Cisco

(400-007)

Cisco Certified Design Expert v3.0

Total: **275 Questions**

Link:

Question: 1

Company XYZ is planning to deploy primary and secondary (disaster recovery) data center sites. Each of these sites will have redundant SAN fabrics and data protection is expected between the data center sites. The sites are 100 miles (160 km) apart and target RPO/RTO are 3 hrs and 24 hrs, respectively. Which two considerations must Company XYZ bear in mind when deploying replication in their scenario? (Choose two.)

A. Target RPO/RTO requirements cannot be met due to the one-way delay introduced by the distance between sites.
B. VSANs must be extended from the primary to the secondary site to improve performance and availability.
C. VSANs must be routed between sites to isolate fault domains and increase overall availability.
D. Synchronous data replication must be used to meet the business requirements.
E. Asynchronous data replication should be used in this scenario to avoid performance impact in the primary site.

Answer: CE

Explanation:

Here's a justification for why options C and E are the correct choices for the given scenario:

Option C: VSANs must be routed between sites to isolate fault domains and increase overall availability.

Fault Domain Isolation: Routing VSANs (Virtual SANs) between the primary and secondary data centers creates distinct fault domains. This means that a failure in one data center's VSAN fabric won't directly propagate to the other, preventing a catastrophic, site-wide outage. By keeping these fabrics separate and independent, each site's storage resources are isolated, reducing the risk of a single failure impacting both environments.

Increased Availability: By having separate VSANs linked by routing, the secondary site's infrastructure is available to assume operations, if the primary site fails. This provides resilience by ensuring that data and storage resources remain available even during disasters or severe disruptions at either location.

Avoidance of Stretching: Stretching VSANs across geographically dispersed locations, (as option B suggests) creates a single failure domain. All elements of the fabric are interlinked and impacted by latency. This increases the risk of issues on one side impacting the other side and is not recommended with distances like 160km.

Option E: Asynchronous data replication should be used in this scenario to avoid performance impact in the primary site.

Latency Considerations: With a 160 km distance between sites, there's significant latency introduced in data transfer. Synchronous replication (as proposed in D), which requires confirmation that writes are complete at both sites before acknowledging the application write, would severely degrade primary site performance. Each write operation would incur the latency of the network roundtrip.

Asynchronous Replication Suitability: Asynchronous replication allows writes to be committed to the primary site's storage without waiting for confirmation from the secondary. This minimizes the performance impact on the primary site. Changes are replicated to the secondary site as they become available, utilizing bandwidth efficiently, this better suits long distance replication.

RPO Alignment: While asynchronous replication introduces the possibility of some data loss, the 3-hour Recovery Point Objective (RPO) specified in the problem statement allows for asynchronous replication. Asynchronous replication often can achieve this RPO when carefully configured and managed.

Authoritative Link for Replication:https://www.vmware.com/topics/glossary/content/replication

Why Options A, B, and D are Incorrect:

A is incorrect: While distance introduces latency, the given RPO/RTO targets are achievable with appropriately configured asynchronous replication.

B is incorrect: Extending VSANs between sites would create a single point of failure, and with significant distance and latency impacting performance and availability, the approach is not recommended.

D is incorrect: Synchronous replication is not suitable for this scenario as the distance creates high latency, impacting application performance. It would also require extremely low latency links which are extremely expensive to operate at distances of 160km.

In conclusion, option C and E are the most appropriate considerations for the scenario, aligning with best practices for disaster recovery and site resilience.

Question: 2

An architect receives a business requirement from a CTO that states the RTO and RPO for a new system should be as close as possible to zero. Which replication method and data center technology should be used?

A. synchronous replication over geographically dispersed dual data centers via MPLS
B. synchronous replication over dual data centers via Metro Ethernet
C. asynchronous replication over geographically dispersed dual data centers via CWDM
D. asynchronous replication over dual data centers via DWDM

Answer: A

Explanation:

The correct answer is **B. synchronous replication over dual data centers via Metro Ethernet.**

Here's why:

The requirement for near-zero RTO (Recovery Time Objective) and RPO (Recovery Point Objective) demands a replication method that ensures minimal data loss and rapid recovery. **Synchronous replication** writes data to both the primary and secondary locations simultaneously. This ensures that if the primary site fails, the secondary site already holds the most current data, minimizing RPO (data loss). Source: VMware, Understanding Data Replication for Disaster Recovery and Business Continuity

Asynchronous replication introduces latency as data is written to the secondary location after being written to the primary location. This leads to a higher RPO, contradicting the requirement for minimal data loss.

Geographically dispersed data centers (options A & C) introduce significant latency due to the physical distance, which renders synchronous replication ineffective. Synchronous replication requires low-latency, high-bandwidth connections, typically found within a metropolitan area. Source: Dell Technologies, Storage replication: synchronous vs. asynchronous

Metro Ethernet provides the low-latency, high-bandwidth connections needed to support synchronous replication between dual data centers in close proximity. Options C & D involving CWDM and DWDM are more appropriate for longer distances and are typically paired with asynchronous replication.

Therefore, synchronous replication over dual data centers using Metro Ethernet provides the best approach to achieving near-zero RTO and RPO by minimizing data loss and providing a readily available, up-to-date secondary site. Options A,C and D cannot guarantee data integrity and low latency making it an unacceptable choice.

Question: 3

What are two primary design constraints when a robust infrastructure solution is created? (Choose two.)

A. component availability
B. monitoring capabilities
C. project time frame
D. staff experience
E. total cost

Answer: CE

Explanation:

The correct answer identifies **project time frame (C)** and **total cost (E)** as primary design constraints for robust infrastructure solutions. These are fundamental limitations that directly dictate the scope and feasibility of any project. The time frame dictates how quickly the solution needs to be implemented, influencing technology choices and resource allocation. A tight schedule may require simpler, more readily available options, potentially sacrificing complexity or advanced features. Simultaneously, total cost is a critical constraint, impacting everything from hardware and software selection to personnel and operational expenses. Budget limitations can force compromises in redundancy, scalability, or performance, leading to trade-offs between robustness and affordability.

Component availability (A), while important for implementation, is a secondary concern, as it's a factor considered within the cost and time frame limitations. Monitoring capabilities (B) and staff experience (D), while crucial for long-term maintenance and effectiveness, are not primary constraints that initially define the boundaries of the design. These become important once a solution has been designed and deployed. The project's initial design must accommodate the schedule and allocated budget. For instance, if the time to market is critical, a fast deployment method with higher costs may need to be selected, or vice versa.

Similarly, a budget can dictate the technologies or vendors selected. Hence the primary constraints are the overarching limits within which infrastructure is architected.

Authoritative Links for Further Research:

NIST Special Publication 800-145 (The NIST Definition of Cloud Computing): https://csrc.nist.gov/publications/detail/sp/800-145/final - Though not directly focused on design constraints, it provides a foundational understanding of cloud computing and its complexities, highlighting trade-offs often involved.

Cloud Design Patterns:https://docs.microsoft.com/en-us/azure/architecture/patterns/ - Microsoft's documentation on cloud design patterns often highlights the influence of cost and time on design choices. **AWS Well-Architected Framework:**https://aws.amazon.com/architecture/well-architected/ - While focusing on best practices, the framework demonstrates how cost optimization and performance impact infrastructure design.

Question: 4

Which network management framework can be used to develop a network architecture that contains business requirements analysis, gap analysis, and network diagrams as artifacts to be used for design and implementation later?

A. FCAPS
B. Cobit
C. TOGAF
D. ITIL

Answer: C

Explanation:

The correct answer is **C. TOGAF (The Open Group Architecture Framework)** because it provides a comprehensive methodology for developing and managing enterprise architecture, which directly aligns with the requirements described in the question. TOGAF's Architecture Development Method (ADM) specifically includes phases for business requirements analysis, gap analysis, and documenting the architecture through diagrams. The ADM's preliminary phase defines the architecture vision based on business needs. The subsequent architecture vision phase develops a high-level view of the desired architecture, informed by the identified business requirements and gaps in the existing infrastructure. Phases like 'Architecture Definition' and 'Technology Architecture' explicitly focus on detailed planning, logical and physical architecture representation including network diagrams. These defined architectures can be reused and implemented.

Unlike FCAPS, which is geared towards fault management, configuration, accounting, performance, and security; or COBIT, which focuses on IT governance and management objectives; or ITIL, which provides a framework for IT service management, TOGAF is explicitly a methodology for enterprise architecture. TOGAF produces architecture artifacts, including network diagrams, along with analysis documents, all directly applicable to the network design and implementation process. Therefore, TOGAF is the suitable framework for the given scenario.

TOGAF Official WebsiteTOGAF Architecture Development Method (ADM) Overview

Question: 5

Which two types of planning approaches are used to develop business-driven network designs and to facilitate the design decisions? (Choose two.)

A. strategic planning approach
B. business optimization approach
C. tactical planning approach
D. modular approach
E. cost optimization approach

Answer: AC

Explanation:

The correct answer is A (strategic planning approach) and C (tactical planning approach) because these are the fundamental approaches used to align network design with overarching business goals.

Strategic planning is a high-level, long-term approach that focuses on defining the overall vision and direction for the organization. In the context of network design, this involves understanding the business's long-term goals, such as expansion, new service offerings, or increased efficiency. Strategic planning helps determine the key functionalities and capabilities the network must possess to support these goals, influencing high-level decisions regarding technology adoption and resource allocation.

Tactical planning, conversely, is a shorter-term, more operational approach that translates the strategic vision into specific, actionable plans. In network design, this involves detailing how the high-level strategic requirements will be implemented. This includes decisions on specific network technologies, protocols, equipment selection, and the design's physical layout. Tactical planning ensures that the network design is achievable, aligns with budget constraints, and addresses immediate business needs.

Neither the 'business optimization approach' nor the 'cost optimization approach' are specific planning methodologies, rather they are considerations that might arise during the strategic or tactical planning process. The 'modular approach' is a design principle or technique used during implementation. Therefore,

strategic and tactical planning approaches, working in tandem, provide the necessary frameworks to create business-driven network designs.

For further information on strategic and tactical planning, refer to resources like the following:

CIO.com: https://www.cio.com/article/229790/strategic-vs-tactical-planning-whats-the-difference.html
Business.org:https://www.business.org/strategy/planning/strategic-planning-vs-tactical-planning/These links provide more detailed explanations of the concepts in a general business context, which directly apply to network design in business environments.

Question: 6

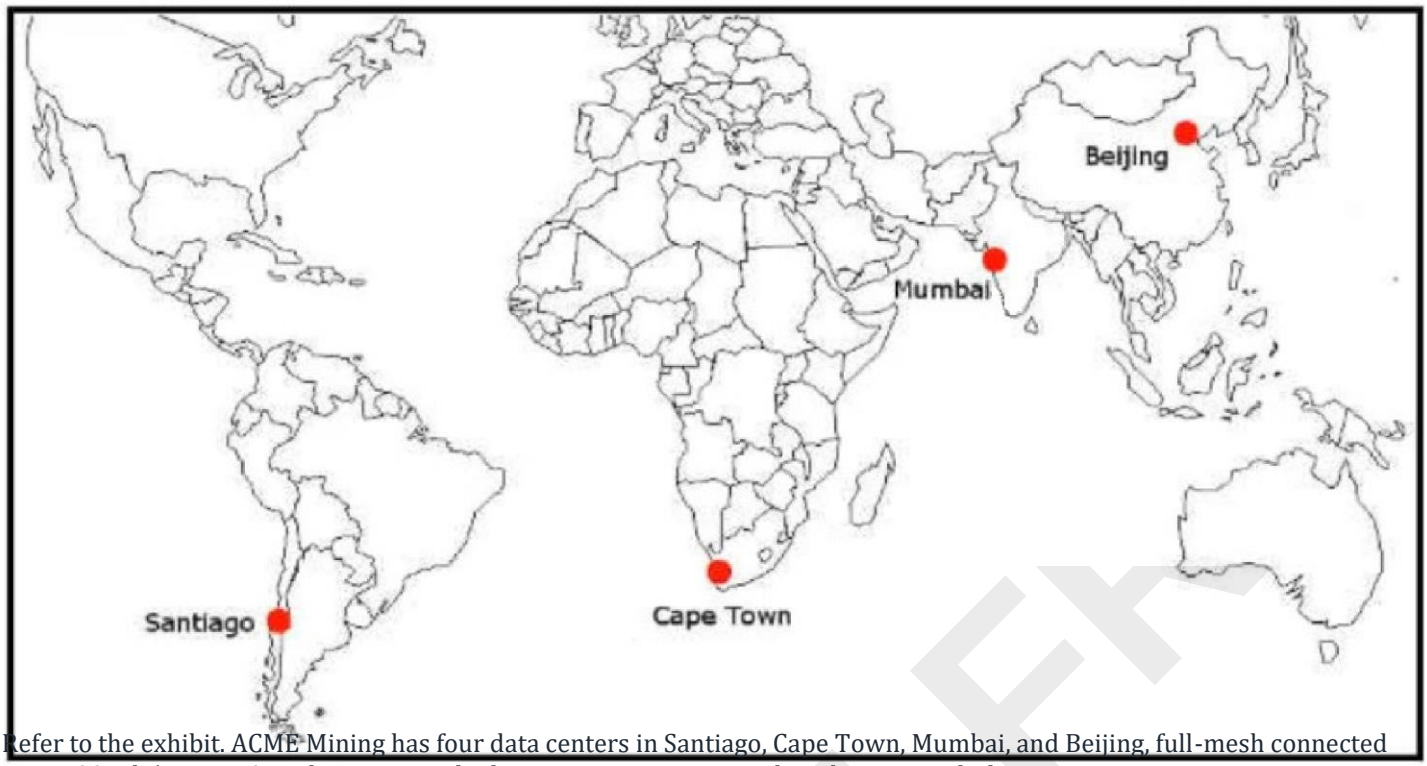

Refer to the exhibit. ACME Mining has four data centers in Santiago, Cape Town, Mumbai, and Beijing, full-mesh connected via a 400 Mb/s EVP-LAN. They want to deploy a new mission-critical application with these requirements:
* cluster heartbeat 2 MB/s continuous (250 KB/s)
* cluster heartbeat one-way maximum latency 100 ms
These are the current ping tests results between the four data centers:

	Santiago	Cape Town	Mumbai	Beijing
Santiago	-	280 ms	378 ms	409 ms
Cape Town	280 ms	-	185 ms	445 ms
Mumbai	383 ms	176 ms	-	443 ms
Beijing	430 ms	448 ms	442 ms	-

Which hosting data center pair can host the new application?

A. Mumbai and Beijing
B. Cape Town and Mumbai
C. Cape Town and Beijing
D. Santiago and Mumbai

E. Santiago and Beijing
F. Santiago and Cape Town

Answer: B

Explanation:

Cape Town and Mumbai is a correct answer.

Question: 7

CONNECTIVITY	CAPEX	OPEX ANNUAL	INSTALLATION FEE	TERM
DWDM over dark fiber	$250,000	$100,000	$30,000	12 months
CWDM over dark fiber	$150,000	$100,000	$25,000	18 months
MPLS wires only	$50,000	$80,000	$5,000	24 months
Metro Ethernet	$45,000	$100,000	$5,000	36 months

Refer to the table. A customer investigates connectivity options for a DCI between two production data centers to aid a large-scale migration project. The solution must provide a single 10G connection between locations and be able to run its own varying QoS profiles without service provider interaction based on the migration stages. All connectivity methods are at 10 Gbps. Which transport technology costs the least if the connectivity is required for just one year?

A. DWDM over dark fiber
B. Metro Ethernet
C. MPLS wires only
D. CWDM over dark fiber

Answer: D

Explanation:

Note the requirement, "be able to run its own varying QoS profiles without service provider interaction based on the migration stages".Totals for each solution.dwdm 380000cwdm 325000mpls 215000metro-e 350000At first glance MPLS is best, however QOS will need SP integration. Thus CWDM (D).

Question: 8

CONNECTIVITY	CAPEX	OPEX ANNUAL	INSTALLATION FEE	TERM
DWDM over dark fiber	$250,000	$100,000	$30,000	60 months
CWDM over dark fiber	$150,000	$100,000	$25,000	60 months
MPLS wires only	$50,000	$150,000	$75,000	12 months
Metro Ethernet	$45,000	$125,000	$5,000	36 months

Refer to the table. A customer investigates connectivity options for a DCI between two production data centers. The solution must provide dual 10G connections between locations with no single points of failure for Day-1 operations. It must also include an option to scale for up to 20 resilient connections in the second year to accommodate isolated SAN over IP and isolated dedicated replication IP circuits. All connectivity methods are

duplex 10 Gbps. Which transport technology costs the least over two years in this scenario?

A. CWDM
B. DWDM
C. MPLS
D. Metro Ethernet

Answer: B

Explanation:

the hint is the requirement 'option to scale for up to 20 resilient connections in the second year'. this effectively multiplies the estimated cost with 20 times.DWDM will leverage multiple lambdas across 2 x fibres to provide that scale without additional cost for each connection added. effectively scaling out is amortised into the original contract

CWDM supports up to 18 wavelength channels transmitted through a fiber at the same time. To achieve this, the different wavelengths of each channel are 20nm apart. DWDM, supports up to 80 simultaneous wavelength channels, with each of the channels only 0.8nm apart. CWDM technology offers a convenient and cost-efficient solution for shorter distances of up to 70 kilometers. For distances between 40 and 70 kilometers, CWDM tends to be limited to supporting eight channels.

Question: 9

What are two examples of business goals to be considered when a network design is built? (Choose two.)

A. integrate endpoint posture
B. ensure faster obsolescence
C. minimize operational costs
D. reduce complexity
E. standardize resiliency

Answer: CD

Explanation:

The correct answers are C (minimize operational costs) and D (reduce complexity). Business goals are overarching objectives that drive organizational strategy and decision-making. Network design, being a critical part of business infrastructure, must align with these goals. Minimizing operational costs (C) directly impacts the bottom line by reducing expenditures related to network maintenance, power consumption, staffing, and upgrades. A well-designed network should aim for efficiency and cost-effectiveness. Reducing complexity (D) simplifies management, troubleshooting, and scalability. A complex network can be costly to operate and maintain due to increased chances of errors, higher troubleshooting times, and the need for specialized personnel. Simpler designs often translate to more reliable and efficient networks, saving time and resources. Options A, B, and E, although valuable for network operations, are not primary business-level drivers for network design. Integrating endpoint posture (A) is a security concern, not a broad business objective. Faster obsolescence (B) is the opposite of a good business goal, where they aim for a longer use cycle. Standardizing resiliency (E) is a network requirement that is necessary to implement but not a business goal.

Cost Optimization: Information technology Infrastructure Library (ITIL) emphasizes the importance of cost optimization in service management. https://www.axelos.com/best-practice-solutions/itil
Complexity Management: A complex network increases risk and costs, which is avoided by IT architecture

principles. https://www.opengroup.org/architecture/togaf

Question: 10

CONNECTIVITY	CAPEX	OPEX ANNUAL	INSTALLATION FEE	TERM
DWDM over dark fiber	$200,000	$100,000	$30,000	12 months
CWDM over dark fiber	$150,000	$100,000	$25,000	18 months
MPLS wires only	$50,000	$180,000	$5,000	12 months
Metro Ethernet	$65,000	$100,000	$5,000	36 months

Refer to the table. A customer investigates connectivity options for a DCI between two production data centers to aid a large-scale migration project. The migration is estimated to take 20 months to complete but might extend an additional 10 months if issues arise. All connectivity options meet the requirements to migrate workloads. Which transport technology provides the best ROI based on cost and flexibility?

A. DWDM over dark fiber
B. MPLS
C. CWDM over dark fiber
D. Metro Ethernet

Answer: D

Explanation:

Metro Ethernet is a correct answer.

Question: 11

SDWAN networks capitalize the usage of broadband Internet links over traditional MPLS links to offer more cost benefits to enterprise customers. However, due to the insecure nature of the public Internet, it is mandatory to use encryption of traffic between any two SDWAN edge devices installed behind NAT gateways.
Which overlay method can provide optimal transport over unreliable underlay networks that are behind NAT gateways?

A. DTLS
B. TLS
C. IPsec
D. GRE

Answer: C

Explanation:

The correct answer is C. IPsec (Internet Protocol Security). Here's a detailed justification:

SD-WAN deployments often utilize the public internet as an underlay, offering cost advantages over traditional MPLS. However, the public internet's inherent insecurity necessitates encryption, especially when edge devices reside behind NAT gateways. DTLS (Datagram Transport Layer Security) and TLS (Transport Layer Security) are primarily designed for securing transport layer protocols like TCP, not the entire IP packet

required in SD-WAN overlays. While DTLS might be suitable for specific applications using UDP, it's not the standard for wide-scale IP packet encryption. GRE (Generic Routing Encapsulation) is a tunneling protocol, not an encryption protocol, making it unsuitable for securing traffic over the public internet; it can encapsulate but doesn't inherently encrypt. IPsec, on the other hand, is an industry-standard suite of protocols designed for securing IP traffic, offering both authentication and encryption at the network layer.

This makes it ideal for creating secure tunnels over unreliable underlay networks. IPsec operates in either tunnel or transport mode; tunnel mode is typically used for creating end-to-end VPN tunnels, encrypting the entire IP packet including header. This is crucial when dealing with NAT, as it protects routing information behind the NAT gateway. IPsec's ability to handle NAT traversal through NAT-T (NAT Traversal) makes it uniquely suited for SD-WAN deployments with edge devices behind NAT gateways. In such scenarios, IPsec establishes secure, encrypted channels, enabling optimal and secure communication over public internet.

Thus, IPsec provides the most suitable combination of security, NAT traversal, and optimal transport over unreliable underlay networks compared to DTLS, TLS or GRE.

Authoritative Links for Further Research:

IETF IPsec Standards:https://datatracker.ietf.org/doc/search/?name=ipsec&rfcs=on
Cisco SD-WAN Security:https://www.cisco.com/c/en/us/td/docs/solutions/CVD/SDWAN/cisco-sdwan-security-design.html
IPsec Overview:https://www.cloudflare.com/learning/network-layer/what-is-ipsec/

Question: 12

Company XYZ wants to use the FCAPS ISO standard for network management design. The focus of the design should be to monitor and keep track of any performance issues by continuously collecting and analyzing statistical information to monitor, correct, and optimize any reduced responsiveness across the network. Which layer accomplishes this design requirement?

A. security management
B. performance management
C. accounting management
D. fault management

Answer: B

Explanation:

The correct answer is **B. performance management**. The question explicitly focuses on monitoring, analyzing, and optimizing network responsiveness using statistical data. This aligns perfectly with the core functions of performance management within the FCAPS (Fault, Configuration, Accounting, Performance, Security) model.

Performance management is dedicated to assessing and maintaining the efficiency and effectiveness of network resources. It involves collecting metrics like latency, throughput, and packet loss, then using this data to identify bottlenecks, optimize configurations, and proactively address potential performance issues.

Option A, security management, primarily concentrates on protecting network assets from unauthorized access and threats. Option C, accounting management, tracks resource utilization for billing and cost allocation. Option D, fault management, focuses on detecting, isolating, and resolving network failures. While these layers are all vital in a comprehensive network management strategy, the direct focus on analyzing statistics to improve network responsiveness is distinctly within the realm of performance management. Therefore, option B is the most accurate choice for the described design requirement.

For further research on FCAPS, refer to:

ITU-T Recommendation M.3400:https://www.itu.int/rec/T-REC-M.3400-200004-I/en (This document defines the FCAPS model).

Network Management: Principles and Practices (Mani Subramanian): A comprehensive resource on network management concepts including FCAPS. (Search for this book title on academic platforms)

Question: 13

Company XYZ has implemented policy-based routing in their network. Which potential problem must be kept in mind about network reconvergence and PBR?

A. It can limit network scalability.
B. It can create microloops during reconvergence.
C. It reduces convergence time.
D. It increases convergence time.

Answer: B

Explanation:

Policy-Based Routing (PBR), while powerful for traffic manipulation, can introduce complexities during network reconvergence, particularly concerning microloops. Unlike traditional routing protocols that rely on destination-based forwarding, PBR makes decisions based on source, protocol, or other criteria, potentially leading to inconsistencies as the network adapts to changes. During reconvergence after a link or node failure, routes might update at different times across the network. PBR policies, if not carefully designed, could result in packets being forwarded through alternative paths based on outdated information. This can create temporary forwarding loops, where packets circulate in a closed path before being discarded or eventually reaching their destination. These microloops are typically short-lived, but can lead to packet loss, latency, and overall network instability during reconvergence.

While PBR itself doesn't inherently impact scalability (option A) or generally reduce convergence time (option C), its implementation needs careful planning. Using simple default routes would be much faster to converge.
Option D, stating that PBR increases convergence time, is not directly true. While PBR can potentially slow down convergence, the main concern during reconvergence is microloops caused by misconfigurations or policies that redirect packets without proper consideration of the overall network topology. The fundamental issue is the policy layer being separate from the actual routing protocols, thus when network changes occur the policy layer can cause misdirections until the routing protocol learns the changes and the policy layer adapts. In summary, the potential for PBR to introduce microloops during reconvergence, due to its reliance on policy-based forwarding rather than purely destination-based routing, is a key challenge to understand when implementing the feature.

Further Research:

Cisco Documentation on Policy-Based Routing: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pbr/configuration/15-sy/irp-pbr-15-sy-book/pbr-oview.html
"Troubleshooting Policy-Based Routing" - Cisco Support Forum: [Search for "Cisco Policy Based Routing Troubleshooting Microloops" on your preferred search engine]
"Policy-Based Routing with a Focus on Avoidance of Loops": [Search in academic databases for papers discussing PBR loop prevention methods]

Question: 14

SD-WAN can be used to provide secure connectivity to remote offices, branch offices, campus networks, data centers, and the cloud over any type of IP-based underlay transport network. Which two statements describe SD-WAN solutions? (Choose two.)

