Cisco

(300-420)

Designing Cisco Enterprise Networks (ENSLD)

Total: **303 Questions**

Link:

Question: 1

Which two BGP features will result in successful route exchanges between eBGP neighbors sharing the same AS number? (Choose two.)

A. advertise-best-external
B. bestpath as-path ignore
C. client-to-client reflection
D. as-override
E. allow-as-in

Answer: DE

Explanation:

The correct answer is **D. as-override** and **E. allow-as-in**. Let's examine why.

Standard eBGP (External Border Gateway Protocol) operates on the principle that neighbor routers must reside in different Autonomous Systems (AS). When two eBGP routers are configured with the same AS number, BGP's built-in loop prevention mechanisms prevent route advertisements. BGP by default will drop any update it receives with its own AS number already in the AS path.

D. as-override: This feature allows a BGP router to replace the AS number of outgoing updates with the local router's AS number. This effectively hides the shared AS number and deceives the receiving router into thinking it's an external route from a different AS, thus enabling route advertisement and exchange. The router receiving the update will not reject it because it does not see its own AS number in the path.

E. allow-as-in: This command allows a BGP router to accept route updates even if its own AS number is already in the AS path. This directly overrides the default loop-prevention behavior and permits BGP to establish neighborship and exchange routes even when the same AS number is shared between eBGP neighbors. This is particularly useful in lab or test environments or special operational scenarios.

The options A. advertise-best-external and B. bestpath as-path ignore do not impact the core issue of identical AS numbers between eBGP peers. Option C. client-to-client reflection is relevant to iBGP (Internal BGP), not eBGP, and is also not directly related to AS number misconfigurations.

Therefore, the combination of **as-override** and **allow-as-in** are the two BGP features that, when configured on respective peers, will allow for successful route exchange when eBGP neighbors inadvertently share the same AS number. They allow you to essentially bypass the inherent loop prevention mechanisms built into BGP.

Authoritative links for further research:

Cisco Documentation on BGP AS-override: https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/202116-Configure-BGP-AS-Override.html
Cisco Documentation on BGP Allow-AS-in: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/15-mt/irg-15-mt-book/bgp-adv-features.html

Question: 2

A customer with an IPv4 only network topology wants to enable IPv6 connectivity while preserving the IPv4 topology services. The customer plans to migrate IPv4 services to the IPv6 topology, then decommission the IPv4 topology. Which topology supports these requirements?

A. dual stack

B. 6VPE
C. 6to4
D. NAT64

Answer: A

Explanation:

The correct answer is A, dual stack. Dual stack allows a network to run both IPv4 and IPv6 simultaneously. This is crucial for the customer's migration plan because it allows them to introduce IPv6 connectivity without disrupting existing IPv4 services. Services can be migrated to the IPv6 network gradually while the IPv4 network remains operational. This approach minimizes service interruption and allows for thorough testing of the new IPv6 infrastructure before fully switching over. The customer's plan to first enable IPv6 alongside IPv4 and then migrate services, is perfectly aligned with the nature of dual-stack deployment. Options B, C and D do not align with this approach. 6VPE (IPv6 VPN over MPLS) is a tunneling mechanism for IPv6 over an IPv4 MPLS backbone. While useful, it doesn't provide a true dual-stack environment. 6to4 is another IPv6 transition mechanism that encapsulates IPv6 packets within IPv4, typically for connecting islands of IPv6 networks. It is not suitable for a complete migration scenario. Lastly, NAT64 translates IPv6 to IPv4, and is appropriate for communication between IPv6-only hosts and IPv4 servers, but it does not address the goal of fully transitioning services to a native IPv6 network, whilst still allowing existing IPv4 connectivity. Therefore, dual-stack networking provides the most appropriate method for a staged migration of services from IPv4 to IPv6 while preserving existing IPv4 functionality.

Here are some authoritative links for further research:

Cisco - IPv6 Transition/Migration Options:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-mt/ipv6-15-mt-book/ipv6-trans.html (Provides comprehensive information on different IPv6 migration strategies including dual-stack.)
IETF RFC 4213 - Basic Transition Mechanisms for IPv6:https://datatracker.ietf.org/doc/html/rfc4213 (Details the architecture and basic transition mechanisms, including dual-stack.)
Wikipedia - IPv6 Transition Mechanisms:https://en.wikipedia.org/wiki/IPv6_transition_mechanism (Offers a broad overview of various IPv6 migration approaches.)

Question: 3

DRAG DROP -
An engineer is designing an addressing plan for a small business using a single /24 network. Each department must have its own subnet. Drag and drop the subnets from the left onto the requirements of the department they fulfill on the right. Not all options are used.
Select and Place:

Answer Area

10.1.1.16/27	5 hosts for Human Resources
10.1.1.96/26	18 hosts for Facilities
10.1.1.96/28	32 hosts for Engineering
10.1.1.112/29	12 hosts for Finance
10.1.1.8/28	
10.1.1.0/26	
10.1.1.64/27	

Answer:

Answer Area

10.1.1.16/27	10.1.1.112/29
10.1.1.96/26	10.1.1.64/27
10.1.1.96/28	10.1.1.0/26
10.1.1.112/29	10.1.1.96/28
10.1.1.8/28	
10.1.1.0/26	
10.1.1.64/27	

Question: 4

A company is running BGP on a single router, which has two connections to the same ISP. Which BGP feature ensures traffic is load balanced across the two links to the ISP?

A. Multihop
B. Multipath Load Sharing
C. Next-Hop Address Tracking
D. AS-Path Prepending

Answer: B

Explanation:

The correct answer is **B. Multipath Load Sharing**.

Here's why:

Multipath Load Sharing in BGP allows a router to utilize multiple paths to the same destination (in this case,

the ISP) if they meet certain criteria. When the router receives multiple routes for the same prefix from the same AS (as is the case with two connections to the same ISP), BGP will normally choose the "best" path.

Multipath Load Sharing relaxes this restriction and permits the use of multiple equal-cost paths for forwarding traffic. This distributes traffic across multiple links, achieving load balancing.

A. **Multihop** is used to establish BGP sessions with routers not directly connected. It's irrelevant for load balancing across directly connected links.

C. **Next-Hop Address Tracking** is a feature that tracks the availability of the next-hop IP address for a specific BGP route, not for load balancing.

D. **AS-Path Prepending** manipulates the AS-PATH attribute to influence path selection by making a path less desirable. It is not used for load balancing across multiple paths.

By enabling BGP multipath, the router can utilize both connections to the ISP, effectively load balancing outbound traffic without relying on equal-cost path metrics for IGP like OSPF or EIGRP. This improves bandwidth utilization and overall network resilience.

Relevant concepts:

BGP Path Selection: BGP has a defined process for selecting the best path, but multipath bypasses this selection.

Equal-Cost Multipath (ECMP): While not direct ECMP in the IGP sense, multipath is a BGP feature that achieves similar load balancing.

Authoritative Links:

Cisco Documentation on BGP Multipath
RFC 4271 - A Border Gateway Protocol 4 (BGP-4) (Section 9.1.2 discusses path selection)

Question: 5

Company A recently acquired another company. Users of the newly acquired company must be able to access a server that exists on Company A's network, both companies use overlapping IP address ranges. Which action conserves IP address space and provides access to the server?

A. Use a single IP address to create overload NAT
B. Use a single IP address to create a static NAT entry
C. Build one-to-one NAT translation for every user that needs access
D. Re-IP overlapping address space in the acquired company

Answer: A

Explanation:

The correct answer is A: Use a single IP address to create overload NAT (PAT). Here's why:

Overlapping IP address ranges are a common issue when companies merge. They create routing conflicts as devices in both networks might have the same IP address. Network Address Translation (NAT), specifically Port Address Translation (PAT) which is often called "overload" NAT, provides a solution. PAT allows multiple devices behind a router to share a single public IP address by using different port numbers. In this scenario, the newly acquired company's users can access the server in Company A by having their private IP addresses translated to a single public IP address as they leave their network. This effectively masks the overlapping IP addresses, and routes traffic correctly. Each user's connection is tracked with a unique port number, allowing the responses from the server to be correctly routed back. Option B (static NAT) is not suitable as each private IP would need a dedicated public IP, wasting space. Option C (one-to-one NAT) would also be inefficient and

impractical for numerous users. Option D (re-IP) while a possible solution, is disruptive and requires significant planning and effort which is not ideal and should be done when other solutions are infeasible. Overload NAT is the most efficient, least resource intensive method for solving this common issue.

Authoritative Links:

Cisco on NAT:https://www.cisco.com/c/en/us/support/docs/ip/network-address-translation-nat/10185-natfaq.html
TechTarget on PAT:https://www.techtarget.com/searchnetworking/definition/port-address-translation-PAT

Question: 6

Which design consideration should be observed when EIGRP is configured on Data Center switches?

A.Perform manual summarization on all Layer 3 interfaces to minimize the size of the routing table.
B.Prevent unnecessary EIGRP neighborships from forming across switch virtual interfaces.
C.Lower EIGRP hello and hold timers to their minimum settings to ensure rapid route reconvergence.
D.Configure multiple EIGRP autonomous systems to segment Data Center services and applications.

Answer: A

Explanation:

Here's a detailed justification for why option A, performing manual summarization, is the most appropriate design consideration when configuring EIGRP on Data Center switches:

In data center environments, where scalability and performance are crucial, a large routing table can strain switch resources and slow down convergence times. EIGRP, while efficient, can still contribute to this issue if not properly managed. Manual summarization on Layer 3 interfaces aggregates contiguous network prefixes into a single, smaller advertisement. This significantly reduces the number of routes advertised and stored, leading to faster lookups, reduced memory consumption, and more stable networks. Option B is less critical as switch virtual interfaces typically don't require extensive adjacency control in a well-planned data center topology. Option C, lowering timers aggressively, can lead to increased network instability due to frequent hello and hold timer expiration causing unnecessary recalculations. Option D, using multiple AS, adds complexity and is generally not required for segmenting data center services, which are more effectively managed with techniques like VLANs and VXLANs. Therefore, manual summarization (A) directly addresses scalability concerns within data centers using EIGRP. It simplifies route distribution, reduces router resource utilization and stabilizes the network, aligning with the core design objectives for a robust data center network.

Authoritative Links for Further Research:

Cisco Design Guides - Data Center Infrastructure:https://www.cisco.com/c/en/us/solutions/data-center/index.html (Explore design recommendations focusing on routing protocols and scalability)
EIGRP Configuration Guides:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/15-mt/ire-15-mt-book.html (Focus on route summarization sections within this guide)
Cisco Validated Designs (CVDs): Search within Cisco documentation for data center focused CVDs that include routing protocol considerations, particularly those that include EIGRP best practices.

Question: 7

Which design consideration must be made when using IPv6 overlay tunnels?

A.Overlay tunnels that connect isolated IPv6 networks are considered a final IPv6 network architecture.
B.Overlay tunnels should only be considered as a transition technique toward a permanent solution.
C.Overlay tunnels should be configured only between border devices and require only the IPv6 protocol stack.
D.Overlay tunneling encapsulates IPv4 packets in IPv6 packets for delivery across an IPv6 infrastructure.

Answer: B

Explanation:

The correct answer is **B. Overlay tunnels should only be considered as a transition technique toward a permanent solution.** Here's why:

Overlay tunnels, particularly IPv6 over IPv4 (or vice-versa), are primarily designed to bridge the gap during network migrations. They allow disparate network technologies (like IPv4 and IPv6) to communicate without requiring an immediate, end-to-end upgrade of the entire infrastructure. Option A is incorrect because overlay tunnels are not a final IPv6 network architecture. They introduce complexity and are not ideal for long-term network design. Option C is incorrect because overlay tunnels can be configured beyond border devices and often involve both IPv4 and IPv6 protocol stacks on supporting devices depending on which protocol is encapsulated. Option D is wrong because typically, it would be an IPv6 packet encapsulated in an IPv4 packet when traversing an IPv4-only network, though the reverse can happen.

Overlay tunnels are like temporary bridges; they add an extra layer of encapsulation and introduce overhead. Relying on them indefinitely is not scalable or efficient. Ideally, the goal should be to transition to a native, end-to-end IPv6 infrastructure rather than maintaining a reliance on tunnels. Overlays can complicate routing and troubleshooting, and they might not offer the same performance as a natively implemented protocol. Therefore, they should be viewed as a step along the way, not the destination itself.

Further research on IPv6 transition technologies and overlay tunnels can be found at:

IETF RFC 4213: Basic Transition Mechanisms for IPv6: https://datatracker.ietf.org/doc/html/rfc4213 **Cisco IPv6 Transition Technologies:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-mt/ipv6-15-mt-book/ipv6-trans.html

Question: 8

When a network is designed using IS-IS, which two circuit types are supported? (Choose two.)

A. nonbroadcast multiaccess
B. multiaccess
C. point-to-multipoint
D. nonbroadcast
E. point-to-point

Answer: BE

Explanation:

The correct answer, point-to-point and nonbroadcast, aligns with IS-IS's operational characteristics. IS-IS, a link-state routing protocol, primarily functions over two types of logical circuits. Point-to-point circuits are direct connections between two routers, where no other intermediate device is involved; these are straightforward and form the backbone of many IS-IS networks. In contrast, nonbroadcast circuits, often found in frame relay or ATM environments, are multiaccess but do not support multicast or broadcast. Thus, IS-IS must adapt its behavior, relying on manually configured neighbor lists. It's crucial to distinguish nonbroadcast from multiaccess, as the latter implies broadcast capabilities, which IS-IS also handles through

a designated intermediate system (DIS). The 'point-to-multipoint' option is specifically designed for IP addressing where one endpoint communicates to many, which isn't a circuit type directly supported by IS-IS link layer. The remaining option, 'multiaccess', is not directly supported by the IS-IS routing protocol, where it depends on whether broadcast is supported or not. IS-IS is designed to work efficiently with point-to-point and nonbroadcast networks, making 'point-to-point' and 'nonbroadcast' the correct circuit types.

For further information, refer to:

1. **Cisco Documentation on IS-IS:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/configuration/15-sy/irs-15-sy-book/irs-isis-basic.html
2. **Juniper Networks Understanding IS-IS:** https://www.juniper.net/documentation/us/en/software/junos/routing/topics/concept/isis-understanding.html

Question: 9

A network solution is being designed for a company that connects to multiple Internet service providers. Which Cisco proprietary BGP path attribute will influence outbound traffic flow?

A.Local Preference
B.MED
C.Weight
D.AS Path
E.Community

Answer: C

Explanation:

The correct answer is **C. Weight**.

Weight is a Cisco proprietary BGP attribute that is locally significant to the router where it's configured. It is the first attribute examined by a router when choosing the best path. A higher weight value indicates a more preferred path for outbound traffic. In scenarios where a company connects to multiple Internet service providers (ISPs), configuring weight allows precise control over which ISP connection is used for egress traffic. Specifically, if a router receives multiple paths for the same destination network from different ISPs, the path with the highest weight is chosen. The other attributes like Local Preference, MED (Multi-Exit Discriminator), AS Path, and Community are considered later in the path selection process or influence inbound traffic. For outbound traffic control, weight takes precedence. Local preference is influential in a particular AS (Autonomous System) and is used to choose between paths to the same destination offered by different routers within the AS. MED is used to influence inbound traffic selection by neighboring ASes. The AS path influences path selection based on the length of the AS Path. The community attribute helps in route filtering and manipulation but does not directly influence outbound traffic like weight. Weight is the primary way to influence outbound routing decisions based on an administrator's specific preference of the ISP connections. Therefore, in the context of managing traffic flow across multiple ISPs, the weight attribute offers a granular and local method for choosing the most preferred path.

Further research:

Cisco documentation on BGP Path Attributes: https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13705-25.html
RFC 4271 (Border Gateway Protocol 4): https://datatracker.ietf.org/doc/html/rfc4271

Question: 10

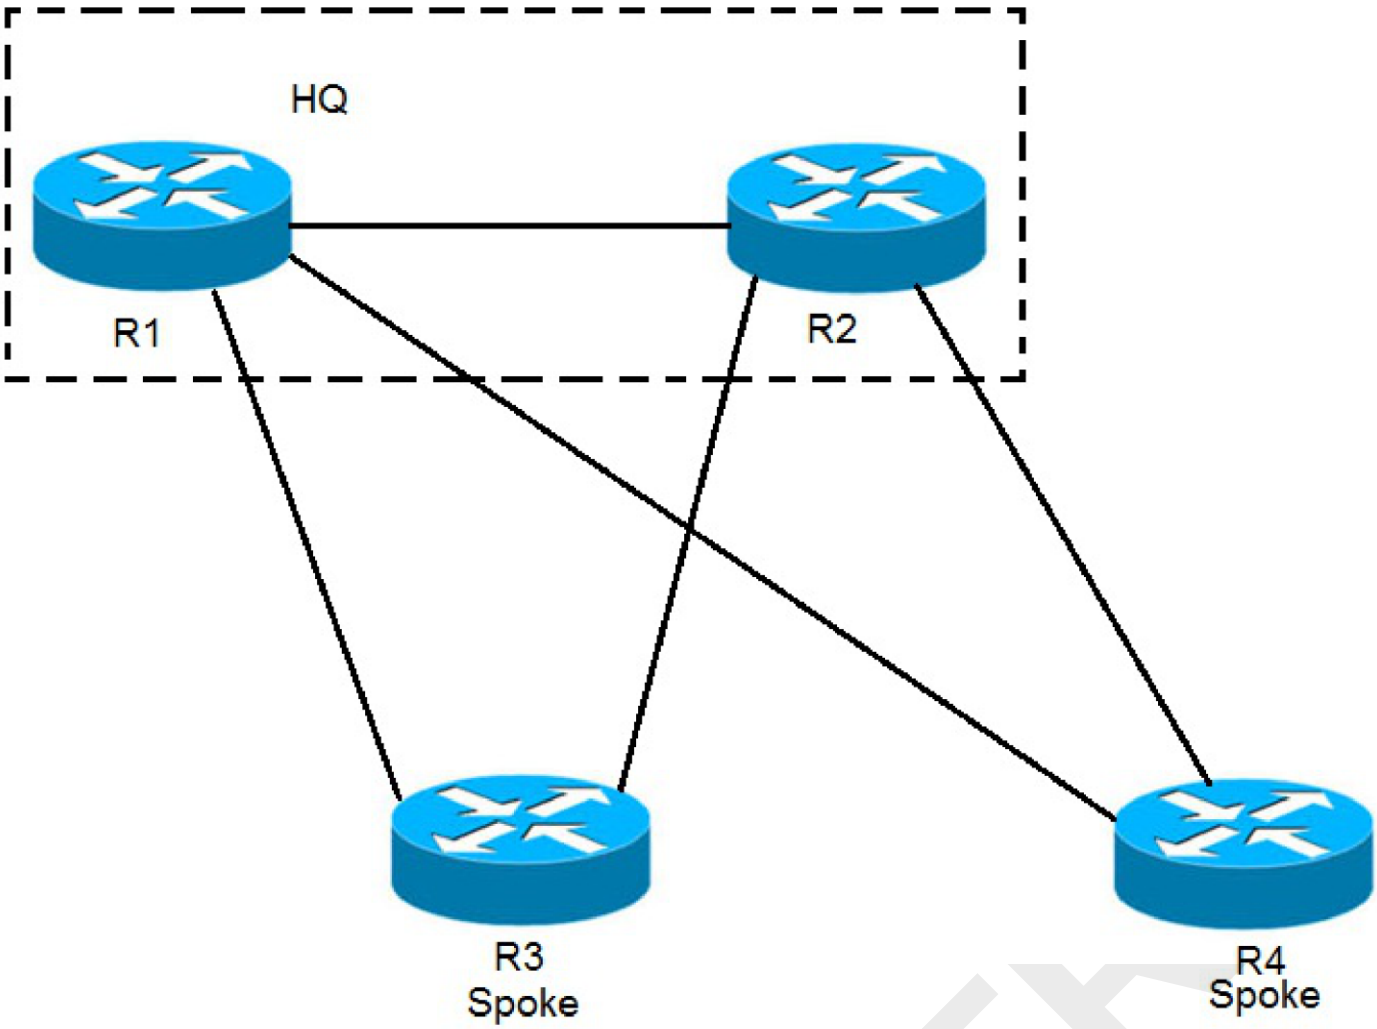

Refer to the exhibit. EIGRP has been configured on all links. The spoke nodes have been configured as EIGRP stubs, and the WAN links to R3 have higher bandwidth and lower delay than the WAN links to R4. When a link failure occurs at the R1-R2 link, what happens to traffic on R1 that is destined for a subnet attached to R2?

A.R1 has no route to R2 and drops the traffic
B.R1 load-balances across the paths through R3 and R4 to reach R2
C.R1 forwards the traffic to R3, but R3 drops the traffic
D.R1 forwards the traffic to R3 in order to reach R2

Answer: A

Explanation:

R1 has no route to R2 and drops the traffic.

Question: 11

A company is using OSPF between its HQ location and a branch office. HQ is assigned area 0 and the branch office is assigned area 1. The company purchases a second branch office, but due to circuit delays to HQ, it decides to connect the new branch office to the existing branch office as a temporary measure. The new branch office is assigned to area 2. Which OSPF configuration enables all three locations to exchange routes?

A.The existing branch office must be configured as a stub area
B.A virtual link must be configured between the new branch office and HQ

C.A sham link must be configured between the new branch office and HQ

D.The new branch office must be configured as a stub area

Answer: B

Explanation:

Here's a detailed justification for why option B, "A virtual link must be configured between the new branch office and HQ," is the correct answer, along with relevant supporting information:

OSPF mandates that all areas must connect to the backbone area (area 0), either directly or via a virtual link. In the scenario, the new branch office (area 2) is connected to the existing branch office (area 1), which is not directly connected to area 0. This violates the OSPF area connection rule. Since we are not changing the branch connection to directly connect to area 0, it is necessary to connect the branch office area 2 to area 0 using virtual link.

A virtual link is a logical tunnel established between two ABRs (Area Border Routers) that are not directly connected. The virtual link uses an existing OSPF area as a transit or intermediate area for OSPF packets. In this scenario, the existing branch office's ABR serves as the transit area (Area 1). By establishing a virtual link between the ABR in Area 2 and an ABR in Area 0, routes can be exchanged between all three areas.

Option A, configuring the existing branch office as a stub area, is incorrect as stub areas do not allow external or summary routes. In this case, we want to exchange routes to all network areas.

Option C, a sham link, is used with multiprotocol label switching (MPLS) and BGP, and is not applicable in this pure OSPF scenario.

Option D, making the new branch office a stub area, while technically possible if no external routes are needed for that area, does not solve the fundamental OSPF connectivity problem of area 2 being disconnected from area 0. Furthermore, it would prevent the new branch office from learning all the OSPF routes.

Therefore, option B, utilizing a virtual link, is the only solution that ensures proper OSPF route exchange between all three areas without changing their underlying topology, making it the correct choice.

Authoritative Links for Further Research:

Cisco OSPF Virtual Links:https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-virtual-links.html
RFC 2328 - OSPF Version 2: (Refer to section 16.2 for detailed information on virtual links): https://www.rfc-editor.org/rfc/rfc2328
OSPF Area Types:https://networklessons.com/ospf/ospf-area-types

Question: 12

Which method will filter routes between EIGRP neighbors within the same autonomous system?

A.distribute-list

B.policy-based routing

C.leak-map

D.route tagging

Answer: A

Explanation:

The correct answer is **A. distribute-list**.

Distribute lists are a fundamental mechanism in EIGRP for filtering routes exchanged between neighbors within the same autonomous system. They use access control lists (ACLs) to define which routes are permitted or denied. When applied to an EIGRP process, the distribute list inspects routes before they are advertised, preventing unwanted routes from being propagated. This filtering occurs at the point of sending or receiving updates, thus controlling the routing table updates of EIGRP neighbors. They're effective for preventing redistribution issues or limiting advertisement based on network boundaries within the same EIGRP autonomous system. Distribute lists directly modify EIGRP routing behavior by preventing route information from being shared between specific neighbors, offering granular control. They provide a straightforward method of controlling routing behavior based on prefixes, enabling network administrators to dictate which routes are allowed or suppressed from EIGRP updates. This makes them ideal for controlling EIGRP route advertisement without needing more complex configurations. Policy-based routing, while powerful, is not used for filtering routing updates directly between EIGRP neighbors; it is used to control how traffic is routed based on factors other than destination IP. Leak-maps primarily deal with route redistribution between routing protocols and are not employed within a single EIGRP AS for filtering updates. Route tagging is a way to mark routes for subsequent filtering or policy application; however, route tagging does not by itself act as a filtering mechanism.

Cisco's Documentation on EIGRP Distribute ListsCisco's Documentation on EIGRP Route Filtering

Question: 13

What are two valid scaling techniques when an EIGRP network is designed that consists of more than 1000 routers? (Choose two.)

A.Use structured hierarchical topology with route summarization
B.Used sub-second timers
C.Use the distribute-list command to filter routes
D.Modify delay parameters on the links
E.Implement multiple EIGRP autonomous systems

Answer: AE

Explanation:

Here's a detailed justification for why options A and E are correct, and why the other options are not, when scaling a large EIGRP network (over 1000 routers):

Correct Options:

A. Use structured hierarchical topology with route summarization: Hierarchical design, similar to a cloud architecture with layers, breaks the network into smaller, manageable domains. This greatly reduces the size of routing tables and update traffic within each area. Route summarization at the boundaries of these areas hides detailed routing information, further reducing routing table size and update overhead. This scalability method significantly improves EIGRP's performance and convergence times.

E. Implement multiple EIGRP autonomous systems: Using multiple EIGRP autonomous systems (AS) creates routing domains. This divides the large network into smaller, independent groups of EIGRP routers, each with its own routing information. Inter-AS routing, like inter-region traffic in a cloud environment, occurs through redistribution, controlling the spread of routing information and limiting the scope of issues to the specific AS.

Incorrect Options:

B. Use sub-second timers: While sub-second timers might improve convergence in smaller networks, they increase the sensitivity of EIGRP to link flaps and lead to instability in a large network by causing frequent updates. This creates unnecessary overhead and can overwhelm routers. It's not a scalable solution.

C. Use the distribute-list command to filter routes: While distribute lists can control routing updates, they primarily aim for traffic engineering or security, not the overall scalability of a large EIGRP network with thousands of devices. They can add complexity and maintenance overhead, and doesn't address core EIGRP scaling issues.

D. Modify delay parameters on the links: Changing delay metrics is not intended for scalability, rather for influencing path selection for traffic. Modifying delays can create inconsistencies and make troubleshooting more difficult, while not significantly improving EIGRP's overall scalability in a very large network.

In Summary: The primary challenges with large EIGRP networks involve routing table size, the number of updates, and the convergence time. Using a structured hierarchical topology and multiple autonomous systems mitigates these by dividing the network into manageable parts, reducing the overall scope of routing information and update activity. This contrasts with the less scalable, and often detrimental, effects of timer modifications and increased use of distribute lists and delay modification.

Authoritative Links for Further Research:

Cisco EIGRP Documentation:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/15-mt/ire-15-mt-book.html (Look for sections on scalability, hierarchical design and summarization).

EIGRP Best Practices: Search for "EIGRP best practices for large networks" for various blogs and articles providing detailed insights.

Cisco's Design Guides: Search for documents regarding network design principles within the Cisco ecosystem, paying particular attention to routing protocols and scalability considerations.

Question: 14

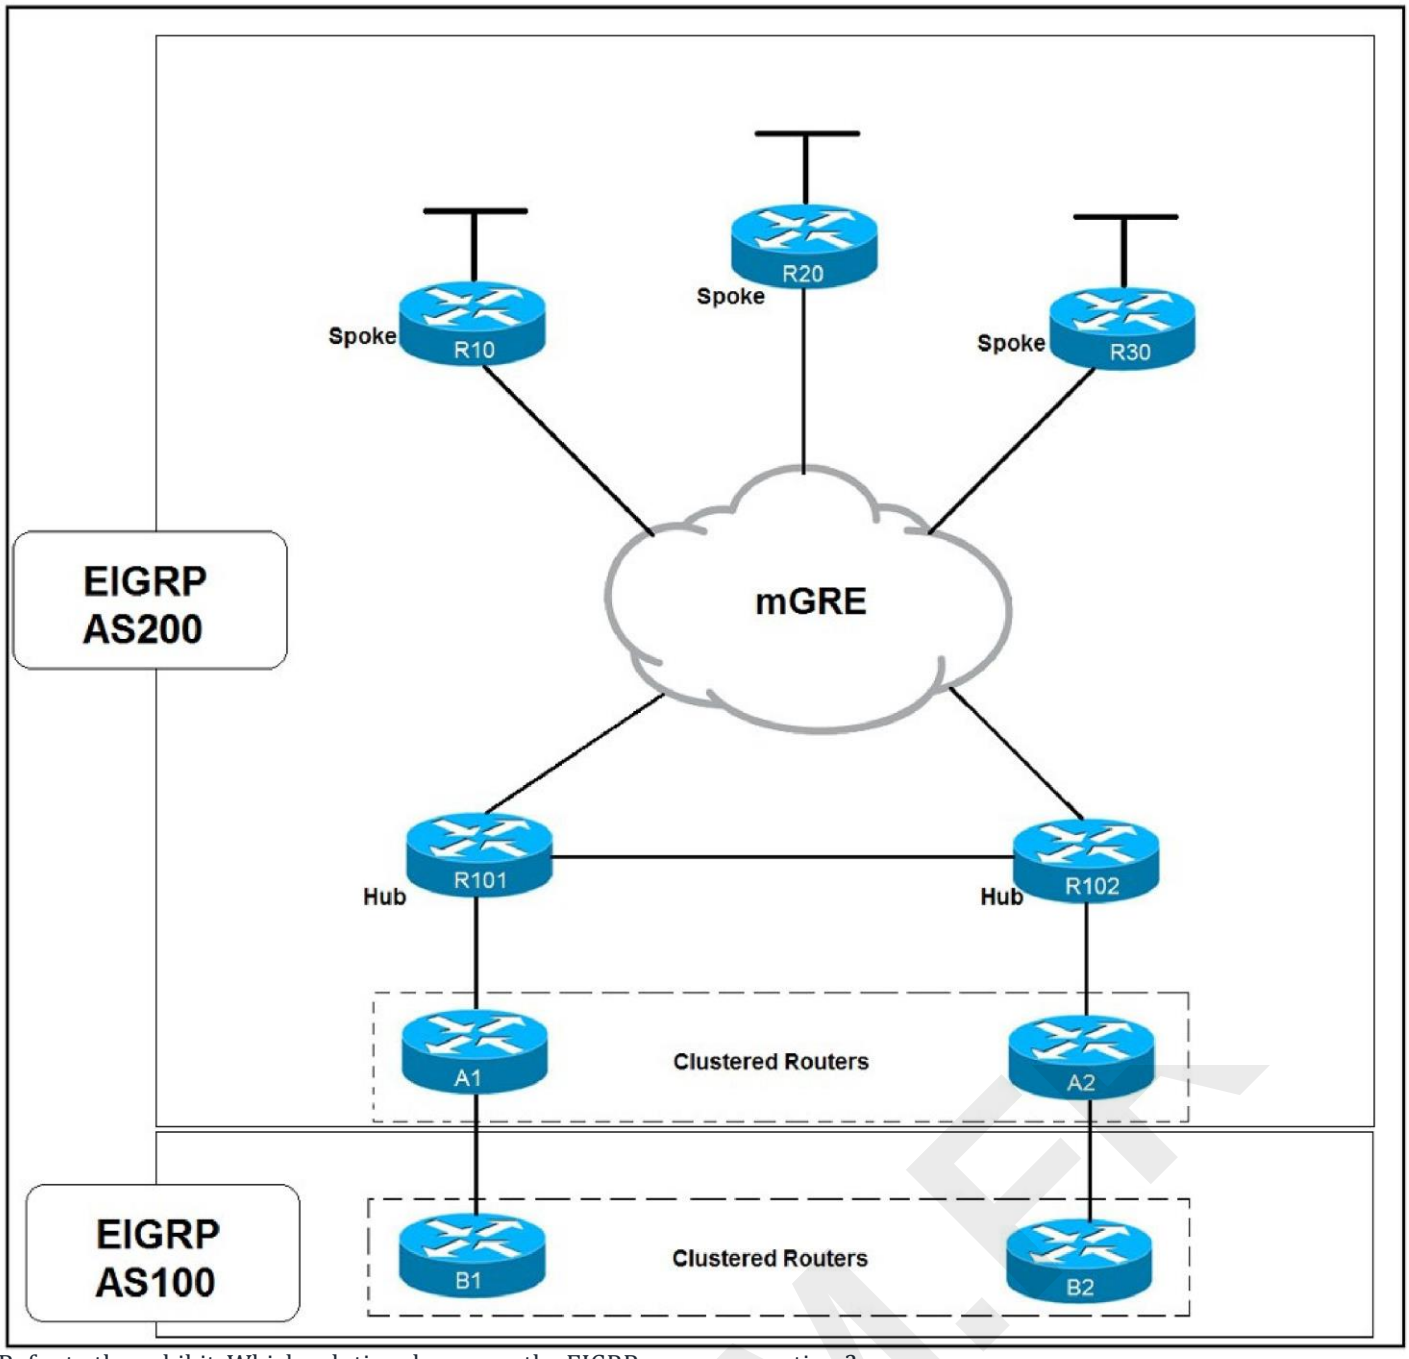

Refer to the exhibit. Which solution decreases the EIGRP convergence time?

