complete your programming course about resources, doubts and more!

# Cisco

(350-401)

Implementing Cisco Enterprise Network Core Technologies
(ENCOR)

Total: **1067 Questions**
Link:

### Question: 1

What is the difference between a RIB and a FIB?

A. The FIB is populated based on RIB content.
B. The RIB maintains a mirror image of the FIB.
C. The RIB is used to make IP source prefix-based switching decisions.
D. The FIB is where all IP routing information is stored.

**Answer: A**

**Explanation:**

The correct answer is A: "The FIB is populated based on RIB content." This accurately reflects the relationship between the Routing Information Base (RIB) and the Forwarding Information Base (FIB) in network devices, such as Cisco routers. The RIB is a database that holds all the routing information learned from various routing protocols (like OSPF, EIGRP, BGP), static routes, and connected interfaces. It's like the "master table" containing all possible routes. The FIB, on the other hand, is a streamlined, optimized version of the RIB, designed for fast packet forwarding. It only contains the best paths to destinations and is used directly by the forwarding engine.

Option B is incorrect because the RIB is the master record, and the FIB is a derivative. The FIB does not mirror the RIB; it only stores the actively used paths. Option C is incorrect because switching decisions are based on destination IP addresses, not source IP prefixes, and this process is done via the FIB, not the RIB. Option D is incorrect because the RIB is where all routing information is stored; the FIB only stores information used for active forwarding. The FIB is designed to be readily accessible by the hardware to ensure speed and is built based on RIB selections. Think of the RIB as a comprehensive map book and the FIB as the single page that shows the route you are taking on a particular journey.

In essence, routing protocols populate the RIB with routing information. The router then runs route selection algorithms to determine the best paths. These best paths are then transferred to the FIB, which the router's forwarding engine uses to make packet-forwarding decisions quickly. This process optimizes for performance; otherwise, the router would have to make many complex routing decisions for each packet if using the entire RIB. This separation enables efficient forwarding of packets and better utilizes router resources.

Authoritative links for further reading:

**Cisco Documentation on RIB and FIB:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_fib/configuration/15-mt/irf-15-mt-book/irf-fib.html
**Explanation on the relationship between RIB and FIB**:
https://networkengineering.stackexchange.com/questions/4205/what-is-the-difference-between-a-rib-and-a-fib

### Question: 2

Which QoS component alters a packet to change the way that traffic is treated in the network?

A. policing
B. classification
C. marking
D. shaping

**Answer: C**

**Explanation:**

The correct answer is **C. marking.** Marking, within the context of Quality of Service (QoS), is the process of altering a packet's header to indicate its relative importance or priority. This alteration allows network devices to identify and apply specific treatments to different traffic flows. Marking involves setting specific bits within Layer 2 (e.g., CoS) or Layer 3 (e.g., DSCP) headers, effectively tagging the packets. These tags act as instructions for intermediate devices. This change in the packet header is the core of differentiating services.

Policing (A) and shaping (D) are traffic management mechanisms used to enforce rate limits and smooth out traffic bursts. They act upon already marked packets, controlling the traffic flow based on its classification. Classification (B) is the process of identifying traffic based on various parameters but does not directly alter the packet header itself. Marking is the critical component that precedes these actions, allowing routers and switches to recognize and prioritize packets. Therefore, while policing, shaping, and classification are vital QoS elements, they rely on marking to achieve differential treatment of packets. Marking is the foundational step in enabling QoS by encoding traffic characteristics directly into the packet header.

For additional research, refer to the following authoritative links:

1. **Cisco QoS - Marking**: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_plc/configuration/15-mt/qos-plc-15-mt-book/qos-mark.html
2. **RFC 2474 - Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers**: https://www.rfc-editor.org/rfc/rfc2474
3. **Wikipedia - Quality of service**: https://en.wikipedia.org/wiki/Quality_of_service

## Question: 3

DRAG DROP -

Drag and drop the descriptions from the left onto the correct QoS components on the right.
Select and Place:

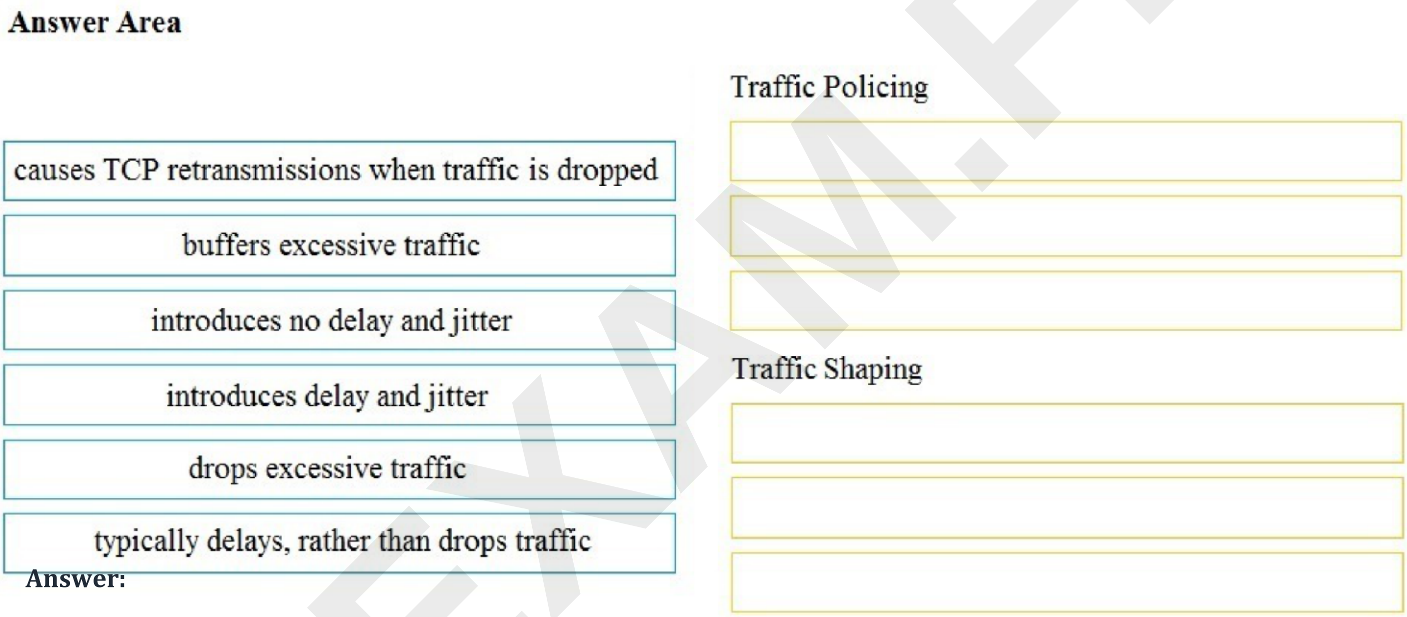

**Answer Area**

causes TCP retransmissions when traffic is dropped

buffers excessive traffic

introduces no delay and jitter

introduces delay and jitter

drops excessive traffic

typically delays, rather than drops traffic

Traffic Policing

Traffic Shaping

**Answer:**

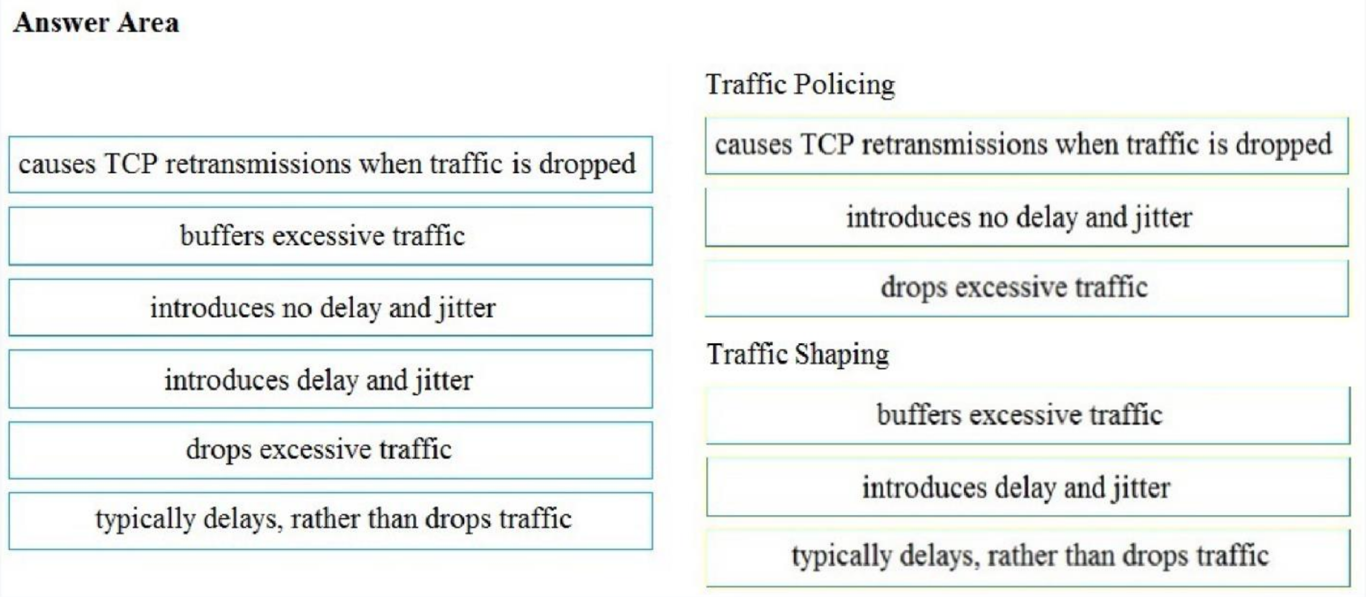

**Answer Area**

causes TCP retransmissions when traffic is dropped

buffers excessive traffic

introduces no delay and jitter

introduces delay and jitter

drops excessive traffic

typically delays, rather than drops traffic

Traffic Policing

causes TCP retransmissions when traffic is dropped

introduces no delay and jitter

drops excessive traffic

Traffic Shaping

buffers excessive traffic

introduces delay and jitter

typically delays, rather than drops traffic

**Explanation:**

Traffic Policing

causes tcp retransmitions

Introduces no delay and jitter

drops excessive traffic

Traffic Shapping

buffer excessive traffic

Introduces delay and jitter

typically delays , rather than drops traffic

**Question: 4** Which statement about Cisco Express Forwarding is true?

A. The CPU of a router becomes directly involved with packet-switching decisions. B. It uses a fast cache that is maintained in a router data plane.

C. It maintains two tables in the data plane: the FIB and adjacency table.

D. It makes forwarding decisions by a process that is scheduled through the IOS scheduler.

**Answer: C**

**Explanation:**

The correct answer is **C. It maintains two tables in the data plane: the FIB and adjacency table.**

Cisco Express Forwarding (CEF) is a layer 3 switching mechanism designed for high-performance packet forwarding in Cisco routers. Unlike process switching where the CPU is directly involved in each packet's forwarding decision, CEF operates in the data plane, also known as the hardware forwarding plane. This allows for faster switching. Crucially, CEF relies on two primary data structures: the Forwarding Information Base (FIB) and the adjacency table. The FIB contains pre-computed routing information, including next-hop IP addresses and outgoing interfaces, derived from the routing table. The adjacency table provides Layer 2 information necessary to encapsulate and send the packet. When a packet arrives, the router consults the FIB to identify the output interface and the corresponding adjacency table entry to determine the specific MAC address needed for delivery. By using these pre-built tables, CEF significantly reduces CPU utilization for packet forwarding. Therefore, the packet switching happens within the data plane using these tables.

Option A is incorrect because CEF bypasses the CPU for the actual forwarding process. Option B is incorrect because the FIB and adjacency tables are not merely caches, they are fundamental structures for CEF. Option D is incorrect because CEF doesn't rely on the IOS scheduler for forwarding decisions; it uses the FIB and adjacency table directly in the data plane.

Here are some authoritative links for further research:

**Cisco Documentation on CEF:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipswitch/configuration/15-mt/ips-15-mt-book/ips-cef.html
**Understanding Cisco Express Forwarding:**https://www.networkworld.com/article/2285894/understanding-cisco-express-forwarding.html

**Question: 5**

What is a benefit of deploying an on-premises infrastructure versus a cloud infrastructure deployment?

A. ability to quickly increase compute power without the need to install additional hardware
B. less power and cooling resources needed to run infrastructure on-premises
C. faster deployment times because additional infrastructure does not need to be purchased
D. lower latency between systems that are physically located near each other

**Answer: D**

**Explanation:**

The correct answer, D, highlights a key advantage of on-premises infrastructure: lower latency for systems located physically close to each other. On-premises deployments, where hardware resides within an organization's physical location, minimize the distance data travels. Shorter distances typically translate to faster communication and reduced latency. This is particularly beneficial for applications requiring real-time responses or high-bandwidth interactions, such as financial trading systems or complex simulations. Cloud infrastructure, conversely, often involves data traveling across a public network to reach remote servers, which can introduce unpredictable latency. While cloud providers strive to minimize latency, they cannot always guarantee the consistent low latencies achievable with an on-premises setup. The cloud's inherent reliance on internet connections and data travel through multiple network hops means there are more points where delay can occur. Options A, B, and C are incorrect. Cloud environments offer easier scalability (A), typically require less internal power and cooling (B) for the organization, and can have faster deployment times in some cases (C) compared to setting up a complete on-premises infrastructure, especially if hardware is required.

For further research:

1. **Latency in Cloud Computing:**https://www.bmc.com/blogs/cloud-latency/ This article explains cloud latency factors and how it differs from on-premises.
2. **On-Premises vs. Cloud:**https://www.spiceworks.com/tech/cloud/articles/on-premise-vs-cloud/ This resource provides a comparison of the pros and cons of both deployment models.
3. **On-Premises Infrastructure:**https://www.ibm.com/topics/on-premises This IBM article offers a definition and benefits of on-premises infrastructure.

**Question: 6**

DRAG DROP -
Drag and drop the characteristics from the left onto the appropriate infrastructure deployment types on the right.
Select and Place:

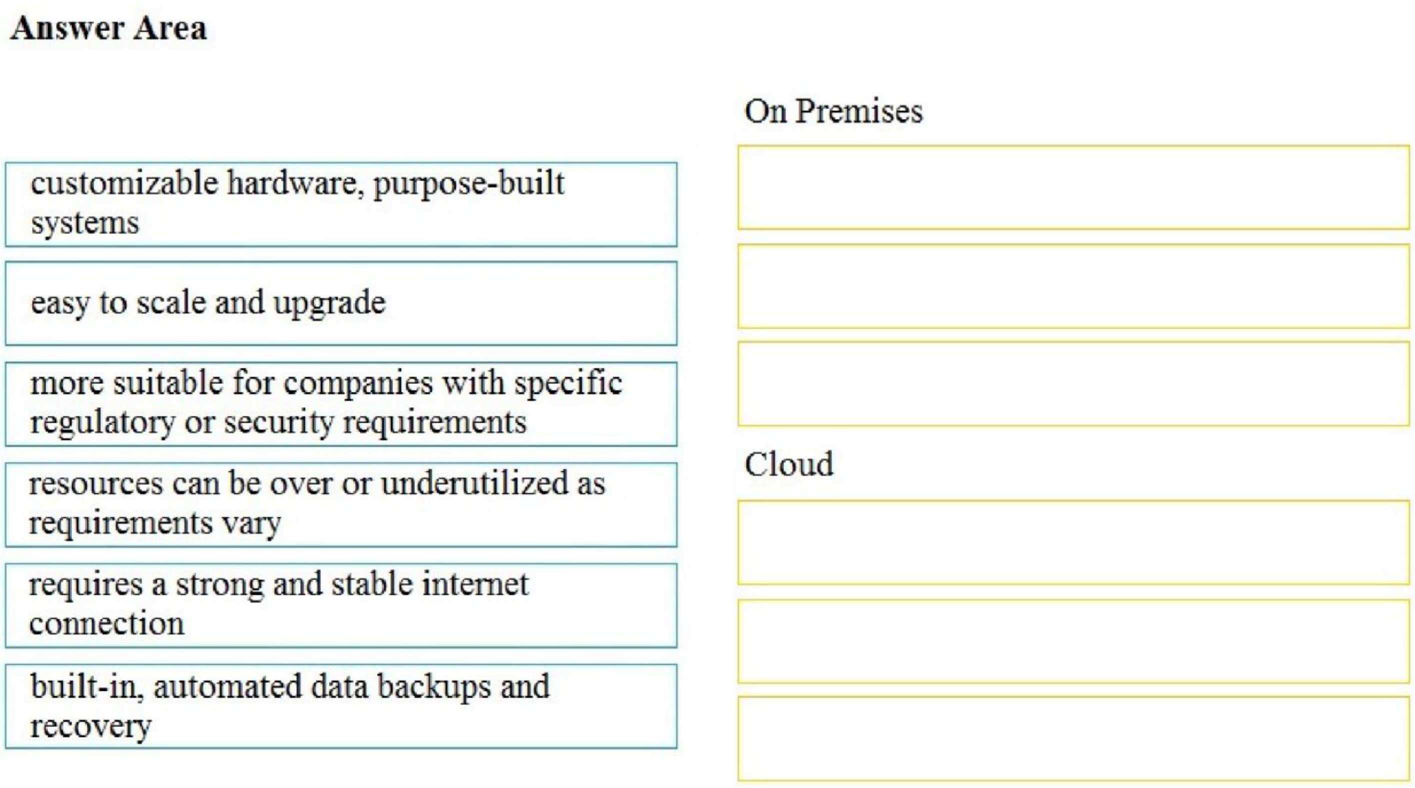

Answer Area

customizable hardware, purpose-built systems

easy to scale and upgrade

more suitable for companies with specific regulatory or security requirements

resources can be over or underutilized as requirements vary

requires a strong and stable internet connection

built-in, automated data backups and recovery

On Premises

Cloud

**Answer:**

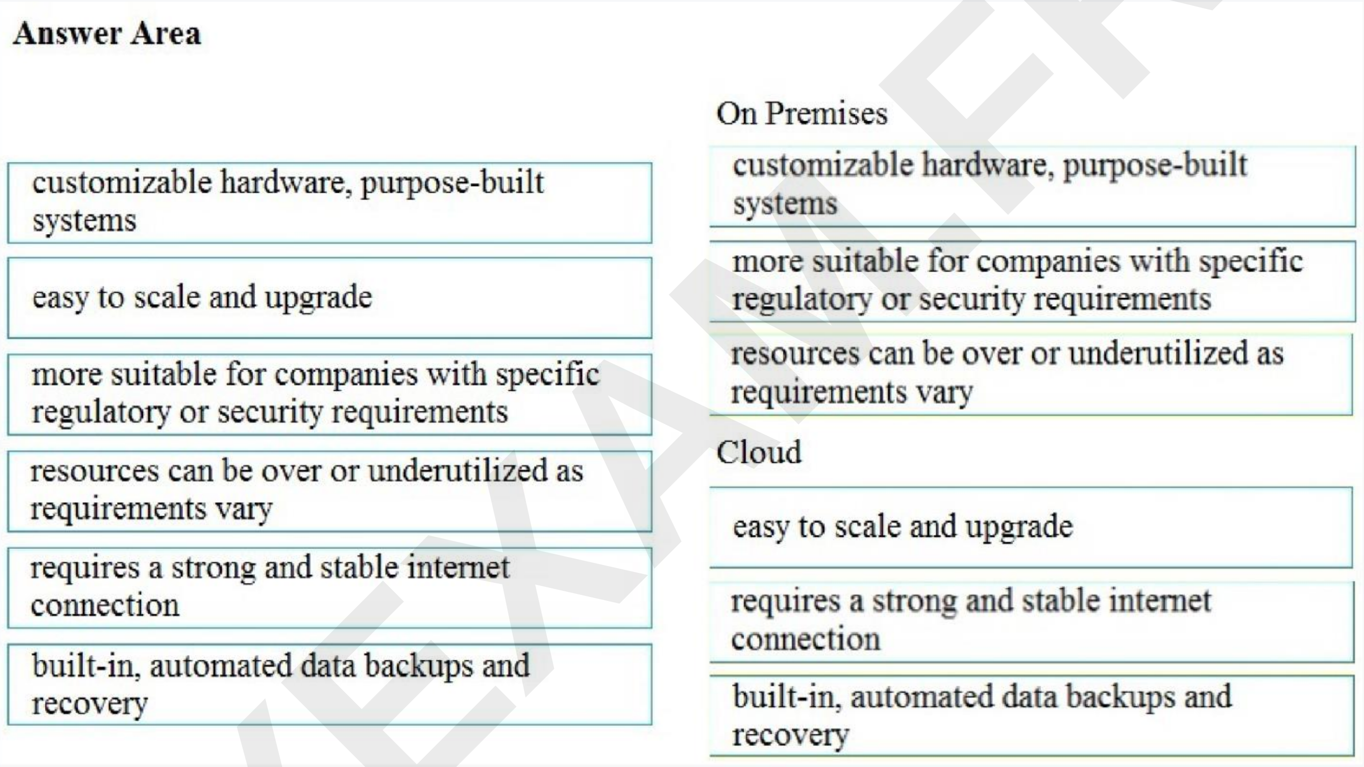

Answer Area

customizable hardware, purpose-built systems

easy to scale and upgrade

more suitable for companies with specific regulatory or security requirements

resources can be over or underutilized as requirements vary

requires a strong and stable internet connection

built-in, automated data backups and recovery

On Premises

customizable hardware, purpose-built systems

more suitable for companies with specific regulatory or security requirements

resources can be over or underutilized as requirements vary

Cloud

easy to scale and upgrade

requires a strong and stable internet connection

built-in, automated data backups and recovery

**Question: 7**

How does QoS traffic shaping alleviate network congestion?

A. It drops packets when traffic exceeds a certain bitrate.

B. It buffers and queues packets above the committed rate.

C. It fragments large packets and queues them for delivery.

D. It drops packets randomly from lower priority queues.

**Answer: B**

**Explanation:**

The correct answer is B. Traffic shaping, a Quality of Service (QoS) mechanism, alleviates network congestion by controlling the rate at which traffic is sent. Unlike traffic policing which discards excess traffic exceeding a defined rate, shaping buffers and queues packets that surpass the configured bandwidth limit. This buffering process smoothes out traffic bursts, preventing sudden spikes from overwhelming network resources. These queued packets are then transmitted at a controlled rate, ensuring a steady stream of data flow and preventing network congestion. This regulated flow prevents packet loss and delays that may arise during periods of high traffic. Essentially, shaping delays excess traffic to conform to a specific rate, rather than immediately discarding it. By regulating the output traffic rate, shaping ensures a smoother and more predictable network behavior during periods of high demand, improving overall network performance.

Furthermore, shaping promotes fair access to bandwidth resources, preventing any single application from monopolizing network capacity.

Authoritative resources for further reading include:

Cisco's documentation on traffic shaping: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_plc/configuration/15-sy/qos-plc-15-sy-book/qos-shpng.html
Juniper's explanation on traffic shaping:
https://www.juniper.net/documentation/us/en/software/junos/qos/topics/concept/traffic-shaping-overview.html
Cloudflare's article on traffic shaping: https://www.cloudflare.com/learning/performance/traffic-shaping/

## Question: 8

An engineer is describing QoS to a client.
Which two facts apply to traffic policing? (Choose two.)

A. Policing should be performed as close to the source as possible.

B. Policing adapts to network congestion by queuing excess traffic.

C. Policing should be performed as close to the destination as possible.

D. Policing drops traffic that exceeds the defined rate.

E. Policing typically delays the traffic, rather than drops it.

**Answer: AD**

**Explanation:**

Here's a detailed justification for why options A and D are the correct facts about traffic policing in QoS, while options B, C, and E are incorrect:

**A. Policing should be performed as close to the source as possible.** This is a fundamental principle in QoS. By policing traffic near its origin, you prevent excessive traffic from consuming bandwidth further down the network path. This reduces congestion and ensures resources are not wasted carrying traffic that will eventually be dropped. It's more efficient to discard excess traffic early rather than after it has traversed multiple hops.

