PassQueen
a.k.a dumpsqueen.com

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

Cisco 350-401

Version Demo

Total Demo Questions: 158

Total Premium Questions: 1588

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Architecture	166
Topic 2, Virtualization	80
Topic 3, Infrastructure	631
Topic 4, Network Assurance	77
Topic 5, Security	298
Topic 6, Automation	333
Topic 7, Mix Questions	3
Total	1588

QUESTION NO: 1

What is the calculation that is used to measure the radiated power of a signal after it has gone through the radio, antenna cable, and antenna?

A. EIRP

B. mW

C. dBm

D. dBi

ANSWER: A

Explanation:

EIRP is the correct calculation because it expresses the effective power radiated by a wireless system after accounting for the complete transmit path. It begins with the radio transmitter output power, subtracts attenuation introduced by antenna cables, connectors, and other RF components, then adds the gain provided by the antenna. In logarithmic units, the common calculation is transmitter power (dBm) − RF-path losses (dB) + antenna gain (dBi) = EIRP (dBm). EIRP is referenced to an ideal isotropic radiator, a theoretical antenna that radiates equally in every direction. This reference makes it possible to compare the radiated strength of different real antenna systems consistently, regardless of their physical design or directionality. In enterprise wireless deployments, EIRP is important for RF design, coverage planning, interference management, and compliance with country-specific regulatory limits, which commonly define permitted transmission levels in terms of maximum EIRP. Cisco documents this power-budget approach for WLAN deployments and identifies EIRP as the value that incorporates transmitter power, cable loss, and antenna gain. See Cisco's EIRP and ERP documentation and the FCC's 47 CFR §15.407 regulations.

QUESTION NO: 2

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

A. port flapping

B. excess jitter

C. misrouted network packets

D. duplicate IP addresses

E. bandwidth-related packet loss

ANSWER: B E

Explanation:

QoS is needed when traffic contention causes measurable service degradation and the network must provide differentiated treatment to important applications. **excess jitter** is a key QoS indicator because it is variation in packet-arrival delay. Real-time media flows, particularly voice and interactive video, require packets to arrive at a consistent rate; excessive delay variation depletes receive buffers and degrades call or video quality. QoS classification, marking, and appropriate queue servicing help give these time-sensitive flows predictable treatment during congestion.

bandwidth-related packet loss indicates that an interface or downstream path has insufficient available capacity at times of offered-load peaks. Packets are then discarded as egress queues fill. QoS addresses this congestion symptom by classifying traffic and applying queuing, scheduling, and congestion-management policies that protect defined critical traffic classes. In a campus design, this can include low-latency queuing for delay-sensitive traffic and managed queue allocation for business-critical applications. QoS does not create additional physical bandwidth, but it determines which traffic receives the most appropriate forwarding behavior when demand exceeds link capacity. Cisco describes congestion management as the use of queueing algorithms to manage packet transmission during congestion and identifies jitter and loss as important performance considerations for real-time traffic. See the Cisco IOS XE QoS Congestion Management Configuration Guide and Cisco's QoS Policing and Shaping Overview.

QUESTION NO: 3

Refer to the exhibit.

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:

The configuration `line vty 0 15 exec-timeout 10 0` achieves the required idle-session timeout because the Cisco IOS `exec-timeout` line command defines the maximum period of inactivity allowed in an EXEC session before the device terminates that session. Its arguments are expressed as minutes followed by optional seconds. Therefore, `10 0` configures a timeout of 10 minutes and zero additional seconds, which is exactly 600 seconds.

Applying the setting to the VTY range from 0 through 15 ensures that the configured remote terminal lines enforce the policy consistently. When a user connected through one of these lines provides no input for the configured period, the router or switch closes the EXEC session. This is a standard management-plane security control because unattended authenticated sessions should not remain available indefinitely. Cisco documents that `exec-timeout` is configured under a line and controls the idle duration before an EXEC session is disconnected. See the Cisco IOS XE Terminal Line Management documentation and the Cisco IOS exec-timeout command reference.

QUESTION NO: 4

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. 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:

Ensure that interfaces Gi1/0/2 and Gi1/0/3 connect to the same neighboring switch. A conventional EtherChannel bundles physical interfaces into one logical link between the same pair of devices. This common-neighbor requirement ensures that the switch can consistently identify all member links as belonging to the same aggregation and can forward traffic over the resulting logical Port-Channel without creating an unintended topology or loop condition. Connecting member interfaces to separate standalone switches does not form a valid EtherChannel unless those switches operate as one logical EtherChannel peer through an appropriate multichassis technology.

Ensure that the neighbor interfaces of Gi1/0/2 and Gi/0/3 are configured as members of the same EtherChannel. EtherChannel configuration must be consistent at both ends: the far-end physical interfaces must participate in one corresponding channel and use compatible channel negotiation and interface settings. This creates a single, mutually recognized logical bundle, allowing LACP or the configured EtherChannel mode to maintain the member links correctly. Cisco documents that channel members must connect to the same device and that EtherChannel parameters must be compatible across the link. See Cisco EtherChannel configuration guidelines and Cisco EtherChannel and LACP troubleshooting guidance.

QUESTION NO: 5

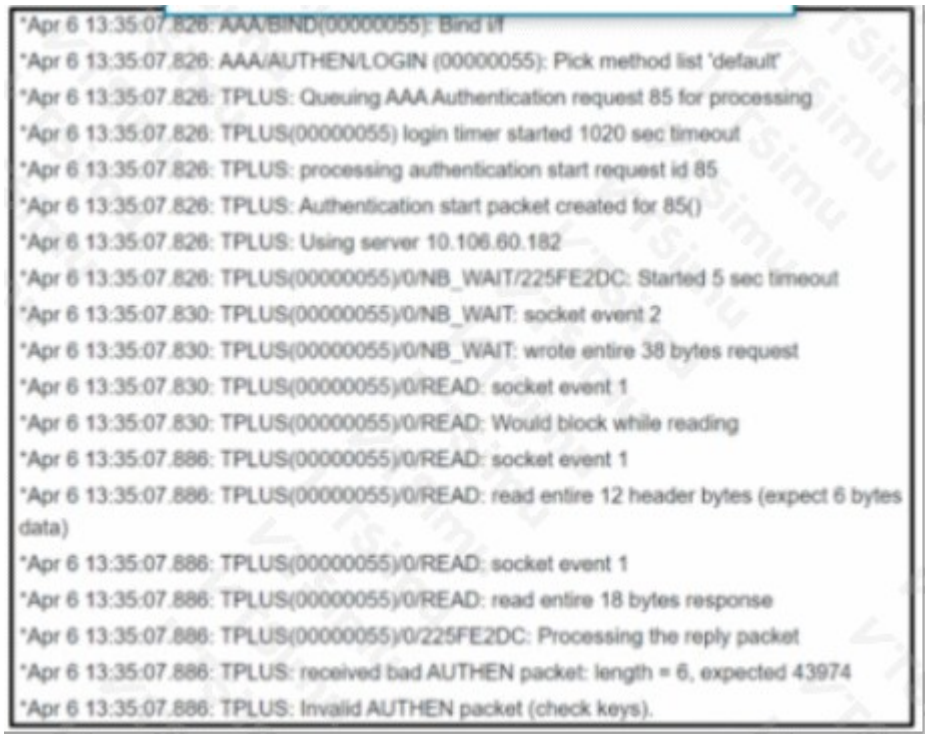

Refer to the exhibit. An engines configured TACACS^ to authenticate remote users but the configuration is not working as expected Which configuration must be applied to enable access?

A)

```
R1(config)# ip tacacs source-interface Gig 0/0
```

B)

```
R1(config)# tacacs server prod
R1(config-server-tacacs)# key cisco123
```

C)

```
R1(config)# aaa authorization exec default group tacacs+ local
```

D)

```
R1(config)# tacacs server prod
R1(config-server-tacacs)# port 1020
```

A. Option A

B. Option B

C. Option C

D. Option D

E. Option E

F. `aaa authentication login default group tacacs+ local`

ANSWER: F

Explanation:

The required configuration is `aaa authentication login default group tacacs+ local`, assuming AAA has already been enabled with `aaa new-model` and the TACACS+ server definition is present. This command creates the default AAA login authentication method list and directs IOS or IOS XE to query the configured TACACS+ server group first whenever a user initiates a login. The `local` keyword provides a local username database fallback when the TACACS+ server is unavailable, helping preserve administrative access during a server or reachability failure. A method list named `default` is applied automatically to login-capable lines that do not explicitly reference a different list, including VTY access used for remote SSH or Telnet administration. Alternatively, a named method list can be bound directly under the relevant VTY lines with `login authentication`, but the default method list is the direct configuration needed to activate TACACS+-based login authentication broadly. Cisco's AAA documentation describes authentication method lists and the special behavior of the default list, while the IOS XE TACACS+ guide documents use of TACACS+ server groups for AAA authentication. See Cisco AAA overview and Cisco IOS XE TACACS+ configuration guide.

QUESTION NO: 6

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

A. FIB is a database of routing prefixes, and the RIB is the information used to choose the egress interface for each packet.

B. The FIB is derived from the data plane, and the RIB is derived from the FIB.

C. The RIB is a database of routing prefixes, and the FIB is the information used to choose the egress interface for each packet.

D. The RIB is derived from the control plane, and the FIB is derived from the RIB.

E. The FIB is derived from the control plane, and the RIB is derived from the FIB.

ANSWER: C D

Explanation:

The RIB is a database of routing prefixes, and the FIB is the information used to choose the egress interface for each packet. The Routing Information Base is maintained by the control plane and represents the router's selected routing information, including routes learned through connected interfaces, static configuration, and routing protocols. Route-selection processes determine the active route for a destination based on attributes such as administrative distance and metric. The Forwarding Information Base is then built from the selected RIB routes and provides an optimized forwarding table for the data plane. It supports rapid destination lookup so packets can be forwarded to the appropriate next hop and egress interface without repeating control-plane route calculations for every packet.

The RIB is derived from the control plane, and the FIB is derived from the RIB. This describes the essential control-plane/data-plane relationship in Cisco Express Forwarding. The control plane learns, calculates, and installs best routes in the RIB. Cisco IOS XE programs the corresponding forwarding entries into the FIB, often together with adjacency information that supplies Layer 2 rewrite details. Consequently, routing changes that affect selected RIB routes can trigger FIB updates, keeping packet forwarding aligned with the current best-path decisions. Cisco documents the CEF FIB as the table used for IP destination-prefix lookup and describes its relationship to routing information in the Cisco Express Forwarding overview and the Cisco IOS XE routing documentation.

QUESTION NO: 7

An engineer must create a new SSID on a Cisco 9800 wireless LAN controller. The client has asked to use a pre-shared key for authentication. Which profile must the engineer edit to achieve this requirement?

A. RF

B. WLAN

C. Policy

D. Flex

ANSWER: B

Explanation:

The **WLAN** profile is where a Cisco Catalyst 9800 controller defines the SSID and its client-facing Layer 2 security configuration. To deploy an SSID that authenticates clients with a pre-shared key, the engineer configures Personal security in the WLAN, selects the applicable WPA2-Personal or WPA3-Personal authentication/key-management settings, and enters the passphrase or PSK. These settings establish how wireless clients authenticate to that particular SSID before receiving network access.

On Catalyst 9800 controllers, a WLAN is created independently and then associated with a policy profile through a policy tag. This separation allows the SSID's identity and security settings to remain part of the WLAN configuration while the assignment of client VLANs, access policies, QoS behavior, and similar forwarding-related parameters is handled elsewhere. Therefore, configuring the WLAN profile is the required step for enabling PSK authentication on the new SSID. Cisco documents WLAN configuration and WLAN security settings in its Catalyst 9800 WLAN Configuration Guide and provides additional deployment documentation through the Catalyst 9800 configuration guides.

QUESTION NO: 8

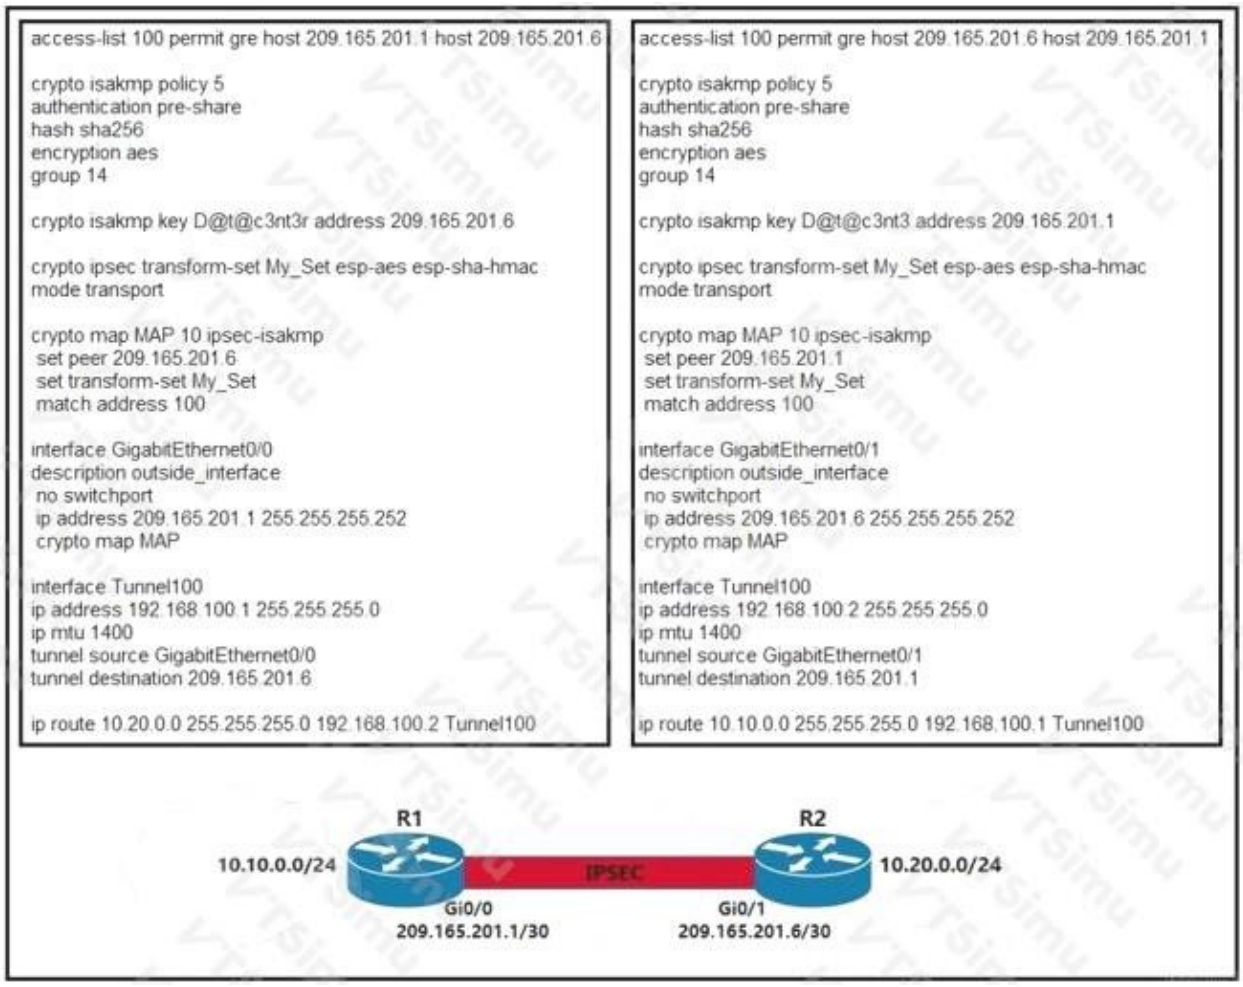

Refer to the exhibit. A network engineer must simplify the IPsec configuration by enabling IPsec over GRE using IPsec profiles. Which two configuration changes accomplish this? (Choose two).

A. Create an IPsec profile, associate the transform-set ACL, and apply the profile to the tunnel interface.

B. Apply the crypto map to the tunnel interface and change the tunnel mode to tunnel mode ipsec ipv4.

C. Remove all configuration related to crypto map from R1 and R2 and eliminate the ACL.

D. Create an IPsec profile, associate the transform-set, and apply the profile to the tunnel interface.

E. Remove the crypto map and modify the ACL to allow traffic between 10.10.0.0/24 to 10.20.0.0/24.

ANSWER: C D

Explanation:

GRE tunnel protection with an IPsec profile is configured directly on the GRE tunnel interface. The profile contains the IPsec Phase 2 parameters by referencing an existing transform set, and it is enabled beneath the tunnel interface with the `tunnel protection ipsec profile` command. This binds IPsec protection specifically to GRE packets carried by that tunnel, providing the intended GRE-over-IPsec design without requiring policy selection based on a crypto access list.

Accordingly, the legacy crypto-map-based configuration on the physical outside interfaces is removed from both routers. A crypto map and its associated “interesting traffic” ACL are used to identify traffic for policy-based IPsec. With tunnel-interface IPsec protection, the GRE tunnel itself defines the traffic to protect, so those crypto-map-related elements are no longer part of the configuration. The IKE policy, peer authentication material, transform set, and related IPsec security-association settings remain available for use by the IPsec profile.

Cisco documents this deployment as applying an IPsec profile to a GRE tunnel interface and describes profiles as the mechanism for tunnel protection. See Cisco’s IPsec Profiles Configuration Guide and its GRE over IPsec configuration overview.

QUESTION NO: 9

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 database backups use external storage rather than relying solely on storage internal to the Cisco DNA Center appliance. An **NFS share** is the supported repository type for Assurance backup data. Cisco DNA Center mounts or accesses the configured NFS export and writes the Assurance backup artifacts to that shared location. This provides persistent storage that can be retained independently of the appliance and used as part of a restoration workflow.

A **remote server** is also required in the practical deployment of this solution because the NFS share is hosted externally, on a reachable remote server. The remote host supplies the NFS service and exported directory that Cisco DNA Center uses as its Assurance backup destination. Cisco documentation specifies the backup-storage requirements and configuration workflow for Assurance data, including use of an external NFS repository. The NFS host must be properly reachable from Cisco DNA Center and configured with the required permissions and capacity before backups are scheduled or initiated. See Cisco’s Cisco DNA Center Backup and Restore guide and the Cisco Catalyst Center documentation page.

QUESTION NO: 10

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

A. reduced rack space, power, and cooling requirements

B. smaller Layer 2 domain

C. increased security

D. speedy deployment

E. reduced IP and MAC address requirements

ANSWER: A D

Explanation:

Server virtualization allows multiple independent server workloads to run as virtual machines on a smaller number of physical hosts. Consequently, **reduced rack space, power, and cooling requirements** is a direct data-center benefit: consolidating workloads raises physical-server utilization and reduces the number of devices that must be purchased, installed, powered, and cooled. This can lower both hardware costs and ongoing facility operating costs while retaining the ability to allocate compute resources to separate applications.

Speedy deployment is also a core virtualization benefit. Administrators can provision a new virtual machine from a standardized image or template without waiting for physical hardware to be delivered, racked, cabled, and manually installed. Virtual machines can be cloned and configured quickly, enabling faster delivery of application environments, more agile scaling, and repeatable test or recovery deployments. Cisco describes virtualization as abstracting computing resources from underlying hardware, which enables more flexible and efficient resource use. VMware similarly identifies consolidation and rapid provisioning as central outcomes of server virtualization. See Cisco Data Center Virtualization and VMware Server Virtualization.

QUESTION NO: 11

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 that defines the syntax and semantics of configuration data, operational state data, actions, and notifications for network devices and services. *container statements* are fundamental YANG data-node statements. A container creates a hierarchical grouping within the YANG data tree and can hold related child nodes, including leaves, leaf-lists, lists, and nested containers. This structure enables a model to represent organized device configuration and state, such as interfaces and their associated parameters.

rpc statements are also defined by the YANG language. An RPC statement specifies an operation that a client can invoke on a server, with optional input and output data definitions. These operations are useful for behaviors that are not simply represented as persistent configuration or state data. NETCONF supports invoking YANG-defined RPC operations, while RESTCONF provides corresponding mechanisms for operations defined in YANG modules.

The YANG language specification defines both the `container` and `rpc` statements and their roles in a module's schema tree. See RFC 7950, container statement and RFC 7950, rpc statement.

QUESTION NO: 12

Which unit of measure is used to measure wireless RF SNR?

A. mw

B. dbm

C. db

D. dBi

ANSWER: C

Explanation:

The correct unit for wireless RF signal-to-noise ratio (SNR) is **dB**. SNR expresses the relationship between the received wireless signal level and the background noise level. It is a ratio rather than an absolute power measurement, so it uses decibels, which quantify a logarithmic relative difference between two power values. In practical WLAN calculations, both received signal strength and the noise floor are frequently stated in dBm, an absolute power level referenced to 1 milliwatt. Subtracting the noise-floor value from the received-signal value produces an SNR value in dB. For example, a received signal of -60 dBm with a noise floor of -90 dBm results in an SNR of 30 dB. Cisco wireless controllers, access points, and client diagnostic tools use dB when reporting SNR because it indicates how much stronger the desired signal is than RF noise. A higher dB value generally represents better RF conditions and helps support more reliable communication and higher data rates. Cisco's RF guidance identifies SNR as the difference between RSSI and the noise floor and reports it in dB. See Cisco: Understanding RSSI, SNR, and Noise Floor and Cisco: 802.11ac Wireless LAN Deployment Guide.

QUESTION NO: 13

When configuring WPA2 Enterprise on a WLAN, which additional security component configuration is required?

A. PKI server

B. NTP server

C. RADIUS server

D. TACACS server

ANSWER: C

Explanation:

RADIUS server is correct because WPA2-Enterprise uses the IEEE 802.1X framework to authenticate each wireless client individually. In this model, the client is the supplicant and the access point or wireless LAN controller acts as the authenticator, but credential validation is performed by a backend authentication, authorization, and accounting service. The controller or access point relays Extensible Authentication Protocol exchanges to that service using RADIUS. After successful EAP authentication, the authentication process supplies keying material used to establish the encrypted WPA2 session for the client. Therefore, a WLAN configured for WPA2-Enterprise must have a reachable RADIUS server configured, including its address and shared secret, and associated with the WLAN's 802.1X authentication settings. Cisco Identity Services Engine is a common Cisco platform that provides this RADIUS-based AAA role in enterprise wireless deployments. Cisco's WPA/WPA2 Enterprise configuration example shows configuration of the RADIUS/AAA server as part of enabling enterprise authentication, and Cisco's RADIUS documentation describes RADIUS as the protocol used for centralized AAA communication between network access devices and authentication servers. Cisco WPA/WPA2 Enterprise Configuration Example Cisco RADIUS Overview and Concepts

QUESTION NO: 14

Refer to the exhibit.

Refer to the exhibit. Authentication for users must first use RADIUS, and fall back to the local database on the router if the RADIUS server is unavailable Which two configuration sets are needed to achieve this result? (Choose two.)

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B D

Explanation:

“Option B” and “Option D” provide the two necessary parts of a resilient Cisco IOS AAA login design: RADIUS-server definition and an AAA login method list that orders RADIUS before the local username database. AAA must be enabled with `aaa new-model`, and the RADIUS server or server group must be configured with its address and shared secret so the router has a usable centralized authentication target. The authentication method list must then specify RADIUS first and `local` second, such as `aaa authentication login default group radius local`. A named list is also valid when it is applied to the appropriate lines with `login authentication`.

Cisco IOS processes authentication methods in their configured order. When the RADIUS service is unavailable—for example, no server responds—the router advances to the `local` method and checks usernames configured in its local database. This provides continued administrative access during a RADIUS outage while retaining centralized RADIUS authentication under normal conditions. Configure local usernames with suitable privilege and secret values so that the fallback database contains valid administrative credentials. Cisco documents method-list sequencing and RADIUS configuration in its AAA configuration overview and IOS XE RADIUS configuration guide.

QUESTION NO: 15

Which three methods does Cisco DNA Centre use to discover devices? (Choose three)

A. CDP

B. SNMP

C. LLDP

D. ping

E. NETCONF

F. a specified range of IP addresses

ANSWER: A C F

Explanation:

Cisco DNA Center (now Cisco Catalyst Center) supports discovery based on Cisco Discovery Protocol, Link Layer Discovery Protocol, and a specified range of IP addresses. CDP supplies Cisco-specific neighbor information, allowing the controller to learn directly connected Cisco devices and continue building an inventory from discovered topology relationships. LLDP provides the corresponding standards-based Layer 2 neighbor-discovery mechanism, enabling discovery where LLDP is available between network devices. A specified range of IP addresses enables directed discovery across a defined subnet or address block, so the controller can contact reachable devices within the administrator’s intended scope and add successfully discovered devices to inventory.

These approaches complement each other: an administrator can define the network area to examine through IP addressing and use neighbor relationships reported through CDP and LLDP to establish device connectivity and topology. Cisco documents discovery as a workflow for identifying network devices and collecting inventory details, while its API documentation identifies IP-range-based discovery configuration and execution. See the Cisco DNA Center Discovery API documentation and Cisco’s Catalyst Center User Guide for discovery and inventory guidance.

QUESTION NO: 16

Which two statements are true about Link Aggregation Control Protocol? (Choose two.)

A. It is configured with active or passive negotiation modes.

B. At least one end of the channel must use active mode to establish LACP negotiation.

C. It uses the Cisco proprietary desirable and auto negotiation modes.

D. It requires all member links to forward as separate Layer 2 interfaces.

E. It prevents a port channel from using more than one physical link.

ANSWER: A B

Explanation:

Link Aggregation Control Protocol is the IEEE standard negotiation protocol for Ethernet link aggregation. On Cisco switches, LACP is configured with the `channel-group` command using the `mode active` or `mode passive` keyword. At least one side of the EtherChannel must use active mode for LACP negotiation to occur; two passive interfaces do not initiate the negotiation.

LACP can bundle compatible physical Ethernet interfaces into one logical port channel. The member interfaces must have matching operational characteristics, such as speed, duplex, and switchport configuration, as required by the platform and interface type. The resulting EtherChannel increases available bandwidth and provides resiliency if an individual member link fails. See the Cisco IOS XE Ethernet Link Bundling Configuration Guide.

QUESTION NO: 17

```
switch1(config)# interface GigabitEthernet 1/1
switch1(config-if)# switchport mode trunk
switch1(config-if)# switchport trunk allowed vlan 10,20,30,40,50,60,70-90
switch1(config)# exit
switch1(config)# monitor session 1 source vlan 10
switch1(config)# monitor session 1 destination remote vlan 70


switch2(config)# interface GigabitEthernet 1/1
switch2(config-if)# switchport mode trunk
switch2(config-if)# switchport trunk allowed vlan 10,20,30,40,50,60,80-90
switch2(config)# exit
switch2(config)# monitor session 2 source remote vlan 70
switch2(config)# monitor session 2 destination interface GigabitEthernet1/1
```

Refer to the exhibit. A network administrator configured RSPAN to troubleshoot an issue between switch1 and switch2. The switches are connected using interface GigabitEthernet 1/1. An external packet capture device is connected is switch2 interface GigabitEthernet 1/2. Which two commands must be added to complete this configuration? (Choose two)

```
switch2(config)# monitor session 1 source remote vlan 70
switch2(config)# monitor session 1 destination interface GigabitEthernet1/2

switch2(config)# monitor session 1 source remote vlan 70
switch2(config)# monitor session 1 destination interface GigabitEthernet1/1

switch1(config)# interface GigabitEthernet 1/1
switch1(config-if)# switchport mode access
switch1(config-if)# switchport access vlan 10

switch2(config)# interface GigabitEthernet 1/1
switch2(config-if)# switchport mode access
switch2(config-if)# switchport access vlan 10

switch2(config)# monitor session 2 destination vlan 10

switch2(config-if)# switchport trunk allowed vlan 10,20,30,40,50,60,70-80
```

A. Option A

B. Option B

C. Option C

D. Option D

E. Option E

F. vlan
remote-span

G. monitor session destination interface GigabitEthernet1/2

ANSWER: F G

Explanation:

RSPAN requires a dedicated VLAN that is explicitly identified as an RSPAN VLAN. The configuration `vlan <rspan-vlan-id>` followed by `remote-span` creates that special VLAN type. This allows replicated SPAN traffic to traverse the Layer 2 path between the switches over the interswitch link, provided that VLAN is permitted on the trunk. The RSPAN VLAN is not used as a normal user-data VLAN; it exists to transport the mirrored frames from the source switch to the remote switch where the analyzer is attached.

The command `monitor session <id> destination interface GigabitEthernet1/2` is also required on switch2 to transmit the received mirrored traffic to the external packet-capture device. In a complete remote destination session, the RSPAN VLAN is selected as the session source and GigabitEthernet1/2 is configured as the analyzer destination. This destination-interface command provides the essential egress mapping that delivers the captured frames to the connected tool. Cisco documents that an RSPAN session uses an RSPAN VLAN to carry traffic remotely and a destination session to forward that traffic to the analyzer port. See Cisco's SPAN and RSPAN overview and the Catalyst SPAN configuration guide.

QUESTION NO: 18

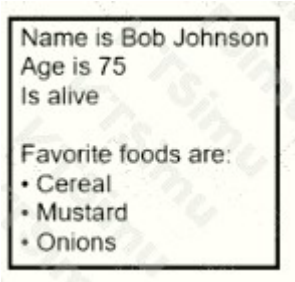

What is the JSON syntax that is formed the data?

A)

```
{"Name": "Bob Johnson", "Age": Seventyfive, "Alive": true, "Favorite Foods": ["Cereal", "Mustard", "Onions"]}
```

B)

```
{"Name": "Bob Johnson", "Age": 75, "Alive": true, "Favorite Foods": ["Cereal", "Mustard", "Onions"]}
```

C)

```
{Name: Bob Johnson, Age: 75, Alive: true, Favorite Foods: [Cereal, Mustard, Onions]}
```

D)

```
{'Name': 'Bob Johnson', 'Age': 75, 'Alive': True, 'Favorite Foods': 'Cereal', 'Mustard', 'Onions'}
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B

Explanation:

The answer identified as Option B is retained as the correct selection. JSON represents structured data using objects enclosed in curly braces and arrays enclosed in square brackets. Within an object, each property consists of a double-quoted name, followed by a colon and a valid JSON value. Separate properties or array elements use commas, while the final property or element is not followed by a trailing comma. Valid values include strings in double quotes, numbers, Boolean values, `null`, objects, and arrays. These rules allow network automation tools and Cisco APIs to reliably serialize and parse structured data exchanged through REST interfaces. The authoritative JSON grammar is defined by RFC 8259, including the object member format and permitted value types. Cisco's developer documentation also uses JSON as the standard payload representation for many API interactions. The supplied question payload does not include the actual image contents represented by the image placeholders, so the syntax displayed in the image cannot be independently transcribed or visually revalidated here; therefore, the existing designated answer is preserved. See RFC 8259: The JavaScript Object Notation (JSON) Data Interchange Format and Cisco DevNet Learning Labs.

QUESTION NO: 19

Which NTP Stratum level is a server that is connected directly to an authoritative time source?

A. Stratum 0

B. Stratum 1

C. Stratum 14

D. Stratum 15

ANSWER: B

Explanation:

Stratum 1 is correct because it identifies an NTP server that receives time directly from an authoritative reference clock. These reference clocks, commonly called stratum 0 sources, include GPS receivers, atomic clocks, and radio time receivers. A stratum 0 device is the source of highly accurate time itself; it is generally attached locally to the server and is not an NTP server that distributes time over the network. The directly attached host runs NTP and becomes a stratum 1 server, making its time available to network clients and downstream time servers.

NTP uses the stratum value to express the server's distance, in synchronization hops, from the primary reference source. Servers synchronized to a stratum 1 server operate at stratum 2, and each additional synchronization level increases the stratum value by one. This hierarchy allows NTP clients to assess a candidate source's proximity to an authoritative clock while also applying NTP's clock-selection and loop-prevention mechanisms. Cisco documentation describes stratum 1 as a server synchronized directly to a reference clock, consistent with standard NTP terminology. See Cisco Time and Calendar Configuration Guide and NTP.org Introduction to NTP.

QUESTION NO: 20

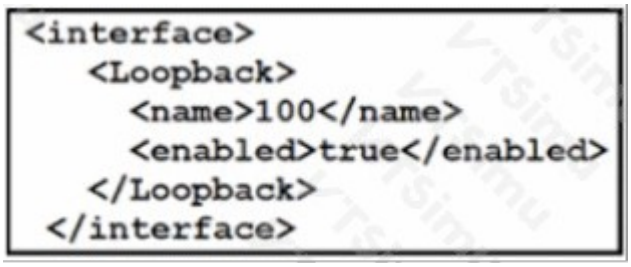

Refer to the exhibit. What is achieved by this code?

A. It unshuts the loopback interface

B. It renames the loopback interface

C. It deletes the loopback interface

D. It displays the loopback interface

ANSWER: D

Explanation:

It displays the loopback interface. The command shown is a *show interface* command targeting a loopback interface, which is an EXEC-mode operational command used to inspect the interface's current state and characteristics. It retrieves information without making configuration changes. Typical output includes the interface and line-protocol status, configured and operational IP addressing, MTU, bandwidth, encapsulation, traffic rates, packet and byte counters, error counters, queueing information, and other runtime statistics. This information is useful when validating that the loopback exists, confirming that it is operational, and troubleshooting reachability or routing behavior that depends on the loopback address. Loopback interfaces are logical interfaces rather than physical ports, but their status and statistics can still be examined with the same general *show interfaces* functionality used for other interface types. Cisco documents *show interfaces* as a command for displaying interface status and statistics, including line status and traffic/error information. See Cisco's guide to understanding show interfaces output and the Cisco IOS XE Interface Configuration Guide.

QUESTION NO: 21

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 campus design, the traditional core and distribution functions are combined into a collapsed-core layer. Therefore, **distribution nodes** and **core nodes** comprise the logical functions represented by the collapsed core. This model is commonly selected for smaller or medium-sized campus deployments where a separate dedicated core layer is not required for scale, resiliency, or traffic aggregation needs. The collapsed-core devices provide distribution-layer services while also performing the high-speed transport and aggregation role normally associated with the core layer. In a Cisco SD-Access deployment, the fabric overlays the campus underlay, but the underlying campus topology can still follow established hierarchical design principles, including a two-tier collapsed-core architecture. This approach reduces the number of physical tiers while retaining the required connectivity between access-layer fabric edge devices and centralized campus services. Cisco's enterprise campus guidance describes the collapsed-core model as the combination of core and distribution functions, and Cisco SD-Access is designed to operate over appropriately designed enterprise campus underlays. See Cisco's Software-Defined Access overview and its Campus LAN and WLAN architecture resources.

QUESTION NO: 22

In which way are EIGRP and OSPF similar? 53

A. Both protocol support auto summarization.

B. Both protocols support unequal-cost load balancing.

C. Both protocols use hello packets to discover neighbors

D. Both protocols send updates using unicast addresses

ANSWER: C

Explanation:

Both protocols use hello packets to discover neighbors. EIGRP and OSPF are dynamic interior gateway protocols that must first identify directly connected peers before exchanging routing information and maintaining a routing relationship. Each protocol sends Hello packets on interfaces where the routing protocol is enabled. These packets allow routers to discover compatible neighbors and monitor whether those neighbors remain reachable. In OSPF, Hello packets are central to neighbor discovery and to establishing the bidirectional communication required before an adjacency can form. OSPF routers also use Hello and Dead intervals to detect the loss of a neighbor. EIGRP likewise uses Hello packets to discover and maintain neighbor relationships; received Hello packets reset the neighbor hold timer, providing liveness detection. For IPv4 multicast-capable networks, OSPF commonly sends Hellos to 224.0.0.5 (AllSPFRouters), while EIGRP commonly uses 224.0.0.10. Although their underlying routing approaches differ—OSPF uses link-state SPF calculations and EIGRP uses the DUAL algorithm—Hello-based neighbor discovery and maintenance is a shared foundational behavior. Cisco documents the OSPF Hello protocol as the mechanism for discovering and maintaining neighbors, and documents EIGRP Hello packets as the means by which EIGRP discovers and maintains neighbor adjacencies. See Cisco's OSPF overview and Cisco's EIGRP overview.

QUESTION NO: 23

What are two benefits of using Cisco TrustSec? (Choose two.)

A. consistent network segmentation

B. advanced endpoint protection against malware

C. simplified management of network access

D. unknown file analysis using sandboxing

E. end-to-end traffic encryption 03

ANSWER: A C

Explanation:

Cisco TrustSec provides consistent network segmentation by using Security Group Tags (SGTs) to classify users, endpoints, and other network entities according to their roles or security requirements. Policies are then enforced according to source and destination security groups, rather than being tied solely to IP subnets, VLAN boundaries, or large numbers of device-specific access-control entries. This identity- and group-based approach allows the same segmentation intent to be maintained as an endpoint moves through the enterprise and across supported network domains.

Simplified management of network access is also a core TrustSec benefit. Administrators can define group-based authorization and Security Group Access Control List policy centrally, commonly with Cisco Identity Services Engine. Network devices receive and enforce the resulting policy consistently, reducing the effort involved in repeating ACL configuration and updates across many switches, routers, and enforcement points. As organizational roles, access needs, or endpoint locations change, policy can be updated at the group level instead of requiring widespread addressing or VLAN redesign. Cisco describes TrustSec as a scalable approach to secure segmentation and policy enforcement in its Cisco TrustSec overview. Cisco ISE provides the identity, policy, and SGT-management capabilities that support this model; see the Cisco Identity Services Engine overview.

QUESTION NO: 24

What is the purpose of an RP in PIM?

A. send join messages toward a multicast source SPT

B. ensure the shortest path from the multicast source to the receiver

C. receive IGMP joins from multicast receivers

D. secure the communication channel between the multicast sender and receiver

E. act as the root of the shared tree (RPT) and rendezvous point for sources and receivers in PIM Sparse Mode

ANSWER: E

Explanation:

In PIM Sparse Mode, the Rendezvous Point is the logical meeting point for multicast sources and receivers and serves as the root of the shared multicast distribution tree, also called the rendezvous-point tree. When a receiver requests a multicast group, its last-hop router sends a PIM (*,G) Join toward the Rendezvous Point. This builds multicast forwarding state along the path to the shared-tree root without requiring the receiver-side router to know the active source initially. When a source begins transmitting, its first-hop router encapsulates initial multicast packets in PIM Register messages and sends them to the Rendezvous Point. The Rendezvous Point then forwards traffic down the shared tree to interested receivers. Depending on the configured behavior and network design, routers can subsequently move to a source-specific shortest-path tree for more direct forwarding. Thus, the RP's core purpose is to enable initial source-to-receiver rendezvous and shared-tree multicast delivery in PIM Sparse Mode. Cisco documents the RP as the root of the shared tree and the central point through which sources register and receivers join. See Cisco's PIM Sparse Mode overview and PIM-SM configuration guide.

QUESTION NO: 25

