DUMPSARENA

Cisco Certified Network Associate

Cisco 200-301

Version Demo

Total Demo Questions: 239

Total Premium Questions: 2393

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Network Fundamentals	451
Topic 2, Network Access	578
Topic 3, IP Connectivity	478
Topic 4, IP Services	268
Topic 5, Security Fundamentals	372
Topic 6, Automation and Programmability	242
Topic 7, Mix Questions	4
Total	2393

QUESTION NO: 1

What are two characteristics of a controller-based network? (Choose two.)

A. It uses Telnet to report system issues.

B. The administrator can make configuration updates from the CLI.

C. It uses northbound and southbound APIs to communicate between architectural layers.

D. It decentralizes the control plane, which allows each device to make its own forwarding decisions.

E. It moves the control plane to a central point.

ANSWER: C E

Explanation:

It uses northbound and southbound APIs to communicate between architectural layers because controller-based networking separates applications, the controller, and infrastructure devices into distinct functional layers. Northbound APIs enable applications, orchestration tools, and management systems to express intent or request network services from the controller. Southbound APIs enable the controller to communicate policies, configurations, and operational instructions to the network infrastructure. This API-driven model supports automation and consistent, scalable policy deployment.

It moves the control plane to a central point because a controller-based architecture logically centralizes control and policy decisions in the controller. The controller maintains a broader network view and translates centralized intent into device-specific actions. Network devices continue to perform data-plane forwarding, but they rely on the controller-based system for coordinated policy and control functions. “Central” describes logical centralization; the controller platform itself can be implemented for scale and high availability rather than as a single physical device. Cisco identifies centralized controllers and programmable APIs as foundational elements of software-defined networking and controller-based enterprise automation. See Cisco’s Software-Defined Networking overview and Enterprise Network Automation overview.

QUESTION NO: 2

Refer to the exhibit.

Which two commands when used together create port channel 10? (Choose two.)

A. channel-group 10 mode active

B. int range g0/0-1 channel-group 10 mode desirable

C. channel-group 10 mode passive

D. int range g0/0-1 channel-group 10 mode auto

E. int range g0/0-1 channel-group 10 mode on

ANSWER: A C

Explanation:

`channel-group 10 mode active` and `channel-group 10 mode passive` form an LACP EtherChannel negotiation pair for Port-Channel 10. These commands are applied under the physical member interfaces, such as the GigabitEthernet interfaces intended to participate in the bundle. The `channel-group 10` portion places each selected physical interface into EtherChannel group 10; when the group is created, Cisco IOS creates the corresponding logical Port-Channel10 interface. The mode keywords select Link Aggregation Control Protocol operation. An interface in active mode initiates LACP negotiation by transmitting LACP packets, while an interface in passive mode listens and responds to LACP packets from its peer. Therefore, placing the member links on one side in active mode and the matching links on the peer side in passive mode allows the two devices to negotiate and establish the same LACP port channel. LACP is the IEEE standards-based EtherChannel negotiation protocol and supports interoperable link aggregation when both ends are configured consistently. Cisco documents the active/passive LACP mode pairing and the use of the `channel-group` interface command in its

EtherChannel configuration guidance: Cisco EtherChannel Configuration Guide. For LACP protocol behavior, see Cisco EtherChannel and LACP overview.

QUESTION NO: 3

Which two characteristics are representative of virtual machines (VMs)? (Choose two.)

A. multiple VMs operate on the same underlying hardware

B. Each VMs operating system depends on its hypervisor

C. A VM on a hypervisor is automatically interconnected to other VMs

D. A VM on an individual hypervisor shares resources equally

E. Each VM runs independently of any other VM in the same hypervisor

ANSWER: A E

Explanation:

Virtualization enables a hypervisor to abstract the physical resources of a server, such as CPU, memory, storage, and network interfaces, into virtual hardware that can be assigned to separate virtual machines. Therefore, *multiple VMs operate on the same underlying hardware* is a core VM characteristic. This consolidation allows an organization to use a physical server more efficiently while supporting distinct workloads, operating systems, and applications.

Each VM runs independently of any other VM in the same hypervisor is also representative of virtual machines. Every VM is an isolated software-defined computer with its own guest operating system, virtual CPU, memory allocation, virtual disk, and virtual network interface. The hypervisor maintains separation between these guest environments, so an application or operating-system event within one VM is contained within that VM rather than being inherently part of another guest environment. This isolation is a key reason virtualization supports workload separation and flexible deployment on shared infrastructure.

Cisco describes virtualization as using software to create multiple simulated environments from one physical system, while VMware explains that VMs are isolated computing environments running on a physical host. See Cisco: What Is Virtualization? and VMware: Virtual Machine.

QUESTION NO: 4

Refer to the exhibit.

```
R1# show ip route
....
D        172.16.32.0/27 [90/2888597172]  via 20.1.1.1
O        172.16.32.0/19  [110/292094]   via 20.1.1.10
R        172.16.32.0/24  [120/2]  via 20.1.1.3
```

An engineer executed the script and added commands that were not necessary for SSH and now must remove the commands.

A. metric

B. cost

C. longest prefix

D. administrative distance

E. The required SSH command removal cannot be determined because the exhibit containing the script is unavailable.

ANSWER: E

Explanation:

The required SSH-related commands cannot be identified from the supplied question data because the referenced exhibit, which should contain the script and configuration commands, is not available in the JSON content. The visible answer choices concern route selection—metric, cost, longest prefix, and administrative distance—and none identifies a command that can be removed from an SSH configuration. Determining whether a command is unnecessary requires seeing the actual script. A basic Cisco IOS SSH setup normally includes a hostname, an IP domain name, locally configured user credentials, generated RSA keys, SSH version 2, and VTY lines configured to authenticate locally and accept SSH access. Additional commands may be appropriate depending on the intended security policy, such as restrictions on VTY access or SSH timeouts. Without the missing exhibit, it is not technically valid to identify a specific command for removal. Cisco's SSH configuration guidance describes the required configuration elements and provides the context needed to assess extra commands when the script is available. See Cisco Secure Shell configuration guide and Cisco IOS XE SSH Version 2 documentation.

QUESTION NO: 5

What is the role of disaggregation in controller-based networking?

A. It divides the control-plane and data-plane functions.

B. It summarizes the routes between the core and distribution layers of the network topology.

C. It enables a network topology to quickly adjust from a ring network to a star network

D. It streamlines traffic handling by assigning individual devices to perform either Layer 2 or Layer 3 functions.

ANSWER: A

Explanation:

It divides the control-plane and data-plane functions. In controller-based networking, disaggregation separates capabilities that have traditionally been tightly integrated within each network device. The control plane—the functions that establish topology information, calculate paths, and apply network policy—can be implemented centrally or logically centrally by a controller. The data plane remains on the network infrastructure and forwards packets at high speed according to the policies and forwarding information it receives. This separation gives administrators a centralized point from which to automate configuration, deploy consistent policy, and maintain network-wide visibility, while switches and routers continue performing efficient packet forwarding. Cisco describes software-defined networking as an architecture that separates the network control plane from the forwarding plane, enabling centralized control and programmability. In Cisco controller-based enterprise solutions, controllers use this architectural model to translate intent and policy into configurations and forwarding behavior across the infrastructure. The result is a more programmable and operationally consistent network, particularly as the number of devices and policies grows. See Cisco's Software-Defined Networking overview and the Cisco DevNet explanation of control-plane and data-plane separation.

QUESTION NO: 6

Which two wireless security stewards use Counter Mode Cipher Block Chaining Message Authentication Code Protocol for encryption and data integrity'? (Choose two.)

A. WPA2

B. WPA3

C. Wi-Fi 6

D. WEP

E. WPA

F. IEEE 802.11i (RSN)

ANSWER: A B F

Explanation:

Counter Mode Cipher Block Chaining Message Authentication Code Protocol (CCMP) is the AES-based IEEE 802.11 data-confidentiality and integrity protocol. It combines AES operating in Counter Mode to encrypt wireless data with CBC-MAC to authenticate frames and protect their integrity. IEEE 802.11i, also called Robust Security Network (RSN), defines CCMP as the replacement for the legacy WEP and TKIP protection mechanisms. WPA2 certification is based on RSN/IEEE 802.11i security and uses AES-CCMP for protected WLAN traffic. WPA3 also supports and uses CCMP-128 in its standard security profiles; for example, WPA3-Personal protects data with CCMP-128 after SAE authentication establishes session keys. Some WPA3 enterprise configurations can additionally use stronger GCMP-based suites, but that does not remove CCMP support and use in WPA3. Therefore, WPA2, WPA3, and IEEE 802.11i (RSN) are all technically valid statements among the supplied choices. The wording “Choose two” is inaccurate because it omits that the underlying IEEE 802.11i standard itself also defines and uses CCMP. NIST describes CCMP as the AES-based data-confidentiality and integrity protocol defined for RSNs, and the Wi-Fi Alliance identifies WPA2 and WPA3 as modern Wi-Fi security generations. See NIST SP 800-97 and Wi-Fi Alliance security overview.

QUESTION NO: 7

Which type of information resides on a DHCP server?

A. a list of the available IP addresses in a pool

B. a list of public IP addresses and their corresponding names

C. usernames and passwords for the end users in a domain

D. a list of statically assigned MAC addresses

ANSWER: A

Explanation:

A DHCP server maintains an address pool, also called a scope, containing the IP addresses that it can allocate dynamically to clients. As clients request network configuration, the server selects an available address from that pool, records the resulting lease, and provides associated parameters such as the subnet mask, default gateway, DNS server addresses, and lease time. Maintaining the pool and lease state lets the server determine which addresses remain available and prevents the same dynamically assigned address from being issued to multiple clients simultaneously.

This is a fundamental part of DHCP operation: the server is responsible for managing address availability and binding a leased address to a client for a defined period. Cisco IOS DHCP configuration similarly defines a DHCP pool and its network range, from which the device assigns addresses to DHCP clients. The DHCP protocol specification describes the server’s allocation of network addresses and maintenance of client bindings. See Cisco’s DHCP overview and configuration concepts and RFC 2131, Dynamic Host Configuration Protocol.

QUESTION NO: 8

What is a service that is provided by a wireless controller?

A. It mitigates threats from the internet.

B. It manages interference in a dense network.

C. It provides Layer 3 routing between wired and wireless devices.

D. It issues IP addresses to wired devices.

ANSWER: B

Explanation:

It manages interference in a dense network. A wireless LAN controller provides centralized control and management for lightweight access points, including Radio Resource Management (RRM). RRM continuously evaluates the radio-frequency environment across the wireless deployment and helps maintain reliable service when many access points and clients operate close together. Its functions include dynamic channel assignment, transmit-power control, coverage optimization, and detection of RF conditions that can degrade wireless performance. In a dense deployment, neighboring access points using overlapping or congested channels can create co-channel interference and reduce client throughput. The controller uses information reported by managed access points to select appropriate RF settings and adapt to environmental changes, rather than requiring each access point to be planned and tuned independently. This centralized RF coordination is a core wireless-controller service and supports consistent WLAN performance as client density and interference conditions change. Cisco describes RRM as an automatic system for optimizing RF parameters such as channel and transmit power across a wireless network. See Cisco's Radio Resource Management configuration guide and its wireless LAN controller overview.

QUESTION NO: 9

What are two benefits of controller-based networking compared to traditional networking?

A. controller-based increases network bandwidth usage, while traditional lightens the load on the network.

B. controller-based inflates software costs, while traditional decreases individual licensing costs

C. Controller-based reduces network configuration complexity, while traditional increases the potential for errors

D. Controller-based provides centralization of key IT functions. While traditional requires distributed management functions

E. controller-based allows for fewer network failure, while traditional increases failure rates.

ANSWER: C D

Explanation:

Controller-based reduces network configuration complexity, while traditional increases the potential for errors because a controller gives administrators a centralized system for defining and deploying network intent, policy, and configuration. Rather than repeatedly configuring devices individually, teams can use automation, templates, and programmatic interfaces to apply consistent settings at scale. This reduces repetitive work and helps prevent configuration drift caused by manual, device-by-device changes.

Controller-based provides centralization of key IT functions. While traditional requires distributed management functions because the controller provides a logical central point for management, policy enforcement, network visibility, and orchestration. Network devices still forward traffic in the data plane, but the controller can collect telemetry and coordinate configurations across the infrastructure. Cisco describes controller-based networking as separating management and control functions from individual network devices and enabling centralized policy and automation. These capabilities are central benefits of the controller-based model covered in CCNA automation and programmability objectives. See Cisco's Intent-Based Networking overview and the CCNA exam topics.

QUESTION NO: 10

A Cisco engineer is configuring a factory-default router with these three passwords:

â€¢ The user EXEC password for console access is p4ssw0rd1

â€¢ The user EXEC password for Telnet access is s3cr3t2

â€¢ The password for privileged EXEC mode is pr1v4t3p4ss Which command sequence must the engineer configured

A)

```
enable secret priv4t3p4ss
!
line con 0
password login p4ssw0rd1
!
line vty 0 15
password login s3cr3t2
login
```

B)

```
enable secret privilege 15 priv4t3p4ss
!
line con 0
password p4ssw0rd1
login
!
line vty 0 15
password s3cr3t2
login
```

C)

```
enable secret priv4t3p4ss
!
line con 0
password p4ssw0rd1
login
!
line vty 0 15
password s3cr3t2
login
```

D)

```
enable secret priv4t3p4ss
!
line con 0
```

A. Option A

B. Option B

C. Option C

D. Option D

E. line console 0
password p4ssw0rd1
login
exit
line vty 0 4
password s3cr3t2
login
exit
enable secret pr1v4t3p4ss

ANSWER: E

Explanation:

The required configuration applies a password to each relevant line type and enables password checking on those lines with the `login` command. `line console 0` selects the physical console line, so assigning `p4ssw0rd1` there establishes the requested user EXEC authentication for direct console connections. `line vty 0 4` selects the default virtual terminal lines used for remote terminal sessions, including Telnet when Telnet is permitted, and assigning `s3cr3t2` with `login` establishes the requested remote user EXEC password.

For privileged EXEC access, `enable secret pr1v4t3p4ss` is the appropriate command. Cisco documents that `enable secret` configures the privileged EXEC secret and stores it using a one-way hash rather than the legacy reversible representation used by `enable password`. This is the Cisco-recommended method for protecting access to enable mode. The sequence can be entered from global configuration mode, moving into and out of each line subconfiguration mode as shown. Cisco's password configuration guidance demonstrates the use of `password` plus `login` under console and VTY lines: Cisco: Password Configuration. Cisco also documents the privileged EXEC secret command in its IOS XE password-encryption configuration guide.

QUESTION NO: 11

```
R2#show ip route
C        192.168.1.0/26 is directly connected, FastEthernet0/1
```

Refer to the exhibit. Which two prefixes are included in this routing table entry? (Choose two.)

A. 192.168.1.17

B. 192.168.1.61

C. 192.168.1.64

D. 192.168.1.127

E. 192.168.1.254

F. 192.168.1.31

ANSWER: A F

Explanation:

The routing-table entry represents the network 192.168.1.16/28. A /28 prefix uses the mask 255.255.255.240, leaving four host bits and creating blocks of 16 addresses in the last octet. Therefore, the subnet beginning at 192.168.1.16 covers the complete address interval from 192.168.1.16 through 192.168.1.31. 192.168.1.17 falls within that interval, so traffic destined for that address matches the route. 192.168.1.31 is also within the numerical range covered by the prefix. Although it is the directed broadcast address for this subnet and ordinarily is not assigned to an individual host, it remains an address contained by the 192.168.1.16/28 network range. Route-prefix matching is determined by applying the prefix mask to the destination address; it is not limited to assignable unicast host addresses. Cisco routing-table documentation describes routes as prefix-based entries and shows how Cisco IOS displays network prefixes and masks. The subnetting calculation follows the IPv4 prefix-length rules defined in RFC 4632. For Cisco IOS routing-table interpretation, see Cisco: Understanding IP Routing Table.

QUESTION NO: 12

When implementing a router as a DHCP server, which two features must be configured ' ? (Choose two)

A. relay agent information

B. database agent

C. address pool

D. smart-relay

E. manual bindings

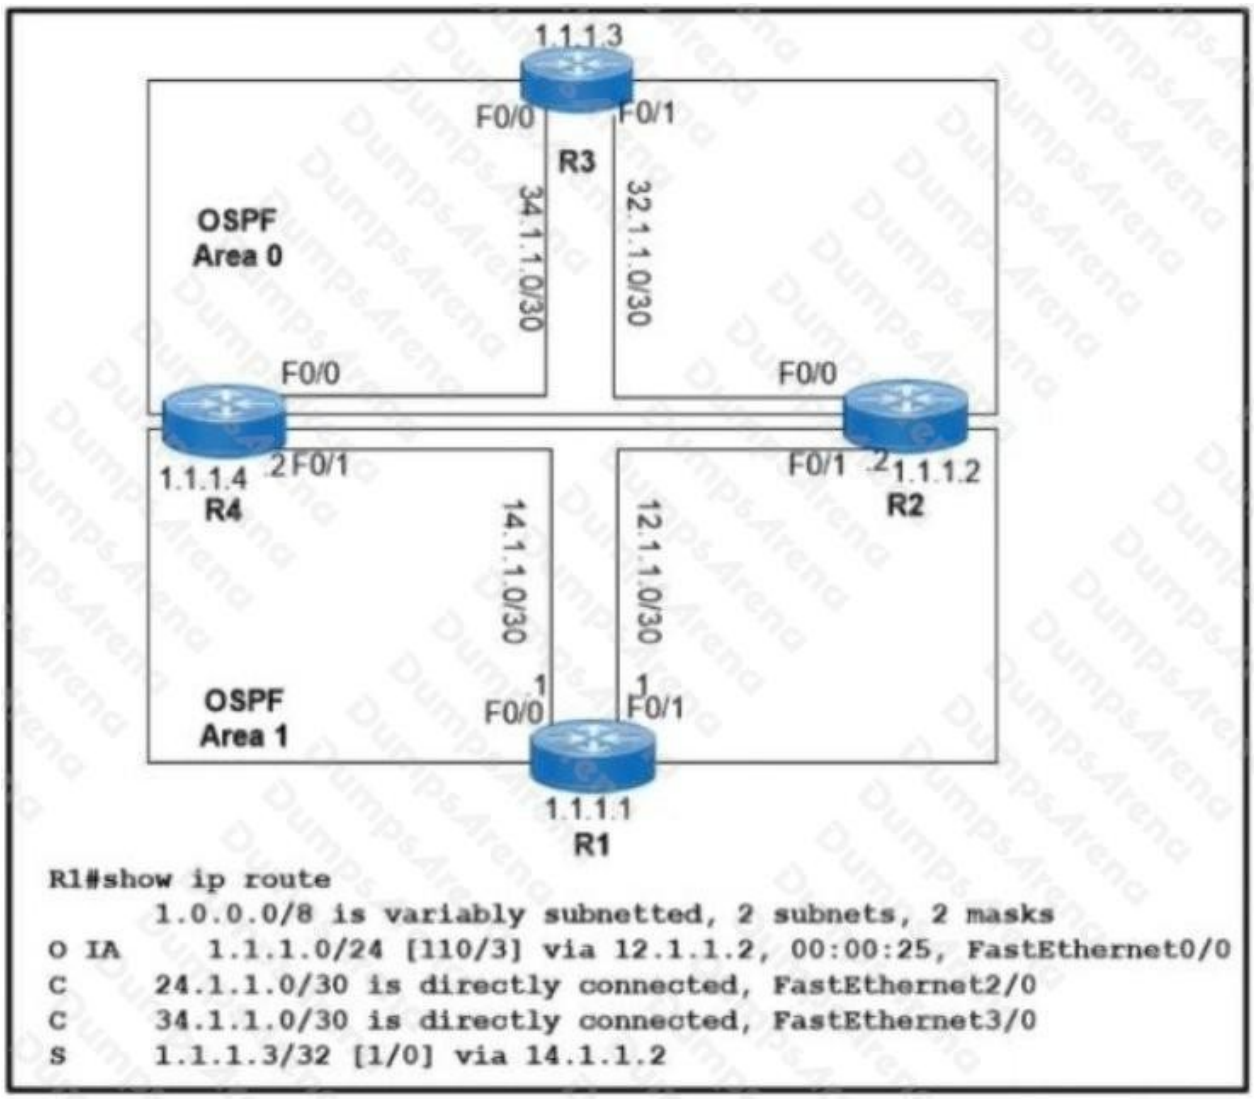
ANSWER: C E

Explanation:

address pool and **manual bindings** are the DHCP-server configuration features identified here. A Cisco IOS router supplies DHCP service through a named address pool, created with the `ip dhcp pool` command. The pool is the central configuration object for defining the client address range or network and the DHCP options delivered to clients, such as the default gateway, DNS servers, domain name, and lease duration. Without an address pool, the router has no scope from which it can allocate addresses to DHCP clients.

manual bindings provide deterministic address assignment for clients that require a fixed address. In IOS, a manual binding is configured under a DHCP pool using a client identifier or hardware address together with a specific address. This associates the known client with its intended IP address while allowing the router to remain the DHCP authority for that assignment. Cisco documents manual bindings as part of IOS DHCP server configuration and uses them when an administrator needs reservations for devices such as infrastructure endpoints or servers. See Cisco's Configuring the Cisco IOS DHCP Server and the Cisco IOS DHCP Command Reference.

QUESTION NO: 13



Refer to the exhibit. Which two values does router R1 use to determine the best path to reach destinations in network 1.0.0.0/8? (Choose two.)

A. lowest cost to reach the next hop

B. highest administrative distance

C. lowest metric

D. highest metric

E. longest prefix match

F. lowest administrative distance

ANSWER: C F

Explanation:

Router R1 uses the **lowest administrative distance** and then the **lowest metric** to select the preferred route when it has learned routes for the same destination prefix from more than one source. Administrative distance represents the router's trust in the source of routing information. A smaller value is preferred, so a route learned from the most trusted source is installed ahead of routes to the same prefix learned from less trusted sources.

If candidate routes to the same prefix have equal administrative distance, R1 compares their routing metrics. The metric is the protocol-specific measurement used to identify the most favorable path, such as a composite bandwidth-and-delay calculation for EIGRP or cost for OSPF. The route with the lowest metric is preferred and placed in the routing table. Thus, for the displayed routes to the identical 1.0.0.0/8 destination network, administrative distance establishes the preferred route source and metric resolves a tie among equally trusted route candidates. Cisco describes this route-selection behavior in its Routing Protocols and Administrative Distance documentation and in the Cisco IP Routing route-process documentation.

QUESTION NO: 14

What are two benefits of controller-based networking compared to traditional networking?

A. controller-based increases network bandwidth usage, while traditional lightens the load on the network.

B. controller-based inflates software costs, while traditional decreases individual licensing costs

C. Controller-based reduces network configuration complexity, while traditional increases the potential for errors

D. Controller-based provides centralization of key IT functions. While traditional requires distributes management function

E. controller-based allows for fewer network failure, while traditional increases failure rates.

ANSWER: C D

Explanation:

Controller-based networking reduces network configuration complexity by separating centralized control and management functions from the individual forwarding devices. Administrators can define intent, policy, and configuration centrally, then use automation to deploy consistent settings to many devices rather than manually configuring each one. This consistency reduces the number of repetitive device-level tasks and helps limit configuration drift and human-entry errors as the network grows or changes. Cisco describes Cisco Catalyst Center as a platform for centralized network management and automation, including provisioning, policy, assurance, and lifecycle operations.

Controller-based networking also provides centralization of key IT functions. A controller maintains a network-wide view that can be used to coordinate configuration, access policy, telemetry, and operational workflows across the infrastructure. This centralized approach allows teams to apply standardized configurations and policies broadly while retaining visibility into overall network state. In contrast to independently managed devices, the controller-driven model supports more uniform operations and faster, repeatable changes. Cisco's enterprise automation architecture emphasizes intent-based management, centralized policy, and automated provisioning as core capabilities. See Cisco Catalyst Center and Cisco Intent-Based Networking.

QUESTION NO: 15

What are two functions of a server on a network? (Choose two.)

A. handles requests from multiple workstations at the same time

B. achieves redundancy by exclusively using virtual server clustering

C. housed solely in a data center that is dedicated to a single client

D. runs the same operating system in order to communicate with other servers

E. runs applications that send and retrieve data for workstations that make requests

ANSWER: A E

Explanation:

A server is a networked system that supplies resources or services to other devices, commonly called clients. “handles requests from multiple workstations at the same time” is correct because server software is designed to listen for, receive, process, and respond to concurrent requests from many client devices. For example, a file server can serve files to several users, while a web server can return content to many browsers during the same period. Supporting shared access to centrally available resources is a fundamental server function.

“runs applications that send and retrieve data for workstations that make requests

” is also correct. Servers host applications and network services that perform work in response to client requests, such as delivering web content, storing and retrieving files, resolving names, assigning network configuration, authenticating users, or providing database records. The service application processes the request and returns the requested data or result across the network. Cisco describes servers as devices that provide services and resources to other devices, including file, web, email, and application services. See Cisco’s overview of servers and Cisco Unified Computing System information.

QUESTION NO: 16

Which security method is used to prevent man-in-the-middle attacks?

A. authentication

B. anti-replay

C. authorization

D. accounting

ANSWER: A

Explanation:

authentication is the security method used to prevent man-in-the-middle attacks because it verifies that a communicating endpoint is genuinely the identity it claims to be. A man-in-the-middle attack depends on the attacker successfully impersonating one endpoint to the other, or impersonating both endpoints, so that the victim devices establish communications through the attacker. Proper authentication prevents that impersonation by requiring proof of identity, such as a trusted digital certificate and a corresponding private key, a validated public-key fingerprint, or credentials used in a mutually authenticated protocol.

For example, TLS uses certificates to authenticate a server and can also authenticate a client. IPsec uses Internet Key Exchange authentication to establish security associations between verified peers. When certificate validation, trusted certificate authorities, hostname verification, and strong cryptographic credentials are correctly configured, an attacker cannot create a trusted session merely by intercepting traffic. Authentication is therefore foundational to establishing an encrypted session with the intended peer rather than an impostor. Cisco discusses the role of authentication in its AAA model at Cisco AAA overview; its secure-communications overview also describes IPsec peer authentication and protected exchanges: Cisco IPsec overview.

QUESTION NO: 17

Which two syslog severity levels identify the most critical events? (Choose two.)

A. emergency

B. alert

C. warning

D. notification

E. debugging

ANSWER: A B

Explanation:

Emergency is severity level 0, the highest-priority syslog level. It indicates that the system is unusable and requires immediate attention.

Alert is severity level 1 and indicates that immediate action is required. Syslog uses lower numerical values for more severe events; therefore, emergency and alert are more critical than levels such as error, warning, notification, informational, and debugging. Cisco IOS supports the standard syslog severity scale described in the Cisco IOS XE System Message Logging Configuration Guide.

QUESTION NO: 18

Which two features introduced in SNMPv2 provides the ability to retrieve large amounts of data in one request

A. Get

B. GetNext

C. Set

D. GetBulk

E. Inform

F. The max-repetitions field of a GetBulk request

ANSWER: D F

Explanation:

GetBulk is the SNMPv2 operation designed to retrieve substantial management information efficiently in a single request/response exchange. Rather than requiring a manager to send repeated requests while traversing a MIB table, a GetBulk request can ask the agent to return lexicographically subsequent variable bindings. This substantially reduces protocol overhead and is particularly useful for collecting rows from SNMP tables, such as interface or routing information.

The **max-repetitions field of a GetBulk request** is the associated SNMPv2 capability that controls how many successive variable bindings the agent should attempt to return for repeating objects. Together with the non-repeaters value, it enables a manager to combine individual scalar-object retrieval with repeated table traversal in one GetBulk PDU. The actual amount returned remains limited by the agent, message-size constraints, and any applicable access-control policy, but max-repetitions is the mechanism that requests the bulk quantity of repeated data.

These behaviors are specified for the GetBulkRequest-PDU in RFC 3416, Protocol Operations for SNMP and in the original SNMPv2 protocol-operations specification, RFC 1905.

QUESTION NO: 19

An engineer must configure an OSPF neighbor relationship between router R1 and R3 The authentication configuration has been configured and the connecting interfaces are in the same 192.168 1.0/30 sublet. What are the next two steps to complete the configuration? (Choose two.)

A. configure the hello and dead timers to match on both sides

B. configure the same process ID for the router OSPF process

C. configure the same router ID on both routing processes

D. Configure the interfaces as OSPF active on both sides.

E. configure both interfaces with the same area ID

ANSWER: A E

Explanation:

For two OSPF routers connected on the same Layer 3 subnet to form a neighbor adjacency, the interfaces must participate in the same OSPF area and use compatible OSPF interface parameters. Configuring both interfaces with the same area ID satisfies the area-membership requirement. OSPF Hello packets carry the area ID, and a router accepts Hellos only when the received area information agrees with the local interface's OSPF configuration. Therefore, both ends of this point-to-point link must be assigned to the same area, such as area 0.

Configuring the hello and dead timers to match on both sides satisfies another essential adjacency requirement. OSPF sends Hello packets periodically to discover and maintain neighbors; the dead interval determines how long a router waits without receiving Hellos before declaring a neighbor down. The Hello and dead intervals advertised and configured on the connected interfaces must agree for the adjacency to progress normally. With the interfaces addressed in the same /30 subnet and authentication already made compatible, matching the timers and assigning a common area completes the stated requirements for the OSPF neighbor relationship. Cisco documents the required matching OSPF parameters in its OSPF Neighbor Relationship Troubleshooting guide and its IOS XE OSPF configuration guide.

QUESTION NO: 20

What is represented by the word "LB20" within this JSON schema?

```
1 [
2 {"load balancer": "LB20", "interface": "te4/3"},
3 {"firewall": "FW49", "interface": "ge4/14"},
4 {"IDS": "IPS_frankfurt", "interface": "e9/7"}
5 ]
```

A. value

B. array

C. object

D. key

ANSWER: D

Explanation:

The correct answer is **key**. JSON represents object members as a name/value pair using the form `"name": value`. In the schema, `"LB20"` is the quoted name positioned before a colon, so it identifies the associated member and acts as the key. The data that follows the colon is the value assigned to that key; that value can itself be a primitive value, an array, or another JSON object.

This naming structure is particularly useful in network automation data, where an identifier such as `LB20` may label a device, resource, interface, or other uniquely named item. Software reading the JSON can use that key to locate the corresponding data within the enclosing object. RFC 8259 defines a JSON object as an unordered collection of name/value pairs and specifies that a name is a string, making the quoted `LB20` a JSON member name (commonly called a key). See RFC 8259: The JSON Data Interchange Format and MDN Web Docs: Working with JSON.

QUESTION NO: 21

Which command should you enter to verify the priority of a router in an HSRP group?

A. show hsrp

B. show sessions

C. show interfaces

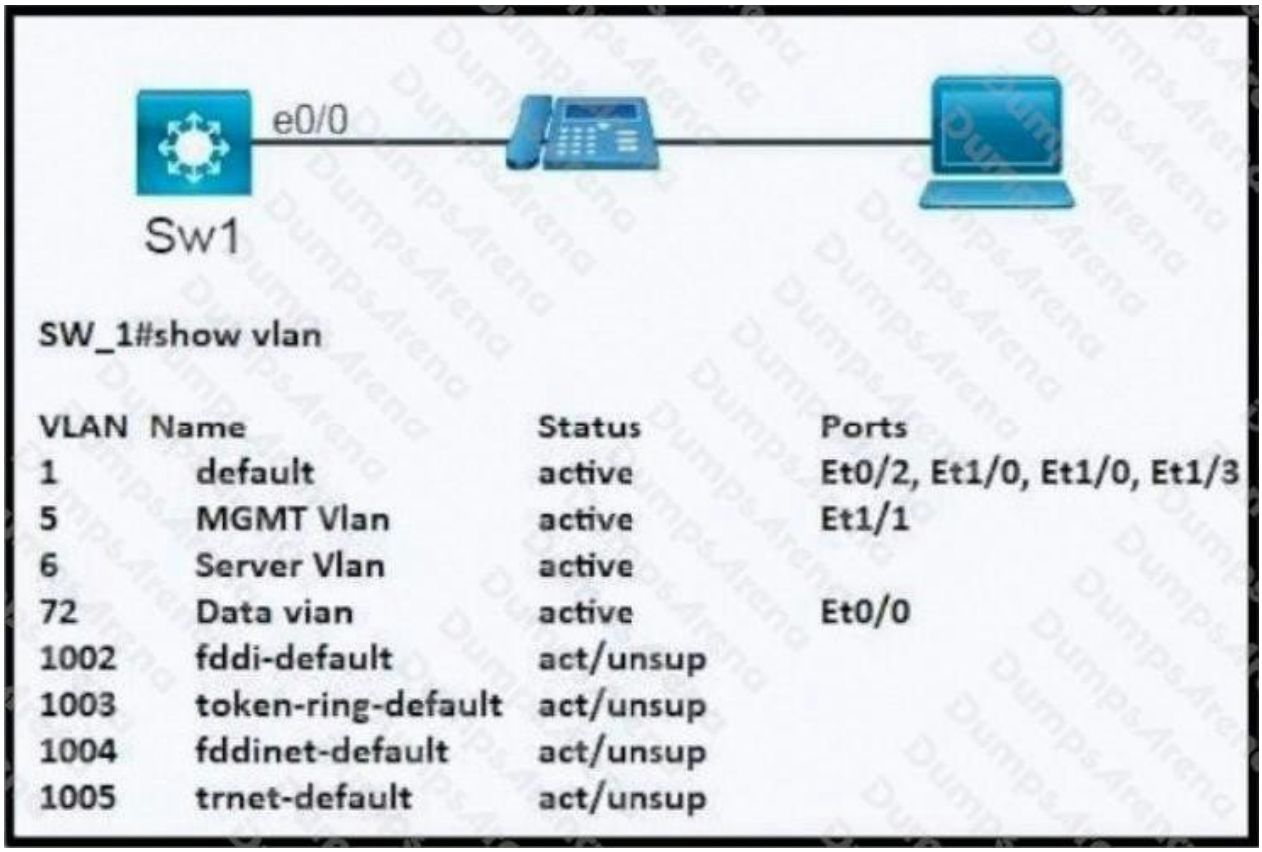
D. show standby

ANSWER: D

Explanation:

The correct command is **show standby**. In Cisco IOS, HSRP is configured and monitored per interface and HSRP group, and this command displays the operational and configuration details needed to verify a router's role in that group. Its output includes the locally configured *Priority* value, as well as the current HSRP state and information identifying the active and standby routers. This directly lets an administrator confirm the priority that participates in the HSRP election. HSRP uses a default priority of 100, and a router with a higher priority is preferred to become the active gateway, subject to the configured preemption behavior. The command can also be narrowed to a specific interface or group when needed, such as **show standby GigabitEthernet0/0** or **show standby brief**, but **show standby** is the standard verification command requested here. Cisco documents HSRP monitoring with the **show standby** command and shows that its output reports the local priority and peer role information. See the Cisco HSRP overview and configuration documentation and the Cisco IOS HSRP configuration guide.

QUESTION NO: 22



Refer to the exhibit. VoIP is being implemented in the network using VLAN ID 73 and named "VoIP". Each user needs a Cisco IP phone at their desk. Switchport e0/0 has been configured as an access port in the data VLA

N. Cisco Discovery Protocol is enabled globally. Which command sequence completed the configuration?

A. vlan73 name VoIP e0/0
switchport voice vlan dot1p

B. vlan 73 name VoIP e0/0
switchport trunk allowed vlan 72,73 switchport voice vlan 73

C. vlan 73 name VoIP e0/0
switchport mode trunk channel-group 73 mode active

D. vlan 73 name VoIP e0/0
switchport voice vlan 73

E. Cisco Discovery Protocol is enabled globally. Which command sequence completed the configuration?

ANSWER: D

Explanation:

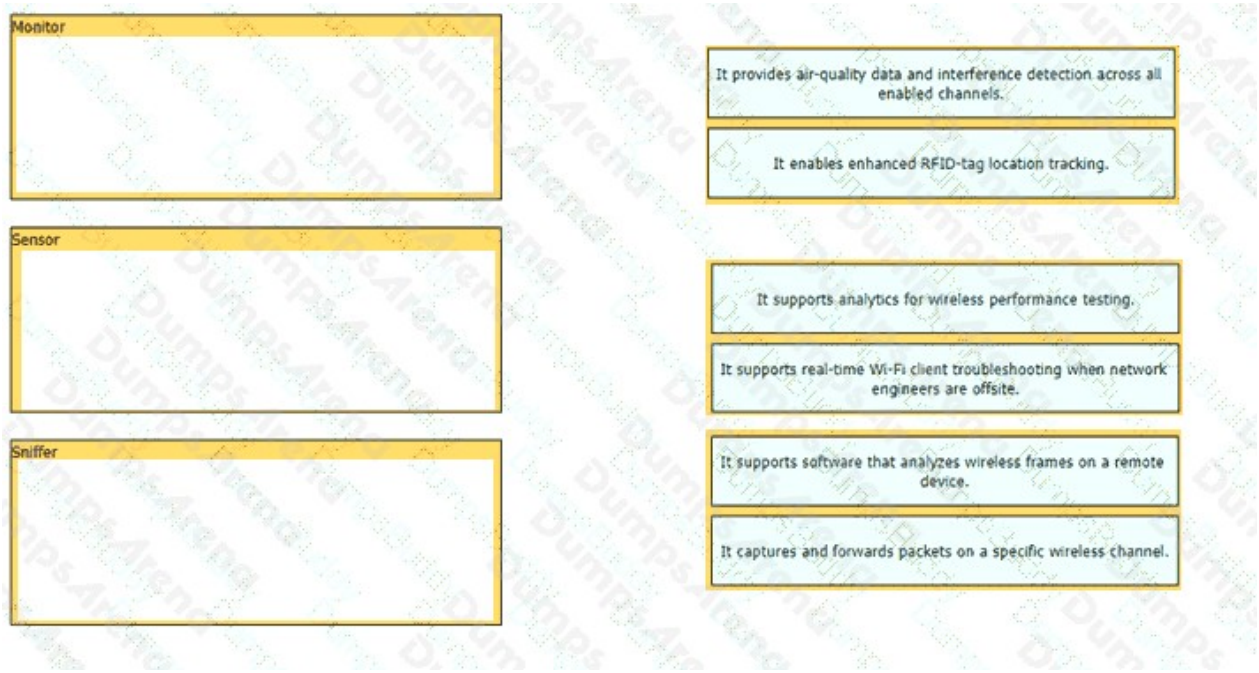
`vlan 73 name VoIP e0/0`
`switchport voice vlan 73` completes the required configuration because it creates VLAN 73, assigns it the requested VoIP name, and designates that VLAN as the voice VLAN on the existing access interface. A Cisco IP phone connected to an access switchport commonly provides a pass-through connection for the user's PC. In this design, the PC sends untagged data traffic on the port's configured data access VLAN, while the phone sends its voice traffic with an IEEE 802.1Q tag for the configured voice VLAN.

Because Cisco Discovery Protocol is enabled, the switch advertises the voice VLAN assignment to the attached Cisco IP phone. The phone can then learn that it must place voice traffic in VLAN 73 without requiring manual voice-VLAN configuration at the endpoint. The access-port configuration already stated in the question remains appropriate; the voice VLAN is an additional switchport setting rather than a replacement for the data VLAN configuration. Cisco documents the `switchport voice vlan vlan-id` interface command as the mechanism used to configure the VLAN that carries voice

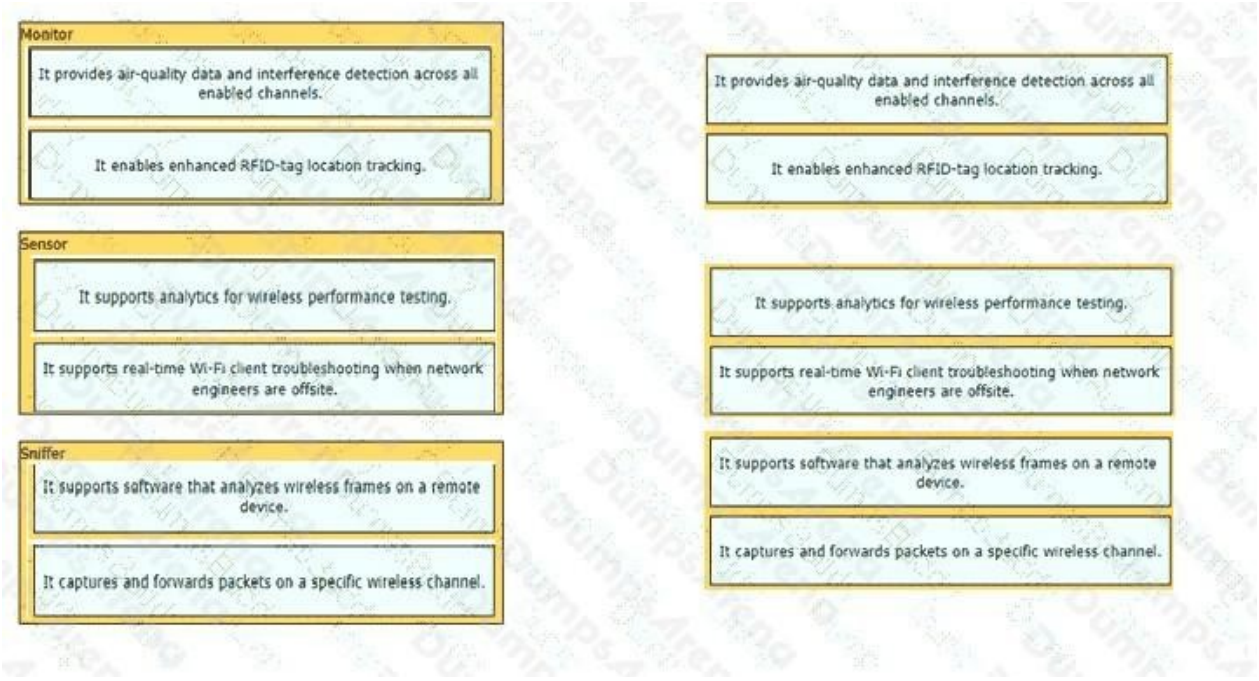
traffic for an IP phone. See Cisco's voice VLAN configuration guidance and the Cisco IOS LAN switching command reference.

QUESTION NO: 23 - (DRAG DROP)

Drag and drop the statements about access-point modes from the left onto the corresponding modes on the right.



ANSWER:



Explanation:

Cisco wireless access-point modes specialize an AP for different operational tasks instead of ordinary client access. In monitor mode, the AP listens across the RF environment rather than serving wireless clients. This lets the controller collect air-quality information and identify interference across enabled channels. Monitor-mode APs also contribute radio measurements used by location services, including more accurate tracking of RFID tags. These functions depend on broad RF observation and reporting, which is the core purpose of monitor mode. Cisco's documentation describes monitor-mode

APs as devices that scan channels and send radio-frequency and security information to the controller; see Cisco Wireless Controller Access Point Modes.

Sensor mode is intended for proactive assurance and testing of the wireless experience. A sensor AP can perform scheduled tests and provide telemetry used for wireless performance analytics. That same test-and-observability role supports remote troubleshooting of Wi-Fi client experience when an engineer is not physically onsite. Cisco Catalyst wireless assurance uses sensor-based testing to verify client onboarding and network-service reachability, helping operations teams detect and investigate experience problems remotely. Cisco provides an overview of this capability in its Catalyst 9100 Access Points data sheet.

Sniffer mode focuses on packet-level troubleshooting. The AP captures 802.11 traffic on a selected wireless channel and forwards those captured frames to a remote device running frame-analysis software, such as Wireshark or another compatible analyzer. This provides the detailed packet evidence required to inspect wireless exchanges, timing, management traffic, retransmissions, and other protocol behavior. Thus, the two wireless analytics and remote-troubleshooting statements belong with Sensor, the RF-quality and RFID-location statements belong with Monitor, and both packet-capture statements belong with Sniffer.

QUESTION NO: 24

Which AP mode is used for capturing wireless traffic and forwarding that traffic to a PC that is running a packet analyzer?

A. bridge

B. monitor

C. rouge detector

D. sniffer

ANSWER: D

Explanation:

sniffer is the AP mode intended for wireless packet capture and analysis. When an access point is placed into sniffer mode, it listens to traffic on a selected 802.11 radio channel and captures the over-the-air wireless frames. The AP then encapsulates those frames and forwards them across the wired network to a computer running compatible packet-analysis software, such as Wireshark with the appropriate Cisco capture support. This enables an administrator to inspect wireless management, control, and data traffic while troubleshooting client associations, authentication issues, roaming behavior, retransmissions, and RF-related problems.

A key characteristic of this operating mode is that the radio used for capture is dedicated to packet collection rather than normal client service. The analyzer PC must have network reachability to the AP or controller path used for the capture stream, and the capture configuration must select the relevant radio/channel so that the desired RF traffic is observed. Cisco documents sniffer mode as an access-point mode that captures packets and forwards them to a remote packet analyzer. See Cisco's Sniffer Mode on Cisco Access Points documentation and its Access Point Modes overview.

QUESTION NO: 25

How does MAC learning function on a switch?

A. protects against denial of service attacks

B. sends frames with unknown destinations to a multicast group

C. adds unknown source MAC addresses to the address table

D. sends a retransmission request when a new frame is received

ANSWER: C

Explanation:

MAC learning dynamically builds the switch's MAC address table, also called a CAM table, from the source address of each Ethernet frame received on a switch port. When a frame arrives, the switch reads its source MAC address and associates that address with the ingress interface and the frame's VLAN. If the source address does not yet have a dynamic entry, the switch creates one; subsequent frames from that source refresh the entry's aging timer. This enables the switch to make efficient forwarding decisions later: when it receives a frame addressed to a learned destination MAC address, it can send that frame only through the interface associated with that destination in the relevant VLAN. The address table is therefore populated from previously unseen source MAC addresses, rather than from destination addresses. Cisco describes dynamic MAC address learning as the process in which a switch learns source MAC addresses from incoming frames and adds them to its address table. See Cisco's LAN Switching and MAC Address Learning overview and the Catalyst MAC Address Table Configuration Guide.

QUESTION NO: 26

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

A. They prevent loops in the Layer 2 network.

B. They allow encrypted traffic.

C. They are able to bundle multiple ports to increase bandwidth.

D. They enable automatic failover of the default gateway.

E. They allow multiple devices to serve as a single virtual gateway for clients in the network.

ANSWER: D E

Explanation:

They enable automatic failover of the default gateway because First Hop Redundancy Protocols maintain a shared virtual gateway address while routers coordinate which device is currently forwarding traffic. End hosts use that virtual IP address as their configured default gateway, rather than needing to know the address of an individual physical router. When the active router or its tracked connectivity fails, an eligible standby router assumes the forwarding role for the virtual gateway. This provides gateway availability with little or no host-side reconfiguration.

They allow multiple devices to serve as a single virtual gateway for clients in the network because protocols such as Hot Standby Router Protocol, Virtual Router Redundancy Protocol, and Gateway Load Balancing Protocol present a common first-hop gateway identity to the LAN. This design removes a single physical default-router failure as a dependency for attached clients. Cisco describes HSRP as a way to provide network redundancy for IP networks by allowing several routers to work together and appear as a single virtual router to hosts. The virtual-gateway model is the core FHRP benefit tested in CCNA. See Cisco's HSRP Overview and Basic Concepts and the Cisco IOS XE FHRP Configuration Guide.