A. Control and data forwarding planes are kept separate.
B. Solutions allow for variations of commodity and specialized switching hardware.
C. SD-WAN networks are inherently protected against slow performance.
D. Solutions include centralized orchestration, control, and zero-touch provisioning.
E. Improved operational efficiencies result in cost savings.

Answer: AD

Explanation:

Here's the justification for choosing options A and D as correct descriptions of SD-WAN solutions:

Option A: Control and data forwarding planes are kept separate. This is a fundamental principle of Software-Defined Networking (SDN), and SD-WAN is a specific application of SDN principles in wide-area networking. The separation allows for centralized control over the network, making management and policy enforcement more flexible and efficient. The control plane dictates how data should flow, while the data plane is responsible for the actual data forwarding. This separation facilitates greater programmability and allows for network behavior adjustments without needing physical configuration changes on individual network devices.

Option D: Solutions include centralized orchestration, control, and zero-touch provisioning. This aligns with the core goals of SD-WAN: simplifying and automating wide-area network operations. Centralized orchestration enables administrators to manage the entire network from a single point, apply policies consistently, and easily monitor performance. Zero-touch provisioning streamlines the deployment of new devices by automating the configuration process, reducing the need for manual intervention and shortening the time to bring new sites online. This improves scalability and reduces operational complexities.

Why other options are incorrect:

Option B: While SD-WAN solutions can leverage commodity hardware, the emphasis is more on software abstraction rather than reliance on variations of switching hardware. SD-WAN's flexibility stems from its ability to use any IP underlay, not specific hardware types.

Option C: SD-WAN provides tools for managing performance issues, not inherent protection. It can dynamically steer traffic and prioritize applications based on network conditions, but it doesn't magically prevent performance problems caused by underlying network issues. It enables better traffic management and optimization.

Option E: While improved operational efficiencies can lead to cost savings, it's an outcome rather than a defining characteristic of SD-WAN solutions themselves. The core function is not cost saving.

Authoritative Links for further research:

1. **Cisco SD-WAN Overview:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/sd-wan/index.html
2. **VMware SD-WAN:**https://www.vmware.com/products/sd-wan.html
3. **Fortinet SD-WAN:**https://www.fortinet.com/products/sd-wan
4. **Juniper SD-WAN:**https://www.juniper.net/us/en/products-services/sd-wan/
5. **Nokia SD-WAN:**https://www.nokia.com/networks/solutions/sd-wan/These links provide comprehensive details on various SD-WAN solutions, including their architectural underpinnings and core features, thus solidifying the choice of options A and D.

Question: 15

Company XYZ is in the process of identifying which transport mechanism(s) to use as their WAN technology. Their main two requirements are:
* a technology that could offer DPI, SLA, secure tunnels, privacy, QoS, scalability, reliability, and ease of

management
* a technology that is cost-effective
Which WAN technology(ies) should be included in the design of company XYZ?

A. Both technologies should be used. Each should be used to back up the other one; where the primary links are MPLS, the Internet should be used as a backup link with IPsec (and vice versa).

B. MPLS meets all these requirements and it is more reliable than using the Internet. It is widely used with clearly defined best practices and an industry standard.

C. Software-defined WAN should be the preferred choice because it complements both technologies, covers all the required features, and it is the most cost- effective solution.

D. Internet should be the preferred option because it is cost effective and supports BFD, IP SLA, and IPsec for secure transport over the public Internet.

Answer: C

Explanation:

The correct answer is C because Software-Defined WAN (SD-WAN) is the most suitable solution for Company XYZ's requirements. While MPLS (B) offers many features like QoS, SLAs, and reliability, it can be less cost-effective and less flexible than SD-WAN. The Internet (D), while cost-effective, lacks inherent features like guaranteed QoS and SLAs. SD-WAN (C) addresses this by providing a centralized management plane that can intelligently steer traffic across multiple transport links (including MPLS, Internet, and even cellular), optimizing for cost, performance, and security. SD-WAN overlays a virtualized network on top of the chosen transport, enabling features like DPI, QoS, secure tunnels (IPsec or similar), and easy management, as well as providing scalability and reliability features through path selection and failover mechanisms. By leveraging both Internet and private circuits as needed, SD-WAN can fulfill all of XYZ's requirements, making it the superior choice, especially compared to relying solely on MPLS (B) or the Internet (D). A hybrid approach with SD-WAN on top of MPLS and Internet can also be implemented, providing the benefits of both without their limitations. Option A is too simplistic as SD-WAN would provide a far more dynamic and automated path selection and backup capability than a static primary/backup approach. The ease of management, which is a key requirement, is a core strength of SD-WAN.

Further research:

1. **Cisco SD-WAN Documentation:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/sd-wan/index.html
2. **VMware SD-WAN Documentation:**https://www.vmware.com/products/sd-wan.html
3. **Fortinet SD-WAN:**https://www.fortinet.com/products/sd-wan
4. **Gartner Magic Quadrant for SD-WAN:** (Search on Gartner website for the latest report) Provides analysis and comparison of SD-WAN vendors.
5. **MEF Standards for SD-WAN:**https://www.mef.net/ (Look for SD-WAN related standards) Provides industry-standard definitions and specifications for SD-WAN.

Question: 16

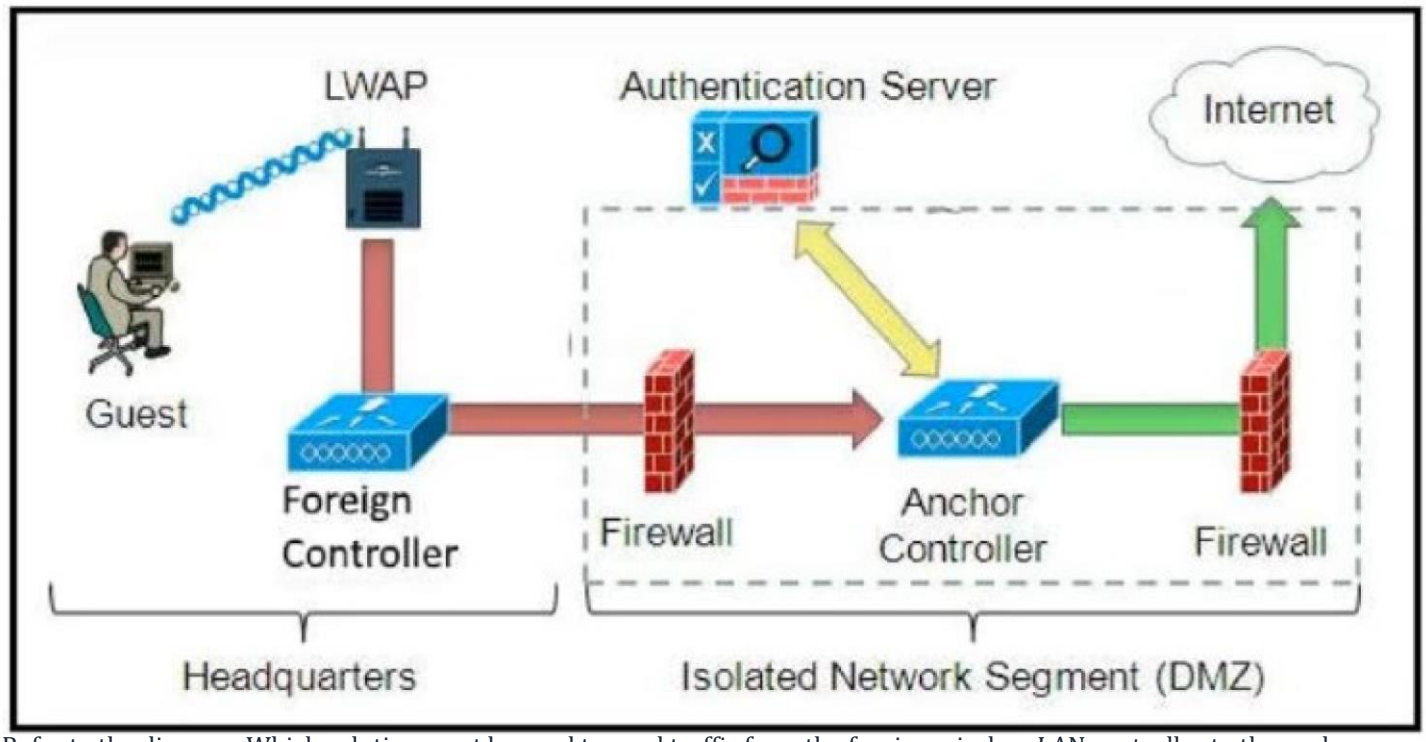

Refer to the diagram. Which solution must be used to send traffic from the foreign wireless LAN controller to the anchor wireless LAN controller?

A. Send packets without encapsulation to the anchor controller over the routed network.
B. Encapsulate packets into an EoIP tunnel and send them to the anchor controller.
C. Send packets from the foreign controller to the anchor controller via Layer 3 MPLS VPN or VRF-Lite.
D. Send packets from the foreign controller to the anchor controller via IPinIP or IPsec tunnel.

Answer: B

Explanation:

Encapsulate packets into an EoIP tunnel and send them to the anchor controller.

Question: 17

The Company XYZ network is experiencing attacks against their router. Which type of Control Plane Protection must be used on the router to protect all control plane IP traffic that is destined directly for one of the router interfaces?

A. Control Plane Protection transit subinterface
B. Control Plane Protection host subinterface
C. Control Plane Protection CEF-exception subinterface
D. Control Plane Protection main interface

Answer: B

Explanation:

The correct answer is **B. Control Plane Protection host subinterface**. Control Plane Protection (CoPP) is a vital security mechanism for network devices like routers, designed to safeguard the control plane from malicious or unintentional traffic overload. CoPP implements policies to rate-limit and prioritize traffic destined for the router's control plane, which handles critical functions such as routing protocol processing and management access.

When traffic is explicitly destined for an interface of the router itself, it's classified as "host" traffic. The "host subinterface" configuration specifically targets this type of traffic. It allows you to apply stricter security policies to traffic attempting direct interaction with the router's interfaces, unlike transit traffic which passes through the router. Transit subinterface (option A) is focused on traffic that traverses the router, not traffic destined for the router itself. CEF-exception (option C) deals with packets that bypass the normal CEF forwarding path, often due to special processing needs. The main interface (option D) offers a broader scope and might not offer the granularity needed for specific host traffic protection.

By applying CoPP through the host subinterface, the router can effectively mitigate attacks like denial-of-service (DoS) or reconnaissance attempts targeting its management interfaces or routing processes. This fine-grained control ensures that the router remains stable and functional even under attack. It is a crucial layer of defence to ensure the continued operability of network infrastructure, highlighting the need for a well-defined policy focused on local control plane security. This configuration ensures that the router processes only legitimate traffic directed towards its own interfaces, reinforcing overall network resilience.

Authoritative Links for further research:

Cisco Control Plane Protection:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/cpp/configuration/15-mt/sec-cpp-15-mt-book.html
Cisco Security Configuration Guide:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/configuration/xe-16/sec-xe-16-book.html

Question: 18

An architect designs a multi-controller network architecture with these requirements: *
Achieve fast failover to control traffic when controllers fail.
* Yield a short distance and high resiliency in the connection between the switches and the controller. * Reduce connectivity loss and enable smart recovery to improve the SDN survivability.
* Improve connectivity by adding path diversity and capacity awareness for controllers.
Which control plane component of the multi-controller must be built to meet the requirements?

A. control node reliability
B. control path reliability
C. controller state consistency
D. controller clustering

Answer: B

Explanation:

The correct answer is **B. control path reliability**. Here's why:

The requirements explicitly focus on ensuring the network remains operational even when controllers fail.
This involves establishing resilient connections between switches and controllers, minimizing connectivity loss during failures, and promoting smart recovery. Control path reliability directly addresses this concern by focusing on the communication pathways between network devices (switches) and the control plane (controllers). It involves establishing redundant paths and mechanisms for rapid failover.

Fast failover and short distance/high resilience: These requirements point to the need for multiple, geographically diverse paths from switches to controllers, which is directly addressed by control path reliability.
Reduced connectivity loss and smart recovery: Improved control path reliability enables rapid detection of failures and automatic rerouting of communication to healthy controllers, minimizing service disruptions.
Path diversity and capacity awareness: These are key features of a robust control path. Control path reliability ensures that there are multiple paths to handle traffic.

Option A, control node reliability, deals with the stability and redundancy of the controllers themselves, not the communication paths to them. Option C, controller state consistency, addresses how controllers maintain synchronized information, but not the control paths to switches. Option D, controller clustering, helps with the overall availability of controllers but doesn't specifically focus on optimizing the connectivity between the switches and controllers. While all options are relevant for a highly available network, the focus of the requirements on communication paths makes control path reliability the best fit.

Authoritative links for further research:

1. **ONF (Open Networking Foundation):** While specific documentation may vary, ONF's work on SDN architectures and control plane design often touches upon the concept of control path reliability.
   https://opennetworking.org/
2. **IETF (Internet Engineering Task Force):** Explore IETF RFCs related to networking protocols and architectures. Search for topics like "High Availability," "Resiliency," and "Path Redundancy".
   https://www.ietf.org/
3. **Research papers on SDN architectures:** You can find academic articles exploring SDN control plane design, including topics of reliability and redundancy. Google Scholar or ACM Digital Library are good places to search.

These resources will offer further insights into the significance and implementation of control path reliability in multi-controller networks.

Question: 19

Which two control plane policer designs must be considered to achieve high availability? (Choose two.)

A. Control plane policers are really needed only on externally facing devices.

B. Control plane policers can cause the network management systems to create false alarms.

C. Control plane policers require that adequate protocols overhead are factored in to allow protocol convergence.

D. Control plane policers must be processed before a forwarding decision is made.

E. Control plane policers are enforced in hardware to protect the software path, but they are hardware platform-dependent in terms of classification ability.

Answer: DE

Explanation:

Here's a detailed justification for why options D and E are the correct choices for control plane policer designs that ensure high availability:

D. Control plane policers must be processed before a forwarding decision is made. This is fundamental for high availability because control plane policing protects the device's CPU from being overwhelmed by excessive traffic directed at it, rather than through it. If a forwarding decision (based on the forwarding plane) is made before control plane filtering, a flood of malicious or malformed packets could overwhelm the CPU and crash the device, preventing it from performing its routing and control functions, thus severely impacting availability. Processing policers before routing ensures the device can remain functional and maintain stability even under attack.

E. Control plane policers are enforced in hardware to protect the software path, but they are hardware platform-dependent in terms of classification ability. This is crucial for high availability because hardware enforcement (typically ASICs) is essential for performance. Software-based policing would be too slow to protect the control plane under load. However, the fact that hardware capabilities differ across platforms implies that policer implementations will vary, and one size will not fit all. This variability dictates that careful

consideration is required when designing and deploying control plane policies across different devices to ensure the same levels of protection and performance. Platform-specific limitations need to be accounted for in order to maintain optimal security and availability on every device within the network. Ignoring these platform differences can lead to weaknesses or inconsistent protections.

Options A, B, and C are not correct. Option A is incorrect because control plane policers are essential on all devices, not just external ones. Any device could become overwhelmed from internal attacks. Option B is incorrect because while poorly designed policers can cause false alarms, they can be tuned properly to prevent this, so its not an inherent aspect of policers. Lastly, Option C is incorrect because protocol overhead is a factor that is considered when configuring policers; while this does relate to a network, the question asks specifically about control plane policer designs for HA.

Authoritative Links for Further Research:

Cisco: Control Plane Protection:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_plcng/configuration/15-sy/qos-plcng-15-sy-book/qos-polic-control-plane.html
Juniper Networks: Control Plane Rate Limiting:
https://www.juniper.net/documentation/us/en/software/junos/security/topics/concept/security-control-plane-rate-limiting.html

These links provide detailed technical information about control plane protection, rate-limiting, and the underlying mechanisms. These resources will greatly help further your understanding of the concept.

Question: 20

A small organization of 20 employees is looking to deliver a network design service for modernizing customer networks to support advanced solutions.
* Project scope and weekly progress should be visualized by the management.
* Always consider feedback and make changes accordingly during the project.
* Should consider flexibility to change scope at the point of time.
Which project methodology meets the requirements and have the least impact on the outcome?

A. LEAN
B. Six-Sigma
C. Scrum
D. Kanban

Answer: D

Explanation:

The correct answer is **D. Kanban**. Here's why:

Kanban excels in scenarios requiring flexibility, transparency, and iterative progress, perfectly aligning with the organization's needs. Its visual nature, using a board to track tasks ("to do," "in progress," "done," etc.), provides management with immediate insight into project status and workflow. This visibility facilitates early identification of bottlenecks and enables informed decisions regarding resource allocation. Kanban's core principle revolves around continuous flow, making it inherently adaptable to changes in scope or requirements. Feedback is easily incorporated by adjusting the board's columns or prioritizing tasks within the workflow. Unlike Lean (A), which focuses on eliminating waste, or Six-Sigma (B), which aims for process improvement via statistical analysis, Kanban is primarily concerned with workflow optimization and transparency, ideal for a dynamic project. Scrum (C), while also iterative, is more prescriptive with time-boxed sprints, which may not suit the need for frequent scope adjustments and rapid feedback incorporation, inherent to the project in question. Kanban, therefore, offers the least disruptive approach and best

accommodates the organization's need for adaptable, transparent, and continuous delivery with regular feedback integration. The focus on visualizing workflow and limiting work in progress makes Kanban highly effective in a small team environment, minimizing overhead and maximizing flexibility.

Supporting Links:

Atlassian - What is Kanban?:https://www.atlassian.com/agile/kanban
Kanbanize - Kanban Explained:https://kanbanize.com/kanban-resources/getting-started/what-is-kanban

Question: 21

Which two impacts of adding the IP event dampening feature to a network design are true? (Choose two.)

A. It switches traffic immediately after a link failure
B. It improves overall network stability
C. It speeds up link failure detection
D. It protects against routing loops
E. It reduces the utilization of system processing resources

Answer: BE

Explanation:

Let's break down why options B and E are the correct impacts of IP event dampening, while A, C, and D are incorrect.

IP event dampening is a mechanism designed to mitigate the effects of rapidly flapping interfaces or routes. A flapping interface is one that repeatedly transitions between up and down states. This constant state change can trigger numerous routing updates and recalculations, which can destabilize the network and consume excessive processing resources.

Option B, "It improves overall network stability," is correct because dampening suppresses repetitive updates caused by flapping. By doing so, it prevents these problematic state changes from constantly impacting routing protocols, thus creating a more stable network environment. Instead of immediately reacting to every change, dampening adds a waiting period, allowing the network to settle before further recalculations are triggered.

Option E, "It reduces the utilization of system processing resources," is also correct. When a route or interface flaps without dampening, each transition necessitates recalculations by routing protocols and updates to the routing table. These repetitive actions drain CPU resources. Dampening slows down the propagation of these changes, thereby conserving computational power and improving the overall system performance.

Option A, "It switches traffic immediately after a link failure," is incorrect. Dampening adds a delay, preventing immediate traffic switching. It's designed to reduce the reaction speed to flapping, not to accelerate it.

Option C, "It speeds up link failure detection," is also incorrect. Dampening does not influence the speed at which failures are detected. Detection speed is typically handled by technologies like BFD (Bidirectional Forwarding Detection). Dampening only controls the reaction to the detected event.

Finally, option D, "It protects against routing loops," is incorrect. While dampening can indirectly reduce the likelihood of loops arising from flapping links, its primary goal is not loop prevention. Routing loops are more directly addressed by techniques such as split horizon and route poisoning, or better convergence times using modern link-state routing protocols.

In summary, IP event dampening prioritizes network stability and efficient resource usage by controlling the

dissemination of network changes, not by immediately reacting to or accelerating detection of those changes.

For further research, you can refer to Cisco's documentation on IP event dampening:

Cisco Documentation: Explore Cisco's detailed guides on configuring and using IP Event Dampening, which are readily available through their online support and documentation platforms. You can find these by searching "[Cisco] IP Event Dampening" or by navigating to the specific documentation for their devices that support this feature.

Question: 22

You have been asked to design a high-density wireless network for a university campus. Which two principles would you apply in order to maximize the wireless network capacity? (Choose two.)

A. Choose a high minimum data rate to reduce the duty cycle.

B. Make use of the 5-GHz band to reduce the spectrum utilization on 2.4 GHz when dual-band clients are used.

C. Enable 802.11n channel bonding on both 2.4 GHz and 5 GHz to increase the maximum aggregated cell throughput.

D. Increase the number of SSIDs to load-balance the client traffic.

E. Implement a four-channel design on 2.4 GHz to increase the number of available channels.

Answer: AB

Explanation:

Here's a justification for why options A and B are the correct choices for maximizing wireless network capacity in a high-density environment, along with explanations of why the other options are incorrect:

Justification for A and B:

A. Choose a high minimum data rate to reduce the duty cycle: A higher minimum data rate forces clients to communicate more efficiently. Lower data rates take up more airtime due to slower transmissions, leading to a higher duty cycle (percentage of time the channel is busy). By setting a higher minimum rate (e.g., requiring 802.11n speeds) slower devices that may occupy the channel longer are effectively excluded or forced to use faster modulation and coding schemes. This reduces airtime consumption for all devices and increases overall network capacity. Think of it like a busy highway; by encouraging vehicles to travel at a more efficient speed, more vehicles can traverse the same highway in less time.

Reference: Cisco: High-Density Wireless Network Design: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/7-6/High_Density_Design_Guide/High_Density_Deployment_Guide.html (Section on Data Rates and Modulation)

B. Make use of the 5-GHz band to reduce the spectrum utilization on 2.4 GHz when dual-band clients are used: The 5 GHz band offers more channels than 2.4 GHz, resulting in less congestion and interference. It is also less susceptible to interference from other non-Wi-Fi sources. Encouraging dual-band clients to use 5 GHz frees up the limited channels available in the 2.4 GHz band for older devices that do not support the 5GHz range. This effectively increases the available bandwidth by reducing contention and interference on the 2.4 GHz spectrum. Think of this like creating separate lanes for faster vehicles which avoids them from congesting the slower lanes.

Reference: Aruba Networks: 802.11ac and the 5GHz band: https://www.arubanetworks.com/techdocs/ArubaOS_8/Content/8.0/UG/WLAN_AP/5GHz_band.htm

Why C, D, and E are incorrect:

C. Enable 802.11n channel bonding on both 2.4 GHz and 5 GHz to increase the maximum aggregated cell throughput: While channel bonding can increase maximum throughput, it's problematic in high-density scenarios. Bonding channels creates a wider channel which makes them susceptible to interference. The wider the channel, the more likely interference from neighboring cells become an issue. In a high-density environment, the potential for co-channel interference on these wider channels increases significantly, negating the benefits of channel bonding. The limited number of channels in 2.4 GHz makes the issue more pronounced than in 5 GHz. Wider channels reduce the number of non-overlapping channels available in 5GHz, which again makes it less suitable for a high density environment.

D. Increase the number of SSIDs to load-balance the client traffic: While segmenting networks using SSIDs is useful for security or different categories of users, increasing the number of SSIDs does not inherently increase network capacity. Each SSID requires additional overhead in the form of beacon frames, which consumes airtime and resources. More SSIDs will increase the overhead and decrease effective throughput. A more effective way to load-balance clients is by moving clients to less congested radios or bands.

E. Implement a four-channel design on 2.4 GHz to increase the number of available channels: 2.4 GHz only offers three non-overlapping channels (1, 6, and 11). Using a four-channel design creates overlapping channels that cause more interference, degrading performance. Therefore, this is not feasible and will reduce rather than increase capacity.

Question: 23

Which optimal use of interface dampening on a fast convergence network design is true?

A. when the switch hardware is faster than the debounce timer down detection
B. when numerous adjacent flaps of very short duration occur
C. when occasional flaps of long duration occur
D. when the router hardware is slower than the carrier delay down detection

Answer: B

Explanation:

Interface dampening is a mechanism used in network routing to mitigate the negative impact of flapping interfaces on network stability. Flapping refers to an interface rapidly transitioning between up and down states. Option B, "when numerous adjacent flaps of very short duration occur," is the correct scenario for optimal use because dampening is specifically designed to address situations where interfaces are rapidly changing state. When multiple short-duration flaps happen, they can cause routing protocols to repeatedly recalculate routes, consuming CPU resources and potentially causing network instability. Interface dampening works by assigning a penalty to each flap. As the penalty accumulates, the interface is increasingly suppressed (dampened). This suppression prevents the routing protocol from reacting to each individual flap, allowing the network to stabilize. Only after the penalty decays below a certain threshold will the interface be brought back into the active routing process.

Option A describes a situation where switch hardware speed outweighs the debounce timer, but this does not directly indicate a scenario for optimal dampening usage. Option C, where occasional long-duration flaps occur, suggests underlying problems with the interface that dampening may not address. Dampening is more suited to prevent continuous instabilities arising from short flaps rather than masking long-duration flaps. Option D involves router hardware speed vs carrier delay, which again, isn't the primary use case of interface dampening. Dampening works on the frequency of state changes rather than relative hardware speeds. Ultimately, interface dampening is best applied when interfaces are exhibiting rapid, continuous, albeit short-term, state changes to prevent the instability that they can cause on a network.

Authoritative Links for Further Research:

Cisco Documentation on Interface Dampening:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/15-mt/iro-ospf-15-mt-book/ospf-flaps.html (This provides a Cisco-specific view of dampening, especially in the context of OSPF, but the principles apply more broadly.)
Juniper Networks Documentation on Interface Dampening:
https://www.juniper.net/documentation/us/en/software/junos/routing/topics/topic-map/interface-damping.html (Another vendor's perspective, which reinforces the broad applicability of the concept.)

Question: 24

A healthcare customer requested that SNMP traps must be sent over the MPLS Layer 3 VPN service. Which protocol must be enabled?

