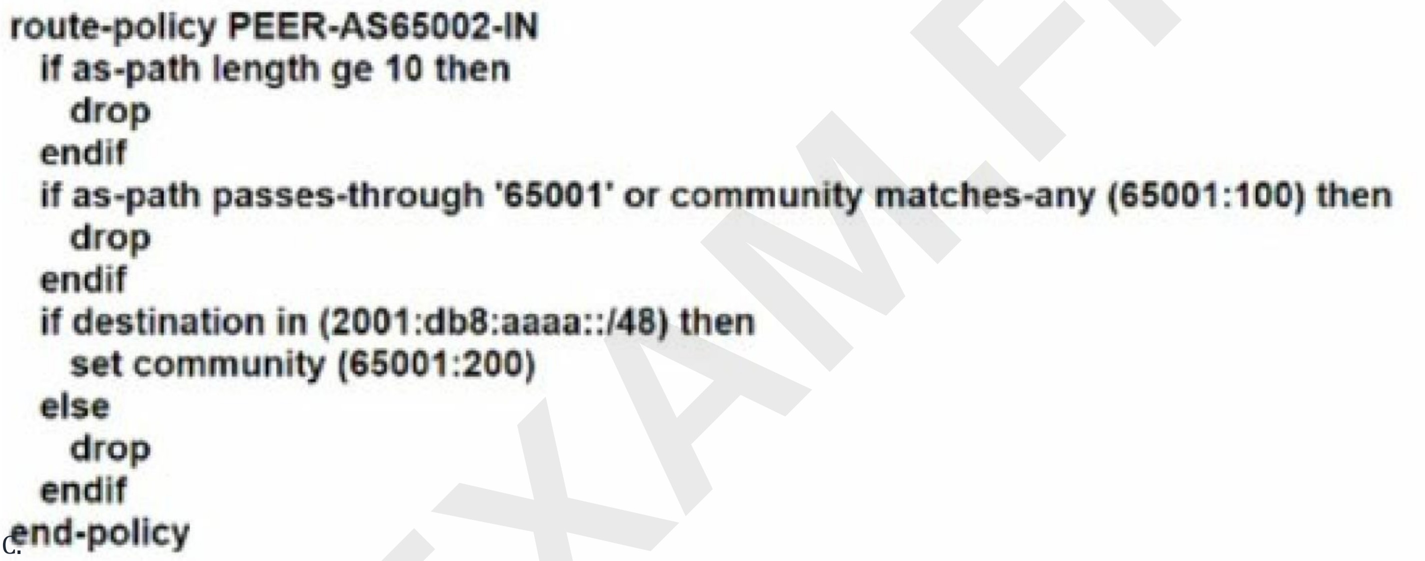

```
route-policy PEER-AS65002-IN
  if as-path length ge 10 then
    drop
  endif
  if as-path passes-through '65001' or community matches-any (65001:100) then
    drop
  endif
  if destination in (2001:db8:aaaa::/48) then
    set community (65001:200)
  else
    drop
  endif
end-policy
```

C.

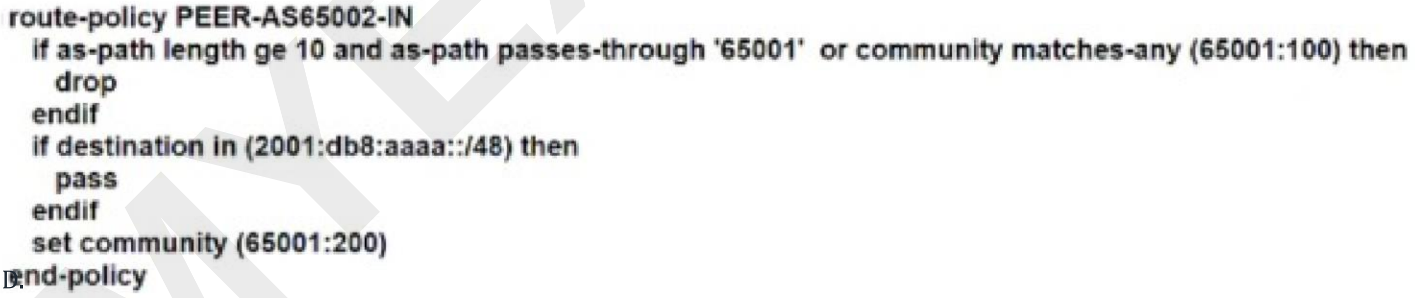

```
route-policy PEER-AS65002-IN
  if as-path length ge 10 and as-path passes-through '65001'  or community matches-any (65001:100) then
    drop
  endif
  if destination in (2001:db8:aaaa::/48) then
    pass
  endif
  set community (65001:200)
end-policy
```

