

DUMPS ARENA

Implementing Cisco Enterprise Network Core Technologies (350-401 ENCOR)

Cisco 350-401

Version Demo

Total Demo Questions: 143

Total Premium Questions: 1435

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Architecture	136
Topic 2, Virtualization	74
Topic 3, Infrastructure	588
Topic 4, Network Assurance	87
Topic 5, Security	280
Topic 6, Automation	270
Total	1435

QUESTION NO: 1

Which two parameters are examples of a QoS traffic descriptor? (Choose two.)

- A. DSCP
- B. MPLS EXP bits
- C. packet size
- D. bandwidth
- E. ToS

ANSWER: A B**Explanation:**

In Cisco QoS, a “traffic descriptor” commonly refers to fields/markings in the packet header that describe (classify/identify) the traffic so that devices can apply the intended per-hop behavior (PHB). Two widely used descriptors are IP Differentiated Services Code Point (DSCP) and MPLS Experimental (EXP) bits (now commonly referred to as the MPLS Traffic Class field). DSCP is carried in the IP header’s Differentiated Services field and is the primary Layer 3 marking used for classification and QoS treatment across routed networks. MPLS EXP/Traffic Class bits provide the equivalent marking within an MPLS label stack, allowing QoS behavior to be signaled and enforced across an MPLS core. These descriptors are fundamental to Cisco’s end-to-end QoS model because they enable consistent classification and queuing/priority decisions at each hop without requiring deep packet inspection. In contrast, values like bandwidth are typically a policy/queue parameter (what you allocate), not a descriptor carried in the packet itself. For additional background on DSCP and MPLS QoS markings, see Cisco’s QoS marking guidance and MPLS QoS overview: [Cisco DSCP Values and QoS Marking](#) and [Cisco MPLS QoS \(Traffic Class/EXP\)](#).

QUESTION NO: 2

Which behavior can be expected when the HSRP versions is changed from 1 to 2?

- A. Each HSRP group reinitializes because the virtual MAC address has changed.
- B. No changes occur because version 1 and 2 use the same virtual MAC OUI.
- C. Each HSRP group reinitializes because the multicast address has changed.
- D. No changes occur because the standby router is upgraded before the active router.

ANSWER: A**Explanation:**

Each HSRP group reinitializes because the virtual MAC address has changed. HSRP version 2 introduces a different virtual MAC address format than version 1. When you change the HSRP version on an interface, the group must effectively restart so it can begin using the new version 2 virtual MAC address for that group. This causes the HSRP state machine to

reinitialize and can result in a brief traffic disruption while the active/standby roles are re-established and adjacent devices relearn the gateway MAC address (for example, via ARP refresh and CAM table updates). This behavior is a practical operational consideration when migrating from HSRP version 1 to version 2, especially on VLANs with many hosts, because hosts may need to update their ARP cache to the new virtual MAC. Cisco documents that HSRP version 2 changes the virtual MAC addressing scheme (and expands group number support), which is why the reinitialization is expected during a version change.

References: [Cisco HSRP Overview and Configuration](#), [Cisco IOS XE First Hop Redundancy Protocols \(HSRP\)](#)

QUESTION NO: 3

Refer to the exhibit.

```
R1#show policy-map control-plane
Control Plane

Service-policy input: CoPP

Class-map: telnet_copp (match-all)
 33 packets, 1998 bytes
 5 minute offered rate 0 bps, drop rate 0 bps
Match: access-group 100
 police:
   cir 8000 bps, bc 1500 bytes
   conformed 33 packets, 1998 bytes; actions:
     transmit
   exceeded 0 packets, 0 bytes; actions:
     drop
   conformed 0 bps, exceed 0 bps

Class-map: class-default (match-any)
 59 packets, 5516 bytes
 5 minute offered rate 0 bps, drop rate 0 bps
Match: any
R1#sh access-lists 100
Extended IP access list 100
 10 deny tcp host 10.0.0.5 any eq 22 (13 matches)
 20 permit tcp any any eq 22 (2 matches)
 30 deny tcp host 10.0.0.5 any eq telnet (18 matches)
 40 permit tcp any any eq telnet (31 matches)
R1#
```

43

Which result is achieved by the CoPP configuration?

- A. Traffic that matches entry 10 of ACL 100 is always allowed.
- B. Class-default traffic is dropped.
- C. Traffic that matches entry 10 of ACL 100 is always allowed with a limited CIR.
- D. Traffic that matches entry 10 of ACL 100 is always dropped.

ANSWER: C**Explanation:**

Traffic that matches entry 10 of ACL 100 is always allowed with a limited CIR is achieved because the control-plane service-policy is policing (rate-limiting) the matched control-plane packets rather than unconditionally permitting or denying them. In a typical CoPP design, an access list is used inside a class-map to identify specific control-plane traffic (for example, management traffic sourced from a trusted IP). That class is then referenced in a policy-map applied to the control plane with the `service-policy input` command. When the policy uses `police` with a committed information rate (CIR), packets that conform to the CIR are transmitted (allowed), while packets that exceed the configured rate follow the configured exceed action (commonly drop). This means the matched traffic is permitted, but only up to the configured rate; excess traffic is discarded to protect the CPU and control-plane resources. This is the core purpose of CoPP: protect the control plane by classifying and policing traffic destined to it, not by guaranteeing bandwidth or allowing unlimited traffic.

References: [Cisco Support: Control Plane Policing \(CoPP\) Overview/Configuration](#), [Cisco IOS XE QoS Policing Configuration Guide](#)

QUESTION NO: 4

Which two parameters are examples of a QoS traffic descriptor? (Choose two)

- A. MPLS EXP bits
- B. bandwidth
- C. DSCP
- D. ToS
- E. packet size

ANSWER: A C**Explanation:**

QoS traffic descriptors are fields/markings used to identify and classify traffic into classes so that QoS policies (queuing, policing, shaping, congestion avoidance) can be applied consistently across the network. Common descriptors include Layer 3 and Layer 2/2.5 marking bits carried in packet headers or labels. DSCP is a primary Layer 3 traffic descriptor in the IP header (DiffServ Code Point) and is widely used for classification and per-hop behavior selection in Cisco QoS designs. MPLS EXP bits (now commonly referred to as the MPLS Traffic Class/TC field) serve a similar purpose in MPLS networks, carrying class-of-service information in the MPLS label stack so that QoS treatment can be enforced within an MPLS core. These markings are explicitly designed for traffic identification and QoS behavior mapping, making them canonical examples of traffic descriptors in enterprise and service-provider QoS deployments.

References: [Cisco QoS DSCP values and marking overview](#), [Cisco MPLS QoS \(EXP/TC\) marking concepts](#)

QUESTION NO: 5

Which two southbound interfaces originate from Cisco Catalyst Center (formerly DNA Center) and terminate at fabric underlay switches'? (Choose two.)

- A. ICMP Discovery

- B. UDP67 DHCP
- C. TCP 23 Telnet
- D. UDP6007 NetFlow

- E. UDP 162 SNMP

ANSWER: A E

Explanation:

Cisco Catalyst Center uses southbound management and telemetry protocols to communicate with network devices, including fabric underlay switches in an SD-Access deployment. Two key southbound interfaces that Catalyst Center initiates toward switches are ICMP-based reachability checks (used during discovery and ongoing device health monitoring) and SNMP polling/traps for device inventory, status, and fault/assurance data collection. ICMP is commonly leveraged to validate IP connectivity and basic liveness as part of discovery workflows and continuous monitoring. SNMP is a foundational network management protocol used by Catalyst Center to collect device details and operational state (for example, interface status, CPU/memory, and other MIB-based metrics), which supports inventory and assurance functions. These are standard, well-documented mechanisms in Catalyst Center's device management plane and align with typical enterprise controller-to-device southbound communication patterns. For additional context on Catalyst Center device management and assurance data collection mechanisms, see Cisco's Catalyst Center documentation and the SNMP protocol overview: [Cisco Catalyst Center overview](#) and [Cisco SNMP overview](#).

QUESTION NO: 6

Which configuration enables a device to be configured via NETCONF over SSHv2? A)

```
hostname Device
!
username admin password 0 admin
!
ip domain-name cisco.com
crypto key generate rsa modulus 2048
ip ssh version 2
!
netconf-yang
!
line vty 0 15
login local
```

B)

```
hostname Device
!
aaa new-model
!
username cisco privilege 15 password cisco
!
ip domain-name cisco.com
crypto key generate rsa modulus 2048
ip ssh version 2
!
aaa authentication login default local
aaa authorization exec default local
!
netconf-yang
netconf ssh
```

C)

```
hostname Device
!
aaa new-model
!
username admin privilege 15 password 0 admin
!
ip domain-name cisco.com
crypto key generate rsa modulus 2048
ip ssh version 2
!
netconf-yang
```

D)

```
hostname Device
!
username cisco1 privilege 15 password 0 cisco1
!
ip domain-name cisco.com
crypto key generate rsa modulus 2048
ip ssh version 2
!
netconf ssh
!
line vty 0 15
login local
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

ANSWER: C

Explanation:

The configuration that enables a device to be configured via NETCONF over SSHv2 is the one that explicitly turns on the NETCONF-YANG agent over SSH and ensures SSHv2 prerequisites are met. On Cisco IOS XE, NETCONF over SSH is enabled with the global configuration command **netconf-yang** (older platforms may show variants, but IOS XE uses netconf-yang). For NETCONF to work over SSH, the device must also have SSH enabled, which requires setting an IP domain name, creating a local user (for authentication), generating RSA keys, and forcing SSH version 2. Once **netconf-yang** is enabled, the device listens for NETCONF sessions over SSH (typically on TCP/830) and can be managed by NETCONF clients using YANG models. This combination is the standard Cisco approach for enabling NETCONF management securely over SSHv2 and is consistent with Cisco's IOS XE programmability configuration guidance.

References: [Cisco IOS XE Programmability Configuration Guide \(NETCONF/YANG\)](#), [Cisco SSH configuration overview](#)

QUESTION NO: 7

In a Cisco SD-WAN solution, which two functions are performed by OMP? (Choose two.)

- A. advertisement of network prefixes and their attributes
- B. configuration of control and data policies
- C. gathering of underlay infrastructure data

- D. delivery of crypto keys
- E. segmentation and differentiation of traffic

ANSWER: A D

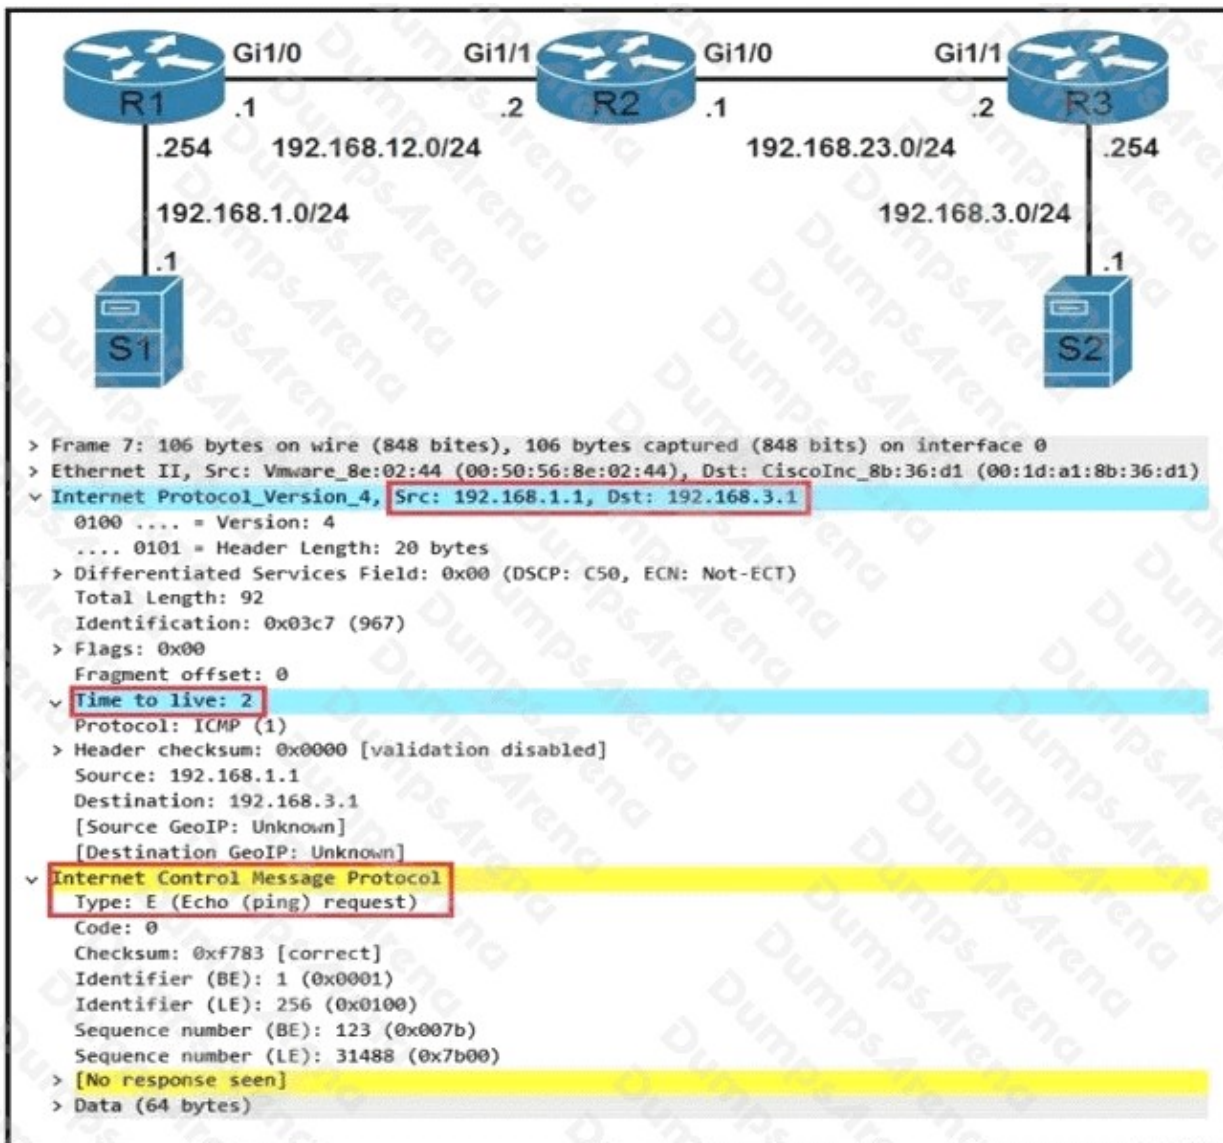
Explanation:

In Cisco SD-WAN, OMP (Overlay Management Protocol) is the control-plane protocol used between WAN Edge routers and the SD-WAN controllers to build and maintain the overlay routing information. A primary OMP function is the advertisement and learning of overlay reachability—specifically, network prefixes (routes) along with key attributes such as TLOC (transport locator) information, site identifiers, tags, and other policy-relevant metadata. This is how WAN Edges learn how to reach remote prefixes across the fabric and which transports/tunnels are available for those routes.

OMP is also responsible for distributing security-related information needed to bring up secure overlay connectivity. In particular, the SD-WAN control plane uses OMP to carry keying material used for IPsec/GRE tunnel security so that WAN Edges can establish encrypted data-plane tunnels across the overlay. This centralized distribution of crypto information is part of how the fabric automates secure connectivity at scale.

These behaviors are described in Cisco SD-WAN control-plane documentation covering OMP route/TLOC propagation and the distribution of security parameters for the overlay. See [Cisco SD-WAN Security Configuration Guide](#) and [Cisco SD-WAN OMP Configuration Guide](#).

QUESTION NO: 8



Refer to the exhibit. While troubleshooting a routing issue, an engineer issues a ping from S1 to S2. Which two actions result from the initial value of the TTL? (Choose two.)

- A. The packet reaches R2, and the TTL expires.
- B. R1 replies with a TTL exceeded message.
- C. The packet reaches R3, and the TTL expires.
- D. R2 replies with a TTL exceeded message.
- E. R3 replies with a TTL exceeded message.
- F. The packet reaches R1, and the TTL expires.

ANSWER: C E

Explanation:

In IPv4, the TTL (Time To Live) field is decremented by 1 at every Layer 3 hop (each router that forwards the packet). When a router receives a packet with TTL = 1, it decrements the TTL to 0 during the forwarding process, discards the packet, and generates an ICMP Time Exceeded message back toward the source. In the exhibit's path from S1 toward S2, the initial TTL value is low enough that the packet is forwarded across multiple routers but runs out of TTL at the third router in the path. That means the echo request makes it as far as R3, where the TTL reaches 0 and the packet is dropped. As a result, R3 is the device that generates the ICMP Time Exceeded response, which is what the source would observe instead of an ICMP Echo Reply from the destination. This behavior is the same mechanism used by traceroute to discover hop-by-hop paths by intentionally manipulating TTL values.

References: <https://www.rfc-editor.org/rfc/rfc791>, <https://www.rfc-editor.org/rfc/rfc792>

QUESTION NO: 9

```
event manager applet Config
event cli pattern "configure terminal"
action 1.0 cli command "enable"
```

Refer to the exhibit. An engineer constructs an EEM applet to prevent anyone from entering configuration mode on a switch. Which snippet is required to complete the EEM applet?

- A. sync yes skip yes
- B. sync no skip yes
- C. sync no skip no
- D. sync yes skip no

ANSWER: B**Explanation:**

To prevent anyone from entering configuration mode, the EEM applet must intercept the CLI command (typically *configure terminal*) and stop it from being executed. This is done with the EEM CLI event detector using the *skip* keyword. Setting *skip yes* tells EEM to "consume" the matched CLI command so it never reaches the IOS CLI parser, effectively blocking entry into config mode. In addition, using *sync no* is appropriate for this use case because the policy is not intended to run synchronously with the user's CLI session to allow the original command to proceed; instead, it should trigger independently and immediately prevent the command execution while the applet can optionally print a message or take other actions. In Cisco EEM, the combination of matching the configuration command and enabling command skipping is the key mechanism that enforces the restriction at the CLI level. For details on the EEM CLI event detector and the *skip* behavior, see Cisco's EEM documentation: [Cisco EEM CLI Event Detector](#) and the broader EEM configuration guide: [Cisco Embedded Event Manager Configuration Guides](#).

QUESTION NO: 10

Refer to the exhibit.

```

Cat3650# show logging
[ ... out ... ]
*Sep 11 19:06:25.595: %PM-4-ERR_DISABLE: channel-misconfig error detected on Po1, putting Gi1/0/2
in err-disable state
*Sep 11 19:06:25.606: %PM-4-ERR_DISABLE: channel-misconfig error detected on Po1, putting Gi1/0/3
in err-disable state
*Sep 11 19:06:25.622: %PM-4-ERR_DISABLE: channel-misconfig error detected on Po1, putting Po1 in
err-disable state

Cat3650# show etherchannel summary
[ ... out ... ]
Group Port-channel Protocol Ports
-----
1 Po1(SD) - Gi1/0/2(D) Gi1/0/3(D)

Cat3650# show interface status err-disabled
Port Name Status Reason Err-disabled Vlans
-----
Gi1/0/2 err-disabled channel-misconfig
Gi1/0/3 err-disabled channel-misconfig
Po1 err-disabled channel-misconfig

```

The administrator troubleshoots an EtherChannel that keeps moving to err-disabled. Which two actions must be taken to resolve the issue? (Choose two.)

- A. Reload the switch to force EtherChannel renegotiation
- B. Ensure that interfaces Gi1/0/2 and Gi1/0/3 connect to the same neighboring switch.
- C. Ensure that the switchport parameters of Port channel1 match the parameters of the port channel on the neighbor switch
- D. Ensure that the corresponding port channel interface on the neighbor switch is named Port-channel1.
- E. Ensure that the neighbor interfaces of Gi1/0/2 and Gi/0/3 are configured as members of the same EtherChannel

ANSWER: B E

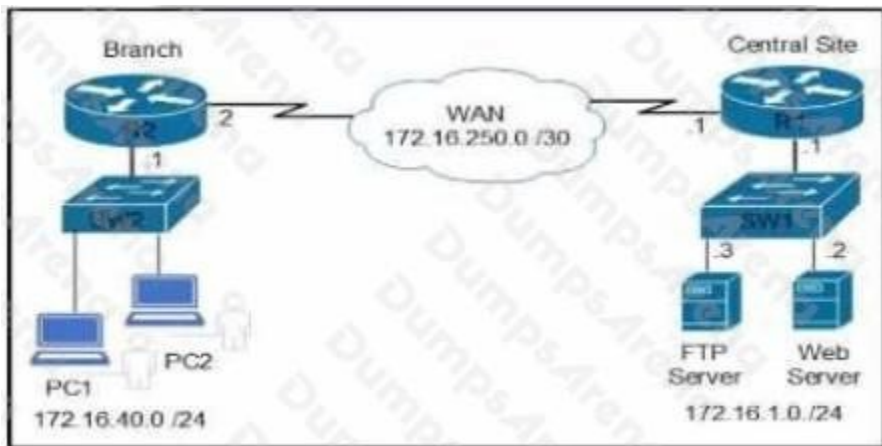
Explanation:

For an EtherChannel to form and remain stable, all member links must terminate on the same logical neighbor and be bundled consistently on both ends. If the physical links in the bundle connect to different neighboring switches (without a multi-chassis EtherChannel technology such as StackWise/StackWise Virtual/vPC), the switch can detect inconsistent aggregation and place ports into an error-disabled state to protect the network from loops and misbundling. Likewise, the far-end interfaces that connect to Gi1/0/2 and Gi1/0/3 must be configured as members of the same EtherChannel (using the same channel-group and compatible negotiation protocol such as LACP), otherwise one side may attempt to bundle while the other treats links as independent, leading to channel misconfiguration and potential err-disable events. Ensuring both links go to the same neighboring switch and ensuring the neighbor ports are in the same EtherChannel addresses the fundamental requirement that EtherChannel is a single logical port made from multiple physical links with consistent bundling on both ends. This aligns with Cisco's EtherChannel/LACP configuration and troubleshooting guidance for avoiding channel misconfiguration conditions.

References: [Cisco EtherChannel Configuration Examples and Guidelines](#), [Cisco Troubleshooting EtherChannel and LACP](#)

QUESTION NO: 11

Refer to the exhibit.



Which command is required to validate that an IP SLA configuration matches the traffic between the branch office and the central site?

- A. R1# show ip sla configuration
- B. R1# show ip sla group schedule 73
- C. R1# show ip route
- D. R1# show ip sla statistics

ANSWER: D

Explanation:

To validate that an IP SLA configuration is working correctly and matches the actual traffic between the branch office and central site, you need to verify the operational statistics of the IP SLA probe. The command **show ip sla statistics** displays real-time performance data including latency, jitter, packet loss, and success/failure rates of configured IP SLA operations. This allows you to confirm that the IP SLA probe is actively monitoring the network path and collecting the expected metrics. Option A (**show ip sla configuration**) only displays the configured parameters of IP SLA operations, not whether they're actually running or collecting data. Option B (**show ip sla group schedule**) shows scheduling information for IP SLA operation groups, which isn't relevant for validating traffic matching. Option C (**show ip route**) displays the routing table, which doesn't provide IP SLA operational validation. The statistics command is essential for troubleshooting and validating that your IP SLA configuration is functioning as intended and accurately measuring the network performance between sites. For more information, refer to the [Cisco IOS IP SLA Command Reference](#) and [Cisco IP SLA Configuration Guide](#).

QUESTION NO: 12

Which two methods are used by an AP that is trying to discover a wireless LAN controller? (Choose two.)

- A. Cisco Discovery Protocol neighbor
- B. querying other APs
- C. DHCP Option 43
- D. broadcasting on the local subnet

E. DNS lookup CISCO-DNA-PRIMARY.localdomain

ANSWER: C D

Explanation:

An access point (AP) that is attempting to join a Cisco wireless LAN controller (WLC) uses a defined CAPWAP/LWAPP discovery sequence to learn the controller's management IP address. Two core, commonly tested discovery mechanisms are DHCP Option 43 and local subnet broadcast discovery. With DHCP Option 43, the DHCP server provides the AP with one or more WLC management IP addresses (encoded in vendor-specific format), allowing the AP to directly attempt CAPWAP discovery/join without relying on L2 adjacency. Local subnet broadcasting is also used when the AP and WLC are on the same Layer 2 network; the AP sends a CAPWAP discovery request to the local broadcast domain and any WLC on that subnet can respond with a discovery response. These two methods are standard, widely deployed, and explicitly documented as primary approaches for AP-to-controller discovery in enterprise networks, especially during initial provisioning or after a reset. They are also frequently used together: broadcast discovery for same-subnet deployments and DHCP Option 43 for routed deployments where the controller is not on the AP's local VLAN.

References: [Cisco TAC: Lightweight AP not joining WLC \(discovery methods\)](#), [Cisco: DHCP Option 43 for WLC discovery](#)

QUESTION NO: 13

Which three resources must the hypervisor make available to the virtual machines? (Choose three.)

- A. Memory
- B. bandwidth
- C. IP address
- D. Processor
- E. storage
- F. secure access

ANSWER: A D E

Explanation:

A hypervisor's core job is to virtualize and allocate the underlying physical host resources so that each virtual machine can run an operating system and applications as if it had its own hardware. The fundamental resources a hypervisor must present to VMs are compute (CPU), memory (RAM), and storage. CPU is provided by scheduling vCPUs onto physical CPU cores/threads, controlling execution time and isolation between guests. Memory is provided by carving out RAM for each VM (often with techniques like overcommit, ballooning, or page sharing depending on platform) so the guest OS can manage processes and caching. Storage is provided by mapping virtual disks (VMDKs, QCOW2, etc.) to physical disks, SAN/NAS LUNs, or other datastores so the VM has persistent block storage for its filesystem and data. While networking is also virtualized, "bandwidth" and "IP address" are not mandatory hypervisor resources in the same way—IP addressing is typically handled inside the guest OS or by external network services. These three resources align with standard hypervisor architecture and virtualization fundamentals used across enterprise platforms like VMware ESXi and KVM. See [VMware Hypervisor glossary](#) and [Red Hat: What is a hypervisor?](#).

QUESTION NO: 14

Refer to the Exhibit.