**D. Policing drops traffic that exceeds the defined rate.** This is the core function of traffic policing. It enforces a specific traffic rate by discarding packets that exceed the configured limit. Unlike traffic shaping, which delays excess traffic, policing strictly drops it, often described as a "leaky bucket" mechanism.

Option B is incorrect because policing doesn't queue excess traffic. Instead, it actively discards it. Queuing is a characteristic of traffic shaping, not policing.

Option C is incorrect because policing at the destination isn't effective as it would mean the traffic has already consumed network resources unnecessarily. The primary goal is to control traffic flow at the source, not the destination.

Option E is incorrect because, as previously mentioned, policing drops traffic rather than delays it. Delaying traffic is a characteristic of traffic shaping, which buffers traffic exceeding the configured limit.

Therefore, A and D accurately describe the behavior of traffic policing.

Here are some links for further research:

1. **Cisco on QoS Policing and Shaping:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/configuration/15-mt/qos-15-mt-book/qos-polic-shape.html
2. **Juniper on Traffic Policing:** https://www.juniper.net/documentation/en_US/junos/topics/concept/firewall-policer-concepts.html 3. **Wikipedia on Traffic Policing:**https://en.wikipedia.org/wiki/Traffic_policing

**Question: 9**

Which component handles the orchestration plane of the Cisco SD-WAN?

A. vBond

B. vSmart

C. vManage

D. WAN Edge

**Answer: A**

**Explanation:**

The correct answer is **A. vBond**. The Cisco SD-WAN architecture operates with three distinct planes: management, control, and orchestration. The orchestration plane, facilitated by vBond, is responsible for the initial authentication and authorization of all SD-WAN components. This includes devices such as vEdge routers (WAN Edge) and controllers (vSmart and vManage). vBond essentially acts as the "introducer" or initial point of contact for all devices joining the SD-WAN fabric. It orchestrates the secure handshake process between different components by providing the necessary certificates and IP address information for subsequent communication. vBond also maintains a list of available vSmart controllers. It does not route traffic itself, nor does it directly manage policies or configurations.

vSmart, option B, is responsible for the control plane, where routing policy and transport decisions are made.

vManage, option C, handles the management plane, providing a centralized interface for configuration, monitoring, and troubleshooting the SD-WAN fabric. WAN Edge routers, option D, are the physical or virtual devices that reside at the edge of the network and handle the actual data forwarding. Therefore, the orchestration function is solely the responsibility of vBond, making option A the accurate choice.

Cisco SD-WAN OverviewCisco SD-WAN ComponentsCisco SD-WAN Architecture

**Question: 10**

What are two device roles in Cisco SD-Access fabric? (Choose two.)

A. edge node
B. vBond controller
C. access switch
D. core switch
E. border node

**Answer: AE**

**Explanation:**

Okay, let's break down why edge node and border node are the correct device roles in Cisco SD-Access fabric, while the other options are not.

Cisco SD-Access, a key component of Cisco's DNA Center architecture, automates network segmentation and policy. It organizes the network into a hierarchical structure. An **edge node (A)**, is where end-user devices connect to the SD-Access fabric. These devices can be wired or wireless and interface directly with the network. They provide access to the fabric for these connected users or endpoints. Therefore, this role is essential. A **border node (E)** is responsible for connecting the SD-Access fabric to external networks, like a traditional enterprise network, a data center, or the internet. It handles traffic going in and out of the SD-Access domain. Border nodes provide connectivity to resources outside the SD-Access fabric, making it a crucial function.

The **vBond controller (B)** is part of the vManage overlay network used in Cisco SD-WAN, not SD-Access. It is a control plane component that facilitates communication between other SD-WAN devices. **Access switch (C)** is a more general term describing a switching device located at the access layer, which can be part of a variety of network architectures, not exclusively SD-Access. **Core switch (D)** is also a general term relating to a high-capacity switch in the core layer of a traditional network design and is not a designated role within the SD-Access fabric. These roles are more related to traditional network design than the SD-Access specific architecture. Cisco SD-Access heavily relies on the fabric concept, which incorporates roles like edge, border, and control plane nodes, often virtualized.

In summary, Edge and border nodes are the key points for device connection and outside connectivity in a Cisco SD-Access fabric.

**Authoritative Links:**

1. **Cisco SD-Access Solution Overview:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/software-defined-access/index.html (This page offers a high-level overview of the solution.)
2. **Cisco SD-Access Design Guide:** https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html (This link leads to a comprehensive design guide that goes into device roles and functionalities.)

**Question: 11**

What is the role of the vSmart controller in a Cisco SD-WAN environment?

A. It performs authentication and authorization.
B. It manages the control plane.

C. It is the centralized network management system.

D. It manages the data plane.

**Answer: B**

**Explanation:**

The vSmart controller in a Cisco SD-WAN architecture is primarily responsible for managing the control plane. This crucial role involves establishing secure, dynamic tunnels and overlay connections between edge routers (vEdges) within the SD-WAN fabric. It facilitates the exchange of routing information and policies, enabling efficient path selection and traffic engineering. vSmart essentially acts as the brain of the SD-WAN, distributing routing updates and policies to the vEdges, ensuring consistent and coordinated behavior across the network. It does not handle data forwarding directly, which is the responsibility of vEdges in the data plane. Authentication and authorization are part of the control plane functions, but not the primary role of vSmart. The centralized network management function is attributed to vManage. Therefore, option B accurately describes the vSmart's core function within Cisco SD-WAN. The vSmart's ability to manage the control plane allows the network to be agile and dynamically responsive to changing network conditions.

Authoritative Links:

**Cisco SD-WAN Documentation:**https://www.cisco.com/c/en/us/solutions/enterprise/sd-wan/index.html (Search for vSmart and its function)
**Cisco SD-WAN Components Overview:**
https://www.cisco.com/c/en/us/td/docs/solutions/CVD/SDWAN/sdwan-design-guide.html (Look for "Control Plane Components")
**Cisco vSmart Controller Details:**https://www.cisco.com/c/en/us/products/routers/vsmart-controller/index.html

## Question: 12

When a wired client connects to an edge switch in a Cisco SD-Access fabric, which component decides whether the client has access to the network?

A. edge node

B. Identity Services Engine

C. RADIUS server

D. control-plane node

**Answer: B**

**Explanation:**

The correct answer is **B. Identity Services Engine (ISE)**. In a Cisco SD-Access fabric, the Identity Services Engine is the central policy engine responsible for authentication and authorization of users and devices. When a wired client connects to an edge switch, that switch doesn't make access decisions directly. Instead, it forwards authentication requests to ISE. ISE evaluates the client's credentials (e.g., username/password, machine certificate) against predefined policies and determines if the client is authorized to access the network, and at what level. This process typically involves communication using protocols like 802.1X. The edge node merely facilitates communication but does not enforce the access control. A RADIUS server, while often used for authentication in traditional networks, is integrated with ISE to provide underlying
authentication services. The control-plane node manages the overall network but doesn't handle user/device authentication. ISE's role is crucial because it ensures consistent policy application throughout the fabric, enabling centralized control and visibility. Furthermore, ISE supports advanced features such as profiling, posture assessment, and dynamic access control based on various contextual factors. Without ISE, the SD-

Access fabric wouldn't be able to provide secure, role-based access control.

Authoritative links for further research:

**Cisco SD-Access Design Guide:**https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html
**Cisco Identity Services Engine (ISE):**https://www.cisco.com/c/en/us/products/security/identity-services-engine/index.html
**802.1X Authentication:**https://en.wikipedia.org/wiki/IEEE_802.1X

**Question: 13**

Which benefit is offered by a cloud infrastructure deployment but is lacking in an on-premises deployment?

A. virtualization
B. supported systems
C. storage capacity
D. efficient scalability

**Answer: D**

**Explanation:**

The correct answer is D, efficient scalability, because cloud infrastructure is inherently designed for rapid and flexible scaling. Cloud platforms offer on-demand resource allocation, allowing users to easily increase or decrease computing power, storage, or network bandwidth based on real-time needs. This elasticity contrasts with on-premises infrastructure, where scaling typically involves procuring, installing, and configuring physical hardware, a process that can be time-consuming and costly. Virtualization (A) is not exclusive to the cloud; on-premises deployments can also utilize virtualization technologies. Supported systems (B) depend on specific vendors and aren't inherently a differentiator. While on-premises systems can be scaled for storage capacity (C), they are less efficient and have less flexible options compared to the rapid, on-demand scaling provided by cloud platforms. The cloud leverages a shared resource pool, allowing for significant economies of scale and simplified management when adjusting resources. This enables organizations to adapt quickly to changing demands, optimize costs by only paying for what they use, and improve agility. Efficient scalability is a core tenet of cloud computing.

For further reading on cloud scalability, consider these resources:

**Amazon Web Services (AWS) Documentation on Autoscaling:**https://aws.amazon.com/autoscaling/ **Microsoft Azure Documentation on Scalability:**https://learn.microsoft.com/en-us/azure/architecture/guide/design-principles/scalability
**Google Cloud Platform (GCP) Documentation on Autoscaling:**
https://cloud.google.com/compute/docs/autoscaler/
**National Institute of Standards and Technology (NIST) Definition of Cloud Computing:**
https://csrc.nist.gov/publications/detail/sp/800-145/final (Specifically, the characteristic of “rapid elasticity”.)

**Question: 14**

Which action is the vSmart controller responsible for in a Cisco SD-WAN deployment?

A. onboard WAN Edge nodes into the Cisco SD-WAN fabric
B. gather telemetry data from WAN Edge routers

C. distribute policies that govern data forwarding performed within the Cisco SD-WAN fabric

D. handle, maintain, and gather configuration and status for nodes within the Cisco SD-WAN fabric

**Answer: C**

**Explanation:**

The correct answer is **C. distribute policies that govern data forwarding performed within the Cisco SD-WAN fabric.**

Here's a detailed justification: In Cisco SD-WAN, the vSmart controller serves as the centralized policy engine. It doesn't directly handle the day-to-day operations of edge devices, but it dictates how they should behave. Specifically, vSmart's primary role involves the creation and dissemination of control plane policies. These policies determine crucial aspects such as path selection, traffic engineering, quality of service (QoS), and security rules. The vSmart controller uses the Overlay Management Protocol (OMP) to distribute these policies to the vEdge and cEdge routers in the SD-WAN fabric. These edge routers, in turn, enforce these policies.

Option A is incorrect because the vManage orchestrator is primarily responsible for onboarding WAN Edge nodes through Zero Touch Provisioning (ZTP). Option B is incorrect because vManage is responsible for collecting telemetry data from the WAN Edge routers and displaying it in the GUI. Option D is also incorrect because vManage handles device configuration and status within the SD-WAN. The vSmart controller's key function is about policy distribution, not direct configuration, onboarding, or data collection. It acts as the brain of the control plane, ensuring consistent and centrally managed behavior across the entire SD-WAN.

Therefore, the vSmart controller doesn't onboard devices or collect data directly; its responsibility lies in defining and distributing the forwarding rules that the WAN Edge routers must obey. It's crucial for maintaining consistent policy enforcement across the SD-WAN environment.

For further reading, refer to these resources:

**Cisco SD-WAN Design Guide:**https://www.cisco.com/c/en/us/td/docs/solutions/CVD/SDWAN/cisco-sdwan-design-guide.html (Look for sections explaining vSmart controller functionality and OMP)
**Cisco SD-WAN Configuration Guides:** (Search on the Cisco website for relevant guides, focusing on the vSmart controller section).

## Question: 15

Where is radio resource management performed in a Cisco SD-Access wireless solution?

A. DNA Center

B. control plane node

C. wireless controller

D. Cisco CMX

**Answer: B**

**Explanation:**

Radio Resource Management (RRM) in a Cisco SD-Access wireless solution, specifically the functions related to channel assignment, power level adjustment, and interference mitigation, is primarily performed at the **control plane node (B)**. While various components interact, the control plane node, often a WLC in a traditional Cisco wireless context but acting in a different capacity within SD-Access, holds the intelligence to make centralized RRM decisions. This is a core function of the software-defined networking (SDN) approach where control is centralized and abstracted from the underlying infrastructure. The control plane node gathers network data, analyzes wireless conditions, and calculates optimal RRM parameters based on policies and inputs. This contrasts with individual access points (APs) making localized decisions, as seen in autonomous deployments. The control plane then pushes these calculated configurations to the access points. This centralization is key to ensuring consistent and efficient RF management across the entire SD-Access wireless fabric. DNA Center (A) provides network management and orchestration, including policy definition, but does not directly perform real-time RRM calculations. Wireless Controllers (C), in traditional architectures, are responsible for RRM, but in the SD-Access fabric, their function is more of an access point termination point; the brains are at the Control Plane. Cisco CMX (D) focuses on location analytics and does not participate in RRM. Therefore, the control plane node is where the core RRM functions are executed within the SD-Access wireless environment.

**Authoritative Links:**

Cisco SD-Access Wireless Design Guide: https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/sda-wireless-design-guide-2021oct.html (Focus on control plane functionality, especially in the context of wireless)
Cisco DNA Center Documentation: https://www.cisco.com/c/en/us/products/cloud-systems-management/dna-center/index.html (General overview; emphasizes management, not direct RRM execution).

**Question: 16**

DRAG DROP -
Drag and drop the characteristics from the left onto the infrastructure types on the right.
Select and Place:

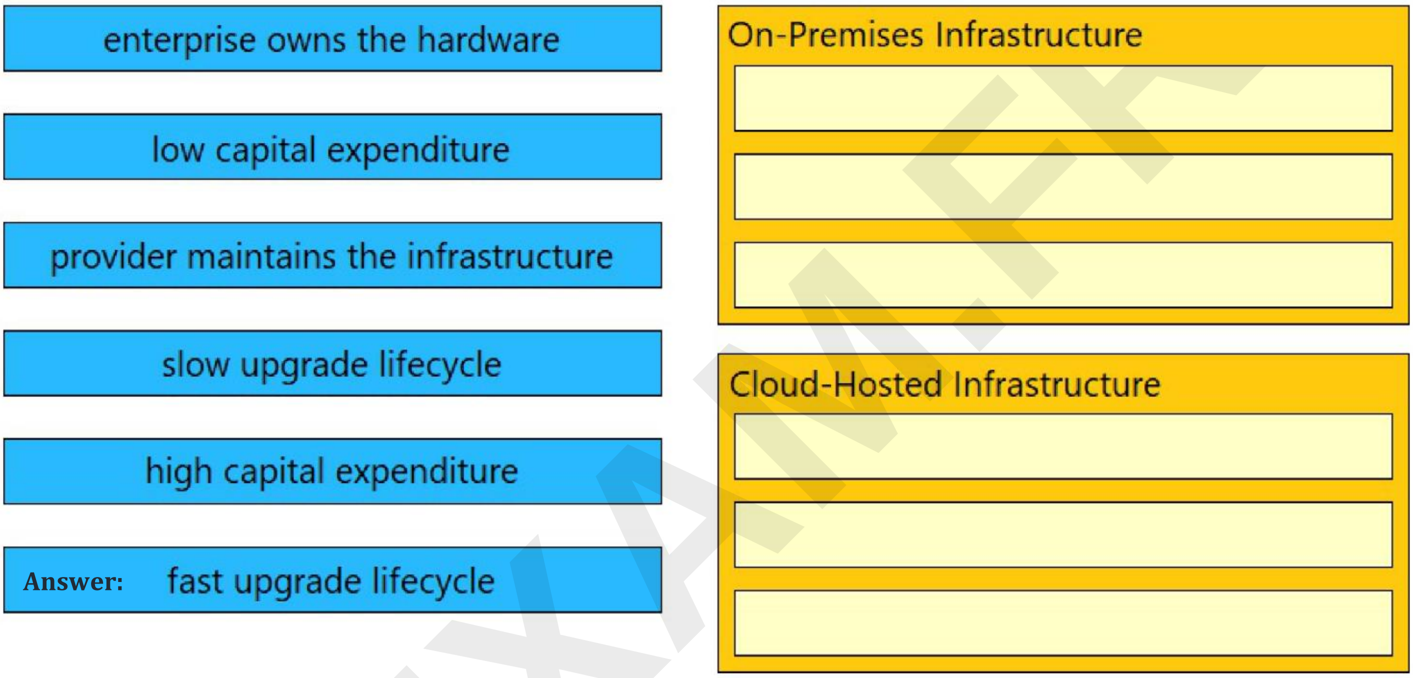

enterprise owns the hardware

low capital expenditure

provider maintains the infrastructure

slow upgrade lifecycle

high capital expenditure

fast upgrade lifecycle

On-Premises Infrastructure

Cloud-Hosted Infrastructure

**Answer:**

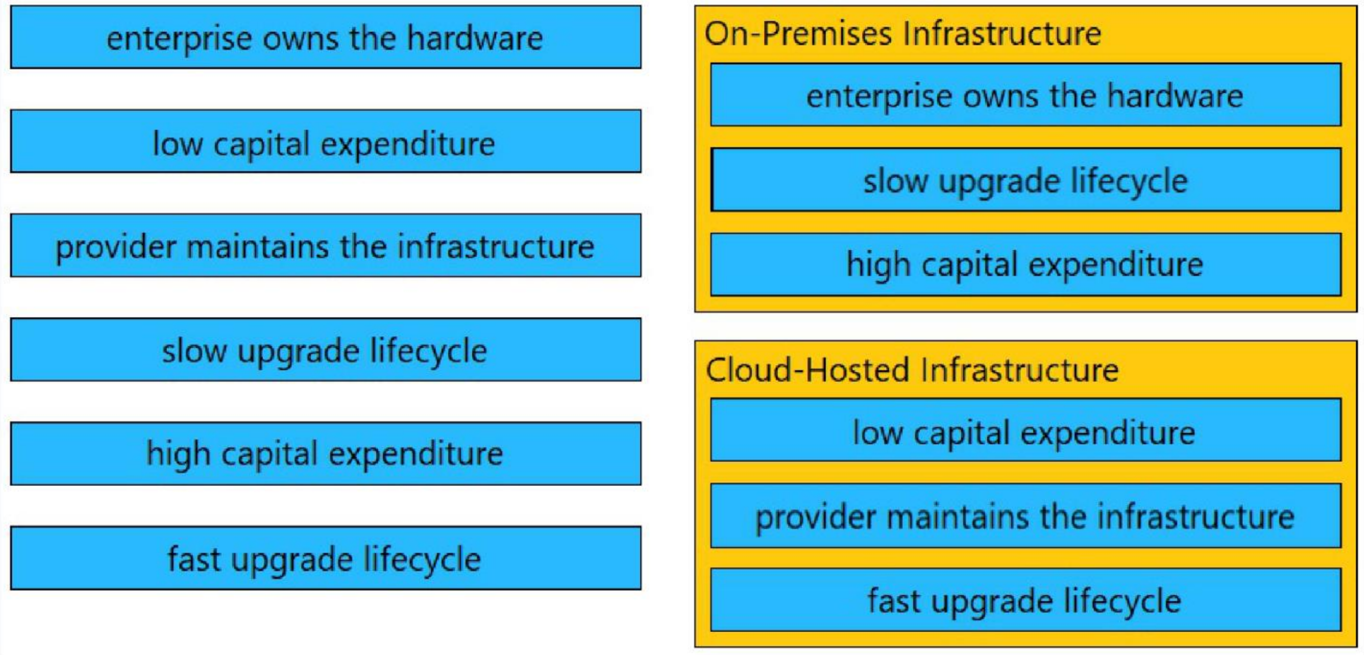

enterprise owns the hardware

low capital expenditure

provider maintains the infrastructure

slow upgrade lifecycle

high capital expenditure

fast upgrade lifecycle

On-Premises Infrastructure

- enterprise owns the hardware
- slow upgrade lifecycle
- high capital expenditure

Cloud-Hosted Infrastructure

- low capital expenditure
- provider maintains the infrastructure
- fast upgrade lifecycle

**Question: 17**

How does the RIB differ from the FIB?

A. The FIB maintains network topologies and routing tables. The RIB is a list of routes to particular network destinations.

B. The FIB includes many routes to a single destination. The RIB is the best route to a single destination.

C. The RIB is used to create network topologies and routing tables. The FIB is a list of routes to particular network destinations.

D. The RIB includes many routes to the same destination prefix. The FIB contains only the best route.

**Answer: D**

**Explanation:**

Here's a breakdown of why option D is the correct answer regarding the difference between the Routing Information Base (RIB) and the Forwarding Information Base (FIB) in Cisco networking:

The RIB, also known as the routing table, acts as a database storing all learned routes to various network destinations. These routes are gathered from different routing protocols, static configurations, and connected interfaces. The RIB often contains multiple paths to the same destination, each potentially learned through a different mechanism. These paths are often differentiated by attributes like metric, administrative distance and source. The router uses a best-path selection algorithm to choose the optimal route from the RIB to install into the FIB. This best path is selected based on factors like lowest metric, preference, and highest bandwidth.

The FIB, on the other hand, is a streamlined table optimized for rapid packet forwarding. It doesn't store all possible routes; rather, it contains only the active, best route to each destination prefix chosen from the RIB. This optimized view is crucial for efficient packet forwarding with minimal lookup time. The FIB is structured for high-speed hardware lookup, enabling routers to make fast forwarding decisions without parsing complex tables or evaluating multiple routing protocols per packet. Thus, the FIB directly maps network prefixes to next-hop information, allowing for immediate lookup during packet transmission.

Therefore, option D accurately portrays the core difference: The RIB holds multiple, potential routes to a destination (even the same prefix), while the FIB holds only the single, most optimal route. This design enables efficient route processing and fast forwarding.

Here are some authoritative links for further research:

**Cisco: Understanding the Routing Table:**https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/15264-route-table.html
**Cisco: Forwarding Information Base (FIB) Overview:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipswitch/configuration/15-mt/ips-15-mt-book/ips-cef.html#GUID-D187E614-C2DF-4E10-8F12-A3D91435B06B
**Juniper Networks: Routing Table Fundamentals:**
https://www.juniper.net/documentation/us/en/software/junos/routing/topics/concept/routing-table-fundamentals.html

**Question: 18**

Which technology is used to provide Layer 2 and Layer 3 logical networks in the Cisco SD-Access architecture?

A. underlay network
B. VPN routing/forwarding
C. easy virtual network
D. overlay network

**Answer: D**

**Explanation:**

The correct answer is D, overlay network. In Cisco SD-Access, an overlay network is the technology that provides Layer 2 and Layer 3 logical network segmentation. An overlay network sits on top of the physical underlay network, decoupling the logical topology from the physical one. This decoupling allows for the creation of multiple independent virtual networks over a single physical infrastructure. These virtual networks are referred to as Virtual Networks (VNs) in SD-Access, and they can each have different routing and forwarding behaviors, independent of the underlay. An overlay network like Virtual Extensible LAN (VXLAN), which SD-Access heavily utilizes, encapsulates Layer 2 frames inside UDP packets, allowing them to traverse the underlay network while maintaining Layer 2 isolation and segmentation between different VNs. VPN Routing/Forwarding (VRF), while a valid segmentation technique, is leveraged within the overlay to provide unique routing domains but it is not the overarching technology providing the logical networks as a whole.

Easy Virtual Network (EVN), a Cisco proprietary segmentation technology, is less directly involved in SD-Access overlay implementation compared to VXLAN. The underlay network, conversely, provides the basic IP connectivity upon which the overlay functions but does not offer logical network segmentation itself. Thus, overlay networks, with their encapsulation and logical network construction, are central to enabling Layer 2 and Layer 3 segmentation in Cisco SD-Access architectures.

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

Cisco SD-Access Solution Overview: https://www.cisco.com/c/en/us/solutions/enterprise/software-defined-access/index.html
VXLAN Technology Overview: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/virtualized-infrastructure-data-center/white-paper-c11-729371.html

**Question: 19**

What is the difference between CEF and process switching?

A. CEF processes packets that are too complex for process switching to manage.
B. Process switching is faster than CEF.
C. CEF uses the FIB and the adjacency table to make forwarding decisions, whereas process switching punts

each packet.