QUESTION NO: 27

After you deploy a new WLAN controller on your network, which two additional tasks should you consider? (Choose two.)

A. deploy load balancers

B. configure additional vlans

C. configure multiple VRRP groups

D. deploy POE switches

E. configure additional security policies

ANSWER: B D

Explanation:

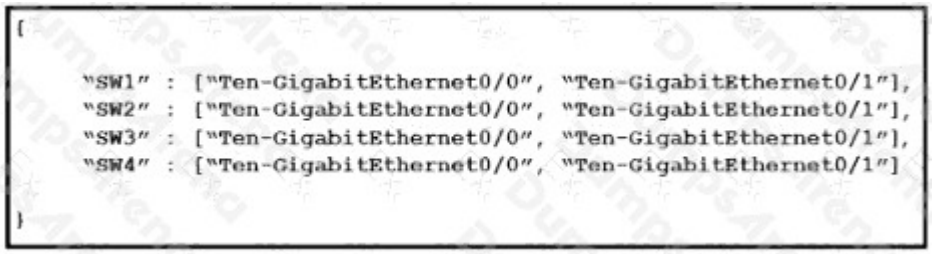
configure additional vlans is appropriate because wireless networks normally use VLANs to provide Layer 2 and Layer 3 segmentation for separate WLANs, such as corporate-user, guest, voice, and IoT services. On a controller-based Cisco wireless deployment, a WLAN is associated with a policy profile and VLAN mapping. The corresponding VLAN must exist across the switching infrastructure, be carried where required, and have the necessary IP addressing, DHCP scope, gateway, routing, and security services. This preparation allows clients joining each WLAN to receive the correct network placement and connectivity.

deploy POE switches is also an important deployment consideration because enterprise access points commonly receive operating power from their connected Ethernet switch port using Power over Ethernet. The access-layer switch must provide adequate PoE capacity for the number and model of access points, including their feature and radio-power requirements. Planning the PoE budget prevents access points from failing to power up or operating with reduced functionality. Cisco documents that PoE-capable switches can supply power to compatible powered devices, including wireless access points. Together, VLAN provisioning and sufficient PoE switching prepare the wired network to transport wireless client traffic and physically support the AP infrastructure managed by the new controller.

References: Cisco Enterprise Wireless; Cisco Catalyst 9300 Series Data Sheet.

QUESTION NO: 28

Refer to the exhibit.



How many JSON objects are represented?

A. 1

B. 2

C. 3

D. 4

ANSWER: D

Explanation:

The correct answer is **4**. In JSON, an object is a collection of name/value pairs delimited by curly braces (`{` and `}`). The exhibit contains four separately delimited brace-enclosed structures, so it represents four JSON objects. These objects may be nested within another structure or included as elements of an array, but each complete curly-brace pair that encloses a set of properties defines an object. Arrays, which use square brackets (`[` and `]`), organize values or objects but are not themselves JSON objects. Correctly distinguishing objects from arrays is important when interpreting API responses and configuring automation workflows, both of which are covered by the CCNA programmability and automation objectives. The JSON specification defines an object as an unordered collection of zero or more name/value pairs and identifies curly braces as its delimiters. See the JSON syntax definition in RFC 8259 and Cisco's overview of REST APIs and JSON data.

QUESTION NO: 29

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

A. All switches must be configured with the same VTP domain name

B. All switches must be configured to perform trunk negotiation

C. All switches must be configured with a unique VTP domain name

D. The VTP server must have the highest revision number in the domain

E. All switches must use the same VTP version

ANSWER: A E

Explanation:

“All switches must be configured with the same VTP domain name” is correct for switches that are intended to participate in the same VTP management domain. The VTP domain name identifies the VLAN database scope for VTP advertisements. A switch processes and synchronizes VLAN information only when the received VTP advertisement belongs to its configured VTP domain. Using a consistent, case-sensitive domain name on participating switches therefore establishes the common administrative boundary in which VLAN additions, deletions, and changes can be advertised across trunk links.

“All switches must use the same VTP version” is also the expected CCNA answer. A consistent VTP version provides predictable advertisement processing and VLAN database synchronization throughout the domain. This is especially important when using features introduced or changed by newer versions, such as the extended VLAN support and enhanced configuration controls associated with VTP version 3. Standardizing the version across participating switches avoids interoperability ambiguity and ensures that every device interprets VTP advertisements according to the same protocol behavior. Cisco recommends planning VTP mode, domain, version, and revision handling carefully before enabling VTP in a production network. See Cisco’s VTP overview and configuration-revision guidance and its VTP version 2 overview.

QUESTION NO: 30

When configuring a WLAN with WPA2 PSK in the Cisco Wireless LAN Controller GUI, which two formats are available to select? (Choose two.)

A. decimal

B. ASCII

C. hexadecimal

D. binary

E. base64

ANSWER: B C

Explanation:

For a WPA2-Personal WLAN, Cisco wireless controllers support defining the pre-shared key in either ASCII or hexadecimal format. An ASCII key is a user-entered passphrase, which WPA2 derives into the 256-bit Pairwise Master Key through the standard PBKDF2-based process. WPA/WPA2 passphrases use 8 through 63 printable ASCII characters. This representation is normally preferred for deployments because it is practical to enter and distribute while still allowing a suitably long, random passphrase.

A hexadecimal key represents the already-derived raw 256-bit pre-shared key. It must contain exactly 64 hexadecimal characters, with each pair of characters representing one byte. This format is useful where an organization needs to supply precise keying material rather than have the controller and clients derive it from a passphrase. Cisco documentation for controller-based WLAN security identifies the PSK input choices as ASCII and hexadecimal and describes the respective passphrase and 64-character hexadecimal-key requirements. See Cisco’s WPA/WPA2 WLAN security overview and the Cisco Wireless Controller WLAN security configuration guide.

QUESTION NO: 31

Refer to the exhibit.

An administrator must configure interfaces Gi1/1 and Gi1/3 on switch SW11 PC-1 and PC-2 must be placed in the Data VLAN and Phone-1 must be placed in the Voice VLAN Which configuration meets these requirements?

A. Option A

B. Option B

C. Option C

D. Option D

E. interface range gi1/1, gi1/3 switchport mode access switchport access vlan switchport voice vlan

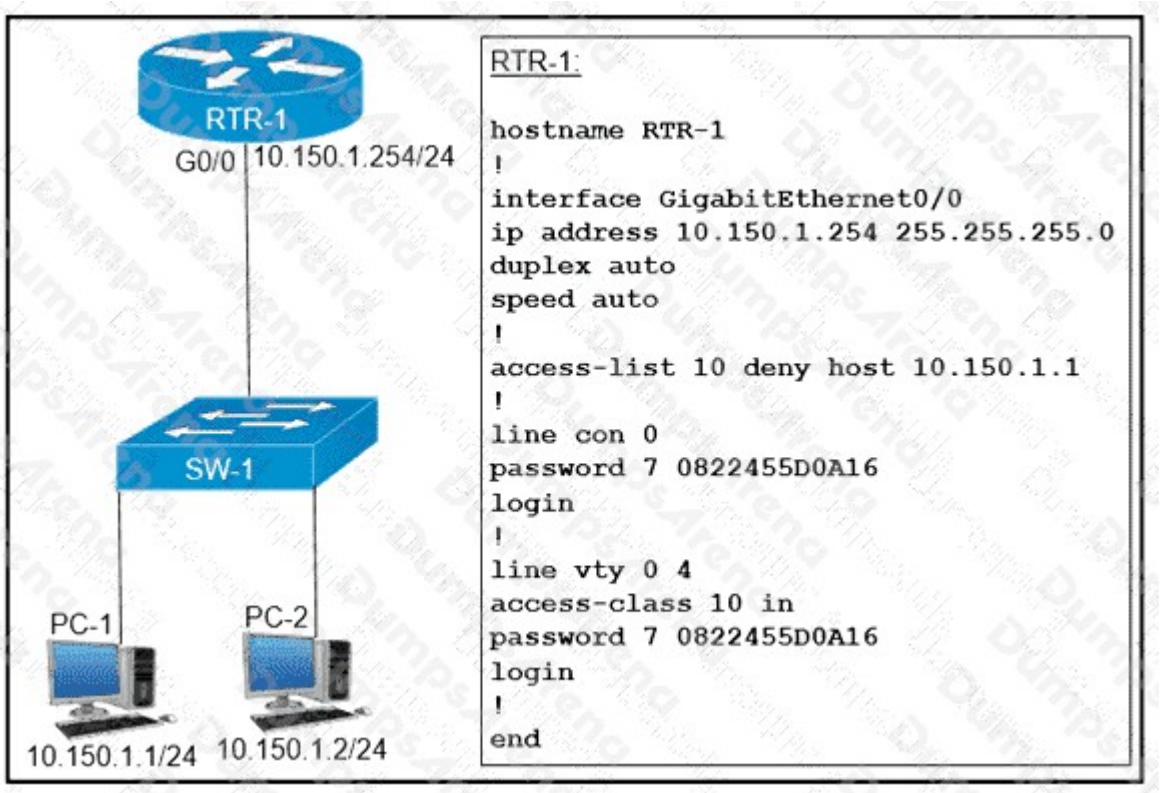
ANSWER: E

Explanation:

The required configuration makes both Gi1/1 and Gi1/3 Layer 2 access ports, assigns their untagged data traffic to the Data VLAN, and configures the Voice VLAN for IP-phone traffic. An access port supports a directly connected PC in its access VLAN. When an IP phone is connected, the switch advertises the configured auxiliary voice VLAN to the phone through CDP (or LLDP-MED where applicable). The phone then sends its voice traffic with an IEEE 802.1Q tag for that voice VLAN, while traffic from a PC connected through the phone remains untagged and is classified into the configured access/data VLAN. Applying the same settings with `interface range gi1/1, gi1/3` ensures both required switch interfaces have the intended data VLAN assignment and voice capability. Replace the placeholders with the actual VLAN IDs shown in the missing exhibit. Cisco documents the access VLAN command in its VLAN trunk and access-port configuration guide and describes voice-VLAN operation in its Voice VLAN configuration documentation.

QUESTION NO: 32

Refer to the exhibit.



An access list is created to deny Telnet access from host PC-1 to RTR-1 and allow access from all other hosts A Telnet attempt from PC-2 gives this message: " % Connection refused by remote host " Without allowing Telnet access from PC-1, which action must be taken to permit the traffic?

A. Add the access-list 10 permit any command to the configuration

B. Remove the access-class 10 in command from line vty 0.4.

C. Add the ip access-group 10 out command to interface g0/0.

D. Remove the password command from line vty 0 4.

ANSWER: A

Explanation:

Add the access-list 10 permit any command to the configuration. The ACL is applied inbound to the VTY lines with `access-class 10 in`, so it controls which source addresses may establish remote-management sessions to the router. ACL entries are evaluated from top to bottom, and processing stops at the first matching entry. The intended deny entry for PC-1 therefore remains first, preserving the requirement that PC-1 cannot use Telnet. However, every Cisco IPv4 ACL ends with an implicit `deny any` if no earlier entry matches. Consequently, PC-2 does not match the explicit PC-1 denial but is still denied by that implicit final rule, which results in the refused connection. Adding `access-list 10 permit any` after the PC-1 deny statement creates an explicit match for all remaining IPv4 source addresses. PC-2 and other hosts can then pass the VTY access-class check, while PC-1 continues to match the earlier deny statement and remains blocked. Cisco documents both the top-down, first-match ACL evaluation model and the implicit deny at the end of an ACL. See Cisco: Configuring Commonly Used IP ACLs and Cisco IOS IP ACL configuration guide.

QUESTION NO: 33

Which two must be met before SSH can operate normally on a Cisco IOS switch? (Choose two)

A. The switch must be running a k9 (crypto) IOS image

B. The ip domain-name command must be configured on the switch

C. IP routing must be enabled on the switch

D. A console password must be configured on the switch

E. Telnet must be disabled on the switch

ANSWER: A B

Explanation:

SSH on Cisco IOS requires an IOS image with cryptographic capabilities, traditionally identified by the *k9* designation. That capability supplies the cryptographic functions needed for the switch to act as an SSH server and establish encrypted management sessions. On current platforms, cryptographic support is generally included in the applicable IOS or IOS XE software package, but the underlying requirement remains that the installed software supports SSH and RSA cryptography.

The command `ip domain-name` must also be configured before generating the RSA key pair used by the IOS SSH server. IOS combines the configured hostname and domain name to create the key name during `crypto key generate rsa`. The RSA keys provide the server identity and cryptographic material required for SSH operation. Consequently, configuring a domain name and using an image with crypto support are prerequisite preparation steps for enabling SSH. Cisco's SSH configuration documentation describes generating RSA keys after setting the hostname and domain name, and identifies cryptographic software support as necessary for SSH functionality. See Cisco: Configure SSH on Routers and Switches and Cisco IOS XE Secure Shell Configuration Guide.

QUESTION NO: 34

Refer to the exhibit.

An engineer is configuring an EtherChannel using LACP between Switches 1 and 2 Which configuration must be applied so that only Switch 1 sends LACP initiation packets?

A. Swrtch2(config-if)#channel-group 1 mode passive

B. Switch2(config-if)#channel-group 1 mode active

C. Switch2(config-if)#channel-group 1 mode passive

D. Switch2(config-if)#channel-group 1 mode active

ANSWER: C

Explanation:

`Switch2(config-if)#channel-group 1 mode passive` is correct because LACP passive mode waits for LACP protocol packets from its neighbor and does not initiate LACP negotiation itself. To ensure that only Switch 1 originates the LACP exchange, Switch 1 must use LACP active mode while Switch 2 uses passive mode on the corresponding physical member interfaces. The active side transmits LACP Data Units (LACPDUs) to begin negotiation; the passive side responds after receiving them, allowing the devices to form the EtherChannel through LACP. This active/passive pairing is a valid LACP operational combination and provides the requested one-sided initiation behavior. The interfaces placed in the channel must also have compatible operational characteristics, such as matching speed, duplex, switchport mode, and VLAN/trunk settings, so that the bundled links can operate correctly after LACP establishes the channel. Cisco documents that `channel-group ... mode active` enables active LACP negotiation and `mode passive` enables passive LACP negotiation. See Cisco's EtherChannel configuration guide and the Cisco interface command reference.

QUESTION NO: 35

What are two descriptions of three-tier network topologies? (Choose two)

A. The core and distribution layers perform the same functions

B. The access layer manages routing between devices in different domains

C. The network core is designed to maintain continuous connectivity when devices fail.

D. The core layer maintains wired connections for each host

E. The distribution layer runs Layer 2 and Layer 3 technologies

ANSWER: C E

Explanation:

The network core is designed to maintain continuous connectivity when devices fail because the core layer is the campus backbone and is engineered for highly available, resilient, high-speed transport. Redundant paths and infrastructure help preserve connectivity and enable rapid recovery when an individual device or link fails. The core should forward traffic efficiently and reliably between distribution blocks, supporting uninterrupted communications across the enterprise campus.

The distribution layer runs Layer 2 and Layer 3 technologies because it aggregates access-layer switches and forms an important boundary for forwarding and policy decisions. At Layer 2, distribution devices can provide VLAN-related aggregation and loop-avoidance design boundaries. At Layer 3, they commonly provide inter-VLAN routing, route summarization, redundant default-gateway services, and policy enforcement such as filtering or quality-of-service classification. This combination lets the distribution layer connect endpoint-access networks to the resilient campus backbone while applying scalable control and routing functions. Cisco's campus design guidance describes the core as a resilient transport layer and the distribution layer as the aggregation and policy boundary between access and core. See Cisco's Campus Design Zone and Hierarchical Network Design documentation.

QUESTION NO: 36

Which two statements describe the use of YANG in network automation? (Choose two.)

A. It defines structured configuration and operational data models.

B. It can be used with NETCONF and RESTCONF.

C. It is a packet-forwarding protocol for switches.

D. It replaces TCP as the transport protocol for APIs.

E. It is an encryption algorithm used by SSH.

ANSWER: A B

Explanation:

YANG is a data modeling language used to define the configuration and operational data exposed by network devices. A YANG model describes data elements, their hierarchy, data types, and constraints, which allows management applications to interact with device data in a consistent structured format.

YANG models are commonly used with management protocols such as NETCONF and RESTCONF. NETCONF provides operations for retrieving and modifying configuration data, while RESTCONF exposes modeled data through a RESTful HTTP interface. RFC 7950 defines the YANG 1.1 data modeling language. See RFC 7950 and Cisco IOS XE YANG Data Models documentation.

QUESTION NO: 37

Which technology is appropriate for communication between an SDN controller and applications running over the network?

A. REST API

B. OpenFlow

C. Southbound API

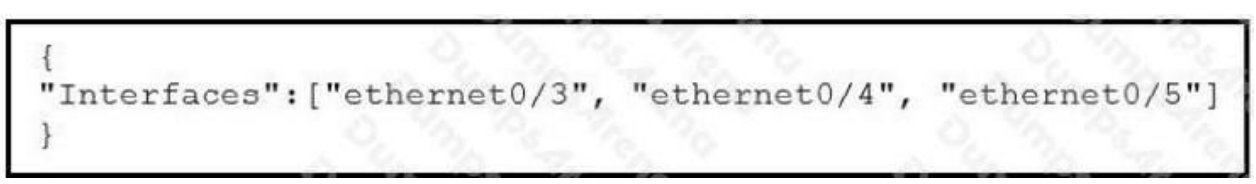
D. NETCONF

ANSWER: A

Explanation:

REST API is appropriate because applications communicate with an SDN controller through the controller's northbound interface. A RESTful API exposes controller resources and services over HTTP methods such as GET, POST, PUT, and DELETE, allowing orchestration platforms and custom applications to request topology information, obtain network state, and apply policy or automation intent programmatically. This model gives application developers a broadly interoperable, web-based interface while allowing the controller to translate application requirements into network configuration and forwarding behavior. In Cisco SDN architectures, the controller provides APIs so software can consume network capabilities and automate operations without needing to interact directly with each network device. REST APIs are especially useful for integrating applications with controller services because their resource-oriented approach and common data formats, such as JSON, are widely supported. Cisco's SDN overview describes the controller as the central software component that exposes capabilities to applications, and Cisco DevNet provides practical material on consuming RESTful interfaces for network automation. See Cisco Software-Defined Networking overview and Cisco DevNet Learning.

QUESTION NO: 38



Refer to the exhibit. Which type of JSON data is shown?

A. Boolean

B. string

C. object

D. sequence

ANSWER: C

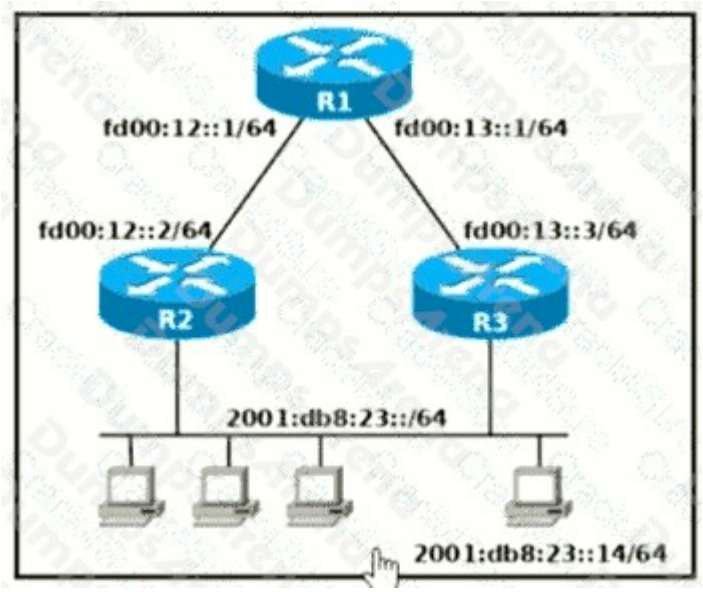
Explanation:

The JSON data shown is an **object**. A JSON object is delimited by opening and closing curly braces ({ }) and represents an unordered collection of named members. Each member consists of a property name enclosed in double quotation marks, followed by a colon and its associated value. This is the standard JSON representation for structured entities and is frequently used by REST APIs, including network automation APIs, to return attributes such as device names, interface properties, configuration values, and operational state.

JSON defines an object as a collection of zero or more name/value pairs. The value for an object member can itself be a string, number, Boolean, null, array, or another object, allowing API responses to represent complex hierarchical data. Identifying brace-enclosed name/value members as an object is therefore essential when reading or processing API payloads in Cisco automation workflows. The formal JSON specification is available in RFC 8259. For additional practical documentation and examples of JSON syntax and data structures, see the MDN JSON guide.

QUESTION NO: 39

Refer to the exhibit.



Which two commands, when configured on router R1, fulfill these requirements? (Choose two.)

Packets towards the entire network 2001:db8:2::/64 must be forwarded through router R2.

Packets toward host 2001:db8:23::14 preferably must be forwarded through R3.

A. Ipv6 route 2001:db8:23::/128 fd00:12::2

B. Ipv6 route 2001:db8:23::14/128 fd00:13::3

C. Ipv6 route 2001:db8:23::14/64 fd00:12::2

D. Ipv6 route 2001:db8:23::/64 fd00:12::2

E. Ipv6 route 2001:db8:23::14/64 fd00:12::2 200

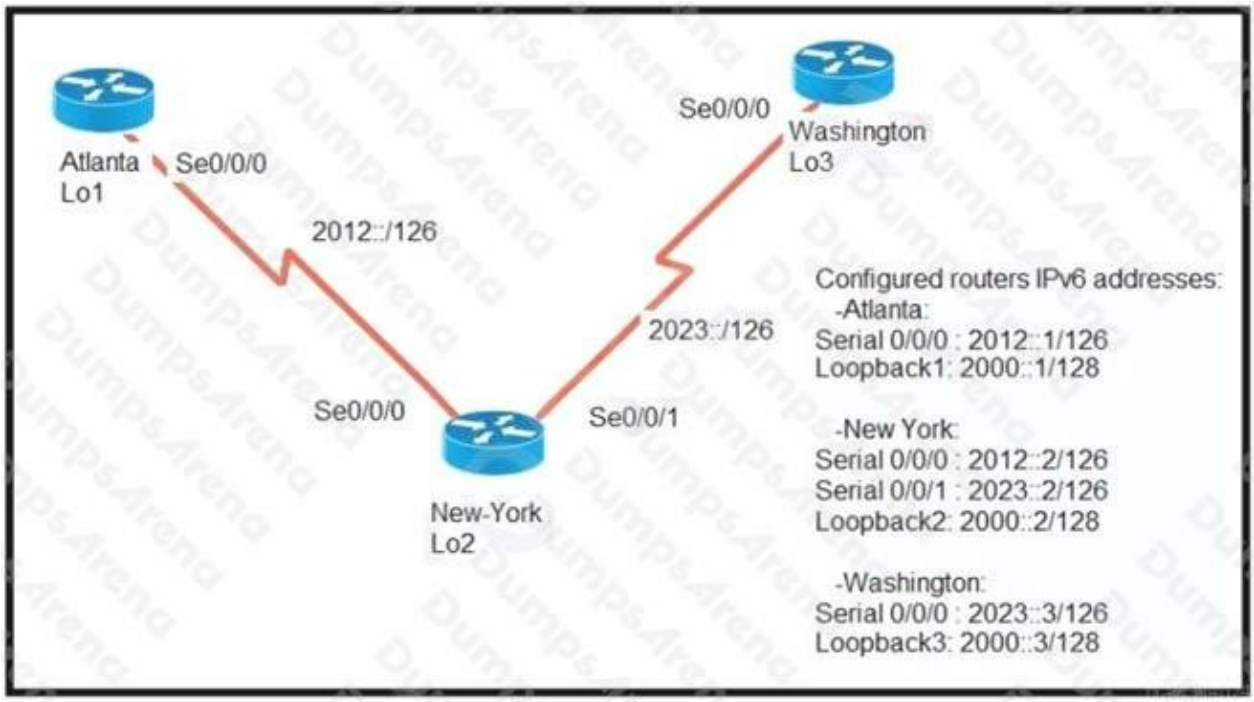
ANSWER: B D

Explanation:

The command `Ipv6 route 2001:db8:23::/64 fd00:12::2` installs a static route for the complete destination subnet through R2, using R2's IPv6 next-hop address. This provides the required network-wide forwarding path. The command `Ipv6 route 2001:db8:23::14/128 fd00:13::3` installs a host-specific IPv6 static route to the individual destination through R3. A `/128` prefix identifies exactly one IPv6 address, so it is more specific than the subnet's `/64` route.

Cisco routers select routes using longest-prefix matching before considering route preference values. Consequently, packets addressed to `2001:db8:23::14` match the `/128` host route and are sent to R3, while packets for other addresses in `2001:db8:23::/64` match the `/64` route and are sent to R2. This combination supplies both the general subnet path and the intended exception for the named host. Cisco documents IPv6 static-route configuration and next-hop addressing in its IPv6 Configuration Guide; route selection based on the most specific prefix is also covered in Cisco's static routing documentation.

QUESTION NO: 40



Refer to the exhibit. The New York router is configured with static routes pointing to the Atlanta and Washington sites.

Which two tasks must be performed so that the Se0/0/0 interfaces on the Atlanta and Washington routers can reach one another? (Choose two.)

A. Configure the ipv6 route 2023::/126 2012::1 command on the Atlanta router.

B. Configure the ipv6 route 2012::/126 2023::2 command on the Washington router.

C. Configure the ipv6 route 2012::/126 2023::1 command on the Washington router.

D. Configure the ipv6 route 2023::/126 2012::2 command on the Atlanta router.

E. Configure the ipv6 route 2012::/126 s0/0/0 command on the Atlanta router.

ANSWER: C D

Explanation:

The required tasks are to configure `ipv6 route 2012::/126 2023::1` on the Washington router and `ipv6 route 2023::/126 2012::2` on the Atlanta router. Atlanta and Washington are separate edge routers, with New York providing the transit path between their serial networks. New York already has routes toward both sites, but each edge router must also know how to forward traffic to the remote edge router's serial-link prefix.

Atlanta reaches the Washington serial network, `2023::/126`, by forwarding packets to New York's address on the Atlanta–New York link, `2012::2`. This creates a valid next-hop static route toward Washington. Similarly, Washington reaches the Atlanta serial network, `2012::/126`, by forwarding packets to New York's address on the Washington–New York link, `2023::1`. With these routes in place, traffic in either direction is forwarded through New York, and New York can use its existing static routes to complete delivery.

Cisco IOS IPv6 static routes identify a destination IPv6 prefix and a reachable next-hop IPv6 address, allowing routers to install forwarding information for remote networks. Cisco documents IPv6 static-route configuration and next-hop behavior at Cisco IPv6 Static Routing Configuration and Cisco IOS XE IPv6 Unicast Routing.

QUESTION NO: 41

Which command enables HTTP access to the Cisco WLC?

A. config network secureweb enable

B. config certificate generate web admin

C. config network webmode enable

D. config network telnet enable

ANSWER: C

Explanation:

config network webmode enable is the AireOS Cisco Wireless LAN Controller CLI command that enables the controller web server for standard HTTP GUI access. The *webmode* setting governs availability of the WLC's nonsecure web-management service, so enabling it permits an administrator to reach the web interface with an HTTP URL directed at the controller management address. This command applies to AireOS-based WLC platforms and software, which use the `config network` command hierarchy for management-protocol settings.

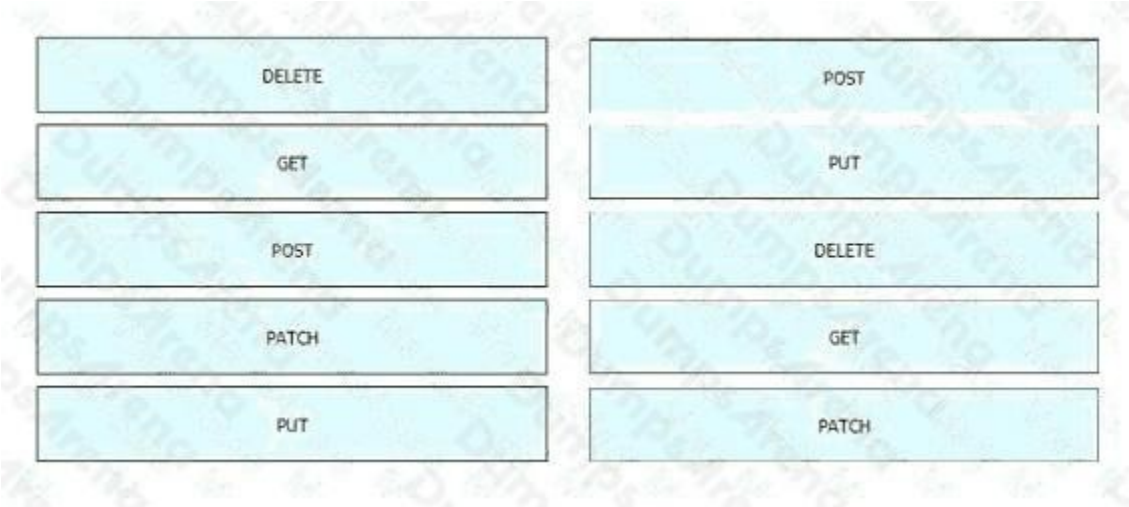
Although HTTP access can be useful for legacy or lab environments, Cisco management best practice is to use encrypted HTTPS access whenever possible because HTTP sends management traffic, including credentials and session information, without transport encryption. In production, administrators should limit management access to trusted networks and use the secure web-management configuration and valid certificates as appropriate. Cisco's Wireless LAN Controller command-reference resources document the available AireOS CLI commands, while its controller configuration-guide resources provide the related web-management deployment and security guidance. See Cisco Wireless LAN Controller Command References and Cisco Wireless LAN Controller Configuration Guides.

QUESTION NO: 42 - (DRAG DROP)

Drag and drop the HTTP methods used with REST-Based APIs from the left onto the descriptions on the right.



ANSWER:



Explanation:

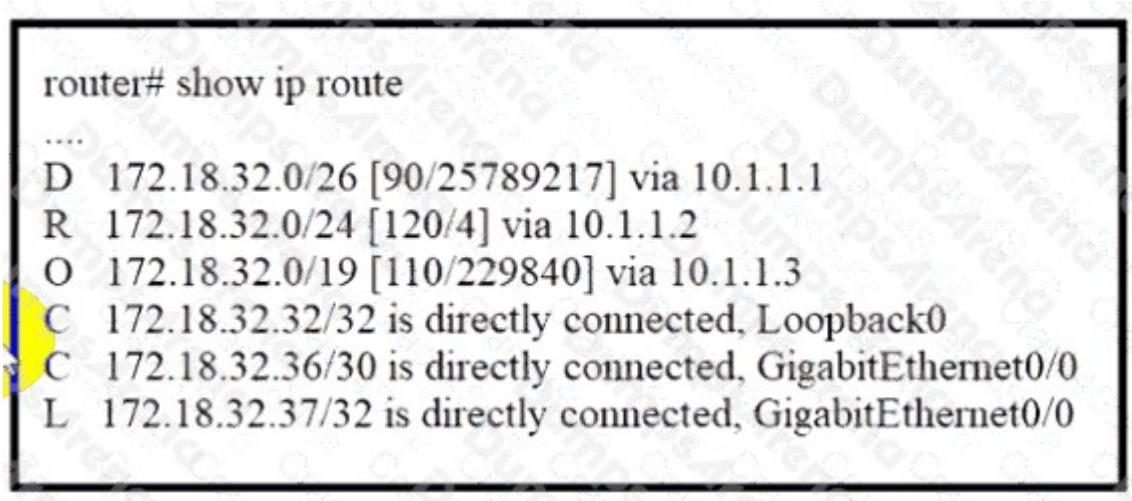
REST-based APIs use HTTP methods to express the intended operation on a resource identified by a URI. **POST** is the appropriate method when the client submits data to create a new resource under a collection endpoint. A successful creation commonly returns the URI of the newly created resource in the HTTP `Location` response header, which directly matches the creation description. The HTTP semantics specification describes POST as requesting that the target resource process the enclosed representation according to its own resource-specific semantics, including resource creation.

PUT is used to create or completely replace the representation of a resource at a known target URI. The request body contains the desired full state of that resource, so it matches the description concerning creation or replacement using information in the request body. **DELETE** requests removal of the resource identified by the target URI. **GET** retrieves a representation of the target resource without changing server state; when that target is a collection URI, the returned representation can contain the list of resource URIs or resource records. **PATCH** applies a set of partial modifications to an existing resource, with the requested changes carried in the request body rather than requiring a complete replacement representation.

These method meanings are defined by HTTP and are the conventions used by REST API designs. See the HTTP Semantics RFC 9110 for GET, POST, PUT, and DELETE, and RFC 5789 for the PATCH method.

QUESTION NO: 43

Refer to the exhibit.



A packet sourced from 172.18.33.2 is destined for 172.18.32.38. Where does the router forward the packet?

A. GigabitEthernet0/0

B. Loopback0

C. 10.1.1.1

D. 10.1.1.3

ANSWER: A

Explanation:

The router forwards the packet through **GigabitEthernet0/0**. Forwarding decisions are based primarily on the destination address, so the source address, 172.18.33.2, does not determine the outgoing route in this case. The router compares destination 172.18.32.38 with every matching routing-table prefix and selects the most-specific matching route, a process called longest-prefix match. The route associated with GigabitEthernet0/0 has the longest applicable prefix for this destination and therefore takes precedence over less-specific entries that may point to next-hop addresses. Administrative distance and route metric are considered only when multiple routes exist for the same prefix length; they do not override a more-specific route. After selecting this route, the router uses the Layer 2 information for GigabitEthernet0/0 to deliver the packet toward its destination. Cisco documents that routers choose the longest matching prefix when more than one route matches a destination address. See Cisco: IP Routing Explained and Cisco IOS IP Routing Configuration Guide.

QUESTION NO: 44

The SW1 interface g0/1 is in the down/down state. What are two reasons for the interface condition? (Choose two.)

A. There is a protocol mismatch

B. There is a duplex mismatch

C. The interface is shut down

D. The interface is error-disabled

E. There is a speed mismatch

ANSWER: C D

Explanation:

An interface that has been shut down is a valid cause of a down/down condition because the Cisco IOS `shutdown` command disables the port administratively. The interface cannot establish its physical Ethernet link while it is disabled, and its line protocol remains down as a consequence. Administratively disabling interfaces is a normal operational control used when a port is unused, during maintenance, or when connectivity must be intentionally prevented. Restoring service requires entering interface configuration mode and issuing `no shutdown`, provided the physical connection is otherwise healthy.

An interface that is error-disabled is also unavailable for forwarding and has its protocol down. Cisco switches place a port into the err-disabled state when a protective feature detects a condition requiring the port to be disabled, such as a port-security violation, BPDU Guard event, or a unidirectional-link condition. The port must be examined and the triggering condition resolved before it can be returned to operation; recovery may be performed manually or, where appropriate, through configured errdisable recovery. Cisco documents the relationship between interface status and physical/link operation in its Ethernet interface troubleshooting guide and describes disabled-port handling in its errdisable recovery documentation.

QUESTION NO: 45

Which two statements about eBGP neighbor relationships are true? (Choose two.)

A. The two devices must reside in different autonomous systems

B. Neighbors must be specifically declared in the configuration of each device

C. They can be created dynamically after the network statement is configured

D. The two devices must reside in the same autonomous system

E. The two devices must have matching timer settings

ANSWER: A B

Explanation:

An eBGP neighbor relationship is a BGP session between routers whose configured BGP autonomous system numbers differ. The local router's BGP process uses its own autonomous system number, while the neighbor configuration includes a `remote-as` value representing the peer's autonomous system. This differing-AS relationship is what identifies the session as external BGP. Cisco's BGP configuration guidance demonstrates eBGP peering by configuring a neighbor address and a remote autonomous system that is distinct from the local BGP autonomous system.

Neighbors must be specifically declared in the configuration of each device because ordinary BGP peering is configured explicitly, rather than discovered automatically. Each router needs sufficient neighbor configuration to initiate and accept the TCP-based BGP session, including the peer address and remote autonomous system. Once IP reachability exists and both peers have compatible BGP configuration, they establish the TCP connection and exchange BGP OPEN messages to form the relationship. This explicit configuration model provides administrative control over which external routers may exchange routing information.

See Cisco's BGP configuration fundamentals and the Cisco IOS XE BGP overview.

QUESTION NO: 46

Two switches have been implemented and all interfaces are at the default configuration level. A trunk link must be implemented between two switches with these requirements:

• using an industry-standard trunking protocol

• permitting VLANs 1 -10 and denying other VLANs

How must the interconnecting ports be configured?

A)

```
switchport mode dynamic
channel-protocol lacp
switchport trunk allowed vlans 1-10
```

B)

```
switchport mode trunk
switchport trunk encapsulation dot1q
switchport trunk allowed vlans 1-10
```

C)

```
switchport mode trunk
switchport trunk allowed vlans 1-10
switchport trunk native vlan 11
```