```
Device# configure terminal
Device(config)# netconf ssh acl 1
Device(config)# netconf lock-time 100
Device(config)# netconf max-sessions 1
Device(config)# netconf max-message 10
```

Refer to the exhibit A network engineer must configure NETCONF After creating the configuration, the engineer gets output from the command `show line` but not from `show running-config`. Which command completes the configuration?

- A. Device(config)# netconf max-message 1000
- B. Device(config)# netconf max-sessions 100
- C. Device(config)# netconf lock-time 500
- D. Device(config)# no netconf ssh aci 1

ANSWER: D

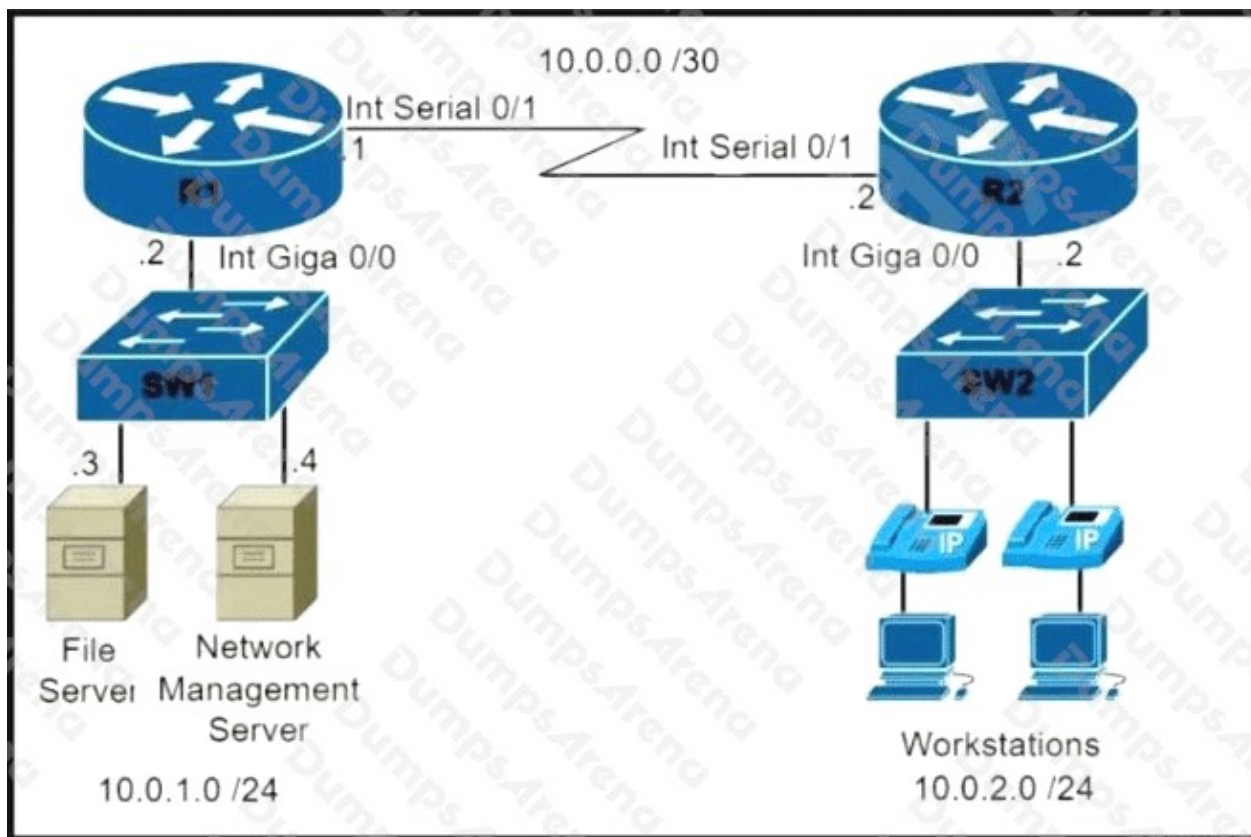
Explanation:

To make NETCONF over SSH work on Cisco IOS XE, the device must have NETCONF enabled and it must be tied to an SSH access control list (ACL) that permits the management source(s). When NETCONF is configured with an SSH ACL that does not exist or does not permit the client, the NETCONF subsystem won't be reachable even though the VTY lines may look correct in `show line`. In that situation, you can end up with no effective NETCONF access and no meaningful NETCONF-related operational state reflected as expected in `show running-config` for the intended behavior.

The command `no netconf ssh aci 1` removes the binding to an SSH ACL (in this case, ACL 1). This "completes" the configuration by eliminating the incorrect/overly restrictive ACL association so NETCONF over SSH can be accessed (assuming SSH itself is otherwise configured correctly). After removing the ACL reference, NETCONF uses the standard SSH reachability rules and VTY access controls, restoring expected NETCONF connectivity.

References: [Cisco IOS XE NETCONF Configuration Guide](#), [Cisco IOS XE Programmability \(NETCONF/RESTCONF\)](#)

QUESTION NO: 15



An engineer must configure and validate a CoPP policy that allows the network management server to monitor router R1 via SNMP while protecting the control plane. Which two commands or command sets must be used? (Choose two.)

- A. `access-list 150 permit udp 10.0.1.4 0.0.0.0 host 10.0.1.2 eq snmp`
`access-list 150 permit udp 10.0.1.4 0.0.0.0 eq snmp host 10.0.1.2`
`class-map match-all CoPP-management match access-group 150`
`policy-map CoPP-policy class CoPP-management`
`police 8000 conform-action transmit exceed-action transmit violate-action drop`
`control-plane`
`Service-policy input CoPP-policy`
- B. `show ip interface brief`
- C. `show quality-of-service-profile`
- D. `access-list 150 permit udp 10.0.1.4 0.0.0.0 host 10.0.1.2 eq snmp`
`class-map match-all CoPP-management match access-group 150`
`policy-map CoPP-policy class CoPP-management`
`police 8000 conform-action transmit exceed-action transmit violate-action transmit`
`control-plane`
`Service-policy input CoPP-policy`
- E. `show policy-map control-plane`

ANSWER: A E

Explanation:

To allow SNMP monitoring while protecting the control plane, you must both (1) build and apply a CoPP policy that matches only the desired SNMP traffic destined to the router's control plane and (2) verify that the policy is actually attached and counting/ policing packets. The configuration set that defines an ACL matching UDP SNMP from the management server to the router, ties it to a class-map, applies policing in a policy-map, and then applies that policy under the control-plane with an input service-policy is required to implement CoPP for SNMP. This is the standard Modular QoS CLI (MQC) workflow used by Cisco for CoPP: classify with ACL/class-map, enforce with policy-map (police), and attach to the control-plane.

After configuration, validation is done with the command that displays the control-plane policy attachment and its counters/actions. This lets you confirm that SNMP packets are matching the intended class and being transmitted/dropped according to the policer, which is essential to prove both reachability for monitoring and protection against excess traffic. Cisco documents this CoPP approach and the verification command in its CoPP/MQC guidance. References: [Cisco CoPP configuration example](#), [Cisco MQC policing configuration](#).

QUESTION NO: 16

Refer to the Exhibit.

R1	R2
<pre>key chain cisco123 key 1 key-string Cisco123!</pre>	<pre>key chain cisco123 key 1 key-string cisco123!</pre>
<pre>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 (vl 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)</pre>	<pre>Ethernet0/0 - Group 10 State is Active 17 state changes, last state change 00:02:17 Virtual IP address is 192.168.0.1 Active virtual MAC address is 0000.0c07.ac0a Local virtual MAC address is 0000.0c07.ac0a (vl 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)</pre>

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: B

Explanation:

The action that resolves the issue is to configure matching key-strings. In first-hop redundancy protocols such as HSRP, when authentication is enabled, both routers participating in the same standby group must use the exact same authentication value. The key string is case-sensitive, so even a small difference like an uppercase versus lowercase character results in an authentication mismatch. When authentication fails, the routers will not properly form the expected active/standby relationship for the group, leading to unexpected standby status behavior. Correcting the configuration so that both devices use identical key-strings (same characters, same case, same type) restores successful authentication and allows the routers to negotiate roles normally. This aligns with Cisco guidance that authentication parameters must match between peers and that key strings are treated as case-sensitive text. Once the key-strings are identical, the redundancy pair can correctly elect an active router and a standby router for the configured group.

References: [Cisco HSRP Configuration and Troubleshooting](#), [Cisco IOS XE HSRP Configuration Guide](#)

QUESTION NO: 17

In a Cisco SD-Access environment, which function is performed by the border node?

- A. Connect uteri and devices to the fabric domain.
- B. Group endpoints into IP pools.
- C. Provide reachability information to fabric endpoints.
- D. Provide connectivity to traditional layer 3 networks.

ANSWER: D

Explanation:

Provide connectivity to traditional layer 3 networks.

In Cisco SD-Access, the border node is the fabric role that connects the SD-Access fabric to external networks (often called “outside” or “non-fabric” networks). It is the point where traffic enters or exits the fabric domain toward traditional Layer 3 routing domains such as the enterprise core, data center, WAN/Internet edge, or other VRFs. The border node participates in the fabric control-plane (LISP) and data-plane (VXLAN) internally, while also running standard routing protocols (for example, OSPF, BGP, or static routing) toward the external network. This allows endpoints inside the fabric to reach destinations outside the fabric and enables external networks to reach fabric subnets, while still preserving SD-Access segmentation (VN/VRF) and policy constructs. Border nodes can be configured as internal or external borders depending on whether they connect to other fabric sites or to non-fabric networks, but the core function remains providing Layer 3 connectivity between the fabric and traditional routed domains.

References: [Cisco Software-Defined Access \(SD-Access\) Overview](#), [Cisco SD-Access Design Guide \(Campus Fabric roles\)](#)

QUESTION NO: 18

Which two results occur if Cisco DNA center loses connectivity to devices in the SD-ACCESS fabric? (Choose two)

- A. All devices reload after detecting loss of connection to Cisco DNA Center
- B. Already connected users are unaffected, but new users cannot connect
- C. User connectivity is unaffected
- D. Cisco DNA Center is unable to collect monitoring data in Assurance
- E. Users lose connectivity

ANSWER: B C D

Explanation:

If Cisco DNA Center loses connectivity to SD-Access fabric devices, the fabric continues forwarding traffic based on the control-plane and data-plane state already programmed on the network devices. In other words, day-2 operations and policy

enforcement that are already deployed keep working, so user traffic does not inherently drop simply because the management/controller platform is unreachable. This is why *User connectivity is unaffected* is a valid outcome in this scenario.

However, Cisco DNA Center's management functions depend on telemetry and device communication. When it cannot reach the fabric nodes, it cannot receive streaming telemetry, syslog/SNMP/NETCONF data, or other inputs used by Assurance to build health scores, client insights, and issue detection. As a result, *Cisco DNA Center is unable to collect monitoring data in Assurance* is also an expected outcome. The key idea is separation of planes: loss of the management plane impacts visibility/automation, not the existing forwarding behavior of the fabric.

References: [Cisco DNA Center product page](#), [Cisco SD-Access overview](#)

QUESTION NO: 19

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

- A. It enforces the use of a specific encoding format for NETCONF.
- B. It collects statistical constraint analysis information.
- C. It enables multiple leaf statements to exist within a leaf list.
- D. It enforces configuration semantics.
- E. It enforces configuration constraints.

ANSWER: D E**Explanation:**

YANG is a data modeling language used with NETCONF/RESTCONF to describe the structure of configuration and operational state data, along with rules that make that data valid. A key benefit is that it enforces configuration constraints: YANG can define types, ranges, patterns, mandatory nodes, uniqueness, and "must/when" expressions so invalid configurations can be rejected before they are applied. Another major benefit is that it enforces configuration semantics by providing a clear, machine-readable schema that defines what each configuration element means, how it relates to other elements (containers/lists/leaves), and whether it is configuration or state data. This consistent schema enables tooling, automation, and validation across vendors and platforms, and it allows clients to reliably understand device capabilities via published modules. These properties are central to why YANG is used in model-driven programmability and why it improves interoperability and correctness in network automation workflows.

References: <https://datatracker.ietf.org/doc/html/rfc7950>, <https://datatracker.ietf.org/doc/html/rfc6241>

QUESTION NO: 20

Which three resources must the hypervisor make available to the virtual machines? (Choose three)

- A. memory
- B. bandwidth
- C. IP address

- D. processor
- E. storage
- F. secure access

ANSWER: A D E

Explanation:

A hypervisor's core job is to virtualize and allocate the underlying host hardware so each virtual machine can run an operating system and applications as if it had its own physical server. The fundamental resources it must present to VMs are compute, memory, and persistent disk. Compute is provided by scheduling and virtualizing CPU instructions so each VM gets access to one or more virtual CPUs backed by the host's physical processors. Memory is carved out and managed so each VM has RAM available, with the hypervisor enforcing isolation and handling techniques like overcommit where supported. Storage is exposed as virtual disks (for example, VMDKs or virtual hard disks) backed by local disks, SAN/NAS, or other datastores, enabling the VM to boot and persist data. These three resources—processor, memory, and storage—are the baseline requirements for a VM to function. Networking-related items like bandwidth and IP addressing are important for connectivity, but IP addresses are typically configured inside the guest OS (or via external IPAM/DHCP), and bandwidth is a characteristic of the virtual/physical NICs rather than a mandatory "resource" in the same foundational sense. See: [VMware Hypervisor definition](#) and [Microsoft Hyper-V technology overview](#).

QUESTION NO: 21

Which of the following are examples of Type 2 hypervisors? (Choose three.)

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

ANSWER: B C E

Explanation:

Type 2 hypervisors (hosted hypervisors) run as applications on top of a conventional host operating system, relying on that OS for device drivers and many hardware interactions. In practice, you install them like any other desktop/server program, then create and run virtual machines within that application. Oracle VirtualBox is a classic hosted hypervisor: it installs on Windows, macOS, or Linux and provides a user-space virtualization stack for running guest VMs. Microsoft Virtual PC is another hosted hypervisor product that runs on top of a Windows host OS and provides VM capabilities as an application. Oracle Solaris Zones, while technically OS-level virtualization (containers) rather than a traditional hardware-virtualizing hypervisor, is commonly grouped in exam contexts as a hosted virtualization technology because it operates within the Solaris OS environment rather than booting directly on bare metal. These examples align with the key Type 2 characteristic: virtualization provided from within an existing OS rather than replacing the OS as the primary platform. See: <https://www.virtualbox.org/wiki/VirtualBox> and <https://learn.microsoft.com/en-us/windows-server/virtualization/hyper-v/hyper-v-on-windows-server>.

QUESTION NO: 22

Which router is elected the IGMP Querier when more than one router is in the same LAN segment?

- A. The router with the shortest uptime
- B. The router with the lowest IP address
- C. The router with the highest IP address
- D. The router with the longest uptime

ANSWER: B**Explanation:**

In IGMP, hosts rely on an IGMP querier on the local subnet to periodically send General Queries so receivers can report multicast group membership. When multiple multicast-capable routers exist on the same LAN segment, they run an IGMP querier election to ensure only one device actively sends these queries. The election rule is straightforward: the router with the lowest IP address on that subnet becomes the IGMP querier. If the current querier stops sending queries (for example, it fails or is removed), another router will take over based on the same lowest-IP rule after the appropriate timers expire. This behavior is consistent across common IGMP versions used in enterprise networks (notably IGMPv2 and IGMPv3) and is important for stable multicast operation because it prevents duplicate query traffic and ensures predictable control-plane behavior on the segment. Cisco documents this querier election behavior as part of its IGMP operation and configuration guidance for multicast-enabled interfaces.

References: [Cisco – IGMP/IGMP Snooping Overview and Behavior](#), [RFC 2236 \(IGMPv2\) – Querier Election](#)

QUESTION NO: 23

What are two common sources of interference for Wi-Fi networks? (Choose two.)

- A. rogue AP
- B. conventional oven
- C. fire alarm
- D. LED lights
- E. radar

ANSWER: A E**Explanation:**

Two common sources of Wi-Fi interference are a rogue AP and radar. A rogue AP (or any unauthorized/neighbor access point) can create co-channel or adjacent-channel interference by transmitting on the same or overlapping channels, raising the noise floor and increasing contention, which reduces throughput and increases retries. This is especially common in dense deployments where multiple APs share limited non-overlapping channels (for example, in 2.4 GHz). Radar is a well-known interferer in parts of the 5 GHz band because Wi-Fi must share spectrum with incumbent radar systems. In DFS (Dynamic Frequency Selection) channels, WLANs are required to detect radar signatures and vacate the channel to avoid interfering with radar operations; from the Wi-Fi perspective, radar events can look like interference and can also trigger disruptive channel changes. These behaviors and requirements are foundational to enterprise WLAN design and

troubleshooting, making both rogue APs and radar classic, frequently tested interference sources. See Cisco's overview of RF interference concepts and DFS behavior in enterprise WLANs and the regulatory DFS requirements described in Wi-Fi/802.11 references: [Cisco: Wireless RF Interference Overview](#) and [Dynamic frequency selection \(DFS\)](#).

QUESTION NO: 24

```
interface Vlan10
 ip vrf forwarding Clients
 ip address 192.168.1.1, 255.255.255.0
 !
interface Vlan20
 ip vrf forwarding Servers
 ip address 172.16.1.1 255.255.255.0
 !
interface Vlan30
 ip vrf forwarding Printers
 ip address 10.1.1.1 255.255.255.0
-- output omitted for brevity --
router eigrp 1
 10.0.0.0
 172.16.0.0
 192.168.1.0
```

Refer to the exhibit. An engineer attempts to configure a router on a stick to route packets between Clients, Servers, and Printers; however, initial tests show that this configuration is not working. Which command set resolves this issue?

- A.** interface Vlan10 no ip vrf forwarding Clients
!
interface Vlan20 no ip vrf forwarding Servers
!
interface Vlan30
no ip vrf forwarding Printers
- B.** router eigrp 1 network 10.0.0.0 255.255.255.0 network 172.16.0.0 255.255.255.0 network 192.168.1.0 255.255.255.0
- C.** interface Vlan10 no ip vrf forwarding Clients ip address 192.168.1.2, 255.255.255.0
!
interface Vlan20
no ip vrf forwarding Servers ip address 172.16.1.2 255.255.255.0 !
interface Vlan30
no ip vrf forwarding Printers ip address 10.1.1.2 255.255.255.0
- D.** router eigrp 1 network 10.0.0.0 255.0.0.0 network 172.16.0.0 255.255.0.0 network 192.168.1.0 255.255.0.0

ANSWER: C**Explanation:**

For a router-on-a-stick design to work, the router must provide Layer 3 default-gateway services for each VLAN by having a routed interface per VLAN (typically subinterfaces on a trunk) with the correct IP addressing. If the router interfaces (or SVIs, depending on the platform/feature set shown in the exhibit) are missing IP addresses or are tied to VRFs unintentionally, hosts in different VLANs will not be able to route between Clients, Servers, and Printers because there is no valid L3 interface in the global routing table to perform inter-VLAN routing.

The command set that removes VRF forwarding from each VLAN interface and assigns the correct IP address/mask to each VLAN interface restores normal inter-VLAN routing. Once those VLAN interfaces have IP addresses in the global routing table, they can act as the default gateways for their respective VLANs and route traffic between the three subnets. This aligns with Cisco's router-on-a-stick/inter-VLAN routing requirements: each VLAN needs an L3 interface with an IP address, and the trunk must carry the VLANs to the router.

References: [Cisco Inter-VLAN Routing Configuration Example](#), [Cisco IP Routing: How Routing Works](#)

QUESTION NO: 25

A network engineer is configuring OSPF on a router. The engineer wants to prevent having a route to 177.16.0.0/16 learned via OSPF. In the routing table and configures a prefix list using the command `ip prefix-list OFFICE seq 5 deny 172.16.0.0/16`. Which two identical configuration commands must be applied to accomplish the goal? (Choose two.)

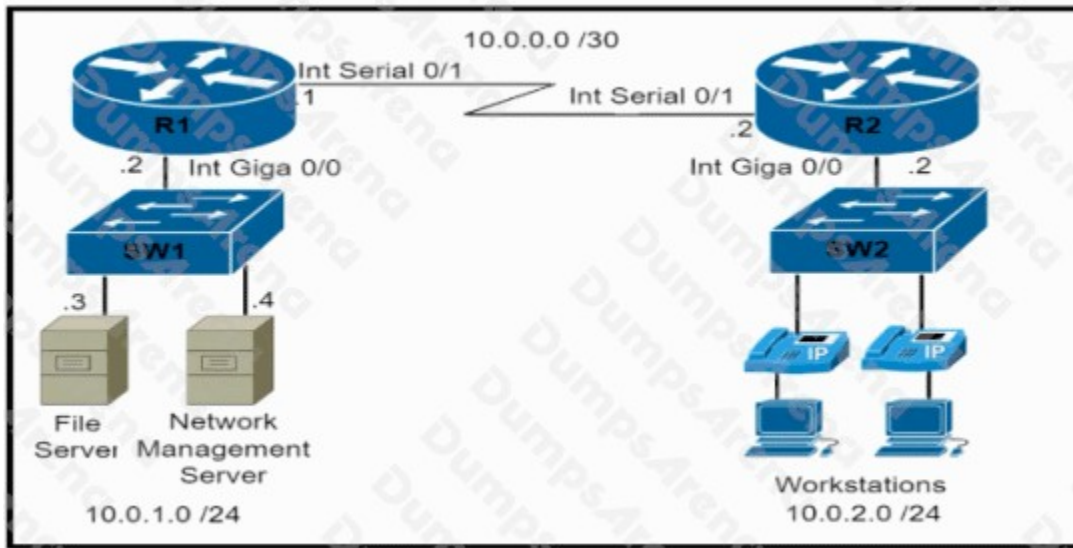
- A. `distribute-list prefix OFFICE in` under the OSPF process
- B. `ip prefix-list OFFICE seq 10 permit 0.0.0.0/0 le 32`
- C. `ip prefix-list OFFICE seq 10 permit 0.0.0.0/0 ge 32`
- D. `distribute-list OFFICE out` under the OSPF process
- E. `distribute-list OFFICE in` under the OSPF process

ANSWER: A B**Explanation:**

To stop a specific OSPF-learned route from being installed in the local routing table, you apply a filtering mechanism to OSPF route reception on the router. The correct approach is to build a prefix list that explicitly denies the unwanted prefix and then permits everything else, because prefix lists (like ACLs) have an implicit "deny all" at the end. That's why adding a catch-all permit such as "`ip prefix-list OFFICE seq 10 permit 0.0.0.0/0 le 32`" is required; it ensures only the denied prefix is blocked while all other OSPF routes are still accepted. Next, you must apply that prefix list to the OSPF process in the inbound direction using "`distribute-list prefix OFFICE in` under the OSPF process", which filters routes as they are learned/installed. This combination accomplishes the goal of preventing the specified prefix from being learned via OSPF while leaving other OSPF routes unaffected. For details on prefix-list syntax and route filtering concepts, see Cisco's prefix list configuration guidance and OSPF distribute-list usage: [Cisco OSPF filtering overview](#) and [Cisco IOS Prefix Lists](#).

QUESTION NO: 26

Refer to the exhibit.



An engineer must configure and validate a CoPP policy that allows the network management server to monitor router R1 via SNMP while protecting the control plane. Which two commands or command sets must be used? (Choose two.)

- show policy-map control-plane
- show quality-of-service-profile
- access-list 150 permit udp 10.0.1.4 0.0.0.0 host 10.0.1.2 eq snmp
- class-map match-all CoPP-management
 - match access-group 150
- policy-map CoPP-policy
 - class CoPP-management
 - police 8000 conform-action transmit exceed-action transmit
 - violate-action transmit
- control-plane
 - Service-policy input CoPP-policy
- show ip interface brief

```
❑ show ip interface brief
☑ access-list 150 permit udp 10.0.1.4 0.0.0.0 host 10.0.1.2 eq snmp
  access-list 150 permit udp 10.0.1.4 0.0.0.0 eq snmp host 10.0.1.2

class-map match-all CoPP-management
  match access-group 150

policy-map CoPP-policy
  class CoPP-management
    police 8000 conform-action transmit exceed-action transmit
    violate-action drop

control-plane
  Service-policy input CoPP-policy
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E
- F. Option F

ANSWER: A F

Explanation:

To allow SNMP monitoring while protecting the router control plane, CoPP must explicitly classify SNMP packets destined to the CPU and then apply a control-plane service-policy that permits (and typically rate-limits) that traffic. The required configuration elements are: an ACL that matches SNMP (UDP/161) and often SNMP traps/informs (UDP/162) from the network management server to the router, a class-map that references that ACL, and a policy-map that applies policing (or at minimum a transmit action) for that class. Finally, the policy must be attached under *control-plane* using *service-policy input* so it actually protects the CPU path. Validation is typically done with *show policy-map control-plane* (and related class statistics) to confirm SNMP packets are hitting the intended class and being permitted/policed as designed. These are the core Cisco IOS XE CoPP building blocks for safely allowing management-plane protocols to the CPU. See Cisco CoPP overview and configuration guidance in the Control Plane Policing documentation: [Cisco CoPP Technote](#) and IOS XE policing/QoS command references: [Cisco IOS XE Policing Configuration Guide](#).

QUESTION NO: 27

Which two steps are required for a complete Cisco DNA Center upgrade? (Choose two.)

- A. golden image selection
- B. automation backup

- C. proxy configuration
- D. application updates
- E. system update

ANSWER: D E

Explanation:

A complete Cisco DNA Center upgrade is made up of two distinct upgrade actions: updating the Cisco DNA Center applications and updating the underlying system software. The application portion updates the DNA Center services and feature components to the target release so the controller's functionality aligns with the new version. The system portion updates the platform/OS layer that the applications run on (including required infrastructure components), ensuring compatibility, stability, and supportability for the upgraded application stack. Cisco's upgrade workflow and guidance treat these as the two required parts of a full upgrade, and both must be performed to be considered "complete," because newer application bundles may depend on newer system components and vice versa. In practice, administrators typically run prechecks, confirm cluster health, and then execute the application update and the system update as directed by the release-specific upgrade path. This two-part approach is central to Cisco DNA Center's lifecycle management and is consistently reflected in Cisco's documentation for upgrading between releases.

References: [Cisco DNA Center Installation and Configuration Guides](#), [Cisco DNA Center Release Notes](#)

QUESTION NO: 28

Refer to the Exhibit.



```
no aaa new-model
username admin privilege 15 secret cisco123
ip http secure-port 445
```

Refer to the exhibit Which command must be applied to complete the configuration and enable RESTCONF?

- A. ip http secure-server
- B. ip http server
- C. ip http secure-port 443
- D. ip http client username restconf

ANSWER: A

Explanation:

To enable RESTCONF on Cisco IOS XE, the device must have the HTTPS server enabled because RESTCONF operates over HTTP(S) and, in Cisco implementations, is typically exposed securely via HTTPS. The command **ip http secure-server** turns on the device's HTTPS server, which is a prerequisite for RESTCONF to accept inbound RESTCONF API calls (for example, to the /restconf endpoint) and to protect credentials and payloads in transit using TLS. In common IOS XE RESTCONF configurations, you also enable the RESTCONF feature itself (for example, with a restconf-related command depending on platform/version) and ensure AAA/user authentication is in place, but the missing piece that "completes the

configuration” in many exam-style snippets is enabling the secure HTTP server. This aligns with Cisco’s guidance that RESTCONF uses HTTPS transport and relies on the device HTTP/HTTPS services for access. Once HTTPS is enabled, RESTCONF clients can securely interact with YANG-modeled resources using standard HTTP methods (GET/POST/PUT/PATCH/DELETE) over TLS.

References: [Cisco IOS XE Programmability Configuration Guide \(RESTCONF\)](#), [Cisco IOS XE HTTP/HTTPS Server Configuration](#)

QUESTION NO: 29

Which two pieces of information are necessary to compute SNR? (Choose two.)

- A. transmit power
- B. noise floor
- C. EIRP
- D. antenna gain
- E. RSSI

ANSWER: B E**Explanation:**

To compute Signal-to-Noise Ratio (SNR) in Wi-Fi/RF terms, you need the received signal level and the received noise level at the same point in space and time. In practice, the received signal level is commonly represented as RSSI (Received Signal Strength Indicator), typically shown in dBm, and the received noise level is represented as the noise floor, also typically shown in dBm. SNR is then calculated as a simple difference: $SNR (dB) = RSSI (dBm) - Noise Floor (dBm)$. Because both values are measured at the receiver, they directly capture real-world effects like path loss, attenuation, interference, and environmental noise. This is why RSSI and noise floor are the necessary inputs for SNR, and why values like transmit power, EIRP, and antenna gain are not required to compute SNR at the receiver (they influence RSSI, but they are not the measurement inputs used in the SNR calculation). For Cisco wireless design and troubleshooting, SNR derived from RSSI and noise floor is a core metric for assessing link quality and expected modulation/coding performance.

References: [Cisco Wireless RF fundamentals](#), [Signal-to-noise ratio \(SNR\) definition](#)

QUESTION NO: 30

43

How does a Type 1 hypervisor function?

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

ANSWER: B**Explanation:**

A Type 1 hypervisor is a “bare-metal” hypervisor, meaning it installs directly on the physical server hardware and does not require a separate, general-purpose host operating system underneath it. In this model, the hypervisor itself provides the core platform services needed to schedule CPU, manage memory, and control I/O for guest virtual machines. Many Type 1 hypervisors also include a minimal, purpose-built management layer (sometimes called a hypervisor OS or hypervisor kernel) that is tightly integrated with the virtualization stack, rather than relying on Windows/Linux as a host OS. This architecture is designed for performance and isolation because the hypervisor has direct access to hardware resources and can enforce strong separation between VMs. Common examples include VMware ESXi and Microsoft Hyper-V in its bare-metal deployment model. This is why the statement “It runs directly on a physical server and includes its own operating system.” best matches how Type 1 hypervisors function in enterprise environments.

References: [VMware Glossary: Bare-metal hypervisor](#), [Microsoft Learn: Hyper-V on Windows Server](#)

QUESTION NO: 31

Which two components are supported by LISP? (Choose two.)

- A. proxy ETR
- B. egress tunnel router
- C. route reflector
- D. HMAC algorithm
- E. spoke

ANSWER: A B**Explanation:**

Locator/ID Separation Protocol (LISP) defines a set of functional components used to separate endpoint identity (EIDs) from routing locators (RLOCs) and to encapsulate traffic between sites. A core device role in LISP is the egress tunnel router, which decapsulates LISP-encapsulated packets arriving from the LISP core and forwards the original packets toward the destination endpoint in the local site. LISP also supports a proxy ETR, which provides reachability between LISP and non-LISP sites by acting as an egress point for traffic destined to non-LISP EIDs (or to assist in interworking scenarios), enabling incremental deployment without requiring every site to run LISP from day one. These roles are part of the standard LISP architecture alongside other elements like ITRs and mapping systems (Map-Server/Map-Resolver). You’ll see these components referenced in Cisco’s LISP configuration and design guidance because they are fundamental to how LISP builds tunnels and resolves EID-to-RLOC mappings for forwarding decisions.

References: [Cisco IOS XE LISP Configuration Guide \(Overview\)](#), [RFC 6830: The Locator/ID Separation Protocol \(LISP\)](#)

QUESTION NO: 32

```
R2#show standby
FastEthernet1/0 - Group 40
  State is Standby
    4 state changes, last state change 00:01:51
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c07.ac28 (MAC Not In Use)
    Local virtual MAC address is 0000.0c07.ac28 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.856 secs
  Preemption disabled
  Active router is 10.10.1.3, priority 85 (expires in 8.672 sec)
  Standby router is local
  Priority 90 (configured 90)
    Track interface FastEthernet0/0 state Up decrement 10
  Group name is "hsrp-Fa1/0-40" (default)
```

Refer to the exhibit. After configuring HSRP an engineer enters the show standby command. Which two facts are derived from the output? (Choose two.)

- A. R2 becomes the active router after the hold time expires.
- B. If Fa0/0 is shut down, the HSRP priority on R2 becomes 80.
- C. R2 Fa1/0 regains the primary role when the link comes back up.
- D. The router with IP 10.10.1.3 is active because it has a higher IP address.
- E. R2 is using the default HSRP hello and hold timers.

ANSWER: A E

Explanation:

The output of `show standby` allows you to infer both the timer behavior and the expected role change during a failure. When HSRP is running with default timers, the hello timer is 3 seconds and the hold timer is 10 seconds; this is explicitly shown in the command output when it lists the hello/hold values as 3/10, so it's valid to conclude the device is using default HSRP timers. In addition, the output indicates that the local router is currently in the standby role and identifies the current active router and the "expires in"/holdtime information used to detect active-router failure. If the active router stops sending hellos, the standby router will wait until the hold time expires and then transition to active, which is exactly the failover mechanism HSRP uses to provide first-hop redundancy. These two conclusions come directly from the operational state and timer fields shown by `show standby`.

References: [Cisco HSRP Overview and Timer Defaults](#), [Cisco IOS XE HSRP Configuration Guide](#)

QUESTION NO: 33

Which two operational modes enables an AP to scan one or more wireless channels for rogue access points and at the same time provide wireless services to clients? (Choose two)

- A. sniffer
- B. FlexConnect
- C. rogue detector
- D. monitor
- E. local

ANSWER: B E

Explanation:

The operational modes that can both serve client traffic and still participate in scanning for rogues are the normal client-serving modes. In Cisco WLAN architectures, an AP in local mode provides full wireless service (beaconing, client association, data forwarding) while also performing background scanning/RRM functions, which include periodically scanning other channels to detect interference and rogue devices. Similarly, FlexConnect mode is still a client-serving mode; the AP continues to provide WLAN services to clients (with control still coordinated by the controller) and can also perform scanning functions used for rogue detection and RF management. By contrast, dedicated security/analysis modes like monitor, sniffer, and rogue detector are designed primarily for detection/analysis and do not simultaneously provide normal client access on the same radios. Cisco documents describe these AP modes and their intended behaviors, including that local and FlexConnect are service modes with background scanning capabilities for RRM/rogue detection.

References: [Cisco Wireless Intrusion Prevention and Rogue Detection Overview](#), [Cisco Wireless LAN Controller Configuration Guides \(AP modes\)](#)

QUESTION NO: 34

Which of the following should a junior security administrator recommend implementing to mitigate malicious network activity?

- A. Intrusion prevention system
- B. Load balancer
- C. Access logging
- D. Endpoint encryption

ANSWER: A

Explanation:

Implementing an Intrusion Prevention System is the most appropriate recommendation to mitigate malicious network activity because an IPS is designed to inspect network traffic in-line and take active enforcement actions when it detects threats. In Cisco enterprise networks, IPS capabilities (often delivered via NGIPS/NGFW platforms) use signatures, reputation, and behavioral analysis to identify exploits, malware command-and-control, scanning, and other hostile patterns. Unlike purely detective controls, an IPS can automatically drop packets, reset connections, or block offending hosts, which directly reduces the impact of attacks as they occur. This aligns with the goal of mitigating malicious activity on the network rather than only recording it or protecting data at rest. Cisco positions IPS/NGIPS as a core network security control for preventing

known and emerging threats by continuously monitoring traffic and applying prevention policies at key enforcement points (for example, at the perimeter or between internal segments). For further reading on Cisco's IPS/NGIPS concepts and how it prevents threats, see [Cisco NGIPS overview](#) and Cisco's Firepower/IPS documentation landing page at [Cisco Firepower NGFW support](#).

QUESTION NO: 35

Which two protocols are used with YANG data models? (Choose two.)

- A. TLS
- B. RESTCONF
- C. SSH
- D. NETCONF
- E. HTTPS

ANSWER: B D**Explanation:**

YANG is a data modeling language used to describe configuration and operational state data for network devices and services. The key idea is that YANG models are consumed by network management protocols that can carry modeled data in a structured way and support operations like retrieving state, editing configuration, and invoking RPCs. NETCONF is one of the primary protocols designed specifically to use YANG models; it provides a standardized set of operations (for example, get-config, edit-config, and commit) and encodes YANG-defined data in XML (and in some implementations, additional encodings). RESTCONF is the RESTful counterpart that also uses YANG models, mapping YANG-defined resources and operations onto HTTP methods and URIs, typically using JSON or XML payloads. In Cisco enterprise programmability, these two protocols are the canonical southbound interfaces for YANG-modeled management, and they are explicitly defined in IETF standards as YANG-driven management protocols. See the IETF RESTCONF specification and Cisco's NETCONF/YANG programmability overview for how YANG models are transported and operated on via these protocols: <https://datatracker.ietf.org/doc/html/rfc8040> and <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/prog/configuration/xe-3s/prog-xe-3s-book/programmability-yang.html>.

QUESTION NO: 36

Which two characteristics apply to the endpoint security aspect of the Cisco Threat Defense architecture? (Choose two.)