A. syslog
B. SNMPv3
C. SNMPv2
D. syslog TLS
E. SSH

Answer: B

Explanation:

The correct answer is SNMPv3. Here's why: The customer requires secure transmission of SNMP traps over an MPLS Layer 3 VPN. SNMPv2, while simpler, lacks robust security features like encryption and authentication, making it unsuitable for sensitive network information crossing VPN boundaries. Syslog, while used for logging, doesn't provide the structured data format for network device information that SNMP offers. Syslog TLS (Transport Layer Security) adds security to syslog, but it still doesn't serve the purpose of SNMP which is to monitor and manage network devices. SSH (Secure Shell) is used for secure remote access, not primarily for device monitoring data. SNMPv3 is designed with enhanced security, including encryption (using protocols like DES, 3DES, and AES) and authentication (using techniques like HMAC-MD5 and HMAC-SHA) via user-based security model (USM). These features are critical for securely transmitting management data, like SNMP traps, across a VPN where confidentiality and integrity are paramount. By enabling SNMPv3, the healthcare customer can ensure their sensitive network monitoring data is protected during transit over the MPLS Layer 3 VPN service. SNMPv3 provides encryption, authentication, and data integrity, which are necessary for secure trap transmission, meeting the requirements of the healthcare organization. Therefore, SNMPv3 is the appropriate protocol.

For further research, refer to the following authoritative links:

Cisco's Documentation on SNMPv3:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/15-mt/snmp-15-mt-book/snmp-snmpv3.html
RFC 3411: An Architecture for Describing SNMP Management Frameworks:
https://datatracker.ietf.org/doc/html/rfc3411
RFC 3414: User-based Security Model (USM) for SNMPv3:https://datatracker.ietf.org/doc/html/rfc3414

Question: 25

DRAG DROP -
Drag and drop the end-to-end network virtualization elements from the left onto the correct network areas on the right.
Select and Place:

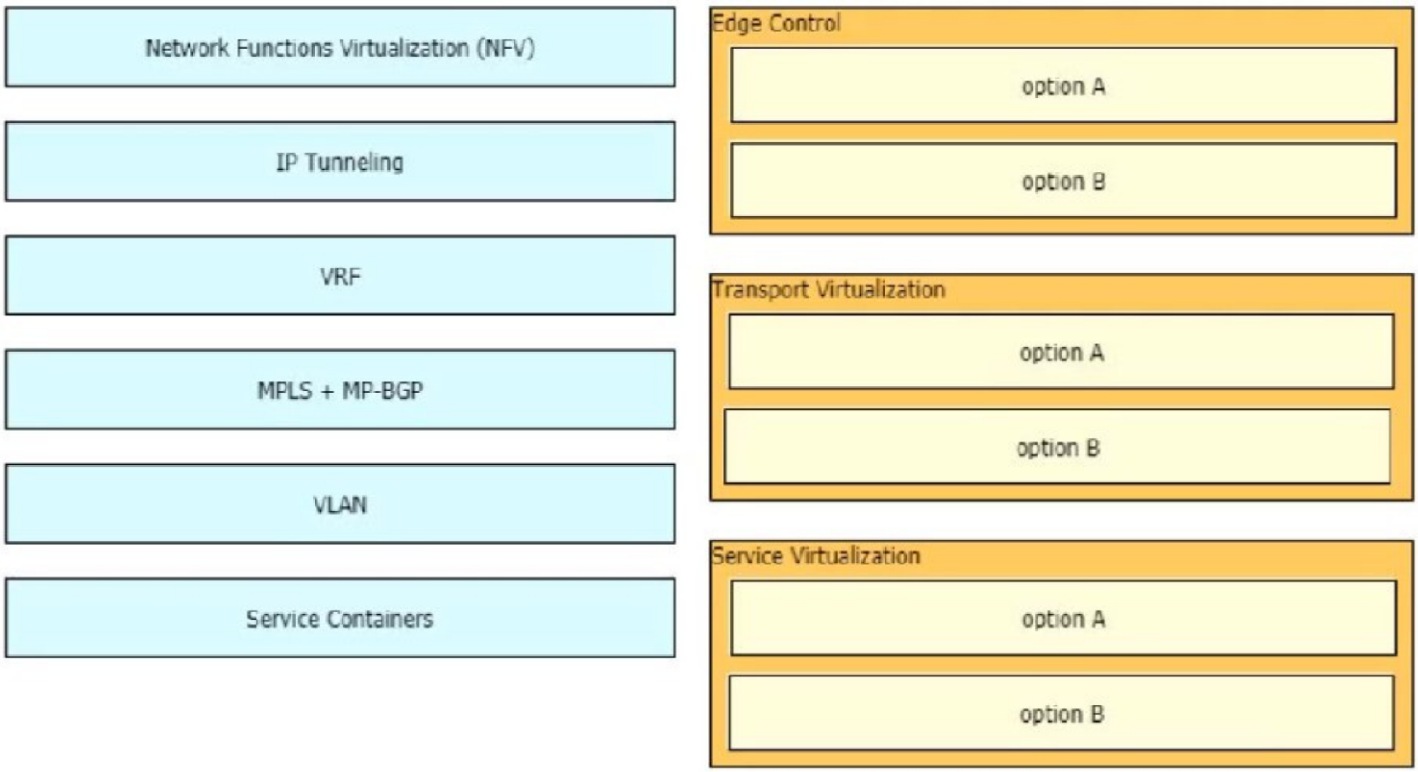

Answer:

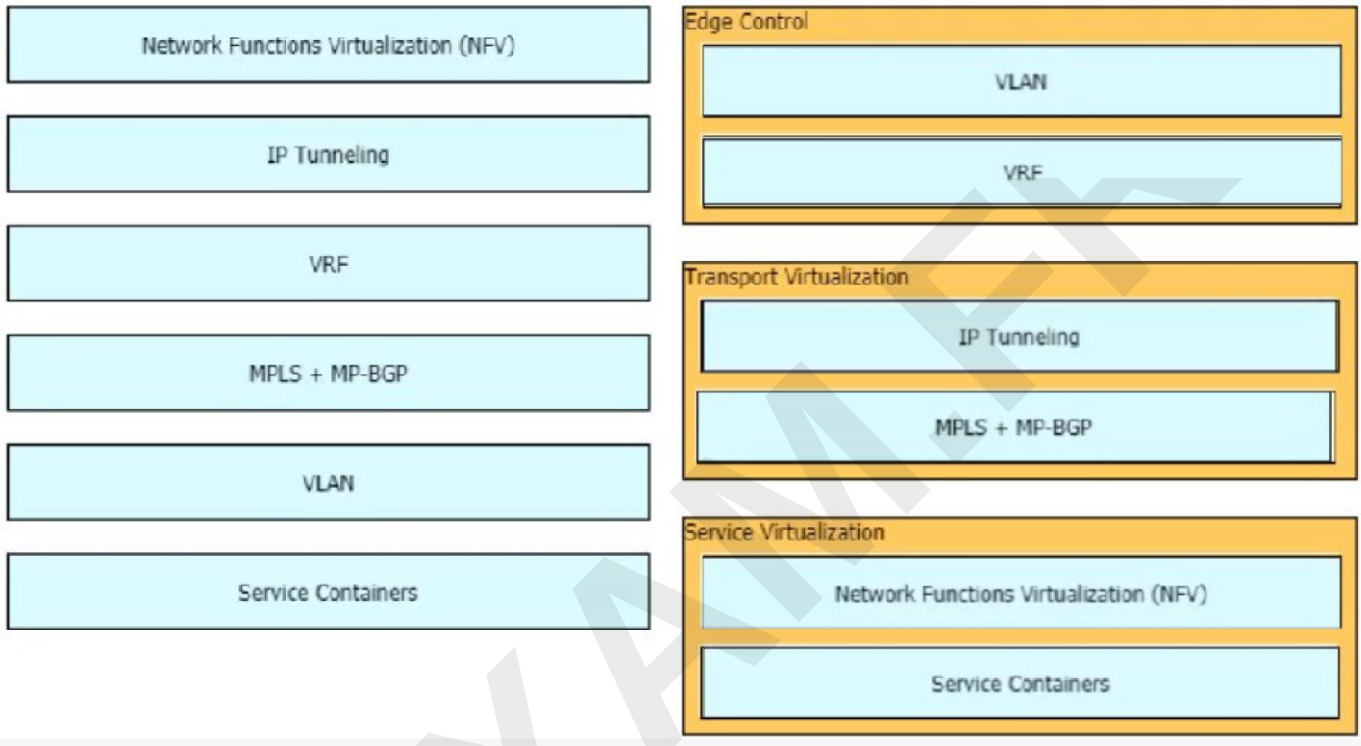

Question: 26

Which management category is not part of FCAPS framework?

A. Performance
B. Authentication
C. Security
D. Fault-management
E. Configuration

Answer: B

Explanation:

The correct answer is **B. Authentication**. FCAPS is a management framework for network management defined by ISO (International Organization for Standardization). It outlines five key functional areas: **Fault, Configuration, Accounting, Performance, and Security**. These areas collectively encompass essential aspects for effectively managing and maintaining networks, including cloud environments. Authentication, while a critical security component, isn't a core category within the FCAPS framework. It falls under the broader Security category of FCAPS. Fault management addresses issues like error detection and recovery. Configuration management involves network device setup and changes. Accounting tracks resource usage for billing or auditing. Performance focuses on network speed and responsiveness. Security addresses data confidentiality, integrity, and availability. Authentication is a critical mechanism within security, but not a separate, high-level category in FCAPS. Therefore, while authentication is very important for network security, it is not a management category that falls within FCAPS's key areas.

https://www.ciscopress.com/articles/article.asp?p=2756241&seqNum=3

https://www.geeksforgeeks.org/f-c-a-p-s-network-management/

Question: 27

A BGP route reflector in the network is taking longer than expected to converge during large network changes. Troubleshooting shows that the router cannot handle all the TCP acknowledgements during route updates. Which action can be performed to tune the device performance?

A. Decrease the size of the small buffers.
B. Increase the size of the large buffers.
C. Increase the keepalive timers for each BGP neighbor.
D. Increase the size of the hold queue.

Answer: D

Explanation:

The correct answer is **D. Increase the size of the hold queue.**

The issue stems from the BGP route reflector struggling to process TCP acknowledgements during significant network changes, indicating a bottleneck in handling the influx of routing information. The hold queue is a buffer that stores incoming BGP messages before they are processed. When this queue is too small, the router cannot keep up with the rate of incoming updates, leading to dropped messages and increased convergence times. Increasing the hold queue provides more space for incoming BGP messages, allowing the router to gracefully absorb the burst of updates without being overwhelmed.

Option A, decreasing small buffer sizes, would exacerbate the problem, making less memory available for packet processing. Option B, increasing large buffer sizes, might help with larger packet handling but is less relevant to the specific issue of message queuing. Option C, increasing keepalive timers, slows down BGP peer liveness detection and would not address the issue of queue overflow.

By increasing the hold queue, we're directly addressing the identified congestion point, enabling more efficient BGP convergence. This is analogous to widening a highway to accommodate more traffic during peak hours. The hold queue provides a temporary buffer, allowing the CPU time to be used for processing routes instead of dropping messages and re-transmitting. Efficient queue management is a crucial aspect of robust network performance, particularly under heavy loads or during periods of change.

Further research on BGP route reflectors and their tuning can be found at these authoritative sources:

Cisco's documentation on BGP: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/15-mt/irg-15-mt-book/irg-bgp-basic.html
Understanding BGP Hold Timers and Keepalives: https://www.networkcomputing.com/networking/bgp-hold-timers-and-keepalives-how-they-work
RFC 4271 on BGP: https://datatracker.ietf.org/doc/html/rfc4271

Question: 28

Which two conditions must be met for EIGRP to maintain an alternate loop-free path to a remote network? (Choose two.)

A. The Reported Distance from a successor is higher than the local Feasible Distance.
B. A feasible successor must be present.
C. The Reported Distance from a successor is lower than the local Feasible Distance.
D. The Feasible Distance from a successor is lower than the local Reported Distance.
E. The feasibility condition does not need to be met.

Answer: BC

Explanation:

Here's a detailed justification for why options B and C are the correct conditions for EIGRP to maintain an alternate loop-free path, also known as a feasible successor:

EIGRP, or Enhanced Interior Gateway Routing Protocol, employs sophisticated mechanisms to prevent routing loops. A key part of this is the concept of feasible successors, which represent backup paths. For an alternate path to qualify as a feasible successor, two primary conditions must be met.

First, **a feasible successor must be present (Option B).** This directly states that the router needs an alternative path (a potential backup route) calculated by another router in the EIGRP domain. Without any discovered alternatives, no feasible successor can exist. This is foundational to maintaining network resiliency.

Second, **the Reported Distance from the potential successor must be lower than the local Feasible Distance (Option C).** The Reported Distance is the metric advertised by a neighbor, and the Feasible Distance is the best path metric to a destination known by the router. This condition ensures that the potential backup path is genuinely better than any previously existing path to the destination. This criteria, known as the Feasibility Condition, avoids loops by preventing a router from selecting a path back through itself or its successors.

Option A is incorrect because the Reported Distance of a successor needs to be lower, not higher, than the current Feasible Distance for it to be considered a feasible successor. Option D is also incorrect as the Feasible Distance from a successor is not considered. Lastly, Option E is incorrect because the feasibility condition must be met to guarantee a loop-free path. The Feasibility Condition is the core principle underpinning EIGRP loop prevention.

In summary, having an alternative path that also passes the Feasibility Condition is mandatory for EIGRP to maintain a loop-free alternate path.

Authoritative Links for further research:

Cisco's EIGRP documentation:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/15-mt/ire-15-mt-book/ire-basic.html (Focus on the "Feasibility Condition" and "Feasible Successor" sections)

Cisco CCNA Study Guide: Various CCNA study materials will delve into the EIGRP Feasibility Condition with explanations and examples. Search for "EIGRP Feasible Successor" in these guides.

Question: 29

DRAG DROP -
Drag and drop the FCAPS network management reference models from the left onto the correct definitions on the right.
Select and Place:

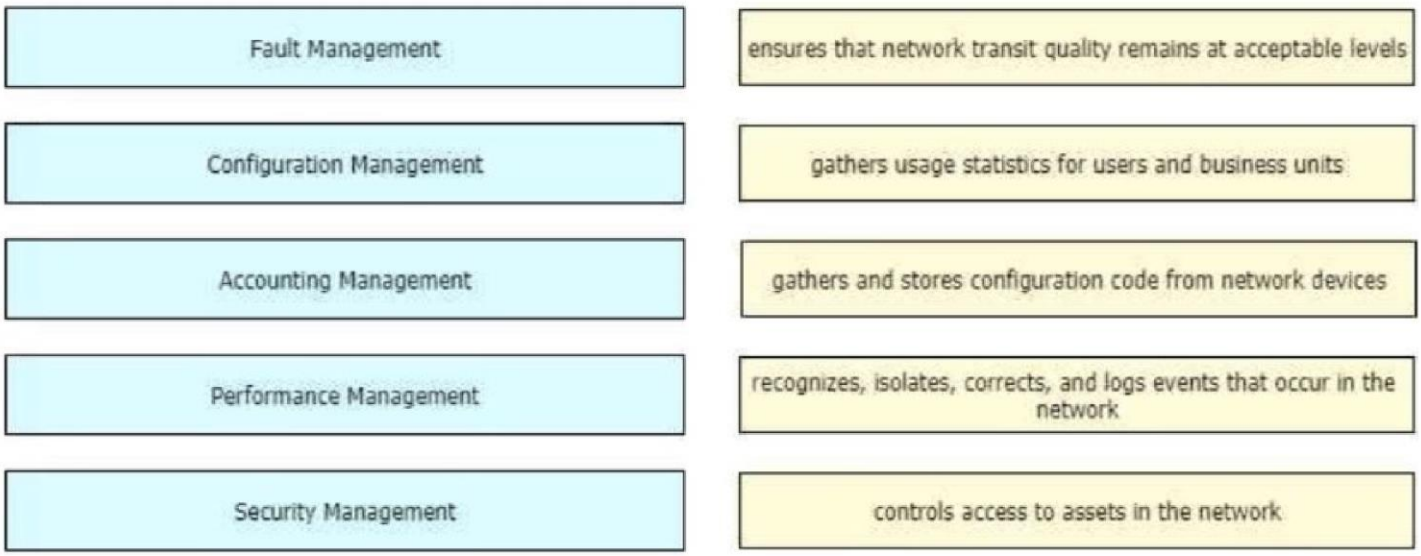

Answer:

Question: 30

Which undesired effect of increasing the jitter compensation buffer is true?

A. The overall transport delay decreases and quality improves.
B. The overall transport jitter increases and quality issues can occur.
C. The overall transport delay increases and quality issues can occur.
D. The overall transport jitter decreases and quality improves.

Answer: C

Explanation:

The correct answer is C: "The overall transport delay increases and quality issues can occur." Jitter compensation buffers, crucial in real-time communication like VoIP and video conferencing, are designed to smooth out variations in packet arrival times (jitter). When this buffer is increased, it means packets are held for a longer period before being released for playout. This holding period directly contributes to an overall increase in transport delay or latency. While buffering helps address jitter, excessively large buffers can introduce unacceptable delays, making real-time conversations difficult and causing lip-sync problems in video. Furthermore, a very large buffer can sometimes become full, leading to discarded packets and quality issues. These discarded packets result in choppy audio or video. Therefore, while a small buffer can help with jitter, an excessively large one introduces latency and potential data loss, impacting overall perceived quality. Finding the right balance for the buffer is crucial.

Authoritative Links for Further Research:

1. **Cisco's Documentation on Jitter Buffers:** https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/srnd/12_5_1/design/cucm_b_1251-srnd/cucm_b_1251-srnd_chapter_010010.html#con_652093 (Look for sections on voice quality and jitter buffer configurations)
2. **IEEE Xplore - Jitter Compensation:** Search IEEE Xplore for papers on "jitter compensation algorithms" or "adaptive jitter buffer" for deeper technical insights. https://ieeexplore.ieee.org/
3. **Cloudflare Learning - Jitter:**https://www.cloudflare.com/learning/performance/what-is-jitter/ (Provides a general overview of jitter)

Question: 31

What is the most important operational driver in building a resilient and secure modular network design?

A. Minimize app downtime
B. Reduce the frequency of failures requiring human intervention
C. Increase time spent on developing new features
D. Dependencies on hardware or software that is difficult to scale

Answer: B

Explanation:

The correct answer is **B. Reduce the frequency of failures requiring human intervention.** This aligns directly with the goals of resilience and operational excellence in network design. A modular network, built for resilience, aims to self-heal and maintain service even during disruptions. Minimizing the need for human intervention translates to faster recovery times and reduced operational costs. Options A and C, while important, are secondary to the core objective of a resilient design. Reducing app downtime (A) is a consequence of resilience, not the driving operational principle itself. Increasing feature development time (C) is unrelated to the immediate need for stability and operational efficiency. Dependence on hard-to-scale components (D) is the opposite of the design goal and introduces fragility, actively undermining both resilience and security. Resilient systems utilize automation, redundancy, and graceful degradation to ensure service continuity with minimal human involvement. By focusing on limiting interventions, the network is inherently more stable, secure (as less manual configuration reduces error opportunities), and cost-effective. This allows resources to be allocated to other areas, such as improving features, once stability and resilience are established. For further reading, consider exploring principles of Site Reliability Engineering (SRE) which heavily emphasizes automation and incident minimization: https://sre.google/ and network automation best practices: https://www.cisco.com/c/en/us/solutions/enterprise-networks/network-automation.html.

Question: 32

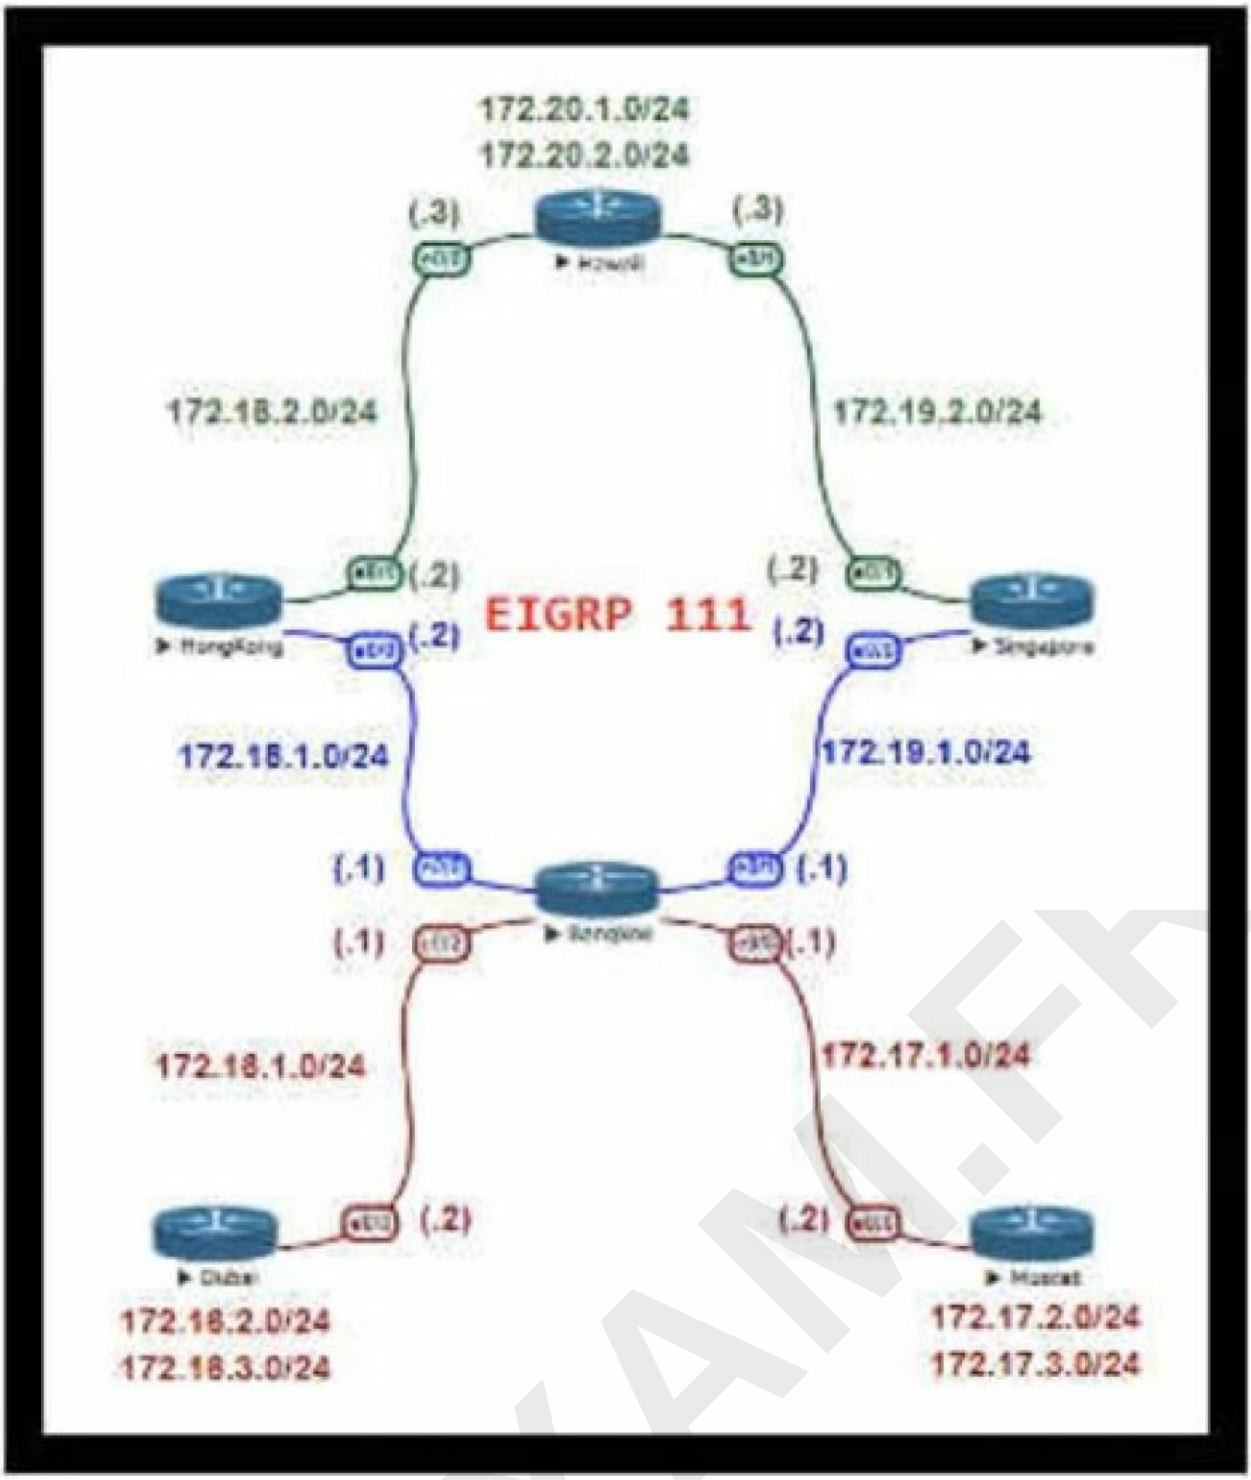

Refer to the exhibit. For Company XYZ, Bangkok is using ECMP to reach the 172.20.2.0/24 network. The company wants a design that would allow them to forward traffic from 172.16.2.0/24 toward 172.20.2.0/24 via the Singapore router as the preferred route. The rest of the traffic should continue to use ECMP. Which technology fulfills this design requirement?

A. policy-based routing
B. unequal-cost load balancing using variance
C. route summarization
D. LFA

Answer: A

Explanation:

policy-based routing is a correct answer.

Question: 33

Company XYZ is running OSPF in their network. They have merged with another company that is running EIGRP as the routing protocol. Company XYZ now needs the two domains to talk to each other with redundancy, while maintaining a loop free environment. The solution must scale when new networks are added into the network in the near future. Which technology can be used to meet these requirements?