D)

```
switchport mode dynamic desirable
channel-group 1 mode desirable
switchport trunk encapsulation isl
switchport trunk allowed vlan except 11-4094
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: D

Explanation:

The configuration shown in **Option D** is correct because it explicitly establishes a trunk on the inter-switch link, uses IEEE 802.1Q encapsulation, and restricts the VLANs that can traverse that trunk to VLANs 1 through 10. IEEE 802.1Q is the open, industry-standard VLAN-tagging method used for Ethernet trunks. Configuring the interfaces as trunks ensures the link carries traffic for multiple VLANs rather than acting as a single access VLAN connection. Explicit trunk configuration is also appropriate when the switches begin with default interface settings, because it avoids relying on Dynamic Trunking Protocol negotiation to form the trunk.

The allowed-VLAN list is essential: `switchport trunk allowed vlan 1-10` permits only the stated inclusive VLAN range on the trunk. VLANs not included in that list are not forwarded across the inter-switch connection, satisfying the requirement to deny all other VLANs. Both ends of a trunk must use compatible trunking settings and matching permitted VLAN requirements so that VLAN traffic is carried consistently between the switches. Cisco documents trunk configuration and VLAN pruning behavior in its VLAN configuration guide, and IEEE 802.1Q is described by the IEEE 802.1Q standard.

QUESTION NO: 47

What is a function of a firewall on an enterprise network?

A. It allows and denies ingress and egress traffic.

B. It serves as a default gateway to hosts on the internet.

C. It processes traffic based on stateless inspection.

D. It acts as the intermediary device between the enterprise and its ISP.

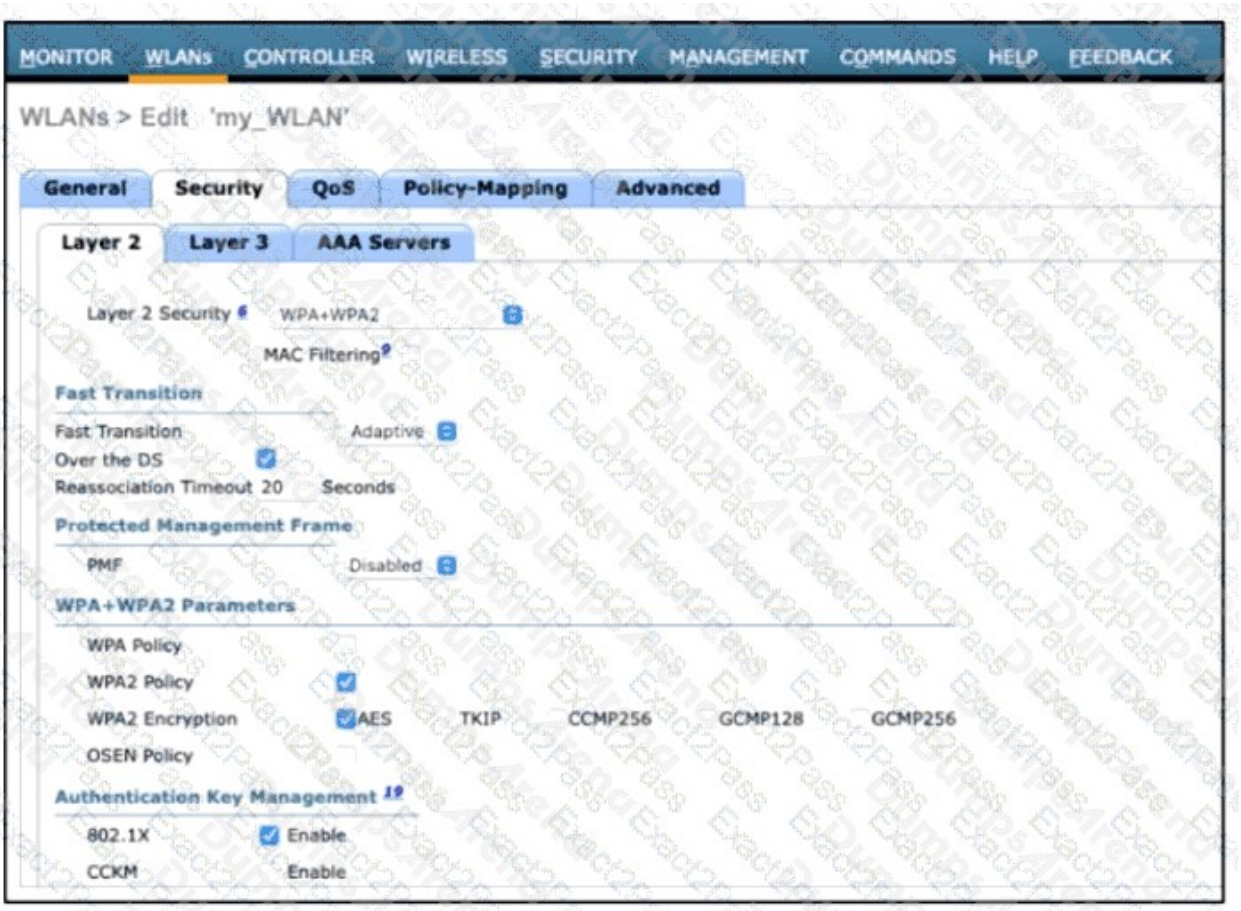
ANSWER: A

Explanation:

It allows and denies ingress and egress traffic. This is a fundamental firewall function because a firewall enforces an organization's security policy at a boundary between networks or security zones. It evaluates traffic entering the protected network (ingress) and traffic leaving it (egress), then permits, denies, or otherwise controls the traffic according to configured rules. Those rules can use criteria such as source and destination IP address, protocol, port number, connection state, user identity, or application characteristics, depending on the firewall's capabilities.

Controlling both directions is important in an enterprise environment. Inbound policy reduces exposure to unauthorized connections and services, while outbound policy can restrict access to unapproved destinations, limit data exfiltration paths, and allow only required business traffic. Stateful firewalls additionally track active connections so that return traffic for a permitted session can be handled appropriately. Cisco describes firewalls as security devices that monitor and control incoming and outgoing network traffic according to security rules. See Cisco: What Is a Firewall? and Cisco: Network Security.

QUESTION NO: 48



Refer to the exhibit. A network engineer is configuring a wirelesss LAN with Web Passthrough Layer 3 Web Policy. Which action must the engineer take to complete the configuration?

A. Set the Layer 2 Security to 802.1X.

B. Enable TKIP and CCMP256 WPA2 Encryption.

C. Set the Layer 2 Security to None.

D. Enable the WPA Policy.

ANSWER: C

Explanation:

Web Passthrough is a Layer 3 web policy that presents a web page to a client but does not require the client to enter credentials for network authentication. It is commonly used for guest-access terms-and-conditions, acceptable-use notices, or a click-through splash page. Because the WLAN is using a web-based Layer 3 policy rather than 802.1X or WPA/WPA2 enterprise authentication, the WLAN's Layer 2 security setting must be configured as **None**. This allows the wireless client to associate to the SSID without Layer 2 authentication or encryption being imposed before the controller can intercept the client's initial web request and display the configured passthrough page. After the user accepts the page, the controller permits the client's traffic according to the WLAN and any applicable access-control policies. This setup is specifically appropriate where the desired workflow is acknowledgement of a web page, rather than verification of a username and password. Cisco's wireless controller documentation describes web passthrough as a web-authentication method and identifies Layer 2 security configuration as part of WLAN security setup. See Cisco's Web Authentication configuration guidance and the Cisco web authentication configuration example.

QUESTION NO: 49

What is the function of northbound API?

A. It upgrades software and restores files.

B. It relies on global provisioning and configuration.

C. It supports distributed processing for configuration.

D. It provides a path between an SDN controller and network applications.

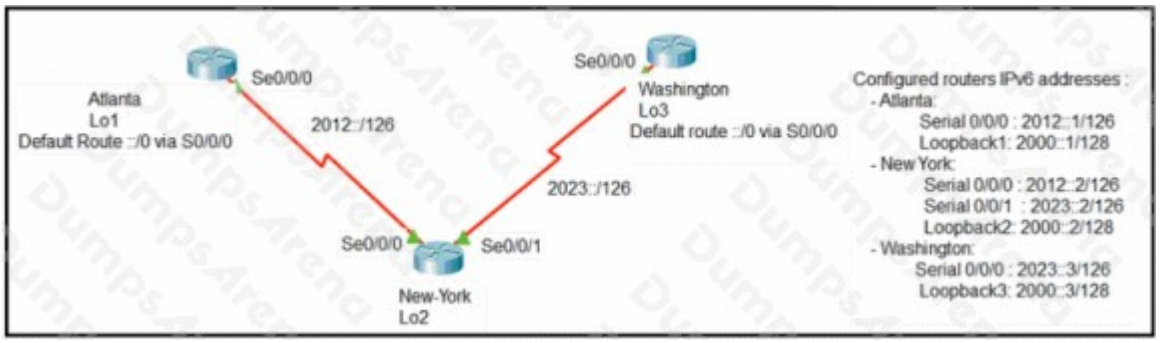
ANSWER: D

Explanation:

A northbound API provides the interface between an SDN controller and the applications or orchestration systems that need to request network services and express network intent. Through this interface, an application can communicate desired outcomes—such as connectivity requirements, security policy, quality-of-service needs, or automation workflows—to the controller. The controller then translates that high-level intent into appropriate network behavior and device-level configuration through its control-plane logic and, where applicable, southbound interfaces. This abstraction is central to software-defined networking because application developers do not need to interact directly with the individual switches, routers, or their device-specific command-line interfaces. Cisco describes controller APIs as enabling applications to program and automate network infrastructure, while its SDN architecture materials distinguish northbound communication with applications from controller-to-device communication. Therefore, the correct function is: “It provides a path between an SDN controller and network applications.” See Cisco’s Software-Defined Networking overview and the Cisco DNA Center API overview.

QUESTION NO: 50

Refer to Exhibit.



The loopback1 interface of the Atlanta router must reach the loopback3 interface of the Washington router. Which two static host routes must be configured on the NEW York router? (Choose two)

A. ipv6 route 2000::1/128 2012::1

B. ipv6 route 2000::3/128 2023::3

C. ipv6 route 2000::3/128 s0/0/0

D. ipv6 route 2000::1/128 2012::2

E. ipv6 route 2000::1/128 s0/0/1

ANSWER: A B

Explanation:

The required routes are `ipv6 route 2000::1/128 2012::1` and `ipv6 route 2000::3/128 2023::3`. A static host route uses a `/128` prefix, which identifies one specific IPv6 address rather than an entire subnet. New York therefore needs a route for Atlanta’s Loopback1 address, `2000::1`, and a separate route for Washington’s Loopback3 address, `2000::3`.

The next-hop address `2012::1` is the adjacent Atlanta-facing router address on the 2012:: link, so it directs traffic for Atlanta's loopback toward Atlanta. Likewise, `2023::3` is the adjacent Washington-facing router address on the 2023:: link, directing traffic for Washington's loopback toward Washington. These entries give New York explicit forwarding information for both remote loopback endpoints, which is necessary for traffic crossing the transit router in either direction.

Cisco IOS supports IPv6 static routes in the form `ipv6 route prefix/prefix-length next-hop-address`; a globally routable next-hop IPv6 address can be used directly when it is reachable through a connected interface. Cisco also documents static routing behavior and route configuration concepts in its IPv6 Routing Configuration Guide and static routing documentation.

QUESTION NO: 51

Which two network actions occur within the data plane? (Choose two.)

A. Add or remove an 802.1Q trunking header.

B. Make a configuration change from an incoming NETCONF RPC.

C. Run routing protocols.

D. Match the destination MAC address to the MAC address table.

E. Reply to an incoming ICMP echo request.

ANSWER: A D

Explanation:

The data plane, also called the forwarding plane, performs the per-frame and per-packet work needed to move transit traffic through a network device. It uses forwarding information that has already been prepared by the control plane and applies that information at high speed, commonly in switching or forwarding hardware. **Add or remove an 802.1Q trunking header.** is a data-plane operation because it modifies the Layer 2 encapsulation as traffic is forwarded between VLAN-aware interfaces. For example, a switch can add a VLAN tag when sending a frame onto a trunk or remove that tag when sending the frame out an access port.

Match the destination MAC address to the MAC address table. is also a fundamental Layer 2 data-plane action. The switch examines the destination MAC address, looks it up in its forwarding/MAC address table, and uses the resulting forwarding decision to select the appropriate egress interface or interfaces. This lookup and forwarding process must occur for each received frame, making it a core forwarding-plane responsibility. Cisco describes the forwarding plane as the component that moves packets based on information installed by control-plane processes; see Cisco IOS IP Forwarding Configuration Guide and Cisco LAN Switching technical resources.

QUESTION NO: 52

Which two tasks must be performed to configure NTP to a trusted server in client mode on a single network device? (Choose two)

A. Enable NTP authentication.

B. Verify the time zone.

C. Disable NTP broadcasts

D. Specify the IP address of the NTP server

E. Set the NTP server private key

ANSWER: A D

Explanation:

“Specify the IP address of the NTP server” is required because an IOS/IOS XE device operating as an NTP client must be configured with a synchronization source. The `ntp server` command identifies the server by IPv4 address, IPv6 address, or hostname and causes the device to form a normal unicast NTP association with that source. This is the essential client-mode configuration action that tells the device where to obtain time.

“Enable NTP authentication.” is also required when the NTP source is described as trusted. Cisco NTP authentication protects time synchronization by validating NTP packets with a configured shared key. In practical IOS configuration, this involves enabling authentication globally with `ntp authenticate`, defining the applicable authentication key, identifying that key as trusted with `ntp trusted-key`, and associating the key with the configured NTP server. These steps ensure that the client accepts time information only from a source that can authenticate using the expected key material. Accurate, authenticated time is important for correlated logging, certificate validation, and security-event analysis. Cisco’s NTP configuration guidance documents the use of `ntp server` for a unicast source and the authentication and trusted-key commands for authenticated NTP deployments. See Cisco’s NTP configuration and troubleshooting guide and the Cisco IOS XE NTP Configuration Guide.

QUESTION NO: 53

What are two reasons for an engineer to configure a floating static route? (Choose two.)

A. to enable fallback static routing when the dynamic routing protocol fails

B. to route traffic differently based on the source IP of the packet

C. to automatically route traffic on a secondary path when the primary path goes down

D. to support load balancing via static routing

E. to control the return path of traffic that is sent from the router

ANSWER: A C

Explanation:

A floating static route provides a backup route for a destination. An engineer configures it with an administrative distance that is higher than that of the preferred, primary route source. Consequently, as long as a route learned through a dynamic routing protocol is available, the router selects that lower-administrative-distance route and keeps the floating static route out of the active routing table. If the dynamic routing protocol fails to advertise the destination or the primary route is withdrawn, the router can install the static route and continue forwarding traffic.

Therefore, *to enable fallback static routing when the dynamic routing protocol fails* is a valid use: the floating route preserves connectivity when the dynamically learned route is no longer usable. Likewise, *to automatically route traffic on a secondary path when the primary path goes down* accurately describes the failover role of a floating static route. This approach is commonly used where a primary WAN connection is backed up by another link, while avoiding use of the backup path during normal operation. Cisco documents that a floating static route is created by assigning a higher administrative distance than the route it is intended to back up. See Cisco's floating static-route documentation and the Cisco IOS XE static-route configuration guide.

QUESTION NO: 54

Which two features introduced in SNMPv2 provide the ability to retrieve large amounts of data in one request and acknowledge a trap using PDUs? (Choose two.)

A. Get

B. GetNext

C. Set

D. GetBulk

E. Inform

ANSWER: D E

Explanation:

GetBulk is the SNMPv2 PDU designed for efficient retrieval of substantial management data, particularly rows in tabular MIB objects. A manager can specify non-repeaters and the maximum number of repetitions, allowing an agent to return multiple lexicographically successive variable bindings in one response. This substantially reduces the number of request/response exchanges needed to walk a table compared with repeatedly sending GetNext requests.

Inform, formally the InformRequest PDU, provides a reliable notification mechanism. The receiving SNMP manager responds with a Response PDU, so the sender can determine whether the notification was received. This acknowledgement capability is useful when delivery confirmation for an operational event matters. SNMPv2 introduced both GetBulkRequest and InformRequest as protocol operations, addressing scalable data collection and acknowledged event notification respectively.

RFC 3416 defines the SNMP protocol operations and specifies that a GetBulkRequest retrieves large blocks of management information efficiently, while an InformRequest expects a Response PDU from its receiver. Cisco's SNMP configuration guidance also distinguishes informs from traps by noting that informs are acknowledged notifications. See RFC 3416: Version 2 of the Protocol Operations for SNMP and Cisco IOS XE SNMP Support.

QUESTION NO: 55

What is a characteristic of private IPv4 addressing?

A. Reduces the forwarding table on network routers

B. Used on the external interface of a firewall

C. Used by ISPs when only one IP is needed to connect to the internet

D. Address space which is isolated from the internet

ANSWER: D

Explanation:

Address space which is isolated from the internet is correct because private IPv4 addresses are intended for use within an organization's internal network rather than as globally reachable addresses on the public internet. RFC 1918 reserves three IPv4 blocks for this purpose: 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. These addresses may be reused independently by different organizations because they are not globally unique and public internet routers do not advertise or forward routes for them. This isolation supports internal addressing plans while preserving scarce globally routable IPv4 address space.

When hosts using private addresses need access to internet resources, an edge device commonly performs Network Address Translation, translating the private source address to a public, globally routable address. The host can then communicate outward without its private address becoming reachable or routable across the public internet. This distinction between private, internally used addressing and public, internet-routable addressing is a foundational IPv4 addressing concept tested in CCNA. See RFC 1918: Address Allocation for Private Internets and Cisco's Network Address Translation overview.

QUESTION NO: 56

Which IPsec mode encapsulates the entire IP packet?

A. tunnel

B. Q-in-Q

C. SSL VPN

D. transport

ANSWER: A

Explanation:

The correct answer is **tunnel**. In IPsec tunnel mode, the complete original IP packet—including its original IP header and payload—is encapsulated within a new IP packet. IPsec applies protection to the original packet, then adds a new outer IP header for delivery between the tunnel endpoints. This approach conceals the original source and destination IP addresses from networks that carry the encrypted traffic and is commonly used for site-to-site VPNs, where security gateways protect traffic between separate private networks.

This behavior differs from protecting only the upper-layer payload: tunnel mode carries the original packet as the protected inner packet and uses the newly added outer header to route it across an untrusted network. Cisco describes IPsec tunnel mode as encrypting and encapsulating the entire original IP datagram, which makes it appropriate when traffic must traverse a VPN tunnel between peers. See Cisco's IPsec virtual tunnel interface documentation and the Cisco IPsec VPN overview for details on IPsec tunneling and packet protection.

QUESTION NO: 57

Which QoS queuing method discards or marks packets that exceed the desired bit rate of traffic flow?

A. shaping

B. policing

C. CBWFQ

D. LLQ

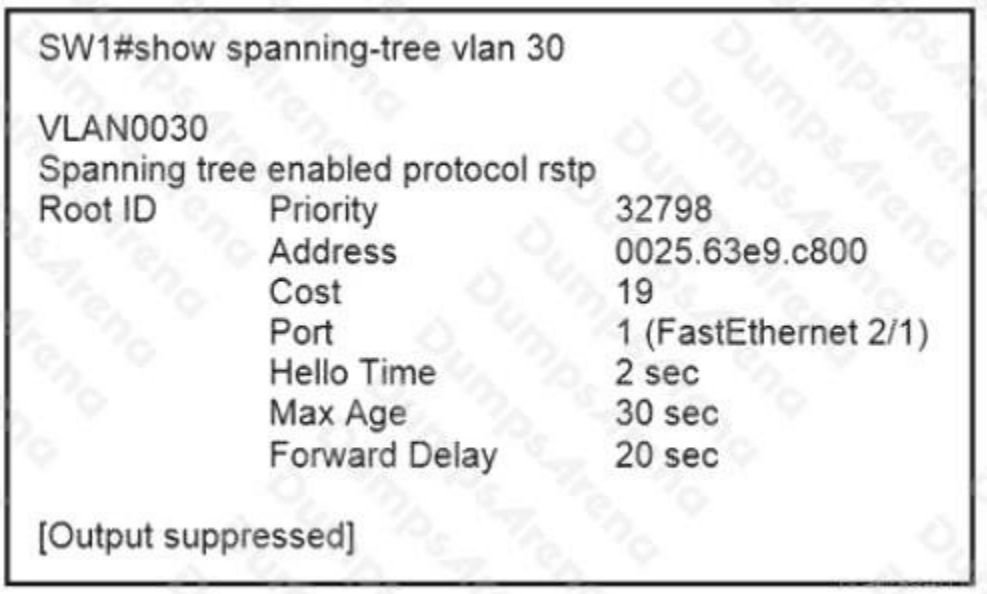
ANSWER: B

Explanation:

Policing is the correct answer because it enforces a configured traffic-rate limit by measuring packets against a defined rate, typically with a token-bucket algorithm. Traffic that conforms to the configured rate is forwarded according to policy. When packets exceed the allowed rate, policing can take an immediate action such as discarding the excess packets or marking them by changing QoS fields such as IP precedence, DSCP, or CoS. This behavior makes policing appropriate when a device must strictly enforce a traffic contract or limit the rate at a network boundary. Cisco distinguishes policing from mechanisms that schedule packets for transmission: policing is a rate-enforcement mechanism and does not ordinarily buffer excess traffic for later transmission. Its configured actions for conforming and exceeding traffic provide the discard-or-mark behavior described in the question. Cisco's QoS documentation describes policing as controlling traffic by dropping or remarking packets that exceed the configured rate, while the Cisco IOS XE configuration guide documents the available policing actions and rate parameters. See Cisco: Policing versus Shaping and Cisco IOS XE QoS Policing and Shaping Configuration Guide.

QUESTION NO: 58

Refer to the exhibit. What two conclusions should be made about this configuration? (Choose two.)



A. The root port is FastEthernet 2/1

B. The designated port is FastEthernet 2/1

C. The spanning-tree mode is PVST+

D. This is a root bridge

E. The spanning-tree mode is Rapid PVST+

ANSWER: A E

Explanation:

The correct conclusions are that the root port is FastEthernet 2/1 and that the spanning-tree mode is Rapid PVST+. In Spanning Tree Protocol, every nonroot switch selects one root port for each VLAN or spanning-tree instance. This is the local interface with the best path toward the root bridge, determined primarily by the lowest accumulated root-path cost. The exhibit identifies FastEthernet 2/1 with the root-port role, so it is the interface this switch uses to reach the root bridge for the displayed instance.

The configuration/output also indicates that the switch is using Rapid PVST+. Rapid PVST+ is Cisco's per-VLAN implementation of Rapid Spanning Tree Protocol and provides faster convergence than the original IEEE 802.1D STP behavior. Cisco IOS displays this operating mode as `rapid-pvst` in relevant spanning-tree command output. Because the mode applies per VLAN, the switch independently maintains the STP topology and port roles for each VLAN while using rapid-transition mechanisms. Cisco documents STP port roles and root-port selection in its Spanning Tree Protocol overview and describes rapid convergence behavior in its Rapid Spanning Tree Protocol overview.

QUESTION NO: 59

What is the root port in STP?

A. It is the port with the highest priority toward the root bridge.

B. It is the port on the root switch that leads to the designated port on another switch.

C. It is the port that is elected only when the root bridge has precisely one port on a single LAN segment.

D. It is the port on a switch with the lowest cost to reach the root bridge.

ANSWER: D

Explanation:

In STP, the root port is the single port selected on each non-root bridge that provides that bridge's best path toward the root bridge. The fundamental selection criterion is the lowest root path cost: a switch adds the cost of its receiving port to the root path cost advertised in received BPDUs, and the port yielding the lowest total is chosen as the root port. This role gives every non-root switch a loop-free forwarding path in the spanning-tree topology. The root bridge itself does not elect a root port, because it is the reference point to which all other switches calculate their paths. Where equal path costs exist, STP applies defined BPDU comparison tie-breakers, including the upstream bridge ID, upstream port ID, and local port ID, to select one deterministic root port. Therefore, "It is the port on a switch with the lowest cost to reach the root bridge." accurately expresses the core CCNA-level definition. Cisco's STP documentation describes a root port as the port offering the lowest-cost path to the root bridge and explains its role in the elected spanning-tree topology. See Cisco's Spanning Tree Protocol overview and the Cisco IOS XE STP overview.

QUESTION NO: 60

PC1 tries to send traffic to newly installed PC2. The PC2 MAC address is not listed in the MAC address table of the switch, so the switch sends the packet to all ports in the same VLAN Which switching concept does this describe?

A. MAC address aging

B. MAC address table

C. frame flooding

D. spanning-tree protocol

ANSWER: C

Explanation:

frame flooding, specifically unknown-unicast flooding, is the correct switching concept. A Layer 2 switch learns MAC addresses by examining the source MAC address of frames received on each port and recording that address-to-port association in its MAC address table for the relevant VLAN. When PC1 sends an Ethernet frame to PC2 and the switch has no entry for PC2's destination MAC address, it cannot select one specific egress interface. The switch therefore floods the frame through all forwarding ports in that VLAN except the port on which it received the frame. This gives PC2 an opportunity to receive the frame and reply. When PC2 transmits a reply, the switch observes PC2's source MAC address on its connected port and adds that information to the MAC address table. Subsequent frames addressed to PC2 can then be forwarded only through that learned port rather than being flooded. Flooding is VLAN-scoped, so the frame is not forwarded into other VLANs. Cisco describes this destination lookup behavior as forwarding unknown unicast frames out all ports in the VLAN other than the incoming port. See Cisco's MAC Address Table documentation and the Cisco guide to configuring MAC address tables.

QUESTION NO: 61

A frame that enters a switch fails the Frame Check Sequence. Which two interface counters are incremented? (Choose two.)

A. input errors

B. frame

C. giants

D. CRC

E. runts

ANSWER: A D

Explanation:

When a switch receives an Ethernet frame, it calculates a cyclic redundancy check over the received frame and compares the result with the Frame Check Sequence carried in the Ethernet trailer. A mismatch means the frame was corrupted in

transit or otherwise arrived with an invalid FCS. Cisco interface statistics record that specific condition in the *CRC* counter. Thus, the *CRC* counter directly identifies frames whose FCS validation failed.

The *input errors* counter is also incremented because it is the aggregate receive-error counter. Cisco documents input errors as including receive-side problems such as CRC errors, frame errors, overruns, ignored packets, and related conditions. Therefore, a frame that fails FCS contributes both to the detailed CRC statistic and to the interface's overall input-error total. Monitoring both values is useful operationally: the CRC count identifies the error type, while input errors shows its contribution to the broader receive-error rate. Cisco notes that increasing CRC/input-error values commonly warrant investigation of the physical Ethernet path, including cabling, connectors, and duplex-related conditions. See Cisco's Understanding Ethernet Interface Counters and Troubleshooting CRC/Input Errors.

QUESTION NO: 62



Refer to the exhibit. What are the two steps an engineer must take to provide the highest encryption and authentication using domain credentials from LDAP?

(Choose two.)

A. Select PSK under Authentication Key Management.

B. Select Static-WEP + 802.1X on Layer 2 Security.

C. Select WPA+WPA2 on Layer 2 Security.

D. Select 802.1X from under Authentication Key Management.

E. Select WPA Policy with TKIP Encryption.

ANSWER: C D

Explanation:

Selecting **WPA+WPA2 on Layer 2 Security** enables the WLAN to use WPA2 security, which supports AES-CCMP encryption. AES-CCMP is the robust encryption suite associated with WPA2 Enterprise deployments and provides the appropriate modern Layer 2 protection for this configuration. In a production deployment intended to provide the strongest available protection, the WPA2 policy is used with AES rather than legacy cipher suites.

Selecting **802.1X from under Authentication Key Management** configures Enterprise authentication rather than a shared wireless password. With 802.1X, the wireless client authenticates through EAP to a RADIUS/AAA service. That AAA service can validate user identity against an LDAP-backed directory, such as Microsoft Active Directory. Successful authentication also results in per-user, dynamically derived encryption keys, tying network access and key establishment to the authenticated domain identity.

Together, WPA2-capable Layer 2 security and 802.1X provide the required Enterprise WLAN design: centralized directory-based user authentication through RADIUS and strong encryption using dynamically managed session keys. Cisco documents WPA/WPA2 Enterprise operation and its relationship to 802.1X in its WPA and WPA2 wireless security overview; Cisco also describes the wireless 802.1X authentication workflow.

QUESTION NO: 63 - (SIMULATION)

Configure IPv4 and IPv6 connectivity between two routers. For IPv4, use a /28 network from the 192.168.1.0/24 private range. For IPv6, use the first /64 subnet from the 2001:0db8:aaaa::/48 subnet.

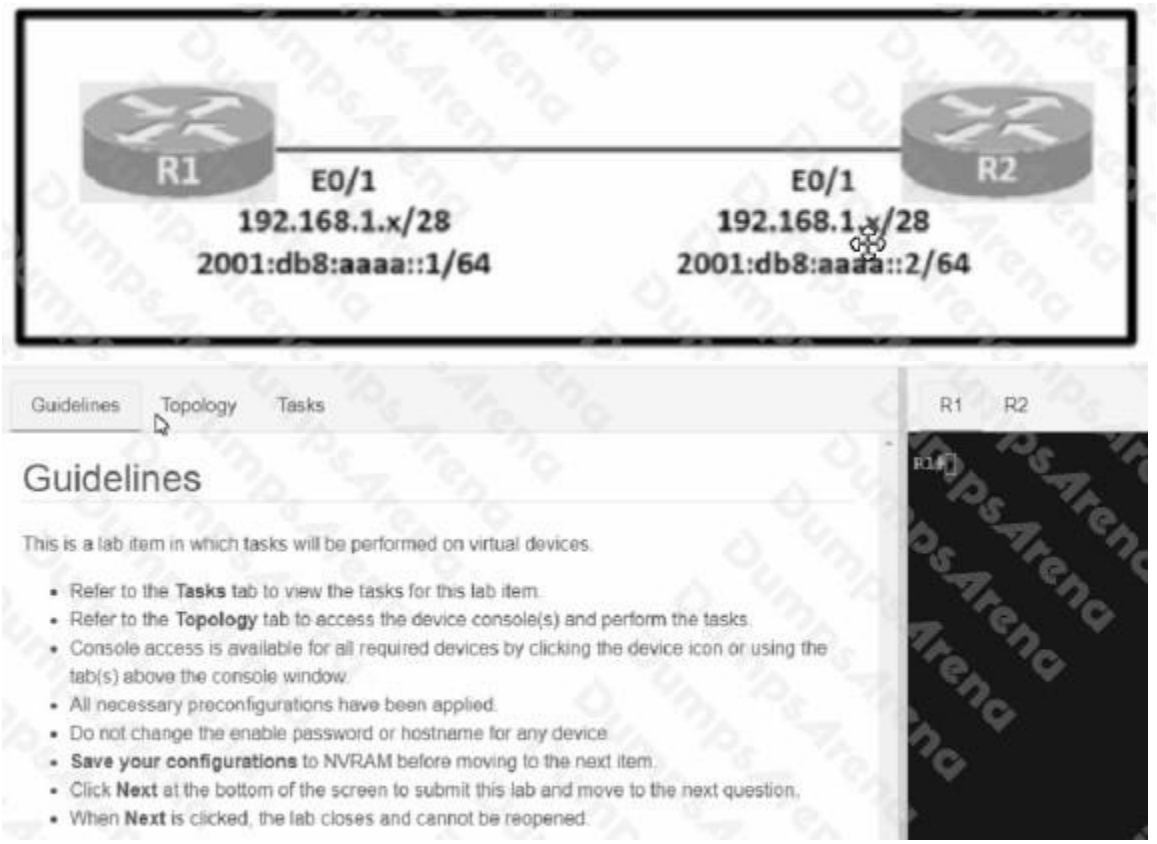
1. Using Ethernet0/1 on routers R1 and R2, configure the next usable/28 from the 192.168.1.0/24 range. The network 192.168.1.0/28 is unavailable.

2. For the IPv4 /28 subnet, router R1 must be configured with the first usable host address.

3. For the IPv4 /28 subnet, router R2 must be configured with the last usable host address.

4. For the IPv6 /64 subnet, configure the routers with the IP addressing provided from the topology.

5. A ping must work between the routers on the IPv4 and IPv6 address ranges.



ANSWER: See the explanation for the answer

Explanation:

The unavailable IPv4 subnet is `192.168.1.0/28`, so the next /28 subnet is `192.168.1.16/28` (mask `255.255.255.240`). Its usable host range is `192.168.1.17` through `192.168.1.30`.

Configure R1 with the first IPv4 host and the topology-provided IPv6 address:

```
R1# configure terminal
R1(config)# ipv6 unicast-routing
R1(config)# interface ethernet0/1
R1(config-if)# ip address 192.168.1.17 255.255.255.240
R1(config-if)# ipv6 address 2001:db8:aaaa::1/64
R1(config-if)# no shutdown
R1(config-if)# end
R1# copy running-config startup-config
```

Configure R2 with the last IPv4 host and the topology-provided IPv6 address:

```
R2# configure terminal
R2(config)# ipv6 unicast-routing
R2(config)# interface ethernet0/1
R2(config-if)# ip address 192.168.1.30 255.255.255.240
R2(config-if)# ipv6 address 2001:db8:aaaa::2/64
R2(config-if)# no shutdown
R2(config-if)# end
R2# copy running-config startup-config
```

Verify from R1:

```
R1# ping 192.168.1.30
R1# ping ipv6 2001:db8:aaaa::2
```

Alternatively, from R2, ping `192.168.1.17` and `2001:db8:aaaa::1`.

QUESTION NO: 64



Refer to the exhibit. Users need to connect to the wireless network with IEEE 802.11r-compatible devices. The connection must be maintained as users travel between floors or to other areas in the building. What must be the configuration of the connection?

A. Disable AES encryption.

B. Enable Fast Transition and select the FT 802.1x option.

C. Enable Fast Transition and select the FT PSK option.

D. Select the WPA Policy option with the CCKM option.

ANSWER: B

Explanation:

Enable Fast Transition and select the FT 802.1x option. is the required configuration for an enterprise WLAN using 802.1X authentication. IEEE 802.11r, also called Fast BSS Transition, is designed to reduce the interruption that occurs when a wireless client roams from one access point to another in the same mobility domain. Rather than completing the full authentication and key-establishment process after associating to the next access point, an 802.11r-capable client can derive and use roaming keys in advance. This substantially reduces handoff time, which helps preserve active voice, video, and other real-time application sessions while a user moves between floors or coverage areas.

On Cisco wireless controllers, Fast Transition must be enabled on the WLAN and matched to its authentication model. FT 802.1X is the Fast Transition method for a WLAN whose users authenticate through 802.1X/EAP with an AAA or RADIUS service. It preserves enterprise identity-based authentication while enabling the standardized 802.11r roaming process for compatible clients. Cisco documents that Fast Transition supports distinct modes for 802.1X and PSK WLAN security configurations; selecting the 802.1X mode is therefore the appropriate enterprise configuration here. See Cisco's 802.11r Fast Transition overview and the Cisco WLAN security configuration guide.

QUESTION NO: 65

Refer to the exhibit. An LACP EtherChannel between two directly connected switches is in the configuration process.

Which command must be configured on switch SW2's Gi0/1-2 interfaces to establish the channel to SW1?

A. channel-group 1 mode desirable

B. channel-group 1 mode on

C. channel-group 1 mode auto

D. channel-group 1 mode active

ANSWER: D

Explanation:

`channel-group 1 mode active` is the required configuration because it enables Link Aggregation Control Protocol negotiation on SW2 and actively sends LACP protocol data units on Gi0/1 and Gi0/2. An LACP EtherChannel forms when LACP negotiation is enabled and the directly connected peer is participating in LACP; active mode initiates that negotiation and interoperates with a peer configured for either active or passive LACP mode. The channel-group number identifies the local EtherChannel interface, so using group 1 creates or associates the member links with Port-channel1 on SW2. LACP is the IEEE standard link-aggregation protocol, standardized as IEEE 802.1AX, and Cisco IOS uses the active and passive channel-group modes to select it. Correctly bundling both physical interfaces provides a single logical port channel, allowing traffic to be distributed across the eligible member links while presenting one logical Layer 2 interface to switching protocols. Cisco documents active mode as the setting that actively attempts to form an LACP EtherChannel; see Cisco EtherChannel configuration guide and the Cisco EtherChannel overview.

QUESTION NO: 66

Which two statements describe Link Aggregation Control Protocol? (Choose two.)

A. It is defined by IEEE 802.3ad and IEEE 802.1AX standards.

B. It supports active and passive negotiation modes.

C. It uses desirable and auto negotiation modes.

D. It requires all member links to use different switchport settings.

E. It operates only between Cisco switches.

ANSWER: A B

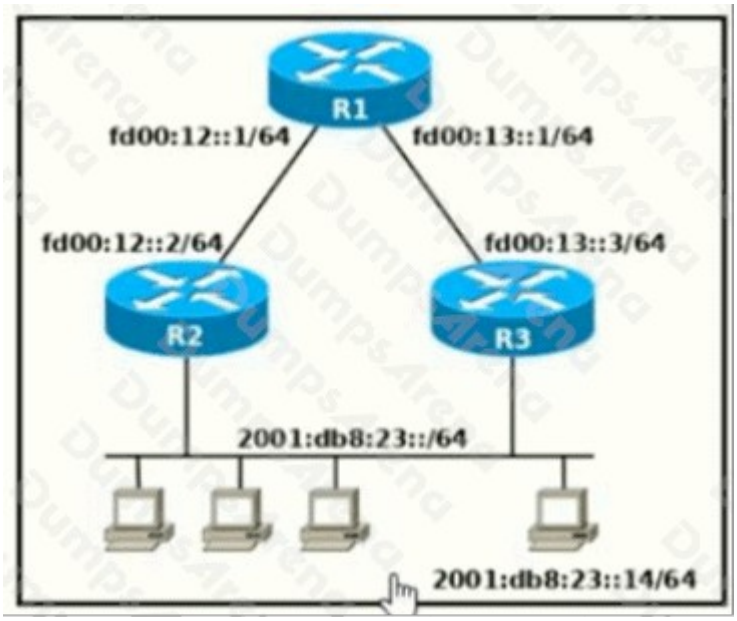
Explanation:

LACP is defined by IEEE 802.3ad and IEEE 802.1AX standards. It is an open standards-based negotiation protocol used to form and manage an EtherChannel bundle between compatible devices. This permits interoperability between equipment that supports the applicable standard.

It supports active and passive negotiation modes. An interface configured for active mode initiates LACP negotiation, while an interface in passive mode waits to receive LACP packets. At least one side of an LACP EtherChannel must operate in active mode for negotiation to occur. Cisco documents LACP operation and negotiation modes in its EtherChannel configuration and troubleshooting guide.

QUESTION NO: 67

Refer to the exhibit.



Which two commands, when configured on router R1, fulfill these requirements? (Choose two.)

Packets towards the entire network 2001:db8:2::/64 must be forwarded through router R2.

Packets toward host 2001:db8:23::14 preferably must be forwarded through R3.

A. Ipv6 route 2001:db8:23::/128 fd00:12::2

B. Ipv6 route 2001:db8:23::14/128 fd00:13::3

C. Ipv6 route 2001:db8:23::14/64 fd00:12::2

D. Ipv6 route 2001:db8:23::/64 fd00:12::2

E. Ipv6 route 2001:db8:23::14/64 fd00:12::2 200

F. ipv6 route 2001:db8:2::/64 fd00:12::2

ANSWER: B F

Explanation:

The command `ipv6 route 2001:db8:2::/64 fd00:12::2` installs a static route for the required destination network and identifies R2's IPv6 next-hop address, `fd00:12::2`. Thus, traffic whose destination falls within `2001:db8:2::/64` is forwarded from R1 toward R2. The command `ipv6 route 2001:db8:23::14/128 fd00:13::3` creates an IPv6 host route, because a `/128` prefix identifies exactly one IPv6 address. It forwards traffic for the specified host through R3 at next hop `fd00:13::3`.

These static routes satisfy two distinct destination requirements. IPv6 forwarding uses longest-prefix matching: when a packet matches multiple installed routes, the route with the greatest prefix length is selected. Consequently, a host-specific `/128` route is the proper way to provide preferred forwarding for one named IPv6 host while retaining separate routing for a broader prefix. Cisco IOS XE documents the `ipv6 route` static-route syntax and use of a next-hop IPv6 address in the Cisco IOS XE IPv6 Static Routes Configuration Guide. Cisco also describes longest-prefix matching as the forwarding rule used to select the most specific route in its route-selection documentation.

QUESTION NO: 68

Which two HTTP methods are suitable for actions performed by REST-based APIs? (Choose two.)

A. REMOVE

B. REDIRECT

C. POST

D. GET

E. POP

ANSWER: C D

Explanation:

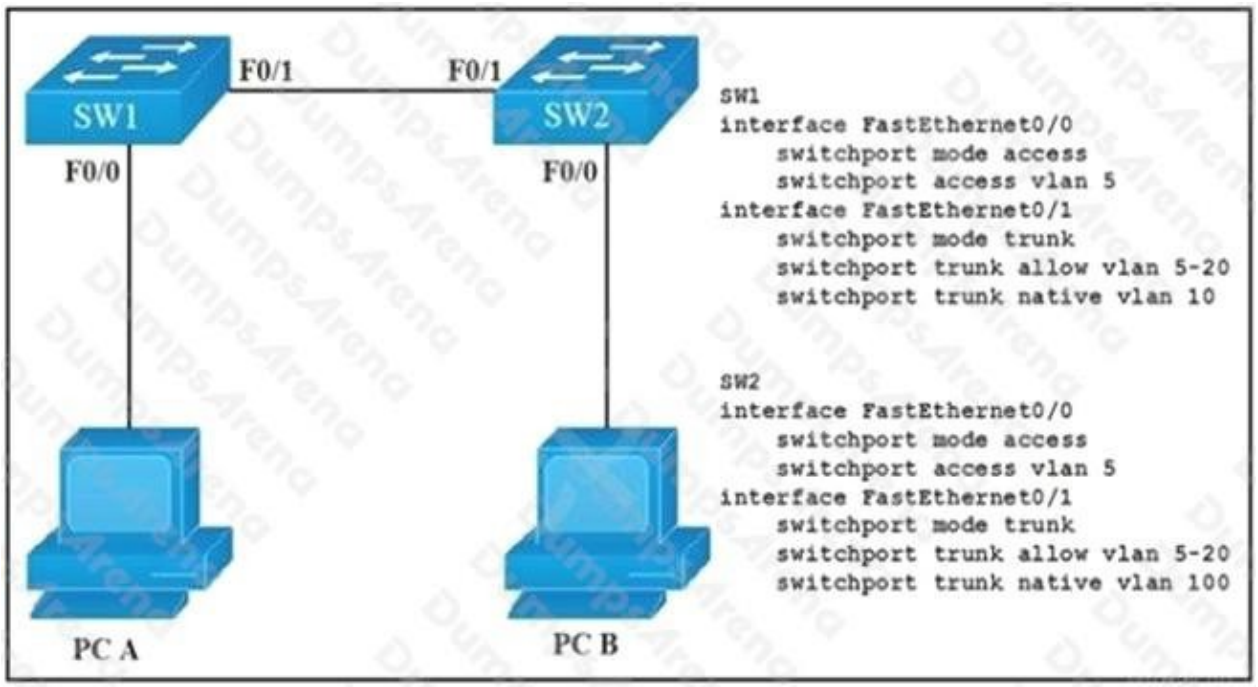
POST and **GET** are standard HTTP request methods and are fundamental to REST-based API interactions. GET retrieves a current representation of a resource, such as obtaining device inventory, interface details, or configuration data through a network controller API. HTTP defines GET as a safe method, meaning its intended semantics are read-only and do not request a state change on the origin server. Its repeated use is also idempotent, which supports reliable retrieval operations.

POST sends a representation or data to a target resource for processing according to that resource's own semantics. REST APIs commonly use POST to create subordinate resources, submit commands, initiate workflows, or provide input for processing. For example, an API client might POST a JSON payload to create an object or begin an automation task. Unlike GET, POST is not generally idempotent, because submitting the same request more than once can create multiple resources or trigger processing repeatedly.

These methods are part of the standardized HTTP method set used by REST APIs to express client intent and operate on resources over HTTP. Their semantics are defined by the HTTP specification; RESTful API designs apply those semantics consistently to resource-oriented endpoints. See RFC 9110, HTTP request methods and MDN HTTP request methods reference.

QUESTION NO: 69

Refer to the exhibit. How will switch SW2 handle traffic from VLAN 10 on SW1?



A. It sends the traffic to VLAN 10.

B. It sends the traffic to VLAN 100.

C. It drops the traffic.

D. It sends the traffic to VLAN 1.

ANSWER: B

Explanation:

It sends the traffic to VLAN 100. The trunk has a native-VLAN mismatch: VLAN 10 is the native VLAN on SW1, while VLAN 100 is the native VLAN on SW2. Under IEEE 802.1Q trunking, frames belonging to the transmitting switch's native VLAN are normally transmitted without an 802.1Q VLAN tag. Therefore, traffic sourced in VLAN 10 on SW1 crosses the trunk as untagged traffic.

When SW2 receives an untagged frame on that 802.1Q trunk, it classifies the frame into its own locally configured native VLAN. Because SW2 uses VLAN 100 as its native VLAN, it associates the received traffic with VLAN 100. Native VLAN configuration is local to each end of a trunk; an untagged frame does not carry information identifying the native VLAN used by the sending switch. Consistently configuring the same native VLAN at both trunk endpoints is therefore important for correct VLAN separation and predictable forwarding behavior. Cisco documents that the native VLAN is the VLAN used for untagged traffic on an 802.1Q trunk and recommends ensuring the native-VLAN configuration matches across the trunk. See Cisco's 802.1Q native VLAN documentation and Cisco's trunking basics documentation.

QUESTION NO: 70

While examining excessive traffic on the network, it is noted that all incoming packets on an interface appear to be allowed even though an IPv4 ACL is applied to the interface.

Which two misconfigurations cause this behavior? (Choose two)

A. The packets fail to match any permit statement

B. A matching permit statement is too high in the access list

C. A matching permit statement is too broadly defined

D. The ACL is empty

E. A matching deny statement is too high in the access list

ANSWER: B C

Explanation:

IPv4 access control lists are evaluated sequentially, from the first access control entry to the last. Cisco IOS stops processing as soon as a packet matches an entry, and the action associated with that first match is applied. Therefore, *A matching permit statement is too high in the access list* can cause unexpectedly broad inbound access when it is positioned before entries intended to restrict more specific traffic. For example, an early permit for a subnet or protocol can match packets before later filtering rules are reached.

A matching permit statement is too broadly defined produces the same observed result when its source, destination, wildcard mask, or protocol criteria encompass substantially more traffic than intended. A broad permit such as one using `any` for relevant fields may match nearly every incoming packet. Because first-match processing ends ACL evaluation immediately, the router does not continue to search for a more restrictive entry later in the ACL. Correct ACL design requires placing specific matches before general matches and carefully validating wildcard masks and address scope. Cisco documents the top-down, first-match ACL evaluation process and IPv4 ACL configuration behavior at Cisco IOS Access Control Lists and Cisco IPv4 ACL Configuration Guide.

QUESTION NO: 71

What is a benefit of VRRP?

A. It provides traffic load balancing to destinations that are more than two hops from the source.

B. It provides the default gateway redundancy on a LAN using two or more routers.

C. It allows neighbors to share routing table information between each other.

D. It prevents lloops in a Layer 2 LAN by forwarding all traffic to a root bridge, which then makes the final forwarding decision.

ANSWER: B

Explanation:

It provides the default gateway redundancy on a LAN using two or more routers. VRRP (Virtual Router Redundancy Protocol) is a first-hop redundancy protocol that lets several physical routers operate as one virtual default gateway for hosts on the same LAN. The routers are configured with a shared virtual IP address, and clients use that virtual address as their default gateway rather than the address of an individual router. One VRRP router is elected as the master and forwards traffic sent to the virtual gateway. The other participating routers serve as backups and monitor the master's availability. If the master fails, an eligible backup assumes the master role and takes ownership of the virtual gateway, allowing hosts to continue forwarding off-network traffic without requiring their default-gateway configuration to be changed. This capability improves LAN gateway availability and minimizes disruption caused by failure of a single router or its interface. Cisco documents VRRP as a mechanism that eliminates the single point of failure inherent in a static default gateway and provides transparent failover for end hosts. See Cisco's VRRP configuration guide and the IETF's VRRP specification (RFC 5798).

QUESTION NO: 72

Which two goals reasons to implement private IPv4 addressing on your network? (Choose two.)

A. Comply with PCI regulations

B. Conserve IPv4 address

C. Reduce the size of the forwarding table on network routers

D. Reduce the risk of a network security breach

E. Comply with local law

ANSWER: B D

Explanation:

Private IPv4 addressing is used to conserve the limited supply of globally routable IPv4 addresses. RFC 1918 reserves the 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16 ranges for private internets, allowing an organization to assign addresses to internal endpoints without obtaining a unique public IPv4 address for every device. When those internal devices need Internet connectivity, Network Address Translation commonly translates private source addresses to one or a small number of public addresses at the network edge.

Private IPv4 addressing can also reduce the risk of a network security breach by limiting direct Internet reachability to internal hosts. RFC 1918 addresses are not intended to be routed on the public Internet, so unsolicited traffic from external networks cannot be delivered directly to an internal private address. This creates an addressing boundary that is typically enforced with NAT and edge firewall policy. It is a useful layer that reduces exposure, although secure network design still requires appropriate access controls, segmentation, firewall rules, and monitoring.

See RFC 1918: Address Allocation for Private Internets and Cisco's NAT overview.

QUESTION NO: 73

How does the dynamically-learned MAC address feature function?

A. The CAM table is empty until ingress traffic arrives at each port

B. Switches dynamically learn MAC addresses of each connecting CAM table.

C. The ports are restricted and learn up to a maximum of 10 dynamically-learned addresses

D. It requires a minimum number of secure MAC addresses to be filled dynamically

ANSWER: A

Explanation:

The CAM table is empty until ingress traffic arrives at each port is the best description of dynamic MAC address learning. A switch builds its forwarding database by inspecting the source MAC address on every Ethernet frame received on an ingress interface. It associates that source address with the port on which the frame arrived, creating a dynamic MAC-to-port entry in the CAM table. Until a connected device sends a frame, the switch has not seen that device's source address and therefore has no dynamic entry identifying the device's location. As traffic continues, the switch refreshes the entry's aging timer whenever it sees frames sourced from that MAC address. This learned information enables efficient unicast forwarding: once the destination MAC address is known, frames can be sent only through the associated outgoing interface rather than being flooded. Dynamic entries are temporary and are normally removed after their aging interval expires if the switch no longer receives traffic from the device. In practice, a table can also contain manually configured static entries, but dynamically learned entries are populated only as ingress frames are received. Cisco's overview of this source-address learning and CAM-table behavior is available in Ethernet Switching: How MAC Address Tables Work and the Cisco IOS XE MAC Address Table Configuration Guide.

QUESTION NO: 74

Refer to the exhibit.



Which set of commands must be applied to the two switches to configure an LACP Layer 2 EtherChannel?

A)

```
SW1(config)#interface range f0/13 -14
SW1(config-if-range)#channel-group 1 mode active

SW2(config)#interface range f0/13 -14
SW2(config-if-range)#channel-group 1 mode passive
```

B)

```
SW1(config)#interface range f0/13 -14
SW1(config-if-range)#channel-group 1 mode desirable

SW2(config)#interface range f0/13 -14
SW2(config-if-range)#channel-group 1 mode passive
```

C)

```
SW1(config)#interface range f0/13 -14
SW1(config-if-range)#channel-group 1 mode auto

SW2(config)#interface range f0/13 -14
SW2(config-if-range)#channel-group 1 mode passive
```

D)

```
SW1(config)#interface range f0/13 -14
SW1(config-if-range)#channel-group 1 mode on