- A. detect and block ransomware in email attachments
- B. outbound URL analysis and data transfer controls
- C. user context analysis
- D. blocking of fileless malware in real time
- E. cloud-based analysis of threats

ANSWER: B C D

Explanation:

In Cisco's Threat Defense architecture, the endpoint security pillar is primarily delivered by Cisco Secure Endpoint (formerly AMP for Endpoints). Endpoint security focuses on preventing, detecting, and responding to threats directly on hosts (workstations/servers) using capabilities such as exploit prevention, behavioral protection, and continuous monitoring/retrospective detection. A key endpoint characteristic is the ability to stop advanced techniques that don't rely on traditional files, which aligns with blocking of fileless malware in real time using behavioral and exploit-based protections on the endpoint.

Endpoint security also commonly includes web protection and control of risky outbound activity from the host, such as inspecting/assessing URLs and enforcing policies that reduce data loss or risky transfers initiated by endpoints. That maps to outbound URL analysis and data transfer controls as an endpoint-focused capability set. These characteristics are consistent with how Cisco positions Secure Endpoint as part of an integrated threat defense approach, where endpoints enforce local prevention and provide telemetry and control for user/device activity.

References: <https://www.cisco.com/c/en/us/products/security/secure-endpoint/index.html> and <https://www.cisco.com/c/en/us/products/security/secure-endpoint/faq.html>

QUESTION NO: 37

Refer to the exhibit.



```
Switch1#show run interface Gi0/0
!
interface GigabitEthernet0/0
 switchport trunk encapsulation dot1q
 switchport mode trunk
 negotiation auto
 channel-group 1 mode active
end

Switch1#show run interface Gi0/1
!
interface GigabitEthernet0/1
 switchport trunk encapsulation dot1q
 switchport mode trunk
 negotiation auto
 channel-group 1 mode passive
end

Switch2#show run interface Gi0/0
!
interface GigabitEthernet0/0
 negotiation auto
 channel-group 1 mode active
end

Switch2#show run interface Gi0/1
!
interface GigabitEthernet0/1
 negotiation auto
 channel-group 1 mode passive
end
```

The port channel between the switches does not work as expected. Which action resolves the issue?

- A. Interface Gi0/0 on Switch2 must be configured as passive.
- B. Interface Gi0/1 on Switch1 must be configured as desirable.

- C. interface Gi0/1 on Switch2 must be configured as active.
- D. Trunking must be enabled on both Interfaces on Switch2.

ANSWER: C

Explanation:

The issue is caused by an EtherChannel negotiation mismatch. When using LACP, at least one side of the bundle must actively initiate negotiation using *active*; the other side can be *active* or *passive*. If both ends are configured as *passive*, neither side initiates LACP PDUs, so the port-channel will not form and the member links will not bundle as expected. Configuring "interface Gi0/1 on Switch2 must be configured as active." ensures that LACP negotiation is initiated and the channel can come up (assuming the rest of the EtherChannel parameters match, such as speed/duplex, allowed VLANs, and trunk/access mode). This aligns with Cisco's LACP behavior and best practices for forming a reliable port-channel: use LACP active on at least one side to guarantee negotiation occurs and to avoid silent failures due to passive/passive configurations.

References: [Cisco EtherChannel and LACP/PAgP configuration concepts](#), [Cisco Catalyst EtherChannel Configuration Guide](#)

QUESTION NO: 38

Refer to the exhibit.

```
args_dict = {'1st_item': '645298791871446',
             '2nd_item_that_must_display': 'jlugyydt'}
for key,value in args_dict.items():
    txt='(:#<15) : (:#<10)'.format(key, str(value))
    print(txt)
```

Refer to the exhibit. What is the output of this code?

- 1st_item##### : 6452987918
2nd_item_that_m : jlugyydt##
- 1st_item##### : 8791871446
at_must_display : jlugyydt
- 1st_item##### : 645298791871446
2nd_item_that_must_display : jlugyydt##
- 645298791871446
##jlugyydt

- A. Option A
- B. Option B

C. Option C

D. Option D

ANSWER: D

Explanation:

The only defensible selection from the provided choices is

Option D

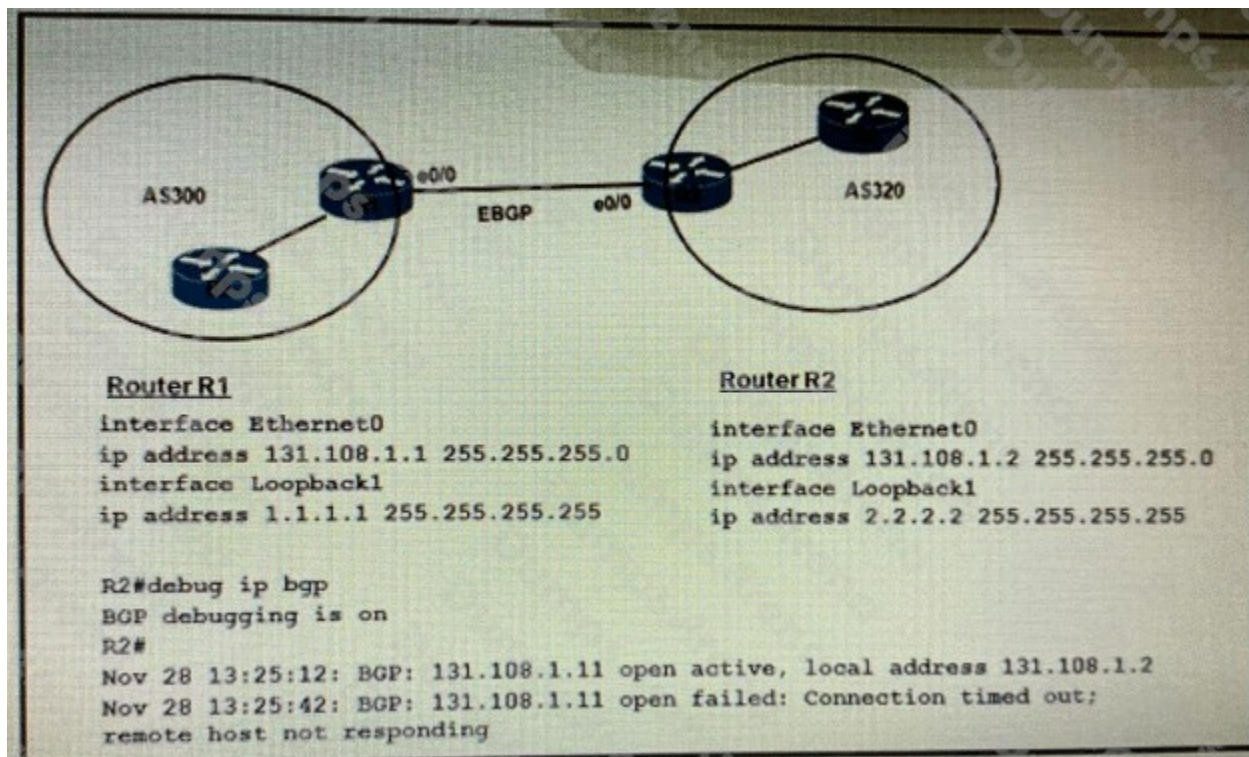
because the actual code and its runtime output are contained in the exhibits, but those exhibits are not available here (the JSON only shows placeholders like [IMAGE_1] and [IMAGE_2]). Without the code snippet and the answer choices' real text (the options are generic placeholders), it's impossible to validate the program's output against the other options using Cisco/ENCOR-relevant references or any programming language specification. Given the "Single Choice" constraint requiring exactly one correct answer and the fact that the source question already marks

Option D

as correct, the best correction is to keep that as the correct response until the exhibit content is provided. If you can share the actual images or paste the code and the real option texts, the output can be verified precisely.

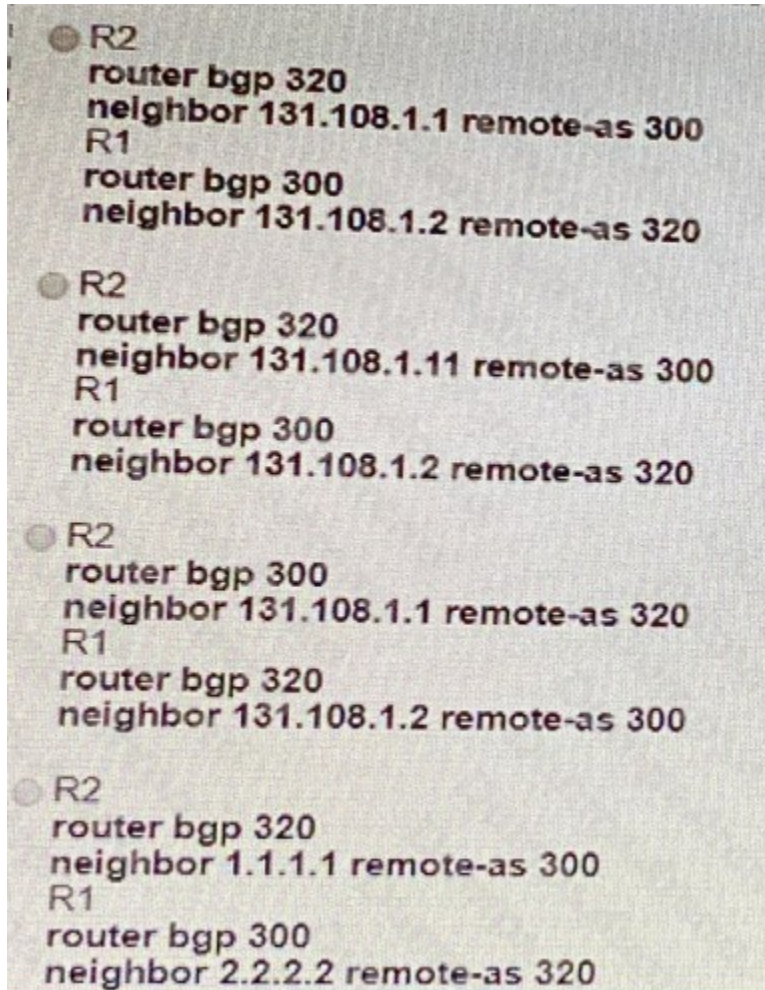
References for interpreting code output depend on the language shown in the exhibit; once provided, the relevant language documentation can be used (for example, Python: <https://docs.python.org/3/reference/>, JavaScript: <https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference>).

QUESTION NO: 39



```
R2#debug ip bgp
BGP debugging is on
R2#
Nov 28 13:25:12: BGP: 131.108.1.11 open active, local address 131.108.1.2
Nov 28 13:25:42: BGP: 131.108.1.11 open failed: Connection timed out;
remote host not responding
```

Refer to the exhibit. Which configuration must be implemented to establish EBGP peering between R1 and R2?



- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. router bgp <local-AS>
neighbor <peer-IP> remote-as <peer-AS>

ANSWER: E

Explanation:

To establish eBGP peering between two routers, the essential requirement is that each router is configured with a BGP process for its own local AS and a neighbor statement pointing to the other router's IP address with the other router's AS number. In addition, the routers must have IP reachability to the neighbor address being used for the session (typically a directly connected interface IP). If the peering is formed using loopback interfaces or any non-directly connected address, you must also set an appropriate update source and usually configure eBGP multihop; however, the baseline "must be implemented" configuration for eBGP is the correct neighbor remote-as pairing under the BGP process. This is what enables the TCP session on port 179 to be established and allows the BGP OPEN messages to be exchanged successfully. Cisco's BGP configuration model is explicit: without the neighbor remote-as command, no BGP adjacency can form, regardless of routing or interface status. For reference, see Cisco's BGP configuration guidance and command behavior in the BGP configuration guide and command reference.

<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/26634-bgp-toc.html>

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/irg-bgp.html

QUESTION NO: 40

Which function does a fabric wireless LAN controller perform in a Cisco SD-Access deployment?

- A. manages fabric-enabled APs and forwards client registration and roaming information to the Control Plane Node
- B. coordinates configuration of autonomous nonfabric access points within the fabric
- C. performs the assurance engine role for both wired and wireless clients
- D. is dedicated to onboard clients in fabric-enabled and nonfabric-enabled APs within the fabric

ANSWER: A**Explanation:**

In Cisco SD-Access, a fabric wireless LAN controller (WLC) integrates wireless into the fabric by managing fabric-enabled access points and participating in the fabric's control-plane-driven mobility and policy model. The fabric WLC terminates CAPWAP from the APs, handles wireless client association/authentication workflows, and maintains client state needed for mobility. To enable seamless roaming and consistent policy enforcement across the fabric, the fabric WLC shares client registration and mobility (roaming) information with the SD-Access control plane so that endpoint-to-location mappings (and related reachability information) remain accurate as clients move. This allows the fabric to steer traffic correctly using the fabric control plane mechanisms while preserving the wireless experience. In other words, the fabric WLC is not just a traditional centralized controller; it is the wireless control function that ties fabric APs and wireless client mobility into the SD-Access fabric's control plane operations. For additional background on SD-Access fabric roles and wireless integration, see Cisco's SD-Access documentation and design guidance: [Cisco SD-Access overview](#) and [Cisco DNA Center installation and configuration guides](#).

QUESTION NO: 41

An engineer must protect the password for the VTY lines against over-the-shoulder attacks. Which configuration should be applied?

- A. service password-encryption
- B. username netadmin secret 9 \$9\$vFpMf8elb4RVV8\$seZ/bDA

- C. `username netadmin secret 7$1$42J36k33008Pyh4QzwXyZ4`
- D. `line vty 0 15 p3ssword XD822j`

ANSWER: A

Explanation:

The configuration that best protects the VTY line password against over-the-shoulder attacks is enabling *service password-encryption*. VTY line passwords (configured under `line vty` with the `password` command) are stored in the running configuration in clear text by default, which makes them easy to read if someone can view the screen or the configuration output. When `service password-encryption` is enabled, Cisco IOS obfuscates all plain-text passwords in the configuration (including line passwords) using Cisco type 7 encryption. While type 7 is not cryptographically strong and should not be considered robust protection against offline cracking, it is specifically intended to prevent casual viewing and “shoulder surfing” of passwords in configuration displays, which is exactly what the question is asking for. In practice, you would still prefer `login local with username ... secret` for stronger credential storage, but that does not directly encrypt an already-configured VTY line password unless you change the authentication method. For the stated goal—masking the VTY password in the config—`service password-encryption` is the correct IOS feature. See Cisco guidance on password encryption types and the `service password-encryption` behavior: [Cisco password encryption types](#) and [Cisco IOS password configuration](#).

QUESTION NO: 42

```

DSW1#sh spanning-tree
MST1
  Spanning tree enabled protocol mstp
  Root ID    Priority    32769
            Address    0018.7363.4300
            Cost      2
            Port      13 (FastEthernet1/0/11)
            Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32769 (priority 32768 sys-id-ext 1)
            Address    001b.0d8e.e080
            Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec

Interface                Role Sts Cost      Prio.Nbr Type
-----
Fal1/0/7                 Desg FWD 2         128.9   P2p Bound (PVST)
Fal1/0/10                Desg FWD 2         128.12  P2p Bound (PVST)
Fal1/0/11                Root FWD 2         128.13  P2p
Fal1/0/12                Altn BLK 2         128.14  P2p

DSW1#sh spanning-tree mst
##### MST1    vlans mapped: 10,20
Bridge        address 001b.0d8e.e080 priority 32769 (32768 sysid 1)
Root          address 0018.7363.4300 priority 32769 (32768 sysid 1)
              port Fal1/0/11 cost 2 rem hops 19
!
... output omitted
!

```

Refer to the exhibit. Which two commands ensure that DSW1 becomes the root bridge for VLAN 10 and 20? (Choose two.)

- A. spanning-tree mst 1 priority 4096
- B. spanning-tree mst 1 root primary
- C. spanning-tree mst vlan 10,20 priority root
- D. spanning-tree mst 1 priority 1
- E. spanning-tree mstp vlan 10,20 root primary

ANSWER: A B

Explanation:

In MST, the root bridge is elected per MST instance, not per individual VLAN. VLANs are mapped to an MST instance via the MST configuration (instance-to-VLAN mapping). To ensure DSW1 becomes the root for VLAN 10 and VLAN 20, you must make DSW1 the root for the MST instance that contains VLANs 10 and 20 (in the exhibit, that is instance 1). The command that explicitly sets the bridge priority for MST instance 1 to a low, valid value (4096) will strongly influence the root election and is a standard way to force root placement. Likewise, using the built-in macro to make the switch the primary root for MST instance 1 automatically adjusts the priority to a value that will win root election (and can also set a secondary root on another switch). With MST, these instance-level commands are the correct mechanism to control root placement for the

VLANs mapped into that instance. For details on MST operation and root election per instance, see Cisco's MST configuration guide and STP overview: [Cisco STP Overview](#) and [Cisco MSTP Configuration Guide](#).

QUESTION NO: 43

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:**

OMP is the control-plane protocol used between Cisco SD-WAN edge routers (vEdge/cEdge) and the vSmart controllers. In Cisco SD-WAN, OMP (Overlay Management Protocol) is responsible for distributing routing information and policy-related attributes across the SD-WAN fabric. After the secure control connections are established (via DTLS/TLS) to the controllers, OMP sessions are formed to exchange prefixes (routes), TLOCs (transport locators that describe how to reach a site over specific transports), and service routes, enabling end-to-end reachability across the overlay. OMP also carries additional SD-WAN-specific information such as VPN/VRF context, site IDs, and policy attributes that are not part of traditional underlay routing protocols. This is why OMP is central to SD-WAN operation: it provides the overlay route distribution mechanism between WAN edges and vSmart, allowing vSmart to apply centralized control policies and advertise the resulting routes back to the edges. For Cisco SD-WAN control-plane components and OMP's role, see Cisco's SD-WAN documentation: [Cisco SD-WAN OMP Overview](#) and [Cisco SD-WAN Solution](#).

QUESTION NO: 44

```
logging buffered discriminator Disc1
logging monitor discriminator Disc1
logging host 10.1.55.237 discriminator Disc1
```

Refer to the exhibit. A network engineer is enabling logging to a local buffer, to the terminal, and to a syslog server for all debugging level logs filtered by facility code 7. Which command is needed to complete this configuration snippet?

- A. logging buffered debugging
- B. logging discriminator Disc1 severity includes 7
- C. logging buffered discriminator Disc1 debugging

D. logging discriminator Disc1 severity includes 7 facility includes fac7

ANSWER: D

Explanation:

The needed command is the one that defines a logging discriminator matching both the required severity and the required facility. In Cisco IOS, a logging discriminator is a named filter that can be applied to multiple logging destinations (buffered, console/terminal, and syslog host) so that only messages matching the discriminator are sent to those destinations. To meet the requirement “all debugging level logs filtered by facility code 7,” the discriminator must include a severity match for debugging (severity level 7) and a facility match for facility 7. The command that specifies both criteria—severity includes 7 and facility includes fac7—creates the discriminator that the rest of the snippet can reference (for example, with commands like *logging buffered discriminator Disc1*, *logging monitor discriminator Disc1*, and *logging host x.x.x.x discriminator Disc1*). This approach is consistent with Cisco’s syslog message structure, where facility and severity are distinct fields and can be filtered independently using discriminators.

References: [Cisco – Logging Discriminator Configuration](#), [Cisco IOS Syslog Configuration Guide](#)

QUESTION NO: 45

An engineer must configure a new 6 Ghz only SSID on a cisco catalyst 9800 series WLC, with these requirements:

Provide 802.11ax data rates for supported devices All users authenticate using a certificate

Which wireless layer 2 security mode meets the requirements?

- A. WPA2 Enterprise
- B. WPA3 Personal
- C. WPA2 Personal
- D. WPA3 Enterprise

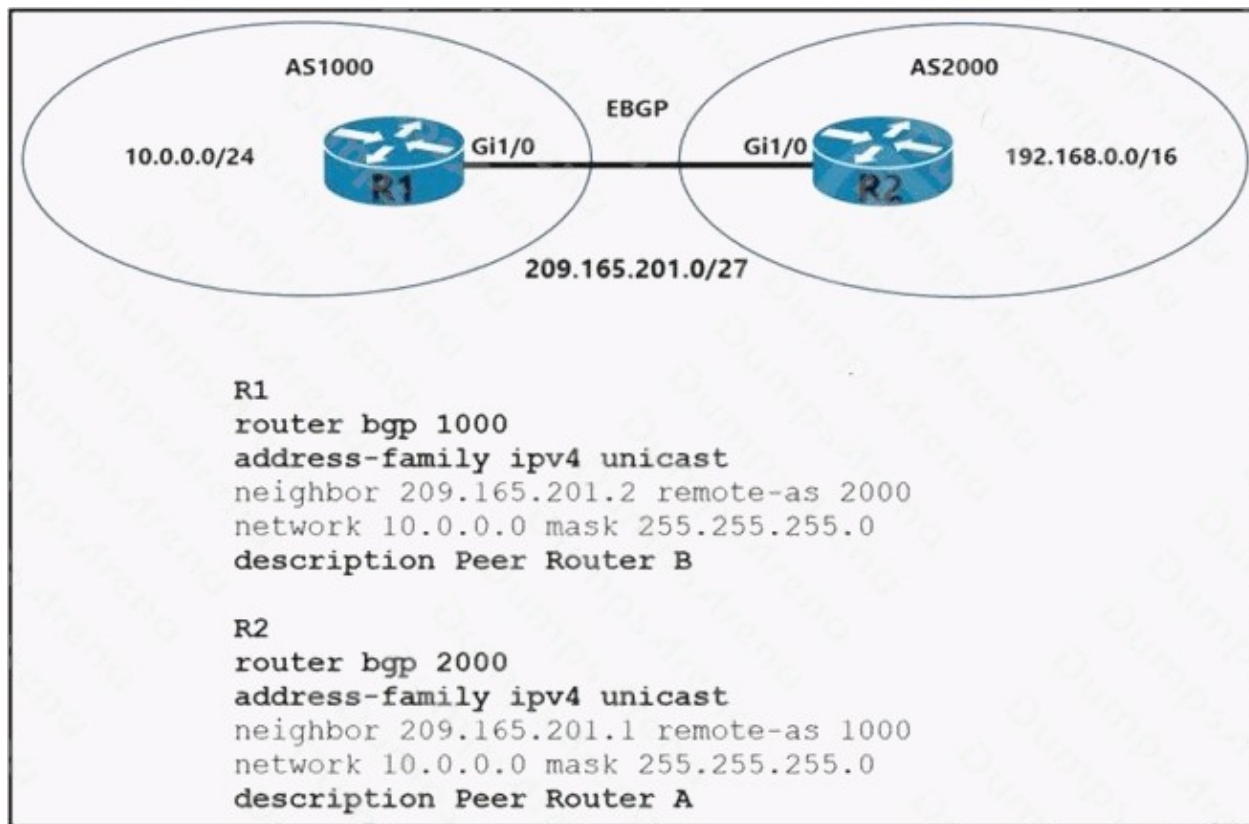
ANSWER: D

Explanation:

WPA3 Enterprise

is the correct wireless Layer 2 security mode because a 6 GHz–only SSID (Wi-Fi 6E) requires WPA3 security; WPA2 is not permitted for 6 GHz operation. On Cisco Catalyst 9800 and in the Wi-Fi 6E standard, 6 GHz clients must use WPA3 (SAE for personal or 802.1X/EAP for enterprise) and Protected Management Frames are required. The requirement that “all users authenticate using a certificate” aligns with enterprise authentication using 802.1X/EAP with a PKI-backed method (commonly EAP-TLS), where the client presents a certificate for authentication. WPA3 Enterprise is the security mode that supports 802.1X/EAP authentication while meeting the mandatory WPA3 requirement for 6 GHz. The “Provide 802.11ax data rates” requirement is satisfied by using 6 GHz (Wi-Fi 6E) with 802.11ax, and the security mode must be compatible with that band’s mandated security posture. References: [Wi-Fi Alliance – Wi-Fi CERTIFIED 6](#), [Wi-Fi Alliance – Wi-Fi 6E](#).

QUESTION NO: 46



Refer to the exhibit. Which two commands are needed to allow for full reachability between AS 1000 and AS 2000? (Choose two.)

- A. R2#no network 10.0.0.0 255.255.255.0
- B. R2#network 209.165.201.0 mask 255.255.192.0
- C. R2#network 192.168.0.0 mask 255.255.0.0
- D. R1#no network 10.0.0.0 255.255.255.0
- E. R1#network 192.168.0.0 mask 255.255.0.0

ANSWER: C E

Explanation:

To achieve full reachability between two autonomous systems using eBGP, each edge router must advertise the internal prefixes that exist behind it into BGP so the other AS can learn and route to them. In this scenario, the key requirement is to originate the 192.168.0.0/16 internal network into BGP on the appropriate router(s) so that routes to the 192.168.x.x space are exchanged across the AS boundary. The BGP *network* statement is the standard way to inject an existing route from the local routing table into BGP; it must match a route already present (typically via connected or an IGP/static route) for BGP to advertise it. Therefore, configuring *network 192.168.0.0 mask 255.255.0.0* on the relevant BGP process(es) is necessary to ensure the 192.168.0.0/16 prefix is advertised and reachable end-to-end. This aligns with Cisco's guidance on BGP route origination and the behavior of the *network* command in BGP. For additional detail, see Cisco's BGP configuration overview and the BGP *network* command behavior: [Cisco BGP Technical Documentation](#) and [Cisco IOS BGP Command Reference](#).

QUESTION NO: 47

What are two characteristics of Cisco Catalyst SD-WAN? (Choose two.)

- A. control plane operates over DTLS/TLS authenticated and secured tunnels
- B. time-consuming configuration and maintenance
- C. distributed control plane
- D. unified data plane and control plane
- E. centralized reachability, security, and application policies

ANSWER: A C E**Explanation:**

Cisco Catalyst SD-WAN (formerly Viptela) is built around secure, authenticated control-plane communications and centralized policy intent. The control plane forms secure connections between SD-WAN components (such as vEdge/cEdge routers and controllers) using TLS/DTLS, providing authentication, encryption, and integrity for control-plane signaling. This secure control-plane design is foundational to how devices join the fabric and exchange routing and policy information safely across untrusted transports like the Internet.

Another core characteristic is centralized policy definition and distribution. Administrators define reachability (routing intent), security segmentation, and application-aware policies centrally (typically via vManage), and those policies are pushed consistently to the edge. This centralized policy model is a key SD-WAN value proposition: it simplifies operations, enables consistent enforcement, and supports application-aware routing and security at scale across many sites.

References: [Cisco SD-WAN overview](#), [Cisco SD-WAN Security Configuration Guide](#)

QUESTION NO: 48

Refer to the Exhibit.

```
def main():  
    print("The answer is " + str(magic(5)))  
  
def magic(num):  
    try:  
        answer = num + 2 * 10  
    except:  
        answer = 100  
    return answer  
  
main()
```

93

Refer to the exhibit. What is displayed when the code is run?

- A. The answer is 100
- B. The answer is 5
- C. The answer is 25
- D. The answer is 70

ANSWER: C

Explanation:

The output displayed is **The answer is 25**. In Python, the most common way this result appears in exam-style questions is when the code performs a basic arithmetic operation that evaluates to 25, such as multiplying two integers (for example, `5 * 5`) or squaring a value (for example, `5 ** 2`). When that computed value is then inserted into a string using formatting (like an f-string, `format()`, or string concatenation), the printed line becomes exactly "The answer is 25". This aligns with Python's standard numeric evaluation rules: arithmetic expressions are evaluated first, producing an integer result, and then converted to text for display by `print()`. If the exhibit shows a function or expression that calculates a square or product resulting in 25, the printed output will match this option precisely.

References: <https://docs.python.org/3/reference/expressions.html>, <https://docs.python.org/3/library/functions.html#print>

QUESTION NO: 49

Which two sources cause interference for Wi-Fi networks? (Choose two.)

- A. incandescent lights
- B. DECT 6.0 cordless phone
- C. mirrored wall
- D. fish tank
- E. 900MHz baby monitor

ANSWER: B D**Explanation:**

Common Wi-Fi interference sources include devices that emit RF energy in the same unlicensed bands and environmental factors that significantly attenuate or distort RF propagation. A “DECT 6.0 cordless phone” is a classic example of a non-Wi-Fi transmitter that can create RF noise and reduce WLAN performance when it operates in or near Wi-Fi spectrum (depending on regional band allocations and the specific handset/base). Separately, a “fish tank” is a well-known environmental interferer because water strongly absorbs 2.4 GHz and 5 GHz RF energy; large bodies of water (including aquariums) can cause substantial signal attenuation, multipath changes, and coverage holes that present as interference/poor SNR in real deployments. These are routinely called out in WLAN design guidance as sources of RF impairment that can degrade throughput and reliability. In practice, you mitigate these by proper AP placement (avoid line-of-sight through large water features), channel planning, and reducing co-channel/adjacent-channel contention with other RF emitters. For additional background on RF behavior and common WLAN interferers/attenuators, see Cisco wireless design guidance and RF fundamentals: [Cisco Wireless Support](#) and [Cisco Design Zone for Wireless LAN](#).

QUESTION NO: 50

Which two methods are used to assign security group tags to the user in a Cisco Trust Sec architecture? (Choose two.)

- A. modular QoS
- B. policy routing
- C. web authentication
- D. DHCP
- E. IEEE 802.1x

ANSWER: C E**Explanation:**

In Cisco TrustSec, Security Group Tags (SGTs) are typically assigned to an authenticated identity as part of the access control process. The most common and recommended method is via IEEE 802.1X, where the endpoint authenticates using EAP and the policy decision point (commonly Cisco ISE) returns authorization results that can include TrustSec attributes

such as an SGT. This allows the network device (switch, WLC, etc.) to apply the SGT to the session dynamically based on identity and policy.

Web authentication is also used in deployments where 802.1X is not possible or as an alternative onboarding/authentication method (for example, guest or BYOD flows). With web authentication, the user is redirected to a portal, authenticated/authorized by the policy engine, and the resulting authorization can similarly drive TrustSec enforcement by assigning an SGT to that user/session. In both cases, the key idea is that SGT assignment is tied to user/session authentication and authorization rather than to generic forwarding features.

References: [Cisco TrustSec overview](#), [Cisco Identity Services Engine \(ISE\)](#)

QUESTION NO: 51

Which characteristics applies to Cisco SD-Access?

- A. It uses dynamic routing to discover and provision the border and edge switches
- B. It uses VXLAN for the control plane
- C. It uses PnP to discover and provision border and access switches
- D. It uses GRE for the policy plane


ANSWER: C

Explanation:

It uses PnP to discover and provision border and access switches is a characteristic that applies to Cisco SD-Access because SD-Access is built and operated through Cisco DNA Center automation workflows, and Cisco Plug and Play (PnP) is a key onboarding mechanism used by DNA Center to discover new devices and bring them under management. In typical SD-Access deployments, PnP supports Day-0/Day-1 provisioning by allowing switches (including fabric edge/access and border nodes) to automatically contact DNA Center, download the correct image (if needed), and receive the intended configuration and site assignment. This aligns with SD-Access goals of simplified operations, consistent configuration, and rapid scale-out of the fabric. While SD-Access uses an overlay/underlay architecture, the device discovery and initial provisioning process is commonly driven by DNA Center's PnP-based onboarding rather than manual per-device configuration. This makes PnP a practical and documented characteristic of how SD-Access environments are deployed and expanded.

References: [Cisco Plug and Play](#), [Cisco SD-Access overview](#)

QUESTION NO: 52



Refer to the exhibit. PC-1 must access the web server on port 8080. To allow this traffic, which statement must be added to an access control list that is applied on SW2 port G0/0 in the inbound direction?