A.Enable subsecond timers
B.Increase the hold time value
C.Increase the dead timer value
D.Enable stub routing on the spokes

Answer: D

Explanation:

Enable stub routing on the spokes.

Question: 15

A router running ISIS is showing high CPU and bandwidth utilization. An engineer discovers that the router is configured as L1/L2 and has L1 and L2 neighbors.
Which step optimizes the design to address the issue?

A. Make this router a DIS for each of the interfaces
B. Disable the default behavior of advertising the default route on the L1/L2 router
C. Configure the router to be either L1 or L2
D. Configure each interface as either L1 or L2 circuit type

Answer: D

Explanation:

Justification:

The core issue is the router's dual L1/L2 role within the ISIS network, leading to excessive CPU and bandwidth consumption. Operating as both L1 and L2 involves processing and propagating both L1 and L2 link-state PDUs (LSPs), increasing overhead.

Option A, making the router a Designated Intermediate System (DIS) for all interfaces, intensifies the problem. DIS responsibilities involve generating and distributing summary LSPs, further burdening the router. This would increase, not decrease CPU and bandwidth usage.

Option B, disabling default route advertisement, tackles a different issue. While important for route control, it doesn't address the core problem of excessive ISIS processing.

Option C, forcing the entire router to be either L1 or L2, lacks granularity. Some interfaces might function efficiently in L1, while others in L2. This approach removes needed flexibility and may not fully resolve the problem.

Option D, configuring interfaces as either L1 or L2, directly addresses the issue. ISIS L1 routers handle local area routing within an area. L2 routers handle routing between areas. A router in an area should not have a mix of L1 and L2, or a mix of interfaces with L1 or L2 roles, so the router won't process both types of LSPs for all neighbors. Thus, by assigning interface-specific roles (L1 or L2), the router avoids unnecessary L1/L2 processing duplication per interface. If we have L2 routers connecting to other L2 routers, this should not use L1, which could cause the L1 router to form an adjacency with the L2 router and cause the L1 router to process L2 routing information it doesn't need, creating overhead. This targeted approach minimizes resource consumption, making option D the most effective solution for optimizing the design.

Authoritative Links:

Cisco: Understanding IS-IS:https://www.cisco.com/c/en/us/support/docs/ip/integrated-intermediate-system-to-intermediate-system-is-is/13692-isis-basic.html
Juniper: Understanding IS-IS Concepts:
https://www.juniper.net/documentation/us/en/software/junos/routing/topics/concept/isis-concepts.html

These links provide in-depth information about IS-IS operation, level concepts (L1, L2), and the implications of mixing L1/L2 behavior on a single device.

Question: 16

Which two routing protocols allow for unequal cost load balancing? (Choose two.)

A.EIGRP
B.IS-IS
C.BGP
D.OSPF
E.RIPng

Answer: AC

Explanation:

The correct answer is A (EIGRP) and C (BGP). EIGRP (Enhanced Interior Gateway Routing Protocol) inherently supports unequal-cost load balancing through its variance command. This command allows traffic to be distributed across paths with different metrics, as long as the metric is within the defined variance multiplier of the feasible successor. This makes EIGRP a very flexible routing protocol when multiple paths with varying qualities exist. BGP (Border Gateway Protocol) also enables unequal-cost load balancing, albeit in a more policy-driven manner. While BGP typically chooses a single best path, path manipulation via attributes such as weight, local preference, and AS path prepending can result in traffic being distributed across multiple paths even with different cost metrics. Unlike distance-vector protocols like RIPng or link-state protocols such as OSPF and IS-IS, EIGRP and BGP have built-in mechanisms for influencing traffic flow beyond simply selecting the best route based on cost. OSPF and IS-IS perform equal-cost load balancing, meaning traffic is only distributed across paths with exactly the same metric. RIPng only uses hop count for path selection and load balancing. This difference stems from the underlying design principles of each protocol, with EIGRP and BGP focusing on more sophisticated routing decisions.

Further reading:

EIGRP Variance:https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/7664-56.html
BGP Path Selection:https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html

Question: 17

Which two steps can be taken to improve convergence in an OSPF network? (Choose two.)

A.Use Bidirectional Forwarding Detection
B.Merge all the areas into one backbone area
C.Tune OSPF parameters
D.Make all non-backbone areas stub areas
E.Span the same IP network across multiple areas.

Answer: AC

Explanation:

Here's a detailed justification for why options A and C are correct choices for improving OSPF convergence, along with why other options are incorrect:

A. Use Bidirectional Forwarding Detection (BFD): BFD is a lightweight protocol designed to quickly detect link failures. By using BFD with OSPF, routers can more rapidly identify adjacency losses and trigger the recalculation of routing tables. This significantly reduces the time OSPF takes to converge after a topology change, compared to relying solely on OSPF's native hello timers. This approach improves responsiveness and reduces downtime.https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/15-sy/iro-15-sy-book/ospf-bfd.html

C. Tune OSPF parameters: OSPF convergence is impacted by various parameters such as hello intervals, dead intervals, and retransmission intervals. By reducing hello and dead intervals, routers can detect failures quicker. Adjusting retransmission intervals can speed up the dissemination of link-state updates. However, this must be done with caution, as too aggressive tuning can lead to increased CPU and bandwidth utilization.
Therefore, parameter tuning provides a balanced approach to speeding up convergence in

OSPF.https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html#anc2

Why the other options are incorrect:

B. Merge all the areas into one backbone area: While a single-area OSPF network might seem simpler, it can drastically reduce scalability and increase the amount of link-state information that every router must store and process. This can slow convergence down because every router is involved in every SPF calculation. It also creates a single point of failure.

D. Make all non-backbone areas stub areas: Stub areas reduce link-state database size by limiting the types of LSAs they accept, but this can also limit routing information. Forcing all non-backbone areas to be stub areas will significantly limit the path options. It can also cause issues with external routes and thus reduces the ability to converge effectively.

E. Span the same IP network across multiple areas: Spanning the same IP subnet across multiple OSPF areas creates routing complexity and can lead to routing loops and suboptimal convergence behavior. Subnets should exist within one area for proper topology summarization and stable operations.

In summary, the correct steps for improved convergence in OSPF are to use BFD for faster link failure detection and to tune OSPF parameters for quicker responsiveness to topology changes, while carefully considering the ramifications of changes to the OSPF domain architecture.

Question: 18

Which OSPF area blocks LSA Type 3, 4 and 5, but allows a default summary route?

A.normal
B.stub
C.NSSA
D.totally stubby

Answer: D

Explanation:

Let's analyze OSPF area types and their LSA handling to understand why "totally stubby" is the correct answer.

Normal Area: A normal area allows all LSA types, including Type 1, 2, 3, 4, and 5. This makes it the most flexible but also the most complex in terms of routing updates. Therefore, it does not block LSA Type 3, 4 and 5, and is not a correct option.

Stub Area: A stub area blocks external routes (Type 5 LSAs) and inter-area routes (Type 3 & 4 LSAs), but it still allows summarized inter-area routes (Type 3) from the ABR. This is not the answer since it does not block all mentioned LSAs.

NSSA (Not-So-Stubby Area): NSSA is a stub area variation that allows importing external routes as Type 7 LSAs, which are then translated to Type 5 LSAs by the ABR. While it blocks Type 3 and 4, it does not provide a default summary and is not the correct option.

Totally Stubby Area: A totally stubby area blocks all inter-area routes (Type 3 & 4 LSAs), external routes (Type 5 LSAs), and replaces them with a default route injected by the ABR. This is exactly the required behavior stated in the question.

Therefore, a "totally stubby area" is the only type that satisfies the conditions of blocking LSA Types 3, 4, and

5 while allowing a default summary route. The ABR in a totally stubby area summarizes all external routes to a single default route, reducing routing table size and complexity within the area. This makes it the optimal choice for areas with single-exit points and a need for simplified routing.

In conclusion, 'totally stubby' areas provide the necessary functionality for simplifying routing within a specific area, blocking external and inter-area routes and relying on a default route.

Authoritative Links:

Cisco - OSPF Areas
RFC 2328 - OSPF Version 2 - Original OSPF specification, useful for deep understanding.
RFC 3101 - The OSPF Not-So-Stubby Area (NSSA) Option - For understanding the behavior of NSSAs.

Question: 19

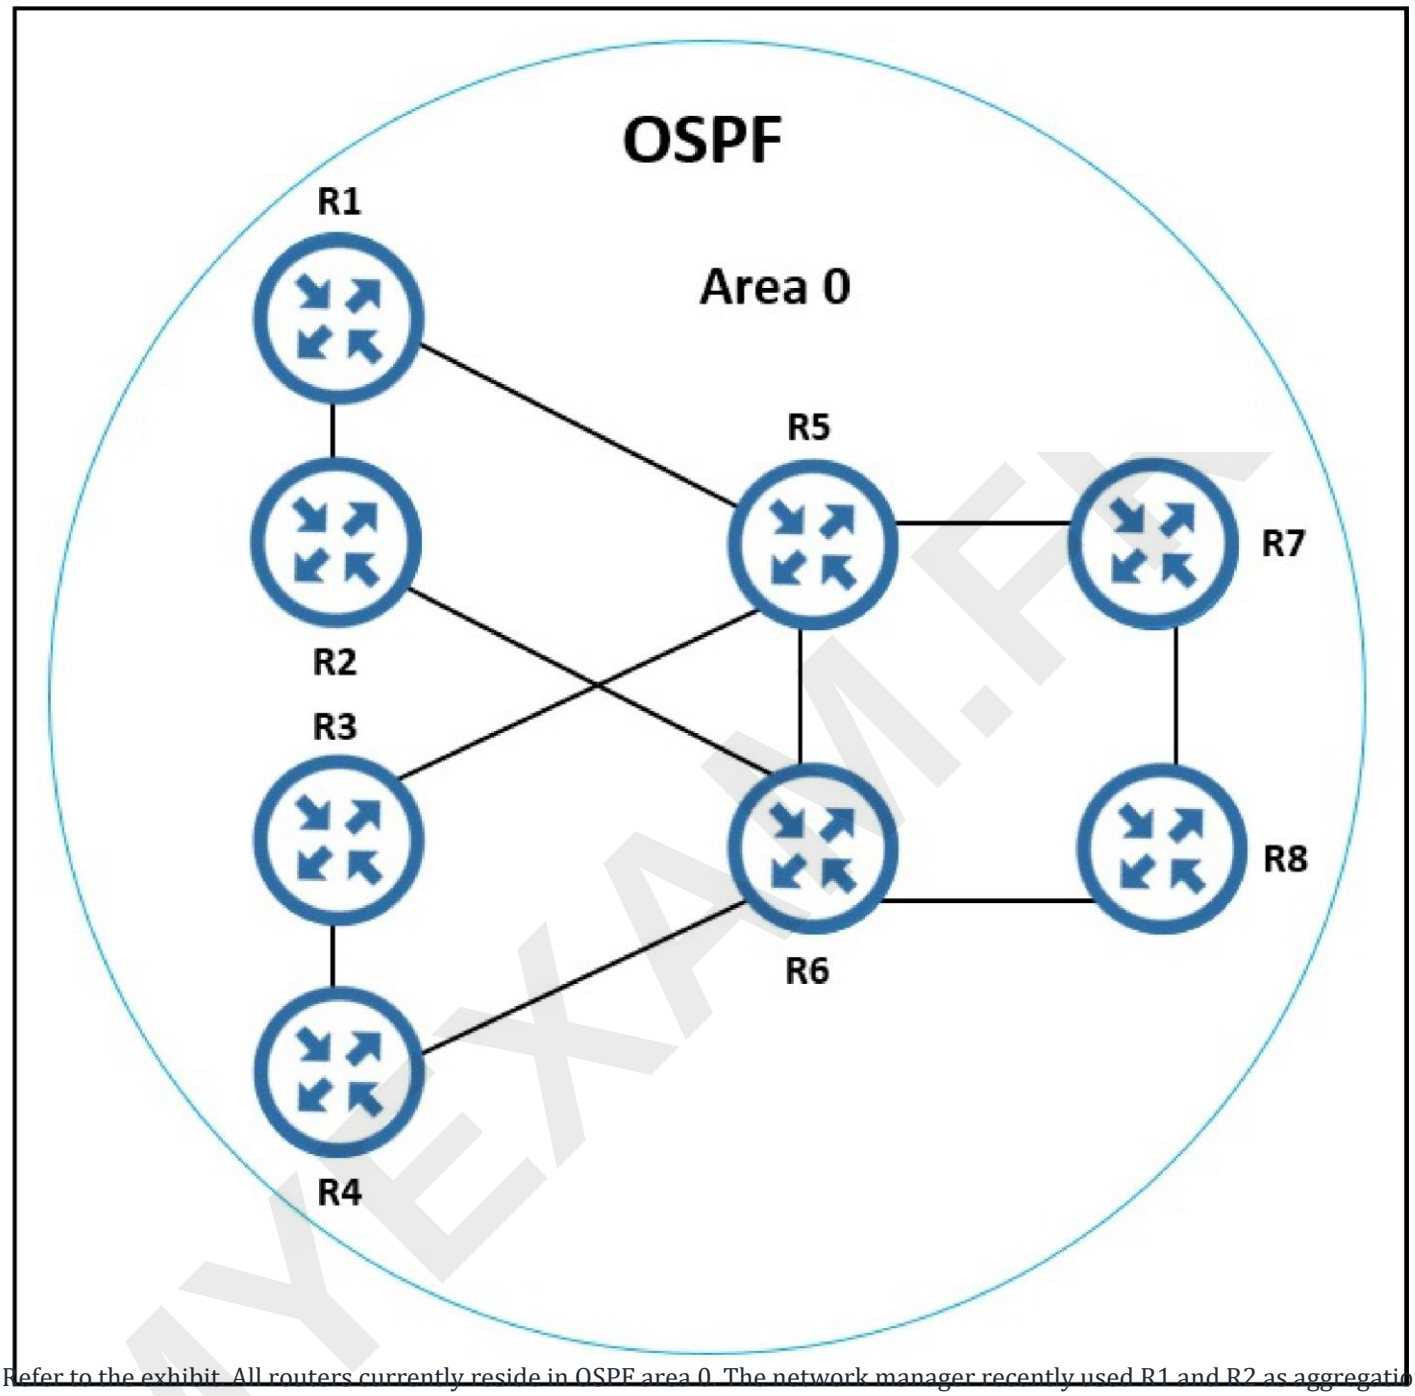

Refer to the exhibit. All routers currently reside in OSPF area 0. The network manager recently used R1 and R2 as aggregation routers for remote branch locations and R3 and R4 as aggregation routers for remote office locations. The network has since been suffering from outages, which are causing frequent SPF runs. To enhance stability and

introduce areas to the OSPF network with the minimal number of ABRs possible, which two solutions should the network manager recommend? (Choose two.)

A.a new OSPF area for R1 and R2 connections, with R1 and R2 as ABRs
B.a new OSPF area for R3 and R4 connections, with R5 and R6 as ABRs
C.a new OSPF area for R3 and R4 connections, with R3 and R4 as ABRs
D.a new OSPF area for R1, R2, R3, and R4 connections, with R1, R2, R3, and R4 as ABRs
E.a new OSPF area for R1 and R2 connections, with R5 and R6 as ABRs

Answer: BE

Explanation:

B.a new OSPF area for R3 and R4 connections, with R5 and R6 as ABRs.

E.a new OSPF area for R1 and R2 connections, with R5 and R6 as ABRs.

Question: 20

An engineer must design a solution to provide backup connectivity between two sites. The engineer plans to use an Internet connection, but company policy requires the connection to be encrypted. Additionally, there are several applications that utilize multicast to deliver video streams between the sites. Which technology should the design include?

A.GRE over IPsec
B.IPsec direct encapsulation
C.GETVPN
D.DMVPN

Answer: A

Explanation:

The correct answer is A, GRE over IPsec. Here's why:

The scenario requires encrypted backup connectivity between two sites over the internet, specifically supporting multicast. IPsec alone (option B) does not inherently support multicast routing. It primarily provides secure point-to-point tunnels. While IPsec can encrypt traffic, including multicast encapsulated within other protocols, it needs a mechanism to effectively route multicast packets. Direct IPsec encapsulation lacks this routing capability for multicast.

DMVPN (option D), though excellent for dynamic VPN setups, typically involves complex hub-and-spoke topologies and isn't primarily designed for simple site-to-site backup connections. It's overkill for this specific scenario and may introduce unnecessary overhead.

GRE (Generic Routing Encapsulation) tunnels, on the other hand, can encapsulate a variety of network layer protocols, including multicast, within IP packets. By combining GRE with IPsec, we can achieve both: GRE provides the multicast routing mechanism, while IPsec provides the required encryption for the traffic traversing the public internet. The GRE packets, including multicast, become the payload of the IPsec tunnel, hence achieving both the required routing and encryption. Thus, GRE over IPsec provides a suitable solution for transmitting multicast over an encrypted internet connection. This approach is widely used and a best practice for similar scenarios.

Authoritative Links:

Cisco Documentation on GRE:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/15-sy/iap-15-sy-book/iap-gre.html
Cisco Documentation on IPsec:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_ipsec/configuration/15-sy/sec-ipsec-15-sy-book.html
Cisco Documentation on DMVPN:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsec/configuration/15-sy/sec-ipsec-15-sy-book/sec-dmvpn.html

Question: 21

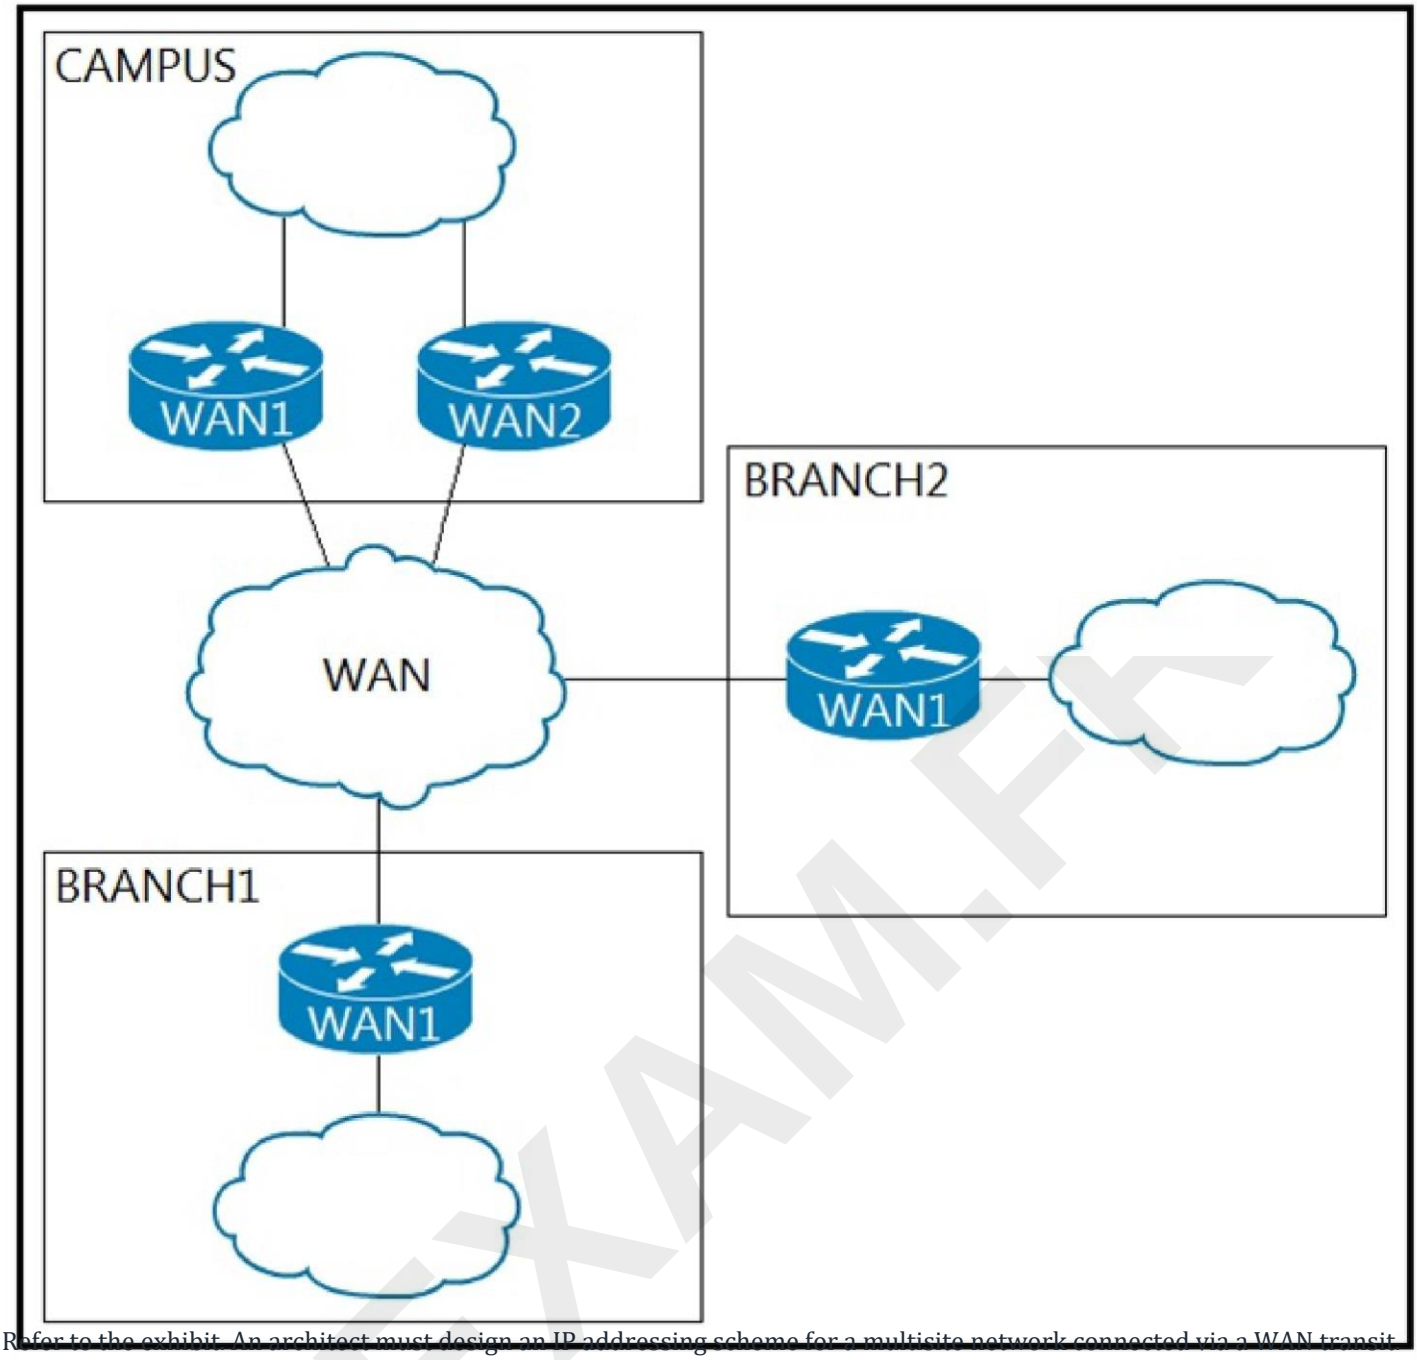

Refer to the exhibit. An architect must design an IP addressing scheme for a multisite network connected via a WAN transit. The campus site must accommodate
12,000 devices, and the branch sites must accommodate 1,000 devices. Which address scheme optimizes network device resources, contains convergence events to the different blocks of the network, and ensures the network's future growth?

A.¢ Campus: 10.0.0.0/18 ¢ Branch1: 10.0.192.0/21 ¢ Branch2: 10.0.200.0/21
B.¢ Campus: 10.0.0.0/16 ¢ Branch1: 10.255.0.0/20 ¢ Branch2: 10.255.16.0/20
C.¢ Campus: 10.0.0.0/10 ¢ Branch1: 10.64.0.0/10 ¢ Branch2: 10.128.0.0/10
D.¢ Campus: 10.0.0.0/20 ¢ Branch1: 10.0.64.0/21 ¢ Branch2: 10.0.128.0/21

Answer: A
Explanation:

¢ Campus: 10.0.0.0/18 ¢ Branch1: 10.0.192.0/21 ¢ Branch2: 10.0.200.0/21.

Question: 22

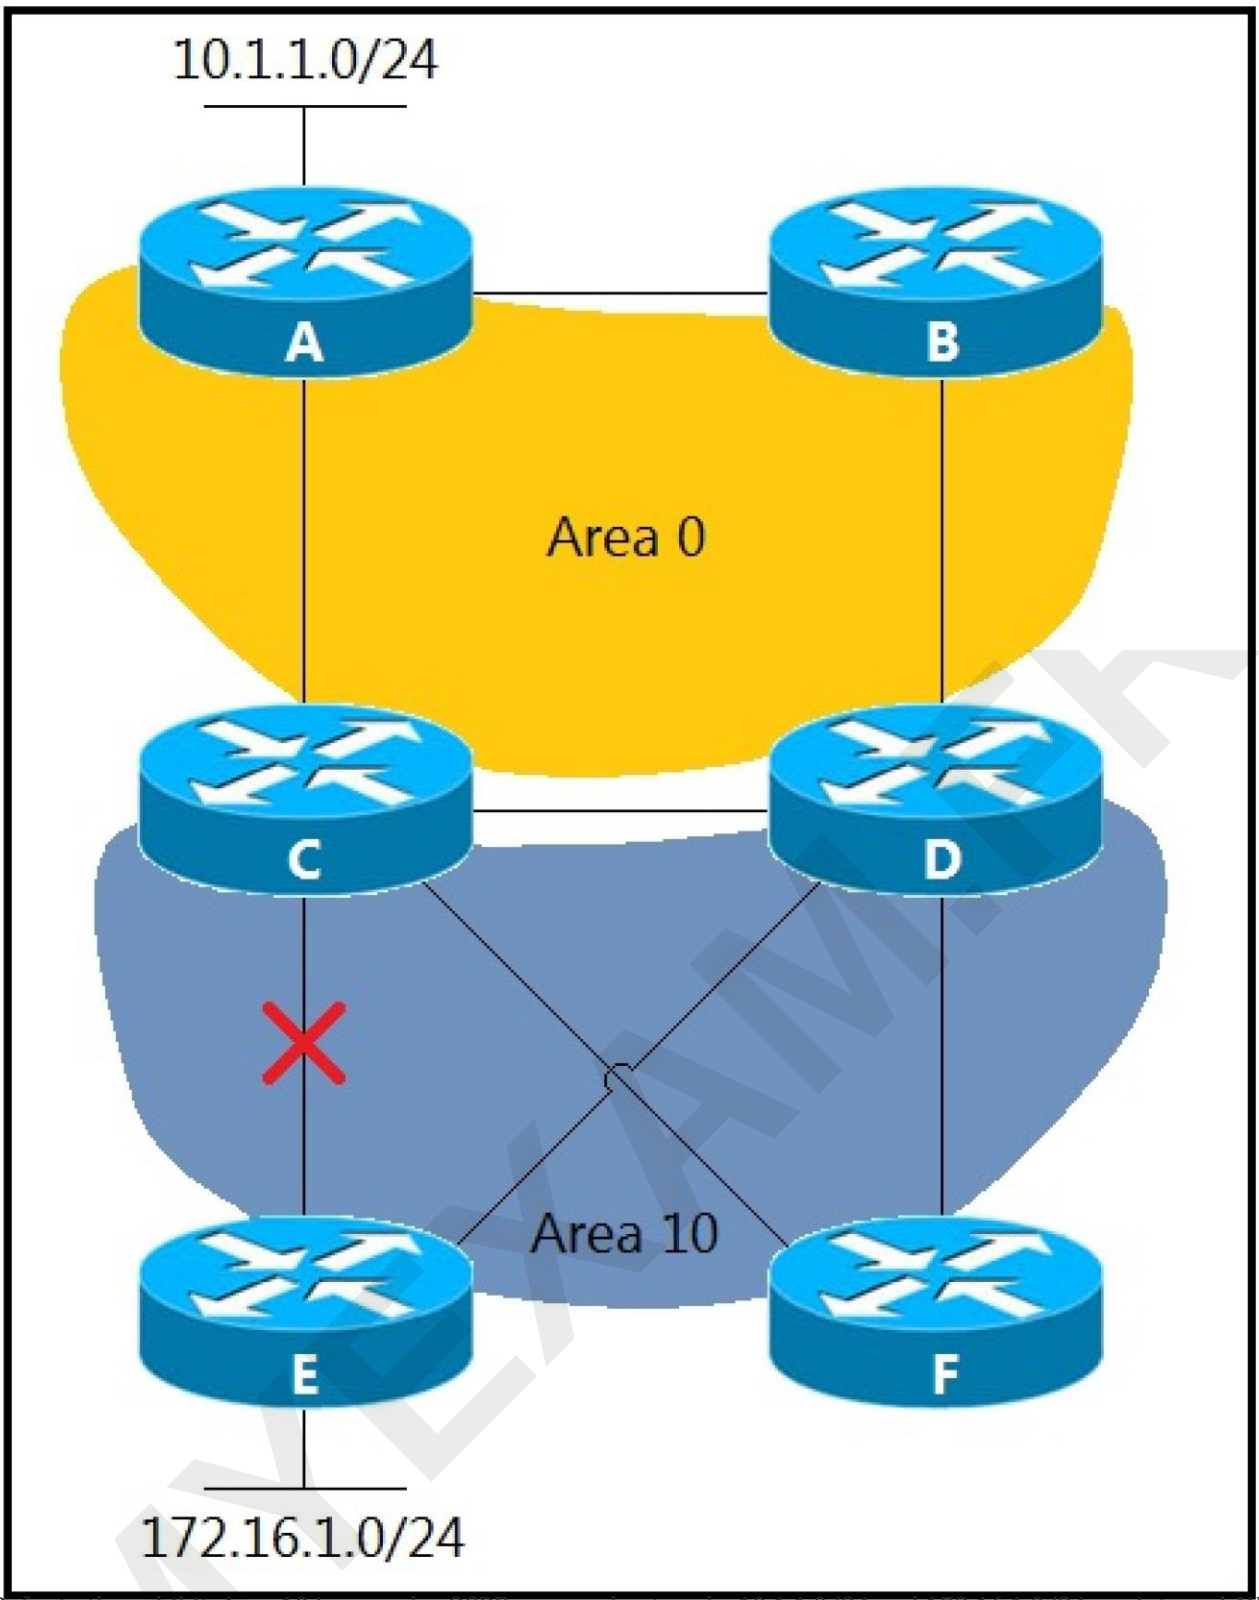

Refer to the exhibit. Area 10 is a regular OSPF area, and networks 10.1.1.0/24 and 172.16.1.0/24 are internal. Which design provides optimal routing between both networks when the link between routers C and E fails?