D.

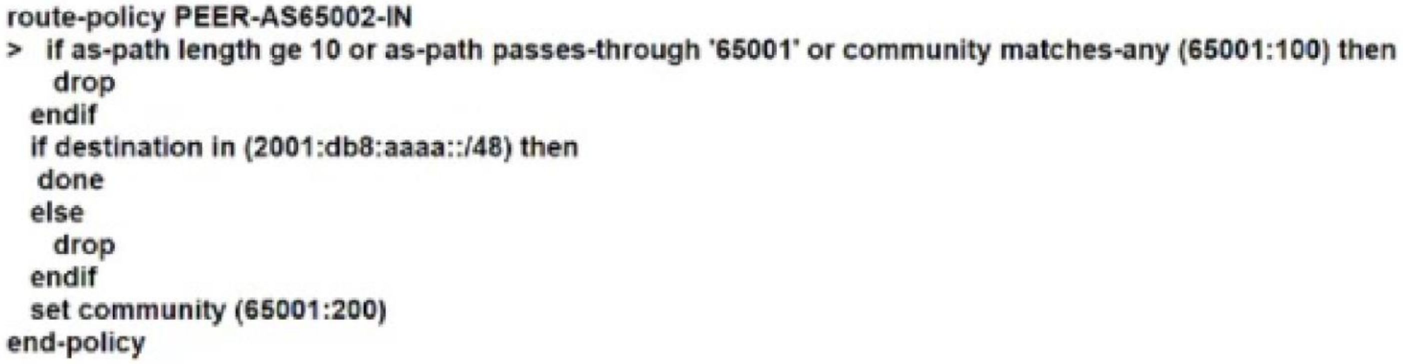

```
route-policy PEER-AS65002-IN
>  if as-path length ge 10 or as-path passes-through '65001' or community matches-any (65001:100) then
     drop
   endif
   if destination in (2001:db8:aaaa::/48) then
    done
   else
     drop
   endif
   set community (65001:200)
end-policy
```

Answer: B

Explanation:

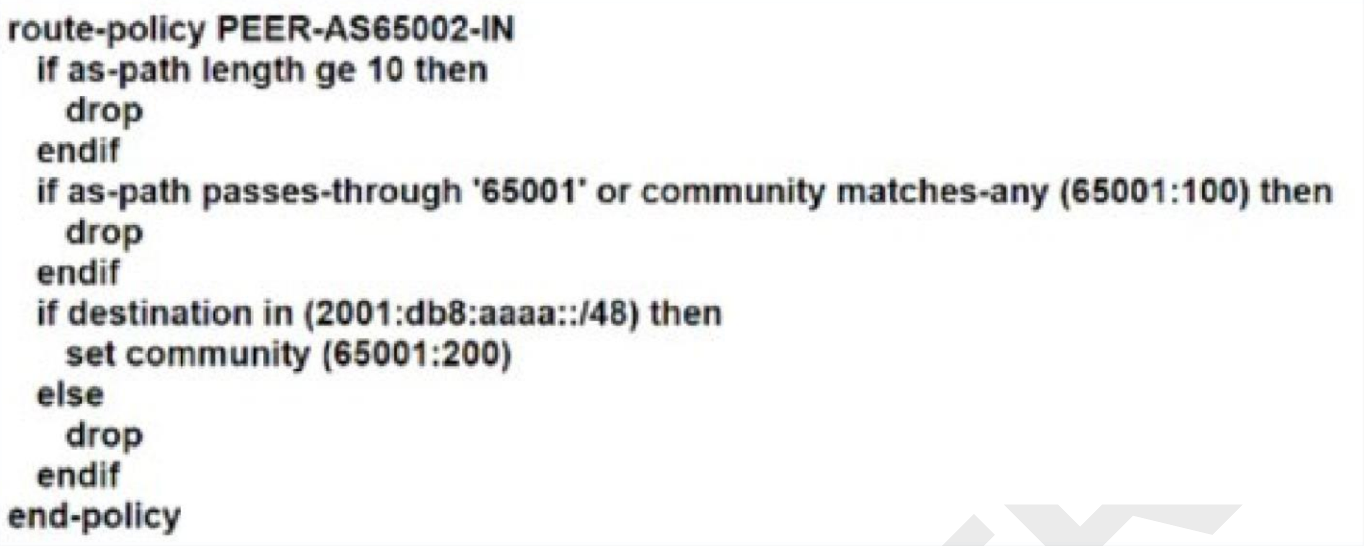

```
route-policy PEER-AS65002-IN
  if as-path length ge 10 then
    drop
  endif
  if as-path passes-through '65001' or community matches-any (65001:100) then
    drop
  endif
  if destination in (2001:db8:aaaa::/48) then
    set community (65001:200)
  else
    drop
  endif
end-policy
```