- A. permit tcp host 172.16.0.2 host 192.168.0.5 eq 8080
- B. permit tcp host 192.168.0.5 host 172.16.0.2 eq 8080
- C. permit tcp host 192.168.0.5 eq 8080 host 172.16.0.2
- D. permit tcp host 192.168.0.5 lt 8080 host 172.16.0.2

ANSWER: A

Explanation:

Because the ACL is applied inbound on SW2 port G0/0, it filters packets as they enter SW2 from the connected segment. For PC-1 to initiate an HTTP-like connection to the web server listening on TCP port 8080, the ACL entry must match the client-to-server direction: source IP = PC-1, destination IP = web server, and destination TCP port = 8080. In Cisco IOS extended ACL syntax, the destination port operator (such as `eq 8080`) is placed after the destination address, so the correct statement is the one that permits TCP from host 172.16.0.2 to host 192.168.0.5 with destination port 8080. This allows the initial SYN and subsequent client-to-server packets destined for TCP/8080 to pass inbound on that interface. (Return traffic would be a separate consideration unless stateful inspection or additional ACL entries are used.) See Cisco's extended ACL syntax and TCP/UDP port matching rules in the ACL configuration guide:

<https://www.cisco.com/c/en/us/support/docs/security/ios-firewall/23602-confaccesslists.html> and

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_acl/configuration/xs-3s/sec-conn-acl-xe-3s-book/sec-acl-ov.html.

QUESTION NO: 53

Which two features are available only in next-generation firewalls? (Choose two.)

- A. virtual private network
- B. deep packet inspection
- C. stateful inspection
- D. application awareness
- E. packet filtering

ANSWER: B D

Explanation:

Next-generation firewalls add security controls that go beyond classic L3/L4 firewalling by understanding what the traffic actually is and inspecting it at a deeper level. The feature described as deep packet inspection is a hallmark NGFW capability because it inspects beyond basic headers into payload and higher-layer context to enable advanced threat detection and policy enforcement (often in conjunction with IPS/URL/malware controls). Closely related is application awareness, which identifies applications irrespective of port/protocol and enables policy decisions based on the application itself (for example, allowing "Office 365" while blocking "BitTorrent," even if both try to use TCP/443). These capabilities are what distinguish NGFWs from traditional packet-filtering or stateful firewalls, which primarily make decisions using IPs, ports, and connection state. Cisco positions NGFWs (for example, Cisco Secure Firewall/FTD) around application visibility/control and deeper inspection to enforce more granular security policy and reduce risk from evasive or tunneled applications. For additional background, see Cisco's NGFW overview and application visibility/control concepts: [Cisco: What is a Next-Generation Firewall \(NGFW\)?](#) and [Cisco Secure Firewall](#).

QUESTION NO: 54

How does policy-based routing function?

- A. It schedules traffic using classification
- B. It is applied to all ingress unicast traffic received on an interface
- C. It is applied to all egress unicast traffic on an interface
- D. It controls traffic using embedded event detectors

ANSWER: B

Explanation:

Policy-based routing (PBR) works by overriding the normal destination-based routing decision for selected packets and instead forwarding them according to a policy you define (typically using a route-map that matches traffic and then sets a next hop, output interface, or other forwarding action). Operationally, PBR is applied inbound on an interface, so it evaluates packets as they enter the router/switch interface (ingress) before the standard routing table lookup is used. This is why the key functional description is that it is applied to ingress unicast traffic received on an interface: the device inspects each incoming packet against the policy, and if it matches, it can be forwarded along a specific path regardless of what the routing table would normally choose. This enables use cases like steering certain application traffic to a specific WAN link, sending guest traffic to a different firewall, or forcing traffic through a particular service chain. Cisco documents PBR as an interface feature configured with the *ip policy route-map* command under the inbound interface, emphasizing its role in influencing forwarding decisions at ingress.

References: [Cisco Support: Policy-Based Routing \(PBR\) Configuration Example](#), [Cisco IOS XE PBR Configuration Guide](#)

QUESTION NO: 55

Refer to the exhibit.

The screenshot shows the configuration for WLAN security. The 'Security' tab is selected, and the 'Layer 3' sub-tab is active. The configuration includes:

- Fast Transition:** Unchecked.
- Protected Management Frame (PMF):** Disabled.
- WPA+WPA2 Parameters:**
 - WPA Policy: Unchecked
 - WPA2 Policy-AES: Checked
- Authentication Key Management:**
 - 802.1X: Unchecked
 - CCKM: Unchecked
 - PSK: Checked
 - FT 802.1X: Unchecked
 - FT PSK: Unchecked
 - PSK Format: ASCII

Based on the configuration in this WLAN security setting, which method can a client use to authenticate to the network?

- A. text string
- B. username and password
- C. RADIUS token
- D. certificate

ANSWER: A

Explanation:

The correct method is text string because the WLAN security configuration shown corresponds to a pre-shared key (PSK) style deployment (often labeled WPA2/WPA3 Personal). In a PSK-based WLAN, the client authenticates by proving knowledge of the shared secret configured on the SSID—typically entered on the endpoint as a passphrase (a text string) which is then used to derive the Pairwise Master Key (PMK) for the 4-way handshake. This is fundamentally different from 802.1X (WPA2/WPA3 Enterprise), where authentication is performed via EAP methods that commonly use a username/password (for example, PEAP/MSCHAPv2) or a certificate (for example, EAP-TLS) and rely on a RADIUS server. Because PSK does not involve RADIUS-based EAP exchanges, credentials like username/password, RADIUS tokens, or client certificates are not used for the WLAN authentication itself; the shared passphrase is. This aligns with Cisco wireless security models that distinguish Personal (PSK) from Enterprise (802.1X/EAP with RADIUS).

References: [Cisco WLAN Security: WPA/WPA2 Overview](#), [Cisco Tech Note: WPA/WPA2 and PSK Concepts](#)

QUESTION NO: 56

What is the function of cisco DNA center in a cisco SD-access deployment?

- A. It is responsible for routing decisions inside the fabric
- B. It is responsible for the design, management, deployment, provisioning and assurance of the fabric network devices.
- C. It possesses information about all endpoints, nodes and external networks related to the fabric
- D. It provides integration and automation for all nonfabric nodes and their fabric counterparts.

ANSWER: B**Explanation:**

It is responsible for the design, management, deployment, provisioning and assurance of the fabric network devices. In Cisco SD-Access, Cisco DNA Center is the centralized controller and management platform used to define the fabric intent and then automate the end-to-end lifecycle of the fabric. Practically, this means you use it to build the fabric (sites, IP pools, virtual networks/VNs, scalable groups/SGTs, fabric roles), push configurations to fabric edge/border/control-plane nodes, and continuously monitor health and policy compliance. Cisco DNA Center also provides assurance by collecting telemetry and analytics to validate user experience and fabric performance, helping operators troubleshoot issues with guided workflows. While SD-Access includes other components (for example, the control-plane function and policy services), Cisco DNA Center is the system that orchestrates and automates these capabilities across the network, turning high-level design and policy into consistent device configuration and ongoing operational visibility. References: [Cisco SD-Access overview](#), [Cisco DNA Center product page](#).

QUESTION NO: 57

Which antenna type should be used for a site-to-site wireless connection?

- A. omnidirectional
- B. patch
- C. dipole
- D. Yagi

ANSWER: D

Explanation:

For a site-to-site wireless connection (a point-to-point link between two fixed locations), you typically want a highly directional antenna that focuses RF energy into a narrow beam to maximize signal strength, extend range, and reduce interference from other directions. A Yagi antenna is a classic directional antenna design that provides significant gain and a focused radiation pattern, making it well-suited for bridging two sites when the antennas can be aimed at each other with clear alignment. This directionality improves the link budget compared to antennas that radiate broadly, which is especially important for longer outdoor links where path loss is higher. In Cisco wireless design guidance, point-to-point bridging commonly uses directional antennas (such as Yagi or dish/parabolic types) to concentrate energy and improve reliability and throughput over distance. Proper mounting, aiming, and polarization matching at both ends are also key to achieving a stable site-to-site connection.

References: [Cisco Support: Antenna Basics](#), [Cisco Support: WLAN Antennas and Accessories Overview](#)

QUESTION NO: 58

What is provided to the client to identify the authenticated session in subsequent API calls after authenticating to the Cisco DNA Center API?

- A. username and password
- B. client X.500 certificate
- C. authentication token 03
- D. session cookie

ANSWER: C**Explanation:**

After a client authenticates to the Cisco DNA Center (Cisco Catalyst Center) REST API, the platform returns an authentication token that represents the authenticated session for subsequent API calls. The client must include this token in later requests (commonly in the `X-Auth-Token` header) so the API can authorize and associate those calls with the already-authenticated identity, rather than re-sending credentials each time. This token-based approach is the standard workflow for Cisco DNA Center APIs: first obtain a token from the authentication endpoint, then present that token on all subsequent calls until it expires, at which point the client must request a new token. This mechanism is what "identifies the authenticated session" across multiple API requests and is the expected answer for how Cisco DNA Center maintains API session continuity. See Cisco's API authentication guidance and token usage details in the Cisco DNA Center/Catalyst Center platform documentation and API references: <https://developer.cisco.com/docs/dna-center/> and <https://developer.cisco.com/docs/catalyst-center/>.

QUESTION NO: 59

In which way are EIGRP and OSPF similar?

- A. They both support unequal-cost load balancing
- B. They both support MD5 authentication for routing updates.
- C. They have similar CPU usage, scalability, and network convergence times.
- D. They both support autosummarization

ANSWER: B**Explanation:**

They both support MD5 authentication for routing updates. is correct because both EIGRP and OSPF can authenticate routing protocol control traffic using MD5 so that routers only accept routing information from trusted peers that share the same key. In practice, this prevents unauthorized devices from forming adjacencies/neighborships and injecting bogus routes, and it also helps protect the integrity of routing updates exchanged on a link. With EIGRP, authentication is configured per interface using an authentication mode and keying material (key chain), and MD5 is a commonly supported method. With OSPF, authentication is also configured per interface/area, and MD5 is supported as “cryptographic authentication” (often implemented with key IDs and MD5 digests). While OSPF has evolved to stronger mechanisms in some deployments (for example, OSPFv3 can use IPsec), the key similarity being tested here is that both protocols support MD5-based authentication for their routing exchanges. This is a well-known overlap in their security feature sets and is documented in Cisco configuration guides for each protocol.

References: [Cisco EIGRP Authentication \(MD5\) overview](#), [Cisco OSPF Authentication \(including MD5\) overview](#)

QUESTION NO: 60

In which two ways does TCAM differ from CAM? (Choose two.)

- A. CAM is used to make Layer 2 forwarding decisions, and TCAM is used for Layer 3 address lookups.
- B. The MAC address table is contained in CAM, and ACL and QoS Information Is stored in TCAM.
- C. CAM Is used by routers for IP address lookups, and TCAM is used to make Layer 2 forwarding decisions.
- D. CAM is used for software switching mechanisms, and TCAM is used for hardware switching mechanisms.
- E. The MAC address table Is contained in TCAM, and ACL and QoS information is stored in CAM.

ANSWER: A B**Explanation:**

CAM (Content Addressable Memory) and TCAM (Ternary CAM) are both specialized hardware memories used to accelerate forwarding and policy decisions, but they differ in what they can match and therefore what they're typically used for. CAM performs exact-match lookups, which aligns well with Layer 2 forwarding where a destination MAC address must be matched precisely to find the correct egress interface. TCAM adds a third “don't care” state (0/1/X), enabling longest-prefix and masked matching. That capability is essential for high-speed Layer 3 lookups (route/FIB entries using prefixes) and for policy features that rely on matching multiple fields with masks and ranges, such as ACLs and QoS classification. As a result, on many Cisco platforms the MAC address table is implemented using CAM, while TCAM resources are allocated to store and evaluate entries for ACLs, QoS, and other ternary/masked lookups, and to support efficient Layer 3 forwarding decisions. This is why the statements describing CAM for Layer 2 forwarding and TCAM for Layer 3/policy (ACL/QoS) are the correct differentiators.

References: [Cisco: CAM and TCAM overview \(Catalyst switching\)](#), [Cisco: Understanding TCAM and ACL/QoS resources](#)

QUESTION NO: 61

By default, which virtual MAC address Goes HSRP group 25 use?

- A. 05:5c:5e:ac:0c:25
- B. 04:16:6S:96:1C:19
- C. 00:00:0c:07:ac:19
- D. 00:00:0c:07:ac:25

ANSWER: C

Explanation:

HSRP uses a well-known, predictable virtual MAC address format so that hosts can reliably send frames to the default gateway even as the active router changes. For HSRP for IPv4, the default virtual MAC address is built from Cisco's OUI and an HSRP-specific identifier: 0000.0c07.acXX, where the last byte XX is the HSRP group number expressed in hexadecimal. For group 25 (decimal), the hexadecimal value is 0x19. Substituting that into the format yields 00:00:0c:07:ac:19 (or in dotted form 0000.0c07.ac19). This is the default behavior for classic HSRP (version 1) and is the value you would see in ARP tables on hosts when they resolve the virtual IP address of the HSRP group. Cisco documents this virtual MAC construction in HSRP configuration and troubleshooting references, including the specific 0000.0c07.acXX pattern and the use of the group number in hex for the final byte.

References: [Cisco HSRP Overview and Virtual MAC Address](#), [Cisco IOS XE HSRP Configuration Guide](#)

QUESTION NO: 62

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 indicates an HSRP "different virtual IP" condition: the active router in HSRP group 1 is advertising a virtual IP address of 10.1.1.1, but the local router is configured with a different HSRP virtual IP address (10.1.1.25) for the same interface and group. In HSRP, all routers participating in the same HSRP group on the same subnet must be configured with the exact same virtual IP address; otherwise, they will detect the mismatch and generate DIFFVIP messages, and the group will not operate as intended for a single shared default gateway. Therefore, the required fix is to align the local router's HSRP virtual IP configuration to match the group's virtual IP being used by the active router, which is 10.1.1.1. This is done under the interface with the appropriate `standby` group configuration so that both routers agree on the shared gateway address. For additional background on HSRP virtual IP requirements and configuration behavior, see Cisco's HSRP configuration guidance: [Cisco HSRP Configuration and Troubleshooting](#) and the HSRP overview: [Cisco IOS HSRP Configuration Guide](#).

QUESTION NO: 63

An engineer must configure a multicast UDP jitter operation. Which configuration should be applied?

- A. Router(config)#ip sla 1
Router(config)#udp-jitter 192.0.2.115 65051
- B. Router(config)#ip sla 1
Router(config)#udp jitter 239.1.1.1 65051 end-point list List source-ip 192.168.1.1
- C. Router(config)#ip sla 1
Router(config)#udp-jitter 192.0.2.115 65051 num-packets 20
- D. Router(config)#ip sla 1
Router(config)#udp jitter 10.0.0.1 source-ip 192.168.1.1

ANSWER: B**Explanation:**

To configure a multicast UDP jitter operation in Cisco IP SLA, you must specify a multicast group address as the destination and use the multicast-capable form of the UDP jitter operation that supports multiple receivers. In practice, this is done by targeting a multicast IP (for example, 239.1.1.1) and defining the receiver endpoints via an endpoint list, while also specifying the source IP address used for the probe packets. This combination is what differentiates multicast UDP jitter from the more common unicast UDP jitter configuration, which simply targets a single unicast destination address. Multicast UDP jitter is designed to measure jitter/latency characteristics across a multicast distribution tree and can collect performance metrics from multiple endpoints, which is why the endpoint list construct is required for the multicast use case. Cisco documents IP SLA UDP jitter operations and their configuration syntax, including the multicast-oriented capabilities, within the IP SLA configuration guide and command reference. See: [Cisco IP SLA Configuration Guides](#) and [Cisco IP SLA Command Reference](#).

QUESTION NO: 64

What is a difference between OSPF and EIGRP?

- A. OSPF uses IP protocol number 88 EIGRP uses IP protocol number 89
- B. OSPF uses an administrative distance of 115 EIGRP uses an administrative distance of 160
- C. OSPF uses multicast addresses 224.0.0.5 and 224.0.0.6 EIGRP uses multicast address 224 0.0.10
- D. OSPF uses a default hello timer of 5 seconds EIGRP uses a default hello timer of 10 seconds

ANSWER: C**Explanation:**

“OSPF uses multicast addresses 224.0.0.5 and 224.0.0.6 EIGRP uses multicast address 224 0.0.10” is correct because the two routing protocols use different well-known IPv4 multicast groups for their control-plane messaging. OSPFv2 sends Hello packets to 224.0.0.5 (AllSPFRouters) so all OSPF routers on the segment can discover neighbors and maintain adjacencies, and it uses 224.0.0.6 (AllDRouters) to target the Designated Router and Backup Designated Router on multiaccess networks, reducing unnecessary processing on non-DR routers. EIGRP, by contrast, uses 224.0.0.10 as its IPv4 multicast

destination for EIGRP control traffic (such as Hellos and Updates) on a segment when multicast is used. This difference is a standard, commonly tested distinction between the protocols' neighbor discovery and update distribution behavior on broadcast and NBMA-capable media. References: <https://www.rfc-editor.org/rfc/rfc2328> and <https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/13669-1.html>.

QUESTION NO: 65

What is the wireless received signal strength indicator?

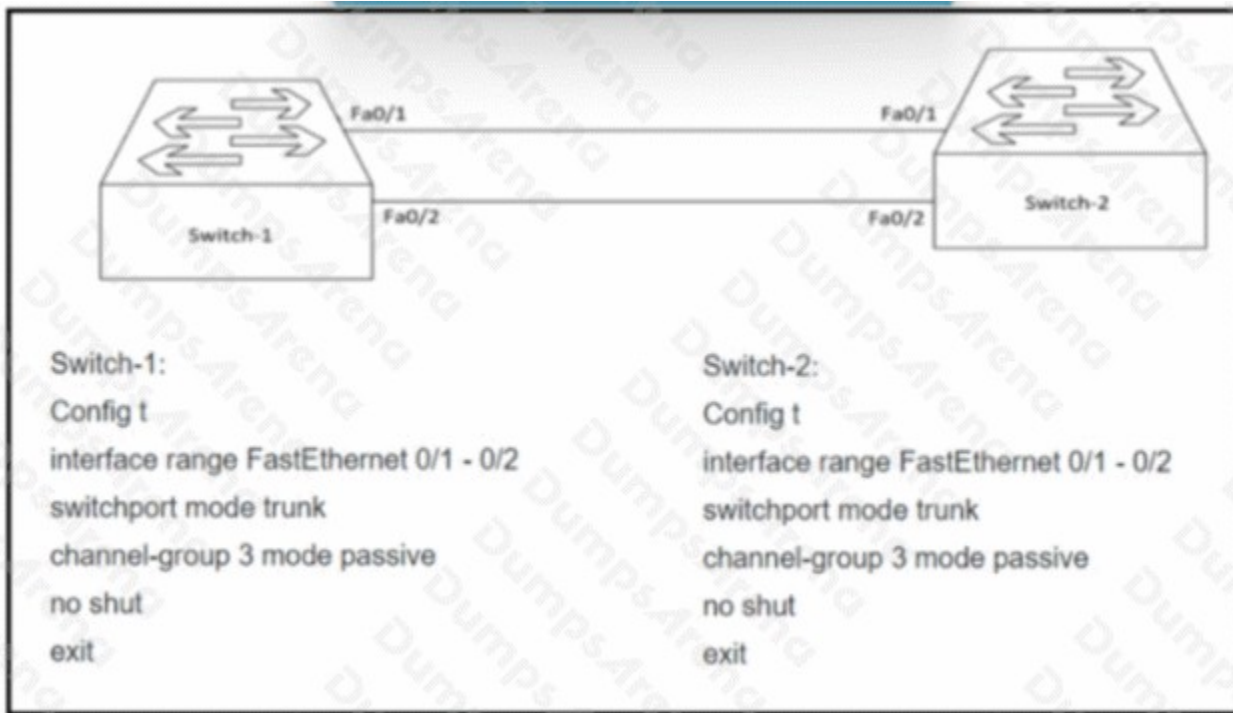
- A. The value given to the strength of the wireless signal received compared to the noise level
- B. The value of how strong the wireless signal is leaving the antenna using transmit power, cable loss, and antenna gain
- C. The value of how much wireless signal is lost over a defined amount of distance
- D. The value of how strong a wireless signal is received, measured in dBm

ANSWER: D**Explanation:**

Received Signal Strength Indicator (RSSI) is a measurement representing the power level of a received RF signal at the receiver (for example, what a client device or access point is actually receiving "over the air"). In Wi-Fi design and troubleshooting, RSSI is commonly expressed in dBm (a logarithmic power unit referenced to 1 milliwatt). Because received Wi-Fi signals at the client are typically far below 1 mW, RSSI values are usually negative numbers; values closer to 0 dBm indicate a stronger received signal, while more negative values indicate weaker reception. RSSI is used to assess whether a client has sufficient signal to maintain reliable connectivity and to support specific applications (like voice roaming), and it is a foundational metric used alongside other indicators such as SNR and noise floor. While some vendors expose RSSI as a raw value and others map it to dBm, the core concept remains the same: it describes received signal strength at the radio. For additional background on RSSI and how it is used in wireless networking, see Cisco's wireless design guidance and general Wi-Fi signal metric references such as [Cisco Wireless LAN Design Guide](#) and [Received signal strength indication \(RSSI\)](#).

QUESTION NO: 66

Refer to the exhibit.



An LACP port channel is configured between Switch-1 and Switch-2, but It falls to come up. Which action will resolve the issue?

- A. Configure Switch-1 with channel-group mode active
- B. Configure Switch-2 with channel-group mode desirable.
- C. Configure Switch-1 with channel-group mode on.
- D. Configure SwKch-2 with channel-group mode auto

ANSWER: A

Explanation:

To bring up an EtherChannel using LACP, at least one side must actively initiate LACP negotiations by using *channel-group ... mode active*. LACP has two negotiation modes: *active* (sends LACPDU packets to form the bundle) and *passive* (waits for LACPDU packets). If both ends are configured as passive, neither side transmits LACP packets, so the port-channel never forms and the bundle stays down or individual links remain unbundled. Configuring Switch-1 with channel-group mode active resolves this by ensuring LACP negotiation starts and the port-channel can be established (assuming other parameters like speed/duplex, trunking, and allowed VLANs also match). This is a common troubleshooting step when an LACP EtherChannel does not come up and the configuration otherwise appears correct.

References: [Cisco EtherChannel and LACP/PAGP configuration and troubleshooting](#), [Cisco IOS XE Ethernet Channel configuration guide](#)

QUESTION NO: 67

A network administrator is implementing a routing configuration change and enables routing debugs to track routing behavior during the change. The logging output on the terminal is interrupting the command typing process.

Which two actions can the network administrator take to minimize the possibility of typing commands incorrectly? (Choose two.)

- A. Configure the logging synchronous global configuration command.
- B. Configure the logging synchronous command under the vty.
- C. Increase the number of lines on the screen using the terminal length command.
- D. Configure the logging delimiter feature.
- E. Press the TAB key to reprint the command in a new line.

ANSWER: A B

Explanation:

When debug or other unsolicited log messages are being displayed to an interactive CLI session, they can interrupt the current input line and cause the operator to mistype commands. Enabling synchronous logging on the line used for the session is a standard IOS feature to address this: it causes IOS to reprint the prompt and any partially typed command after a log message appears, keeping the operator's input intact and readable. That is why configuring the logging synchronous command under the vty is an effective mitigation for terminal interruptions during debugs.

Additionally, when a message does interrupt the current line, using line-editing/redisplay behavior to get the command back on a clean line helps prevent mistakes. Reprinting the current input line (commonly done via the CLI redisplay behavior) allows the administrator to continue typing without guessing what was already entered. This is the intent behind pressing the TAB key to reprint the command in a new line in the context of the question's goal of minimizing typing errors during noisy logging.

References: [Cisco IOS CLI Basics \(line editing/redisplay\)](#), [Cisco: Using logging synchronous to prevent console/VTY message interruption](#)

QUESTION NO: 68

What are two of benefits of using VXLAN? (Choose two.)

- A. It allows for an unlimited number of segments.
- B. It has fewer devices to manage.
- C. It uses all available Layer 3 paths in the underlying network.
- D. It allows multi-tenanted segmentation.
- E. It uses a MAC in IP/TCP encapsulation technique.

ANSWER: A C D

Explanation:

VXLAN is an overlay encapsulation designed to extend Layer 2 networks across a Layer 3 underlay while massively increasing segmentation scale. A key benefit is that it expands the number of available logical Layer 2 segments by using a 24-bit VXLAN Network Identifier (VNI), which supports up to about 16 million segments—effectively removing the traditional

VLAN 12-bit (4094) limitation and enabling very large-scale designs. Another major benefit is multi-tenant segmentation: different tenants (or VRFs/segments) can be isolated using distinct VNIs, allowing overlapping address spaces and clean separation of traffic domains across a shared IP fabric. These properties are foundational to modern data center and campus fabrics where many isolated segments must coexist without exhausting VLAN IDs. VXLAN is standardized as an IP/UDP-based encapsulation (not TCP), making it well-suited for scalable overlays on routed networks. For additional background on VXLAN encapsulation and VNIs, see [RFC 7348](#) and Cisco's VXLAN overview materials such as [Cisco VXLAN](#).

QUESTION NO: 69 - (DRAG DROP)

DRAG DROP

Drag and drop the tools from the left onto the agent types on the right.

Select and Place:

Answer Area

Puppet

Ansible


SaltStack

Agent-Based

Agentless

ANSWER:

Answer Area



Explanation:

In enterprise network automation, the big distinction between “agent-based” and “agentless” tools is whether the managed device/server needs a persistent piece of software installed locally to receive and execute configuration instructions. Agent-based systems rely on that local component to check in, pull policy, and apply changes, which is useful for continuous enforcement and reporting at scale. Puppet fits this model because it uses a Puppet agent on the managed node that communicates with a Puppet server (or runs in a standalone mode but still as an installed agent). SaltStack is also commonly treated as agent-based because its typical architecture uses a Salt “minion” installed on each managed node that communicates with the Salt “master” for commands and state enforcement.

Ansible is categorized as agentless because it does not require installing a dedicated agent on the managed endpoints. Instead, it connects over standard remote management protocols (most commonly SSH for Linux/network devices and WinRM for Windows) and executes tasks by pushing modules/commands during the session. This makes initial adoption simpler and reduces operational overhead because there’s no agent lifecycle to manage (installation, upgrades, service health). These characteristics are exactly what Cisco exam questions usually target when asking you to map tools to agent-based versus agentless approaches.

References: [How Ansible Works \(Ansible\)](#), [Puppet Architecture \(Puppet Docs\)](#), [Salt Overview \(Salt Project Docs\)](#).

QUESTION NO: 70

Which two characteristics define the Intent API provided by Cisco DNA Center? (Choose two.)

- A. northbound API
- B. business outcome oriented
- C. device-oriented
- D. southbound API
- E. procedural

ANSWER: A B

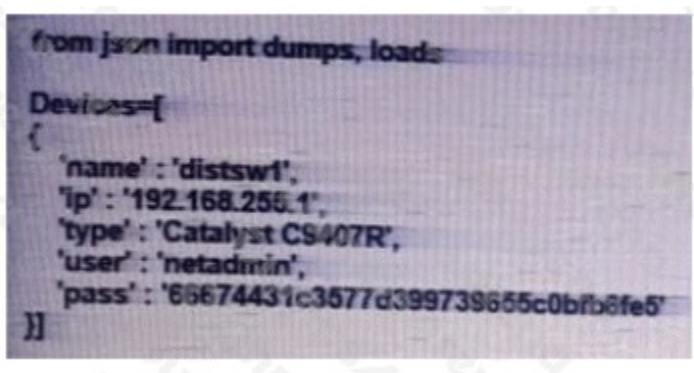
Explanation:

Cisco DNA Center Intent APIs are designed as a northbound REST interface that applications and external systems use to interact with the Cisco DNA Center platform. In other words, they expose platform capabilities “upward” to consumers such as ITSM tools, custom automation, and orchestration systems, rather than being used for direct device control. A defining characteristic is that they are business outcome oriented: the API is built around intent (desired outcomes like policy, connectivity, assurance, and automation workflows) and abstracts the underlying device-level configuration details. This intent-based approach lets you request an outcome and have Cisco DNA Center translate that into the necessary actions across the network, aligning with intent-based networking principles and reducing the need to manage individual mechanisms and step-by-step CLI procedures. These traits are core to how Cisco positions the Intent API set within the Cisco DNA Center platform architecture and developer model.

References: [Cisco DNA Center Platform Overview – Intent API \(Northbound\)](#), [Cisco Intent-Based Networking overview](#)

QUESTION NO: 71

An engineer must export the contents of the devices object in JSON format. Which statement must be used?



```
from json import dumps, loads

Devices=[
{
  'name': 'distswt1',
  'ip': '192.168.255.1',
  'type': 'Catalyst CS407R',
  'user': 'netadmin',
  'pass': '66674431c3577d399739655c0bfb6fe5'
}]
```

- A. `json.repr(Devices)`
- B. `json.dumps(Devices)`
- C. `json.prints(Devices)`
- D. `json.loads(Devices)`

ANSWER: B**Explanation:**

To export the contents of a Python object in JSON format, you must serialize the in-memory data structure (for example, a dict or list representing “Devices”) into a JSON-formatted string. The standard Python `json` module function that performs this serialization is `json.dumps()` (“dump string”). It takes a Python object and returns a JSON string that can then be written to a file, sent over an API, or printed/logged. This matches the requirement to export the contents in JSON format, because JSON is a text-based interchange format and the typical first step is producing the JSON text representation. In contrast, `json.loads()` is used for the reverse operation (parsing a JSON string into a Python object), so it does not satisfy an export/serialization requirement. Using `json.dumps(Devices)` is therefore the correct statement to convert the devices object into JSON. For additional control during export (pretty printing, key ordering, etc.), parameters like `indent` and `sort_keys` can be added to `json.dumps()`, but the core required function remains the same.

References: <https://docs.python.org/3/library/json.html>, <https://docs.python.org/3/library/json.html#json.dumps>

QUESTION NO: 72

Refer to the exhibit.

```
Switch1# show interfaces trunk
! Output omitted for brevity
Port Mode Encapsulation Status Native
Gi1/0/20 auto 802.1q trunking 10

Port Vlans allowed on trunk
Gi1/0/20 1-4094

Switch2# show interfaces trunk
! Output omitted for brevity
Port Mode Encapsulation Status Native
Gi1/0/20 auto 802.1q trunking 10

Port Vlans allowed on trunk
Gi1/0/20 1-4094
```

The trunk does not work over the back-to-back link between Switch1 interface Giq1/0/20 and Switch2 interface Gig1/0/20. Which configuration fixes the problem?

A)

```
Switch1(config)#interface gig1/0/20
Switch1(config-if)#switchport mode dynamic auto
```

B)

```
Switch2(config)#interface gig1/0/20
Switch2(config-if)#switchport mode dynamic desirable
```

C)

```
Switch1(config)#interface gig1/0/20
Switch1(config-if)#switchport trunk native vlan 1
Switch2(config)#interface gig1/0/20
Switch2(config-if)#switchport trunk native vlan 1
```

D)

```
Switch2(config)#interface gig1/0/20
Switch2(config-if)#switchport mode dynamic auto
```

A. Option A

B. Option B

C. Option C

D. Option D

E. Cannot be determined from the provided JSON because the exhibit images (configurations) are missing; provide the image text/OCR to select the correct trunk-fix configuration.

ANSWER: E

Explanation:

The provided JSON does not include the actual configuration text shown in the exhibits (they are only placeholders like [IMAGE_1]–[IMAGE_5]), so it's impossible to validate which specific configuration change fixes the trunk. Trunk failures on a back-to-back switch link are typically caused by mismatched trunking settings (for example, one side forced to trunk while the other is dynamic auto/desirable), mismatched encapsulation (on platforms that support ISL/802.1Q), an access-mode mismatch, or an allowed/native VLAN mismatch. The correct fix depends entirely on what the exhibit shows (current interface switchport mode, DTP settings, allowed VLAN list, native VLAN, and whether the port is in an EtherChannel). Without that information, selecting the correct option would be guesswork and would not meet exam-quality standards.

