CISCO 350-401 Exam Questions

Total Questions: 340+
Demo Questions: 30
Version: Updated for 2025

Prepared and Verified by Cert Empire – Your Trusted IT Certification Partner

For Access to the full set of Updated Questions – Visit: 350-401 Exam Dumps by Cert Empire

Question: 1

What are two benefits of YANG? (Choose two.)

A: It enforces the use of a specific encoding formatfor NETCONF.

B: It collectsstatistical constraint analysisinformation.

C: It enables multiple leaf statements to exist within a leaf list.

D: It enforces configuration semantics.

E: It enforces configuration constraints.

Correct Answer:

D, E

Explanation:

YANG (Yet Another Next Generation) is a data modeling language used to define the content of configuration and state data for network management protocols like NETCONF and RESTCONF. Its primary benefits are to ensure data is well-formed, valid, and semantically correct.

YANG enforces configuration semantics (D) by defining the meaning of data elements, their data types (e.g., integer, string, boolean), and the relationships between them. It also enforces configuration constraints (E) through various statements like range, length, pattern, and must, which impose specific rules that the data must follow to be considered valid. These features ensure that network devices receive consistent, predictable, and valid configurations.

Why Incorrect Options are Wrong:

A: This is incorrect. YANG is explicitly designed to be encoding-independent. While it is commonly used with NETCONF (which uses XML), the same YANG model can be used with RESTCONF, which supports both XML and JSON.

B: This is incorrect. YANG is a language for modeling data; it does not have a built-in mechanism to collect or analyze statistical information. That is the role of the network management system or the device agent.

C: This describes a specific language feature (leaf-list), which allows for a list of simple values. While a useful feature, it is a specific construct rather than a high-level benefit like ensuring overall data validity and meaning.

References:

1. RFC 7950: The YANG 1.1 Data Modeling Language, Section 1, Introduction: "YANG is an encoding-independent data modeling language... YANG is used to model both configuration and state data... YANG models can define constraints to be enforced on the data, giving formally defined meaning to the data." This supports options D and E and refutes option A.

2. Cisco IOS XE Programmability Configuration Guide, "YANG Data Modeling" chapter: "YANG is a data modeling language that is used to model configuration and state data... It can be used to define the structure, syntax, and semantics of data... YANG is independent of the management protocol, such as NETCONF or RESTCONF." This confirms that YANG defines semantics (D) and is protocol/encoding independent (refuting A).

3. RFC 7950: The YANG 1.1 Data Modeling Language, Section 9.4.5, The must Statement: "The must statement, which is a substatement to any data definition statement, takes as an argument a string that is an XPath 1.0 expression... The constraint is enforced when a datastore is validated." This is a direct example of how YANG enforces constraints (E).

Question: 2

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

Correct Answer:

A, E

Explanation:

In the Cisco Software-Defined Access (SD-Access) architecture, network devices are assigned specific roles within the fabric. The edge node is a fabric switch that connects wired endpoints (e.g., user computers, phones, IoT devices) to the SD-Access fabric. The border node is a fabric switch that provides the connectivity between the SD-Access fabric and any external network, such as a traditional enterprise core, WAN, or data center. It is responsible for translating fabric-specific policies and protocols to those used in the external networks.

Why Incorrect Options are Wrong:

B: The vBond controller is a fundamental component of the Cisco SD-WAN (Viptela) solution, responsible for orchestration and authentication, not the SD-Access fabric.

C: "Access switch" is a generic network term. In the context of SD-Access, the specific and precise term for this function is "edge node."

D: "Core switch" is a generic network term. While a device in the underlay may perform core functions, it is not a defined SD-Access fabric role.

References:

1. Cisco. (2021). Cisco SD-Access Solution Design Guide (CVD). In the "Fabric Device Roles" section of Chapter 2, "Solution Components," the guide explicitly lists the main device roles: "The main device roles in the SD-Access fabric are: Control plane node, Border node, Edge node". This directly validates options A and E.

2. Cisco. (2023). Cisco Campus LAN and Wireless LAN Design Guide. In the chapter "SD-Access Fabric Design," the section "Fabric Roles" details the functions of the Border Node and Edge Node. It describes the Edge Node as the point of connection for endpoints and the Border Node as the gateway to networks outside the fabric.

3. Cisco. (2023). Cisco DNA Center User Guide, 2.3.5. In the "About the Fabric" section, the documentation describes the components of an SD-Access fabric, defining the roles of "Fabric edge nodes" and "Fabric border nodes" consistent with the correct answer.

Question: 3

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.

Correct Answer:

A, D

Explanation:

Traffic policing is a Quality of Service (QoS) mechanism used to enforce a traffic rate limit. A key characteristic of policing is that it drops packets that exceed the configured rate, although it can also be configured to re-mark them. This action is immediate and does not involve queuing. For maximum efficiency, policing should be implemented at the network ingress, as close to the traffic source as possible. This prevents non-conforming traffic from consuming bandwidth across the network core, preserving resources for compliant traffic.

Why Incorrect Options are Wrong:

B: Queuing excess traffic to smooth out bursts is a function of traffic shaping, not policing.

C: Applying policing near the destination is inefficient, as the unwanted traffic has already traversed and consumed network resources.

E: Delaying traffic by buffering it in a queue is characteristic of traffic shaping, not policing.

References:

1. Cisco, "Policing and Shaping Overview," Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.2. In the "Policing Versus Shaping" section, it states, "When the traffic rate reaches the configured maximum rate, excess traffic is dropped (or remarked)." It also clarifies, "In contrast, traffic shaping retains excess packets in a queue..." This directly supports option D and refutes B and E.

2. Cisco, "Enterprise QoS Solution Reference Network Design Guide Version 3.3," Chapter 2: QoS Design Principles. The guide states, "Classification, marking, and policing functions are most effectively and efficiently performed at the network edge." This principle supports applying policing as close to the source as possible (Option A) rather than the destination (Option C).

3. Cisco, "Quality of Service (QoS) Configuration Guide, Cisco IOS XE Gibraltar 16.12.x (Catalyst 9300 Switches)," Chapter: Configuring Policing. The overview section states, "A policer determines whether a packet is in or out of profile and specifies the actions to be taken on the packet. These actions, which include dropping the packet or rewriting a QoS field, are applied to all packets of a particular traffic stream." This confirms the primary actions of policing.

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.

Correct Answer:

C

Explanation:

Cisco Express Forwarding (CEF) is a high-performance, proprietary Layer 3 switching technology. Its architecture is designed to offload packet forwarding from the main CPU. To achieve this, CEF pre-populates and maintains two key data structures in the data plane: the Forwarding Information Base (FIB) and the Adjacency Table. The FIB is derived from the IP routing table (RIB) and contains all known routes, while the Adjacency Table contains the corresponding Layer 2 next-hop information (e.g., MAC addresses). This allows for extremely fast lookups and packet forwarding without involving the router's main processor for each packet.

Why Incorrect Options are Wrong:

A: This describes process switching, the slowest method, where the CPU must perform a route lookup and forward each packet individually. CEF is designed to avoid this.

B: This describes fast switching, an older method that builds a route cache after the first packet of a flow is process-switched. CEF is more efficient as its tables are pre-populated.

D: This is another characteristic of process switching, where the forwarding decision is handled by a software process managed by the Cisco IOS scheduler, leading to high CPU utilization.

References:

1. Cisco Systems, "Cisco Express Forwarding (CEF) White Paper," Document ID: 13732. In the "CEF Components" section, it states, "CEF uses two main components: a Forwarding Information Base (FIB) and adjacency tables."

2. Cisco Systems, "Troubleshooting Cisco Express Forwarding (CEF)-Related Drops," Document ID: 116244. The "CEF Components" section details, "CEF maintains two sets of tables in the data plane: the Forwarding Information Base (FIB) and the Adjacency table."

3. Rami, R. S. (2020). CCNP and CCIE Enterprise Core ENCOR 350-401 Official Cert Guide. Cisco Press. Chapter 4, "Understanding Cisco Express Forwarding," section "CEF Components," states, "CEF uses two primary data structures that are stored in the data plane: the Forwarding Information Base (FIB) and the adjacency table." The chapter also explicitly contrasts CEF with process switching (described in options A and D) and fast switching (described in option B).