D. CEF is more CPU-intensive than process switching.

**Answer: C**

**Explanation:**

The correct answer is **C: CEF uses the FIB and the adjacency table to make forwarding decisions, whereas process switching punts each packet.**

Here's a detailed justification:

Cisco Express Forwarding (CEF) is a crucial packet-forwarding mechanism within Cisco routers. It dramatically improves forwarding performance compared to older methods like process switching. CEF utilizes two key data structures: the Forwarding Information Base (FIB) and the adjacency table. The FIB mirrors the routing table, containing optimal routes for each destination network. The adjacency table stores precomputed Layer 2 information (e.g., MAC addresses) for next-hop devices. This allows the router to quickly determine both the exit interface and the necessary Layer 2 rewrite information to forward a packet.

Process switching, on the other hand, handles each packet individually. Upon receiving a packet, the router's CPU must analyze the packet header and consult the routing table. This per-packet lookup and decision-making process is CPU-intensive and causes significant performance overhead, especially with high traffic volumes. It's often referred to as “punt-to-CPU” behavior, as each packet is handled by the CPU’s software.

Therefore, CEF offers a more efficient and scalable approach by pre-calculating forwarding paths and storing them in specialized data structures. This reduces the processing burden on the CPU and provides faster forwarding. Option A is incorrect because process switching doesn't handle complexity, it simply does so inefficiently. Option B is incorrect because CEF is considerably faster than process switching. Option D is incorrect as process switching is more CPU intensive because it analyzes each packet individually, while CEF only updates when the FIB/adjacency tables need to be updated.

For further research, consult the following authoritative Cisco documentation:

**Cisco Express Forwarding (CEF):**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipswitch/configuration/15-sy/ips-15-sy-book/ips-cef.html
**How does process switching work?:**https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/ios-software-releases-122-mainline/13081-process-switched.html

## Question: 20

What are two considerations when using SSO as a network redundancy feature? (Choose two.)

A. requires synchronization between supervisors in order to guarantee continuous connectivity

B. the multicast state is preserved during switchover

C. must be combined with NSF to support uninterrupted Layer 3 operations

D. both supervisors must be configured separately

E. must be combined with NSF to support uninterrupted Layer 2 operations

**Answer: AC**

**Explanation:**

The correct answer is **A and C.**

**A: requires synchronization between supervisors in order to guarantee continuous connectivity.** Single

Session Object (SSO), also known as Stateful Switchover (SSO) in Cisco context, relies on active-standby redundancy architecture. The active supervisor handles all control plane operations, while the standby maintains a synchronized mirror copy. This synchronization is crucial to ensure a seamless transition if the active supervisor fails. Without proper synchronization, the standby cannot take over the active role smoothly, leading to service disruption. The standby must have an identical control-plane state to ensure uninterrupted forwarding.

**C: must be combined with NSF to support uninterrupted Layer 3 operations.** Non-Stop Forwarding (NSF) complements SSO by focusing on uninterrupted forwarding. While SSO guarantees a smooth control-plane switchover, it does not ensure packet forwarding continuity on its own. When a failover occurs, routing protocols need time to reconverge, which results in brief packet loss. NSF, typically combined with protocols like OSPF and BGP, allows the forwarding plane to continue operating using the last known state while the control plane is converging, minimizing packet loss. It is the combined use of SSO (for the control plane) and NSF (for the forwarding plane) that results in a highly available Layer 3 network.

**Why the other options are incorrect:**

**B: the multicast state is preserved during switchover:** While SSO aims to preserve most state, multicast state isn't always preserved across a switchover, and depends on the specific platform and software version. Multicast recovery is complex and often requires protocols like PIM to re-establish the multicast distribution tree.

**D: both supervisors must be configured separately:** One of the key benefits of SSO is that the standby supervisor doesn't need to be configured independently; instead it automatically synchronizes its configuration with the active.

**E: must be combined with NSF to support uninterrupted Layer 2 operations:** NSF is generally associated with Layer 3 operations. For Layer 2 operations, features such as Etherchannel and Spanning-tree variations provide redundancy rather than NSF. While SSO helps preserve L2 tables, it is the combination of those features, not NSF, which is crucial for uninterrupted L2.

**Authoritative Links for Further Research:**

**Cisco Documentation on Stateful Switchover:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ha/configuration/16-6/ha-16-6-book/ha-stateful-switch.html
**Cisco Documentation on Nonstop Forwarding:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ha/configuration/15-sy/ha-15-sy-book/ha-nsf.html

These links provide detailed information regarding Cisco's implementations of SSO and NSF and their importance in achieving high availability in enterprise networks.

## Question: 21

DRAG DROP -
Drag and drop the characteristics from the left onto the correct infrastructure deployment type on the right.
Select and Place:

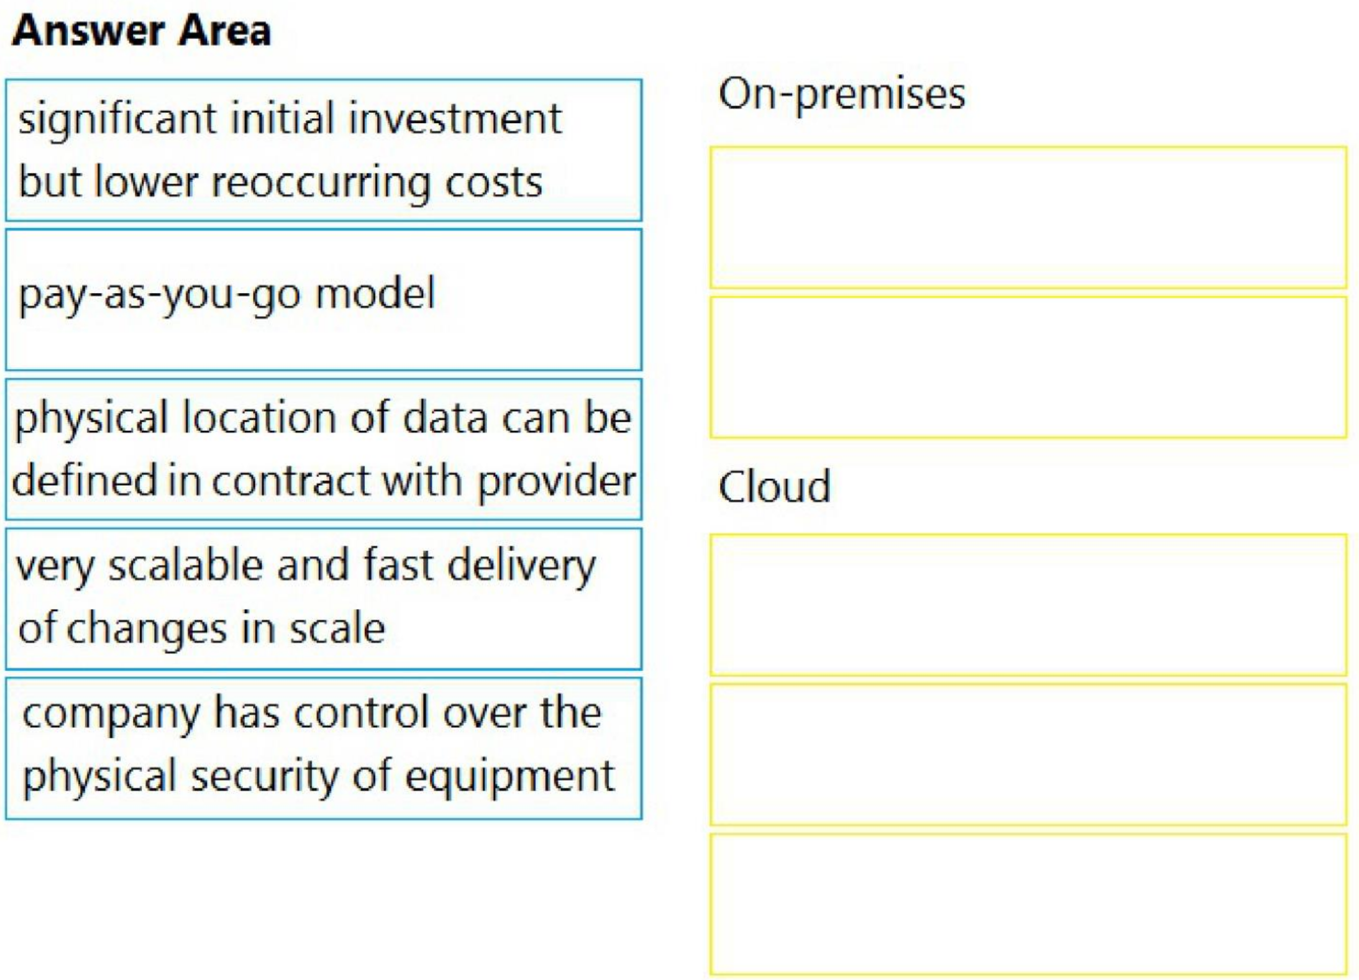

**Answer Area**

significant initial investment but lower reoccurring costs

pay-as-you-go model

physical location of data can be defined in contract with provider

very scalable and fast delivery of changes in scale

company has control over the physical security of equipment

On-premises

Cloud

**Answer:**

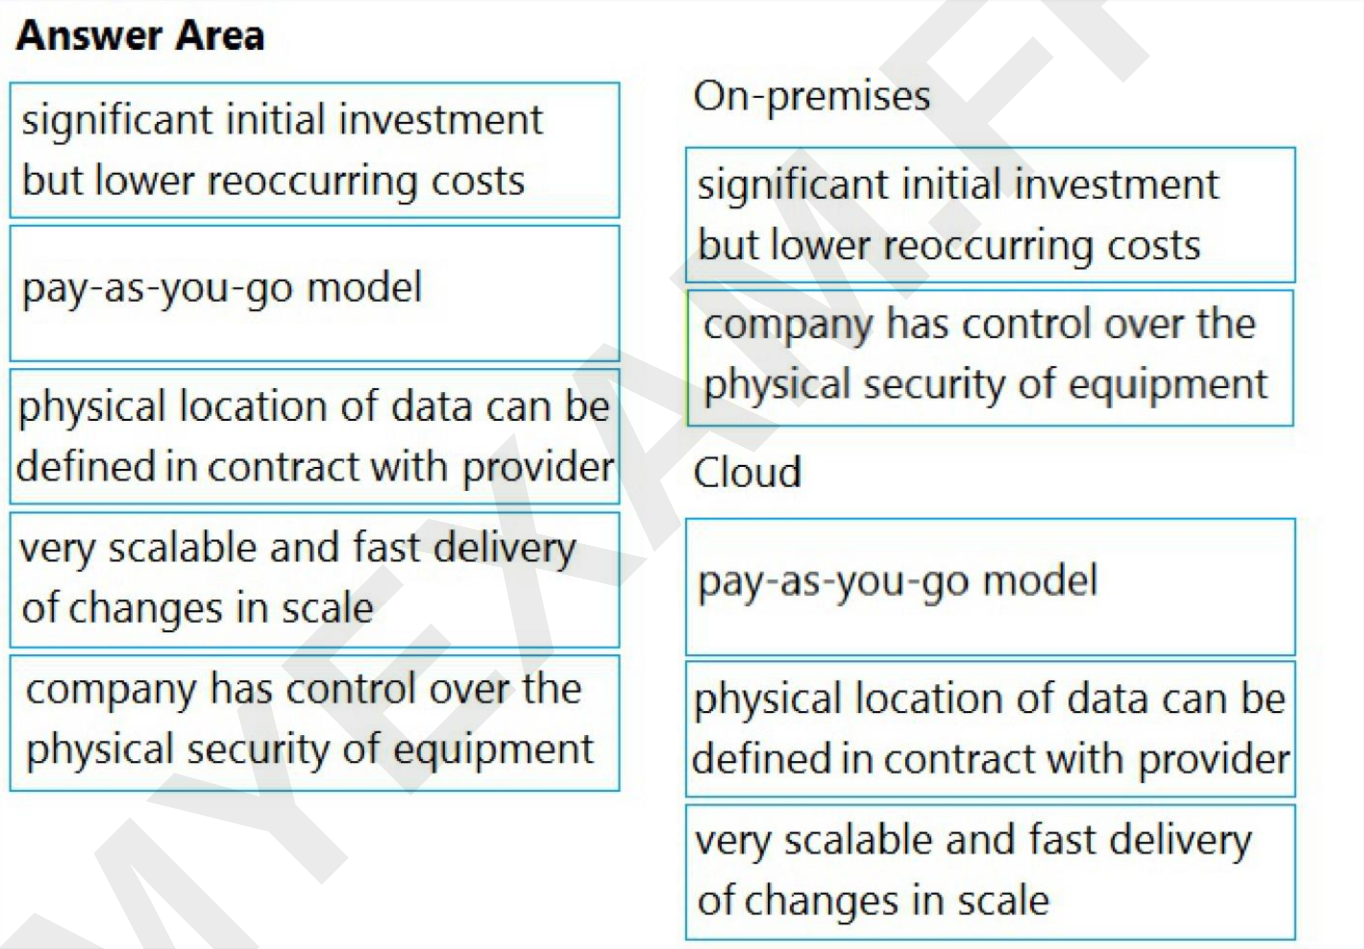

**Answer Area**

significant initial investment but lower reoccurring costs

pay-as-you-go model

physical location of data can be defined in contract with provider

very scalable and fast delivery of changes in scale

company has control over the physical security of equipment

On-premises

- significant initial investment but lower reoccurring costs
- company has control over the physical security of equipment

Cloud

- pay-as-you-go model
- physical location of data can be defined in contract with provider
- very scalable and fast delivery of changes in scale

**Question: 22**

In a Cisco SD-Access fabric, which control plane protocol is used for mapping and resolving endpoints?

A. DHCP
B. VXLAN
C. SXP
D. LISP

**Answer: D**

**Explanation:**

The correct answer is D. LISP (Locator/ID Separation Protocol) is the core control plane protocol responsible for mapping and resolving endpoints within a Cisco SD-Access fabric. In SD-Access, endpoint identifiers (EIDs) which are typically IP addresses are separated from their routing locators (RLOCs) which are typically the IP addresses of fabric edge nodes. This separation allows for mobility and policy enforcement within the fabric. LISP maintains this mapping within its database. When a host sends traffic, the edge node consults the LISP mapping system to find the RLOC associated with the destination EID. This RLOC is then used to encapsulate the traffic in VXLAN. DHCP (Option A) assigns IP addresses but doesn't participate in mapping endpoint locations within the fabric. VXLAN (Option B) is the data plane encapsulation used for transport after the LISP lookup is complete. SXP (Option C) is used for Security Group Tag (SGT) exchange and does not perform the EID to RLOC mapping function. The key function of LISP is to allow for routing based on location rather than just the IP address. This feature enables the mobility and policy enforcement benefits provided by SD-Access.

Further research on this can be found at the following authoritative links:

**Cisco SD-Access Overview:**https://www.cisco.com/c/en/us/solutions/enterprise/software-defined-access/index.html
**LISP Overview:**https://www.cisco.com/c/en/us/solutions/enterprise/campus-lan-switching/lisp.html **Cisco SD-Access Design Guide:** (Search Cisco documentation for the most recent SD-Access design guide)

**Question: 23**

What are two differences between the RIB and the FIB? (Choose two.)

A. FIB is a database of routing prefixes, and the RIB is the information used to choose the egress interface for each packet.
B. The FIB is derived from the data plane, and the RIB is derived from the FIB.
C. The RIB is a database of routing prefixes, and the FIB is the information used to choose the egress interface for each packet.
D. The RIB is derived from the control plane, and the FIB is derived from the RIB.
E. The FIB is derived from the control plane, and the RIB is derived from the FIB.

**Answer: CD**

**Explanation:**

Let's analyze the differences between the Routing Information Base (RIB) and the Forwarding Information Base (FIB) within a Cisco network, focusing on their roles and data sources.

**C. The RIB is a database of routing prefixes, and the FIB is the information used to choose the egress interface for each packet.** This is correct. The RIB, often referred to as the routing table, is essentially a comprehensive listing of all known network routes. It stores prefixes (network destinations) along with associated attributes like next-hop addresses, metric, and administrative distances. The FIB, on the other hand, is a streamlined table optimized for efficient packet forwarding. It contains a subset of the RIB's information that is directly used to make routing decisions for every incoming packet, mapping destinations to egress interfaces.

**D. The RIB is derived from the control plane, and the FIB is derived from the RIB.** This statement is also correct. The control plane is responsible for building and maintaining routing tables by executing routing protocols and network configurations. The RIB is populated by the control plane. Then, the FIB is derived from the RIB's selected best routes, formatted for rapid lookup during data forwarding. This ensures that forwarding decisions are made quickly using an optimized data structure. The control plane essentially computes the routes, whereas the FIB focuses on data plane operations.

**Incorrect options:**

**A.** The description is reversed. The RIB holds network prefixes, while the FIB is concerned with egress interface selection.

**B.** The FIB is not derived from the data plane. The FIB is derived from the RIB which is created within the Control Plane. The data plane handles packet forwarding using the information in FIB.

**E.** The FIB is not derived from the control plane. Instead, it's derived from the RIB. The RIB is generated via the control plane.

**In summary:** The RIB acts as a comprehensive database created by the control plane containing routing information, whereas the FIB is a data-plane-optimized subset used for actual packet forwarding, generated from the RIB. The RIB contains routing prefixes, while the FIB links prefixes to egress interfaces.

**Authoritative Links:**

**Cisco - Routing Information Base:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/15-sy/irr-15-sy-book/ip6-route-rib.html
**Cisco - Forwarding Information Base:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/15-sy/irr-15-sy-book/ip6-fwd-fib.html

## Question: 24

Which two network problems indicate a need to implement QoS in a campus network? (Choose two.)

A. port flapping
B. excess jitter
C. misrouted network packets
D. duplicate IP addresses
E. bandwidth-related packet loss

**Answer: BE**

**Explanation:**

Here's a detailed justification for why options B (excess jitter) and E (bandwidth-related packet loss) indicate a need for QoS implementation in a campus network:

Quality of Service (QoS) mechanisms are designed to prioritize and manage network traffic based on specific criteria. Jitter, defined as variations in packet arrival time, directly impacts real-time applications like voice and video. Excessive jitter leads to degraded user experience, causing distorted audio or choppy video playback. This occurs when different packets belonging to the same flow experience varying delays across the network, often due to network congestion. QoS can address this by prioritizing real-time traffic, ensuring timely delivery and smoother user experience.

Bandwidth-related packet loss arises when the network's capacity is exceeded. When the rate of incoming packets surpasses the available bandwidth on a link or device, routers or switches start dropping packets, resulting in data loss. This impacts all applications, but can severely hinder time-sensitive applications. QoS techniques such as traffic shaping and prioritization, allow allocating bandwidth to critical applications while dropping low-priority packets. This ensures sufficient bandwidth is available for crucial flows, thus minimizing packet loss.

Options A, C, and D are not primarily addressed by QoS. Port flapping (A) indicates a physical layer issue or misconfiguration and is typically resolved through troubleshooting. Misrouted network packets (C) are a routing issue, and duplicate IP addresses (D) are a network addressing issue; both are not directly related to traffic prioritization or management provided by QoS. These problems can lead to network performance issues, but not specifically resolved with QoS mechanism.

In summary, excess jitter and bandwidth-related packet loss are direct symptoms of network congestion and a lack of prioritization, making them prime candidates for QoS solutions in a campus network. By implementing QoS, a network can manage these issues and provide reliable and consistent user experience.

**Authoritative Links for Further Research:**

**Cisco QoS:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/quality-of-service-qos/index.html
**Understanding Jitter:**https://www.cloudflare.com/learning/performance/what-is-jitter/
**Packet Loss and its effect:**https://www.cloudflare.com/learning/performance/what-is-packet-loss/

## Question: 25

In a Cisco SD-Access wireless architecture, which device manages endpoint ID to edge node bindings?

A. fabric control plane node
B. fabric wireless controller
C. fabric border node
D. fabric edge node

**Answer: A**

**Explanation:**

The correct answer is **A. fabric control plane node.** In a Cisco Software-Defined Access (SD-Access) architecture, the fabric control plane node, often a Cisco DNA Center or ISE server, acts as the central intelligence and policy engine. It's responsible for maintaining the endpoint ID to edge node bindings, essentially mapping where a specific device is connected to the network. This binding information is crucial for routing traffic efficiently and enforcing consistent policies. The fabric control plane distributes these bindings to the fabric edge nodes, enabling them to forward traffic based on the correct virtual network and policy attributes. The fabric wireless controller (B) manages wireless access points, not endpoint-to-edge bindings directly. Fabric border nodes (C) handle traffic between the SD-Access fabric and external networks, while fabric edge nodes (D) are the access layer switches where endpoints physically connect. These edge nodes enforce policies learned from the control plane, but don't manage the initial binding process. Therefore, the fabric control plane node serves as the single source of truth for mapping endpoint IDs and their network location within the SD-Access fabric. It maintains this mapping, ensuring consistent and scalable endpoint management.

Authoritative links for further research:

Cisco SD-Access Solution Overview: https://www.cisco.com/c/en/us/solutions/enterprise/software-defined-access/index.html
Cisco DNA Center Documentation: https://www.cisco.com/c/en/us/support/cloud-systems-management/dna-center/series.html

**Question: 26**

DRAG DROP -
Drag and drop the QoS mechanisms from the left onto their descriptions on the right.
Select and Place:

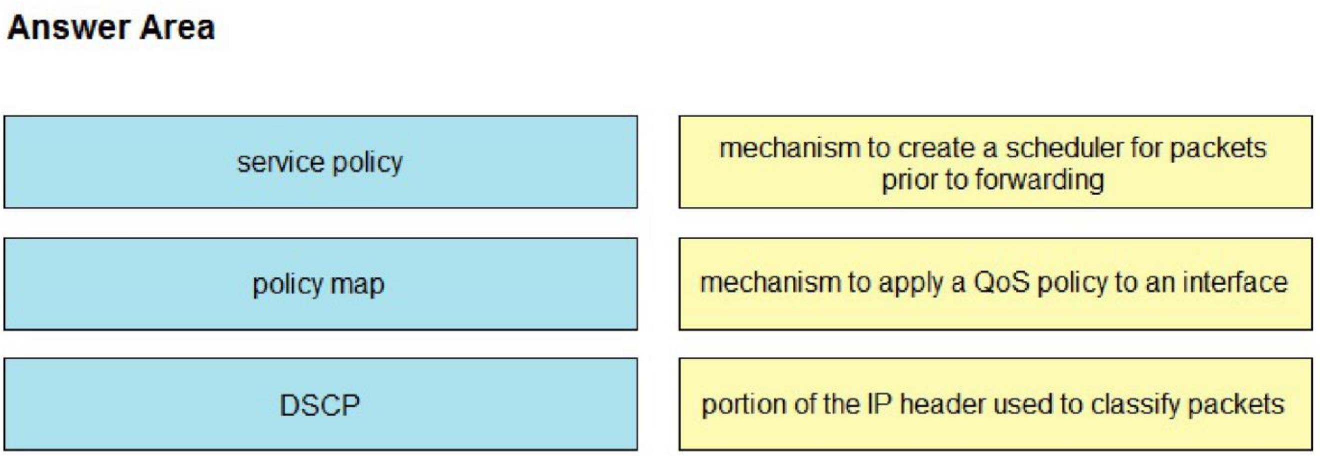

**Answer Area**

| | |
|---|---|
| service policy | mechanism to create a scheduler for packets prior to forwarding |
| policy map | mechanism to apply a QoS policy to an interface |
| DSCP | portion of the IP header used to classify packets |

**Answer:**

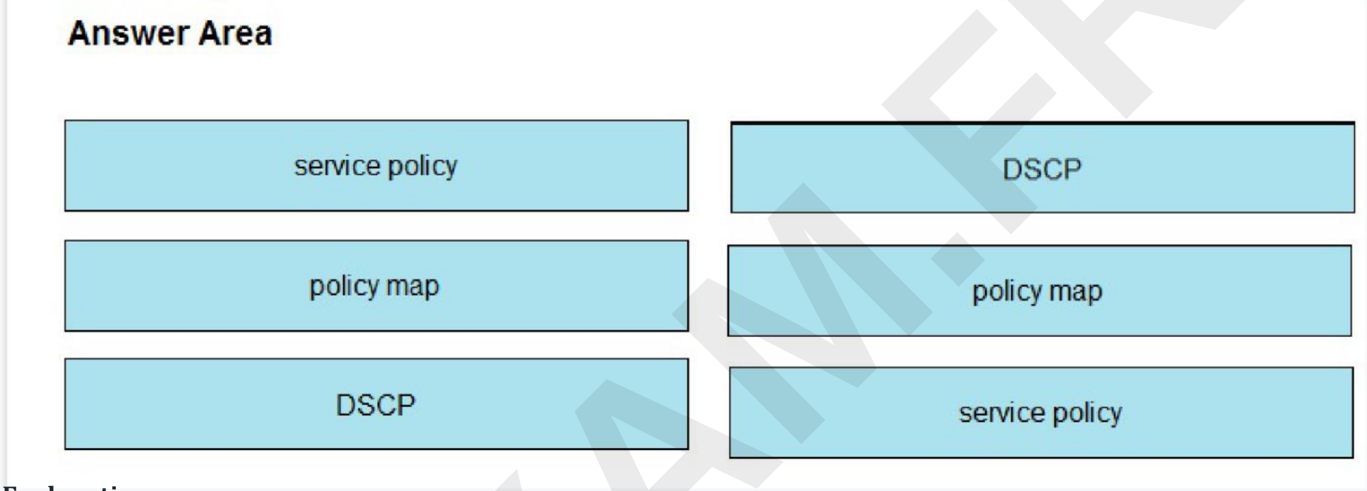

**Answer Area**

| | |
|---|---|
| service policy | DSCP |
| policy map | policy map |
| DSCP | service policy |

**Explanation:**

DSCP == portion of the IP header used to classify the packets

policy map == mechanism to create a scheduler for packets prior to forwarding

service policy == mechanism to apply a QoS policy to an interface

**Question: 27**

Which control plane protocol is used between Cisco SD-WAN routers and vSmart controllers?