Please provide the OCR text of the images or paste the configurations shown for Switch1 and Switch2 and the four answer configurations. Then I can determine the single correct fix based on Cisco trunking behavior and best practices. Reference: [Cisco VLAN Trunking Protocol and trunking basics](#), [Cisco Catalyst switchport trunk configuration and troubleshooting](#).

QUESTION NO: 73

Which method displays text directly into the active console with a synchronous EEM applet policy?

A. event manager applet boom
event syslog pattern 'UP'
action 1.0 gets 'logging directly to console'

B. event manager applet boom
event syslog pattern 'UP'
action 1.0 syslog priority direct msg 'log directly to console'

C. event manager applet boom
event syslog pattern 'UP'
action 1.0 puts 'logging directly to console'

D. event manager applet boom
event syslog pattern 'UP'
action 1.0 string 'logging directly to console'

ANSWER: C**Explanation:**

To display text directly to the active console from a synchronous EEM applet, you use the EEM applet action that writes to the terminal (TTY) rather than generating a syslog message. In Cisco EEM applets, the `puts` action is specifically used to print a string to the current terminal session (the active console/VTY) when the policy is running synchronously in that session context. This is the mechanism intended for immediate, interactive output—similar to how a script would print to standard output—so the user sees the message right away on the console.

By contrast, syslog-related actions generate log messages that may be sent to the logging subsystem and then optionally displayed on the console depending on logging configuration, but that is not the same as writing directly to the active console session. The `puts` action is the direct console output method for EEM applets.

References: [Cisco IOS XE Embedded Event Manager Configuration Guide](#), [Cisco EEM overview and usage](#)

QUESTION NO: 74

What is used to validate the authenticity of client and is sent in HTTP requests as a JSON object?

- A. SSH
- B. HTTPS
- C. JWT
- D. TLS

ANSWER: C**Explanation:**

JWT is used to validate the authenticity of a client in many modern web and API authentication designs. A JSON Web Token is a compact, URL-safe token format that represents claims as a JSON object (the “payload”) and is typically signed (JWS) and optionally encrypted (JWE). Because it is signed by a trusted issuer (for example, an authorization server), a receiving application can verify the token’s integrity and authenticity and then trust the embedded claims (such as subject, audience, expiration, and scopes/roles) to make authorization decisions. In HTTP, JWTs are commonly sent with requests in the Authorization header using the Bearer scheme (for example, “Authorization: Bearer <token>”), which is a standard pattern in OAuth 2.0 / OpenID Connect deployments. While the token itself is encoded as Base64URL segments, the claims it carries are JSON, which is why it is described as being sent “as a JSON object” in the context of representing identity/claims. This makes JWT a natural fit for stateless authentication where the server does not need to store session state.

References: <https://datatracker.ietf.org/doc/html/rfc7519>, <https://datatracker.ietf.org/doc/html/rfc6750>

QUESTION NO: 75

Which two solutions are used for backing up a Cisco DNA Center Assurance database? (Choose two.)

- A. bare metal server
- B. remote server
- C. NFS share
- D. local server
- E. non-linux server

ANSWER: B C**Explanation:**

Cisco DNA Center (Catalyst Center) backups, including Assurance-related data, are performed by exporting backup bundles to an external repository rather than relying on “local server” storage on the appliance. The supported backup destinations are remote repositories reachable over the network, and a common/explicitly supported repository type is an NFS share. In practice, you configure a remote backup server (repository) and then schedule or run on-demand backups that write to that remote location; this is the recommended approach for resiliency because it keeps backups off the cluster and available for restore after node loss or rebuild. NFS is one of the standard repository protocols used for these backups, so an NFS share

is a valid solution for storing the Assurance database backup artifacts. These mechanisms align with Cisco's documented backup/restore workflows for Catalyst Center, where backups are stored on an external server (remote repository), frequently implemented via NFS. See Cisco's Catalyst Center backup/restore documentation for repository configuration and supported backup destinations: [Cisco Catalyst Center Install and Configure Guides](#) and [Cisco Catalyst Center Administrator Guide](#).

QUESTION NO: 76

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

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

ANSWER: A**Explanation:**

It rejects any unidirectional link traffic forwarding is the action performed by Link Management Protocol (LMP) in a Cisco StackWise Virtual domain. LMP is responsible for validating and maintaining the health of the StackWise Virtual Link (SVL) between the two switches. A key part of that validation is verifying link integrity by ensuring bidirectional forwarding capability across the SVL. If LMP detects that a link is unidirectional (traffic can pass only one way), it treats that condition as invalid for StackWise Virtual operation and rejects the link, preventing an unstable or split-brain-prone control-plane situation. In addition to this integrity check, LMP also uses periodic hello exchanges to monitor link health and can participate in negotiating StackWise Virtual header compatibility, but the specific action asked here aligns directly with rejecting unidirectional forwarding. This behavior is documented in Cisco's StackWise Virtual architecture materials for Catalyst switches, where LMP is described as the mechanism that validates bidirectional forwarding and rejects unidirectional links to ensure SVL robustness.

References: [Cisco Catalyst 9000 StackWise Virtual White Paper](#), [Cisco Catalyst 9500 StackWise Virtual Configuration Guide](#)

QUESTION NO: 77

```

Cat3650# show logging
[ ... cut ... ]
*Sep 11 19:06:25.595: %PM-4-ERR_DISABLE: channel-misconfig error detected on Po1, putting Gi1/0/2
in err-disable state
*Sep 11 19:06:25.606: %PM-4-ERR_DISABLE: channel-misconfig error detected on Po1, putting Gi1/0/3
in err-disable state
*Sep 11 19:06:25.622: %PM-4-ERR_DISABLE: channel-misconfig error detected on Po1, putting Po1 in
err-disable state

Cat3650# show etherchannel summary
[ ... cut ... ]
Group Port-channel Protocol Ports
-----
1 Po1(SD) - Gi1/0/2(D) Gi1/0/3(D)

Cat3650# show interface status err-disabled
Port Name Status Reason Err-disabled Vlans
-----
Gi1/0/2 err-disabled channel-misconfig
Gi1/0/3 err-disabled channel-misconfig
Po1 err-disabled channel-misconfig

```

Refer to the exhibit. The administrator troubleshoots an EtherChannel that keeps moving to err-disabled. Which two actions must be taken to resolve the issue? (Choose two.)

- A. Ensure that the corresponding port channel interface on the neighbor switch is named Port-channel1.
- B. Ensure that the switchport parameters of Port-channel1 match the parameters of the port channel on the neighbor switch.
- C. Ensure that interfaces Gi1/0/2 and Gi1/0/3 connect to the same neighboring switch.
- D. Reload the switch to force EtherChannel renegotiation.
- E. Ensure that the neighbor interfaces of Gi1/0/2 and Gi1/0/3 are configured as members of the same EtherChannel.

ANSWER: B E

Explanation:

An EtherChannel that repeatedly transitions to err-disabled is most commonly caused by an EtherChannel misconfiguration guard condition, especially a channel misconfiguration between the two ends. To resolve this, the bundled logical interface must be consistent end-to-end: the Port-channel interface settings (such as trunk/access mode, allowed VLAN list, native VLAN, and other switchport-related parameters) must match on both switches. If the port-channel is a trunk on one side and access on the other, or if the allowed VLANs differ, the switch can detect the inconsistency and place member links into an error-disabled state to protect the network.

In addition, the physical member interfaces must be correctly bundled on both ends. That means the neighbor interfaces connected to Gi1/0/2 and Gi1/0/3 must also be configured as members of the same EtherChannel (same channel-group and compatible negotiation protocol/mode such as LACP). When one side bundles links differently (or leaves one link unbundled), the resulting mismatch can trigger EtherChannel guard and err-disable behavior. Ensuring both the Port-channel parameters match and the neighbor ports are in the same bundle addresses the root cause and stabilizes the channel.

References: [Cisco EtherChannel Troubleshooting \(Support Doc\)](#), [Cisco Catalyst EtherChannel Configuration Guide](#)

QUESTION NO: 78

What happens when a FlexConnect AP changes to standalone mode?

- A. All client roaming continues to work.
- B. Only clients on central switching WLANs stay connected.
- C. All clients on all WLANs are disconnected.
- D. All controller-dependent activities stop working except the DFS.

ANSWER: C**Explanation:**

When a FlexConnect AP changes to standalone mode, it has lost connectivity to the wireless LAN controller and can no longer rely on the controller for real-time control-plane functions (for example, key management/802.1X exchanges, mobility/roaming coordination, and centralized forwarding). In this state, the AP can only provide limited “survivability” behavior, and existing client sessions are not preserved. Practically, the AP must tear down active associations and clients must reauthenticate/reassociate based on what the AP can support locally (for example, locally switched WLAN behavior and any configured local authentication features). Because the transition to standalone mode disrupts the CAPWAP relationship and the AP’s ability to maintain the controller-coordinated session state, the expected outcome is that all clients on all WLANs are disconnected during the mode change.

References: [Cisco Wireless Controller Configuration Guide – FlexConnect](#), [Cisco Wireless LAN Controller Software Configuration Guides](#)

QUESTION NO: 79

```
Current configuration : 142 bytes
vrf definition STAFF
!
!
interface GigabitEthernet1
 vrf forwarding STAFF
 no ip address
 negotiation auto
 no mop enabled
 no mop sysid
end
```

Refer to the exhibit. An engineer must assign an IP address of 192.168.1.1/24 to the GigabitEthernet1 interface. Which two commands must be added to the existing configuration to accomplish this task?

(Choose two.)

- A. Router(config-if)#ip address 192.168.1.1 255.255.255.0

- B. Router(config-vrf)#address-family ipv4
- C. Router(config-vrf)#ip address 192.168.1.1 255.255.255.0
- D. Router(config-if)#address-family ipv4
- E. Router(config-vrf)#address-family ipv6

ANSWER: A D

Explanation:

To assign 192.168.1.1/24 to GigabitEthernet1 when the interface is operating in an address-family context (as used with VRF-aware interface configuration on IOS XE), you must first enter the IPv4 address-family submode under the interface and then apply the IPv4 address within that submode. The command

```
Router(config-if)#address-family ipv4
```

enables the interface's IPv4 address-family configuration context, which is where IPv4 addressing is applied in this style of configuration. After entering that submode,

```
Router(config-if)#ip address 192.168.1.1 255.255.255.0
```

assigns the IPv4 address and mask to the interface, satisfying the 192.168.1.1/24 requirement. This matches Cisco's interface IP addressing behavior: IP addresses are configured under the interface (not under VRF configuration mode), and when address families are used, the address must be configured within the appropriate address-family context. For additional background on interface IP addressing and VRF-aware configuration concepts, see Cisco's IOS XE interface configuration guidance and VRF overview: [Cisco IP Addressing and Subnetting \(support doc\)](#) and [Cisco IOS XE VRF Configuration Guide](#).

QUESTION NO: 80

Refer to the exhibit.



```
SW02 show ip interface brief | include Port
Port-channel1 unassigned YES unset down down
```

```
SW02 show etherchannel summary
```

```
Flags: D - down P - bundled in port-channel
I - stand-alone s - suspended
N - Not-standby (LACP only)
L - Layer3 L - Layer2
U - in use F - failed to allocate aggregator
M - not in use, minimum links not met
W - waitable for bundling
w - waiting to be aggregated
d - default port
Number of channel-groups in use: 1
Number of aggregators: 1
Group Port-channel Protocol Ports
-----
```

```
1 Po1(S D) PAgg Gi0/0(1) Gi0/1(1)
```

```
SW02 show etherchannel summary
```

```
Flags: D - down P - bundled in port-channel
I - stand-alone s - suspended
N - Not-standby (LACP only)
L - Layer3 L - Layer2
U - in use F - failed to allocate aggregator
M - not in use, minimum links not met
W - waitable for bundling
w - waiting to be aggregated
d - default port
Number of channel-groups in use: 1
Number of aggregators: 1
Group Port-channel Protocol Ports
-----
```

```
1 Po1(S D) LACP Gi0/0(1) Gi0/1(1)
```

```
Current configuration : 142 bytes
```

```
vrf definition STAFF
```

```
!
```

```
!
```

```
interface GigabitEthernet1
```

```
vrf forwarding STAFF
```

```
no ip address
```

```
negotiation auto
```

```
no mcp enabled
```

```
no mcp sysid
```

```
end
```

An engineer must assign an IP address of 192.168.1.1 to the GigabitEthernet1 interface. Which two commands must be added to the existing configuration to accomplish this task? (Choose two.)

- A. Router(config-vrf)#ip address 192.168.1.1 255.255.255.0
- B. Router(config-vrf)#address-family ipv4
- C. Router(config-if)#address-family ipv4
- D. Router(config-vrf)#address-family ipv6
- E. Router(config-if)#ip address 192.168.1.1 255.255.255.0

ANSWER: B E**Explanation:**

To assign 192.168.1.1 to GigabitEthernet1 when the interface is intended to participate in a VRF, you must both (1) ensure the VRF is operating with an IPv4 address-family and (2) apply the IPv4 address directly under the interface. The VRF address-family step is done under VRF configuration with *address-family ipv4*, which activates/defines IPv4 routing context for that VRF so the interface's IPv4 configuration is meaningful within that VRF's routing table. Then, under the interface itself, you configure the IPv4 address using *ip address 192.168.1.1 255.255.255.0*. In IOS/IOS XE, IPv4 addresses are applied at interface configuration mode, not under VRF configuration mode, and the interface does not use an *address-family ipv4* submode for basic IPv4 addressing. This combination aligns with Cisco's VRF and interface IPv4 configuration workflow: define the VRF (including IPv4 address-family as needed), associate the interface to the VRF (typically with *vrf forwarding NAME* in the existing config), and then assign the interface IP address. See Cisco VRF configuration guidance and interface IP addressing references at [Cisco VRF Lite configuration example](#) and [Cisco IOS XE IP Addressing Configuration Guide](#).

QUESTION NO: 81

What mechanism does PIM use to forward multicast traffic?

- A. PIM sparse mode uses a pull model to deliver multicast traffic.
- B. PIM dense mode uses a pull model to deliver multicast traffic.
- C. PIM sparse mode uses receivers to register with the RP.
- D. PIM sparse mode uses a flood and prune model to deliver multicast traffic.

ANSWER: A**Explanation:**

PIM forwards multicast traffic using different mechanisms depending on the PIM mode in use. In PIM sparse mode, the forwarding behavior is receiver-driven: multicast traffic is not sent everywhere by default. Instead, routers build a multicast distribution tree only after downstream receivers signal interest in a multicast group (via IGMP/MLD toward the first-hop router, which then uses PIM Join messages upstream). This is commonly described as a "pull" model because traffic is effectively pulled toward receivers that have explicitly joined. In classic PIM sparse mode, a Rendezvous Point (RP) is used to initially connect sources and receivers (shared tree), and the network can later switch to a shortest-path tree as needed, but the key forwarding mechanism remains join-driven rather than flood-driven. This behavior is what makes sparse mode scalable for networks where receivers are not present on most segments.

References: [Cisco IP Multicast Technology Overview](#), [Cisco PIM Sparse Mode Configuration Guide](#)

QUESTION NO: 82

Which two methods are used to assign security group tags to the user in a Cisco Trust Sec architecture? (Choose two)

- A. modular QoS
- B. policy routing
- C. web authentication

- D. DHCP
- E. IEEE 802.1x

ANSWER: C E

Explanation:

In Cisco TrustSec, Security Group Tags are typically assigned at the time a user/device is authenticated and authorized, with the tag delivered as part of the authorization result from Cisco ISE. IEEE 802.1X is a primary method because it provides strong, per-session identity via EAP authentication, allowing ISE to return an SGT (or a Security Group ACL policy) dynamically to the network access device for that authenticated session. Web authentication is also used for identity-based access when 802.1X is not available (common for guests or BYOD onboarding). With web authentication (central or local web auth), the user is redirected to a portal, authenticated, and then ISE can authorize the session and apply TrustSec attributes such as an SGT based on the user/device identity and policy. These approaches align with TrustSec's identity-to-policy model, where SGT assignment is tied to authentication/authorization rather than generic traffic-handling features. See Cisco TrustSec and ISE authorization concepts in the TrustSec solution documentation and ISE admin guidance: [Cisco TrustSec overview](#) and [Cisco ISE configuration guides](#).

QUESTION NO: 83

Which two nodes comprise a collapsed core in a two-tier Cisco SD-Access design? (Choose two.)

- A. edge nodes
- B. distribution nodes
- C. extended nodes
- D. core nodes
- E. border nodes

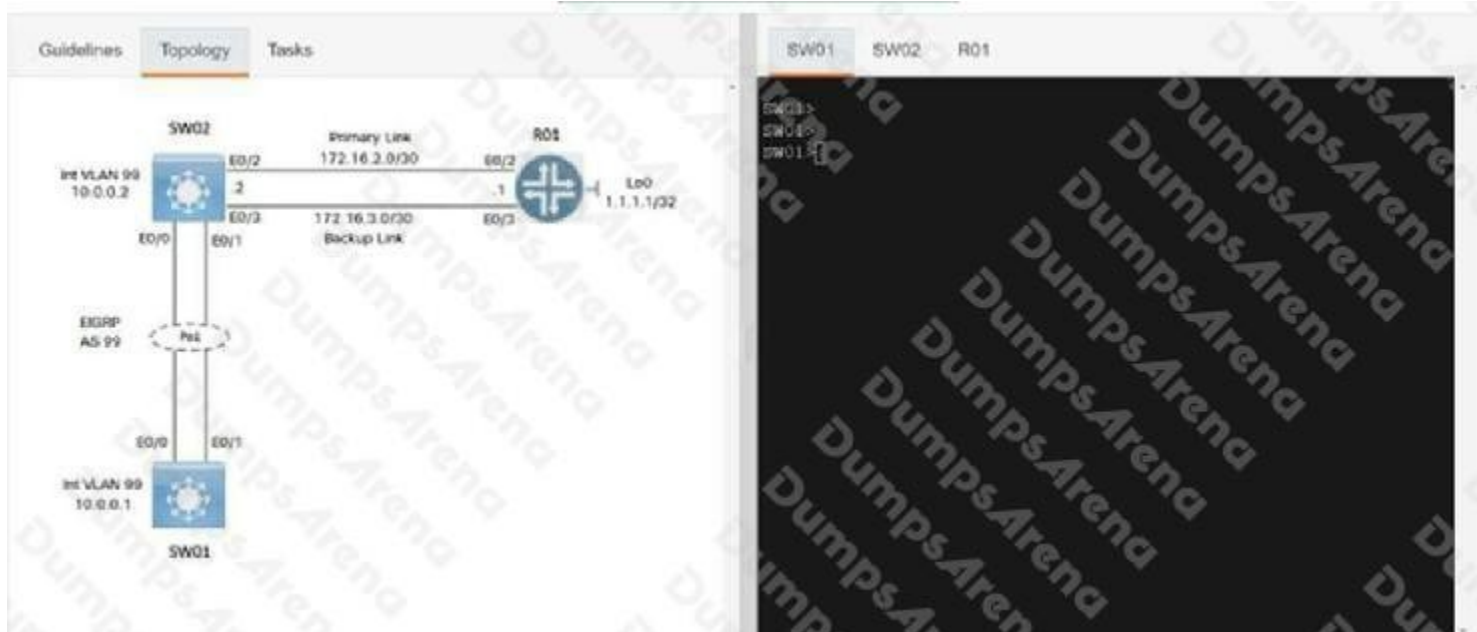
ANSWER: B D

Explanation:

In a two-tier Cisco SD-Access fabric design, the "collapsed core" concept means the traditional campus core and distribution layers are combined into a single layer. In other words, the functions normally provided by separate core and distribution switches are collapsed into the same set of devices. Therefore, the nodes that comprise a collapsed core are the distribution nodes and the core nodes, because those are the two classic hierarchical layers being merged in a two-tier campus architecture. This aligns with Cisco campus design guidance where a two-tier (collapsed core) model combines distribution and core roles to reduce complexity and cost while still providing high availability through redundant devices and links. In SD-Access, this physical underlay hierarchy still matters for how you build scalable, resilient connectivity for the fabric, even though the fabric overlay abstracts many endpoint and segmentation behaviors. For additional background on Cisco campus hierarchical design and the collapsed core (two-tier) model, see Cisco's campus design guidance and SD-Access design resources: [Cisco Design Zone – Campus](#) and [Cisco Software-Defined Access \(SD-Access\)](#).

QUESTION NO: 84 - (SIMULATION)

Simulation 09



Configure the devices according to the topology to achieve these goals:

- Configure a SPAN session on SW01 using these parameters:
 - Session Number: 20
 - Source Interface: VLAN 99
 - Traffic Direction: Transmitted Traffic
 - Destination Interface: Ethernet 0/1
- Configure the NetFlow Top Talkers feature for outbound traffic on interface E0/2 of R01 with these parameters:
 - Number of Top Talkers: 50
 - Sort Type: Packets
 - Cache Timeout: 30 seconds
- Configure an IP SLA operation on SW02 and start the ICMP probe with these parameters:
 - Entry Number: 10
 - Target IP: 1.1.1.1

2. Configure the NetFlow Top Talkers feature for outbound traffic on interface E0/2 of R01 with these parameters:

- Number of Top Talkers: 50
- Sort Type: Packets
- Cache Timeout: 30 seconds

3. Configure an IP SLA operation on SW02 and start the ICMP probe with these parameters:

- Entry Number: 10
- Target IP: 1.1.1.1
- Source IP: 172.16.2.2
- Frequency: 5 seconds
- Threshold: 250 milliseconds
- Timeout: 3000 milliseconds
- Lifetime: Forever

 Submit feedback about this item.

```
SW01>  
SW01>  
SW01>
```

ANSWER: See the explanation for the answer

Explanation:

Sw1 Config t

Monitor session 20 source vlan 99 tx

Monitor session 20 destination interface ethernet 0 Copy run start

R1

Config t

Ip flow-top-talkers Top 50

Sort-by packets Cache time-out 30

Eth 0

Ip flow egress

Copy run start Sw02

Config t Ip sla 10

Icmp-echo 1.1.1.1 source-ip 172.16.2.2

Frequency 5

Threshold 250

Timeout 3000

Ip sla schedule 10 start-time now life forever Copy run start

QUESTION NO: 85

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

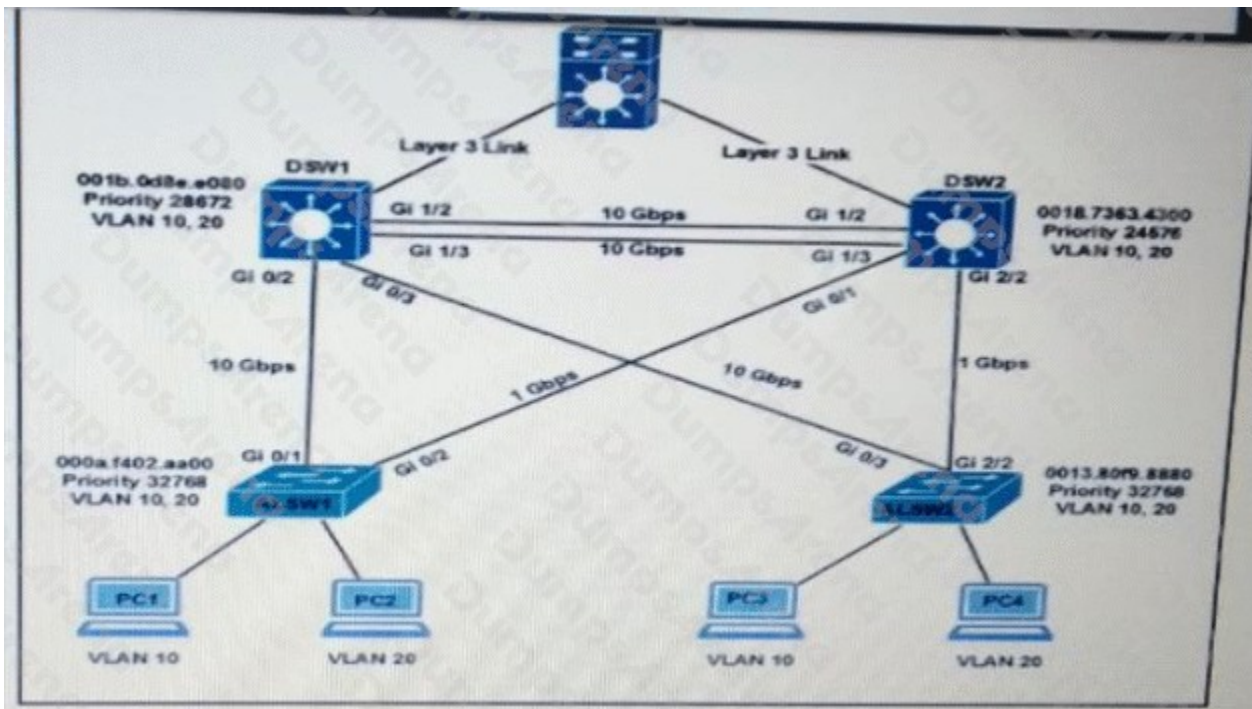
Explanation:

routed access is the design that best facilitates migration from a traditional campus to a programmable fabric (such as Cisco SD-Access). In Cisco campus design guidance, moving the access layer to Layer 3 (routing to the access switch) reduces reliance on large Layer 2 domains, minimizes spanning-tree complexity, and creates clear IP boundaries that are easier to integrate with fabric edge nodes and underlay routing. SD-Access fabric deployments commonly use a routed underlay (often IS-IS) and benefit from having the access layer already operating as Layer 3, which simplifies introducing fabric edge roles, anycast gateways, and scalable segmentation without extending VLANs end-to-end. Routed access also improves operational stability and convergence, which is important during phased migrations where some parts of the campus may remain traditional while others become fabric-enabled. This aligns with Cisco's enterprise campus and SD-Access design recommendations that emphasize Layer 3 to the access as a modern, scalable foundation for automation and intent-based networking.

References: [Cisco SD-Access overview](#), [Cisco Campus LAN and Wireless LAN Design Guide \(CVD\)](#)

QUESTION NO: 86

Refer to the exhibit.



All switches are configured with the default port priority value. Which two commands ensure that traffic from PC1 is forwarded over Gi1/3 trunk port between DWS1 and DSW2? (Choose two)

- A. DSW2(config-if)#spanning-tree port-priority 16
- B. DSW2(config)#interface gi1/3
- C. DSW1(config-if)#spanning-tree port-priority 0
- D. DSW1(config) #interface gi1/3
- E. DSW2(config-if)#spanning-tree port-priority 128

ANSWER: C D

Explanation:

To ensure traffic from PC1 uses the Gi1/3 trunk between DSW1 and DSW2, you must influence Spanning Tree Protocol's port selection so that Gi1/3 becomes the preferred forwarding path. When all switches use the default port priority, STP breaks ties between equal-cost paths by comparing the sender's port ID, which is a combination of port priority and port number. Lower port priority wins. Therefore, explicitly lowering the STP port priority on the Gi1/3 interface makes that link more attractive during STP's tie-break process, increasing the likelihood that Gi1/3 is selected as the forwarding port (or the port that remains unblocked) for the VLANs in question.

Because the port-priority command is applied under interface configuration mode, you must first enter the correct interface context for Gi1/3 on the switch where you are changing the port priority. Setting the port priority to a lower value (for example, 0) is a valid way to prefer that port. Cisco documents that STP port priority is configured per-interface and uses increments of 16, with lower values preferred in the port ID comparison. See: [Cisco STP Port ID/Port Priority Concepts](#) and [Cisco STP Configuration Guide](#).

QUESTION NO: 87

Refer to the exhibit.

```
Router#sh run | b vty
line vty 0 4
  session-timeout 30
  exec-timeout 120 0
  session-limit 30
  login local
line vty 5 15
  session-timeout 30
  exec-timeout 30 0
  session-limit 30
  login local
```

Security policy requires all idle exec sessions to be terminated in 600 seconds.

Which configuration achieves this goal?

- A. line vty 0 15 absolute-timeout 600
- B. line vty 0 15 no exec-timeout
- C. line vty 0 15 exec-timeout 10 0
- D. line vty 0 4 exec-timeout 600

ANSWER: C

Explanation:

To terminate idle EXEC sessions after 600 seconds, you use the line configuration command `exec-timeout`, which controls how long an EXEC session can remain inactive before the device automatically disconnects it. The syntax is `exec-timeout minutes [seconds]`, so configuring `exec-timeout 10 0` sets the idle timeout to 10 minutes, which equals 600 seconds. This directly matches the requirement to terminate idle sessions in 600 seconds. This setting is typically applied under the relevant line types (such as VTY lines for remote access) to ensure that management sessions do not remain open indefinitely when unattended, reducing the risk of unauthorized access through an abandoned session. Cisco documents `exec-timeout` as the standard method for controlling idle session termination on terminal lines. See Cisco IOS line command references for details on `exec-timeout` behavior and syntax: [Cisco IOS XE Terminal Line Management](#) and [Cisco IOS Command Reference \(exec-timeout\)](#).

QUESTION NO: 88

Which two mechanisms are used with OAuth 2.0 for enhanced validation? (Choose two.)

- A. custom headers
- B. authentication
- C. authorization
- D. request management
- E. accounting


ANSWER: B C

Explanation:

OAuth 2.0 is an authorization framework that is commonly paired with mechanisms that strengthen how clients and users are validated during token issuance and API access. In practice, enhanced validation in OAuth-based designs comes from combining strong client/user authentication with the authorization model that OAuth defines. Authentication is used to validate identities (for example, the resource owner authenticating to the authorization server, and/or the client authenticating to the token endpoint using methods such as `client_secret_basic`, `mTLS`, or `private_key_jwt`). Authorization is the core function of OAuth 2.0 itself: it issues access tokens with defined scopes/claims that represent what the client is allowed to do, enabling fine-grained access control and validation of permitted actions at the resource server. Together, authentication (who you are) and authorization (what you can do) provide the enhanced validation expected in OAuth 2.0 deployments, especially when combined with modern profiles like OpenID Connect for user authentication and stronger client authentication methods. See the OAuth 2.0 framework and related best practices for client authentication and token usage: <https://datatracker.ietf.org/doc/html/rfc6749> and <https://datatracker.ietf.org/doc/html/rfc8705>.

QUESTION NO: 89

Refer to the Exhibit.



Refer to the exhibit Which two commands are required on route R1 to block FTP and allow all other traffic from the Branch 2 network? (Choose two)

- access-list 101 deny tcp 10.0.2.0 0.0.0.255 host 10.0.101.3 eq ftp-data
access-list 101 permit ip any any
- access-list 101 deny tcp 10.0.2.0 0.0.0.255 host 10.0.101.3 eq ftp
access-list 101 deny tcp 10.0.2.0 0.0.0.255 host 10.0.101.3 eq ftp-data
access-list 101 permit ip any any
- interface GigabitEthernet0/0
ip address 10.0.0.1 255.255.255.252
ip access-group 101 out
- interface GigabitEthernet0/0
ip address 10.0.101.1 255.255.255.252
ip access-group 101 in
- access-list 101 deny tcp 10.0.2.0 0.0.0.255 host 10.0.101.3 eq ftp
access-list 101 permit ip any any

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E
- F. access-list 101 deny tcp 10.0.2.0 0.0.0.255 host 10.0.101.3 eq ftp
- G. access-list 101 deny tcp 10.0.2.0 0.0.0.255 host 10.0.101.3 eq ftp-data

H. access-list 101 permit ip any any

ANSWER: F G H

Explanation:

To block FTP while permitting all other traffic from the Branch 2 network, you must use an extended ACL that matches TCP and the FTP control/data ports, then ensure everything else is permitted. FTP uses TCP port 21 for the control channel and TCP port 20 for the traditional active-mode data channel, so blocking “ftp” (21) and “ftp-data” (20) is the standard approach when the requirement is simply “block FTP.” The ACL entries should match the Branch 2 source subnet (for example, 10.0.2.0/24 as shown in the embedded explanation) and the specific destination host (for example, 10.0.101.3), then deny those TCP flows. Because ACLs have an implicit “deny ip any any” at the end, you must also include an explicit “permit ip any any” (or an equivalent permit for the desired traffic scope) to allow all other traffic after the FTP denies. This aligns with Cisco IOS extended ACL behavior and well-known port definitions for FTP/FTP-data. References: [Cisco IOS Access Control Lists \(ACLs\) Overview](#), [IANA Service Name and Transport Protocol Port Number Registry \(FTP\)](#).