Question: 5

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.

Correct Answer:

A

Explanation:

The Routing Information Base (RIB) is a control plane construct that contains all the routing information learned by the router from various sources, such as static routes and dynamic routing protocols. The control plane processor analyzes the RIB to determine the best path to each destination network. This best-path information is then used to populate the Forwarding Information Base (FIB). The FIB is a data plane structure, optimized for fast lookups, that the router uses to make actual packet forwarding decisions. Thus, the FIB is directly derived from the content of the RIB.

Why Incorrect Options are Wrong:

B: This reverses the relationship. The FIB is a subset of the RIB (containing only the best paths), not the other way around. The information flows from RIB to FIB.

C: Forwarding (switching) decisions are made using the FIB, not the RIFurthermore, standard IP forwarding is based on the destination prefix, not the source prefix.

D: The RIB is where all IP routing information is stored. The FIB contains only the active, best-path routes selected from the RIB for the purpose of high-speed forwarding.

References:

1. Cisco Press, CCNP and CCIE Enterprise Core ENCOR 350-401 Official Cert Guide. In Chapter 10, "Understanding Cisco Express Forwarding (CEF)," the text states, "CEF uses the routing information base (RIB) to populate the forwarding information base (FIB). The FIB contains the destination IP network, the next-hop IP address, and the outgoing interface." This directly supports that the FIB is populated from the RIB.

2. Cisco, IP Switching: Cisco Express Forwarding Configuration Guide, Cisco IOS Release 15M&T. In the "Cisco Express Forwarding Overview" section, it is explained that "Information from the IP routing table is used to populate the FIB." The IP routing table is another name for the RIThis confirms the unidirectional population from the RIB to the FIB.

3. Cisco, Cisco SD-WAN: WAN Edge Onboarding and Provisioning. In the "Control Plane and Data Plane" section, the document describes the fundamental architecture: "The control plane of a router builds and maintains the routing and forwarding tables (RIB and FIB)... The data plane uses the tables created by the control plane to forward traffic." This highlights the distinct roles and the dependency of the FIB (data plane) on the RIB (control plane).

Question: 6

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

Correct Answer:

A, C

Explanation:

Stateful Switchover (SSO) is a redundancy feature that provides high availability by using two supervisor engines. The core mechanism of SSO involves the active supervisor continuously synchronizing its state information—including configuration and Layer 2 forwarding tables—with the standby supervisor. This synchronization is essential for a rapid and stateful failover. While SSO effectively maintains Layer 2 operations during a switchover, it does not preserve Layer 3 routing protocol adjacencies on its own. To achieve uninterrupted Layer 3 packet forwarding and prevent routing protocol reconvergence, SSO must be used in conjunction with Non-Stop Forwarding (NSF).

Why Incorrect Options are Wrong:

B: SSO by itself does not preserve the Layer 3 multicast state. This functionality requires the combination of SSO with NSF.

D: A key aspect of the SSO design is that the configuration is synchronized from the active to the standby supervisor, not configured separately on each.

E: SSO is specifically designed to provide uninterrupted Layer 2 operations. It does not require NSF for this purpose; NSF is for Layer 3.

References:

1. Cisco IOS XE Bengaluru 17.6.x (Catalyst 9600 Switches) - High Availability Configuration Guide, "Stateful Switchover (SSO)" Chapter: "Stateful Switchover (SSO) works by establishing one of the supervisor engines as the active supervisor engine while the other supervisor engine is designated as the standby supervisor engine... The active supervisor engine has the SSO function, which synchronizes the startup-config and running-config files, and all other system state information from the active supervisor engine to the standby supervisor engine." This supports option A.

2. Cisco IOS XE Bengaluru 17.6.x (Catalyst 9600 Switches) - High Availability Configuration Guide, "NSF and SSO" Chapter: "When a networking device restarts, all routing peers of the device detect that the device went down and then came back up. This transition results in a routing flap, which could spread across multiple routing domains. Routing flaps caused by a restarting device can be prevented by using NSF. NSF allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a switchover. With NSF, routing peers do not experience routing flaps." This supports option C, highlighting the need for NSF for uninterrupted Layer 3 operations.

3. Cisco IOS XE Software-Defined WAN Configuration Guide, Release 17.x - "High Availability" Chapter, "Stateful Switchover" section: "SSO provides redundancy for the control plane. In the event of a failure of the active RP [Route Processor], the standby RP immediately takes over control plane functions. SSO synchronizes all the necessary state information between the RPs. However, SSO does not provide any redundancy for the forwarding plane. Nonstop forwarding (NSF) must be used to provide redundancy for the forwarding plane." This further clarifies the distinct roles of SSO (control plane) and NSF (forwarding plane, especially Layer 3).

Question: 7

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

Correct Answer:

C

Explanation:

In a Cisco Software-Defined Access (SD-Access) wireless solution, the Wireless LAN Controller (WLC) is the component responsible for performing Radio Resource Management (RRM). The WLC collects radio frequency (RF) metrics from all its associated Access Points (APs) within the fabric site. It then uses this data to run RRM algorithms, such as Dynamic Channel Assignment (DCA) and Transmit Power Control (TPC), to optimize the wireless network's performance and minimize interference. While Cisco DNA Center is used to provision and manage the RRM policies on the WLC, the WLC itself is the engine that executes the real-time RRM calculations and adjustments.

Why Incorrect Options are Wrong:

A: DNA Center: This is the central management, automation, and assurance platform. It configures RRM policies but does not perform the real-time RRM computations.

B: control plane node: This component manages the LISP and VXLAN overlay protocols for the fabric, mapping endpoint identifiers to locations, which is unrelated to RF management.

D: Cisco CMX: The Connected Mobile Experiences (CMX) platform is used for location-based services and analytics, not for controlling the RF environment through RRM.

References:

1. Cisco SD-Access Wireless Design Guide: In the section "Role of the WLC in SD-Access Wireless," the guide explicitly states: "The WLC is the single point of management and control for the APs and wireless clients. It is responsible for the following functions: ... Radio Resource Management (RRM) for all APs in the fabric site."

2. Cisco Catalyst 9800 Series Wireless Controller Software Configuration Guide, Cisco IOS XE Bengaluru 17.6.x, "Radio Resource Management" Chapter: This guide details the WLC's role and configuration for RRM, stating, "The Radio Resource Management (RRM) software in the controller is a built-in RF engineer that is embedded in the Cisco wireless network." This confirms the RRM function resides on the controller.

3. Deploying SD-Access for Wireless with Cisco Catalyst 9800 WLCs and Cisco DNA Center, Cisco Live Presentation BRKEWN-2001 (2022): Slide 13, "SD-Access Wireless Architecture," shows the WLC as the component handling "Wireless Control Plane" functions, which traditionally include RRM, while DNA Center is positioned as the "Management/Orchestration" layer.

Question: 8

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.

Correct Answer:

B

Explanation:

The Cisco vSmart controller is the central brain of the Cisco SD-WAN overlay network and is responsible for managing the control plane. It establishes secure Overlay Management Protocol (OMP) sessions with all WAN Edge routers to disseminate routing information, security policies, and data policies. By centralizing the control plane, the vSmart controller ensures that all network-wide policies are consistently enforced across the entire SD-WAN fabric, dictating how data plane traffic is handled by the edge devices without participating in the data forwarding itself.

Why Incorrect Options are Wrong:

A: Authentication and authorization of devices joining the fabric are primary functions of the vBond orchestrator, which acts as the initial point of contact.

C: The centralized network management system (NMS) with a graphical user interface for configuration, monitoring, and troubleshooting is the role of vManage.

D: The data plane, which involves forwarding user traffic through secure IPsec tunnels, is managed directly by the WAN Edge (vEdge/cEdge) routers.