A. TCP

B. OMP

C. UDP

D. BGP

**Answer: B**

**Explanation:**

The correct answer is **B. OMP (Overlay Management Protocol).** In Cisco SD-WAN, OMP is the control plane protocol that governs the communication between vSmart controllers and vEdge/cEdge routers. It's crucial for establishing and maintaining the SD-WAN fabric. OMP is not a typical routing protocol like BGP or OSPF but rather a specialized protocol designed for SD-WAN overlay networking. It's responsible for distributing routing information, security policies, and other control plane data within the SD-WAN domain. Routers learn about the SD-WAN network topology and policies via OMP updates received from the vSmart controllers. This allows them to dynamically build their forwarding tables without direct peer-to-peer relationships. OMP operates over a secure DTLS (Datagram Transport Layer Security) connection, ensuring the confidentiality and integrity of control plane communication. Therefore, options A (TCP), C (UDP), and D (BGP) are incorrect as they serve different purposes in networking. OMP is uniquely designed for Cisco SD-WAN control plane functionality.

**Authoritative Links:**

**Cisco SD-WAN Overlay Management Protocol (OMP):**
https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/routing/vedge-routing-config-book/sdwan-routing-config-book_chapter_0110.html
**Cisco SD-WAN Control Plane Overview:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/sd-wan/index.html#~stickynav=2

**Question: 28**

In a three-tier hierarchical campus network design, which action is a design best-practice for the core layer?

A. provide QoS prioritization services such as marking, queueing, and classification for critical network traffic

B. provide redundant Layer 3 point-to-point links between the core devices for more predictable and faster convergence

C. provide advanced network security features such as 802.1X, DHCP snooping, VACLs, and port security

D. provide redundant aggregation for access layer devices and first-hop redundancy protocols such as VRRP

**Answer: B**

**Explanation:**

The best practice for the core layer in a three-tier hierarchical campus network is to provide redundant Layer 3 point-to-point links between core devices (Option B). The core layer's primary function is to facilitate fast and reliable transport of data between different parts of the network. Redundant links ensure high availability and prevent single points of failure. Layer 3 links, using routing protocols, allow for efficient path selection and fast convergence during failures. This directly contributes to the core's role of being the network backbone. Options A, C, and D, while important, are typically functionalities associated with the distribution and access layers. QoS prioritization (A) and advanced security features (C) are usually implemented closer to the network's edge, where traffic originates and terminates. Aggregation and first-hop redundancy protocols (D) are primarily handled in the distribution layer. Core layer focuses on high-speed, efficient forwarding. The use of point-to-point links further simplifies routing and enhances overall network performance and stability in the core.

For further research:

Cisco's Hierarchical Network Design Model: https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-network-design-guide.html
Understanding the Three-Tier Model: https://www.networkworld.com/article/2229418/the-3-tier-network-model-defined.html

## Question: 29

What is a VPN in a Cisco SD-WAN deployment?

A. common exchange point between two different services
B. attribute to identify a set of services offered in specific places in the SD-WAN fabric
C. virtualized environment that provides traffic isolation and segmentation in the SD-WAN fabric
D. virtual channel used to carry control plane information

**Answer: C**

**Explanation:**

The correct answer is C: "virtualized environment that provides traffic isolation and segmentation in the SD-WAN fabric." In a Cisco SD-WAN deployment, a VPN (Virtual Private Network) is not a physical connection or a simple access point; rather, it's a logical construct. It functions as a virtualized space within the SD-WAN fabric that ensures traffic isolation and segmentation. This means different VPNs can be established for various departments, user groups, or applications, keeping their data separate and secure within the same network infrastructure. These VPNs are not related to traditional IPsec or VPN tunnels used for connecting remote networks; instead, they are logical overlays. This approach enhances security and enables granular traffic management. Each VPN is assigned a unique VPN ID, which allows the SD-WAN controller to enforce specific policies, including routing, QoS, and security rules, for traffic within that VPN. This segmentation is vital for enterprise-level networks where different departments may need separate network environments with unique requirements. Option A refers more to a peering point or an interface and is incorrect. Option B describes a feature used in service provider networks, not SD-WAN VPNs. Option D refers to a control plane channel but not the data segmentation provided by VPNs.

Authoritative links:

**Cisco SD-WAN Configuration Guide:** https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/sdwan-xe-config/m-sdwan-xe-config-book.html (Refer to sections on VPNs and segmentation)
**Cisco SD-WAN Design Guide:**https://www.cisco.com/c/en/us/solutions/enterprise/software-defined-wan-sd-wan/design-guide.html (Focus on the design principles and multi-tenancy aspects of SD-WAN)

## Question: 30

Which function does a fabric edge node perform in an SD-Access deployment?

A. Connects endpoints to the fabric and forwards their traffic.
B. Encapsulates end-user data traffic into LISP.
C. Connects the SD-Access fabric to another fabric or external Layer 3 networks.
D. Provides reachability between border nodes in the fabric underlay.

**Answer: A**

**Explanation:**

The correct answer is **A. Connects endpoints to the fabric and forwards their traffic.** In a Cisco SD-Access (Software-Defined Access) deployment, the fabric edge node serves as the first point of attachment for user endpoints like PCs, printers, and IP phones. These nodes are responsible for directly connecting devices to the network. Upon device connection, the edge node authenticates and authorizes them, mapping them to appropriate virtual networks (VN) or segments. Subsequently, the edge node forwards the traffic to the fabric control plane (control node) for further processing based on policy. This initial connection and traffic forwarding constitutes the fundamental role of the fabric edge node. It doesn't typically encapsulate traffic into LISP; that's primarily the role of a fabric border node when connecting to external networks. Neither does it offer reachability between border nodes in the underlay; this is handled by the underlay network. The primary function focuses on endpoint attachment and initial forwarding within the SD-Access fabric. Edge nodes bridge the physical infrastructure with the logical segmentation defined by SD-Access, creating a virtualized access layer.

**Further Research Links:**

**Cisco SD-Access Solution Overview:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/software-defined-access/index.html
**Cisco SD-Access Design Guide:**https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html
**Cisco SD-Access Component Overview:**https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/2-2-3/b_cisco_dna_center_2_2_3_user_guide/m_sda_fabric_components.html

## Question: 31

What is the role of a fusion router in an SD-Access solution?

A. acts as a DNS server

B. provides additional forwarding capacity to the fabric

C. performs route leaking between user-defined virtual networks and shared services

D. provides connectivity to external networks

**Answer: C**

**Explanation:**

The correct answer is **C. performs route leaking between user-defined virtual networks and shared services**. In an SD-Access (Software-Defined Access) solution, the fusion router plays a crucial role in enabling communication between different Virtual Networks (VNs). SD-Access uses VNs to segment the network and isolate traffic based on user groups or applications. However, certain shared services like DNS, DHCP, or internet access need to be available to all VNs. This is where the fusion router comes in. It's not primarily for added forwarding capacity (option B), though it does handle routing. It doesn't function as a DNS server (option A), though it may route traffic to one. While it does connect to external networks (option D), its primary role isn't simply providing that. Instead, the fusion router allows controlled route leaking; it selectively permits routes to and from shared services to be advertised to specific VNs, and vice-versa. This controlled process prevents VNs from seeing other VNs' routes and ensures proper security and isolation. This is critical for maintaining the segmentation of traffic within the SD-Access fabric and ensuring services are reachable where needed. The fusion router also performs network address translation (NAT) for internet access if required. This role is not just about connectivity; it's about managed and controlled connectivity.

Further research can be conducted using these links:

**Cisco SD-Access Solution Design Guide:** https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/sda-design-guide-2022jul.html (Focus on sections detailing the Control Plane and Fusion Router)
**Cisco SD-Access Architecture White Paper:** https://www.cisco.com/c/en/us/solutions/collateral/enterprise/software-defined-access/white-paper-c11-740986.html (This document outlines the high-level architecture, including the role of segmentation).

**Cisco SD-Access Data Sheet:**https://www.cisco.com/c/en/us/products/collateral/software/dna-center/datasheet-c78-739374.html (This provides a general overview of SD-Access).

## Question: 32

Which action is the vSmart controller responsible for in an SD-WAN deployment?

A. onboard vEdge nodes into the SD-WAN fabric
B. gather telemetry data from vEdge routers
C. distribute security information for tunnel establishment between vEdge routers
D. manage, maintain, and gather configuration and status for nodes within the SD-WAN fabric

**Answer: C**

**Explanation:**

The vSmart controller in a Cisco SD-WAN deployment is primarily responsible for the control plane functions, including policy enforcement, route distribution, and security. Option C correctly identifies its role in distributing security information. Specifically, the vSmart controller pushes the necessary data for IPsec tunnel establishment, such as cryptographic keys and security parameters, to vEdge routers. This allows vEdges to securely form encrypted tunnels for data transport. While the vSmart doesn't handle data directly, it dictates how vEdges should communicate securely. Option A is incorrect as vEdge onboarding is typically managed by vManage. Option B is incorrect; telemetry is gathered by vManage, not vSmart. While vSmart manages configurations, option D is too broad; its main focus is establishing control plane intelligence and security, not general status gathering. Therefore, distributing security information for tunnel establishment (C) is the most accurate description of the vSmart's function.

**Authoritative Links:**

**Cisco SD-WAN Design Guide:**https://www.cisco.com/c/en/us/td/docs/solutions/CVD/SDWAN/cisco-sdwan-design-guide.html (Specifically, look for sections detailing vSmart controller functions)
**Cisco SD-WAN Documentation:**https://www.cisco.com/c/en/us/solutions/enterprise/sd-wan/index.html (Navigate to the relevant technical documentation sections about the control plane and vSmart role) **Cisco vSmart Controller Explained:** (Search for relevant articles or white papers explaining the vSmart controller and its security function.) (e.g., on the Cisco website itself, or resources like network blogs, TechTarget.)

These resources will provide a more in-depth understanding of Cisco SD-WAN and the specific roles of its components.

## Question: 33

What is one fact about Cisco SD-Access wireless network deployments?

A. The access point is part of the fabric overlay.
B. The wireless client is part of the fabric overlay.

C. The access point is part of the fabric underlay.

D. The WLC is part of the fabric underlay.

**Answer: A**

**Explanation:**

The correct answer is A: The access point is part of the fabric overlay in Cisco SD-Access wireless deployments. This is because Cisco SD-Access architecture separates the underlay (physical network infrastructure) from the overlay (virtualized network). In this model, the access point (AP), while physically connected to the underlay, logically operates within the virtualized fabric overlay. The fabric overlay provides services such as segmentation, policy enforcement, and user identity management. Wireless clients, conversely, are not part of the fabric overlay itself; their traffic is encapsulated and tunneled through the overlay from the access point. Options C and D are incorrect since both the access point and Wireless LAN Controller (WLC) are part of the underlay at the physical level. The WLC traditionally manages the access points, so it's part of the infrastructure itself; it's not part of the virtualized overlay, though it provides data for the overlay operations. Option B is incorrect as clients connect through the overlay not being part of it. Fabric overlay means that APs operate within the controlled virtualized overlay environment, allowing for consistent policy enforcement across wired and wireless networks. Cisco's SD-Access solution allows for consistent policies, regardless of how users and devices connect. The overlay uses technologies like VXLAN, allowing for the creation of virtual networks independent of the underlying physical infrastructure. The WLC communicates with Cisco DNA Center for the orchestration and enforcement of policies on the fabric overlay.

This architecture simplifies network management by providing a centralized view and control of the entire network. In summary, in SD-Access, the access point is an endpoint of the fabric overlay rather than a physical part of the underlay, unlike the WLC, which is infrastructure component within the underlay.

**Authoritative Links:**

**Cisco SD-Access Solution Overview:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/software-defined-access/index.html
**Cisco DNA Center SD-Access White Paper:**
https://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise-networks/software-defined-access/white-paper-c11-741639.pdf
**Cisco SD-Access Design Guide:**https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/sda-design-guide.html

## Question: 34

In a Cisco SD-Access solution, what is the role of a fabric edge node?

A. to connect external Layer 3 networks to the SD-Access fabric

B. to connect wired endpoints to the SD-Access fabric

C. to advertise fabric IP address space to external networks

D. to connect the fusion router to the SD-Access fabric

**Answer: B**

**Explanation:**

The correct answer is **B. to connect wired endpoints to the SD-Access fabric.**

In a Cisco SD-Access (Software-Defined Access) solution, a fabric edge node serves as the point of attachment for wired endpoints, such as computers, printers, and IP phones, into the SD-Access fabric. These edge nodes are typically Layer 2 switches that provide access ports for devices and enforce policies defined within the SD-Access control plane. They receive instructions from the DNA Center, the central management platform, which dictates how traffic is handled for connected endpoints. The edge node does not connect the fabric to external networks (option A), which is the role of border nodes. Nor does it advertise the fabric's IP space to external networks (option C); this is also the responsibility of border nodes. Similarly, while a fusion router might be part of the SD-Access infrastructure, edge nodes don't directly connect to it (option D). Edge nodes primarily function as the physical gateway for wired client devices joining the fabric and allow for policy enforcement at the access layer. They bridge the gap between the physical endpoint and the virtualized network of SD-Access.

**Authoritative Links for Further Research:**

**Cisco SD-Access Solution Overview:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/software-defined-access/index.html
**Cisco DNA Center Documentation:**https://www.cisco.com/c/en/us/support/cloud-systems-management/dna-center/tsd-products-support-series-home.html
**SD-Access Fabric Components:** Search for "SD-Access Fabric Components" on the Cisco website for in-depth explanations.

## Question: 35

What are two reasons a company would choose a cloud deployment over an on-prem deployment? (Choose two.)

A. Cloud costs adjust up or down depending on the amount of resources consumed. On-prem costs for hardware, power, and space are on-going regardless of usage.

B. Cloud resources scale automatically to an increase in demand. On-prem requires additional capital expenditure.

C. In a cloud environment, the company is in full control of access to their data. On-prem risks access to data due to service provider outages.

D. In a cloud environment, the company controls technical issues. On-prem environments rely on the service provider to resolve technical issues.

E. Cloud deployments require long implementation times due to capital expenditure processes. On-prem deployments can be accomplished quickly using operational expenditure processes.

**Answer: AB**

**Explanation:**

The correct answer is **A and B**. Here's why:

**A. Cloud costs adjust up or down depending on the amount of resources consumed. On-prem costs for hardware, power, and space are on-going regardless of usage.** This highlights the core principle of cloud elasticity and pay-as-you-go models. Cloud providers offer dynamic resource allocation; you pay only for what you use, scaling down during low demand and up during peaks. On-premises infrastructure, conversely, incurs fixed costs for hardware, maintenance, and utilities, regardless of actual resource utilization. This makes cloud deployment financially advantageous for workloads with fluctuating needs, avoiding wasted resources and expenses.

**B. Cloud resources scale automatically to an increase in demand. On-prem requires additional capital expenditure.** This underscores the scalability benefits of the cloud. Cloud platforms are designed to automatically adjust resources based on demand through features like auto-scaling. This ensures applications can handle sudden traffic spikes without performance degradation. On-premises scaling involves manual procurement, installation, and configuration of additional hardware, which is time-consuming and requires upfront capital investment. The automatic scaling capabilities of cloud infrastructure provide agility and cost efficiency.

Option C is incorrect because in a cloud environment the company is responsible for protecting their data and their access, while in on-prem they maintain that control. Option D is incorrect because in a cloud environment the cloud provider is responsible for the infrastructure and most technical issues, and in on-prem the company has the responsibility. Option E is incorrect because cloud deployments can be very rapid, while on-prem takes time for procurement and installation.

For further research, consider exploring these links:

**NIST Definition of Cloud Computing:**https://csrc.nist.gov/publications/detail/sp/800-145/final (This defines the essential characteristics of cloud computing, including elasticity and on-demand self-service.) **AWS Cloud Economics:**https://aws.amazon.com/economics/ (A good example of cloud pricing and cost optimization.)
**Azure Cloud Adoption Framework: Business Justification:**https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/strategy/business-justification (Explores the business value of cloud adoption.)These resources provide deeper insights into cloud computing principles, benefits, and cost models, further clarifying why A and B are the preferred reasons for choosing a cloud deployment.

**Question: 36**

What is the difference between the MAC address table and TCAM?

A. TCAM is used to make L2 forwarding decisions. CAM is used to build routing tables.

B. Router prefix lookups happen in CAM. MAC address table lookups happen in TCAM.

C. The MAC address table supports partial matches. TCAM requires an exact match.

D. The MAC address table is contained in CAM. ACL and QoS information is stored in TCAM.

**Answer: D**

**Explanation:**

The correct answer is **D**, "The MAC address table is contained in CAM. ACL and QoS information is stored in TCAM." Let's break down why.

**CAM (Content Addressable Memory)** is specialized hardware memory optimized for rapid lookup operations. In networking, CAM is primarily used to store the MAC address table, also known as the forwarding table or switching table. This table maps MAC addresses to specific switch ports, enabling a switch to quickly determine the outgoing port for a given destination MAC address. When a frame arrives, the switch uses the destination MAC to lookup the relevant port in the CAM table.

**TCAM (Ternary Content Addressable Memory)**, on the other hand, provides a more versatile lookup mechanism. It supports "ternary" states: 0, 1, and "don't care" or "X". This ability allows TCAM to perform lookups based on patterns, not just exact matches. Thus, TCAM can be leveraged for complex tasks like access control lists (ACLs), quality of service (QoS) policies, and route lookups. ACLs, for instance, are based on rules that match source IP addresses, destination IP addresses, ports, and other criteria, which need flexible pattern matching. Similarly, QoS rules require identifying different types of traffic and assigning them appropriate priority.

Option A is incorrect because CAM is primarily used for L2 forwarding based on MAC addresses. Routing tables, which are used for L3 decisions, are not built by CAM. Option B is incorrect because prefix lookups (used for routing) happen in a separate table structure within a router's hardware, not primarily within CAM and MAC address lookups are also not directly happening within TCAM. Option C is incorrect because neither MAC address tables nor TCAM, directly support "partial matches" within the context of forwarding and access control lists. Both require exact or predefined pattern matches. MAC address entries within the CAM table are exact matches to the MAC addresses.

In summary, CAM's primary role is to store the MAC address table for rapid L2 forwarding, while TCAM handles more complex packet filtering and traffic management tasks by storing ACLs and QoS rules, hence making D the correct answer.

For further research, consider the following:

**Cisco Documentation on CAM and TCAM:** https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SX/configuration/guide/swcg/camtcam.html
**Wikipedia on Content-Addressable Memory:**https://en.wikipedia.org/wiki/Content-addressable_memory
**TechTarget definition of TCAM:**https://www.techtarget.com/searchnetworking/definition/ternary-content-addressable-memory-TCAM

**Question: 37**

Which controller is the single plane of management for Cisco SD-WAN?

A. vBond
B. vSmart
C. vManage
D. vEdge

**Answer: C**

**Explanation:**

The correct answer is **C. vManage**. Cisco SD-WAN architecture employs several key components, each with a specific function. vManage acts as the central management system, providing a single pane of glass for configuring, monitoring, and troubleshooting the entire SD-WAN fabric. Unlike vBond (responsible for orchestration and initial authentication) and vSmart (the brain of the network, enforcing policies), vManage is the graphical interface and management plane. It allows administrators to define policies, deploy configurations, and view the real-time status of devices and circuits. vEdge routers are the physical or virtual devices deployed at the branch and data center locations, forwarding traffic according to policies defined by vSmart. vManage centralizes these functions, eliminating the need to individually manage each component.

Think of it as the control tower for the SD-WAN system, giving visibility and control over the entire deployment, making it the single point for management. Its graphical interface simplifies complex tasks associated with network management and monitoring. This centralized approach is a key tenet of software-defined networking (SDN), promoting efficiency and agility.

**Authoritative Links:**

**Cisco SD-WAN Overview:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/sd-wan/index.html **Cisco SD-WAN Components:**https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/sdwan-xe-gs-book/m-sdwan-components.html
**vManage Explained:**https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/sdwan-xe-gs-book/m-sdwan-vmanage.html

**Question: 38**

A company plans to implement intent-based networking in its campus infrastructure.
Which design facilitates a migration from a traditional campus design to a programmable fabric design?

A. two-tier

B. Layer 2 access
C. three-tier
D. routed access

**Answer: D**

**Explanation:**

The correct answer is **D. routed access.**

Here's why: Intent-based networking (IBN) aims to automate network operations based on high-level business intent. A crucial element is having a programmable and flexible underlying infrastructure. Routed access aligns perfectly with this goal.

Traditional three-tier campus designs (core, distribution, access) often rely on Layer 2 spanning tree protocols, creating complexity, limiting scalability, and hindering automated changes. In contrast, a routed access design moves Layer 3 closer to the edge, placing routing functionality at the access layer. This removes the constraints of Layer 2, enabling a more granular and adaptable network topology. Each access switch becomes a routing hop, enabling more streamlined traffic flow and facilitating a fabric-like model with clearly defined and isolated routing domains. This design approach also simplifies policy implementation, network segmentation, and allows for more scalable growth through easily added access nodes. With routed access, the network becomes inherently more modular and lends itself better to automation, which is vital for IBN. By using routing protocols like OSPF or BGP at the access, you're establishing a more dynamic and flexible network. Two-tier and Layer 2 access models don't provide the necessary agility required for IBN and the transition to a programmable fabric.

**Further Research:**

**Cisco SD-Access Solution Overview:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/software-defined-access/index.html (Cisco's approach to IBN and how it leverages routed access) **Software-Defined Access (SD-Access):**https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/sda-design-guide-2019jul.html (Detailed design guide highlighting benefits of routed access in an IBN environment)
**Routed Access vs. Layer 2 Access:**https://www.networkcomputing.com/networking/routed-access-vs-layer-2-access-which-right (Comparison of the two models)

## Question: 39

Which statement about a fabric access point is true?

A. It is in local mode and must be connected directly to the fabric edge switch.
B. It is in local mode and must be connected directly to the fabric border node.
C. It is in FlexConnect mode and must be connected directly to the fabric border node.
D. It is in FlexConnect mode and must be connected directly to the fabric edge switch.

**Answer: A**

**Explanation:**

The correct answer is A: "It is in local mode and must be connected directly to the fabric edge switch."