A. single point route-redistribution with route filtering using route tags
B. multipoint route-redistribution with route filtering using ACLs
C. multipoint route-redistribution with route filtering using route tags
D. single point route-redistribution with route filtering using ACLs

Answer: C

Explanation:

Here's a breakdown of why option C is the correct solution:

Explanation:

The scenario requires integrating two distinct routing domains (OSPF and EIGRP) while maintaining loop prevention and scalability. Route redistribution is necessary for exchanging routing information between the protocols.

Why Multipoint Redistribution is Preferred:

Redundancy and Resilience: Multipoint redistribution involves distributing redistribution points across multiple routers. This avoids a single point of failure; if one redistribution router fails, others can maintain the route exchange. This enhances overall network reliability. Single-point redistribution creates a bottleneck and a single point of failure, making it unsuitable for networks requiring high availability.

Why Route Tags are Preferred over ACLs for Filtering:

Scalability and Management: Route tags are injected into route updates during redistribution, allowing for filtering based on these tags at other redistribution points. This approach is more scalable and easier to manage than ACLs. ACLs can become complex to maintain as networks grow and require changes in multiple locations, whereas route tags offer a centralized method of managing redistribution policies. Route tags are propagated within routing updates, making them easier to track and understand compared to distributed ACLs. This helps in troubleshooting network problems much more effectively.

Why Option C is Correct:

Option C, "multipoint route-redistribution with route filtering using route tags," precisely addresses all requirements. Multipoint redistribution provides the redundancy and resilience that Company XYZ needs.

Route tags provide the scalable, maintainable method for controlling which routes are passed between the OSPF and EIGRP domains. This approach ensures a loop-free environment, prevents excessive updates, and is easy to manage with new networks added in the future.

Why other options are incorrect:

Options A and D: Single-point redistribution creates a single point of failure.

Options A and D: While ACLs are a filtering mechanism, they are not as scalable or manageable for route filtering in this scenario compared to route tags. Option D would introduce a single point of failure.

Authoritative Links for Further Research:

1. **Cisco Documentation on Route Redistribution:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_redist/configuration/15-mt/irr-15-mt-book/irr-redist-cfg.html
2. **Understanding Route Tags:**https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13442-route-tag.html
3. **Redistribution with EIGRP and OSPF:**https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16408-redistribute.html

Question: 34

What are two common approaches to analyzing and designing networks? (Choose two.)

A. three-tier approach
B. top-down approach
C. high-low security approach
D. bottom-up approach
E. left-right approach

Answer: BD

Explanation:

The correct answer identifies two fundamental network design methodologies: top-down and bottom-up approaches. A top-down approach begins with a high-level view of the network's overall purpose and requirements, translating these into progressively detailed designs. This involves understanding business goals, user needs, and application requirements before delving into technology specifics. It ensures the network's design effectively supports strategic objectives and operational requirements. In contrast, a bottom-up approach starts with available technologies and components, focusing on their capabilities and configurations. This approach involves assembling and integrating individual components to achieve a functional network infrastructure. While it allows for quick implementation, a bottom-up strategy may not fully align with long-term goals and scalability. Options A, C, and E are not established network design methodologies. The three-tier approach (A) is a common network architecture and is often incorporated within a design approach, not the approach itself. Security considerations, including high-low security levels (C), are important aspects of design, but are not design approaches in themselves. The left-right (E) approach is not recognized.

For further research, consider these resources:

Cisco Network Design Guide: While specific guides vary, Cisco's documentation often outlines these approaches in the context of network design: https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Oct2017/CVD-CampusNetworkDesignGuide-OCT17.html (This link is a general guide but will lead to related materials).

Network Design Fundamentals: Search for textbooks or academic articles discussing network design methodologies to get a deeper understanding of the theoretical basis for top-down and bottom-up approaches.

Question: 35

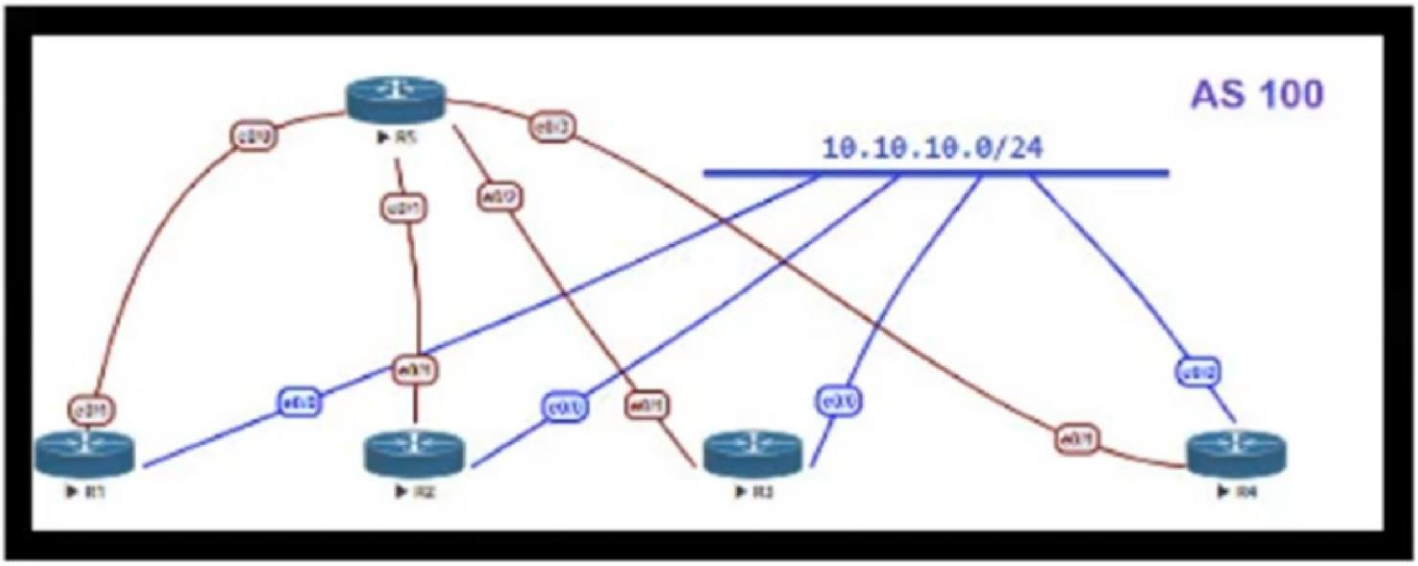

Refer to the exhibit. OSPF is running as the IGP to provide reachability to all AS100 networks. R3 and R4 are the current ABRs at the boundary of OSPF Area0 and Area1. Now BGP must be deployed within AS100 because it will be receiving Internet routes from its eBGP peers (the service provider) connected to R1 and
R2. What is an optimal solution for this deployment to configure BGP relationships and redistribute BGP learned routes into OSPF?

A. R5 should be configured as a route reflector for R1, R2, R3 and R4. BGP routes must be redistributed at R1 and R2 into OSPF.

B. Confederation should be set up with R1, R5, and R3 in one sub AS, with R2 and R4 in another, and redistribution at R1 and R2.

C. R1, R2, R3 and R4 must be set up with a neighbor relationship with R5 only. R5 must not be a route reflector. D. A full mesh should be deployed between all the routers with mutual redistribution to take place at R1 and R2.

Answer: A

Explanation:

R5 should be configured as a route reflector for R1, R2, R3 and R4. BGP routes must be redistributed at R1 and R2 into OSPF.

Question: 36

A multicast network is using Bidirectional PIM. Which two combined actions achieve high availability so that two RPs within the same network can act in a redundant manner? (Choose two.)

A. Advertise the two RP addresses in the routing protocol.

B. Use two phantom RP addresses.

C. Manipulate the multicast routing table by creating static mroutes to the two RPs. D. Control routing to the two RPs through a longest match prefix.

E. Use Anycast RP based on MSDP peering between the two RPs.

F. Manipulate the administrative distance of the unicast routes to the two RPs.

Answer: BD

Explanation:

Here's a detailed justification for why options B and D are the correct choices for achieving high availability with Bidirectional PIM using redundant RPs:

Bidirectional PIM (Bidir PIM) relies on a Rendezvous Point (RP) for shared tree creation. To ensure resilience,

redundant RPs are needed. Let's examine why options B and D achieve this and why the others don't.

Option B: Use two phantom RP addresses. In Bidir PIM, multiple routers can be configured with the same RP address. These aren't physical router interfaces; they are logical addresses, often referred to as "phantom" or "virtual" RPs. By using the same phantom RP address on multiple routers, the network views this as a single, logical RP. When a source sends multicast, it sends to this shared address. The Bidir PIM routers advertise this RP address in their routing protocols. They then participate in an election process (lowest IP address usually wins). If the active RP fails, another router (having the same RP address configured) takes over, providing failover. This approach effectively creates a distributed RP for enhanced reliability.

Option D: Control routing to the two RPs through a longest match prefix. The idea here is to distribute the traffic between the two RPs. If there were two phantom RPs each on different subnets, then a source would send to one or the other based on the route, the RP address being included in the IP destination. With multiple routers using the same phantom address (option B) the routes to the active RP are announced via routing protocol. Option D adds that when more than one router has the same phantom address, longest match routing preference is used to direct sources and receivers to the closest RP. This is important because in PIM-SM, sources and receivers can send to any RP for a particular multicast group. The longest match approach makes sure sources and receivers use the same RP, which speeds up convergence.

Why other options are incorrect:

Option A: Advertise the two RP addresses in the routing protocol. While this is part of the process, it doesn't provide the redundancy in a Bidir-PIM context, where multiple routers share the same logical RP address. If each router advertised a different RP address, the system would not use them for the same multicast group.

Option C: Manipulate the multicast routing table by creating static mroutes to the two RPs. This is a highly inflexible and error-prone solution that is difficult to manage and won't provide dynamic failover. Also, static mroutes are not used in Bidir PIM, so this option is fundamentally wrong.

Option E: Use Anycast RP based on MSDP peering between the two RPs. Anycast RP using MSDP is more relevant in Source-Specific Multicast (SSM) scenarios. MSDP is for sharing source information between RPs (which are separate), which isn't the core concern in Bidir PIM where the goal is having redundant routers using the same RP. Further, Anycast RP is not typically used in Bidir-PIM, rather the same phantom IP address is used.

Option F: Manipulate the administrative distance of the unicast routes to the two RPs. This could influence which router is selected as the active RP but doesn't address the fundamental need for having redundant RPs for the same group. In fact, when the phantom RP address is configured on multiple routers, administrative distance would affect unicast path to RP, which is fine. Longest match is more important.

In summary: The combination of using a single phantom RP address across multiple routers (B) combined with routing to that single address based on the longest match (D) is the correct approach to creating a resilient, redundant RP architecture in Bidir PIM.

Authoritative Links for Further Research:

Cisco's documentation on Bidirectional PIM:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/configuration/15-mt/imc-15-mt-book/imc-bidir-pim.html
Juniper's documentation on Bidirectional PIM:
https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/concept/bidirectional-pim-overview.html

Question: 37

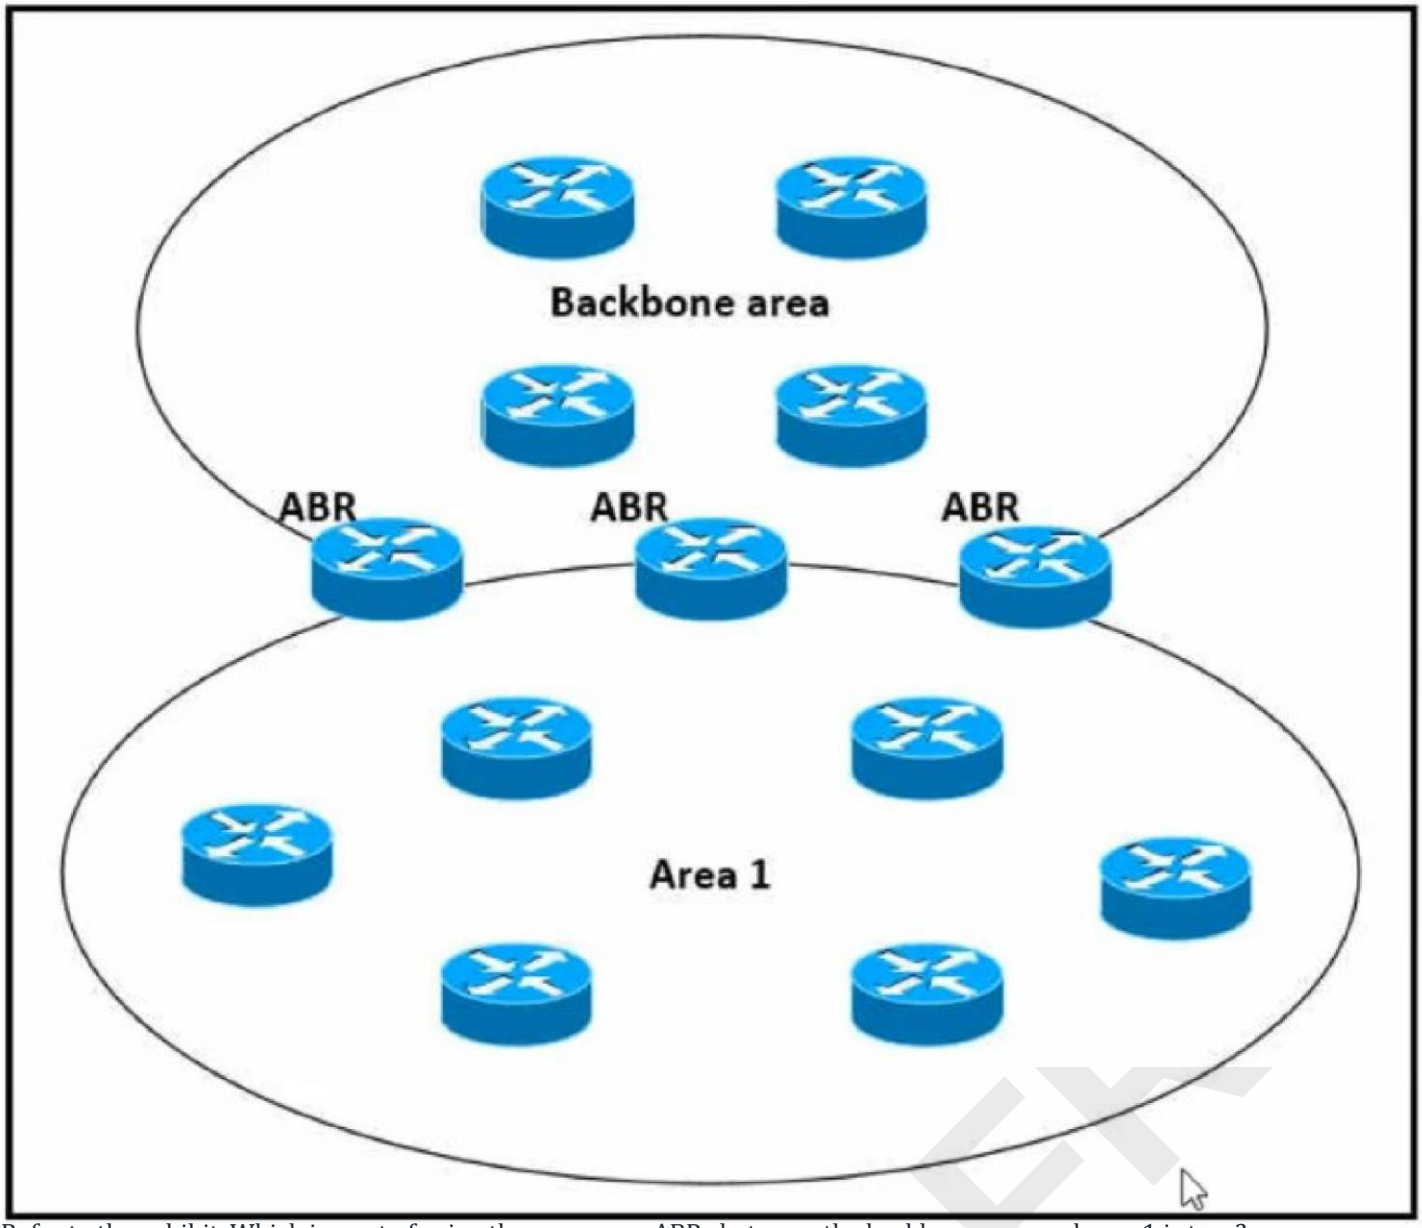

Refer to the exhibit. Which impact of using three or more ABRs between the backbone area and area 1 is true?

A. In a large-scale network, multiple ABRs can create microloops.

B. Multiple ABRs reduce the CPU processing on each ABR due to splitting prefix advertisement between areas. C. Prefixes from the non-backbone area are advertised by one ABR to the backbone area.

D. In a large-scale network, LSA replication, by all ABRs, can cause serious scalability issues.

Answer: D

Explanation:

In a large-scale network, LSA replication, by all ABRs, can cause serious scalability issues.

Question: 38

Company XYZ runs OSPF in their network. A design engineer decides to implement hot-potato routing architecture. How can this implementation be achieved?

A. Enable iBGP and apply prepend to ensure all prefixes will have the same length of the AS path attribute value.

B. Redistribute the external prefixes onto OSPF and ensure that the total metric calculation includes external and internal values.

C. Enable OSPF load-balancing over unequal cost path.

D. Redistribute the external prefixes onto OSPF and ensure the total metric calculation includes only the external value and the value is the same in all ASBRs.

Answer: B

Explanation:

Okay, let's break down why option B is the correct answer for implementing hot-potato routing with OSPF in Company XYZ's network.

Hot-potato routing, also known as exit-point selection, aims to hand off traffic to an external network as quickly as possible, minimizing its travel within the internal network. This is achieved by selecting the closest exit point, usually measured in terms of the routing protocol's metric.

Option B describes redistributing external prefixes into OSPF and ensuring the OSPF metric calculation prioritizes the external metric equally across all Area Border Routers (ASBRs). This setup effectively guides all routers within the OSPF domain to choose the ASBR closest to them for traffic destined for those external networks. Essentially, the metric representing the distance to the outside world is the dominating factor, and because it's the same everywhere, the choice is based solely on the distance to the ASBR.

Option A, involving iBGP and AS path manipulation, pertains to BGP routing, not OSPF and would not accomplish hot-potato routing within OSPF. Option C, unequal-cost load balancing, aims for distribution across paths regardless of the exit point and is not aligned with the principles of hot-potato routing. Option D is partially correct in redistributing and focusing on external metrics but incorrectly indicates that only the external value matters, not the total metric which includes distance to the ASBR. It also includes an illogical requirement that external values are the same in all ASBRs, implying the network has a very specific topology.

By focusing on the total metric, but ensuring the external component is set the same across all ASBRs, OSPF routers effectively choose the ASBR offering the shortest internal path, thus realizing hot-potato routing. This optimizes path selection by quickly exiting the internal OSPF domain rather than traversing it extensively. This approach is critical for optimal traffic flow and resource usage in larger, multi-AS networks, especially for reducing congestion by avoiding long backhauls.

For more information, you can research these topics using these authoritative links:

1. **OSPF Routing Protocol:** Cisco documentation is a great resource. Search for "Cisco OSPF metric calculation" for details on how OSPF computes its routing metrics.
2. **Hot Potato Routing:** Networking books and papers discussing routing concepts will shed more light on its advantages and implementations, like those from "TCP/IP Illustrated."

In summary, redistributing external routes into OSPF and ensuring the total cost calculation prioritizes the external metric equally across all ASBRs is the correct method for implementing hot-potato routing within an OSPF network.

Question: 39

How many fully established neighbour relationships exist on an Ethernet with five routers running OSPF as network type broadcast?

A. 5

B. 6

C. 7

D. 10
E. 20

Answer: C

Explanation:

In a broadcast OSPF network, routers elect a Designated Router (DR) and a Backup Designated Router (BDR). All other routers become DR Others (DROther). The DR and BDR form full neighbor adjacencies with each other. Subsequently, each DROther forms a full adjacency with the DR and BDR, but not with other DROthers. With five routers, one will be DR, one will be BDR, and the remaining three will be DROthers. The DR forms an adjacency with the BDR (1 adjacency). Then, the DR forms an adjacency with each of the three DROthers (3 adjacencies), and the BDR also forms an adjacency with each of the three DROthers (3 adjacencies). Therefore, the total number of full adjacencies becomes 1 + 3 + 3 = 7. Hence, option C is the correct answer.

The underlying principle here is that broadcast OSPF networks optimize traffic by limiting the number of full adjacencies, reducing the amount of LSA (Link State Advertisement) flooding. Instead of each router forming an adjacency with every other router, the DR and BDR act as central hubs for information exchange. This strategy enhances the efficiency of the OSPF protocol by significantly decreasing unnecessary routing updates across the network. Without DR/BDR, the 5 routers would form 10 adjacencies. This highlights how OSPF's DR/BDR election plays a key role in network scalability and performance. Specifically, in networks with many routers, the use of a DR/BDR significantly decreases the complexity of routing operations.

For further research, the following resources are authoritative:

Cisco documentation on OSPF network types: https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html
OSPF DR and BDR Election Process: https://networklessons.com/ospf/ospf-dr-bdr-election-process

Question: 40

How can EIGRP topologies be designed to converge as fast as possible in the event of a point-to-point link failure?

A. Limit the query domain by use of summarization.
B. Limit the query domain by use of default routes.
C. Build neighbor adjacencies in a squared fashion.
D. Limit the query domain by use of distribute lists.
E. Build neighbor adjacencies in a triangulated fashion.

Answer: A

Explanation:

EIGRP convergence speed is primarily impacted by the scope of query propagation when a route is lost. A large query domain, where many routers need to be contacted to find an alternate path, slows down convergence. Summarization, or route aggregation, effectively limits this query domain. By summarizing routes, routers only advertise the aggregate prefixes, masking specific details. When a link fails, a router only needs to query for the summary route, not each individual, underlying prefix. This reduction in scope confines queries to a smaller set of routers, reducing the total number of queries and overall convergence time. Default routes, while offering a simpler routing table, don't specifically limit the EIGRP query domain; instead, they redirect traffic for unknown destinations, not impacting query scope. Building neighbor adjacencies in squared or triangulated fashions doesn't directly address the issue of EIGRP query propagation. Distribute lists, used to filter routing updates, may help limit the routing table size but aren't meant to reduce the size of an EIGRP query domain. Therefore, limiting the query domain using summarization directly reduces the

number of routers that EIGRP must query, leading to the fastest possible convergence in the event of a point-to-point link failure.

Further Research:

1. **Cisco Documentation on EIGRP Summarization:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/15-mt/ire-15-mt-book/ire-sumry.html
2. **EIGRP Query Process Explanation:**https://www.ciscopress.com/articles/article.asp?p=27046&seqNum=5

Question: 41

DRAG DROP -
Drag and drop the multicast protocols from the left onto the correct design situations on the right.
Select and Place:

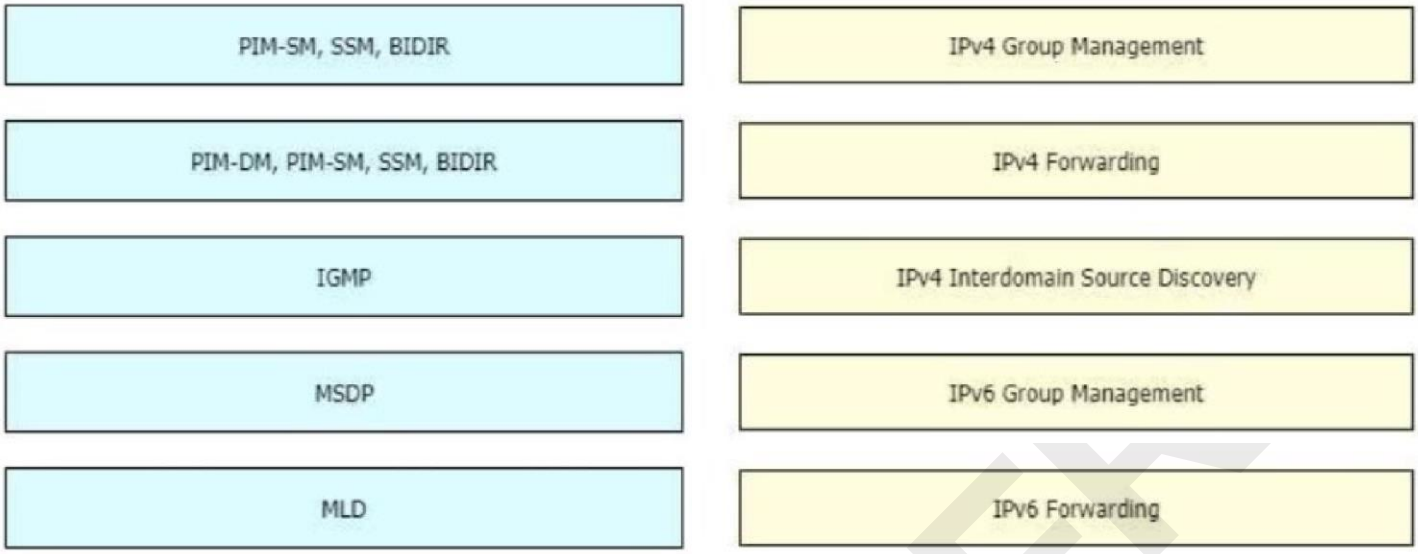

Answer:

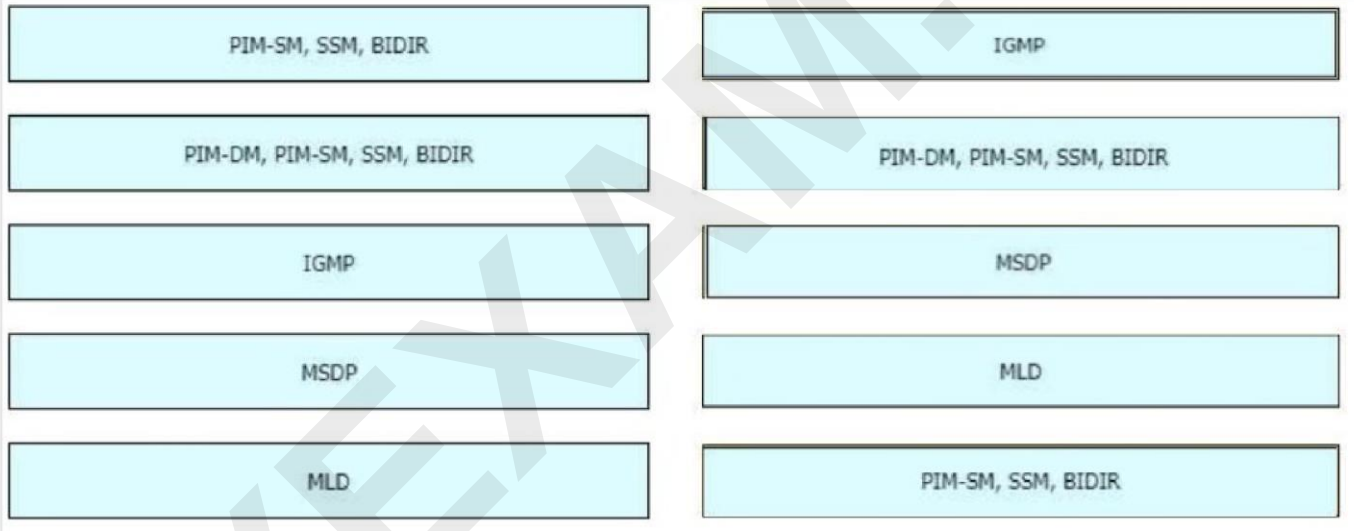

Question: 42

Company XYZ, a global content provider, owns data centers on different continents. Their data center design involves a standard three-layer design with a Layer
3-only core. HSRP is used as the FHRP. They require VLAN extension across access switches in all data centers,

and they plan to purchase a Layer 2 interconnection between two of their data centers in Europe. In the absence of other business or technical constraints, which termination point is optimal for the
Layer 2 interconnection?