References:

1. Cisco Systems, "Cisco SD-WAN Design Guide," September 2021. In the "Cisco SD-WAN Components" section, it states, "The vSmart controller is the centralized brain of the solution and is responsible for the control plane." It also clarifies, "The vBond orchestrator is responsible for the initial authentication and authorization of all elements into the network," and "vManage is the centralized network management system."

2. Cisco Systems, "Cisco SD-WAN Overlay Network Bring-Up Process," July 2022. The document details the bring-up sequence, where the vSmart controller's role is described in Step 6: "The vSmart controller authenticates the WAN Edge router and pushes policies and routing information to the router via OMP." This confirms its control plane function.

3. Cisco Systems, "SD-WAN Design and Deployment Guide for Cisco IOS XE SD-WAN Software," Release 17.x. Chapter: "Cisco SD-WAN Architecture." This guide explicitly defines the roles: "vSmart Controller: This software-based component is responsible for the centralized control plane of the network... WAN Edge: This is the physical or virtual device at a site that is responsible for the data plane..."

Question: 9

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

Correct Answer:

C

Explanation:

In a Cisco SD-WAN deployment, a VPN is a segmentation construct analogous to a Virtual Routing and Forwarding (VRF) instance in a traditional network. It creates a distinct, isolated routing and forwarding domain within the SD-WAN fabric. This allows for the logical separation of traffic for different tenants, departments, or applications, ensuring that traffic from one VPN does not leak into another unless explicitly permitted by policy. Each VPN maintains its own routing table, providing a secure, virtualized environment.

Why Incorrect Options are Wrong:

A: A common exchange point for services is more accurately described as a service insertion point or a regional hub, not a VPN, which is a segmentation container.

B: This is an imprecise description. A VPN is a complete virtualized environment, not merely an "attribute" used to identify a set of services.

D: The virtual channel for control plane information between WAN Edge devices and vSmart controllers is a DTLS or TLS tunnel, not a data plane VPN.

References:

1. Cisco SD-WAN Design Guide, "Segmentation (VPN)" section. The guide states, "Cisco SD-WAN provides the capability to segment the network. Segmentation is the equivalent of VRFs in a traditional network. In Cisco SD-WAN, segments are called VPNs. Each VPN is equivalent to a VRF, which provides a separate routing and forwarding table."

2. Cisco SD-WAN Configuration Guide, Cisco IOS XE Release 17.x, "Configure Segmentation (VPNs)" chapter. The documentation explains, "A key feature of the Cisco SD-WAN solution is the ability to segment the network. In the Cisco SD-WAN overlay network, segmentation is achieved by setting up multiple VPNs, which are equivalent to VRFs in a traditional network. Each VPN is a separate routing domain."

3. Cisco SD-WAN End-to-End Deployment Guide, "Segmentation" section. This guide details the use of VPNs for segmentation: "The Cisco SD-WAN solution provides segmentation by using VPNs. VPNs are used to segment the user network. Each VPN is a private network space and has its own forwarding table."

Question: 10

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

Correct Answer:

A

Explanation:

In a Cisco Software-Defined Access (SD-Access) architecture, the fabric control plane node is responsible for managing the endpoint-to-location mapping system. It runs the Locator/ID Separation Protocol (LISP) Map-Server and Map-Resolver functions. This creates a centralized host tracking database that maps each endpoint's identifier (EID), such as its IP or MAC address, to its current location, which is the Routing Locator (RLOC) of the fabric edge node it is connected to. This database is essential for forwarding traffic within the fabric overlay.

Why Incorrect Options are Wrong:

B: fabric wireless controller: The WLC integrates with the fabric and is responsible for registering wireless clients with the control plane node, but it does not manage the central EID mapping database itself.

C: fabric border node: The border node's primary function is to connect the SD-Access fabric to external networks (e.g., WAN, data center), handling policy and protocol translation, not internal host tracking.

D: fabric edge node: The edge node is the location (RLOC) where endpoints connect. It registers its connected endpoints with the control plane node but does not maintain the fabric-wide mapping database.

References:

1. Cisco Systems, Inc. (2021). Cisco SD-Access Solution Design Guide (CVD). In the "SD-Access Fabric Components and Roles" section, it states, "Control Plane Node: This node

contains the settings, protocols, and tables to provide the endpoint-to-location (EID-to-RLOC) mapping system for the fabric. This function is based on the LISP (Locator/ID Separation Protocol) Map-Server/Map-Resolver functionality."

2. Odom, W., et al. (2020). CCNP and CCIE Enterprise Core ENCOR 350-401 Official Cert Guide. Cisco Press. Chapter 10, "Understanding Cisco SD-Access," in the section "SD-Access Control Plane," explains: "The control plane node holds a database that tracks all user devices in the fabric... The database maps each endpoint's EID to its current RLOC."

3. Cisco Systems, Inc. (2023). Software-Defined Access Solution Test Report. In the "SD-Access Architecture" section, the document details the roles, specifying that the control plane node "maintains an EID-to-RLOC mapping database for all endpoints in the fabric."

Question: 11

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

Correct Answer:

D

Explanation:

In a Cisco Software-Defined Access (SD-Access) fabric, the Locator/ID Separation Protocol (LISP) is the fundamental control plane protocol. LISP separates an endpoint's identity (Endpoint Identifier or EID), typically its IP or MAC address, from its location (Routing Locator or RLOC), which is the IP address of the fabric edge node it is connected to. The control plane node in the fabric maintains a LISP mapping database of all EID-to-RLOC associations. When a fabric device needs to forward traffic to an endpoint, it queries this database using LISP to resolve the destination EID to its current RLOC, enabling traffic routing across the fabric.

Why Incorrect Options are Wrong:

A: DHCP: DHCP is a network management protocol used for dynamically assigning IP addresses to endpoints. It is not the control plane protocol for endpoint mapping in the fabric.

B: VXLAN: VXLAN is the data plane encapsulation protocol. It creates the overlay tunnels used to transport traffic between fabric edge nodes but does not handle control plane mapping and resolution.

C: SXP: The SGT Exchange Protocol (SXP) is used to propagate Security Group Tags (SGTs) for policy enforcement, particularly to devices that are not natively part of the fabric.

References:

1. Cisco. (2021). Cisco SD-Access Solution Design Guide (CVD). "The SD-Access solution uses the Locator/ID Separation Protocol (LISP) as its primary control plane protocol. LISP is

an IETF standard (RFC 6830) that provides a flexible and scalable control plane." Section: SD-Access Fabric Control Plane.

2. Cisco. (2023). Cisco SD-Access Technology Overview. "Control Plane: Based on Locator/ID Separation Protocol (LISP), the control plane maintains an EID-to-RLOC mapping database for all endpoints in the fabric." Section: SD-Access Architecture.

3. Cisco. (2023). Cisco DNA Center User Guide, 2.3.5. "The control plane node contains the settings for the LISP protocol, which is used in the fabric." Chapter: Design Network Settings for a Fabric.

Question: 12

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

Correct Answer:

D

Explanation:

On-premises infrastructure places servers and networking equipment within an organization's own physical facilities. This co-location allows for systems to be interconnected via a high-speed, low-latency Local Area Network (LAN). The physical proximity minimizes signal travel time and the number of network hops between communicating systems, resulting in significantly lower latency than is typically achievable when communicating with resources in a geographically distant public cloud data center over a Wide Area Network (WAN). This is a critical advantage for applications that are highly sensitive to communication delays, such as high-frequency trading or industrial control systems.

Why Incorrect Options are Wrong:

A: This describes rapid elasticity, a primary benefit of cloud computing where resources can be scaled programmatically without physical hardware changes. On-premises scaling requires hardware procurement.

B: On-premises infrastructure requires the organization to bear the full cost and management overhead of power and cooling, which is a significant operational expense, not a benefit.

C: Cloud deployments are significantly faster as they eliminate the hardware procurement, shipping, and installation lifecycle inherent to on-premises infrastructure expansion.

References:

1. Official Vendor Documentation: Cisco, "Cloud, On-Premises, or Hybrid: Which Is Best for You?", Cisco Blogs. This article discusses the trade-offs, noting that on-premises can be superior for "applications that require low latency or have high throughput needs," as the infrastructure is located locally.