Question: 79

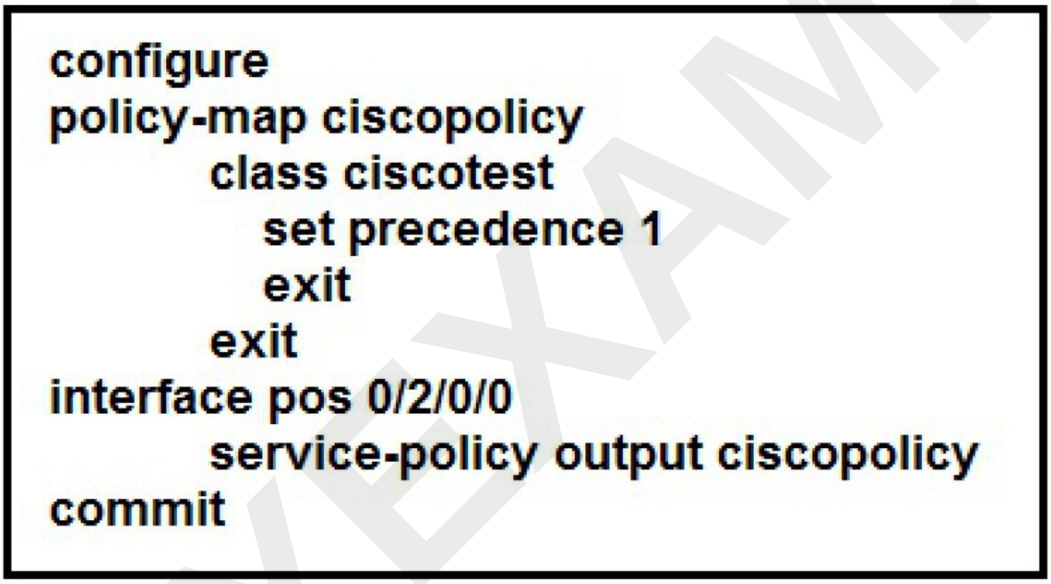

```
configure
policy-map ciscopolicy
        class ciscotest
           set precedence 1
           exit
        exit
interface pos 0/2/0/0
        service-policy output ciscopolicy
commit
```

Refer to the exhibit. An engineer needs to implement this QoS policy on customer's network due to ongoing slow network issues. What will be the effect on the network when the engineer implements this configuration?

A. Traffic that is identified in the ciscopolicy class map will be remarked from IP precedence 1 to DSCP AF11 when it exits the pos0/2/0/0 interface

B. Traffic that is identified in the ciscotest class map will be marked with IP precedence 1 when it exits the

pos0/2/0/0 interface
C. Traffic that is identified in the ciscopolicy class map will be marked with IP precedence 1 when it enters the pos0/2/0/0 interface
D. Traffic that is identified in the ciscotest class map will be remarked from IP precedence 1 to DSCP AF11 when it enters the pos0/2/0/0 interface

Answer: B

Explanation:

Traffic that is identified in the ciscotest class map will be marked with IP precedence 1 when it exits the pos0/2/0/0 interface.