SW2(config)#interface range f0/13 -14
SW2(config-if-range)#channel-group 1 mode passive
```

A. Option A

B. Option B

C. Option C

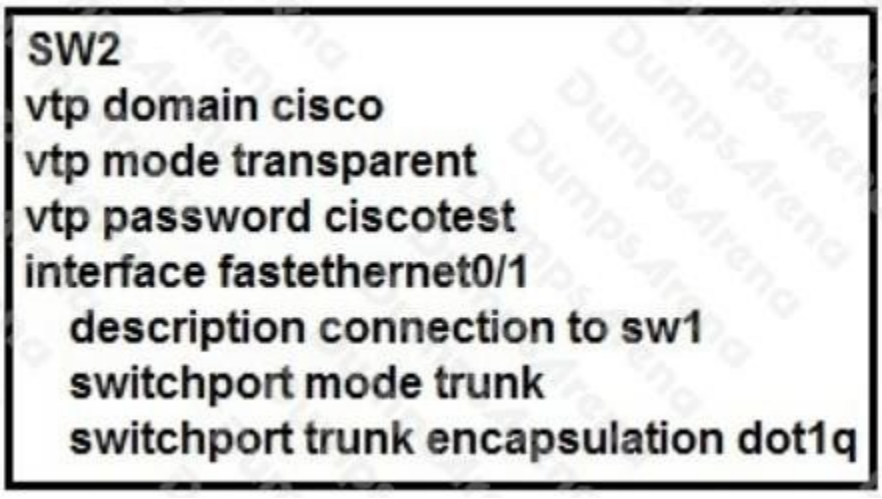
D. Option D

ANSWER: A

Explanation:

The LACP EtherChannel configuration is correct because it forms a single logical Layer 2 port channel using the Link Aggregation Control Protocol. On Cisco switches, LACP is enabled when member interfaces are assigned to the same channel group with a compatible LACP negotiation mode: `active` initiates LACP negotiation, while `passive` waits for LACP packets from its peer. Therefore, using active mode on both switches, or active mode on one side and passive mode on the other, establishes the bundle. The physical interfaces must use the same channel-group number locally, and their operational Layer 2 settings must be compatible. Once negotiation succeeds, the switch creates the logical `Port-channel` interface, which is the interface on which the EtherChannel operates as a Layer 2 switching link. LACP also monitors member-link compatibility and can remove an unsuitable member from the bundle while allowing the remaining operational links to continue forwarding traffic. Cisco documents LACP as an IEEE 802.3ad/802.1AX standards-based EtherChannel negotiation protocol and identifies `active` and `passive` as its supported modes. See Cisco's EtherChannel configuration guide and the Cisco EtherChannel overview.

QUESTION NO: 75



Refer to the exhibit. How does SW2 interact with other switches in this VTP domain?

A. It transmits and processes VTP updates from any VTP clients on the network on its trunk ports.

B. It processes VTP updates from any VTP clients on the network on its access ports.

C. It receives updates from all VTP servers and forwards all locally configured VLANs out all trunk ports.

D. It forwards only the VTP advertisements that it receives on its trunk ports.

ANSWER: D

Explanation:

It forwards only the VTP advertisements that it receives on its trunk ports. This behavior identifies SW2 as operating in VTP transparent mode. A transparent switch does not use received VTP advertisements to synchronize or modify its own VLAN database; VLAN configuration remains locally administered on that switch. It nevertheless relays received VTP advertisements through its trunk interfaces, allowing VTP information originated elsewhere in the domain to traverse SW2 and reach adjacent switches. Because VTP messages are carried over trunk links, forwarding occurs on trunk ports rather

than access ports. The switch is therefore able to sit between other VTP-speaking switches without becoming a participant in VLAN database synchronization. This forwarding behavior is especially useful where a switch must maintain an independent local VLAN database while still providing Layer 2 connectivity between switches that exchange VTP advertisements. Cisco documents that VTP transparent-mode switches do not synchronize their VLAN databases with received advertisements but forward VTP advertisements on trunk links. See Cisco's VTP Concepts and Configuration and the Catalyst VTP Configuration Guide.

QUESTION NO: 76

When DHCP is configured on a router, which command must be entered so the default gateway is automatically distributed?

A. default-router

B. default-gateway

C. ip helper-address

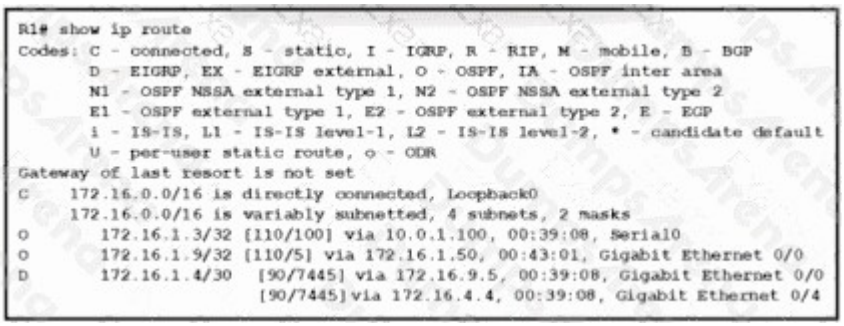
D. dns-server

ANSWER: A

Explanation:

The `default-router` command is configured within a Cisco IOS DHCP pool to provide DHCP clients with the IPv4 default gateway. For example, `ip dhcp pool USERS` followed by `default-router 192.0.2.1` instructs the router's DHCP service to include that gateway address in its lease offer and acknowledgment. A client uses the received address as its next-hop router when sending traffic to destinations outside its own local IP subnet. This setting corresponds to the DHCP Router option (option 3), which allows endpoint configuration to be delivered automatically rather than requiring a manually configured default route on every host. The command can include more than one gateway address when a subnet has multiple routers, with clients receiving the addresses in the configured order. Cisco's DHCP server configuration guide documents `default-router` as the pool command for specifying the default router for DHCP clients. See Cisco's IOS DHCP server configuration guide and the DHCP Router Option specification in RFC 2132.

QUESTION NO: 77



Refer to the exhibit. How does router R1 handle traffic to the 172.16.1.4/30 subnet?

A. It sends all traffic over the path via 172.16.9.5 using 172.16.4.4 as a backup.

B. It sends all traffic over the path via 172.16.4.4.

C. It load-balances traffic over 172.16.9.5 and 172.16.4.4.

D. It sends all traffic over the path via 10.0.1.100

ANSWER: C

Explanation:

It load-balances traffic over 172.16.9.5 and 172.16.4.4. The exhibit identifies two routes from R1 toward the identical destination prefix, 172.16.1.4/30, with next-hop addresses 172.16.9.5 and 172.16.4.4. Because these routes have equal route preference (the same administrative distance and metric), R1 can install both as equal-cost paths in its routing table rather than selecting only one path. This is equal-cost multipath routing (ECMP).

After both paths are installed, Cisco Express Forwarding uses the forwarding information to distribute eligible traffic across the available equal-cost next hops. In typical Cisco IOS operation, the distribution is flow based, using fields in the packet header to select a next hop. This allows separate traffic flows to use different paths while maintaining consistent forwarding for packets belonging to the same flow. The usable result is that R1 forwards traffic for the 172.16.1.4/30 network through both listed next hops as part of its load-sharing behavior. Cisco documents equal-cost route installation and load sharing in its static routing configuration guide and describes CEF forwarding behavior in the Cisco Express Forwarding configuration guide.

QUESTION NO: 78

What are two southbound APIs? (Choose two.)

A. Thrift

B. DSC

C. CORBA

D. NETCONF

E. OpenFlow

ANSWER: D E

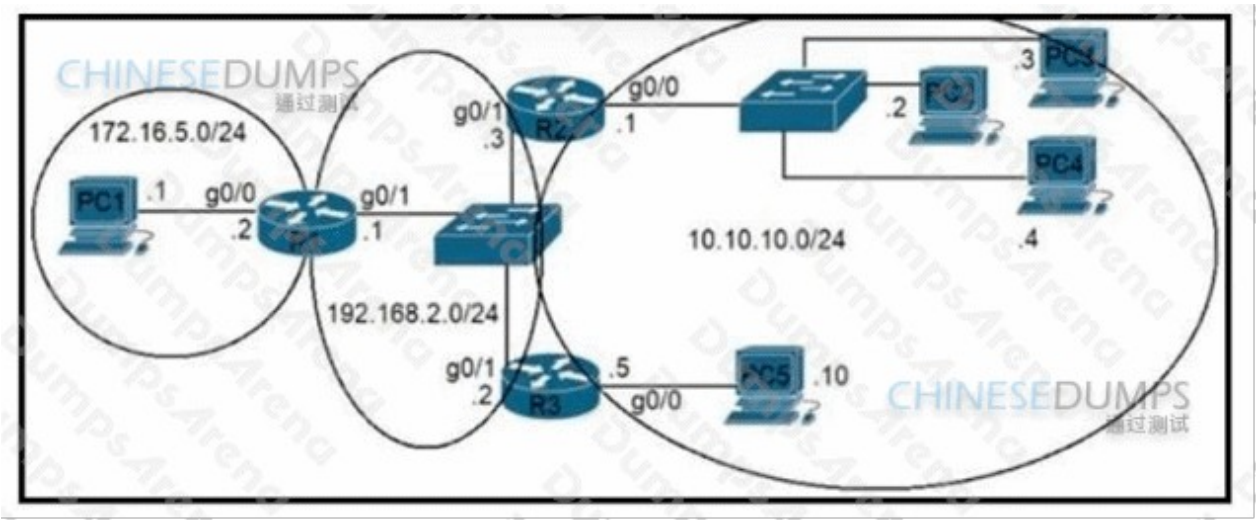
Explanation:

NETCONF and **OpenFlow** are southbound interfaces because they enable an SDN controller or network-management application to communicate with and control network infrastructure devices. Southbound APIs sit between the controller/control layer and the forwarding devices in the data plane. Their purpose is to convey configuration, policy, forwarding instructions, and operational-state requests downward to routers, switches, and other managed devices.

OpenFlow is a foundational SDN protocol designed for a controller to manage a switch's forwarding behavior by installing, modifying, and removing flow-table entries. This direct controller-to-forwarding-device relationship is the classic southbound API model. **NETCONF**, commonly used with YANG data models, provides a standardized mechanism for network automation systems and controllers to retrieve configuration and state data and to make configuration changes on network devices. Cisco platforms widely support NETCONF/YANG for model-driven programmability, making it a relevant southbound interface in the CCNA automation and programmability context.

For protocol details, see IETF RFC 6241: Network Configuration Protocol (NETCONF). For the SDN architecture and controller-to-device interface concept, see the Open Networking Foundation technical library.

QUESTION NO: 79



Refer to the exhibit. The router R1 is in the process of being configured. Routers R2 and R3 are configured correctly for the new environment. Which two commands must be configured on R1 for PC1 to communicate to all PCs on the 10.10.10.0/24 network? (Choose two.)

A. ip route 10.10.10.0 255.255.255.0 192.168.2.3

B. ip route 10.10.10.10 255.255.255.255 192.168.2.2

C. ip route 10.10.10.10 255.255.255.255 g0/1

D. ip route 10.10.10.8 255.255.255.248 g0/1

E. ip route 10.10.10.0 255.255.255.248 192.168.2.2

F. ip route 10.10.10.8 255.255.255.248 192.168.2.3

ANSWER: E F

Explanation:

R1 requires a separate static route for each remote subnet that comprises the addressed 10.10.10.0/24 space in the topology. The command `ip route 10.10.10.0 255.255.255.248 192.168.2.2` installs a route for the 10.10.10.0/29 subnet, whose usable host range is 10.10.10.1 through 10.10.10.6. It forwards packets for that subnet to R2 at the next-hop address 192.168.2.2. The command `ip route 10.10.10.8 255.255.255.248 192.168.2.3` installs the corresponding route for the next /29 subnet, 10.10.10.8/29, whose usable addresses are 10.10.10.9 through 10.10.10.14, and forwards packets to R3 at 192.168.2.3. Together, these routes give R1 reachability to the PCs located on both remote LAN segments. Cisco IOS static-route syntax specifies a destination prefix and mask followed by either a next-hop IP address or an exit interface; using each downstream router's next-hop address clearly identifies the intended path across the shared transit network. Routes are selected using longest-prefix match, so the two /29 entries provide the required specific forwarding information for their respective destination networks. See Cisco's Static Routes Configuration Guide and IP Routing: Static Routes Overview.

QUESTION NO: 80

What is a DNS lookup operation?

A. DNS server pings the destination to verify that it is available

B. serves requests over destination port 53

C. DNS server forwards the client to an alternate IP address when the primary IP is down

D. Resolves a domain name to an IP address, or an IP address to a domain name

ANSWER: D

Explanation:

"Resolves a domain name to an IP address, or an IP address to a domain name" correctly describes a DNS lookup. The Domain Name System is a distributed naming service that maps human-readable domain names, such as `www.example.com`, to the numeric IP addresses that hosts use for network communication. This common forward-lookup process lets an application identify the destination address before it establishes a connection. DNS also supports a reverse lookup, which starts with an IP address and queries the appropriate reverse-mapping namespace for a PTR record that identifies the associated domain name. Therefore, the general term "DNS lookup" encompasses both name-to-address and address-to-name resolution; the query type and requested record determine which form occurs. Cisco network engineers use DNS lookups when validating client name-resolution behavior, troubleshooting connectivity to named resources, and confirming that configured DNS servers return the expected records. DNS protocol fundamentals, including the resolver and name-server roles, are specified in RFC 1034. The DNS record types used for IPv4, IPv6, and reverse mappings are defined in the IANA DNS Parameters registry.

QUESTION NO: 81

Which two statements about the purpose of the OSI model are accurate? (Choose two.)

A. Defines the network functions that occur at each layer

B. Facilitates an understanding of how information travels throughout a network

C. Changes in one layer do not impact other layer

D. Ensures reliable data delivery through its layered approach

ANSWER: A B

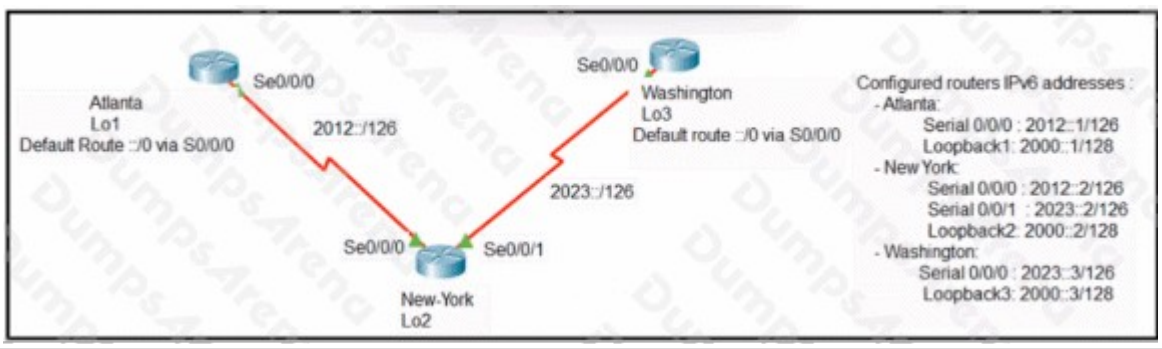
Explanation:

The OSI model is a conceptual reference framework that divides network communication into seven layers, with each layer describing a defined category of networking functions and services. Therefore, "Defines the network functions that occur at each layer" is accurate: the model organizes functions such as physical transmission, framing, logical addressing and routing, end-to-end transport, and application-facing network services into a common layered structure. This common vocabulary enables network professionals and vendors to describe where a networking process belongs without tying the discussion to a single protocol suite or product implementation.

"Facilitates an understanding of how information travels throughout a network" is also accurate because the layered model illustrates the logical progression of data as it is prepared by an endpoint, encapsulated for transmission, sent across media and intermediate devices, then decapsulated at the destination. Using layers helps engineers understand which functions are involved at each stage of communication and provides a systematic framework for studying network operations. Cisco's networking training materials use the OSI model as a foundational way to classify protocols and communication functions. See Cisco's Internetworking Basics and the ISO overview of the OSI Basic Reference Model.

QUESTION NO: 82

Refer to the exhibit.



The loopback1 interface of the Atlanta router must reach the lookback3 interface of the Washington router.

A. ipv6 route 2000::1/128 2012::2

B. ipv6 route 2000::1/128 2012::1

C. ipv6 route 2000:3 123 s0/0/0

D. ipv6 route 2000::3/128 2023::3

E. ipv6 route 2000::1/128 s0/0/1

F. ipv6 route 2000::3/128 2023::2

G. ipv6 route 2000::3/128 s0/0/1

ANSWER: D G

Explanation:

The required destination is Washington's Loopback3 address, `2000::3/128`. A loopback interface configured with a single IPv6 address is reached through a host route using the `/128` prefix length. The command `ipv6 route 2000::3/128 2023::3` correctly installs that host route and identifies the IPv6 address of the Washington router on the Atlanta-to-Washington path as the next hop. Cisco IOS performs recursive lookup for the next-hop address and forwards traffic toward Washington.

`ipv6 route 2000::3/128 s0/0/1` is also a valid static-route form when Serial0/0/1 is Atlanta's outgoing point-to-point interface toward Washington. Cisco IOS supports IPv6 static routes that specify either a next-hop IPv6 address or an exit interface. On a point-to-point serial connection, an exit-interface-only route is sufficient because the interface identifies the single directly connected forwarding path. Both commands therefore create a route from Atlanta to the required Loopback3 host prefix, using the same destination but different valid next-hop specifications.

Cisco documents IPv6 static-route configuration and next-hop/exit-interface usage in IPv6 Static Routing Configuration and the Cisco IOS XE IPv6 Static Routes guide.

QUESTION NO: 83

How is noise defined in Wi-Fi?

A. ratio of signal-to-noise rating supplied by the wirelesss device

B. signals from other Wi-Fi networks that interfere with the local signal

C. measured difference between the desired Wi-Fi signal and an interfering Wi-Fi signal

D. any interference that is not Wi-Fi traffic that degrades the desired signal

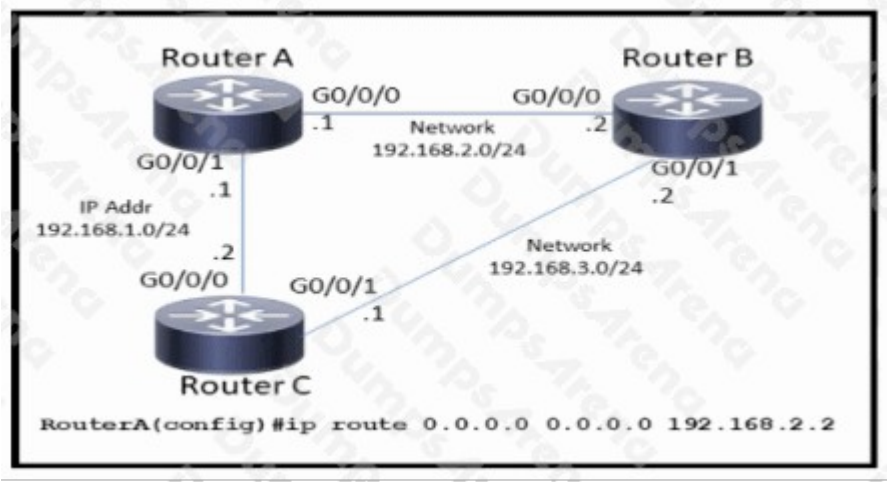
ANSWER: D

Explanation:

Any interference that is not Wi-Fi traffic that degrades the desired signal is the best definition of noise in this Wi-Fi context. In RF terms, noise is unwanted energy received on or near the channel that reduces the receiver's ability to distinguish and correctly decode the intended wireless transmission. It raises the effective noise floor, reducing the signal-to-noise ratio available to a client or access point. Common sources include microwave ovens, Bluetooth devices, cordless phones, radar systems, analog video equipment, and other non-802.11 RF emitters. The practical effect is increased frame corruption, retransmissions, lower data rates, and potentially unstable client connectivity, even when the desired access point's received signal strength appears adequate. Cisco wireless design and troubleshooting therefore assess both received signal strength and the ambient noise floor, often using CleanAir-capable access points or spectrum-analysis tools to identify persistent non-Wi-Fi interference. Cisco describes CleanAir as a capability for detecting, classifying, and mitigating RF interference, while IEEE 802.11 defines the WLAN technology whose frames are distinct from such external RF energy. See Cisco's CleanAir technology overview and the IEEE 802.11 standard overview.

QUESTION NO: 84

Refer to the exhibit.



Which command must be issued to enable a floating static default route on router A?

A. Ip route 0.0.0.0 0.0.0.0 192.168.1.2

B. ip default-gateway 192.168.2.1

C. ip route 0.0.0.0 0.0.0.0 192.168.2.1 10

D. ip route 0.0.0.0 0.0.0.0 192.168.1.2 10

ANSWER: D

Explanation:

`ip route 0.0.0.0 0.0.0.0 192.168.1.2 10` correctly creates a floating static default route on Router A. The destination and mask of `0.0.0.0 0.0.0.0` define an IPv4 default route, which is used when no more-specific route matches a destination. The next-hop address, `192.168.1.2`, identifies the backup-path neighbor shown for Router A in the topology.

The trailing value, `10`, is the static route's administrative distance. Cisco IOS normally assigns static routes an administrative distance of 1. Assigning a higher value makes this route less preferred than the primary static route while that primary route is available. Consequently, the route is installed and used only after the preferred route is no longer present in the routing table, which is the intended floating-static-route behavior. This provides a straightforward backup default path without requiring a dynamic routing protocol.

Cisco documents that administrative distance is used to select among route sources and that a static route can be configured with an administrative-distance value. See Cisco: Understanding Administrative Distance and Cisco IOS XE: Configuring Static Routes.

QUESTION NO: 85

In which way does a spine-and-leaf architecture allow for scalability in a network when additional access ports are required?

A. A spine switch and a leaf switch are added with redundant connections between them.

B. A spine switch is added with at least 40 GB uplinks.

C. A leaf switch is added with a single connection to a core spine switch.

D. A leaf switch is added with connections to every spine switch.

ANSWER: D

Explanation:

A leaf switch is added with connections to every spine switch. In a spine-and-leaf topology, leaf switches provide the attachment points for endpoints, such as hosts, servers, and other access-layer devices. Therefore, adding a leaf increases the number of available access ports without requiring a redesign of the existing fabric. Each new leaf is connected to all spine switches, preserving the architecture's predictable, equal-cost paths between leaves and providing consistent bandwidth and resiliency across the fabric. This scale-out approach means that capacity can be increased incrementally: additional endpoint connectivity is delivered by deploying another leaf, while the spine layer continues to provide the high-speed transit network between all leaves. Cisco describes leaf-and-spine designs as fabrics in which every leaf connects to every spine, creating a nonblocking, multipath topology that can grow by adding leaf switches for endpoint ports. This design is commonly used in modern data centers because it supports efficient east-west traffic and straightforward expansion. For further detail, see Cisco's Spine-Leaf Network overview and the Cisco Data Center Networking resources.

QUESTION NO: 86 - (HOTSPOT)

All physical cabling is in place. Router R4 and PCI are fully configured and

inaccessible. R4 ' s WAN interfaces use .4 in the last octet for each subnet.

Configurations should ensure that connectivity is established end-to-end.

1 . Configure static routing to ensure RI prefers the path through R2 to

reach only PCI on R4 ' s LAN

2. Configure static routing that ensures traffic sourced from RI will take

an alternate path through R3 to PCI in the event of an outage along

the primary path

3. Configure default routes on RI and R3 to the Internet using the least number of hops

Guidelines

This is a lab item in which tasks will be performed on virtual devices.

â€¢ Refer to the Tasks tab to view the tasks for this lab item.

â€¢ Refer to the Topology tab to access the device console(s) and perform the tasks.

â€¢ Console access is available for all required devices by clicking the device icon or using

the tab(s) above the console window.

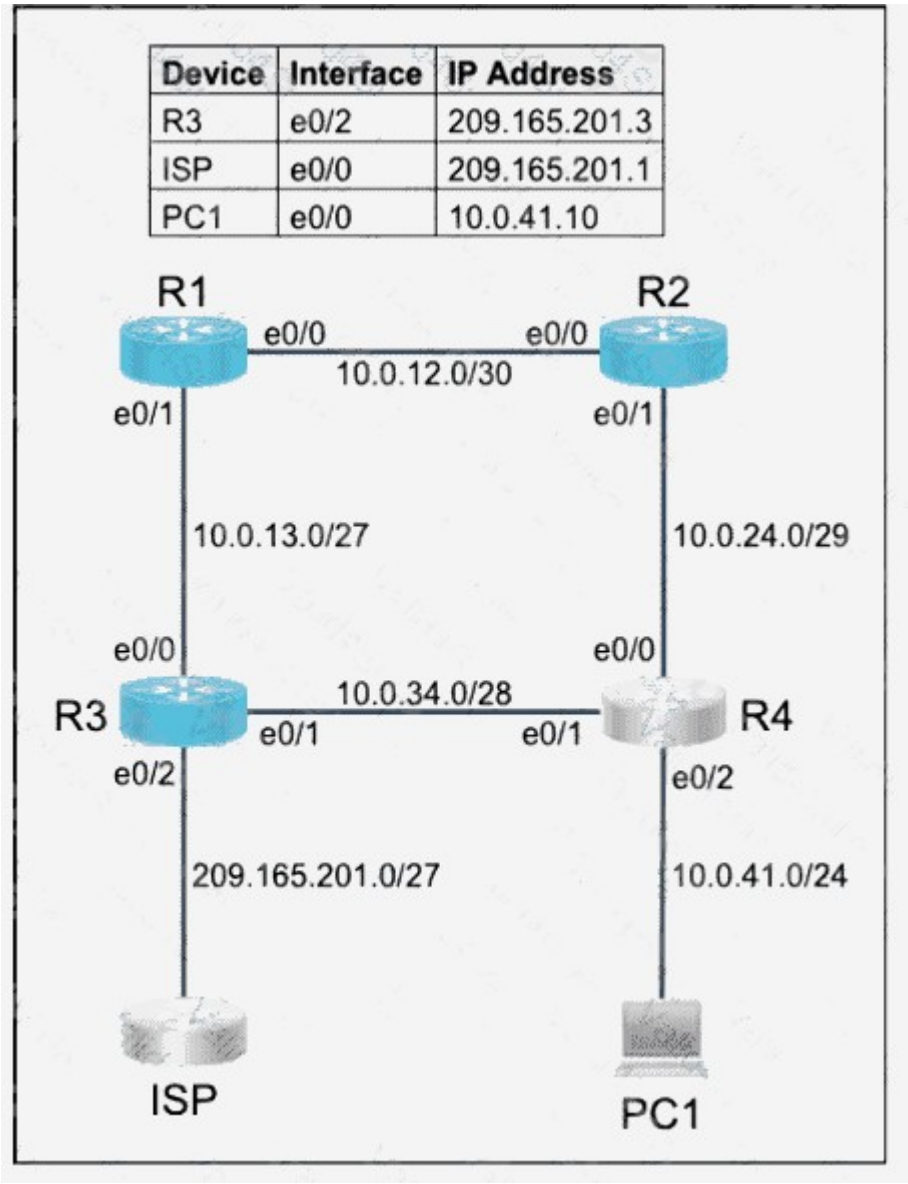
â€¢ All necessary preconfigurations have been applied.

â€¢ Do not change the enable password or hostname for any device.

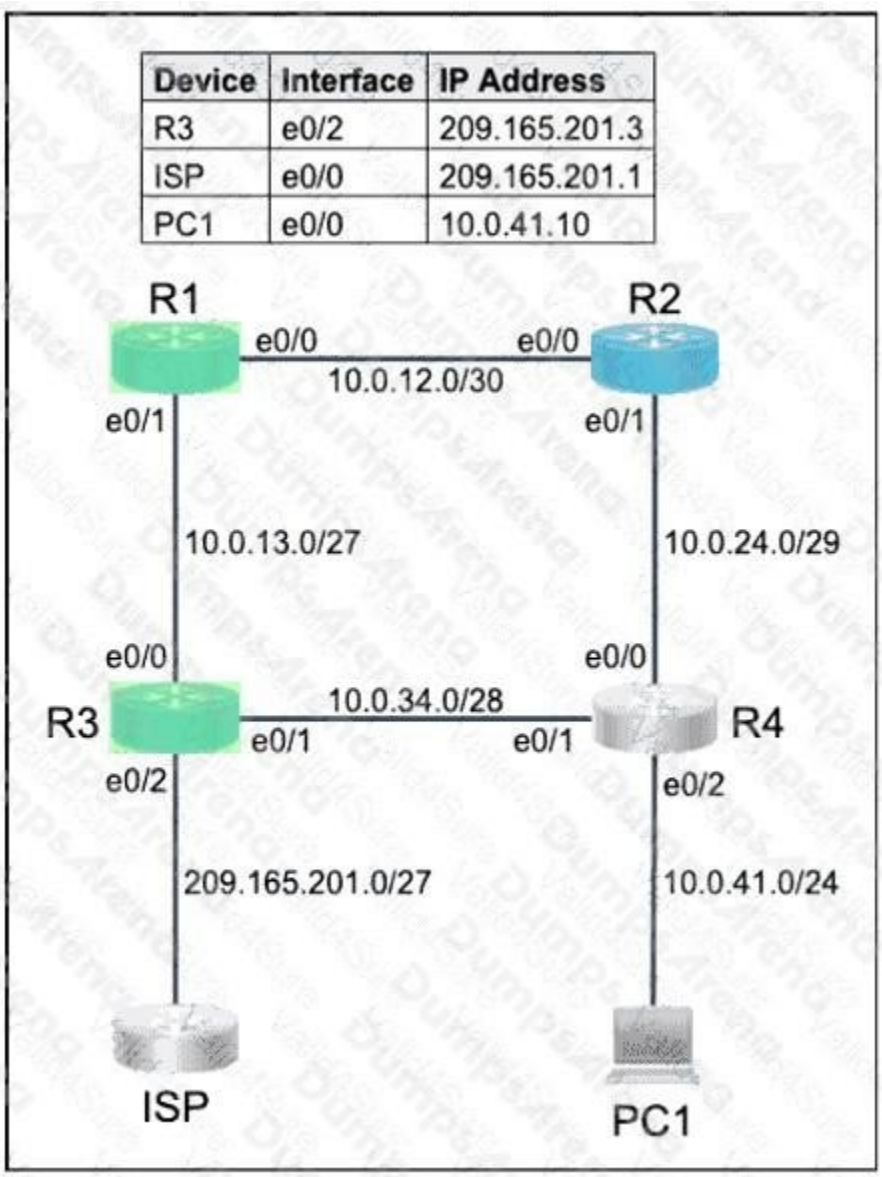
â€¢ Save your configurations to NVRAM before moving to the next item.

â€¢ Click Next at the bottom of the screen to submit this lab and move to the next question.

â€¢ When Next is clicked, the lab closes and cannot be reopened.


Device Interface IP Address
R3 e0/2 209.165.201.3
ISP e0/0 209.165.201.1
PC1 e0/0 10.0.41.10
R1
R2
e0/0
e0/0
10.0.12.0/30
e0/1
e0/1
10.0.13.0/27
10.0.24.0/29
e0/0
e0/0
R3
10.0.34.0/28
e0/1
e0/1
R4
e0/2
e0/2
209.165.201.0/27
10.0.41.0/24
ISP
PC1

ANSWER:



Explanation:

PC1 has the specific address 10.0.41.10, so the route from R1 must be a host route using the 255.255.255.255 mask. This precisely meets the requirement to prefer the R2 path only when the destination is PC1, rather than creating a route that covers every host in R4's 10.0.41.0/24 LAN. R2 is the directly adjacent next hop from R1 on the 10.0.12.0/30 network, using address 10.0.12.2. A static route without an explicitly configured administrative distance uses the normal static-route administrative distance of 1, making this the preferred path.

R3 is R1's other adjacent router, reachable at 10.0.13.3 on the 10.0.13.0/27 link. Installing an identical host route through that next hop with an administrative distance of 2 creates a floating static route. It remains in the routing table as a standby route and is selected when the preferred route through R2 is no longer usable. This gives PC1 a defined alternate path through R3 and then R4.

The Internet-facing link is attached to R3, and the ISP next-hop address shown is 209.165.201.1. R3 therefore sends unknown destinations directly to that ISP with its default route. R1 sends its own unknown destinations to R3 at 10.0.13.3, which is the shortest path from R1 to the Internet exit. The more-specific host routes for PC1 continue to take precedence over R1's default route because routers select the longest matching prefix first. Cisco documents this route-selection behavior and static-route configuration in its static routing configuration guide and default routing documentation.

QUESTION NO: 87

Which two actions are taken as the result of traffic policing? (Choose two.)

A. bursting

B. dropping

C. remarking

D. fragmentation

E. buffering

ANSWER: B C

Explanation:

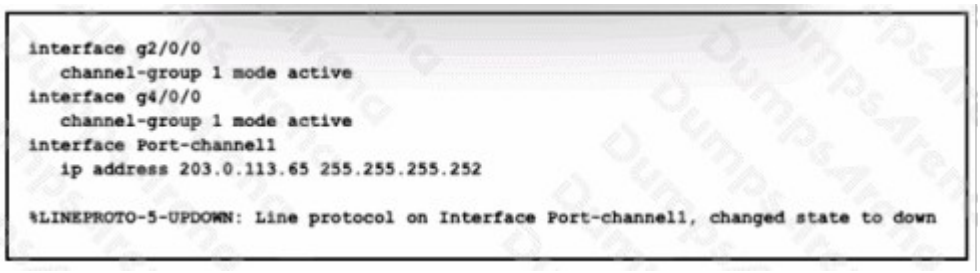
Traffic policing enforces a configured traffic rate by comparing packets with a defined rate limit, commonly implemented with a token-bucket mechanism. When packets conform to the configured parameters, they are forwarded according to the policy. For packets that exceed the permitted rate, a policer applies a configured exceed action immediately rather than delaying traffic to smooth the transmission rate.

dropping is a standard policing result. A device can discard packets identified as exceeding the policer's committed or permitted rate, enforcing the traffic contract and protecting downstream resources from excessive traffic.

remarking is also a valid policing result. Instead of discarding excess traffic, the policer can rewrite its QoS marking, such as an IP DSCP value or Ethernet CoS value. The packet remains eligible for forwarding, but its lower-priority marking enables later QoS nodes to handle it according to the revised service classification. Cisco documents policing as a mechanism that can drop packets or mark them down when traffic exceeds the configured limit. See Cisco's Policing and Shaping Overview and the Cisco IOS XE QoS Policing and Shaping Configuration Guide.

QUESTION NO: 88

Refer to the exhibit.



An engineer is configuring a Layer 3 port-channel interface with LACP. The configuration on the first device is complete, and it is verified that both interfaces have registered the neighbor device in the CDP table. Which task on the neighbor device enables the new port channel to come up without negotiating the channel?

A. Change the EtherChannel mode on the neighboring interfaces to auto.

B. Configure the IP address of the neighboring device.

C. Bring up the neighboring interfaces using the no shutdown command.

D. Configure the EtherChannel mode on the neighboring interfaces to LACP passive.

ANSWER: D

Explanation:

Configure the neighboring member interfaces for LACP passive mode. LACP EtherChannels require compatible LACP channel-group modes at both ends. An interface in *active* mode initiates LACP by transmitting LACP data units, while an interface in *passive* mode listens for and responds to those units. Therefore, when the completed first-device configuration uses LACP active mode, setting the neighbor-facing physical interfaces to LACP passive mode provides the complementary configuration required to establish the Layer 3 port-channel. The neighbor does not need to initiate the exchange itself; it

responds to the LACP packets sent by the active side and joins the member links into the bundle. This preserves dynamic LACP operation, including its ability to identify compatible links and maintain the aggregate. The Layer 3 characteristics of the port-channel do not change the LACP requirement on its physical member interfaces. Cisco documents that LACP supports active and passive modes and that at least one side must be active for the channel to form. See Cisco's EtherChannel configuration and negotiation overview and the IOS XE EtherChannel configuration guide.

QUESTION NO: 89

Which interface is used to send traffic to the destination network?



A. G0/21

B. G0/4

C. G0/5

D. G0/16

ANSWER: D

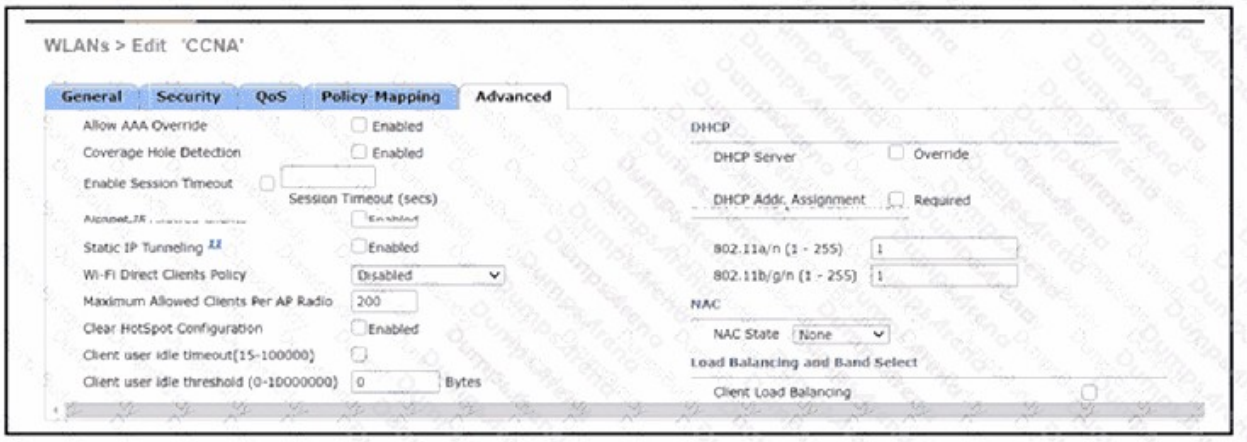
Explanation:

G0/16 is the outbound interface selected to reach the destination network. A router makes this forwarding decision by comparing the packet's destination IP address with all applicable entries in its routing table. When more than one route can match, Cisco routers apply the longest-prefix-match rule: the route with the most specific network prefix is installed in the forwarding decision for that destination. The selected route identifies how the packet leaves the router, either through a directly connected exit interface or through an exit interface used to reach a next-hop router.

In the routing information shown, the most specific matching route for the destination identifies G0/16 as its exit interface. The router therefore encapsulates the packet in a Layer 2 frame appropriate for the network attached to G0/16 and forwards it through that interface. If a next-hop address is associated with the route, the router first resolves that next hop on G0/16, such as by using ARP for IPv4, before transmitting the frame. This behavior is fundamental to Cisco IP routing and forwarding operations. Cisco's overview of route selection and prefix matching is available in Understanding Routing Table Output and the Cisco IP Routing Overview.

QUESTION NO: 90

Refer to the exhibit.



A network engineer configures the CCNA WLAN so that clients must authenticate hourly and to limit the number of simultaneous connections to the WLAN to Which two actions complete this configuration? (Choose two.)

A. Enable the Enable Session Timeout option and set the value to 3600.

B. Set the Maximum Allowed Clients value to 10.

C. Enable the Client Exclusion option and set the value to 3600.

D. Enable the Wi-Fi Direct Clients Policy option.

E. Set the Maximum Allowed Clients Per AP Radio value to 10.

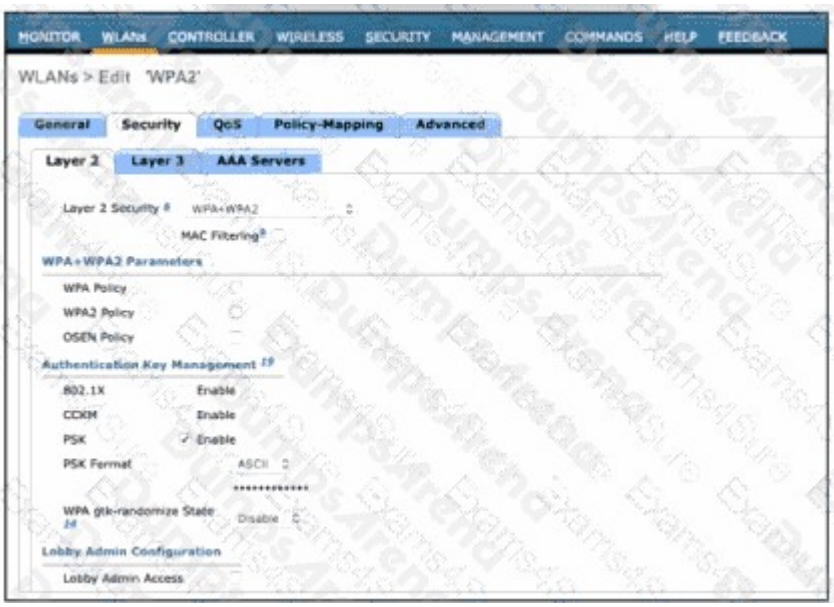
ANSWER: A B

Explanation:

Enable the Enable Session Timeout option and set the value to 3600 to impose an hourly client-session limit. Cisco wireless controllers express the session-timeout value in seconds, so 3,600 seconds equals one hour. When that period expires, the controller ends the client session; the client must then complete the applicable authentication process again before continuing WLAN access. This enforces the stated requirement for recurring hourly authentication rather than merely setting an inactivity-based timer.

Set the Maximum Allowed Clients value to 10 to limit simultaneous associations for this WLAN to ten clients. This WLAN-level client limit applies to the SSID as a whole, which matches a requirement to cap the number of concurrent connections to the WLAN. Together, these settings independently enforce the requested authentication interval and the total concurrent-client capacity. Cisco's controller WLAN configuration documentation describes WLAN advanced parameters, including session controls and client limits, in the Cisco Wireless Controller Configuration Guide. Cisco's Wireless Controller Configuration Guide, Release 8.10 also provides the configuration context for these WLAN settings.

QUESTION NO: 91



Refer to the exhibit. A network engineer is configuring a WLAN to use a WPA2 PSK and allow only specific clients to join. Which two actions must be taken to complete the process? (Choose two.)

A. Enable the 802.1X option for Authentication Key Management

B. Enable the WPA2 Policy option

C. Enable the CCKM option for Authentication Key Management

D. Enable the MAC Filtering option

E. Enable the OSEN Policy option

ANSWER: B D

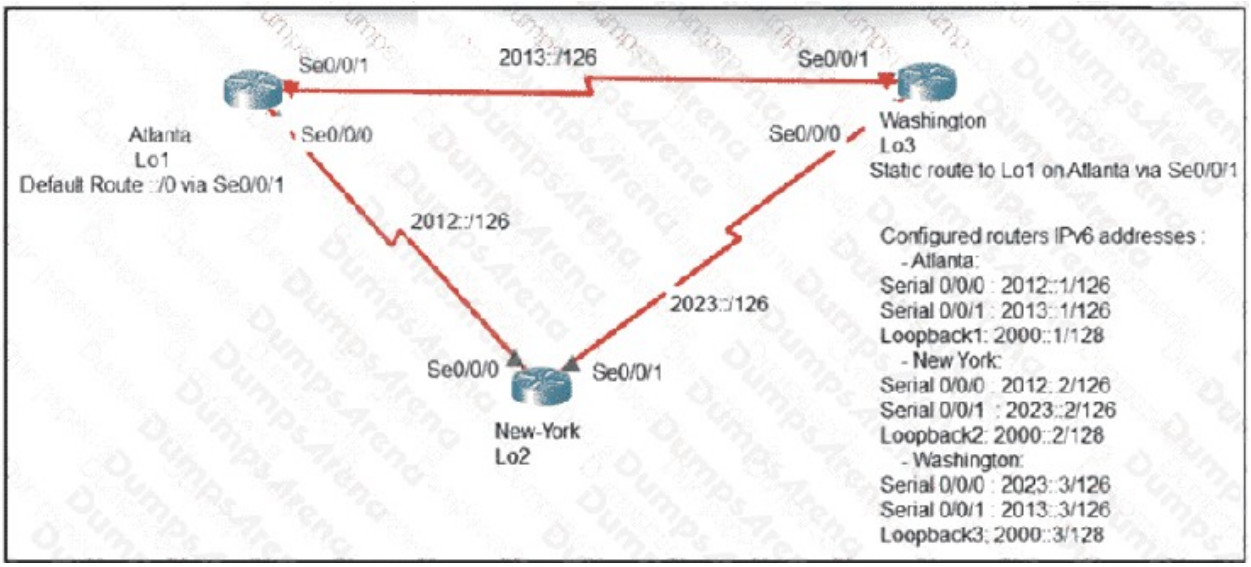
Explanation:

“Enable the WPA2 Policy option” is required because the WLAN must advertise and enforce WPA2 security before it can use a WPA2 pre-shared key. On Cisco wireless controllers, the WPA2 policy enables WPA2 protection for the WLAN; the PSK is then configured under the WPA2 authentication-key-management settings. This provides the WPA2 Personal security model, in which all authorized clients use the configured shared passphrase to derive encryption keys for the WLAN.

“Enable the MAC Filtering option” is also required to restrict association to specific client devices. Enabling MAC filtering on the WLAN allows the controller to check a joining client’s wireless MAC address against the configured permitted-client entries. Consequently, a client needs both the WPA2 PSK and an allowed MAC address to join, satisfying the requirement for WPA2 PSK access plus a specific-client admission list. Cisco documents WLAN security-policy configuration, including WPA2 and PSK settings, in its Wireless Controller WLAN Security Configuration Guide. Cisco also documents client MAC filtering and its use for controlling WLAN access in the MAC Filtering configuration guidance.

QUESTION NO: 92

Refer to the exhibit.



The New York router must be configured so that traffic to 2000::1 is sent primarily via the Atlanta site, with a secondary path via Washington that has an administrative distance of 2. Which two commands must be configured on the New York router? (Choose two.)

A. ipv6 route 2000::1/128 2012::1

B. ipv6 route 2000::1/128 2023::2 5

C. ipv6 route 2000::1/128 2012::1 5

D. ipv6 route 2000::1/128 2023::3 2

E. ipv6 route 2000::1/128 2012::2

ANSWER: D E

Explanation:

The commands `ipv6 route 2000::1/128 2012::2` and `ipv6 route 2000::1/128 2023::3 2` create the required primary and backup IPv6 routes. The `/128` prefix identifies one specific IPv6 destination host, 2000::1. The route using next hop 2012::2 forwards traffic to the directly reachable Atlanta neighbor. Because no administrative distance is explicitly configured, this static route uses the normal IPv6 static-route administrative distance of 1, making it the preferred route while the Atlanta path is available.

The route using next hop 2023::3 forwards through the directly reachable Washington neighbor and explicitly assigns administrative distance 2. This makes it a floating static route: it remains installed as a candidate backup but is not selected while the lower-distance Atlanta static route is valid. If the preferred Atlanta route can no longer be resolved or is removed from the routing table, the Washington route becomes the best available route and is used to reach 2000::1. This design provides deterministic primary-path forwarding and automatic failover without requiring a dynamic routing protocol. Cisco documents IPv6 static-route configuration and next-hop syntax in its IPv6 Static Routes Configuration Guide; Cisco also describes route preference through administrative distance in its Routing Protocol Administrative Distances reference.

QUESTION NO: 93

Which two statements describe wireless channels in the 2.4-GHz band? (Choose two.)

A. Channels 1, 6, and 11 are commonly used as nonoverlapping 20-MHz channels.

B. Non-Wi-Fi devices can cause interference in the band.

C. All 2.4-GHz channels are nonoverlapping.

D. The band supports only 5-GHz Wi-Fi clients.

E. Channel bonding is required for client association.

ANSWER: A B

Explanation:

In the 2.4-GHz band, 20-MHz wireless channels overlap with adjacent channel frequencies. In many regulatory domains, channels 1, 6, and 11 are used for a three-channel plan because their frequencies do not overlap. Using nonoverlapping channels on neighboring access points reduces adjacent-channel interference.

The 2.4-GHz band is also more susceptible to interference from common non-Wi-Fi devices and technologies, including Bluetooth devices, microwave ovens, and other equipment that operates in the same unlicensed spectrum. Cisco discusses channel planning and RF interference in its RF design and wireless interference guidance.

QUESTION NO: 94

What are two differences between optical-fiber cabling and copper cabling? (Choose two)

A. Light is transmitted through the core of the fiber

B. A BNC connector is used for fiber connections

C. The glass core component is encased in a cladding

D. Fiber connects to physical interfaces using Rj-45 connections

E. The data can pass through the cladding

ANSWER: A C

Explanation:

“Light is transmitted through the core of the fiber” is correct because optical networking conveys information as modulated pulses of light rather than as electrical current through metal conductors. A fiber-optic transmitter converts an electrical network signal to optical energy, and the receiver converts the arriving light back to an electrical signal. This light-based signaling is a fundamental physical distinction between optical fiber and copper Ethernet media.

“The glass core component is encased in a cladding” is also correct. An optical fiber is constructed with a central core surrounded by cladding. The core has a higher refractive index than the cladding, a design that guides the light signal along the fiber by total internal reflection. In common networking fiber, the core and cladding are typically made from carefully engineered glass, with protective buffer coatings and an outer jacket added around them. This structure enables fiber links to carry high-bandwidth signals over long distances while remaining resistant to electromagnetic interference. Cisco’s fiber-optics overview describes light propagation and the roles of the core and cladding, while Cisco’s CCNA exam topics include identifying the characteristics of copper and fiber physical interfaces and cabling. See Cisco: Introduction to Fiber Optics and Cisco CCNA Exam Topics.

QUESTION NO: 95

Refer to the exhibit. Which two commands when used together create port channel 10? (Choose two.)