2. Official Vendor Documentation: Cisco CCNA 200-301 Official Cert Guide, Volume 2, by Wendell Odom, Cisco Press, 2019. Chapter 23, "Cloud Architecture," contrasts cloud and on-premises models. It implicitly supports the latency benefit by describing on-premises as a "private data center" where all components are local, in contrast to the remote nature of public cloud services. The fundamental architectural difference (local vs. remote) is the basis for the latency advantage.

3. University Courseware: Patterson, A., & Hennessy, J. L. (2017). Computer Organization and Design RISC-V Edition: The Hardware Software Interface. Morgan Kaufmann. In discussions of data center architecture (Chapter 6), the text highlights the importance of network latency and how physical proximity within a data center is key to performance, a principle that directly applies to the on-premises model's advantage for local workloads.

Question: 13

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

A: vBond

B: vSmart

C: vManage

D: WAN Edge

Correct Answer:

A

Explanation:

The Cisco SD-WAN architecture is composed of four distinct planes: orchestration, management, control, and data. The orchestration plane is responsible for the initial bring-up, authentication, and secure onboarding of all SD-WAN components into the fabric. The vBond orchestrator is the dedicated component that performs this function. It acts as the first point of contact for WAN Edge routers and vSmart controllers, authenticates them, and provides the necessary information for them to connect to each other, thereby "orchestrating" the formation of the secure overlay network.

Why Incorrect Options are Wrong:

B: vSmart: This component operates on the control plane, responsible for distributing routing information and security policies to WAN Edge routers using the Overlay Management Protocol (OMP).

C: vManage: This is the management plane component, providing the centralized Network Management System (NMS) for configuration, monitoring, provisioning, and troubleshooting of the entire SD-WAN fabric.

D: WAN Edge: This device operates on the data plane. It is responsible for forwarding user traffic across the WAN and enforcing the policies dictated by the vSmart controller.

References:

1. Cisco Systems, Inc. (2021). Cisco SD-WAN Design Guide. "Cisco SD-WAN Architecture" chapter, section "Orchestration Plane (vBond)". The guide states, "The orchestration plane is responsible for the automatic onboarding of the SD-WAN routers into the SD-WAN fabric.

The vBond controller, or orchestrator, is the first point of authentication... The vBond orchestrator stitches the SD-WAN fabric together."

2. Cisco Systems, Inc. (2023). Cisco SD-WAN Getting Started Guide. "Cisco SD-WAN Components" chapter, section "vBond Orchestrator". This document defines the vBond orchestrator's role: "Performs initial authentication of WAN Edge devices and vSmart controllers... It is the only device in the overlay that must have a public IP address."

3. Cisco Systems, Inc. (2023). Cisco SD-WAN Configuration Guide, Cisco IOS XE Release 17.x. "System and Interfaces" chapter, section "Cisco SD-WAN Overlay Network Bring-Up Process". This section details the process where a WAN Edge router first establishes a DTLS control connection to the vBond orchestrator to get authenticated and learn the IP addresses of the vManage and vSmart controllers.

Question: 14

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

Correct Answer:

C

Explanation:

Marking is the Quality of Service (QoS) process that explicitly alters a packet to influence its treatment by downstream network devices. This is accomplished by modifying a specific field in the packet header, such as the Differentiated Services Code Point (DSCP) value in the IP header or the Class of Service (CoS) value in the 802.1Q VLAN tag. This "mark" serves as a signal that QoS-enabled devices use to apply appropriate policies, such as priority queuing or selective discarding, as the packet traverses the network.

Why Incorrect Options are Wrong:

A: policing: Policing is a traffic-conditioning mechanism that enforces a rate limit by dropping or re-marking non-conforming packets; its primary function is rate enforcement, not packet alteration for signaling.

B: classification: Classification is the process of identifying and categorizing traffic into different classes based on header information. It reads packets but does not alter them.

D: shaping: Shaping buffers excess packets in a queue to smooth out traffic bursts and enforce a traffic rate, which involves delaying packets, not modifying their headers for downstream treatment.

References:

1. Cisco IOS Quality of Service Solutions Configuration Guide, Release 15M&T, "QoS: Classification and Marking Overview". In the "Marking" section, it states: "Marking a packet is the process of changing a field in the packet. For instance, you can change the

differentiated services code point (DSCP) value in the type of service (ToS) byte of an IP packet."

2. Cisco Press, "End-to-End QoS Network Design: Quality of Service for Rich-Media & Cloud Networks, Second Edition", Chapter 2, "QoS Tools and Architectures", Section "Classification and Marking". The text explains: "Marking is the QoS tool that 'writes' a value into the header of a frame or packet... This marking can then be matched and acted on by subsequent network devices."

3. Cisco SD-WAN Quality of Service (QoS) Configuration Guide, Cisco IOS XE Release 17.x, "Forwarding and QoS Overview". The document describes the QoS process: "Classification assigns packets to a service class. After a packet is classified, you can apply a QoS policy to it. For example, you can subject the packet to a policing or marking action." This clearly separates classification from the action of marking (altering).

Question: 15

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

Correct Answer:

B

Explanation:

In a Cisco Software-Defined Access (SD-Access) fabric, the Identity Services Engine (ISE) serves as the central policy engine. When a client connects to an edge node, the edge node (acting as a RADIUS client) forwards the client's credentials to ISISE, acting as the Policy Decision Point (PDP), authenticates the user/device and evaluates configured policies to authorize access. It then sends a decision back to the edge node, which may include assigning the client to a specific Scalable Group Tag (SGT) and virtual network. The edge node then enforces this policy decision.

Why Incorrect Options are Wrong:

A: edge node: The edge node is the Policy Enforcement Point (PEP) that enforces the access decision, but it does not make the decision itself.

C: RADIUS server: While ISE functions as a RADIUS server, "Identity Services Engine" is the specific Cisco component that makes the policy decision in this architecture, making it the more precise answer.

D: control-plane node: The control-plane node is responsible for the endpoint-to-location (EID-to-RLOC) mapping database using LISP, not for client authentication or authorization policy.

References:

1. Cisco. (2021). SD-Access Solution Design Guide (CVD). "Cisco Identity Services Engine (ISE) is the policy and authentication source for all endpoints in the SD-Access fabric... ISE is the component that makes the dynamic mapping of users and devices to scalable

groups." (See Chapter: "SD-Access Solution Components", Section: "Cisco Identity Services Engine").

2. Cisco. (2023). Cisco SD-Access Fabric Provisioning Configuration Guide, Cisco IOS XE Bengaluru 17.9.x. "Cisco ISE is a required component in a Cisco SD-Access fabric. It performs the authentication and authorization of the endpoints that are connected to the fabric." (See Chapter: "Cisco SD-Access Fabric Provisioning", Section: "Cisco SD-Access Fabric Roles").

3. Cisco. (2023). Authentication and Authorization Basics Configuration Guide, Cisco IOS XE 17. "The AAA server (for example, Cisco ISE) checks the user's credentials... If the credentials are correct, the server returns a response to the network access server that the user is authenticated and authorized." (See Chapter: "Configuring Authentication", Section: "Authentication"). This document describes the fundamental AAA process where the server (ISE) makes the decision.

Question: 16

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.

Correct Answer:

D

Explanation:

The Routing Information Base (RIB) is a control plane construct that contains all routing information learned from all configured routing sources, such as routing protocols (OSPF, BGP), static routes, and directly connected networks. For any given destination prefix, the RIB may contain multiple potential paths. The router's control plane process selects the single best path based on administrative distance and metric. This best path is then installed into the Forwarding Information Base (FIB). The FIB is a data plane structure, optimized for high-speed lookups, containing only the single best route to each destination, which is used for actual packet forwarding.

Why Incorrect Options are Wrong:

A: This is incorrect. The RIB is the primary routing table, not just a list. The FIB is used for forwarding, it does not maintain network topologies.

B: This statement reverses the roles. The RIB can include many routes to a destination, while the FIB contains only the single best route.