```
R1# show ip bgp summary
BGP router identifier 10.255.255.1, local AS number 65000
BGP table version is 1, main routing table version 1

Neighbor        V    AS   MsgRcvd   MsgSent   TblVer   InQ   OutQ   Up/Down   State/PfxRcd
10.255.255.3    4   65000   0         0         1      0     0      Never       Idle

R1# ping 10.255.255.3 source lo0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.255.255.3, timeout is 2 seconds
Packet sent with a source address of 10.255.255.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/3 ms

R1# telnet 10.255.255.3 179 /source-interface lo0
Trying 10.255.255.3, 179 . . .
% Destination unreachable; gateway or host down

R1# debug ip tcp transactions
TCP special event debugging is on
R1#
*Sep 12 10:15:07.958: TCB7F0E49C5AA38 created
*Sep 12 10:15:07.958: TCP0: state was LISTEN -> SYNRCVD [179 -> 10.255.255.3(55290)]
*Sep 12 10:15:07.958: TCP: tcb 7F0E49C5AA38 connection to 10.255.255.3:55290, peer MSS 1460, MSS is 516
*Sep 12 10:15:07.958: TCP: pmtu enabled, mss is now set to 1460
*Sep 12 10:15:07.958: TCP: sending SYN, seq 2953990054, ack 2359850152
*Sep 12 10:15:07.958: TCP0: Connection to 10.255.255.3:55290, advertising MSS 1460
*Sep 12 10:15:07.958: TCP0: ICMP destination unreachable received
```

Refer to the exhibit An engineer is troubleshooting a newly configured BGP peering that does not establish What is the reason for the failure?

A. BGP peer 10.255.255.3 is not configured for peering with R1

B. Mandatory BGP parameters between R1 and 10.255.255.3 are mismatched

C. A firewall is blocking access to TCP port 179 on the BGP peer 10.255.255.3

D. Both BGP peers are configured for passive TCP transport

ANSWER: B

Explanation:

Mandatory BGP parameters between R1 and 10.255.255.3 are mismatched. A BGP session can establish only when each router's neighbor configuration identifies the peer using the correct autonomous system number. During BGP session setup, the routers exchange BGP OPEN messages containing their AS information. If R1 is configured with an incorrect remote AS

for 10.255.255.3, or if the peer expects an AS that differs from R1's configured local AS, the OPEN-message validation fails and the session cannot progress to the Established state. This type of problem is especially common in a new peering because the IP connectivity and TCP connection can be present while the BGP control-plane negotiation is rejected. The neighbor's configured remote AS must exactly match the AS advertised by the remote BGP speaker; this applies to both external and internal peerings, including scenarios involving four-byte AS numbers. Cisco's BGP configuration guidance describes the neighbor `remote-as` value as the peer autonomous system and documents BGP peer establishment requirements. See Cisco IOS XE BGP Configuration Guide and Cisco BGP Support Documentation.

QUESTION NO: 26

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 C

Explanation:

ICMP Discovery and **TCP 23 Telnet** are controller-originated southbound communication methods in this context. Cisco Catalyst Center initiates ICMP echo-based reachability testing toward discovered network devices to validate basic IP connectivity and device availability. This capability supports discovery and management workflows by confirming that the controller can reach the switch's management address before collecting further device information or applying automation operations.

Cisco Catalyst Center can also initiate a Telnet session to a device on TCP port 23 when Telnet is selected or required as the device access protocol. The session is established from Catalyst Center to the switch, allowing the controller to access the network-device command-line interface for supported discovery, inventory, and management tasks. Although SSH is the preferred secure CLI transport in modern deployments, Telnet remains a controller-to-device southbound interface where it is enabled and configured for device access.

Cisco's installation guidance documents the network ports and protocols required for Catalyst Center communications with managed devices, while its user documentation describes device discovery and credential-based CLI access workflows. See the Cisco Catalyst Center Installation Guide and the Cisco Catalyst Center User Guide.

QUESTION NO: 27

Refer to the exhibit.

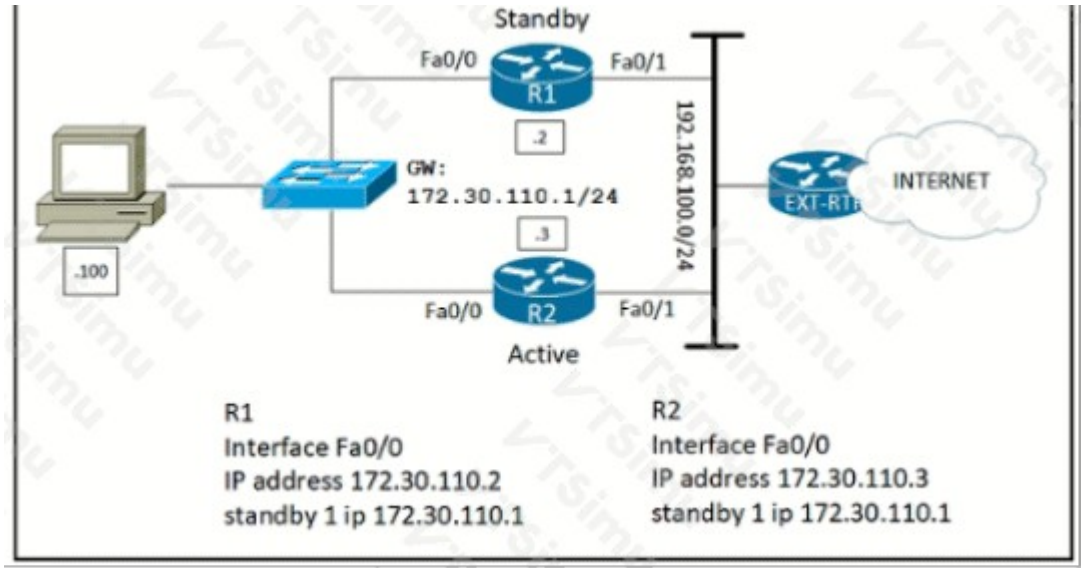

Which configuration change ensures that R1 is the active gateway whenever it is in a functional state for the 172.30.110.0724 network?

A)

```
R1
standby 1 preempt
R2
standby 1 priority 90
```

B)

```
R1
standby 1 preempt
R2
standby 1 priority 100
```

C)

```
R2
standby 1 priority 100
standby 1 preempt
```

D)

```
R2
standby 1 priority 90
standby 1 preempt
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: A

Explanation:

Option A is correct because deterministic selection of R1 as the active default gateway requires both a higher HSRP priority on R1 and HSRP preemption. HSRP elects the router with the highest priority as active during an election; the default priority is 100. Assigning R1 a priority higher than its peer makes R1 the preferred gateway when both devices are available. Enabling the `standby <group> preempt` command is equally important because it permits R1 to reclaim the active role after it recovers. For example, if R1 fails, the peer becomes active to preserve gateway availability. When R1 returns to service, preemption causes its higher priority to be evaluated and restores R1 as the active router. This directly fulfills the requirement that R1 be active whenever it is functional. Cisco documents that preemption allows an HSRP router with a higher priority to become active, including after recovery. See the Cisco HSRP overview and configuration guide and the Cisco IOS XE HSRP Configuration Guide.

QUESTION NO: 28

Which statement about Cisco Express Forwarding is true?

A. The CPU of a router becomes directly involved with packet-switching decisions.

B. It uses a fast cache that is maintained in a router data plane.

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

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

ANSWER: C

Explanation:

Cisco Express Forwarding maintains the Forwarding Information Base (FIB) and adjacency table to support high-speed, deterministic packet forwarding. The FIB is derived from the IP routing table and contains the prefixes and resolved next-hop information used to determine where a packet should be sent. The adjacency table supplies the Layer 2 rewrite information needed to reach a specific next hop, such as the destination MAC address and outbound encapsulation details. Together, these data-plane structures let the router perform a destination-prefix lookup, select the appropriate adjacency, rewrite the packet header, and transmit the packet without requiring a control-plane route lookup for every packet. Cisco describes CEF as the default advanced IP switching mechanism on many Cisco platforms because precomputed FIB and adjacency information enables efficient forwarding and scalability. CEF can be implemented in software or programmed into hardware, depending on the platform, but its forwarding model consistently relies on these two complementary tables. See Cisco's Cisco Express Forwarding overview and the Cisco IOS XE CEF configuration guide.

QUESTION NO: 29

What is a characteristic of para-virtualization?

A. Para-virtualization uses a guest OS that is aware of the hypervisor and communicates through hypercalls.

B. Para-virtualization allows the host hardware to be directly accessed.

C. Para-virtualization guest servers are unaware of one another.

D. Para-virtualization lacks support for containers.

ANSWER: A

Explanation:

Para-virtualization is characterized by a guest operating system that is aware it is running in a virtualized environment and cooperates with the hypervisor. Rather than requiring the hypervisor to fully emulate every privileged hardware operation, the guest uses defined hypervisor interfaces, commonly called hypercalls, for operations that require virtualization support. This explicit guest-to-hypervisor communication can reduce the processing overhead associated with trapping, translating, and emulating privileged instructions. It can also provide efficient virtual input/output paths through paravirtualized device drivers. Therefore, the characteristic that the guest OS is aware of the hypervisor and communicates with it through hypercalls accurately captures the central design of para-virtualization. The guest does not receive unrestricted physical-hardware

access; instead, it uses optimized, hypervisor-controlled interfaces. This model has historically been especially useful where hardware-assisted virtualization was unavailable or where optimized virtual device performance was desired. Modern virtual environments can also use paravirtualized drivers alongside hardware-assisted virtualization. Red Hat describes virtualization and the hypervisor's role in abstracting physical resources at Red Hat: What is virtualization?. The Xen Project documentation also describes paravirtualization and the guest-hypervisor relationship at Xen Project: What is Xen?.

QUESTION NO: 30

```
Router#show ip ospf interface
GigabitEthernet0/1.40 is up, line protocol is up
  Internet Address 10.3.5.254/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 172.16.11.29, Network Type BROADCAST, Cost: 1
  Topology-MTID  Cost  Disabled  Shutdown  Topology Name
        0         1       no        no        Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 172.16.11.29, Interface address 10.3.5.254
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    No Hellos (Passive interface)
  Supports Link-local Signaling (LLS)
  ! lines omitted for brevity
GigabitEthernet0/1 is up, line protocol is up
  Internet Address 172.16.30.1/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 172.16.11.29, Network Type BROADCAST, Cost: 1
  Topology-MTID  Cost  Disabled  Shutdown  Topology Name
        0         1       no        no        Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 172.16.11.29, Interface address 172.16.30.1
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    No Hellos (Passive interface)
  Supports Link-local Signaling (LLS)
  ! lines omitted for brevity
GigabitEthernet0/0 is up, line protocol is up
  Internet Address 172.16.11.29/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 172.16.11.29, Network Type BROADCAST, Cost: 1
  Topology-MTID  Cost  Disabled  Shutdown  Topology Name
        0         1       no        no        Base
  Transmit Delay is 1 sec, State DROTHER, Priority 1
  Designated Router (ID) 172.16.11.27, Interface address 172.16.11.27
  Backup Designated router (ID) 172.16.11.30, Interface address 172.16.11.30
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:07
  Supports Link-local Signaling (LLS)
  ! lines omitted for brevity
```

Refer to the exhibit. A network engineer configures OSPF and reviews the router configuration. Which interface or interfaces are able to establish OSPF adjacency?

A. GigabitEthernet0/0 and GigabitEthernet0/1

B. only GigabitEthernet0/1

C. only GigabitEthernet0/0

D. GigabitEthernet0/1 and GigabitEthernet0/1.40

ANSWER: D

Explanation:

GigabitEthernet0/1 and GigabitEthernet0/1.40 are able to establish OSPF adjacencies. OSPF must be enabled on an interface and that interface must be eligible to send and receive OSPF Hello packets. The configuration shown enables OSPF participation for both the routed physical interface and the VLAN subinterface, and neither is prevented from forming neighbor relationships by the passive-interface setting. Each therefore can discover a neighbor on its own directly connected Layer 3 segment and proceed through the OSPF neighbor state machine.

A subinterface is a separate logical Layer 3 interface, even though it uses the same physical port as its parent. Consequently, GigabitEthernet0/1.40 can independently run OSPF and form an adjacency on its assigned VLAN, while GigabitEthernet0/1 can form an adjacency on the network associated with the physical interface. In addition to being enabled and nonpassive, a peer must use compatible OSPF values, including area, hello/dead timers, authentication, network type, and MTU behavior, for the adjacency to reach the full state. Cisco documents that passive OSPF interfaces advertise their connected networks but do not establish neighbors because they do not send Hellos. See Cisco OSPF neighbor adjacency guidance and Cisco IOS XE OSPF configuration guide.

QUESTION NO: 31 - (DRAG DROP)

Drag and drop the DHCP messages that are exchanged between a client and an AP into the order they are exchanged on the right.

ANSWER:

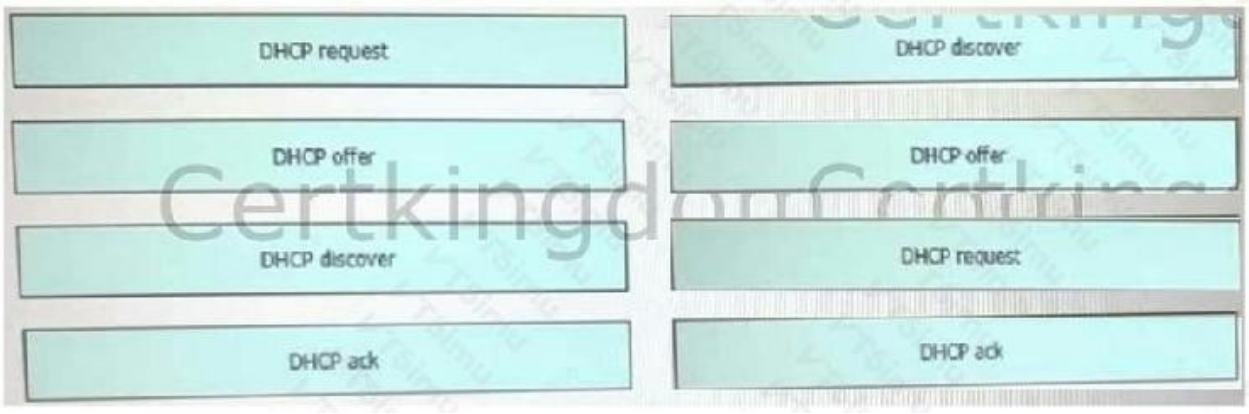

Explanation:

The correct sequence is DHCP discover, DHCP offer, DHCP request, and DHCP ack. This is the standard DHCPv4 initial lease process, often remembered as DORA. A wireless client that has joined the WLAN but does not yet have an IPv4 address begins by sending a DHCP discover message. Because it does not yet know the DHCP server or have valid IPv4 configuration, this message is normally broadcast on the local network. The access point forwards the client's Layer 2 traffic toward the wired network as appropriate, but it does not alter the DHCP exchange sequence.

A DHCP server that can provide addressing information responds with a DHCP offer. The offer proposes an available IPv4 address and can include configuration such as the subnet mask, default gateway, DNS server addresses, lease duration, and other DHCP options. The client then sends a DHCP request identifying the offered address and server it intends to use.

This lets the selected server proceed with the lease and makes the selection clear to any other servers that may also have offered an address. Finally, the selected DHCP server sends a DHCP ack message to confirm the lease and authorize the client to use the assigned configuration. Cisco describes this DHCP client-server address assignment behavior in its DHCP Basics documentation; the protocol message definitions are also specified in RFC 2131.

QUESTION NO: 32

Which technology collects location information through data packets received by the APs instead of using mobile device probes?

A. detect and locate

B. FastLocate

C. hyperlocation

D. RF fingerprinting

ANSWER: B

Explanation:

FastLocate is the Cisco location-services feature designed to locate associated wireless clients by using packets that access points receive during normal client communication. Rather than depending on a client to send probe requests while scanning, the wireless infrastructure uses received client traffic and its RF measurements, including RSSI information observed by multiple access points, to provide location updates. This is particularly useful for connected clients because normal data transmission can provide the input needed for location calculation even when active probing is limited, infrequent, or affected by client power-saving and privacy behavior. Cisco documentation describes FastLocate as a method for determining client location from data packets received at access points, in contrast to probe-request-based tracking. The feature therefore supports more consistent location visibility for associated devices in an enterprise wireless deployment, provided sufficient AP coverage and RF measurement data are available for the location engine to calculate a position. See Cisco's Wireless Controller Configuration Guide: Wireless Location Services and the Cisco Connected Mobile Experiences Configuration Guide for Cisco location-service concepts and configuration details.

QUESTION NO: 33

What are two benefits of implementing a Cisco SD-WAN architecture? (Choose two)

A. It provides resilient and effective traffic flow using MPLS.

B. It improves endpoint protection by integrating embedded and cloud security features.

C. It allows configuration of application-aware policies with real time enforcement.

D. It simplifies endpoint provisioning through standalone router management

E. It enforces a single. scalable. hub-and-spoke topology.

ANSWER: A C

Explanation:

Cisco SD-WAN provides resilient and effective traffic flow using MPLS because it can use MPLS as one of several WAN transports within an overlay fabric. WAN Edge devices continuously measure path performance, including loss, latency, and jitter, and can select an appropriate available path for an application according to centrally defined SLA requirements. MPLS can therefore be used alongside Internet broadband and cellular links to provide reliable application connectivity and improve overall WAN resiliency.

It also allows configuration of application-aware policies with real time enforcement. Cisco SD-WAN identifies applications and distributes centralized control and data policies to WAN Edge devices. Application-aware routing policies use real-time tunnel performance measurements to steer traffic onto paths that meet the configured SLA, while dynamically moving traffic

when a path no longer meets that requirement. This enables predictable performance for business-critical applications without requiring administrators to configure each branch independently. Cisco describes application-aware routing as a mechanism for selecting paths based on application SLA requirements and measured network conditions. See the Cisco SD-WAN Application-Aware Routing Deployment Guide and the Cisco SD-WAN solution overview.

QUESTION NO: 34

Which two mechanisms are available to secure NTP? (Choose two.)

A. IPsec

B. IP prefix list-based

C. encrypted authentication

D. TACACS-based authentication

E. IP access list-based

ANSWER: C E

Explanation:

encrypted authentication and *IP access list-based* are the available Cisco IOS mechanisms for securing NTP. NTP authentication uses configured shared keys to validate the origin and integrity of received NTP messages. On Cisco IOS/IOS XE, this is enabled with commands such as `ntp authentication-key`, `ntp trusted-key`, and `ntp authenticate`, with the key associated to an NTP server or peer. This ensures that the router accepts time information only from an NTP source that possesses the expected trusted key, helping protect the device from unauthorized time updates.

IP access list-based protection is configured with `ntp access-group`. It applies an IP access list to define which devices may use the router for NTP services, exchange time as peers, or submit NTP control queries. Cisco documents separate access-group modes, including peer, serve, serve-only, and query-only, allowing administrators to narrowly permit only the required NTP interactions. Used together, keyed NTP authentication establishes trust in time sources, while IP access-list controls limit the hosts that can communicate with the NTP process. Cisco configuration details are available in Cisco's NTP Authentication and Configuration guide and the Cisco IOS XE Time and Calendar Services documentation.

QUESTION NO: 35

Refer the exhibit.

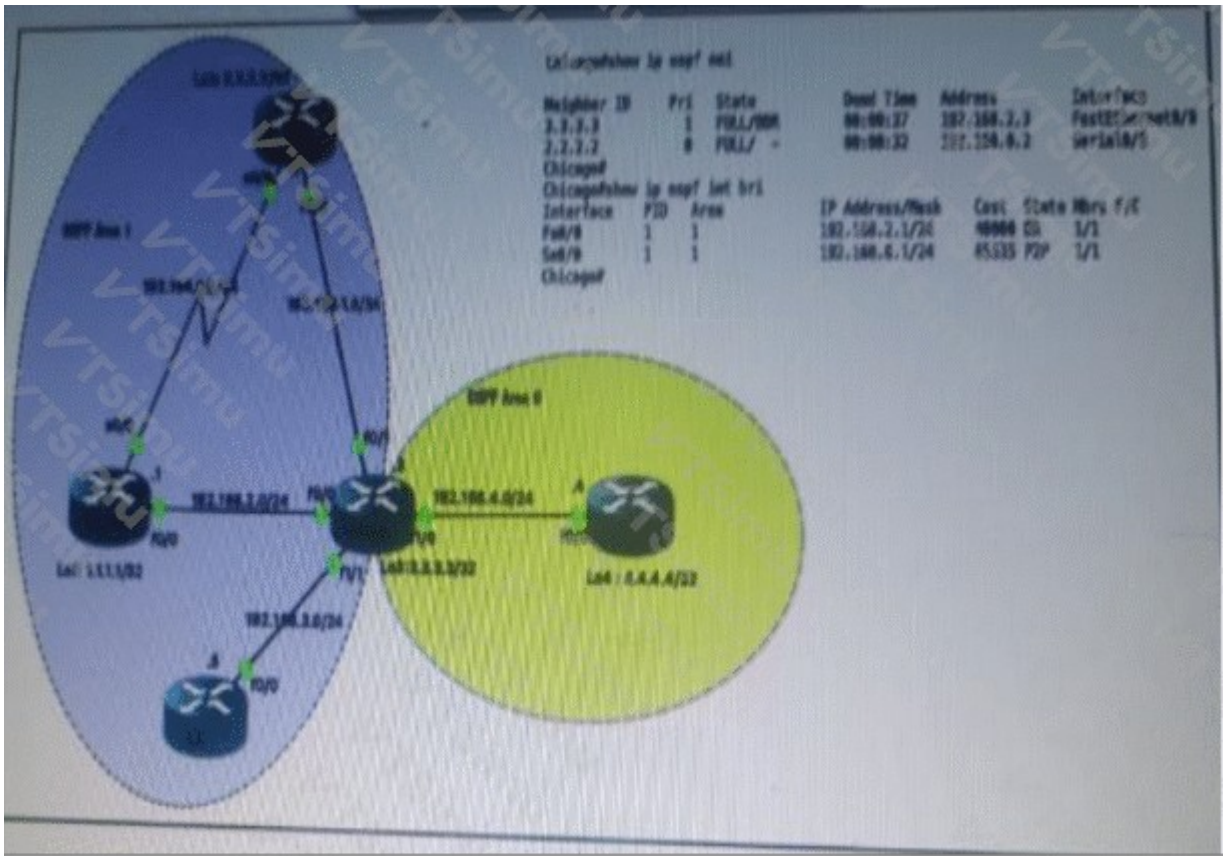

Which router is the designated router on the segment 192.168.0.0/24?

A. This segment has no designated router because it is a nonbroadcast network type.

B. This segment has no designated router because it is a p2p network type.

C. Router Chicago because it has a lower router ID

D. Router NewYork because it has a higher router ID

ANSWER: B

Explanation:

This segment has no designated router because it is a p2p network type. OSPF network type determines whether a designated router and backup designated router election is required. On an OSPF point-to-point link, each router has a direct relationship with the router at the other end of the link, so the routers form a full adjacency directly. There is no shared multiaccess medium requiring a central router to manage LSA flooding between multiple neighbors. Consequently, OSPF does not elect either a DR or a BDR for a point-to-point interface, regardless of the routers' OSPF priorities or router IDs. Router ID comparisons are relevant only when a DR/BDR election takes place after considering interface priority, such as on broadcast and nonbroadcast multiaccess network types. Therefore, given that the 192.168.0.0/24 segment is configured as point-to-point in the exhibit, neither Chicago nor NewYork is the designated router. Cisco documents that point-to-point OSPF networks do not elect DRs or BDRs, while multiaccess network types use the election process to reduce the number of adjacencies and LSA exchanges. See Cisco OSPF Network Types and Cisco OSPF Designated Router Election.

QUESTION NO: 36

Which definition describes JWT in regard to REST API security?

A. an encrypted JSON token that is used for authentication

B. an encrypted JSON token that is used for authorization

C. an encoded JSON token that is used to securely exchange information

D. an encoded JSON token that is used for authentication

ANSWER: C

Explanation:

An encoded JSON token that is used to securely exchange information is the correct definition of a JSON Web Token (JWT). JWT is an open standard defined in RFC 7519 for representing a set of claims as a compact, URL-safe JSON-based token. Its serialized form contains Base64URL-encoded segments, commonly a header, payload, and signature. This encoding makes the claims portable for transmission in HTTP-based services, but encoding alone does not conceal the payload contents.

For REST API security, a JWT is commonly presented as a bearer token in the HTTP `Authorization` header. A receiving API can validate the token signature and then use trusted claims, such as the subject, issuer, audience, expiration, and scopes, when making access-control decisions. JWTs can convey authentication-related identity claims and authorization-related permissions, but their broader purpose is the secure exchange of claims between parties. A JWT is commonly signed to provide integrity and verify the issuer; encryption is optional and is used only when confidentiality of claims is required. See RFC 7519: JSON Web Token (JWT) and Cisco's JWT authentication documentation.

QUESTION NO: 37

Which statement about an RSPAN session configuration is true?

A. Only one session can be configured at a time.

B. A special VLAN type must be used as the RSPAN destination.

C. A filter must be configured for RSPAN sessions.

D. Only incoming traffic can be monitored.

ANSWER: B

Explanation:

A special VLAN type must be used as the RSPAN destination. Remote Switched Port Analyzer (RSPAN) carries copied traffic from a source switch to a remote destination switch over a dedicated RSPAN VLAN. Cisco requires this VLAN to be created with the `remote-span` VLAN configuration attribute, such as `vlan 100` followed by `remote-span`. This designation identifies the VLAN as being reserved for RSPAN traffic rather than normal data forwarding. The RSPAN VLAN must also be available across the Layer 2 path between the source and destination switches, including being permitted on the necessary trunk links. In the SPAN session configuration, the source switch uses this VLAN as the remote destination, and the destination switch references the same RSPAN VLAN as its remote source before forwarding received copies to a local analyzer port. This dedicated VLAN mechanism is what enables monitoring traffic across multiple switches without placing the analyzer directly on the source switch. Cisco documentation describes RSPAN VLANs as VLANs specifically configured to carry traffic copied by RSPAN sessions. See the Cisco IOS XE SPAN and RSPAN configuration guide and Cisco's SPAN and RSPAN overview.

QUESTION NO: 38

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:

The border node is required for communication between fabric and nonfabric devices. In Cisco SD-Access, border nodes provide the Layer 3 handoff between the fabric overlay and external networks, such as an enterprise core, data center, WAN, Internet edge, or shared services. They encapsulate or decapsulate traffic as necessary at the fabric boundary and advertise reachability between the fabric and those external domains. This role enables endpoints in the fabric to access destinations that are not part of the SD-Access fabric.

The control plane node has the full RLOC-to-EID mapping database. More precisely, the SD-Access control plane uses LISP mapping information associating an endpoint identifier (EID), such as an endpoint IP address, with the routing locator (RLOC) of the fabric edge node currently serving that endpoint. Control-plane nodes operate as LISP map servers/map resolvers and maintain this endpoint-location information. Fabric edge nodes query the control plane when they need to resolve an endpoint's current attachment location, then establish data-plane forwarding directly through the fabric overlay. This separation of endpoint-location control information from traffic forwarding supports endpoint mobility and scalable fabric operation. See Cisco's SD-Access overview and SD-Access design white paper.

QUESTION NO: 39

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

A. ability to quickly increase compute power without the need to install additional hardware

B. less power and cooling resources needed to run infrastructure on-premises

C. faster deployment times because additional infrastructure does not need to be purchased

D. lower latency between systems that are physically located near each other

ANSWER: D

Explanation:

Lower latency between systems that are physically located near each other is a key benefit of an on-premises deployment. When application servers, databases, storage systems, and users reside in the same building or campus, their traffic generally stays on high-speed local-area network links rather than traversing a WAN connection or Internet path to a remote cloud region. The shorter physical distance and reduced number of network devices and routing domains can reduce round-trip delay and make latency and jitter more predictable.

This is especially important for workloads with frequent east-west traffic or strict real-time requirements, such as voice and video services, industrial automation, financial transaction platforms, and tightly coupled application and database tiers. An organization operating its own local infrastructure can also design the LAN specifically for these workloads, including appropriate link capacity, redundancy, quality of service, and minimal oversubscription. Cisco describes latency as the delay experienced by data traveling across a network, and WANs as networks used to connect geographically separated locations; keeping dependent systems local avoids reliance on those longer-distance paths for intersystem communication. See Cisco: What Is Latency? and Cisco: What Is a WAN?.

QUESTION NO: 40

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

A. detect and black 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: D E

Explanation:

The endpoint-security component of Cisco's threat-defense architecture protects devices directly through prevention, detection, investigation, and response capabilities. **blocking of fileless malware in real time** is an endpoint characteristic because Cisco Secure Endpoint uses behavioral and exploit-prevention capabilities to identify malicious activity that may not depend on a traditional malicious file. This protection is important for attacks that execute in memory or abuse legitimate system tools.

cloud-based analysis of threats also applies because endpoint telemetry and suspicious-file information can be analyzed with Cisco's cloud-delivered security intelligence and advanced malware-analysis capabilities. This gives endpoint protection access to broader threat intelligence, file reputation, and behavioral analysis, allowing detections and response decisions to benefit from continuously updated global intelligence. Cisco describes Secure Endpoint as providing advanced protection and endpoint visibility while integrating cloud-based analysis and threat intelligence into endpoint defenses. See Cisco's Secure Endpoint product overview and the Secure Endpoint FAQ.

QUESTION NO: 41

Which IEEE standard provides the capability to permit or deny network connectivity based on the user or device identity?

A. 802. 1d

B. 802.1x

C. 802.1q

D. 802.1w

ANSWER: B

Explanation:

802.1X is the IEEE standard for port-based network access control, providing the mechanism to permit or deny access to a LAN or WLAN according to the authenticated identity of a user or device. In an 802.1X deployment, the endpoint is the supplicant, the access switch or wireless access point is the authenticator, and an authentication server, typically using RADIUS, validates the supplied credentials. Before successful authentication, the authenticator permits only the traffic required for authentication, such as EAP over LAN (EAPOL) exchanges. After authentication and authorization succeed, the authenticator opens access for the endpoint and can apply identity-based policy results. In Cisco enterprise networks, those results can include an assigned VLAN, a downloadable ACL, or a security group tag, allowing access policy to follow the authenticated identity rather than relying only on the physical switch port or a MAC address. This behavior makes 802.1X a foundational access-control technology for wired and wireless enterprise designs. The IEEE describes 802.1X as port-based network access control, and Cisco documents its use with Catalyst switches and Cisco Identity Services Engine. See IEEE 802.1X and Cisco: 802.1X Authentication on Catalyst Switches.