```
Switch#show etherchannel summary
[output omitted]

Group       Port-channel        Protocol        Ports
------------+-------------------+---------------+-----------------------------------
10          Po10(SU)            LACP            Gi0/0(P)    Gi0/1(P)
20          Po20(SU)            LACP            Gi0/2(P)    Gi0/3(P)
```

A. int range g0/0-1 channel-group 10 mode active

B. int range g0/0-1 channel-group 10 mode desirable

C. int range g0/0-1 channel-group 10 mode passive

D. int range g0/0-1 channel-group 10 mode auto

E. int range g0/0-1 channel-group 10 mode on

ANSWER: A C

Explanation:

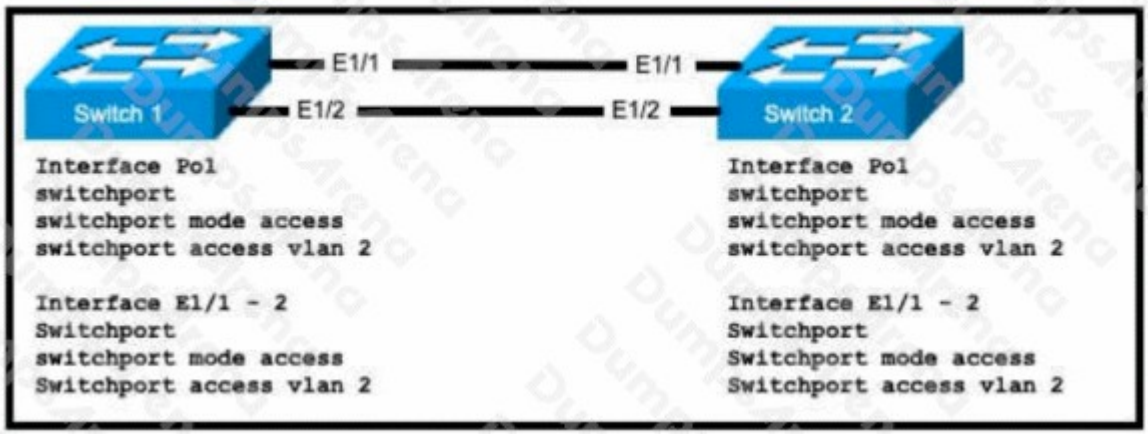
`interface range g0/0-1` followed by `channel-group 10 mode active` configures the two member interfaces to use Link Aggregation Control Protocol negotiation and assigns them to Port-Channel 10. The active mode initiates Link Aggregation Control Protocol negotiations by sending protocol data units to the peer.

`interface range g0/0-1` followed by `channel-group 10 mode passive` also assigns those interfaces to channel group 10 using Link Aggregation Control Protocol. Passive mode listens for and responds to negotiation from an active peer. Consequently, configuring active at one end of the bundled links and passive at the other end creates the negotiated EtherChannel represented by Port-Channel 10. The member links must have compatible operational settings, such as speed, duplex, and switchport characteristics, before they can be successfully bundled.

Cisco documents that Link Aggregation Control Protocol EtherChannels use active and passive negotiation modes, and a channel forms when at least one endpoint actively initiates negotiation. See Cisco's EtherChannel configuration and negotiation overview and the IOS XE EtherChannel Configuration Guide.

QUESTION NO: 96

Refer to the exhibit.



An engineer is configuring an EtherChannel using LACP between Switches 1 and 2 Which configuration must be applied so that only Switch 1 sends LACP initiation packets?

A. Switch 1 (config-if)#channel-group 1 mode on
Swrtch2(config-if)#channel-group 1 mode passive

B. Switch1(config-if)#channel-group 1 mode passive
Switch2(config-if)#channel-group 1 mode active

C. Switch1{config-if)£channel-group 1 mode active
Switch2(config-if)#channel-group 1 mode passive

D. Switch1(config-if)#channel-group 1 mode on
Switch2(config-if)#channel-group 1 mode active

ANSWER: C

Explanation:

Switch1{config-if)£channel-group 1 mode active
Switch2(config-if)#channel-group 1 mode passive is correct because LACP active mode initiates LACP negotiation. An interface configured as active periodically transmits LACP data units (LACPDUs) to discover and negotiate with a compatible peer. Therefore, configuring Switch 1 for active mode makes it the switch that sends the initiation packets required to establish the LACP EtherChannel.

An LACP interface in passive mode does not begin negotiation by itself. It listens for received LACPDUs and responds when an active LACP neighbor contacts it. Consequently, the active/passive combination successfully forms an LACP EtherChannel while ensuring that initiation originates only from Switch 1. This behavior is useful when the design specifically requires one side to control LACP negotiation initiation while retaining standards-based dynamic EtherChannel operation.

Cisco documents that LACP supports active and passive modes, and that active mode sends LACP packets whereas passive mode waits to receive them. See Cisco's EtherChannel configuration and negotiation overview and the Catalyst EtherChannel configuration guide.

QUESTION NO: 97

Refer to the exhibit.

An extended ACL has been configured and applied to router R2 The configuration failed to work as intended Which two

changes stop outbound traffic on TCP ports 25 and 80 to 10.0.20 0 26 from the 10.0.10 0/26 subnet while still allowing all other traffic? (Choose

two)

A. Add a " permit ip any any " statement to the beginning of ACL 101 for allowed traffic.

B. Add a " permit ip any any " statement at the end of ACL 101 for allowed traffic

C. The source and destination IPs must be swapped in ACL 101

D. The ACL must be configured the Gi0/2 interface inbound on R1

E. The ACL must be moved to the Gi0/1 interface outbound on R2

ANSWER: B C

Explanation:

Adding a `permit ip any any` statement at the end of ACL 101 preserves connectivity for traffic that does not match either of the intended TCP deny entries. Cisco IOS evaluates access control entries sequentially, from the first entry to the last, and stops at the first match. Traffic reaching the end of an ACL encounters the implicit deny, so an explicit final permit is required when the policy is to block only selected traffic while allow all remaining IP traffic.

The source and destination IPs must also be swapped in ACL 101 because an extended IPv4 ACL identifies the source address before the destination address. The required policy concerns TCP sessions originating in `10.0.10.0/26` and

destined for `10.0.20.0/26`, specifically destination TCP ports 25 and 80. Therefore, each deny entry must express `10.0.10.0/26` as its source and `10.0.20.0/26` as its destination, followed by the relevant destination-port match. With those address roles corrected and the final permit present, the ACL selectively prevents the specified outbound SMTP and HTTP traffic while permitting all other traffic. Cisco documents ACL processing order, the implicit deny behavior, and extended ACL source/destination syntax in its Configuring Commonly Used IP ACLs guide and IPv4 ACL Configuration Guide.

QUESTION NO: 98

What are two advantages of implementing a controller-based architecture instead of traditional network architecture? (Choose two.)

A. It allows for seamless connectivity to virtual machines.

B. It increases security against denial-of-service attacks.

C. It supports complex and high-scale IP addressing schemes.

D. It enables configuration task automation.

E. It provides increased scalability and management options.

ANSWER: D E

Explanation:

Controller-based architecture centralizes network control, policy definition, and visibility in a software controller while network devices primarily enforce the resulting configuration and policy. This model enables configuration task automation because administrators and applications can use controller APIs and workflows to define intent once and consistently deploy configuration across many devices. It reduces repetitive device-by-device changes, improves consistency, and makes routine provisioning or policy updates faster.

It also provides increased scalability and management options. A controller offers a centralized point for inventory, topology awareness, telemetry, assurance, policy administration, and orchestration. Rather than requiring operators to individually coordinate changes throughout the infrastructure, the controller can translate high-level business or network intent into device-specific configuration. This supports more consistent operations as the number of endpoints, devices, sites, and services grows. Cisco describes controller-based networking as separating management and control functions from the underlying infrastructure and using automation to simplify operations. Cisco Catalyst Center is an example of a controller platform that provides automation and assurance capabilities for enterprise networks. See Cisco Software-Defined Networking overview and Cisco Catalyst Center.

QUESTION NO: 99

When configuring IPv6 on an interface, which two IPv6 multicast groups are joined? (Choose two.)

A. 2000::/3

B. 2002::5

C. FC00::/7

D. FF02::1

E. FF02::2

F. FF02::1:FFxx:xxxx

ANSWER: D F

Explanation:

An IPv6-enabled interface joins the link-local all-nodes multicast group, **FF02::1**. This group represents every IPv6 node on the local link and is a fundamental multicast membership for IPv6 operation. The **FF02** prefix identifies a multicast address with link-local scope, ensuring that traffic sent to the group remains on the directly connected network segment. RFC 4291 specifies that every IPv6 node must join the all-nodes multicast group on each interface.

The interface also joins one or more solicited-node multicast groups in the **FF02::1:FFxx:xxxx** format. A solicited-node address is formed from the low-order 24 bits of each unicast or anycast IPv6 address assigned to the interface. Neighbor Discovery uses this membership for efficient neighbor solicitation processing, including resolving a neighbor's link-layer address and performing duplicate address detection. Rather than requiring every node to process a request sent to all nodes, solicited-node multicast limits delivery to a much smaller set of likely recipients. Cisco documents this multicast membership and its Neighbor Discovery purpose, while RFC 4291 defines both multicast address forms and required memberships. See RFC 4291, section 2.7 and Cisco IOS XE IPv6 addressing documentation.

QUESTION NO: 100

Refer to the exhibit.

```
Switch#show etherchannel summary
[output omitted]

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------
10     Po10(SU)         LACP      Gi0/0(P)    Gi0/1(P)
20     Po20(SU)         LACP      Gi0/2(P)    Gi0/3(P)
```

Which two commands were used to create port channel 10? (Choose two)

- int range g0/0-1
  channel-group 10 mode active
- int range g0/0-1
  channel-group 10 mode desirable
- int range g0/0-1
  channel-group 10 mode passive
- int range g0/0-1
  channel-group 10 mode auto
- int range g0/0-1
  channel-group 10 mode on

A. Option A

B. Option B

C. Option C

D. Option D

E. Option E

F. `interface port-channel 10` and `channel-group 10 mode active` (or another `channel-group 10 mode ...` variant) were used to create Port-channel10.

G. `channel-group 10 mode active`

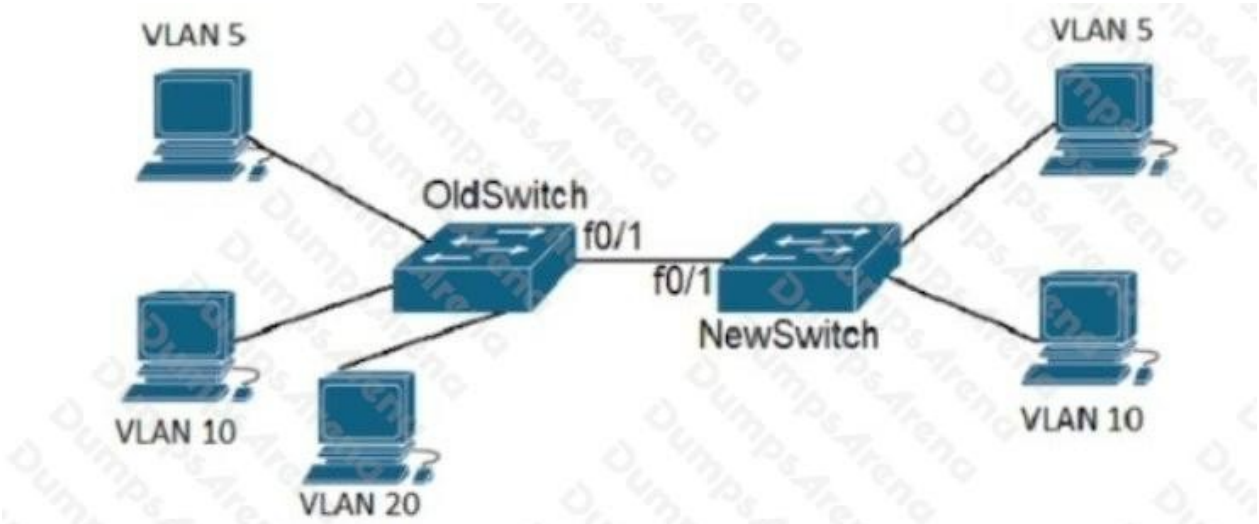
H. `channel-group 10 mode passive`

ANSWER: G H

Explanation:

`channel-group 10 mode active` and `channel-group 10 mode passive` are the complementary Cisco IOS commands used on the member interfaces at opposite ends of an LACP EtherChannel. The `channel-group 10` portion assigns physical switch ports to EtherChannel group 10 and causes the logical `Port-channel10` interface to be created automatically when it does not already exist. The `active` and `passive` keywords select Link Aggregation Control Protocol negotiation: an active interface initiates LACP negotiation, while a passive interface listens for and responds to LACP packets. Because at least one side is active, the two bundled physical links can negotiate and form the LACP port channel identified as group 10. These commands are entered under each physical member interface, or under an interface range containing the intended member ports; Layer 2 trunk and VLAN attributes are then normally applied to the resulting port-channel interface. Cisco documents that the `channel-group` command creates the EtherChannel and that the active/passive modes are the LACP negotiation modes. See the Cisco EtherChannel Configuration Notes and the Cisco IOS XE EtherChannel Configuration Guide.

QUESTION NO: 101




```
OldSwitch(config)#interface fastEthernet 0/1
OldSwitch(config-if)#switchport mode trunk
OldSwitch(config-if)#switchport trunk allowed vlan 5,10
OldSwitch(config-if)#switchport trunk native vlan 15
**output suppressed**

NewSwitch(config)#interface fastEthernet 0/1
NewSwitch(config-if)#switchport mode trunk
NewSwitch(config-if)#switchport trunk encapsulation isl
NewSwitch(config-if)#switchport trunk allowed vlan 5,10
NewSwitch(config-if)#switchport trunk native vlan 15
```

Refer to the exhibit A new VLAN and switch are added to the network. A remote engineer configures OldSwitch and must ensure that the configuration meets these requirements:

• accommodates current configured VLANs

• expands the range to include VLAN 20

• allows for IEEE standard support for virtual LANs

Which configuration on the NewSwitch side of the link meets these requirements?

A. switch port mode dynamic channel group 1 mode active
switchport trunk allowed vlan 5,10,15, 20

B. no switchport mode trunk switchport trunk encapsulation isl switchport mode access vlan 20

C. switchport nonegotiate
no switchport trunk allowed vlan 5,10 switchport trunk allowed vlan 5,10,15,20

D. no switchport trunk encapsulation isl switchport trunk encapsulation dot1q switchport trunk allowed vlan add 20

ANSWER: D

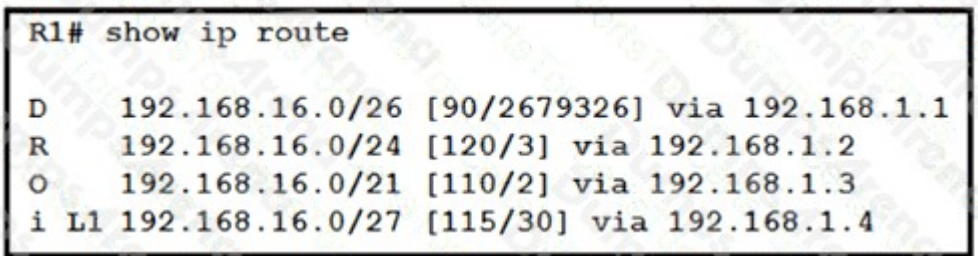
Explanation:

`no switchport trunk encapsulation isl switchport trunk encapsulation dot1q switchport trunk allowed vlan add 20` is correct because it makes IEEE 802.1Q (dot1q) the trunk encapsulation used on the interswitch link. IEEE 802.1Q is the standards-based VLAN tagging method required for a trunk that carries multiple VLANs between switches. The configuration explicitly removes ISL encapsulation and selects dot1q, ensuring compatibility with the IEEE VLAN trunking standard.

The `switchport trunk allowed vlan add 20` command is equally important: `add` appends VLAN 20 to the interface's existing allowed-VLAN list rather than replacing that list. Thus, VLANs already permitted across the trunk remain available while the new VLAN is added. This fulfills both the requirement to preserve the current VLAN carriage and the requirement to extend trunk access to VLAN 20. Cisco documents that 802.1Q trunks identify VLAN traffic with tags and that the allowed-VLAN configuration controls which VLANs can traverse a trunk. See Cisco's IEEE 802.1Q VLAN Trunking overview and the Cisco VLAN trunk configuration guide.

QUESTION NO: 102

Refer to the exhibit.



Which route does R1 select for traffic that is destined to 192.168.16.2?

A. 192.168.16.0/21

B. 192.168.16.0/24

C. 192.168.16.0/26

D. 192.168.16.0/27

ANSWER: D

Explanation:

192.168.16.0/27 is selected because it is the most specific route that matches destination address 192.168.16.2. A /27 prefix uses the mask 255.255.255.224 and covers addresses from 192.168.16.0 through 192.168.16.31, so the destination is within that route's address range. The other displayed prefixes also encompass the destination, but Cisco routers apply longest-prefix matching first: among all matching entries, the route with the greatest prefix length is installed as the forwarding choice. Because /27 contains more network bits than /26, /24, and /21, it identifies the narrowest, most specific destination network. This decision is made before administrative distance and routing-protocol metrics are considered; those values matter only when routes represent the same prefix length and network. Therefore, R1 forwards packets for 192.168.16.2 according to the 192.168.16.0/27 entry. This behavior is a core IPv4 forwarding requirement and is described in the route-selection discussion in RFC 1812, Section 5.2.4. It is also fundamental knowledge within Cisco's CCNA 200-301 exam routing objectives.

QUESTION NO: 103

What makes Cisco DNA Center different from traditional network management applications and their management of networks?

A. Its modular design allows the implementation of different versions to meet the specific needs of an organization.

B. It only supports auto-discovery of network elements in a greenfield deployment.

C. It omits support high availability of management functions when operating in cluster mode.

D. It abstracts policy from the actual device configuration.

ANSWER: D

Explanation:

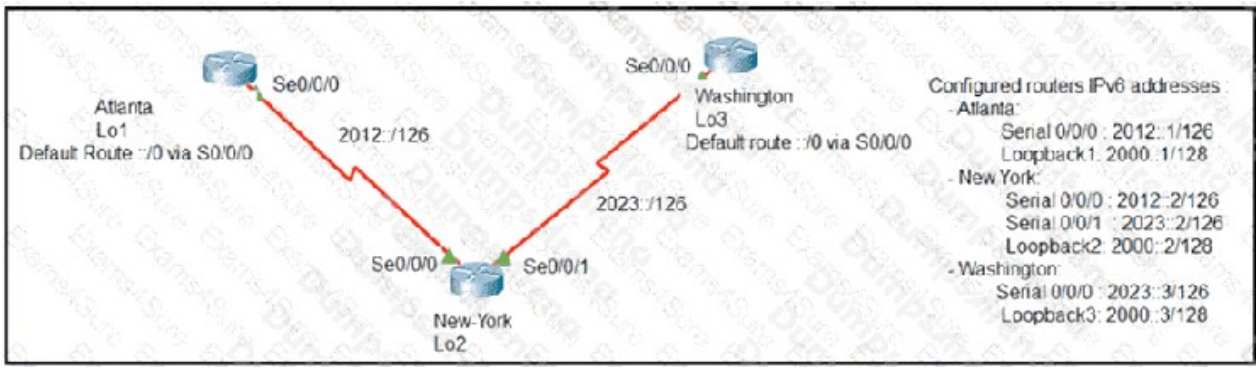
Cisco DNA Center, now known as Cisco Catalyst Center, differentiates itself through intent-based networking and policy abstraction. Rather than requiring administrators to define and maintain detailed, device-specific configuration commands as the primary management method, an administrator can express the desired network outcome centrally—for example, segmentation, user and device access policy, or application experience requirements. The platform translates that intent into appropriate configurations for the supported network infrastructure and continuously evaluates operational telemetry to provide assurance that the intended outcome is being achieved.

This approach connects network design, policy, provisioning, automation, and assurance in a single controller-driven workflow. Policy is therefore expressed at a higher level than individual switch, router, or wireless-controller commands. Cisco describes Catalyst Center as a network management system that simplifies network operations through automation and assurance, helping organizations implement intent consistently while using analytics to identify and resolve issues. This policy-to-configuration abstraction is the defining distinction from traditional management applications that focus primarily on individual devices and their configurations.

See Cisco's Cisco Catalyst Center overview and the Cisco Catalyst Center Administrator Guide.

QUESTION NO: 104

Refer to the exhibit.



The loopback1 interface of the Atlanta router must reach the loopback3 interface of the Washington router.

A. ipv6 route 2000::1/128 2012::2

B. ipv6 route 2000::1/128 2012::1

C. ipv6 route 2000:3 123 s0/0/0

D. ipv6 route 2000::3/128 2023::3

E. ipv6 route 2000::1/128 s0/0/1

ANSWER: B D

Explanation:

The required static routes are `ipv6 route 2000::1/128 2012::1` and `ipv6 route 2000::3/128 2023::3`. The transit router requires a host route for each remote loopback address so that it can forward traffic in both directions between Atlanta's Loopback1 address, `2000::1/128`, and Washington's Loopback3 address, `2000::3/128`. A `/128` prefix is the IPv6 host-route prefix length, making each route an exact match for its respective loopback interface rather than a route to a broader address range.

The next-hop addresses identify the adjacent router on the appropriate point-to-point IPv6 segment. For traffic destined to Atlanta's loopback, the route uses `2012::1`, which is Atlanta's reachable next-hop address on the link toward Atlanta. For traffic destined to Washington's loopback, the route uses `2023::3`, the reachable next-hop address on the link toward Washington. With these two routes installed, the transit device has forwarding information for traffic toward either endpoint, allowing packet delivery and return traffic between the loopbacks. Cisco documents the `ipv6 route` command and IPv6 static-route configuration behavior in its IOS IP Routing command reference. Static routing is also within the IP Connectivity domain of the CCNA exam topics.

QUESTION NO: 105

What is a role of access points in an enterprise network?

A. connect wireless devices to a wired network

B. support secure user logins to devices or the network

C. integrate with SNMP in preventing DDoS attacks

D. serve as a first line of defense in an enterprise network

ANSWER: A

Explanation:

Access points connect wireless devices to a wired network. An enterprise access point provides the radio interface through which Wi-Fi clients, such as laptops, phones, and handheld scanners, join an IEEE 802.11 wireless LAN. After a client associates to the SSID, the access point forwards the client's traffic between the wireless medium and the wired Ethernet network. This function enables wireless users to reach enterprise resources, services, and external networks through the existing LAN infrastructure.

In a controller-based Cisco wireless deployment, the access point may establish a control-and-data tunnel to a wireless LAN controller, which centrally applies WLAN policies and manages the access point. In other deployment models, such as Cisco Embedded Wireless Controller or autonomous-style local forwarding designs, traffic handling can differ; however, the fundamental AP role remains providing wireless access and bridging or forwarding client traffic toward the wired network. Cisco describes access points as devices that create wireless LAN connectivity for client devices, while its WLAN overview explains that a wireless LAN connects users through radio access to a network infrastructure. See Cisco Wireless Access Points and Cisco: What Is a WLAN?.

QUESTION NO: 106

What is the RFC 4627 default encoding for JSON text?

A. UCS-2

B. GB18030

C. UTF-8

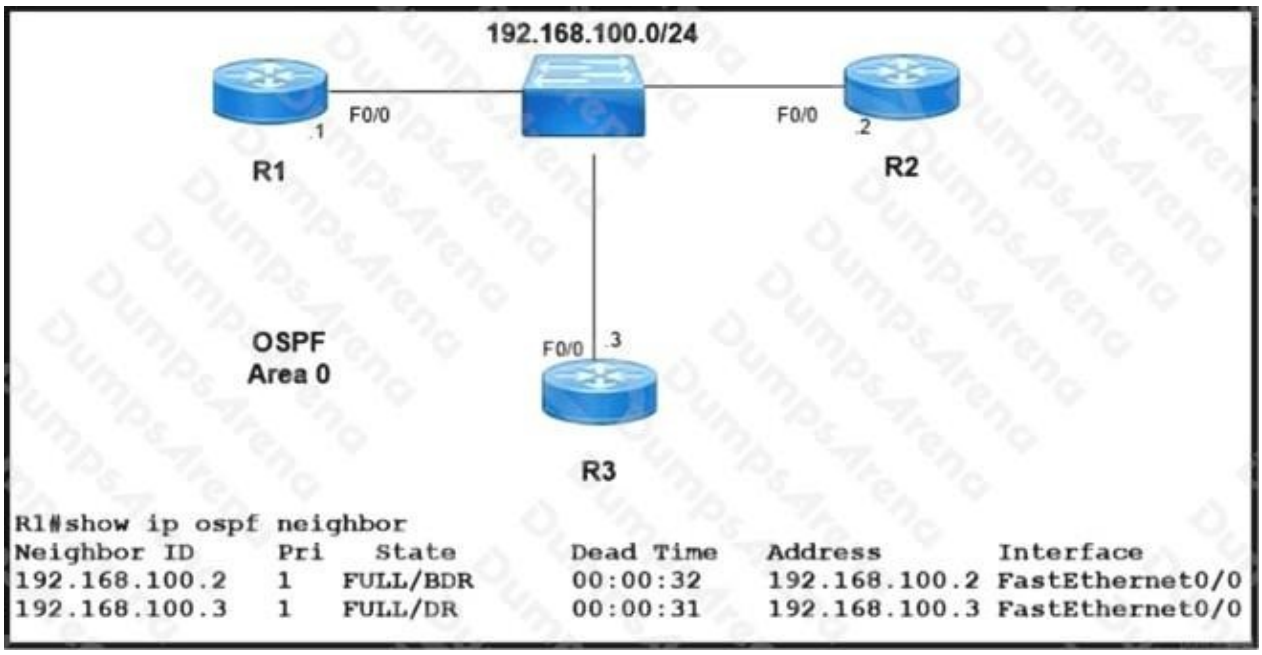
D. Hex

ANSWER: C

Explanation:

UTF-8 is the default encoding specified by RFC 4627 for JSON text. The RFC defines JSON as an interchange format based on Unicode and states that a JSON text *shall* be encoded in Unicode. It then establishes UTF-8 as the default encoding, while allowing UTF-16 and UTF-32 representations in situations where the encoding is indicated by a byte-order mark. This default was important for interoperability: implementations receiving JSON without an explicit encoding signal could consistently treat the content as UTF-8. UTF-8 also represents all Unicode characters while remaining compatible with the ASCII range used heavily by JSON syntax, including braces, brackets, quotation marks, commas, colons, digits, and the literal names used for Boolean and null values. Consequently, UTF-8 is the practical and standards-based baseline for JSON exchanged between independent systems. Later JSON guidance in RFC 8259 further strengthens this operational convention by requiring JSON exchanged between systems outside a closed ecosystem to use UTF-8. Therefore, UTF-8 is the correct answer for the RFC 4627 default encoding. See RFC 4627, Section 3 and RFC 8259, Section 8.1.

QUESTION NO: 107



Refer to the exhibit. Which two configurations must the engineer apply on this network so that R1 becomes the DR? (Choose two.)

A. R3(config)#interface fastethernet 0/0 R3(config-if)#ip ospf priority 0

B. R1(config)#router ospf 1
R1(config-router)#router-id 192.168.100.1

C. R1(config)#interface fastethernet 0/0 R1(config-if)#ip ospf priority 200

D. R1(config)#interface fastethernet 0/0 R1(config-if)#ip ospf priority 0

E. R3(config)#interface fastethernet 0/0 R3(config-if)#ip ospf priority 200

ANSWER: A C

Explanation:

On an OSPF broadcast multiaccess network, such as the Ethernet segment in the exhibit, routers elect a designated router and backup designated router to reduce adjacency overhead. OSPF compares each participating interface's priority first; the router with the highest nonzero priority is elected DR. The router ID is used only as a tiebreaker when eligible routers have the same priority. The configuration `R1(config)#interface fastethernet 0/0 R1(config-if)#ip ospf priority 200` gives R1 a priority substantially higher than the default OSPF interface priority of 1, making R1 the preferred eligible DR candidate on that segment. The configuration `R3(config)#interface fastethernet 0/0 R3(config-if)#ip ospf priority 0` removes R3 from both DR and BDR eligibility. A priority of zero does not prevent normal OSPF neighbor adjacency or route exchange, but it ensures that interface cannot be selected for either election role. Together, these settings establish R1 as the intended DR and exclude R3 from competing for the role. OSPF DR/BDR elections are nonpreemptive, so if an election has already occurred, the engineer must reset the OSPF adjacency or process, as appropriate for the maintenance window, for the changed election parameters to take effect. Cisco documents the election rules and the priority-zero behavior in its OSPF DR/BDR election guidance and IOS OSPF command reference.

QUESTION NO: 108

An engineer must configure a core router with a floating static default route to the backup router at 10.200.0.2. Which command meets the requirements?

A. ip route 0.0.0.0 0.0.0.0 10.200.0.2 1

B. ip route 0.0.0.0 0.0.0.0 10.200.0.2 10

C. ip route 0.0.0.0 0.0.0.0 10.200.0.2

D. ip route 0.0.0.0 0.0.0.0 10.200.0.2 floating

ANSWER: B

Explanation:

`ip route 0.0.0.0 0.0.0.0 10.200.0.2 10` configures an IPv4 static default route whose next-hop address is the backup router, 10.200.0.2. The two zero masks define the default prefix, so this route can provide forwarding for destinations that do not have a more-specific route in the routing table. The final value, `10`, explicitly sets the route's administrative distance. A static route normally has an administrative distance of 1; assigning 10 makes this route less preferred than a primary route with an administrative distance lower than 10. Consequently, when that preferred primary route is present, this backup default route is not selected. If the primary route is withdrawn or becomes unusable, the router can install and use this static default route, which is the intended floating-static-route behavior. Cisco IOS supports adding an administrative-distance value at the end of the `ip route` command specifically to control route preference. In a production design, the configured distance must always be greater than the actual administrative distance of the intended primary route; here, 10 provides the expected backup value where the primary default route is a normal static route. See Cisco's Static Routes Configuration Guide and its Administrative Distance overview.

QUESTION NO: 109

A network engineer starts to implement a new wireless LAN by configuring the authentication server and creating the dynamic interface. What must be performed next to complete the basic configuration?

A. Create the new WLAN and bind the dynamic interface to it.

B. Configure high availability and redundancy for the access points.

C. Enable Telnet and RADIUS access on the managoment interface.

D. Install the management interface and add the management IP.

ANSWER: A

Explanation:

Create the new WLAN and bind the dynamic interface to it. is the required next step. On an AireOS-based Cisco Wireless LAN Controller, the RADIUS server configuration supplies the external authentication method, while a dynamic interface defines the client VLAN, gateway, and controller port mapping used for a specific wireless network. These elements do not make a client-facing SSID available by themselves. The WLAN must be created with its SSID and security settings, then mapped to the appropriate dynamic interface through the WLAN's interface setting. This association ensures that clients joining the SSID are placed into the intended VLAN and that their authentication requests use the configured AAA/RADIUS infrastructure. In a basic enterprise WLAN deployment, the WLAN is also enabled after its required parameters are configured so that access points can advertise it. Cisco's controller configuration documentation describes WLAN creation as the step where the WLAN is assigned an interface/VLAN mapping, while its security configuration can reference the previously defined RADIUS server. See the Cisco Wireless Controller Configuration Guide, Release 8.10 and Cisco's WLAN security and RADIUS configuration overview.

QUESTION NO: 110

What are two differences between WPA2 and WPA3 wireless security? (Choose two.)

A. WPA3 um AES for stronger protection than WPA2 which uses SAE

B. WPA2 uses 1 M-bit key encryption and WPA3 requires 256-brt key encryption

C. WPA3 uses AES for stronger protection than WPA2 which uses TKIP WPA3 uses

D. SAE tor stronger protection than WPA2 which uses AES

E. WPA2 uses 12B-M key encryption and WPA3 supports 128 bit and 192 bit key encryption

F. WPA3-Personal uses SAE (Simultaneous Authentication of Equals), while WPA2-Personal uses a pre-shared key (PSK) authentication method.

G. WPA3-Enterprise supports an optional 192-bit security suite, while WPA2 commonly uses 128-bit AES-CCMP security.

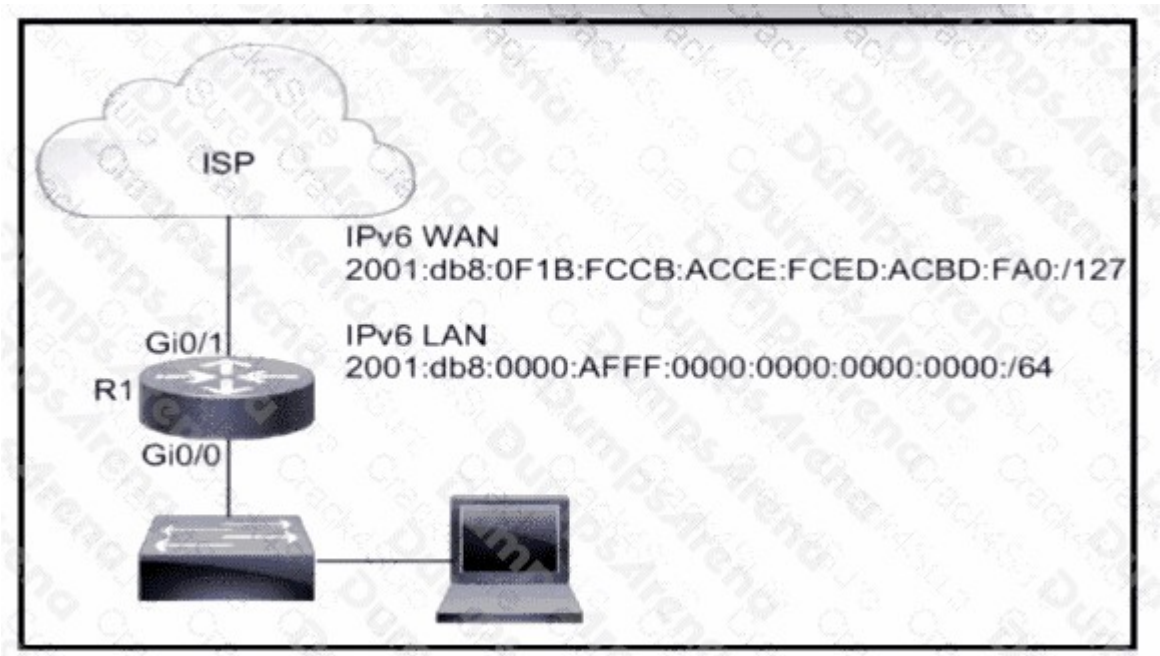
ANSWER: F G

Explanation:

WPA3-Personal uses Simultaneous Authentication of Equals (SAE), whereas WPA2-Personal uses a pre-shared key (PSK) passphrase with the traditional four-way handshake. SAE is a password-authenticated key exchange that provides stronger protection against offline password-guessing attempts. An attacker capturing WPA3-SAE authentication traffic cannot use that capture to repeatedly test password guesses offline in the manner possible with a captured WPA2-Personal handshake. SAE also provides forward secrecy for the session keys, so compromise of a password later does not expose previously captured session traffic.

WPA3-Enterprise also defines an optional 192-bit security mode for environments needing higher-assurance wireless protection. This mode uses the Commercial National Security Algorithm (CNSA) Suite–aligned cryptographic requirements. WPA2 commonly uses AES-CCMP with 128-bit cryptographic strength; WPA3 does not require 192-bit mode in every deployment, but its Enterprise specification makes that stronger suite available. These differences distinguish the authentication improvement in WPA3-Personal from the additional high-security cryptographic option in WPA3-Enterprise. Cisco summarizes WPA3 and its SAE and Enterprise enhancements in its WPA3 and Enhanced Open overview. The Wi-Fi Alliance also describes WPA3 security capabilities at its Wi-Fi security page.

QUESTION NO: 111



Refer to the exhibit. IPv6 must be implemented on R1 to the ISP The uplink between R1 and the ISP must be configured with a manual assignment, and the LAN interface must be self-provisioned Both connections must use the applicable IPv6 networks Which two configurations must be applied to R1? (Choose two.)

A. interface Gi0/1
ipv6 address 2001:db8:0F1B:FCCB:ACCE:FCED:ABCD:FA02/127

B. interface Gi0/0
ipv6 address 2001:db8:1:AFFF::/64 eui-64

C. interface Gi0/1
ipv6 address 2001:db8:0F1B:FCCB:ACCE:FCED:ABCD:FA00/127

D. interface Gi0/0
ipv6 address 2001:db8:0:AFFF::/64 eui-64

E. interface Gi0/0
ipv6 address 2001:db8:0F1B:FCCB:ACCE:FCED:ABCD:FA03/127

ANSWER: C D

Explanation:

`interface Gi0/1`
`ipv6 address 2001:db8:0F1B:FCCB:ACCE:FCED:ABCD:FA00/127` is the required manually assigned address for the R1-to-ISP point-to-point uplink. A /127 prefix provides exactly two usable IPv6 interface addresses, which is appropriate for a point-to-point link. The selected address is in the applicable uplink subnet and places R1 on the required two-address network. Cisco supports /127 prefixes on point-to-point IPv6 links, consistent with the operational guidance defined in RFC 6164.

`interface Gi0/0`
`ipv6 address 2001:db8:0:AFFF::/64 eui-64` is required on the LAN interface because EUI-64 self-provisions the interface identifier from the interface MAC address. The command combines the given /64 LAN prefix with an automatically derived 64-bit interface ID, rather than requiring an administrator to choose a complete host address manually. A /64 prefix is the standard LAN prefix length for IPv6 stateless address configuration and EUI-64 addressing. Cisco documents the `ipv6 address ... eui-64` command and its automatic interface-ID generation behavior in its IPv6 Addressing Configuration Guide.

QUESTION NO: 112



Refer to the exhibit. IPv6 must be implemented on R1 to the IS

P. The uplink between R1 and the ISP must be configured with a manual assignment, and the LAN interface must be self-provisioned. Both connections must use the applicable IPv6 networks. Which two configurations must be applied to R1? (Choose two.)

A. interface Gi0/0
ipv6 address 2001:db8:0F1B:FCCB:ACCE:FCED:ABCD:FA03:/127

B. interface Gi0/0
ipv6 address 2001:db8:0:AFFF::/64 eui-64

C. interface Gi0/1
ipv6 address 2001:db8:0F1B:FCCB:ACCE:FCED:ABCD:FA02:/127

D. interface Gi0/0
ipv6 address 2001:db8:1:AFFF::/64 eui-64

E. interface Gi0/1
ipv6 address 2001:db8:0F1B:FCCB:ACCE:FCED:ABCD:FA00:/127

F. The uplink between R1 and the ISP must be configured with a manual assignment, and the LAN interface must be self-provisioned. Both connections must use the applicable IPv6 networks. Which two configurations must be applied to R1? (Choose two.)

G. interface Gi0/1
ipv6 address 2001:db8:0F1B:FCCB:ACCE:FCED:ABCD:FA02/127

ANSWER: B G

Explanation:

The required configuration combines an explicitly configured IPv6 address on the ISP-facing point-to-point link with an automatically generated interface identifier on the LAN-facing interface. `interface Gi0/1` `ipv6 address 2001:db8:0F1B:FCCB:ACCE:FCED:ABCD:FA02/127` supplies a complete address and prefix length directly, so it is a manual IPv6 assignment. The `/127` prefix is appropriate for the two-address point-to-point subnet shown: the address ending in `FA02` belongs to the same /127 pair as the peer address ending in `FA03`.

`interface Gi0/0`
`ipv6 address 2001:db8:0:AFFF::/64 eui-64` configures the applicable LAN /64 while directing Cisco IOS to generate the final 64-bit interface identifier with modified EUI-64 processing. This fulfills the self-provisioned requirement

without manually entering the complete host portion of the LAN address. A /64 prefix is the standard prefix size for EUI-64 interface-ID construction and IPv6 LAN addressing. Cisco IOS supports the `eui-64` keyword on an IPv6 interface address command for this purpose. See Cisco's IPv6 configuration guidance and RFC 6164 for IPv6 /127 point-to-point prefixes.

QUESTION NO: 113

hostname CPE

service password-encryption

ip domain name ccna.cisco.com ip name-server 198.51.100.210

crypto key generate rsa modulus 1024 username admin privilege 15 secret s0m3s3cr3t line vty 0 4

transport input ssh

login local

Refer to the exhibit. An engineer executed the script and added commands that were not necessary for SSH and now must remove the commands. Which two commands must be executed to correct the configuration? (Choose two.)

A. no ip name-serveer 198.51.100.210

B. no login local

C. no service password-encryption

D. no ip domain mame ccna.cisco.com

E. no hostname CPE

ANSWER: A C

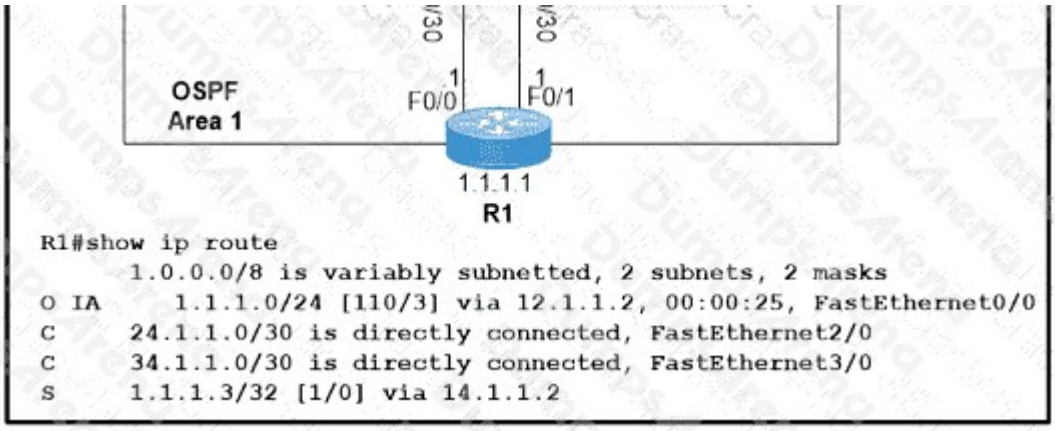
Explanation:

The commands `no ip name-serveer 198.51.100.210` and `no service password-encryption` remove configuration that is not required to provide SSH access to the device. SSH requires an RSA key pair, and Cisco IOS uses the configured hostname and domain name when generating those keys. It also requires an authentication method on the VTY lines; here, `login local` uses the configured local username and secret. Restricting VTY transport to SSH with `transport input ssh` ensures remote access uses SSH rather than Telnet.

An `ip name-server` command only identifies a DNS resolver that the router can query when it must translate hostnames into IP addresses. SSH server operation does not depend on the router having a DNS resolver configured, particularly when administrators connect directly to the device IP address. The `service password-encryption` command applies Cisco IOS reversible type 7 obfuscation to eligible plaintext passwords in the configuration. It does not enable, secure, encrypt, or otherwise support the SSH protocol; the local `secret` already uses a password-hash mechanism independent of that service. Therefore, removing these two commands leaves all essential SSH prerequisites and VTY access controls intact. See Cisco's SSH configuration example and the Cisco IOS SSH configuration guide.

QUESTION NO: 114

Refer to the exhibit.



Which two values does router R1 use to determine the best path to reach destinations in network 10.0.0.0/8? (Choose two.)

A. longest prefix match

B. highest administrative distance

C. highest metric

D. lowest metric

E. lowest cost to reach the next hop

ANSWER: A D

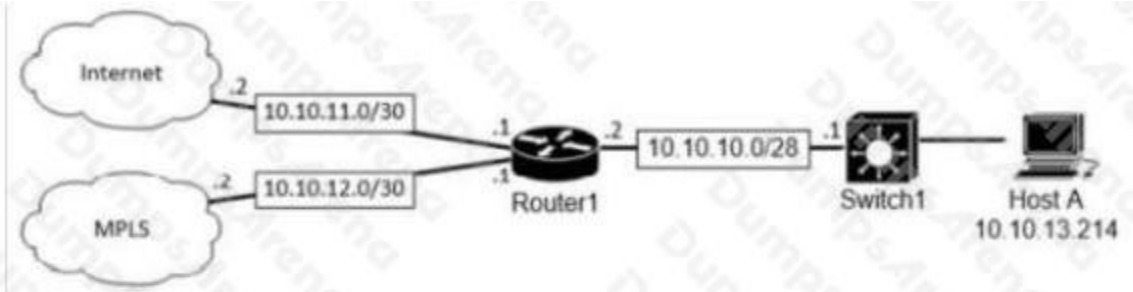
Explanation:

Router R1 first applies **longest prefix match** when forwarding a packet. The destination IP address is compared with all matching entries in the routing table, and the route with the most specific matching prefix length is selected. For example, a route for a subnet within 10.0.0.0/8 is preferred over the less-specific 10.0.0.0/8 route whenever the destination address matches that subnet. This forwarding behavior ensures traffic follows the most precise destination information available.

When candidate routes have the same destination prefix length and are eligible for comparison, the router selects the route with the **lowest metric**. A metric is the route-selection value used by a routing protocol to represent the relative desirability of paths learned through that protocol. Lower metric values represent more preferable paths, so the router installs or uses the path with the smallest metric among equivalent prefix candidates. Cisco documents route selection and longest-prefix forwarding behavior in its Routing Protocols and Administrative Distance guidance and describes routing-table route preference concepts in the Cisco IOS IP Routing Configuration Guide.

QUESTION NO: 115

Refer to the exhibit.