A. at the core layer, to offer the possibility to isolate STP domains

B. at the aggregation layer because it is the Layer 2 to Layer 3 demarcation point

C. at the access layer because the STP root bridge does not need to align with the HSRP active node

D. at the core layer because all external connections must terminate there for security reasons

Answer: B

Explanation:

The optimal termination point for the Layer 2 interconnection is at the **aggregation layer**. This is because the aggregation layer traditionally serves as the Layer 2 to Layer 3 boundary in a three-tier network design. Connecting the Layer 2 link here allows for seamless VLAN extension between the access switches in the two data centers without requiring core layer involvement in Layer 2 forwarding. Option A, terminating at the core, would likely create a large STP domain which could increase convergence time in the event of failures and is not recommended. Option C, connecting at the access layer, might not be practical as it adds unnecessary complexity and could lead to suboptimal routing paths. Option D is incorrect; while security is vital, that does not dictate where a Layer 2 link should terminate, and core layer connections are not always necessary.

Terminating at the aggregation layer provides the best balance between network simplicity, scalability, and efficiency. By keeping the Layer 2 extension contained within the aggregation and access layers, the core remains Layer 3 focused, which aligns with the provided network design. This approach also allows the organization to manage the STP domain more effectively, as it is limited to the access and aggregation layers.

Authoritative Links for Further Research:

Cisco Enterprise Networks: https://www.cisco.com/c/en/us/solutions/enterprise-networks/index.html Cisco Design Guides: https://www.cisco.com/c/en/us/solutions/enterprise/design-zone/index.html
Understanding Layer 2 and Layer 3 Networking: https://www.cloudflare.com/learning/network-layer/what-is-layer-2/ and https://www.cloudflare.com/learning/network-layer/what-is-layer-3/

Question: 43

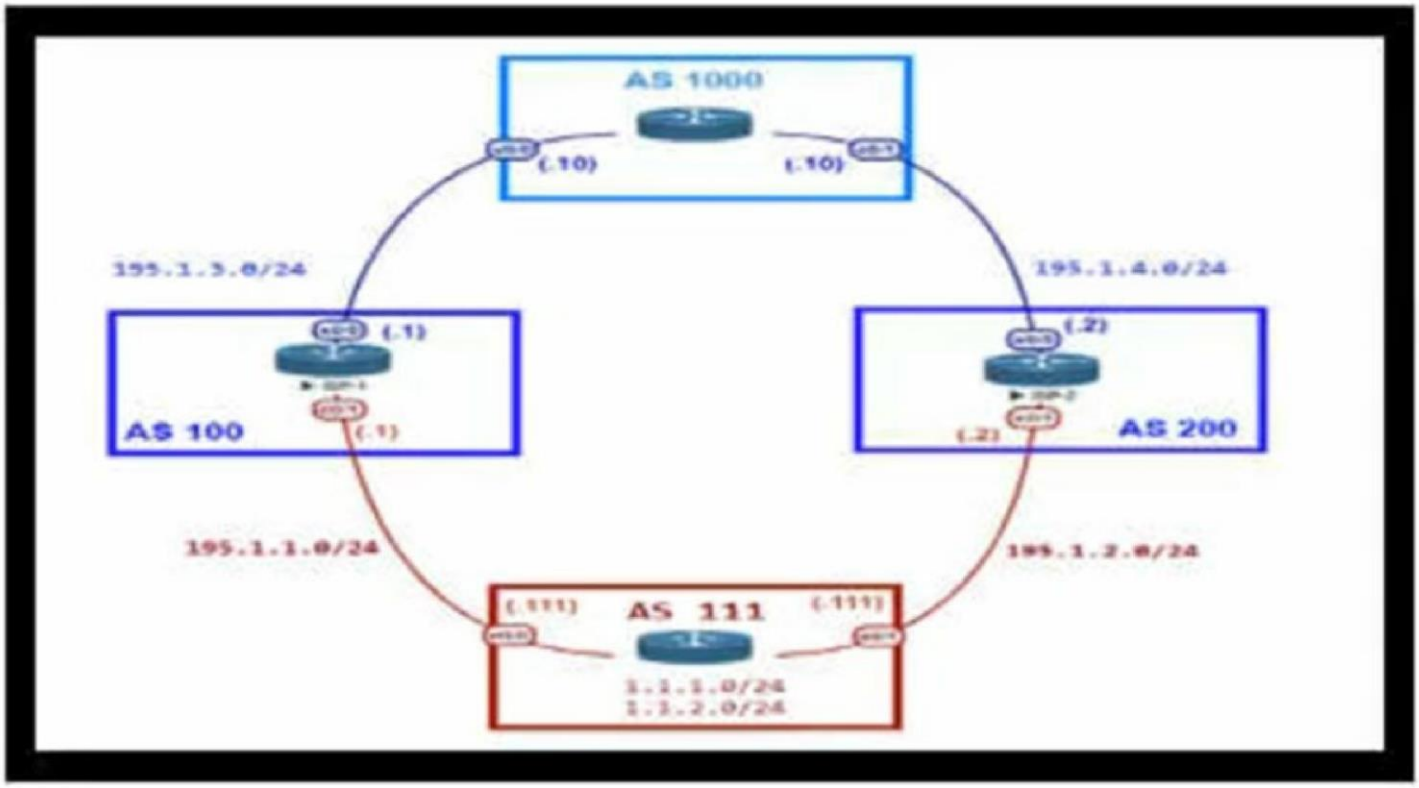

Refer to the exhibit. An engineer is designing the network for a multihomed customer running in AS 111 AS 111 does not have any other ASs connected to it.
Which technology is more comprehensive to use in the design to make sure that the AS is not being used as a transit AS?

A. Include an AS path access list to send routes to the neighboring ASs that only have AS 111 in the AS path field.

B. Configure the AS-set attribute to allow only routes from AS 111 to be propagated to the neighbor ASs. C. Use the local preference attribute to configure your AS as a "non-transit" AS.

D. Include a prefix list to only receive routes from neighboring ASs.

Answer: A

Explanation:

Include an AS path access list to send routes to the neighboring ASs that only have AS 111 in the AS path field.

Question: 44

Which interconnectivity method offers the fastest convergence in the event of a unidirectional issue between three Layer 3 switches connected together with routed links in the same rack in a data center?

A. Fiber Ethernet connectivity with UDLD enabled
B. Copper Ethernet connectivity with UDLD enabled
C. Fiber Ethernet connectivity with BFD enabled
D. Copper Ethernet connectivity with BFD enabled

Answer: C

Explanation:

The correct answer is **C. Fiber Ethernet connectivity with BFD enabled**. Here's why:

Convergence speed, especially in the face of a unidirectional link failure, is critical in data center

environments to maintain service availability. The primary goal is swift detection and reaction to such issues. UDLD (Unidirectional Link Detection) can detect unidirectional link problems but relies on periodic hello messages, which results in slower detection times compared to BFD.

BFD (Bidirectional Forwarding Detection) is specifically designed for rapid failure detection. It uses aggressive hello timers (often in milliseconds) and is independent of the routing protocol. BFD can detect a loss of connectivity significantly faster than UDLD or standard routing protocol timers. This rapid detection allows routing protocols to quickly react and converge on a new path.

Fiber Ethernet connectivity, as opposed to copper, inherently has less physical issues that can cause unidirectional failures. However, this doesn't address faster convergence of a unidirectional failure itself like the use of BFD would. While the link medium isn't as directly relevant to convergence speed as the BFD vs UDLD protocol, when considering the context, the question is related to a potential unidirectional issue, the use of fiber would lead to less likelihood of a failure of this kind. Copper, in this case, doesn't give any convergence benefits over fiber.

Therefore, the combination of fiber (for reliability in the link itself) and BFD (for rapid failure detection) provides the fastest convergence in the given scenario. BFD's microsecond timing capability ensures that even transient or minute unidirectional issues are rapidly identified and addressed, allowing the network to swiftly reroute traffic. In contrast, UDLD and traditional routing timers would introduce significant delays in such a critical data center environment.

Authoritative Links for Further Research:

Cisco - Bidirectional Forwarding Detection:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/15-mt/ipapp-15-mt-book/ipapp-bidir-forward.html
Cisco - Unidirectional Link Detection:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/configuration/xe-16/ir-xe-16-book/ir-unidir-link-detec.html
Juniper Networks - Understanding Bidirectional Forwarding Detection
https://www.juniper.net/documentation/en_US/junos/topics/concept/bfd-overview.html

Question: 45

You want to mitigate failures that are caused by STP loops that occur before UDLD detects the failure or that are caused by a device that is no longer sending
BPDUs. Which mechanism do you use along with UDLD?

A. BPDU guard
B. root guard
C. loop guard
D. BPDU filtering

Answer: C

Explanation:

The correct answer is **C. loop guard**. Here's a detailed justification:

Loop guard is specifically designed to prevent network loops that can arise when a switch stops receiving Bridge Protocol Data Units (BPDUs) on a port, either due to a unidirectional link failure or a misconfigured or malfunctioning device. UDLD, while excellent at detecting unidirectional link issues, operates separately from STP's loop prevention mechanisms. If a device stops sending BPDUs on a normally forwarding port (for example, due to a software issue), STP may not recalculate, leading to a potential loop if an alternate path exists. In such cases, STP assumes that a port that has not received BPDU information is actually the best path. Loop guard steps in by disabling a port that stops receiving BPDUs. It does not act on the root path or

block new root advertisement as BPDU guard and root guard do respectively. BPDU filtering disables STP operation, removing the benefits of STP and other mechanisms like UDLD and thus makes it unsuitable.

Loop guard works at the interface level. When the port is receiving BPDUs, the port will operate as normal. However, if BPDU messages aren't received for a time that exceeds max aging time, the loop guard will put the port into an inconsistent state, effectively shutting down the port and breaking the potential for a STP loop. Thus, combining UDLD, which prevents unidirectional links, and loop guard, which manages non-BPDU receiving ports, provides the strongest mitigation to STP loop issues in the network.

Here are some authoritative links for further research:

Cisco Documentation on Loop Guard:https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/24064-loopguard.html
Cisco Documentation on UDLD:https://www.cisco.com/c/en/us/support/docs/lan-switching/ethernet/10590-udld.html
Understanding Spanning Tree Protocol:https://www.networkworld.com/article/3238282/lan-wan/spanning-tree-protocol-stp-explained.html

Question: 46

Company XYZ needs advice in redesigning their legacy Layer 2 infrastructure. Which technology should be included in the design to minimize or avoid convergence delays due to STP or FHRP and provide a loop-free topology?

A. Use BFD.
B. Use switch clustering in the core/distribution layer.
C. Use spanning-tree PortFast.
D. Use switch clustering in the access layer.

Answer: B

Explanation:

The correct answer is B: Use switch clustering in the core/distribution layer. Switch clustering, often involving technologies like Virtual Switching System (VSS) or Multi-chassis EtherChannel (MEC), creates a single logical switch from multiple physical switches. This eliminates the need for traditional Spanning Tree Protocol (STP) at the core/distribution layer, as the clustered switches act as a unified entity with a single control plane. Thus, convergence delays due to STP are avoided. Similarly, First Hop Redundancy Protocol (FHRP) convergence delays can also be minimized, as the clustered switches maintain synchronized forwarding tables and state information. This achieves a loop-free topology because the clustered devices present themselves as a single logical entity to other switches. By avoiding STP, the network can utilize all paths for forwarding, increasing bandwidth and resilience. Option A, BFD, is a failure detection protocol, but it doesn't address STP convergence delays or loop avoidance. PortFast, option C, speeds up convergence for access ports, but doesn't address the core/distribution issues. Switch clustering in the access layer (option D) doesn't alleviate the problem in the core/distribution, where STP is most problematic. Therefore, switch clustering in the core/distribution is the most effective approach for minimizing convergence delays related to STP and FHRP while creating a loop-free topology.

Authoritative Links:

Cisco Virtual Switching System (VSS):
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SX/configuration/guide/vss.html
(Although documentation for older 6500 platform, the conceptual understanding is still valid)
Multi-chassis EtherChannel (MEC):https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9000/16-

6/configuration/guide/b_166_config_9000/b_166_config_9000_chapter_01111.html

Question: 47

DRAG DROP -
Drag and drop the multicast protocols from the left onto the correct design situations on the right. Not all options are used.
Select and Place:

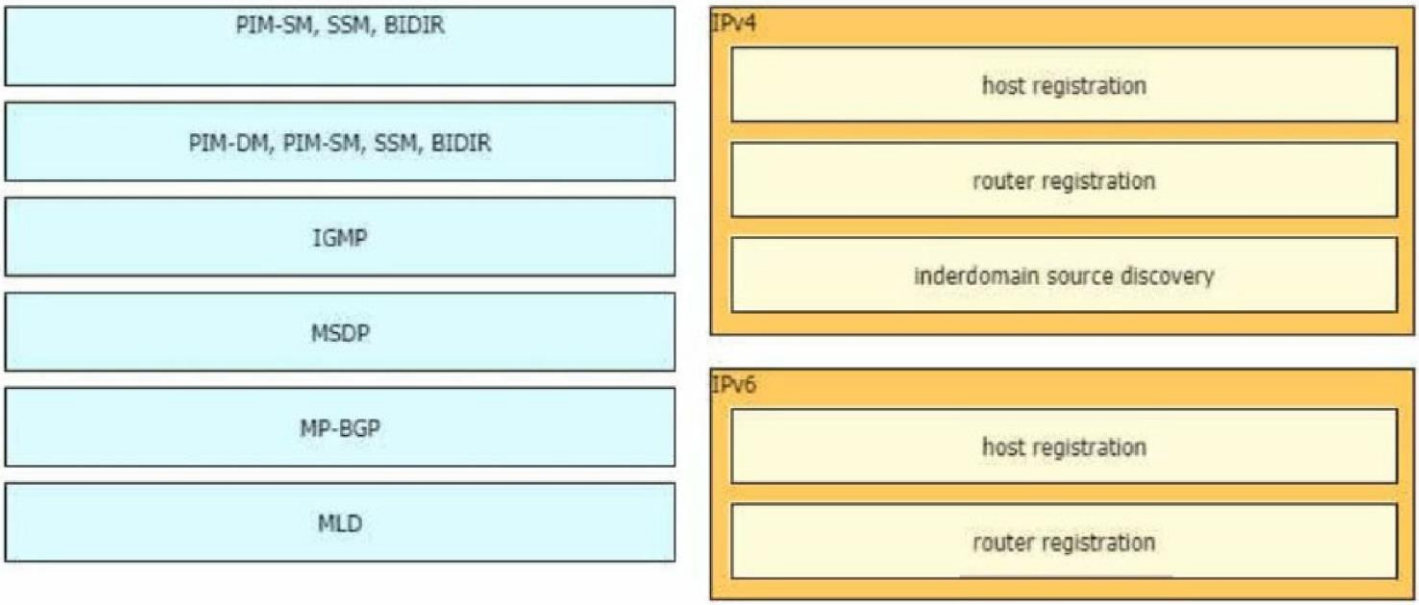

Answer:

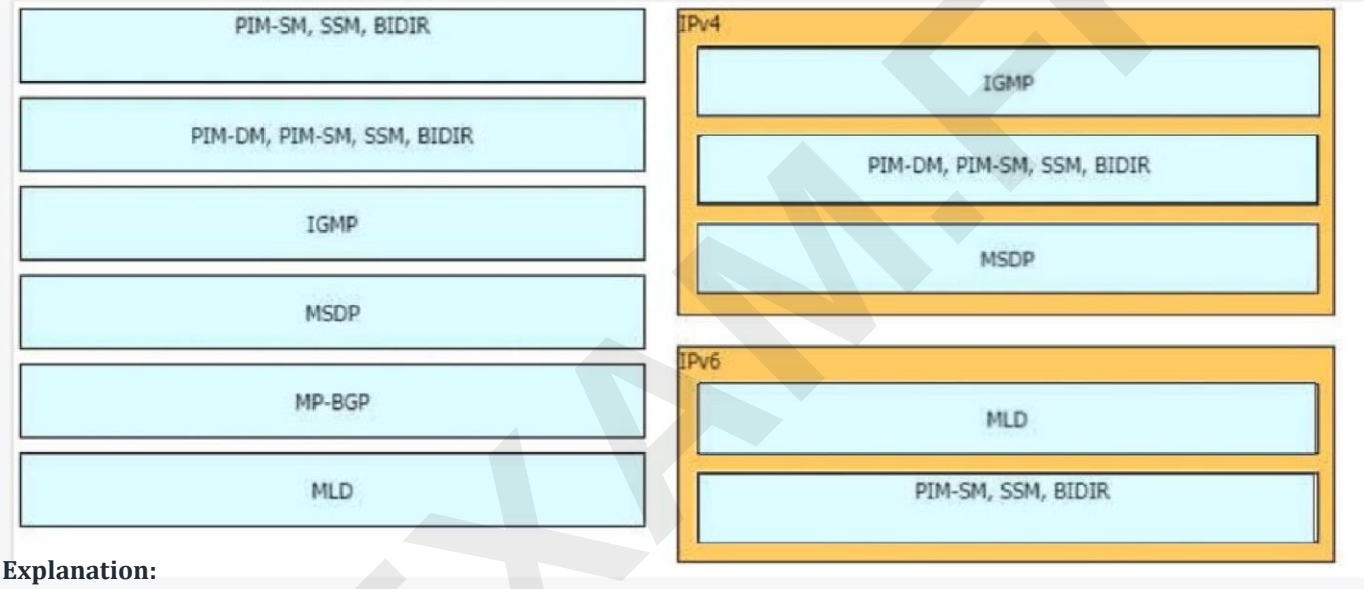

Explanation:

(PIM-DM PIM-SM SSM BIDIR for IPv4, PIM-SM SSM BIDIR for IPv6)

Question: 48

Which function is performed at the access layer of the three-layer hierarchical network design model?

A. fast transport

B. reliability
C. fault isolation
D. redundancy and load balancing
E. QoS classification and marking boundary

Answer: E

Explanation:

The access layer in the three-layer hierarchical model focuses on connecting end-user devices to the network. Its primary function is to provide network access and enforce basic policies. QoS classification and marking are crucial at this layer because it's the point where traffic enters the network. By classifying packets based on application or user type, and marking them with appropriate priority tags, the network can then prioritize critical traffic as it traverses higher layers. This early classification and marking allows the distribution and core layers to make intelligent forwarding and queuing decisions, ensuring a smooth user experience. While access layer devices might participate in some transport functions (A), reliability (B) and fault isolation (C) are mainly concerns of the distribution layer, and redundancy and load balancing (D) are primarily implemented at both the distribution and core layers. Therefore, only QoS classification and marking boundary (E) aligns with the essential function of the access layer.

Further reading can be found here:

Cisco's Hierarchical Network Design Model: https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Aug2018/CVD-CampusLANDesign-AUG2018.html QoS Overview: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/configuration/15-mt/qos-15-mt-book/qos-overview.html

Question: 49

Which two features control multicast traffic in a VLAN environment? (Choose two.)

A. RGMP
B. PIM snooping
C. MLD snooping
D. pruning
E. IGMP snooping

Answer: CE

Explanation:

The correct answer is **C. MLD snooping and E. IGMP snooping**. These two features are crucial for controlling multicast traffic within a VLAN environment.

IGMP Snooping (Internet Group Management Protocol Snooping) operates at the Layer 2 switch level. It listens to IGMP messages exchanged between hosts and multicast routers. By doing so, the switch learns which ports have hosts interested in specific multicast groups. It then selectively forwards multicast traffic only to those ports, preventing unnecessary flooding across the VLAN. This optimizes bandwidth usage and reduces network congestion. Without IGMP snooping, a switch would treat multicast like broadcast, sending it to every port regardless of whether a host is interested or not.

MLD Snooping (Multicast Listener Discovery Snooping) functions similarly to IGMP snooping, but it manages multicast traffic for IPv6. It listens to MLD messages to understand which interfaces on the switch are interested in particular IPv6 multicast groups. Just like IGMP snooping for IPv4, MLD snooping prevents

unnecessary forwarding of multicast traffic, ensuring that only intended recipients receive the data.

Both IGMP and MLD snooping enhance network efficiency in multicast environments. They prevent flooding, conserve bandwidth, and allow switches to intelligently manage multicast data based on host interest.

The other options are incorrect: RGMP isn't a standard multicast protocol, PIM snooping doesn't exist (PIM is a routing protocol), and pruning is a broader multicast concept not specific to VLAN control.

Authoritative Links:

Cisco on IGMP Snooping:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960/software/release/15-0_2_se/configuration/guide/scg2960/swigmp.html
Cisco on MLD Snooping:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/configuration/15-mt/imc-15-mt-book/imc-mld-snoop.html

Question: 50

In an OSPF network with routers connected together with Ethernet cabling, which topology typically takes the longest to converge?

A. squared
B. ring
C. partial mesh
D. triangulated
E. full mesh

Answer: B

Explanation:

The ring topology in OSPF networks typically experiences the longest convergence times compared to other common topologies. This is because in a ring, a single link failure breaks the path between multiple routers, forcing OSPF to recalculate routes through a longer, alternate path. In contrast, a full mesh offers multiple redundant paths so a single failure has minimal impact. A partial mesh provides some redundancy but not as much as a full mesh or as direct as triangulated. A squared topology, while having a structured arrangement, typically has faster convergence compared to a ring. Convergence in OSPF involves several steps: detecting the link failure, flooding the Link State Advertisement (LSA), updating the link-state database, and finally recomputing the Shortest Path First (SPF) algorithm. The time required for this process is directly influenced by the number of routers and links involved in recalculations. In a ring, all routers must be part of the calculation following a break, whereas a mesh or triangulated topology may not require as many. The "hop count" tends to be longer in a ring path, further contributing to delayed convergence. OSPF's reliance on Dijkstra's algorithm for SPF makes it converge faster in topologies with more direct, readily available, alternative routes, which rings lack. This slow convergence in ring topology is a known disadvantage, making it less suitable for time-sensitive applications relying on network stability.

Authoritative links for further research:

1. **OSPF Overview:**https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html
2. **OSPF Convergence:**https://www.juniper.net/documentation/en_US/junos/topics/concept/ospf-convergence.html
3. **Network Topologies:**https://www.geeksforgeeks.org/types-of-network-topology/

Question: 51

An enterprise network has two core routers that connect to 200 distribution routers and uses full-mesh iBGP peering between these routers as its routing method.
The distribution routers are experiencing high CPU utilization due to the BGP process. Which design solution is the most cost effective?

A. Increase the memory on the core routers.
B. Increase bandwidth between the core routers.
C. Implement eBGP between the core and distribution routers.
D. Increase the memory on the distribution routers.
E. Implement route reflectors on the two core routers.

Answer: E

Explanation:

The problem describes a full-mesh iBGP network with 202 routers (2 core, 200 distribution) leading to high CPU on the distribution routers. Full-mesh iBGP requires each router to peer with every other iBGP router, resulting in n*(n-1)/2 adjacencies. With 202 routers, this leads to a massive number of BGP sessions and route advertisements, causing significant processing overhead, especially on the 200 distribution routers.

Increasing memory on the core (A) or distribution routers (D), or increasing bandwidth between cores (B) would not address the root cause: the exponential scaling of iBGP sessions. These options merely provide more resources to handle the inefficient design. Switching to eBGP between core and distribution routers (C) introduces complexity and is not ideal for inter-domain routing within the same autonomous system.

Implementing route reflectors (RRs) on the two core routers (E) is the most cost-effective solution. RRs allow a single point of BGP peering instead of a full mesh. Distribution routers peer only with the RRs, reducing the number of BGP sessions. The RRs reflect routes to other distribution routers, thus alleviating the CPU load and simplifying iBGP. This is a standard practice for scaling iBGP in large networks.

Cisco Route Reflector Configuration GuideUnderstanding Route Reflectors

Question: 52

Which purpose of a dynamically created tunnel interface on the design of IPv6 multicast services is true?