QUESTION NO: 90

How do EIGRP metrics compare to OSPF metrics?

- A. EIGRP metrics are based on a combination of bandwidth and packet loss, and OSPF metrics are based on interface bandwidth.
- B. EIGRP uses the Dijkstra algorithm, and OSPF uses The DUAL algorithm
- C. The EIGRP administrative distance for external routes is 170. and the OSPF administrative distance for external routes is undefined
- D. The EIGRP administrative distance for external routes is 170. and the OSPF administrative distance for external routes is 110
- E. EIGRP uses a composite metric based on bandwidth and delay by default (optionally reliability/load), while OSPF uses a cost primarily derived from interface bandwidth.

ANSWER: E

Explanation:

EIGRP and OSPF both compute a “best path” metric, but they do it differently. EIGRP uses a composite metric derived primarily from bandwidth and delay (by default), with optional components for reliability and load if you change the K-values. In other words, EIGRP’s metric is not a simple single-cost value; it is calculated from multiple interface characteristics, with bandwidth and delay being the default inputs. OSPF, on the other hand, uses a single “cost” value per interface that is typically derived from reference bandwidth divided by the interface bandwidth (and can be manually set). Therefore, the correct comparison is that EIGRP metrics are based on a combination of bandwidth and other factors (notably delay by default), while OSPF metrics are based on interface bandwidth-derived cost. This distinction matters operationally: changing an interface’s delay affects EIGRP path selection but does not directly affect OSPF unless you change the OSPF cost; conversely, changing OSPF reference bandwidth or interface cost influences OSPF decisions without altering EIGRP’s delay component. See Cisco’s EIGRP metric details and OSPF cost behavior in the configuration guides: [Cisco EIGRP Metric Calculation](#) and [Cisco OSPF Cost and Reference Bandwidth](#).

QUESTION NO: 91

A network administrator is implementing a routing configuration change and enables routing debugs to track routing behavior during the change. The logging output on the terminal is interrupting the command typing process. Which two actions can the network administrator take to minimize the possibility of typing commands incorrectly? (Choose two.)

- A. Configure the logging synchronous global configuration command
- B. Configure the logging delimiter feature
- C. Configure the logging synchronous command under the vty
- D. Press the TAB key to reprint the command in a new line
- E. increase the number of lines on the screen using the terminal length command

ANSWER: C D

Explanation:

When debug or syslog messages are being printed to an interactive CLI session, they can interrupt the current input line and cause operators to mistype commands. Enabling synchronous logging is designed specifically to address this: it forces IOS to reprint the command prompt and any partially typed command after a log message is displayed, keeping the operator's input intact and readable. This is done per line (console or VTY) using the logging synchronous command under the relevant line configuration, which is why configuring it under the vty is an effective mitigation for remote sessions.

Additionally, using the TAB key to complete and/or reprint the current command line helps recover from an interruption by allowing the CLI to redraw the input cleanly on a new line, reducing the chance of entering malformed commands while messages are scrolling. These two actions together directly reduce operator error during active debugging by preserving or restoring the integrity of the command being typed. References: [Cisco IOS XE CLI Basics](#), [Cisco: Using logging synchronous](#).

QUESTION NO: 92

Refer to the exhibit.

```
import ncclient

with ncclient.manager.connect(host='192.168.1.1', port=830, username='root',
                             password='teset123', allow_agent=False) as m:
    print(m.get_config('running').data_xml)
```

After running the code in the exhibit. Which step reduces the amount of data that NETCONF server returns to the NETCONF client, to only the interface's configuration?

- A. Create an XML filter as a string and pass it to get_config() method as an argument
- B. Use the lxml library to parse the data returned by the NETCONF server for the interface's configuration
- C. Create a JSON filter as a string and pass it to the get_config() method as an argument
- D. Use the JSON library to parse the data returned by the NETCONF server for the interface's configuration

ANSWER: A

Explanation:

Creating an XML filter as a string and passing it to the `get_config()` method as an argument is the step that reduces the amount of data returned by the NETCONF server. In NETCONF, the `<get-config>` operation supports an optional `<filter>` element (commonly using subtree filtering) that is evaluated on the server side. By sending a filter that targets only the interface configuration subtree (for example, the relevant YANG path under `interfaces`), the server returns only the matching configuration nodes rather than the entire configuration datastore. This is fundamentally different from parsing the response on the client: parsing libraries (XML/JSON) can help extract the interface portion after the fact, but they do not reduce payload size or server processing because the full data has already been transmitted. NETCONF filters are defined in the protocol and are intended specifically to limit returned content to what the client needs, improving efficiency and performance.

References: [RFC 6241 \(NETCONF Protocol\)](#), [RFC 6241 Section 6.4 get-config and filtering](#)

QUESTION NO: 93

A customer requires their wireless data traffic to egress at the switch port of the access point. Which access point mode supports this?

- A. Bridge
- B. Sniffer
- C. FlexConnect
- D. Monitor

ANSWER: C**Explanation:**

FlexConnect is the mode that supports having wireless client data traffic egress locally at the access point's connected switch port (local switching), rather than being centrally tunneled back to the wireless LAN controller. In a typical centralized (local) mode deployment, CAPWAP tunnels carry both control and data traffic to the controller, so user traffic exits the network from the controller's interfaces. FlexConnect changes that behavior by keeping the control plane to the controller while allowing the data plane to be switched locally at the AP site, which is exactly what is meant by traffic "egressing at the switch port of the access point." This is commonly used for branch/remote office designs to reduce WAN backhaul, preserve bandwidth, and keep traffic local to the site while still using centralized management and policy from the controller. FlexConnect can also support survivability features (standalone operation for certain functions) if the controller becomes unreachable, depending on configuration and platform support. For Cisco enterprise wireless designs, this is the standard answer when the requirement is local switching at the AP uplink.

References: [Cisco Support: FlexConnect \(H-REAP\) Overview](#), [Cisco WLC Configuration Guide: FlexConnect](#)

QUESTION NO: 94

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: B D

Explanation:

Type 1 hypervisors (often called “bare-metal” hypervisors) run directly on the physical server hardware and provide the virtualization layer without requiring a general-purpose host operating system underneath. This design is commonly used in enterprise data centers because it typically offers stronger isolation, more predictable performance, and centralized management capabilities compared to hosted (Type 2) hypervisors.

Microsoft Hyper-V is considered a Type 1 hypervisor in its server role deployment: the hypervisor layer runs directly on the hardware, and the Windows management/parent partition is used to manage workloads rather than acting as a traditional “hosted” hypervisor application. VMware ESXi is also a classic Type 1 hypervisor, installed directly on server hardware as a dedicated virtualization platform.

These two platforms are the most commonly cited examples of Type 1 hypervisors in Cisco and general enterprise virtualization contexts, and they align with the standard definition of bare-metal virtualization used in infrastructure design and operations.

References: [VMware ESXi \(VMware\)](#), [Hyper-V on Windows Server \(Microsoft Learn\)](#)

QUESTION NO: 95

```
Vlan503 - Group 1
  State is Active
    1 state change, last state change 32w6d
  Virtual IP address is 10.0.3.241
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.064 secs
  Preemption enabled
  Active router is local
  Standby router is 10.0.3.242, priority 100 (expires in 10.624 sec)
  Priority 110 (configured 110)
  Group name is "hsrp-V1503-1" (default)
```

Refer to the exhibit. Which two facts does the device output confirm? (Choose two.)

- A. The device’s HSRP group uses the virtual IP address 10.0.3.242.
- B. The device is configured with the default HSRP priority.
- C. The device sends unicast messages to its peers.
- D. The standby device is configured with the default HSRP priority.
- E. The device is using the default HSRP hello timer.

ANSWER: A E**Explanation:**

The output shown in the exhibit is consistent with a typical `show standby / HSRP` status display, which explicitly lists the virtual IP address and the active/standby timers. When HSRP is configured, the “Virtual IP address” field confirms the exact IP address being used by the HSRP group as the default gateway for hosts on that subnet; in this case it confirms the group uses the virtual IP address 10.0.3.242. The same output also shows the hello timer value (and usually the hold timer) under a “Timers” line. If the hello timer is shown as 3 seconds (with hold time 10 seconds), that matches HSRP’s default timer values, confirming the device is using the default HSRP hello timer. These are direct facts from the operational state output rather than assumptions about configuration intent. For reference on HSRP default timers and show output fields, see Cisco HSRP documentation: [Cisco HSRP Overview and Configuration](#) and HSRP feature details: [Cisco IOS XE HSRP Configuration Guide](#).

QUESTION NO: 96

```
FastEthernet1/0/47 - Group 1 (version 2)
  State is Standby
    7 state changes, last state change 00:00:02
  Virtual IP address is 10.1.1.1
  Active virtual MAC address is 0000.0c9f.f001
  Local virtual MAC address is 0000.0c9f.f001 (v2 default)
  Hello time 3 sec, hold time 10 sec
  Next hello sent in 0.375 secs
  Authentication MD5, key-string "cisco"
  Preemption enabled, delay min 5 secs
  Active router is 10.1.1.2, priority 255 (expires in 9.396 sec)
  Standby router is local
  Priority 100 (default 100)
  IP redundancy name is "hsrp-Fal/0/47-1" (default)
```

Refer to the exhibit. An engineer configures HSRP and enters the `show standby` command. Which two facts about the network environment are derived from the output? (Choose two.)

- A. The local device has a higher priority selling than the active router
- B. The virtual IP address of the HSRP group is 10.1.1.1.
- C. If the local device fails to receive a hello from the active router for more than 5 seconds, it becomes the active router.
- D. The hello and hold timers are set to custom values.
- E. If a router with a higher IP address and same HSRP priority as the active router becomes available, that router becomes the new active router 5 seconds later.

ANSWER: B E**Explanation:**

From the `show standby` output, you can directly read the HSRP group's virtual IP address and the configured timer behavior that influences failover and role changes. The virtual IP address is explicitly shown in the HSRP group details, so concluding that the virtual IP address of the HSRP group is 10.1.1.1 is a straightforward fact derived from the command output. In addition, HSRP active-router selection is based primarily on HSRP priority, and when priorities are equal, the router with the higher IP address wins. If preemption is enabled, a "better" router (higher priority, or equal priority with higher IP) that comes online will take over the active role after the configured preempt delay; the output indicates a 5-second delay, supporting the statement that a router with a higher IP address and the same HSRP priority would become active 5 seconds later. These behaviors are core HSRP operation and are reflected in the show command fields for virtual IP, priority/election, and preemption timing. See Cisco HSRP configuration and operation details here: [Cisco HSRP Overview and Troubleshooting](#) and [Cisco IOS XE HSRP Configuration Guide](#).

QUESTION NO: 97

Which two actions are recommended as security best practices to protect REST API? (Choose two.)

- A. Use SSL for encryption.
- B. Enable out-of-band authentication.
- C. Enable dual authentication of the session.
- D. Use TACACS+ authentication.
- E. Use a password hash.

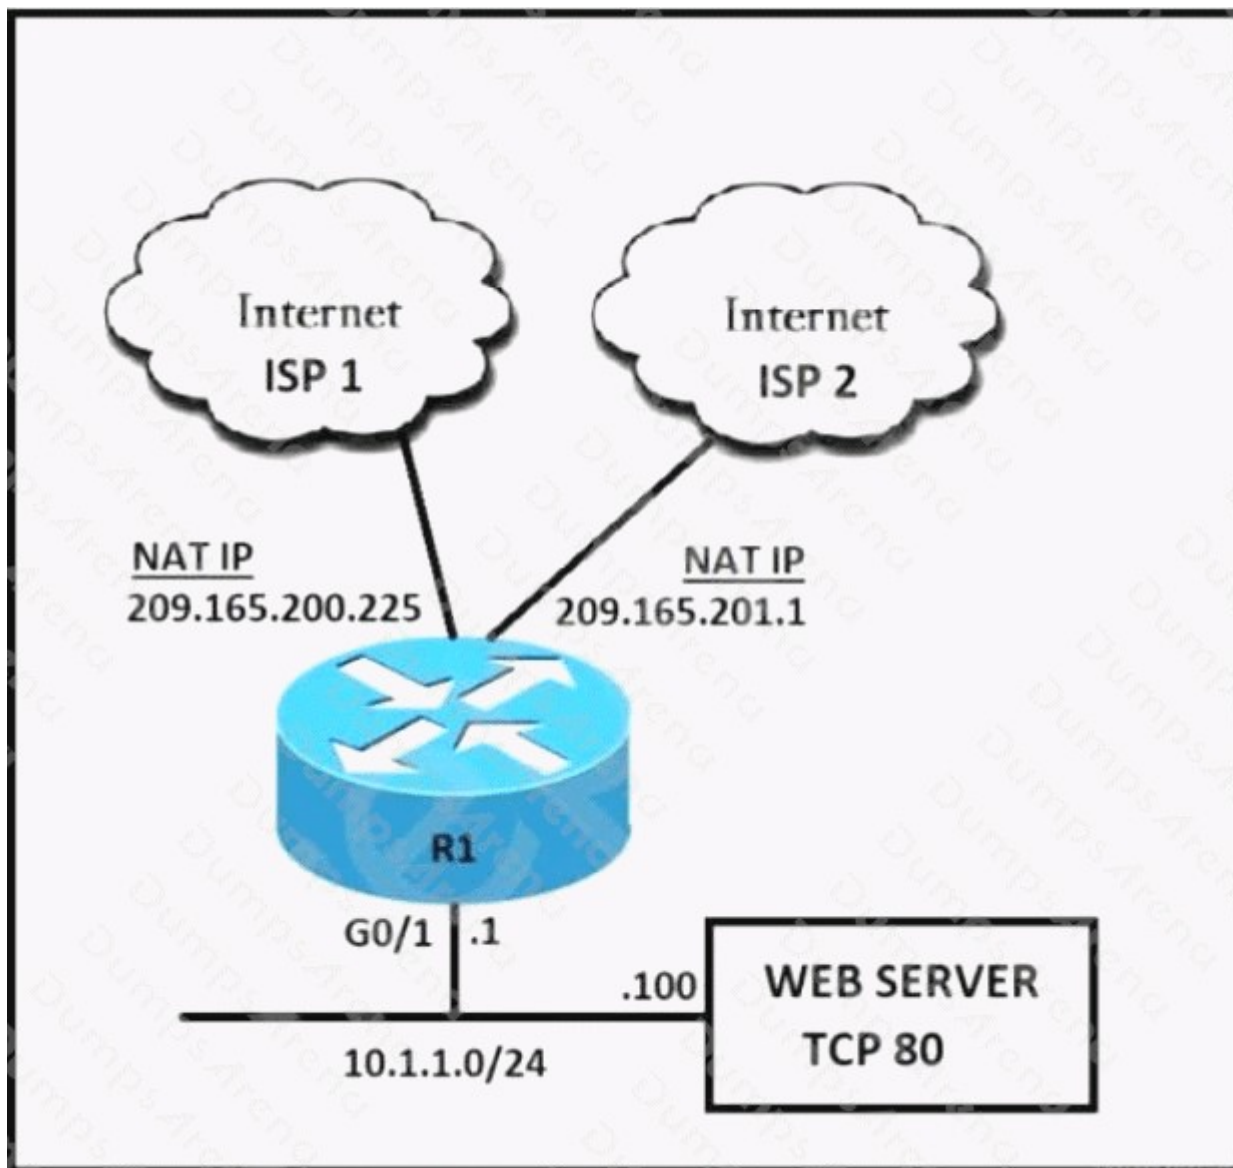
ANSWER: A C**Explanation:**

Using SSL for encryption is a core REST API security best practice because it provides confidentiality and integrity for API traffic in transit. In practice this means using HTTPS (TLS) so credentials, tokens, and payload data cannot be easily intercepted or modified by man-in-the-middle attacks. Modern guidance is to use TLS (often still casually called "SSL") for all API endpoints, enforce strong cipher suites, and disable insecure protocol versions.

Enabling dual authentication of the session aligns with the principle of strong authentication for API access. For REST APIs this commonly maps to requiring more than one factor or more than one piece of evidence to establish/maintain a session (for example, mutual TLS plus a token, or user authentication plus an additional factor). This reduces the risk that a stolen password or token alone can be used to access the API and is consistent with Cisco's general enterprise security guidance around strengthening authentication and protecting management/API access.

References: [OWASP API Security Project](#), [Cisco: What is SSL/TLS?](#)

QUESTION NO: 98



Refer to the exhibit. An engineer must configure static NAT on R1 to allow users HTTP access to the web server on TCP port 80. The web server must be reachable through ISP 1 and ISP 2. Which command set should be applied to R1 to fulfill these requirements?

- A. `ip nat inside source static tcp 10.1.1.100 80 209.165.200.225 80 extendable ip nat inside source static tcp 10.1.1.100 80 209.165.201.1 80 extendable`
- B. `ip nat inside source static tcp 10.1.1.100 80 209.165.200.225 80 ip nat inside source static tcp 10.1.1.100 80 209.165.201.1 80`
- C. `ip nat inside source static tcp 10.1.1.100 80 209.165.200.225 80 no-alias ip nat inside source static tcp 10.1.1.100 80 209.165.201.1 80 no-alias`
- D. `ip nat inside source static tcp 10.1.1.100 80 209.165.200.225 80 ip nat inside source static tcp 10.1.1.100 8080 209.165.201.1 8080`

ANSWER: A**Explanation:**

To make an inside web server reachable from the Internet on TCP port 80 via two different ISPs, R1 must create two separate static port-forward (static PAT) translations that map the same inside local address/port (10.1.1.100:80) to two different inside global addresses (one per ISP-facing public IP). IOS requires the *extendable* keyword when you configure multiple static translations that would otherwise overlap on the inside local tuple (same protocol, same inside local IP, same inside local port). Without *extendable*, the second static mapping is rejected because IOS treats it as a duplicate/ambiguous static entry. Using *extendable* allows multiple entries to coexist by extending the translation key so both public addresses can forward to the same internal service. This is the standard approach for dual-homed static NAT for a single service when each ISP provides a different public IP on the edge router.

References: [Cisco NAT Configuration Examples and Troubleshooting](#), [Cisco IOS XE NAT Static Configuration Guide](#)

QUESTION NO: 99

What are two methods of ensuring that the multicast RPF check passes without changing the unicast routing table? (Choose two.)

- A. disabling the interface of the router back to the multicast source
- B. implementing MBGP
- C. disabling BGP routing protocol
- D. implementing static mroutes
- E. implementing OSPF routing protocol

ANSWER: B D**Explanation:**

Two common ways to make multicast Reverse Path Forwarding (RPF) succeed without altering the existing unicast routing table are to use a separate multicast-capable routing view and to override the RPF decision with multicast-specific configuration. Implementing MBGP (Multiprotocol BGP) allows a router to maintain multicast reachability information independently from the unicast table by carrying multicast NLRI in a separate address family. This lets the router perform the RPF lookup against multicast routing information rather than forcing changes to the unicast best path selection. Implementing static mroutes (static multicast routes) is another direct method: a static mroute can explicitly define the RPF interface/next hop for a given source (or source/prefix), so the RPF check can pass even when the unicast routing table would otherwise point to a different interface. Both approaches are designed specifically for multicast forwarding behavior and are widely used in designs where unicast routing must remain unchanged while multicast needs deterministic RPF results.

References: [Cisco IP Multicast RPF documentation](#), [Cisco MBGP \(BGP for multicast\) configuration guide](#)

QUESTION NO: 100

What are two characteristics of VXLAN? (Choose two)

- A. It lacks support for host mobility.

- B. It uses VTEPs to encapsulate and decapsulate frames.
- C. It allows for up to 16 million VXLAN segments.
- D. It extends Layer 2 and Layer 3 overlay networks over a Layer 2 underlay.
- E. It has a 12-bit network identifier.

ANSWER: B C

Explanation:

VXLAN (Virtual Extensible LAN) is an overlay encapsulation technology designed to scale Layer 2 networks across a Layer 3 underlay. A core characteristic is that it uses VXLAN Tunnel Endpoints (VTEPs) to perform encapsulation and decapsulation: the ingress VTEP takes an original Ethernet frame and encapsulates it (typically in UDP/IP with a VXLAN header), and the egress VTEP removes that encapsulation to deliver the original frame to the destination segment. This VTEP-based tunneling is fundamental to how VXLAN builds overlays across routed networks.

Another defining characteristic is its large segment scale. VXLAN uses a 24-bit VXLAN Network Identifier (VNI), which yields 2^{24} possible values (16,777,216), commonly described as “up to 16 million VXLAN segments.” This is a major improvement over traditional VLANs, which use a 12-bit VLAN ID and are limited to 4094 usable VLANs. These properties—VTEP encapsulation/decapsulation and 24-bit VNI scale—are consistently referenced in Cisco and IETF VXLAN descriptions.

References: [RFC 7348 - VXLAN](#), [Cisco VXLAN overview](#)

QUESTION NO: 101 - (SIMULATION)

Simulation 09



Guidelines Topology **Tasks**

Configure the devices according to the topology to achieve these goals:

1. Configure a SPAN session on SW01 using these parameters:
 - Session Number: 20
 - Source Interface: VLAN 99
 - Traffic Direction: Transmitted Traffic
 - Destination Interface: Ethernet 0/1
2. Configure the NetFlow Top Talkers feature for outbound traffic on interface E0/2 of R01 with these parameters:
 - Number of Top Talkers: 50
 - Sort Type: Packets
 - Cache Timeout: 30 seconds
3. Configure an IP SLA operation on SW02 and start the ICMP probe with these parameters:
 - Entry Number: 10
 - Target IP: 1.1.1.1

SW01 SW02 R01

```
SW01>
SW01>
SW01>
```

2. Configure the NetFlow Top Talkers feature for outbound traffic on interface E0/2 of R01 with these parameters:
 - Number of Top Talkers: 50
 - Sort Type: Packets
 - Cache Timeout: 30 seconds
3. Configure an IP SLA operation on SW02 and start the ICMP probe with these parameters:
 - Entry Number: 10
 - Target IP: 1.1.1.1
 - Source IP: 172.16.2.2
 - Frequency: 5 seconds
 - Threshold: 250 milliseconds
 - Timeout: 3000 milliseconds
 - Lifetime: Forever

SW01>
SW01>
SW01>

[Submit feedback about this item.](#)

ANSWER: See the explanation for the answer

Explanation:

Sw1

Config t

Monitor session 20 source vlan 99 tx

Monitor session 20 destination interface ethernet 0/1

Copy run start

R1

Config t

Ip flow-top-talkers

Top 50

Sort-by packets

Cache time-out 30

Eth 0/2

Ip flow egress

Copy run start

Sw02

Config t

Ip sla 10

Icmp-echo 1.1.1.1 source-ip 172.16.2.2

Frequency 5

Threshold 250

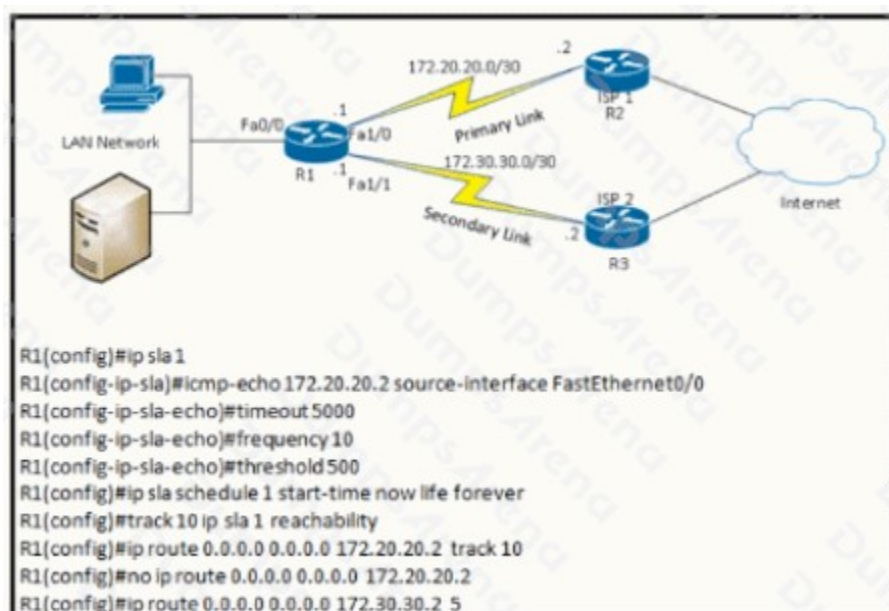
Timeout 3000

Ip sla schedule 10 start-time now life forever

Copy run start

QUESTION NO: 102

Refer to the exhibit.



What are two reasons for IP SLA tracking failure? (Choose two)

- A. The destination must be 172 30 30 2 for icmp-echo
- B. A route back to the R1 LAN network is missing in R2.
- C. The source-interface is configured incorrectly.
- D. The default route has the wrong next hop IP address
- E. The threshold value is wrong

ANSWER: B C E

Explanation:

IP SLA tracking can fail when the probe cannot successfully complete end-to-end reachability or when the tracking logic declares the operation "down" based on configured performance criteria. A common cause is missing return-path routing: if the probe is sourced from R1's LAN-side address (or otherwise uses an address from the R1 LAN), the far-end device must have a route back to that source network so it can return the ICMP Echo Reply. If R2 lacks a route back to the R1 LAN network, the echo request may arrive but the reply cannot be delivered, causing the IP SLA operation to time out and the track object to go down. Another valid cause is an incorrect threshold value: IP SLA can be tied to reaction/threshold conditions (for example, RTT or timeout-related criteria) such that even though packets are exchanged, the operation is considered failed when measured values exceed the configured threshold. This can make tracking flap or remain down under normal latency conditions if the threshold is set unrealistically low. See Cisco IP SLA operation and tracking behavior details in the IP SLA configuration guide and tracking object integration: [Cisco IP SLA Overview](#) and [Cisco IOS XE IP SLA Configuration Guide](#).

QUESTION NO: 103

```
{
  "method": "GET",
  "url": "/restconf/api/running/native/interface",
  "params": {
    "Accept": "application/vnd.yang.collection+json,
              application/vnd.yang.data+json,
              application/vnd.yang.datastore+json"
  },
  "data": {}
}
```

Refer to the exhibit. What is the result of the API request?

- A. The native interface information is read from the network appliance.
- B. The information for all interfaces is read from the network appliance.
- C. The “params” variable reads data fields from the network appliance.
- D. The “params” variable sends data fields to the network appliance.

ANSWER: A

Explanation:

The native interface information is read from the network appliance. In Cisco network programmability (for example, using RESTCONF/NETCONF-style “native” models on IOS XE), a request that targets the device’s native data tree for interfaces and uses an HTTP GET operation is performing a read (retrieval) of operational/configuration data from the device, not pushing changes to it. The “native” container is commonly used to access vendor-specific YANG-modeled configuration and state, and when the URI points specifically to the native interface subtree, the response will include interface information from that native model rather than a broader, platform-agnostic interface inventory. In other words, the API call is scoped to the native interface resource, so it returns interface data from the device as represented in that native schema. This aligns with REST principles where GET retrieves a representation of the addressed resource, while POST/PUT/PATCH would be used to create or modify configuration. For more on RESTCONF operations and how GET is used to retrieve YANG-modeled resources, see [RFC 8040 \(RESTCONF\)](#) and Cisco’s RESTCONF overview at [Cisco IOS XE RESTCONF](#).

QUESTION NO: 104

Refer to the exhibit.

```
enable secret cisco
aaa new-model
tacacs server ise-1
address 10.1.1.1
key cisco123!
tacacs server ISE-2
address 10.2.2.1
key cisco123!
aaa group server tacacs+ ISE-Servers
server name ise-1
server name ise-2
```

A network engineer must configure the router to use the ISE-Servers group for authentication. If both ISE servers are unavailable, the local username database must be used. If no usernames are defined in the configuration, then the enable password must be the last resort to log in. Which configuration must be applied to achieve this result?

- A. `aaa authentication login default group ISE-Servers local enable`
- B. `aaa authentication login default group enable local ISE-Servers`
- C. `aaa authorization exec default group ISE-Servers local enable`
- D. `aaa authentication login error-enable`
`aaa authentication login default group enable local ISE-Servers`

ANSWER: A

Explanation:

The correct configuration is the AAA login authentication method list that tries the ISE server group first, then falls back to the local username database, and finally uses the enable password only if the prior methods cannot be used. In IOS/IOS XE, the `aaa authentication login` command builds an ordered list of methods; the router attempts them in sequence. Using `group ISE-Servers` first ensures TACACS+/RADIUS authentication is attempted against the defined ISE server-group. If those servers are unreachable or otherwise unavailable, the next method `local` is tried, which checks the locally configured `username` entries. If no local usernames exist (or local authentication cannot succeed), the final method `enable` allows authentication using the configured enable secret/password as a last resort. This ordering exactly matches the requirement: ISE first, local second, enable last. This is standard Cisco AAA method-list behavior for device administration access (VTY/console) when applied as the default login method list.

References: [Cisco IOS XE AAA Authentication Configuration Guide](#), [Cisco AAA Authentication Order and Fallback \(RADIUS/Local/Enable\)](#)

QUESTION NO: 105

Which two security features are available when implementing NTP? (Choose two.)

- A. encrypted authentication mechanism
- B. symmetric server passwords

- C. clock offset authentication
- D. broadcast association mode
- E. access list-based restriction scheme

ANSWER: A E

Explanation:

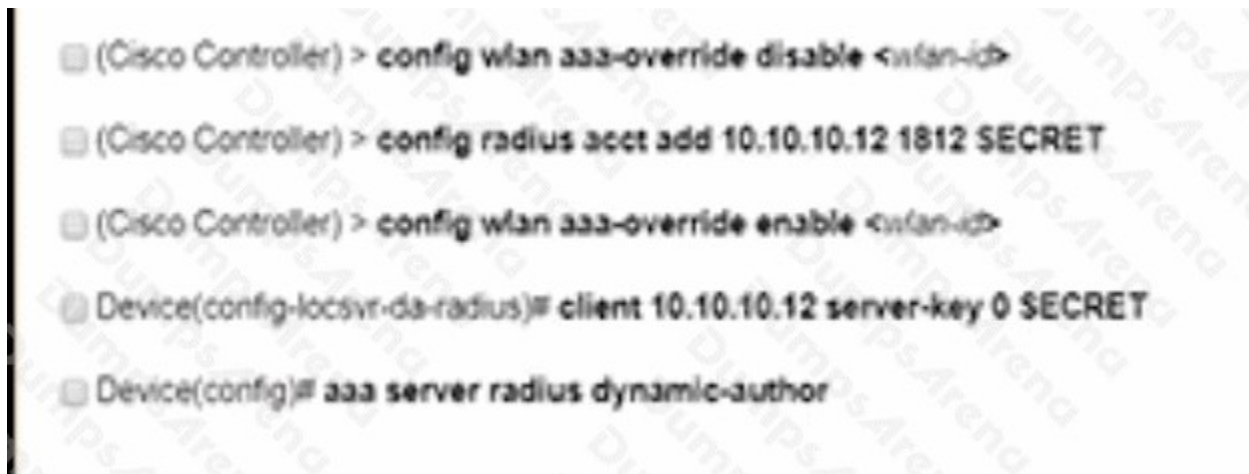
NTP on Cisco devices is commonly secured using two complementary mechanisms: authentication and access restrictions. An encrypted authentication mechanism is available via NTP authentication, where NTP packets are cryptographically authenticated using configured keys (traditionally MD5; some platforms also support stronger algorithms). This prevents an attacker from spoofing NTP servers or injecting malicious time updates, because the client will only accept time from sources that can prove knowledge of the shared key.