```
Router1#show ip route
Gateway of last resort is 10.10.11.2 to network 0.0.0.0

     209.165.200.0/27 is subnetted, 1 subnets
B       209.165.200.224 [20/0] via 10.10.12.2, 03:22:14
     209.165.201.0/27 is subnetted, 1 subnets
B       209.165.201.0 [20/0] via 10.10.12.2, 02:26:33
     209.165.202.0/27 is subnetted, 1 subnets
B       209.165.202.128 [20/0] via 10.10.12.2, 02:26:03
    10.0.0.0/8 is variably subnetted, 8 subnets, 4 masks
C      10.10.10.0/28 is directly connected, GigabitEthernet0/0
C      10.10.11.0/30 is directly connected, FastEthernet2/0
C      10.10.12.0/30 is directly connected, GigabitEthernet0/1
O      10.10.13.0/25 [110/2] via 10.10.10.1, 00:00:04, GigabitEthernet0/0
O      10.10.13.128/28 [110/2] via 10.10.10.1, 00:00:04, GigabitEthernet0/0
O      10.10.13.144/28 [110/2] via 10.10.10.1, 00:00:04, GigabitEthernet0/0
O      10.10.13.160/29 [110/2] via 10.10.10.1, 00:00:04, GigabitEthernet0/0
O      10.10.13.208/29 [110/2] via 10.10.10.1, 00:00:04, GigabitEthernet0/0
S*   0.0.0.0/0 [1/0] via 10.10.11.2
```

What is the prefix length for the route that router1 will use to reach host A?

A. /25

B. /27

C. /28

D. /29

ANSWER: D

Explanation:

The prefix length is **/29**. Router1 performs a longest-prefix-match lookup for the destination address of host A. When more than one routing-table entry can match a destination, Cisco routers select the matching entry with the greatest number of leading network bits. This is the most specific route because it represents the smallest address range among the applicable entries.

A /29 mask has 29 network bits and contains eight IPv4 addresses in its block, so it is more specific than the broader matching prefixes shown in the exhibit. After router1 identifies that host A belongs to the /29 destination range, it uses the forwarding information associated with that route to send the packet to the appropriate next hop or outgoing interface. This lookup behavior applies regardless of whether the route was learned dynamically or configured statically; route source and administrative distance are considered when selecting between routes for the same prefix, while the longest matching prefix determines which destination prefix is used for forwarding.

Cisco describes this forwarding behavior in its route selection and longest-match documentation and in the Cisco IOS IP routing configuration guide.

QUESTION NO: 116

Refer to the exhibit.

Between which zones do wireless users expect to experience intermittent connectivity?

A. between zones 1 and 2

B. between zones 2 and 5

C. between zones 3 and 4

D. between zones 3 and 6

ANSWER: D

Explanation:

between zones 3 and 6 is correct because the RF coverage shown in the exhibit creates an unreliable transition area between those two zones. As a wireless client moves through this boundary, its signal from the current access point can fall below a usable level before it has a sufficiently strong, stable signal from the next access point. This produces a coverage gap or marginal-overlap area, where frames may be retried or lost and the client may repeatedly attempt to reassociate or roam. The resulting user experience is intermittent connectivity, such as brief drops in data traffic, inconsistent application performance, or disrupted voice calls. A wireless design should provide sufficient overlap at the target receive-signal level so clients can roam predictably while maintaining service. Cisco's wireless design guidance emphasizes using an RF site survey to validate coverage, capacity, interference, and overlap requirements before deployment. See Cisco's wireless LAN site-survey guidance and the Cisco Enterprise Mobility Design Guide.

QUESTION NO: 117

What are two examples of multifactor authentication? (Choose two.)

A. single sign-on

B. soft tokens

C. passwords that expire

D. shared password repository

E. unique user knowledge

ANSWER: B E

Explanation:

Multifactor authentication requires a user to prove identity with factors from at least two distinct categories. The standard categories include something the user knows, something the user has, and something the user is. "Soft tokens" represent the possession category: an authenticator application or software token on a registered device generates or receives a time-based, one-time verification code. The device or enrolled application is the item the user possesses.

"Unique user knowledge" represents the knowledge category, such as a unique password, PIN, or passphrase known by the individual user. When unique user knowledge is used together with a soft token, the authentication process requires both knowledge and possession evidence. This combination meets the core MFA requirement because the factors are independent categories rather than multiple instances of the same category. For example, entering a password and then supplying a current code from an authenticator app is a common MFA workflow used to protect network, cloud, and administrative access.

Cisco describes MFA as using multiple verification methods to validate user identity, while NIST explains the importance of combining different authentication factors. See Cisco Multi-Factor Authentication and NIST Multi-Factor Authentication guidance.

QUESTION NO: 118

Which port type supports the spanning-tree portfast command without additional configuration?

A. Layer 3 main interfaces

B. Layer 3 subinterfaces

C. trunk ports

D. access ports

ANSWER: D

Explanation:

The correct answer is **access ports**. PortFast is designed for an edge-facing Layer 2 switchport that connects directly to a single end device, such as a workstation, printer, or IP phone. When PortFast is enabled on an access interface with the interface command `spanning-tree portfast`, the port bypasses the normal spanning-tree listening and learning transition delays and moves to the forwarding state immediately when the physical link comes up. This allows a directly attached host to obtain network connectivity promptly, including timely DHCP access.

An access port carries traffic for one VLAN and is the standard switchport mode for this classic PortFast use case. Administrators should deploy PortFast only where an unintended Layer 2 connection cannot create a bridging loop. Cisco also recommends enabling BPDU Guard on PortFast-enabled edge ports so that the interface is placed into an error-disabled state if it receives BPDUs, helping protect the spanning-tree topology. Cisco documents PortFast configuration on access interfaces and its role in allowing end-device ports to transition directly to forwarding: Cisco: Using PortFast and Other Commands to Fix Workstation Startup Connectivity Delays. See also the Cisco IOS LAN Switching Command Reference.

QUESTION NO: 119

What are two features of the DHCP relay agent? (Choose two.)

A. assigns DNS locally and then forwards request to DHCP server

B. permits one IP helper command under an individual Layer 3 interface

C. allows only MAC-to-IP reservations to determine the local subnet of a client

D. minimizes the necessary number of DHCP servers

E. configured under the Layer 3 interface of a router on the client subnet

ANSWER: D E

Explanation:

A DHCP relay agent enables DHCP clients on one IP subnet to obtain addresses and other DHCP options from a server located on a different subnet. This permits DHCP services to be centralized rather than deploying and maintaining a separate DHCP server for every client VLAN or network. Therefore, it **minimizes the necessary number of DHCP servers** while still allowing each subnet to receive an appropriate address scope and configuration.

On Cisco IOS and IOS XE devices, relay operation is configured on the Layer 3 interface that receives DHCP broadcasts from the client subnet. This can be a routed physical interface or a switched virtual interface for a VLAN. The interface is configured with an `ip helper-address` pointing to the DHCP server. When a client sends a DHCP broadcast, the relay agent forwards the request toward the specified server and includes subnet-identifying relay information so the server can select the suitable scope. Thus, being **configured under the Layer 3 interface of a router on the client subnet** accurately describes the normal Cisco implementation.

Cisco documents the use of `ip helper-address` to forward DHCP/BOOTP requests received on an interface to a remote server in its DHCP Relay Agent and IP Helper Address documentation and its IOS XE DHCP Relay Agent Configuration Guide.

QUESTION NO: 120

Which two HTTP methods are suitable for actions performed by REST-based APIs? (Choose two.)

A. REMOVE

B. REDIRECT

C. OPOST

D. GET

E. UPOP

F. POST

ANSWER: D F

Explanation:

GET and **POST** are standardized HTTP request methods and are suitable for REST-based API operations. GET requests retrieve a current representation of a resource, such as a device inventory item or a list of interfaces. HTTP defines GET as a safe method, meaning its intended semantics are read-only and do not request a state change on the origin server. This makes it the normal REST API method for reading resources.

POST requests submit a representation to a resource for processing. In REST APIs, POST is commonly used to create a subordinate resource in a collection, such as adding a new item to an existing collection, or to invoke an operation that requires server-side processing. POST is not required to be idempotent: sending the same POST more than once can result in multiple created resources or multiple processing actions. API designers communicate the precise resource paths, request bodies, response codes, and authentication requirements through the API specification, but both methods are fundamental parts of HTTP-based REST interfaces.

The HTTP Semantics specification defines the meaning and properties of these request methods. See RFC 9110: HTTP Semantics and MDN Web Docs: HTTP request methods.

QUESTION NO: 121

Refer to the exhibit. User traffic originating within site 0 is failing to reach an application hosted on IP address 192.168 0 10. Which is located within site A What is determined by the routing table?

A. The default gateway for site B is configured incorrectly

B. The lack of a default route prevents delivery of the traffic

C. The traffic is blocked by an implicit deny in an ACL on router2

D. The traffic to 192 168 010 requires a static route to be configured in router 1.

ANSWER: B

Explanation:

The lack of a default route prevents delivery of the traffic. A router makes its forwarding decision by comparing a packet's destination address with entries in its routing table and selecting the longest-prefix match. To reach the application at 192.168.0.10, the router must have either a route that specifically matches the destination network or a usable gateway of last resort, represented by a default route, 0.0.0.0/0. When the routing table has no matching route for that destination and no default route, the router has no valid next hop or outgoing interface on which to forward the packet. It therefore discards the traffic rather than forwarding it toward site A. This conclusion is based directly on the forwarding information available in the routing table: a default route provides reachability for destinations not covered by more-specific entries. Cisco documents that route selection uses longest-prefix matching and that a default route is used when no more-specific route matches. See Cisco's IP Routing: Route Concepts and Understanding Routing Table Concepts.

QUESTION NO: 122

What is a characteristic of encryption in wireless networks?

A. provides increased protection against spyware

B. uses policies to prevent unauthorized users

C. converts electrical current to radio waves

D. protects the confidentiality of data if it is intercepted as it transits a network

ANSWER: D

Explanation:

Encryption in a wireless network protects the confidentiality of data transmitted over the radio medium. Wi-Fi signals can be received by devices within RF range, so wireless frames may be captured while they are in transit. Encryption transforms protected data into ciphertext that cannot be meaningfully read by a party that does not possess the required cryptographic keys. Thus, encryption protects the contents of traffic if it is intercepted; it does not make radio transmissions physically impossible to capture.

In modern Wi-Fi security, WPA2 and WPA3 use robust encryption mechanisms to protect data traffic after a client has successfully authenticated and established the required security keys. For example, CCMP uses AES to provide confidentiality for protected wireless frames, helping ensure that passive capture of those frames does not disclose their payload. This is a fundamental reason wireless networks should use WPA2 or WPA3 rather than open networks or obsolete security mechanisms. Cisco's overview of WPA/WPA2 describes how these standards secure WLAN communications, while the Wi-Fi Alliance explains the role of Wi-Fi security protections for user data. See Cisco WPA/WPA2 Overview and Wi-Fi Alliance: Wi-Fi Security.

QUESTION NO: 123

Which two circumstances can prevent two routers from establishing an OSPF neighbor adjacency? (Choose two.)

A. mismatched autonomous system numbers

B. an ACL blocking traffic from multicast address 224.0.0.10

C. mismatched process IDs

D. mismatched hello timers and dead timers

E. use of the same router ID on both devices

ANSWER: D E

Explanation:

OSPF routers must exchange and accept Hello packets before they can progress through the neighbor state machine and establish an adjacency. Hello and Dead intervals are parameters carried in OSPF Hello packets, and neighboring routers on the same network segment must use matching values. With mismatched hello timers and dead timers, each router detects the mismatch during Hello processing and does not form the required neighbor relationship.

Use of the same router ID on both devices can also prevent a valid, stable OSPF adjacency. The router ID is the unique 32-bit identifier OSPF uses to identify a router in its link-state database and in OSPF protocol exchanges. It must be unique within the OSPF routing domain. A duplicate router ID creates an identity conflict that can cause OSPF neighbor processing and adjacency maintenance to fail or reset, as the routers cannot be reliably distinguished as separate OSPF nodes.

Cisco lists matching Hello/Dead intervals among the required OSPF neighbor parameters and identifies duplicate router IDs as an OSPF neighbor troubleshooting condition. See Cisco: OSPF Neighbor Relationship and Cisco: OSPF Design Guide.

QUESTION NO: 124

Which two wireless security standards use counter mode cipher block chaining Message Authentication Code Protocol for encryption and data integrity? (Choose two.)

A. Wi-Fi 6

B. WPA3

C. WEP

D. WPA2

E. WPA

ANSWER: B D

Explanation:

WPA2 and **WPA3** use CCMP (Counter Mode with Cipher Block Chaining Message Authentication Code Protocol) to protect wireless data. CCMP is an IEEE 802.11i security mechanism based on the Advanced Encryption Standard (AES). It combines AES operating in Counter Mode to provide confidentiality for transmitted frames with CBC-MAC-based integrity protection and origin authentication. This combination protects the contents of wireless traffic and helps detect unauthorized modification or replay of frames.

WPA2 certification made AES-CCMP the required robust security mechanism for protected Wi-Fi networks. Consequently, WPA2 networks configured for AES use CCMP for their data confidentiality and integrity services. WPA3 retains CCMP-128 as its baseline data-protection cipher suite for its standard 128-bit security mode. Although WPA3 introduces stronger authentication, such as Simultaneous Authentication of Equals (SAE) for personal networks, its protected data traffic continues to rely on AES-CCMP in this mode. Therefore, both WPA2 and WPA3 are correctly associated with CCMP-based encryption and integrity protection. See the Wi-Fi Alliance's Wi-Fi security overview and Cisco's wireless security guidance.

QUESTION NO: 125

What are two differences between WPA2 and WPA3 wireless security? (Choose two.)

A. WPA2 uses 192-bit key encryption, and WPA3 requires 256-bit key encryption.

B. WPA3 uses AES for stronger protection than WPA2, which uses SAE.

C. WPA2 uses 128-bit key encryption, and WPA3 supports 128-bit and 192-bit key encryption.

D. WPA3 uses SAE for stronger protection than WPA2, which uses AES.

E. WPA3 uses AES for stronger protection than WPA2, which uses TKIP.

ANSWER: C D

Explanation:

WPA2 uses 128-bit key encryption, and WPA3 supports 128-bit and 192-bit key encryption is correct when describing the security levels available in the standards. WPA2-Enterprise commonly uses the 128-bit AES-CCMP security suite. WPA3 retains 128-bit operation for ordinary deployments and adds the WPA3-Enterprise 192-bit security suite for environments that require higher cryptographic assurance. This 192-bit mode is a complete suite with stronger algorithms and parameters, rather than merely a change to a single encryption-key length.

WPA3 uses SAE for stronger protection than WPA2, which uses AES is also correct as a high-level comparison of the principal authentication enhancement. WPA2-Personal normally uses a pre-shared key with the four-way handshake, while WPA3-Personal uses Simultaneous Authentication of Equals, a password-authenticated key exchange. SAE improves password-based Wi-Fi authentication by making captured authentication exchanges substantially less useful for offline password-guessing attacks and by providing forward secrecy properties. AES remains an important data-protection algorithm in WPA2 and WPA3, while SAE identifies WPA3-Personal's authentication improvement. The Wi-Fi Alliance describes WPA3's SAE-based password protection and optional 192-bit Enterprise mode at Wi-Fi CERTIFIED WPA3 and summarizes current Wi-Fi security technologies at Wi-Fi Security.

QUESTION NO: 126

Why is UDP more suitable than TCP for applications that require low latency, such as VoIP?

A. UDP uses sequencing data for packets to arrive in order, and TCP offers the capability to receive packets in random order.

B. TCP uses congestion control for efficient packet delivery, and UDP uses flow control mechanisms for the delivery of packets

C. UDP reliably guarantees delivery of all packets, and TCP drops packets under heavy load.

D. TCP sends an acknowledgment for every packet that is received, and UDP operates without acknowledgments.

ANSWER: D

Explanation:

TCP sends an acknowledgment for every packet that is received, and UDP operates without acknowledgments. This distinction helps explain why UDP is commonly used for real-time media such as Voice over IP. TCP establishes a connection and provides reliable, ordered byte-stream delivery. Its reliability mechanisms include sequence numbers, acknowledgments, retransmission of data that is not acknowledged, and flow and congestion control. Although these features are valuable for applications that require complete and correctly ordered data, they can add delay. In particular, retransmitting lost data and preserving in-order delivery can make already-delayed voice data arrive too late to be useful to the listener.

UDP provides a connectionless datagram service with a much smaller transport-layer header and no transport-layer acknowledgments or retransmissions. A VoIP endpoint can therefore continue sending audio datagrams at the required real-time cadence rather than waiting for recovery of a missing datagram. VoIP applications generally address occasional loss through techniques such as jitter buffers, codecs designed to tolerate loss, and packet-loss concealment. This prioritizes timely playback over perfect delivery, which is the appropriate tradeoff for interactive voice. The UDP specification is defined in RFC 768; TCP acknowledgment and reliable-delivery behavior are specified in RFC 9293.

QUESTION NO: 127

What are two reasons to implement DHCP in a network? (Choose two.)

A. manually control and configure IP addresses on network devices

B. control the length of time an IP address is used by a network device

C. reduce administration time in managing IP address ranges for clients

D. dynamic control over the best path to reach an IP address

E. access a website by name instead of by IP address

ANSWER: B C

Explanation:

DHCP is implemented to centrally automate host IP configuration and to manage the temporary assignment of addresses efficiently. The benefit described by “control the length of time an IP address is used by a network device” comes from DHCP leases. A DHCP server assigns an address for a defined lease period, and the client must renew that lease to continue using it. When a device leaves the network or no longer renews, its address can eventually return to the available pool. This is especially useful for guest, wireless, lab, and other networks in which endpoints frequently connect and disconnect.

“Reduce administration time in managing IP address ranges for clients” is also a core DHCP benefit. Administrators define address pools and configuration options centrally, then clients automatically receive consistent settings such as an IP address, subnet mask or prefix length, default gateway, DNS server addresses, and lease information. Centralized automated delivery avoids configuring those settings individually on every endpoint, reduces repetitive effort, and helps maintain consistent client configuration across a network. Cisco describes DHCP as a protocol that dynamically assigns IP

addresses and other network configuration parameters to hosts. The DHCP lease behavior is defined in RFC 2131; Cisco configuration guidance is available in the Cisco IOS XE DHCP Server Configuration Guide.

QUESTION NO: 128

The OSPF Hello protocol performs which of the following tasks? (Choose two.)

A. It provides dynamic neighbor discovery.

B. It detects unreachable neighbors in 90 second intervals.

C. It maintains neighbor relationships.

D. It negotiates correctness parameters between neighboring interfaces.

E. It uses timers to elect the router with the fastest links as the designated router.

F. It broadcasts hello packets throughout the internetwork to discover all routers that are running OSPF.

ANSWER: A C

Explanation:

OSPF Hello packets provide dynamic neighbor discovery and maintain neighbor relationships. An OSPF router sends Hello packets on each OSPF-enabled interface, allowing routers attached to the same link or network segment to identify one another and begin the neighbor-state process. The packet includes information such as the router ID, area ID, network mask where applicable, Hello and Dead intervals, authentication information, and a neighbor list. Seeing its own router ID in a neighbor's Hello packet confirms bidirectional communication, which is a prerequisite to forming an adjacency where required.

After discovery, periodic Hello exchanges maintain the neighbor relationship. Each router expects to receive Hellos before the configured Dead interval expires. Receipt of a Hello refreshes the neighbor's inactivity timer; failure to receive them causes the neighbor to be declared down, enabling OSPF to react to a lost neighbor and recalculate routes when necessary. The specific Hello and Dead timer values are determined by the OSPF network type and interface configuration, so they are not inherently fixed at one universal interval. Cisco documents Hello packets as the mechanism used to establish and maintain neighbor adjacencies, and RFC 2328 defines their role in discovering and monitoring neighboring routers. See Cisco's OSPF neighbor relationship documentation and RFC 2328.

QUESTION NO: 129

Which function does the range of private IPv4 addresses perform?

A. allows multiple companies to each use the same addresses without conflicts

B. provides a direct connection for hosts from outside of the enterprise network

C. ensures that NAT is not required to reach the Internet with private range addressing

D. enables secure communications to the Internet for all external hosts

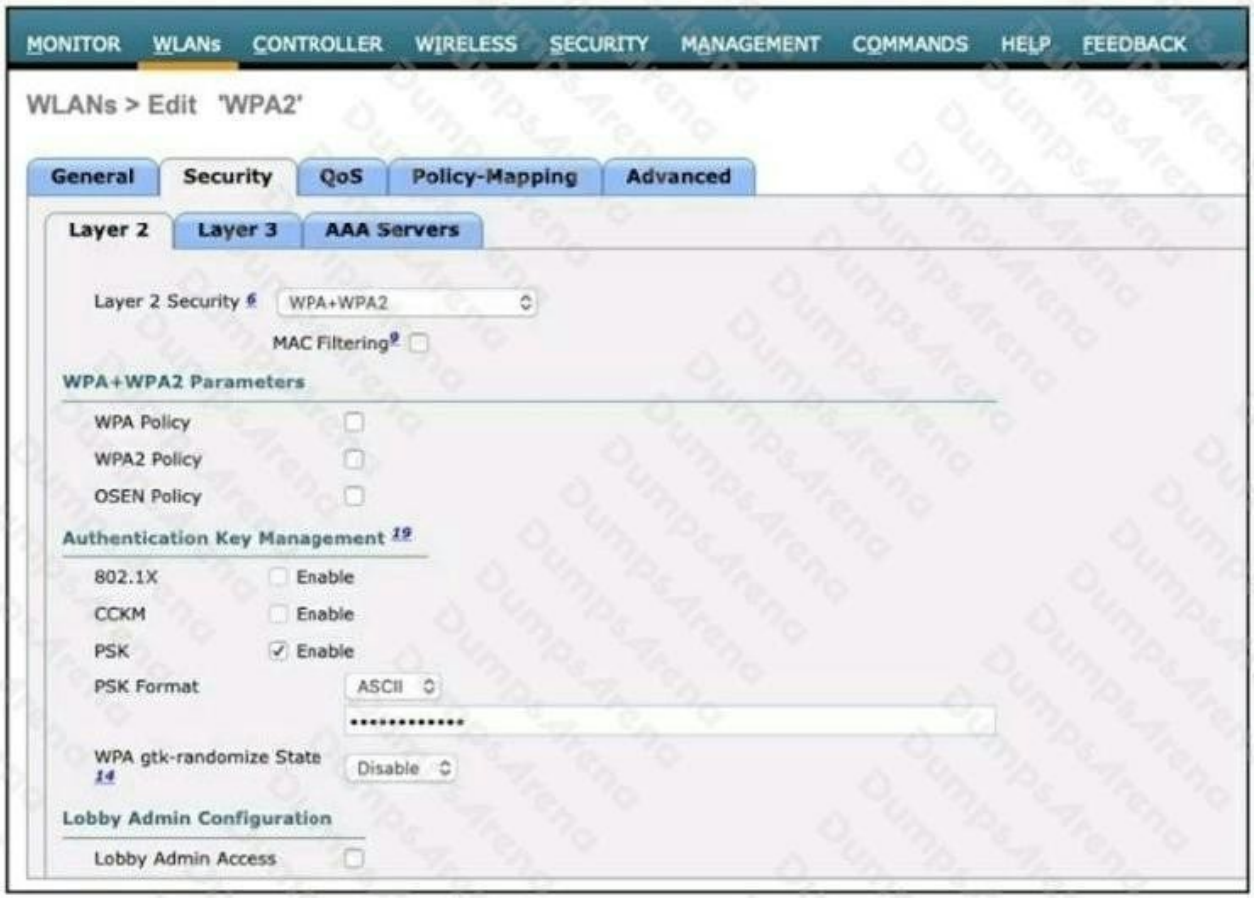
ANSWER: A

Explanation:

The correct answer is **allows multiple companies to each use the same addresses without conflicts**. Private IPv4 address space is reserved for use inside private networks under RFC 1918. The reserved blocks are 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. Unlike publicly assigned IPv4 addresses, these addresses do not need to be globally unique. An organization can therefore assign them to internal devices such as user workstations, servers, printers, and network infrastructure without obtaining an equivalent number of public IPv4 addresses.

Because RFC 1918 addresses are intended for private internets and are not globally routed across the public Internet, separate organizations can independently use identical private addressing schemes. For example, many organizations may use 192.168.1.0/24 internally without creating a routing conflict on the Internet. When privately addressed hosts need connectivity to public networks, an edge device commonly uses Network Address Translation to represent internal addresses with one or more publicly routable addresses. This address-reuse capability conserves the limited public IPv4 address space while supporting independent internal network designs. See RFC 1918, Address Allocation for Private Internets and Cisco's Network Address Translation overview.

QUESTION NO: 130



Refer to the exhibit. A network engineer is configuring a WLAN to use a WPA2 PSK and allow only specific clients to join. Which two actions must be taken to complete the process? (Choose two.)

A. Enable the OSEN Policy option.

B. Enable the 802.1X option for Authentication Key Management.

C. Enable the WPA2 Policy option.

D. Enable the MAC Filtering option.

E. Enable the CCKM option for Authentication Key Management.

ANSWER: C D

Explanation:

Enabling the **WPA2 Policy** is required because it activates WPA2 security for the WLAN. WPA2-Personal uses a pre-shared key, so after WPA2 is enabled the WLAN can be configured with PSK-based authentication and the required passphrase. This provides the WPA2 encryption and authentication framework requested for client access.

Enabling the **MAC Filtering** option is required to limit association to a defined set of client devices. MAC filtering applies an allow or deny decision based on each wireless client's MAC address. With an allow list populated for the WLAN, only clients

whose addresses are explicitly permitted can proceed to join it. This restriction is separate from the shared WPA2 passphrase: knowing the PSK alone does not satisfy the requirement to permit only specified devices when MAC filtering is enabled and configured with approved client addresses.

Therefore, WPA2 policy supplies the requested WPA2-PSK WLAN security, while MAC filtering supplies the requested client-specific admission control. Cisco documents WPA/WPA2 WLAN security behavior at Cisco WPA/WPA2 Security Overview and describes MAC-based client access control at Cisco MAC Authentication Overview.

QUESTION NO: 131

Which QoS feature drops traffic that exceeds the committed access rate?

A. policing

B. FIFO

C. shaping

D. weighted fair queuing

ANSWER: A

Explanation:

policing is the QoS feature that enforces a configured traffic rate and takes action on packets that exceed that rate. A policer classifies and measures traffic against a defined rate, commonly the committed information rate (CIR), using token-bucket parameters such as committed burst size. Packets that conform to the configured rate can be transmitted normally. Traffic that exceeds the allowed rate is commonly dropped immediately; depending on the policy configuration, it may instead be remarked to a lower priority value. Because this enforcement occurs without buffering excess packets for later transmission, policing is particularly appropriate where a provider or network edge must strictly limit traffic to an agreed access or service rate.

Cisco describes policing as a mechanism that limits traffic by dropping or remarking packets that exceed the configured rate, whereas rate enforcement that delays packets is handled through shaping. This behavior makes policing the direct match for a requirement to drop traffic above a committed access rate. See Cisco's Policing and Shaping Overview and the Cisco IOS XE QoS Policing and Shaping Configuration Guide.

QUESTION NO: 132

Which protocol should be used to transfer large files on a company intranet that allows TCP 20 and 21 through the firewall?

A. SMTP

B. REST API

C. TFTP

D. FTP

ANSWER: D

Explanation:

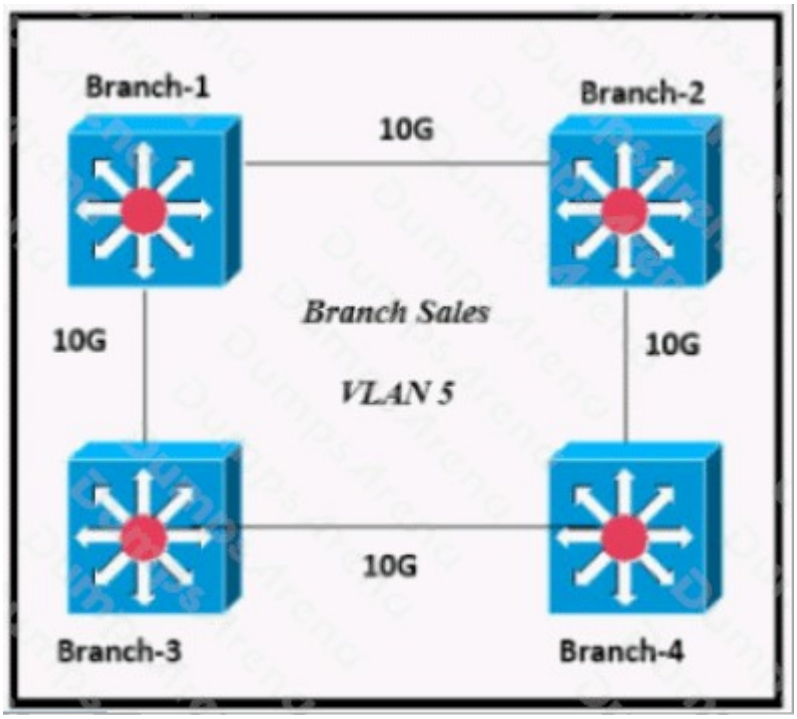
FTP is the appropriate protocol because it is specifically designed for file transfers and uses TCP, which provides reliable, connection-oriented delivery. Classic FTP uses TCP port 21 for its control connection, where the client authenticates and issues commands, and active-mode FTP uses TCP port 20 as the server-side source port for the data connection. Therefore, a firewall policy permitting TCP ports 20 and 21 directly accommodates traditional active FTP operation for transferring files on the intranet.

TCP's reliability is particularly relevant for large transfers: it provides sequencing, acknowledgments, retransmission of lost segments, and flow control so that the transferred file can be delivered accurately. FTP also supports standard file-transfer

operations such as uploading, downloading, directory navigation, and transfer-type selection. RFC 959 defines FTP's separate control and data connections, including the use of port 21 for the server control connection and port 20 for active-mode data transfer. See RFC 959: File Transfer Protocol and the IANA Service Name and Port Number Registry. In environments requiring encryption, FTPS or SFTP would generally be preferable, but the explicitly permitted TCP ports identify classic FTP here.

QUESTION NO: 133

Refer to the exhibit.



Only four switches are participating in the VLAN spanning-tree process.

Branch-1 priority 614440

Branch-2: priority 39082416

Branch-3: priority 0

Branch-4: root primary

Which switch becomes the permanent root bridge for VLAN 5?

A. Branch-1

B. Branch-2

C. Branch-3

D. Branch-4

ANSWER: C

Explanation:

Branch-3 becomes the permanent root bridge for VLAN 5 because it has a spanning-tree bridge priority of 0, the lowest configurable priority value. STP elects the root bridge by comparing Bridge IDs (BIDs). The BID includes the bridge priority, the extended system ID associated with the VLAN in PVST+/Rapid PVST+, and the switch MAC address. Priority is evaluated before the MAC address, so a switch configured with priority 0 has the decisive lowest priority component among

the participating switches. Its MAC address would be relevant only if another participating switch had the same effective BID priority.

The `root primary` macro is intended to make a switch the preferred root by assigning a low priority automatically. It cannot establish a lower STP priority than 0, however. Therefore, Branch-3's explicitly configured priority remains the lowest possible value and ensures its root-bridge role for VLAN 5 while it remains active in that VLAN's spanning-tree domain. Cisco documents that the root bridge is selected from the switch with the lowest Bridge ID and that bridge priority values are configurable in increments of 4096. See Cisco's Spanning Tree Protocol root bridge documentation and the Cisco IOS XE STP configuration guide.

QUESTION NO: 134

What is the main purpose of SSH management access?

A. To support DES 56-bit and 3DES (168-bit) ciphers

B. To enable secured access to the inbound management interface

C. To validate management access with username and domain name only

D. To allow passwords protected with HTTPS encryption to be sent

ANSWER: B

Explanation:

To enable secured access to the inbound management interface is correct because Secure Shell (SSH) provides an encrypted, authenticated remote command-line management session to a Cisco device. Administrators use SSH to connect to the device's virtual terminal (VTY) lines and perform administrative tasks without exposing credentials, commands, or device output in plaintext across the network. SSH is the secure replacement for Telnet, whose traffic is unencrypted. On Cisco IOS, enabling SSH normally involves configuring a hostname and domain name, generating RSA keys, creating or integrating an authentication method, and permitting SSH on the VTY lines. SSH's purpose is therefore secure remote management access; the domain name is only one prerequisite used when generating cryptographic keys, rather than an access-validation mechanism by itself. Cisco recommends using SSH version 2, which provides stronger security than SSH version 1. Restricting remote administrative access to SSH, potentially together with access-control lists and centralized AAA, helps protect the device management plane from credential disclosure and session interception. Cisco's SSH configuration guidance is available in the Cisco IOS XE Secure Shell Configuration Guide. The CCNA exam topics also identify configuration and verification of device access using SSH as a core administration skill: Cisco CCNA exam topics.

QUESTION NO: 135

What is a similarity between global and unique local IPv6 addresses?

A. They use the same process for subnetting.

B. They are part of the multicast IPv6 group type.

C. They are routable on the global internet.

D. They are allocated by the same organization.

ANSWER: A

Explanation:

They use the same process for subnetting. Both global unicast addresses and unique local addresses are unicast address types that an organization can divide hierarchically into prefixes for sites, routing domains, and individual network links. In normal IPv6 LAN design, each link receives a /64 subnet. The first 64 bits form the network prefix, including the allocated or locally selected prefix and subnet ID, while the remaining 64 bits identify an interface. This consistent /64 structure supports standard IPv6 host addressing mechanisms, including Stateless Address Autoconfiguration (SLAAC), and permits the same

practical planning workflow: assign a larger organizational prefix, allocate subnet-ID values to networks, and use one /64 for each broadcast-domain equivalent or routed link.

The distinction is the source and scope of the top-level prefix, not the subnetting method. Global unicast space is intended for globally routable addressing, whereas unique local addressing is intended for private organizational use. RFC 4193 describes the unique-local prefix format, including its 40-bit global ID and 16-bit subnet ID, which together produce a /64 subnet prefix. Cisco's IPv6 addressing overview also describes the conventional use of /64 prefixes on IPv6 LANs. See RFC 4193 and Cisco IPv6 Addressing.

QUESTION NO: 136

A network engineer must configure two new subnets using the address block 10.70.128.0/19 to meet these requirements:

• The first subnet must support 24 hosts

• The second subnet must support 472 hosts

• Both subnets must use the longest subnet mask possible from the address block

Which two configurations must be used to configure the new subnets and meet a requirement to use the first available address in each subnet for the router interfaces? (Choose two)

A. interface vlan 1234ip address 10.70.159.1 255.255.254.0

B. interface vlan 1148ip address 10.70.148.1 255.255.254.0

C. interface vlan 4722ip address 10.70.133.17 255.255.255.192

D. interface vlan 3002ip address 10.70.147.17 255.255.255.224

E. interface vlan 155ip address 10.70.155.65 255.255.255.224

ANSWER: B E

Explanation:

The configuration **interface vlan 1148ip address 10.70.148.1 255.255.254.0** satisfies the subnet requiring 472 hosts. A /23 mask (255.255.254.0) provides 512 total IPv4 addresses and 510 usable host addresses, making it the longest possible prefix that accommodates 472 hosts. Within the supplied 10.70.128.0/19 block, 10.70.148.0 is a valid /23 network boundary, and 10.70.148.1 is its first usable host address, as required for the router interface.

The configuration **interface vlan 155ip address 10.70.155.65 255.255.255.224** satisfies the subnet requiring 24 hosts. A /27 mask (255.255.255.224) creates 32 addresses, with 30 usable host addresses. It is therefore the longest mask that still supports at least 24 hosts. The relevant /27 subnet begins at 10.70.155.64, so 10.70.155.65 is the first usable address and is appropriate for the router interface. Both selected networks fall within the original /19 allocation, whose address range is 10.70.128.0 through 10.70.159.255, and the two allocations do not overlap. Cisco's IPv4 addressing documentation discusses prefix lengths and subnet-mask representation: Cisco: IP Addressing and Subnetting. Cisco's CCNA exam topics also include IPv4 addressing and subnetting: Cisco CCNA Exam Topics.

QUESTION NO: 137

Which field within the access-request packet is encrypted by RADIUS?

A. authorized services

B. password

C. authenticator

D. username

ANSWER: B

Explanation:

password is correct because, in a RADIUS Access-Request using password authentication, the user password is carried in the User-Password attribute and is specifically protected before transmission. RADIUS combines the shared secret with the Access-Request packet's Request Authenticator to transform the password value. The process uses an MD5-based masking algorithm defined by the RADIUS protocol: the password is padded as required, then XORed with MD5-derived blocks based on the shared secret and authenticator. A RADIUS server that possesses the same shared secret can reverse this operation and recover the supplied password for authentication.

This protection applies to the password attribute rather than to the complete RADIUS packet. Accordingly, network administrators should still protect RADIUS traffic with an appropriate trusted network design or with RADIUS over TLS where available, since the protocol's original User-Password protection is not equivalent to modern full-packet encryption. RFC 2865 defines the User-Password attribute encoding and states that it is hidden using the shared secret and Request Authenticator. Cisco's RADIUS documentation likewise describes RADIUS as protecting the user password rather than encrypting all Access-Request information. See RFC 2865, User-Password attribute and Cisco RADIUS overview.

QUESTION NO: 138

Why would a network administrator choose to implement automation in a network environment?

A. To simplify the process of maintaining a consistent configuration state across all devices

B. To centralize device information storage

C. To implement centralized user account management

D. To deploy the management plane separately from the rest of the network

ANSWER: A

Explanation:

To simplify the process of maintaining a consistent configuration state across all devices is correct because network automation enables administrators to define and apply standardized configurations repeatedly across routers, switches, and other managed infrastructure. Rather than relying on individual manual CLI changes, automation tools can use templates, structured data, APIs, and scripts to deploy the intended configuration uniformly. This reduces configuration drift, minimizes the risk of human entry errors, and makes routine provisioning and large-scale changes substantially faster. Automation can also validate device state against the intended standard, allowing administrators to identify and remediate deviations efficiently. These capabilities improve operational reliability and scalability, particularly when the network contains many devices or requires frequent changes. Cisco identifies improved consistency, faster service delivery, and reduced manual effort as key benefits of network automation. The CCNA curriculum likewise expects candidates to understand how automation and programmability support more efficient network operations and configuration management. See Cisco's Network Automation overview and the CCNA exam topics.

QUESTION NO: 139

Which IPsec encryption mode is appropriate when the destination of a packet differs from the security termination point?

A. transport

B. main

C. aggressive

D. tunnel

ANSWER: D

Explanation:

tunnel is the appropriate IPsec mode when the packet's ultimate destination is different from the device at which IPsec processing terminates. Tunnel mode protects and encapsulates the complete original IP packet, including its original IP header and payload. It then adds a new outer IP header whose source and destination identify the IPsec peers, such as two VPN gateways. After the protected packet reaches the receiving security gateway, that gateway removes the IPsec encapsulation and forwards the original packet toward its actual destination. This separation between the IPsec tunnel endpoint and the final destination is the normal design for site-to-site VPNs, where hosts on one private network communicate securely with hosts on another private network through edge routers or firewalls. Cisco describes tunnel mode as encapsulating the original packet and using a new IP header for delivery between security gateways. The IPsec architecture also defines tunnel mode for cases in which an SA protects an entire inner IP packet rather than only upper-layer traffic. See Cisco: IPsec Modes (Tunnel vs. Transport) and RFC 4301: Security Architecture for the Internet Protocol.

QUESTION NO: 140

The SW1 interface g0/1 is in the down/down state. Which two configurations are valid reasons for the interface conditions?(choose two)

A. There is a duplex mismatch

B. There is a speed mismatch

C. There is a protocol mismatch

D. The interface is shut down

E. The interface is error-disabled

ANSWER: D E

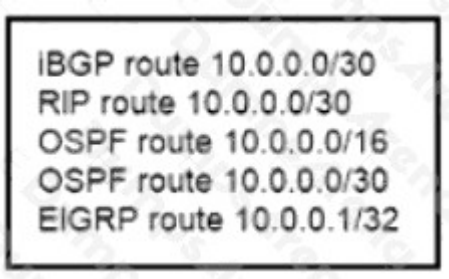
Explanation:

The interface is shut down and The interface is error-disabled are valid configuration-related reasons that a switch port does not establish an operational Ethernet link. Applying the `shutdown` command intentionally disables the interface. In detailed interface output, Cisco IOS identifies this specifically as *administratively down, line protocol is down*; it is therefore a deliberate administrative form of a down/down condition. Restoring operation requires entering interface configuration mode and issuing `no shutdown`, assuming a viable physical connection is present.

An error-disabled interface is also taken out of service by the switch. Cisco switches can place a port into err-disabled status after detecting a feature violation or protection event, such as a port-security violation, BPDU Guard event, UDLD condition, or EtherChannel inconsistency. Detailed output can describe the line protocol as down with an err-disabled reason, and interface-status output commonly displays `err-disabled`. The port remains disabled until the underlying condition is corrected and the port is recovered manually or through configured errdisable recovery. Cisco's interface troubleshooting guidance and errdisable documentation describe these operational states and recovery behavior: Cisco Ethernet interface troubleshooting and Cisco errdisable troubleshooting.

QUESTION NO: 141

Refer to the exhibit.



A router received these five routes from different routing information sources.

Which two routes does the router install in its routing table? (Choose two)

A. RIP route 10.0.0.0/30

B. iBGP route 10.0.0.0/30

C. OSPF route 10.0.0.0/30

D. EIGRP route 10.0.0.1/32

E. OSPF route 10.0.0.0/16

ANSWER: C D E

Explanation:

The router installs the OSPF route 10.0.0.0/30 because it competes with the RIP and iBGP routes for the identical 10.0.0.0/30 prefix. Cisco routers use administrative distance to choose between routes to the same prefix learned from different sources. OSPF has the default administrative distance of 110, which is lower than RIP's default of 120 and iBGP's default of 200. Consequently, the OSPF path is the selected route for that /30 destination.

The EIGRP route 10.0.0.1/32 is installed as a distinct host route. It has a different prefix from 10.0.0.0/30 and is therefore not an alternative path for the same routing-table entry. The OSPF route 10.0.0.0/16 is likewise a distinct, less-specific prefix and is also installed when no better route to that exact /16 prefix exists. The routing table can contain overlapping prefixes simultaneously; longest-prefix matching is applied later, when forwarding a packet. Thus, traffic specifically destined for 10.0.0.1 matches the /32 entry, traffic for other addresses in the /30 matches the /30 entry, and applicable destinations elsewhere within 10.0.0.0/16 can match the /16 entry. Cisco documents administrative distance behavior and route selection at Cisco's administrative distance reference and longest-prefix forwarding at Cisco's routing overview.

QUESTION NO: 142

Which plane is centralized in software-defined networking?

A. application

B. services

C. data

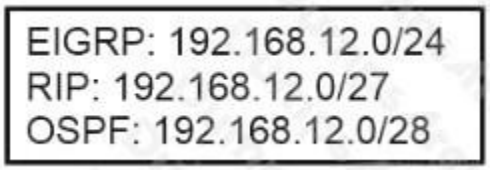
D. control

ANSWER: D

Explanation:

The control plane is centralized—more precisely, logically centralized—in software-defined networking (SDN). SDN separates the decision-making function that determines how traffic should be handled from the data plane, which performs packet forwarding on switches and routers. A software-based SDN controller provides the control-plane logic and maintains a network-wide view, enabling it to calculate policy and forwarding behavior consistently across the infrastructure. The controller communicates desired forwarding rules to network devices through southbound interfaces, while applications and orchestration systems interact with the controller through northbound interfaces. "Logically centralized" does not require one physical controller: resilient SDN deployments can use a distributed controller cluster while presenting a unified control-plane function and policy model. Cisco describes SDN as an architecture that separates the control plane from the forwarding plane, allowing centralized control and programmability. This separation lets administrators automate network changes and apply intent and policy consistently without configuring each forwarding device independently. The forwarding hardware remains responsible for rapidly processing and transmitting packets according to the rules it receives. See Cisco's SDN overview and the Open Networking Foundation's SDN architecture and standards resources.

QUESTION NO: 143



Refer to the exhibit. How does the router manage traffic to 192.168.12.16?

A. It chooses the EIGRP route because it has the lowest administrative distance.

B. It load-balances traffic between all three routes.

C. It chooses the OSPF route because it has the longest prefix inclusive of the destination address.

D. It selects the RIP route because it has the longest prefix inclusive of the destination address.

ANSWER: D

Explanation:

It selects the RIP route because it has the longest prefix inclusive of the destination address. Cisco routers make forwarding decisions by looking up the destination address in the routing table and choosing the matching route with the greatest prefix length. This is the longest-prefix-match rule: a route describing a smaller, more specific address range takes precedence over routes that describe broader ranges, provided the destination falls within that range. For destination 192.168.12.16, the RIP entry shown in the exhibit is the most specific matching network, so it is the route used to forward packets.

Administrative distance does not override this forwarding behavior across routes with different prefix lengths. Administrative distance is used when the router must choose which route source to install for the same destination prefix; it is not used to prefer a less-specific network over a more-specific matching network. Similarly, equal-cost multipath forwarding applies to eligible routes for the same prefix rather than to overlapping routes of different prefix lengths. The router therefore forwards traffic for this particular host address using the RIP next hop associated with the most-specific matching prefix. Cisco documents the route-lookup process and longest-match behavior in its routing table and route selection guidance and its IP routing overview.

QUESTION NO: 144

Which three statements about MAC addresses are correct? (Choose three.)

A. To communicate with other devices on a network, a network device must have a unique MAC address

B. The MAC address is also referred to as the IP address

C. The MAC address of a device must be configured in the Cisco IOS CLI by a user with administrative privileges

D. A MAC address contains two main components, the first of which identifies the manufacturer of the hardware and the second of which uniquely identifies the hardware

E. An example of a MAC address is 0A:26:B8:D6:65:90

F. A MAC address contains two main components, the first of which identifies the network on which the host resides and the second of which uniquely identifies the host on the network

ANSWER: A D E

Explanation:

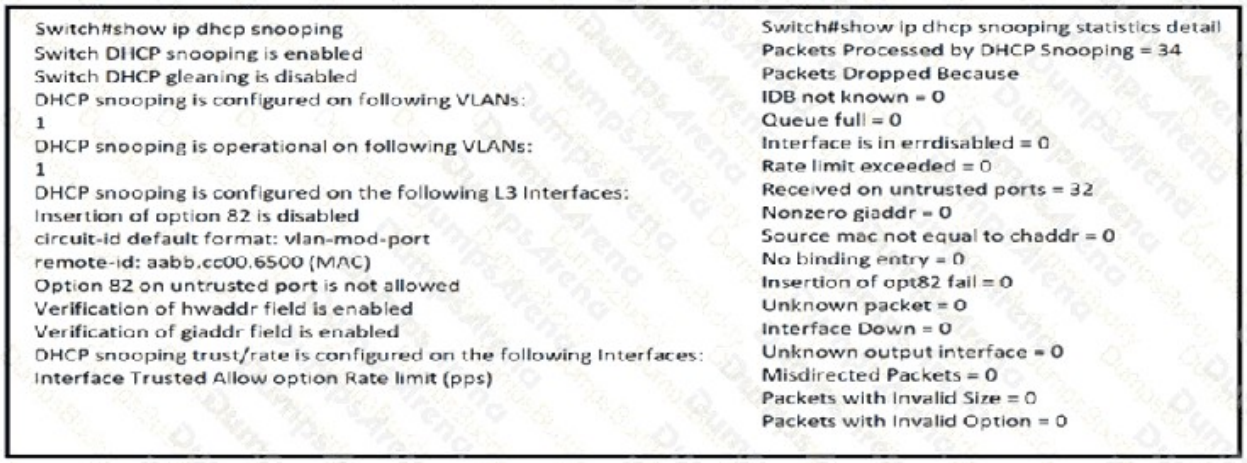
“To communicate with other devices on a network, a network device must have a unique MAC address” is correct in the normal Ethernet operating model. A network interface uses its MAC address as a Layer 2 source identifier and as a destination identifier for Ethernet frame delivery on the local network. MAC addresses must be unique within the relevant Layer 2 broadcast domain so switches can build reliable forwarding-table entries and frames can be delivered to the intended interface.

“A MAC address contains two main components, the first of which identifies the manufacturer of the hardware and the second of which uniquely identifies the hardware” correctly describes the traditional 48-bit universally administered MAC-address structure. The first 24 bits are the Organizationally Unique Identifier, assigned to an organization by IEEE, and the remaining 24 bits are assigned by that organization to identify an interface.

“An example of a MAC address is 0A:26:B8:D6:65:90” is also correct. It has six octets, each represented by two hexadecimal digits, which is a standard colon-separated notation for a 48-bit MAC address. See Cisco’s Ethernet and MAC addressing overview and IEEE’s Registration Authority tutorial.

QUESTION NO: 145

Refer to the exhibit.



The DHCP server and clients are connected to the same switch. What is the next step to complete the DHCP configuration to allow clients on VLAN 1 to receive addresses from the DHCP server?

A. Configure the ip dhcp snooping trust command on the interface that is connected to the DHCP client.

B. Configure the ip dhcp relay information option command on the interface that is connected to the DHCP client.

C. Configure the ip dhcp snooping trust command on the interface that is connected to the DHCP server.

D. Configure the Ip dhcp relay information option command on the interface that is connected to the DHCP server.

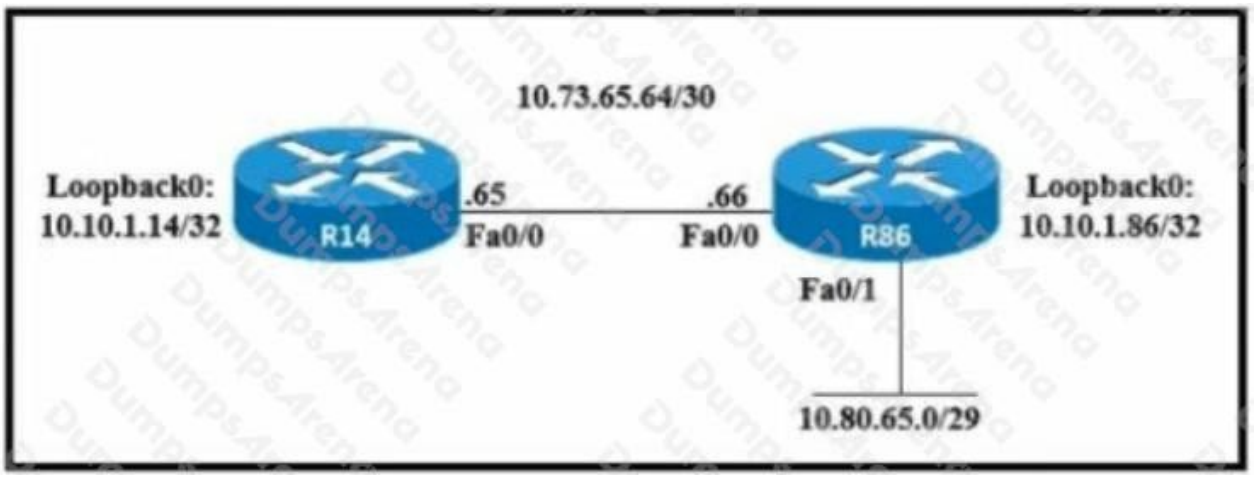
ANSWER: C

Explanation:

Configure the ip dhcp snooping trust command on the interface that is connected to the DHCP server. DHCP snooping treats switch interfaces as untrusted by default once DHCP snooping is enabled for a VLAN. An untrusted interface can forward DHCP client-originated messages, such as DHCPDISCOVER and DHCPREQUEST, but DHCP server messages, including DHCPOFFER and DHCPACK, must enter the switch through a trusted interface. Therefore, the port facing the legitimate DHCP server must be explicitly trusted so that the server’s replies can be forwarded to clients in VLAN 1 and leases can be completed.

This configuration also establishes the trusted boundary required for DHCP snooping’s protection model: DHCP responses received on the trusted server-facing port are accepted, while the switch can continue to restrict unauthorized DHCP-server behavior on client-facing access ports. The command is applied under the relevant Layer 2 interface, for example `interface gigabitEthernet...`, after DHCP snooping has been enabled globally and for VLAN 1. Cisco documents that trusted ports are the interfaces from which valid DHCP server messages are expected. See Cisco’s DHCP Snooping configuration guide and the Cisco DHCP snooping overview.

QUESTION NO: 146



Refer to the exhibit. An engineer must configure a floating static route on an external EIGRP network. The destination subnet is the /29 on the LAN interface of R86. Which command must be executed on R14?

A. ip route 10.80.65.0 255.255.248.0 10.73.65.66 1

B. ip route 10.80.65.0 255.255.255.240 fa0/1 89

C. ip route 10.80.65.0 255.255.255.248 10.73.65.66 171

D. ip route 10.73.65.66 0.0.0.224 10.80.65.0 255

ANSWER: C

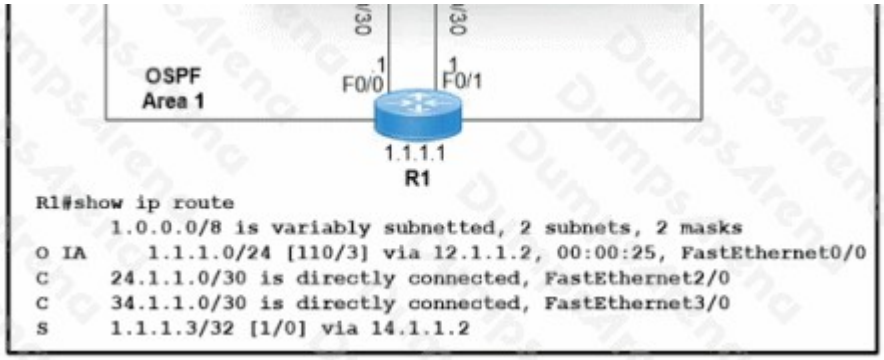
Explanation:

`ip route 10.80.65.0 255.255.255.248 10.73.65.66 171` correctly creates the required floating static route. The specified destination is the /29 LAN subnet behind R86, and a /29 prefix is represented by the mask `255.255.255.248`. The route therefore identifies the entire destination network as `10.80.65.0/29`, rather than a host address or a differently sized subnet. It uses `10.73.65.66` as the next-hop address toward that remote LAN, consistent with the path shown in the exhibit.

For a route to float, its administrative distance must be greater than that of the normally preferred learned route. Cisco assigns an administrative distance of 170 to routes learned as EIGRP external routes. Assigning an administrative distance of 171 means the static route is retained as a backup candidate but is not selected while the external EIGRP route is available. If the EIGRP external route is withdrawn, the router can install this static route and continue forwarding traffic toward the /29 destination. Cisco documents the default administrative distances, including 170 for EIGRP external routes, at Cisco Administrative Distance. Static-route command syntax and administrative-distance configuration are described in Cisco IOS XE Static Routes.

QUESTION NO: 147

Refer to the exhibit.



Which two values does router R1 use to determine the best path to reach destinations in network 1,0.0.0/8? (Choose two.)

A. longest prefix match

B. highest administrative distance

C. highest metric

D. lowest metric

E. lowest cost to reach the next hop

ANSWER: A D

Explanation:

longest prefix match is the first route-selection value R1 applies when forwarding a packet. R1 compares the packet destination address with the prefixes in its routing table and selects the matching route with the greatest prefix length. A more-specific prefix represents a narrower range of addresses, so it takes precedence over a less-specific route that also includes the destination. This lookup behavior is fundamental to Cisco IP forwarding.

After routes of the applicable prefix length have been considered, **lowest metric** is used to select the preferred path among comparable routes supplied by the same routing source. A metric is the routing protocol's path-quality measurement. For example, OSPF uses cumulative interface cost, while EIGRP calculates a composite metric. In either case, the route with the lower calculated metric is preferred. This enables R1 to select the most favorable path when multiple valid candidate paths exist for the same destination prefix. Cisco documents that routing decisions use prefix specificity and then route preference attributes, including protocol metrics where applicable. See Cisco's route selection and administrative distance overview and Cisco's OSPF metric documentation.

QUESTION NO: 148

Which two functions are performed by the core layer in a three-tier architecture? (Choose two.)

A. Provide uninterrupted forwarding service

B. Inspect packets for malicious activity

C. Ensure timely data transfer between layers

D. Provide direct connectivity for end user devices

E. Police traffic that is sent to the edge of the network

ANSWER: A C

Explanation:

Provide uninterrupted forwarding service and ensure timely data transfer between layers are core-layer functions in Cisco's hierarchical campus design. The core is the high-speed, highly available campus backbone that interconnects distribution blocks. Its design priority is dependable, rapid packet transport, with redundant paths and devices supporting continued forwarding during a link or device failure. The core therefore must be engineered to minimize disruption and converge quickly so traffic can continue moving across the campus.

The core also provides the fast transport needed for predictable, timely communication among the network's distribution components. Cisco design guidance emphasizes maintaining a simple core and avoiding unnecessary processing that could introduce delay or reduce backbone resiliency. High-capacity links, resilient Layer 3 forwarding, and a topology designed for fast recovery allow the core to deliver traffic efficiently between distribution-layer blocks. These responsibilities make uninterrupted forwarding and timely interlayer data transfer fundamental outcomes of the core layer's role in a three-tier architecture. See Cisco's Campus Network Design Fundamentals and the Cisco Press overview of hierarchical network design.

QUESTION NO: 149

Refer to the exhibit Routers R1 R2 and R3 use a protocol to identify their neighbors' IP addresses hardware platforms, and software versions. A network engineer must configure R2 to avoid sharing any neighbor information with R3, and maintain its relationship with R1. What action meets this requirement?

A. Configure the no cdp enable command on gO/2.

B. Configure the no cdp run command globally.

C. Configure the no lldp run command globally.

D. Configure the no lldp receive command on gQV1.

ANSWER: A

Explanation:

Configure the no cdp enable command on gO/2. The neighbor-discovery details described—IP addressing, hardware platform, and Cisco software version—are Cisco Discovery Protocol information. CDP operates at Layer 2 and is enabled on individual interfaces by default when CDP is enabled globally. Applying `no cdp enable` under R2's interface connected to R3 disables CDP only on that link. As a result, R2 no longer sends CDP advertisements to R3 or processes CDP advertisements received through that interface, preventing CDP neighbor information from being exchanged with R3.

Because this is an interface-level configuration, CDP remains operational on R2's other interfaces, including the interface facing R1. This preserves R2's existing CDP neighbor relationship and discovery capability with R1 while meeting the requirement to isolate R3 specifically. Cisco documents that CDP can be disabled per interface with the `no cdp enable` interface command, allowing selective control of CDP operation without affecting all router links. See Cisco's CDP overview and configuration guidance and the Cisco IOS CDP command reference.

QUESTION NO: 150

SIP-based Call Admission Control must be configuredd in the Cisco WLC GUI. SIP call-snooping ports are configured. Which two actions must be completed next? (Choose two.)

A. Set the QoS level to silver or greater for voice traffic.

B. Set the QoS level to platinum for voice traffic.

C. Enable Media Session Snooping on re WLAN.

D. Enable traffic shaping for the LAN interface of the WL

E. configure two different QoS rotes tor data and voice traffic.

ANSWER: B C

Explanation:

For SIP-based Call Admission Control on a Cisco wireless LAN controller, voice traffic must be assigned the Platinum QoS level, and Media Session Snooping must be enabled on the WLAN. Platinum is the controller QoS profile intended for latency-sensitive voice applications. It applies the wireless voice treatment needed to prioritize real-time packets and supports the admission-control behavior expected for voice clients using that WLAN.

Media Session Snooping lets the controller inspect SIP signaling and associate the negotiated media streams with the call. With the SIP snooping ports already configured, enabling this WLAN feature gives the controller the SIP-session visibility required to identify calls and make admission decisions for those calls. Together, the Platinum WLAN QoS classification and Media Session Snooping provide both the appropriate wireless treatment for voice and the SIP/media awareness necessary for SIP CAC. Cisco wireless controller configuration documentation describes WLAN QoS profiles and voice-related WLAN settings, while Cisco's voice-over-WLAN guidance emphasizes prioritization and admission control for dependable real-time wireless voice service. See the Cisco Wireless Controller Configuration Guide and Cisco's Voice over Wireless LAN design guidance.

QUESTION NO: 151 - (SIMULATION)

SIMULATION

-

Guidelines

-

This is a lab item in which tasks will be performed on virtual devices

• Refer to the Tasks tab to view the tasks for this lab item.

• Refer to the Topology tab to access the device console(s) and perform the tasks.

• Console access is available for all required devices by clicking the device icon or using the tab(s) above the console window.

• All necessary pre-configurations have been applied.

• Do not remove any existing configurations from the devices, only those necessary to make the appropriate

changes required to fulfill the listed tasks.

• Do not change the enable password or hostname for any device.

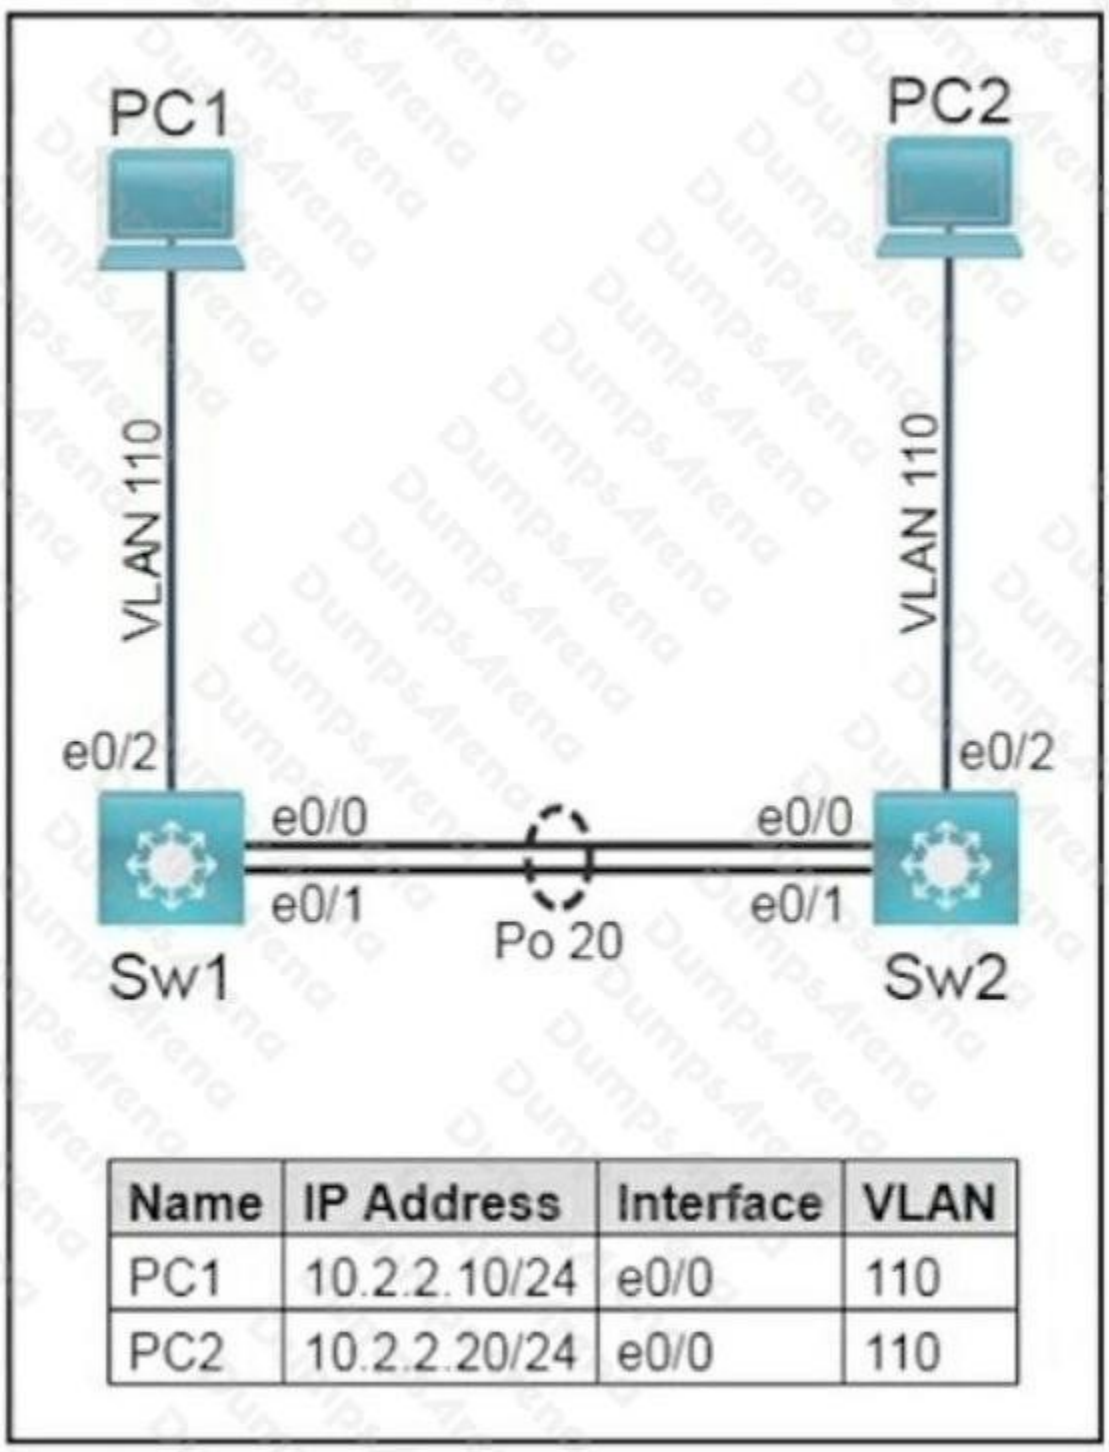
• Save your configurations to NVRAM before moving to the next item.

• Click Next at the bottom of the screen to submit this lab and move to the next question.

• When Next is clicked, the lab closes and cannot be reopened.

Topology

-



Name	IP Address	Interface	VLAN
PC1	10.2.2.10/24	e0/0	110
PC2	10.2.2.20/24	e0/0	110

Tasks

-

Task 1

-

Configure trunks between Sw1 and Sw2 on ports E0/0 and E0/1 using the IEEE standard frame tagging method.

• Add VLAN 99 as untagged on the trunk ports.

• Only extend VLAN 110 and the untagged VLAN across the trunk.

• Verify that PC1 is capable of pinging PC2.

Task 2

-

On Sw1 and Sw2, use IEEE 802.3ad link aggregation.

• Combine E0/0 and E0/1 into a single logical link while leaving the trunk configurations intact.

• Assign number 20 to the link.

• Both links must negotiate aggregation.

ANSWER: See the explanation for the answer

Explanation:

Configure the same LACP EtherChannel and trunk settings on **both Sw1 and Sw2**. IEEE 802.3ad is LACP, so use `channel-group 20 mode active` on both ends. VLAN 99 is the native (untagged) VLAN, and only VLANs 99 and 110 are allowed.

On **Sw1** and then on **Sw2**, enter:

```
enable
configure terminal
vlan 99
 name NATIVE
exit
interface range ethernet 0/0 - 1
 switchport trunk encapsulation dot1q
 switchport mode trunk
 switchport trunk native vlan 99
 switchport trunk allowed vlan 99,110
 channel-group 20 mode active
exit
interface port-channel 20
 switchport trunk encapsulation dot1q
 switchport mode trunk
 switchport trunk native vlan 99
 switchport trunk allowed vlan 99,110
exit
end
copy running-config startup-config
```

If the switch does not accept `switchport trunk encapsulation dot1q`, omit that command; on platforms that support only 802.1Q, it is already the enforced encapsulation.

Verify the EtherChannel and trunk state:

```
show etherchannel summary
show interfaces trunk
```

The expected EtherChannel state is Port-Channel 20 with both `e0/0` and `e0/1` bundled using LACP (typically shown as `(SU)` for Po20 and `(P)` for each member). PC1 at `10.2.2.10/24` should then successfully ping PC2 at `10.2.2.20/24`.

QUESTION NO: 152

Refer to the exhibit. What is the administrative distance for the advertised prefix that includes the host IP address 192.168.20.1?

A. 0

B. 192.168.10.2

C. 24

D. 1

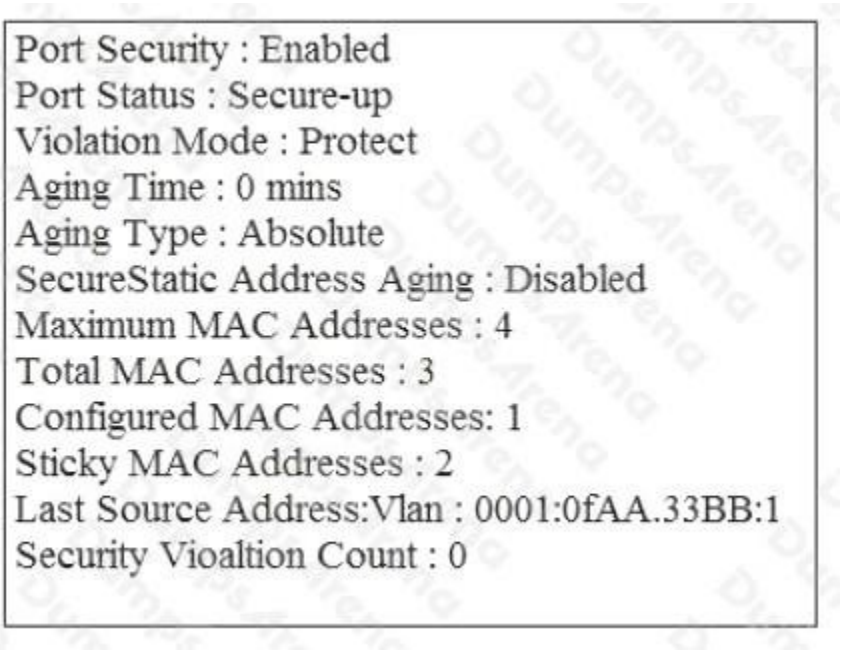
ANSWER: D

Explanation:

The administrative distance is **1**. The host address 192.168.20.1 belongs to the advertised network prefix 192.168.20.0/24. In Cisco IOS route output, a route is commonly displayed with brackets in the form `[administrative-distance/metric]`. Thus, when the route entry for that prefix shows a value such as `[1/0]`, the first number, 1, is the administrative distance, while the second number, 0, is the route metric. An administrative distance of 1 is Cisco IOS's default preference value for a static route. Administrative distance is a locally significant value used by a router to select among routes to the same destination that were learned from different sources; lower values are preferred. Cisco documents the default administrative distance for static routes as 1 and describes how route-source preference is used during route selection. See Cisco's Administrative Distance documentation and the Cisco IOS ip route command reference.

QUESTION NO: 153

Refer to the exhibit. Which two events occur on the interface, if packets from an unknown Source address arrive after the interface learns the maximum number of secure MAC address? (Choose two.)



A. The security violation counter dose not increment

B. The port LED turns off

C. The interface is error-disabled

D. A syslog message is generated

E. The interface drops traffic from unknown MAC address

ANSWER: D E

Explanation:

The configured port-security violation mode is *restrict*. Once the interface has learned its allowed maximum number of secure MAC addresses, a frame received with a previously unknown source MAC address causes a port-security violation. In restrict mode, the switch continues operating the interface but discards traffic received from the unauthorized source. Therefore, *The interface drops traffic from unknown MAC address* is an expected event.

Restrict mode also records the violation and sends a notification through the system logging mechanism. Thus, *A syslog message is generated* is also expected. This behavior is useful when administrators need to prevent unauthorized endpoints from communicating while retaining the port's operational state and receiving an audit trail of the violation. Cisco documents that the restrict violation action drops packets from unknown source addresses and generates an SNMP trap and syslog

message. The port-security feature and violation actions are described in Cisco's Port Security configuration and behavior documentation and the Catalyst switch port-security configuration guide.

QUESTION NO: 154

Which three describe the reasons large OSPF networks use a hierarchical design? (Choose three.)

A. to speed up convergence

B. to reduce routing overhead

C. to lower costs by replacing routers with distribution layer switches

D. to decrease latency by increasing bandwidth

E. to confine network instability to single areas of the network

F. to reduce the complexity of router configuration

ANSWER: A B E

Explanation:

Large OSPF deployments use a hierarchical, multi-area design to make the link-state routing domain scalable and stable. The goal **to reduce routing overhead** is achieved because routers retain detailed topology information primarily for their own area, while Area Border Routers advertise summarized interarea reachability. This reduces the scope of link-state advertisements, lowers LSDB size, and limits CPU and memory demand for SPF processing.

The goal **to speed up convergence** follows from limiting the topology affected by a change. When a failure occurs within an area, the associated LSA flooding and shortest-path recalculation are generally restricted to that area rather than forcing every router in the OSPF domain to process the event. Smaller topology databases and a reduced recalculation scope support faster, more predictable recovery.

Hierarchical areas also exist **to confine network instability to single areas of the network**. Frequent link changes or routing events in one area can remain local, preventing unnecessary control-plane churn throughout the backbone and other areas. This containment improves overall routing-domain stability. Cisco describes OSPF areas as a method to reduce SPF calculations and LSA flooding while isolating topology changes. See the Cisco OSPF Design Guide and the Cisco IOS XE OSPF Configuration Guide.

QUESTION NO: 155

Which virtual MAC address is used by VRRP group 1?

A. 0050.0c05.ad81

B. 0007.c061.bc01

C. 0000.5E00.0101

D. 0500.3976.6401

ANSWER: C

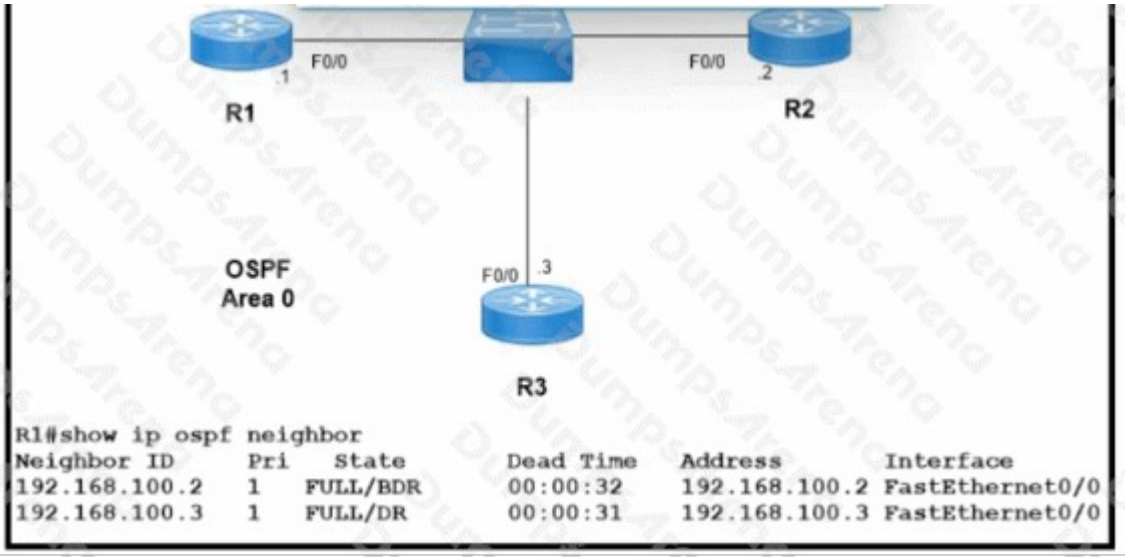
Explanation:

The virtual MAC address used by VRRP group 1 is **0000.5E00.0101**. VRRP identifies each virtual router with a Virtual Router Identifier (VRID), often called a VRRP group number in Cisco configuration and exam terminology. For IPv4 VRRP, the standards-defined virtual MAC address format is **00:00:5E:00:01:VRID**, where the final byte represents the VRID in hexadecimal. Group 1 has a VRID of decimal 1, which is hexadecimal **01**. Substituting that value produces **00:00:5E:00:01:01**, written in Cisco dotted-hex notation as **0000.5E00.0101**. This stable virtual MAC is associated with the virtual IPv4 default-gateway address. Whichever VRRP router is acting as the master uses it, allowing end hosts to keep

sending traffic to the same Layer 2 destination even if responsibility moves to another router. The Internet standard specifies this MAC-address construction, including the IANA OUI prefix **00:00:5E** and the IPv4 VRRP identifier bytes **00:01**. See RFC 5798, section 7.3 and Cisco's VRRP configuration guide.

QUESTION NO: 156

Refer to the exhibit.



Which two configurations must the engineer apply on this network so that R1 becomes the DR? (Choose two.)

A)

```
R1(config)#router ospf 1
R1(config-router)#router-id 192.168.100.1
```

B)

```
R1(config)#interface fastethernet 0/0
R1(config-if)#ip ospf priority 200
```

C)

```
R3(config)#interface fastethernet 0/0
R3(config-if)#ip ospf priority 0
```

D)

```
R1(config)#interface fastethernet 0/0
R1(config-if)#ip ospf priority 0
```

E)