A. Move the link between routers C and D to area 10.
B. Create an OSPF virtual link between routers E and F.
C. Create a tunnel between routers E and F in area 10.
D. Make area 10 a not-so-stubby area.

Answer: A

Explanation:

Move the link between routers C and D to area 10.

Question: 23

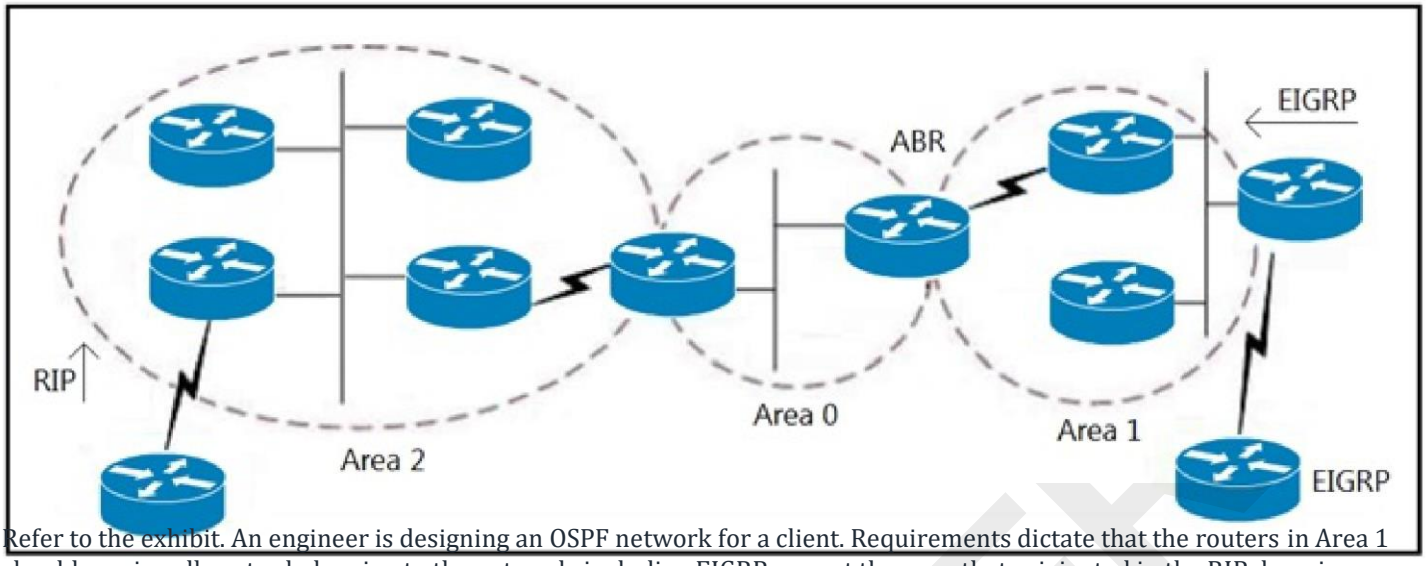

Refer to the exhibit. An engineer is designing an OSPF network for a client. Requirements dictate that the routers in Area 1 should receive all routes belonging to the network, including EIGRP, except the ones that originated in the RIP domain. Which action should the engineer take?

A.Make area 1 a NSSA.
B.Make area 1 a stub.
C.Make area 1 a standard OSPF area.
D.Make the area 1 routers part of area 0.

Answer: A

Explanation:

Correct answer is A:Make area 1 a NSSA.

Question: 24

An engineer is tasked with designing a dual BGP peering solution with a service provider. The design must meet these conditions:
* The routers will not learn any prefix with a subnet mask greater than /24.
* The routers will determine the routes to include in the routing table based on the length of the mask alone. * The routers will make this selection regardless of the service provider configuration.
Which solution should the engineer include in the design?

A.Use a route map and access list to block the desired networks, and apply the route map to BGP neighbors

inbound.

B.Use a route map and prefix list to block the desired networks, and apply the route map to BGP neighbors outbound.

C.Use an IP prefix list to block the desired networks and apply the IP prefix list to BGP neighbors outbound.

D.Use an IP prefix list to block the desired networks and apply the IP prefix list to BGP neighbors inbound.

Answer: D

Explanation:

The correct solution is **D. Use an IP prefix list to block the desired networks and apply the IP prefix list to BGP neighbors inbound.**

Here's why:

The problem requires the engineer to filter incoming BGP routes based solely on the subnet mask length, specifically preventing the learning of any prefix longer than /24. This filtering action needs to happen before the routes are considered for inclusion in the BGP routing table, regardless of the service provider's configuration.

Let's analyze each option:

A: Using an access list is not ideal because access lists typically match based on source/destination IP addresses and ports, not prefix length. While a route-map could use an access list, a prefix list is designed explicitly for matching based on prefix length making the filtering process more efficient. Applying the route map inbound means the filtering will happen before the route is considered for BGP best-path selection. **B:** Applying the filter outbound would not address the requirement, since that applies only to routes that the router is advertising and not those received from peers. Filtering after the routes have already been considered and chosen would mean the routers have already learned and selected best paths.

C: While a prefix list can correctly match prefix length, applying it outbound suffers the same problem as option B.

Option D utilizes the prefix-list to match based on prefix length, which is the exact method required. By applying the prefix list inbound to the BGP neighbor the routes with prefix lengths /25-/32 will never be considered by the router.

Key BGP Concepts:

Prefix List: A list of prefixes with specified mask lengths, used for route filtering. It is designed specifically for matching based on prefix length, which is essential in this scenario.

Inbound Route Filtering: Filtering routes received from a BGP peer before they are considered for inclusion in the BGP RIB (Routing Information Base) or the routing table. This is the right place to ensure no prefixes longer than /24 are learned.

Outbound Route Filtering: Filtering routes that a router intends to advertise to its BGP peers. This is not relevant to the problem requirements.

In summary, applying an IP prefix list inbound is the proper mechanism to control route advertisements based on prefix length before the routes are considered. This achieves the requirement of preventing the router from learning any prefixes with a subnet mask greater than /24.

Authoritative Links for further research:

Cisco Prefix Lists: https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13741-prefix-list.html
BGP Route Filtering: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/15-mt/irg-15-mt-book/irg-route-filter.html

Question: 25

An engineer is designing an EIGRP network for a small branch office site where there is only one Layer 3 router. The engineer wants the router to advertise the local LAN network to remote EIGRP neighbors without sending any unnecessary multicast messages on the local LAN. Which action should the engineer take?

A.Use a static default route for this site instead of EIGRP
B.Advertise the local LAN using the network command and the passive-interface feature
C.Redistribute the local LAN network using the redistribute connected command
D.Advertise the local LAN subnet as a stub network

Answer: B

Explanation:

The correct answer is **B. Advertise the local LAN using the network command and the passive-interface feature.** Here's why:

The goal is to advertise the local LAN subnet using EIGRP while preventing unnecessary multicast traffic on that LAN. EIGRP normally sends multicast hello packets to discover neighbors and maintain adjacencies.
Using the network command under the EIGRP configuration enables EIGRP on the specified interfaces and networks. The passive-interface command, applied to the local LAN interface, prevents EIGRP from sending hello packets on that interface, thus suppressing multicast traffic on the LAN. Importantly, this does not prevent the router from advertising the LAN subnet via other interfaces where EIGRP is enabled. In the described scenario, this is perfect because the branch office router only has one active EIGRP interface facing the WAN, so it would still advertise the local LAN out that interface.

Option A is incorrect because it replaces dynamic routing with a static route, failing to meet the stated requirement of using EIGRP. Option C is unsuitable as redistribute connected is typically used to introduce connected interfaces into EIGRP from other routing protocols which is not the case here since the local LAN is configured under the EIGRP instance itself. Option D, advertising the LAN as a stub network, won't prevent multicast on the local network, and it's specifically meant for more remote sites.

Therefore, the combination of the network command to include the local LAN in EIGRP and the passive-interface command on the LAN interface provides the precise desired behavior – advertising the LAN and eliminating unnecessary multicast traffic.

Authoritative Links:

Cisco Documentation on EIGRP Passive Interface:
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16438-passive.html
Cisco Documentation on EIGRP Configuration:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/15-mt/ire-15-mt-book/ire-eigrp-cfg.html

Question: 26

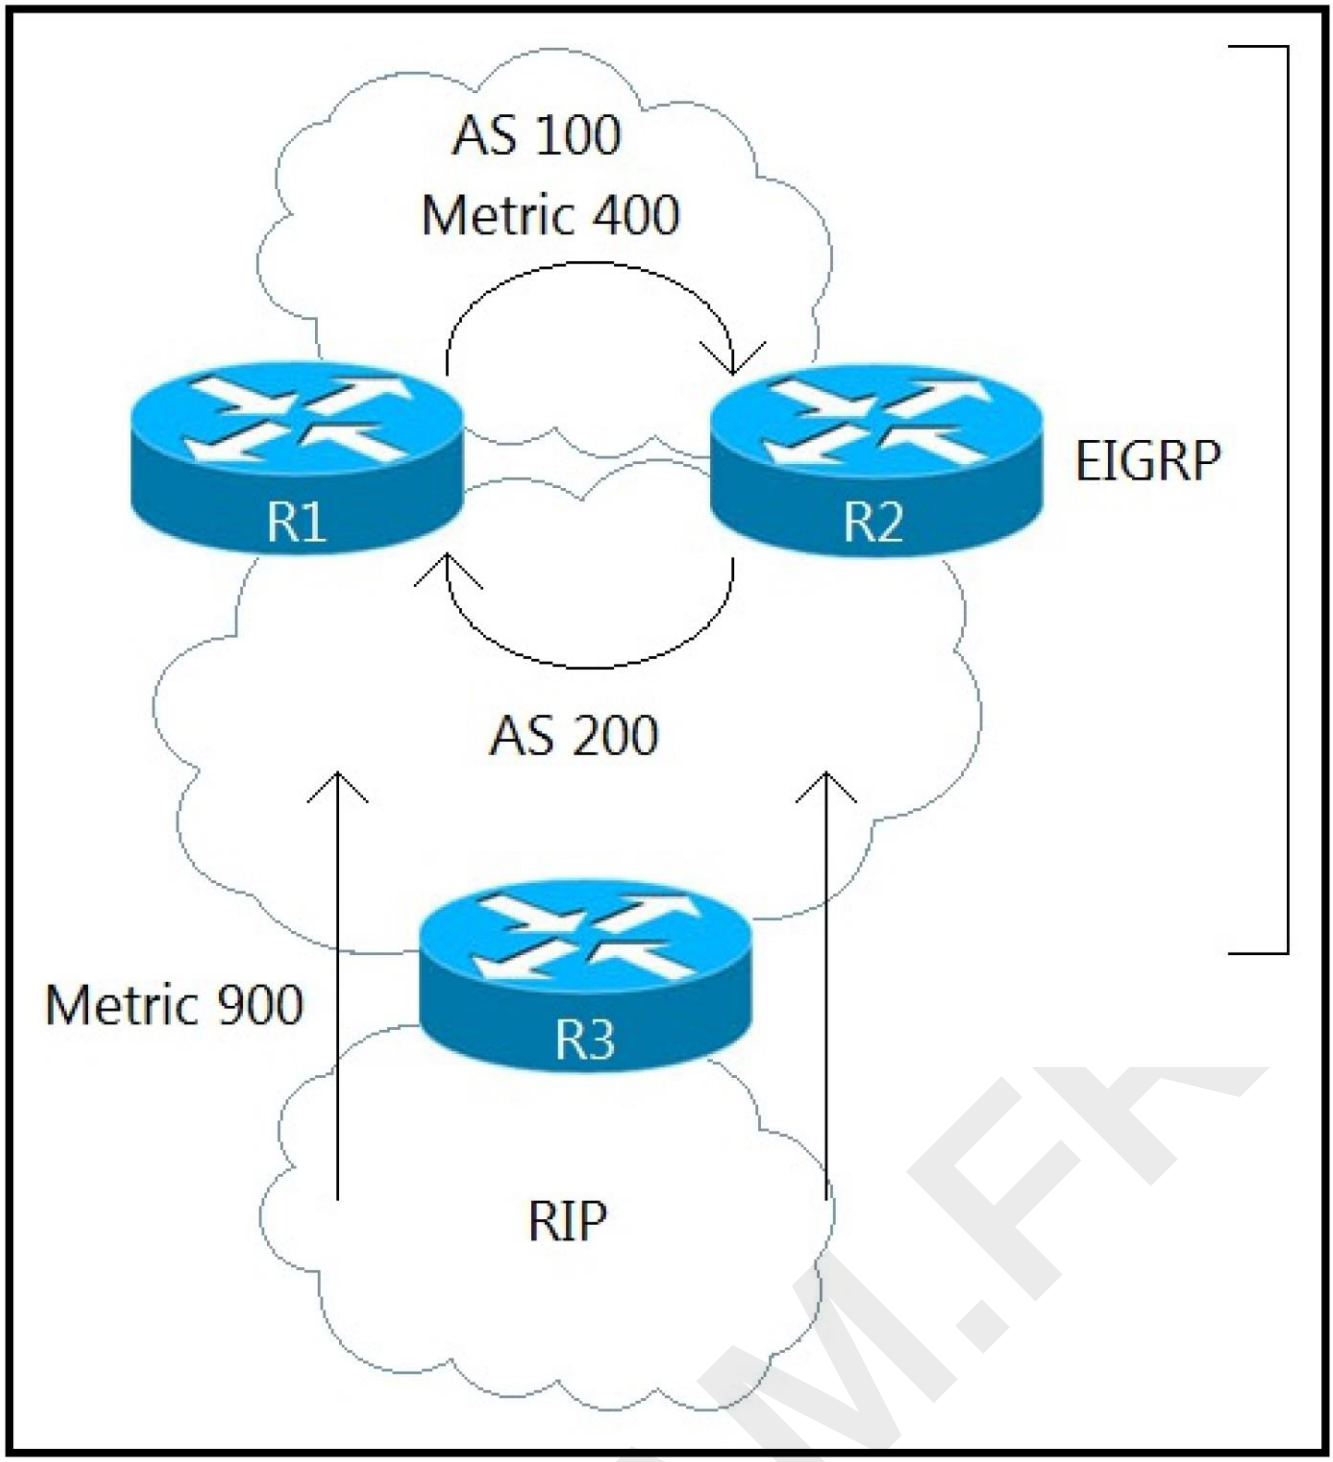

Refer to the exhibit. An architect must design a solution to connect the network behind R3 with the EIGRP network. Which mechanism should be included to avoid routing loops?

A.down bit
B.split-horizon
C.route tags
D.summarization

Answer: C

Explanation:

Correct answer is C:route tags.

Question: 27

An architect is creating a migration strategy for a large organization in which the choice made by the application between IPv6 and IPv4 is based on the DNS request. Which migration strategy does the architect choose?

A.AFT for public web presence
B.host-initiated tunnels
C.dual-stack
D.site-to-site IPv6 over IPv4 tunnels

Answer: C

Explanation:

The correct answer is **C. dual-stack**.

Here's why: The scenario describes a migration strategy where applications dynamically choose between IPv6 and IPv4 based on DNS resolution. This is a core characteristic of **dual-stack** implementations. In dual-stack, both IPv4 and IPv6 are enabled on devices and network infrastructure. When an application performs a DNS lookup, it receives both IPv4 and IPv6 addresses if available. The application then can choose which IP version to use. This method allows a gradual migration, as applications don't need to be all switched over
simultaneously, and the selection is often automated and managed by the applications.

Option A, AFT (Address Family Translation), is a technique for bridging IPv4 and IPv6 networks but doesn't align with the described scenario's DNS-driven choice for applications; it is more of a translation technique at the network layer. Option B, Host-initiated tunnels, involves the creation of tunnels by each endpoint, which is more complex and less scalable than dual-stack and is not related to DNS selection. Option D, Site-to-site IPv6 over IPv4 tunnels, establishes an IPv6 connection through an existing IPv4 network and also doesn't address the applications' DNS-based IP selection.

Dual-stack enables a seamless transition for applications, allowing them to use either IPv4 or IPv6 as necessary and making it the most suitable strategy in this specific scenario.

Authoritative links for further research:

Cisco: Understanding Dual Stack IPv6
IETF:RFC 4213: Basic Transition Mechanisms for IPv6 (While it's an older RFC, it still serves as a foundational document for understanding IPv6 transition mechanisms, including dual-stack).
Wikipedia:IPv6 Transition Mechanisms - Provides an overview of various IPv6 transition technologies, including dual-stack.

Question: 28

An engineer is creating a design to enable IPv6 to run on an existing IPv4 IS-IS network. The IPv4 and IPv6 topologies will match exactly, and the engineer plans to use the same IS-IS router levels for each protocol per interface. Which IS-IS design is required?

A.multi topology without enabling transition feature
B.multi topology with transition feature enabled
C.single topology without enabling transition feature
D.single topology with transition feature enabled

Answer: C

Explanation:

Here's a detailed justification for why option C, "single topology without enabling transition feature," is the

correct IS-IS design for the scenario:

The question specifies an environment where IPv4 and IPv6 topologies will be identical, using the same IS-IS router levels per interface. This key requirement eliminates the need for multiple IS-IS topologies. IS-IS, by default, operates as a single topology protocol. The 'multi topology' approach (options A and B) is intended for scenarios where different network topologies exist for distinct protocols or services, like for multicast or when there's no consistent mapping between IPv4 and IPv6 paths. Given that both the IPv4 and IPv6 topologies are mirrored here, there's no topological differentiation necessitating a multi-topology setup.

Transition features, generally employed when migrating from IPv4 to IPv6 to manage dual stacks and differing topologies, are also not necessary in this simple, mirrored architecture. The "transition feature" refers to capabilities like using IPv4 transport for IPv6 packets during a migration. Since both protocols will run alongside each other in the same topology, enabling a transition feature is inappropriate. Therefore, a single topology, the standard IS-IS behavior, is sufficient, making Option C the optimal and most efficient solution. Using a single topology keeps the configuration simpler and reduces overhead since the same set of link-state database (LSDB) information is used for both IPv4 and IPv6 routing.

Further reading:

Cisco documentation on IS-IS and IPv6:

Understanding IS-IS for IPv6
RFC 5308:

Routing IPv6 with IS-IS

Question: 29

An engineer must connect a new remote site to an existing OSPF network. The new site consists of two low-end routers, one for WAN, and one for LAN. There is no demand for traffic to pass through this area. Which area type does the engineer choose to provide minimal router resource utilization, while still allowing for full connectivity to the rest of the network?

A. not so stubby
B. totally not so stubby
C. totally stubby area
D. stubby area

Answer: C

Explanation:

The correct answer is **C. totally stubby area**. Here's the justification:

OSPF areas are designed to manage the complexity of large networks. Different area types offer varying levels of route information, which directly impacts router resource usage. A totally stubby area minimizes resource utilization by blocking all external routes and summary routes, only allowing intra-area routes and a default route. This behavior is crucial in the scenario, where the remote site is not intended to serve as a transit point for other traffic.

Let's examine the options:

A. Not So Stubby Area (NSSA): While NSSAs are designed to limit route information, they still allow external routes to be advertised, making it less efficient compared to a totally stubby area for the specific situation.

B. Totally Not So Stubby Area (TNSSA): This area type blocks external and summary routes, similar to a totally stubby area. However, TNSSA allows the import of Type-7 LSAs that describe external routes within

the NSSA, which is unnecessary and adds processing load in this specific scenario.

D. Stubby Area: Stubby areas block external routes but allow summary routes. This still introduces more information than a totally stubby area, requiring slightly more router resources.

A totally stubby area only needs to maintain intra-area routes and a default route, thus drastically reducing the size of its routing table, memory consumption, and the overall CPU load associated with running OSPF. This makes it ideal for resource-constrained environments like the specified low-end routers in a remote site that do not carry transit traffic. The routers will be able to reach the rest of the network via the default route. In addition, the question states that no traffic will be routed through the remote site, thereby a totally stubby area is the best option since no external or summary routes are needed.

In summary, a totally stubby area best aligns with the requirements by providing the lowest router resource utilization while still ensuring full network connectivity via the default route, precisely addressing the needs of the new remote site.

Further Research:

Cisco Official Documentation on OSPF Areas: https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html
RFC 2328: OSPF Version 2: https://datatracker.ietf.org/doc/html/rfc2328

Question: 30

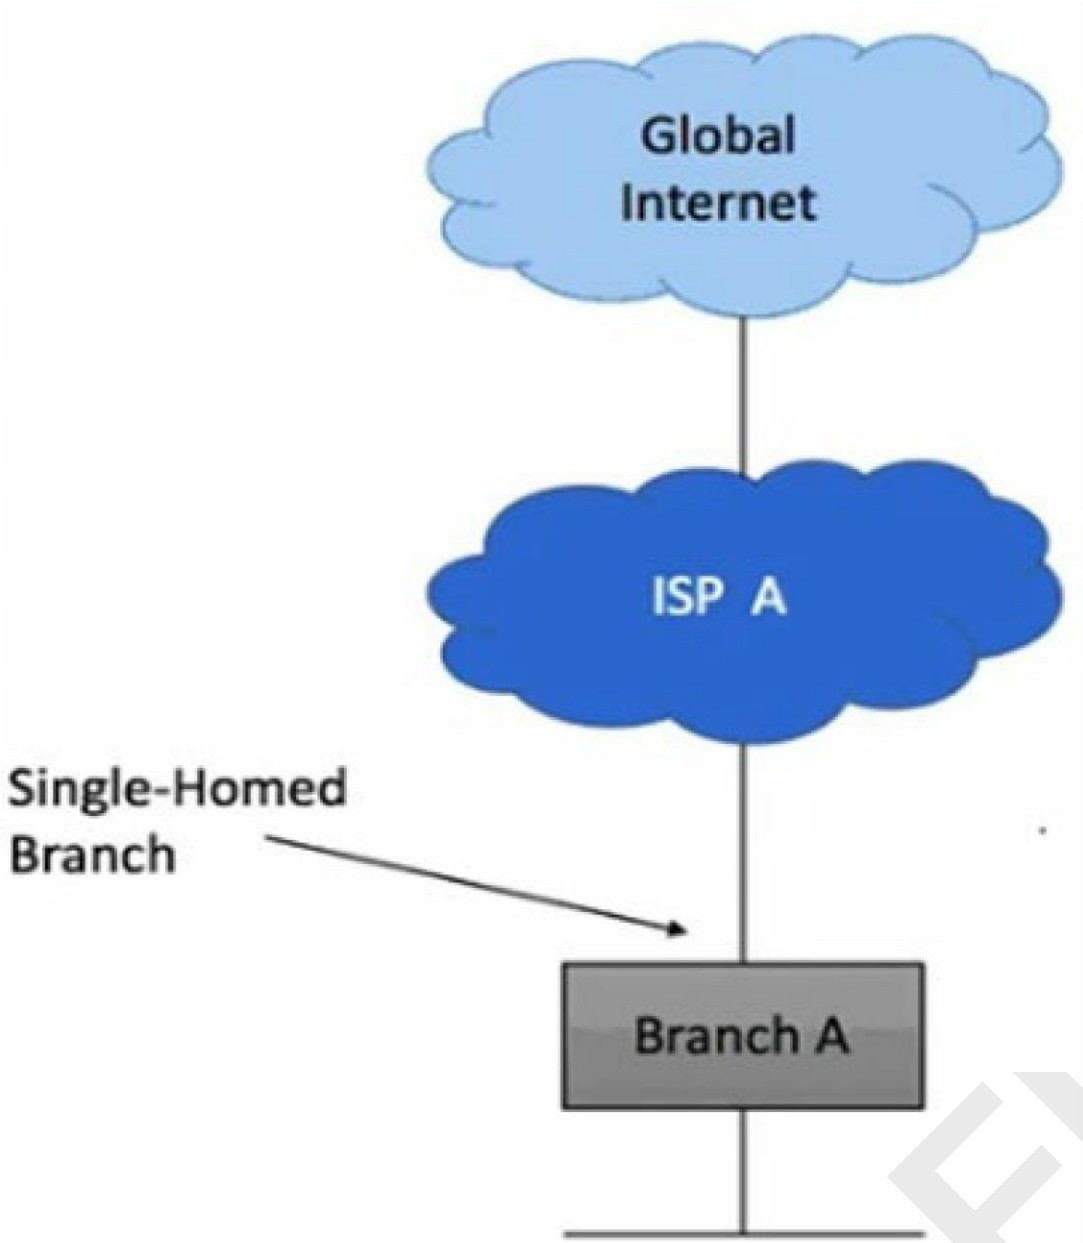

Refer to the exhibit. An architect is designing a BGP solution to connect a remote branch to a service provider. There are several prefixes within the branch that the company does not want to be advertised to the Internet. Which solution should the architect use to accomplish this?

A.Attach the No-Export community with the prefixes to exclude. B.Use the BGP No-Advertise community for the prefixes to exclude. C.Set the BGP Internet community for all prefixes.

D.Implement the NOPEER community.

Answer: A

Explanation:

Attach the No-Export community with the prefixes to exclude.

Question: 31

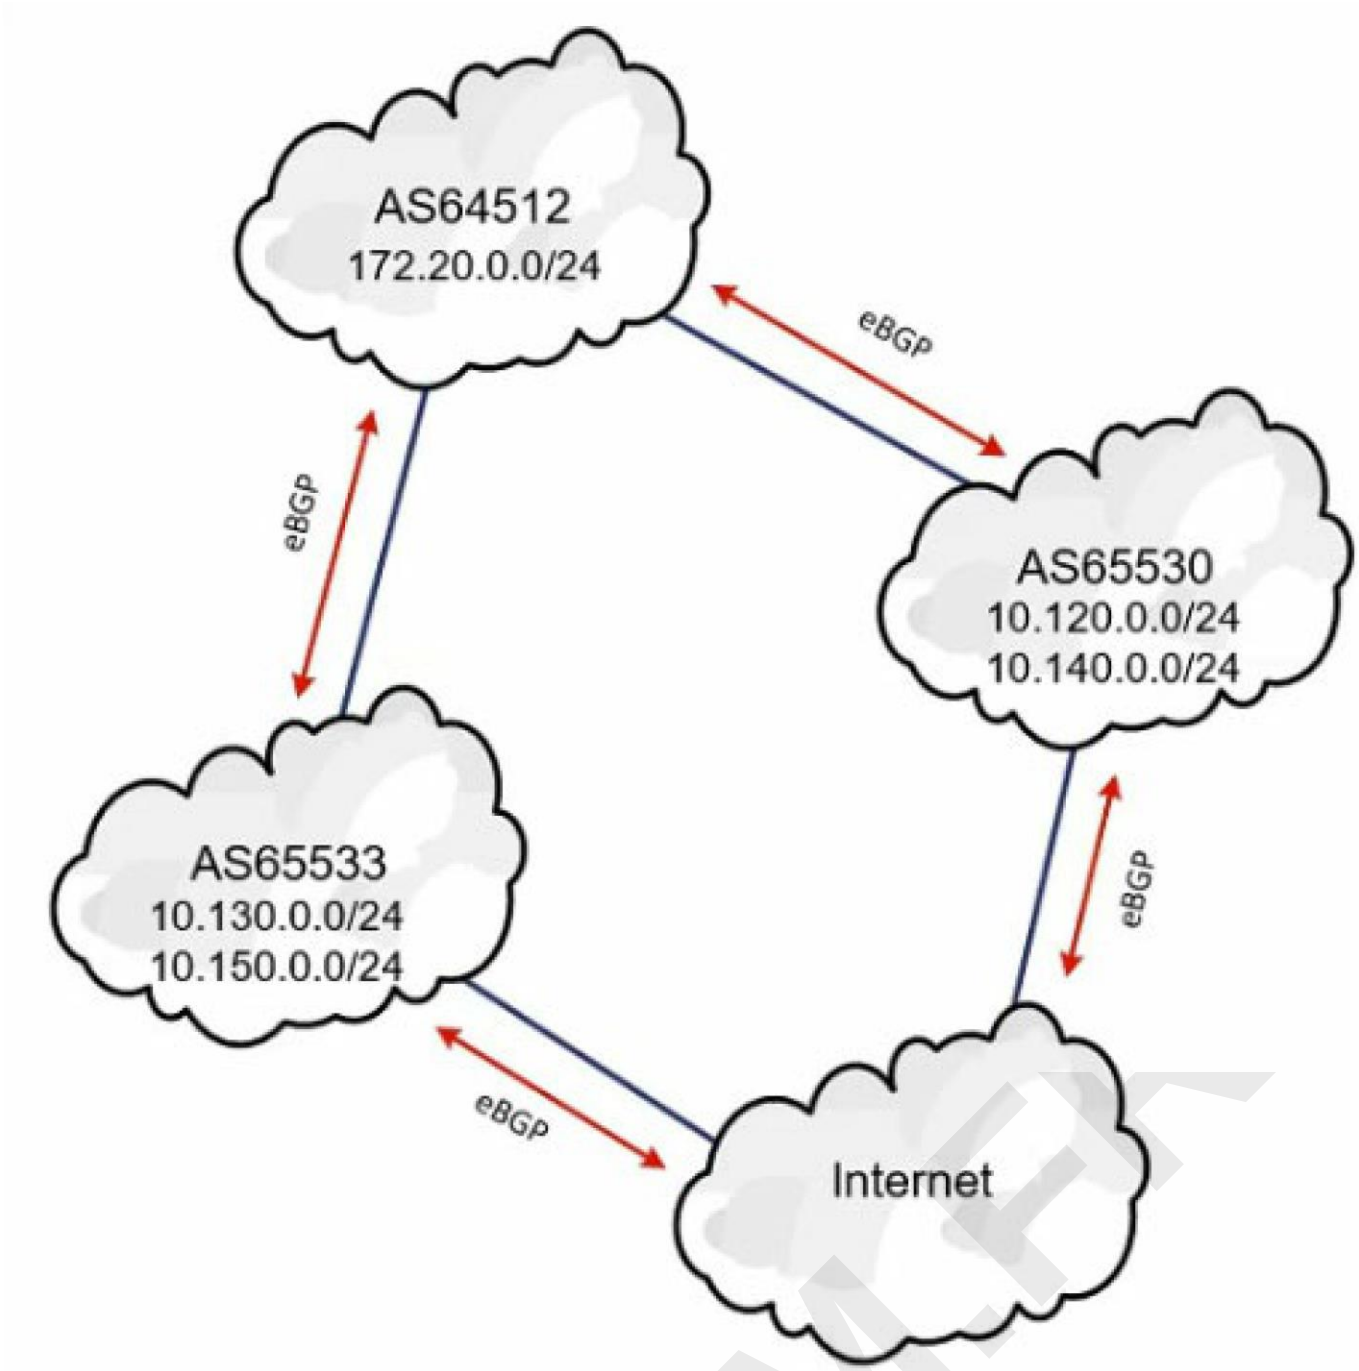

Refer to the exhibit. AS65533 and AS65530 are announcing a partial Internet routing table as well as their IP subnets. An architect must create a design that ensures AS64512 does not become a transit AS. Which filtering solution must the architect choose?