C: The RIB is the routing table; it doesn't "create" it. Routing protocols populate the RIThe description of the FIB is too generic.

References:

1. Cisco Systems, Inc. (2022). Cisco Express Forwarding Overview.

Section: "Components of CEF"

Content: This document explains that the Forwarding Information Base (FIB) is "conceptually similar to a routing table or information base. It maintains a mirror image of the forwarding information contained in the IP routing table [RIB]." It clarifies that when the IP routing table is updated, the changes are propagated to the FIB, which contains the next-hop address information for each destination. This supports that the FIB is derived from the RIB.

2. Cisco Systems, Inc. (2013). Troubleshooting Cisco Express Forwarding (CEF)-Related Drops.

Section: "CEF Components"

Content: This guide states, "The routing table (RIB) is handled by the route-processor, and it is where all IP routing information is stored... The route-processor computes the best route for a destination and installs this route in the Forwarding Information Base (FIB)." This directly confirms that the RIB contains all routes, from which the best is selected for the FIB.

3. Medhi, D., & Ramasamy, K. (2017). Network Routing: Algorithms, Protocols, and Architectures (2nd ed.). Morgan Kaufmann.

Chapter 6: Router Architectures, Section 6.2.2: Forwarding Engine

Content: University-level networking textbooks describe the separation of the control plane and data plane. The control plane runs routing protocols and builds the Routing Information Base (RIB). The data plane uses a Forwarding Information Base (FIB), which is a copy of the forwarding entries from the RIB, optimized for fast lookups. This academic source validates the relationship where the RIB holds potentially multiple routes and the FIB holds the selected, optimal forwarding paths.

Question: 17

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.

Correct Answer:

C

Explanation:

Cisco Express Forwarding (CEF) is a high-performance packet-forwarding mechanism. It pre-populates a Forwarding Information Base (FIB) from the routing table and an adjacency table from Layer 2 information (like the ARP cache). When a packet arrives, CEF performs a highly efficient lookup in these tables to make a forwarding decision, minimizing CPU involvement. In contrast, process switching is the slowest method, where each packet is "punted" to the main CPU. The CPU must then perform a full routing table lookup and process the packet individually, which is a highly intensive and slow operation.

Why Incorrect Options are Wrong:

A: This is reversed. Process switching is the fallback mechanism for packets that are too complex for CEF or require special handling.

B: This is incorrect. CEF is the fastest switching method, while process switching is the slowest.

D: Process switching is far more CPU-intensive than CEF because the main processor must handle every packet individually.

References:

1. Cisco. (n.d.). Cisco Express Forwarding (CEF) Overview. Cisco Technology White Paper. Retrieved from Cisco.com. In the "CEF Operation" section, it states, "CEF uses a Forwarding Information Base (FIB) and an adjacency table to make packet-forwarding decisions... When a network device receives a packet, it searches the FIB for a matching entry."

2. Cisco. (n.d.). Troubleshooting High CPU Utilization on Cisco Routers. Cisco Technical Support & Documentation. Retrieved from Cisco.com. The document explains, "Process switching means that the router cannot use a faster switching method... The packet is punted to the process level, and the CPU has to do all the work." This confirms process switching is CPU-intensive.

3. Graziani, R., & Vinit, (2020). CCNP and CCIE Enterprise Core ENCOR 350-401 Official Cert Guide. Cisco Press. Chapter 5, "Advanced IP Forwarding and Routing," section "Packet-Forwarding Mechanisms." The text explicitly states, "With process switching, the router processor must get involved with every packet... This process is very slow and CPU-intensive." It contrasts this with CEF, which "builds its own tables... the Forwarding Information Base (FIB) and the adjacency table."

Question: 18

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

Correct Answer:

B, E

Explanation:

Quality of Service (QoS) is a suite of technologies used to manage network traffic and ensure the performance of critical applications during periods of network congestion. The primary problems that QoS addresses are packet loss, delay, and jitter. Excess jitter, the variation in packet delay, severely degrades real-time applications like voice and video. Bandwidth-related packet loss occurs when network links are congested and router/switch buffers overflow, forcing the device to drop packets. QoS mechanisms, such as queuing, scheduling, and classification, are implemented specifically to prioritize sensitive traffic, minimize jitter, and prevent packet loss for critical flows by managing bandwidth contention.

Why Incorrect Options are Wrong:

A: port flapping: This is a physical or data-link layer instability issue, typically caused by faulty hardware or configuration mismatches, not a traffic congestion problem that QoS resolves.

C: misrouted network packets: This is a Layer 3 routing protocol or configuration error. QoS manages traffic along a given path; it does not correct the routing path itself.

D: duplicate IP addresses: This is an IP address assignment or configuration issue (e.g., DHCP or static addressing error) and is unrelated to traffic prioritization or congestion management.

References:

1. Cisco, "Enterprise QoS Solution Reference Network Design Guide, Version 3.3," Chapter: QoS Overview. This guide states, "Without QoS, packets are processed on a best-effort basis... When congestion occurs, packets are dropped, and the retransmission of these packets adds to the congestion and delay... QoS is a critical, fundamental technology that allows for the successful deployment of converged networks." It explicitly identifies packet loss, delay, and jitter as the primary issues QoS is designed to solve.

2. Cisco, "Campus QoS Design Simplified," White Paper, Page 3. This document explains, "The main goal of QoS is to provide better and more predictable network service by providing dedicated bandwidth, controlling jitter and latency (required by some real-time and interactive traffic), and reducing packet loss." This directly links the need for QoS to controlling jitter and reducing packet loss.

3. Cisco, "Quality of Service (QoS) Networking," Cisco Press, Chapter 1: The Need for QoS. This chapter details the problems in a best-effort network, stating, "When network congestion occurs, routers and switches can do two things: drop packets or queue (delay) them. Both of these actions result in degraded application performance." This confirms that bandwidth-related packet loss (dropping) and jitter (a component of delay) are the core problems addressed by QoS.

Question: 19

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.

Correct Answer:

B

Explanation:

Traffic shaping alleviates network congestion by regulating the rate of traffic sent to the network. When the traffic volume exceeds the configured rate (committed information rate), the shaping mechanism delays the excess packets by placing them in a buffer or queue. These buffered packets are then transmitted later when network capacity is available, effectively smoothing out traffic bursts into a steady, predictable stream. This prevents the sudden spikes in traffic that can overwhelm downstream network devices and cause congestion.

Why Incorrect Options are Wrong:

A: Dropping packets that exceed a certain bitrate is characteristic of traffic policing, not shaping. Shaping delays packets, while policing drops or re-marks them.

C: Packet fragmentation is a Layer 3 function to accommodate smaller Maximum Transmission Unit (MTU) sizes and is not the primary mechanism of traffic shaping.

D: Randomly dropping packets from queues is a congestion avoidance mechanism, such as Weighted Random Early Detection (WRED), not a traffic shaping function.

References:

1. Cisco Systems, Inc., QoS: Congestion Management Configuration Guide, Cisco IOS XE Gibraltar 16.12.x, "Congestion Management Overview". This guide explicitly differentiates shaping from policing: "Shaping implies the existence of a queue and of sufficient memory to buffer delayed packets, while policing does not... Packets that exceed the shaping-configured rate are queued. In contrast, a policer will drop the packets that exceed the policer-configured rate."

2. Cisco Systems, Inc., Quality of Service (QoS) Networking, Cisco Press, ISBN: 1-58705-177-3, Chapter 11, "Shaping and Policing," Section "Shaping." The text states, "Shaping buffers excess traffic packets so that they conform to the desired rate... Shaping smooths traffic."

3. Kurose, J. F., & Ross, K. W. (2017). Computer Networking: A Top-Down Approach (7th ed.). Pearson. Chapter 3, "The Network Core: Data Plane," Section 3.7, "Scheduling." The principles of shaping are described as regulating the average rate and burstiness of data entering the network, which involves buffering when the rate is exceeded.