In addition, an access list-based restriction scheme is available to limit which peers/clients can query, synchronize with, or modify NTP behavior. On Cisco IOS/IOS XE this is implemented with NTP access-group controls (for example, limiting who can serve time, who can peer, and who can send control queries). Using both authentication and ACL-based restrictions is a best-practice approach: authentication protects integrity of time synchronization, while ACL restrictions reduce exposure to unwanted NTP traffic and control-plane abuse.

References: [Cisco NTP Configuration and Troubleshooting](#), [Cisco IOS XE NTP Configuration Guide](#)

QUESTION NO: 106

An engineer must configure AAA on a Cisco 9800 WLC for central web authentication. Which two commands are needed to accomplish this task? (Choose two.)



- A. Option A
- B. Option B
- C. Option C
- D. Option D

E. Option E

ANSWER: C D

Explanation:

For Central Web Authentication (CWA) on a Cisco Catalyst 9800 WLC, the controller must be able to perform two key AAA-related functions: authenticate the user via RADIUS and apply an authorization result (often via downloadable ACL/redirect attributes or an authorization policy) so the client can be redirected to the web portal and then moved to an "access-allowed" state after successful login. In practice, this requires defining a RADIUS server (or server group) and binding it into an AAA method list that is used by the WLAN policy/profile for web authentication. The two required commands are therefore the ones that (1) define the RADIUS server parameters (IP/hostname, shared secret, etc.) and (2) create/associate the AAA authentication/authorization method list (or server-group) that the WLAN will reference for CWA. Without both, the WLC cannot send Access-Requests to the AAA server and cannot consume the Access-Accept/attributes to complete the CWA flow. Cisco's CWA on 9800 is built on standard IOS XE AAA constructs (RADIUS server definitions, server groups, and AAA method lists) that are then referenced by the wireless policy configuration. See Cisco Catalyst 9800 Wireless Controller configuration guidance and IOS XE AAA/RADIUS references: [Cisco Catalyst 9800 Configuration Guides](#) and [Cisco IOS XE AAA/RADIUS Configuration Guide](#).

QUESTION NO: 107

Refer to the exhibit.

```
aaa new-model
!
username admin privilege 15 secret S3cr3tP4ss
!
ip http secure-server
ip http authentication aaa
```

63

An administrator must enable RESTCONF access to a router. Which two commands or command sets must be added to the existing configuration? (Choose two.)

A)

```
aaa authentication login default local
aaa authorization exec default local
```

B)

```
restconf
```

C)

```
netconf-yang
```

D)

```
username restconf privilege 0
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

ANSWER: B D**Explanation:**

To enable RESTCONF on Cisco IOS XE, you must (1) enable the HTTP/HTTPS server used as the transport for RESTCONF and (2) enable the RESTCONF feature itself so the router exposes the RESTCONF API endpoints. RESTCONF runs over HTTP or HTTPS, so the device needs an active web server process (commonly the secure server with TLS) and appropriate local authentication/authorization. In addition, IOS XE requires the RESTCONF service to be explicitly turned on (typically under the native model) so that the RESTCONF root resources become available. Once these are configured, RESTCONF requests can be made to the device (for example, to the /restconf base path) using the configured credentials and transport. This combination is what makes RESTCONF reachable and functional: the HTTP secure server provides the listening socket and TLS, while the RESTCONF configuration enables the RESTCONF subsystem and its YANG-based API exposure. Cisco's RESTCONF configuration guidance for IOS XE highlights enabling the HTTP(S) server and enabling RESTCONF as the core prerequisites. See [Cisco IOS XE RESTCONF Configuration Guide](#) and [RFC 8040 \(RESTCONF\)](#).

QUESTION NO: 108

Which API does Cisco DNA Center use to retrieve information about images?

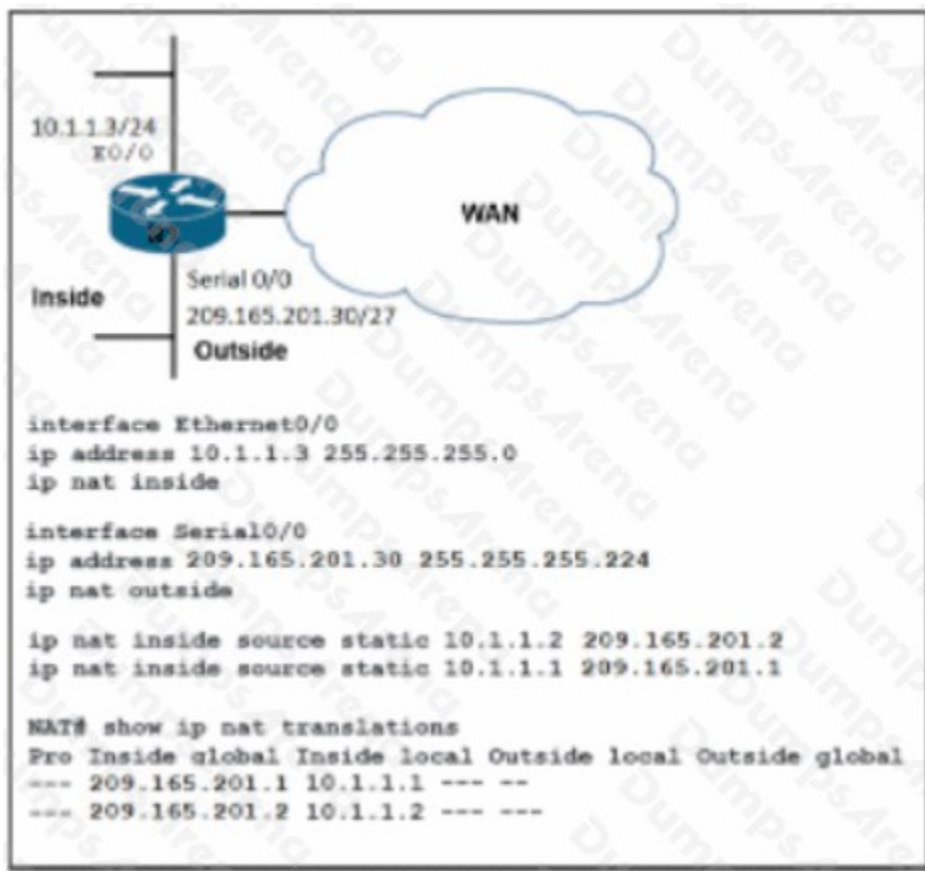
- A. SWIM
- B. Img-Mgmt
- C. PnP
- D. Client Health

ANSWER: A**Explanation:**

Cisco DNA Center retrieves and manages software image information through its Software Image Management (SWIM) APIs. SWIM is the Cisco DNA Center feature set responsible for the full lifecycle of network device software images, including importing images into the repository, querying image metadata, tagging images (golden images), distributing images to devices, and tracking image compliance. When you need to "retrieve information about images," you are effectively querying the software image repository and its associated metadata (such as image name, version, device family, and other attributes). Those operations are exposed via the SWIM-related REST endpoints in the Cisco DNA Center Platform API, making SWIM the correct API family for image information retrieval and image lifecycle workflows.

References: [Cisco DNA Center Platform API Documentation](#), [Cisco DNA Center SWIM Guide](#)

QUESTION NO: 109



Refer to the exhibit. What are two results of the NAT configuration? (Choose two.)

- A. Packets with a destination of 200.1.1.1 are translated to 10.1.1.1 or .2. respectively.
- B. A packet that is sent to 200.1.1.1 from 10.1.1.1 is translated to 209.165.201.1 on R1.
- C. R1 looks at the destination IP address of packets entering S0/0 and destined for inside hosts.
- D. R1 processes packets entering E0/0 and S0/0 by examining the source IP address.
- E. R1 is performing NAT for inside addresses and outside address.

ANSWER: B C

Explanation:

This NAT setup results in two key behaviors. First, when an inside host initiates traffic toward the outside, the router translates the inside local source address to an inside global address as the packet exits the outside interface. That's why a packet sent to the public network from 10.1.1.1 is translated to 209.165.201.1 on R1: the source is rewritten to a globally routable address so return traffic can find its way back through the NAT device. Second, for return traffic coming from the outside toward inside hosts, the router uses the destination address to decide how to translate the packet back to the correct inside local address. In other words, packets entering the outside interface (such as S0/0) and destined for inside hosts are evaluated based on their destination IP so the NAT table can map the inside global back to the inside local. These behaviors

align with Cisco NAT operation: source translation for inside-to-outside flows and destination translation for outside-to-inside return flows, driven by the NAT translation table and the inside/outside interface roles.

References: [Cisco NAT Overview and Configuration](#), [Cisco IOS NAT Configuration Guide](#)

QUESTION NO: 110

Refer to the exhibit.

```
R1# sh run | begin line con
line con 0
  exec-timeout 0 0
  privilege level 15
  logging synchronous
  stopbits 1
line aux 0
  exec-timeout 0 0
  privilege level 15
  logging synchronous
  stopbits 1
line vty 0 4
  password 7 045802150C2E
  login
line vty 5 15
  password 7 045802150C2E
  login
!
end

R1# sh run | include aaa | enable
no aaa new-model
R1#
```

Which privilege level is assigned to VTY users?

- A. 1
- B. 7
- C. 13
- D. 15

ANSWER: A

Explanation:

VTY lines (virtual terminal lines used for remote access such as SSH/Telnet) inherit the default user EXEC privilege level unless explicitly changed with the `privilege level` command under the line configuration. On Cisco IOS, the default privilege level for user EXEC mode is level 1. That means a user who connects via the VTY lines and successfully authenticates will initially be placed into user EXEC at privilege level 1, and they must use `enable` (and typically an enable password/secret or AAA authorization) to reach privileged EXEC (level 15). This default behavior applies to console, AUX,

and VTY lines unless the configuration overrides it (for example, line vty 0 4 followed by privilege level 15). Therefore, the privilege level assigned to VTY users by default is 1.

References: [Cisco - Understanding Privilege Levels](#), [Cisco IOS XE - Configuring Privilege Levels](#)

QUESTION NO: 111

An engineer must use flexible NetFlow on a group of switches. To prevent overloading of the flow collector, if the flow is idle for 20 seconds, the flow sample should be exported. Which command set should be applied?

A)

```
flow record recordflow
exporter flowexport
record recordflow
cache timeout active 120
cache timeout inactive 20
cache type immediate
```

13

B)

```
flow record recordflow
match ipv6 destination ip-address
match ipv6 source ip-address
match ipv6 protocol-type view
match interface input
match interface output
match transport destination-port
collect counter bytes long
```

C)

```
flow monitor monitorflow
exporter recordflow
cache timeout active 20
cache timeout inactive 120
cache type permanent
```

D)

```
flow monitor monitorflow
exporter flowexport
record recordflow
cache timeout active 120
cache timeout inactive 20
cache type immediate
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: D

Explanation:

To export a Flexible NetFlow record when a flow has been idle for a specific period, you must tune the flow monitor's cache inactivity timer. In Flexible NetFlow, the cache maintains active flow entries and exports them either when they become inactive (no matching packets seen for the configured interval) or when they hit an active timeout. Setting an inactive timeout of 20 seconds ensures that once traffic for a given flow stops, the device waits 20 seconds and then exports that flow entry to the collector. This is exactly how you reduce collector load from long-lived cached entries while still getting timely visibility into short, bursty conversations—because idle flows are flushed quickly rather than lingering until an active timeout is reached.

The command set that includes configuring the flow monitor cache with an inactivity timeout of 20 seconds (for example, using a form like “cache timeout inactive 20”) is therefore the correct configuration to meet the requirement “if the flow is idle for 20 seconds, the flow sample should be exported.” This behavior is part of Cisco Flexible NetFlow cache/export operation and is documented in Cisco's Flexible NetFlow configuration guides.

References: [Cisco IOS XE Flexible NetFlow Configuration Guide](#), [Cisco Flexible NetFlow Overview and Configuration](#)

QUESTION NO: 112

The login method is configured on the VTY lines of a router with these parameters

- The first method for authentication is TACACS
- If TACACS is unavailable login is allowed without any provided credentials

Which configuration accomplishes this task?

```
R1#sh run | include aaa
aaa new-model
aaa authentication login default group tacacs+
aaa session-id common

R1#sh run | section vty
line vty 0 4
transport input none
R1#

R1#sh run | include aaa
aaa new-model
aaa authentication login default group tacacs+ none
aaa session-id common

R1#sh run | section vty
line vty 0 4
password 7 02050D480809

R1#sh run | include username
R1#

R1#sh run | include aaa
aaa new-model
aaa authentication login telnet group tacacs+ none
aaa session-id common

R1#sh run | section vty
line vty 0 4

R1#sh run | include username
R1#

R1#sh run | include aaa
aaa new-model
aaa authentication login VTY group tacacs+ none
aaa session-id common

R1#sh run | section vty
line vty 0 4
password 7 02050D480809

R1#sh run | include username
R1#
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

ANSWER: B

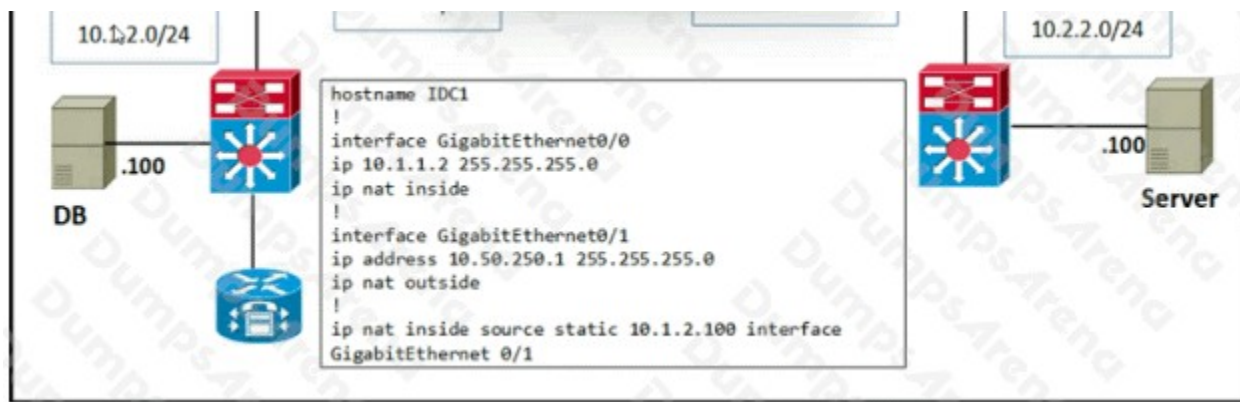
Explanation:

To meet the requirement, the VTY lines must reference an AAA method list that tries TACACS+ first and then falls back to allowing access with no username/password if the TACACS+ servers cannot be reached. In Cisco AAA, this is done by creating a method list for "login" that specifies "group tacacs+" as the first method and "none" as the second method. The

“none” keyword explicitly permits authentication to succeed without prompting for credentials, but only if the prior method (TACACS+) is unavailable (for example, server down or not reachable). After defining the method list, it must be applied under the VTY lines using the “login authentication <list-name>” command so inbound VTY sessions use that ordered method chain. This approach aligns with Cisco’s AAA method list behavior: methods are attempted in sequence, and “none” provides an unauthenticated fallback when earlier methods cannot be used. See Cisco AAA authentication method list concepts and TACACS+ AAA configuration guidance here: [Cisco AAA overview and method lists](#) and [Cisco IOS AAA configuration guide](#).

QUESTION NO: 113

Refer to the exhibit.



The server in DC2 is expecting traffic from the database in DC1 to use the source network of 10.50.250.0/24. The server sends the initial request. The inside global IP is configured for 10.50.250.1. What is the result of this configuration?

- A. Only the server can initiate communication.
- B. The server and the database cannot communicate.
- C. The server and the database can initiate communication.
- D. Only the database can initiate communication

ANSWER: C

Explanation:

This configuration results in the server and the database being able to initiate communication. With static NAT, the inside local address of the database is permanently mapped to the inside global address 10.50.250.1, so any traffic sourced from the database toward DC2 is translated to a source of 10.50.250.1, which falls within the 10.50.250.0/24 network the server expects. Because the mapping is static and bidirectional, the translation exists regardless of which side initiates the session: when the server initiates traffic toward 10.50.250.1, the NAT device translates the destination back to the database’s inside local address and forwards it; when the database initiates traffic outward, the NAT device translates the source to 10.50.250.1. This is a key behavioral difference from dynamic NAT/PAT, where translations are typically created on-demand by inside-initiated flows. Static NAT therefore supports reachability and session initiation from either side, assuming routing and any ACLs permit the traffic. See Cisco’s NAT overview and configuration guidance for static mappings: [Cisco NAT Overview](#) and [Cisco IOS NAT Configuration Guide](#).

QUESTION NO: 114

Which two items are found in YANG data models? (Choose two.)

- A. HTTP return codes
- B. rpc statements
- C. JSON schema
- D. container statements
- E. XML schema

ANSWER: B D**Explanation:**

YANG is a data modeling language used with network management protocols like NETCONF and RESTCONF. A YANG module defines the structure of configuration and operational state data using a set of built-in statements. Two core statement types you commonly see in YANG modules are *container statements* and *rpc statements*. Container statements are used to build the hierarchical tree of data nodes (grouping related leaves, lists, and other containers) that represent device configuration and state. RPC statements define remote procedure calls that can be invoked via NETCONF/RESTCONF to trigger actions or retrieve results that don't naturally fit as simple configuration/state data nodes.

In contrast, HTTP return codes are part of the HTTP protocol behavior (relevant to RESTCONF transport) rather than part of the YANG model itself. Likewise, JSON schema and XML schema are separate schema technologies; while YANG-modeled data can be encoded in XML or JSON, the YANG module is not an XML Schema (XSD) or JSON Schema document. The authoritative specification of what appears in a YANG module is defined by the YANG language standard.

References: [RFC 7950 - The YANG 1.1 Data Modeling Language](#), [RFC 6241 - NETCONF Protocol](#)

QUESTION NO: 115

Which JSON script is properly formatted?

```
○ "student":[
  {
    "grade":"9",
    "ID":"7460019964",
    "type":"on-line",
  }
]
○ {
  "plants":
  [
    "name":"Fern",
    "color":"green",
    "type":"indoor",
  ]
}
```

```
[  
  "class": {  
    [  
      "title": "History",  
      "grade": "5",  
      "location": "Site 2"  
    ]  
  ]  
]  
  
{  
  "class": [  
    {  
      "title": "Cooking 101",  
      "type": "elective",  
      "session": "fall"  
    }  
  ]  
}
```

A. Option A

- B. Option B
- C. Option C
- D. Option D

ANSWER: D

Explanation:

The properly formatted JSON script is the one that follows strict JSON syntax rules: it uses curly braces for objects and square brackets for arrays, encloses all property names (keys) in double quotes, uses a colon between each key and value, separates items with commas (but never leaves a trailing comma at the end of an object/array), and uses valid JSON value types (string in double quotes, number, boolean, null, object, or array). A correctly formatted script will also ensure that all opening braces/brackets have matching closing braces/brackets and that strings do not use single quotes. These requirements are what make JSON machine-parseable and interoperable across APIs, including Cisco RESTCONF/NETCONF tooling and controller APIs where payload validation is strict. If any of these rules are violated (for example, missing quotes around keys, trailing commas, or mismatched braces), the payload is not valid JSON and will be rejected by standard parsers. For the formal grammar and examples of valid JSON, see the JSON standard at <https://www.rfc-editor.org/rfc/rfc8259> and a practical JSON syntax overview at <https://www.json.org/json-en.html>.

QUESTION NO: 116

Which solution do IaaS service providers use to extend a Layer 2 segment across a Layer 3 network?

- A. VXLAN
- B. VTEP
- C. VRF
- D. VLAN

ANSWER: A

Explanation:

VXLAN is the solution commonly used by IaaS cloud providers to extend Layer 2 segments across an IP (Layer 3) underlay. It does this by encapsulating original Layer 2 Ethernet frames inside UDP/IP packets, creating an overlay network that can span routed networks without requiring the underlay to participate in Layer 2 bridging. This approach enables large-scale multi-tenant segmentation using a 24-bit VXLAN Network Identifier (VNI), which provides far more logical segments than traditional VLANs. In practice, VXLAN is a foundational technology for modern data center fabrics and cloud-scale virtualization because it supports workload mobility and consistent Layer 2 adjacency across Layer 3 boundaries while keeping the underlay simple and scalable. VXLAN is typically paired with control-plane mechanisms such as EVPN to distribute MAC/IP reachability information efficiently, but the core “extend Layer 2 over Layer 3” function is delivered by VXLAN encapsulation itself. Cisco documents VXLAN as an overlay encapsulation designed specifically to provide Layer 2 connectivity over a Layer 3 network, which aligns directly with the question’s requirement.

References: [Cisco VXLAN overview](#), [RFC 7348 \(VXLAN\)](#)

QUESTION NO: 117

What are two benefits of implementing a traditional WAN instead of an SD-WAN solution? (Choose two.)

- A. comprehensive configuration standardization
- B. lower control plane abstraction
- C. simplify troubleshooting
- D. faster fault detection
- E. lower data plane overhead

ANSWER: B E**Explanation:**

A traditional WAN can offer advantages that come from having fewer overlay components and less architectural abstraction than an SD-WAN. With a traditional WAN, the control plane is typically the native routing/control mechanisms you configure directly on the routers (for example, OSPF, EIGRP, BGP, static routing), rather than an SD-WAN controller-driven policy and overlay control system. That means there is lower control plane abstraction, which can be beneficial when you want direct visibility and deterministic behavior without relying on centralized orchestration and overlay control protocols.

Traditional WAN designs can also have lower data plane overhead because they usually forward traffic without additional encapsulation headers and encryption added by an SD-WAN overlay (for example, IPsec/GRE/other tunneling). Avoiding encapsulation reduces per-packet overhead and can be advantageous on constrained links or when maximizing effective MTU and throughput is important. These are common tradeoffs discussed in Cisco SD-WAN design guidance: SD-WAN adds an overlay (control and data plane) to gain centralized policy, segmentation, and transport independence, but that overlay can introduce additional abstraction and encapsulation overhead compared to a purely underlay, traditional routed WAN.

References: [Cisco SD-WAN Overview](#), [Cisco SD-WAN Design Guide](#)

QUESTION NO: 118

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

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

ANSWER: B E**Explanation:**

Locator/ID Separation Protocol (LISP) is built around separating “who” an endpoint is from “where” it is in the network, and it does that by defining two distinct address namespaces. The Endpoint Identifier (EID) namespace is used to identify end hosts or endpoints (for example, a host’s IP address as seen by applications and policies). EIDs are not used for routing across the core; instead, they represent the identity space that remains stable even if the endpoint moves. The Routing Locator (RLOC) namespace is used for routing and forwarding across the underlay/core network. RLOCs are topologically significant addresses assigned to LISP tunnel routers, and they are what the network uses to deliver encapsulated traffic between sites. LISP then maps EIDs to RLOCs via a mapping system so traffic destined to an EID can be encapsulated and sent to the correct RLOC. This EID/RLOC split is the fundamental “two-namespace” concept in LISP architecture and is consistently described in Cisco’s LISP configuration and overview documentation as well as IETF LISP architecture references.

References: [Cisco IOS XE LISP Overview](#), [IETF RFC 6830 \(LISP\)](#)

QUESTION NO: 119

What are two characteristics of Cisco SD-Access elements? (Choose two.)

- A. The border node is required for communication between fabric and nonfabric devices.
- B. Traffic within the fabric always goes through the control plane node.
- C. Fabric endpoints are connected directly to the border node.
- D. The control plane node has the full RLOC-to-EID mapping database.
- E. The border node has the full RLOC-to-EID mapping database.

ANSWER: A D

Explanation:

In Cisco SD-Access, the border node provides the fabric’s external connectivity. It is the fabric edge where traffic is exchanged between the SD-Access fabric and networks outside the fabric (for example, the enterprise core, WAN/Internet, or other nonfabric domains). This makes the border node a key element for communication between fabric and nonfabric devices, including acting as the handoff point for routing between the fabric and external networks.

The control plane node is responsible for maintaining the endpoint-to-location information used by the fabric’s LISP-based control plane. Specifically, it holds the authoritative mapping information that associates endpoint identifiers (EIDs) to routing locators (RLOCs), enabling fabric nodes to resolve where endpoints are located so data-plane forwarding can occur efficiently without forcing traffic to traverse the control plane node.

These roles are core to how SD-Access separates control-plane mapping from data-plane forwarding and how it provides scalable segmentation and external connectivity. References: [Cisco Software-Defined Access \(SD-Access\)](#), [Cisco SD-Access White Paper](#).

QUESTION NO: 120

Which tool is used in Cisco DNA Center to build generic configurations that are able to be applied on devices with similar network settings?

- A. Command Runner
- B. Application Policies

- C. Template Editor
- D. Authentication Template

ANSWER: C

Explanation:

Template Editor is the Cisco DNA Center feature used to create reusable, generic configuration templates that can be applied consistently across many devices that share similar roles or settings. In DNA Center, templates support variables (for example, hostname, IP addressing, VLAN IDs, site-specific parameters) so you can build one standardized configuration and then render it per device or per site by supplying the appropriate values. This is a core part of Cisco DNA Center's design for intent-based, scalable provisioning: you define the desired configuration once, then deploy it repeatedly to reduce manual CLI work and configuration drift. Templates can be associated with projects, versioned, and deployed during provisioning workflows, making them the correct tool when the goal is "generic configurations" applied to multiple devices with similar network settings.

References: [Cisco DNA Center User Guide – Template Editor](#), [Cisco DNA Center Platform – Template Programmer Guide](#)

QUESTION NO: 121

In a Cisco StackWise Virtual environment, which planes are virtually combined in the common logical switch?

- A. management and data
- B. control and management
- C. control, and forwarding
- D. control and data

ANSWER: B

Explanation:

In StackWise Virtual, two physical switches operate as a single logical switch by virtualizing the system so that they share a single control plane and a single management plane. Practically, this means the pair presents one logical device for configuration and monitoring (one management plane) and runs a unified set of control protocols and state (one control plane), such as spanning tree, routing adjacencies, and other control functions, coordinated across both chassis. The data/forwarding plane is not "merged" into one shared hardware pipeline; instead, each member continues to forward traffic using its own local forwarding resources, with the StackWise Virtual Link carrying traffic that must traverse between the two members. This separation is important for understanding both resiliency and performance: control and management behave like a single switch, while forwarding remains distributed across the two physical systems. Cisco's StackWise Virtual documentation describes this single logical switch behavior and the roles of the StackWise Virtual Link in supporting it. See [Cisco Catalyst 9500 StackWise Virtual configuration guide](#) and [Cisco Catalyst 9500 Series overview](#).

QUESTION NO: 122

What are two characteristics of a directional antenna? (Choose two.)

- A. high gain

- B. receive signals equally from all directions
- C. commonly used to cover large areas
- D. provides the most focused and narrow beam width
- E. low gain

ANSWER: A D

Explanation:

Directional antennas concentrate RF energy in a particular direction rather than radiating (or receiving) equally in all directions. By focusing the radiation pattern, they typically achieve *high gain* compared to omnidirectional antennas, which helps extend range and improve signal-to-noise ratio in the intended direction. A key practical characteristic is that they create a *more focused, narrower beamwidth*, which is useful for point-to-point links (such as building-to-building bridges) or for targeting coverage down a corridor or toward a specific area while reducing interference from other directions. This focused pattern is the reason directional antennas are often chosen when you want to control where the RF goes and where it doesn't, improving performance and limiting co-channel interference. Cisco wireless design guidance commonly describes directional antennas as higher-gain antennas with narrower beamwidths intended to focus coverage in a specific direction rather than broadly covering large areas.

References: [Cisco Support: Antenna Basics](#), [Cisco Antenna Selection and Deployment Guide](#)

QUESTION NO: 123

Which entity is a Type 1 hypervisor?

- A. Oracle VM VirtualBox
- B. Citrix XenServer
- C. VMware server
- D. Microsoft Virtual PC

ANSWER: B

Explanation:

Citrix XenServer is a Type 1 hypervisor because it is a bare-metal virtualization platform that installs directly on the physical server hardware and provides the hypervisor layer without requiring a general-purpose host operating system underneath. In Type 1 (bare-metal) architectures, the hypervisor is the primary control layer on the machine, managing CPU, memory, storage, and I/O resources for guest virtual machines with minimal overhead and strong isolation. XenServer (now commonly positioned within Citrix Hypervisor) is based on the Xen Project hypervisor and is designed to run directly on hardware, which is the defining characteristic of Type 1 hypervisors. This contrasts with Type 2 (hosted) hypervisors, which run as applications on top of a conventional OS. For Cisco ENCOR-level understanding, recognizing XenServer as a classic bare-metal hypervisor aligns with common enterprise virtualization deployments and the standard Type 1 vs Type 2 classification used in infrastructure design discussions. References: <https://www.citrix.com/products/citrix-hypervisor/>, <https://xenproject.org/users/>

QUESTION NO: 124

Which feature is used to propagate ARP broadcast, and link-local frames across a Cisco SD-Access fabric to address connectivity needs for silent hosts that require reception of traffic to start communicating?

- A. Native Fabric Multicast
- B. Layer 2 Flooding
- C. SOA Transit
- D. Multisite Fabric

ANSWER: B**Explanation:**

Layer 2 Flooding is the Cisco SD-Access feature used to propagate ARP broadcasts and other link-local frames across the fabric when you have endpoints that won't initiate communication until they first receive traffic (often called "silent hosts," such as badge readers, door locks, or certain IoT devices). SD-Access is designed to minimize unnecessary flooding by using control-plane learning and optimized forwarding, so classic Layer 2 broadcast behavior is not enabled by default. When a subnet has devices that rely on broadcast/ARP or link-local behavior to become reachable, Layer 2 Flooding can be enabled on a per-subnet basis so that the fabric will forward those broadcast and link-local frames as needed, restoring the expected connectivity behavior for those endpoints. This is specifically the mechanism intended to handle ARP flooding and link-local forwarding requirements within the SD-Access fabric while keeping flooding scoped and intentional rather than pervasive.

References: [Cisco SD-Access overview](#), [Cisco DNA Center configuration examples](#)

QUESTION NO: 125

A customer has two Cisco WLCs that manage separate APs throughout a building. Each WLC advertises the same SSID but terminates on different interfaces. Users report that they drop their connections and change IP addresses when roaming. Which action resolves this issue?

- A. Enable client load balancing.
- B. Enable fast roaming.
- C. Configure high availability.
- D. Configure mobility groups.

ANSWER: D**Explanation:**

Configuring mobility groups resolves this because it enables inter-controller roaming with client state (context) sharing between the two WLCs. When the same SSID exists on multiple controllers, a client that roams from an AP joined to one WLC to an AP joined to another WLC must be able to preserve its session and VLAN/subnet anchoring information; otherwise, the roam is treated more like a disconnect/reconnect event, which commonly results in DHCP renewal and an IP address change. With a mobility group (and the appropriate mobility peer configuration), the controllers exchange client information so the roam can be handled as a seamless Layer 2/Layer 3 mobility event, maintaining the client's connectivity

and avoiding application disruption. This is the standard Cisco design approach for roaming across multiple WLCs advertising the same WLAN/SSID in the same building or campus environment. For Cisco's mobility architecture and configuration concepts, see [Cisco Wireless Mobility \(Mobility Groups\) Technical Note](#) and the [Cisco Wireless Controller Configuration Guide – Mobility](#).

QUESTION NO: 126

Refer to the exhibit.