A.no-advertise
B.next-hop
C.no-export
D.maximum-prefix

Answer: C

Explanation:

Correct answer is C:no-export.

Question: 32

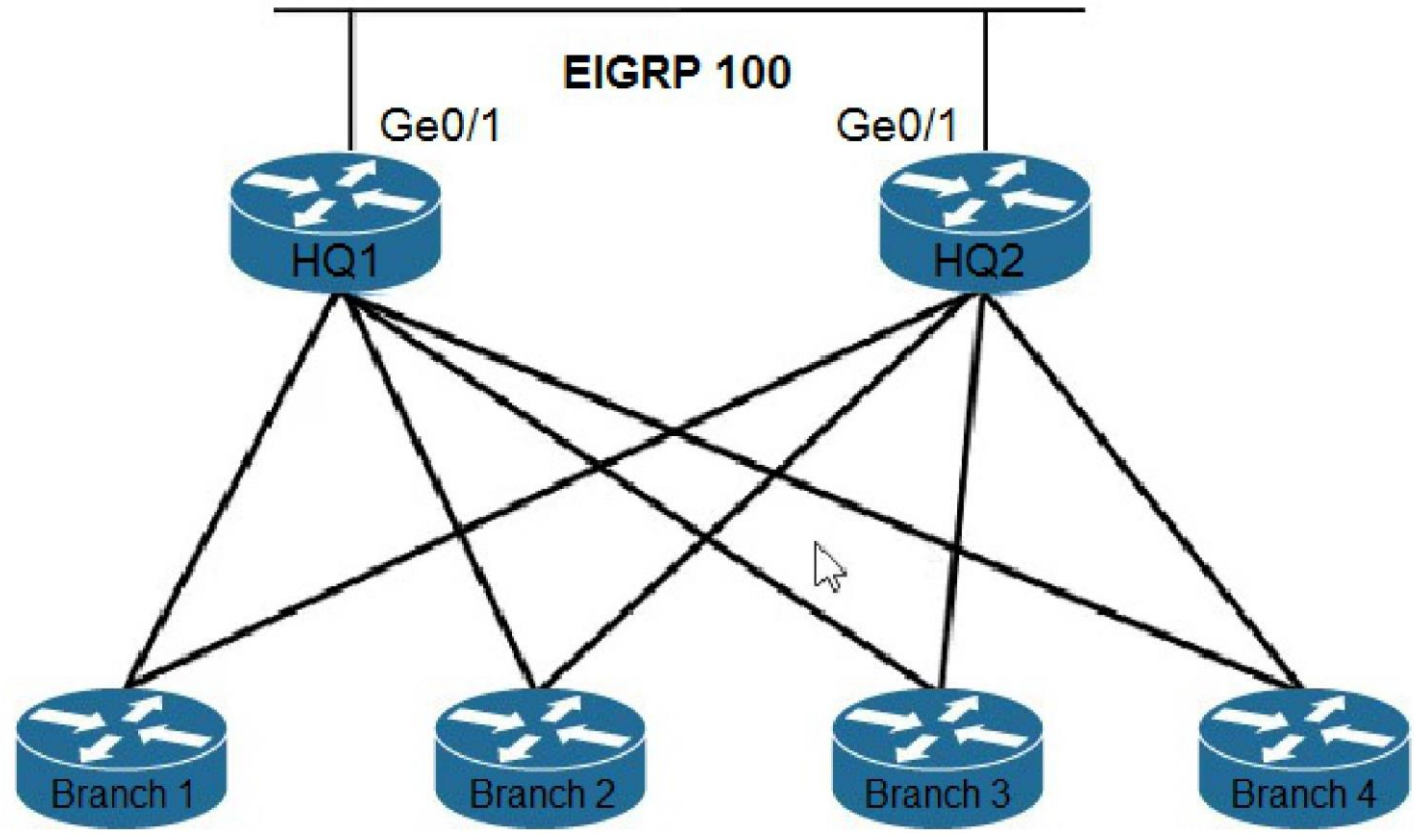

Refer to the exhibit. An architect must create a stable and scalable EIGRP solution for a customer. The design must:
➪ conserve bandwidth, memory, and CPU processing
➪ prevent suboptimal routing
avoid any unnecessary queries
Which two solutions must the architect select? (Choose two.)

A. route summarization
B. prefix lists
C. distribute lists
D. stub routing
E. static redistribution

Answer: AD

Explanation:

A. route summarization.

D. stub routing.

Question: 33

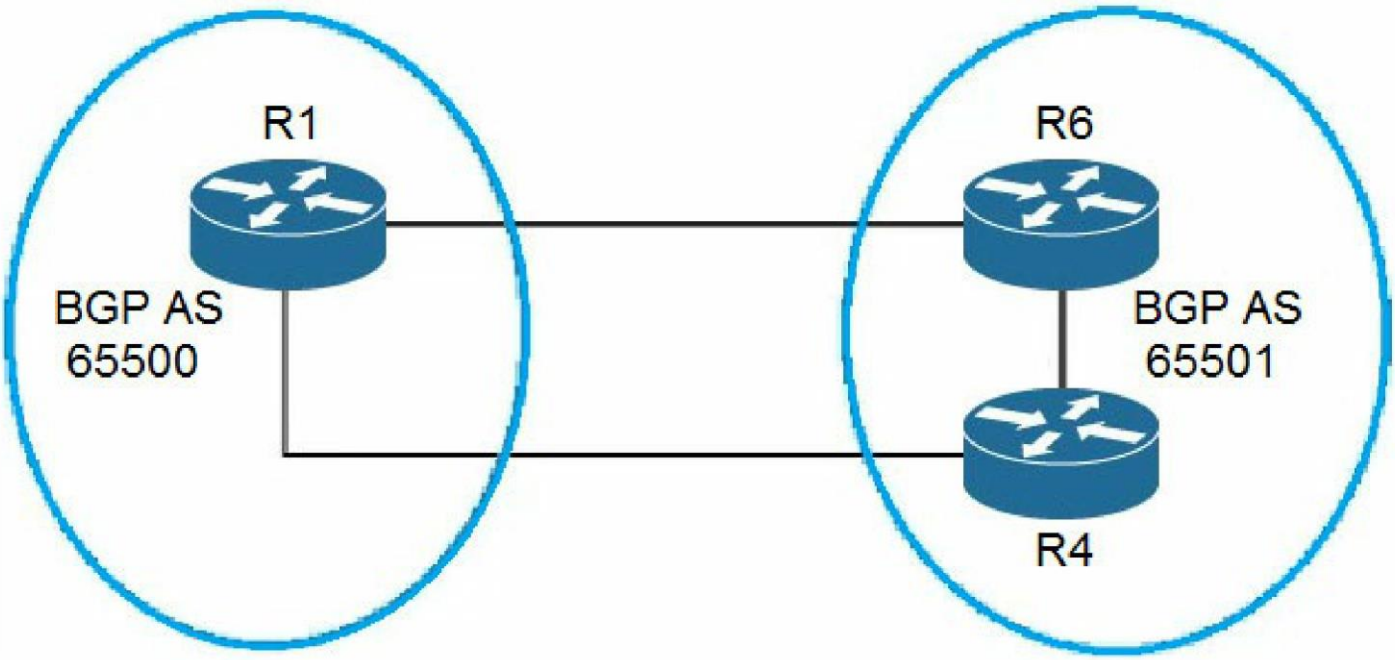

Refer to the exhibit. An architect must design a solution to connect the two ASs. To optimize bandwidth, the design will implement load sharing between router R6 and router R1. Which solution should the design include?

A.Use update-source to specify the Loopback interface.

B.Use next-hop-self attributes only for routes that are learned from eBGP peers.

C.Configure the eBGP TTL to support eBGP multihop.

D.Use maximum-paths to install multiple paths in the routing table.

Answer: D

Explanation:

D. Use maximum-paths to install multiple paths in the routing table: This configuration allows multiple paths to be installed in the routing table for the same destination prefix, enabling load sharing between R1 and R6.

By using the maximum-paths command, traffic can be load balanced between the two routers, optimizing bandwidth utilization.

Question: 34

A customer's environment includes hosts that support IPv6-only. Several of these hosts must communicate with a public web server that has only IPv4 domain name resolution. Which solution should the customer use in this environment?

A.utilize NAT64 to translate the addresses

B.implement NAT44 at the edge of the customer network

C.use 6to4 and a tunnel to translate the addresses

D.implement 6PE to resolve hostname resolution

Answer: A

Explanation:

The correct solution is **A. utilize NAT64 to translate the addresses.** Here's why:

The core challenge is that IPv6-only hosts need to reach an IPv4-only web server. NAT64 is specifically designed to translate IPv6 addresses to IPv4 addresses, enabling communication between these disparate

address families. It operates by taking an IPv6 packet from an IPv6-only host and translating the source IPv6 address into an IPv4 address, allowing it to traverse the IPv4 internet and reach the destination web server. Conversely, responses from the web server, which are IPv4 packets, are translated back to IPv6 for the IPv6-only host.

Option B, NAT44, translates IPv4 addresses to other IPv4 addresses. It's irrelevant in this scenario because the issue is IPv6 to IPv4 communication. Option C, 6to4, uses a type of tunneling and is generally considered an obsolete and insecure solution, it mainly helps with IPv6 addressing across IPv4 networks, but doesn't inherently offer a translation service to directly reach an IPv4-only destination. Option D, 6PE (IPv6 Provider Edge), focuses on providing IPv6 connectivity through a service provider's MPLS infrastructure, which doesn't inherently resolve the IPv4-only destination issue.

NAT64 is the industry-standard mechanism recommended for IPv6 to IPv4 communication specifically for this type of interoperability scenario. It ensures that IPv6-only devices can seamlessly access IPv4 resources without the need for double stacking on the endpoints or cumbersome tunneling solutions. It is crucial for facilitating the smooth transition from IPv4 to IPv6 networks.

Authoritative Links:

RFC 6145:https://www.rfc-editor.org/rfc/rfc6145 (Describes IP/ICMP Translation Algorithm)
RFC 6146:https://www.rfc-editor.org/rfc/rfc6146 (Describes Stateless IP/ICMP Translation)
Cisco Documentation on NAT64:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-mt/ipv6-15-mt-book/ip6-nat-npt.html

Question: 35

A company is planning to open two new branches and allocate the 2a01:c30:16:7009::3800/118 IPv6 network for the region. Each branch should have the capacity to accommodate a maximum of 200 hosts. Which two networks should the company use? (Choose two.)

A.2a01:0c30:0016:7009::3a00/120
B.2a01:0c30:0016:7009::3b00/121
C.2a01:0c30:0016:7009::3a80/121
D.2a01:0c30:0016:7009::3c00/120
E.2a01:0c30:0016:7009::3b00/120

Answer: AE

Explanation:

Here's the justification for the correct answer (A and E) to the IPv6 subnetting question:

The initial network is 2a01:c30:16:7009::3800/118. To accommodate 200 hosts per branch, we need to determine the required prefix length. IPv6 uses prefix lengths rather than traditional subnet masks. A /120 prefix provides 2^(128-120) = 2^8 = 256 addresses, which is sufficient for the 200 hosts and some overhead for broadcast (not used in IPv6) and future expansion. Options B and C have /121 prefix lengths, leading to only 2^7=128 addresses, which is insufficient. Option D also has a /120 prefix but its subnet assignment conflicts with A.

Options A and E, both using /120 prefixes, correctly divide the original /118 network into two unique subnets. Option A, 2a01:0c30:0016:7009::3a00/120, has the lowest address range within the allocated /118. Option E, 2a01:0c30:0016:7009::3c00/120, follows immediately afterward in the address space allowing for no conflict with Option A. To calculate the actual address range, you should look at the /120 after the '::'.

In IPv6, it is common practice to increment the hexadecimal address by 1 to allocate the next network. This is

a more common and easily managed system than incrementing the decimal part, in IPv4 addressing. Options A and E thus provide non-overlapping, appropriately sized subnets from the original allocated IPv6 network, making them ideal for the two new branches.

Authoritative Links:

1. **RFC 4291 - IP Version 6 Addressing Architecture:**https://datatracker.ietf.org/doc/html/rfc4291 (Section 2.5 provides information about prefix lengths and subnetting)
2. **Cisco Documentation on IPv6 Addressing:** https://www.cisco.com/c/en/us/support/docs/ip/ipv6/113328-ipv6-addr-plan.html (Provides an overview of IPv6 addressing and subnetting considerations)
3. **Cloudflare: Understanding IPv6 Subnetting:**https://www.cloudflare.com/learning/network-layer/what-is-ipv6-subnetting/ (Provides an explanation on IPv6 Subnetting)

Question: 36

Refer to the exhibit.

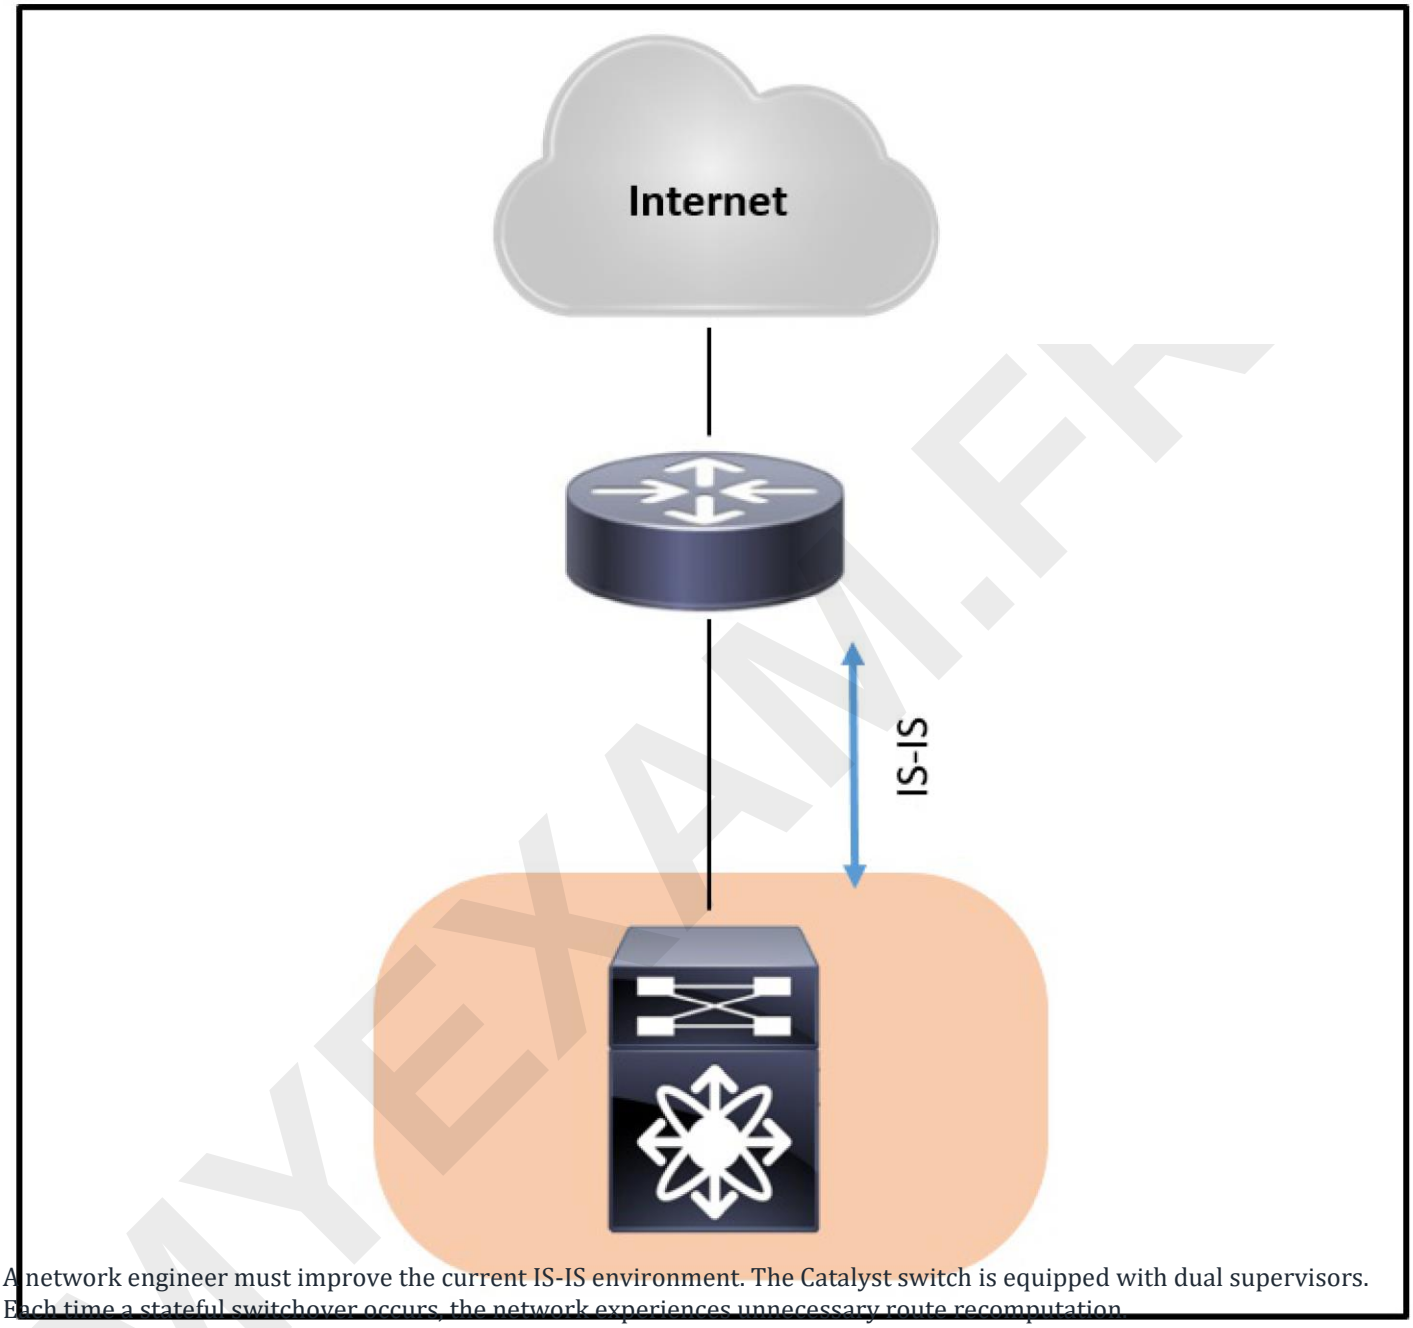

A network engineer must improve the current IS-IS environment. The Catalyst switch is equipped with dual supervisors. Each time a stateful switchover occurs, the network experiences unnecessary route recomputation.

Which solution addresses this issue if the upstream router does not understand graceful restart messaging?

A.Enable IS-IS remote LFA FRR on both devices.
B.Enable NSR on the switch.
C.Enable NSF on the switch.
D.Configure ISIS aggressive timers on both devices.

Answer: B

Explanation:

Enable NSR on the switch.

Question: 37

Refer to the exhibit.

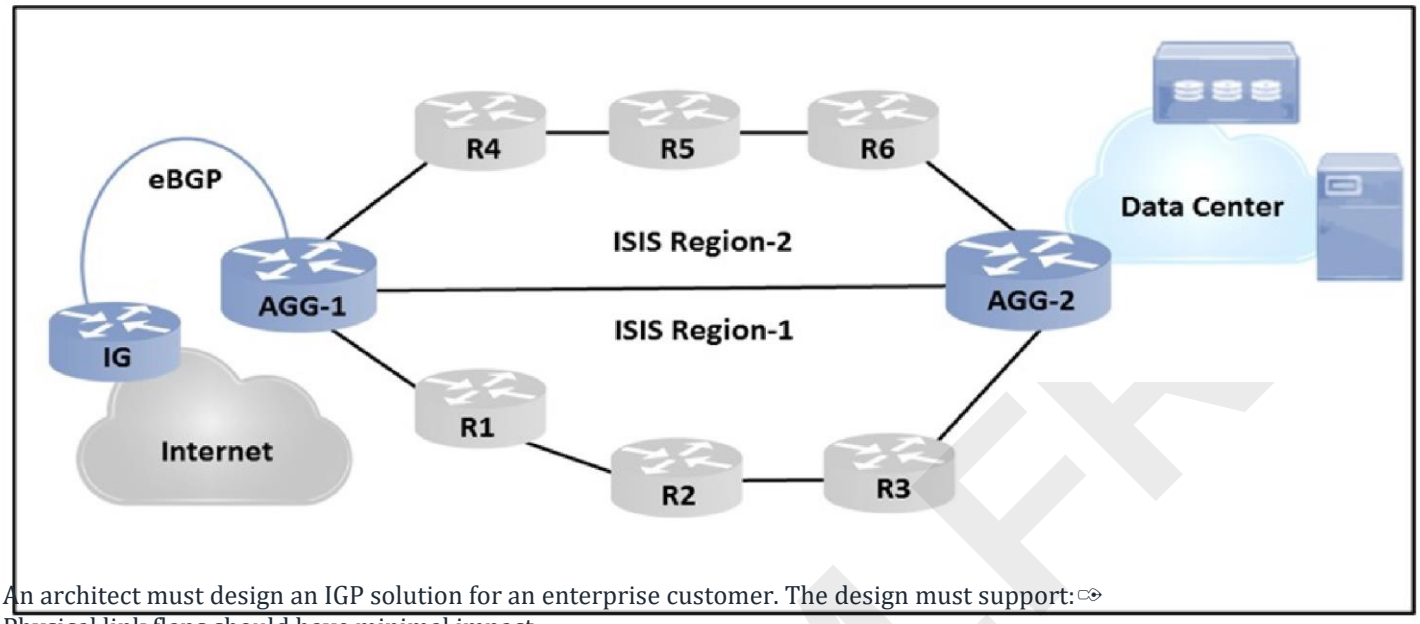

An architect must design an IGP solution for an enterprise customer. The design must support:☞
Physical link flaps should have minimal impact.
☞ Access routers should converge quickly after a link failure.
Which two ISIS solutions should the architect include in the design? (Choose two.)

A.Use BGP to IS-IS redistribution to advertise all Internet routes in the Level 1 area.
B.Advertise the IS-IS interface and loopback IP address toward the Internet and data center.
C.Reduce SPF and PRC intervals to improve convergence time.
D.Configure all access and aggregate routers to establish Level 1 / Level 2 adjacencies across the network. E.Configure access routers to establish a Level 1 adjacency and aggregate routers to establish a Level 1 / Level 2 adjacency.

Answer: CE

Explanation:

C.Reduce SPF and PRC intervals to improve convergence time.

E.Configure access routers to establish a Level 1 adjacency and aggregate routers to establish a Level 1 / Level 2 adjacency.

Question: 38

Refer to the exhibit.

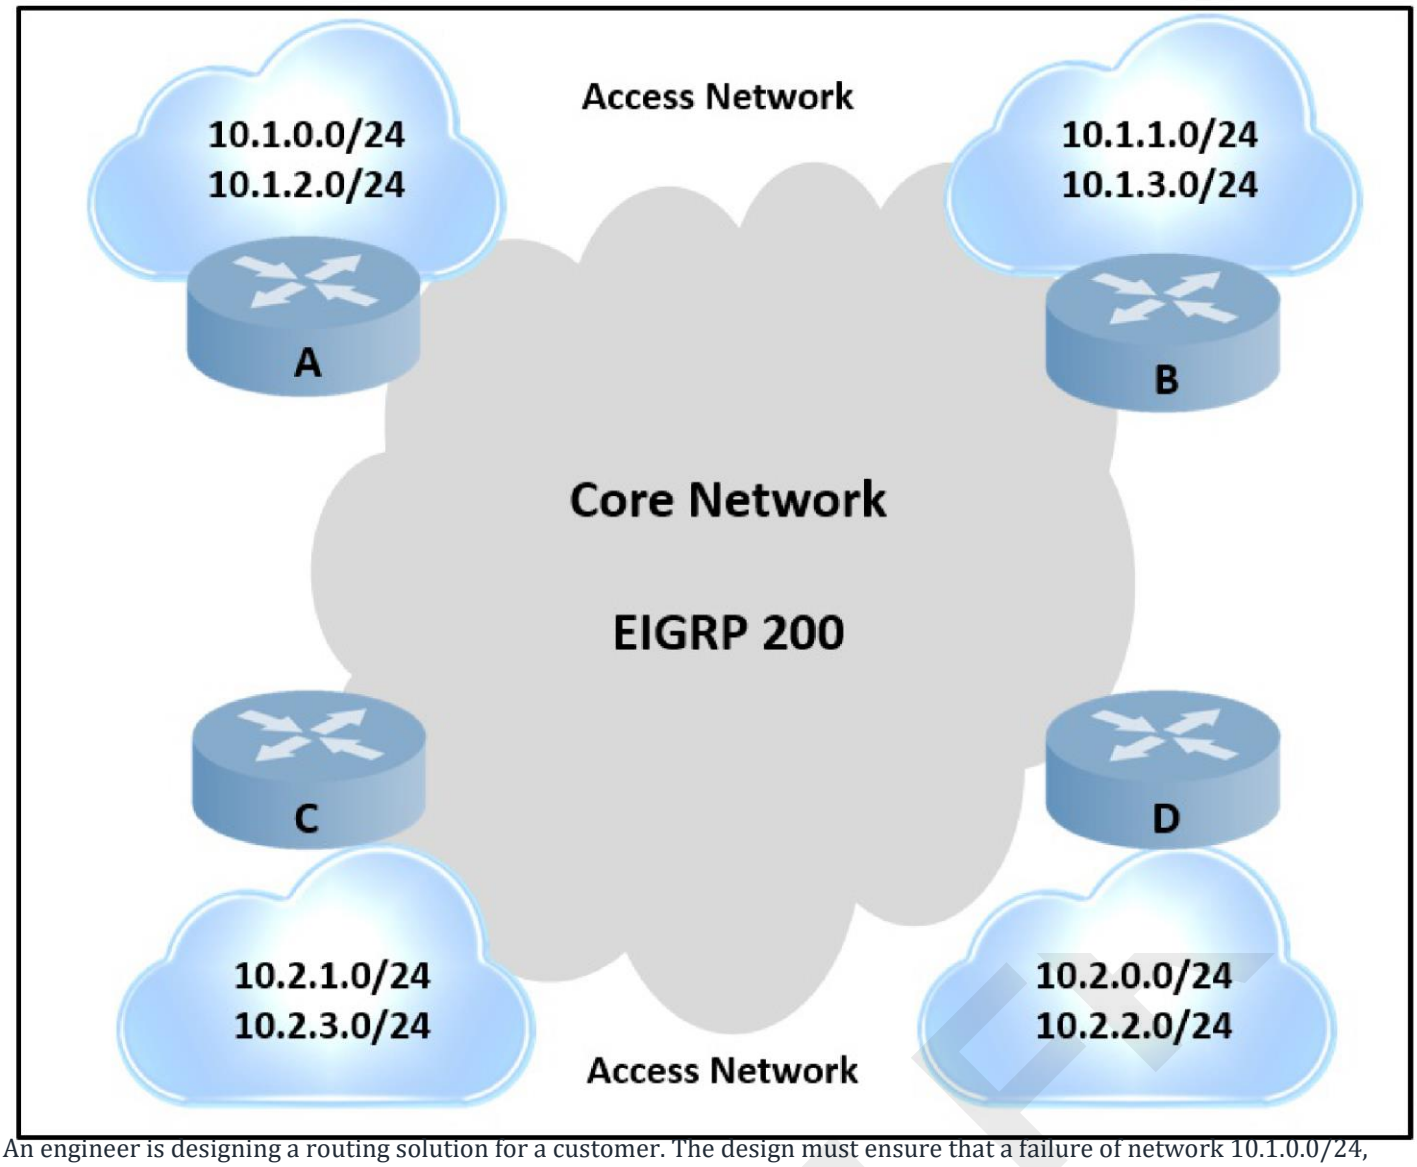

An engineer is designing a routing solution for a customer. The design must ensure that a failure of network 10.1.0.0/24, 10.1.2.0/24, 10.2.1.0/24, or 10.2.3.0/24 does not impact the core. It also requires fast convergence time during any link failover in the core or access networks.
Which solution must the engineer select?