Question: 20

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

Correct Answer:

D

Explanation:

The primary benefit of a cloud infrastructure that is fundamentally lacking in a traditional on-premises deployment is efficient scalability, often termed "rapid elasticity." Cloud platforms are designed to allow for the automated, on-demand provisioning and de-provisioning of resources (compute, storage, etc.) to match workload demands precisely. This capability minimizes waste and optimizes costs. In contrast, scaling an on-premises environment is a slow, capital-intensive process that requires manual procurement, installation, and configuration of physical hardware, making it inherently inefficient and unable to respond rapidly to dynamic changes.

Why Incorrect Options are Wrong:

A: virtualization: Virtualization is a core enabling technology for both cloud and modern on-premises data centers; it is not an exclusive benefit of the cloud.

B: supported systems: Both deployment models offer vendor-supported systems. On-premises support comes from hardware/software vendors, while cloud support comes from the service provider.

C: storage capacity: Both models provide storage. The key differentiator is the efficiency and speed of scaling that capacity, which is an aspect of scalability.

References:

1. Mell, P., & Grance, T. (2011). The NIST Definition of Cloud Computing (Special Publication 800-145). National Institute of Standards and Technology.

Page 2, Section 2, "Essential Characteristics": The document defines "Rapid elasticity" as a key characteristic of cloud computing. It states, "Capabilities can be elastically provisioned and released... to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited..." This directly contrasts with the fixed, manually scaled nature of on-premises infrastructure.

2. Cisco. (n.d.). What Is Cloud Computing?. Cisco.com.

Section: "Benefits of cloud computing": The official Cisco documentation highlights "Greater elasticity" as a primary benefit. It explains, "Cloud computing lets you scale elastically... it delivers the right amount of IT resources... right when they are needed." This efficiency is contrasted with the on-premises need to over-provision hardware in anticipation of demand spikes.

3. Armbrust, M., Fox, A., Griffith, R., Joseph, D., Katz, R., Konwinski, A., ... & Zaharia, M. (2010). A view of cloud computing. Communications of the ACM, 53(4), 50-58.

Section 2.1, "Elasticity and the Illusion of Infinite Resources": This foundational academic paper on cloud computing states, "Cloud Computing offers the illusion of infinite computing resources available on demand... This elasticity of resources, without paying a premium for large scale, is unprecedented in the history of IT." This underscores that efficient, on-demand scalability is a defining and unique advantage over traditional on-premises models.

Question: 21

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

Correct Answer:

D

Explanation:

In the Cisco SD-Access architecture, the overlay network is the virtualized network built on top of the physical underlay. It is responsible for creating logical Layer 2 and Layer 3 networks, known as Virtual Networks (VNs). The overlay uses a combination of protocols, primarily VXLAN (Virtual Extensible LAN) for the data plane to create Layer 2 segments over the Layer 3 underlay, and LISP (Locator/ID Separation Protocol) for the control plane. Each VN is an isolated partition, typically mapped to a VRF (VPN Routing/Forwarding) instance for Layer 3 segmentation, providing multi-tenancy and policy enforcement independent of the physical network topology.

Why Incorrect Options are Wrong:

A: underlay network: The underlay is the physical network of routers and switches that provides IP connectivity and transport for the overlay. It does not create the logical networks.

B: VPN routing/forwarding: VRF is a technology used within the overlay to provide Layer 3 segmentation. However, "overlay network" is the comprehensive architectural term that includes both L2 and L3 virtualization.

C: easy virtual network: EVN is a simplified VRF-Lite virtualization technology. Cisco SD-Access uses a more advanced architecture based on VXLAN, LISP, and standard VRFs, not EVN.

References:

1. Cisco. (2021). Cisco SD-Access Solution Design Guide (CVD). "The overlay network is the virtualized network that is built on top of the physical underlay. In the SD-Access solution, the overlay network is created through a set of network protocols and features (for example, LISP, VXLAN, and Cisco TrustSec). The overlay network provides services such as host mobility, IP portability, segmentation, and other services." (Chapter 2: SD-Access Architecture, Section: Overlay Network).

2. Cisco. (2023). SD-Access for Distributed Campus Design Guide. "The SD-Access fabric is composed of a physical underlay network and a virtual overlay network. The underlay provides the physical connectivity, while the overlay provides the logical connectivity, segmentation, and services." (Chapter 2: SD-Access Fabric Architecture, Section: Underlay and Overlay Networks).

3. Avramov, R., & Edgeworth, R. (2020). CCNP and CCIE Enterprise Core ENCOR 350-401 Official Cert Guide. Cisco Press. In the chapter on SD-Access, the text explains that the overlay network is responsible for virtualization and segmentation, creating virtual networks (VNs) that provide both Layer 2 and Layer 3 services using VXLAN and LISP. (Chapter 26: Cisco SD-Access Solution).

Question: 22

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

onboard WAN Edge nodes into the Cisco SD-WAN fabric

gather telemetry data from WAN Edge routers

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

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

Correct Answer:

C

Explanation:

The vSmart controller is the centralized control plane "brain" of the Cisco SD-WAN fabric. Its primary responsibility is to establish Overlay Management Protocol (OMP) adjacencies with all WAN Edge routers. It receives routing information from these routers, applies centralized control policies (e.g., for traffic engineering, topology) and data policies (e.g., application-aware routing), and then advertises the resulting calculated routing paths and policies back to the WAN Edge routers. This distribution of policy and routing information directly governs how data plane traffic is forwarded across the overlay network.

Why Incorrect Options are Wrong:

A: Onboarding and authenticating WAN Edge nodes into the fabric is the primary function of the vBond orchestrator, which acts as the initial point of contact.

B: Gathering telemetry data, statistics, and logs from WAN Edge routers for monitoring and analytics is a management plane function performed by the vManage NMS.

D: Handling, maintaining, and pushing configuration templates to nodes are management plane tasks performed by the vManage NMS, not the vSmart controller.

References:

1. Cisco SD-WAN Design Guide, "SD-WAN Controller and WAN Edge Platform Selection" section, "Controller Roles" subsection. The guide states, "The vSmart controller is the centralized brain of the solution. It is responsible for the centralized control plane of the network... it runs control plane policies, such as service chaining, traffic engineering, and

per-VPN topology. It advertises routing, security, and policy information to the WAN Edge routers."

2. Cisco SD-WAN End-to-End Deployment Guide, "Chapter: Cisco SD-WAN Overlay Network Bring-Up," "Cisco SD-WAN Components" section. This document describes the vSmart controller's role: "The Cisco vSmart Controller is the centralized brain of the Cisco SD-WAN solution that is responsible for the centralized control plane of the network. It distributes control and data policies to be enforced by the WAN Edge routers."

3. Cisco SD-WAN Getting Started Guide, "Chapter: Cisco SD-WAN Architecture," "Cisco vSmart Controller" section. This guide specifies, "The Cisco vSmart Controller... is responsible for enforcing the policies that you configure through Cisco vManage. These policies govern the data and control traffic in the Cisco SD-WAN overlay network."

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.

Correct Answer:

C, D

Explanation:

The Routing Information Base (RIB) is a control plane construct that contains all routing information learned from various routing protocols and static configurations. The router's processor uses the RIB to determine the best path to each destination.

The Forwarding Information Base (FIB) is a data plane structure derived from the RIIt contains only the best, loop-free routes, optimized for high-speed packet forwarding. The FIB is used by the forwarding engine (e.g., ASICs) to make immediate decisions on where to send incoming packets, specifying the egress interface and next-hop information.

Why Incorrect Options are Wrong:

A: This option incorrectly reverses the roles; the RIB is the prefix database, and the FIB is used for forwarding decisions like choosing an egress interface.

B: This option incorrectly states the derivation. The RIB is built by the control plane, not derived from the FIB, which is a data plane component.

E: This option incorrectly reverses the derivation process. The FIB is derived from the RIB, not the other way around.