Let's break down why:

In a Cisco Software-Defined Access (SD-Access) fabric, access points (APs) typically operate in local mode, also sometimes called centralized mode. This means that all wireless traffic is tunneled back to the Wireless

LAN Controller (WLC) for centralized processing and forwarding. The WLC is usually deployed outside of the fabric.

Fabric edge switches are the designated access points to the fabric for end-user devices, including wireless clients. Therefore, APs connecting to the fabric must do so through a fabric edge node to ensure proper integration with the SD-Access fabric architecture. The edge switch acts as the first hop within the fabric for client traffic coming from the AP. It's important to understand that this differs from older wireless deployment models where APs connect to general access layer switches and do not have the same fabric-aware characteristics.

Option B is incorrect because fabric border nodes connect the SD-Access fabric to external networks and are not the intended connection point for access layer devices like APs.

Options C and D are incorrect as FlexConnect mode, while being another supported mode for APs, is primarily utilized in branch locations with less stable connectivity to the WLC. It is not the standard mode of operation within an SD-Access fabric where the WLC resides outside the fabric itself. FlexConnect allows for some local data switching if the WLC connection fails, but this is not the default or recommended setup for an SD-Access fabric AP deployment.

To summarize, the design philosophy of SD-Access prioritizes central control and policy enforcement. Therefore, APs within the SD-Access fabric are in local mode and must connect directly to a fabric edge switch to enable these features correctly.

Relevant links:

**Cisco SD-Access Design Guide:**https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/sda-design-guide.html (Focus on the access layer and wireless integration sections)
**Cisco Wireless LAN Controller Configuration Guide:**
https://www.cisco.com/c/en/us/support/wireless/wireless-lan-controller-software/products-installation-and-configuration-guides-list.html (Specifically look for information on AP modes like local and FlexConnect) **Cisco SD-Access Solution Overview:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/software-defined-access/index.html (General information on SD-Access architecture)

### Question: 40

A customer requests a network design that supports these requirements:✑
FHRP redundancy
✑ multivendor router environment
✑ IPv4 and IPv6 hosts
Which protocol does the design include?

A. VRRP version 2
B. VRRP version 3
C. GLBP
D. HSRP version 2

**Answer: B**

**Explanation:**

The correct answer is **B. VRRP version 3**. Here's why:

The question specifies a need for First Hop Redundancy Protocol (FHRP) support in a multi-vendor environment with both IPv4 and IPv6 hosts. VRRP (Virtual Router Redundancy Protocol) is designed for this purpose, allowing multiple routers to act as a single virtual router, ensuring continuous network connectivity in case of router failure. While VRRP version 2 only supports IPv4, VRRP version 3 is backward compatible with VRRPv2 and importantly supports both IPv4 and IPv6 addresses. HSRP (Hot Standby Router Protocol), while a popular FHRP, is Cisco proprietary and is not suitable in a multivendor scenario. GLBP (Gateway Load Balancing Protocol) also does not satisfy the vendor independence requirement as it too is Cisco-proprietary. Therefore, VRRPv3 stands out as the sole option that fulfills all requirements: FHRP redundancy, multivendor compatibility, and IPv4 & IPv6 support. The choice highlights the importance of considering protocol standards when designing heterogeneous networks. VRRP's standardized approach ensures interoperability across devices from different manufacturers.

Further research:

**RFC 5798:** Virtual Router Redundancy Protocol (VRRP) Version 3: https://www.rfc-editor.org/rfc/rfc5798 **Cisco Documentation on VRRP:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_fhr/configuration/15-e/fhr-15-e-book/fhr-vrrp.html

**Question: 41**

While configuring an IOS router for HSRP with a virtual IP of 10.1.1.1, an engineer sees this log message. Jan 1 12:12:12.111 : %HSRP-4-DIFFVIP1: GigabitEthernet0/0 Grp 1 active routers virtual IP address 10.1.1.1 is different to the locally configured address 10.1.1.25
Which configuration change must the engineer make?

A. Change the HSRP group configuration on the local router to 1.
B. Change the HSRP virtual address on the local router to 10.1.1.1.
C. Change the HSRP virtual address on the remote router to 10.1.1.1.
D. Change the HSRP group configuration on the remote router to 1.

**Answer: B**

**Explanation:**

The log message "%HSRP-4-DIFFVIP1: GigabitEthernet0/0 Grp 1 active routers virtual IP address 10.1.1.1 is different to the locally configured address 10.1.1.25" indicates a Hot Standby Router Protocol (HSRP) configuration mismatch. HSRP uses a virtual IP address (in this case, intended to be 10.1.1.1) shared between redundant routers. The active router possesses this virtual IP, providing a single gateway for client devices.

The message states the active router claims 10.1.1.1 as its virtual IP, but the local router is configured for 10.1.1.25. This discrepancy prevents HSRP from functioning correctly. Option B is correct because the local router must be configured to use the same virtual IP as the active router, i.e., 10.1.1.1. Changing the local router's virtual IP address to 10.1.1.1 resolves this mismatch, allowing both routers to agree on the virtual IP. Options A and D are incorrect because the group number does not affect the IP address conflict and option C is incorrect as we are seeing the issue on our local router therefore this is the router that needs to be amended. Only by ensuring consistency of virtual IP address can the HSRP group operate correctly and provide failover.
https://www.cisco.com/c/en/us/support/docs/ip/hot-standby-router-protocol-hsrp/10584-hsrp-virtualip.htmlhttps://www.geeksforgeeks.org/hot-standby-router-protocol-hsrp/

**Question: 42**

A network administrator has designed a network with two multilayer switches on the distribution layer, which act as default gateways for the end hosts. Which two technologies allow every end host in a VLAN to use both gateways? (Choose two.)

A. VRRP

B. GLBP
C. VSS
D. MHSRP
E. HSRP

**Answer: BC**

**Explanation:**

The correct answer is **B. GLBP** and **C. VSS**.

Here's why:

**B. Gateway Load Balancing Protocol (GLBP):** GLBP is a Cisco proprietary protocol that allows multiple routers or multilayer switches to share the forwarding load for a single virtual IP address. Unlike other redundancy protocols, GLBP actively uses all available gateways. Each gateway participates in forwarding traffic, ensuring optimal use of resources. GLBP achieves this by designating an Active Virtual Gateway (AVG) responsible for assigning virtual MAC addresses to the other members, called Active Virtual Forwarders (AVFs). End devices use the shared virtual IP address as their default gateway, and GLBP handles the distribution of traffic across the AVFs using those virtual MAC addresses.

**C. Virtual Switching System (VSS):** VSS is a technology that combines two Cisco Catalyst switches into a single logical unit. This allows both switches to act as a single default gateway, presenting a single logical path to the network. VSS creates a single control plane and data plane, providing both redundancy and load-sharing. End hosts simply see a single virtual switch and send all traffic to the VSS domain. VSS actively forwards traffic through both physical switches, leveraging their aggregate capacity and providing redundancy. VSS eliminates the complexity of configuring routing protocols for redundancy, making it simpler to manage and maintain.

**Why other options are incorrect:**

**A. VRRP (Virtual Router Redundancy Protocol):** VRRP is a redundancy protocol that allows multiple routers or switches to share a virtual IP address. However, only one device actively forwards traffic at a time while the others are in standby mode. VRRP provides redundancy but doesn't share the load across multiple gateways.
**D. MHSRP (Multiple HSRP):** MHSRP is not a standard term in Cisco networking. The commonly known protocol is HSRP, but it is not a correct answer for load balancing multiple gateways for hosts.
**E. HSRP (Hot Standby Router Protocol):** HSRP is a redundancy protocol similar to VRRP. It involves one active gateway forwarding traffic, while another standby gateway waits for the active gateway to fail. HSRP doesn't provide load sharing and therefore isn't an answer to the question.

**In Summary:**GLBP achieves load balancing at the default gateway level by actively utilizing multiple switches, and VSS combines multiple physical switches into a single logical entity that also actively uses both for forwarding. Thus, they offer both redundancy and load sharing, enabling all end hosts within a VLAN to utilize the resources of multiple gateway devices, distributing the load effectively.

**Authoritative Links for further research:**

**GLBP:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/15-mt/ipapp-15-mt-book/ipapp-glbp.html
**VSS:**https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-6500-series-switches/white_paper_c11-525418.html

**Question: 43**

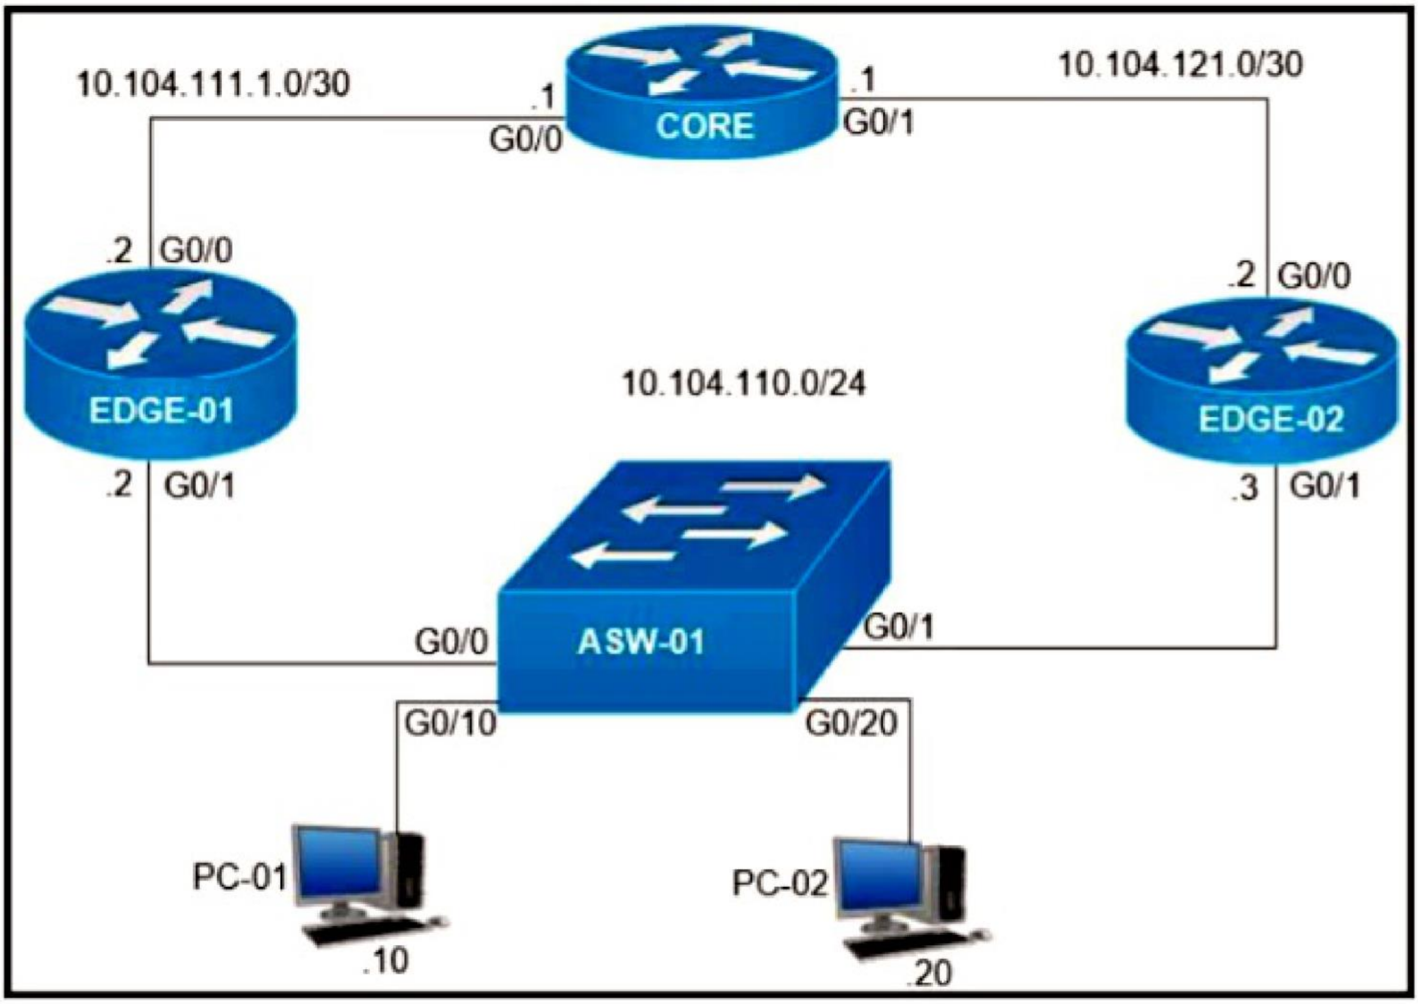

Refer to the exhibit. On which interfaces should VRRP commands be applied to provide first hop redundancy to PC-01 and PC-02?

A. G0/0 and G0/1 on Core
B. G0/0 on Edge-01 and G0/0 on Edge-02
C. G0/1 on Edge-01 and G0/1 on Edge-02
D. G0/0 and G0/1 on ASW-01

**Answer: C**

**Explanation:**

G0/1 on Edge-01 and G0/1 on Edge-02

**Question: 44**

Under which network conditions is an outbound QoS policy that is applied on a router WAN interface most beneficial?

A. under traffic classification and marking conditions
B. under interface saturation conditions
C. under all network conditions
D. under network convergence conditions

**Answer: B**

**Explanation:**

The correct answer is **B. under interface saturation conditions**.

Outbound QoS (Quality of Service) policies on a router's WAN interface are most effective when the link is experiencing congestion or saturation. QoS mechanisms, such as traffic shaping and prioritization, become crucial when the available bandwidth is insufficient to handle all traffic demands simultaneously. When an interface is saturated, packets may be dropped or experience excessive delays, negatively impacting application performance. An outbound QoS policy helps to mitigate these issues by intelligently managing and prioritizing outbound traffic.

For example, if a router is connected to a WAN link with limited bandwidth, and multiple applications are generating traffic concurrently (e.g., voice, video, and data), the link can easily become saturated. Without QoS, all traffic might compete equally for bandwidth, leading to poor performance for latency-sensitive applications like voice. By implementing an outbound QoS policy, the router can prioritize voice packets over less critical data, ensuring that voice calls remain clear and free of interruptions.

In contrast, A (traffic classification and marking) is a prerequisite for QoS, not the condition under which it's most beneficial. C (all network conditions) is incorrect because QoS isn't always needed. If the interface is not saturated, implementing QoS adds processing overhead without providing tangible benefits. Similarly, D (network convergence conditions) is incorrect as convergence is a network routing and switching issue, not where outbound QoS is beneficial. Hence, the most significant impact of outbound QoS is seen when bandwidth is a constraint.

**Authoritative Links:**

Cisco QoS: https://www.cisco.com/c/en/us/solutions/enterprise-networks/quality-of-service-qos/index.html Understanding QoS: https://www.networkcomputing.com/networking/understanding-qos-quality-service ITU-T Recommendation Y.1541: https://www.itu.int/rec/T-REC-Y.1541/en

## Question: 45

An engineer must configure HSRP group 300 on a Cisco IOS router. When the router is functional, it must be the active HSRP router. The peer router has been configured using the default priority value. Which command set is required?

A. standby version 2 standby 300 priority 110 standby 300 preempt

B. standby 300 priority 110 standby 300 timers 1 110

C. standby version 2 standby 300 priority 90 standby 300 preempt

D. standby 300 priority 90 standby 300 preempt

**Answer: A**

**Explanation:**

The correct answer is A because it accurately configures HSRP (Hot Standby Router Protocol) with the required specifications. HSRP employs a priority system to determine the active router. The default priority for HSRP is 100. To ensure the router becomes the active router when functional, its priority must be higher than the default priority of the peer router. Option A sets the priority to 110, making it a higher value. The standby preempt command is crucial; it ensures the router actively takes over the active role when it comes online, if it is the highest priority router. Without preempt enabled, the router will not assume active status even if it has higher priority when it comes online after an active router failure. standby version 2 is

recommended for HSRP to support more multicast addresses, although it doesn't affect the priority behavior. Option B is incorrect because the timer command does not affect HSRP priority. Option C is incorrect because a priority of 90 is lower than the default and will not make this the active router. Option D is incorrect because 90 is less than the default and it will not be elected as the active router.

Here's a more detailed breakdown:

**standby version 2**: This configures HSRP version 2, recommended over version 1 because it uses a multicast address of 224.0.0.102 instead of 224.0.0.2, allowing for better scaling and support of more groups. **standby 300 priority 110**: This sets the priority for HSRP group 300 to 110. Since the peer router has the default priority of 100, this router will become the active one in normal operation.
**standby 300 preempt**: This enables the preempt feature, which is essential for the requirement. Preempt allows the router with the highest priority to immediately become the active router when it comes online.

**Authoritative Links for Further Research:**

**Cisco HSRP Configuration Guide:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/15-sy/iap-15-sy-book/config-hsrp.html
**HSRP on Cisco Learning Network:**https://learningnetwork.cisco.com/s/article/hsrp-fundamentals-configuration-and-troubleshooting

## Question: 46

What is the function of a fabric border node in a Cisco SD-Access environment?

A. To collect traffic flow information toward external networks.
B. To connect the Cisco SD-Access fabric to another fabric or external Layer 3 networks.
C. To attach and register clients to the fabric.
D. To handle an ordered list of IP addresses and locations for endpoints in the fabric.

**Answer: B**

**Explanation:**

The correct answer is B, "To connect the Cisco SD-Access fabric to another fabric or external Layer 3 networks." A fabric border node acts as the gateway between the SD-Access fabric, which is a logically defined network within an organization, and the outside world, including other fabrics or traditional layer 3 networks. It's crucial for communication between the internal fabric and external destinations. These nodes provide routing and forwarding functions, translating addresses and policies between the SD-Access environment and the external world. They encapsulate and decapsulate traffic using protocols like VXLAN to maintain the separation of the fabric. Without border nodes, the SD-Access fabric would be an isolated environment. Option A is incorrect because while border nodes may collect flow information, that is not its primary function; that is more the job of analytics engines. Option C describes the function of edge nodes, not border nodes. Option D describes the function of the control plane, specifically the DNA center. Therefore, B accurately describes the core purpose of a fabric border node.

Further research:

**Cisco SD-Access Solution Design Guide:**
https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html (Look for sections on fabric architecture and border nodes)
**Cisco SD-Access Documentation:**https://www.cisco.com/c/en/us/solutions/enterprise/software-defined-access/index.html
**VXLAN Technology:**https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/white-paper-c11-729818.html (Understanding VXLAN is key to understanding fabric encapsulation)

## Question: 47

In a wireless Cisco SD-Access deployment, which roaming method is used when a user moves from one AP to another on a different access switch using a single
WLC?

A. Layer 3
B. inter-xTR
C. auto anchor
D. fast roam

**Answer: B**

**Explanation:**

The correct answer is **B. inter-xTR**. Here's why:

In a Cisco SD-Access wireless deployment utilizing a single Wireless LAN Controller (WLC), roaming between APs connected to different access switches involves inter-xTR (inter-eXtended Trusted Relay) mobility. This occurs because SD-Access relies on a fabric overlay, and the access switches act as fabric edge nodes. When a wireless client moves between APs on different edge nodes (access switches), the fabric needs to maintain the client's session and policy.

Inter-xTR ensures seamless mobility by using a combination of techniques. The original access switch, designated as the "anchor" xTR, handles the initial wireless association. As the client roams to an AP on a new access switch, the new switch becomes the "new" xTR. The WLC does not act as a centralized mobility anchor; instead, it coordinates the mobility handoff. The fabric control plane manages the mobility update and forwarding changes through LISP (Locator/ID Separation Protocol) updates. This ensures that the client’s traffic is correctly routed and policy is maintained without re-authentication. Layer 3 roaming (option A) typically involves a change in IP subnet and requires a different handling mechanism. Auto-anchor (option C) is more related to guest access scenarios or scenarios where the anchor controller is explicitly defined. Fast roam (option D) refers to faster 802.11 handoffs, which speeds up mobility but does not address the underlying control and data plane considerations of SD-Access. The inter-xTR method is specifically designed for fabric mobility in SD-Access architectures, enabling efficient routing and policy enforcement as users move across different access points and switches within the fabric.

**Authoritative Links:**

**Cisco SD-Access Wireless Design Guide:**
https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-wireless-design-guide.html (Specifically, look at the mobility sections and the discussion of fabric edge nodes.)
**Cisco DNA Center User Guide:**https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/2-3-3/user_guide/b_cisco_dna_center_ug_2_3_3/b_cisco_dna_center_ug_2_3_3_chapter_0100.html (Look at the sections on SD-Access Fabric and Wireless integration.)

## Question: 48

What is the recommended MTU size for a Cisco SD-Access Fabric?

A. 4464
B. 17914
C. 9100
D. 1500

**Answer: C**

**Explanation:**

The recommended Maximum Transmission Unit (MTU) size for a Cisco Software-Defined Access (SD-Access) Fabric is 9100 bytes, also known as a jumbo frame. This recommendation stems from the need to accommodate the additional overhead introduced by the VXLAN (Virtual Extensible LAN) encapsulation used within the SD-Access fabric. Standard Ethernet frames typically have an MTU of 1500 bytes, which becomes insufficient when VXLAN adds its header. VXLAN headers can range from 50 to 70 bytes and additional header information needs to be added at certain points, requiring a larger MTU to avoid fragmentation.

Fragmentation can negatively impact network performance, leading to increased latency and reduced throughput. Using jumbo frames, with an MTU of 9100 bytes, allows for the original 1500-byte IP packet, plus VXLAN headers and other necessary headers, without fragmentation. This enhances efficiency, reduces processing overhead, and allows for higher bandwidth utilization in the underlay and overlay networks of the SD-Access fabric. The increased payload size also minimizes the ratio of header to payload data, further optimizing data transmission. While an option like 17914 bytes exists, it's generally not supported across all transport links in common enterprise deployments, making 9100 a safer and more universally applicable recommendation. Choosing 9100 aligns with best practices for SD-Access implementation, ensuring efficient data transport within the virtualized network.

Cisco SD-Access Design GuideCisco Enterprise Network Fabric Design Guide

**Question: 49**

What is the function of the fabric control plane node in a Cisco SD-Access deployment?

A. It is responsible for policy application and network segmentation in the fabric.

B. It performs traffic encapsulation and security profiles enforcement in the fabric.

C. It holds a comprehensive database that tracks endpoints and networks in the fabric.

D. It provides integration with legacy nonfabric-enabled environments.

**Answer: C**

**Explanation:**

The correct answer is C. The fabric control plane node in a Cisco SD-Access deployment, often implemented using Cisco DNA Center and its associated services, serves as the central intelligence repository for the entire fabric. This node maintains a comprehensive database that tracks all endpoints (like user devices and servers) and network resources within the fabric. This database includes information about user identities, device locations, and network associations, crucial for implementing policy and ensuring consistent access control.

The control plane doesn't directly handle data plane operations like traffic encapsulation or forwarding; instead, it dictates how the fabric components should behave based on defined policies. Options A, B, and D describe functions performed by other elements within the SD-Access architecture. Option A describes the function of the policy plane within Cisco DNA Center. Option B relates more to the role of fabric edge nodes or border nodes. Option D highlights the role of border nodes and their ability to interface with legacy networks.

In summary, the fabric control plane node acts as the brain of the SD-Access network, providing the necessary information for policy enforcement, network segmentation, and overall fabric operation.

Cisco SD-Access ArchitectureCisco DNA Center

**Question: 50**

What is the data policy in a Cisco SD-WAN deployment?

A. list of ordered statements that define node configurations and authentication used within the SD-WAN overlay

B. set of statements that defines how data is forwarded based on IP packet information and specific VPNs

C. detailed database mapping several kinds of addresses with their corresponding location

D. group of services tested to guarantee devices and links liveliness within the SD-WAN overlay

**Answer: B**

**Explanation:**

The correct answer is **B. a set of statements that defines how data is forwarded based on IP packet information and specific VPNs.** This aligns with the core function of a data policy within Cisco SD-WAN. Data policies dictate how traffic is handled within the overlay network. They operate on packet attributes such as source/destination IP addresses, ports, and DSCP values, and are applied to specific VPNs (Virtual Routing and Forwarding instances) to govern traffic flow. Essentially, they're access control lists and policy routing rules for the SD-WAN overlay. Option A describes control policies, which handle node configurations and authentication, not data forwarding. Option C describes a location database, which relates IP addresses to physical locations but isn't a traffic-handling policy. Option D pertains to monitoring and health checks, which, while important, aren't data forwarding policies. Data policies ensure application traffic takes the appropriate path based on defined criteria, enabling granular control and optimization within the SD-WAN. They are a crucial component for implementing QoS, security, and application-aware routing within the SD-WAN.