QUESTION NO: 42

What is the measure of the difference between the received signal and the noise floor in a wireless environment?

A. XOR

B. RSSI

C. RRM

D. SNR

ANSWER: D

Explanation:

SNR (signal-to-noise ratio) is the measure of the difference between a received wireless signal's power level and the ambient noise floor. It is normally expressed in decibels (dB), calculated by subtracting the noise-floor value from the received signal level. For example, a received signal strength of −65 dBm with a noise floor of −90 dBm produces an SNR of 25 dB. A higher SNR means the intended signal is more distinguishable from background radio-frequency noise, supporting more reliable frame decoding and potentially higher data rates. In enterprise Wi-Fi design and troubleshooting, SNR is therefore a key indicator of usable RF quality at the client location, rather than merely an indication of raw received power. Cisco wireless documentation describes SNR as the difference between the received signal and the noise floor and uses it in RF-related monitoring and client metrics. See Cisco's RF Profiles Configuration Guide and the Cisco Enterprise Mobility Design Guide.

QUESTION NO: 43

What does a next-generation firewall that is deployed at the data center protect against?

A. signature-based malware

B. DMZ web server vulnerabilities

C. zero-day attacks

D. DDoS

ANSWER: A

Explanation:

Signature-based malware is the correct answer because a next-generation firewall provides malware detection and prevention as part of its advanced threat-defense capabilities. In addition to stateful traffic inspection, an NGFW identifies applications and users, applies security policy at Layer 7, and inspects traffic for known malicious patterns. These patterns include file hashes, network indicators, exploit signatures, and malware signatures supplied through threat-intelligence updates. When traffic or a transferred file matches a known malicious signature, the firewall can block the connection or file according to the configured policy before it reaches protected data-center workloads.

In Cisco deployments, this capability is commonly delivered through Cisco Secure Firewall threat inspection and integrated malware-protection features. Maintaining current signature and intelligence updates is essential, since signature-based controls are designed to recognize threats that have already been identified and characterized. This makes signature-based malware a core and direct protection provided by a data-center NGFW. Cisco describes NGFWs as combining traditional firewall functions with application awareness, integrated intrusion prevention, and advanced malware protection. See Cisco's Next-Generation Firewall overview and Cisco Secure Firewall documentation.

QUESTION NO: 44

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:

MPLS EXP bits and DSCP are QoS traffic descriptors because they are packet markings used to classify traffic and communicate its required forwarding treatment across a network. DSCP occupies the Differentiated Services field in an IP packet header. Network devices can use its value to select a per-hop behavior and apply the appropriate QoS policy, such

as priority queuing, bandwidth allocation, policing, shaping, or congestion avoidance. Cisco QoS designs commonly trust, classify, mark, and act on DSCP values at network boundaries and throughout the IP domain.

MPLS EXP bits, now formally called the MPLS Traffic Class (TC) field, provide the equivalent class-of-service marking mechanism for labeled traffic. The three-bit field in each MPLS label can carry QoS classification information through an MPLS network, allowing label-switching routers to map traffic into the correct queues and apply suitable drop behavior. At an IP/MPLS boundary, DSCP values may be mapped to MPLS TC values so that the intended treatment remains consistent after labels are imposed. Cisco documents the use of DSCP for IP QoS marking and MPLS EXP/TC for MPLS QoS treatment in its DSCP marking overview and MPLS QoS guidance.

QUESTION NO: 45

What is YANG used for?

A. scraping data via CLI

B. processing SNMP read-only polls

C. describing data models

D. providing a transport for network configuration data between client and server

ANSWER: C

Explanation:

YANG is used for **describing data models**. Defined by the IETF, YANG is a data-modeling language that provides a machine-readable schema for the configuration data, operational state data, notifications, and remote procedure calls supported by network devices and services. A YANG module defines the hierarchical structure of data, including containers, lists, leaf values, data types, default values, and validation constraints. This shared schema enables automation tools and network devices to interpret data consistently and validate it before configuration changes are applied.

YANG is protocol-independent: it defines what the managed data looks like rather than how that data is transported. Management interfaces such as NETCONF and RESTCONF use YANG models to retrieve and modify structured device data. Cisco IOS XE supports YANG models, including Cisco-native and standards-based models, to enable model-driven programmability and automation. See IETF RFC 7950: The YANG 1.1 Data Modeling Language and Cisco IOS XE Programmability Guide: YANG.

QUESTION NO: 46

Refer to the exhibit.

```
from requests.auth import HTTPBasicAuth
import requests

ROUTER="192.0.2.1"
BASEURI="restconf/data/Cisco-IOS-XE-native:native"
SUBURI="router/router-bgp"

response = requests.get(
        f"https://{ROUTER}/{BASEURI}/{SUBURI}",
        auth=HTTPBasicAuth('admin', 'S3cr3tP4ss'),
        verify=False)

print(f"HTTP Response: {response.status_code} {response.reason}\n")
print(response.text)
```

```
admin@linux:/tmp$ python3 get-bgp.py
HTTP Response: 404 Not Found

<errors xmlns="urn:ietf:params:xml:ns:yang:ietf-restconf">
  <error>
        <error-message>uri keypath not found</error-message>
        <error-tag>invalid-value</error-tag>
        <error-type>application</error-type>
  </error>
</errors>

admin@linux:/tmp$
```

An engineer is creating a Pytnon script to fetch the BGP configuration from a device using RESTCONF. What does the output Indicate?

A. The BGP data resource identifier in the URL is incorrect.

B. There is no BGP process running on the device

C. RESTCONF is not enabled on the device. 23

ANSWER: A

Explanation:

The BGP data resource identifier in the URL is incorrect. The RESTCONF output shown in the exhibit indicates that the server received and processed the request but could not resolve the requested URI into a valid YANG-modeled data resource. RESTCONF URLs are schema-driven: the path after the RESTCONF data root must match the exact container, list, module namespace, and—when applicable—list keys exposed by the device's supported YANG models. A message such as `uri keypath not found` therefore identifies a resource-path problem rather than a general reachability issue.

To retrieve BGP configuration successfully, the script must use the RESTCONF data endpoint and the precise resource identifier for the BGP model implemented by that IOS XE device. The appropriate path can be verified by reviewing the device YANG capabilities, retrieving its YANG library information, or consulting the model documentation for the installed software release. RESTCONF defines the relationship between request URIs and YANG data resources, including the requirement that the target URI identify an existing resource. See RFC 8040, RESTCONF Protocol and Cisco's IOS XE RESTCONF overview.

QUESTION NO: 47 - (SIMULATION)

Simulation 07

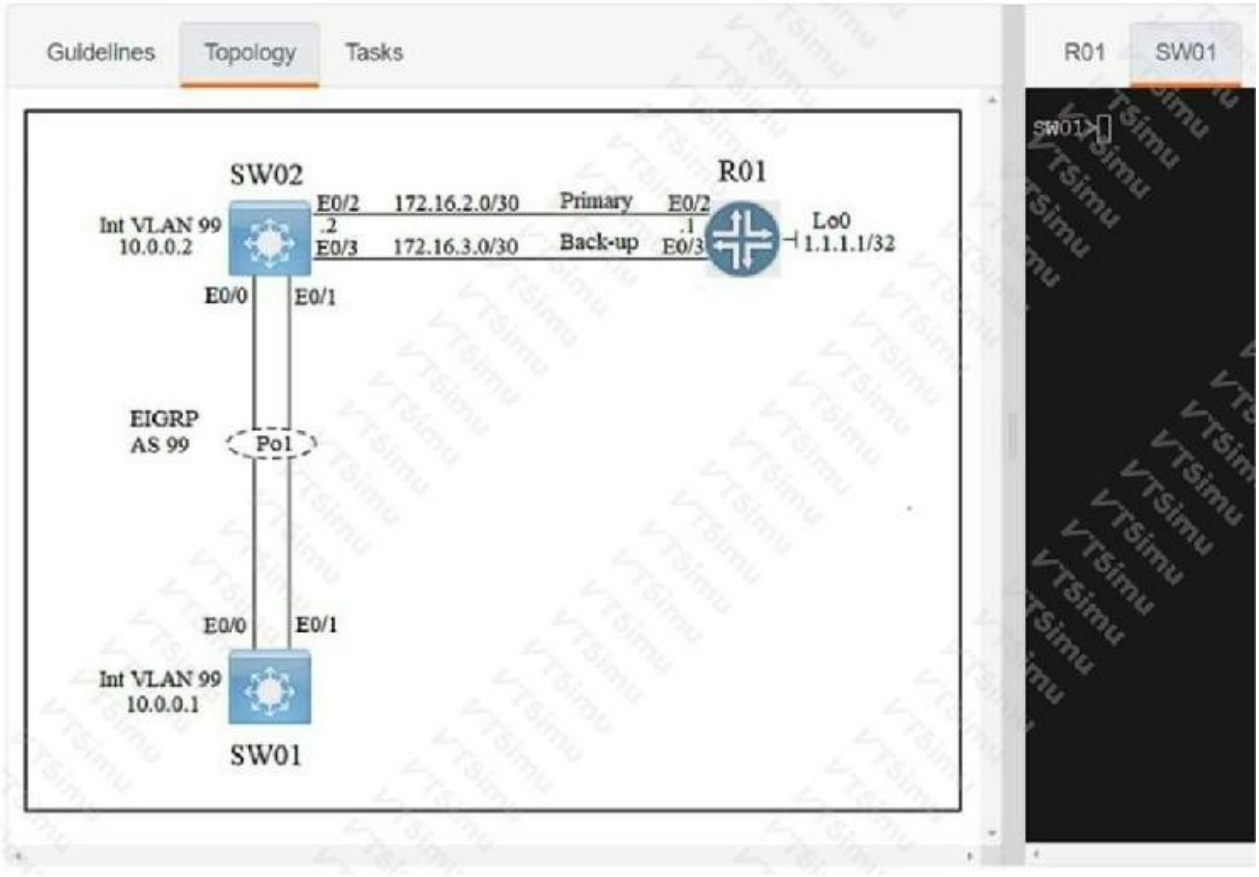

ANSWER: See the explanation for the answer

Explanation:

Configure SW01 to archive configuration-command changes (including the user associated with the command) and send the archive messages to syslog:

Configure R01 to enable the already configured NetFlow Top Talkers feature for traffic entering interface E0/2:

`logging enable` enables configuration change logging under the configuration archive. The log records the command and, when available from the login/AAA session, the user who entered it. `notify syslog` sends those archive log events to syslog. `ip flow ingress` enables NetFlow collection for all inbound packets on R01 E0/2.

QUESTION NO: 48

Refer to the Exhibit.

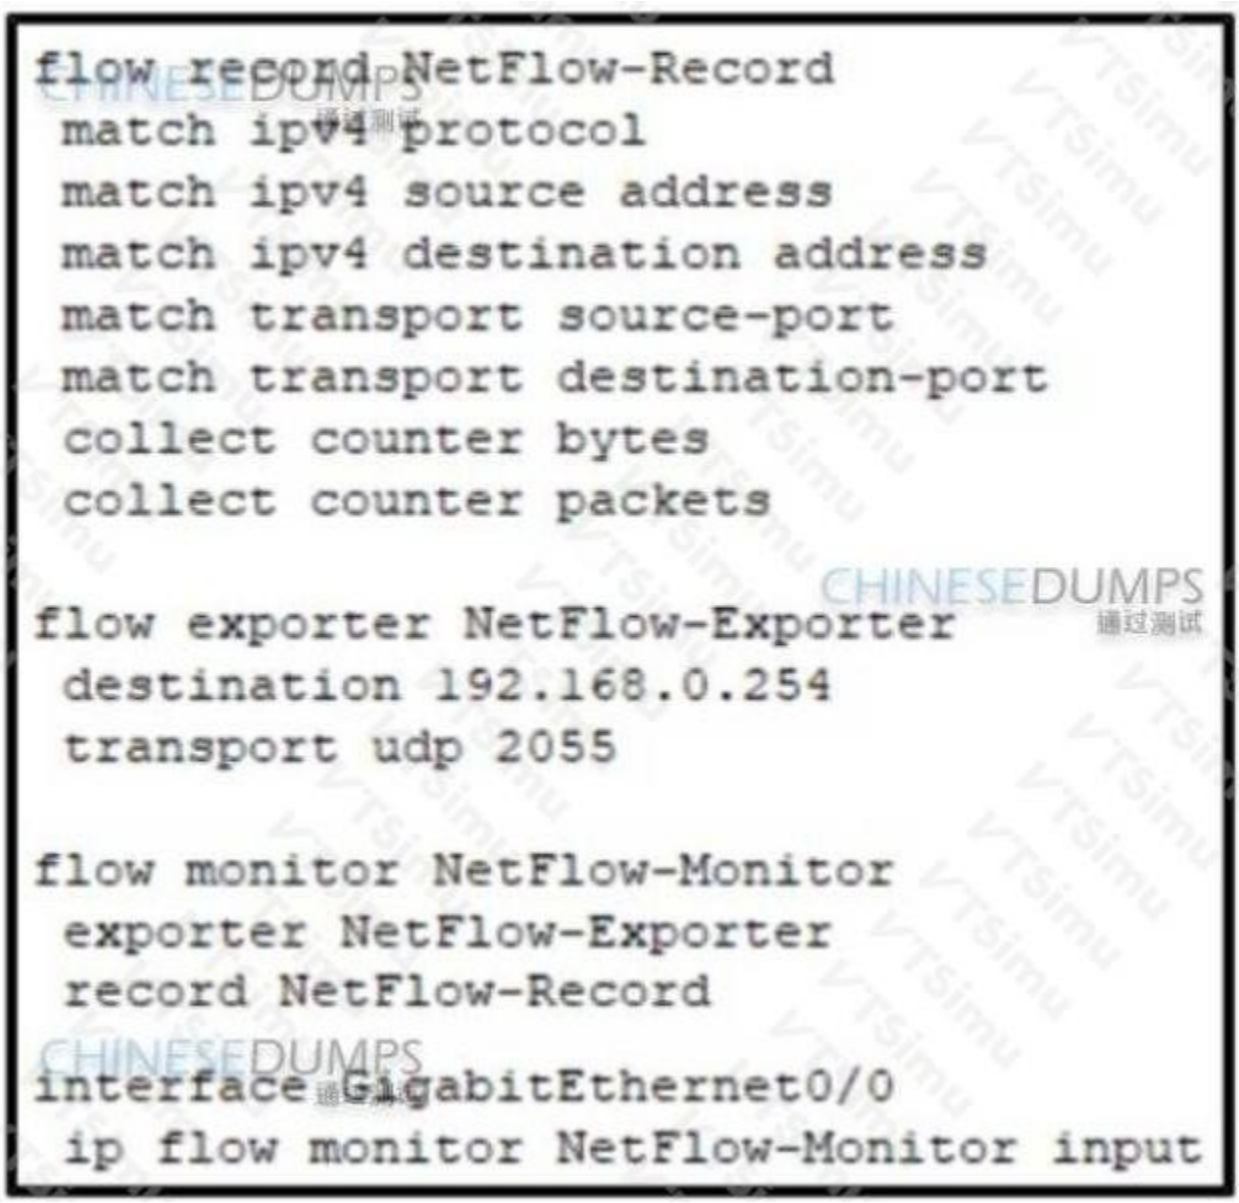

Refer to the exhibit. The DevOps team noticed missing NetFlow data during peak utilization times for 93 remote branches. Which configuration allows for this issue to be minimized or resolved?

A. Configure NetFlow on the in and outbound directions.

B. Change the transport type from UDP to TCP.

C. Configure long byte counters when specifying a flow record.

D. Change the flow monitor to IPv6 from IPv4.

ANSWER: C

Explanation:

Configure long byte counters when specifying a flow record. Flexible NetFlow can export byte and packet counters using either normal or long counter fields. Normal byte counters are 32-bit values, whereas long byte counters are 64-bit values. At periods of high utilization, especially when traffic from many remote branches is aggregated or flows remain active for

substantial periods, a 32-bit byte counter can wrap much sooner. That rollover makes the reported traffic volume appear incomplete or missing to the monitoring platform, even though the router is still observing and exporting flows.

Including the long counter field in the flow record, such as `collect counter bytes long`, greatly extends the measurable byte range and preserves accurate usage values during sustained high-volume traffic. This is particularly important for telemetry systems that use exported counters to calculate utilization, capacity trends, or application consumption. The collector can then receive a counter value that remains valid over the relevant flow lifetime instead of interpreting a wrapped value as missing or reduced data. Cisco documents the available Flexible NetFlow counter collection fields and their use in flow records in its Flexible NetFlow Configuration Guide and the Flexible NetFlow Command Reference.

QUESTION NO: 49

Refer to the exhibit.

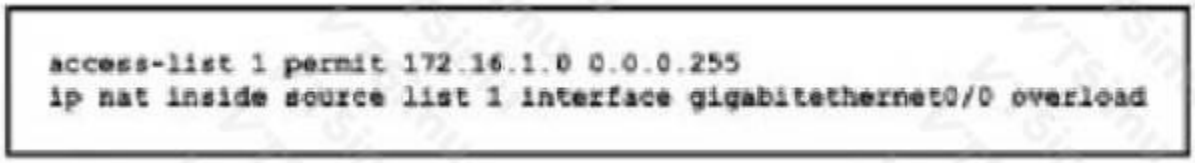

The inside and outside interfaces u configuration of this device have been correctly identified. What is the effect of this configuration?

A. dynamic NAT

B. NAT64

C. PAT

D. static NAT

ANSWER: C

Explanation:

PAT is correct. Port Address Translation, also called NAT overload, enables multiple inside hosts to share a single public IPv4 address for connections to external networks. On Cisco IOS, this behavior is configured when an inside-source NAT rule matches a group of private addresses, commonly through an access control list, and translates them using one global address or the outside interface address with the `overload` keyword. The router differentiates simultaneous connections by assigning and tracking unique Layer 4 source port values, along with protocol and address information, in its NAT translation table. As a result, many internal users can access the Internet at the same time even though their traffic appears externally to originate from one public IPv4 address. Return traffic is matched to the appropriate translation entry and delivered to the original inside host. This conserves scarce public IPv4 addresses while maintaining separate sessions for each host and application flow. Cisco identifies this port-multiplexing behavior as NAT overload/PAT. See Cisco's NAT overload configuration and operation overview and the Cisco IOS NAT Configuration Guide.

QUESTION NO: 50

Which two functions are performed by IKEv2 when establishing an IPsec VPN? (Choose two.)

A. authenticates VPN peers

B. negotiates IPsec security associations

C. forwards encrypted user packets between tunnel endpoints

D. performs Network Address Translation for protected traffic

E. assigns IP addresses to tunnel interfaces through DHCP

ANSWER: A B

Explanation:

IKEv2 authenticates the VPN peers before IPsec security associations are established. Authentication can use methods such as pre-shared keys or digital certificates, allowing each peer to verify the identity of the other peer before protected traffic is accepted.

IKEv2 also negotiates the IPsec security associations and cryptographic parameters used to protect data traffic. The peers agree on parameters such as encryption algorithms, integrity algorithms, lifetimes, and traffic selectors. IPsec Encapsulating Security Payload then uses the negotiated security associations to protect the actual user traffic. Cisco documents IKEv2 policy and IPsec negotiation in the Cisco IOS XE Internet Key Exchange Version 2 Configuration Guide.

QUESTION NO: 51

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

F. control-plane nodes

ANSWER: E F

Explanation:

In a two-tier Cisco SD-Access fabric, the collapsed core is formed by the border-node and control-plane-node functions. These functions are commonly colocated on the same highly available pair of devices, rather than deployed as separate dedicated tiers. The border function provides connectivity between the SD-Access fabric and external networks, such as the enterprise network, data center, WAN, Internet edge, or shared services. It is therefore the fabric's ingress and egress point for traffic that must leave or enter the fabric domain.

The control-plane function provides the LISP-based mapping service used by SD-Access. It maintains endpoint-to-location mappings and responds to mapping requests, allowing fabric edge devices to determine where endpoint traffic should be encapsulated and forwarded. Combining the border and control-plane functions into the collapsed core reduces the number of infrastructure nodes required while retaining these essential external-connectivity and overlay-control roles. Cisco recommends redundancy for these colocated functions to provide availability and resilience. Cisco documents the fabric border and control-plane roles in its SD-Access Design Guide and describes the SD-Access fabric architecture at Cisco Software-Defined Access.

QUESTION NO: 52

What is a TLOC in a Cisco Catalyst SD-WAN deployment?

A. component set by the administrator to differentiate similar nodes that offer a common service

B. value that identifies a specific tunnel within the Cisco Catalyst SD-WAN overlay

C. identifier that represents a specific service offered by nodes within the Cisco Catalyst SD-WAN overlay

D. attribute that acts as a next hop for network prefixes

ANSWER: B

Explanation:

The value that identifies a specific tunnel within the Cisco Catalyst SD-WAN overlay is correct because a TLOC, or Transport Locator, identifies the reachable transport-side tunnel endpoint advertised by a WAN Edge device. TLOC information enables other WAN Edge devices to determine where and how to establish overlay data-plane tunnels. In the Cisco Catalyst SD-WAN control plane, this information is distributed through Overlay Management Protocol route advertisements, allowing the fabric to build connectivity across the available underlay transports.

A TLOC is associated with a transport interface and is characterized by values including the device system IP address, transport color, and encapsulation type. The transport color conveys the underlay transport classification, such as MPLS, public internet, or LTE, while the encapsulation specifies the tunnel mechanism used for the transport. A WAN Edge with more than one active WAN transport can advertise multiple TLOCs. This gives the fabric multiple potential tunnel endpoints and supports path selection based on policy, transport availability, and performance. Cisco documents TLOCs as transport locations used by OMP to advertise and resolve overlay reachability: Cisco SD-WAN OMP Configuration Guide and Cisco SD-WAN Solution Overview.

QUESTION NO: 53

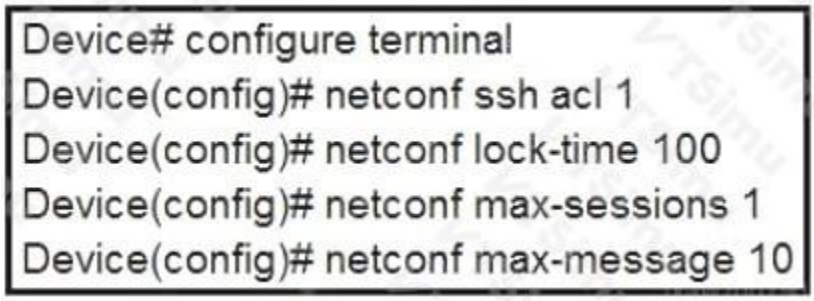

Refer to the exhibit. A network engineer must configure NETCONF. After creating the configuration, the engineer gets output from the command show line, but not from show running-config. Which command completes the configuration?

A. Device(config)# netconf max-sessions 100

B. Device(config)# no netconf ssh acl 1

C. Device(config)# netconf lock-time 500

D. Device(config)# netconf max-message 1000

ANSWER: D

Explanation:

Device(config)# netconf max-message 1000 completes the required NETCONF configuration by explicitly setting the largest NETCONF message that the device will process to 1000 bytes. NETCONF exchanges XML-encoded RPC requests and replies, and the device needs a defined message-size limit to control the resources consumed while receiving and parsing those protocol messages. This setting is a global NETCONF operational parameter, rather than a line-specific setting, so it applies to NETCONF protocol processing regardless of the VTY line used for SSH access. Cisco IOS XE supports configuring NETCONF message limits with the `netconf max-message` command as part of its NETCONF configuration. Setting the value to 1000 establishes the requested limit and finalizes the configuration shown in the exhibit. NETCONF itself is defined as an RPC-based network-management protocol that uses structured configuration and operational-data exchanges; message handling is therefore a core part of a usable deployment. See Cisco's IOS XE NETCONF documentation and RFC 6241 for NETCONF protocol and implementation details.

QUESTION NO: 54

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

A. Use a password hash

B. Use SSL for encryption

C. Enable dual authentication of the session

D. Use TACACS+ authentication

E. Enable out-of-band authentication 33

ANSWER: B C

Explanation:

Using SSL for encryption—implemented in current deployments as HTTPS with TLS—protects REST API traffic while it travels between a client and the API endpoint. TLS provides confidentiality for credentials, tokens, request payloads, and responses, while also providing integrity protection against undetected modification in transit. Proper certificate validation also lets an API client authenticate the server, helping prevent interception and man-in-the-middle attacks. REST APIs should therefore be exposed through HTTPS using supported TLS versions and securely managed certificates.

Enabling dual authentication of the session adds stronger identity assurance before sensitive API operations are allowed. This can involve validating both parties to the session through mutual authentication, or requiring an additional authentication factor for the API consumer, depending on the platform's authentication design. The security objective is to ensure that possession of a single compromised password or credential is not sufficient to establish an authorized API session. Strong session authentication is especially important for APIs because they often provide direct programmable access to network configuration and operational data.

These practices align with API-security guidance to protect communications with TLS and enforce robust authentication controls. See the OWASP API Security Project and Cisco's Cisco Developer Documentation.

QUESTION NO: 55

What are two characteristics of vManage APIs? (Choose two.)

A. Southbound API is based on OMP and DTLS.

B. Southbound API is based on NETCONF and XML.

C. Northbound API is based on RESTCONF and JSON.

D. Southbound API is based on RESTCONF and JSON.

E. Northbound API is RESTful, using JSON.

ANSWER: B E

Explanation:

Southbound API is based on NETCONF and XML. Cisco SD-WAN Manager (formerly vManage) uses NETCONF as its southbound management interface for communication with managed SD-WAN components and devices. NETCONF is an IETF network-configuration protocol whose operations and data encoding are XML-based. This interface supports centrally driven configuration deployment and operational management, allowing SD-WAN Manager to translate intent configured in the management platform into device-level configuration and management exchanges.

Northbound API is RESTful, using JSON. The SD-WAN Manager northbound interface is intended for external systems, such as orchestration platforms, automation tools, and custom applications. Cisco exposes REST API endpoints over HTTPS, with JSON commonly used for request and response bodies. This model enables programmatic access to management functions, including retrieving operational data, monitoring status, and automating configuration workflows without requiring direct interaction with the graphical interface.

These two interface directions separate external application integration from the platform's managed-device communication model. Cisco's SD-WAN documentation describes REST APIs for SD-WAN Manager automation, while NETCONF provides the structured southbound protocol used for controller and device management interactions. See Cisco SD-WAN Developer Center and Cisco SD-WAN API documentation.

QUESTION NO: 56

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:

FlexConnect and **local** are client-serving AP modes that can also contribute to wireless security monitoring. An AP operating in local mode normally provides centrally managed WLAN service, including client association and traffic forwarding, while periodically leaving its serving channel for short background scans. These off-channel scans provide RF information to the controller and help identify rogue APs, interfering devices, and other RF conditions without dedicating the radio exclusively to monitoring.

A FlexConnect AP also provides WLAN access to clients, particularly at remote sites where traffic can be centrally switched or locally switched according to the WLAN configuration. When operating with controller connectivity, FlexConnect APs can participate in controller-coordinated RF and rogue-detection functions while continuing to serve clients. This allows a deployment to retain client connectivity and visibility into nearby wireless activity instead of requiring every AP radio to be dedicated solely to monitoring.

Cisco describes local mode as the standard client-serving mode with background scanning capabilities, and its wireless security documentation explains the use of AP scanning and controller-based rogue detection in WLAN deployments. See Cisco AP Modes Configuration Guide and Cisco Wireless Intrusion Prevention and Rogue Detection Overview.

QUESTION NO: 57

Which lag/profile on a Cisco Catalyst 9800 Series WLC must be modified to allow Cisco ISE to dynamically assign VLANs to users on an 802.1X-based SSID?

A. interface lag

B. site tag

C. WLAN profile

D. policy profile

ANSWER: D

Explanation:

policy profile is correct because it contains the client policy settings that control how a Catalyst 9800 Wireless LAN Controller handles authorization results for a WLAN. For Cisco ISE to place authenticated 802.1X users into VLANs dynamically, ISE returns standard RADIUS tunnel attributes, commonly including Tunnel-Type, Tunnel-Medium-Type, and Tunnel-Private-Group-ID. The controller must be configured to honor these RADIUS-provided values rather than exclusively using a locally configured VLAN. On Catalyst 9800 controllers, this is enabled through the policy profile by configuring AAA override. Once AAA override is enabled and the policy profile is mapped to the applicable WLAN through the policy tag, the controller can apply the VLAN supplied in the user's ISE authorization result at connection time. The policy profile also defines the VLAN behavior used for centrally switched clients, making it the appropriate location for this configuration. Cisco documents AAA override and dynamic VLAN assignment as policy-profile-based configuration on Catalyst 9800 platforms.

See the Cisco configuration guide for AAA Override and Dynamic VLAN and the Cisco Catalyst 9800 Wireless Controller Configuration Guide.

QUESTION NO: 58

Refer to the exhibit. An engineer must adjust the configuration so that Router A becomes the active router. Which commands should be applied to router A? (Choose two)

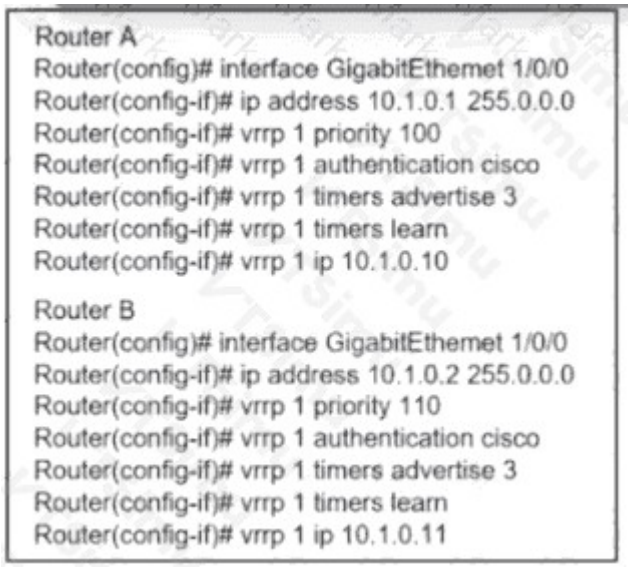

A. vrrp 1 priority 90

B. vrrp 1 timers advertise 1

C. vrrp 1 ip 10.1.0.11

D. ip address 10.1.0.11 255.0.0.0

E. vrrp 1 priority 120

ANSWER: C E

Explanation:

The commands `vrrp 1 ip 10.1.0.11` and `vrrp 1 priority 120` configure Router A as a member of VRRP group 1 using the required shared virtual IPv4 address, then give it the election value needed to become the active, or master, router. A VRRP group must use the same virtual IP address on every participating router. Hosts use that virtual address as their default gateway, while the active VRRP router answers ARP requests and forwards traffic sent to the virtual address. Therefore, Router A must be configured with `10.1.0.11` as the group virtual address rather than merely assigning that address as a regular interface address.

VRRP elects the router with the highest priority as the master. The default priority is 100, so configuring Router A with priority 120 makes it preferable to a peer using the default or a lower priority. Cisco IOS XE supports the VRRP group IP address and priority configuration under the relevant Layer 3 interface; when Router A participates in group 1 with the higher priority, it is selected to provide the active default-gateway function. Cisco documents VRRP configuration and election behavior in its VRRP Configuration Guide and the command syntax in the First Hop Redundancy Protocols Command Reference.

QUESTION NO: 59 - (SIMULATION)

SIMULATION 12:-

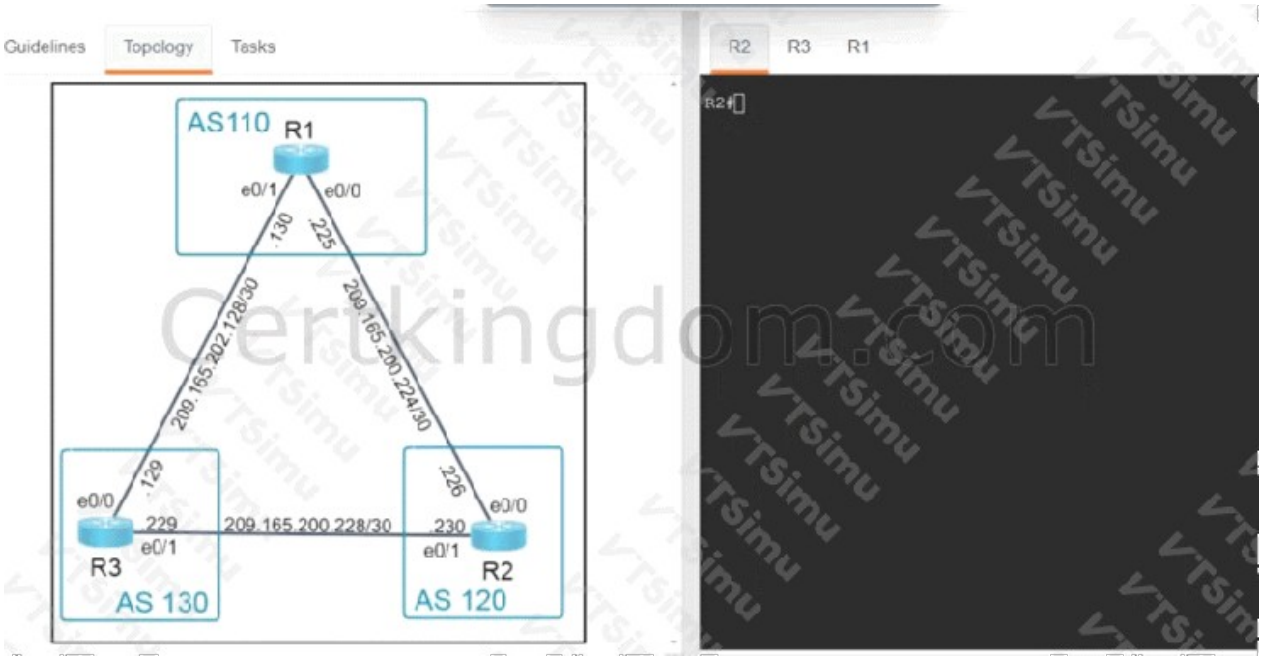

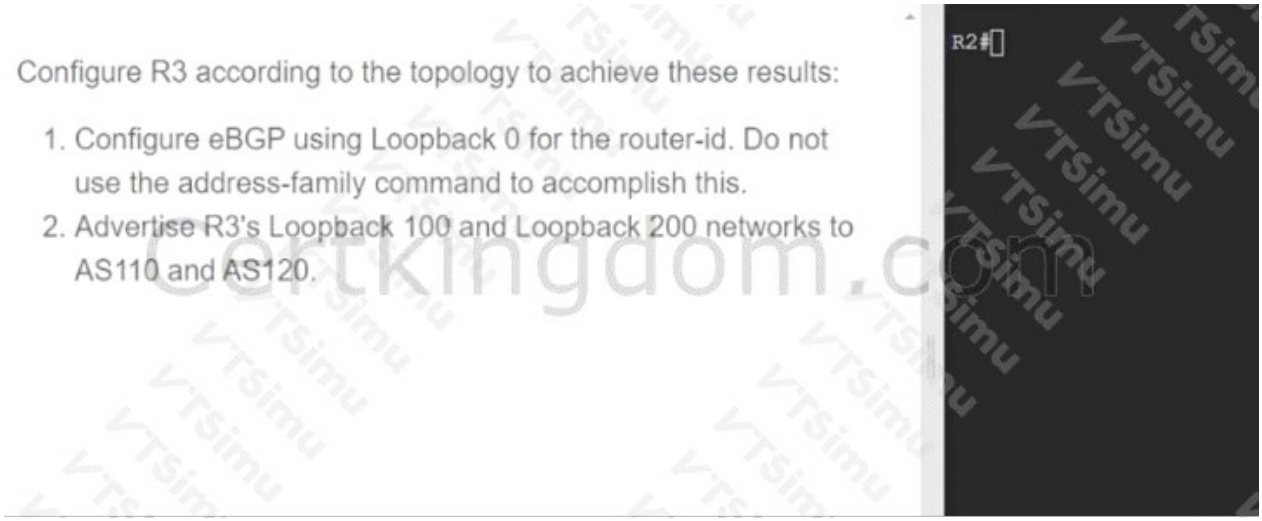

ANSWER: See the explanation for the answer

Explanation:

On **R3**, configure BGP in AS 130. The eBGP peer toward R1 (AS 110) is `209.165.202.130`, and the peer toward R2 (AS 120) is `209.165.200.230`.