References:

1. Cisco Systems, "Cisco Express Forwarding Overview": This document states, "Information from the IP routing table (also called the Routing Information Base, or RIB) is

used to populate the FIB." It further clarifies that the FIB is used to make forwarding decisions, which supports the distinction between the RIB as a database and the FIB as a forwarding tool. (Reference: Cisco.com, Document ID: 13733)

2. Cisco Systems, "Troubleshooting Cisco Express Forwarding (CEF) - Related Drops": This guide explains the architecture: "The control plane is responsible for building and maintaining the Routing Information Base (RIB) and Forwarding Information Base (FIB)... The data plane is responsible for forwarding packets based on the FIB table created by the control plane." This directly supports that the RIB is a control plane function and the FIB is derived from it for data plane use. (Reference: Cisco.com, Document ID: 116308, "CEF Architecture" section)

3. Medhi, D., & Ramasamy, K. (2017). Network Routing: Algorithms, Protocols, and Architectures (2nd ed.). Morgan Kaufmann. Chapter 5, Section 5.2, "Router Architecture," describes the separation of the control plane and data plane. It explains that the control plane runs routing protocols to build the RIB, and the forwarding plane uses a forwarding table (FIB) derived from the RIB for packet switching. This academic source confirms the relationship described in options C and D.

Question: 24

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

A: TCP

B: OMP

C: UDP

D: BGP

Correct Answer:

B

Explanation:

The Overlay Management Protocol (OMP) is the specific control plane protocol used within the Cisco SD-WAN fabric. It operates between the vSmart controllers and the WAN Edge routers (cEdge/vEdge), as well as between the vSmart controllers themselves. OMP is responsible for distributing routing information (OMP routes, TLOC routes, service routes), security policies, and data plane security keys across the overlay network. This protocol runs inside secure DTLS or TLS tunnels, which are established between the network devices and the controllers.

Why Incorrect Options are Wrong:

A: TCP: TCP is the transport layer protocol over which the secure TLS session for OMP runs, but it is not the control plane protocol itself.

C: UDP: UDP is a transport layer protocol. While DTLS (which runs over UDP) can be used for control connections, UDP itself is not the control plane protocol.

D: BGP: BGP is a standard routing protocol that can be used on the service-side (LAN) of the WAN Edge routers but is not the native control plane protocol between routers and vSmart controllers.

References:

1. Cisco Systems, Inc. (2023). Cisco SD-WAN Design Guide. "Control Plane" section. Retrieved from Cisco Design Zone. This guide states, "The Overlay Management Protocol (OMP) is the control protocol that runs between the Cisco vSmart Controllers and the WAN

Edge routers where routing, policy, and management information is exchanged over a secure connection."

2. Cisco Systems, Inc. (2023). Cisco SD-WAN Overlay Management Protocol (OMP) Configuration Guide, Cisco IOS XE Release 17.x. "Overlay Management Protocol Overview" chapter. This document defines OMP: "The Overlay Management Protocol (OMP) is the protocol responsible for establishing and maintaining the Cisco SD-WAN control plane."

3. Cisco Systems, Inc. (2021). Cisco SD-WAN End-to-End Deployment Guide. "Control Plane Connectivity" section. This guide details the process: "WAN Edge routers establish DTLS/TLS connections to the vSmart controller and form OMP peering."

Question: 25

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.

Correct Answer:

A

Explanation:

The fabric edge node functions as the access layer switch in a Cisco SD-Access architecture. Its primary responsibility is to connect wired endpoints, such as user computers, IP phones, and IoT devices, to the network fabric. Upon connection, the edge node identifies and authenticates the endpoint, registers its location (EID-to-RLOC mapping) with the control plane node, and applies appropriate policies. It then encapsulates the endpoint's data traffic into the VXLAN overlay for transport across the fabric.

Why Incorrect Options are Wrong:

B: The edge node encapsulates data traffic using VXLAN. LISP is the control plane protocol used for the endpoint-to-location mapping, not for data plane encapsulation.

C: This describes the function of a fabric border node, which is responsible for connecting the SD-Access fabric to external Layer 3 networks like the WAN or data center.

D: Providing reachability between fabric nodes is the function of the underlay network routing protocol (e.g., IS-IS), not a specific role of the edge node itself.

References:

1. Cisco. (2023). Cisco SD-Access Solution Design Guide (CVD). "Fabric Roles and Terminology" section. This guide states, "Fabric edge node: This is the equivalent of an access layer switch in a traditional campus LAN design. The fabric edge nodes are responsible for connecting and authenticating endpoints and forwarding their traffic to and from the fabric."

2. Odunsi, R., & Edgeworth, (2023). Implementing and Operating Cisco Enterprise Network Core Technologies (ENCOR) v1.2 Official Cert Guide. Cisco Press. Chapter 23, "Cisco SD-Access Solution," section "SD-Access Fabric Roles." The text specifies, "Fabric edge nodes are the switches that endpoints connect to... The edge node is responsible for identifying and authenticating endpoints... and then encapsulating and forwarding traffic for the connected endpoints."

3. Cisco. (2021). SD-Access for Distributed Campus Design Guide. "SD-Access Fabric Building Blocks" section. This document details the roles, stating, "Fabric Edge Nodes: Onboard and provide access for wired user and endpoint devices to the SD-Access fabric."

Question: 26

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

Correct Answer:

B

Explanation:

The core layer in a three-tier hierarchical design serves as the high-speed backbone of the network. Its primary responsibilities are reliability, scalability, and fast convergence. A key best practice is to use redundant Layer 3 point-to-point links between core switches. This design avoids Layer 2 loops and the associated complexities of Spanning Tree Protocol (STP), allowing for the use of routing protocols that provide faster, more predictable convergence and enable load balancing through equal-cost multipathing (ECMP). This keeps the core simple, fast, and highly available.

Why Incorrect Options are Wrong:

A: QoS classification and marking are best performed at the network edge (access or distribution layer), closer to the traffic source, not in the high-speed core.

C: Features like 802.1X, DHCP snooping, and port security are access-layer security functions applied where end-user devices connect to the network.

D: Aggregating access layer devices and providing first-hop redundancy protocols (like HSRP/VRRP) are primary functions of the distribution layer, not the core layer.

References:

1. Rybaczyk, A., et al. (2023). CCNP and CCIE Enterprise Core ENCOR 350-401 Official Cert Guide (2nd ed.). Cisco Press. Chapter 2, "Enterprise Network Design," section "Core Layer," states, "The core layer should be a high-speed, Layer 3-switched environment... For fast convergence, the core should be configured for Layer 3 switching between the core switches and between the core and distribution switches."

2. Cisco. (2022). Campus Wired LAN Design Guide. Cisco Validated Design (CVD). In the "Hierarchical Network Design" section, the role of the core layer is defined as providing "high-speed transport" and "high availability." The guide recommends Layer 3 links for "fast convergence and simplified troubleshooting."

3. Cisco. (2022). Campus Wired LAN Design Guide. Cisco Validated Design (CVD). The "Access Layer" section details the implementation of features like 802.1X, port security, and DHCP snooping. The "Distribution Layer" section describes its role in aggregating access layer switches and hosting first-hop redundancy protocols (FHRPs).

Question: 27

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

A: vBond

B: vEdge

C: vSmart

D: vManage

Correct Answer:

D

Explanation:

Cisco vManage is the centralized Network Management System (NMS) for the Cisco SD-WAN solution. It provides a graphical user interface (GUI) that serves as a single pane of glass for configuration, monitoring, and management of the entire SD-WAN fabric. It is the authoritative entity for all policy and configuration, representing the management plane of the architecture. The other controllers have distinct roles in the orchestration and control planes.

Why Incorrect Options are Wrong:

A: vBond: This is the orchestrator. Its primary role is to authenticate new devices and orchestrate connectivity between vSmart controllers and vEdge routers, not to manage the fabric.

B: vEdge: This is a data plane device (a WAN Edge router) that forwards traffic. It is managed by the control and management planes, but is not a controller itself.

C: vSmart: This is the control plane component. It is responsible for distributing routing information, security policies, and data policies to the vEdge routers via the OMP protocol.

References:

1. Cisco SD-WAN Design Guide, "Cisco SD-WAN Components" section. This document states, "Cisco vManage is the centralized network management system that provides a GUI interface to easily monitor, configure, and maintain all Cisco SD-WAN devices and links... It provides a single pane of glass for day-0, day-1, and day-2 operations."