A. Add aggregation layer between core and access networks.
B. Enable graceful restart on routers A and C.
C. Enable FRR for the connected networks of routers A and C.
D. Enable summarization on routers A and C.

Answer: C

Explanation:

Enable FRR for the connected networks of routers A and C.

Question: 39

Refer to the exhibit.

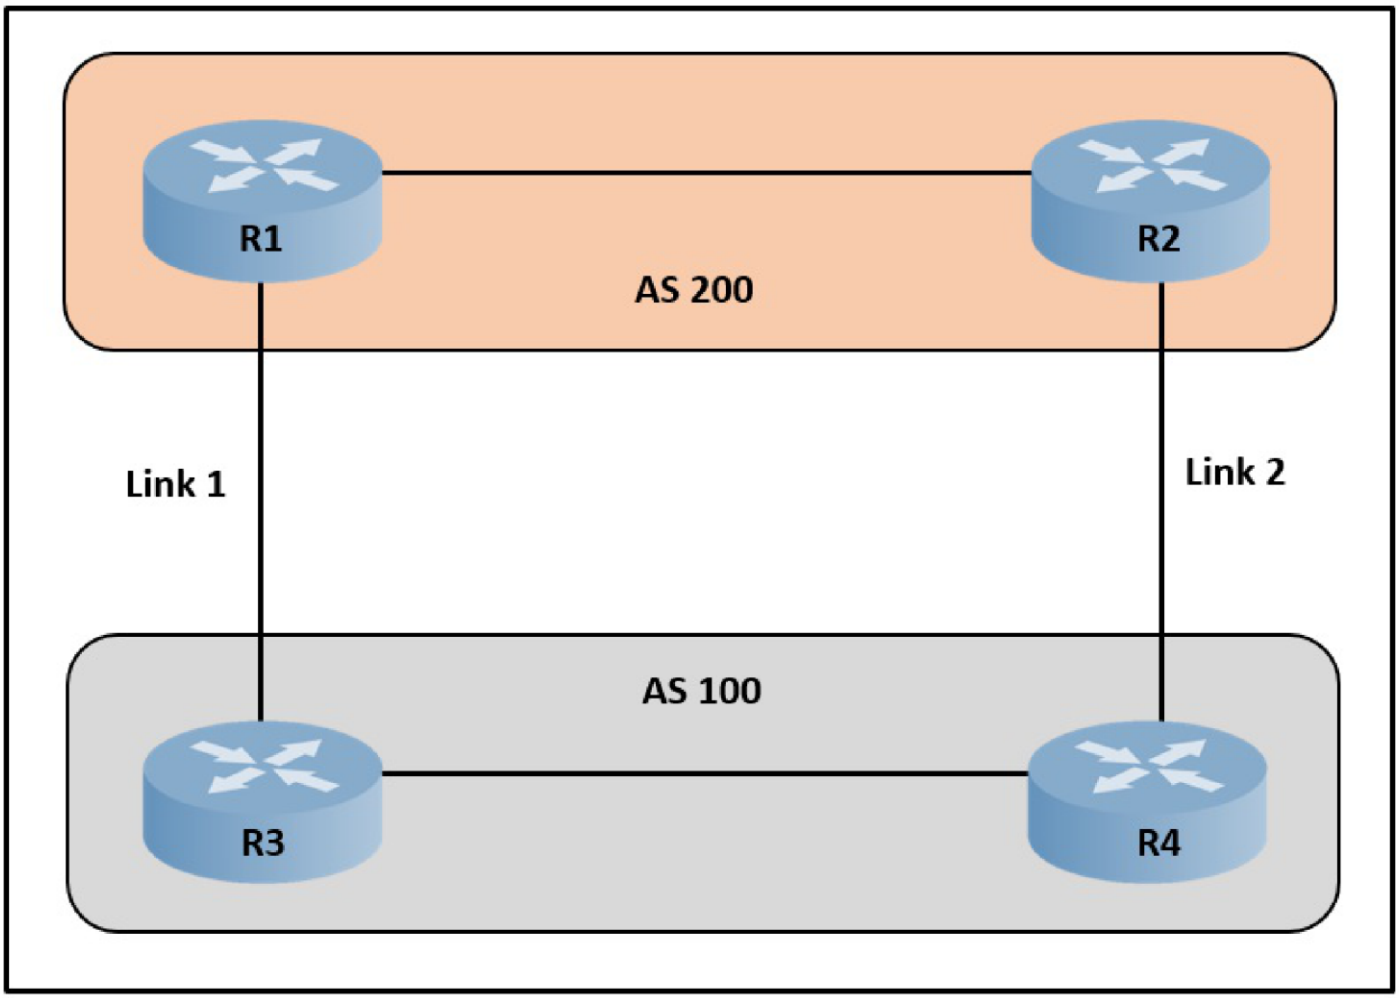

A network engineer is designing a network for AS100. The design should ensure that all traffic enters AS100 via link 1 unless there is a network failure. In the event of a failure, link 2 should function as the path for incoming traffic. Which solution should the design include?

A.Modify the next-hop attribute on R3.
B.Use AS-Path prepending on R3.
C.Modify the next-hop attribute on R4.
D.Use AS-Path prepending on R4.

Answer: D

Explanation:

Use AS-Path prepending on R4.

Question: 40

Refer to the exhibit.

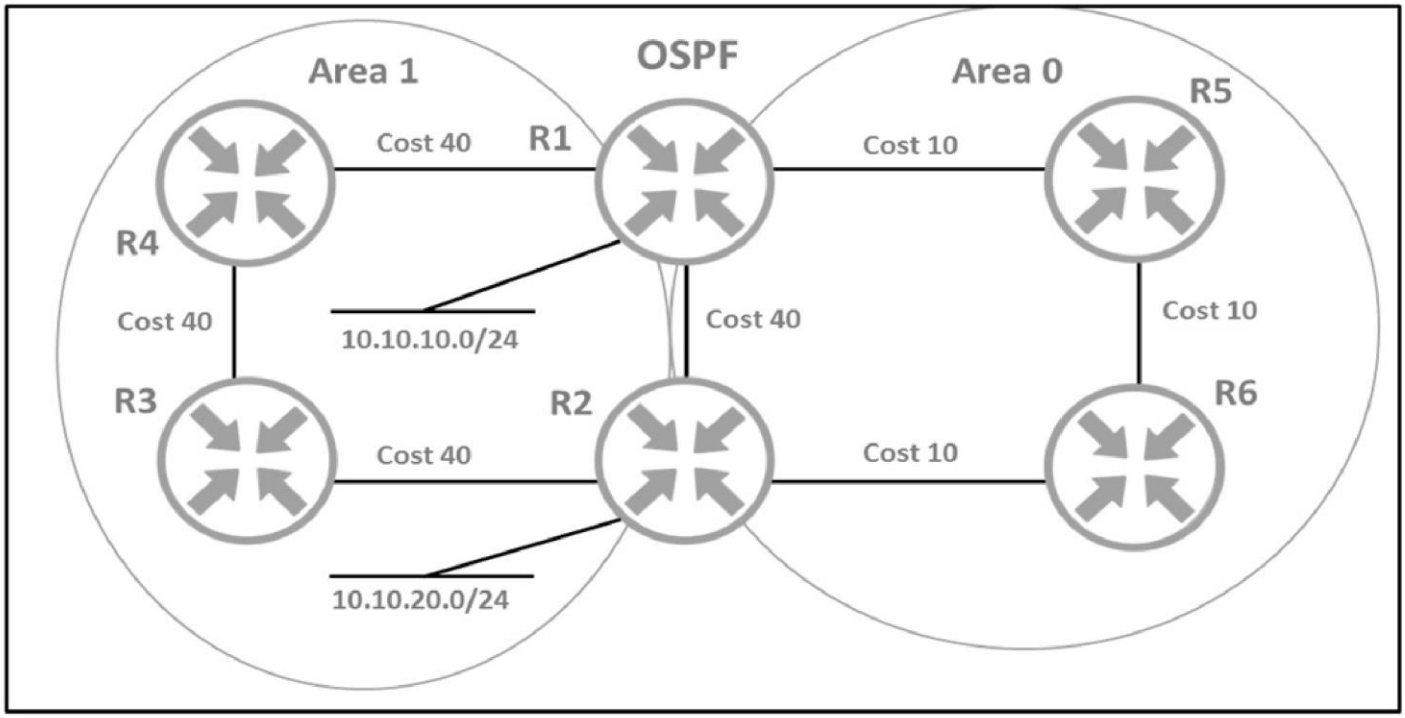

An architect must design a solution that uses the direct link between R1 and R2 for traffic from 10.10.10.0/24 toward network 10.10.20.0/24. Which solution should the architect include in the design?

A. Configure the OSPF cost of the link to a value lower than 30.

B. Lower the Administrative Distance for OSPF area 0.

C. Place the link into area 2 and install a new link between R1 and R2 in area 0. D. Configure the link to provide multiarea adjacency.

Answer: D

Explanation:

Configure the link to provide multiarea adjacency.

Question: 41

Refer to the exhibit.

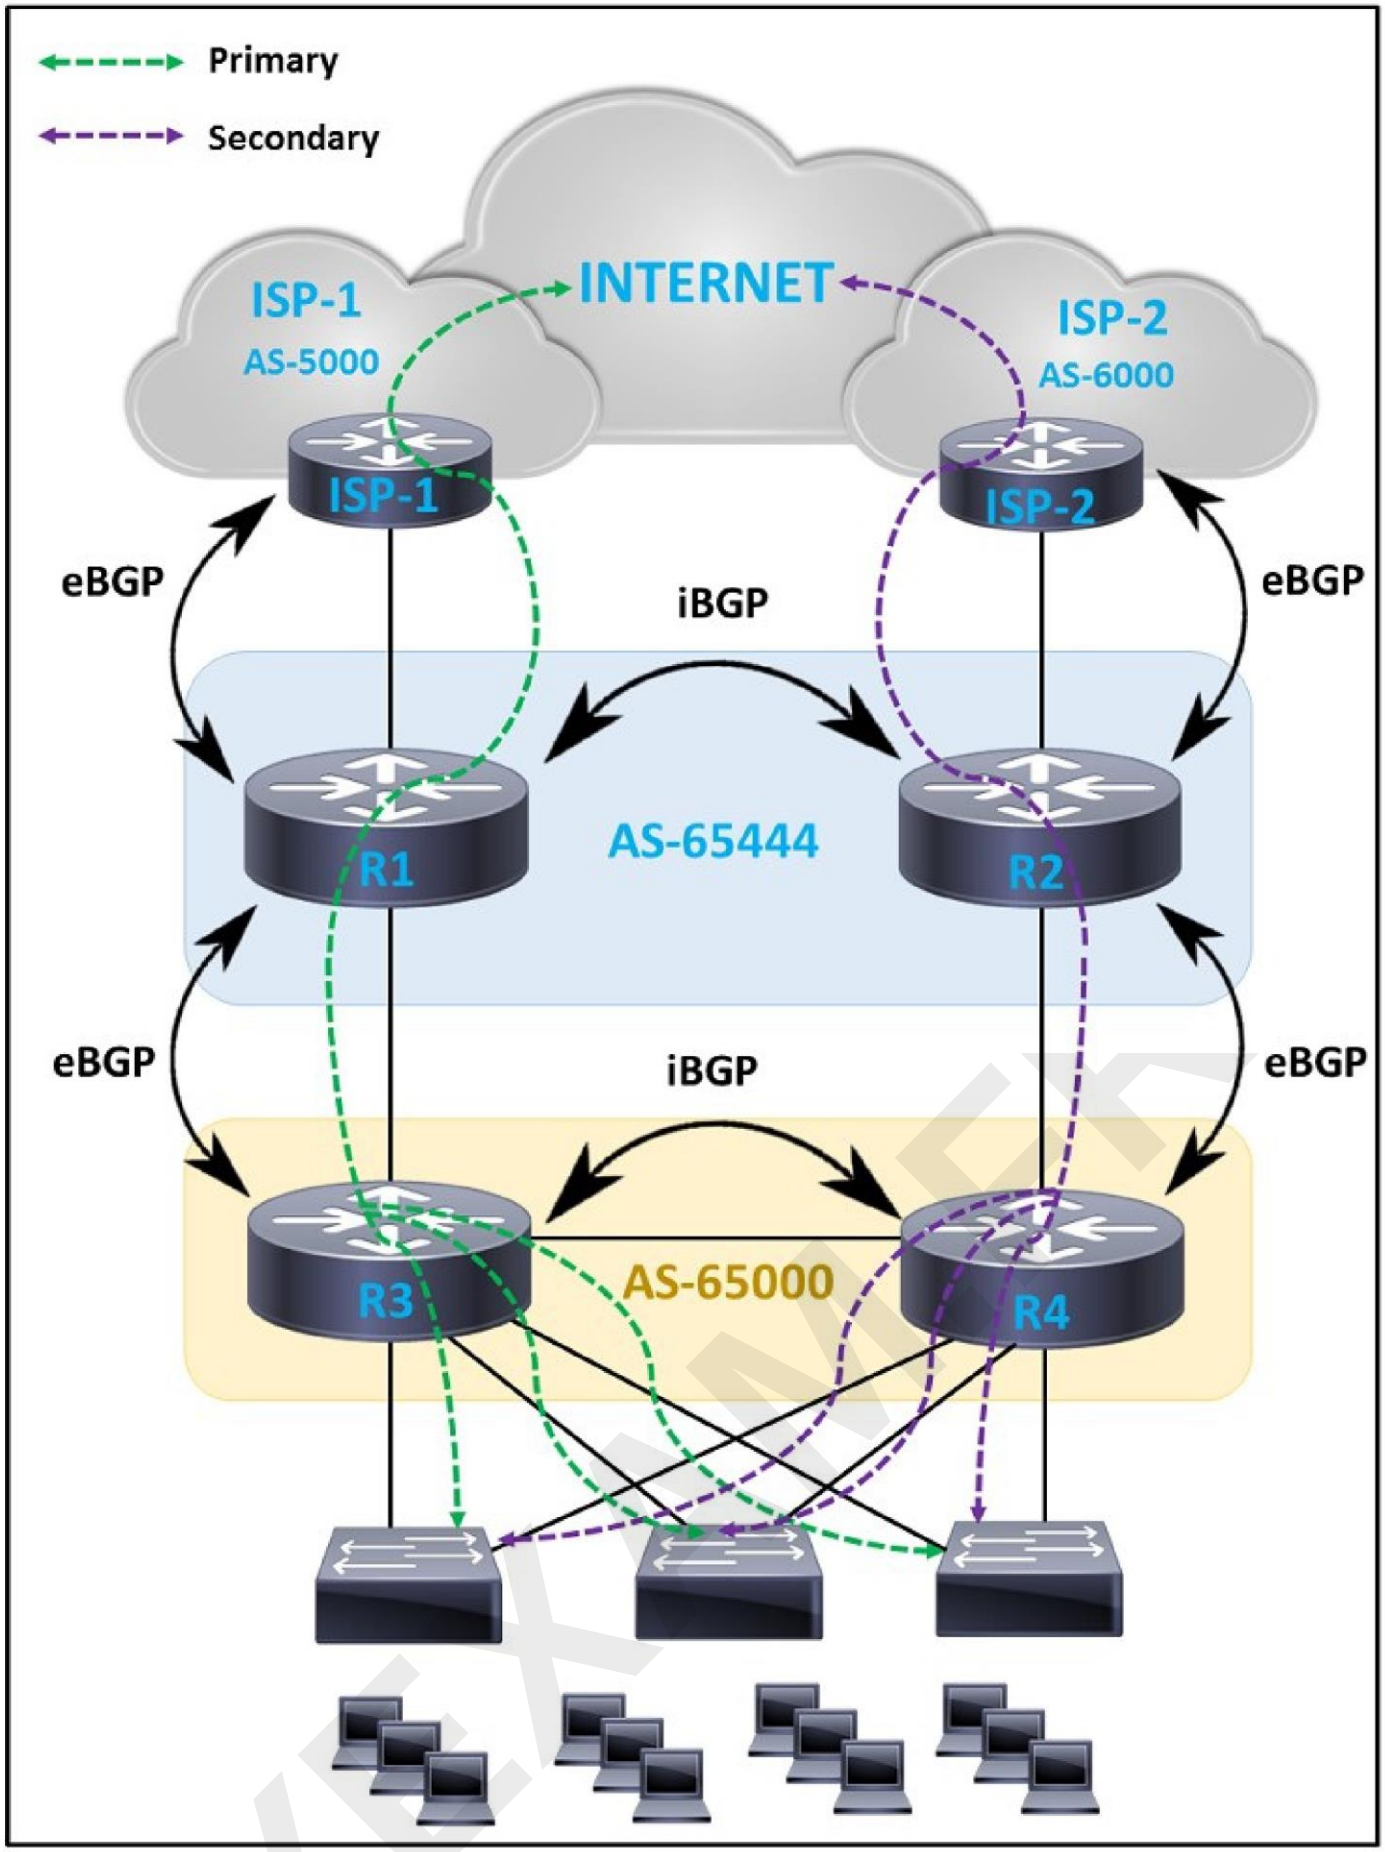

An engineer must design a WAN solution so that ISP-1 is always preferred over ISP-2. The path via ISP-2 is considered as a backup and must be used only when the path to ISP-1 is down. Which solution must the engineer choose?

A.R1: - Routes advertised to ISP-1: 0x AS-path prepend - Routes received from ISP-1: HIGH local-preference -Routes advertised to R2: no action - Routes received from R2: community NO-EXPORT R2: - Routes advertised to ISP-2:5x AS-path prepend - Routes received from ISP-2: LOW local-preference - Routes advertised to R1: community NO-ADVERTISE - Routes received from R1: no action

B.R1: - Routes advertised to ISP-1: 0x AS-path prepend - Routes received from ISP-1: HIGH local-preference -Routes advertised to R2: community NO-EXPORT - Routes received from R2: no action R2: - Routes advertised to ISP-2: 5x AS-path prepend - Routes received from ISP-2: LOW local-preference - Routes advertised to R1: no action - Routes received from R1: no action

C.R1: - Routes advertised to ISP-1: 0x AS-path prepend - Routes received from ISP-1: LOW local-preference -Routes advertised to R2: community NO-ADVERTISE - Routes received from R2: no action R2: - Routes advertised to ISP-2: 5x AS-path prepend - Routes received from ISP-2: HIGH local-preference - Routes advertised to R1: no action - Routes received from R1: community NO-ADVERTISE

D.R1: - Routes advertised to ISP-1: 5x AS-path prepend - Routes received from ISP-1: LOW local-preference -Routes advertised to R2: community NO-ADVERTISE - Routes received from R2: no action R2: - Routes advertised to ISP-2: 0x AS-path prepend - Routes received from ISP-2: HIGH local-preference - Routes advertised to R1: community NO-EXPORT - Routes received from R1: no action

Answer: B

Explanation:

R1: - Routes advertised to ISP-1: 0x AS-path prepend - Routes received from ISP-1: HIGH local-preference -Routes advertised to R2: community NO-EXPORT - Routes received from R2: no action R2: - Routes advertised to ISP-2: 5x AS-path prepend - Routes received from ISP-2: LOW local-preference - Routes advertised to R1: no action - Routes received from R1: no action.

Question: 42

Which feature must be incorporated into the campus LAN design to enable Wake on LAN?

A.dynamic ARP Inspection Snooping on layer 2 devices

B.directed broadcasts on layer 3 devices

C.proxy ARP on layer 3 devices

D.DHCP Snooping on layer 2 devices

Answer: B

Explanation:

The correct answer is **B. directed broadcasts on layer 3 devices**. Wake-on-LAN (WoL) relies on sending a "magic packet" to a sleeping device. This packet, typically a UDP broadcast on a specific port, must reach the target device despite its inactive network interface. Standard Layer 2 broadcasts are confined to a single broadcast domain. Therefore, they won't cross VLAN boundaries or reach a device on a different subnet. To enable WoL across different subnets, layer 3 devices, like routers, must be configured to forward directed broadcasts. Directed broadcasts are specific broadcasts aimed at a subnet's broadcast address. By enabling directed broadcasts, the router receives the initial broadcast and then forwards it to the target subnet, allowing the magic packet to reach the sleeping device. This method is preferred over simply using general broadcasts which could become network disruptive.

Options A, C, and D are not related to enabling WoL functionalities across networks: Dynamic ARP Inspection (DAI) is a security feature designed to prevent man-in-the-middle attacks by validating ARP packets; Proxy ARP enables a router to respond to ARP requests on behalf of another device; and DHCP Snooping is a security feature that inspects DHCP traffic, preventing unauthorized DHCP servers. None of these three features facilitates the routing or forwarding of magic packets necessary for WoL across Layer 3 boundaries.

Further Research:

1. Cisco documentation on Directed Broadcasts: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/configuration/15-mt/iad-15-mt-book/iad-directed-bcast.html

2. Wake-on-LAN Fundamentals: https://en.wikipedia.org/wiki/Wake-on-LAN
3. Understanding Broadcasts: https://www.geeksforgeeks.org/broadcast-address-in-computer-network/

Question: 43

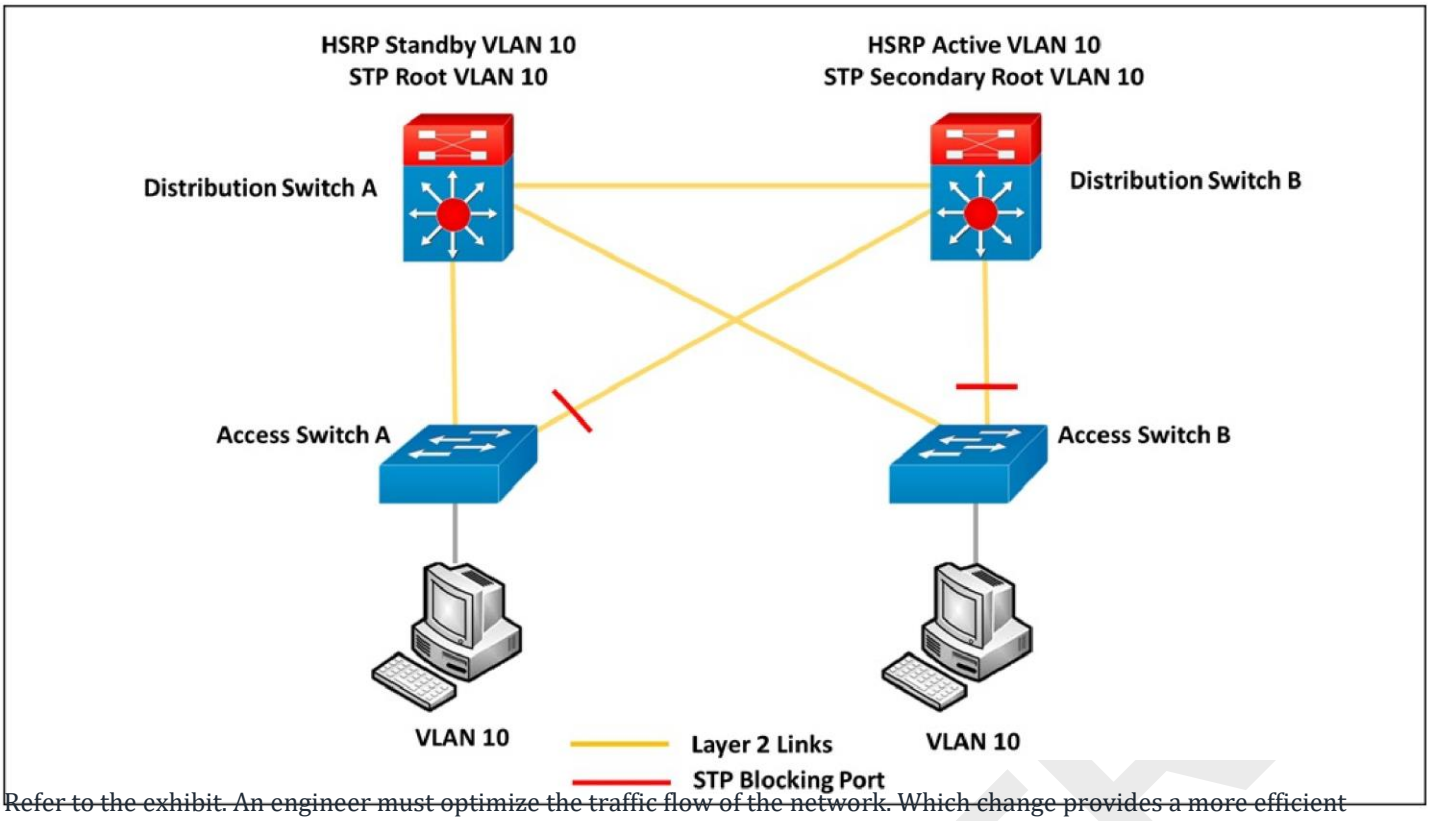

~~Refer to the exhibit. An engineer must optimize the traffic flow of the network. Which change provides a more efficient~~ design between the access layer and the distribution layer?

A.Add a link between access switch A and access switch B
B.Reconfigure the distribution switch A to become the HSRP Active
C.Change the link between distribution switch A and distribution switch B to be a routed link
D.Create an EtherChannel link between distribution switch A and distribution switch B

Answer: B

Explanation:

Reconfigure the distribution switch A to become the HSRP Active.

Question: 44

Which first hop redundancy protocol ensures that load balancing occurs over multiple routers using a single virtual IP address and multiple virtual MAC addresses?

A.GLBP
B.IRDP
C.VRRP
D.HSRP

Answer: A

Explanation:

The correct answer is A, **GLBP (Gateway Load Balancing Protocol)**. GLBP is designed to provide first-hop redundancy while actively load balancing traffic across multiple routers. Unlike other protocols, GLBP uses a single virtual IP address but distributes traffic by having each router act as a "forwarder" for a subset of the traffic using different virtual MAC addresses. This allows all participating routers to be actively forwarding, optimizing network resource utilization and reducing the impact of a single router failure. HSRP (Hot Standby Router Protocol) and VRRP (Virtual Router Redundancy Protocol) both use a single virtual MAC address and elect one router as active, while the others remain in standby. They do not offer active load balancing. ICMP Router Discovery Protocol (IRDP) is used by hosts to discover available routers; it's not a redundancy protocol and does not offer load balancing. GLBP's use of multiple virtual MAC addresses associated with a single virtual IP is the key difference that provides its load-balancing feature, which directly addresses the question's requirement. Each router, using its unique virtual MAC, forwards frames based on a round-robin or weighted round-robin algorithm. This prevents any one router from becoming a bottleneck. This approach makes GLBP a superior choice when active load-balancing is a primary requirement for first-hop redundancy, allowing network architects to leverage the full bandwidth of multiple gateway routers.

Further Reading:

Cisco's GLBP Documentation: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_fh/configuration/15-sy/fhp-15-sy-book/fhp-gateway-lb.html
Wikipedia's GLBP Entry: https://en.wikipedia.org/wiki/Gateway_Load_Balancing_Protocol

Question: 45

A company with multiple service providers wants to speed up BGP convergence time in the event a failure occurs with their primary link. Which approach achieves this goal and does not impact router CPU utilization?

A.Utilize BFD and tune the multiplier to 50
B.Lower the BGP hello interval
C.Decrease the BGP keepalive timer
D.Utilize BFD and keep the default BGP timers

Answer: D

Explanation:

The correct answer is **D. Utilize BFD and keep the default BGP timers.**

Bidirectional Forwarding Detection (BFD) is a low-overhead protocol designed to quickly detect failures in the path between two forwarding engines. Unlike BGP timers which rely on keepalives exchanged periodically, BFD proactively monitors the connectivity with rapid, sub-second detection capabilities. This rapid detection allows the routing protocol (in this case, BGP) to react much faster to a link failure, thus significantly speeding up convergence time. BFD operates independently of BGP keepalives. Tuning BFD timers to faster intervals and a lower multiplier than default may increase BFD's resource utilization, potentially impacting router CPU.

Keeping the BFD timers at their default values is a good balance of rapid failure detection and low resource consumption. Options A, B, and C all impact BGP timers. Tuning BGP timers may speed up failure detection somewhat but at the expense of increased resource consumption as those timers are processed by the CPU and are not designed to be sub-second like BFD. Also, while BFD is designed for low-overhead operation, drastically shortening the BFD timer values like changing the multiplier to 50 can inadvertently increase CPU load. In this scenario, it is best practice to rely on the default values, which still allow for rapid failure detection but do not over-tax the system.

Here are some authoritative links for further research:

Cisco's BFD Configuration Guide:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/15-sy/bfd-15-sy-book.html
RFC 5880: Bidirectional Forwarding Detection:https://datatracker.ietf.org/doc/html/rfc5880
Understanding BFD:https://networklessons.com/bgp/bidirectional-forwarding-detection-bfd

Question: 46

An engineer is designing an enterprise campus network. The LAN infrastructure consists of switches from multiple vendors, and Spanning Tree must be used as a
Layer 2 loop prevention mechanism. All configured VLANs must be grouped in two STP instances. Which standards-based Spanning Tree technology supports this design solution?

A.MSTP
B.RSTP
C.Rapid PVST
D.STP

Answer: A

Explanation:

The correct answer is A. MSTP (Multiple Spanning Tree Protocol). MSTP is the only standards-based Spanning Tree protocol among the options that allows for the grouping of VLANs into multiple independent spanning tree instances. This is crucial for the scenario described, which requires two STP instances and the flexibility to assign VLANs to them. RSTP (Rapid Spanning Tree Protocol) and STP (Spanning Tree Protocol) operate on a per-VLAN basis (PVST for Cisco proprietary) or a single instance for all VLANs, limiting scalability and the ability to segment traffic effectively. Rapid PVST+ is a Cisco proprietary enhancement and therefore would not work in a multivendor enviroment. MSTP, defined in IEEE 802.1s, creates regions, and each region can run multiple instances. Each VLAN is mapped to an instance, and each instance computes a separate spanning tree, allowing for load balancing and redundancy while preventing loops. This meets the requirement of having different spanning-tree topologies for different VLAN groups, making it the ideal solution for the described campus network with multi-vendor infrastructure. This enables more efficient resource utilization and greater flexibility in controlling data flow. Other STP variations like RSTP and PVST+ do not have the same capability to map multiple VLANs to a single spanning tree instance.

Further Research:

IEEE 802.1s:https://standards.ieee.org/ieee/802/1479/
Cisco MST Overview:https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SX/configuration/guide/mst_cfg.html
Juniper MST Overview:https://www.juniper.net/documentation/en_US/junos/topics/concept/spanning-tree-mstp-understanding.html

Question: 47

A network engineer must segregate three interconnected campus networks using IS-IS routing. A two-layer hierarchy must be used to support large routing domains and to avoid more specific routes from each campus network being advertised to other campus network routers automatically. Which two actions does the engineer take to accomplish this segregation? (Choose two.)