Replace the bracketed values with the addresses already configured on R3. For example, if the loopbacks are configured as host routes, the final two commands use a `255.255.255.255` mask:

Use `show ip interface brief` or `show running-config | section interface Loopback` on R3 if needed to obtain the exact Loopback0, Loopback100, and Loopback200 addresses/masks. Do not enter `address-family`; the required IPv4 BGP network advertisements are configured directly under `router bgp 130`.

The `network` prefix and mask must exactly match an installed route for the respective loopback, or BGP will not originate it.

QUESTION NO: 60

Refer to the exhibit.

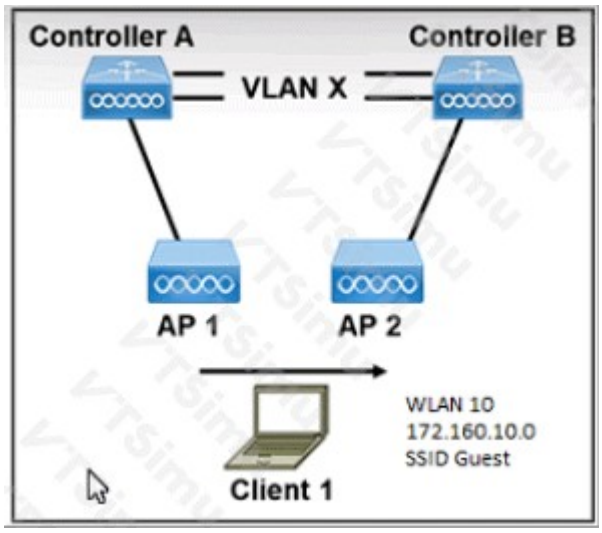

Both controllers are in the same mobility group. Which result occurs when client 1 roams between APs that are registered to different controllers in the same WLAN?

A. Client 1 contact controller B by using an EoIP tunnel.

B. CAPWAP tunnel is created between controller A and controller B.

C. Client 1 users an EoIP tunnel to contact controller A.

D. The client database entry moves from controller A to controller B.

ANSWER: D

Explanation:

The client database entry moves from controller A to controller B. A roam between access points joined to separate controllers is an intercontroller mobility event. Because the controllers belong to the same mobility group, they exchange mobility information and client context so that the destination controller can assume responsibility for the client after it associates to its new access point. For a Layer 2 roam in which the client remains on the same WLAN and subnet, the client keeps its existing IP addressing and the active client record is transferred to the controller serving the destination access point. This preserves the client's authentication, security, QoS, and session context as part of a seamless roam rather than requiring the endpoint to establish an entirely new wireless session. Cisco describes this behavior as intercontroller Layer 2 roaming: the client database entry is moved from the original controller to the new controller when the client roams within the same mobility group. The mobility-group relationship is therefore essential because it enables the controllers to recognize one another as mobility peers and coordinate the client-context handoff. See Cisco's Wireless LAN Controller Mobility overview and Mobility Groups configuration guide.

QUESTION NO: 61

Refer to the exhibit.

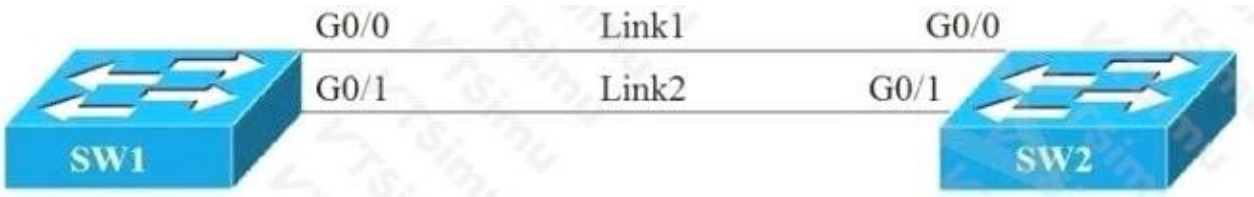


```
SW2#show spanning-tree

 VLAN0001

   Spanning tree enabled protocol ieee

   Root ID     Priority    32769
               Address     5000.0005.0000
               Cost        4
               Port        1 (GigabitEthernet0/0)
               Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

   Bridge ID   Priority    32769 (priority 32768 sys-id-ext 1)
               Address     5000.0006.0000
               Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
               Aging Time  300 sec

Interface           Role  Sts  Cost  Prio.Nbr   Type
------------------  ----  ---  ----  --------  --------------------

Gi0/0               Root  FWD   4    128.1      P2p
Gi0/1               Altn  BLK   4    32.2       P2p
```

Link1 is a copper connection and Link2 is a fiber connection. The fiber port must be the primary port for all forwarding. The output of the show spanning-tree command on SW2 shows that the fiber port is blocked by spanning tree. An engineer enters the spanning-tree port-priority 32 command on G0/1 on SW2, but the port remains blocked.

Which command should be entered on the ports that are connected to Link2 to resolve the issue?

A. Enter spanning-tree port-priority 4 on SW2.

B. Enter spanning-tree port-priority 32 on SW1.

C. Enter spanning-tree port-priority 224 on SW1.

D. Enter spanning-tree port-priority 64 on SW2.

ANSWER: B

Explanation:

`Enter spanning-tree port-priority 32 on SW1.` is correct because the STP port-priority value that matters to SW2's root-port decision is the sender port ID advertised by SW1 in its BPDUs. With two equal-cost parallel paths to the same upstream bridge, STP uses the upstream sender port ID as a tie-breaker. The sender port ID consists of the port priority and port number; a lower priority produces a lower, more-preferred port ID. Applying a priority of 32 to SW1's interface connected to the fiber Link2 causes SW1 to advertise a more-preferred port ID over that fiber segment. SW2 can then select its fiber-facing interface as its root port, allowing it to forward rather than remain in the alternate/blocking role. Configuring a port priority on SW2 does not change the port ID that SW2 receives from SW1, so it does not influence this particular upstream-port tie-break. Cisco documents the STP root-port selection sequence, including the sender bridge ID and sender port ID comparisons, in its Spanning Tree Protocol overview. The valid port-priority range and configuration behavior are also described in the Cisco IOS XE Spanning Tree Configuration Guide.

QUESTION NO: 62

Which security mechanism is offered on a next-generation firewall but not on a traditional firewall?

A. integrated IPS

B. stateful firewall policies

C. SSL VPN

D. NAT

ANSWER: A

Explanation:

integrated IPS is correct because a next-generation firewall combines traditional stateful firewalling with deeper, content-aware security inspection capabilities, including intrusion prevention. An integrated intrusion prevention system examines traffic for known attack signatures, protocol anomalies, and malicious activity, then can automatically block detected threats inline. This capability allows the firewall to enforce security policy based not only on addresses, ports, and connection state, but also on threats found within permitted traffic flows.

Cisco describes its next-generation firewall approach as integrating intrusion prevention with firewall controls and application visibility, enabling inspection and threat prevention across network traffic. In Cisco Secure Firewall deployments, IPS functions are provided through Snort-based inspection policies that detect and prevent exploits and other malicious network activity. This consolidated, inline inspection is a defining capability expected of an NGFW architecture and provides protection that goes beyond conventional packet filtering and stateful connection tracking. See Cisco's overview of next-generation firewalls and its NGFW technical overview for further details.

QUESTION NO: 63

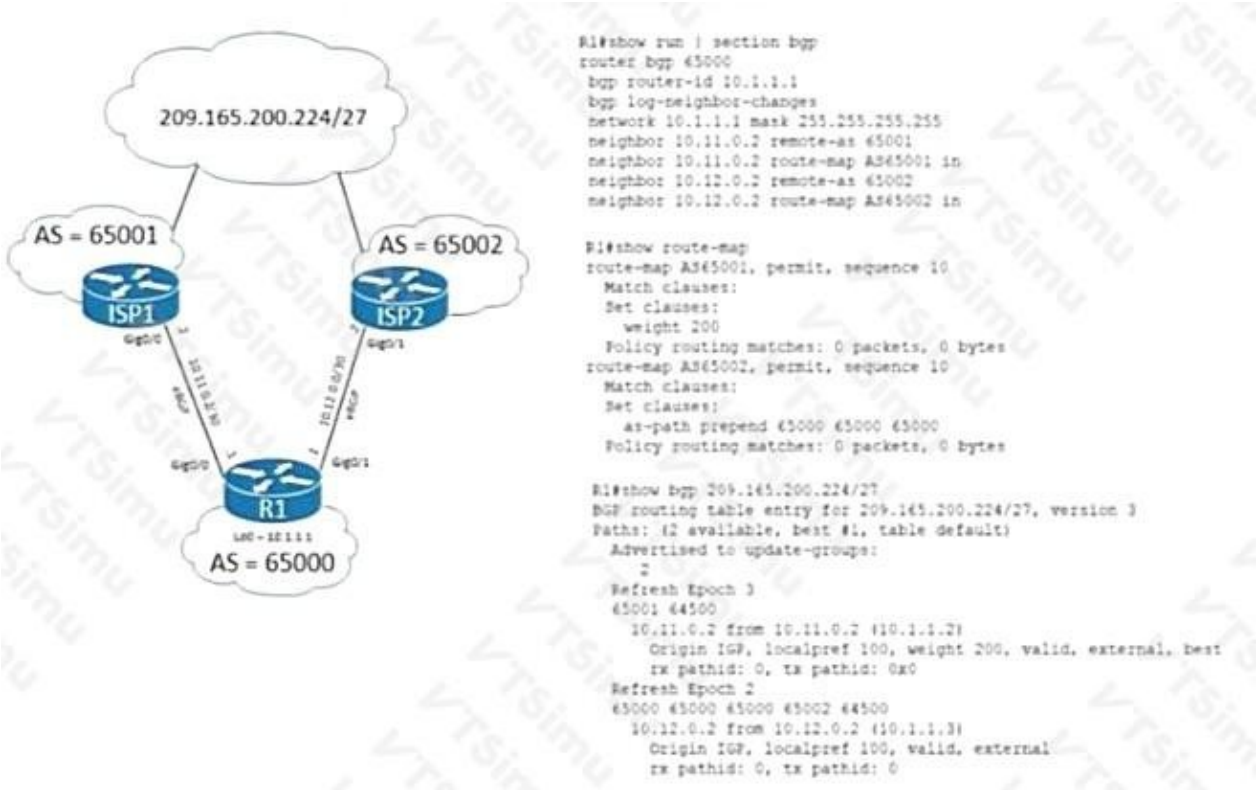

Refer to the exhibit. A client has two directly connected eBGP peering links with diverse ISPs. Both providers advertise the same public prefix 209.165.200.224/27 to R1 without any route manipulation.

Traffic leaves R1 outbound via ISP1 but returns inbound via ISP2. Which configuration prevents asymmetrical routing and makes ISP1 the preferred path inbound and outbound?

A.
```
R1# config t
R1(config)# router bgp 65000
R1(config-router)# neighbor 10.11.0.2 route-map AS65001 out
```

B.
```
R1# config t
R1(config)# route-map AS65002 permit 10
R1(config-route-map)# set weight 100
```

C.
```
R1# config t
R1(config)# router bgp 65000
R1(config-router)# neighbor 10.12.0.2 route-map AS65002 out
```

D.
```
R1# config t
R1(config)# route-map AS65001 permit 10
R1(config-route-map)# set local-preference 100
```

A. Option A

B. Option B

C. Option C

D. Option D

E. Set higher local preference for routes learned from ISP1 and prepend the AS-path when advertising 209.165.200.224/27 to ISP2

ANSWER: E

Explanation:

Set higher local preference for routes learned from ISP1 and prepend the AS-path when advertising 209.165.200.224/27 to ISP2 is the appropriate dual-direction BGP traffic-engineering policy. Local preference is an attribute evaluated within the customer's autonomous system, and BGP selects the path with the highest local preference before considering later attributes. Assigning a higher local preference to routes received from ISP1 therefore causes R1 to select ISP1 as its preferred egress path for destinations reachable through both providers.

For return traffic, R1 must influence route selection performed by networks outside its autonomous system. Prepending its own autonomous system number to advertisements sent through ISP2 makes the route to 209.165.200.224/27 via ISP2 appear to have a longer AS path. Because AS-path length is a BGP best-path consideration, remote networks that receive otherwise comparable routes through both providers generally prefer the shorter advertisement through ISP1. This combines local outbound control with an Internet-facing inbound-routing signal. AS-path prepending is an influence mechanism rather than an absolute guarantee, since remote routing policies can differ, but it is the standard approach when provider-specific BGP community controls are not being used. Cisco documents local preference and AS-path length in its BGP best-path process and describes AS-path prepending for multihomed inbound traffic engineering: Cisco BGP Best Path Selection and Cisco BGP Traffic Engineering.

QUESTION NO: 64 - (HOTSPOT)

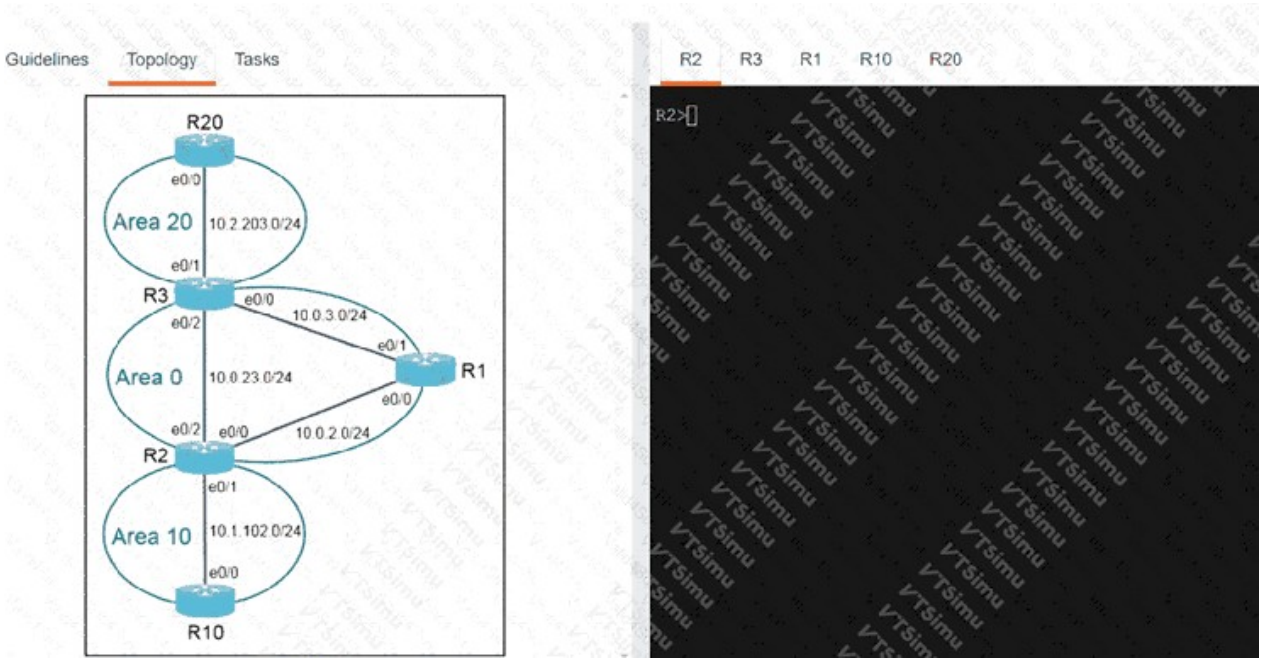

ANSWER:

Explanation:

The correct configuration is to enter `ip ospf priority 255` under R2's Ethernet0/1 interface. The topology shows that R2 Ethernet0/1 connects to R10 and belongs to Area 10, so this is the OSPF interface where the designated-router election must be influenced. OSPF performs a DR and BDR election only on multiaccess network types, such as the Ethernet broadcast segment illustrated between R2 and R10. The election uses the OSPF interface priority before using the router ID as a tie-breaker. Cisco IOS permits an OSPF interface priority from 0 through 255, and the default is 1. Setting the priority to 255 gives R2 the highest possible election preference while leaving its router ID unchanged, exactly as required. The router ID is not being modified; only the election behavior of R2 on the Area 10-facing interface is being configured. The displayed verification output from R10 confirms the intended result: its neighbor R2 has priority 255 and is in the `FULL/DR` state, demonstrating that R2 won the DR role on that segment. Cisco documents the OSPF interface priority behavior and its role in DR election in the Cisco IOS OSPF configuration guide. Because OSPF elections are nonpreemptive, if an election had already occurred, the adjacency or OSPF process might need to be reset for the changed priority to take effect immediately; however, the configuration command itself is the correct required action.

QUESTION NO: 65

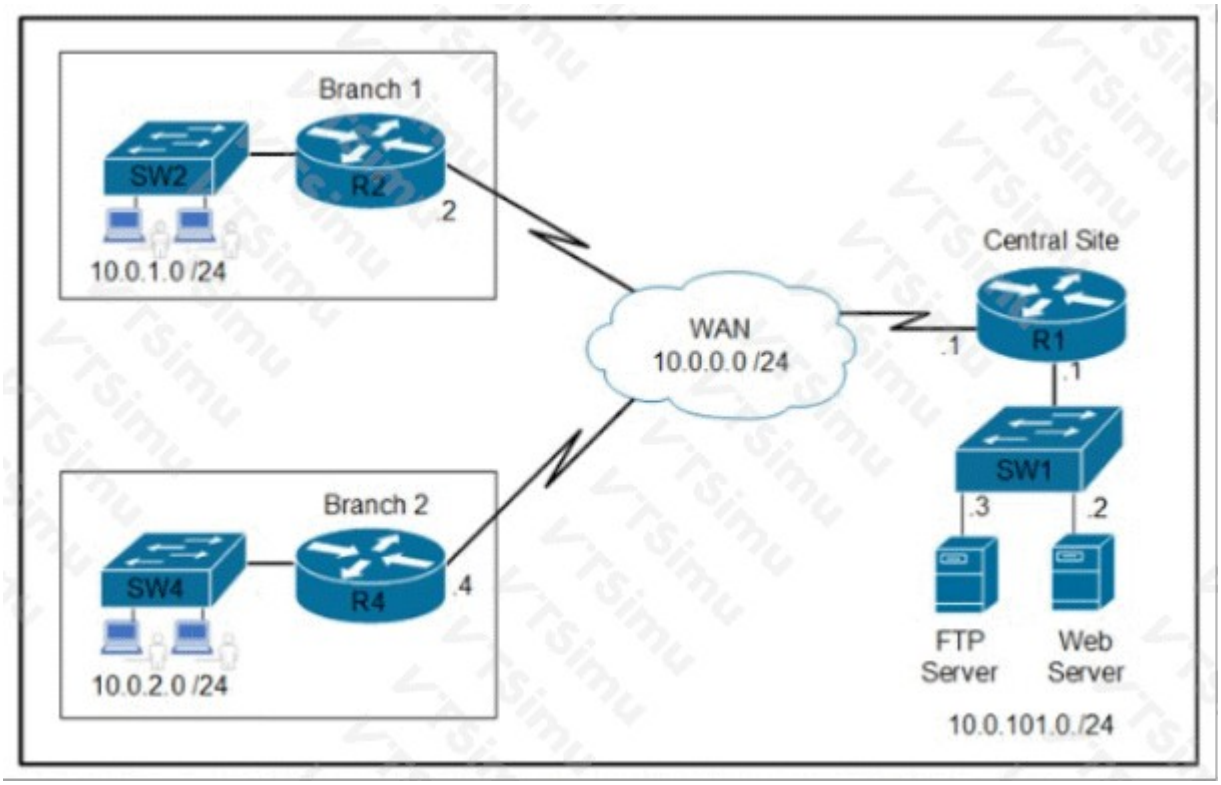

Refer to the exhibit Which two commands are required on route» R1 to block FTP and allow all other traffic from the Branch 2 network' (Choose two)

```
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:

Filtering FTP from Branch 2 while preserving all other connectivity requires an extended IPv4 ACL, because the policy must match both the source network and the TCP FTP service port. The required configuration components are the ACL entry that denies TCP traffic sourced from the Branch 2 subnet when its destination port is FTP, an explicit IP permit for the same source so that non-FTP traffic is allowed, and application of that ACL to the R1 interface in the direction in which Branch 2 traffic enters the router. Cisco IOS evaluates ACL entries sequentially from top to bottom and stops at the first match. Consequently, the FTP deny must precede the broad permit. The explicit permit is essential because every ACL ends with

an implicit deny; without an appropriate permit, other Branch 2 IP traffic would also be dropped. Applying the ACL inbound on the interface facing Branch 2 filters traffic as close to its source as possible and prevents unwanted FTP sessions from being routed farther into the network. Cisco documents extended ACL matching, processing order, and interface attachment in its Configuring Commonly Used IP ACLs and Cisco IOS XE ACL Configuration Guide.

QUESTION NO: 66

Which method requires a client to authenticate and has the capability to function without encryption?

A. open

B. WEP

C. WebAuth

D. PSK

ANSWER: C

Explanation:

WebAuth requires the wireless client to authenticate through a web-based captive portal and can be used on a WLAN that has no Layer 2 Wi-Fi encryption. A client first associates with the open wireless network, receives network access sufficient for DHCP, DNS, and portal access, and is redirected to a login page when attempting to browse. The user then supplies credentials, accepts an acceptable-use policy, or completes a guest-access workflow. After successful authentication, the controller applies the configured authorization policy and permits broader network access. This capability makes WebAuth a common design for guest WLANs, where an organization needs an authentication or terms-acceptance step without requiring every visitor to preconfigure a WPA personal key or an enterprise supplicant profile. Although the WLAN can operate without 802.11 encryption, deployments should use HTTPS for the portal and should carefully limit preauthentication access because an open WLAN does not protect wireless traffic at Layer 2. Cisco classifies web authentication as a Layer 3 security feature and documents its use with open WLAN configurations and web redirection. See Cisco's Web Authentication Overview and the Cisco Wireless Controller Security Configuration Guide.

QUESTION NO: 67

A customer wants to connect a device to an autonomous Cisco AP configured as a WGB. The WGB is configured properly: however, it fails to associate to a CAPWAP- enabled AP. Which change must be applied in the advanced WLAN settings to resolve this issue?

A. Enable Aironet IE.

B. Enable passive client.

C. Disable AAA override.

D. Disable FlexConnect local switching.

ANSWER: A

Explanation:

Enable Aironet IE. is required for an autonomous Cisco access point operating as a workgroup bridge to associate with a controller-based CAPWAP access point. A workgroup bridge uses Cisco proprietary Aironet information elements to identify and negotiate the capabilities needed for WGB operation on the WLAN. The controller must include these IEs in the WLAN beacon and probe-response information so that the autonomous AP can recognize the WLAN as supporting the Cisco extensions required for a workgroup bridge association.

On Cisco wireless LAN controllers, Aironet IE advertisement is configured in the WLAN's Advanced settings. Enabling it allows the CAPWAP AP to advertise the appropriate Cisco-specific information to the WGB, permitting the association and the intended bridging of traffic for the wired device connected behind it. This requirement applies even when the autonomous

AP's radio, SSID, security, and WGB configuration are otherwise correct, because the needed interoperability capability is advertised by the infrastructure WLAN rather than configured solely on the bridge.

Cisco documents WGB interoperability requirements and configuration considerations in its Wireless Workgroup Bridge documentation and Workgroup Bridge FAQ.

QUESTION NO: 68

A network engineer is configuring OSPF on a router. The engineer wants to prevent having a route to 172.16.0.0/16 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 additional configuration commands must be applied to accomplish the goal? (Choose two.)

A. ip prefix-list OFFICE seq 10 permit 0.0.0.0/0 le 32

B. distribute-list prefix OFFICE in under the OSPF process

C. distribute-list OFFICE in under the OSPF process

D. distribute-list OFFICE out under the OSPF process

E. ip prefix-list OFFICE seq 10 permit 0.0.0.0/0 ge 32

ANSWER: A B

Explanation:

The required configuration is `ip prefix-list OFFICE seq 10 permit 0.0.0.0/0 le 32` together with an inbound OSPF distribute list that references the prefix list: `distribute-list prefix OFFICE in` under the OSPF routing process. The existing sequence 5 entry exactly denies the target prefix, 172.16.0.0/16. Prefix lists are processed in sequence order and have an implicit deny when no entry matches. Therefore, the later permit entry is necessary to allow every other IPv4 prefix, from a default route through a /32 host route, to continue through the filter. This preserves installation of other valid OSPF-learned routes while excluding the specified route from the local routing table.

Applying the prefix list in the inbound direction is the key OSPF behavior: the router filters routes that OSPF has learned before those routes are installed in its local routing information base. Cisco documents that OSPF supports `distribute-list prefix` filtering, and that inbound filtering affects route installation locally rather than the OSPF link-state database or route advertisements to neighbors. See Cisco's OSPF distribute-list guidance and the Cisco IOS IP routing command reference.

QUESTION NO: 69

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, or bare-metal, hypervisor. A Type 1 hypervisor installs directly on a server's physical hardware rather than operating as an application within a conventional, general-purpose host operating system. It supplies the virtualization layer that allocates processor capacity, memory, storage access, and network I/O to guest virtual machines. This architecture is designed for enterprise server virtualization, where centralized management, workload isolation, and efficient use of physical infrastructure are required.

XenServer is built on the Xen Project hypervisor, which is a well-established bare-metal virtualization platform. Its control domain and management components operate as part of the Xen-based virtualization architecture, while guest systems run as virtual machines on the underlying hardware. Therefore, it fits the defining Type 1 characteristic: the hypervisor is the primary software layer responsible for managing hardware resources and hosting guest workloads. The Xen Project describes Xen as an open-source hypervisor that enables multiple operating systems to run on the same hardware, and Citrix's hypervisor documentation identifies its Xen-based virtualization platform for server workloads. See the Xen Project users overview and Citrix Hypervisor documentation.

QUESTION NO: 70

Which framework is used for third-party authorization?

A. API keys

B. custom tokens

C. OAuth

D. SOAP

ANSWER: C

Explanation:

OAuth is the framework used for third-party authorization, more specifically delegated authorization. It enables a resource owner to allow an application to access protected resources held by another service without disclosing the resource owner's password or other credentials to that application. Instead of receiving the user's credentials, the application obtains an access token from an authorization server after the user grants consent for defined permissions, known as scopes. The application presents that token to the resource server, such as an API, which evaluates the token and permits only the authorized actions or data access.

This token-based delegation model is common when a user permits an external application to access a cloud service, API, or enterprise resource on the user's behalf. OAuth 2.0 defines the roles, authorization flows, token concepts, and interactions that support this model. It is an authorization framework rather than an authentication protocol, although it is frequently combined with OpenID Connect when an application also needs to identify and authenticate the user. See RFC 6749: The OAuth 2.0 Authorization Framework and OAuth 2.0 overview.

QUESTION NO: 71

What is two characteristic of Cisco DNA Center and vManage northbound APIs?

A. They exchange XML-formatted content

B. They exchange JSON-kxmatted content.

C. They implement the RESTCONF protocol.

D. They implement the NETCONF protocol.

E. They exchange protobuf-formatted content.

F. They implement REST APIs.

ANSWER: B F

Explanation:

Cisco DNA Center and Cisco SD-WAN vManage provide northbound interfaces for applications, orchestration systems, and automation workflows through HTTP-based REST APIs. These interfaces use resource-oriented URIs and standard HTTP methods such as GET, POST, PUT, and DELETE, allowing external software to retrieve operational information and submit

configuration or management requests. This REST API model is the common northbound integration approach documented for both platforms.

They also exchange JSON-formatted content. JSON is used to represent request bodies and API responses, making the interfaces easy to consume from common automation languages and tools. For example, clients can authenticate, invoke an API endpoint over HTTPS, and parse the returned JSON data to incorporate inventory, topology, policy, device, or operational information into an external workflow. Cisco's developer documentation for both products presents API calls and payload examples in this format. See the Cisco DNA Center developer documentation and the Cisco SD-WAN vManage developer documentation.

QUESTION NO: 72

```
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:

"Ensure that the switchport parameters of Port-channel1 match the parameters of the port channel on the neighbor switch." is required because an EtherChannel is a single logical Layer 2 interface between the two devices. Its operational switching characteristics must be compatible at both ends. In particular, trunking operation and its associated settings, such as the allowed VLAN range and native VLAN, must be aligned on the paired port-channel interfaces. Consistent port-channel configuration prevents the switches from detecting an unsafe or inconsistent bundled-link condition.

"Ensure that the neighbor interfaces of Gi1/0/2 and Gi1/0/3 are configured as members of the same EtherChannel." is also required. Every physical link intended to participate in the bundle must terminate on interfaces that are members of the corresponding bundle at the peer. The member ports must use compatible EtherChannel negotiation and operational settings so that they can be selected into the same logical channel. Once the peer-side interfaces are correctly grouped and the port-channel switching parameters are consistent, the bundle can form reliably rather than repeatedly being protected by

an error-disabled condition. Cisco's EtherChannel troubleshooting guidance and configuration documentation describe these consistency requirements: Cisco EtherChannel Troubleshooting and Cisco EtherChannel Configuration Guide.

QUESTION NO: 73

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:

Implementing MBGP and implementing static mroutes are multicast-specific mechanisms that can control the Reverse Path Forwarding decision independently of the normal unicast routing table. MBGP carries multicast reachability in the IPv4 multicast BGP address family, allowing the router to maintain a multicast routing information base that is distinct from the unicast information. When multicast BGP routes are available, the router can use this multicast topology for source-path RPF lookups, preserving the existing unicast best-path decision while selecting an RPF path appropriate for multicast traffic.

Static mroutes explicitly install multicast-routing information for a source or source prefix, including the upstream RPF next hop or incoming interface. This supplies a deterministic RPF path for multicast forwarding without requiring an alteration to unicast routes. Static mroutes are particularly useful when the desired multicast upstream path differs from the route selected by unicast routing, or when a multicast source is reachable through a topology that must be treated differently for multicast distribution. Cisco documents both multicast BGP and static multicast routes as sources of multicast routing information used in RPF processing. See Cisco's multicast RPF documentation and the Cisco MBGP configuration guide.

QUESTION NO: 74

Which two advanced security features are available in next-generation firewalls but were not provided by standard firewalls? (Choose two.)

A. Stateful traffic inspection

B. Intrusion prevention

C. Application control

D. Remote access VPN

E. Network telemetry

ANSWER: B C

Explanation:

Intrusion prevention and **Application control** are defining next-generation firewall capabilities because they extend policy enforcement beyond the connection-oriented, stateful inspection performed by traditional firewalls. Intrusion prevention integrates deep packet inspection with threat detection and prevention. It evaluates traffic against threat signatures, vulnerability information, and protocol analysis to identify exploit attempts and malicious activity, then can block that traffic inline before it reaches protected hosts.

Application control, often called application visibility and control, identifies the actual application in use and enables policy based on that application rather than relying only on IP addresses, ports, and protocols. This is important because many applications can use common ports or attempt to evade port-based controls. An administrator can therefore create granular security policy that permits approved business applications while controlling or blocking undesired application behavior. Cisco characterizes NGFW solutions as combining traditional firewall functions with application awareness/control and integrated intrusion prevention. See Cisco's overview of next-generation firewalls and its network security overview.

QUESTION NO: 75

Refer to the exhibit.

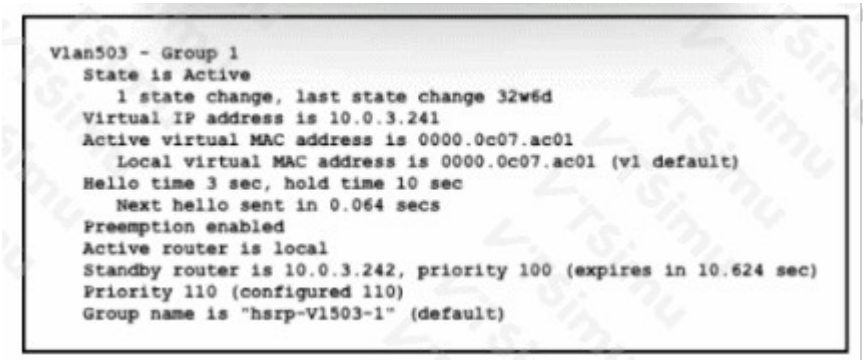

Which two facts does the device output confirm? (Choose two.)

A. The device sends unicast messages to its peers

B. The device's HSRP group uses the virtual IP address 10.0.3.242

C. The standby device is configured with the default HSRP priority.

D. The device is using the default HSRP hello timer

E. The device is configured with the default HSRP priority

ANSWER: C D

Explanation:

The device is using the default HSRP hello timer because the displayed HSRP timer values show a hello interval of 3 seconds. HSRP uses a 3-second hello timer by default; the associated default hold timer is 10 seconds. These values control how frequently HSRP peers send hello packets and how long a router waits before declaring that an expected peer is unavailable. The displayed timer information therefore confirms that no nondefault hello interval is in operation for this HSRP group.