2. Cisco SD-WAN End-to-End Deployment Guide, "Chapter: Cisco SD-WAN Solution Components", Section: "vManage". This guide describes vManage as the "centralized network management system" and the "single pane of glass for monitoring and management".

3. Cisco SD-WAN Getting Started Guide, "Cisco SD-WAN Overlay Network Components" section. This guide explicitly defines the roles: "Cisco vManage (management plane)", "Cisco vSmart Controller (control plane)", "Cisco vBond Orchestrator (orchestration plane)", and "Cisco vEdge router (data plane)".

Question: 28

Which tag defines the properties to be applied to each specific WLAN?

A: RF tag

B: policy tag

C: AP tag

D: site tag

Correct Answer:

B

Explanation:

In the Cisco Catalyst 9800 tag-based configuration model, the policy tag is the central component that defines the properties applied to a specific WLAN. It acts as a link between a WLAN Profile and a Policy Profile. The WLAN Profile defines the SSID name and Layer 2 security settings, while the Policy Profile defines network policies such as VLAN assignment, Quality of Service (QoS), and Access Control Lists (ACLs). By assigning a policy tag to an Access Point (AP), the controller dictates which WLANs are broadcast and what specific set of properties and policies are applied to clients connecting to that WLAN.

Why Incorrect Options are Wrong:

A: RF tag: This tag applies radio-specific configurations like data rates, power levels, and channel settings to a group of APs, not WLAN properties.

C: AP tag: This is not a valid tag type in the Catalyst 9800 configuration model. The primary tags are Policy, Site, and RF.

D: site tag: This tag applies location-specific configurations, such as the AP Join Profile or FlexConnect profile, which are common to all APs at a physical site.

References:

1. Cisco Systems, Inc., "Wireless Configuration Model," Cisco Catalyst 9800 Series Wireless Controller Configuration Guide, Release IOS XE Cupertino 17.9.x. This guide states, "Policy Tag: This tag is the link between a WLAN profile and a policy profile. It is the policy tag that is assigned to an AP." (Section: Tags in the New Configuration Model).

2. Cisco Systems, Inc., "New Configuration Model for Cisco Catalyst 9800 Series Wireless Controllers," Cisco Catalyst 9800 Series Wireless Controllers Deployment Guide. The document explains, "The policy tag is a pointer to a policy profile, which in turn is a pointer to different profiles such as VLAN, QoS, Session timeout, and so on. The policy tag also points to one or more WLAN profiles." (Section: Policy Tag).

3. Gooley, J., & Hucaby, (2020). CCNP and CCIE Enterprise Core ENCOR 350-401 Official Cert Guide. Cisco Press. In Chapter 27, "Wireless LAN Configuration," the section on the "Cisco Catalyst 9800 WLC Configuration Model" details that the policy tag maps WLAN profiles to policy profiles, which are then assigned to APs.

Question: 29

A customer has several small branches and wants to deploy a Wi-Fi solution with local management using CAPWAP. Which deployment model meets this requirement?

A: local mode

B: SD-Access wireless

C: autonomous

D: Mobility Express

Correct Answer:

D

Explanation:

Cisco Mobility Express is a wireless LAN solution specifically designed for small to medium-sized deployments. In this model, a capable Access Point (AP) is designated as the Master AP, which runs an embedded Wireless LAN Controller (WLC) function. This Master AP manages other subordinate APs at the same site using the Control and Provisioning of Wireless Access Points (CAPWAP) protocol. This architecture provides a controller-based feature set with localized management at the branch, eliminating the need for a dedicated hardware WLC appliance, which perfectly aligns with the customer's requirements for a small branch deployment with local management using CAPWAP.

Why Incorrect Options are Wrong:

A: local mode: This is an AP operational mode that requires a dedicated, typically centralized, WLIt does not provide the "local management" sought for a small branch without a controller appliance.

B: SD-Access wireless: This is a highly centralized, policy-driven architecture for large campus networks, managed by Cisco DNA Center. It is the opposite of a locally managed solution for small branches.

C: autonomous: Autonomous APs are managed individually and locally, but they do not use the CAPWAP protocol to communicate with a controller, which violates a key requirement of the question.

References:

1. Cisco Mobility Express Deployment Guide: "Cisco Mobility Express is a virtual wireless LAN controller function embedded on a Cisco Aironet® 1850 or 1830 Series Access Point... The master AP, with the embedded Cisco Mobility Express Wireless LAN Controller, can manage up to 25 access points... The master AP and the subordinate APs communicate using the Control and Provisioning of Wireless Access Points (CAPWAP) protocol."

Source: Cisco, "Cisco Mobility Express Deployment Guide," Release 8.8, Page 7.

2. Cisco Wireless Controller Configuration Guide: "In the local mode, an access point creates a CAPWAP tunnel to the controller... All traffic is sent to the controller." This describes the dependency on an external controller, contrasting with the Mobility Express model.

Source: Cisco, "Cisco Wireless Controller Configuration Guide, Release 8.5," Chapter: Configuring Access Point Modes, Section: Modes of Operation.

3. Cisco SD-Access Wireless Design and Deployment Guide: "Cisco SD-Access wireless is a solution that integrates the wireless network with the Cisco SD-Access architecture. It is managed by the Cisco DNA Center..." This highlights its centralized management nature.

Source: Cisco, "SD-Access Wireless Design and Deployment Guide," Cisco DNA Center 2.3.3, Chapter: Introduction.

4. Cisco Aironet Access Point FAQ: "What is the difference between a 'lightweight' and an 'autonomous' access point? ... Autonomous Cisco Aironet APs have the Cisco IOS Software on board and can be configured manually, one device at a time... Lightweight APs... use CAPWAP to communicate with a Cisco wireless LAN controller (WLC)." This confirms autonomous APs do not use CAPWAP.

Source: Cisco, "Cisco Aironet Access Point FAQ," Section: General.

Question: 30

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: run on a virtual server and includes its own operating system.

Correct Answer:

B

Explanation:

A Type 1 hypervisor, also known as a bare-metal hypervisor, is a virtualization layer that is installed and runs directly on the physical hardware of a host computer. It does not require a pre-existing host operating system. Instead, the hypervisor itself contains the necessary kernel and management components to control the hardware and manage guest virtual machines. This direct access to hardware resources results in higher efficiency and performance. Prominent examples include VMware ESXi, Microsoft Hyper-V, and Xen.

Why Incorrect Options are Wrong:

A: The dependency on a previously installed operating system is the defining characteristic of a Type 2 (hosted) hypervisor, not a Type 1 hypervisor.

C: A hypervisor runs on a physical server to create virtual servers, not on a virtual server itself. This option also incorrectly describes a Type 2 dependency.

D: A hypervisor's fundamental role is to run on physical hardware, not on a virtual server. This premise is incorrect.

References:

1. Ramel, (2020). CCNP and CCIE Enterprise Core ENCOR 350-401 Official Cert Guide. Cisco Press. Chapter 23, "Virtualization," section "Hypervisors," states, "A Type 1 hypervisor, also called a bare-metal hypervisor, is installed directly on the physical server hardware. The hypervisor is the operating system."

2. VMware. (2023). vSphere Basics. VMware vSphere Documentation. In the "What is a Hypervisor?" section, it defines a Type 1 hypervisor: "A Type 1 hypervisor runs directly on the host's hardware to manage guest operating systems." It explicitly identifies VMware ESXi as a Type 1 hypervisor.

3. Microsoft. (2023). Hyper-V Technology Overview. Microsoft Learn. The documentation describes the Hyper-V architecture where the "Windows hypervisor" is a layer of software that runs directly above the hardware, controlling parent and child partitions, which aligns with the Type 1 definition.

4. Goldberg, R. P. (1974). "Survey of Virtual Machine Research". Computer, 7(6), 34–45. https://doi.org/10.1109/MC.1974.6323581. This foundational academic paper establishes the classification, describing Type 1 (what it calls a "VMM" or Virtual Machine Monitor) as running directly on the bare hardware.