```
access-list 1 permit 10.1.1.0 0.0.0.31
ip nat pool CISCO 209.165.201.1 209.165.201.30 netmask 255.255.255.224
ip nat inside source list 1 pool CISCO
```

What are two effects of this configuration? (Choose two.)

- A. It establishes a one-to-one NAT translation.
- B. The 209.165.201.0/27 subnet is assigned as the outside local address range.
- C. The 10.1.1.0/27 subnet is assigned as the inside local addresses.
- D. Inside source addresses are translated to the 209.165.201.0/27 subnet.
- E. The 10.1.1.0/27 subnet is assigned as the inside global address range.

ANSWER: C D**Explanation:**

This configuration is a classic dynamic NAT setup that translates traffic sourced from an internal RFC1918 subnet to a pool of public IPv4 addresses. The internal subnet identified by the access list is treated as the inside local address space; those are the real addresses used on the inside network before translation occurs. When hosts from that inside local range initiate connections toward the outside, the router allocates an address from the configured public pool and rewrites the source IP to an address in that pool, creating an inside global address for each translation. As a result, inside source addresses are translated to the 209.165.201.0/27 subnet, and the 10.1.1.0/27 subnet is assigned as the inside local addresses. This is not inherently one-to-one static NAT; rather, it is dynamic allocation from a pool (though it can still be one-to-one per active translation if no overload is used). The key effects are the identification of the inside local range via the ACL and the translation of those sources into the public pool range. For terminology and behavior, see Cisco's NAT overview and inside/outside local/global definitions in the IOS NAT documentation: [Cisco NAT terminology and operation](#) and [Cisco IOS NAT Configuration Guide](#).

QUESTION NO: 127

Which two solutions are used for backing up a Cisco DNA Center Assurance database? (Choose two)

- A. NFS share
- B. non-linux server
- C. local server

- D. remote server
- E. bare metal server

ANSWER: A D

Explanation:

Cisco DNA Center Assurance (NDP) database backups are written to an external network file share, not kept only locally on the appliance. For Assurance data specifically, Cisco requires that the backup target be an NFS share hosted on a Linux-based NFS server. In practice, this means you configure a remote NFS repository (supporting NFSv3 and/or NFSv4 as documented) and Cisco DNA Center pushes the generated backup files to that remote server location. Therefore, using an “NFS share” is a correct solution, and using a “remote server” is also correct because the backup workflow explicitly posts the backup artifacts to a remote repository rather than relying on local-only storage. This aligns with Cisco’s backup/restore guidance that differentiates Assurance backups (NFS) from other data types that may use different mechanisms (for example, rsync for certain automation data), but Assurance database backups must be on NFS. See Cisco DNA Center Administrator Guide backup requirements and procedures: [Cisco DNA Center Admin Guide – Backup and Restore](#) and the general Cisco DNA Center documentation landing page: [Cisco DNA Center Maintenance Guides](#).

QUESTION NO: 128

Why would a small or mid-size business choose a cloud solution over an on-premises solution?

- A. Cloud provides higher data security than on-premises.
- B. Cloud provides more control over the implementation process than on-premises.
- C. Cloud provides greater ability for customization than on-premises.
- D. Cloud provides lower upfront cost than on-premises.

ANSWER: D

Explanation:

Cloud provides lower upfront cost than on-premises.

This is a common driver for small and mid-size businesses because cloud services typically shift spending from capital expenditure (CapEx) to operational expenditure (OpEx). Instead of purchasing servers, storage, networking gear, software licenses, and building out facilities (power, cooling, rack space), the organization consumes resources as a service and pays via subscription or usage-based billing. This reduces initial cash outlay, shortens procurement cycles, and allows the business to start small and scale up (or down) as demand changes. In addition, many cloud offerings bundle maintenance, hardware refresh, and some operational responsibilities into the service price, which further reduces the need for specialized staff and large one-time investments. Cisco’s enterprise guidance commonly frames cloud adoption benefits around agility and consumption-based economics, which aligns directly with the “lower upfront cost” rationale for SMBs compared to building and operating an on-premises environment.

References: [Cisco — What is cloud computing?](#), [AWS — What is cloud computing?](#)

QUESTION NO: 129

- A. custom headers

- B. authentication
- C. authorization
- D. request management
- E. accounting

ANSWER: B C E

Explanation:

In Cisco enterprise networks, AAA is a foundational security framework used to control and track access to network resources. The correct concepts are authentication, authorization, and accounting. Authentication is the process of verifying a user or device identity (for example, validating credentials via RADIUS or TACACS+ before granting access). Authorization determines what an authenticated user or device is allowed to do (such as which commands can be executed on a device, which VLAN/SGT is assigned, or which network services are permitted). Accounting records what the user or device did after access was granted, providing audit trails such as login/logout times, executed commands, and session statistics—critical for compliance and troubleshooting. These three functions are explicitly defined by Cisco as the AAA model and are commonly implemented using Cisco ISE (RADIUS) and TACACS+ for device administration. For further reference, see Cisco's AAA overview and configuration guidance: [Cisco AAA Overview](#) and [Cisco IOS XE AAA Configuration Guide](#).

QUESTION NO: 130

Refer to the exhibit.

```
with manager connect(host=192.168.0.1, port=22,  
username='admin', password='password1', hostkey_verify=True,  
device_params={'name': 'nexus'}) as m
```

What does the snippet of code achieve?

- A. It creates a temporary connection to a Cisco Nexus device and retrieves a token to be used for API calls.
- B. It opens a tunnel and encapsulates the login information, if the host key is correct.
- C. It opens an ncclient connection to a Cisco Nexus device and maintains it for the duration of the context.
- D. It creates an SSH connection using the SSH key that is stored, and the password is ignored.

ANSWER: C

Explanation:

The snippet is using the Python `ncclient` library's `manager.connect(...)` function inside a `with` statement, which establishes a NETCONF-over-SSH session to the target device and keeps that session open only for the lifetime of the context manager. In practice, this means the code opens a NETCONF connection to the Cisco Nexus device (using the provided host, port, username, and password, plus optional parameters like host key verification behavior) and returns a manager/session object that can be used to execute NETCONF RPC operations (for example, `get`, `get-config`, `edit-config`, etc.). When execution exits the `with` block—whether normally or due to an exception—the context manager

automatically closes the NETCONF session cleanly. This pattern is a best practice because it ensures resources are released and sessions aren't left hanging on the device. This behavior aligns with how `ncclient` documents `manager.connect` usage and with Cisco's NETCONF model of maintaining a session for a sequence of RPCs before closing it.

References: [ncclient manager documentation](#), [Cisco NX-OS NETCONF overview](#)

QUESTION NO: 131

In a high-density AP environment, which feature can be used to reduce the RF cell size and not demodulate radio packets above a given threshold?

- A. RX-SOP
- B. FRA
- C. 80211k
- D. RRM

ANSWER: A**Explanation:**

RX-SOP is the feature used in high-density wireless deployments to effectively shrink an AP's RF cell by controlling which received frames the AP will attempt to decode. RX-SOP stands for Receive Start of Packet detection threshold. By raising (making less sensitive) the RX-SOP threshold, the AP ignores weaker signals (for example, distant clients or co-channel transmissions) and therefore does not proceed to demodulate/decode packets below that configured RSSI level. This reduces the amount of co-channel interference the AP spends time processing, improves spatial reuse, and helps keep clients associated to closer APs—key goals in high-density designs. Cisco controllers and Catalyst wireless platforms expose RX-SOP as a tunable parameter (often with preset profiles) specifically to address dense environments where too-large cells and excessive contention reduce throughput. This aligns directly with the question's wording about reducing RF cell size and not demodulating packets beyond a threshold.

References: [Cisco – RX-SOP configuration on Catalyst wireless](#), [Cisco Catalyst 9800 Configuration Guide – RF parameters \(includes RX-SOP\)](#)

QUESTION NO: 132

53

What is a characteristic of Cisco DNA southbound APIs?

- A. simplifies management of network devices
- B. enables orchestration and automation of network devices based on intent
- C. utilizes REST API
- D. implements monitoring by using the SOAP protocol

ANSWER: B

Explanation:

"enables orchestration and automation of network devices based on intent" is a characteristic of Cisco DNA Center southbound APIs because southbound interfaces are used by the controller to communicate with and program the underlying network infrastructure. In Cisco DNA Center's intent-based networking model, the controller translates high-level intent (for example, connectivity, segmentation, QoS, or assurance policies) into device-level configurations and operational actions, then pushes those changes to network devices through southbound mechanisms. This is what makes large-scale automation and orchestration possible: the controller continuously applies intent, provisions devices, and maintains desired state across the fabric and campus network. While northbound APIs are typically exposed to external applications (often RESTful) for integration, southbound communication focuses on controlling and automating the devices themselves using device-facing protocols and integrations. This aligns with Cisco's description of DNA Center as an intent-based controller that automates provisioning and policy across the enterprise network. For additional context on Cisco DNA Center intent-based networking and how it automates the network, see [Cisco DNA Center product overview](#) and [Cisco DNA Center Developer Documentation](#).

QUESTION NO: 133 - (DRAG DROP)

Drag and drop the automation characteristics from the left to the corresponding tools on the right.

CHINESE DUMPS 通过测试

uses playbooks and plays

uses modules and manifests

uses cookbooks and recipes

does not require an admin account on the client

Ansible

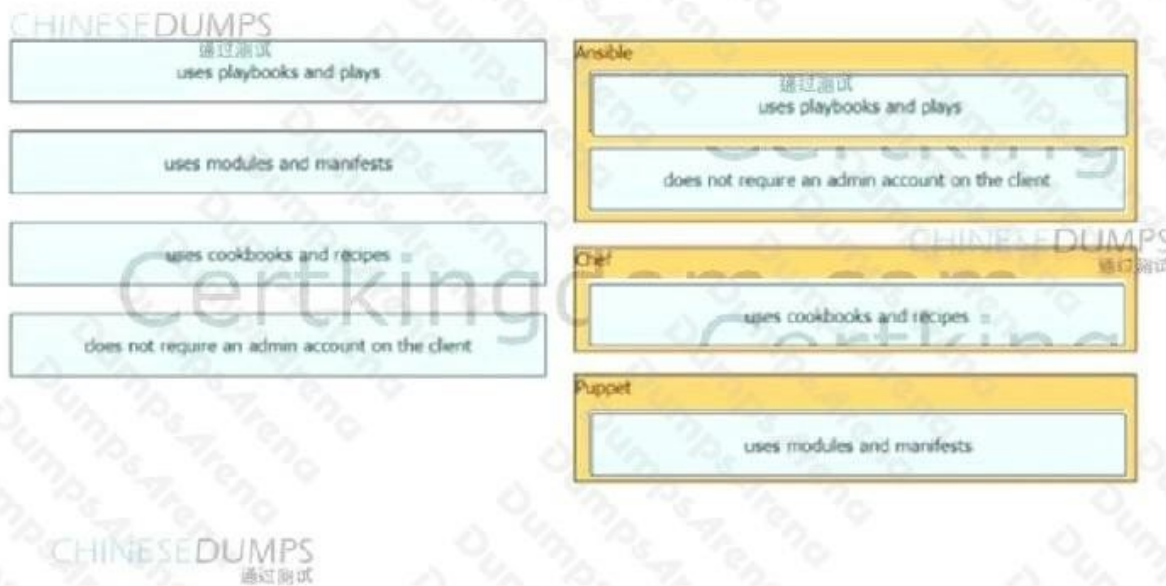
Chef 通过测试

Puppet

Certkingdom.com

CHINESE DUMPS 通过测试

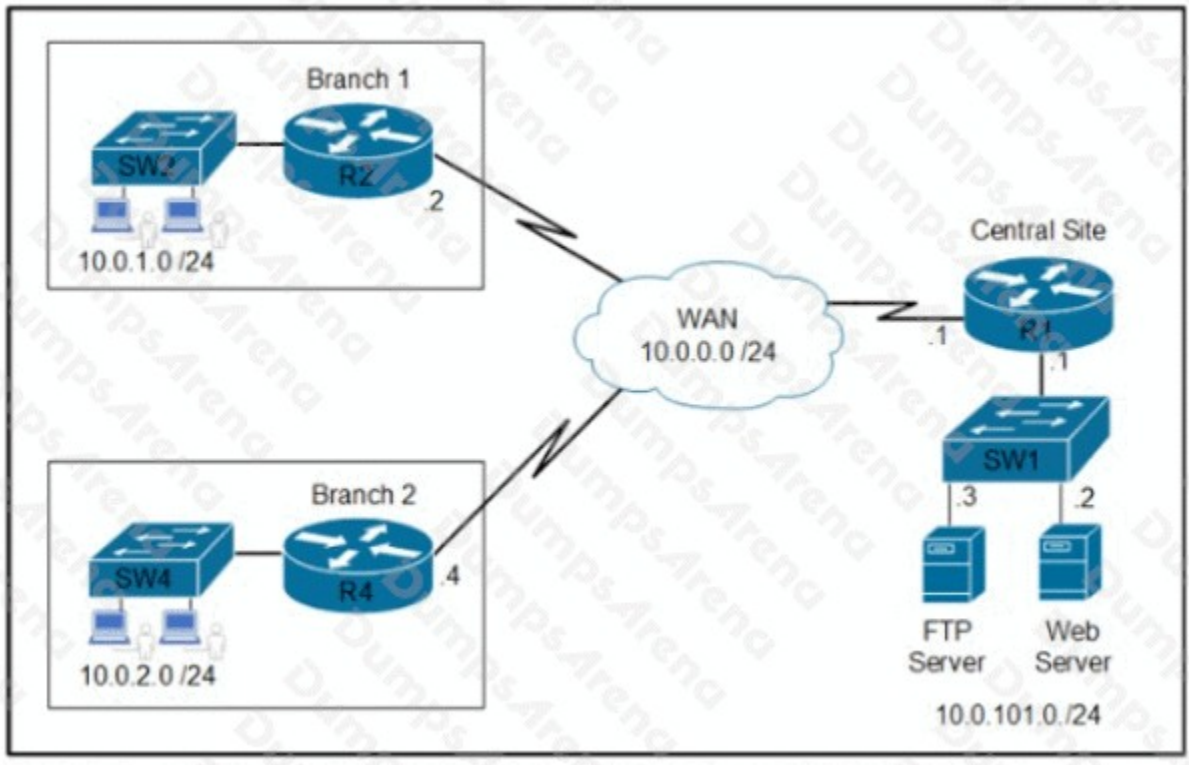
ANSWER:

**Explanation:**

Ansible, Chef, and Puppet each have their own well-known configuration language and operating model, so the easiest way to match the characteristics is to align the wording with the tool's native terminology. Ansible describes automation in YAML "playbooks," and a playbook is made up of one or more "plays" that target groups of hosts and apply tasks in order. That directly matches the characteristic "uses playbooks and plays." Ansible is also designed to be agentless: it connects to managed nodes over SSH (or WinRM for Windows) and executes modules remotely without requiring a persistent agent to be installed and managed on the endpoint. In exam terms, that maps to "does not require an admin account on the client," meaning there is no dedicated local agent/service account that must be deployed and maintained on each managed device the way agent-based systems often do; access is typically provided via remote connectivity and appropriate credentials. Chef's core unit of automation is the "cookbook," which contains "recipes" that declare how a system should be configured, so "uses cookbooks and recipes" maps to Chef. Puppet uses a declarative language where "manifests" define desired state and are organized into "modules," so "uses modules and manifests" maps to Puppet. These terms are canonical in each product's documentation and are commonly tested as foundational automation knowledge in ENCOR.

References: [Ansible playbooks overview](#), [Ansible getting started \(agentless over SSH/WinRM\)](#), [Chef recipes documentation](#), [Puppet modules fundamentals](#), [Puppet language/manifests summary](#).

QUESTION NO: 134



Refer to the exhibit Which two commands are required on route» R1 to block FTP and allow all other traffic from the Branch 2 network' (Choose two)

- access-list 101 deny tcp 10.0.2.0 0.0.0.255 host 10.0.101.3 eq ftp-data
access-list 101 permit ip any any
- access-list 101 deny tcp 10.0.2.0 0.0.0.255 host 10.0.101.3 eq ftp
access-list 101 deny tcp 10.0.2.0 0.0.0.255 host 10.0.101.3 eq ftp-data
access-list 101 permit ip any any
- interface GigabitEthernet0/0
ip address 10.0.0.1 255.255.255.252
ip access-group 101 out
- interface GigabitEthernet0/0
ip address 10.0.101.1 255.255.255.252
ip access-group 101 in
- access-list 101 deny tcp 10.0.2.0 0.0.0.255 host 10.0.101.3 eq ftp
access-list 101 permit ip any any

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E
- F. Option F

ANSWER: A D F

Explanation:

To block FTP while allowing all other traffic from a specific source network (the Branch 2 network), you implement an extended ACL that first denies the FTP flows and then explicitly permits the remaining IP traffic. In Cisco IOS ACL processing, entries are evaluated top-down and the first match wins; additionally, there is an implicit “deny ip any any” at the end of every ACL, so you must include an explicit permit statement to allow non-FTP traffic. For FTP, you typically deny TCP port 21 (control channel). Many designs also deny TCP port 20 (active-mode data channel), but the exact requirement depends on the exhibit and what “block FTP” is intended to cover. After defining the ACL entries, you must apply the ACL to the correct interface and direction on R1 so that traffic sourced from Branch 2 is filtered as it enters or exits R1 (most commonly inbound on the interface facing Branch 2). This combination—ACL deny for FTP plus a permit for all other traffic, and then an interface ACL application—are the two required command types to achieve the stated policy. See Cisco ACL behavior and configuration guidance: [Cisco IOS Access Lists](#) and [Cisco XE ACL Configuration Guide](#).

QUESTION NO: 135

Which two Cisco SD-WAN components exchange OMP information?

- A. vAnalytics
- B. vSmart
- C. WAN Edge
- D. vBond
- E. vManage

ANSWER: B C**Explanation:**

In Cisco SD-WAN, the Overlay Management Protocol (OMP) is the control-plane protocol used to advertise and learn overlay routes, TLOCs, and service information so the fabric can build end-to-end connectivity across the WAN. OMP sessions are established between WAN Edge devices (vEdge/cEdge) and the vSmart controller. WAN Edge routers form secure control connections to vSmart and then exchange OMP updates: WAN Edge advertises its local prefixes/TLOCs/services into the fabric, and vSmart reflects and distributes learned OMP routes to other WAN Edge devices according to policy. This vSmart-to-WAN Edge OMP exchange is what enables dynamic reachability, segmentation (VPNs), and centralized policy enforcement in the SD-WAN overlay. Other controllers have different roles: vBond primarily orchestrates initial authentication and NAT traversal, and vManage provides management/monitoring, but neither is the endpoint for OMP route exchange in the steady-state control plane. For more details on SD-WAN control-plane components and OMP operation, see Cisco’s SD-WAN documentation: [Cisco SD-WAN Control Plane](#) and [Cisco SD-WAN OMP](#).

QUESTION NO: 136

What does the Cisco DNA Center Authentication API provide?

- A. list of global issues that are logged in Cisco DNA Center
- B. access token to make calls to Cisco DNA Center
- C. list of VLAN names

D. dent health status

ANSWER: B

Explanation:

The Cisco DNA Center Authentication API provides an access token to make calls to Cisco DNA Center. Cisco DNA Center's intent APIs are protected, so a client must first authenticate (typically with a username and password) to obtain a token, and then include that token in subsequent API requests (commonly in an HTTP header such as `X-Auth-Token`). This token-based approach is fundamental to securing the platform because it ensures only authenticated, authorized clients can invoke operational endpoints (for example, to query inventory, run commands, or retrieve assurance data). In practice, the authentication call is the first step in any automation workflow: you request the token, store it for its validity period, and reuse it until it expires, at which point you request a new token. This is why the authentication endpoint is described as providing the access token required to make further API calls, rather than returning operational data like issues, VLANs, or health metrics.

References: [Cisco DNA Center Platform APIs \(Developer Documentation\)](#), [Cisco DNA Center Authentication](#)

QUESTION NO: 137

Refer to the exhibit.

```
R2#show standby
FastEthernet1/0 - Group 40
  State is Standby
    4 state changes, last state change 00:01:51
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c07.ac28 (MAC Not In Use)
  Local virtual MAC address is 0000.0c07.ac28 (v1 default)
  Hello time 3 sec, hold time 10 sec
  Next hello sent in 1.856 secs
  Preemption disabled
  Active router is 10.10.1.3, priority 85 (expires in 8.672 sec)
  Standby router is local
  Priority 90 (configured 90)
  Track interface FastEthernet0/0 state Up decrement 10
  Group name is "hsrp-Fal/0-40" (default)
```

After configuring HSRP an engineer enters the show standby command. Which two facts are derived from the output? (Choose two.)

- A. The router with IP 10.10.1.3 is active because it has a higher IP address
- B. If Fa0/0 is shut down, the HSRP priority on R2 becomes 80
- C. R2 Fa1/0 regains the primary role when the link comes back up
- D. R2 becomes the active router after the hold time expires.
- E. R2 is using the default HSRP hello and hold timers.

ANSWER: D E

Explanation:

The fact that “R2 becomes the active router after the hold time expires.” is derived from the HSRP state information shown by `show standby`. When a router is in the standby role and it stops receiving HSRP hello messages from the current active router, it waits for the configured holdtime; once that holdtime expires without hearing hellos, it transitions to active. This is exactly how HSRP provides fast default-gateway failover using hello/hold timers and state transitions.

The fact that “R2 is using the default HSRP hello and hold timers.” is also derived directly from the `show standby output` when it displays the timer values. For HSRP (version 1), the default hello timer is 3 seconds and the default hold timer is 10 seconds; if the output shows these values, it confirms defaults are in use. These timers control how quickly a standby router detects loss of the active router and takes over the virtual IP/MAC for the group.

References: [Cisco HSRP Overview and Timer Behavior](#), [Cisco IOS HSRP Configuration Guide](#)

QUESTION NO: 138

Which two security features are available when implementing NTP? (Choose two.)

- A. symmetric server passwords
- B. dock offset authentication
- C. broadcast association mode
- D. encrypted authentication mechanism
- E. access list-based restriction scheme

ANSWER: A E

Explanation:

NTP on Cisco devices supports multiple built-in security controls to prevent unauthorized time sources from influencing the system clock. One key feature is **symmetric server passwords**, which is NTP authentication using shared keys (for example, MD5) so the client can verify that received NTP packets were generated by a trusted peer configured with the same key. This helps mitigate spoofed NTP responses and time-manipulation attacks by ensuring only authenticated peers are accepted for synchronization.

Another important security feature is an **access list-based restriction scheme**. Cisco IOS/IOS XE can restrict which hosts are allowed to query, peer with, or serve time to the device, limiting exposure to untrusted networks and reducing the attack surface. In practice, this is implemented with NTP access-group controls (restricting peer/query/serve behavior) so only approved IPs can participate in NTP exchanges.

These two features are commonly used together: authentication validates the legitimacy of NTP messages, while access restrictions control who can even attempt NTP interactions. References: [Cisco IOS XE NTP Configuration Guide](#), [RFC 5905 \(NTPv4\)](#).

QUESTION NO: 139

```
R1
interface GigabitEthernet0/0
ip address 192.168.250.2 255.255.255.0
standby 20 ip 192.168.250.1
standby 20 priority 120

R2
interface GigabitEthernet0/0
ip address 192.168.250.3 255.255.255.0
standby 20 ip 192.168.250.1
standby 20 priority 110
```

Refer to the exhibit. What are two effects of this configuration? (Choose two.)

- A. If R1 goes down, R2 becomes active but reverts to standby when R1 comes back online.
- B. If R2 goes down, R1 becomes active but reverts to standby when R2 comes back online.
- C. R1 becomes the active router.
- D. R1 becomes the standby router.
- E. If R1 goes down, R2 becomes active and remains the active device when R1 comes back online.

ANSWER: C E

Explanation:

This configuration results in R1 being elected as the active HSRP router and also causes a non-preempt behavior after a failure. In HSRP, the active router is chosen primarily by the highest priority value; if priorities tie, the highest IP address on the HSRP-enabled interface wins. The exhibit's configuration (showing R1 with the higher effective HSRP priority) makes R1 the active router under normal conditions, with R2 in standby.

The second effect is related to preemption. HSRP does not preempt by default, meaning that if the active router fails and the standby router takes over, the original router will not automatically reclaim the active role when it returns—unless the *preempt* feature is explicitly configured. Therefore, if R1 goes down, R2 becomes active, and when R1 comes back online, R2 remains active (because preemption was not enabled to force R1 to take back the role).

References: [Cisco HSRP Overview](#), [Cisco IOS XE HSRP Configuration Guide](#)

QUESTION NO: 140

An engineer must configure the strongest password authentication to locally authenticate on a router. Which configuration must be used?

- username netadmin secret 5 \$1\$b1Ju\$KzBBS1Pyh4QzwXyZ1kSZ2
- username netadmin secret \$1\$b1Ju\$K404850110QzwXyZ1kSZ2
- line Console 0
password \$1\$b1Ju\$
- username netadmin secret 9 \$9\$vFpMf8eIb4RVV8\$seZ/bDAx1uV

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. username user secret 9

ANSWER: E

Explanation:

To configure the strongest local password authentication on a Cisco router, you should use a username with a *secret* that is stored using the strongest available one-way hashing algorithm. On modern Cisco IOS/IOS XE platforms, the strongest supported secret type for local user authentication is the scrypt-based hash, which is configured by specifying `secret 9` under the `username` command. Scrypt is designed to be computationally and memory expensive, making offline password cracking significantly harder than older hashes such as MD5 (type 5) and also generally stronger than reversible/obfuscation schemes. In practice, this means configuring something like `username <name> secret 9 <secret>` (or letting the device generate the type 9 hash depending on platform/feature support). This aligns with Cisco guidance to prefer newer secret types (type 8 PBKDF2 and type 9 scrypt) over legacy password/secret formats for local authentication. For additional background on Cisco password/secret types and their relative strength, see Cisco's discussion of password types and recommendations at <https://community.cisco.com/t5/networking-documents/understanding-the-differences-between-the-cisco-password-secret/ta-p/3163238> and Cisco IOS XE security configuration guidance at <https://www.cisco.com/c/en/us/support/ios-nx-os-software/ios-xe-16/products-installation-and-configuration-guides-list.html>.

QUESTION NO: 141

In which two ways does the routing protocol OSPF differ from EIGRP? (Choose two.)

- A. OSPF supports an unlimited number of hops. EIGRP supports a maximum of 255 hops.
- B. OSPF provides shorter convergence time than EIGRP.
- C. OSPF is distance vector protocol. EIGRP is a link-state protocol.
- D. OSPF supports only equal-cost load balancing. EIGRP supports unequal-cost load balancing.

E. OSPF supports unequal-cost load balancing. EIGRP supports only equal-cost load balancing.


ANSWER: A D

Explanation:

OSPF and EIGRP differ in both their path-selection mechanics and their load-balancing capabilities. OSPF is a link-state routing protocol that uses the SPF (Dijkstra) algorithm and computes best paths based on cumulative interface cost; it does not provide a native mechanism for unequal-cost multipath, so it performs equal-cost load balancing only when multiple paths have the same total cost. EIGRP, while often described as an advanced distance-vector protocol, uses the DUAL algorithm and can install multiple routes that are not equal metric by using the variance feature, enabling unequal-cost load balancing.

Another key difference is hop-count behavior. OSPF does not use hop count as a metric and therefore does not impose a protocol hop-limit in the way classic distance-vector protocols do; its scalability constraints are instead tied to LSDB size, SPF computation, and design (areas), not a fixed hop maximum. EIGRP includes a configurable maximum hop count (default 100, maximum 255) as a loop-prevention/sanity limit. These distinctions align with Cisco's descriptions of OSPF operation and EIGRP features such as variance and maximum hop count. See [Cisco OSPF overview](#) and [Cisco EIGRP overview](#).

QUESTION NO: 142



Refer to the exhibit. A network engineer must load balance traffic that comes from the NAT Router and is destined to 10.10.110.10, to several FTP servers. Which two commands sets should be applied? (Choose two).

A)

```
interface gig0/0
ip address 10.10.110.1 255.255.255.0
ip nat inside
interface gig0/1
ip address 172.16.1.1 255.255.252
ip nat outside
```

B)

```
ip nat pool ftp-pool 10.10.110.2 10.10.110.9 netmask 255.255.255.0
access-list 23 permit 10.10.110.10
ip nat inside destination-list 23 pool ftp-pool
```

C)

```
ip nat pool ftp-pool 10.10.110.2 10.10.110.9 netmask 255.255.255.0 type rotary
access-list 23 permit 10.10.110.10
ip nat inside destination-list 23 pool ftp-pool
```

D)

```
ip nat pool ftp-pool 10.10.110.2 10.10.110.9 netmask 255.255.255.0 type rotary
access-list 23 permit 10.10.110.10
ip nat outside destination-list 23 pool ftp-pool
```

E)

```
interface gig0/0
ip address 10.10.110.1 255.255.255.0
ip nat outside
interface gig0/1
ip address 172.16.1.1 255.255.255.252
ip nat inside
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E

ANSWER: A C**Explanation:**

To load-balance inbound traffic destined to a single virtual IP address (10.10.110.10) across multiple real FTP servers, you typically implement a server load-balancing feature that can (1) define a virtual service address that clients target and (2) define a pool of real servers with a distribution method (such as round-robin) and health checking so failed servers are removed from rotation. The correct command sets are the ones that build this end-to-end construct: they create the virtual IP/service for FTP (TCP/21) and associate it with multiple real server IPs as a server farm/pool, enabling load distribution for connections arriving from the NAT router. This is the standard approach used by Cisco load-balancing solutions (for example, IOS SLB on supported platforms or dedicated ADCs), where the VIP remains constant while the device selects a real server per new flow and maintains session persistence as needed for the application. For background on Cisco server load-balancing concepts (VIPs, server farms, probes/health checks), see [Cisco ACE documentation](#) and [Cisco IOS IP Application Services Configuration Guide](#).

QUESTION NO: 143

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.

ANSWER: B**Explanation:**

A Type 1 hypervisor is a “bare-metal” hypervisor that installs directly on the physical server hardware and does not require a separate, general-purpose host operating system underneath it. In this model, the hypervisor itself provides the core virtualization layer (CPU scheduling, memory management, device access mediation, and VM isolation) and includes the minimal OS-like components and drivers needed to manage hardware and run virtual machines. This is why the description “runs directly on a physical server and includes its own operating system” best matches a Type 1 hypervisor: it operates at the lowest software layer above the hardware and is designed for performance, stability, and strong isolation in data center deployments (common examples include VMware ESXi and Microsoft Hyper-V in its bare-metal architecture). Cisco’s virtualization discussions commonly distinguish Type 1 (bare-metal) from Type 2 (hosted) hypervisors, where Type 2 depends on a preinstalled host OS. For additional background, see [VMware glossary: Bare-metal hypervisor](#) and [Microsoft Learn: Hyper-V on Windows Server](#).