The standby device is configured with the default HSRP priority because its reported priority is 100. Cisco HSRP uses 100 as the default priority value. Priority is the election value used to select the active router, with the higher value preferred when routers are otherwise eligible. The local router's explicitly displayed configured priority can differ from the standby router's value, so the output can identify the standby peer as retaining the standard priority while the local device is using a different configured priority. Cisco documents HSRP timer defaults and priority behavior in its HSRP overview and the Cisco IOS XE FHRP configuration guide.

QUESTION NO: 76

Which QoS feature uses the IP Precedence bits in the ToS field of the IP packet header to partition traffic into different priority levels?

A. marking

B. shaping

C. policing

D. classification

ANSWER: D

Explanation:

classification is correct because it is the QoS process that identifies traffic and assigns packets to traffic classes based on defined criteria in packet headers. For IPv4, one such criterion is IP Precedence, which occupies the three most significant bits of the legacy Type of Service (ToS) byte. These three bits provide eight possible precedence values, allowing a network device to distinguish traffic categories and place packets into priority groupings before applying the appropriate QoS treatment.

On Cisco platforms, classification is commonly configured with Modular QoS CLI class maps. A class map can match the IP Precedence value carried by packets, such as with the `match ip precedence` command. After the packets are classified, a policy map can apply the intended forwarding behavior to that class, including a suitable queueing, bandwidth, shaping, or policing policy. Thus, reading IP Precedence values to separate traffic into priority levels is specifically an identification and grouping function, which is the purpose of classification. Cisco also describes IP Precedence as the upper three bits associated with the DSCP/ToS field and supports matching it in QoS class maps. See Cisco's DSCP and IP Precedence reference and Cisco IOS XE MQC configuration guide.

QUESTION NO: 77

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

A. Cisco DNA Center is unable to collect monitoring data in Assurance.

B. All devices reload after detecting loss of connection to Cisco DNA Center.

C. Already connected users are unaffected, but new users cannot connect

D. Users lose connectivity.

E. User connectivity is unaffected.

ANSWER: A E

Explanation:

Cisco DNA Center is the management, automation, and assurance platform for an SD-Access deployment; it is not in the forwarding path for endpoint traffic. Fabric edge nodes, control-plane nodes, and border nodes retain the control-plane and data-plane information needed for the operational fabric to continue carrying traffic. Therefore, an outage affecting connectivity between Cisco DNA Center and fabric devices does not by itself disrupt established fabric forwarding, and user connectivity is unaffected as long as the fabric infrastructure and its internal control-plane connectivity remain healthy.

Assurance relies on device communication and telemetry to gather and update operational information, including health metrics, events, and related monitoring data. When Cisco DNA Center cannot reach the devices, it cannot collect current monitoring data from them, so Assurance data becomes unavailable or stale until management connectivity is restored. The outage also prevents normal centralized management operations, such as deploying new configuration changes, but the existing operational state on the devices remains in effect.

Cisco describes SD-Access as a fabric architecture with distributed network functions, while Cisco DNA Center provides centralized automation and assurance capabilities. See the Cisco SD-Access overview and Cisco DNA Center product information.

QUESTION NO: 78

Which action limits the total amount of memory and CPU that is used by a collection of VMs?

A. Place the collection of VMs in a resource pool.

B. Place the collection of VMs in a vApp.

C. Limit the amount of memory and CPU that is available to the cluster.

D. Limit the amount of memory and CPU that is available to the individual VMs.

ANSWER: A

Explanation:

Place the collection of VMs in a resource pool. A vSphere resource pool is a logical container that manages CPU and memory resources for its member virtual machines collectively. Administrators can configure CPU and memory reservations, shares, and limits on the resource pool itself. In particular, a pool-level limit establishes the maximum aggregate CPU or memory capacity that all VMs in that pool can consume, even if the individual workloads request more resources. The resource pool then allocates its available resources among its child VMs according to demand and configured shares, while ensuring the group does not exceed the pool's configured ceiling. This is useful when a cluster serves separate departments, applications, or environments and each group must be constrained to a defined portion of cluster capacity without requiring fixed allocations for every VM. Resource pools can also be nested, enabling hierarchical allocation policies while retaining overall control at each pool boundary. VMware documents resource pools as the mechanism for partitioning and controlling CPU and memory resources among groups of virtual machines; pool reservations, limits, and shares apply to the resources available to the pool and its descendants. See the vSphere Resource Management documentation and the VMware resource pool guidance.

QUESTION NO: 79

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 are designed to concentrate transmitted and received RF energy in one intended direction rather than distributing it uniformly around the antenna. This concentration produces **high gain** in the direction of the main lobe. Antenna gain does not create additional transmitter power; instead, it focuses the available RF energy into a smaller spatial area, improving the effective signal level and link budget along the intended path. This makes directional antennas appropriate for targeted wireless coverage and point-to-point links, where the antenna can be aimed accurately.

A directional antenna also **provides the most focused and narrow beam width**. Its radiation pattern has a comparatively narrow main lobe, allowing RF coverage to be directed toward a specific client area, corridor, outdoor location, or remote bridge endpoint. A narrow beamwidth also limits energy radiated outside the desired coverage zone and can reduce reception of signals arriving from off-axis directions. Cisco's antenna documentation describes directional antenna types, such as patch and Yagi antennas, as providing focused coverage patterns with higher gain than omnidirectional designs. See Cisco's Aironet antenna product information and the Cisco wireless antenna reference.

QUESTION NO: 80

What are two benefits of virtual switching when compared to hardware switching? (Choose two.)

A. increased MTU size

B. VM-level isolation

C. extended 802.1Q VLAN range

D. hardware independence

E. increased flexibility

ANSWER: B E

Explanation:

VM-level isolation is a principal benefit because a virtual switch connects and controls traffic at each virtual machine's virtual NIC. This allows administrators to apply port-group policies, VLAN assignments, traffic controls, and security settings directly to individual workloads. As a result, separate VMs can share physical uplinks while retaining distinct Layer 2 connectivity and policy boundaries. This workload-level control is especially useful in multitenant and highly consolidated server environments.

Increased flexibility is also a core advantage of virtual switching. A virtual switch is implemented in software and can be provisioned, reconfigured, and managed as part of the virtualized compute environment. Network connectivity and policies can therefore be applied rapidly as VMs are created, moved, cloned, or removed, without recabling server interfaces or manually changing a physical switch port for every workload change. Centralized management and automation can further make these changes repeatable at scale. Cisco describes virtualization as abstracting resources to make infrastructure more agile, while VMware explains that virtual switches provide network connectivity between virtual machines and physical networks. See Cisco: What Is Virtualization? and VMware: Virtual Switch.

QUESTION NO: 81

What is a client who is running 802.1x for authentication reffered to as?

A. supplicant

B. NAC device

C. authenticator

D. policy enforcement point

ANSWER: A

Explanation:

A client running IEEE 802.1X authentication is called a **supplicant**. The supplicant is the endpoint seeking access to the LAN or WLAN, such as a workstation, IP phone, or other network-attached device. It runs the 802.1X client software and uses Extensible Authentication Protocol over LAN (EAPOL) to exchange authentication information with the network-access device. Depending on the selected EAP method, the supplicant provides an identity and credentials, such as a username and password or a device/user certificate. After the authentication exchange succeeds, the network grants access according to the applicable authorization policy. Cisco describes the three principal 802.1X roles as the supplicant, authenticator, and authentication server; the supplicant is specifically the client requesting network access. This terminology applies to both wired switch-port access control and enterprise wireless access control. See Cisco's 802.1X overview and terminology and the IEEE 802.1 working group's 802.1X information.

QUESTION NO: 82

Which JSON syntax is valid?

A. {"switch": "name": "dist1", "interfaces": ["gig1", "gig2", "gig3"]}

B. {/"switch/": {/"name/": "dist1", /"interfaces/": ["gig1", "gig2", "gig3"]}}

C. {"switch": {"name": "dist1", "interfaces": ["gig1", "gig2", "gig3"]}}

D. {'switch': ('name': 'dist1', 'interfaces': ['gig1', 'gig2', 'gig3'])}

ANSWER: C

Explanation:

`{"switch": {"name": "dist1", "interfaces": ["gig1", "gig2", "gig3"]}}` is valid JSON. JSON represents data using objects, arrays, strings, numbers, Boolean values, and null. An object is delimited by curly braces and consists of zero or more name/value members. In this example, the outer object has one member named `switch`. Its value is a nested object that contains a `name` member with the string value `dist1` and an `interfaces` member whose value is an array of interface-name strings. Each property name and each string value is enclosed in double quotation marks, each property name is separated from its value by a colon, and the two members in the nested object are separated by a comma. The array is correctly enclosed in square brackets and its string elements are comma-separated. This nested object-and-array structure is commonly used by network automation tools and REST APIs to model device attributes and interface lists. The formal JSON grammar is defined in RFC 8259; see also the MDN JSON guide.

QUESTION NO: 83

Refer to the Exhibit.

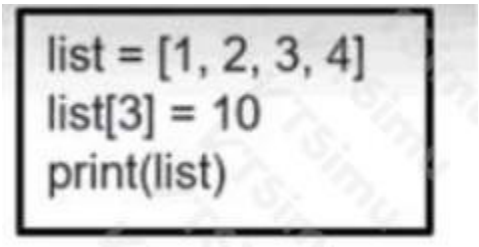

Refer to the exhibit. What is the value of the variable list after the code is run?

A. [1, 2, 10]

B. [1, 2, 3, 10]

C. [1, 2, 10, 4]

D. [1, 10, 10, 10]

ANSWER: C

Explanation:

`[1, 2, 10, 4]` is the resulting value because the code performs an indexed assignment on an existing Python list. A Python list is mutable, meaning that an individual item can be replaced without creating a new list or changing the other items. The list initially contains four values, and the assignment updates the item at the relevant position to `10`. Python list indexes begin at zero, so an index of `2` identifies the third item in the list. Consequently, the original third value is replaced with `10`, while the first, second, and fourth values remain unchanged. The list also keeps its original four-item length because item assignment replaces an element rather than inserting or deleting one. This fundamental behavior is important when working with Python-based network automation scripts, where lists commonly store interface names, IP addresses, VLANs, or device data that may need targeted updates. Python documents lists as mutable sequence types and supports assignment to individual list indexes. See the Python tutorial on data structures and the Python sequence types documentation.

QUESTION NO: 84

An engineer measures the Wi-Fi coverage at a customer site. The RSSI values are recorded as follows: LocationA -72dBm

Location B:-75 dBm Location C; -65 dBm Location D -80 dBm

Which two statements does the engineer use to explain these values to the customer? (Choose two.)

A. The signal strength at location B is 10 dB better than location C.

B. The signal strength at location C is too weak to support web surfing

C. Location D has the strongest RF signal strength.

D. The RF signal strength at location B is 50% weaker than location A.

E. The RF signal strength at location C is 10 times stronger than location B

F. The RF signal strength at location C is stronger than location A.

ANSWER: E F

Explanation:

“The RF signal strength at location C is 10 times stronger than location B” is correct because RSSI is expressed in dBm, a logarithmic power measurement. Values closer to zero represent higher received signal power. The difference between −65 dBm at location C and −75 dBm at location B is 10 dB. A 10 dB increase corresponds to a tenfold increase in received power, so the signal measured at location C is approximately ten times the power measured at location B.

“The RF signal strength at location C is stronger than location A” is also correct. Location C measures −65 dBm, while location A measures −72 dBm. Since −65 dBm is closer to zero than −72 dBm, location C has the stronger received RF signal. The 7 dB difference is substantial: it represents roughly five times as much received power at location C. These comparisons describe received signal power only; practical wireless performance also depends on factors such as noise level, signal-to-noise ratio, interference, client capability, and channel utilization. Cisco describes RSSI as a received signal-power indication and explains the importance of evaluating it alongside SNR for wireless connectivity and performance. See Cisco: Understanding RSSI and SNR and Cisco wireless design resources.

QUESTION NO: 85

Which two statements are true about Cisco EtherChannel? (Choose two.)

A. It combines multiple physical links into one logical port-channel interface.

B. It can use LACP to negotiate the channel.

C. It requires all member links to operate at different speeds.

D. It creates one independent STP topology for every physical member link.

E. It eliminates the need for Layer 2 loop prevention.

ANSWER: A B

Explanation:

It combines multiple physical links into one logical port-channel interface. EtherChannel presents the bundled links to Spanning Tree Protocol and to the routing process as a single logical interface. This reduces the number of logical links that STP must evaluate while allowing the member links to provide aggregate bandwidth.

It can use LACP to negotiate the channel. Link Aggregation Control Protocol is the IEEE standard negotiation protocol for link aggregation. Cisco devices can also use the Cisco-proprietary Port Aggregation Protocol on platforms that support it, or a channel can be configured without a negotiation protocol. See the Cisco IOS XE EtherChannel Configuration Guide.

QUESTION NO: 86

In a LISP topology, which component does the ITR send a resolution request to when it does not know the path to a specific EID?

A. Map resolver

B. Routing locator

C. Proxy ingress tunnel router

D. Proxy egress tunnel router

ANSWER: A

Explanation:

Map resolver is correct. In the Locator/ID Separation Protocol control plane, an ingress tunnel router must obtain an Endpoint ID to Routing Locator mapping before it can encapsulate traffic for an EID whose mapping is not already present in its map cache. The ITR creates a LISP Map-Request and sends it to a Map Resolver. The Map Resolver is the control-plane component that accepts Map-Requests from ITRs and determines where to forward the request, typically by consulting its mapping knowledge or forwarding it toward an authoritative Map Server. The authoritative mapping system ultimately returns a Map-Reply containing the reachable RLOC information; the ITR installs this result in its cache and uses the RLOC as the outer destination when it encapsulates data-plane packets. This process is central to LISP's separation of endpoint identity, represented by the EID, from network location, represented by the RLOC. Cisco's LISP documentation describes the Map Resolver role as receiving Map-Requests from ITRs and forwarding them toward the appropriate mapping authority. See the Cisco LISP overview and RFC 6833: LISP Map-Server.

QUESTION NO: 87

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:

Traditional WAN architectures use the routers' native routing and forwarding functions directly, which provides **lower control plane abstraction**. Routing decisions and reachability are established with familiar mechanisms such as static routes, OSPF, EIGRP, or BGP, rather than through a controller-managed overlay control plane and centrally distributed policies. This can be beneficial when an organization requires direct device-level visibility and operational control over routing behavior, with fewer overlay-specific components involved in path selection and policy enforcement.

Traditional WANs can also provide **lower data plane overhead**. In a conventional routed WAN, packets generally traverse the network without the additional headers associated with an SD-WAN overlay tunnel. SD-WAN commonly uses secure overlay tunnels across available transports, and encapsulation plus encryption consumes packet space and can reduce the effective payload size available within a link MTU. Avoiding that overlay encapsulation preserves more usable bandwidth per packet and reduces the need to account for tunnel-related MTU adjustments. These characteristics are tradeoffs: SD-WAN adds centralized management, policy control, and transport independence, while a traditional WAN can remain simpler at the native routing and packet-forwarding layers. Cisco describes SD-WAN as an overlay architecture with centralized management and secure transport tunnels in its Cisco SD-WAN overview and Cisco SD-WAN Design Guide.

QUESTION NO: 88

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:

LISP, or Locator/ID Separation Protocol, is based on separating endpoint identity from network location. The **EID** (Endpoint Identifier) namespace identifies an endpoint or a set of endpoints. An EID can be represented by an IPv4 or IPv6 address prefix, a MAC address, or other identifier types. It is the address space used by communicating hosts and represents the destination identity that an ingress tunnel router needs to reach.

The **RLOC** (Routing Locator) namespace identifies the topologically meaningful location of a LISP tunnel router in the underlying transport network. After the ingress tunnel router obtains an EID-to-RLOC mapping, it encapsulates packets addressed to the EID and forwards them toward the selected RLOC. This separation enables endpoint addressing to remain independent of the routing topology and supports mobility, multihoming, and more scalable routing designs. The EID-to-RLOC mapping relationship is therefore the central architectural function of LISP. Cisco's LISP overview and the IETF LISP specification define these as the two namespaces used by the protocol. See Cisco IOS XE LISP Overview and RFC 6830.

QUESTION NO: 89

Refer to the Exhibit.

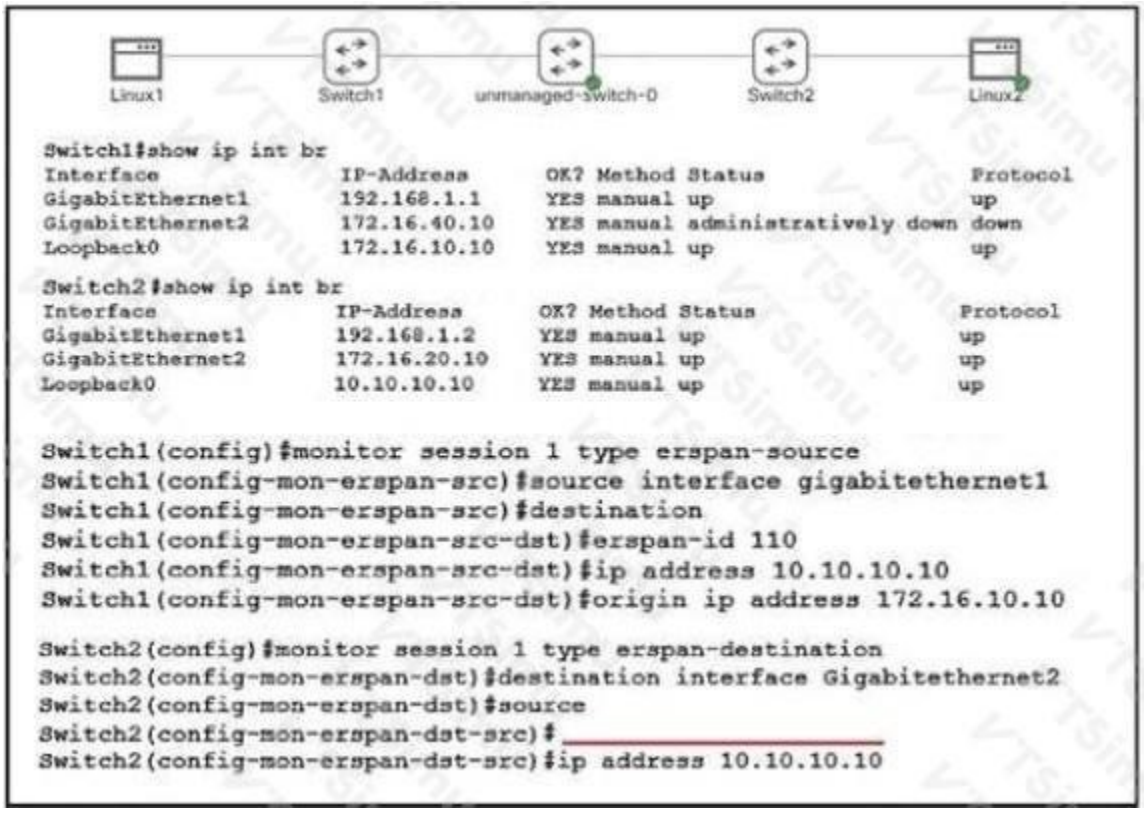

Refer to the exhibit. An engineer must configure an ERSPAN tunnel that mirrors traffic from linux1 on Switch1 to Linux2 on Switch2. Which command must be added to the destination configuration to enable the ERSPAN tunnel?

A. (config-mon-erspan-dst-src)# origin ip address 172.16.10.10

B. (config-mon-erspan-dst-src)# erspan-id 172.16.10.10

C. (config-mon-erspan-dst-src)# no shut

D. (config-mon-erspan-dst-src)# erspan-id 110

ANSWER: A

Explanation:

`origin ip address 172.16.10.10` is required in the ERSPAN destination session's source configuration because it identifies the IP address of the remote switch that originates the encapsulated ERSPAN traffic. ERSPAN transports mirrored packets across an IP network using GRE encapsulation, and the receiving switch must match incoming ERSPAN packets to the configured destination session. The origin IP address, together with the configured ERSPAN ID, provides that identification and allows Switch2 to accept the mirrored flow sent from Switch1. In this topology, 172.16.10.10 is the originator address shown for the ERSPAN source, so it must be configured as the origin address at the collector-side destination session. Once received and matched, the destination session forwards the decapsulated mirrored traffic to the configured destination interface connected to Linux2. Cisco's ERSPAN configuration model specifically includes an origin IP address under an ERSPAN destination source definition to identify the remote ERSPAN source. See Cisco's ERSPAN configuration example and the Cisco SPAN and ERSPAN configuration guide.

QUESTION NO: 90

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:

The router with the lowest IP address is elected as the IGMP querier on a shared LAN segment. In IGMPv2 and IGMPv3, multicast-capable routers send General Membership Queries to discover whether receivers still want traffic for particular multicast groups. Because multiple routers sending periodic queries on the same segment would be unnecessary, routers use a querier-election process. Each router observes queries received on the interface and compares the source IP address of the querying router with its own interface address. The router using the numerically lowest IP address becomes, and remains, the active querier. Other multicast routers operate as non-queriers while they continue monitoring for queries from the elected device. If they no longer receive queries within the applicable other-querier-present interval, they reevaluate the election and the lowest-addressed available router assumes the querier role. This predictable behavior ensures one router maintains membership-query activity for the subnet while allowing resilient takeover if that router fails. Cisco multicast deployments therefore require careful consideration of interface addressing when a specific router should serve as querier. The election behavior is specified in the IGMPv2 standard and is also described in Cisco's multicast documentation. See RFC 2236, section 4 and Cisco IGMP Snooping Overview and Behavior.

QUESTION NO: 91

Refer to the exhibit. A network engineer attempts to connect to the Router1 console port.

Which configuration is needed to allow Telnet connections?

A. Router1(config)# line vty 0 15
Router1(config-line)# transport output telnet

B. Router1(config)# telnet client

C. Router1(config)# line console 0
Router1(config-line)# transport output telnet

D. Router1(config)# access-list 100 permit tcp any any eq telnet
Router1(config)# line console 0
Router1(config-line)# access-class 100 out

E. Router1(config)# line vty 0 15
Router1(config-line)# transport input telnet

ANSWER: E

Explanation:

`Router1(config)# line vty 0 15`
`Router1(config-line)# transport input telnet` is the required configuration because incoming Telnet management sessions are handled by the router's virtual terminal (VTY) lines. The `transport input` command specifies which remote-access protocols the VTY lines accept; explicitly configuring `telnet` permits inbound Telnet sessions on those lines. The range `vty 0 15` configures all available VTY lines on platforms that support sixteen simultaneous VTY sessions, ensuring a Telnet connection can be assigned to an enabled line. Telnet uses TCP port 23 and provides remote CLI access, unlike the physical console port, which is intended for direct, local console connections. In a complete production configuration, the VTY lines must also have an authentication method configured, such as a local username database with `login local`, and Cisco recommends SSH instead of Telnet because Telnet sends credentials and session data in clear text. Nevertheless, when the requirement is specifically to permit Telnet, allowing Telnet under VTY line configuration is the applicable IOS/IOS XE mechanism. Cisco documents that `transport input` defines the protocols accepted for inbound connections on a line: Cisco IOS XE transport input Command Reference. General VTY access configuration is also covered in Cisco's VTY access documentation.

QUESTION NO: 92

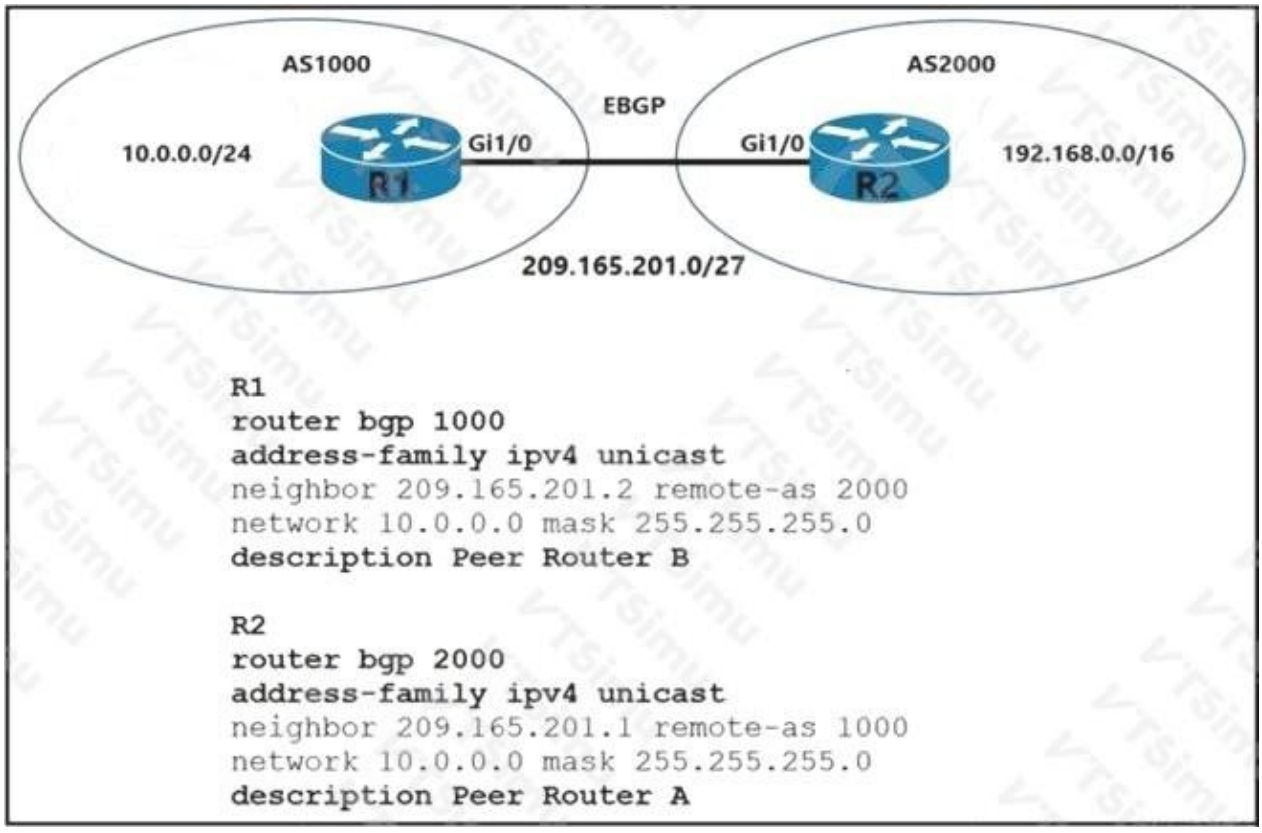

Refer to the exhibit. Which two commands are needed to allow for full reachability between AS 1000 and AS 2000? (Choose two.)

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:

The commands *R2#network 192.168.0.0 mask 255.255.0.0* and *R1#network 192.168.0.0 mask 255.255.0.0* originate the 192.168.0.0/16 prefix into the respective routing processes, allowing that reachability information to be exchanged across the external peer relationship between AS 1000 and AS 2000. For end-to-end communication, each autonomous system must learn a valid route for the destination networks located behind the remote edge router. Advertising the internal 192.168.0.0/16 route from both relevant edge routers supplies the necessary reachability information in each direction.

A *network* statement does not itself create a route. It instructs the routing process to advertise a prefix only when an exactly matching prefix and mask are already present in the local routing table. Therefore, the 192.168.0.0/16 route must exist locally, such as through a connected interface, static route, or interior routing protocol, before it is originated. Once those matching routes are present and advertised, each autonomous system can install the learned remote route and forward traffic toward the external peer. Cisco documents this route-origination behavior in its BGP route advertisement guidance and the Cisco IOS BGP network command reference.

QUESTION NO: 93

Which two features does the Cisco SD-Access architecture add to a traditional campus network? (Choose two.)

A. software-defined segmentation

B. private VLANs

C. SD-WAN

D. modular QoS

E. identity services

ANSWER: A E

Explanation:

Cisco SD-Access adds software-defined segmentation and identity services to traditional campus networking. Software-defined segmentation is implemented through the SD-Access fabric using virtual networks for macro-segmentation and scalable group tags for micro-segmentation. This lets an organization separate users, devices, and applications according to centrally defined policy, while maintaining consistent policy enforcement across the fabric rather than depending on individually configured VLANs and access-control lists. Segmentation can apply across both wired and wireless access and is not constrained by the physical placement of endpoints.

Identity services provide the identity, authentication, authorization, and contextual information needed to apply those policies. Cisco Identity Services Engine integrates with SD-Access to classify users and endpoints, assign scalable group tags, and supply identity-based policy information. The fabric can then enforce communication rules based on group identity and policy intent, rather than only on IP addressing or switch port location. Together, these capabilities make access control more automated, scalable, and consistent in an enterprise campus. Cisco describes SD-Access fabric segmentation and policy capabilities in its SD-Access overview; Cisco ISE identity and policy capabilities are described on the Cisco Identity Services Engine product page.

QUESTION NO: 94

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:

OMP (Overlay Management Protocol) is the Cisco SD-WAN overlay control-plane protocol used by WAN Edge devices and vSmart controllers to exchange the information required to establish reachability across the fabric. **advertisement of network prefixes and their attributes** is a core OMP function: WAN Edges advertise and receive OMP routes for VPN prefixes, together with relevant route attributes. OMP also distributes Transport Locator (TLOC) information, enabling devices to associate reachable prefixes with the transport endpoints and tunnels through which those prefixes can be reached. This control-plane exchange lets each WAN Edge select valid overlay paths according to OMP path selection and centralized policy.

delivery of crypto keys is also performed through OMP. Cisco SD-WAN uses centrally generated and distributed IPsec keying material so WAN Edges can create authenticated, encrypted IPsec tunnels for the data plane. vSmart controllers distribute this information securely over their authenticated control-plane connections, allowing encrypted overlay tunnels to be formed dynamically without manual pairwise key configuration. Together, route/TLOC advertisement and key distribution provide the reachability and security information needed for scalable SD-WAN overlay operation. Cisco documents OMP route and TLOC propagation in its SD-WAN OMP Configuration Guide and describes the SD-WAN security/control-plane architecture in the SD-WAN Security Configuration Guide.

QUESTION NO: 95

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:

The 10.1.1.0/27 subnet is assigned as the inside local addresses because the NAT access list identifies that private source network as traffic eligible for inside-source translation. An inside local address is the address assigned to a host on the internal network before NAT processing. In this configuration, hosts using addresses in the permitted 10.1.1.0/27 range match the ACL when they initiate applicable traffic through the router.

Inside source addresses are translated to the 209.165.201.0/27 subnet because the NAT rule associates the ACL-selected inside sources with the configured NAT pool. The pool supplies the translated, publicly routable address space used for the inside global side of NAT. As translations are created, the router substitutes an eligible pool address for the originating 10.1.1.x source address, allowing return traffic to be mapped back to the corresponding inside host. This is the normal operation of inside source dynamic NAT using an ACL and address pool. Cisco defines inside local addresses as the

pretranslation internal host addresses and inside global addresses as the translated addresses representing those hosts externally. See Cisco NAT terminology and operation and the Cisco IOS NAT Configuration Guide.

QUESTION NO: 96

Refer to the exhibit.

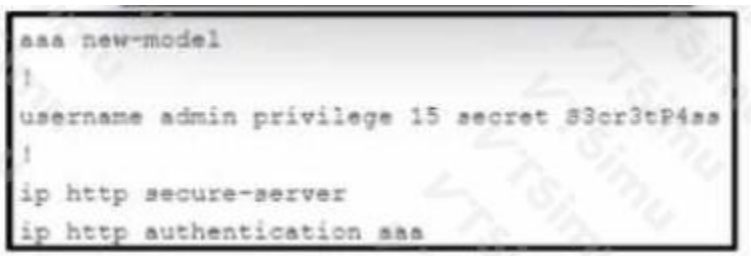

63

An administrator must enaWe RESTCONP access to a router. Which two commands or command sets must be added to the existing configuraton? (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:

RESTCONF on Cisco IOS XE requires both the RESTCONF service and an HTTP-based transport service. The command set represented by **Option B** supplies the required HTTP or HTTPS server capability, allowing the router to listen for RESTCONF client requests. In production, HTTPS is normally preferred because it protects API credentials and configuration data while in transit. RESTCONF does not create a separate standalone network listener; it uses the IOS XE HTTP/HTTPS infrastructure.

The command set represented by **Option D** enables RESTCONF itself. This activates the RESTCONF interface and exposes the device's YANG-modeled configuration and operational data through RESTCONF resource paths, normally rooted at `/restconf`. Both functions must be present: the HTTP(S) service makes the API reachable, and RESTCONF

activation makes the API resources available. The existing local username configuration can then be used with the configured HTTP authentication method, subject to the device's AAA and authorization settings.

Cisco documents enabling the HTTP secure server and enabling RESTCONF as core IOS XE RESTCONF configuration steps. See the Cisco IOS XE RESTCONF Configuration Guide and the protocol definition in RFC 8040.

QUESTION NO: 97

When does a Cisco StackWise primary switch lose its role?

A. when a stack member fails

B. when the stack primary is reset

C. when a switch with a higher priority is added to the stack

D. when the priority value of a stack member is changed to a higher value

ANSWER: B

Explanation:

when the stack primary is reset is correct because the StackWise primary (also called the stack master or active switch, depending on platform and software terminology) owns the stack's control-plane role while it remains operational and connected to the stack. A reset makes that current primary unavailable to continue performing the role. The remaining operational stack members then perform the stack-master election process and select a new primary according to the applicable election criteria, including stored configuration and stack-member priority.

Stack mastership is not normally preemptive. In other words, the existing primary does not surrender its role during normal operation merely because election-related attributes elsewhere in the stack change. A master-election event is required when the current primary becomes unavailable, such as through a reset. This preserves control-plane stability and prevents routine stack changes from unnecessarily disrupting the active management role. After the reset switch returns, it rejoins as a member unless another election is triggered under the platform's supported behavior. Cisco's Stack Manager documentation describes the election process and the active switch's role, while the Catalyst 9300 configuration-guide portal provides current StackWise platform documentation: Cisco Stack Manager Configuration Guide and Cisco Catalyst 9300 Configuration Guides.

QUESTION NO: 98

Refer io me exhibit.