A. Designate two IS-IS routers as BDR routers at the edge of each campus, and configure one BDR for all Level 1 routers and one BDR for all Level 2 routers.

B. Designate two IS-IS routers from each campus to act as Level 1/Level 2 backbone routers at the edge of each campus network.

C. Assign the same IS-IS NET value for each campus, and configure internal campus routers with Level 1/Level 2 routing.

D. Utilize different MTU values for each campus network segment. Level 2 backbone routers must utilize a larger MTU size of 9216.

E. Assign a unique IS-IS NET value for each campus, and configure internal campus routers with Level 1 routing.

Answer: BE

Explanation:

The correct actions to segregate campus networks using IS-IS with a two-layer hierarchy are **B and E**. Option B is correct because designating Level 1/Level 2 routers at the edge of each campus creates a clear boundary. These routers act as intermediaries, connecting the internal Level 1 area of a campus to the Level 2 backbone.

This design prevents detailed routing information from one campus from being unnecessarily distributed to other campuses. Each campus becomes a Level 1 area, while the interconnected routers form a Level 2 backbone, fulfilling the two-layer hierarchy requirement. Option E is also correct; assigning unique NET values for each campus defines separate IS-IS routing domains. Level 1 routers within each campus will then only form adjacencies and exchange routes within their specific area defined by their NET value. This prevents automatic route summarization at the Level 2 backbone as the Level 1 areas are distinct. Option A is incorrect, IS-IS BDR's don't differentiate between levels, they manage the election process in multi-access network segment and not area segregation, Option C is incorrect as using same NET values would not create separate IS-IS areas. Option D is also incorrect because, though MTU can affect IS-IS peering, it's irrelevant to route separation and network segregation. IS-IS is designed to handle differences in MTU. The use of a Level 1 and Level 2 structure is important for scalability and managing routing complexity. Cisco's IS-IS configuration guide provides more details on IS-IS area design and how to configure different levels.

Question: 48

Which consideration must be taken into account when using the DHCP relay feature in a Cisco SD-Access Architecture?

A. DHCP-relay must be enabled on fabric edge nodes to provide the correct mapping of DHCP scope to the local anycast gateway.

B. A DHCP server must be enabled on the border nodes to allow subnets to span multiple fabric edges.

C. DHCP servers must support Cisco SD-Access extensions to correctly assign IPs to endpoints in an SD-Access fabric with anycast gateway.

D. DHCP Option-82 must be enabled to map the circuit IP option to the access fabric node where the DHCP discover originated.

Answer: D

Explanation:

The correct answer is D. DHCP Option 82 is crucial in Cisco SD-Access when using DHCP relay because it allows the DHCP server to identify the specific network segment from which a DHCP request originated. In an SD-Access fabric with an anycast gateway, multiple fabric edge nodes might share the same subnet. Without Option 82, the DHCP server wouldn't know which edge node and thus which VLAN a client belongs to, potentially assigning an incorrect IP address. Option 82 embeds information about the fabric edge node where the DHCP request was first received. This information allows the DHCP server to apply the correct scope or policy associated with that specific edge and ensure that IP assignments align with the intended network segmentation. Specifically, in Cisco SD-Access, Option 82 is instrumental in mapping the client's

location to the correct Virtual Network (VN), and the corresponding IP pool from which to draw an IP address.
DHCP relay itself simply forwards the request; it doesn't provide the context needed for SD-Access segmentation. While a DHCP server must be reachable (answer B is incorrect because it doesn't have to be hosted on border nodes) and DHCP scopes must be correctly configured, they alone don't provide the crucial mapping that Option 82 provides. Option 82, therefore, is fundamental to proper DHCP operation and segmentation in an SD-Access fabric. Option A is incorrect because DHCP relay should be configured on the layer 3 interface on fabric edge nodes and it does not provide the mapping of scope to anycast gateway. Option C is incorrect as DHCP servers are standard DHCP servers, no Cisco specific SD-Access extensions are required.

Authoritative Links:

Cisco SD-Access Deployment Guide:https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html (Look for sections discussing DHCP relay and Option 82 within the SD-Access architecture.)
Cisco DHCP Option 82 Configuration Guide:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/15-sy/dhcp-15-sy-book/dhcp-relay-agent-information.html (Focuses on the general use of Option 82, applicable to SD-Access environments)

Question: 49

Which function are fabric intermediate nodes responsible for in an SD-Access Architecture?

A.mapping EIDs to RLOCs
B.encapsulating user traffic in a VXLAN header including the SGT
C.registering new endpoints in the HTDB
D.transporting IP packets between edge nodes and border nodes

Answer: D

Explanation:

The correct answer is **D. transporting IP packets between edge nodes and border nodes.**

In Cisco's Software-Defined Access (SD-Access) architecture, fabric intermediate nodes act as transit points within the underlay network. They are responsible for forwarding traffic based on the IP address of the destination, not for higher-layer functions like endpoint registration or VXLAN encapsulation. The fabric intermediate nodes primarily concern themselves with moving the IP packets between the edge and border nodes, analogous to a typical routing function in a traditional network. They do not participate in the SD-Access overlay functions such as EID-to-RLOC mapping, which is the domain of the control plane elements like the Cisco Identity Services Engine (ISE) or DNAC. Similarly, VXLAN encapsulation and decapsulation are primarily handled by the fabric edge nodes. The intermediate nodes are essentially the plumbing, carrying the packets from one section of the fabric to another. Therefore, transporting IP packets is their core
responsibility, facilitating the communication flow within the SD-Access fabric.

Here are some authoritative links for further research on Cisco SD-Access:

Cisco SD-Access Solution Design Guide:https://www.cisco.com/c/en/us/solutions/enterprise/design-zone/networking/sda-solution-design.html
Cisco DNA Center User Guide:https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/2-3-3/user_guide/b_cisco_dna_center_user_guide_2_3_3/m_sda_fabric_provisioning.html
Cisco SD-Access Fabric Design White Paper: (Search on Cisco's website using keywords like "Cisco SD-Access Fabric Design White Paper").

These links provide detailed information on the roles and functions of different nodes within an SD-Access architecture.

Question: 50

How do endpoints inside an SD-Access network reach resources outside the fabric?

A.a VRF fusion router is used to map resources in one VN to another VN

B.Fabric borders use VRFs to map VNs to VRFs

C.SD-Access transit links are used to transport encapsulated traffic from one fabric to another

D.A fabric edge is used to de-encapsulate VXLAN traffic to normal IP traffic then transported over the outside network

Answer: B

Explanation:

The correct answer is **B. Fabric borders use VRFs to map VNs to VRFs.** This is the primary mechanism for connecting SD-Access fabric endpoints to external resources.

SD-Access uses Virtual Networks (VNs) to segment traffic within the fabric. To reach resources outside, the traffic needs to be routed out of these VNs. Fabric border nodes are specifically designed for this purpose.

They act as gateways between the fabric and the external network. Border nodes maintain a VRF (Virtual Routing and Forwarding) instance for each VN present in the fabric. When traffic destined for an external network arrives at the border node, it de-encapsulates the VXLAN traffic. Then, based on the VN, it uses the corresponding VRF to route the traffic out onto the external network. This mapping of VN to VRF at the border ensures seamless connectivity between the fabric and the outside world. Each VRF on the border node would then have its associated routing protocols and routing table. Option A is incorrect because VRF fusion is related to the intercommunication of VRFs, not primarily the function of reaching resources outside the fabric.

Option C is incorrect because SD-Access transit links are used to connect different fabrics or provide core transit. Option D is incorrect as the edge node is not the correct element to de-encapsulate traffic to normal IP traffic for onward transport on external networks; that is the purpose of the border node.

Authoritative Links:

Cisco SD-Access Design Guide: (Search on Cisco's website for the latest version). This guide provides comprehensive information on SD-Access architecture and its components, including border nodes. **Cisco DNA Center Documentation:** (Search on Cisco's website for the latest version). This documentation will cover the process of configuring and managing border nodes within SD-Access.

In summary, the fabric border node, utilizing VRFs to map VNs, is the critical element enabling SD-Access endpoints to reach resources outside the fabric.

Question: 51

When vEdge router redundancy is designed, which FHRP is supported?

A. HSRP

B. OMP

C. GLBP

D. VRRP

Answer: D

Explanation:

The correct answer is D, VRRP (Virtual Router Redundancy Protocol). When designing for vEdge router redundancy in a Cisco SD-WAN (formerly Viptela) environment, VRRP is the supported First Hop Redundancy Protocol (FHRP). While other FHRPs like HSRP (Hot Standby Router Protocol) and GLBP (Gateway Load Balancing Protocol) are widely used in traditional Cisco networks, they are not directly supported for vEdge router redundancy within the SD-WAN fabric. vEdge routers primarily rely on the OMP (Overlay Management Protocol) for route distribution and control plane functions within the SD-WAN overlay. While OMP provides a level of path redundancy, VRRP is essential for ensuring seamless failover of the default gateway role at the LAN edge when a vEdge router experiences an issue. VRRP allows multiple vEdge routers to share a common virtual IP address, with one acting as the active router while others stand by as backups. In case of an active router failure, one of the backup routers quickly takes over, minimizing service disruption. Therefore, VRRP is the designated FHRP for achieving high availability for the LAN side connectivity of vEdge routers within the Cisco SD-WAN architecture. Cisco's official documentation consistently highlights VRRP as the supported FHRP mechanism.

Authoritative Link:

Cisco SD-WAN Design Guide: https://www.cisco.com/c/en/us/td/docs/solutions/CVD/SDWAN/cisco-sdwan-design-guide.html (Search for "VRRP" within the document)

Question: 52

What is the purpose of an edge node in an SD-Access network fabric?

A. Edge nodes identify and authenticate endpoints and register endpoint information with control plane nodes.
B. Edge nodes track endpoint IDs to location mappings, along with IPv4, IPv6, or MAC addresses.
C. Edge nodes are the gateway between the fabric domain and network outside of the fabric.
D. Edge nodes resolve lookup requests from edge and border nodes to locate destination endpoint IDs.

Answer: A

Explanation:

The correct answer is **A. Edge nodes identify and authenticate endpoints and register endpoint information with control plane nodes.**

In a Cisco SD-Access fabric, edge nodes serve as the access layer devices directly connected to user endpoints (e.g., PCs, phones, IoT devices). These nodes are responsible for the initial interaction with these endpoints. Their primary role is to discover and authenticate devices upon connection to the network. Once authenticated, the edge node registers relevant endpoint information, including identity, location, and MAC address, with the control plane nodes, typically acting as database or management servers. This registration process allows the fabric to maintain a comprehensive inventory of connected devices and their locations, enabling features like policy enforcement and microsegmentation. In essence, edge nodes are the first point of contact in the fabric for endpoint awareness.

Option B describes a function more associated with the control plane, which maintains the endpoint-to-location mappings. Option C defines the function of border nodes, which act as the gateways to external networks. Option D describes the role of control plane nodes, which manage the endpoint lookups. Edge nodes initiate the process, not resolve lookup requests, which are handled by the control plane.

For further research, refer to the following Cisco documentation:

Cisco DNA Center SD-Access Solution Overview:https://www.cisco.com/c/en/us/solutions/enterprise/dna-center/sd-access-solution.html
Cisco SD-Access Design Guide:https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.htmlThese resources provide in-depth explanations of SD-Access architecture and its components.

Question: 53

Which component of Cisco SD-Access integrates with Cisco DNA Center to perform policy segmentation and enforcement through the use of security group access control lists and security group tags?

A. Cisco Application Policy Infrastructure Controller Enterprise Module
B. Cisco Network Data Platform
C. Cisco Identity Services Engine
D. Cisco TrustSec

Answer: C

Explanation:

The correct answer is **C. Cisco Identity Services Engine (ISE)**. Cisco ISE is the cornerstone for identity and access control within Cisco SD-Access, particularly when it comes to policy segmentation and enforcement. It achieves this by integrating with Cisco DNA Center. ISE centrally manages user and device identities, assigning them to security groups. These security groups are tagged using Security Group Tags (SGTs). When a user or device attempts to access network resources, ISE communicates the assigned SGT to network devices. These devices, using SGT-based access control lists (SGACLs), then enforce policies based on the SGTs, rather than solely relying on IP addresses. This approach enables micro-segmentation, granular access control, and dynamic policy changes based on user roles and device contexts. This enhances security by limiting lateral movement within the network and simplifying policy management.

The other options are incorrect:

A. Cisco Application Policy Infrastructure Controller Enterprise Module (APIC-EM): APIC-EM was a predecessor to DNA Center and primarily focused on network programmability rather than identity management and policy enforcement based on SGTs.

B. Cisco Network Data Platform (NDP): NDP is focused on collecting and analyzing network telemetry data, rather than policy management and enforcement.

D. Cisco TrustSec: While Cisco TrustSec is the overall framework that utilizes SGTs, ISE is the component responsible for the core functionality of assigning those tags and communicating them to the network infrastructure.

In summary, while other components play a role within the SD-Access ecosystem, ISE's role in identity management, security group assignment, and SGACL enforcement is directly tied to policy segmentation.

Authoritative Links for further research:

Cisco Identity Services Engine (ISE):https://www.cisco.com/c/en/us/products/security/identity-services-engine/index.html
Cisco SD-Access:https://www.cisco.com/c/en/us/solutions/enterprise-networks/software-defined-access/index.html
Cisco TrustSec:https://www.cisco.com/c/en/us/solutions/enterprise-networks/trustsec/index.html

Question: 54

Which design element should an engineer consider when multicast is included in a Cisco SD-Access architecture?

A. PIM SSM must run in the underlay.
B. Multicast clients reside in the underlay, and the multicast source is outside the fabric or in the overlay.
C. Rendezvous points must be used in a PIM SSM deployment.
D. Multicast traffic is transported in the overlay and the EID space for wired and wireless clients.

Answer: D

Explanation:

The correct answer is **D. Multicast traffic is transported in the overlay and the EID space for wired and wireless clients.** This is because Cisco SD-Access leverages a VXLAN-based overlay network to provide segmentation and abstraction of the physical infrastructure. In this architecture, endpoint identifiers (EIDs), which represent client IP addresses, are mapped to routable locators (RLOCs) in the underlay. Multicast traffic, like unicast, is encapsulated within VXLAN headers in the overlay. This means that multicast groups are defined within the EID space, allowing for multicast to be distributed to end devices wherever they are within the fabric. Wired and wireless clients, regardless of their physical location, can participate in the same multicast groups as their EIDs are used for group membership and traffic delivery within the overlay.

Option A is incorrect. While PIM-SM (Protocol Independent Multicast-Sparse Mode) can be used in the underlay for infrastructure multicast (like routing updates), the overlay itself does not depend on the underlay using PIM-SSM (Source Specific Multicast). The overlay handles multicast distribution based on the EIDs and group membership within the fabric. Option B is also incorrect, as multicast clients and sources can exist both inside and outside the fabric, in both the overlay and potentially connected to the underlay, but within the fabric they leverage EIDs and the VXLAN overlay. SD-Access doesn't strictly dictate where multicast entities can or cannot be. Option C is incorrect. Rendezvous Points (RPs) are used with PIM-SM, not PIM-SSM, and while PIM-SM is an option in the underlay to support the overall network, Cisco SD-Access does not mandate the use of Rendezvous Points.

Authoritative Links:

Cisco SD-Access Design Guide:https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/sda-design-guide.html (Specifically, the sections on Multicast and Overlay Networking)
Cisco Live Presentations on SD-Access: (Search on Cisco Live website for relevant presentations related to Multicast in SD-Access).
VXLAN Technology Overview:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/virtual-extensible-lan-vxlan/index.html (To understand how VXLAN encapsulation supports overlay networks)

Question: 55

What is the role of a control-plane node in a Cisco SD-Access architecture?

A. fabric device that connects wired endpoints to the SD-Access fabric
B. map system that manages endpoint to device relationships
C. fabric device that connects APs and wireless endpoints to the SD-Access fabric
D. map system that manages External Layer 3 networks

Answer: B

Explanation:

The correct answer is **B. map system that manages endpoint to device relationships.** In Cisco SD-Access, the control-plane node, often a Cisco Identity Services Engine (ISE) instance, acts as the central authority for mapping users and endpoints to their respective locations within the fabric. It doesn't directly handle data forwarding; instead, it dictates policies and defines how traffic should be handled within the SD-Access network. The control-plane node uses the LISP (Locator/ID Separation Protocol) database to maintain the mapping between endpoint identifiers (EIDs) and Routing Locators (RLOCs), which are IP addresses of fabric devices. This mapping is crucial for enabling location-independent routing and policy enforcement within the SD-Access fabric. When a new endpoint connects, the control plane (ISE) authenticates it, and subsequently registers the endpoint's EID (e.g., MAC or IP address) to its RLOC (the IP address of the access switch it connected to) in the LISP database. Fabric devices, acting as data plane nodes, then query this control-plane node for the appropriate EID-to-RLOC mapping to forward traffic correctly. Options A and C describe fabric devices (edge nodes and wireless access points respectively), and option D refers to border nodes. These devices are parts of the data plane, not the control plane. Therefore, the central mapping function that manages endpoint-to-device associations and that is the control plane is best defined by option B.

Authoritative Links:

Cisco SD-Access Solution Overview:https://www.cisco.com/c/en/us/solutions/enterprise-networks/software-defined-access/index.html
Cisco Identity Services Engine (ISE) Documentation:
https://www.cisco.com/c/en/us/support/security/identity-services-engine/tsd-products-support-series-home.html
LISP (Locator/ID Separation Protocol) RFC:https://www.rfc-editor.org/rfc/rfc6830

Question: 56

How is end-to-end microsegmentation enforced in a Cisco SD-Access architecture?

A.VLANs are used to segment traffic at Layer 2.
B.5-tuples and ACLs are used to permit or deny traffic.
C.SGTs and SGTACLs are used to control access to various resources.
D.VRFs are used to segment traffic at Layer 3.

Answer: C

Explanation:

The correct answer is **C. SGTs and SGTACLs are used to control access to various resources.**

Here's why: In a Cisco SD-Access architecture, microsegmentation isn't achieved through traditional methods like VLANs or VRFs, which provide macro-segmentation. Instead, SD-Access uses Security Group Tags (SGTs) to categorize users and devices based on their roles or functions. These SGTs are propagated throughout the network using Cisco's proprietary protocol, Secure Group Tag eXchange Protocol (SXP). Security Group ACLs (SGTACLs), based on SGTs instead of IP addresses, are then applied at network boundaries like switches and firewalls to enforce policies, controlling which groups can communicate with others. This provides granular access control down to the individual user or device level, allowing for true end-to-end microsegmentation, independent of IP addressing or VLANs.

Option A (VLANs) provides Layer 2 segmentation, not microsegmentation. Option B (5-tuples and ACLs) is a more traditional network security approach based on IP addresses and port numbers, not the identity-based microsegmentation used in SD-Access. Option D (VRFs) offers Layer 3 segmentation, creating separate routing domains but not the fine-grained control that SGTs provide.

For further research on Cisco SD-Access and microsegmentation with SGTs, consult the following links:

1. **Cisco SD-Access Solution Guide:**https://www.cisco.com/c/en/us/solutions/enterprise/design-zone/solution-overviews/sda-solution.html
2. **Cisco TrustSec:**https://www.cisco.com/c/en/us/solutions/enterprise/trustsec/index.html
3. **Understanding and Deploying Security Group Tags (SGTs) in Cisco DNA Center:** https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/2-3-3/configuration/b_cisco_dna_center_configuration_guide_2_3_3/b_cisco_dna_center_configuration_guide

Question: 57

Which two border nodes are available in the Cisco SD-Access architecture? (Choose two.)

A.extended border
B.edge border
C.internal border
D.anywhere border
E.intermediate border

Answer: CD

Explanation:

The correct border nodes within Cisco SD-Access architecture are **internal border** and **anywhere border**.

Internal border nodes connect fabric sites to each other, facilitating communication between different SD-Access fabric domains. These nodes route traffic between the virtualized network segments within the SD-Access deployment. They handle the internal routing and policy enforcement required for seamless connectivity between different fabric areas. An internal border manages the exchange of route information and maintains reachability within the overall SD-Access fabric, supporting network segmentation.

Anywhere border nodes, on the other hand, connect the SD-Access fabric to the external world, such as a WAN, Internet, or legacy networks. These nodes function as gateways, handling route redistribution and policy enforcement between the fabric and external entities. They act as the primary entry and exit points, ensuring secure and controlled traffic flow across network boundaries. They are the critical interfaces where the SD-Access fabric interacts with external networks.

Extended, edge, and intermediate borders are not standard components in Cisco's SD-Access architecture. Instead, the architecture is designed with a clear separation between internal fabric operations managed by the internal border and external connectivity controlled by the anywhere border. The Cisco SD-Access design prioritizes simplicity and efficiency with these core border node types.

Further Research:

Cisco SD-Access Design Guide
Cisco DNA Center Documentation

Question: 58

Which control-plane protocol is used to map an endpoint to a location in a Cisco SD-Access network?

A.FabricPath
B.IS-IS

C.LISP
D.MP-BGP

Answer: C

Explanation:

The correct answer is **C. LISP (Locator/ID Separation Protocol)**. In a Cisco SD-Access (Software-Defined Access) network, LISP is the core control-plane protocol responsible for mapping endpoint identifiers (EIDs), which are typically IP addresses of devices, to routing locators (RLOCs), which are the physical network locations where these devices are connected. This mapping is crucial for achieving network virtualization and location independence within the SD-Access fabric. When an endpoint moves, its EID remains the same, but its RLOC changes, allowing seamless mobility without requiring changes to the endpoint's IP address. The LISP control plane uses a map server/map resolver system to maintain this mapping database and enables efficient routing to endpoints regardless of their location within the network. FabricPath (A), IS-IS (B), and MP-BGP (D) serve other purposes. FabricPath is a Layer 2 protocol used for creating a scalable switching fabric; IS-IS is an interior gateway routing protocol commonly used within a network's core, and MP-BGP is a protocol used for inter-domain routing. While these protocols might play roles in an SD-Access deployment, they do not directly handle endpoint-to-location mapping. LISP is specifically designed to provide the necessary control-plane functionality to enable this essential capability of SD-Access.

Authoritative Links:

1. **Cisco SD-Access Solution Guide:**https://www.cisco.com/c/en/us/solutions/enterprise/design-zone/sd-access/index.html - This is the main page for Cisco's SD-Access solutions, which often refer to the role of LISP
2. **LISP Technology Overview:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_lisp/configuration/16-12/irl-16-12-book/lisp-overview.html - Cisco's official documentation explaining the function and benefits of LISP.

Question: 59

Which feature is required for graceful restart to recover from a processor failure?

A.Cisco Express Forwarding
B.Virtual Switch System
C.Stateful Switchover
D.Bidirectional Forwarding Detection

Answer: C

Explanation:

The correct answer is **C. Stateful Switchover (SSO)**.

Graceful restart, in the context of networking, aims to maintain forwarding and routing information during a control plane failure, minimizing disruption. Stateful Switchover (SSO) is specifically designed to enable this by mirroring the control plane state between redundant processors. When a failure occurs on the active processor, the standby processor, having a copy of the operational data, can immediately take over, ensuring continuous packet forwarding. This eliminates the need for lengthy convergence times associated with routing protocols re-establishing adjacencies.

Cisco Express Forwarding (CEF) (A) is a switching mechanism but doesn't inherently provide the state redundancy required for graceful restart after a processor failure. While CEF is crucial for efficient

forwarding, it relies on a working control plane. Virtual Switch System (VSS) (B) enhances availability through aggregation and redundancy but doesn't directly manage control plane state in the way SSO does. Bidirectional Forwarding Detection (BFD) (D) is a protocol for fast fault detection but does not provide the redundancy necessary for graceful recovery from a processor failure; rather it detects issues and triggers convergence.

Therefore, SSO is the specific feature that provides the necessary control plane state replication and failover mechanism to enable graceful restart after a processor failure. It guarantees seamless transfer of operations to the redundant processor with minimal service disruption.

Authoritative Links:

Cisco Documentation on Stateful Switchover:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mcl/configuration/16-6/mcl-16-6-book/mcl-stateful-switchover.html
Cisco Design Guide on High Availability:https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Campus/HA_Campus_DG/HA_Campus_DG

Question: 60

An architect is designing a network that will utilize the spanning tree protocol to ensure a loop-free topology. The network will support an engineering environment where it is necessary for end-users to connect their own network switches for testing purposes. Which feature should the architect include in the design to ensure the spanning-tree topology is not affected by these rogue switches?

A.BPDU Skew Detection
B.BPDU guard
C.loop guard
D.root guard

Answer: D

Explanation:

The correct answer is **D. Root guard.** Root guard is designed specifically to protect the Spanning Tree Protocol (STP) root bridge from being influenced by rogue or misconfigured switches. When a switch configured with root guard receives a superior Bridge Protocol Data Unit (BPDU) on a designated port (a port that should not be the root port), it will immediately move that port to a root-inconsistent state, effectively blocking the rogue switch from becoming the root. This prevents unauthorized switches from disrupting the stable STP topology by potentially becoming the root bridge, a scenario that could lead to network instability, suboptimal path selection, and even temporary network outages. While BPDU guard is used to protect against situations where a BPDU is received on a port that is not expected to have one, like end-user ports, it does not directly protect against unauthorized root bridge elections. Loop guard primarily focuses on ensuring a loop-free topology in situations where an STP port transitions to a forwarding state due to a unidirectional link, while BPDU skew detection is focused on identifying issues where there may be a misconfiguration that is causing BPDU's to not be sent or processed correctly. These three features are relevant for different types of issues in the network that are related to STP, but in the specific scenario of protecting against a rogue switch trying to become root, the best feature is root guard.

Authoritative Links for Further Research:

Cisco Documentation on Root Guard:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3850/software/release/16-12/configuration/guide/b_1612_stg_3850_cg/b_1612_stg_3850_cg_chapter_0110.html **Cisco Documentation on Spanning Tree Port Security:**

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3850/software/release/16-12/configuration/guide/b_1612_stg_3850_cg/b_1612_stg_3850_cg_chapter_0111.html

Question: 61

An engineer is designing a Layer 3 campus network running EIGRP between the core, aggregation, and access layers. The access layer switches will be connected to the aggregation layer using Layer 3 copper connections. The engineer wants to improve the convergence time for access layer switch failures.
Which technique must the design include?

A. enabling BFD for EIGRP on the access layer uplinks
B. reducing the EIGRP Hello / Hold timer values
C. EIGRP summarization from core to aggregation layer
D. EIGRP summarization from access to aggregation layer

Answer: A

Explanation:

The correct answer is **A. enabling BFD for EIGRP on the access layer uplinks.**

Bidirectional Forwarding Detection (BFD) is a lightweight, low-overhead protocol that provides rapid failure detection between two forwarding engines. In the context of EIGRP, enabling BFD on the access layer uplinks drastically reduces the time it takes for EIGRP to recognize a link failure. Without BFD, EIGRP relies on its hello and hold timers for adjacency maintenance. While adjusting these timers can speed up convergence to some degree, it comes at the cost of increased control plane overhead. BFD detects failures much faster, typically within milliseconds, compared to EIGRP's default timers, which are in the order of seconds. This fast failure detection capability allows EIGRP to trigger routing recalculation and convergence quicker, improving the overall network stability and reducing downtime experienced by end users when access layer switches fail.

Option B, reducing EIGRP timers, can expedite failure detection, but it also increases network traffic due to frequent hello packet exchanges and places a higher processing burden on network devices. Option C, EIGRP summarization from the core to the aggregation, reduces the routing table size and improves routing efficiency, but it does not directly impact the convergence time on access layer failures. Similarly, Option D, EIGRP summarization from access to aggregation, helps reduce routing table size in the aggregation layer, but it's unrelated to the goal of quickly detecting access layer failures. BFD specifically addresses this quick link failure detection need. BFD allows EIGRP to react rapidly to link failures, facilitating faster convergence and better network performance, which is essential in a campus network environment.

Authoritative Links:

Cisco Documentation on BFD: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/15-sy/ire-15-sy-book/ire-bfd.html
EIGRP Best Practices: https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/13676-eigrp-best.html

Question: 62

An existing network solution is using BFD in echo mode. Several network devices are experiencing high CPU utilization, which an engineer has determined is related to the BFD feature. Which solution should the engineer leverage to reduce the CPU load?

A.Implement slow timers between peers with low CPU resources.

B.Implement BFD asynchronous mode between peers with low CPU resources.

C.Enable BFD multi-hop on the devices with low CPU resources.

D.Utilize carrier delay on all routers in the network.

Answer: A

Explanation:

The correct answer is A, implementing slow timers between peers with low CPU resources, because BFD echo mode, while offering fast failure detection, can be CPU intensive due to the constant transmission and processing of echo packets. When devices experience high CPU utilization directly related to BFD, decreasing the frequency of these echo packets reduces the load. This is achieved by increasing the minimum transmit and receive intervals, effectively slowing down the BFD timers. Option B, switching to asynchronous mode, mitigates echo's CPU load but also relies on periodic hello packets; it doesn't inherently address high CPU usage if the asynchronous timers are too aggressive. Option C, enabling multi-hop, is irrelevant as it deals with BFD sessions over multiple hops which doesn't directly relate to CPU usage of a direct session. Option D, utilizing carrier delay, manages interface flapping and doesn't specifically reduce BFD processing. By using slow timers on resource constrained devices specifically, you maintain the fast fault detection capability on the fast links with fast timers and reduce overhead on slow resources with slow timers. This provides a more balanced implementation and is the most accurate solution for addressing the stated issue with minimal impact.

For further research, consider exploring these resources:

Cisco's BFD Configuration Guide: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/15-sy/iap-15-sy-book/bfd.html
Understanding BFD:
https://www.juniper.net/documentation/us/en/software/junos/routing/topics/concept/bfd-understanding.html

Question: 63

How is a sub-second failure of a transport link detected in a Cisco SD-WAN network?

A. Hellos are sent between the WAN Edge routers and the vSmart controller.

B. BFD runs on the IPsec tunnels between WAN Edge routers.

C. BGP is used between WAN Edge routers and the vSmart controller.

D. Link state change messages are sent between vSmart controllers.

Answer: B

Explanation:

The correct answer is B, BFD runs on the IPsec tunnels between WAN Edge routers. Bidirectional Forwarding Detection (BFD) is a lightweight, low-overhead protocol designed for fast failure detection in network paths. In a Cisco SD-WAN network, BFD sessions are established over the IPsec tunnels that connect WAN Edge routers, enabling rapid detection of transport link failures. This is crucial for maintaining application uptime and optimizing performance. When a link fails, BFD detects the disruption almost immediately, usually within sub-second timescales. This triggers a fast failover mechanism, redirecting traffic to alternative paths, thus minimizing downtime. The use of BFD on IPsec tunnels is a standard practice for efficient fault detection in various networking scenarios that require quick adaptation. Unlike relying solely on routing protocols, BFD operates at the data plane and doesn't depend on convergence time of routing protocols, providing faster detection capabilities. Options A, C and D use control plane protocol, not used for sub-second detection.
Hellos (Option A), are part of the control plane, and are related to the control plane between vEdge and

vSmart and does not detect data plane failure in sub-second time. BGP (Option C) is a routing protocol, and is also used for control plane, and slower than BFD in convergence. Link state change message (Option D) does not provide fast sub second detection of link failure.