A. first-hop router registration to the RP
B. multicast source registration to the RP
C. multicast client registration to the RP
D. transport of all IPv6 multicast traffic

Answer: D

Explanation:

Let's break down why option D, "transport of all IPv6 multicast traffic," is the correct answer regarding dynamically created tunnel interfaces in IPv6 multicast design. Dynamically created tunnel interfaces, often seen in sparse-mode multicast implementations like PIM-SM (Protocol Independent Multicast - Sparse Mode), are not primarily for direct registration of first-hop routers, sources, or clients with the Rendezvous Point (RP).

Those registrations are handled via other PIM messages. Instead, these tunnels function as data transport paths. When a first-hop router receives multicast traffic from a source, it encapsulates that multicast traffic

within an IP tunnel and forwards it to the RP using these interfaces. Subsequently, the RP uses these tunnels to distribute the multicast traffic down the shared tree to interested receivers. The same tunnel interface can also be used to transport traffic from sources using the shortest-path tree after the receiver performs a join. Hence, these dynamically established tunnels act as conduits for all IPv6 multicast traffic flowing through the shared or shortest-path trees, facilitating efficient distribution of multicast streams. Therefore, the primary purpose of a dynamically created tunnel interface in IPv6 multicast design is, indeed, the transport of all IPv6 multicast traffic, not the specific registrations detailed in the other options.

For further research, refer to the following authoritative Cisco documentation on multicast:

1. **Cisco Multicast Deployment Guide:** https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/WAN_and_MAN/Multicast_Ent.html 2. **Understanding PIM Sparse Mode:**https://www.cisco.com/c/en/us/support/docs/ip/protocol-independent-multicast-pim/13695-pim-sparse.html

Question: 53

Company XYZ network runs IPv4 and IPv6 and they want to introduce a multidomain, multicast-based network. The new design should use a flavor of PIM that forwards traffic using SPT. Which technology meets this requirement?

A. PIM-SSM
B. PIM-SM
C. BIDIR-PIM
D. PIM-DM

Answer: A

Explanation:

The correct answer is **A. PIM-SSM (Protocol Independent Multicast - Source Specific Multicast)**.

PIM-SSM is the only PIM variant that inherently builds and forwards multicast traffic based solely on the shortest path tree (SPT). This behavior is critical for multicast forwarding in multidomain environments where optimal paths are crucial. In PIM-SSM, a receiving host must explicitly indicate both the multicast group address (G) it wishes to receive and the source address (S) from which it expects to receive traffic, forming (S,G) channels. Consequently, routers build multicast distribution trees from the receiver to the specific source. PIM-SM (Sparse Mode) and PIM-DM (Dense Mode) initially use shared trees and then transition to shortest path trees. BIDIR-PIM, as its name implies, operates using bidirectional trees, making it unsuitable for pure SPT-based forwarding. Therefore, PIM-SSM's fundamental design perfectly aligns with the requirement for a multidomain, multicast network that relies solely on shortest path trees, establishing a direct path between each receiver and the respective source. Further research can be done at the following links:RFC 4607,Cisco - Understanding PIM-SSM.

Question: 54

Company XYZ has 30 sites running a legacy private WAN architecture that connects to the Internet via multiple high-speed connections. The company is now redesigning their network and must comply with these design requirements:
* Use a private WAN strategy that allows the sites to connect to each other directly and caters for future expansion
* Use the Internet as the underlay for the private WAN
* Securely transfer the corporate data over the private WAN
Which two technologies should be incorporated into the design of this network? (Choose two.)

A. PPTP
B. DMVPN
C. IPsec
D. GET VPN
E. S-VTI

Answer: BC

Explanation:

The correct answer is **B. DMVPN** and **C. IPsec**. Here's why:

DMVPN (Dynamic Multipoint VPN): This technology is ideally suited for creating a scalable, hub-and-spoke or spoke-to-spoke private WAN over the internet. DMVPN allows for dynamic creation of VPN tunnels between sites as needed, rather than requiring pre-configured, static tunnels. This is crucial for scalability and supports future expansion, aligning with the requirement for direct site-to-site connections. In the given scenario, multiple sites needing to connect to each other dynamically, rather than only through a centralized location, DMVPN offers a flexible and efficient solution.

IPsec (Internet Protocol Security): IPsec provides the necessary encryption and authentication mechanisms to securely transfer data over the internet. It is a core component of any secure VPN solution and complements DMVPN perfectly. IPsec secures the tunnels established by DMVPN, ensuring that all data transmitted over the private WAN is protected from eavesdropping and tampering. This is crucial to meeting the security requirement of securely transferring corporate data.

Why the other options are incorrect:

A. PPTP (Point-to-Point Tunneling Protocol): PPTP is an older VPN protocol that is considered insecure and should be avoided. It lacks strong encryption and is vulnerable to known attacks, making it unsuitable for securing sensitive corporate data.

D. GET VPN (Group Encrypted Transport VPN): While GET VPN offers secure, group-based encryption, it's typically used in scenarios with a more complex key management and a pre-defined topology, usually more complex than the use case described. While it is a possibility for private WANs, it is not the best fit for this scenario. DMVPN is a better option for a scalable and dynamic environment where the underlying infrastructure is the public internet.

E. S-VTI (Static Virtual Tunnel Interface): S-VTI is a technology that creates static tunnels between two locations. This option does not meet the requirement of dynamic connectivity and scalability in private WANs, as it requires manually configuring tunnels between each endpoint, which is not an efficient solution.

In summary: DMVPN provides the dynamic and scalable tunnel creation needed for a private WAN over the internet, while IPsec secures the data transmitted through these tunnels. Together, they meet the design requirements for the described scenario.

Authoritative links for further research:

Cisco Documentation on DMVPN:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/vpn/configuration/15-mt/sec-vpn-15-mt-book/sec-dmvpn.html
Cisco Documentation on IPsec:https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike/116096-technote-ipsec-00.html
DMVPN Overview:https://www.firewall.cx/cisco-technical-knowledgebase/cisco-routers/492-cisco-router-dmvpn-basic-config.html

Question: 55

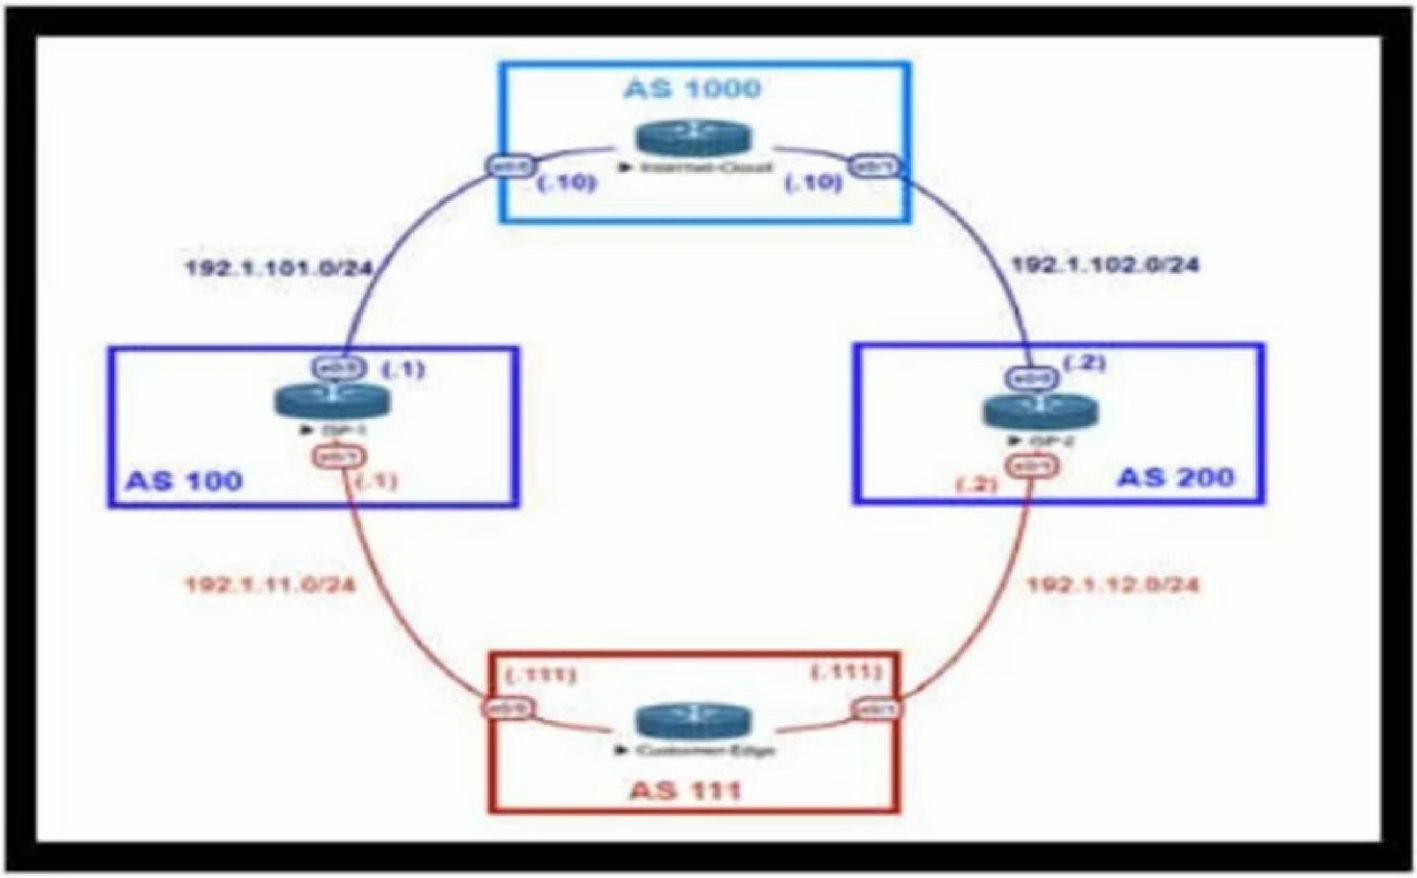

Refer to the exhibit. An engineer is designing the traffic flow for AS 111. Traffic from AS 111 should be preferred via AS 100 for all external routes. A method must be used that only affects AS 111. Which BGP attributes are best suited to control outbound traffic?

A. MED
B. community
C. local preference
D. AS path

Answer: C

Explanation:

local preference is a correct answer.

Question: 56

Which BGP feature provides fast convergence?

A. BGP-LS
B. BGP FlowSpec
C. BGP-EVPN
D. BGP PIC

Answer: D

Explanation:

BGP Prefix Independent Convergence (PIC) is the BGP feature that primarily focuses on achieving fast convergence. Traditional BGP convergence relies on withdrawing and re-announcing routes upon path failures, which can introduce delays. BGP PIC, however, maintains multiple pre-computed paths for each

destination, effectively providing a backup path that can be activated immediately upon the failure of the primary route, thus significantly reducing convergence time. BGP-LS (BGP Link State) is primarily used for network topology discovery and information distribution, not for route convergence speed. BGP FlowSpec is designed for traffic filtering and rate-limiting based on traffic characteristics, and doesn't directly impact convergence. BGP-EVPN (Ethernet VPN) is a technology used for MAC address learning and VPN services in Layer 2 networks, not for standard IP prefix convergence. Therefore, BGP PIC stands out for its
active/standby path model directly aimed at rapid failover and improved network stability. The other options do not provide immediate path switching as is implemented with PIC. In cloud environments, where rapid failure detection and response is critical, PIC is highly valuable for maintaining consistent service availability.

Further information:

Cisco BGP PIC Implementation: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/15-mt/irg-15-mt-book/bgp-pic.html
RFC 7342 - Enhanced Route Refresh Capability for BGP: https://datatracker.ietf.org/doc/html/rfc7342 (although this is not directly about PIC, understanding Route Refresh helps in understanding the need for PIC)

Question: 57

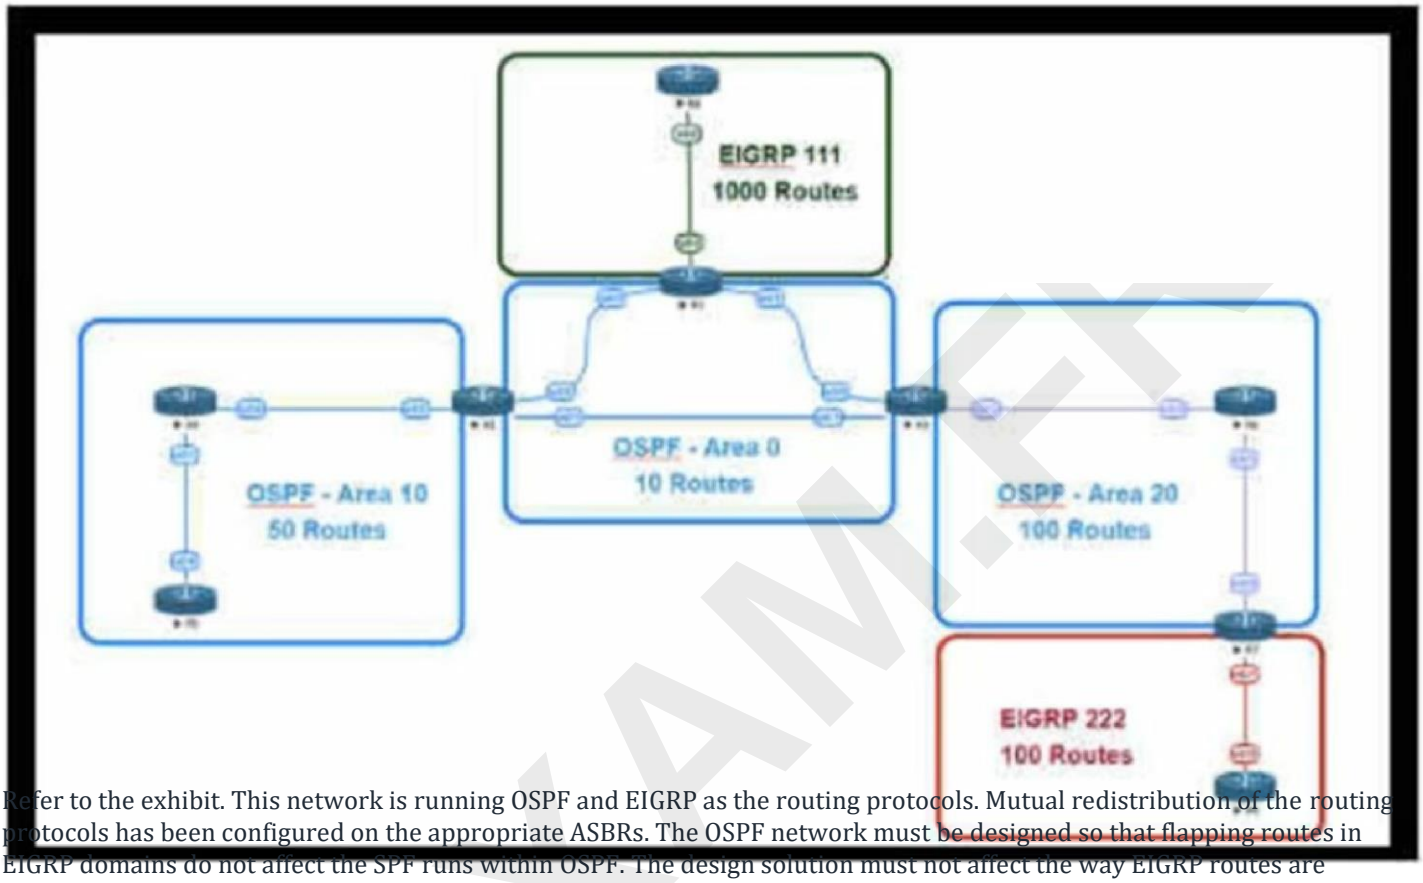

Refer to the exhibit. This network is running OSPF and EIGRP as the routing protocols. Mutual redistribution of the routing protocols has been configured on the appropriate ASBRs. The OSPF network must be designed so that flapping routes in EIGRP domains do not affect the SPF runs within OSPF. The design solution must not affect the way EIGRP routes are propagated into the EIGRP domains. Which technique accomplishes the requirement?

A. route summarization on the ASBR interfaces facing the OSPF domain
B. route summarization on the appropriate ABRs
C. route summarization on EIGRP routers connecting toward the ASBR
D. route summarization on the appropriate ASBRs

Answer: D

Explanation:

route summarization on the appropriate ASBRs

Question: 58

Which two mechanisms avoid suboptimal routing in a network with dynamic mutual redistribution between multiple OSPFv2 and EIGRP boundaries? (Choose two.)

A. AD manipulation
B. matching OSPF external routes
C. route filtering
D. matching EIGRP process ID
E. route tagging

Answer: CE

Explanation:

Here's a detailed justification for why options C (route filtering) and E (route tagging) are the correct choices for avoiding suboptimal routing in a network with dynamic mutual redistribution between OSPFv2 and EIGRP:

Mutual redistribution between routing protocols like OSPFv2 and EIGRP can lead to suboptimal routing due to metric differences and potential routing loops. Route filtering (C) addresses this by controlling which routes are advertised between the protocols. By carefully filtering routes, you can prevent the propagation of less desirable paths and ensure that traffic follows the optimal route. This mechanism allows precise control over the exchange of routing information.

Route tagging (E) is another vital mechanism. When routes are redistributed, they can be assigned tags. These tags can be used to identify the original source of the route, helping to prevent routing loops and enable specific routing policies to be applied. For example, a tagged route received through redistribution can be given a higher metric or be ignored to prevent it from looping back and overriding the original route.

Administrative Distance (AD) manipulation (A) is useful but is often a coarse adjustment mechanism that doesn't offer the granular control of route filtering or tagging. Matching OSPF external routes (B) and matching EIGRP process ID (D) are primarily used for internal identification, not directly to prevent suboptimal routing during redistribution.

Therefore, the combination of route filtering (C) and route tagging (E) is the most effective for mitigating the risk of suboptimal routing resulting from mutual redistribution between OSPFv2 and EIGRP.

Authoritative Links for further research:

Cisco Documentation on Route Redistribution:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/15-sy/irr-15-sy-book.html
Cisco Documentation on Route Tagging:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/15-mt/irr-15-mt-book.html
Cisco Documentation on OSPF and EIGRP:

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html

Question: 59

Company XYZ has a new network based on IPv6. Some of the subnets that they are planning to use will be confidential and need an addressing scheme that confines them to the local campus network. Which type of IPv6 addresses can be used for these networks in the IPv6 addressing design?

A. link-local addresses
B. private addresses
C. unique local addresses
D. local addresses

Answer: C

Explanation:

Unique Local Addresses (ULAs), specified by the prefix fc00::/7 with an allocation of fd00::/8, are designed for private addressing within a limited site or organization, aligning perfectly with the requirement of confining confidential subnets to a local campus network. Unlike Link-Local addresses (fe80::/10), which are used for local network segment communication and are not routable beyond that segment, ULAs are routable within the defined administrative domain of the organization. Private addresses, as seen in IPv4, do not exist in IPv6; instead, ULAs fulfill the need for private addressing. While "local addresses" might seem generally correct, in IPv6 context, "unique local addresses" is the precise term. Therefore, ULAs provide the desired characteristic of private, site-local scope, preventing external internet routing and maintaining the desired confidentiality of the subnets within the company's campus. This makes ULAs the correct choice for isolating confidential network segments within an IPv6 infrastructure, whereas the other options don't meet the given
requirements.

Supporting Links:

RFC 4193 - Unique Local IPv6 Unicast Addresses:https://datatracker.ietf.org/doc/html/rfc4193 **Cisco - Understanding IPv6 Addressing:**https://www.cisco.com/c/en/us/support/docs/ip/ip-version-6-ipv6/113324-ipv6-addr-overview.html
Wikipedia - Unique Local Address:https://en.wikipedia.org/wiki/Unique_local_address

Question: 60

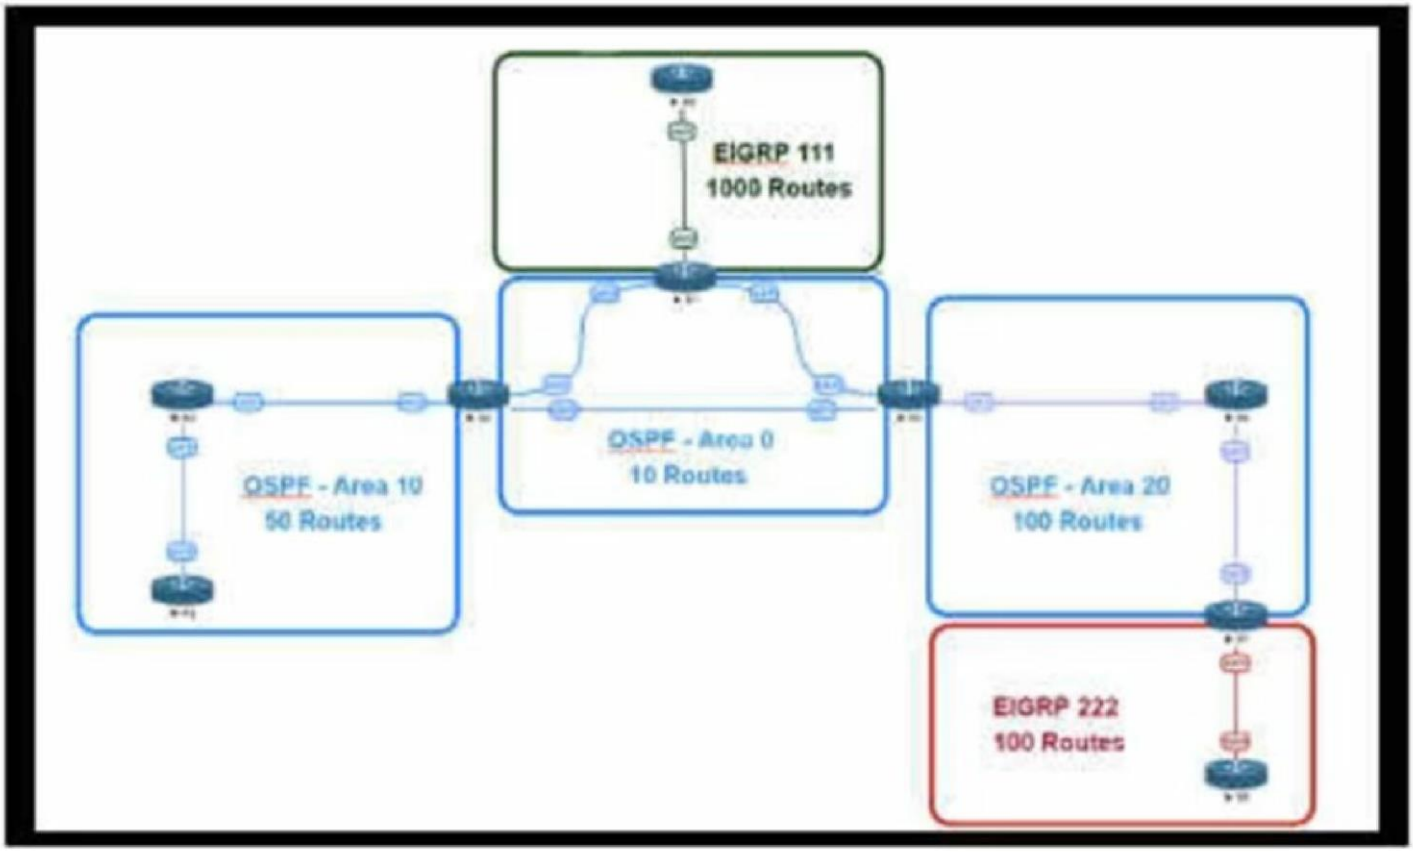

Refer to the exhibit. An engineer is designing a multiarea OSPF network for a client who also has a large EIGRP domain. EIGRP routes are getting redistributed into OSPF. OSPF area 20 has routers with limited memory and CPU resources. The engineer wants to block routes from EIGRP 111 from propagating into area
20 and allow EIGRP 222 routes to flow in. Which OSPF area type fulfills this design requirement?