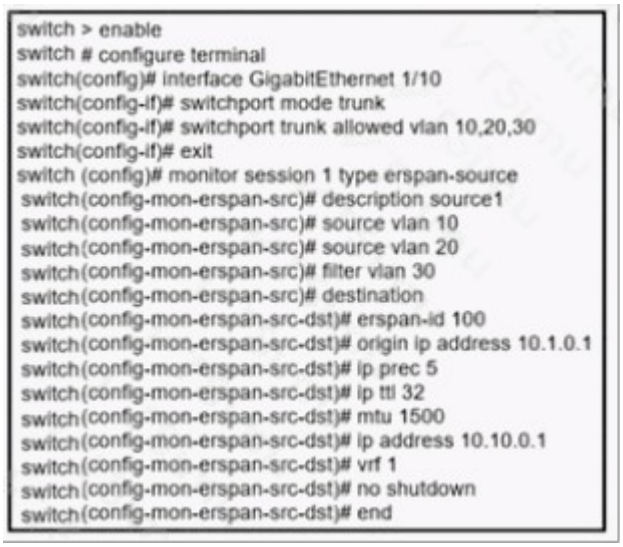

An engineer configures the trunk and proceeds to configure an ESPAN session to monitor VLANs10. 20. and 30. Which command must be added to complete this configuration?

A. Device(config.mon.erspan.stc)# no filter vlan 30

B. Devic(config.mon.erspan.src-dst)# no vrf 1

C. Devic(config.mon.erspan.src-dst)# erspan id 6

D. Device(config.mon-erspan.Src-dst)# mtu 1460

ANSWER: C

Explanation:

`erspan id 6` must be configured in the ERSPAN source session's destination configuration to complete the ERSPAN encapsulation parameters. ERSPAN transports mirrored Layer 2 traffic across a Layer 3 network by encapsulating it in GRE. The ERSPAN identifier is carried in the ERSPAN header and identifies the specific mirrored traffic stream at the receiving endpoint. It must match the ERSPAN ID expected by the remote ERSPAN destination/collector so that the received mirrored packets are associated with the intended monitoring session.

The trunk-facing source and VLAN filter define which traffic is copied into the monitoring session; in this case, the VLAN selection limits the mirrored traffic to the required VLANs. The destination portion then requires the information needed to send that copy through the routed network, including its IP addressing and an ERSPAN ID. Configuring `erspan id 6` under the ERSPAN source destination context supplies that mandatory stream identifier and allows the device to generate properly identified ERSPAN packets for the remote analyzer.

Cisco's ERSPAN configuration documentation describes the ERSPAN source destination parameters and ERSPAN ID requirement in the Cisco IOS XE SPAN and ERSPAN Configuration Guide. Cisco also provides an overview of ERSPAN operation in its SPAN and RSPAN configuration guidance.

QUESTION NO: 99

What is a characteristic of MACsec?

A. 8021AE provides encryption and authentication services

B. 8021AE is butt between the host and switch using the MKA protocol which negotiates encryption keys based on the master session key from a successful 802IX session

C. 802.1AE is built between the host and switch using the MKA protocol using keys generated via the Diffie-Hellman algorithm (anonymous encryption mode)

D. 8021AE is negotiated using Cisco AnyConnect NAM and the SAP protocol

ANSWER: A

Explanation:

8021AE provides encryption and authentication services is correct because MACsec is the IEEE 802.1AE standard for securing Ethernet frames on a Layer 2 link. It protects traffic on a hop-by-hop basis by applying confidentiality through encryption and by providing data-origin authenticity and integrity protection through an integrity check value. MACsec also includes replay protection, allowing a receiving device to identify and discard frames that have been replayed or received outside the expected packet-number sequence. These protections are applied to the MAC service data unit as it crosses a secured Ethernet link, making MACsec useful for protecting traffic between directly connected network devices, such as a host and an access switch or two switches. Cisco platforms commonly use MACsec with IEEE 802.1X and MKA for authentication and key-management deployments, but the fundamental characteristic of the 802.1AE MACsec standard itself is its Ethernet-frame encryption and authentication capability. Cisco's MACsec configuration documentation describes MACsec as providing Layer 2 encryption, integrity, and replay protection; see Cisco IOS XE MACsec Configuration Guide. The IEEE standard is listed at IEEE 802.1AE.

QUESTION NO: 100

Which two methods are used to reduce the AP coverage area? (Choose two.)

A. Reduce channel width from 40 MHz to 20 MHz.

B. Reduce AP transmit power.

C. Enable Fastlane.

D. Increase minimum mandatory data rate.

E. Disable 2.4 GHz and use only 5 GHz.

ANSWER: B D

Explanation:

Reducing AP transmit power is a direct RF cell-sizing control. A lower transmit-power setting reduces the received signal level from the AP as distance increases, so the usable edge of the cell moves closer to the AP. Cisco wireless RF design uses transmit-power tuning to control cell overlap, support predictable roaming, and limit excessive contention in dense deployments. AP power should be designed together with the client side of the link, since client devices often have lower transmit capability than an AP.

Increasing the minimum mandatory data rate also reduces the practical coverage area. Mandatory rates are used for management and control traffic and establish the lowest rate at which a client must be able to communicate. Higher rates require stronger received signal quality and a better signal-to-noise ratio. Clients near the former edge of the cell that cannot reliably sustain that higher mandatory rate therefore cannot remain effectively connected at that distance, pulling the usable cell boundary inward. Cisco recommends careful data-rate configuration as part of RF design because it affects both coverage and roaming behavior. See Cisco's RF configuration documentation and its enterprise wireless resources.

QUESTION NO: 101

```
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
------------------- ---- --- --------- -------- --------------------------------
Fa1/0/7             Desg FWD 2         128.9    P2p Bound (PVST)
Fa1/0/10            Desg FWD 2         128.12   P2p Bound (PVST)
Fa1/0/11            Root FWD 2         128.13   P2p
Fa1/0/12            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    Fa1/0/11       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:

With Multiple Spanning Tree, bridge-root elections occur independently for each MST instance rather than individually for every VLAN. VLANs 10 and 20 are mapped to MST instance 1 in the exhibited MST configuration, so DSW1 must be configured to have the preferred bridge ID for instance 1. The command `spanning-tree mst 1 priority 4096` explicitly assigns DSW1 a low valid STP priority for that instance. Bridge priority values are set in increments of 4096, and a lower bridge ID wins the root-bridge election. Provided no competing bridge has a lower bridge ID, this priority causes DSW1 to be elected root for instance 1 and therefore for its mapped VLANs 10 and 20.

The command `spanning-tree mst 1 root primary` is Cisco's root-primary macro for MST instance 1. It automatically selects an appropriately low priority—typically 24576 or 4096 lower than the current root's priority—to make the local switch the primary root for that instance. Because both commands target the specific MST instance carrying VLANs 10 and 20, each is a valid method for placing the STP root role on DSW1. Cisco documents MST instance mapping and per-instance root operation in its Spanning Tree Protocol overview and MSTP configuration guide.

QUESTION NO: 102

A network engineer configures a WLAN controller with increased security for web access. There is IP connectivity with the WLAN controller, but the engineer cannot start a management session from a web browser. Which action resolves the issue?

A. Disable Adobe Flash Player.

B. Use a private or incognito session.

C. Use a browser that supports 128-bit or larger ciphers.

D. Disable JavaScript on the web browser.

ANSWER: C

Explanation:

Use a browser that supports 128-bit or larger ciphers. On Cisco wireless LAN controllers, enabling stronger security for web administration can restrict HTTPS access to cryptographic capabilities that older or export-grade browser implementations cannot negotiate. Basic IP connectivity only confirms that the controller is reachable at Layer 3; it does not establish that the browser and controller can complete the HTTPS/TLS negotiation required to open the management GUI. A browser with support for strong encryption can negotiate a mutually supported cipher suite and establish the encrypted session needed for web management. Cisco documentation for AireOS controllers notes browser cipher-strength requirements for secure web access, including the need for a browser supporting 128-bit encryption when strong encryption is configured. Current browser TLS support and cipher availability are therefore a management-client compatibility requirement, not a routing or reachability issue. For Cisco guidance on controller web administration and browser requirements, see Cisco Wireless Controller Configuration Guide. Cisco also describes how HTTPS relies on SSL/TLS cipher-suite negotiation in its SSL/TLS interoperability documentation.

QUESTION NO: 103

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

A. Use a virtual switch provided by the hypervisor.

B. Use a virtual switch running as a separate virtual machine.

C. Use VXLAN fabric after installing VXLAN tunneling drivers on the virtual machines.

D. Use a single routed link to an external router on stick.

E. Use a single trunk link to an external Layer2 switch.

ANSWER: A E

Explanation:

Use a virtual switch provided by the hypervisor. A hypervisor virtual switch connects virtual NICs within the host at Layer 2 and applies network policy at virtual ports or port groups. Administrators can assign VM interfaces to VLAN-backed port groups, enforce VLAN isolation, and apply controls such as MAC-address and promiscuous-mode security policies. This gives VMs local Ethernet connectivity without requiring their frames to leave the host when the destination is also local.

Use a single trunk link to an external Layer2 switch. The hypervisor's physical uplink can operate as an IEEE 802.1Q trunk, carrying traffic for multiple VLANs between the host and the physical switching infrastructure. VM port groups map virtual NIC traffic into the appropriate VLANs, maintaining Layer 2 segmentation and allowing the external switch to extend each VLAN to other hosts or network devices. This design provides a controlled and scalable connection between the virtual switching environment and the physical Layer 2 network while preserving VLAN boundaries. Cisco describes VLAN trunk operation at Cisco VLAN Trunking and VLAN Basics, and VMware documents VLAN configuration for virtual networking at vSphere VLAN Configuration.

QUESTION NO: 104

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:

Configuring `logging synchronous` under the VTY line is appropriate when the administrator is connected remotely through Telnet or SSH. This line-configuration command changes how IOS displays unsolicited system messages, including debug and logging output, while a command is being entered. After displaying a message, IOS redraws the command prompt and restores the partially entered command, so the administrator can clearly see the input before pressing Enter. The command is applied in the relevant `line vty` configuration context, such as `line vty 0 4`, rather than as a global configuration command.

Pressing the TAB key is also useful during an interrupted entry. IOS uses TAB for command completion and for redisplaying the command line; this lets the administrator restore a readable command line after output has appeared in the terminal. Used with synchronous logging, it helps preserve operator awareness of the exact command text being entered while routing debugs generate frequent messages. Cisco documents synchronous logging as a feature that prevents system messages from disrupting command entry and describes terminal editing and completion behavior in its IOS CLI documentation. See Cisco: Using logging synchronous and Cisco IOS XE CLI Basics.

QUESTION NO: 105

Which Cisco wireless RRM function dynamically assigns RF channels to access point radios to reduce co-channel interference?

A. Dynamic Channel Assignment

B. Transmit Power Control

C. Coverage Hole Detection

D. Client Load Balancing

ANSWER: A

Explanation:

Dynamic Channel Assignment is the Radio Resource Management function that selects and adjusts wireless channels for access point radios. DCA evaluates RF conditions and channel use, then assigns channels to reduce co-channel interference and improve the overall RF design. In a controller-based Cisco wireless deployment, this function allows the network to adapt to changing RF conditions without requiring administrators to statically plan every channel assignment.

Transmit Power Control is a separate RRM function that adjusts AP transmit power. Coverage Hole Detection identifies client coverage problems, while Client Load Balancing influences client distribution. Cisco describes DCA and the related RRM functions in the Cisco Wireless Controller Configuration Guide.

QUESTION NO: 106

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 E

Explanation:

Cisco Catalyst SD-WAN uses a secure, logically centralized control-plane architecture. The control plane operates over DTLS/TLS authenticated and secured tunnels because WAN Edge devices establish authenticated control connections with Cisco SD-WAN controllers. TLS is used over TCP and DTLS is used over UDP, protecting control-plane communications through encryption, integrity protection, and peer authentication. This allows edge devices and controllers to exchange overlay, route, and policy information securely, including when the underlay is an untrusted Internet connection.

Centralized reachability, security, and application policies are another defining characteristic. Cisco SD-WAN Manager provides centralized management and policy definition, while controllers distribute the required control information to WAN Edge devices. Administrators can centrally define routing and traffic-engineering intent, segmentation and VPN policy, service-chaining behavior, and application-aware forwarding rules. The WAN Edge devices then enforce those policies locally in the data plane, enabling consistent treatment across branch locations while retaining direct forwarding performance. This model supports scalable operations and policy consistency without requiring independent configuration of every individual site.

See Cisco's Catalyst SD-WAN overview and the Cisco SD-WAN Security Configuration Guide for control-plane security and centralized policy architecture details.

QUESTION NO: 107

Which plane transports traffic between sites across the SD-WAN fabric?

A. IPsec

B. Management

C. Data

D. Control

ANSWER: C

Explanation:

The **Data** plane transports actual user and application traffic between sites across a Cisco SD-WAN fabric. WAN Edge devices establish secure data-plane tunnels, typically IPsec tunnels, over the available transport networks. Once traffic enters the SD-WAN overlay, the WAN Edge forwards packets through these tunnels toward the destination site according to routing information and centralized policy. The data plane therefore provides the packet-forwarding function that carries business applications, voice, video, and other payload traffic across the fabric.

Cisco SD-WAN separates this forwarding role from the functions that configure the fabric and distribute routing or policy information. This separation lets the overlay maintain secure connectivity and application-aware forwarding across diverse underlay transports such as Internet, MPLS, and LTE. Cisco documents that the data plane consists of WAN Edge devices and the secure IPsec tunnels used for forwarding traffic between them. See Cisco's Cisco SD-WAN System Overview and the Cisco SD-WAN solution overview.

QUESTION NO: 108

In a campus network design, what are two benefits of using BFD for failure detection'? (Choose two.)

A. BFD speeds up routing convergence time

B. BFD provides path failure detection in less than a second

C. BFD provides fault tolerance by enabling multiple routers to appear as a single virtual router

D. BFD is an efficient way to reduce memory and CPU usage

E. BFD enables network peers to continue forwarding packets in the event of a restart

ANSWER: A B

Explanation:

"BFD speeds up routing convergence time" is correct because Bidirectional Forwarding Detection supplies rapid forwarding-path failure information to routing protocols. When a BFD session detects that an adjacent forwarding path is unavailable, it signals the associated protocol, allowing it to withdraw or recompute routes and select an alternate path much sooner than it could when relying solely on that protocol's standard hello and dead intervals. This is especially valuable in campus networks that require rapid recovery for critical applications.

"BFD provides path failure detection in less than a second" is also correct. BFD is a lightweight, protocol-independent mechanism designed for low-overhead and fast detection between systems. Its detection interval is based on negotiated control-packet transmission intervals and a detection multiplier, permitting subsecond detection when the platform, interface type, and configured timers support it. A routing protocol such as OSPF, EIGRP, IS-IS, or BGP can then use the BFD session state rather than waiting for its normal neighbor timeout. Cisco documents BFD as a method for providing rapid detection of failures in the forwarding path and for improving convergence following a failure. See Cisco's BFD overview and the Cisco IOS XE BFD configuration guide.

QUESTION NO: 109

Refer to this output.

R1# *Feb 14 37:09:53.129: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/1, changed state to up
What is the logging severity level?

A. notification

B. emergency

C. critical

D. alert

ANSWER: A

Explanation:

The correct answer is **notification**. Cisco IOS syslog messages use a standard identifier format: `%FACILITY-SEVERITY-MNEMONIC`. In `%LINEPROTO-5-UPDOWN`, `LINEPROTO` is the facility, `5` is the severity value, and `UPDOWN` is the mnemonic describing the event. Cisco assigns severity 5 to the notifications level. Notifications represent normal, but significant, events that an administrator may want recorded, including operational state changes such as an interface line protocol coming up or going down. The event shown reports that the line protocol on GigabitEthernet0/1 changed to the up state, so its embedded severity value directly identifies it as a notification. Cisco syslog uses levels from 0 through 7, with lower numerical values representing increasingly urgent conditions. Reading the severity digit in the message itself is therefore the reliable way to identify the event's logging level, independently of the destination where the message is displayed or stored. Cisco documents the severity mapping and the format of IOS system messages in its Syslog Messages and Severity Levels guide and its IOS XE System Message Logging documentation.

QUESTION NO: 110

What does the destination MAC on the outer MAC header identify in a VXLAN packet?

A. thee emote spine

B. the next hop

C. the leaf switch

D. the remote switch

ANSWER: B

Explanation:

The next hop is correct because the outer Ethernet header belongs to the underlay transport packet, not to the original Layer 2 frame being carried inside VXLAN. Before encapsulation, the sending VTEP determines how to reach the remote VTEP using the outer destination IP address. Its normal underlay forwarding lookup selects an egress interface and, where applicable, an IP next hop. The sending VTEP then resolves that next hop's Ethernet address through ARP for IPv4 or Neighbor Discovery for IPv6 and places the resolved address in the outer destination MAC field.

This allows the first Ethernet segment of the underlay to deliver the VXLAN-encapsulated UDP packet to the appropriate adjacent device. As the packet crosses routed underlay hops, each router removes and rebuilds the Layer 2 header for its own outgoing link; therefore, the outer destination MAC can change at every hop. The outer IP destination remains the remote VTEP address, while the original source and destination MAC addresses remain preserved in the inner Ethernet frame for overlay forwarding. VXLAN is consequently transported as a MAC-in-UDP encapsulation over conventional Layer 2 and Layer 3 underlay forwarding. See RFC 7348: Virtual eXtensible Local Area Network and Cisco Nexus 9000 VXLAN overview.

QUESTION NO: 111

What are two best practices when designing a campus Layer 3 infrastructure? (Choose two.)

A. Summarize routes from the aggregation layer toward the core layer.

B. Summarize from the access layer toward the aggregation layer.

C. Configure passive-interface on nontransit links.

D. Implement security features at the core.

E. Tune Cisco Express Forwarding load balancing hash for ECMP routing.

ANSWER: A C

Explanation:

Summarize routes from the aggregation layer toward the core layer is a campus-design best practice because the aggregation/distribution layer is the routing boundary between individual access blocks and the highly available core. Advertising summarized access-block prefixes northbound reduces the number of routes carried by the core, limits the propagation of topology changes, improves routing-table scalability, and supports fault isolation. Cisco's routing design guidance describes route aggregation as a way to reduce routing information and contain instability: Cisco route aggregation documentation.

Configure passive-interface on nontransit links is also appropriate. A nontransit interface, such as one facing an end-host subnet or a management segment, does not need to establish routing adjacencies. Making that interface passive still advertises its connected network through the routing protocol while preventing routing-protocol hello packets and neighbor formation on that link. This reduces unnecessary control-plane traffic and exposure, and helps ensure adjacencies exist only where routing devices are intended to peer. Cisco documents that passive interfaces suppress routing-protocol packet transmission while retaining network advertisement behavior: Cisco OSPF passive-interface configuration guide.

QUESTION NO: 112

What is a difference between Chef and other automation tools?

A. Chef is an agenless tool that uses a primary/minion architecture, and SaltStack is an agent-based tool that uses a primary/secondary architecture

B. Chef uses Domain Specific Language, and Puppet uses Ruby.

C. Chef is an agent-based tool that uses cookbooks, and Ansible is an agentless tool that uses playbooks.

D. Chef is an agentless tool that uses playbooks, and Ansible is an agent-based toot that uses cookbooks.

ANSWER: C

Explanation:

Chef is an agent-based configuration-management platform: managed systems run Chef Infra Client, which retrieves assigned policy and performs a convergence run to bring the system to its declared desired state. Chef packages reusable infrastructure and configuration content in *cookbooks*. A cookbook can include recipes, attributes, templates, files, and other resources; recipes are written using Chef's Ruby-based domain-specific language. This model is particularly suited to continuously enforcing configuration policy on nodes that have the Chef client installed.

Ansible commonly operates without a persistent software agent on managed hosts. Its control node connects to targets, typically through SSH for network devices and Linux systems or WinRM for Windows systems, and executes automation defined in YAML *playbooks*. Thus, the distinction accurately combines both the principal execution architecture—Chef client-based versus Ansible agentless—and the principal content constructs—Chef cookbooks versus Ansible playbooks. Cisco automation workflows can use either approach, but recognizing these terms and their deployment models is important when selecting and operating an automation solution. See the Chef Infra Client overview and the Ansible Getting Started documentation.

QUESTION NO: 113

```
client.connect(sd, port=22, username=username, password=password, allow_agent=False)
stdin, stdout, stderr = client.exec_command('show lld neighbors\n')
u = 0
for u in stdout:
    if 'Router' not in u and 'Capability' not in u and 'Repeater' not in u:
        if 'Device ID' not in u and 'displayed' not in u:
            u101 = u.split()
            if len(u101) != 0:
                u2.append(u101)
    if 'displayed:' in u:
        cx = u.split()
        c0 = cx[3]
d1 = {'x0': u1, 'c0': c0}
```

Refer to the exhibit. What is achieved by this Python script?

A. It displays the output from show lldp neighbors into a standard output.

B. It reads the neighbor count from show lldp neighbors into a dictionary list.

C. It displays the Layer 3 neighbors from show lldp neighbors on the terminal screen.

D. It reads the output from show lldp neighbors into an array object.

ANSWER: B

Explanation:

It reads the neighbor count from show lldp neighbors into a dictionary list. The script uses the structured response returned for the `show lldp neighbors` command rather than treating the command result as unstructured terminal text. NX-API responses are encoded as JSON, which Python can deserialize into native dictionary and list objects. The response hierarchy contains dictionary keys for the API envelope and command output, while repeating LLDP neighbor records are represented as a list-like collection. The script traverses that hierarchy to access the LLDP neighbor data and obtains the count associated with the discovered neighbor entries. Storing and accessing the data in this structured form allows subsequent Python logic to inspect LLDP attributes programmatically, such as neighbor device identifiers, local interfaces, and remote ports, without needing to parse formatted CLI columns. This is the expected approach when automating NX-OS operational commands through NX-API: submit the CLI command, decode the JSON payload, and navigate its nested objects. Cisco documents that NX-API provides command output in structured formats, including JSON, for programmatic network automation. See the Cisco NX-API CLI Reference and the Cisco NX-OS Developer Documentation.

QUESTION NO: 114

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 provides a formal, machine-readable data-modeling language for describing network device configuration, operational state, RPC operations, and notifications. “It enforces configuration semantics” is correct because a YANG model defines the

meaning and structure of managed data: for example, its hierarchy, data types, relationships, defaults, and whether a node represents configurable or operational data. This gives controllers and management applications a consistent schema through which they can interpret and automate device behavior. “It enforces configuration constraints” is also correct because YANG supports validation rules such as ranges, string patterns, mandatory data, list keys, cardinality, uniqueness, and conditional `must` and `when` expressions. These constraints allow a server to validate submitted configuration data against the model before accepting or applying it, improving consistency and preventing invalid configurations from entering the device configuration. YANG is encoding-independent: its modeled data can be encoded in XML or JSON when used with management protocols such as NETCONF and RESTCONF. The IETF YANG specification defines both the data-modeling constructs and the validation constraints that enable model-driven network automation. See RFC 7950: The YANG 1.1 Data Modeling Language and RFC 8040: RESTCONF Protocol.

QUESTION NO: 115

Refer to the exhibit.

```
Status Code: 200
Body:
{
    "response": [
        {
            "memorySize": "3735302144",
            "family": "Wireless Controller",
            "role": "ACCESS",
            "description": "Cisco Controller Wireless Version:8.5.140.0",
            "roleSource": "AUTO",
            "lastUpdated": "2021-09-10 13:48:02",
            "deviceSupportLevel": "Supported",
            "softwareType": "Cisco Controller",
            "softwareVersion": "8.5.140.0",
            "macAddress": "ac:4a:56:6c:7c:00",
            "collectionInterval": "Global Default",
            "inventoryStatusDetail": "<status><general code=\"SUCCESS\"/></
            status>",
            "serialNumber": "FOL25040021",
            "lastUpdateTime": 1631281682276,
            "hostname": "c3504.abc.inc",
            "tagCount": "0",