Supporting Information:

Cisco SD-WAN Design Guide:https://www.cisco.com/c/en/us/td/docs/solutions/CVD/SDWAN/cisco-sdwan-design-guide.html (Search for "BFD")
BFD Protocol RFC:https://datatracker.ietf.org/doc/html/rfc5880

Question: 64

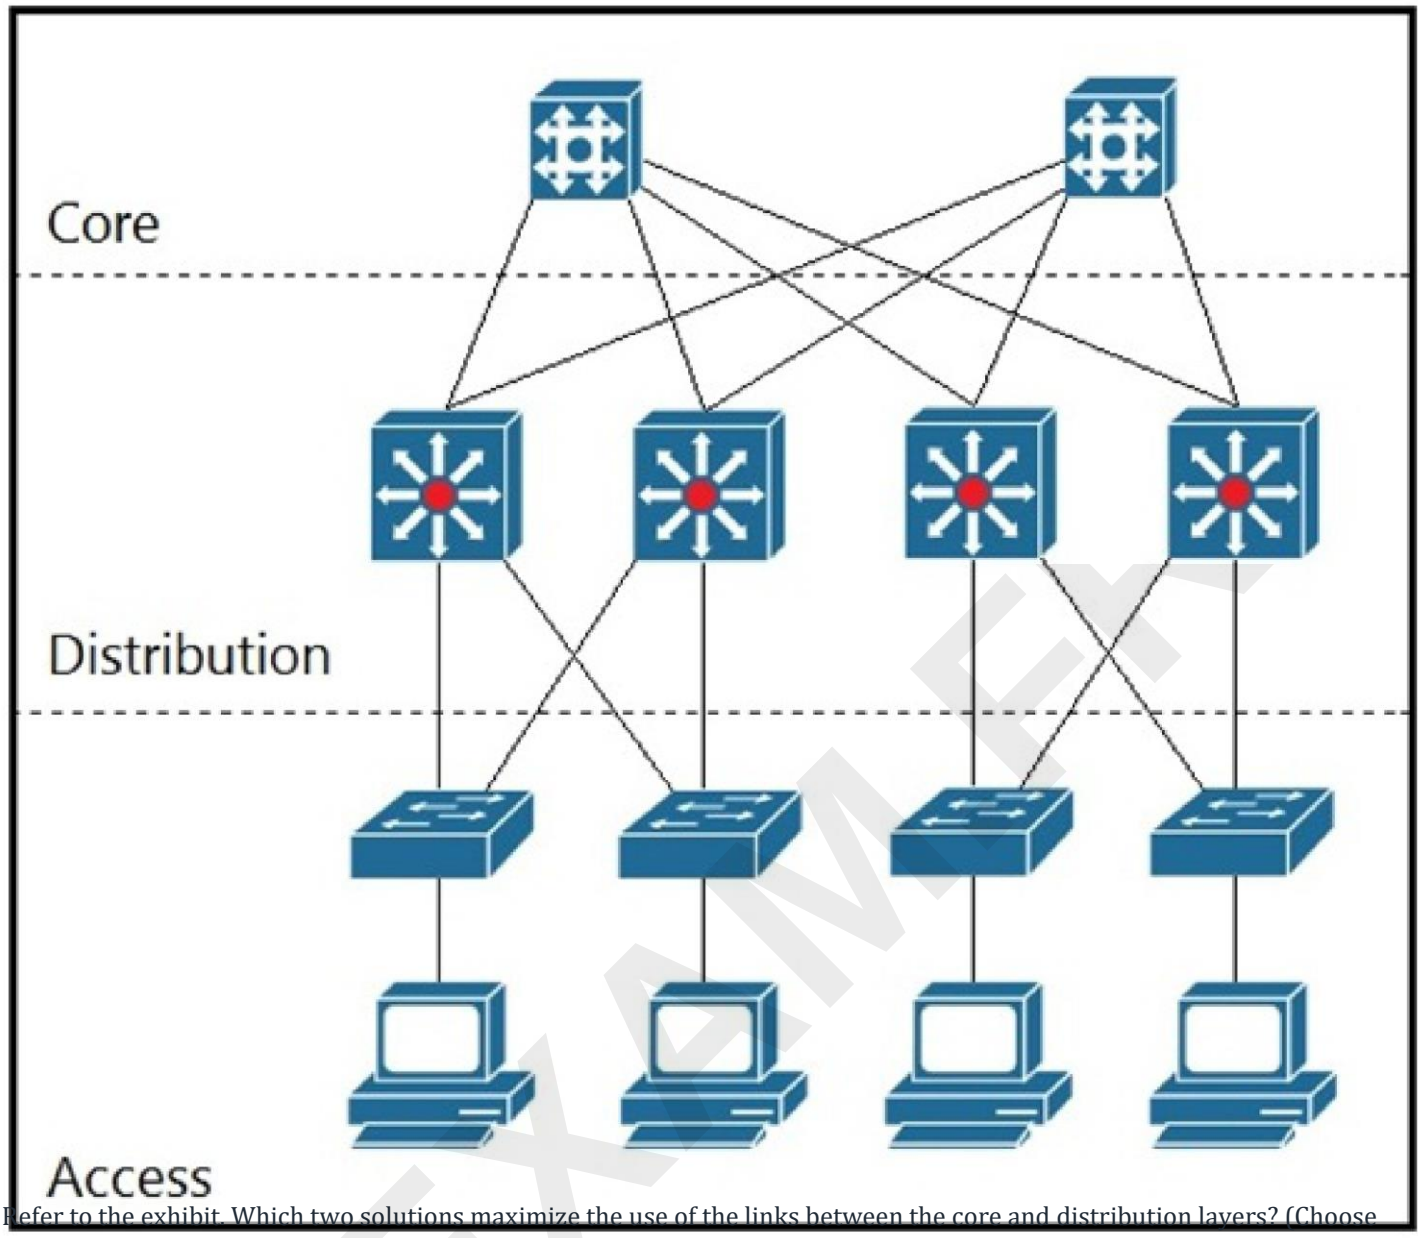

Refer to the exhibit. Which two solutions maximize the use of the links between the core and distribution layers? (Choose two.)

A.use multiple equal-cost links
B.use an IGP
C.use HSRP
D.use RPVSTP+
E.use multiple unequal-cost links

Answer: AB

Explanation:

A.use multiple equal-cost links.

B.use an IGP.

Question: 65

A customer's current Layer 2 infrastructure is running Spanning Tree 802.1d, and all configuration changes are manually implemented on each switch. An architect must redesign the Layer 2 domain to achieve these goals: * reduce the impact of topology changes
* reduce the time spent on network administration
* reduce manual configuration errors
Which two solutions should the architect include in the new design? (Choose two.)

A.Implement Rapid PVST+ instead of STP.

B.Implement MST instead of STP.

C.Use VTP to propagate VLAN information and to prune unused VLANs.

D.Configure broadcast and multicast storm control on all switches.

E.Configure dynamic trunking protocol to propagate VLAN information.

Answer: AC

Explanation:

The correct answer is **A. Implement Rapid PVST+ instead of STP** and **C. Use VTP to propagate VLAN information and to prune unused VLANs.** Let's break down why.

A. Rapid PVST+ vs. STP (802.1d): The original Spanning Tree Protocol (STP) 802.1d has slow convergence times after a topology change (link failure, switch failure). This can result in network disruptions. Rapid PVST+ (Per-VLAN Spanning Tree Plus), an enhanced version of STP, drastically reduces convergence times by employing faster timers and using mechanisms like proposal and agreement handshake. Implementing Rapid PVST+ addresses the requirement to reduce the impact of topology changes because it minimizes the network downtime.https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2sg/configuration/guide/config/stp.html

C. VLAN Trunking Protocol (VTP): VTP automates the management of VLANs across multiple switches.
Instead of manually configuring each switch with VLAN information, VTP allows you to create, modify, and delete VLANs on a VTP server switch, and those changes are propagated automatically to other VTP client switches. This reduces the time spent on network administration, and reduces the manual configuration errors since we will be maintaining only the server configuration. It also manages unused vlans on a trunk linkhttps://www.cisco.com/c/en/us/support/docs/lan-switching/vlan/10559-vtp-10559.html

Why B, D, and E are incorrect:

B. MST (Multiple Spanning Tree): MST addresses Spanning Tree inefficiencies within larger, complex environments. It is beneficial in situations with several VLANs per instance, while in this scenario the problem statement is concerned with STP convergence speed.

D. Broadcast and Multicast Storm Control: While beneficial for network stability, storm control does not directly address the concerns of spanning tree convergence time and manual configuration errors. It would be necessary after setting up the network and is not as crucial as optimizing stp and using vtp.

E. Dynamic Trunking Protocol (DTP): While DTP does automate trunking negotiation, it's primarily for creating trunks automatically and does not play a major role in VLAN propagation or reducing errors made

from manual configurations. It can be useful in optimizing trunks but does not address the points in the problem statement.

In summary, upgrading from STP to Rapid PVST+ will help with topology convergence and implementing VTP would reduce administration efforts and manual configuration errors. These improvements will address all the goals of this design.

Question: 66

Which component is part of the Cisco SD-Access overlay architecture?

A.border node
B.spine node
C.leaf node
D.Cisco DNA Center

Answer: A

Explanation:

The correct answer is **A. border node**.

Cisco SD-Access employs an overlay network architecture built upon a physical underlay. The overlay is where user traffic is logically segmented and transported. Key components within this overlay include border nodes, edge nodes (leaf nodes), and control plane elements. Border nodes are crucial for connecting the SD-Access fabric to external networks, including traditional enterprise networks, the internet, and other SD-Access fabrics. They act as gateways, translating traffic from the VXLAN-encapsulated overlay to routable IP packets on the underlay and beyond. This translation process allows for seamless communication between the SD-Access environment and the external world.

While spine and leaf nodes (options B and C) are fundamental to the underlay fabric's physical topology, they are not considered overlay components in the way a border node is. Spine nodes act as high-speed transit points within the fabric, and leaf nodes connect user devices to the network. Cisco DNA Center (option D) is the centralized management and automation platform for SD-Access, controlling the configuration and operation of the entire system, not a component of the overlay transport itself.

Therefore, among the listed options, only border nodes directly participate in the overlay fabric, acting as a critical interface for traffic leaving or entering the SD-Access environment. They are specifically designed to bridge the gap between the overlay's VXLAN encapsulation and external network protocols.

Authoritative Links for Further Research:

Cisco SD-Access Solution Design Guide:
https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html - This document provides comprehensive details on the architecture and components of Cisco SD-Access.

Cisco SD-Access Fabric Fundamentals:https://www.cisco.com/c/en/us/solutions/enterprise/software-defined-access/index.html - This page offers a high-level overview of SD-Access, including fabric components.

Cisco DevNet SD-Access Learning Labs:https://developer.cisco.com/site/devnet/learning/tracks/network-automation/sd-access.html - These labs provide hands-on experience with SD-Access and its components.

Question: 67

How are wireless endpoints registered in the HTDB in a Cisco SD-Access architecture?

A.Border nodes first register endpoints and then update the HTDB.
B.Fabric WLCs update the HTDB as new clients connect to the wireless network.
C.Fabric APs update the HTDB with the clients' EID and RLOC.
D.Fabric edge nodes update the HTDB based on CAPPWAP messaging from the AP.

Answer: B

Explanation:

The correct answer is B: Fabric WLCs update the HTDB as new clients connect to the wireless network. In Cisco SD-Access, the Host Tracking Database (HTDB) is a centralized repository for endpoint information, crucial for location and policy enforcement. Wireless endpoints, upon associating with a Fabric Access Point (AP) connected to the SD-Access fabric, do not directly register in the HTDB. Instead, the Fabric Wireless LAN Controller (WLC), which manages the APs, monitors client connections. When a new wireless client
associates, the Fabric WLC identifies the endpoint's EID (Endpoint Identifier) and its corresponding RLOC (Routing Locator, representing the location within the fabric), and this information is then communicated to the HTDB. This is accomplished via Control Plane communication between the WLC and the fabric control node(s). Fabric edge nodes receive information from the WLC indirectly and don't perform initial registration themselves. Border nodes are primarily for external network connectivity and do not manage internal client registrations. Therefore, Fabric WLC's role in dynamically updating the HTDB with wireless endpoint information as they connect is the accurate method within an SD-Access environment. This design ensures central visibility and consistent policy application for all wireless devices in the fabric.

Relevant Resources:

Cisco SD-Access Design Guide:https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html (Look for sections on Host Tracking and Wireless integration)
Cisco DNA Center Documentation:https://www.cisco.com/c/en/us/support/cloud-systems-management/dna-center/tsd-products-support-series-home.html (Explore sections related to endpoint management and wireless integration)

Question: 68

What is the purpose of a Cisco SD-Access underlay network?

A.to abstract IP-based connectivity from physical connectivity
B.to emulate LAN segments to transport Layer 2 frames over a Layer 3 network
C.to establish physical connectivity between switches and routers
D.to provide virtualization by encapsulating network traffic over IP tunnels

Answer: C

Explanation:

The correct answer is **C. to establish physical connectivity between switches and routers.**

In Cisco SD-Access, the underlay network serves as the foundational infrastructure providing the basic physical connectivity. It's the traditional Layer 3 IP network consisting of switches and routers interconnected via physical links. This underlay network is the "plumbing" upon which the SD-Access fabric is built. Its primary purpose is to ensure devices can communicate at the IP level, setting the stage for the more complex, overlay functionality of SD-Access. The underlay doesn't handle segmentation, policy enforcement, or virtual network creation; those roles belong to the SD-Access overlay. The underlay needs to be a stable, reliable IP

network providing reachability and transport.

Options A, B, and D describe aspects of the overlay network. Option A, "to abstract IP-based connectivity from physical connectivity," refers to the overlay's ability to create logical networks independent of the physical topology. Option B, "to emulate LAN segments to transport Layer 2 frames over a Layer 3 network," describes a common overlay function using technologies like VXLAN. Option D, "to provide virtualization by
encapsulating network traffic over IP tunnels," is also a function of the overlay using tunneling mechanisms.

In contrast, the underlay's purpose is foundational IP transport; it is the underlying Layer 3 network.
Therefore, C is the only option accurately describing the role of the underlay.

For further research, refer to Cisco's official documentation on SD-Access:

Cisco SD-Access Solution Overview:https://www.cisco.com/c/en/us/solutions/enterprise-networks/software-defined-access/index.html
Cisco DNA Center SD-Access Design Guide:
https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-dna-center-sda-design-guide.html

Question: 69

DRAG DROP -
Drag and drop the components in a Cisco SD-Access architecture from the left onto their descriptions on the right.
Select and Place:

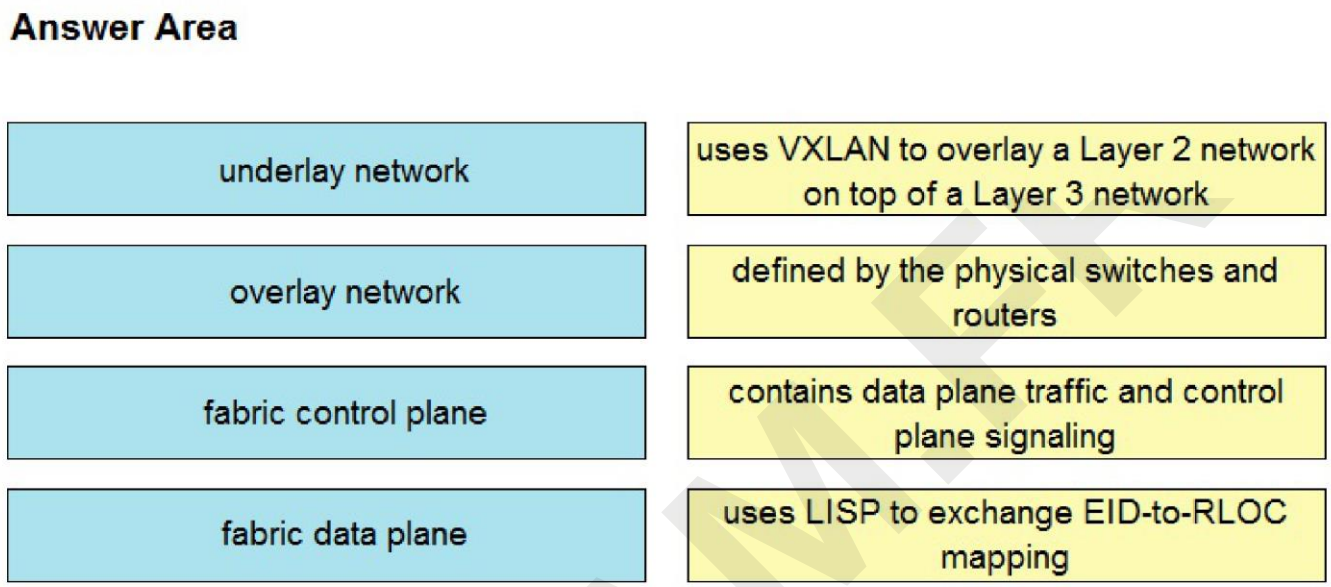

Answer:

Answer Area

overlay network

underlay network

fabric data plane

fabric control plane

Question: 70

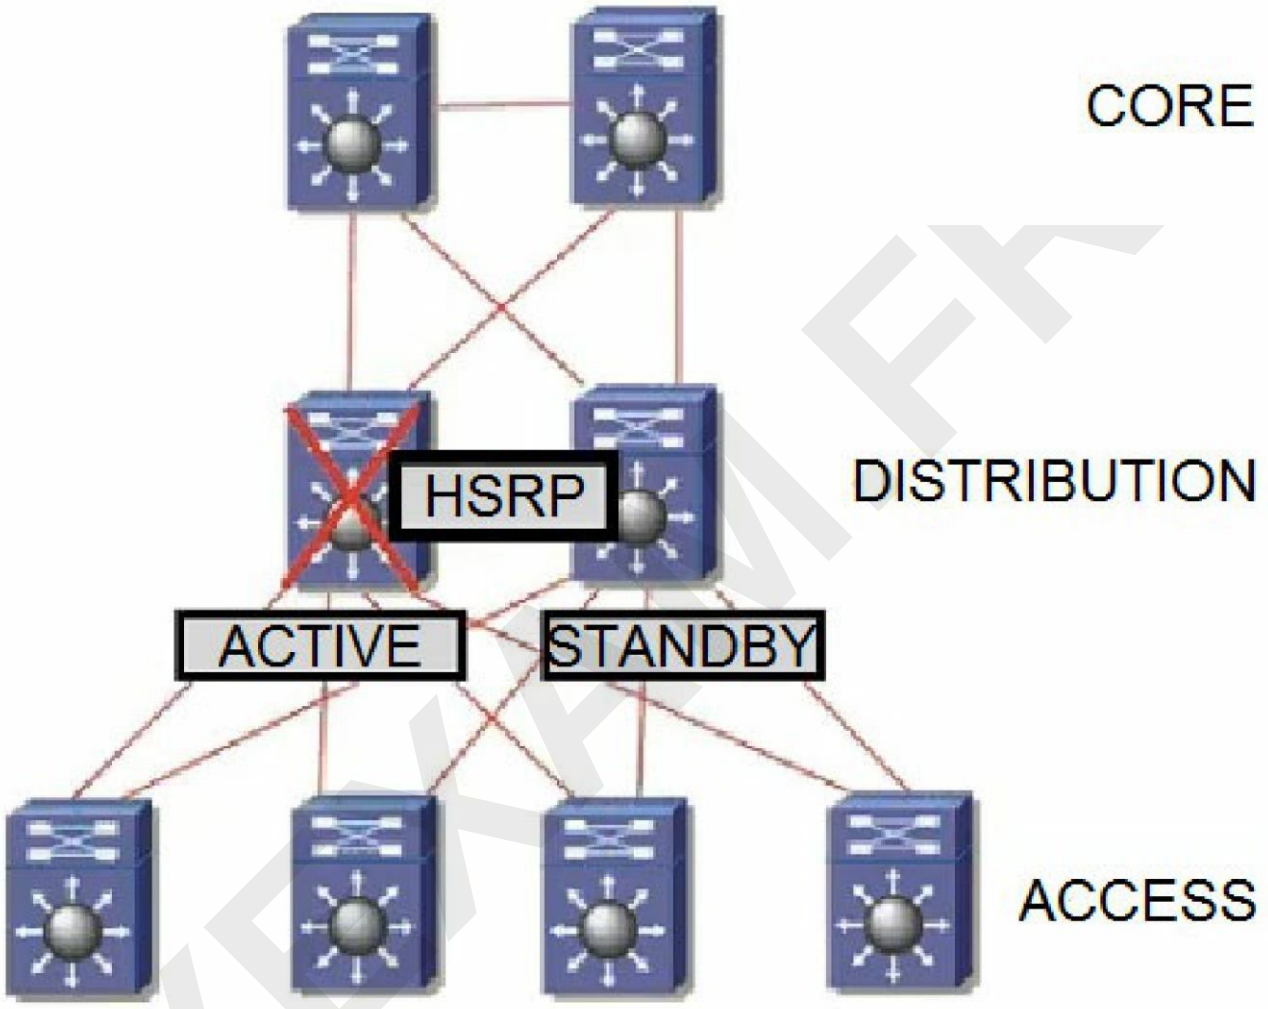

Refer to the exhibit. The distribution switches serve as the Layer 3 boundary. HSRP preemption is enabled. When the primary switch comes back after a failure, traffic is initially dropped. Which solution must be implemented to improve the design?

A. Increase the hello timers on both HSRP devices.

B. Use the preempt delay feature on the backup HSRP device.

C. Use the preempt delay feature on the primary HSRP device. D.
Configure a higher mac-refresh interval on both HSRP devices.

Answer: C

Explanation:

Use the preempt delay feature on the primary HSRP device.

Question: 71

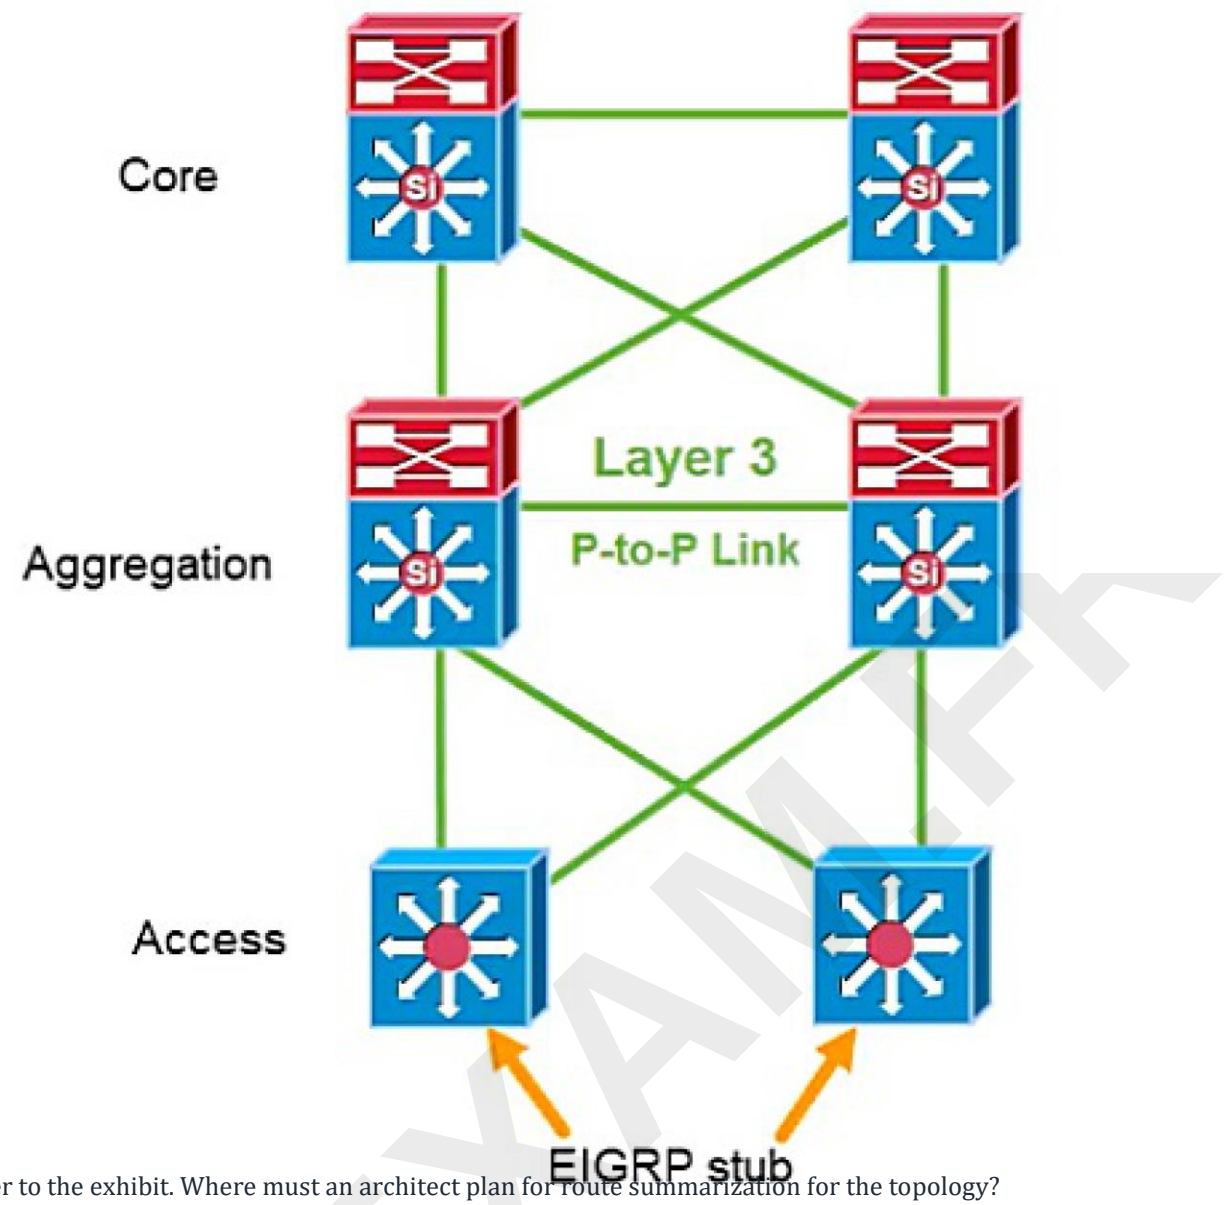

Refer to the exhibit. Where must an architect plan for route summarization for the topology?

A.from the core toward the aggregation and the access toward the aggregation B.from the core toward the aggregation and the aggregation toward the core C.from the aggregation toward the access and the access toward the aggregation D.from the aggregation toward the core and the aggregation toward the access

Answer: D

Explanation:

from the aggregation toward the core and the aggregation toward the access.

Question: 72

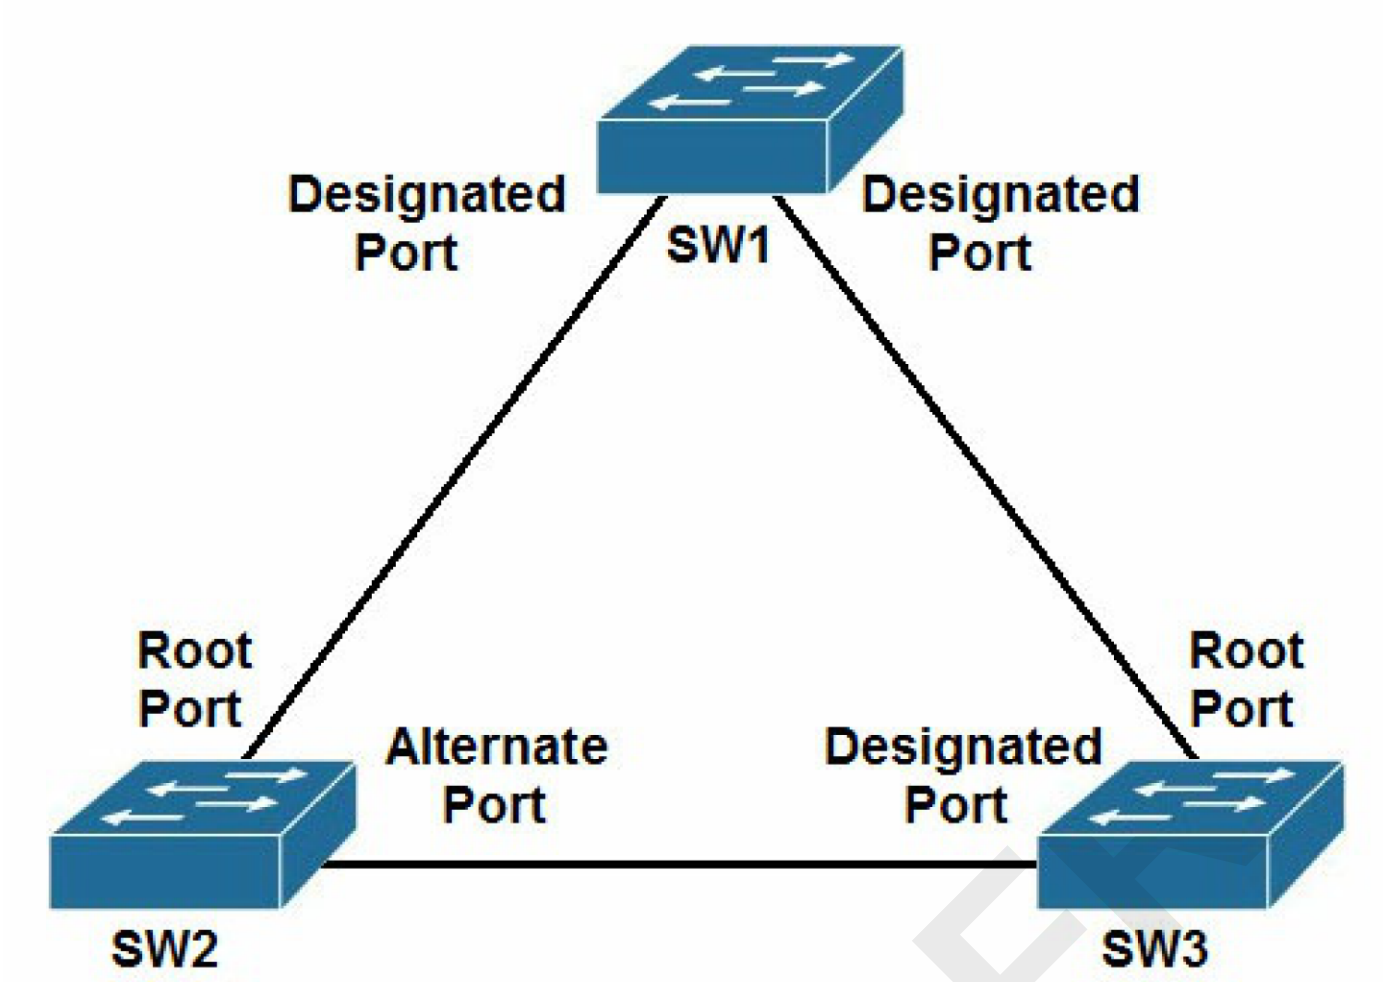

Refer to the exhibit. The connection between SW2 and SW3 is fiber and occasionally experiences unidirectional link failure. An architect must optimize the network to reduce the change of Layer 2 forwarding loops when the link fails. Which solution should the architect include?