A. type 3 LSA filtering on the ABR between area 0 and area 20
B. type 5 LSA filtering on the ASBR between EIGRP 111 and area 0
C. area 20 as a stub area
D. area 20 as a NSSA area

Answer: D

Explanation:

area 20 as a NSSA area

Question: 61

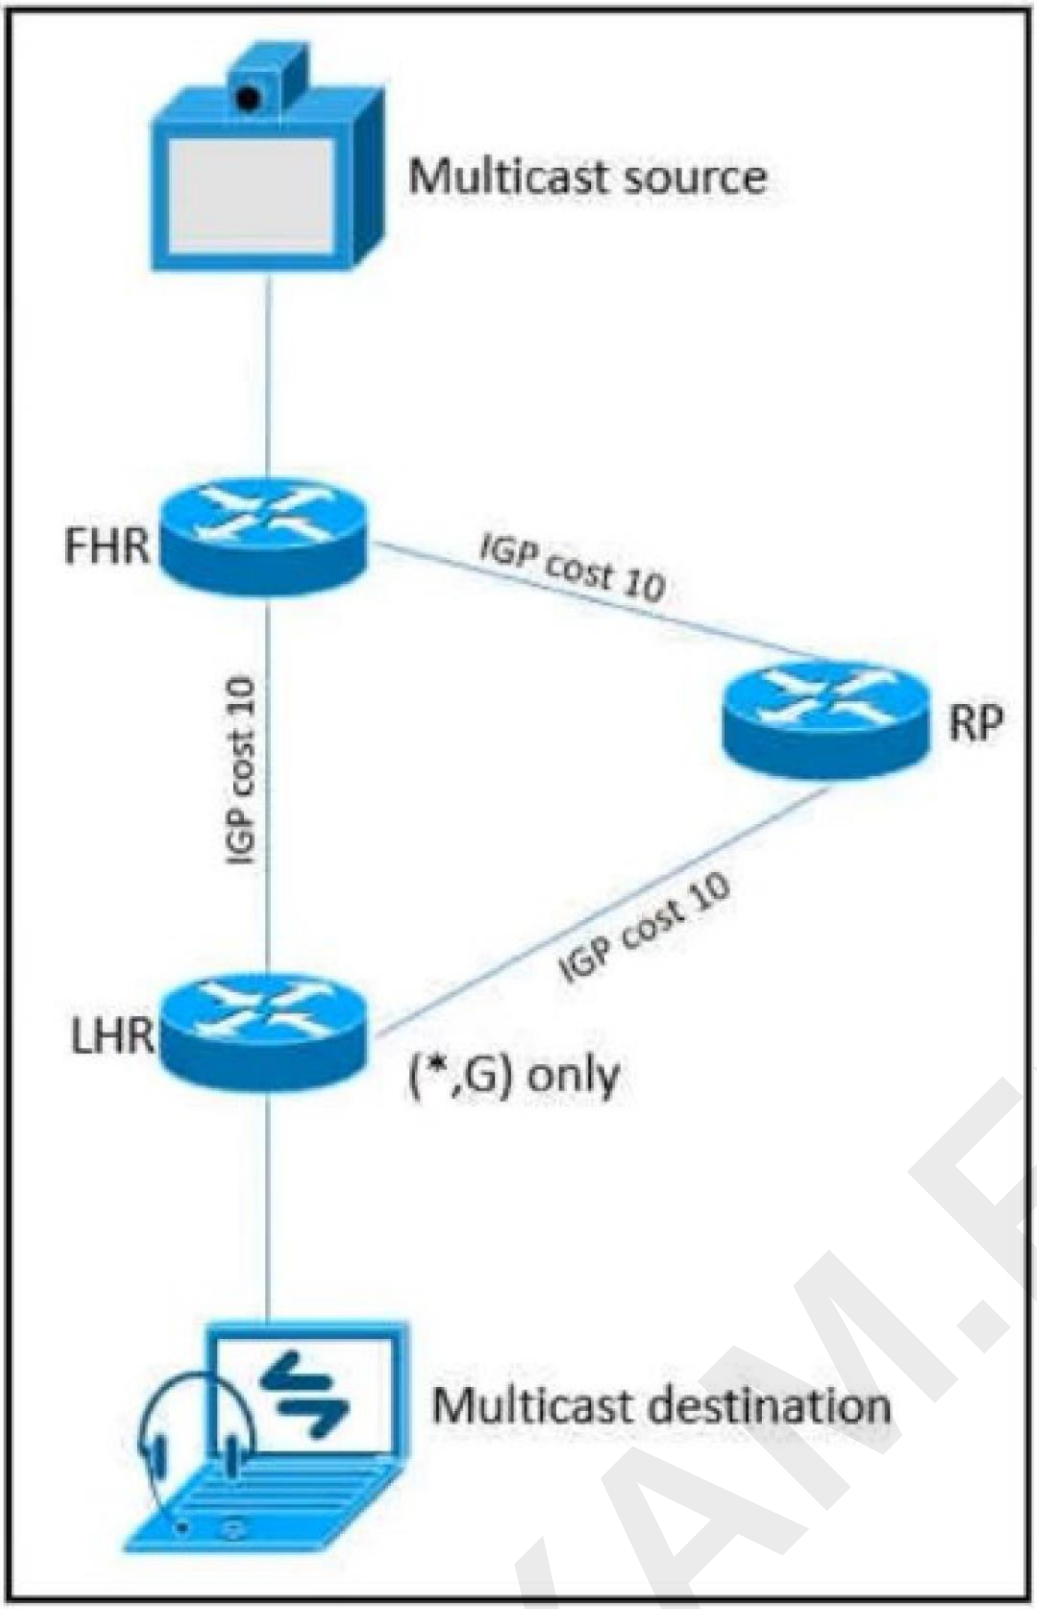

Refer to the exhibit. As part of a redesign project, you must predict multicast behavior. What happens to the multicast traffic received on the shared tree (*, G), if it is received on the LHR interface indicated?

A. It is switched due to a successful RPF check against the routing table.

B. It is switched given that no RPF check is performed.

C. It is dropped due to an unsuccessful RPF check against the multicast receiver. D. It is dropped due to an unsuccessful RPF check against the multicast source.

Answer: A

Explanation:

A. https://www.cisco.com/c/en/us/support/docs/ip/ip-multicast/16450-mcastguide0.html

When a multicast packet arrives on an interface, the RPF process checks to ensure that this incoming interface is the outgoing interface used by unicast routing in order to reach the source of the multicast packet. This RPF check process prevents loops. Multicast routing does not forward a packet unless the source of the packet passes a RPF check. Once a packet passes this RPF check, multicast routing forwards the packet based only upon the destination address.

Question: 62

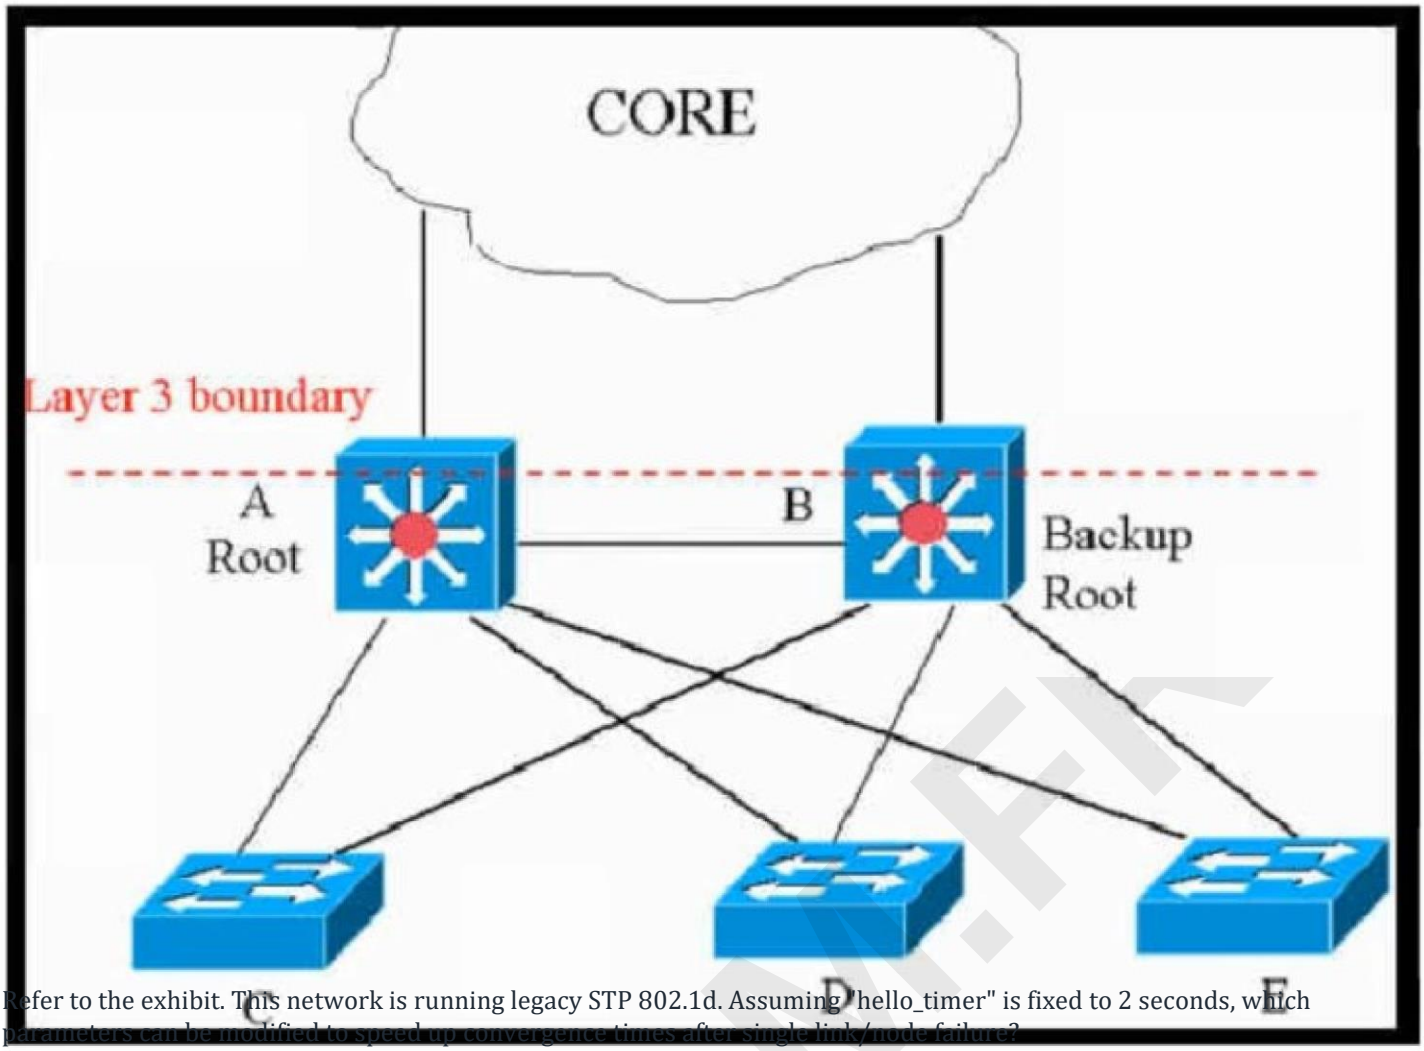

Refer to the exhibit. This network is running legacy STP 802.1d. Assuming "hello_timer" is fixed to 2 seconds, which parameters can be modified to speed up convergence times after single link/node failure?

A. Only the maximum_transmission_halt_delay and diameter parameters are configurable parameters in 802.1d to speed up STP convergence process.

B. The max_age and forward delay parameters can be adjusted to speed up STP convergence process.

C. The transit_delay=5 and bpdu_delay=20 are recommended values, considering hello_timer=2 and specified diameter.

D. Only the transit_delay and bpdu_delay timers are configurable parameters in 802.1d to speed up STP convergence process.

Answer: B

Explanation:

Reference:
https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/19120-122.html

Question: 63

Company XYZ is running a redundant private WAN network using OSPF as the underlay protocol. The current design accommodates for redundancy in the network, but it is taking over 30 seconds for the network to reconverge upon failure. Which technique can be implemented in the design to detect such a failure in a subsecond?

A. fate sharing
B. OSPF LFA
C. flex links
D. STP
E. BFD

Answer: E

Explanation:

The correct answer is **E. BFD (Bidirectional Forwarding Detection)**. BFD is a lightweight, low-overhead protocol specifically designed for fast failure detection in network paths. Unlike routing protocols like OSPF, which rely on timers that can take several seconds to expire before a link failure is detected, BFD actively probes a path at very short intervals (milliseconds) to quickly identify issues. This rapid detection allows for significantly faster failover times, often within subsecond ranges, making it ideal for enhancing network resilience in scenarios requiring minimal interruption. Fate sharing (A) is a design principle, not a protocol for fast detection, while OSPF LFA (B) helps in creating alternate paths but doesn't speed up detection. Flex links (C) offer redundant paths on switches but are not a fast detection mechanism. STP (D), a Layer 2 protocol, is irrelevant to an OSPF routing protocol concern on a WAN. BFD operates independently of the routing protocol, thus enabling fast detection across various link technologies, and is the technique most suitable for the stated scenario.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/15-sy/iap-15-sy-book/ip-bfd.htmlhttps://www.juniper.net/documentation/us/en/software/junos/bfd/topics/topic-map/bfd-overview.html

Question: 64

Which three elements help network designers to construct secure systems that protect information and resources (such as devices, communication, and data) from unauthorized access, modification, inspection, or destruction?
(Choose three.)

A. scalability
B. availability
C. serviceability
D. integrity
E. confidentiality
F. reliability

Answer: BDE

Explanation:

Here's the justification for why availability, integrity, and confidentiality are crucial elements for building secure networks:

Availability (B): Availability ensures that authorized users can access information and resources when

needed. In a secure system, this means that security measures must not hinder legitimate access. Denial-of-service (DoS) attacks, for example, directly target availability by overwhelming systems and preventing access. A secure design must incorporate redundancy, failover mechanisms, and robust infrastructure to maintain availability despite such threats. The concept of high availability (HA) in cloud environments, achieved through features like load balancing and auto-scaling, reinforces its importance.https://cloud.google.com/architecture/reliability/availability

Integrity (D): Integrity focuses on ensuring that data and systems remain accurate, complete, and unaltered during transit and storage. Unauthorized modifications can have severe consequences, from data corruption to misrepresentation. Security measures that protect integrity include checksums, digital signatures, and access controls. A breach of integrity could allow an attacker to manipulate data for malicious purposes.

Technologies such as blockchain, which are widely used for maintaining data integrity, are becoming increasingly important.https://www.nist.gov/itl/applied-cybersecurity/nice/resources/cybersecurity-glossary#data-integrity

Confidentiality (E): Confidentiality is the principle of keeping sensitive information secret from unauthorized individuals. Access controls, encryption, and other privacy mechanisms are used to maintain confidentiality.

Data breaches, caused by inadequate confidentiality measures, can lead to reputational damage and legal issues. Protecting sensitive customer data, trade secrets, and personally identifiable information (PII) often entails specific cloud compliance needs. Strong encryption strategies are commonly used in the cloud to secure data both at rest and in transit.https://www.iso.org/files/live/sites/isoorg/files/store/en/PUB100079.pdfThese three elements form the CIA triad, a foundational concept in information security. Scalability, serviceability, and reliability are important qualities but are not the core elements used to construct secure systems. Instead they focus on a systems efficiency and are often used to achieve availability.

Question: 65

Which relationship between iBGP and the underlying physical topology is true?

A. iBGP full mesh requires an underlying fully meshed network topology.

B. iBGP full mesh requirement does not dictate any specific network topology.

C. iBGP does not work on a ring network topology even with an underlying IGP.

D. iBGP can work only on a ring network topology with a link-state protocol like OSPF or IS-IS.

Answer: B

Explanation:

Okay, here's a detailed justification for why option B is the correct answer, along with supporting information:

The question focuses on the relationship between iBGP (Internal Border Gateway Protocol) and the underlying physical network topology. iBGP's primary role is routing within a single autonomous system (AS). While iBGP does require a full mesh of peerings for proper route propagation (or route reflectors to avoid a full mesh), this requirement is logical and independent of the physical connections between the routers.

Option A is incorrect because iBGP's full mesh requirement refers to the logical peering relationships between routers, not the underlying physical infrastructure. You can have an iBGP full mesh even if routers are not physically connected to each other, so long as reachability exists. For example, using an IGP (like OSPF or EIGRP) provides the underlying reachability needed for BGP peers, which would exist in a typical enterprise network. Option C is incorrect because iBGP absolutely works over ring topologies, and this is a common scenario, especially when used with a properly configured IGP providing underlying IP connectivity between the iBGP speakers. Option D is wrong as iBGP's functioning is not restricted to a ring topology with a

link-state protocol and can work with distance-vector IGPs like RIP or EIGRP as long as basic IP routing is available to reach the BGP peers. The important point is that the underlying network must ensure IP reachability between the iBGP peers. The physical arrangement is irrelevant to iBGP as long as logical connectivity via IP is provided. The physical design can be anything: hub-and-spoke, partial mesh, full mesh, or even a ring topology.

Therefore, option B accurately states that the iBGP full mesh requirement is a logical, peering requirement, and not a physical, topology requirement. The underlying network design, as long as IP connectivity exists, is not dictated by iBGP. iBGP operates at layer 3 and requires only IP reachability provided by underlying protocols.

Authoritative Links for Further Research:

1. **Cisco BGP Configuration Guide:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/15-sy/irg-15-sy-book/irg-basic-bgp.html (Focus on understanding BGP peering)
2. **RFC 4271 - A Border Gateway Protocol 4 (BGP-4):**https://datatracker.ietf.org/doc/html/rfc4271 (The official standard for BGP, giving full technical detail)
3. **Juniper BGP Documentation:** https://www.juniper.net/documentation/us/en/software/junos/bgp/index.html (Provides vendor-specific insights and configuration guidance on BGP)

Question: 66

Which two statements describe the hierarchical LAN design model? (Choose two.)

A. It is a well-understood architecture that provides scalability.

B. It is the best design for modem data centers.

C. Changes, upgrades, and new services can be introduced in a controlled and staged manner.

D. It is the most optimal design but is highly complex.

E. It provides a simplified design.

Answer: AC

Explanation:

The hierarchical LAN design model, commonly structured into core, distribution, and access layers, offers significant advantages in network architecture. Statement A correctly identifies its scalability benefit; the modular structure allows for growth by adding new components at appropriate layers without requiring massive overhauls. This modularity is a key factor in managing the increasing demands of modern networks.

Statement C accurately points out that changes can be implemented in a controlled manner. This layer-by-layer approach allows for testing and upgrades to be performed within specific segments before broader deployment, minimizing disruption. This benefit stems from the isolation provided by the hierarchical design, ensuring localized impacts of alterations. Conversely, statement B is incorrect; while suitable for many networks, it's not the absolute "best" for all modern data centers, some may favor a flatter or spine-leaf architecture. Statement D is also flawed, because while effective, the hierarchical model is well-defined and not inherently "highly complex" compared to other network designs. Finally, statement E is incorrect as it implies oversimplification; while well-structured, a hierarchical design isn't necessarily simplified but rather it's well-organized. The layering enhances manageability but isn't meant to make the design "simpler" in a reductive sense. Therefore, options A and C accurately capture the advantages of the hierarchical LAN model regarding scalability and controlled implementation.

Further research can be found at:https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Oct2019/CVD-

Campus-Design-Guide-2019OCT.htmlhttps://www.networkcomputing.com/networking/hierarchical-network-design-what-it-and-why-you-should-care

Question: 67

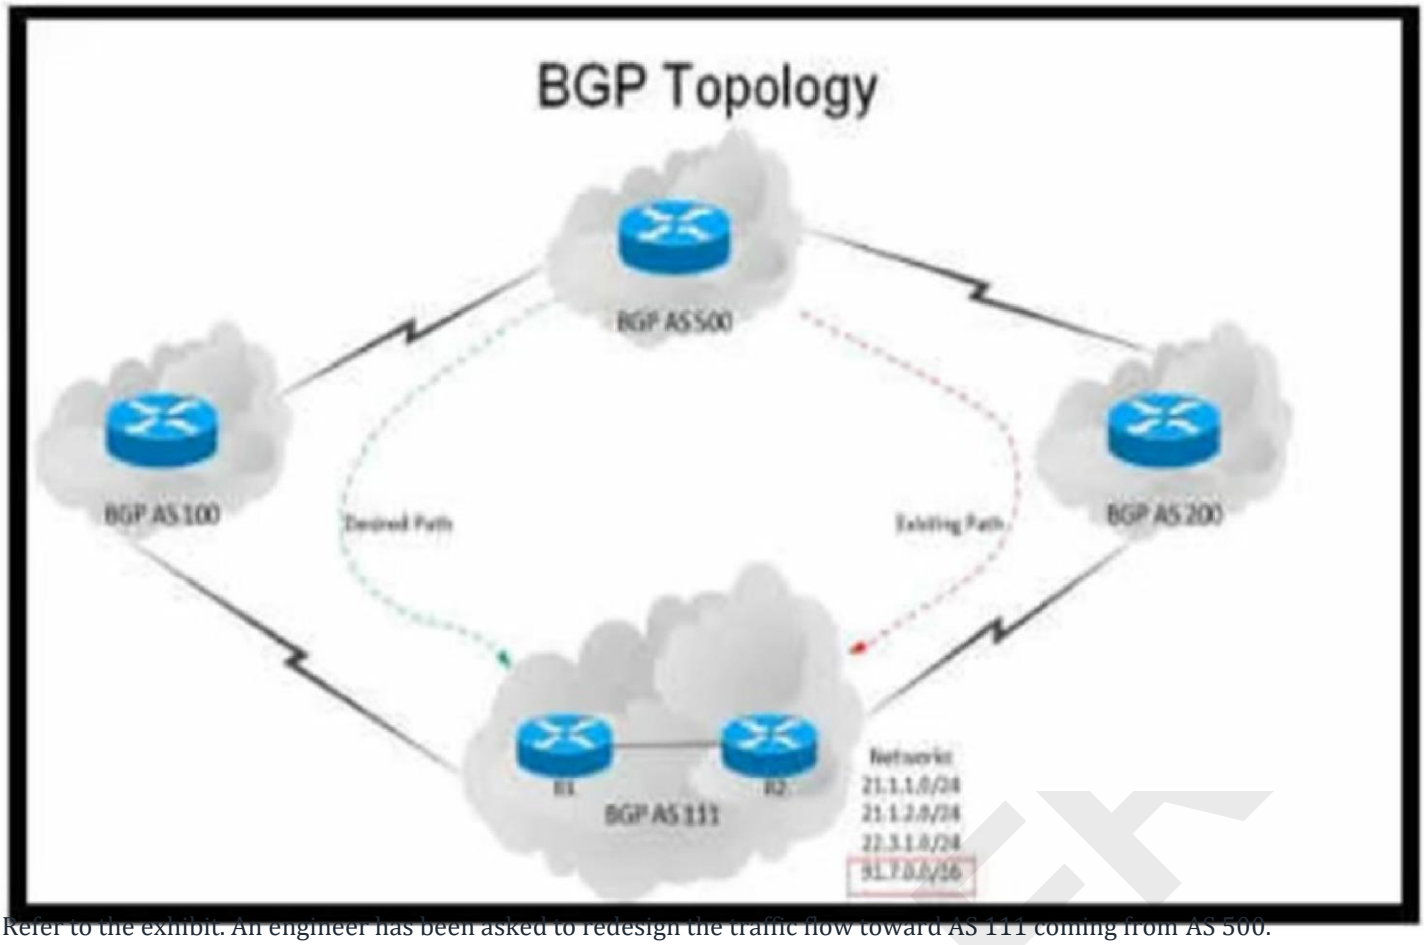

Refer to the exhibit. An engineer has been asked to redesign the traffic flow toward AS 111 coming from AS 500. Traffic destined to AS 111 network 91.7.0.0/16 should come in via AS 100, while traffic destined to all other networks in AS 111 should continue to use the existing path. Which BGP attributes are best suited to control this inbound traffic coming from BGP AS 500 into the 91.7.0.0/16 network?

A. Use local preference on R1 for the networks that AS 500 advertises to AS 111
B. Prepend AS path for the 91.7.0.0/16 network and set it for neighbor in AS 200
C. Use extended community for the 91.7.0.0/16 network, not advertising it to the bi-lateral peer
D. Set higher MED for neighbor in AS 100 to influence incoming traffic for the 91.7.0.0/16 network

Answer: B

Explanation:

B. Prepend AS path for the 91.7.0.0/16 network and set it for neighbor in AS 200

Question: 68

An enterprise that runs numerous proprietary applications has major issues with its on-premises server estate hardware, to the point where business-critical functions are compromised. The enterprise accelerates plans to migrate services to the cloud. Which cloud service should be used if the enterprise wants to avoid hardware issues yet have control of its applications and operating system?

A. SaaS
B. PaaS
C. IaaS
D. hybrid cloud

Answer: C

Explanation:

The correct answer is Infrastructure as a Service (IaaS). IaaS offers the foundational building blocks for cloud computing, providing virtualized computing resources like virtual machines, storage, and networks. This model allows the enterprise to avoid the physical hardware issues they are experiencing on-premises since the cloud provider manages the underlying infrastructure. Crucially, unlike other options, IaaS grants the enterprise full control over their operating systems and applications. Platform as a Service (PaaS) abstracts away the operating system level, making it unsuitable for a scenario where control over the OS is essential. Software as a Service (SaaS) delivers applications over the internet, limiting control over both the OS and the application's underlying structure. A hybrid cloud is a deployment model, not a service type, that integrates on-premises resources with cloud-based resources, so not appropriate in addressing the user's need to migrate away from physical hardware. The enterprise requires the flexibility to customize their virtual environment and choose the OS that fits the specific needs of their proprietary applications, a key
characteristic of IaaS. This provides the agility and flexibility without the burden of managing physical hardware.

For further research, please see:

Microsoft Azure Documentation on IaaS:https://azure.microsoft.com/en-us/resources/cloud-computing-dictionary/what-is-iaas/
AWS Documentation on IaaS:https://aws.amazon.com/what-is/iaas/
Google Cloud Platform Documentation on IaaS:https://cloud.google.com/learn/what-is-iaas

Question: 69

How must the queue sizes be designed to ensure that an application functions correctly?

A. The default queue sizes are good for any deployment as it compensates the serialization delay.
B. The queuing delay on every device in the chain must be exactly the same to the application required delay.
C. Each individual device queuing delay in the chain must be less than or equal to the application required delay.
D. The sum of the queuing delay of all devices plus serialization delay in the chain must be less than or equal to the application required delay.

Answer: D

Explanation:

The correct answer is **D** because it accurately reflects how queuing delays impact application performance in a network path. Let's break down why:

Network applications have specific latency requirements. When data packets traverse a network, they encounter queuing delays at each device (routers, switches, firewalls) where they're placed in queues awaiting processing. These delays accumulate along the path. Simultaneously, serialization delay, the time it takes to transmit a packet onto the physical medium, is also a factor, particularly at lower bandwidth links. Option A is incorrect, because default queue sizes are not universally applicable; they are often generic and may not be optimized for specific application needs. Option B is flawed because forcing identical queuing delay on all devices is impractical and unnecessary. What matters is the total delay, not the individual device

contribution. Option C is incorrect since it only focuses on the local device delay and not the overall path. The application's acceptable delay is a constraint on the end-to-end latency and not on individual delays. The total delay experienced is the cumulative effect of each device's queuing delay combined with the serialization delays. Thus, for an application to function correctly, the sum of all queuing delays along the entire path plus the serialization delays must be less than or equal to the application's acceptable delay limit. If the total delay exceeds the application's tolerance, it may experience performance issues such as high latency, dropped packets, or timeouts. For robust application design, it's crucial to design network architectures considering these cumulative delays and optimize queue sizes and bandwidth appropriately to maintain performance.For further understanding on queuing and network delays, consider these resources:

Cisco's QoS guide:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/configuration/15-mt/qos-15-mt-book.html
Packet Delay Variation (Jitter):https://www.cisco.com/c/en/us/support/docs/quality-of-service-qos/qos-packet-jitter/12402.html
Wikipedia on Queueing Theory:https://en.wikipedia.org/wiki/Queueing_theory

Question: 70

An enterprise requires MPLS connected branches to access cloud-based Microsoft 365 services over an SD-WAN solution. Internet access is available only at dual regional hub sites that are connected to the MPLS network. Which connectivity method provides an optimum access method to the cloud-based services if one ISP suffers loss or latency?