**Authoritative Links:**

Cisco SD-WAN Data Policy Configuration:
https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/policies/ios-xe-17/policies-book-xe-17/data-policies.html
Cisco SD-WAN Policy Overview:
https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/policies/ios-xe-17/policies-book-xe-17/policy-overview.html

## Question: 51

In Cisco SD-WAN, which protocol is used to measure link quality?

A. IPsec

B. OMP

C. RSVP

D. BFD

**Answer: D**

**Explanation:**

The correct answer is BFD (Bidirectional Forwarding Detection). BFD is a lightweight, standardized protocol designed specifically for fast link failure detection. In Cisco SD-WAN, BFD is used to monitor the health and quality of the data plane connections between vEdge routers. It achieves this by exchanging short, frequent control packets (BFD control packets) between the routers. These packets are used to check the reachability and latency of the underlying transport links. If BFD packets are not received within a predefined timeout period, the link is considered down, and the SD-WAN fabric reacts accordingly, such as rerouting traffic over another available path. OMP (Overlay Management Protocol), while a crucial part of SD-WAN, is responsible for the control plane, sharing routing and policy information, not link quality monitoring directly. IPsec (Internet Protocol Security) provides secure tunnel establishment but does not perform link quality measurements in the SD-WAN context. RSVP (Resource Reservation Protocol) is traditionally used for QoS in MPLS networks, and is not directly used for link quality measurements within Cisco SD-WAN. Therefore, BFD's dedicated purpose for quick and reliable connection health assessment makes it the correct answer for link quality measurement in Cisco SD-WAN.

Authoritative links for further research:

Cisco SD-WAN documentation: https://www.cisco.com/c/en/us/solutions/enterprise-networks/sd-wan/index.html
BFD RFC: https://datatracker.ietf.org/doc/html/rfc5880
Cisco SD-WAN BFD configuration guides: Search within Cisco's documentation portal.

## Question: 52

What is used to perform QoS packet classification?

A. the Type field in the Layer 2 frame
B. the Options field in the Layer 3 header
C. the TOS field in the Layer 3 header
D. the Flags field in the Layer 3 header

**Answer: C**

**Explanation:**

The correct answer is C, the TOS (Type of Service) field in the Layer 3 header, which has been superseded by the DSCP (Differentiated Services Code Point) field within the same byte. QoS packet classification relies on examining specific fields within network headers to categorize traffic for differentiated treatment. The TOS/DSCP field within the IP header is specifically designed for this purpose, allowing network devices to identify packets requiring prioritized or specific handling. This field indicates the desired quality of service for a particular packet. By interpreting the value in the TOS/DSCP field, network devices like routers and switches can apply appropriate queuing, scheduling, and other QoS mechanisms to ensure critical traffic receives preferential treatment over less important data. Option A, the Type field in the Layer 2 frame, primarily identifies the Layer 3 protocol being carried (e.g., IP, ARP) and is not used for QoS classification. Options B and D, the Options and Flags fields in the Layer 3 header, serve other functions unrelated to traffic prioritization. Therefore, the TOS/DSCP field is the primary mechanism within the IP header used for QoS packet classification, making C the correct choice. The use of TOS/DSCP facilitates the implementation of differentiated services across a network infrastructure.

Authoritative links:

Cisco - Quality of Service (QoS) Overview
RFC 2474 - Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers

## Question: 53

How do cloud deployments differ from on-premises deployments?

A. Cloud deployments require longer implementation times than on-premises deployments.
B. Cloud deployments are more customizable than on-premises deployments.
C. Cloud deployments have lower upfront costs than on-premises deployments.
D. Cloud deployments require less frequent upgrades than on-premises deployments.

**Answer: C**

**Explanation:**

Cloud deployments typically involve lower upfront costs (Option C) compared to on-premises deployments.

This is primarily due to the pay-as-you-go model common in cloud services. Instead of purchasing and maintaining physical infrastructure, organizations rent computing resources from cloud providers. This eliminates significant capital expenditures associated with hardware, software licenses, and physical facilities needed for on-premises solutions. On-premises deployments, conversely, involve substantial initial investments in infrastructure, requiring considerable upfront capital outlay. Cloud solutions also benefit from economies of scale, allowing providers to offer resources at competitive prices. Consequently, cloud deployments are more budget-friendly when starting, which makes option C the correct one. Cloud deployments do not necessarily have longer implementation times (Option A); in fact, they often have shorter implementation cycles. Customization options (Option B) can vary; while cloud providers offer customization options, on-premises solutions often allow for greater granular control. Cloud deployments may involve frequent upgrades (Option D) depending on the provider's update cycle, not necessarily less frequent ones.

Thus, option C is the most accurate answer as it directly relates to the economic benefits of initial costs.

Relevant Links for further research:

**NIST Definition of Cloud Computing:**https://csrc.nist.gov/publications/detail/sp/800-145/final
**AWS Cost Management:**https://aws.amazon.com/cost-management/
**Azure Cost Management:**https://azure.microsoft.com/en-us/pricing/cost-management/ **GCP Cost Management:**https://cloud.google.com/cost-management

**Question: 54**

Which controller is capable of acting as a STUN server during the onboarding process of Edge devices?

A. vBond
B. vSmart
C. vManage
D. PNP Server

**Answer: A**

**Explanation:**

The correct answer is A, vBond. Here's why: vBond orchestrators, acting as the initial point of contact, facilitate the secure onboarding of Cisco SD-WAN edge devices (such as routers). During this process, edge devices often reside behind NAT firewalls, making direct communication with other controllers challenging. To overcome this, vBond uses the Session Traversal Utilities for NAT (STUN) protocol. The vBond controller acts as a STUN server, allowing edge devices to discover their public IP address and port mapping. This information is then used to establish secure tunnels with other SD-WAN components. vSmart controllers, responsible for policy and control, and vManage, focusing on management and monitoring, don't fulfill the STUN server role. The Plug and Play (PNP) server, while crucial for device discovery, isn't directly involved in the STUN process during SD-WAN onboarding. Therefore, vBond's specific design and function as an orchestrator uniquely equip it to serve as the STUN server. By providing STUN functionality, vBond ensures seamless and secure initial connectivity of edge devices behind NAT, which is essential for a functional SD-WAN fabric. This is a critical initial step in the SD-WAN architecture and a key differentiator for vBond's role.

Authoritative Links:

Cisco SD-WAN Documentation: https://www.cisco.com/c/en/us/solutions/enterprise/sd-wan/index.html

Cisco SD-WAN Architecture Whitepaper: (Search on Cisco site for "Cisco SD-WAN Architecture Whitepaper" for more in depth detail).

RFC 5389 - Session Traversal Utilities for NAT (STUN): https://datatracker.ietf.org/doc/html/rfc5389 (General STUN RFC)

## Question: 55

How is 802.11 traffic handled in a fabric-enabled SSID?

A. centrally switched back to WLC where the user traffic is mapped to a VXLAN on the WLC
B. converted by the AP into 802.3 and encapsulated into VXLAN
C. centrally switched back to WLC where the user traffic is mapped to a VLAN on the WLC
D. converted by the AP into 802.3 and encapsulated into a VLAN

**Answer: B**

**Explanation:**

The correct answer is B. When a fabric-enabled SSID is utilized in a Cisco environment, the access point (AP) plays a critical role in encapsulating wireless traffic. Upon receiving 802.11 traffic, the AP converts it to 802.3 Ethernet frames, which is the standard format used in wired networks. Crucially, in a fabric environment, these 802.3 frames are then encapsulated within a Virtual Extensible LAN (VXLAN) header. This VXLAN encapsulation is what allows the wireless traffic to participate in the fabric, providing benefits like micro-segmentation and enhanced scalability. The VXLAN tunnel effectively transports the traffic across the underlay network to the desired destination. Option A is incorrect because in a fabric, the WLC doesn't handle the VXLAN encapsulation for this type of traffic; it’s the AP's job. Options C and D are incorrect as fabric environments rely on VXLAN encapsulation instead of VLANs for this specific operation. Specifically, VLANs are not extended across the fabric; VXLAN tunnels enable the extension of Layer 2 segments over Layer 3 network and provide flexibility and scalability. Therefore, option B accurately describes how 802.11 traffic is handled within a fabric-enabled SSID environment where the AP performs the necessary conversion and encapsulation.

Further research can be found at:

**Cisco DNA Center Wireless Fabric:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/dna-center/wireless-fabric.html
**Understanding VXLAN:**https://www.cisco.com/c/en/us/solutions/data-center-virtualization/vxlan/index.html **Cisco Catalyst 9800 Wireless Controllers:**https://www.cisco.com/c/en/us/products/wireless/catalyst-9800-series-wireless-controllers/index.html

## Question: 56

Refer to the exhibit.

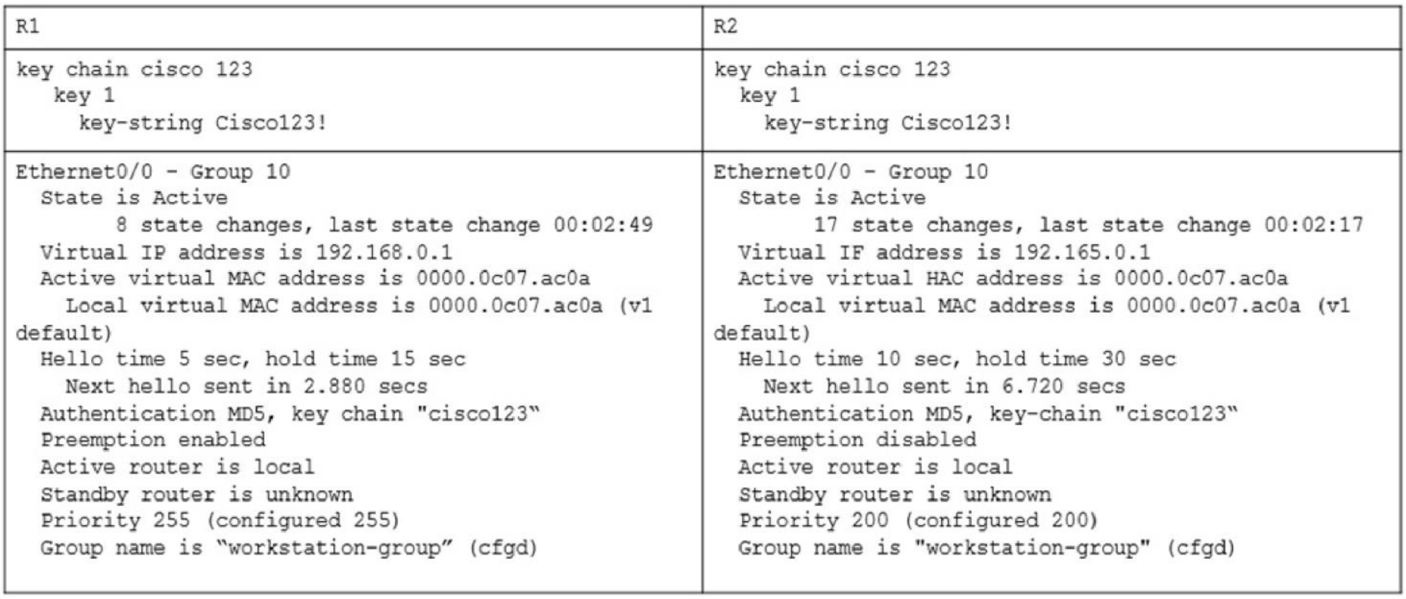

R1

```
key chain cisco 123
   key 1
     key-string Cisco123!
```

```
Ethernet0/0 - Group 10
  State is Active
        8 state changes, last state change 00:02:49
  Virtual IP address is 192.168.0.1
  Active virtual MAC address is 0000.0c07.ac0a
    Local virtual MAC address is 0000.0c07.ac0a (v1
default)
  Hello time 5 sec, hold time 15 sec
    Next hello sent in 2.880 secs
  Authentication MD5, key chain "cisco123"
  Preemption enabled
  Active router is local
  Standby router is unknown
  Priority 255 (configured 255)
  Group name is "workstation-group" (cfgd)
```

R2

```
key chain cisco 123
  key 1
    key-string Cisco123!
```

```
Ethernet0/0 - Group 10
  State is Active
        17 state changes, last state change 00:02:17
  Virtual IF address is 192.165.0.1
  Active virtual HAC address is 0000.0c07.ac0a
    Local virtual MAC address is 0000.0c07.ac0a (v1
default)
  Hello time 10 sec, hold time 30 sec
    Next hello sent in 6.720 secs
  Authentication MD5, key-chain "cisco123"
  Preemption disabled
  Active router is local
  Standby router is unknown
  Priority 200 (configured 200)
  Group name is "workstation-group" (cfgd)
```

An engineer is installing a new pair of routers in a redundant configuration. When checking on the standby status of each router, the engineer notices that the routers are not functioning as expected.
Which action will resolve the configuration error?

A. configure matching hold and delay timers
B. configure matching key-strings
C. configure matching priority values
D. configure unique virtual IP addresses

**Answer: D**

**Explanation:**

The most situable for this question is D (configure unique virtual address) as is one of requirements for a group to work, as seeing in the image. The others are wrong for the following reason. (so Wrong based in the question)A) Not mandatory the timers must match to work on HSRP. on HSRP negotiation the active router will override the standby timers. https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/nb-06-cat-9k-stack-wp-cte-en.pdfB) On the image key string match so is not this reason HSRP cant be seen each otherC) Wrong, priority is for selecing the active router, the best practice is should be different thus (WRONG)D) Different virtual Gateway configured on the same group number indeed will make HSRP routers negotiation fail for that group thus (CORRECT)

VIP does not match. VIP should match while the physicals are unique.

## Question: 57

Refer to the exhibit.

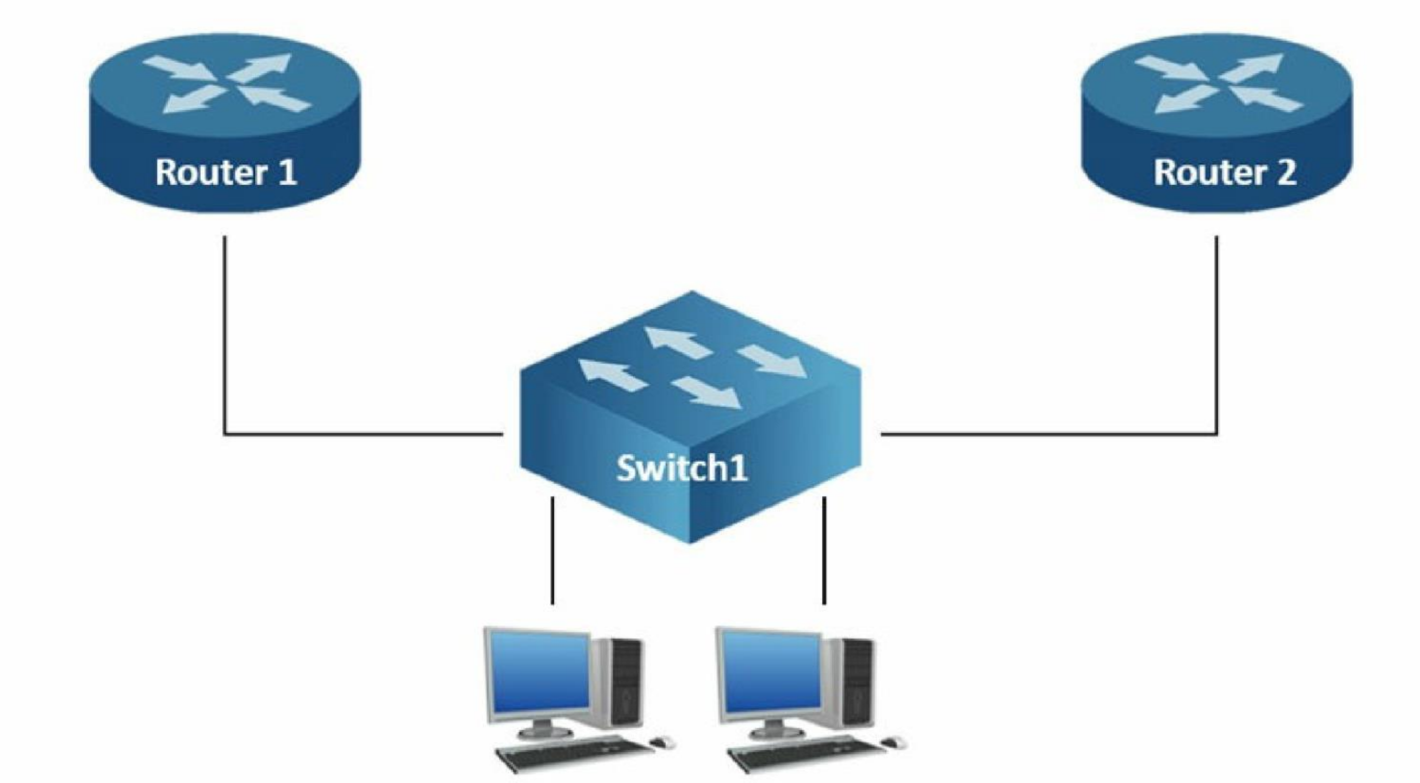

Router1 is currently operating as the HSRP primary with a priority of 110. Router1 fails and Router2 takes over the forwarding role. Which command on Router1 causes it to take over the forwarding role when it returns to service?

A. standby 2 priority

B. standby 2 preempt

C. standby 2 track

D. standby 2 timers

**Answer: B**

**Explanation:**

Reference:
https://www.cisco.com/c/en/us/support/docs/ip/hot-standby-router-protocol-hsrp/13780-6.html

## Question: 58

An engineer has deployed a single Cisco 5520 WLC with a management IP address of 172 16.50.5/24. The engineer must register 50 new Cisco AIR-CAP2802I-
E-K9 access points to the WLC using DHCP option 43. The access points are connected to a switch in VLAN 100 that uses the 172.16.100.0/24 subnet. The engineer has configured the DHCP scope on the switch as follows: Network 172.16.100.0 255.255.255.0

Default Router 172.16.100.1 -

Option 43 ASCII 172.16.50.5 -
The access points are failing to join the wireless LAN controller. Which action resolves the issue?

A. configure option 43 Hex F104.AC10.3205

B. configure option 43 Hex F104.CA10.3205

C. configure dns-server 172.16.50.5

D. configure dns-server 172. 16.100.1

**Answer: A**

**Explanation:**

The issue stems from the incorrect configuration of DHCP option 43, which is used by Cisco access points to discover their Wireless LAN Controller (WLC). DHCP option 43 for Cisco APs requires a specific vendor-specific format, not the ASCII representation of the WLC's IP address as configured in the provided scenario. The correct format is hexadecimal and depends on the number of WLC management IPs to be shared with the APs. Specifically, it starts with a type code 'F1' which denotes a single management IP, followed by the length of the IPv4 address representation which is '04', and the WLC's management IPv4 address in hex format (converted from dotted decimal notation). The WLC address 172.16.50.5 translates to AC.10.32.05 in hex. Therefore the correct hexadecimal value for DHCP option 43 is 'F104.AC10.3205'. The other provided options are not valid: using an incorrect hex representation (option B) or trying to use a DNS server for AP discovery (options C and D) which are not applicable. DHCP option 43 bypasses DNS discovery process. The correct option, A, provides the accurate format which allows the APs to properly discover the WLC and register to it.

**Authoritative Links:**

1. **Cisco Documentation on DHCP Option 43:** https://www.cisco.com/c/en/us/support/docs/wireless/aironet-1100-series-ap/108460-dhcp-option-43-config.html (This document provides a good overview, though the hexadecimal calculation examples are for older models)
2. **Cisco Communities Discussion on Option 43:**https://community.cisco.com/t5/wireless/dhcp-option-43/td-p/2417102 (This forum discussion provides community guidance on configuring option 43 with examples for IPv4)

**Question: 59**

What is the role of vSmart in a Cisco SD-WAN environment?

A. to establish secure control plane connections

B. to monitor, configure, and maintain SD-WAN devices

C. to provide secure data plane connectivity over WAN links

D. to perform initial authentication of devices

**Answer: A**

**Explanation:**

The correct answer is **A. to establish secure control plane connections.** In a Cisco SD-WAN environment, the vSmart controller is the central component responsible for managing the network's control plane. It acts as the brain of the SD-WAN fabric. Crucially, the vSmart establishes secure connections with all other SD-WAN devices (vEdges and cEdges), using DTLS or TLS protocols. These secure connections are critical for the vSmart to distribute routing policies, security policies, and other network configurations. Without these secure control plane connections, the SD-WAN network would not function. The vSmart controller dictates how data traffic should be routed and managed, essentially orchestrating the entire network. It doesn't handle data plane traffic directly, which is the responsibility of the vEdge/cEdge routers. While the vManage tool (option B) handles monitoring, configuration, and maintenance, it interacts with vSmart to achieve these actions. VEdges/cEdges handle the actual data plane connectivity (option C), and initial authentication (option D) occurs through vBond. Therefore, the primary and fundamental role of vSmart is to establish and maintain secure control plane connections.

For further reading on Cisco SD-WAN architecture and the role of vSmart, refer to these resources:

**Cisco SD-WAN Design Guide:**https://www.cisco.com/c/en/us/solutions/enterprise/sd-wan/index.html (Navigate to relevant documentation on architecture)
**Cisco SD-WAN Documentation:**https://www.cisco.com/c/en/us/support/routers/sd-wan/tsd-products-support-series-home.html
**Cisco DevNet Learning Labs:** Search for "Cisco SD-WAN" on the Cisco DevNet website: https://developer.cisco.com/

## Question: 60

Which action is performed by Link Management Protocol in a Cisco StackWise Virtual domain?

A. It determines which switch becomes active or standby.
B. It determines if the hardware is compatible to form the StackWise Virtual domain.
C. It rejects any unidirectional link traffic forwarding.
D. It discovers the StackWise domain and brings up SVL interfaces.

**Answer: C**

**Explanation:**

The correct answer is C: It rejects any unidirectional link traffic forwarding. Link Management Protocol (LMP) in Cisco StackWise Virtual is crucial for maintaining the integrity of the virtual switch by ensuring bidirectional communication across all inter-switch links (ISLs). Specifically, LMP continuously monitors the health of these links. If LMP detects a unidirectional link, meaning traffic flows in only one direction, it immediately blocks traffic on that link. This action is vital to prevent forwarding loops and inconsistent network states, which can lead to severe network disruptions. Unlike options A and B, LMP doesn't handle election of active/standby switches or compatibility checks. Option D, discovering the StackWise domain and bringing up SVL interfaces, is primarily managed by StackWise Virtual technology itself and related configurations, not LMP. LMP's responsibility is solely to ensure reliable bidirectional communication once the stack is operational. In essence, LMP acts as a watchdog for link integrity, enforcing bidirectionality and halting traffic on
problematic links to guarantee a stable and predictable StackWise Virtual domain.

**Authoritative Links for Further Research:**

**Cisco StackWise Virtual Configuration Guide:**
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/16-12/configuration_guide/ha/b_1612_ha_9300_cg/m_stackwise_virtual.html (Search for "Link Management Protocol" within the document)
**Cisco StackWise Virtual Technology Overview:**
https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-6800-series-switches/whitepaper-c11-734098.html (Look for sections on link failure detection mechanisms)

## Question: 61

What are two reasons why broadcast radiation is caused in the virtual machine environment? (Choose two.)

A. vSwitch must interrupt the server CPU to process the broadcast packet.
B. The Layer 2 domain can be large in virtual machine environments.
C. Virtual machines communicate primarily through broadcast mode.
D. Communication between vSwitch and network switch is broadcast based.