        ***Output omitted***

            "lineCardId": "",
            "managedAtleastOnce": true,
            "location": null,
            "type": "Cisco 3504 Wireless LAN Controller",
            "managementState": "Managed",
            "instanceUuid": "6b741b27-f7e7-4470-b6fc-d5168cc59502",
            "instanceTenantId": "5e8e896e4d4add00ca2b6487",
            "id": "6b741b27-f7e7-4470-b6fc-d5168cc59502"
        }
    ],
    "version": "1.0"
}
```

Which HTTP request produced the REST API response that was returned by Cisco DNA Center?

A. fetch /network-device?macAddress=ac:4a:56:6c:7c:00

B. POST/network-device?macAddress=ac:4a:56:6c:7c:00

C. GET/network-device?macAddress=ac:4a:56:6c:7c:00

ANSWER: C

Explanation:

GET/network-device?macAddress=ac:4a:56:6c:7c:00 is correct because it performs a read operation against Cisco DNA Center's network-device inventory resource while using `macAddress` as a query parameter to locate the device associated with that MAC address. REST APIs use the HTTP `GET` method to retrieve a representation of an existing resource without changing the server-side inventory. The returned payload shown in the exhibit is therefore consistent with a device-inventory lookup: Cisco DNA Center receives the request, filters the network-device collection according to the supplied MAC address, and returns matching device information in the response body. Cisco's Intent APIs expose network-device inventory endpoints for retrieving discovered-device details, and their API conventions use query parameters to refine collection results. This request syntax identifies the resource path, `/network-device`, and appends the filtering criterion after the question mark. For additional details on Cisco DNA Center API authentication, endpoint usage, and Intent API workflows, see the Cisco DNA Center Developer documentation and Cisco's Get Network Device API reference.

QUESTION NO: 116

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

A. TCP MSS

B. PMTUD

C. DF bit Clear

D. MTU ignore

E. IP MTU

F. TCP window size

ANSWER: A E

Explanation:

TCP MSS and **IP MTU** are configured proactively to keep GRE-encapsulated traffic below the packet size that would require fragmentation. GRE adds an outer IP header and a GRE header to every tunneled packet. If IPsec also protects the GRE tunnel, additional encapsulation overhead further reduces the effective payload size that can traverse the path without fragmentation.

Configuring an IP MTU on the tunnel interface establishes a smaller maximum IP packet size for traffic sent into that tunnel. This accounts for encapsulation overhead before the router creates the GRE outer packet, helping ensure that the resulting packet fits the available underlying path MTU. TCP MSS adjustment complements this setting for TCP flows. By advertising a reduced maximum TCP segment size during TCP session establishment, it causes endpoints to send smaller TCP segments, preventing otherwise oversized IP packets from entering the tunnel.

Using both mechanisms is especially useful where Path MTU Discovery cannot be relied on consistently, such as when ICMP fragmentation-needed messages are filtered. Cisco documents tunnel-interface IP MTU and TCP MSS adjustment as mechanisms for avoiding fragmentation in GRE tunnel deployments. See Cisco's GRE tunnel MTU and MSS guidance and the Cisco IOS XE tunnel interface configuration guide.

QUESTION NO: 117

When using a Cisco Catalyst 9800 Series WLC, which tag/profile can be applied to APs to change the mode to FlexConnect in a specific location?

A. policy tag

B. site tag

C. AP join profile

D. flex profile

ANSWER: B

Explanation:

site tag is correct because Cisco Catalyst 9800 controllers use tags to apply configuration consistently to APs, and the site tag associates an AP with a site profile. The site profile contains location- and AP-operational settings, including the AP mode. Assigning a site tag whose associated site profile is configured for FlexConnect changes the APs assigned to that tag to operate in FlexConnect mode. This makes site tags particularly useful when APs in a branch or remote location need local switching behavior and survivability characteristics while APs at another location remain in local mode. The same site tag can be assigned to multiple APs, providing a scalable way to apply the appropriate mode for a physical site. Cisco's Catalyst 9800 tag architecture separates this site-specific operational assignment from WLAN-to-policy mapping, allowing administrators to organize AP configuration by location. See Cisco's Catalyst 9800 Tags Configuration Guide and the Catalyst 9800 FlexConnect Configuration Guide for the site-tag and FlexConnect workflow.

QUESTION NO: 118

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:

DHCP Option 43 is a valid controller-discovery mechanism because it lets a DHCP server supply a lightweight access point with one or more controller management IP addresses during address assignment. The AP can then use those addresses to send CAPWAP discovery messages, making this approach especially useful when the AP and wireless LAN controller reside on different IP subnets. Cisco documents the vendor-specific Option 43 format used to provide controller addresses to access points.

Broadcasting on the local subnet is also a valid discovery mechanism. When an AP is on the same Layer 2 broadcast domain as a controller management interface, it sends a CAPWAP discovery request as a local broadcast. A controller on that subnet can respond with a discovery response, allowing the AP to select a controller and proceed with the CAPWAP join process. This method is therefore useful for same-subnet deployments and initial connectivity scenarios. These mechanisms are part of Cisco's documented AP controller-discovery process; DHCP Option 43 supplies a reachable controller address through DHCP, while local broadcast discovery finds controllers available in the AP's own broadcast domain. See Cisco's AP join troubleshooting guidance and DHCP Option 43 configuration documentation.

QUESTION NO: 119

Which two descriptions of FlexConnect mode for Cisco APs are true? (Choose two.)

A. APs that operate in FlexConnect mode cannot detect rogue APs.

B. When connected to the controller, FlexConnect APs can tunnel traffic back to the controller.

C. FlexConnect mode is used when the APs are set up in a mesh environment and used to bridge between each other.

D. FlexConnect mode is a feature that is designed to allow specified CAPWAP-enabled APs to exclude themselves from managing data traffic between clients and infrastructure.

E. FlexConnect mode is a wireless solution for branch office and remote office deployments.

ANSWER: B E

Explanation:

FlexConnect is designed for enterprise branch and remote-office wireless deployments, particularly where the WAN connection to the centralized wireless LAN controller might have limited bandwidth, high latency, or intermittent availability. A FlexConnect access point still joins and is managed by the controller using CAPWAP, allowing the organization to centrally configure WLANs, security policies, and AP settings. This centralized management model makes FlexConnect practical for distributed sites without requiring a controller at every branch.

While the AP has connectivity to its controller, a WLAN on a FlexConnect AP can be configured for central switching. With central switching, client data traffic is encapsulated in the CAPWAP tunnel and sent back to the controller for forwarding. FlexConnect also supports local switching, in which client traffic is switched directly onto the branch LAN, but the availability of local switching does not remove the capability to tunnel traffic to the controller when central switching is configured. Cisco documents FlexConnect as supporting both centrally switched and locally switched WLAN traffic, depending on the WLAN configuration and deployment requirements. See Cisco's FlexConnect configuration guidance and the FlexConnect overview.

QUESTION NO: 120

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 northbound REST APIs that enable external applications, automation tools, and IT service-management platforms to consume Cisco DNA Center capabilities. They sit above the controller and provide a programmatic interface for requesting services, retrieving network information, managing policy, and using assurance and automation workflows. This northbound role is fundamental to the Cisco DNA Center platform architecture because it lets developers integrate business and operational systems without interacting directly with each managed network device.

The APIs are also business outcome oriented. Rather than requiring an application to specify every device-level command and configuration mechanism, an intent-based interface expresses the desired network result. Cisco DNA Center then translates that requested outcome into the appropriate workflows and infrastructure changes. For example, integrations can use the platform to automate provisioning, apply policy, or obtain assurance data in a manner aligned to organizational requirements, while the controller abstracts much of the underlying network complexity. This model supports Cisco's intent-based networking approach by connecting high-level business intent to automated network operations. Cisco documents these APIs as the platform's northbound interface and describes intent-based networking as using automation and analytics to implement business intent. See Cisco DNA Center Platform Overview: Intent API (Northbound) and Cisco Intent-Based Networking overview.

QUESTION NO: 121

Which two Cisco SD-Access components provide communication between traditional network elements and the controller layer? (Choose two.)

A. network underlay

B. network control platform

C. network data platform

D. partner ecosystem

E. fabric overlay

ANSWER: A E

Explanation:

The **network underlay** and **fabric overlay** are the SD-Access network components that together enable connectivity and communication between physical, traditional network infrastructure and the controller-driven fabric environment. The network underlay is the routed physical IP infrastructure formed by the fabric nodes and intermediate devices. It provides the reliable, end-to-end IP transport required for control-plane and data-plane reachability across the campus. This foundation lets fabric devices establish the connectivity needed to participate in the SD-Access domain and communicate with controller-managed services.

The fabric overlay is built on top of that IP underlay and creates the logical SD-Access fabric. It uses encapsulation to carry endpoint traffic across the fabric while separating virtual networks and maintaining scalable segmentation. Through the overlay, the controller layer can apply intent, automate fabric behavior, and coordinate endpoint reachability and policy across the physical infrastructure without requiring the underlying topology to be redesigned for every logical network requirement. Cisco describes SD-Access as an overlay architecture that relies on an IP-based underlay for transport; both layers are therefore essential to bridging conventional network elements with controller-managed fabric operations. See Cisco's SD-Access Design Guide and Cisco SD-Access overview.

QUESTION NO: 122

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:

symmetric server passwords are supported through NTP authentication based on shared symmetric keys. On Cisco IOS and IOS XE, an administrator defines an NTP authentication key and key number, marks the key as trusted, and associates that key with an NTP server or peer. NTP packets include an authentication digest derived from the shared secret, enabling the receiving device to validate that time information came from a configured, trusted NTP association and has not been accepted anonymously. This is an authentication and integrity control for the NTP exchange.

An **access list-based restriction scheme** is also available through NTP access-group configuration. Cisco devices can apply IP access lists to NTP functions such as peer relationships, serving time, and responding to NTP queries. These controls limit which source addresses may establish or use NTP-related services, reducing the set of systems that can participate in time synchronization or obtain NTP information from the device. Combining shared-key authentication with access-group restrictions provides both message-origin validation and policy-based control over which hosts can interact with the NTP service. Cisco documents NTP authentication and access-group operation in the Cisco IOS XE Network Time Protocol Configuration Guide; the protocol's symmetric-key authentication model is also specified in RFC 5905.

QUESTION NO: 123

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

A. Layer 3

B. inter-xTR

C. auto anchor

D. fast roam

ANSWER: B

Explanation:

inter-xTR is the roaming method used when a wireless client moves between access points connected to different SD-Access fabric edge switches while using one wireless LAN controller. In an SD-Access fabric, each fabric edge node performs the xTR (Ingress/Egress Tunnel Router) role. When the client roams to an AP attached to a different fabric edge node, the client's point of attachment changes from one xTR to another, making this an inter-xTR roam. The wireless controller remains responsible for WLAN control-plane functions and client mobility coordination, while the fabric provides the data-plane forwarding needed to maintain the client's connectivity and policy context across the fabric nodes. This model enables users to move through the fabric without requiring a traditional wireless-controller-to-wireless-controller mobility event, because the same WLC manages both APs. Cisco's SD-Access wireless design guidance describes roaming in terms of whether the client remains on the same fabric edge/xTR or moves to another fabric edge/xTR; movement between distinct fabric edge nodes is specifically an inter-xTR event. See Cisco's SD-Access Wireless Design Guide and the Cisco Catalyst 9800 client roaming documentation.

QUESTION NO: 124 - (DRAG DROP)

DRAG DROP

Drag and drop the virtual components from the left onto their descriptions on the right.

Select and Place:

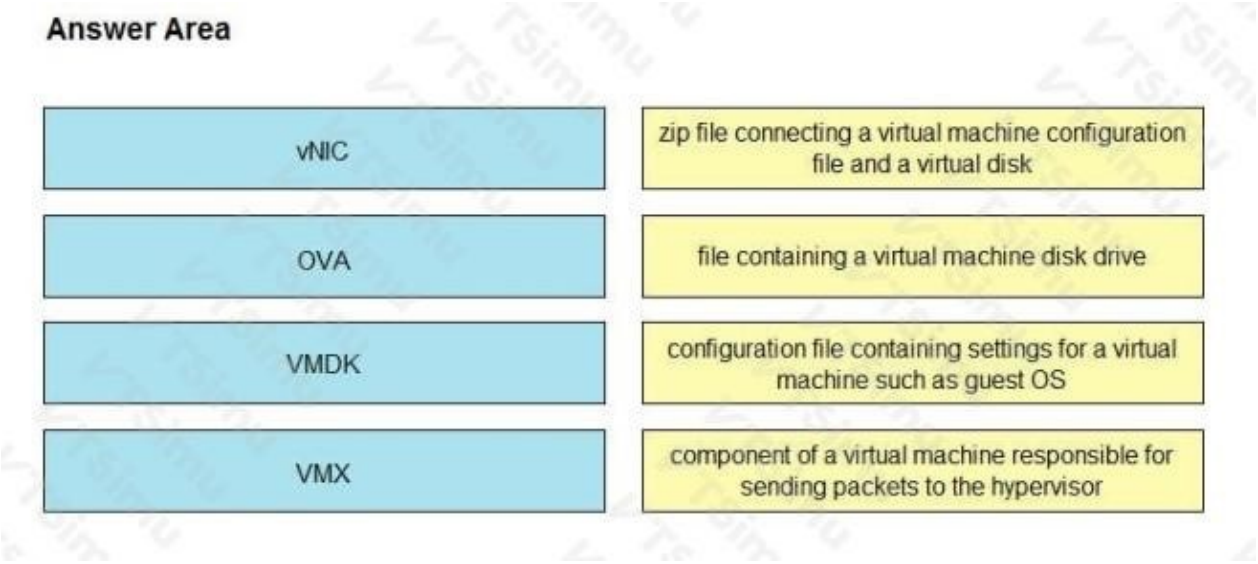

ANSWER:

Answer Area

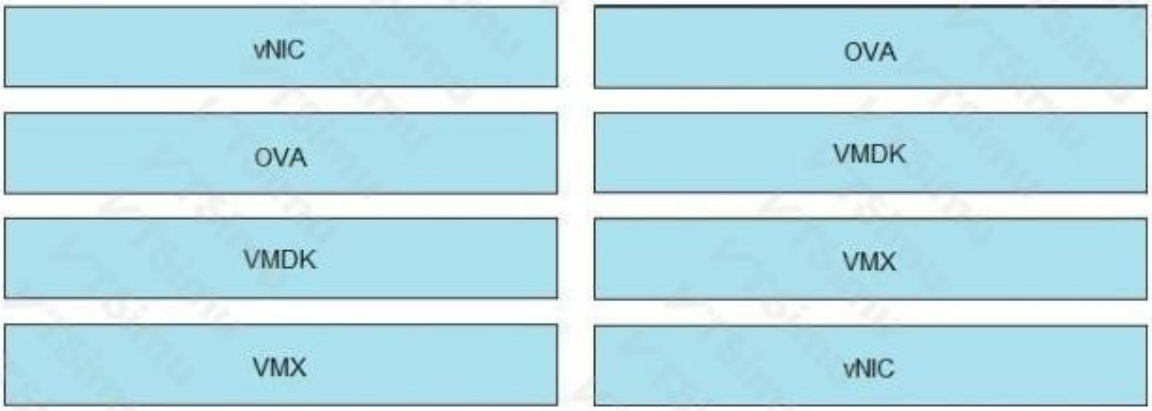

Explanation:

An OVA is the portable, single-file form of an Open Virtualization Format appliance. It packages the virtual machine definition together with its related virtual disks and metadata so that the appliance can be exported, transferred, and deployed as one unit. This makes it the component associated with the archive description that connects virtual-machine configuration information and a virtual disk. VMware's OVF tooling documentation describes OVF deployment and OVA package handling in the vSphere Automation API OVF documentation.

A VMDK is VMware's virtual disk format. It represents the guest machine's virtual hard drive and holds the disk data used by that VM. A VMX file is the primary VMware virtual-machine configuration file. It records the virtual hardware and runtime settings for the guest, including values such as the guest operating system type, memory, processor, storage, and virtual-device settings. VMware documents the role of virtual machine configuration and disk files in its virtual machine files documentation.

A vNIC is the network adapter emulated or presented to a virtual machine. The guest uses this adapter as its network interface, and frames sent through it are handed to the hypervisor's virtual networking layer for switching or forwarding. Thus, the correct arrangement maps the packaged appliance, virtual disk, configuration file, and virtual network interface to the descriptions of their respective functions.

QUESTION NO: 125

An engineer measures the Wi-Fi coverage at a customer site. The RSSI values are recorded as follows:

•LocationA -72dBm

• Location B:-75 dBm

• Location C; -65 dBm

• Location D -80 dBm

Which two statements does the engineer use to explain these values to the customer? (Choose two.)

A. The signal strength at location B is 10 dB better than location C.

B. The signal strength at location C is too weak to support web surfing

C. Location D has the strongest RF signal strength.

D. The RF signal strength at location B is 50% weaker than location A.

E. The RF signal strength at location C is 10 times stronger than location B

ANSWER: D E

Explanation:

RSSI is expressed in dBm, which is a logarithmic measure of received RF power. Consequently, differences in dB represent power ratios rather than linear percentage changes. The statement “The RF signal strength at location B is 50% weaker than location A.” is correct because location B measures −75 dBm and location A measures −72 dBm. This is a 3 dB decrease; a 3 dB reduction corresponds to approximately one-half the received power, since $10^{-3/10}$ is about 0.5.

The statement “The RF signal strength at location C is 10 times stronger than location B” is also correct. Location C is measured at −65 dBm, while location B is −75 dBm, creating a 10 dB difference. A 10 dB increase corresponds to a tenfold increase in received RF power, calculated as $10^{10/10} = 10$. The less-negative value at location C therefore represents substantially greater received signal power than the reading at location B. Cisco wireless design guidance uses RSSI measurements and RF thresholds when assessing WLAN coverage and capacity. See the Cisco Wireless LAN Design Guide and Cisco’s Radio Resource Management white paper for WLAN RF design context.

QUESTION NO: 126

Refer to the exhibit.

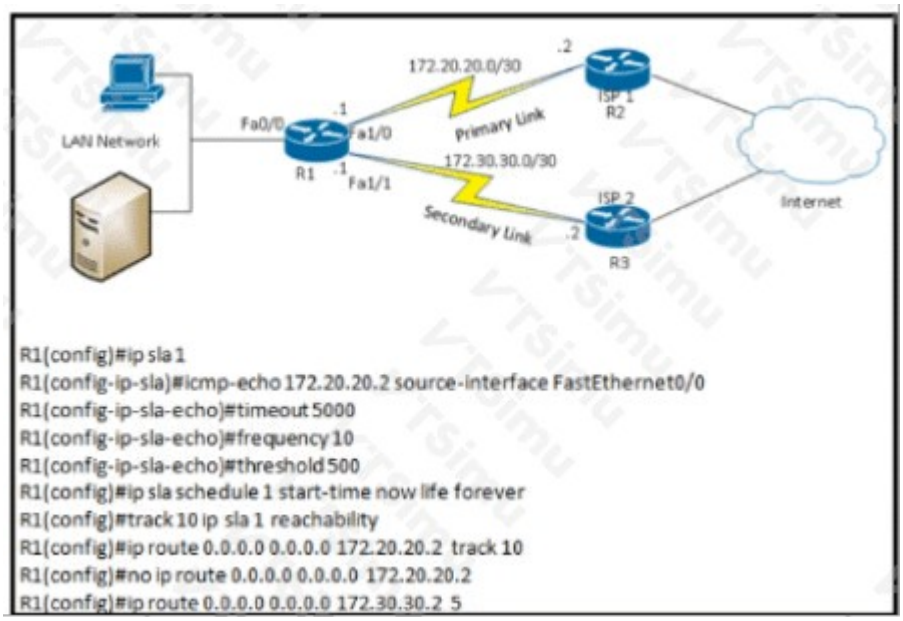

What are two reasons for IP SLA tracking failure? (Choose two)

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

Explanation:

A route back to the R1 LAN network is missing in R2. is a valid cause because an ICMP echo operation requires bidirectional IP connectivity. When R1 sends a probe with an address from its LAN as the source, R2 must be able to route the echo reply back to that source network. Without that return route, R1 does not receive a reply before the operation timeout expires. The IP SLA operation is then unsuccessful, and a track object configured for IP SLA reachability changes state accordingly.

The source-interface is configured incorrectly. is also correct because the source interface controls the source address placed in the probe packet. For a probe intended to test reachability across the R1-to-R2 transit connection, the source

should be an interface/address for which R2 has a valid return path, normally the R1 interface on that transit segment. Selecting the LAN-facing interface can cause replies to be sent toward an unreachable or unrouted source network, preventing successful completion of the ICMP echo exchange. Cisco documents that IP SLA ICMP echo supports specifying a source interface and that tracking evaluates the reachability result of the associated operation. See the Cisco IOS XE IP SLA Configuration Guide and Cisco's IP SLA overview.

QUESTION NO: 127

What is a client considered when it is in web authentication state and roams between two controllers with mobility tunnels?

A. anchor

B. mobile

C. foreign

D. new

ANSWER: C

Explanation:

A client in the web-authentication state that roams between controllers connected by mobility tunnels is considered **foreign** on the controller to which it has roamed. In Cisco wireless mobility architecture, the controller currently serving the client through the associated access point acts in the foreign role, while mobility tunneling preserves the client's session handling across controllers. This distinction is especially important for web authentication because the client has not yet completed the authentication process; the mobility system must maintain the appropriate client state and direct traffic through the established intercontroller mobility relationship. The foreign controller provides the client's local radio connectivity after the roam and exchanges the necessary mobility information with the peer controller so that connectivity and policy processing continue correctly. Cisco's wireless-controller documentation describes mobility tunnels as the mechanism used for client roaming between controllers and identifies the foreign controller as the controller serving the roaming client. See Cisco's Mobility Groups configuration guide and the Cisco Wireless Controller mobility command reference.

QUESTION NO: 128

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 E

Explanation:

Use SSL for encryption. is correct because REST APIs should protect every request and response in transit with HTTPS, which uses TLS (historically referred to as SSL). Transport encryption protects credentials, session tokens, API keys, and application data from interception or modification while traversing untrusted networks. API clients should validate the server certificate and implementations should use current TLS versions and secure cipher suites. Cisco describes SSL/TLS as providing encryption and authentication for data exchanged across a network: Cisco: What Is SSL/TLS?.

Use a password hash. is also correct because an API or its supporting identity service must never retain user passwords in plaintext or reversibly encrypted form. Passwords should be transformed using a purpose-built, salted, adaptive password-hashing function, such as Argon2id, bcrypt, PBKDF2, or scrypt. This limits credential exposure if the authentication datastore is compromised and makes offline password-cracking attempts substantially more costly. The OWASP guidance on secure

password storage explains the use of modern password-hashing algorithms, salts, and appropriate work factors: OWASP Password Storage Cheat Sheet.

QUESTION NO: 129

Refer to the exhibit.

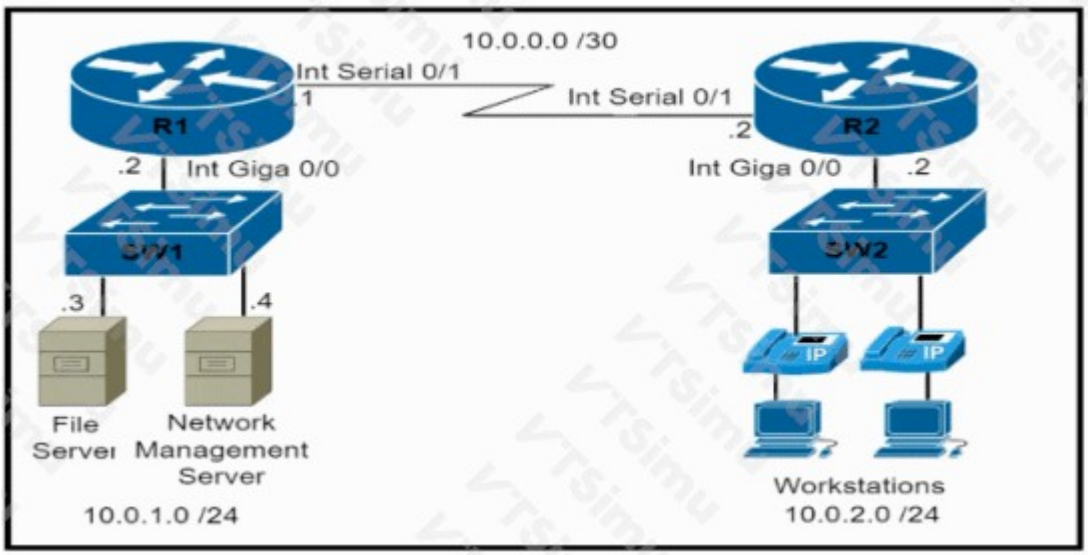

An engineer must configure and validate a CoPP policy that allows the network management server to monitor router R1 via SNMP while protecting the control plane. Which two commands or command sets must be used? (Choose two.)

- [x] 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

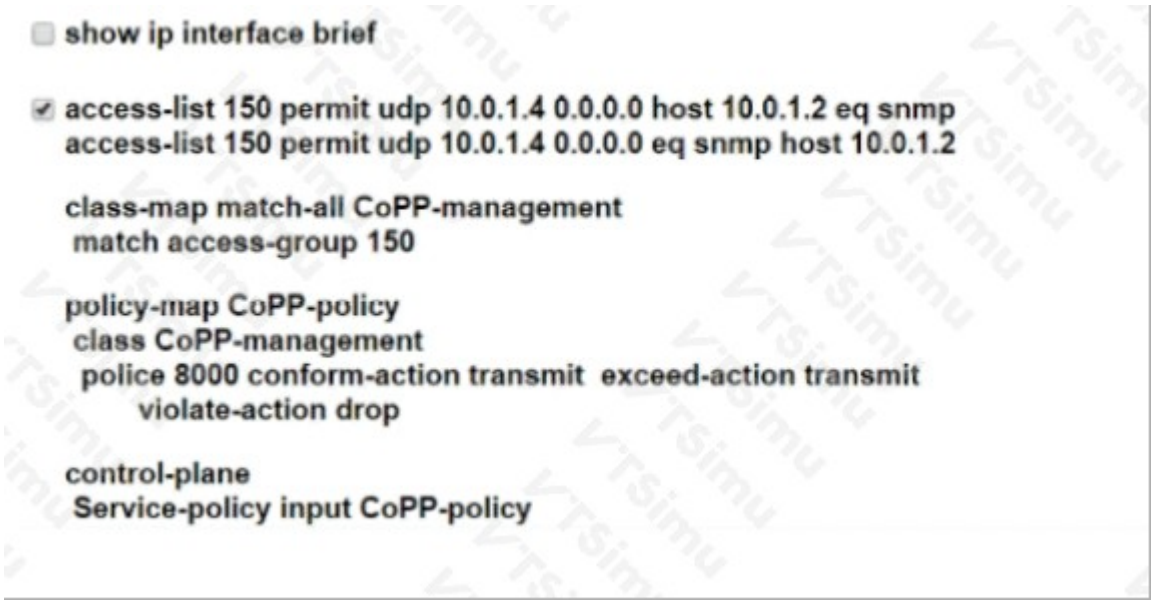

A. Option A

B. Option B

C. Option C

D. Option D

E. Option E

F. Option F

ANSWER: A F

Explanation:

Option A and **Option F** together provide the required CoPP configuration and verification workflow for SNMP access to R1's control plane. CoPP operates on traffic that is destined to the router itself, including SNMP queries sent by a network-management server. The policy must first identify the intended SNMP traffic, normally by matching the management server source and UDP port 161 in an access control list, then associating that match with a class map. A policy map can apply an appropriate policing rate while allowing conforming SNMP packets to reach the CPU. The policy becomes active only when it is applied in the control-plane configuration context with an input service policy. This preserves management reachability while limiting the amount of SNMP traffic that can consume control-plane resources. Validation must inspect the active control-plane policy and its class counters. This confirms that SNMP packets are being classified into the intended class and shows whether packets conform to, exceed, or violate the configured policing parameters. Cisco describes CoPP classification, policy-map actions, and control-plane attachment in its Control Plane Policing configuration guide. The relevant IOS XE QoS and policing behavior is also documented in the Cisco IOS XE QoS Policing Configuration Guide.

QUESTION NO: 130

Refer to the exhibit. Why was the response code generated?

A. The resource was unreachable.

B. Access was denied based on the user permissions.

C. Access was denied based on the credentials.

D. The resource is no longer available on the server.

ANSWER: B

Explanation:

Access was denied based on the user permissions. The exhibited HTTP 403 Forbidden response means that the server received and understood the request but refuses to authorize it. This status is used when the target resource is known to the server, but the authenticated client does not have permission to perform the requested action. In a Cisco API or controller environment, this commonly occurs when a user has successfully logged in but lacks the required role-based access control privilege for a particular REST endpoint or HTTP method. For example, an account might be permitted to view operational data but not change configuration, or it may be restricted from accessing a protected resource entirely. The response therefore indicates an authorization-policy decision, such as an ACL, assigned role, privilege level, or endpoint-specific permission, rather than a connectivity issue. RFC 9110 defines 403 as a refusal to fulfill a request when the server understands the request but will not authorize it. See RFC 9110, HTTP 403 Forbidden and MDN Web Docs: 403 Forbidden.

QUESTION NO: 131

Refer to the exhibit.

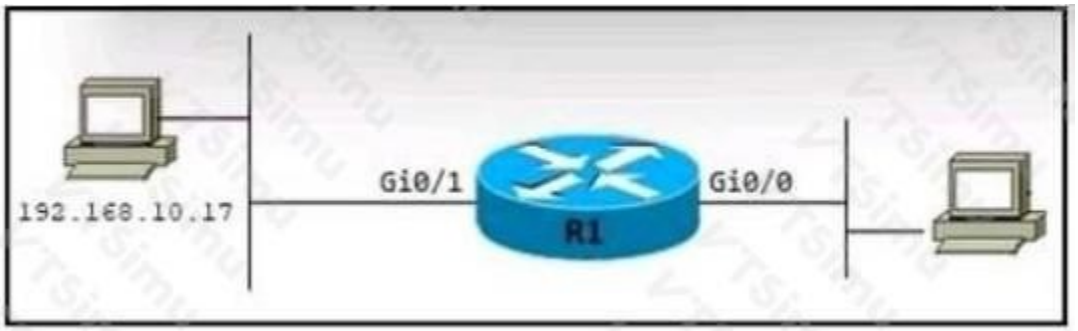

An engineer applies this configuration to R1:

ip nat inside source static 192.168.10.17 192.168.27.42

Which command set should be added to complete the configuration? A)

```
R1(config)# interface GigabitEthernet 0/0
R1(config)# ip pat inside

R1(config)# interface GigabitEthernet 0/1
R1(config)# ip pat outside
```

B)

```
R1(config)# interface GigabitEthernet 0/0
R1(config-if)# ip nat outside

R1(config)# interface GigabitEthernet 0/1
R1(config-if)# ip nat inside
```

C)

```
R1(config)# interface GigabitEthernet 0/0
R1(config)# ip pat outside

R1(config)# interface GigabitEthernet 0/1
R1(config)# ip pat inside
```

D)

```
R1(config)# interface GigabitEthernet 0/0
R1(config-if)# ip nat inside

R1(config)# interface GigabitEthernet 0/1
R1(config-if)# ip nat outside
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B

Explanation:

Option B is correct because a Cisco IOS static NAT statement maps an inside-local address, `192.168.10.17`, to an inside-global address, `192.168.27.42`, but the router also requires NAT roles on the relevant interfaces. The interface connected toward the private network containing `192.168.10.17` must be configured with `ip nat inside`. The interface toward the network where the translated address is reachable must be configured with `ip nat outside`. These interface designations establish the NAT translation boundary and allow R1 to apply the configured static mapping as packets traverse between the inside and outside networks.

Static NAT is bidirectional: packets sourced by the internal host can be translated from the inside-local address to the inside-global address, and traffic sent to the inside-global address can be translated back to the internal host. The translated address must also be routable from the outside network, either as an address on an appropriately connected subnet or through suitable routing. Cisco's NAT configuration guidance describes the required use of `ip nat inside` and `ip nat outside` interface commands alongside static mappings. See Cisco NAT Configuration Examples and Cisco IOS NAT Configuration Guide.

QUESTION NO: 132

How is a data modeling language used?

A. To enable data lo be easily structured, grouped, validated, and replicated

B. To represent finite and well-defined network elements that cannot be changed

C. To model the flows of unstructured data within the infrastructure

D. To provide human readability to scripting languages

ANSWER: A

Explanation:

To enable data lo be easily structured, grouped, validated, and replicated is correct. In Cisco's model-driven programmability architecture, a data modeling language such as YANG defines a formal schema for device configuration and operational state. The schema organizes information into a hierarchy of reusable containers, lists, leaves, and data types, allowing software to work with network data in a consistent, machine-readable form. It also supplies constraints, including required values, value ranges, patterns, and relationships between data elements. These rules allow a server to validate configuration data before applying it, which improves reliability in automated workflows.

Models provide a common structure that clients can use through interfaces such as NETCONF and RESTCONF, rather than requiring automation tools to interpret unstructured CLI output. This makes it practical to group related network data, retrieve only the required portions of a configuration or state tree, and exchange or synchronize modeled data among controllers, automation systems, and devices. YANG is specifically designed to model configuration and state data manipulated by network-management protocols, as defined in RFC 7950: The YANG 1.1 Data Modeling Language. Cisco also describes how YANG models support programmatic configuration and telemetry in its model-driven programmability documentation.

QUESTION NO: 133

What is the intent API in Cisco Catalyst Center (formerly DNA Center)?

A. northbound consumer-facing RESTful API, which enables network discovery and configuration management

B. southbound consumer-facing RESTful API. which enables network discovery and configuration management

C. interface between the controller and the network devices, which enables network discovery and configuration management

D. westbound interface, which allows the exchange of data such as ITSM. IPAM, and reporting information

ANSWER: A

Explanation:

The correct description is **northbound consumer-facing RESTful API, which enables network discovery and configuration management**. Cisco Catalyst Center Intent APIs are the controller's external, application-facing REST interfaces. They let a client application, orchestration tool, or custom automation workflow communicate with Catalyst Center at an intent and service level rather than requiring that application to manage individual network-device CLI commands or device-specific protocols directly.

Through these APIs, consumers can automate operational tasks such as onboarding and discovering network devices, obtaining inventory and topology information, provisioning network settings, managing configurations, and accessing assurance-related data and workflows. Catalyst Center receives the API request, applies its controller logic and automation capabilities, and then performs the required underlying interactions with the managed infrastructure. This model makes the API northbound: it is exposed to systems and applications that sit above and consume services from the controller.

Cisco's developer documentation organizes these REST resources as Catalyst Center platform and Intent APIs, including areas for discovery, network devices, configuration, topology, and automation. See the Cisco Catalyst Center Developer Center and Cisco's Catalyst Center product overview for the platform's API-driven automation model.

QUESTION NO: 134

What is one benefit of adopting a data modeling language?

A. deploying machine-friendly codes to manage a high number of devices

B. augmenting the use of management protocols like SNMP for status subscriptions

C. augmenting management process using vendor centric actions around models

D. refactoring vendor and platform specific configurations with widely compatible configurations

ANSWER: D

Explanation:

Adopting a data modeling language supports refactoring vendor and platform specific configurations with widely compatible configurations because a model defines data using a structured schema rather than relying on device-specific command-line syntax. In Cisco enterprise automation, YANG is the principal data modeling language used to describe configuration data, operational state, available operations, and notifications. Clients can then interact with that modeled data through programmatic interfaces such as NETCONF and RESTCONF. This lets automation workflows use defined data objects, types, constraints, and hierarchy, making those workflows more consistent and reusable across devices that support the same standard or compatible model. It also improves validation: a system can determine whether data conforms to the model before applying a configuration. Cisco's model-driven programmability approach uses YANG models to provide structured access to IOS XE capabilities, helping teams build automation that is less dependent on parsing command output and maintaining unique CLI templates for every platform. The YANG specification describes this purpose as modeling configuration and state data, RPC operations, and notifications for network management protocols. See RFC 7950: The YANG 1.1 Data Modeling Language and Cisco's IOS XE Model-Driven Programmability documentation.

QUESTION NO: 135

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

A. Connects endpoints to the fabric and forwards their traffic.

B. Encapsulates end-user data traffic into LISP.

C. Connects the SD-Access fabric to another fabric or external Layer 3 networks.

D. Provides reachability between border nodes in the fabric underlay.

ANSWER: A

Explanation:

Connects endpoints to the fabric and forwards their traffic. is correct because a fabric edge node is the SD-Access fabric role that provides the access-facing attachment point for wired endpoints and, through fabric wireless integration, wireless clients. It classifies connected users and devices into the appropriate virtual network and applies group-based policy at the point where their traffic enters or leaves the fabric. The fabric edge then transports that traffic across the SD-Access overlay to its destination. In Cisco SD-Access, fabric edge nodes use VXLAN encapsulation for overlay data-plane forwarding and use LISP-based control-plane information to map endpoint identities and locations. This enables endpoint mobility while maintaining consistent segmentation and policy. Thus, the defining function is not merely generic access switching; it is connecting endpoints to the fabric and forwarding their traffic in accordance with the fabric's virtual-network and policy constructs. Cisco describes fabric edge nodes as the fabric devices to which endpoints connect and as the devices that provide endpoint connectivity into the SD-Access fabric. See Cisco's Software-Defined Access overview and the Cisco SD-Access Design Guide.

QUESTION NO: 136

With IGMPv2, which multicast group address does the IGMP querier use to send query messages to all hosts on the LAN?

A. 239.0.0.2

B. 224.0.0.1

C. 239.0.0.1

D. 224.0.0.2

ANSWER: B

Explanation:

224.0.0.1 is the IPv4 link-local multicast address reserved for all hosts on the local subnet. In IGMPv2, an IGMP querier sends a General Membership Query to this address to ask every multicast-capable host on the LAN to report the multicast groups it still wishes to receive. This lets the querier maintain current multicast membership information and enables the network to forward multicast traffic only where active receivers exist. The query has a TTL of 1, so it remains on the local LAN segment and is not forwarded by routers, which is appropriate because IGMP manages host membership between directly connected hosts and their local multicast router or Layer 3 switch interface. RFC 2236 specifies that General Queries are addressed to the all-systems multicast group, 224.0.0.1. The address is also registered by IANA as the “All Hosts” group in the local network control block. See RFC 2236, section 4 and the IANA IPv4 Multicast Address Space Registry.

QUESTION NO: 137

Which two new security capabilities are introduced by using a next-generation firewall at the Internet edge? (Choose two.)

A. DVPN

B. NAT

C. stateful packet inspection

D. application-level inspection

E. integrated intrusion prevention

ANSWER: D E

Explanation:

Next-generation firewalls add context-aware security controls and active threat prevention capabilities beyond the traditional stateful firewall role at an Internet edge. **Application-level inspection** provides visibility into the actual application generating a flow, rather than relying only on IP addresses, port numbers, and protocols. This enables an organization to construct and enforce access policies based on application identity and associated attributes, which is important because many applications can use common ports or dynamically selected ports. **Integrated intrusion prevention** enables the firewall to inspect permitted traffic for exploits, known attack signatures, and malicious activity, then prevent identified threats inline before they reach protected internal resources. Combining application awareness with intrusion prevention allows the edge device to make more informed policy decisions while also detecting and blocking threats carried within otherwise allowed sessions. Cisco identifies application visibility and control, along with integrated intrusion prevention, as core next-generation firewall capabilities. See Cisco’s Next-Generation Firewall overview and its firewall product and technology overview.

QUESTION NO: 138

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.)

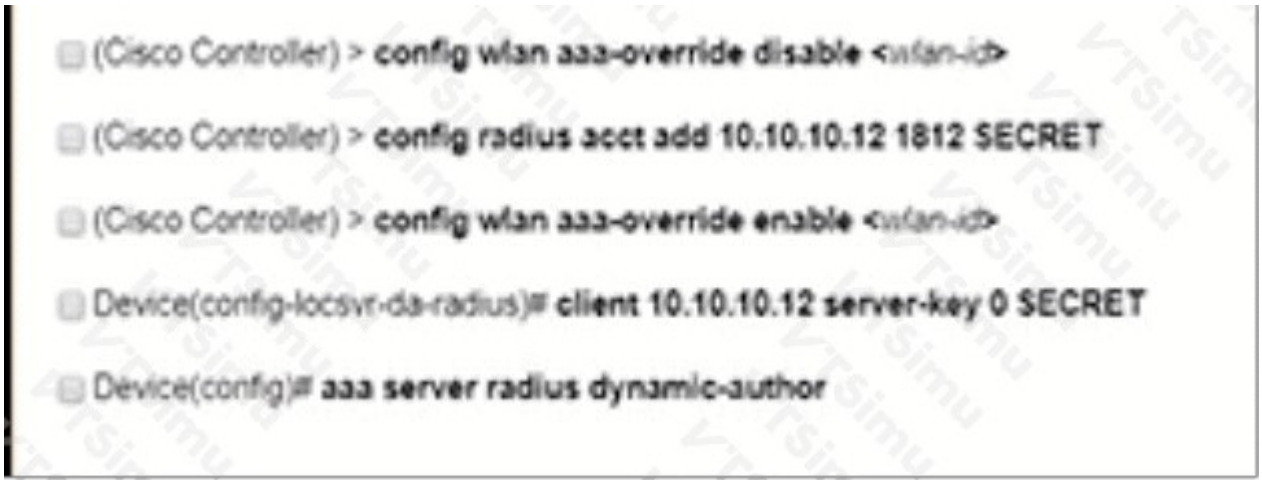

A. Option A

B. Option B

C. Option C

D. Option D

E. Option E

F. aaa authentication login CWA_AUTH group ISE

G. aaa authorization network CWA_AUTH group ISE

ANSWER: F G

Explanation:

Central Web Authentication requires the Catalyst 9800 controller to use an AAA method list for the web-login authentication transaction and another AAA method list for network authorization. The command `aaa authentication login CWA_AUTH group ISE` creates a named login authentication method list that sends the client's portal credentials to the configured RADIUS server group. This lets the external identity service validate the user during the central web portal flow.

The command `aaa authorization network CWA_AUTH group ISE` creates the corresponding network authorization method list. After authentication, the controller obtains the RADIUS authorization result, including policy attributes needed to determine the client's permitted network access. The same method-list name is intentionally used so the wireless web-authentication policy can consistently invoke the matching authentication and authorization services. `ISE` is the configured RADIUS server-group name and can be replaced with the actual server-group name used in the deployment. These AAA constructs must subsequently be associated with the applicable wireless policy/profile and RADIUS server configuration. Cisco documents IOS XE RADIUS server groups and AAA method lists in its IOS XE AAA and RADIUS Configuration Guide; Catalyst 9800 deployment guidance is available in the Catalyst 9800 configuration guides.

QUESTION NO: 139

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

A. Cisco Discovery Protocol neighbour

B. broadcasting on the local subnet

C. DNS lookup cisco-DNA-PRIMARY.localdomain

D. DHCP Option 43

E. querying other APs

Answer: BD

Explanation:

A. Cisco Discovery Protocol neighbour

B. broadcasting on the local subnet

C. DNS lookup cisco-DNA-PRIMARY.localdomain

D. DHCP Option 43

ANSWER: B D

Explanation:

Cisco lightweight access points use CAPWAP discovery mechanisms to locate a wireless LAN controller before joining it. **broadcasting on the local subnet** is a valid discovery method when the access point and controller are reachable in the

same broadcast domain. The access point sends CAPWAP discovery messages as a local broadcast, and an available controller on that subnet can respond with discovery-response information. This method is particularly applicable to deployments in which the AP management network and controller management interface share the same Layer 2 segment.

DHCP Option 43 is also a supported controller-discovery method. A DHCP server can return the management IP address or addresses of Cisco wireless LAN controllers in the vendor-specific Option 43 field. Once the AP receives its addressing information through DHCP, it parses this option and sends CAPWAP discovery requests directly to the supplied controller addresses. This enables APs to discover controllers across routed networks, where local-subnet broadcasts cannot reach the controller. Cisco documents both local broadcast discovery and DHCP Option 43 as CAPWAP AP discovery methods. See Cisco's AP Discovery and Join Process documentation and Cisco's DHCP Option 43 configuration guidance.

QUESTION NO: 140

```
username cisco privilege 15 noescape secret 5 F7u$9cyE438490035m8TQ$nv&6502x
username cisco autocommand show startup-config
aaa authentication login default local-case enable
aaa authorization exec default local
```

An engineer applies this configuration to router R1. How does R1 respond when the user 'cisco' logs in?

A. It displays the startup conftg and then permits the user to execute commands

B. It places the user into EXEC mode and permits the user to execute any command

C. It displays the startup conflg and then terminates the session.

D. It places the user into EXEC mode but permits the user to execute only the show startup-config command

ANSWER: C

Explanation:

It displays the startup config and then terminates the session. The configuration uses the local username's `autocommand` capability to associate `show startup-config` with the user `cisco`. After successful local authentication, IOS automatically runs that specified EXEC command instead of providing an interactive CLI session. This is useful for tightly constrained operational access, such as allowing a user or an automated process to retrieve one item of information without receiving a normal router prompt or broader command access.

After the automatic command completes, the login session is closed. This is the default behavior for an autocommand. Cisco IOS supports a `nohangup` keyword in contexts where an administrator needs to prevent automatic disconnect behavior, but that behavior must be explicitly configured. Because the configured automatic command is `show startup-config`, R1 outputs the saved startup configuration and then ends the connection. Cisco documents the `username ... autocommand` function in the Cisco IOS Security Command Reference; related local authentication and EXEC-access configuration is covered in the Cisco IOS AAA Configuration Guide.