Question: 80

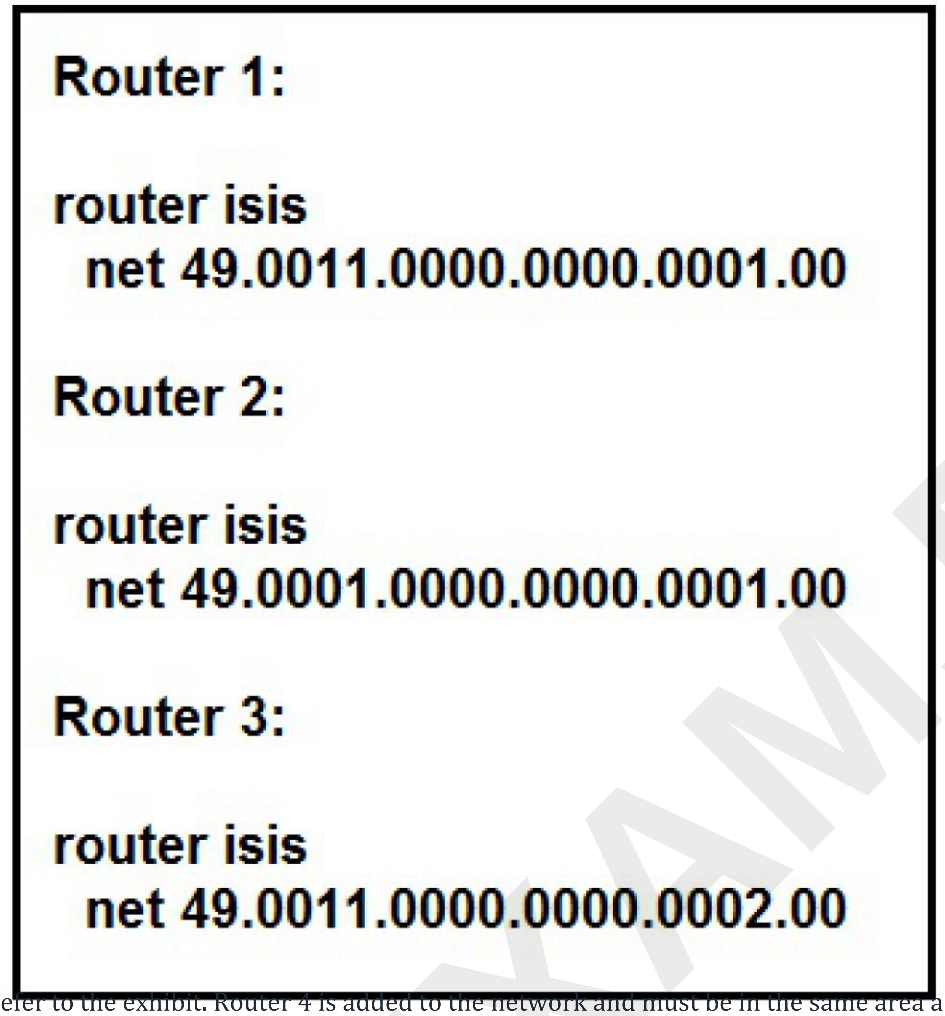

```
Router 1:

router isis
  net 49.0011.0000.0000.0001.00

Router 2:

router isis
  net 49.0001.0000.0000.0001.00

Router 3:

router isis
  net 49.0011.0000.0000.0002.00
```

Refer to the exhibit. Router 4 is added to the network and must be in the same area as router 1. Which NET should the engineer assign?

A. 49.0111.0000.0000.0001.00
B. 49.0001.0000.0000.0004.00
C. 49.0011.0000.0000.0002.00
D. 49.0011.0000.0000.0003.00

Answer: D

Explanation:

49.0011.0000.0000.0003.00.