E. Communication between vSwitch and network switch is multicast based.

**Answer: AB**

**Explanation:**

The correct answer is **A and B**. Let's break down why.

**A. vSwitch must interrupt the server CPU to process the broadcast packet.** Virtual switches (vSwitches) reside within the hypervisor and manage network traffic for virtual machines (VMs) on the same physical host. When a broadcast packet is received by a vSwitch, it must analyze the packet and decide which VM(s) should receive it. This involves interrupting the server's CPU, diverting its processing power from other tasks. Such interruptions are necessary for the vSwitch to efficiently distribute broadcast traffic across VMs, but this also contributes to broadcast radiation. Every broadcast requires these interruptions and the processing of the packet, regardless of whether a VM needs it.

**B. The Layer 2 domain can be large in virtual machine environments.** Unlike traditional physical networks, virtualized environments can have extensive Layer 2 networks spanning many hosts. If a VM transmits a broadcast, it floods that large Layer 2 domain. This increased size means that more VMs and hosts will receive the broadcast, resulting in greater overall broadcast traffic. The larger the virtual Layer 2 domain, the further and more widespread the broadcast will propagate.

**Why other options are incorrect:**

**C. Virtual machines communicate primarily through broadcast mode.** VMs often use unicast for targeted communication, not primarily broadcast. Broadcasts are used for specific purposes like ARP (Address Resolution Protocol) which is required for a VM to find the MAC address of another machine on the local network.

**D. Communication between vSwitch and network switch is broadcast based.** Communication between vSwitches and physical network switches is largely unicast, not broadcast. vSwitches act as virtual network devices and handle traffic according to Layer 2/3 rules.

**E. Communication between vSwitch and network switch is multicast based.** While some multicast is present, like for certain discovery protocols, the bulk of traffic is not multicast-based, it is unicast.

**In summary**, broadcast radiation occurs in VM environments because the processing of broadcast packets by vSwitches requires CPU interruption (A), and the expansive Layer 2 domains in virtualized environments amplify the spread of these broadcasts (B). These two factors contribute to the observed broadcast radiation and potential performance challenges within the virtual infrastructure.

**Authoritative Links for Further Research:**

**VMware vSphere Networking:**https://docs.vmware.com/en/VMware-vSphere/ (Search for vSwitch documentation)
**Microsoft Hyper-V Networking:**https://learn.microsoft.com/en-us/windows-server/virtualization/hyper-v/ (Search for virtual switch documentation)
**Cisco Networking Academy:**https://www.netacad.com/ (For general networking principles, including Layer 2 domains and broadcast behavior)

## Question: 62

Which two GRE features are configured to prevent fragmentation? (Choose two.)

A. TCP window size
B. IP MTU
C. TCP MSS

D. DF bit clear

E. MTU ignore

**Answer: BC**

**Explanation:**

The correct answer is B (IP MTU) and C (TCP MSS). GRE (Generic Routing Encapsulation) tunnels add overhead to the original packet, potentially exceeding the maximum transmission unit (MTU) of the underlying physical network. This leads to fragmentation, which can degrade performance.

IP MTU (Maximum Transmission Unit) plays a crucial role in avoiding fragmentation. Configuring a lower IP MTU on the tunnel interface compared to the physical interface ensures that the encapsulated packet, including the GRE header, fits within the network's MTU limit. This prevents IP-layer fragmentation from occurring before the packet enters the tunnel.

TCP MSS (Maximum Segment Size) is a TCP option that negotiates the largest segment a TCP connection can send without requiring fragmentation. By adjusting TCP MSS, endpoints communicate their maximum acceptable packet size before encapsulation. This effectively manages the size of TCP segments before they are encapsulated into GRE, preventing packets from becoming too large and exceeding the tunnel MTU.

The "DF bit clear" (D) is not a direct mechanism to prevent fragmentation. While it allows fragmentation to occur, it doesn't proactively prevent it. Instead, it instructs routers to fragment packets if necessary, which is usually undesirable. TCP window size (A) is about flow control and does not prevent fragmentation. "MTU ignore" (E) does not exist as a standard configuration option for preventing fragmentation.

In summary, configuring IP MTU and TCP MSS on the tunnel interface and associated endpoints, respectively, proactively addresses the root cause of fragmentation, which is a packet exceeding the network's MTU limit. They control packet size at the IP and TCP layer, thereby preventing fragmentation before it occurs.

**Authoritative Links:**

**Cisco on GRE Tunnel MTU and MSS:**https://www.cisco.com/c/en/us/support/docs/ip/generic-routing-encapsulation-gre/117800-technote-gre-00.html
**RFC 791 on IP Fragmentation:**https://www.rfc-editor.org/rfc/rfc791
**RFC 879 on TCP MSS Option:**https://www.rfc-editor.org/rfc/rfc879

## Question: 63

Which LISP device is responsible for publishing EID-to-RLOC mappings for a site?

A. ETR

B. MR

C. ITR

D. MS

**Answer: A**

**Explanation:**

The correct answer is A, ETR (Egress Tunnel Router). In a LISP (Locator/ID Separation Protocol) environment, the ETR is the device that advertises the Endpoint Identifier (EID) to Routing Locator (RLOC) mappings for a specific site. These mappings essentially inform the rest of the LISP-enabled network how to reach the EIDs located behind the ETR. When a host within a site communicates with a remote host, the ITR (Ingress Tunnel Router) initially encapsulates the packet using the RLOC of the ETR. The ETR then decapsulates the packet, making it accessible within the site's local network. The ETR is the authoritative source of EID-to-RLOC information for that particular site, registering and announcing these mappings to the Map-Server (MS). The Map-Server (MS) acts as a central directory, distributing these mappings to ITRs that request them. The ITR queries the Map-Server (MS) to determine how to reach a particular EID, using the returned RLOC for encapsulation. An MR (Map Resolver) is another component responsible for caching and distributing these mappings but does not initiate the mapping publication. The MS is where all the ETRs register their mappings, but the ETR is the device responsible for creating the mappings. In simpler terms, the ETR acts as the "gatekeeper" for the site's network by advertising how to find its devices. This is fundamental to LISP's function of separating endpoint identities from their network location. This promotes efficient routing and enables mobility of devices within LISP networks.

Relevant Links:

Cisco LISP Overview: https://www.cisco.com/c/en/us/solutions/enterprise-networks/lisp/index.html
LISP RFC 6830: https://datatracker.ietf.org/doc/html/rfc6830
Cisco LISP Configuration Guide: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iprouting/configuration/15-sy/irt-15-sy-book/irt-lisp.html

**Question: 64**

Refer to the exhibit.

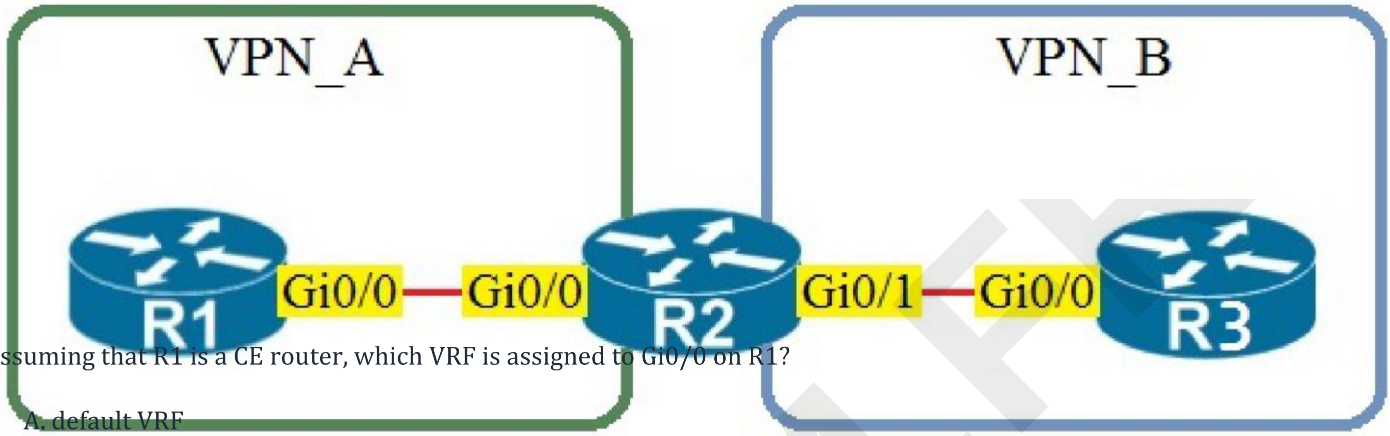

Assuming that R1 is a CE router, which VRF is assigned to Gi0/0 on R1?

A. default VRF
B. VRF VPN_A
C. VRF VPN_B
D. management VRF

**Answer: A**

**Explanation:**

Default VRF:

All Layer 3 interfaces exist in the default VRF until they are assigned to another VRF.

Routing protocols run in the default VRF context unless another VRF context is specified.

The default VRF uses the default routing context for all show commands.

The default VRF is similar to the global routing table concept in Cisco IOS.

**Question: 65**

What are two benefits of virtualizing the server with the use of VMs in a data center environment? (Choose two.)

A. reduced rack space, power, and cooling requirements
B. smaller Layer 2 domain
C. increased security
D. speedy deployment
E. reduced IP and MAC address requirements

**Answer: AD**

**Explanation:**

The correct answers are **A. reduced rack space, power, and cooling requirements** and **D. speedy deployment.**

Virtualization, through the use of Virtual Machines (VMs), allows multiple operating systems and applications to run concurrently on a single physical server. This consolidation directly translates to a smaller physical footprint in the data center, reducing the need for numerous physical servers. Consequently, less hardware means lower power consumption and reduced heat generation, thereby decreasing cooling demands. This leads to significant cost savings and a more efficient use of resources. Virtualization also enhances resource utilization because the same hardware resources are shared by multiple virtual machines, rather than sitting idle when the hosted application is not actively used. The capability to quickly deploy new virtual machines is another significant benefit, allowing for rapid scalability and agility in response to changing business needs.

Setting up a VM is typically faster and easier compared to acquiring, installing, and configuring physical hardware. This quick deployment is important for faster software delivery, development processes and reacting to immediate demands.

Option B is incorrect. Virtualization doesn't inherently reduce the size of a Layer 2 domain; that's typically addressed through network segmentation techniques like VLANs or VXLANs, which can be used independently of virtualization. Option C is incorrect. Virtualization does not increase security by itself.

Although, some hypervisors offer security benefits, they are not the primary driver for choosing VM implementation. It also introduces a new layer of complexity that could become a security vulnerability if not managed correctly. Option E is incorrect. Virtualization doesn’t reduce IP and MAC address requirements; each VM still requires a unique IP address and MAC address to communicate on the network.

Here are some authoritative links for further research:

**VMware's What is Server Virtualization:**https://www.vmware.com/topics/glossary/content/server-virtualization.html
**Red Hat's What is Virtualization:**https://www.redhat.com/en/topics/virtualization/what-is-virtualization
**Microsoft's Virtualization:**https://learn.microsoft.com/en-us/virtualization/

**Question: 66**

Which statement about route targets is true when using VRF-Lite?

A. Route targets control the import and export of routes into a customer routing table.
B. When BGP is configured, route targets are transmitted as BGP standard communities.
C. Route targets allow customers to be assigned overlapping addresses.
D. Route targets uniquely identify the customer routing table.

**Answer: A**

**Explanation:**

Let's analyze the role of route targets within the context of VRF-Lite. VRF-Lite, or Virtual Routing and Forwarding Lite, provides a mechanism to create multiple independent routing domains within a single physical router. Each VRF has its own routing table, allowing for segmentation and isolation of network traffic.

Route targets, crucial for VRF-Lite with MP-BGP (Multiprotocol BGP), aren't about uniquely identifying a customer VRF itself; that's the job of the Route Distinguisher (RD). Rather, route targets dictate which prefixes should be imported into and exported from a specific VRF routing table.

Option A is correct because route targets are attributes, typically configured on BGP routers, that control the import and export of route information. An export route target is attached to a route when it's advertised, and an import route target is configured on a VRF. If the two match, the route will be placed in the associated VRF's routing table. Option B is incorrect; while route targets are BGP extended communities, they are not standard communities. Standard BGP communities are shorter, 32-bit values, whereas extended communities, used for route targets, are longer, 64-bit values.

Option C is incorrect; while VRF-Lite enables overlapping address spaces among customers, the mechanism facilitating this is the VRF itself, not the route target. Route targets contribute to the traffic control of these overlapping addresses, not the address overlap itself. Option D is incorrect; route targets do not uniquely identify a customer VRF's routing table. The Route Distinguisher (RD) is the identifier that makes a VRF routing table unique. Route targets work in conjunction with the RD, but their function is different. Therefore, route targets function primarily as a filtering tool determining which routes are moved to or from a VRF's routing table.

For further research, please refer to the following resources:

Cisco documentation on VRF-Lite: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_vrf/configuration/15-mt/vrf-15-mt-book.html
RFC 4364 (BGP/MPLS IP VPNs): https://datatracker.ietf.org/doc/html/rfc4364

## Question: 67

Which LISP infrastructure device provides connectivity between non-LISP sites and LISP sites by receiving non-LISP traffic with a LISP site destination?

A. PITR

B. map resolver

C. map server

D. PETR

**Answer: A**

**Explanation:**

The correct answer is A, PITR (Proxy Ingress Tunnel Router). A PITR is a crucial component in a LISP (Locator/ID Separation Protocol) deployment, specifically when dealing with traffic originating from non-LISP sites destined for LISP sites. It acts as a gateway, translating non-LISP traffic into LISP-encapsulated traffic.

When a device in a non-LISP site needs to communicate with an endpoint within a LISP site, it sends its packets normally based on IP addresses. The PITR, configured with knowledge of LISP site reachability, intercepts these packets. It then performs the lookup to find the RLOC (Routing Locator) associated with the destination EID (Endpoint Identifier), and encapsulates the original packet with LISP headers. This allows the packet to traverse the LISP core network and reach its intended destination. Conversely, a PETR (Proxy

Egress Tunnel Router) handles traffic in the opposite direction - from LISP to non-LISP sites. The map resolver and map server are involved in the mapping process between EIDs and RLOCs within the LISP infrastructure. These components do not directly handle the encapsulation of non-LISP traffic.

Therefore, the PITR's unique role in acting as an ingress point for non-LISP traffic entering the LISP domain makes it the correct choice. It ensures seamless connectivity between legacy networks and networks utilizing LISP. The PITR facilitates interworking between traditional IP routing and the LISP overlay network.

Cisco's LISP documentation offers a comprehensive overview of LISP components and their functionalities. Additional resources on the LISP protocol can be found on IETF LISP Working Group pages. These links provide further technical details about PITR and its role in LISP deployments.

## Question: 68

Which statement explains why Type 1 hypervisor is considered more efficient than Type2 hypervisor?

A. Type 1 hypervisor is the only type of hypervisor that supports hardware acceleration techniques.

B. Type 1 hypervisor relies on the existing OS of the host machine to access CPU, memory, storage, and network resources.

C. Type 1 hypervisor runs directly on the physical hardware of the host machine without relying on the underlying OS.

D. Type 1 hypervisor enables other operating systems to run on it.

**Answer: C**

**Explanation:**

The correct answer is C, "Type 1 hypervisor runs directly on the physical hardware of the host machine without relying on the underlying OS." Type 1 hypervisors, also known as bare-metal hypervisors, achieve superior efficiency because they install directly onto the host's hardware. This direct access eliminates the overhead of an intermediate operating system, which is required in Type 2 hypervisors. A Type 2 hypervisor, conversely, operates as an application within a conventional OS, which needs to translate virtual machine requests into hardware instructions. This extra layer of abstraction introduces latency and consumes resources. By bypassing the host OS, Type 1 hypervisors can manage hardware resources more directly and efficiently. This results in lower resource consumption, faster performance, and more consistent access to hardware, making virtual machines more performant. Examples of Type 1 hypervisors include VMware ESXi and Microsoft Hyper-V. Therefore, the primary reason for Type 1's greater efficiency is the absence of a host OS intermediary, allowing for a more streamlined interaction with the hardware. This design is crucial for server virtualization, where performance and resource utilization are paramount. The direct interaction ensures better isolation between VMs and reduces resource contention compared to Type 2 architectures.

For further information, refer to these authoritative resources:

1. **VMware Documentation on Hypervisor Architecture:**https://docs.vmware.com/en/VMware-vSphere/index.html (Look for concepts related to ESXi architecture)
2. **Microsoft's Documentation on Hyper-V:**https://docs.microsoft.com/en-us/virtualization/hyper-v-on-windows/ (Examine Hyper-V's architecture as a type-1 hypervisor)
3. **TechTarget's definition of Type 1 and Type 2 hypervisors:** https://www.techtarget.com/searchservervirtualization/definition/hypervisor (Provides a comparison and definitions)

## Question: 69

Which statement about VXLAN is true?

A. VXLAN encapsulates a Layer 2 frame in an IP-UDP header, which allows Layer 2 adjacency across router boundaries.

B. VXLAN uses the Spanning Tree Protocol for loop prevention.

C. VXLAN extends the Layer 2 Segment ID field to 24-bits, which allows up to 4094 unique Layer 2 segments over the same network.

D. VXLAN uses TCP as the transport protocol over the physical data center network.

**Answer: A**

**Explanation:**

The correct answer is **A. VXLAN encapsulates a Layer 2 frame in an IP-UDP header, which allows Layer 2 adjacency across router boundaries.** Here's why:

VXLAN (Virtual Extensible LAN) is a network virtualization technology that overcomes the limitations of traditional VLANs in large-scale cloud and data center environments. It achieves this by encapsulating Layer 2 Ethernet frames within a Layer 4 UDP packet, along with an outer IP header. This encapsulation allows the Layer 2 frames to be transported across a Layer 3 network, essentially creating a virtual Layer 2 network overlay. Routers, which normally do not forward Layer 2 traffic, can now forward these UDP packets based on their IP header information. This effectively extends Layer 2 network segments beyond the boundaries of a single Layer 2 domain, enabling virtual machines or workloads to communicate as if they were on the same LAN, even if they are located in different physical locations or subnets. The use of UDP for transport is crucial for VXLAN's efficiency and scalability.

Option B is incorrect because VXLAN doesn't rely on the Spanning Tree Protocol (STP). Instead, VXLAN leverages a flood-and-learn approach for MAC address discovery or uses control plane protocols like BGP EVPN to distribute MAC address information, avoiding the issues with STP scalability. Option C is wrong; VXLAN extends the Layer 2 segment identifier, known as the VXLAN Network Identifier (VNI), to 24 bits. This allows for approximately 16 million (2^24) unique Layer 2 segments, not 4094. Lastly, option D is incorrect because VXLAN uses UDP, not TCP, as the transport protocol. The choice of UDP facilitates stateless operation and reduces overhead, making it more efficient for transporting encapsulated Layer 2 frames.

For further research, you can refer to these authoritative resources:

1. **RFC 7348: Virtual eXtensible Local Area Network (VXLAN):** https://datatracker.ietf.org/doc/html/rfc7348 (This is the foundational RFC for VXLAN.) 2. **Cisco's VXLAN Overview:**https://www.cisco.com/c/en/us/solutions/data-center-virtualization/virtual-extensible-lan-vxlan/index.html
3. **VMware's Understanding VXLAN:**https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/glossary/GUID-6A98F0E1-4732-4651-BA2C-52F98984B7B7.html (Provides a good overview within the VMware context)

## Question: 70

Which TCP setting is tuned to minimize the risk of fragmentation on a GRE/IP tunnel?

A. MSS

B. MTU

C. MRU

D. window size

**Answer: A**

**Explanation:**

The correct answer is A, MSS (Maximum Segment Size). MSS directly relates to the TCP payload size and is crucial in preventing IP fragmentation, especially when dealing with tunnels like GRE (Generic Routing Encapsulation). GRE adds overhead to the original IP packet, reducing the available space for the encapsulated packet. If the encapsulated packet, including TCP data, is too large, the resulting GRE packet may exceed the MTU (Maximum Transmission Unit) of the underlying network, leading to fragmentation.

Fragmentation introduces complexity, performance overhead, and risks packet loss. By adjusting the TCP MSS, we ensure that TCP segments are small enough before they are encapsulated within the GRE header, thus preventing the resulting GRE packet from exceeding the MTU. This avoids fragmentation at the IP level by proactively managing the TCP segment size. MTU, while related, dictates the overall size limit of an IP packet. MRU (Maximum Receive Unit) concerns PPP interfaces, not TCP. The TCP window size controls the data flow between sender and receiver but doesn't directly influence fragmentation. Tuning MSS prevents IP fragmentation by ensuring the data segments fit within the GRE tunnel's available space, precluding the need for reassembly at intermediate points, streamlining the transfer process.

Authoritative Links for further research:

Cisco: Understanding TCP MSS and its impact on Performance
RFC 793: Transmission Control Protocol - While an old RFC, it remains foundational for understanding TCP. RFC 2784: Generic Routing Encapsulation (GRE) - Provides details on GRE encapsulation.

**Question: 71**

Which statement describes the IP and MAC allocation requirements for virtual machines on Type 1 hypervisors?

A. Virtual machines do not require a unique IP or unique MAC. They share the IP and MAC address of the physical server.

B. Each virtual machine requires a unique IP address but shares the MAC address with the physical server.

C. Each virtual machine requires a unique IP and MAC addresses to be able to reach to other nodes.

D. Each virtual machine requires a unique MAC address but shares the IP address with the physical server.

**Answer: C**

**Explanation:**

The correct answer is C: "Each virtual machine requires a unique IP and MAC addresses to be able to reach to other nodes." This is fundamental to network communication, both within a virtualized environment and when communicating with external networks.

Let's break down why:

**IP Addresses for Identification:** IP addresses serve as logical identifiers for devices on a network. Just like physical computers, virtual machines need unique IP addresses to be individually identified and located within the network. Without a unique IP, the network wouldn't know which VM to send data to. Sharing an IP would lead to communication conflicts and network chaos.

**MAC Addresses for Data Link Layer:** MAC (Media Access Control) addresses are unique hardware identifiers assigned to network interface cards (NICs). In the context of a virtualized environment, each virtual machine gets a virtual network interface (vNIC), which needs its own unique MAC address. This is essential for network switches to correctly deliver data frames to the specific virtual machine within the virtualized environment. Without a unique MAC address, the network switch would not know which VM to forward the data to, leading to data delivery failures.

**Analogy:** Imagine postal delivery; each house needs a unique address (IP) and each person within that house needs a unique name or identifier (MAC within the LAN). Otherwise, mail would be misdirected or lost. The same applies in a network.

**Hypervisors and Virtual NICs:** Type 1 hypervisors (bare-metal) like VMware ESXi or Microsoft Hyper-V abstract the underlying physical hardware. They allow multiple virtual machines to share the physical network interface. The hypervisor creates virtual NICs (vNICs) for each VM, and these vNICs are assigned unique MAC addresses and configured with unique IP addresses by the operating system within the guest VM, acting as distinct network nodes.

**Network Communication:** For a virtual machine to communicate with other virtual machines on the same host, or with other machines in the wider network (physical or virtual), it needs these unique IP and MAC addresses. This allows layer 2 switching and layer 3 routing to function correctly. These unique identifiers enable other devices to correctly find and connect with them on the network.

In summary, assigning unique IP and MAC addresses to each virtual machine is a basic requirement for network functionality. This separation allows the logical segregation and communication for every individual VM which is a requirement for virtualized environments.

**Authoritative Links for further research:**

**VMware: Understanding Virtual Networking:**https://docs.vmware.com/en/VMware-vSphere/7.0/vsphere-esxi-vcenter-server-70-networking-guide.pdf (search for sections on virtual switches and network interface cards)
**Microsoft: Hyper-V Network Virtualization:**https://learn.microsoft.com/en-us/windows-server/networking/sdn/technologies/hyper-v-network-virtualization/hyper-v-network-virtualization **Cisco: Networking Fundamentals:**https://www.cisco.com/c/en/us/support/docs/ip/addressing/13921-address.html (general IP/MAC address information)

## Question: 72

Which two namespaces does the LISP network architecture and protocol use? (Choose two.)