A. Cloud onRamp SWG
B. Cloud onRamp
C. Cloud onRamp gateway site
D. Cloud onRamp SaaS

Answer: D

Explanation:

The correct answer is **D. Cloud onRamp SaaS**. Here's why:

Cloud onRamp SaaS is designed to optimize access to Software-as-a-Service (SaaS) applications like Microsoft 365 over SD-WAN. It dynamically selects the best path for SaaS traffic based on factors like latency, packet loss, and jitter, ensuring consistent performance. In the scenario presented, where internet access is centralized at regional hub sites with dual ISPs, Cloud onRamp SaaS plays a crucial role.

Option A, Cloud onRamp SWG (Secure Web Gateway), focuses on security and filtering of web traffic, not optimal path selection for SaaS. Option B, Cloud onRamp, is a broader term and isn't specific enough to address the SaaS optimization need. Option C, Cloud onRamp gateway site, is often a component of cloud connectivity, but it does not handle the direct SaaS path optimization.

When one ISP suffers from loss or latency at the hub sites, Cloud onRamp SaaS automatically switches to the alternative hub or uses a different path to maintain connectivity and performance. This avoids the branch locations being affected by the ISP issues and ensures uninterrupted access to Microsoft 365. The dynamic path selection and QoS capabilities within Cloud onRamp SaaS make it ideal for this dual-ISP scenario, aligning with the requirements for resilient and optimized cloud-based service access through the SD-WAN solution. It is not solely about internet access, but about ensuring the end user experience for the specific SaaS application. It leverages the best available route, in this case, it's the alternative ISP.

Further research on Cloud onRamp SaaS and its application with Cisco SD-WAN can be found on the following resources:

Cisco SD-WAN Cloud OnRamp:https://www.cisco.com/c/en/us/solutions/enterprise-networks/sd-wan/cloud-onramp.html
Cisco SD-WAN Cloud OnRamp for SaaS:
https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/sdwan-xe-latest/sd-wan-xe-system-configuration/m-cloud-onramp-for-saas-app.html
Cisco SD-WAN Cloud OnRamp for IaaS and SaaS:
https://www.cisco.com/c/en/us/products/collateral/routers/sd-wan/whitepaper-c11-742676.html

Question: 71

As part of workspace digitization, a large enterprise has migrated all their users to Desktop as a Service (DaaS), by hosting the backend system in their on- premises data center. Some of the branches have started to experience disconnections to the DaaS at periodic intervals, however, local users in the data center and head office do not experience this behavior. Which technology can be used to mitigate this issue?

A. traffic policing
B. WRED
C. tail drop
D. traffic shaping

Answer: B

Explanation:

The correct answer is **B. WRED (Weighted Random Early Detection)**. Here's why:

The problem describes intermittent DaaS disconnections specifically at remote branches, while users closer to the data center experience no issues. This suggests network congestion at the WAN links connecting the branches to the data center is the culprit. When traffic bursts occur, especially with DaaS which can involve bursts of user activity and multimedia streams, congestion leads to packet drops. Tail drop (option C) indiscriminately drops all packets when the queue is full, exacerbating the issue by causing multiple simultaneous TCP retransmissions, which further congest the link and potentially leading to more disconnections. Traffic policing (option A) simply enforces a bandwidth limit and drops or marks packets exceeding this limit, not ideal for mitigating congestion in a shared environment. Traffic shaping (option D) is typically used to limit and smooth traffic over time, which might also address the issue, but is less suited to dealing with bursty traffic, and is also a more complex configuration.

WRED is designed to address this specific scenario. It proactively manages queue lengths by selectively dropping packets based on priority, before the queue is full and the device switches to tail drop. It's most effective against TCP flows because it causes them to slow down before the queue is full, preventing more drops from happening. By dropping low priority packets before the queue is full, WRED avoids massive TCP retransmissions. Thus WRED helps reduce congestion, improve network stability, and minimize packet loss, directly addressing the branch DaaS disconnections. This makes it a more refined solution for managing congestion and optimizing performance in congested WAN links compared to traffic policing or tail drop.

Authoritative Links:

Cisco on WRED:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_conmgt/configuration/15-sy/qos-conmgt-15-sy-book/qos-wred.html
Wikipedia on WRED:https://en.wikipedia.org/wiki/Weighted_random_early_detection

Question: 72

A European government passport agency considers upgrading its IT systems to increase performance and workload flexibility in response to constantly changing requirements. The budget manager wants to reduce capital expenses and IT staff and must adopt the lowest-cost technology. Which technology choice is suitable?

A. public cloud
B. hybrid cloud
C. on premises
D. private cloud

Answer: B

Explanation:

The correct answer is **B. hybrid cloud**. Here's why:

A hybrid cloud model best addresses the passport agency's needs by blending the benefits of both public and private cloud solutions. Public cloud (A) provides cost-effectiveness through shared infrastructure and reduced upfront capital expenditure, aligning with the budget manager's goals. It also offers scalability and flexibility for fluctuating workloads. However, public cloud alone might not meet all security or compliance requirements concerning sensitive passport data. On-premises (C) is unsuitable because it involves high capital expenditure for infrastructure, contradicts the cost reduction goal, and requires ongoing IT staff management. While private cloud (D) allows for control and security, it would still incur significant capital and operational costs for infrastructure and staff. Hybrid cloud enables the passport agency to host less sensitive applications and data on the public cloud for cost-efficiency and scalability, while retaining control over sensitive data within a private cloud or on-premises infrastructure, ensuring data security and compliance. This combination allows for optimized resource utilization, cost control, flexibility, and security, which is ideal for the agency's specific needs. The agency can leverage public cloud elasticity to handle fluctuating demand while maintaining control over sensitive information.

Authoritative Links:

NIST Definition of Cloud Computing:https://csrc.nist.gov/publications/detail/sp/800-145/final - Provides a foundational understanding of cloud computing models, including hybrid cloud.
Microsoft Azure Hybrid Cloud:https://azure.microsoft.com/en-us/solutions/hybrid-cloud/ - Offers practical examples and benefits of hybrid cloud implementation.
AWS Hybrid Cloud:https://aws.amazon.com/hybrid/ - Provides information about AWS' approach to hybrid cloud.

Question: 73

Which technology is an open-source infrastructure automation tool that automates repetitive tasks for users who work in networks such as cloud provisioning and intraservice orchestration?

A. Java
B. Ansible
C. Contrail
D. Jinja2

Answer: B

Explanation:

The correct answer is **B. Ansible**. Ansible is an open-source, agentless automation engine that simplifies IT tasks such as configuration management, application deployment, and cloud provisioning. It employs a

declarative language using YAML, enabling users to define desired states for systems, rather than specifying each step to achieve that state. This makes Ansible ideal for automating repetitive network tasks, including cloud resource creation and orchestration of services within a network. Ansible uses SSH or WinRM to connect and execute tasks on target machines. Its simplicity, ease of learning, and idempotent nature make it widely adopted for network automation. It doesn't require agents on managed nodes, reducing overhead.

Java (A) is a programming language and runtime environment, not an infrastructure automation tool. Contrail (C) is a network virtualization and cloud networking platform and provides more of a platform for network services, not automation. Jinja2 (D) is a template engine often used in conjunction with automation tools, but is not an automation engine itself. Therefore, among the given options, Ansible directly fits the description of an open-source automation tool for tasks like cloud provisioning and intraservice orchestration.

Authoritative Links:

Ansible Documentation:https://docs.ansible.com/
Red Hat Ansible Website:https://www.redhat.com/en/technologies/management/ansible
Introduction to Ansible:https://www.networkworld.com/article/3279748/what-is-ansible-the-basics-of-configuration-management.html

Question: 74

A European national bank considers migrating its on-premises systems to a private cloud offering in a non-European location to significantly reduce IT costs. What is a primary factor prior to migration?

A. security
B. cloud connectivity
C. additional latency
D. data governance

Answer: A

Explanation:

The correct answer is **A. security**. Migrating a European national bank's systems, especially to a non-European location, elevates security to a paramount concern. Banks are highly regulated and handle sensitive financial data, making them attractive targets for cyberattacks. A migration must prioritize ensuring that the private cloud infrastructure, regardless of its location, adheres to stringent security standards, including encryption, access control, intrusion detection, and compliance with relevant European regulations like GDPR (General Data Protection Regulation). Data location and residency are key legal aspects that must be addressed; failure to do so could result in hefty fines. Cloud connectivity, latency, and data governance are important, but they are secondary to the core security requirements in this highly regulated industry. Neglecting security can lead to data breaches, financial losses, and severe reputational damage. Therefore, robust security measures must be the foundational pillar of any such migration strategy. The bank must ensure that data is protected throughout the entire lifecycle, including in transit, at rest, and in use.

Supporting Concepts:

Data Security in the Cloud: Securing sensitive data in the cloud requires a multi-layered approach encompassing various aspects such as encryption, access control, and network security.

Link:NIST Special Publication 800-146, Cloud Computing Synopsis
GDPR and Data Residency: The General Data Protection Regulation (GDPR) has strict rules regarding the transfer of personal data outside the European Economic Area (EEA). Any bank processing personal data of European citizens must ensure compliance with these rules.

Link: Official GDPR Website
Financial Industry Regulations: Financial institutions are typically subject to rigorous regulations and compliance requirements aimed at protecting sensitive customer data.

Link: Varies significantly based on national and regional authorities. Research applicable financial regulatory agencies (e.g., European Banking Authority) is advised.

Question: 75

Which two actions ensure voice quality in a branch location with a low-speed, high-latency WAN connection? (Choose two.)

A. Prioritize voice packets.
B. Replace any electrical links with optical links.
C. Increase memory on the branch switch.
D. Fragment data packets.
E. Increase WAN bandwidth.

Answer: AE

Explanation:

The correct answer is A and E because these actions directly address the common challenges associated with low-speed, high-latency WAN connections for voice traffic.

A. Prioritize voice packets (Quality of Service - QoS): Prioritizing voice packets ensures they receive preferential treatment over other types of traffic on the network. This is achieved through QoS mechanisms like queuing, shaping, and marking. By giving voice packets higher priority, we minimize latency and jitter (variation in delay), which are detrimental to voice quality. This mechanism ensures real-time sensitive voice data is delivered promptly, reducing delays that can lead to choppy or distorted audio.

E. Increase WAN bandwidth: Low bandwidth is a major contributor to congestion, which can lead to packet loss and increased latency, both of which badly affect voice quality. Increasing the WAN bandwidth provides more capacity for all traffic, including voice. This reduces contention and allows voice packets to traverse the network without encountering excessive queuing delays, improving call clarity. When a link is saturated, delays become severe, but with more bandwidth the effects of congestion are mitigated and more room exists for real-time traffic to pass without contention.

Why the other options are not the best fit:

B. Replace any electrical links with optical links: While optical links often offer higher bandwidth and lower latency than electrical links, they are not always feasible or cost-effective for addressing branch location issues. The core issue is low speed and not always just the link type. A direct bandwidth upgrade is a more effective and targeted approach to the issue of low speeds.

C. Increase memory on the branch switch: Increasing switch memory might help in handling a higher volume of data, but it does not address the underlying issue of limited WAN bandwidth or priority for voice packets. It won't directly improve voice quality on a choked WAN link.

D. Fragment data packets: Fragmentation might help reduce transmission time on lower MTU links, but can be counterproductive due to the overhead of reassembly, and it doesn't address the core challenges of bandwidth limitation and lack of voice priority that cause quality issues. This is usually not a first response to improve quality.

Authoritative Links for further research:

Cisco QoS:https://www.cisco.com/c/en/us/solutions/enterprise/design-zone/index.html#~quality
Voice over IP (VoIP) Performance:
https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Voice/perf_guide/perfguide.html
Understanding QoS Concepts:
https://www.juniper.net/documentation/us/en/software/junos/qos/topics/topic-map/understanding-qos.html

Question: 76

Which three tools are used for ongoing monitoring and maintenance of a voice and video environment? (Choose three.)

A. active monitoring via synthetic probes to measure loss, latency, and jitter
B. call management analysis to identify network convergence-related failures
C. passive monitoring via synthetic probes to measure loss, latency, and jitter
D. call management analysis to identify CAC failures and call quality issues
E. flow-based analysis to measure bandwidth mix of applications and their flows
F. flow-based analysis with PTP time-stamping to measure loss, latency, and jitter

Answer: ADE

Explanation:

Here's a detailed justification for why options A, D, and E are the correct tools for ongoing monitoring and maintenance of a voice and video environment, and why the others are not:

Justification for A (active monitoring): Active monitoring via synthetic probes is crucial for simulating voice and video traffic. These probes generate test packets, allowing administrators to directly measure network performance metrics like loss, latency, and jitter. These metrics are fundamental indicators of voice and video quality, and real-time visibility is essential for proactive troubleshooting and identifying potential issues before users experience service degradation. This proactive approach aligns with the need for high availability in cloud-based communication systems.

Justification for D (call management analysis for CAC and call quality): Call management analysis is vital for identifying issues related to Call Admission Control (CAC) and overall call quality. CAC ensures the network has sufficient resources to handle new calls, preventing overload and maintaining quality. By examining call data, administrators can identify instances where calls are being rejected due to insufficient bandwidth or resource constraints, which is a crucial indicator for capacity planning. Analysis can also pinpoint poor call quality by looking at metrics such as MOS scores and packet loss reported by endpoints which can be used to pinpoint issues. This provides essential insights into user experience.

Justification for E (flow-based analysis): Flow-based analysis provides visibility into the bandwidth usage of different applications and their flows across the network. Understanding how voice and video applications are consuming bandwidth, alongside other applications, is important for resource allocation and prioritization. This analysis can reveal if voice and video traffic is being bottlenecked or if other applications are interfering with its performance. This awareness allows for proactive adjustments to QoS policies to ensure optimal voice and video performance.

Why B, C, and F are incorrect:

B (Call management analysis for network convergence): While network convergence is important during network design and initial setup, it's not a primary focus for ongoing monitoring and maintenance. Convergence issues are typically addressed during troubleshooting.

C (Passive monitoring via synthetic probes): Passive monitoring using probes would simply listen to network traffic and wouldn't be able to proactively generate test packets to assess specific paths. Synthetic probes

for active testing must be active to test parameters like loss, latency and jitter.

F (Flow-based analysis with PTP): While PTP (Precision Time Protocol) is important for timing synchronization in certain network environments, it is not directly used to measure loss, latency, and jitter within a typical voice and video monitoring framework. PTP primarily aids in time synchronization and its inclusion in this flow analysis method doesn't make it any more appropriate for the requested context. PTP also usually deals with hardware appliances which is not the context.

Authoritative Links for Further Research:

Cisco Collaboration Monitoring:
https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/service/12_5_1/monitor/cucm_b_monitoring-guide-1251/cucm_b_monitoring-guide-1251_chapter_00.html
RFC 6791 - Flow-Based Measurement:https://datatracker.ietf.org/doc/html/rfc6791
VoIP Monitoring Best Practices:https://www.exoprise.com/blog/best-practices-for-voip-monitoring/

In summary, ongoing monitoring of voice and video requires active testing, call management analysis, and flow-based analysis which gives the capability to proactively and reactively maintain the environment and keep the end user experience positive.

Question: 77

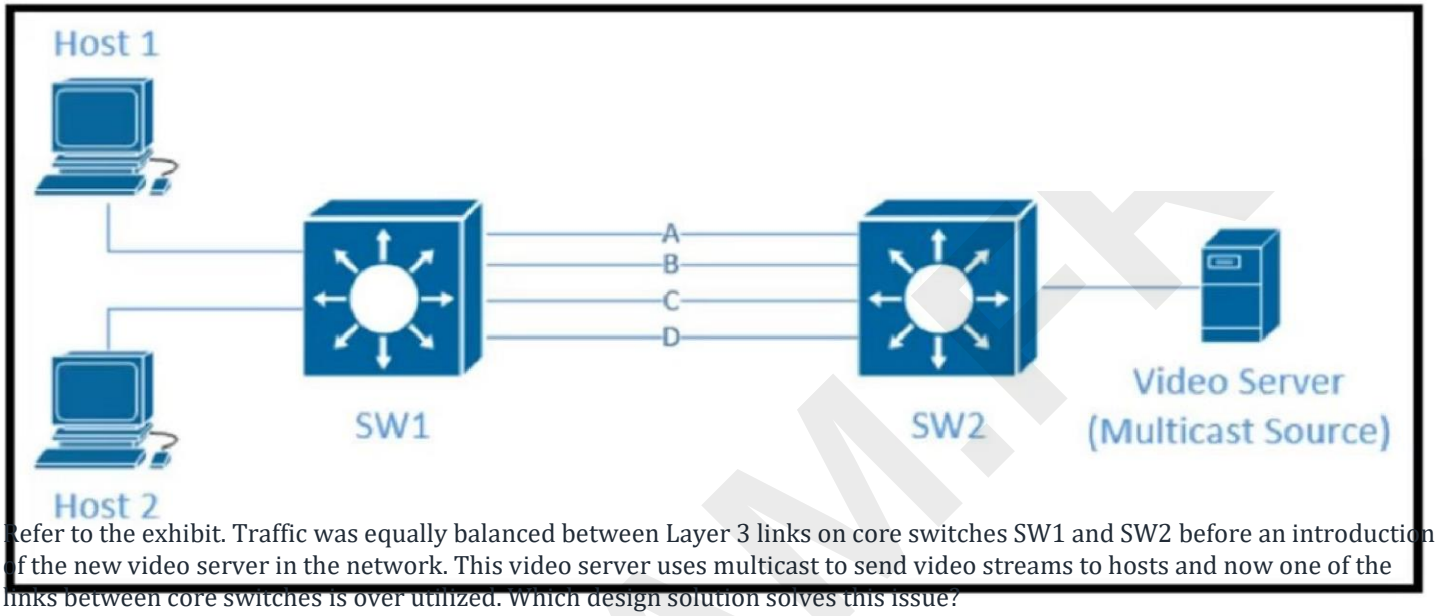

Refer to the exhibit. Traffic was equally balanced between Layer 3 links on core switches SW1 and SW2 before an introduction of the new video server in the network. This video server uses multicast to send video streams to hosts and now one of the links between core switches is over utilized. Which design solution solves this issue?

A. Aggregate links using Layer 2 link aggregation.
B. Add more links between core switches.
C. Apply a more granular load-balancing method on SW2.
D. Filter IGMP joins on an over-utilized link.
E. Apply a more granular load-balancing method on SW1.

Answer: A

Explanation:

With no L2 aggregation the only option is ECMP load splitting, and it's not load balancing.

Question: 78

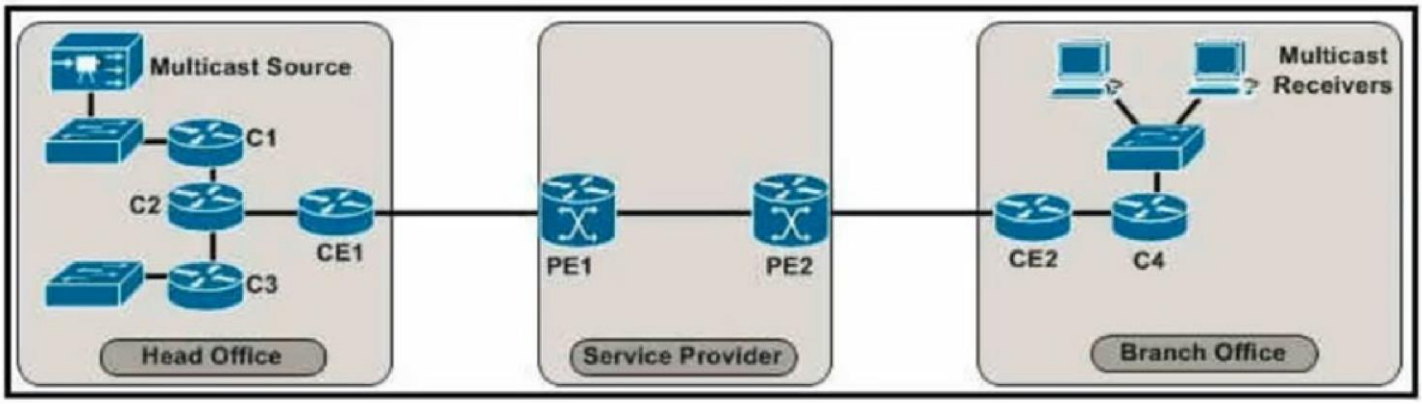

Refer to the exhibit. This enterprise customer wants to stream one-way video from their head office to eight branch offices using multicast. Their current service provider provides a Layer 3 VPN solution and manages the CE routers, but they do not currently support multicast. Which solution quickly allows this multicast traffic to go through while allowing for future scalability?

A. Enable a GRE tunnel between nodes C1 and C4.

B. Enable a GRE tunnel between nodes C2 and C4.

C. Enable a GRE tunnel between nodes CE1 and CE2.

D. Implement hub and spoke MPLS VPN over DMVPN (also known as 2547oDMVPN) between CE1 and CE2. E. The service provider must provide a Draft Rosen solution to enable a GRE tunnel between nodes PE1 and PE2.

Answer: B

Explanation:

B

https://www.cisco.com/c/en/us/support/docs/ip/ip-multicast/43584-mcast-over-gre.html

Question: 79

Which architecture does not require an explicit multicast signaling protocol, such as PIM or P2MP, to signal the multicast state hop-by-hop, but instead uses a link state protocol to advertise the multicast forwarding state?

A. Bi-Directional Implicit Replication
B. Bit Indexed Explicit Replication
C. Binary Intermediate Enhanced Routing
D. Binary Indexed Explicit Routing

Answer: B

Explanation:

The correct answer is **B. Bit Indexed Explicit Replication (BIER)**. BIER is a multicast forwarding architecture that deviates from traditional multicast methods, like PIM, by not relying on explicit signaling protocols to establish multicast state hop-by-hop. Instead, BIER utilizes a bit string within the packet header to represent the set of destination receivers. Routers along the path, referred to as Bit Forwarding Routers (BFRs), inspect this bit string and forward the packet based on their pre-configured bitmask. This pre-configuration means that multicast forwarding state is implicitly derived through link-state protocol advertisements, which distribute information about the BFRs and their associated bit positions. Consequently, BIER eliminates the need for explicit signaling protocols like PIM to build forwarding trees, simplifying multicast network design

and operations. Options A, C, and D do not represent valid or relevant multicast architectures in this context. Implicit replication, while a concept in multicast, is not a specific architecture like BIER. Furthermore, Binary Intermediate Enhanced Routing and Binary Indexed Explicit Routing are not established or recognized multicast terminology.

Authoritative Links:

RFC 8279 - Bit Index Explicit Replication (BIER):https://datatracker.ietf.org/doc/html/rfc8279 - Provides the official specification for BIER.

Cisco White Paper on BIER:https://www.cisco.com/c/en/us/solutions/collateral/service-provider/ios-xr-software/white-paper-c11-742588.html - Offers a practical overview of BIER from a Cisco perspective. **Wikipedia article on BIER:**https://en.wikipedia.org/wiki/Bit_Indexed_Explicit_Replication - Provides a general overview of BIER concepts.

Question: 80

Which two advantages of using DWDM over traditional optical networks are true? (Choose two.)

A. inherent topology flexibility with a service protection provided through a direct integration with an upper layer protocol

B. inherent topology flexibility with built-in service protection

C. ability to expand bandwidth over existing optical infrastructure

D. inherent topology flexibility with intelligent chromatic dispersion

E. inherent topology flexibility and service protection provided without penalty through intelligent oversubscription of bandwidth reservation

Answer: BC

Explanation:

Here's the justification for why options B and C are the correct advantages of using Dense Wavelength Division Multiplexing (DWDM) over traditional optical networks:

B. inherent topology flexibility with built-in service protection: DWDM systems allow for flexible network topologies because each wavelength acts as a separate virtual channel, enabling multiple point-to-point, ring, and mesh configurations on the same fiber. Furthermore, DWDM systems can incorporate built-in protection mechanisms, such as optical switch protection, to automatically reroute traffic in case of a fiber cut or node failure, increasing the overall resilience and availability of the network. This protection isn't tied to upper-layer protocols.

C. ability to expand bandwidth over existing optical infrastructure: DWDM's primary advantage is its ability to vastly increase the capacity of an existing optical fiber. Instead of requiring new fiber installation, DWDM utilizes multiple wavelengths (colors of light) to transmit data simultaneously over the same fiber. This dramatically increases the bandwidth by effectively creating multiple virtual fibers within one physical fiber, making it a cost-effective way to meet increasing bandwidth demands.

Option A is incorrect because DWDM's protection mechanisms are built-in at the optical layer, not integrated directly with upper-layer protocols. Option D is incorrect because while DWDM can mitigate chromatic dispersion, it doesn't inherently introduce intelligent dispersion capabilities through topology flexibility. Option E is incorrect because DWDM doesn't provide service protection through intelligent oversubscription of bandwidth reservations; rather, its protection is built-in using redundancy and failover mechanisms. DWDM's core advantage is to boost bandwidth and enable diverse topologies.

Authoritative Links for further research:

Cisco DWDM Technology Overview:https://www.cisco.com/c/en/us/solutions/collateral/service-provider/optical-networking/whitepaper_c11-445298.html
Juniper Networks Optical Networking:https://www.juniper.net/us/en/products/optical-networking/
IEEE Xplore: DWDM Research Papers:https://ieeexplore.ieee.org/search/searchresult.jsp?newsearch=true&queryText=DWDM