A. Utilize BPDU filter on SW3.
B. Utilize root guard on SW1.
C. Utilize BPDU guard on SW1.
D. Utilize loop guard on SW2.

Answer: D

Explanation:

Utilize loop guard on SW2.

Question: 73

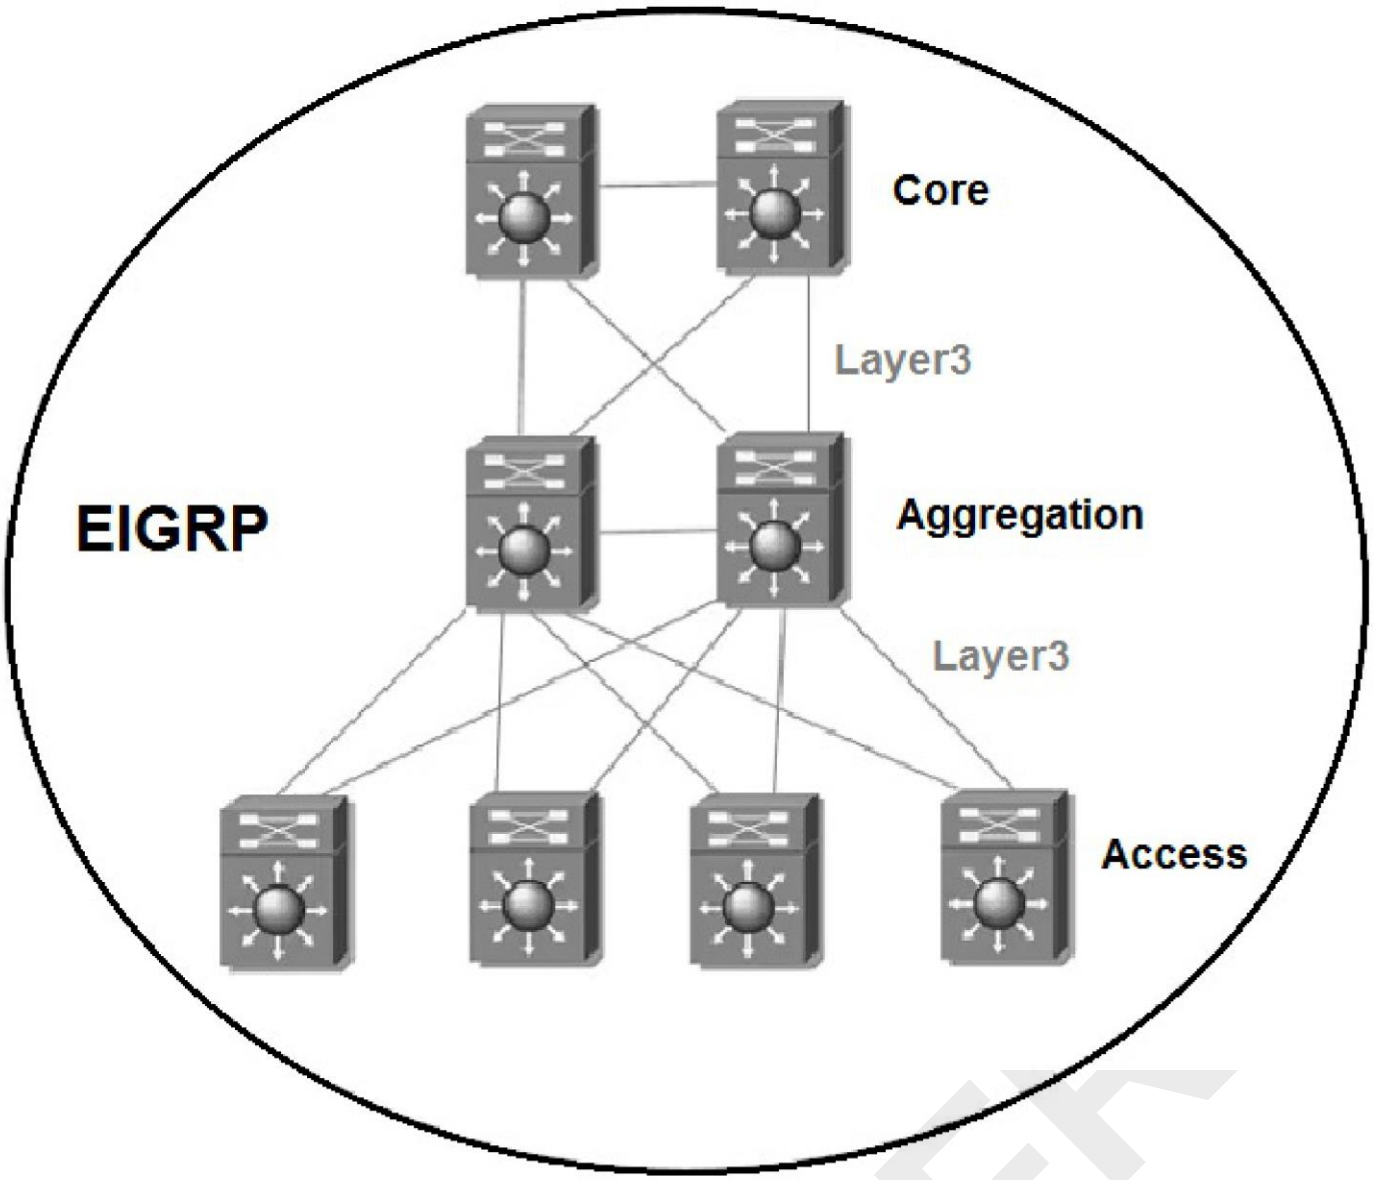

Refer to the exhibit. The full EIGRP routing table is advertised throughout the network. Currently, users experience data loss when any one link in the network fails. An architect must optimize the network to reduce the impact when a link fails. Which solution should the architect include in the design?

A.Run BFD on the inter links between EIGRP neighbors.

B.Summarize the access layer networks from each access layer switch toward the aggregation layer. C.Reduce the default EIGRP hello interval and hold time.

D.Summarize the access layer networks from the aggregation layer toward the core layer.

Answer: D

Explanation:

Summarize the access layer networks from the aggregation layer toward the core layer.

Question: 74

What is the purpose of a control plane node in a Cisco SD-Access network fabric?

A.to maintain the endpoint database and mapping between endpoints and edge nodes
B.to detect endpoints in the fabric and inform the host tracking database of EID-to-fabric-edge node bindings C.to identify and authenticate endpoints within the network fabric
D.to act as the network gateway between the network fabric and outside networks

Answer: A

Explanation:

The correct answer is **A. to maintain the endpoint database and mapping between endpoints and edge nodes.** Here's why:

In a Cisco SD-Access fabric, the control plane node (often the Cisco DNA Center acting as a control plane) is responsible for centralizing policy and managing endpoint-to-location mappings. Specifically, it maintains the **endpoint database**, which tracks all connected devices and their associated location within the fabric. This database is critical for the fabric's operation. This mapping, which is an **EID (Endpoint Identifier) to RLOC (Routing Locator)** association, essentially ties a user's identity to their physical location within the network.

The EID is the endpoint's logical address (like an IP), while the RLOC represents the physical edge node it's connected to. This separation of identity and location is fundamental to SD-Access. The control plane node does not directly detect endpoints (that is handled by the edge nodes) or act as the direct gateway to the external network (that is the border node's responsibility). It also doesn't handle initial authentication directly, but it influences that process based on user profiling defined within its policies. Instead, it acts as the brain of the network, distributing policies and managing the logical addressing scheme across the fabric. This control ensures consistent policy enforcement across all locations.

Here are some authoritative links for further research:

Cisco SD-Access Solution Overview:https://www.cisco.com/c/en/us/solutions/enterprise-networks/software-defined-access/index.html
Cisco DNA Center documentation: Search for “Endpoint Database”, “Control Plane” in relevant documentation pages on Cisco's website.

SD-Access Deep Dive: Review videos or articles that specifically discuss the role of the control plane within the SD-Access architecture.

These resources will provide a deeper understanding of the SD-Access architecture and the role of the control plane.

Question: 75

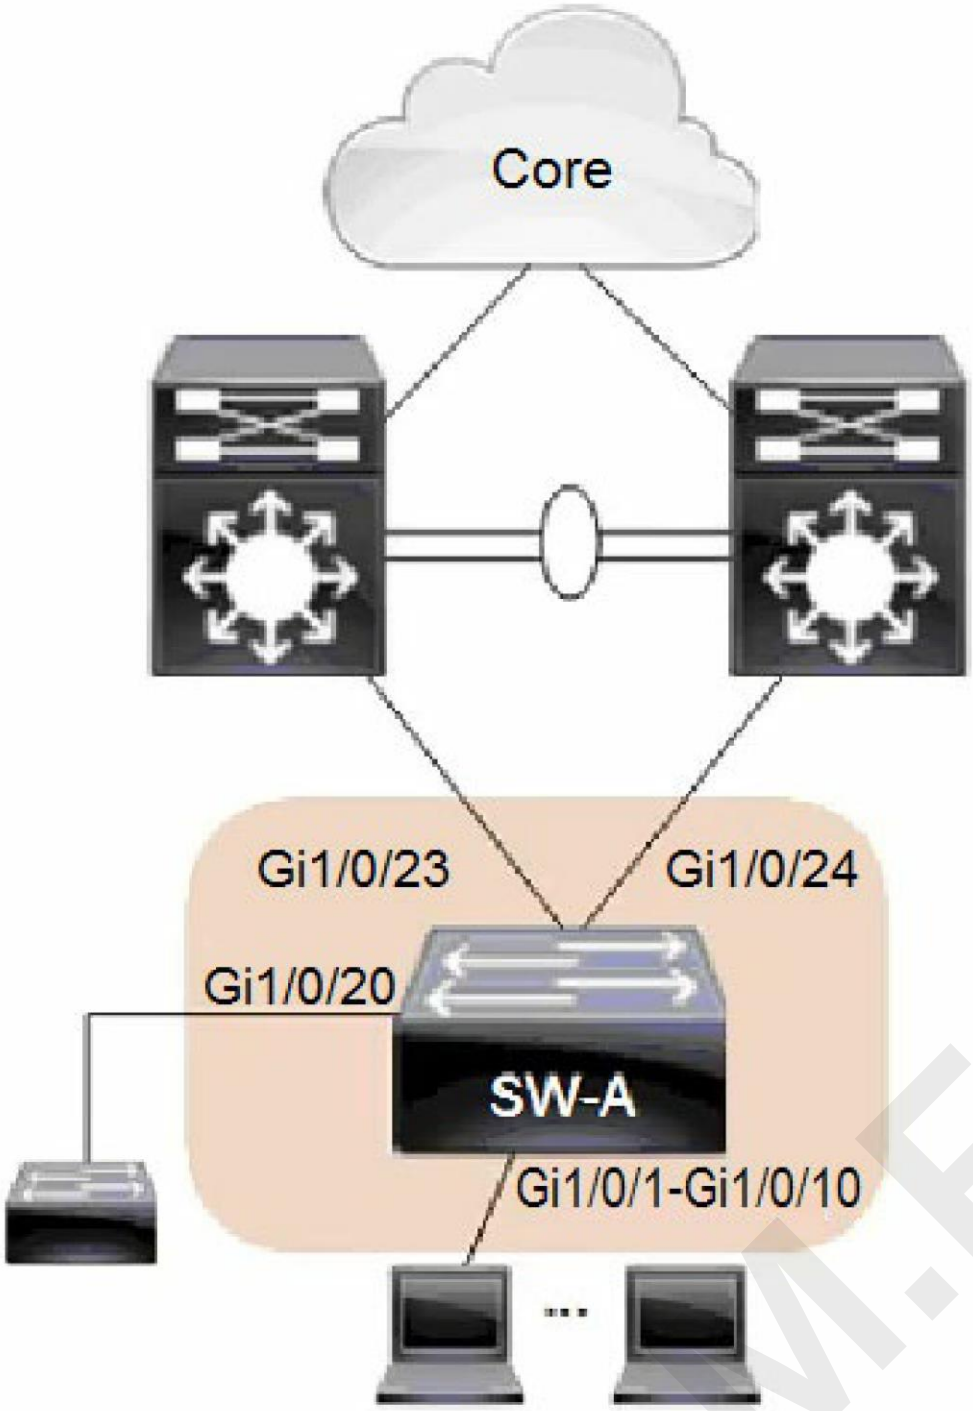

Refer to the exhibit. An architect reviews the low-level design of a company's enterprise network and advises optimizing the STP convergence time. Which functionality must be applied to Gi1/0/1-10 to follow the architect's recommendation?

A.UplinkFast
B.root guard
C.BPDU guard
D.PortFast

Answer: D

Explanation:

Correct answer is D:PortFast.

Question: 76

An engineer must design a large Layer 2 domain that contains hundreds of switches and VLANs. The engineer's primary goals are to:

- Efficiently utilize the bandwidth of all links
- Avoid Layer 2 loops
- Cause minimal impact on switch CPU and memory

Which technology should the engineer include in the design?

A.MST
B.Rapid PVST+
C.RSTP
D.PVST+

Answer: A

Explanation:

The correct answer is A. MST (Multiple Spanning Tree) is the most suitable technology for this scenario due to its ability to address the engineer's stated goals. MST allows for multiple spanning-tree instances, each with its own root bridge and forwarding topology. This allows traffic for different VLANs to be load-balanced across multiple physical links, which efficiently utilizes bandwidth. Importantly, MST avoids Layer 2 loops by enforcing a tree-like topology within each instance. Unlike PVST+ or Rapid PVST+, which create a separate instance per VLAN, MST groups VLANs into instances. This results in fewer spanning-tree processes running on switches, minimizing CPU and memory impact, especially crucial in large Layer 2 domains. Furthermore, MST provides superior scalability compared to per-VLAN STP variants. It is the standard IEEE 802.1s and is widely adopted, making it highly reliable and manageable. This allows for better resource utilization and minimizes network complexity, a key concern when dealing with hundreds of switches and VLANs.

Here are some links for further research:

Cisco Documentation on MST:https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2sg/configuration/guide/config/mst.html
IEEE 802.1s Standard:https://standards.ieee.org/standard/802_1s-2002.html
Understanding Multiple Spanning Tree Protocol:https://www.geeksforgeeks.org/understanding-multiple-spanning-tree-protocol-mstp/

Question: 77

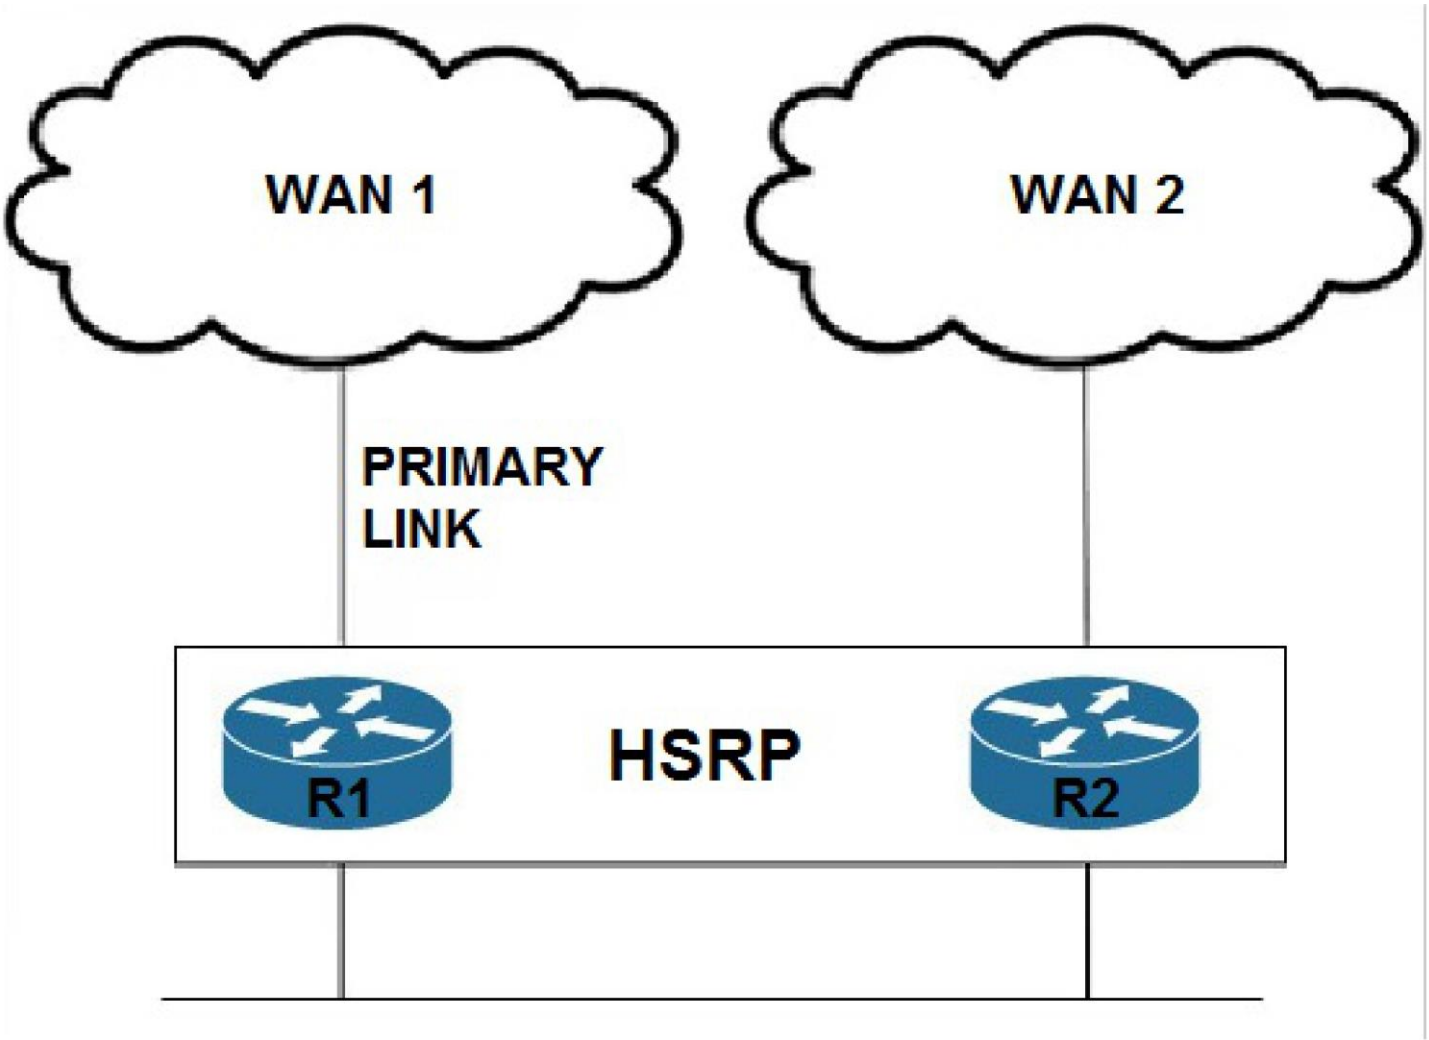

Refer to the exhibit. An engineer must design an automatic failover solution. The solution should allow HSRP to detect a WAN 1 failure and initiate an automatic failover, making router R2 the active HSRP router. Which two solutions should the engineer choose? (Choose two.)

A. implement IP SLA on router R1
B. implement PBR on router R1
C. implement Enhanced Object Tracking on router R1
D. use IP source routing
E. use a floating static route

Answer: AC

Explanation:

A. implement IP SLA on router R1.

C. implement Enhanced Object Tracking on router R1.

Question: 78

Which topology within a network underlay eliminates the need for first hop redundancy protocols while improving fault tolerance, increasing resiliency, and simplifying the network?

A. virtualized topology
B. routed access topology
C. Layer 2 topology
D. logical fabric topology

Answer: B

Explanation:

The correct answer is **B. routed access topology**. Here's why:

A routed access topology eliminates the need for traditional First Hop Redundancy Protocols (FHRPs) like HSRP, VRRP, or GLBP. In this design, each access switch acts as a Layer 3 device and performs routing, typically using a routing protocol like OSPF or BGP, directly to the distribution or core layer. This approach offers significant advantages in terms of fault tolerance and network resiliency. If a link or switch fails, routing protocols automatically converge to an alternative path, minimizing disruptions. This distributed routing model also simplifies the network by removing reliance on single-point-of-failure FHRP gateways. It improves network convergence, as routing protocols can react more quickly to failures compared to the timers in FHRPs. Furthermore, it allows for more efficient utilization of network resources, as all available paths are used simultaneously for forwarding traffic, whereas FHRPs typically only use one path actively. This topology aligns well with modern network best practices, promoting a more scalable and robust network infrastructure.

For further reading, explore these resources:

Cisco Design Zone:https://www.cisco.com/c/en/us/solutions/enterprise/design-zone/index.html - Explore design guides and best practices.

Cisco Validated Designs (CVD):https://www.cisco.com/c/en/us/solutions/enterprise/design-zone/index.html -Find detailed architectural documentation.

OSPF Configuration Guide:https://www.cisco.com/c/en/us/support/ios-xml/ios-xml-routing/configuration/15-mt/ospf-multiarea-support.html - Understand routing protocol fundamentals.

BGP Configuration Guide:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/15-mt/irg-15-mt-book/irg-basic-bgp.html - Learn about advanced routing options.

Question: 79

Refer to the exhibit.

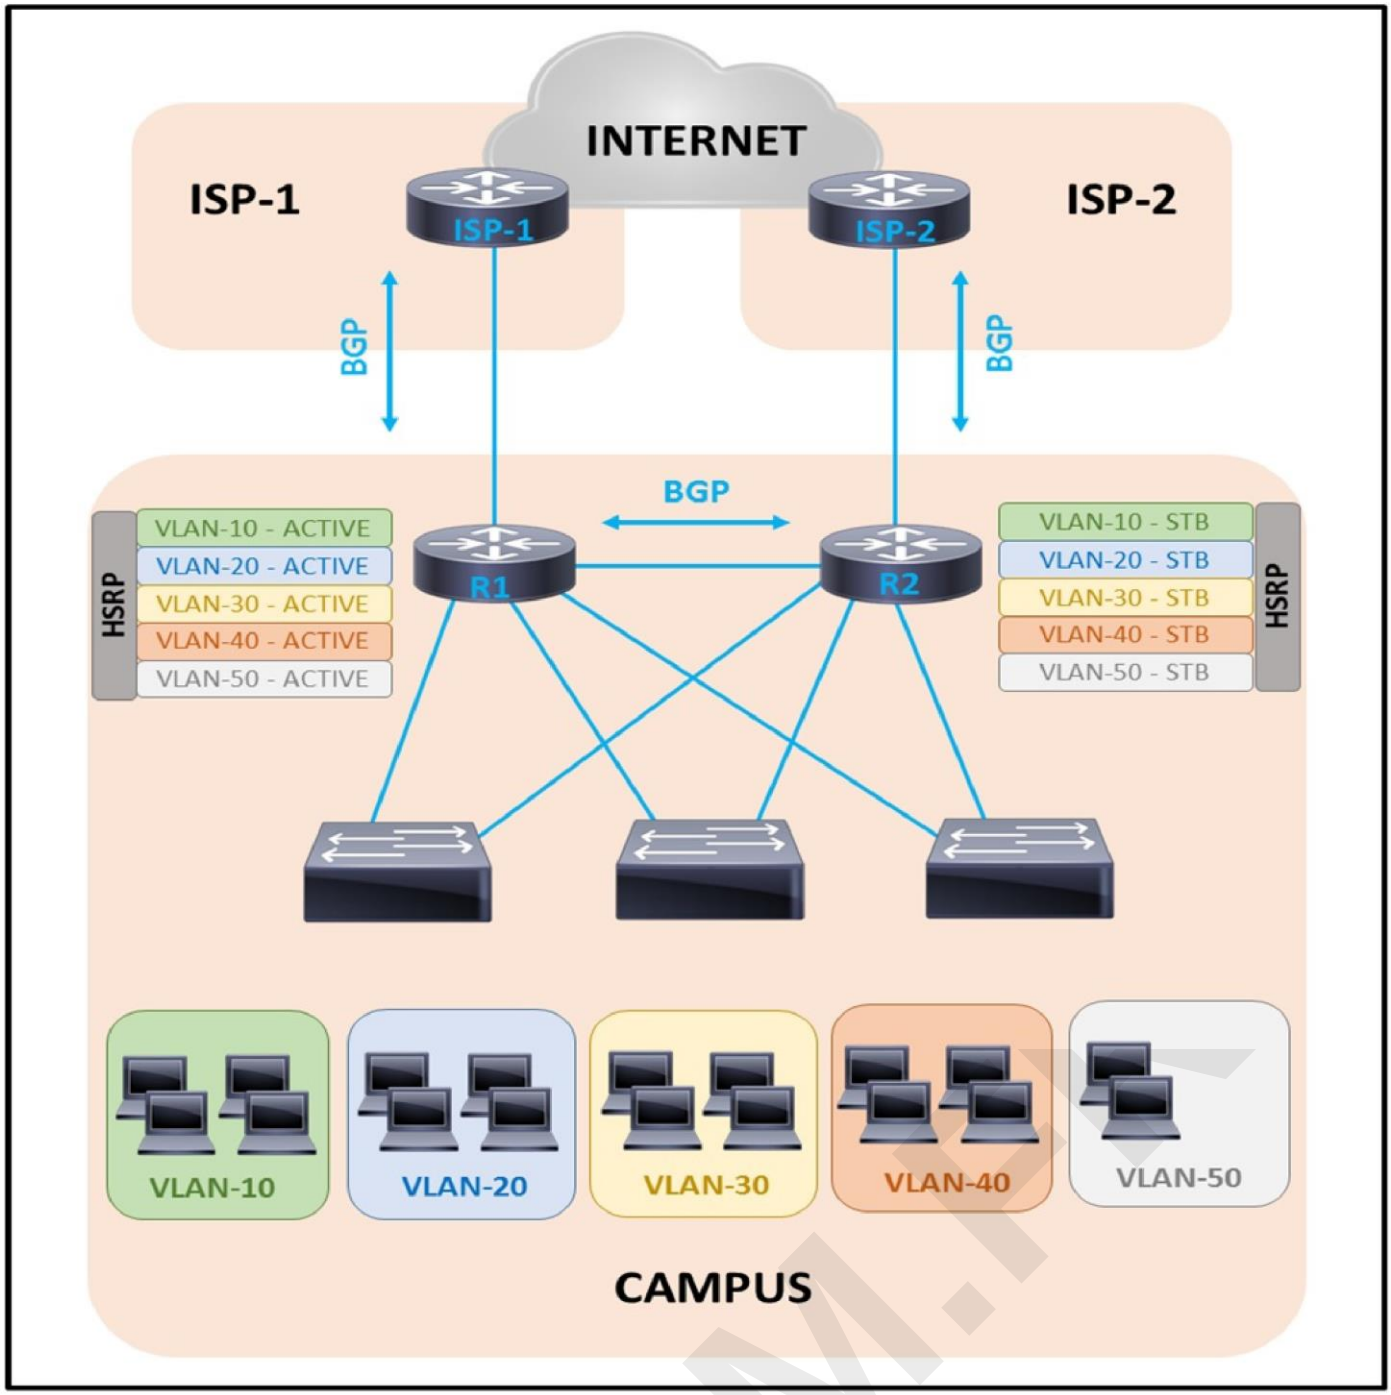

A customer is running HSRP on the core routers. Over time the company has grown and requires more network capacity. In the current environment, some of the downstream interfaces are almost fully utilized, but others are not. Which solution improves the situation?

A. Make router R2 active for half of the VLANs.
B. Add more interfaces to R1 and R2.
C. Configure port channel toward downstream switches.
D. Enable RSTP on the downstream switches.

Answer: A

Explanation:

Make router R2 active for half of the VLANs.

Question: 80

An architect must develop a campus network solution that includes:➪ logically segmented and isolated networks
➪ ability to communicate between network segments when required➪ support for overlapping IP addresses
➪ widely available technologies to avoid purchasing specialized equipment
Which solution must the architect select?

A. VSS with IGP
B. 802.1Q with HSRP
C. vPC with HSRP
D. VRF-Lite with OSPF

Answer: D

Explanation:

The correct answer is D, VRF-Lite with OSPF. Here's why:

VRF-Lite (Virtual Routing and Forwarding Lite) fulfills the requirements for logical segmentation and isolation. Each VRF acts as a separate routing table, isolating traffic within that VRF. When communication is needed between VRFs, it can be achieved through controlled routing or dedicated interfaces. This addresses the need for both isolation and inter-segment communication.

VRF-Lite also supports overlapping IP addresses. Because each VRF has its own routing table, the same IP address space can be used in multiple VRFs without conflicts. This addresses the support for overlapping IPs requirement.

VRF-Lite and OSPF are widely available technologies, being core components of most enterprise-grade network equipment. This avoids the need for specialized hardware and ensures the solution is easily implementable.

Options A, B, and C do not completely address all the requirements as effectively:

VSS (Virtual Switching System) primarily provides high availability and simplified management rather than network segmentation like VRFs.
802.1Q with HSRP (Hot Standby Router Protocol) focuses on VLAN tagging and gateway redundancy but does not inherently provide the required IP overlapping and logical isolation, especially not to the same level as VRFs.
vPC (virtual PortChannel) provides link aggregation and redundancy at Layer 2, not the logical isolation at Layer 3 with address overlapping that the prompt asks for.

Therefore, VRF-Lite with OSPF is the most suitable solution for creating logically segmented, isolated networks with controlled inter-segment communication while supporting overlapping IP addresses and using readily available technology.

Authoritative links for further research:

Cisco documentation on VRF-Lite: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_vrf/configuration/15-sy/ipv-vrf-15-sy-book/ipv-vrf-lite.html
Cisco documentation on OSPF: https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html