QUESTION NO: 141

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:

An encrypted authentication mechanism and an access list-based restriction scheme are the two NTP security features available on Cisco IOS and IOS XE devices. NTP authentication uses configured trusted keys so that an NTP client can validate that time information originates from an authorized NTP source. The shared key is used to generate and verify an authentication digest on NTP messages, protecting the integrity and authenticity of synchronization exchanges. Cisco configurations commonly define an NTP authentication key, identify it as trusted, enable NTP authentication, and associate the key with an NTP server or peer.

An access list-based restriction scheme is implemented with `ntp access-group` controls. Cisco supports distinct access-group roles, including peer, serve, serve-only, and query-only, which let an administrator define which source addresses may form NTP peer relationships, receive time service, or issue NTP control queries. Applying these restrictions reduces the device's exposure to unauthorized NTP clients and control traffic while allowing intended synchronization relationships. Used together, authenticated time sources and NTP access controls provide both trustworthy time updates and controlled access to the NTP service. See Cisco's NTP Configuration and Troubleshooting and the Cisco IOS XE NTP Configuration Guide.

QUESTION NO: 142

Refer to the Exhibit.

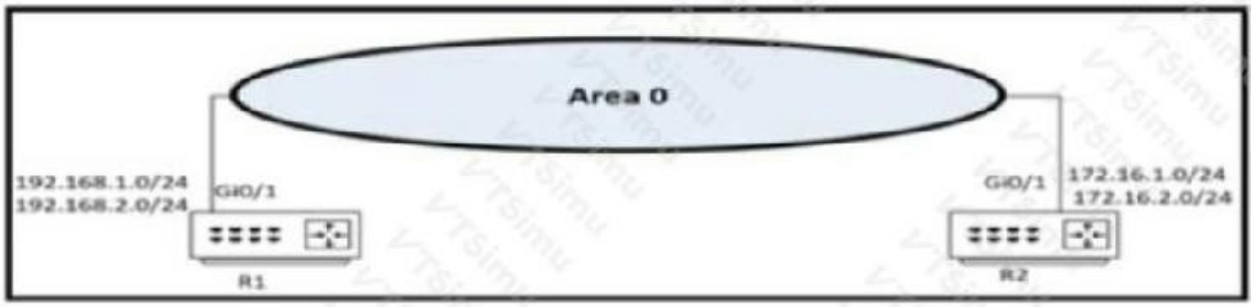

Refer to the exhibit. Which two configurations enable R1 and R2 to advertise routes into OSPF? (Choose two)

A)

```
R2
router ospf 0
network 172.16.1.0  255.255.255.0 area 0
network 172.16.2.0  255.255.255.0 area 0
```

B)

```
R2
router ospf 0
network 172.16.1.0  0.0.0.255 area 0
network 172.16.2.0  255.255.255.0 area 0
```

C)

```
R1
router ospf 0
network 192.168.1.0 0.0.0.255 area 0
network 192.168.2.0 0.0.0.255 area 0
```

D)

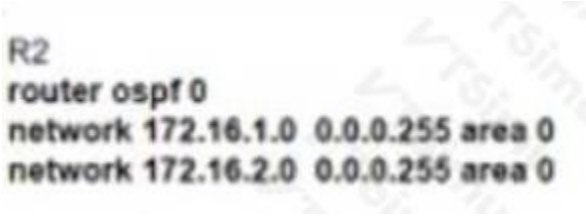

E)

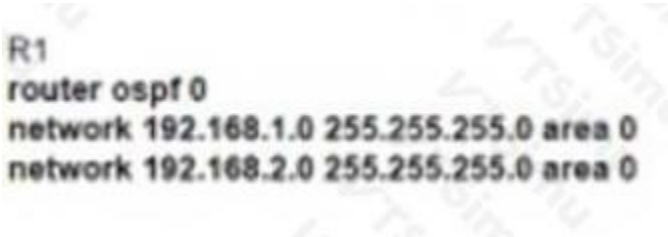

A. Option A

B. Option B

C. Option C

D. Option D
E) Option E

ANSWER: A D

Explanation:

The two indicated configurations are correct because each enables OSPF for the applicable routed interfaces and associated connected IPv4 networks on both routers. Cisco IOS can activate OSPF either with `network` statements under the relevant `router ospf` process or directly under an interface using `ip ospf process-id area area-id`. A network statement does not advertise an arbitrary prefix merely because it is typed; it matches an interface address using its wildcard mask, enables OSPF on that matching interface, and advertises the connected network in the specified area. Interface-based activation accomplishes the same result more explicitly for that interface.

Once OSPF is enabled on the transit-facing interfaces in the same area and the interfaces are operational, R1 and R2 can establish an adjacency. OSPF then originates LSAs describing their participating connected networks, allowing those routes to be advertised throughout the OSPF domain. The OSPF process identifier is locally significant, so matching process numbers are not required between routers; however, the common link must be assigned to a compatible OSPF area. Cisco documents both network-based and interface-based OSPF activation methods in its OSPF configuration documentation and IOS XE OSPF configuration guide.

QUESTION NO: 143

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 S deny 172.16.0.0/16. Winch 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 Ie 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

F. ip prefix-list OFFICE seq 5 deny 177.16.0.0/16

ANSWER: A B F

Explanation:

To prevent the stated `177.16.0.0/16` OSPF route from being installed in the local routing table, the prefix list must explicitly deny `177.16.0.0/16`, and the prefix list must be applied inbound to the OSPF process. The existing deny entry shown in the question is for `172.16.0.0/16`, which does not match the stated route; therefore, an additional deny entry for the stated prefix is necessary. Applying `distribute-list prefix OFFICE in` under the OSPF process causes the router to evaluate received OSPF routes against the named prefix list before installing them in its routing table. This is a local route-installation filter and does not alter OSPF LSAs or the topology database.

A prefix list also requires an explicit permit for the remaining routes. Prefix lists have an implicit deny at the end, so `ip prefix-list OFFICE seq 10 permit 0.0.0.0/0 le 32` permits every IPv4 prefix length from 0 through 32 after the targeted prefix has been denied. Together, the explicit denial of `177.16.0.0/16`, the catch-all permit, and inbound OSPF distribute-list application block only the specified route while retaining all other eligible OSPF routes. Cisco documents this OSPF filtering behavior at Cisco OSPF filtering overview and prefix-list matching at Cisco IOS Prefix Lists.

QUESTION NO: 144

A wireless network engineer must configure a WPA2+WPA3 policy with the Personal security type. Which action meets this requirement?

A. Configure the GCMP256 encryption cipher.

B. Configure the CCMP256 encryption cipher.

C. Configure the CCMP128 encryption cipher.

D. Configure the GCMP128 encryption cipher.

ANSWER: C

Explanation:

Configuring the CCMP128 encryption cipher meets the requirement for a WPA2+WPA3 Personal policy, commonly called WPA3-Personal transition mode. This mode is designed to permit both WPA2-Personal clients and WPA3-Personal clients to associate with the same WLAN during a migration. WPA2-Personal uses a pre-shared key and AES-CCMP protection, while WPA3-Personal introduces Simultaneous Authentication of Equals (SAE) for stronger password-based authentication. CCMP with a 128-bit key is the interoperable AES data-protection cipher used for this mixed-client deployment, allowing WPA2 clients to continue using CCMP while WPA3-capable clients use SAE authentication on the same WLAN. Cisco wireless configuration guidance identifies AES-CCMP as the applicable encryption method for WPA2/WPA3 Personal transition deployments. The transition policy therefore combines WPA2 PSK compatibility and WPA3 SAE capability without requiring a high-assurance 256-bit cipher suite. Cisco's WPA3 documentation describes WPA3-Personal and its SAE authentication method, while the Wi-Fi Alliance explains WPA3's transition support for WPA2 devices. See Cisco Wireless Controller WPA3 Configuration Guide and Wi-Fi Alliance Wi-Fi Security overview.

QUESTION NO: 145

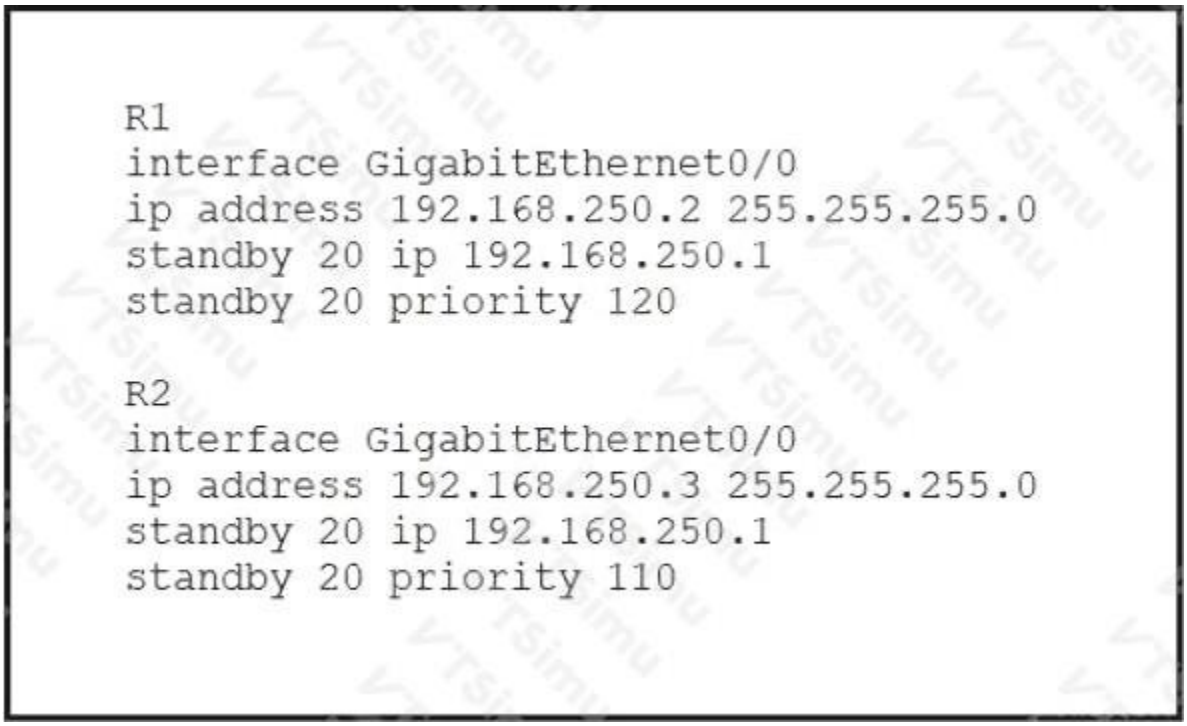

Refer to the exhibit. What are two effects of this configuration? (Choose two.)

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:

R1 becomes the active router because HSRP elects the router with the highest configured HSRP priority as active. The configuration shown assigns R1 the higher effective priority, so it assumes the active forwarding role for the HSRP group, while the peer operates as the standby router. HSRP supplies a shared virtual default-gateway address, allowing hosts to continue forwarding through whichever router currently owns the active role.

If R1 goes down, R2 becomes active and remains the active device when R1 comes back online because HSRP preemption is not enabled. When the active router fails, the standby router transitions to active and begins forwarding packets sent to the HSRP virtual IP and virtual MAC address. Although R1 returns with its higher priority, HSRP does not automatically replace an already active router by default. R2 therefore keeps the active role until it fails, is reset, or another relevant HSRP state change occurs. Cisco documents that preemption must be explicitly configured when a higher-priority router should reclaim the active role after recovery.

References: Cisco: Hot Standby Router Protocol Features and Functionality; Cisco IOS XE HSRP Configuration Guide.

QUESTION NO: 146

Which two statements are true about NETCONF? (Choose two.)

A. It uses XML to encode protocol operations and data.

B. It commonly uses SSH over TCP port 830.

C. It uses JSON-RPC over UDP port 161.

D. It requires HTTP methods such as GET and POST for configuration changes.

E. It can manage only operational state and cannot change configuration.

ANSWER: A B

Explanation:

NETCONF uses a structured remote-procedure-call model to retrieve and modify configuration and operational data. Network management operations are encoded in XML, which provides a structured representation of the data defined by a YANG model. Common operations include retrieving configuration, editing configuration, locking a datastore, and committing configuration changes.

NETCONF commonly runs over SSH, using TCP port 830 as the well-known port. SSH provides secure transport, server authentication, and encryption for the management session. NETCONF is therefore widely used in model-driven automation environments where a client needs structured, transactional interaction with a network device. RFC 6241 defines the NETCONF protocol, and Cisco documents NETCONF support in the Cisco IOS XE Programmability Configuration Guide.

QUESTION NO: 147

```
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 B

Explanation:

The commands `Router(config-vrf)#address-family ipv4` and `Router(config-if)#ip address 192.168.1.1 255.255.255.0` complete IPv4 operation for an interface associated with the configured VRF. The IPv4 address family is configured from VRF-definition configuration mode. This enables the VRF to contain IPv4 routes and connected IPv4 interface information. Once the IPv4 address family exists for that VRF, the interface can receive its IPv4 address in interface configuration mode using the standard `ip address` syntax. The dotted-decimal mask `255.255.255.0` is equivalent to the required `/24` prefix length, so the resulting connected network is 192.168.1.0/24 and the interface address is 192.168.1.1. Cisco IOS XE separates VRF address-family configuration from physical interface

configuration: VRF address-family commands are entered under `vrf definition`, while an interface's IPv4 address is entered under the interface itself. Cisco documents this VRF-definition and address-family workflow in its IOS XE VRF Configuration Guide; standard interface IPv4 addressing syntax is also shown in Cisco's IP Addressing and Subnetting documentation.

QUESTION NO: 148

Which two security mechanisms aie used by Cisco Threat Defense to gain visibility into the most dangerous cyber threats? (Choose two.)

A. Traffic Telemetry

B. VLAN segmentation

C. virtual private networks

D. dynamic enforce policy

E. file reputation

ANSWER: A E

Explanation:

Traffic Telemetry and file reputation provide complementary visibility capabilities in Cisco Firepower Threat Defense. Traffic Telemetry gives the firewall context about communications traversing the network, including observed connections, applications, users, hosts, and traffic characteristics. This context supports identification and investigation of suspicious activity by making network behavior visible to the security platform and its eventing, analysis, and threat-intelligence functions. Cisco describes Secure Firewall as providing comprehensive visibility into network traffic and threats, which is a core purpose of its next-generation firewall inspection capabilities.

File reputation adds intelligence about files detected during inspected transfers. Through file-control and malware-inspection workflows, Firepower Threat Defense can use cloud-delivered threat intelligence to determine a file's disposition and assess whether it is associated with malicious activity. Reputation-based assessment is valuable because it connects an observed file to continuously updated threat knowledge, enabling security teams to identify high-risk malware that might otherwise appear as ordinary encrypted or application-layer traffic. Used together, Traffic Telemetry supplies the network and behavioral context, while file reputation supplies file-specific threat context. Cisco documents these inspection and file-policy capabilities in its Secure Firewall NGFW overview and its file-control configuration guide.

QUESTION NO: 149

Which two statements describe PIM sparse mode? (Choose two.)

A. It uses an explicit join model for multicast receivers.

B. It can use a rendezvous point for source and receiver discovery.

C. It floods multicast traffic throughout the PIM domain before pruning unwanted branches.

D. It requires IGMP snooping to elect the designated router.

E. It uses only source-specific trees and never uses a shared tree.

ANSWER: A B

Explanation:

PIM sparse mode uses an explicit join model. Routers do not forward multicast traffic toward a receiver-facing interface unless downstream routers have sent PIM Join messages indicating that receivers want to receive traffic for the multicast group. This behavior makes sparse mode suitable when multicast receivers are distributed sparsely across the network.

PIM sparse mode can also use a rendezvous point as a meeting location for multicast sources and receivers. Sources initially register traffic with the RP, and receivers join the shared tree rooted at the RP. Routers can subsequently use a shortest-path tree from the source when the network design and configured behavior permit it. Cisco documents the sparse-mode shared-tree and shortest-path-tree operation in the Cisco IOS XE PIM Sparse Mode Configuration Guide.

QUESTION NO: 150

Why does the vBond orchestrator have a public IP?

to enable vBond to team the public IP of WAN Edge devices that are behind NAT gateways or in private address space

A. to facilitate downloading and distribution of operational and security patches

B. to allow for global reachability from all WAN Edges in the Cisco SD-WAN and

C. to facilitate NAT traversal for WAN Edge devices behind NAT gateways or in private address space

D. to Cisco Smart Licensing servers for license enablement

ANSWER: C

Explanation:

To facilitate NAT traversal for WAN Edge devices behind NAT gateways or in private address space is correct because vBond is the Cisco SD-WAN orchestration component that WAN Edge devices contact first when establishing the control plane. A WAN Edge at a branch commonly uses an RFC 1918 address or sits behind a NAT device, so it cannot be directly reached through its local address from the public internet. The WAN Edge can initiate an outbound secure DTLS or TLS connection to the publicly reachable vBond orchestrator. During this exchange, vBond identifies the translated public address and port information associated with the WAN Edge and shares the connectivity information required for the device to establish authenticated control connections with the appropriate controllers. This process enables the SD-WAN overlay to onboard and connect sites even when their WAN interfaces do not have globally routable addresses. Public reachability of vBond is therefore central to initial orchestration and NAT traversal, rather than being a software-distribution or licensing function. Cisco describes vBond as facilitating connections between WAN Edge devices and controllers, including devices located behind NAT. See Cisco's SD-WAN Control Plane documentation and SD-WAN Security Overview.

QUESTION NO: 151

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 correct feature because it controls the receive start-of-packet detection threshold at the access point. In a dense wireless deployment, an AP can hear transmissions from distant clients and neighboring AP cells. Spending airtime and processing resources attempting to receive these weak signals contributes to co-channel contention and reduces efficient spatial reuse. RX-SOP lets the administrator raise the receive threshold so that frames received below the configured RSSI level are not considered valid packets for reception and are not demodulated. The practical result is a smaller effective RF receive cell: the AP focuses on stronger, nearby transmissions while disregarding weaker, more distant ones. This behavior is especially useful where access points are closely spaced and coverage overlap is intentionally high, because it limits unnecessary packet processing and helps optimize capacity. Cisco documents RX-SOP as a high-density

tuning mechanism and provides configurable threshold profiles on supported wireless platforms. See Cisco's RX-SOP configuration guide for Catalyst wireless and the Catalyst 9800 RF parameters configuration guide.

QUESTION NO: 152

A company recently decided to use RESTCONF instead of NETCONF and many of their NETCONF scripts contain the operation

(operation=â€ createâ€).Which RESTCONF operation must be used to replace these statements?

A. POST

B. GET

C. PUT

D. CREATE

ANSWER: A

Explanation:

POST is the appropriate RESTCONF operation for the intended migration from a NETCONF `<edit-config>` request using the `create` operation. RESTCONF uses HTTP methods to apply configuration changes to resources represented by YANG data. A POST request submits a representation of a new resource to a target collection or parent resource, such as creating a new interface list entry or a new routing-protocol instance. This is the resource-creation behavior that corresponds to the NETCONF intent expressed by `operation="create"`.

The client sends the new configuration object as JSON or XML to the applicable RESTCONF data URI. When creation succeeds, the server normally returns an HTTP success response, commonly `201 Created`, and may provide the URI of the newly created resource in the `Location` header. This approach lets scripts move from NETCONF RPC-oriented configuration edits to RESTCONF's HTTP resource model while retaining the intent to add a configuration instance. RFC 8040 defines POST as the RESTCONF method used to create a resource, including behavior for requests sent to datastore resources and list or container targets.

References: RFC 8040, RESTCONF POST; Cisco IOS XE RESTCONF documentation.

QUESTION NO: 153

Which two components are needed when a Cisco SD-Access fabric is designed? (Choose two.)

A. Cisco Catalyst Center (formerly DNA Center) application

B. Firepower Threat Defense

C. Identity Service Engine

D. Cisco Data Center Network Manager

E. Cisco Prime Infrastructure

ANSWER: A C

Explanation:

Cisco Catalyst Center (formerly DNA Center) application and Identity Service Engine are the core management and policy components used in a Cisco SD-Access design. Cisco Catalyst Center provides the automation and orchestration platform for designing, provisioning, and operating the fabric. It defines fabric sites, assigns network-device roles such as fabric edge, control-plane, and border nodes, creates virtual networks, and automates deployment of the underlay and overlay configuration. This centralized workflow is fundamental to the intent-based operational model of SD-Access.

Identity Service Engine supplies identity-aware access control and policy services. It authenticates users and endpoints, gathers contextual identity information, and integrates Security Group Tags and scalable group-based policies with the SD-Access fabric. This allows access and segmentation policies to follow authenticated users and devices consistently, rather than depending solely on IP subnets or physical location. Together, Cisco Catalyst Center (formerly DNA Center) application establishes and automates the fabric infrastructure, while Identity Service Engine supplies the identity and segmentation policy model that the fabric enforces. Cisco documents both platforms as integral elements of SD-Access architecture and deployment. See the Cisco SD-Access overview and Cisco Catalyst Center documentation.

QUESTION NO: 154

Refer to the Exhibit.

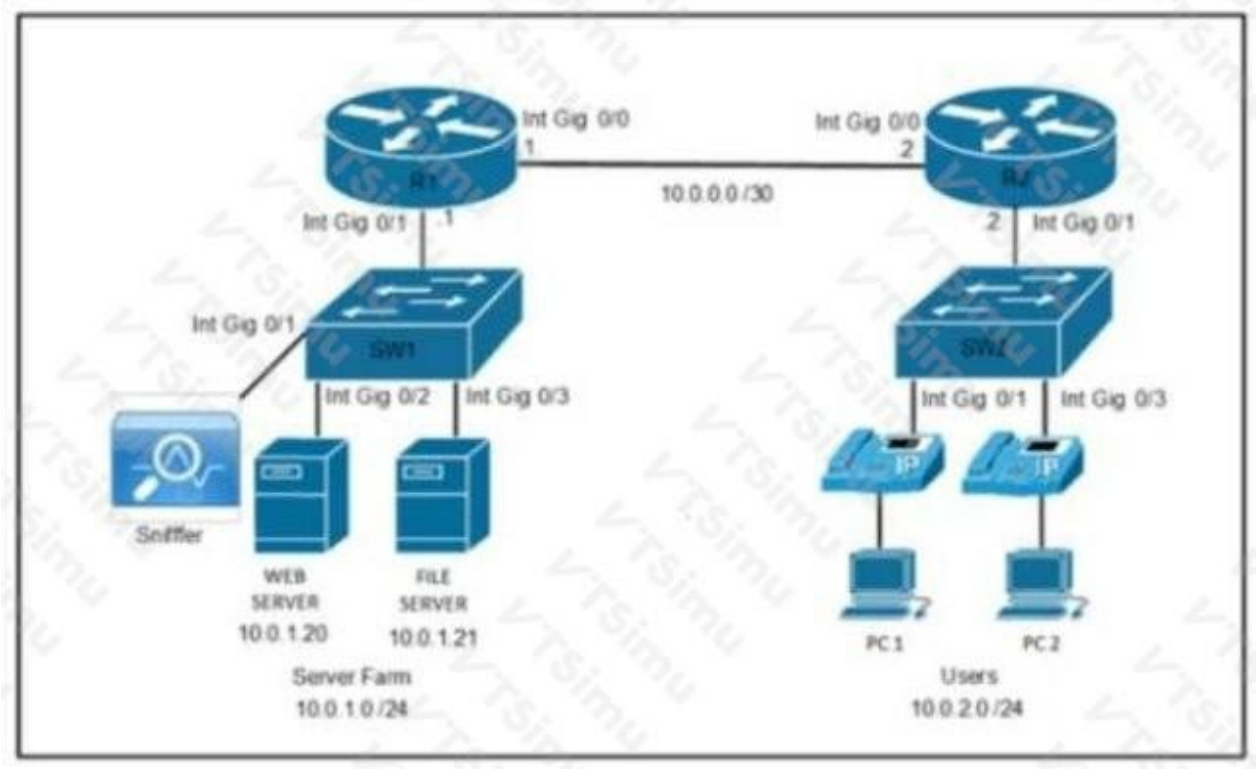

Refer to the exhibit. A network engineer Is troubleshooting an Issue with the file server based on reports of slow file transmissions. Which two commands or command sets are required. In switch SW1 to analyze the traffic from the file server with a packet analyzer? (Choose two.)

A)

```
SW1#show monitor
```

B)

```
SW1(config)# monitor session 1 source interface gigabitethernet0/3
SW1(config)# monitor session 1 destination interface gigabitethernet0/1 encapsulat
```

```
SW1#show ip route
```

D)

```
SW1#show vlan
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: A B

Explanation:

Traffic analysis with a packet analyzer attached to a switch requires a local Switched Port Analyzer (SPAN) session. A complete SPAN session has two essential configuration elements: a source and a destination. The selected command sets provide those elements by identifying the switch interface that carries the file server traffic as the monitored source and identifying the interface connected to the packet analyzer as the SPAN destination. After both parts are configured under the same monitor-session number, SW1 copies frames observed on the source to the destination interface, allowing the analyzer to capture and inspect the file-transfer packets.

The source direction should be chosen to match the traffic under investigation. For example, monitoring traffic received from the server captures frames that the server sends into SW1, while using both directions provides visibility into the complete client-server exchange. The destination interface is dedicated to the analyzer for the duration of the capture; it receives replicated frames rather than forwarding normal switched traffic. This mirrored packet view lets the engineer examine retransmissions, window behavior, packet loss indicators, and throughput-related symptoms without inserting a device inline. Cisco documents the required source and destination components of a local SPAN session in its Catalyst SPAN configuration guide and SPAN configuration and concepts documentation.

QUESTION NO: 155

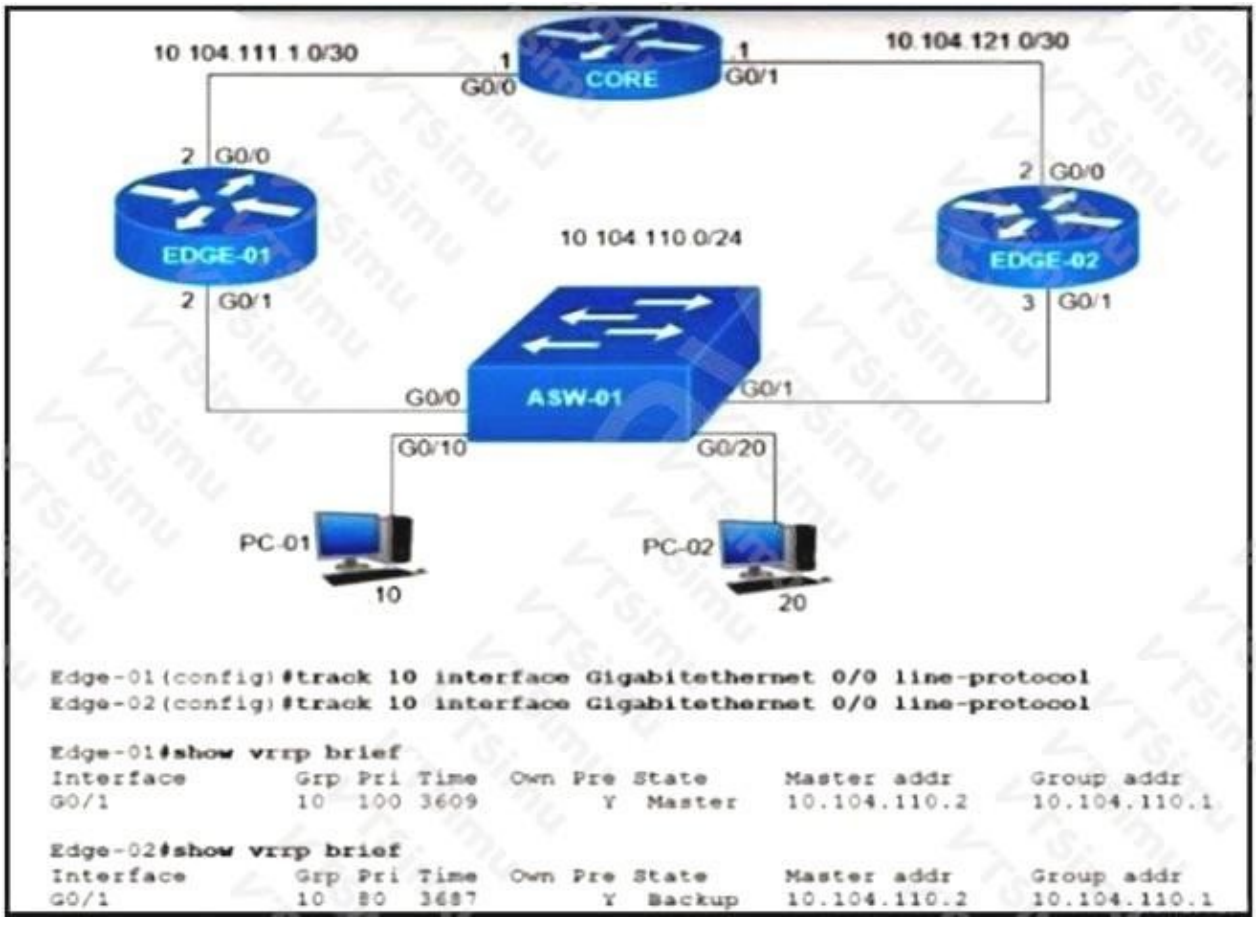

Refer to the exhibit. Object tracking has been configured for VRRP-enabled routers Edge-01 and Edge-02. Which commands cause Edge-02 to preempt Edge-01 in the event that interface G0/0 goes down on Edge-01?

A. Edge-01(config)#interface G0/1
Edge-01(config-if)#vrrp 10 track 10 decrement 30

B. Edge-02(config)#interface G0/1
Edge-02(config-if)#vrrp 10 track 10 decrement 30

C. Edge-02(config)#interface G0/1
Edge-02(config-if)#vrrp 10 track 10 decrement 10

D. Edge-01(config)#interface G0/1
Edge-01(config-if)#vrrp 10 track 10 decrement 10

ANSWER: A

Explanation:

`Edge-01(config)#interface G0/1`
`Edge-01(config-if)#vrrp 10 track 10 decrement 30` is correct because the tracked interface failure occurs on Edge-01, which is initially the VRRP master for group 10. The tracking command must therefore be applied to VRRP group 10 on Edge-01's gateway-facing interface, G0/1. When tracked object 10 reports that G0/0 is down, VRRP reduces Edge-01's configured priority by 30. Based on the priorities shown in the exhibit, this lowers Edge-01's effective priority from 110 to 80, below Edge-02's priority of 100. Edge-02 then has the higher VRRP priority and preempts to become the master, taking responsibility for the virtual IP and forwarding traffic as the active default gateway. Object tracking is used precisely for this purpose: it allows a router's first-hop redundancy priority to reflect the availability of an important upstream path rather than merely the state of its LAN-facing interface. Cisco documents that the VRRP `track` command decrements VRRP priority when the associated object goes down. See the Cisco IOS XE VRRP Configuration Guide and the Cisco First Hop Redundancy Protocol command reference.

QUESTION NO: 156

An engineer configures GigabitEthernet 0/1 for VRRP group 115. The router must assume the primary role when it has the highest priority in the group. Which command set is required to complete this task?

A. Option A

B. Option B

C. Option C

D. Option D

E. `interface GigabitEthernet0/1`
`vrrp 115 priority 110`
`vrrp 115 preempt`

ANSWER: E

Explanation:

The required configuration is `vrrp 115 priority 110` together with `vrrp 115 preempt` under GigabitEthernet0/1. VRRP uses a numerical priority to elect the Master router for a virtual router group; a larger value is preferred. Setting a value such as 110 makes this router preferred over routers using the default VRRP priority of 100. The `preempt` command ensures that, when this router is available and its configured or effective priority exceeds that of the current Master, it takes over the Master role rather than remaining a backup router. This produces the requested outcome: the router becomes primary whenever it has the highest priority in VRRP group 115. In a complete interface configuration, the VRRP virtual IPv4 address for group 115 must also already be configured so that hosts have a shared default-gateway address; however, priority and preemption are the commands that implement the stated role-selection requirement. Cisco documents VRRP priority and preemption behavior in its VRRP overview and configuration guide. The standards-based VRRP election and preemption behavior is also defined in RFC 5798.

QUESTION NO: 157

How are the different versions of IGMP compatible?

A. IGMPv2 is compatible only with IGMPv2.

B. IGMPv3 is compatible only with IGMPv3.

C. IGMPv2 is compatible only with IGMPv1.

D. IGMPv3 is compatible only with IGMPv1

E. IGMPv3 is backward compatible with IGMPv2 and IGMPv1 (and IGMPv2 is backward compatible with IGMPv1).

ANSWER: E

Explanation:

IGMPv3 is backward compatible with IGMPv2 and IGMPv1 (and IGMPv2 is backward compatible with IGMPv1). This compatibility supports incremental migration of multicast-enabled networks without requiring all hosts to be upgraded simultaneously. On a shared LAN, IGMP-capable routers use received IGMP messages to identify the oldest active IGMP version and adjust their interface behavior accordingly. For example, when an IGMPv3 router detects IGMPv2 hosts, it can operate with IGMPv2-compatible query and report behavior for that interface; similarly, the presence of IGMPv1 hosts requires compatibility behavior appropriate to IGMPv1. An IGMPv2 router can also support IGMPv1 hosts through its version-compatibility mechanisms. IGMPv3 adds source-filtering capabilities, including support for source-specific multicast, but those enhanced capabilities are available only when the relevant receivers and router interface can operate as IGMPv3. Compatibility preserves basic multicast group membership operation for older receivers while allowing newer versions to be deployed over time. Cisco documents IGMP version compatibility and router fallback behavior in its multicast configuration guidance. The IGMPv3 protocol specification also defines interactions with earlier versions. See Cisco: IGMPv3 Overview and Compatibility and RFC 3376: IGMPv3.

QUESTION NO: 158

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:

SNR is computed from the signal level received at a receiver and the background noise level at that same receiver. In Cisco wireless terminology, RSSI represents the received signal strength, while the noise floor represents the aggregate ambient noise measured by the radio. These values are normally expressed in dBm, a logarithmic power unit, so their difference produces SNR in dB: SNR = RSSI (dBm) − noise floor (dBm). For example, a received signal of −65 dBm with a noise floor of −95 dBm yields an SNR of 30 dB. This receiver-side calculation is useful because it reflects the actual RF conditions affecting the client or access point at the time of measurement, including propagation loss and environmental noise. Cisco wireless documentation identifies both RSSI and SNR as key RF measurements used to assess wireless link quality, and Cisco's RF fundamentals material describes the noise floor as the baseline against which received signals are evaluated. See Cisco's Wireless LAN RF fundamentals and the Cisco wireless design guide for RF measurement and design context.