A. TLOC
B. RLOC
C. DNS
D. VTEP
E. EID

**Answer: BE**

**Explanation:**

The correct answer is **B. RLOC** and **E. EID**. LISP (Locator/ID Separation Protocol) operates by separating the endpoint identifier (EID) from its routing locator (RLOC). The EID represents the endpoint's identity, typically an IP address, that remains constant regardless of its physical location. The RLOC, on the other hand, is the IP address of the network device where the endpoint connects. This separation enables mobility and multihoming, allowing endpoints to move without requiring network-wide routing updates. The LISP architecture uses a mapping system to translate EIDs to RLOCs. When a packet is destined for a specific EID, the source router consults the mapping system to find the corresponding RLOC. The packet is then encapsulated with the destination RLOC and forwarded through the core network. The receiving router, closest to the destination EID, de-encapsulates the packet and forwards it to the intended endpoint.

Therefore, EID and RLOC are the two core namespaces utilized by LISP to achieve location independence. TLOC is not a direct LISP term but it's found in SD-WAN architectures. VTEP is specific to VXLAN and not LISP. DNS is part of the service resolution but not an integral namespace within LISP architecture itself.

**Authoritative links:**

Cisco LISP Documentation
LISP IETF Draft

**Question: 73**

Which two entities are Type 1 hypervisors? (Choose two.)

A. Oracle VM VirtualBox
B. Microsoft Hyper-V
C. VMware server
D. VMware ESXi
E. Microsoft Virtual PC

**Answer: BD**

**Explanation:**

The correct answer is **B and D: Microsoft Hyper-V and VMware ESXi.**

Type 1 hypervisors, also known as bare-metal hypervisors, run directly on the host's hardware, without an underlying operating system. This architecture provides better performance and security compared to Type 2 hypervisors.

Microsoft Hyper-V is a Type 1 hypervisor that's integrated into Windows Server operating systems, although a standalone version also exists. It manages the allocation of hardware resources directly to virtual machines.
Similarly, VMware ESXi is a Type 1 hypervisor that operates directly on the hardware, often found in enterprise environments for its robustness and performance.

Options A, C, and E are incorrect. Oracle VM VirtualBox and Microsoft Virtual PC are Type 2 hypervisors that require an underlying operating system (like Windows or macOS) to run. VMware Server (now discontinued) was also a Type 2 hypervisor. Type 2 hypervisors run as applications on top of the host OS and therefore have an additional layer, resulting in potential performance overhead. The direct interaction of Type 1 hypervisors with the hardware reduces latency and increases efficiency, critical for demanding workloads like enterprise servers.

Further research:

**Microsoft Hyper-V Documentation:**https://learn.microsoft.com/en-us/virtualization/hyper-v-on-windows/
**VMware ESXi Documentation:**https://www.vmware.com/products/esxi-and-esx.html
**Hypervisor types on TechTarget:**
https://www.techtarget.com/searchservervirtualization/definition/hypervisor

**Question: 74**

DRAG DROP -
Drag and drop the LISP components from the left onto the functions they perform on the right. Not all options are used.
Select and Place:

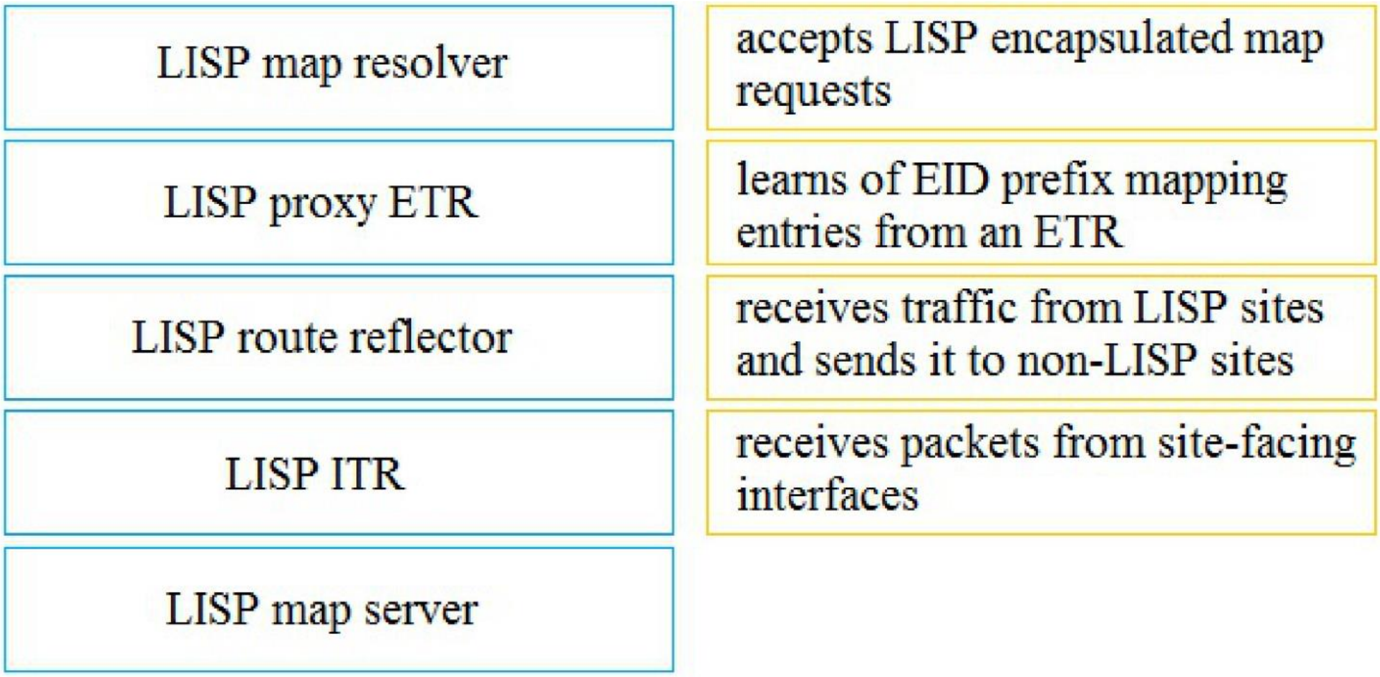

LISP map resolver

LISP proxy ETR

LISP route reflector

LISP ITR

LISP map server

accepts LISP encapsulated map requests

learns of EID prefix mapping entries from an ETR

receives traffic from LISP sites and sends it to non-LISP sites

receives packets from site-facing interfaces

**Answer:**

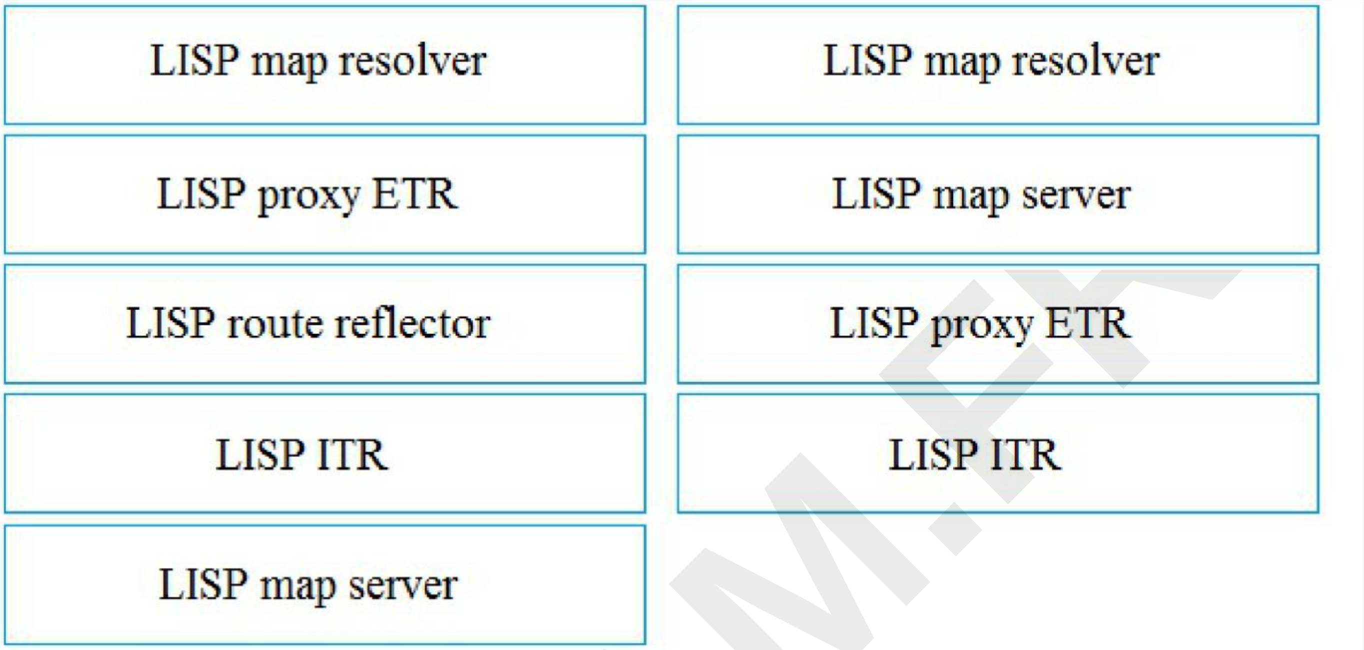

| | |
|---|---|
| LISP map resolver | LISP map resolver |
| LISP proxy ETR | LISP map server |
| LISP route reflector | LISP proxy ETR |
| LISP ITR | LISP ITR |
| LISP map server | |

**Explanation:**

Reference:
https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/DCI/5-0/LISPmobility/DCI_LISP_Host_Mobility/LISPmobile_2.html#:~:text=%E2%80%93%20Proxy%20ITR%20(PITR)%3A%20A,devices%20deployed%20at%20LISP%20sites.

### Question: 75

Which action is a function of VTEP in VXLAN?

A. tunneling traffic from IPv6 to IPv4 VXLANs

B. allowing encrypted communication on the local VXLAN Ethernet segment

C. encapsulating and de-encapsulating VXLAN Ethernet frames

D. tunneling traffic from IPv4 to IPv6 VXLANs

**Answer: C**

**Explanation:**

The correct answer is C, encapsulating and de-encapsulating VXLAN Ethernet frames. A Virtual Tunnel Endpoint (VTEP) is the key component enabling VXLAN (Virtual Extensible LAN). Its primary function is to encapsulate Ethernet frames into VXLAN packets and to de-encapsulate VXLAN packets back into Ethernet frames at the other end. This encapsulation process involves adding a VXLAN header, which includes the VXLAN Network Identifier (VNI), allowing for segmentation and multi-tenancy. When a VTEP receives an Ethernet frame destined for a remote VTEP within the same VXLAN segment, it adds the VXLAN header and an outer IP and UDP header. This new packet is then forwarded over the IP underlay network. Upon reaching the destination VTEP, the outer headers are removed, and the original Ethernet frame is delivered to the intended destination VM or endpoint within that logical VXLAN. Options A and D involve IPv6/IPv4 translation which is not a function of VTEPs. Option B about encrypted communication is also incorrect because while VXLAN itself doesn't inherently provide encryption, separate mechanisms like IPsec can be implemented for that. VTEPs are solely concerned with adding/removing the overlay header to facilitate layer 2 extension over a layer 3 network.

**Authoritative Links for further research:**

**Cisco Documentation on VXLAN:**https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/virtualized-services-data-center/white-paper-c11-687019.html
**RFC 7348 - VXLAN:**https://www.rfc-editor.org/rfc/rfc7348
**VMware Documentation on VXLAN:**https://docs.vmware.com/en/VMware-NSX/3.2/nsx-networking/GUID-C8F88F49-70FF-4501-88A0-70CE28F63D9A.html

## Question: 76

Which two actions provide controlled Layer 2 network connectivity between virtual machines running on the same hypervisor? (Choose two.)

A. Use a virtual switch provided by the hypervisor.
B. Use a virtual switch running as a separate virtual machine.
C. Use VXLAN fabric after installing VXLAN tunneling drivers on the virtual machines.
D. Use a single routed link to an external router on stick.
E. Use a single trunk link to an external Layer2 switch.

**Answer: AE**

**Explanation:**

Let's analyze the options for providing Layer 2 connectivity between VMs on the same hypervisor. Option A, "Use a virtual switch provided by the hypervisor," is correct because hypervisors like VMware ESXi or Hyper-V inherently include a virtual switch. This switch allows VMs within the same hypervisor to communicate at Layer 2 without needing to traverse a physical network. The virtual switch acts as a bridge between virtual network interfaces of VMs, enabling direct communication. For instance, VMs connected to the same port group within a VMware vSwitch can directly communicate with each other. Option E, "Use a single trunk link to an external Layer2 switch," is also a correct approach. While the VMs are on the same hypervisor, they can be connected to an external physical switch via a single trunk port. This approach is used when you need the VMs to belong to different VLANs, or to use features and policies of the external physical switch.

Option B, "Use a virtual switch running as a separate virtual machine," is incorrect because it's an inefficient method for basic Layer 2 connectivity of VMs on the same hypervisor. Running a switch as a VM introduces unnecessary overhead and complexity for simple internal communication. Option C, "Use VXLAN fabric after installing VXLAN tunneling drivers on the virtual machines," while a valid technology for extending Layer 2 over Layer 3, is also unnecessary complex when the VMs are on the same hypervisor. It's designed for cross-hypervisor or cross-datacenter L2 connectivity, not basic intra-hypervisor communication. Finally, Option D, "Use a single routed link to an external router on stick," creates a Layer 3 connection which is unsuitable for simple L2 networking within the hypervisor, and defeats the purpose of local VM to VM communication at layer 2. Therefore, options A and E are the best solutions for controlled L2 connectivity.

Authoritative Link for Virtual Switching:

VMware vSphere Networking Documentation: https://docs.vmware.com/en/VMware-vSphere/index.html
(search for "virtual switch", "vSwitch")
Microsoft Hyper-V Networking Documentation: https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/get-started/create-a-virtual-switch-for-hyper-v-virtual-machines

## Question: 77

What is a Type 1 hypervisor?

A. runs directly on a physical server and depends on a previously installed operating system
B. runs directly on a physical server and includes its own operating system
C. runs on a virtual server and depends on an already installed operating system
D. runs on a virtual server and includes its own operating system

**Answer: B**

**Explanation:**

The correct answer is **B: runs directly on a physical server and includes its own operating system.**

A Type 1 hypervisor, also known as a bare-metal hypervisor, operates directly on the hardware of a physical server without relying on a pre-existing operating system. This direct interaction with the hardware allows for efficient resource management and better performance compared to Type 2 hypervisors. The hypervisor itself includes an operating system layer which is responsible for managing the physical hardware and virtualizing resources for the guest operating systems, which are then installed on top of the hypervisor. This
independence from a host OS allows the hypervisor to be more lightweight and streamlined, contributing to increased stability and security. Examples of Type 1 hypervisors include VMware ESXi, Microsoft Hyper-V (Server), and Citrix XenServer. The key distinguishing factor is the hypervisor's position directly on the bare metal, controlling the hardware. This architecture provides a foundation for hosting and managing multiple virtual machines efficiently.

**Authoritative Links:**

**VMware Documentation:**https://www.vmware.com/topics/glossary/content/hypervisor
**Microsoft Docs:**https://learn.microsoft.com/en-us/windows-server/virtualization/hyper-v/
**TechTarget:**https://www.techtarget.com/searchservervirtualization/definition/hypervisor

## Question: 78

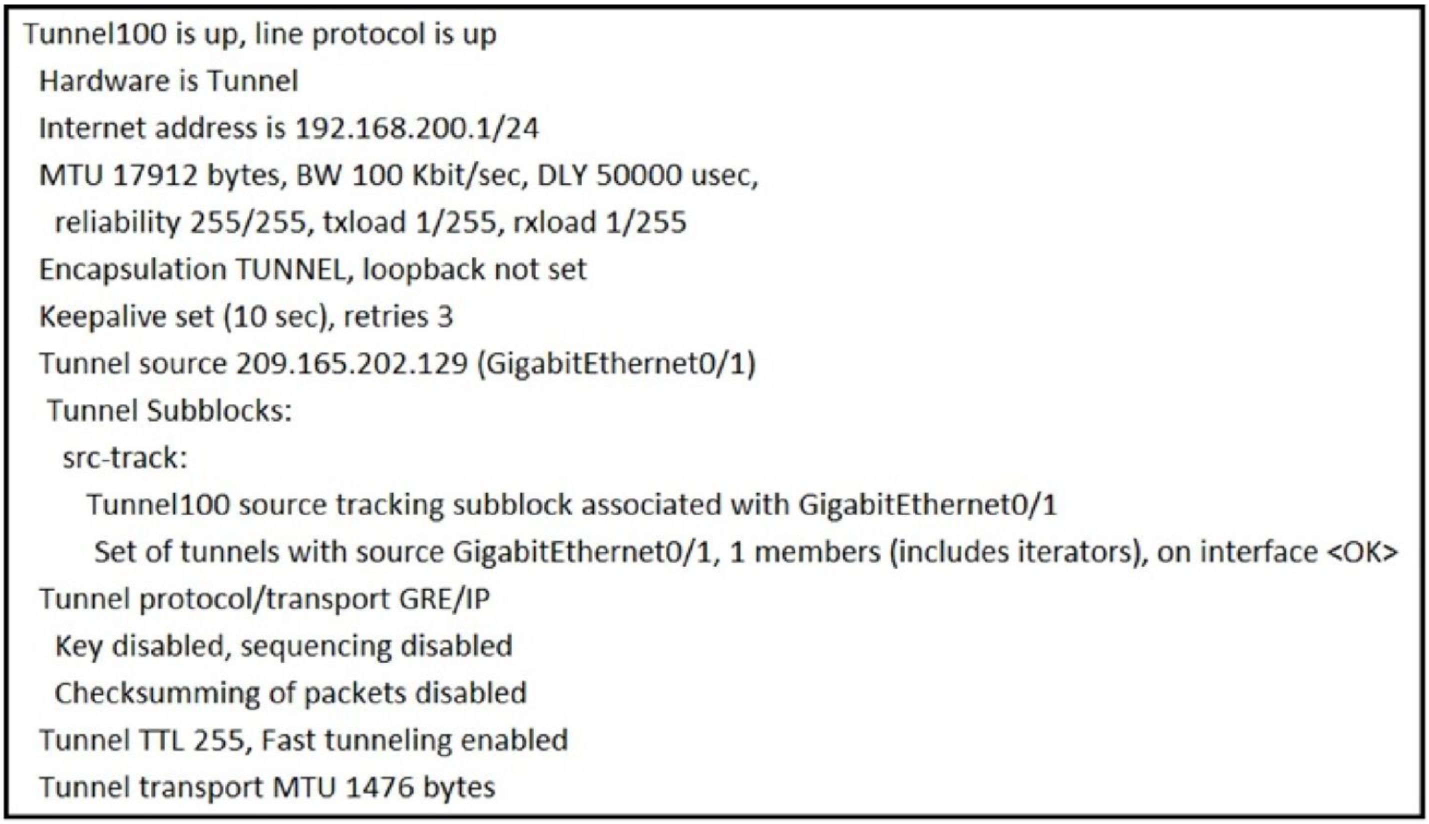

```
Tunnel100 is up, line protocol is up
  Hardware is Tunnel
  Internet address is 192.168.200.1/24
  MTU 17912 bytes, BW 100 Kbit/sec, DLY 50000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive set (10 sec), retries 3
  Tunnel source 209.165.202.129 (GigabitEthernet0/1)
   Tunnel Subblocks:
     src-track:
        Tunnel100 source tracking subblock associated with GigabitEthernet0/1
         Set of tunnels with source GigabitEthernet0/1, 1 members (includes iterators), on interface <OK>
  Tunnel protocol/transport GRE/IP
    Key disabled, sequencing disabled
    Checksumming of packets disabled
  Tunnel TTL 255, Fast tunneling enabled
  Tunnel transport MTU 1476 bytes
```

Refer to the exhibit. A network engineer configures a GRE tunnel and enters the show interface tunnel command. What does the output confirm about the configuration?

A. The keepalive value is modified from the default value.
B. The physical interface MTU is 1476 bytes.

C. The tunnel mode is set to the default.
D. Interface tracking is configured.

**Answer: C**

**Explanation:**

The tunnel mode is set to the default.

**Question: 79**

What it the purpose of the LISP routing and addressing architecture?

A. It creates two entries for each network node, one for its identity and another for its location on the network. B. It allows LISP to be applied as a network virtualization overlay though encapsulation.

C. It allows multiple instances of a routing table to co-exist within the same router.

D. It creates head-end replication used to deliver broadcast and multicast frames to the entire network.

**Answer: A**

**Explanation:**

The correct answer is **A. It creates two entries for each network node, one for its identity and another for its location on the network.** This accurately describes the core function of the Locator/ID Separation Protocol (LISP). LISP separates the identity of a network endpoint (represented by Endpoint Identifiers or EIDs) from its location within the network (represented by Routing Locators or RLOCs).

In traditional IP routing, an IP address serves dual purposes, both identifying the endpoint and indicating its network location. This makes it challenging to move endpoints without reconfiguring routing. LISP overcomes this limitation by introducing two distinct namespaces. EIDs are used to identify the endpoint, and they are kept constant, regardless of where the endpoint physically resides. RLOCs, on the other hand, denote the network location (specifically the location of ingress routers).

This separation allows for endpoint mobility and multi-homing because a change in location (RLOC) does not require a corresponding change in the identifier (EID). The mapping between EIDs and RLOCs is maintained in a distributed database. When a packet needs to reach an endpoint, the ingress router queries this mapping database to find the RLOC(s) associated with the destination EID, encapsulates the original packet with a LISP header, and forwards it to the appropriate RLOC. This creates a network virtualization overlay effect. Option B correctly highlights this overlay aspect of LISP but doesn't explain the fundamental purpose of the
addressing architecture, which is separation of ID and Location. Option C relates to VRF, not LISP specifically. Option D describes some multicast use cases but doesn't pertain to the core LISP function.

**Authoritative Links for Further Research:**

**RFC 6830: The Locator/ID Separation Protocol (LISP):**https://datatracker.ietf.org/doc/html/rfc6830 **Cisco LISP Overview:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/lisp/index.html **LISP Wikipedia:**https://en.wikipedia.org/wiki/Locator/Identifier_Separation_Protocol

**Question: 80**

What function does VXLAN perform in a Cisco SD-Access deployment?

A. policy plane forwarding

B. control plane forwarding

C. data plane forwarding

D. systems management and orchestration

**Answer: C**

**Explanation:**

The correct answer is C, data plane forwarding. VXLAN (Virtual Extensible LAN) acts as the primary data plane encapsulation mechanism within a Cisco SD-Access (Software-Defined Access) fabric. It facilitates the creation of logical Layer 2 networks (VNets) on top of an existing Layer 3 infrastructure. This overlay network allows for flexible segmentation and mobility of devices, regardless of their physical location. When a frame enters the SD-Access fabric, it's encapsulated with VXLAN headers, including a VNI (VXLAN Network Identifier) which uniquely identifies the specific VNet. This encapsulated packet, now with a UDP header as well, is routed through the underlay (physical network) based on IP addressing, until it reaches the destination fabric node where the VXLAN header is stripped and the original frame is forwarded. Essentially, VXLAN handles the forwarding of user traffic based on the VNI, effectively creating the virtualized network that is the cornerstone of SD-Access. Options A, B, and D relate to other aspects of SD-Access. Policy plane is typically handled by Cisco DNA Center, control plane involves the use of protocols like LISP, and system management is associated with DNA Center's administrative functions.

Authoritative Links:

1. **Cisco SD-Access Solution Guide:**https://www.cisco.com/c/en/us/solutions/enterprise/software-defined-access/solution-overview.html (This provides a broad overview of SD-Access, which uses VXLAN)
2. **VXLAN RFC:**https://datatracker.ietf.org/doc/rfc7348/ (This document specifies the VXLAN protocol itself)

3. **Cisco DNA Center Documentation:**https://www.cisco.com/c/en/us/support/cloud-systems-management/dna-center/products-technical-reference-list.html (For information on policy and control plane aspects).