```
R3(config)#interface fastethernet 0/0
R3(config-if)#ip ospf priority 200
```

A. Option A

B. Option B

C. Option C

D. Option D

E. Option E

ANSWER: B C

Explanation:

The selected configurations ensure that R1 participates in a fresh OSPF designated-router election with the election attributes needed to win on the shared multiaccess segment. OSPF elects a DR and backup DR per broadcast or nonbroadcast multiaccess network. The first election criterion is the OSPF interface priority: the eligible router with the highest priority is preferred. A priority of 0 makes an interface ineligible for either DR or BDR, while priorities from 1 through 255 are eligible. If eligible routers have equal interface priorities, OSPF uses the highest router ID as the tie-breaker.

DR and BDR elections are nonpreemptive. Consequently, changing an interface priority alone does not displace a router that is already acting as DR. The configuration that establishes R1's winning election value must therefore be accompanied by the configuration that causes the OSPF adjacency or process to restart and triggers a new election. Once the election occurs, R1 is selected as DR according to the updated priority-based election process. Cisco documents the DR/BDR election behavior at Cisco: OSPF DR and BDR Election; the interface-priority command is documented in the Cisco IOS OSPF command reference.

QUESTION NO: 157

Which configuration ensures that the switch is always the root for VLAN 750?

A. Switch(config)#spanning-tree vlan 750 priority 38003685

B. Switch(config)#spanning-tree vlan 750 root primary

C. Switch(config)#spanning-tree vlan 750 priority 614440

D. Switch(config)#spanning-tree vlan 750 priority 0

ANSWER: D

Explanation:

`Switch(config)#spanning-tree vlan 750 priority 0` assigns the lowest configurable STP bridge priority to VLAN 750. For a given VLAN, STP elects the root bridge using the lowest bridge ID, whose priority component is evaluated before the switch MAC-address component. Cisco IOS supports bridge-priority values from 0 through 61440 in increments of 4096; therefore, zero gives this switch the strongest possible configured priority for the VLAN and is the intended choice when the design requires it to be the root bridge. If another switch were deliberately configured with the same priority, the MAC-address portion of the bridge ID would act as the tie-breaker, so sound network design also requires controlling STP settings on the other switches. In normal administration, assigning priority 0 to the designated root and higher priorities to all other switches establishes deterministic root placement for VLAN 750. Cisco documents the VLAN-specific priority syntax and valid range in its LAN Switching Command Reference. The root-election process is also described in Cisco's Spanning Tree Protocol overview.

QUESTION NO: 158

Which two functions does a WLC perform in the lightweight access-point architecture that an AP performs independently in an autonomous architecture? (Choose two.)

A. managing RF channels, including transmission power

B. handling the association, authentication, and roaming of wireless clients

C. sending and processing beacon frames

D. encrypting and decrypting traffic that uses the WAP protocol family

E. preventing collisions between wireless clients on the same RF channel

ANSWER: A B

Explanation:

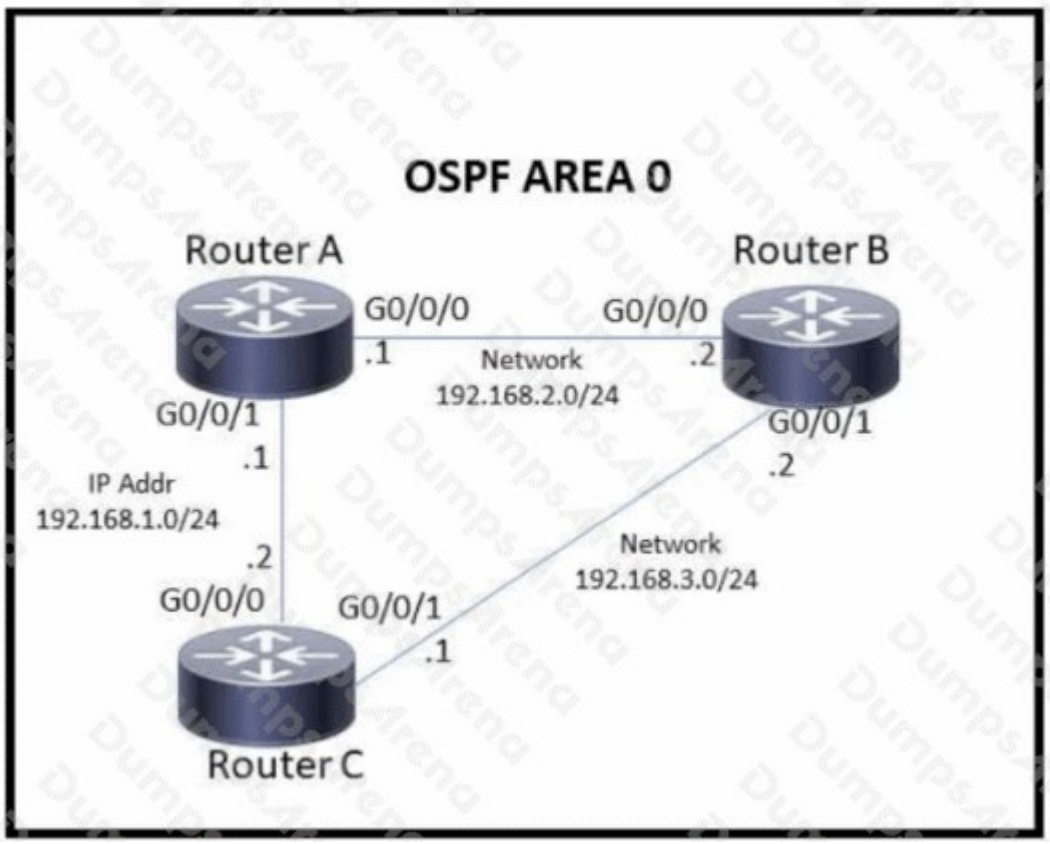
In a controller-based lightweight AP deployment, the wireless LAN controller centralizes important control-plane decisions that an autonomous AP would make for itself. **managing RF channels, including transmission power** is a controller function through Cisco Radio Resource Management (RRM). The controller collects radio measurements from managed APs and uses that information to coordinate channel assignment and transmit-power levels across the WLAN. Centralized RF coordination helps provide suitable coverage and capacity while reducing unnecessary radio-frequency contention between neighboring APs.

The WLC also performs centralized client management represented by **handling the association, authentication, and roaming of wireless clients**. It maintains WLAN configuration and client state, applies policy, and coordinates mobility as a client moves between controller-managed APs. This centralization gives users a consistent WLAN experience and enables the network to make roaming and access-control decisions using information shared across the AP infrastructure, rather than limiting those decisions to an individual AP.

Cisco describes the controller/AP split in its CAPWAP architecture documentation and details how RRM manages channel and power planning. See Cisco CAPWAP architecture documentation and Cisco Radio Resource Management overview.

QUESTION NO: 159

Refer to the exhibit.



Which action must be taken to ensure that router A is elected as the DR for OSPF area 0?

A. Configure the OSPF priority on router A with the lowest value between the three routers.

B. Configure router B and router C as OSPF neighbors of router A.

C. Configure router A's interface on the shared segment with the highest OSPF priority value.

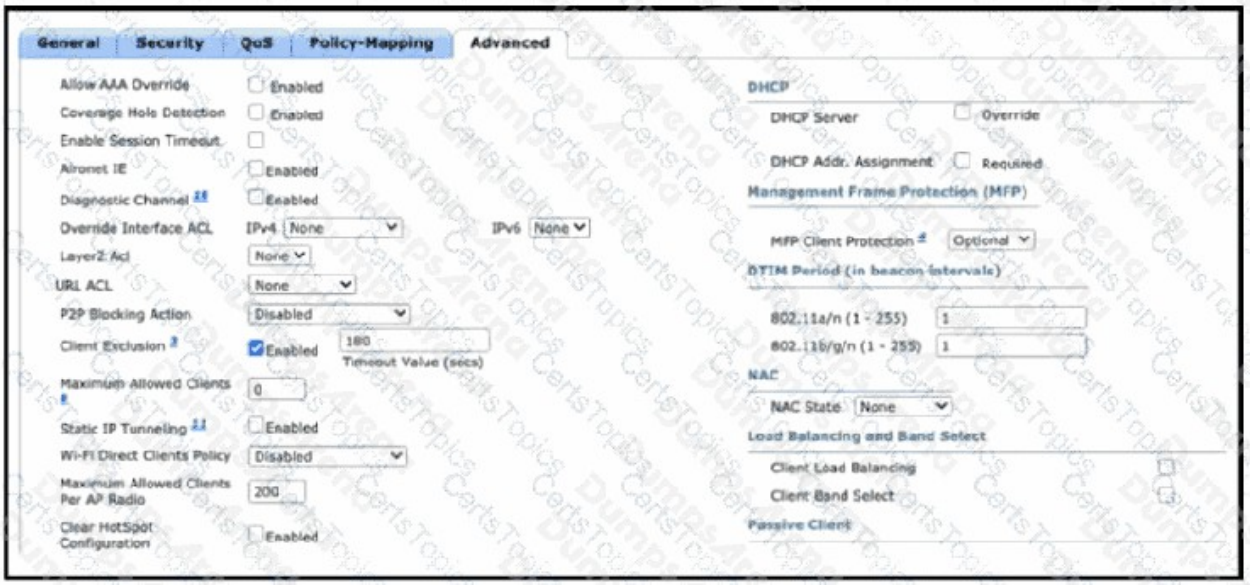
D. Configure router A with a fixed OSPF router ID

ANSWER: C

Explanation:

Configure router A's OSPF interface connected to the shared multiaccess segment with the highest OSPF priority. On broadcast and nonbroadcast multiaccess OSPF networks, routers elect a designated router to reduce the number of required adjacencies and to originate the network LSA for that segment. The election first compares the OSPF interface priority values: the eligible router with the highest priority becomes the DR. The priority is an interface-level setting, so it must be configured on router A's interface participating in the common area 0 segment, rather than as a general area-wide attribute. A value from 1 through 255 is eligible, and configuring a value higher than the priorities on the corresponding interfaces of the other routers makes router A win the election. For example, the interface command `ip ospf priority 255` gives router A the maximum eligible priority. Because OSPF DR election is nonpreemptive, if a DR has already been selected, the election must be triggered again, such as by resetting OSPF adjacencies or restarting the relevant interface/process, for the new priority to take effect in an election. Cisco documents the election criteria and priority behavior in its OSPF Designated Router Election guidance and the interface command details in the Cisco IOS OSPF command reference.

QUESTION NO: 160



Refer to the exhibit. An administrator is configuring a new WLAN for a wirelesss network that has these requirements:

Dual-band clients that connect to the WLAN must be directed to the 5-GHz spectrum.

Wireless clients on this WLAN must be able to apply VLAN settings from RADIUS attributes.

Which two actions meet these requirements? (Choose two.)

A. Enable the Aironet IE option.

B. Enable the Coverage Hole Detection option.

C. Set the MFP Client Protection option to Required

D. Enable the client band select option.

E. Enable the allow AAA Override option

ANSWER: D E

Explanation:

Enabling the client band select option meets the requirement to preferentially place dual-band-capable clients on the 5-GHz radio. Band Select is an RF-management feature that influences association behavior by discouraging capable clients from joining on 2.4 GHz, where fewer nonoverlapping channels and greater interference commonly reduce available capacity. The client ultimately makes the association decision, but this feature provides the controller/AP mechanism to steer eligible dual-band devices toward 5 GHz.

Enabling the allow AAA Override option permits the WLAN to accept authorized RADIUS-returned attributes that override applicable WLAN defaults on a per-client basis. This includes dynamic VLAN assignment, allowing the RADIUS server to return VLAN-related authorization attributes after successful client authentication. The WLAN can therefore place different authenticated users or device groups into the VLAN specified by their RADIUS policy rather than assigning every client to one static WLAN VLAN. Cisco documents AAA Override as the WLAN setting used to apply RADIUS authorization parameters, including interface or VLAN assignment, and documents Band Select as the feature used to encourage dual-band clients to use the 5-GHz band. See Cisco Wireless Controller Band Select documentation and Cisco WLAN security and AAA Override documentation.

QUESTION NO: 161

An engineer must configure R1 for a new user account. The account must meet these requirements:

* It must be configured in the local database.

* The username is engineer.

* It must use the strongest password configurable. Which command must the engineer configure on the router?

A. R1(config)# username engineer algorithm-type scrypt secret test2021

B. R1(config)# username engineer2 secret 5 .password S1$b1Ju$kZbBS1Pyh4QzwXyZ

C. R1(config)# username engineer2 privilege 1 password 7 test2021

D. R1(config)# username englneer2 secret 4 S1Sb1Ju$kZbBS1Pyh4QzwXyZ

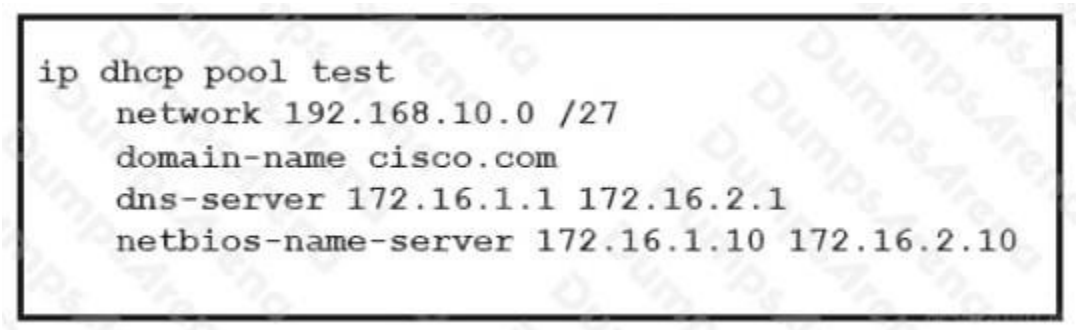
ANSWER: A

Explanation:

`R1(config)# username engineer algorithm-type scrypt secret test2021` creates a local username database entry named `engineer` and protects its secret with the scrypt password-hashing algorithm. The `username` global configuration command is the IOS mechanism for defining locally stored user accounts, which can subsequently be used by local authentication methods. Including `algorithm-type scrypt` explicitly selects Cisco IOS Type 9, scrypt-based secret protection. Scrypt is intentionally computationally and memory intensive, making offline password-guessing attacks substantially more costly than legacy password representations. The value following `secret` is entered as the account password; IOS derives and stores a protected hash rather than retaining the clear-text value in the configuration. This command also uses the required username exactly as specified, rather than creating an account with a different name. Cisco documents the enhanced password-security algorithms, including scrypt, in its Enhanced Password Security configuration guide. The general local-user configuration syntax is documented in the Cisco IOS XE user account configuration guide.

QUESTION NO: 162

Refer to the exhibit. After you apply the given configuration to a router, the DHCP clients behind the device cannot communicate with hosts outside of their subnet. Which action is most likely to correct the problem?



A. Configure the dns server on the same subnet as the clients

B. Activate the dhcp pool

C. Correct the subnet mask

D. Configure the default gateway

ANSWER: D

Explanation:

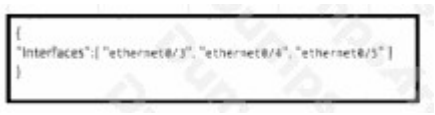
Configure the default gateway is correct because DHCP clients need a default route to send packets to destinations outside their local IP subnet. On a Cisco IOS DHCP server, this setting is supplied in the DHCP pool with the `default-router` command. Its value must be the IP address of the Layer 3 gateway interface serving the clients' subnet, such as the router interface or switched virtual interface associated with that VLAN.

A host uses its subnet mask to determine whether a destination is local. For an off-subnet destination, it sends the packet to its configured default gateway, resolving that gateway's MAC address with ARP and forwarding the frame to it. Without a DHCP-provided default gateway, a client may still communicate with devices in its own subnet, but it has no Layer 3 next hop for remote networks. Adding the appropriate `default-router` value to the pool and renewing client DHCP leases supplies the required gateway information.

Cisco documents `default-router` as the DHCP pool command used to configure the default gateway option delivered to clients. See the Cisco IOS DHCP Server Configuration Guide and Cisco's DHCP configuration example.

QUESTION NO: 163

Refer to the exhibit.



Which types of JSDN data is shown

A. Object

B. Sequence

C. String

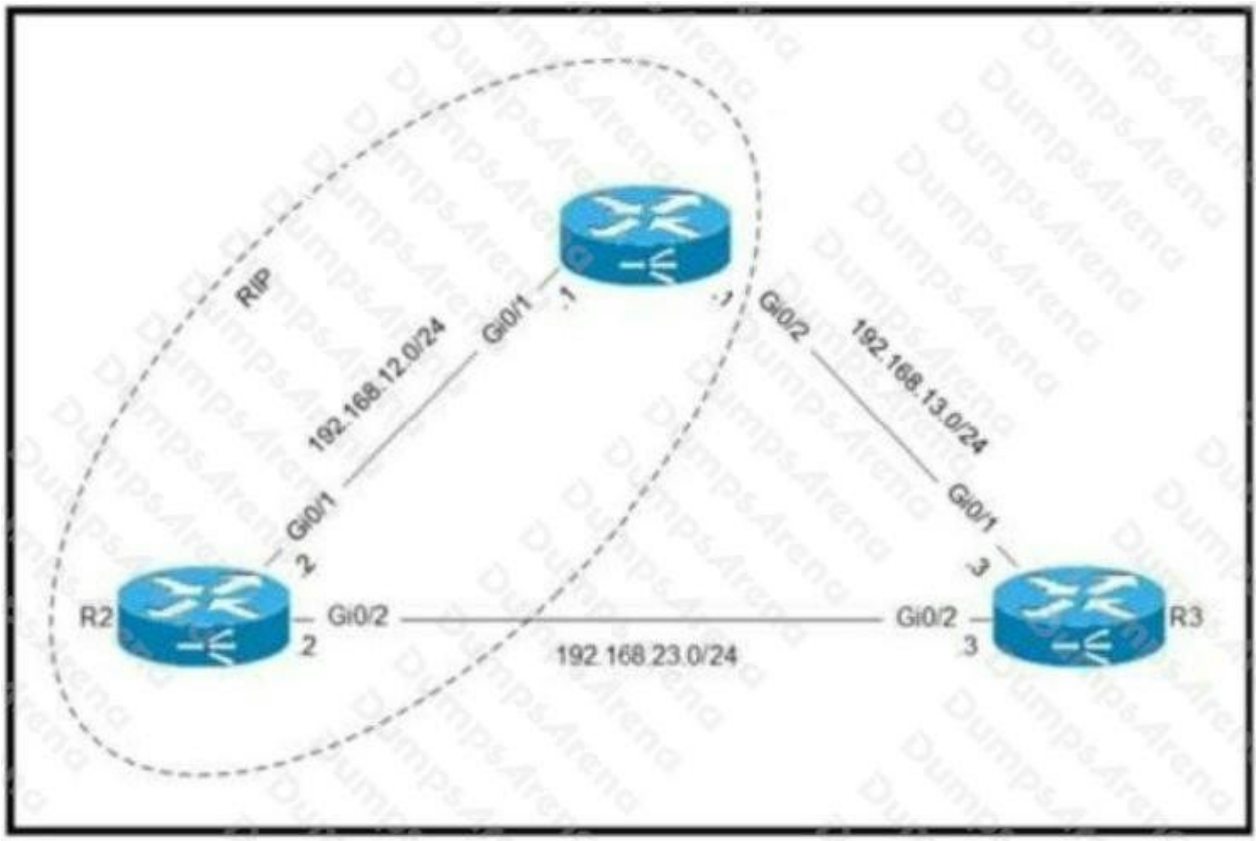
D. boolean

ANSWER: C

Explanation:

String is correct because the value shown in the exhibit is represented as text enclosed in double quotation marks. In JSON, a string is a sequence of zero or more Unicode characters delimited by quotation marks. Strings are used for textual values such as interface names, device hostnames, descriptions, addresses, identifiers, and status labels. This distinction matters in Cisco automation because JSON payloads exchanged with RESTful services, RESTCONF, and controller APIs preserve each value's data type. A client consuming a response must interpret quoted textual content as a string rather than attempt to process it as a numeric or logical value. JSON syntax also requires double quotation marks around both member names and string values, with special characters represented through defined escape sequences when necessary. The formal JSON specification defines a string as quotation marks surrounding zero or more Unicode characters and describes the permitted escape sequences. See RFC 8259, section 7. For practical Cisco DevNet examples of JSON being used as structured API data, see the Cisco DevNet REST API Basics learning module.

QUESTION NO: 164



Refer to the exhibit. Routers R1 and R2 are configured with RIP as the dynamic routing protocol. A network engineer must configure R1 with a floating static route to service as a backup route to network 192.168.23. which command must the engineer configure on R1?

A. ip route 192.168.23.0 255.255.255.0 192.168.13.3 100

B. ip route 192.168.23.0 255.255.255.255 192.168.13.3 121

C. ip route 192.168.23.0 255.255.255.0 192.168.13.3 121

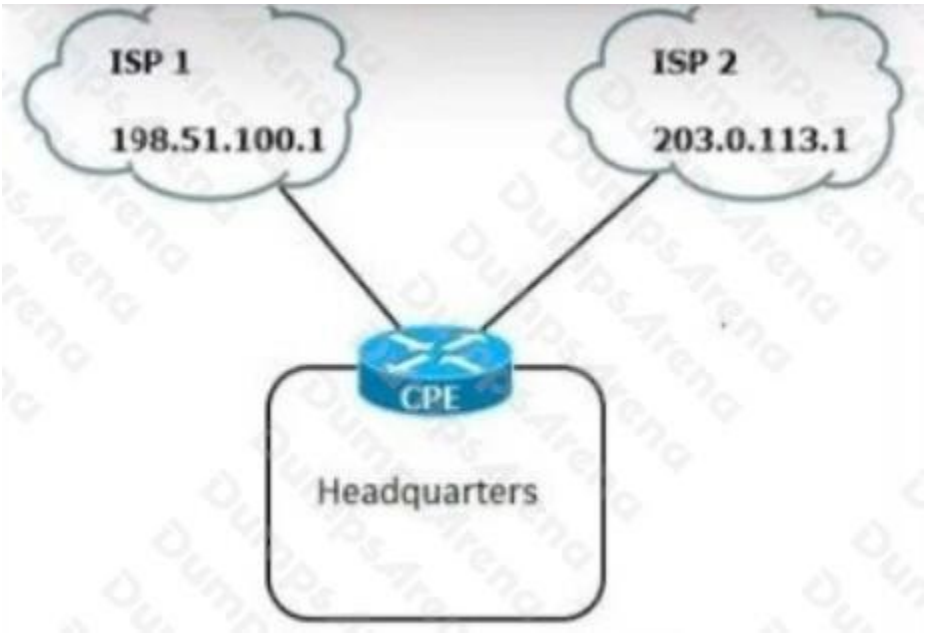
D. ip route 192.168.23.0 255.255.255.0 192.168.13.3

ANSWER: C

Explanation:

`ip route 192.168.23.0 255.255.255.0 192.168.13.3 121` correctly creates a floating static route on R1 for the entire 192.168.23.0/24 destination network. A floating static route is a backup route whose administrative distance is intentionally set higher than that of the routing protocol supplying the primary path. Cisco IOS assigns RIP routes an administrative distance of 120. Configuring the static route with an administrative distance of 121 means that RIP remains the preferred source for this prefix while its learned route is available. If R1 loses the RIP route, the static route becomes eligible for installation in the routing table and forwards traffic through next-hop address 192.168.13.3, which is the alternate path shown in the topology. The destination mask of 255.255.255.0 is also essential because it defines the required 192.168.23.0/24 network rather than an individual host. This configuration provides automatic backup forwarding without requiring manual route changes when the dynamically learned path fails. Cisco documents static-route administrative-distance configuration in its static routing configuration guide; the standard RIP administrative distance is listed in Cisco's administrative distance documentation.

QUESTION NO: 165



Refer to the exhibit. A network administrator configures the CPE to provide internet access to the company headquarters. Traffic must be load-balanced via ISP1 and ISP2 to ensure redundancy.

Which two command sets must be configured on the CPE router? (Choose two.)

A. ip route 0.0.0.0 0.0.0.0 198.51.100.1 255
ip route 0.0.0.0 0.0.0.0 203.0.113.1 255
ip route 128.0.0.0 128.0.0.0 203.0.113.1

B. ip route 0.0.0.0 128.0.0.0 198.51.100.1
ip route 128.0.0.0 128.0.0.0 203.0.113.1
ip route 0.0.0.0 0.0.0.0 198.51.100.1
ip route 0.0.0.0 0.0.0.0 203.0.113.1

C. ip route 0.0.0.0 0.0.0.0 198.51.100.1
ip route 0.0.0.0 0.0.0.0 203.0.113.1

D. ip route 0.0.0.0 128.0.0.0 198.51.100.1
ip route 128.0.0.0 128.0.0.0 203.0.113.1

E. ip route 0.0.0.0 0.0.0.0 198.51.100.1
ip route 0.0.0.0 0.0.0.0 203.0.113.1 2

ANSWER: C D

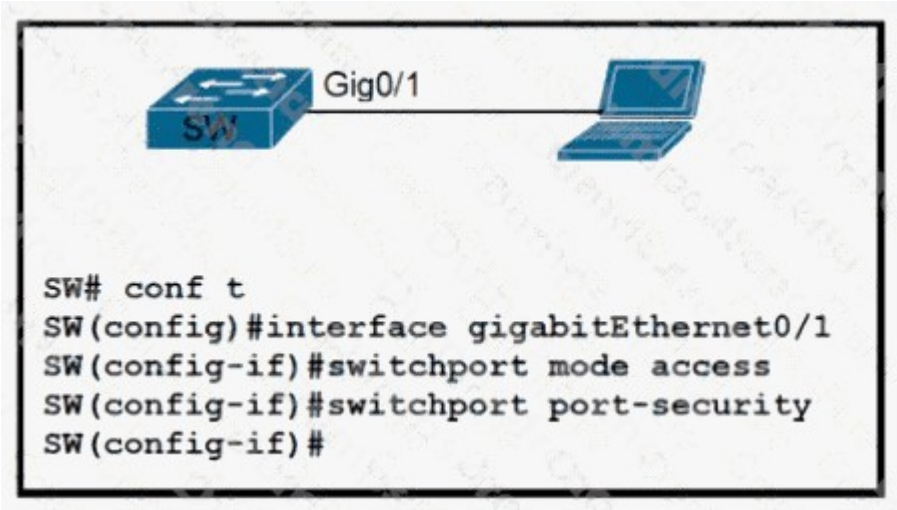
Explanation:

The command set `ip route 0.0.0.0 0.0.0.0 198.51.100.1` and `ip route 0.0.0.0 0.0.0.0 203.0.113.1` creates two static default routes with identical prefix length and administrative distance. Cisco IOS can install equal-cost routes and use Cisco Express Forwarding load sharing to forward Internet-bound traffic through both next hops. Because either default route can provide reachability for destinations not otherwise known in the routing table, this design also supplies an alternate path when one route is withdrawn after the associated path fails.

The command set `ip route 0.0.0.0 128.0.0.0 198.51.100.1` and `ip route 128.0.0.0 128.0.0.0 203.0.113.1` divides the IPv4 address space into two /1 prefixes. The first route directs destinations from 0.0.0.0 through 127.255.255.255 to ISP1, and the second directs destinations from 128.0.0.0 through 255.255.255.255 to ISP2. This is a static, deterministic form of traffic distribution: each ISP carries approximately one half of the possible IPv4 destination space. Cisco documents equal-cost multipath behavior and static-route configuration at Cisco ECMP Load Sharing and Cisco IOS XE Static Routes Configuration Guide.

QUESTION NO: 166

Refer to the exhibit



A network engineer started to configure port security on a new switch. These requirements must be met:

* MAC addresses must be learned dynamically

* Log messages must be generated without disabling the interface when unwanted traffic is seen

Which two commands must be configuredd to complete this task " ? (Choose two)

A. SW(config-if)=switchport port-security mac-address sticky

B. SW(config-if)=switchport port-security violation restrict

C. SW(config-if) switchport port-security mac-address 0010.7B84.45E6

D. SW(config-if) switchport port-security maximum 2

E. SW(config-if)=switchport port-security violation shutdown

ANSWER: A B

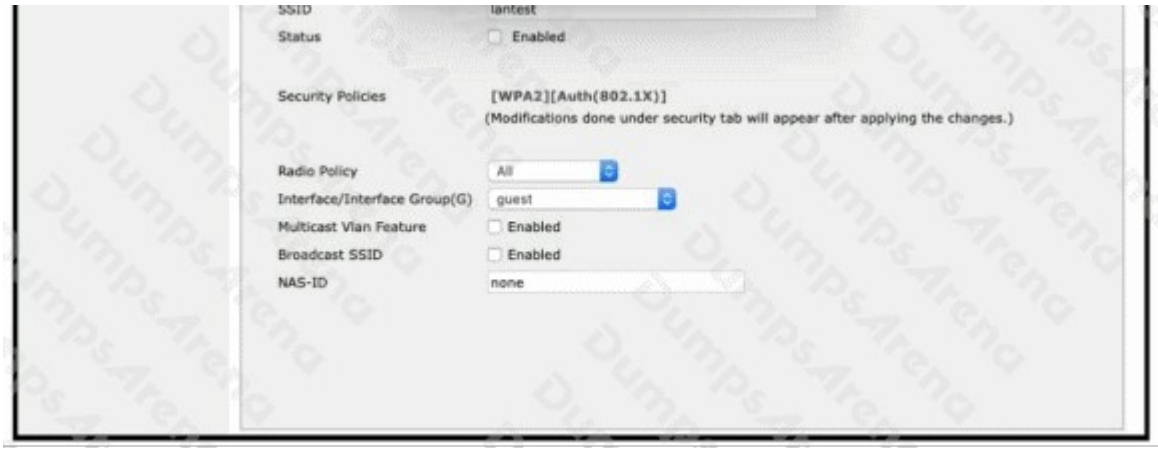
Explanation:

`switchport port-security mac-address sticky` satisfies the dynamic-learning requirement. With sticky secure MAC addressing enabled, the switch learns source MAC addresses dynamically on the secured interface and converts the learned addresses into sticky secure MAC addresses. These addresses are placed in the running configuration and can be retained permanently by saving the configuration. This provides dynamic discovery while still enforcing port-security behavior for the learned endpoints.

`switchport port-security violation restrict` satisfies the requirement to log unwanted traffic while keeping the interface operational. In restrict violation mode, frames from an unauthorized source are dropped, the security-violation counter is incremented, and the switch generates notifications, including syslog messages and SNMP traps when configured. Unlike an interface-disabling response, restrict mode allows the port to continue forwarding traffic from authorized secure MAC addresses after a violation occurs. Together, sticky learning and restrict mode implement secure dynamic endpoint learning plus visible violation reporting without taking the port out of service. Cisco documents sticky MAC learning and the restrict violation action in its Catalyst switch security configuration guide and in the Cisco IOS XE port-security documentation.

QUESTION NO: 167

Refer to the exhibit.



A Cisco engineer creates a new WLAN called lantest. Which two actions must be performed so that only high-speed 2.4-Ghz clients connect? (Choose two.)

A. Enable the Broadcast SSID option

B. Enable the Status option.

C. Set the Radio Policy option to 802 11g Only.

D. Set the Radio Policy option to 802.11a Only.

E. Set the Interface/Interface Group(G) to an interface other than guest

F. Disable 802.11b (1, 2, 5.5, 11 Mbps) data rates / require OFDM (802.11g) rates so only high-speed 2.4-GHz clients can associate.

ANSWER: B C F

Explanation:

Enabling the Status option is required to place the new WLAN into service. A configured WLAN that is administratively disabled is not available for client association, regardless of its radio-policy or data-rate settings. Setting the Radio Policy option to 802 11g Only restricts the WLAN to the 2.4-GHz 802.11g radio policy rather than advertising it on the 5-GHz 802.11a band. To strictly ensure that only high-speed 2.4-GHz clients associate, the controller must also exclude legacy 802.11b modulation and rates. Disabling the 1, 2, 5.5, and 11 Mbps data rates and requiring an OFDM rate prevents 802.11b-only clients from joining while permitting 802.11g-capable clients to connect. Cisco documents that disabling the 802.11b rates prevents 802.11b clients from associating and that mandatory data rates determine the rates clients must support for association. Together, enabling the WLAN, assigning its 2.4-GHz radio policy, and requiring OFDM-capable clients produce an active WLAN intended for high-speed 2.4-GHz operation. See Cisco's WLAN Data Rates and 802.11b Support guidance and the Cisco Wireless Controller WLAN configuration guide.

QUESTION NO: 168

Which two VPN technologies are recommanded by Cisco for multiple branch offices and large-scale deployments? (Choose two.)

A. site-to-site VPN

B. DMVPN

C. GETVPN

D. IPsec remote access

E. clientless VPN

ANSWER: B C

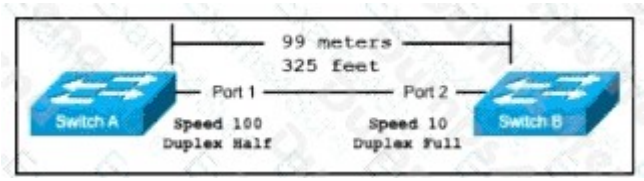
Explanation:

DMVPN and **GETVPN** are Cisco VPN technologies designed to scale beyond a small number of static point-to-point tunnels. DMVPN uses a hub-and-spoke control-plane design with multipoint GRE, NHRP, and IPsec protection. Branch routers can dynamically discover and establish direct, encrypted spoke-to-spoke tunnels when traffic requires it. This substantially reduces the tunnel-configuration burden compared with building a separate static tunnel for every pair of branch sites, making it well suited to organizations with many distributed offices.

GETVPN is intended for large, private WAN environments in which branch sites communicate over a trusted group network, such as an MPLS VPN. Instead of encapsulating traffic in individual tunnels, GETVPN uses group security associations and centrally managed encryption policies. A key server distributes policy and keying information to group members, allowing any authorized site to encrypt traffic to another site while preserving the original IP header. This group-based model supports scalable any-to-any connectivity and centralized security management. Cisco describes DMVPN as a scalable solution for dynamically created IPsec tunnels and GETVPN as a group-encryption technology for private WANs; see Cisco GETVPN overview and Cisco DMVPN configuration guide.

QUESTION NO: 169

Exhibit.



The switches are connected via a Cat5 Ethernet cable that was successfully tested. The Interfaces are configured as access ports and are both in a ' down " status. What is the cause of this issue?

A. The switches are configured with incompatible duplex settings.

B. The distance between the two switches is not supported by Cut5.

C. The speed settings on the switches are mismatched.

D. The portfast command is missing froin the configuration.

ANSWER: C

Explanation:

The speed settings on the switches are mismatched. Ethernet interfaces at either end of a copper link must operate at a mutually supported speed to establish a physical Layer 1 link. If one switch interface is manually configured for a speed that differs from the speed configured on its neighbor, the interfaces cannot complete link signaling successfully, so both ports can remain down even when the cable itself has passed testing. Access-port configuration affects Layer 2 switching behavior after the physical connection is active; it does not bring up a link whose Ethernet speed settings cannot interoperate. Cisco recommends using compatible speed and duplex settings on directly connected interfaces. When speed is configured manually on both ends, the configured values must match. This condition is verified with commands such as `show interfaces status` and `show interfaces`, which display each port's configured and operational speed. See Cisco's troubleshooting guide for Ethernet switching ports and the Cisco IOS interface command reference for interface speed configuration details.

QUESTION NO: 170

Which type does a port become when it receives the best BPDU on a bridge?

A. The designated port

B. The backup port

C. The alternate port

D. The root port

ANSWER: D

Explanation:

The root port is correct. In Spanning Tree Protocol, every nonroot bridge selects one root port: the interface offering that bridge's best path toward the root bridge. The selection is based on the superior BPDU received by the bridge. BPDUs advertise root-path information, and the port receiving the most favorable information is selected as the root port. The comparison begins with the lowest root bridge ID, then considers the lowest root-path cost; if necessary, additional BPDU tie-breakers such as the sending bridge ID and sending port ID determine the winning path.

After election, the root port is a forwarding port in the active topology and is the bridge's direction toward the root bridge. A root bridge itself has no root port, because it is the reference point for the spanning-tree topology. In Rapid STP terminology, the same core role applies: a nonroot switch uses its root port as its best available path to the root. Cisco's STP overview describes root ports as the ports on nonroot switches with the lowest path cost to the root bridge. See Cisco's Spanning Tree Protocol overview and Cisco's Spanning Tree Protocol technology page.

QUESTION NO: 171

An engineer must configure an OSPF neighbor relationship between router R1 and R3 The authentication configuration has been configured and the connecting interfaces are in the same 192.168 1.0/30 sublet. What are the next two steps to complete the configuration? (Choose two.)

A. configure the hello and dead timers to match on both sides

B. configure the same process ID for the router OSPF process

C. configure the same router ID on both routing processes

D. Configure the interfaces as OSPF active on both sides.

E. configure both interfaces with the same area ID

ANSWER: D E

Explanation:

OSPF must be enabled on each of the two connected interfaces before either router sends or processes OSPF Hello packets on that link. Therefore, *Configure the interfaces as OSPF active on both sides.* is required. This can be accomplished by matching the interface address with a `network` command under the OSPF routing process, or by explicitly applying `ip ospf <process-id> area <area-id>` in interface configuration mode. In either approach, the link must participate in OSPF at both ends.

The two interfaces must also belong to the same OSPF area, making *configure both interfaces with the same area ID* required. The area identifier is included in OSPF Hello packets and must agree between directly connected neighbors for an adjacency to progress. With compatible authentication already in place, matching IPv4 addressing on the point-to-point subnet, OSPF enabled on both interfaces, and a common area, the routers have the required configuration elements to establish the neighbor relationship. Cisco documents area-ID agreement and enabled OSPF operation on the connected network as prerequisites for OSPF neighbor formation. See the Cisco OSPF Neighbor Relationship documentation and the Cisco IOS XE OSPF Configuration Guide.

QUESTION NO: 172

Which two primary drivers support the need for network automation? (Choose two.)

A. Eliminating training needs

B. Increasing reliance on self-diagnostic and self-healing

C. Policy-derived provisioning of resources

D. Providing a ship entry point for resource provisioning

E. Reducing hardware footprint

ANSWER: B C

Explanation:

Increasing reliance on self-diagnostic and self-healing is a primary driver for network automation because modern networks operate at a scale and speed that make manual detection, troubleshooting, and remediation increasingly difficult. Automation enables telemetry-driven monitoring and repeatable operational workflows that can identify abnormal conditions, correlate relevant information, and initiate approved corrective actions more quickly and consistently. This supports a more resilient network by reducing the time between fault detection and response.

Policy-derived provisioning of resources is also a central automation driver. Instead of configuring devices individually through repetitive command-line changes, operators can express business and technical intent as policies and use automation tools to deploy the required configurations consistently across the infrastructure. This approach makes it easier to scale services, maintain configuration consistency, and adapt the network as application or organizational requirements change. Cisco describes intent-based networking as translating business intent into network behavior and continuously validating that the network is operating as intended. Cisco also emphasizes automation as a way to simplify operations and manage infrastructure efficiently through software-driven workflows. See Cisco Intent-Based Networking and Cisco Automation Solutions.

QUESTION NO: 173

What facilitates a Telnet connection between devices by entering the device name?

A. SNMP

B. DNS lookup

C. syslog

D. NTP

ANSWER: B

Explanation:

DNS lookup facilitates a Telnet connection when an administrator enters a device name rather than a numeric IP address. Before Telnet can establish its TCP session, the initiating device must translate the entered hostname into the destination device's IPv4 or IPv6 address. DNS provides this hostname-to-address resolution, allowing the Telnet client to identify where it must connect. On Cisco IOS, DNS resolution is enabled with the `ip domain lookup` command and DNS servers can be configured using `ip name-server`. For example, after DNS is configured and contains a record for `router.example.com`, entering `telnet router.example.com` causes IOS to resolve that name and initiate the Telnet session to the resulting address. Cisco IOS can also use locally configured hostname mappings with `ip host`, but DNS lookup is the named service and mechanism described by the question. DNS is therefore essential to the convenience of reaching network devices through recognizable hostnames instead of requiring administrators to remember and enter IP addresses. Cisco documents IOS DNS client operation and the associated configuration commands in its DNS Configuration Guide and its DNS on Cisco IOS technical note.

QUESTION NO: 174

What are two functions of an SDN controller? (Choose two)

A. Layer 2 forwarding

B. coordinating VTNs

C. tracking hosts

D. managing the topology

E. protecting against DDoS attacks

ANSWER: C D

Explanation:

tracking hosts and **managing the topology** are core control-plane functions of an SDN controller. The controller maintains a centralized, logically consistent view of network devices, links, endpoints, and their connectivity. By learning endpoint location and host-related information, it can determine where traffic sources and destinations are attached and apply the intended connectivity or policy consistently as those endpoints move or change.

Topology management supplies the controller with the network graph needed to calculate available paths, assess link and device relationships, and program forwarding behavior through southbound interfaces. This centralized view separates control decisions from the packet-forwarding data plane: network devices forward traffic according to the policies and flow instructions that the controller distributes. Cisco describes SDN as an architecture that centralizes network intelligence and abstracts underlying infrastructure, allowing software-based control of network behavior. Cisco controller platforms likewise use topology and endpoint information to automate policy-driven connectivity and forwarding decisions.

These functions are fundamental to the controller's role of translating high-level connectivity intent into network-wide forwarding state. See Cisco's Software-Defined Networking overview and the Cisco Application Policy Infrastructure Controller overview.

QUESTION NO: 175

Which security program element involves installing badge readers on data-center doors to allow workers to enter and exit based on their job roles?

A. physical access control

B. biometrics

C. role-based access control

D. multifactor authentication

ANSWER: A

Explanation:

physical access control is correct because badge readers at data-center doors control access to a physical facility or restricted area. A physical access control system validates a person's badge or card credential at an entry point and applies the organization's authorization policy before unlocking the door. The policy can assign permissions according to job responsibilities, so a data-center engineer may be authorized to enter the server room while personnel without that operational need are not. This is an important administrative and technical safeguard for protecting network infrastructure, servers, cabling, and other assets from unauthorized physical access. Physical access controls also commonly maintain entry and exit records, allowing security teams to audit who accessed a protected location and when. Cisco's guidance on network security recognizes physical security as a foundational protection for network devices and facilities; controls such as secured wiring closets, locked equipment rooms, and managed entry systems reduce the risk of tampering, theft, or disruption. Badge readers are therefore a direct implementation of physical access control, even when the underlying access permissions are organized around employee roles. See Cisco's Physical Security resources and the Cisco SAFE Security Architecture Guide.

QUESTION NO: 176

Refer to the exhibit. This ACL is configured to allow client access only to HTTP, HTTPS, and DNS services via UDP. The new administrator wants to add TCP access to the DNS service. Which configuration updates the ACL efficiently?

A. 35 permit tcp 10.0.0.0 0.255.255.255 host 198.51.100.11 eq domain

B. ip access-list extended Services

C. permit tcp 10.0.0.0 0.255.255.255 host 198.51.100.11 eq domain

D. ip access-list extended Services

ANSWER: A

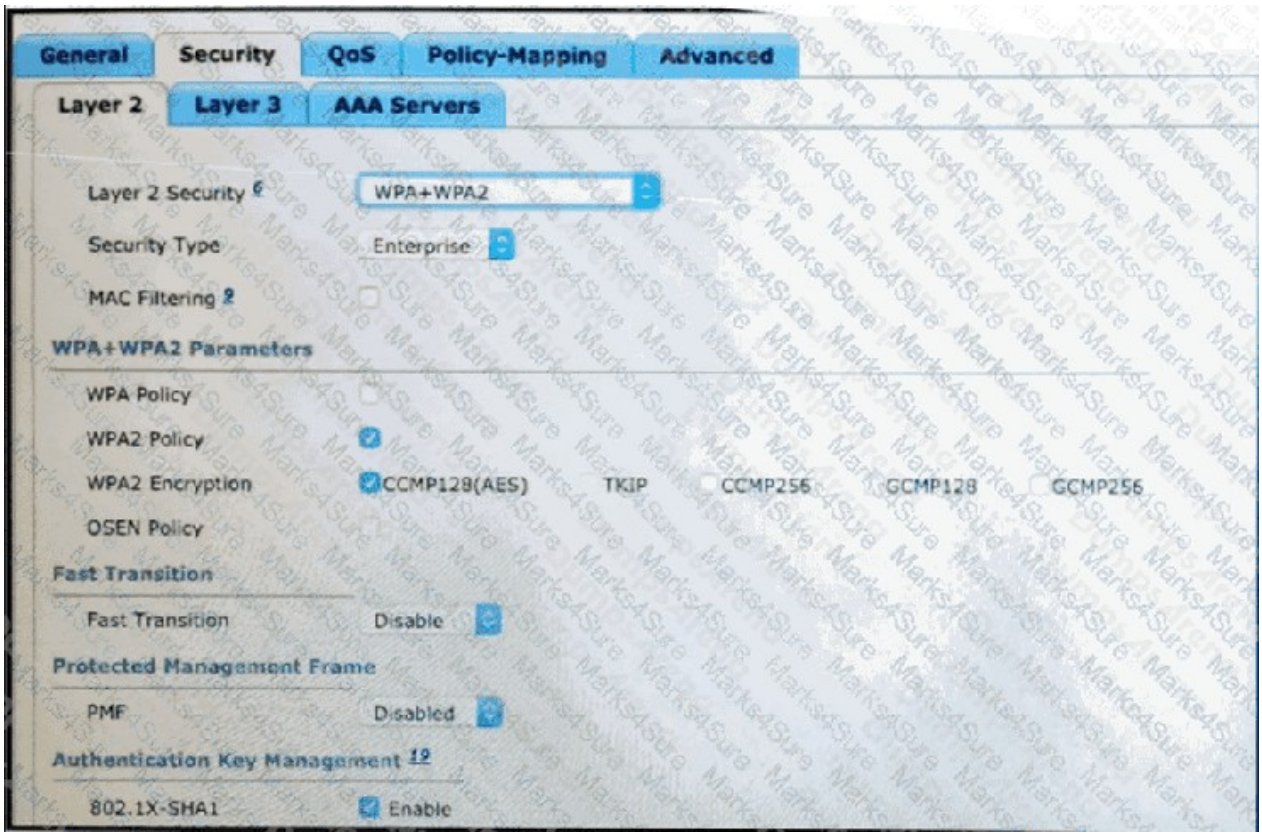
Explanation:

The configuration `35 permit tcp 10.0.0.0 0.255.255.255 host 198.51.100.11 eq domain` efficiently adds the required DNS-over-TCP permission at the appropriate point in the existing extended ACL. DNS uses well-known port 53, represented in Cisco IOS ACL syntax by the `domain` port keyword. Although many ordinary DNS queries use UDP, TCP port 53 is required for some DNS operations, including responses that exceed UDP size limits and zone transfers.

Using sequence number 35 places the new permit entry between the existing service permits and the subsequent catch-all deny entry shown in the ACL. ACL entries are processed in ascending sequence order and evaluation stops at the first matching entry. Explicit sequence numbering therefore allows the administrator to insert only the needed ACE without deleting and recreating the ACL or disrupting its established order. The source wildcard mask `0.255.255.255` matches the 10.0.0.0/8 client network, while the `host` keyword limits the destination to 198.51.100.11. Cisco documents ordered ACL processing and sequence-numbered ACL entries in its IPv4 Access Control Lists Configuration Guide. Port 53 protocol details are defined by RFC 1035.

QUESTION NO: 177

Refer to the exhibit.



What must be configured to enable 802.11w on the WLAN?

A. Set PMF to Required.

B. Enable MAC Filtering.

C. Enable WPA Policy.

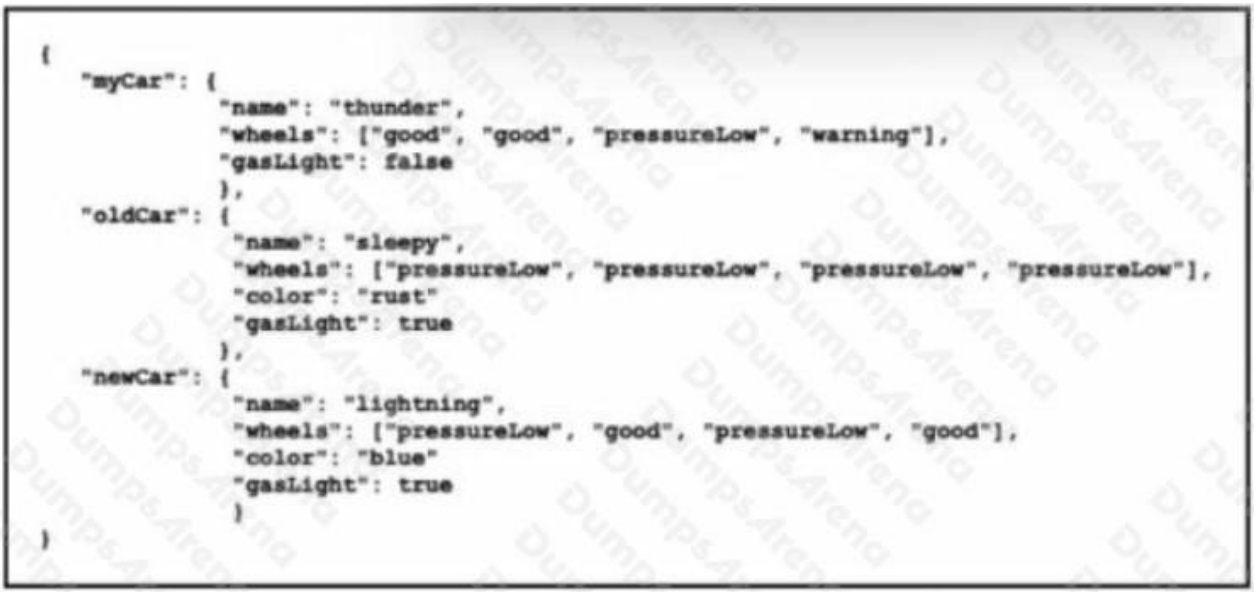
D. Set Fast Transition to Enabled

ANSWER: A

Explanation:

Set PMF to Required. 802.11w is the IEEE amendment that protects selected wireless management frames, and Cisco implements this capability as Protected Management Frames (PMF). Management frames such as deauthentication and disassociation frames can otherwise be spoofed, so PMF provides integrity protection for these robust management frames. On a Cisco WLAN, setting the PMF policy to *Required* enables 802.11w enforcement: clients must support PMF and use it when joining the WLAN. This is commonly used for WLANs that require protected management traffic and is also a security prerequisite for some modern WLAN security configurations. PMF protection is negotiated as part of the wireless security association, so the WLAN configuration—not merely a client-side preference—must specify the PMF policy. Cisco documents PMF as the configuration feature used to support IEEE 802.11w and describes the required policy as allowing only PMF-capable clients to associate. See Cisco's Protected Management Frames configuration guide and the Wi-Fi Alliance's overview of Wi-Fi security protections.

QUESTION NO: 178



Refer to the exhibit. In which structure does the word "warning" directly reside?

A. array

B. object

C. Boolean

D. string

ANSWER: A

Explanation:

The word *"warning"* directly resides in an **array**. In the exhibited JSON, it appears as a quoted value within square brackets. Square brackets define an array in JSON, and each comma-separated item inside those brackets is an array element. Therefore, the immediate JSON container of the `"warning"` value is the array.

JSON represents structured data using objects, arrays, and primitive values. Arrays provide an ordered collection of values and use the syntax `[ ... ]`. A quoted text value such as `"warning"` is a string value, but the question asks for the

structure in which that word directly resides. The relevant structural relationship is that the string is an element of the square-bracketed collection shown in the exhibit.

This distinction is important when interpreting API responses and configuration data: square brackets identify a list whose members can be accessed by position. Cisco automation and REST API workflows commonly return JSON arrays when a response contains multiple related values or records. See the JSON grammar in RFC 8259 and Cisco's overview of JSON data formats.

QUESTION NO: 179

A DHCP pool has been created with the name CONTRO

L. The pool uses the next to last usable IP address as the default gateway for the DHCP clients. The server is located at 172.16.32.15. What is the next step in the process for clients on the 192.168.52.0/24 subnet to reach the DHCP server?

A. ip helper-address 172.16.32.15

B. ip default-gateway 192.168.52.253

C. ip forward-protocol udp 137

D. ip detault-network 192.168.52.253

E. The pool uses the next to last usable IP address as the default gateway for the DHCP clients. The server is located at 172.16.32.15. What is the next step in the process for clients on the 192.168.52.0/24 subnet to reach the DHCP server?

ANSWER: A

Explanation:

`ip helper-address 172.16.32.15` is the required DHCP relay configuration command. DHCP clients initially transmit a DHCPDISCOVER as a local broadcast because they do not yet have an IP address or a route to remote networks. Routers do not forward those broadcasts between subnets by default, so clients on the 192.168.52.0/24 network cannot directly reach a DHCP server at 172.16.32.15.

The command must be applied under the Layer 3 interface or SVI that receives client broadcasts from the 192.168.52.0/24 subnet. The router then acts as a DHCP relay agent: it receives the broadcast, places the receiving interface address in the relay information as needed, and unicasts the request to 172.16.32.15. This enables the DHCP server to select the appropriate scope for the client subnet and return a DHCP offer through the relay. The configured pool can therefore provide addresses and the next-to-last usable address as the clients' default gateway.

Cisco documents this relay behavior and the interface-level use of `ip helper-address` in its DHCP relay overview and DHCP relay agent configuration guide.

QUESTION NO: 180

Refer to the exhibit. An engineer is creating a secure preshared key based SSID using WPA2 for a wireless network running on 2.4 GHz and 5 GHz. Which two tasks must the engineer perform to complete the process? (Choose two.)

A. Select the 802.1 x option for Auth Key Management.

B. Select the WPA Policy option.

C. Select the PSK option for Auth Key Management.

D. Select the AES option for Auth Key Management.

E. Select the AES (CCMP128) option for WPA2/WPA3 Encryption.

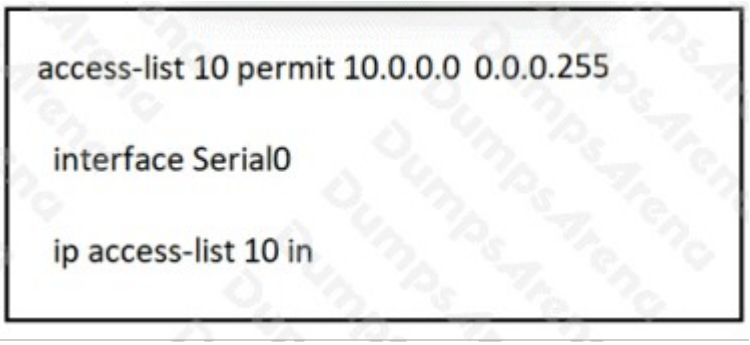
ANSWER: C E

Explanation:

A WPA2 personal SSID authenticates clients with a shared passphrase rather than with an external RADIUS/AAA service and individual user credentials. Therefore, **Select the PSK option for Auth Key Management.** is required: PSK establishes the personal-mode authentication key-management method used by every client joining the SSID. The engineer must also configure **Select the AES (CCMP128) option for WPA2/WPA3 Encryption.** WPA2 protection relies on AES using the CCMP cipher suite to provide confidentiality and message integrity for wireless traffic. This combination represents WPA2-Personal security: a configured pre-shared key for authentication and AES-CCMP for protected over-the-air data frames. The selected security settings apply to the WLAN/SSID and are independent of whether clients associate on the 2.4-GHz or 5-GHz radio, provided that the SSID is enabled on both bands. Cisco documents WPA2-Personal as using a pre-shared key and identifies AES-CCMP as the applicable WPA2 encryption method. See Cisco's wireless security configuration guidance and the Cisco WLAN security overview.

QUESTION NO: 181

Refer to the exhibit.



A network administrator must permit traffic from the 10.10.0.0/24 subnet to the WAN on interlace Seria10. What is the effect of the configuration as the administrator applies the command?

A. The permit command fails and returns an error code.

B. The router accepts all incoming traffic to Seria10 with the last octet of the source IP set to 0.

C. The sourced traffic from IP range 10.0.0.0 -10.0.0.255 is allowed on Seria10.

D. The router fails to apply the access list to the interface.

ANSWER: B

Explanation:

The router accepts all incoming traffic to Seria10 with the last octet of the source IP set to 0. Cisco IOS access control lists use wildcard masks rather than conventional subnet masks. In a wildcard mask, a binary 0 means that the corresponding address bit must match the base address, while a binary 1 means that bit is ignored. Thus, with a source address of `10.10.0.0` and a wildcard mask of `255.255.255.0`, the first three octets are entirely ignored and the final octet must match 0. This produces a match for source addresses in the form `x.x.x.0`, rather than matching every address in `10.10.0.0/24`. Because the access list is applied inbound on the serial WAN interface, packets arriving there whose source address ends in `.0` match the permit entry and can proceed according to the ACL. To represent the intended /24 source network in an IOS ACL, the required wildcard mask would be `0.0.0.255`. Cisco documents wildcard-mask matching and ACL configuration in its ACL samples and access-list configuration guidance.

QUESTION NO: 182

What is a benefit for external users who consume public cloud resources?

A. implemented over a dedicated WAN

B. located in the same data center as the users

C. all hosted on physical servers

D. accessed over the Internet

ANSWER: D

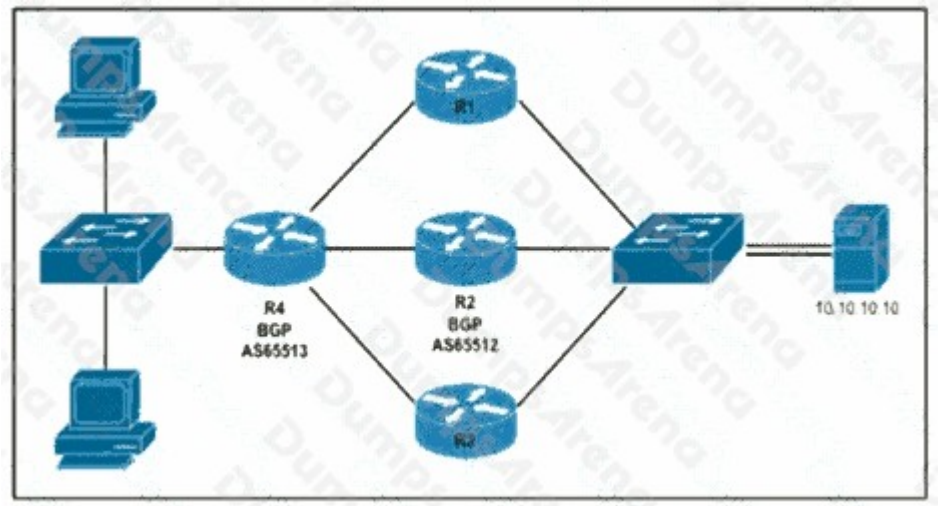
Explanation:

Accessed over the Internet is a central benefit of public cloud resources for external users. Public cloud providers make computing, storage, applications, and other services available through provider-managed infrastructure that customers can reach using standard network mechanisms, typically the Internet. This gives geographically distributed users a practical way to use services without needing to be physically present at a provider facility or connected to an organization's internal network. Users commonly access cloud services through web portals, HTTPS-based applications, APIs, or client software, subject to the provider's identity, authentication, authorization, and security controls.

This characteristic aligns with NIST's essential cloud-computing characteristic of broad network access: cloud capabilities are available over the network and accessed through standard mechanisms that support varied client platforms. For external consumers, Internet accessibility enables rapid consumption of services from different locations and devices, while the provider operates the underlying infrastructure. Connectivity and access can still be restricted by policies, firewalls, identity controls, or private-access designs when required, but Internet-based reachability remains a defining practical advantage of the public-cloud consumption model. See NIST SP 800-145, The NIST Definition of Cloud Computing and Cisco: What Is Cloud Computing?.

QUESTION NO: 183

Refer to the exhibit.



Router R4 is dynamically learning the path to the server. If R4 is connected to R1 via OSPF Area 20, to R2 via R2 BGP, and to R3 via EIGRP 777, which path is installed in the routing table of R4?

A. the path through R1, because the OSPF administrative distance is 110

B. the path through R2. because the IBGP administrative distance is 200

C. the path through R2 because the EBGP administrative distance is 20

D. the path through R3. because the EIGRP administrative distance is lower than OSPF and BGP

ANSWER: C

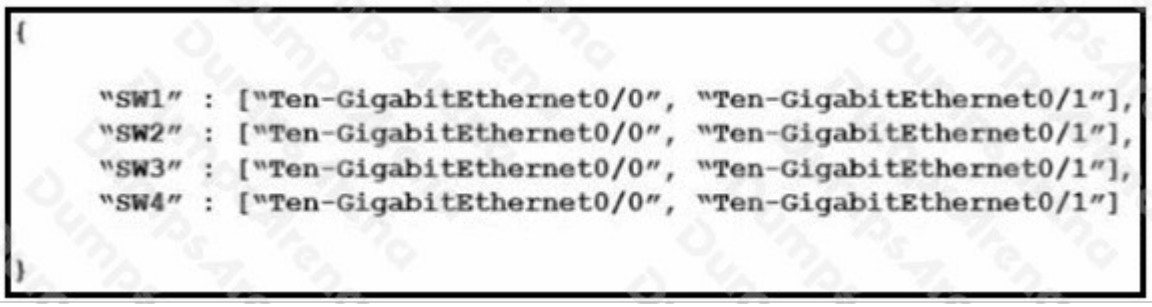
Explanation:

The path through R2 because the EBGP administrative distance is 20 is installed, assuming each routing protocol advertises the same destination prefix with an equal prefix length. When a Cisco router receives identical routes from multiple routing sources, it selects the source with the lowest administrative distance before considering the protocol-specific metric. By

default, externally learned Border Gateway Protocol routes have an administrative distance of 20. This value is lower than the default distance for internal Enhanced Interior Gateway Routing Protocol routes, which is 90, and Open Shortest Path First routes, which is 110. Therefore, R4 considers the route learned from the external BGP neighbor to be the most trustworthy source and installs that route in its Routing Information Base. The BGP path can then be used to populate the forwarding table, subject to normal route validity and next-hop reachability requirements. Administrative distance is locally significant: it is a route-source preference applied by R4 and is not advertised to neighboring routers. Cisco documents the default administrative-distance values in its administrative distance documentation; BGP route selection behavior is further described in the Cisco IOS XE BGP configuration guide.

QUESTION NO: 184

Refer to the exhibit.



How many JSON objects are presented?

A. 1

B. 2

C. 3

D. 4

ANSWER: D

Explanation:

4 is correct because the JSON structure in the exhibit contains four values delimited by curly braces. In JSON, an object is a collection of zero or more name/value members enclosed by { and }. Each separately delimited brace structure counts as an object, including objects that occur as elements within an array or as nested values within another object. Square brackets, by contrast, identify an array and do not themselves create an additional JSON object. Therefore, the count must be based on the four distinct curly-brace-delimited structures shown in the payload, rather than on the number of arrays, fields, or nesting levels. This distinction is important when reading API output and when working with automation tools that parse structured data: an array can hold multiple object records, while each record remains an individual JSON object. The formal JSON grammar defines an object using braces and an array using brackets; see RFC 8259, The JavaScript Object Notation Data Interchange Format. Cisco's DevNet learning material also discusses JSON as a common structured format used when consuming REST API data: Cisco DevNet: Introduction to REST APIs.

QUESTION NO: 185

A client experiences slow throughput from a server that is directly connected to the core switch in a data center. A network engineer finds minimal latency on connections to the server, but data transfers are unreliable, and the output of the show Interfaces counters errors command shows a high FCS-Err count on the interface that is connected to the server. What is the cause of the throughput issue?

A. high bandwidth usage

B. a physical cable fault

C. a speed mismatch

D. a cable that is too long

ANSWER: B

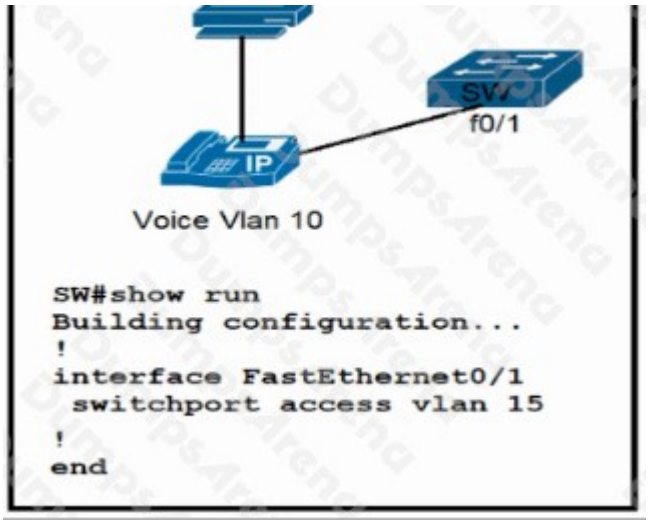
Explanation:

a physical cable fault is the cause of the throughput problem. FCS-Err identifies Ethernet frames received with an invalid Frame Check Sequence, meaning their contents were corrupted in transit and failed the integrity check performed by the receiving interface. This is a Layer 1 symptom: the physical medium or its associated hardware is introducing errors into the signal. Common sources include damaged copper cabling, poor cable termination, a loose connector, contaminated or failing fiber optics, an incompatible optic, or a faulty NIC or switch interface.

Although latency can remain minimal because the server is directly attached to the core switch, corrupted frames cannot be successfully delivered. Higher-layer protocols such as TCP detect the missing or damaged data and retransmit it. Those retransmissions reduce effective application throughput and make transfers appear unreliable without necessarily increasing measured round-trip latency substantially. The engineer should validate the physical link by checking connectors and transceivers, replacing or testing the cable, and, if necessary, moving the connection to known-good ports and hardware. Cisco identifies CRC/FCS errors as an indicator to investigate cabling and other physical-link conditions. See Cisco Ethernet troubleshooting guidance and Cisco interface error counter guidance.

QUESTION NO: 186

Refer to the exhibit.



All VLANs are present in the VLAN database. Which command sequence must be applied to complete the configuration?

A. Interface FastEthernet0/1 switchport trunk native vlan 10 switchport trunk allowed vlan 10,15

B. Interface FastEthernet0/1 switchport mode trunk switchport trunk allowed vlan 10,15

C. interface FastEthernet0/1 switchport mode access switchport voice vlan 10

D. Interface FastEthernet0/1 switchport trunk allowed vlan add 10 vlan 10 private-vlan isolated

ANSWER: B

Explanation:

`Interface FastEthernet0/1 switchport mode trunk switchport trunk allowed vlan 10,15` is correct because FastEthernet0/1 must carry traffic for more than one VLAN across the link. Configuring `switchport mode trunk` administratively places the switch port into trunking mode, allowing it to transmit and receive VLAN-tagged Ethernet frames. This is the required operating mode when the connected link must extend both VLAN 10 and VLAN 15 beyond the local switch.

The `switchport trunk allowed vlan 10,15` command then explicitly limits the VLANs carried by that trunk to VLAN 10 and VLAN 15. Restricting the allowed VLAN list is a sound operational practice: it ensures that only the VLANs required by the topology traverse the link, reduces unnecessary broadcast-domain extension, and helps maintain the intended segmentation. Because the VLANs already exist in the VLAN database, no VLAN-creation commands are needed; the remaining task is to make the interconnecting interface a trunk and permit the required VLANs. Cisco documents trunk configuration and VLAN pruning in its VLAN trunking overview and VLAN trunk configuration guide.

QUESTION NO: 187

What are two characteristics of an SSID? (Choose Two)

A. It can be hidden or broadcast in a WLAN

B. It uniquely identifies an access point in a WLAN

C. It uniquely identifies a client in a WLAN

D. It is at most 32 characters long.

E. IT provides secured access to a WLAN

ANSWER: A D

Explanation:

An SSID is the Service Set Identifier: the name assigned to an IEEE 802.11 wireless LAN that enables users to recognize and select that WLAN. “It can be hidden or broadcast in a WLAN” is correct because an access point can include the SSID in beacon and probe-response management frames, making the WLAN name discoverable, or suppress that name in beacons. In the latter configuration, often called a hidden SSID, a station that already knows the configured network name can still attempt to join the WLAN. Cisco notes that suppressing SSID broadcast is a visibility setting, not an access-control mechanism.

“It is at most 32 characters long.” is also correct in the conventional WLAN configuration sense. The IEEE 802.11 SSID information element supports a name from zero through 32 octets; network platforms commonly present this limit to administrators as a maximum of 32 characters. This bounded identifier allows the same WLAN name to be configured consistently across multiple access points participating in an extended service set, so client devices can recognize the intended wireless network while moving within its coverage area. See Cisco’s discussion of hidden SSIDs and Cisco’s wireless LAN configuration guide.

QUESTION NO: 188

What is the temporary state that switch ports always enter immediately after the boot process when Rapid PVST+ is used?

A. forwarding

B. listening

C. learning

D. discarding

ANSWER: D

Explanation:

The correct answer is **discarding**. Rapid PVST+ is Cisco’s implementation of Rapid Spanning Tree Protocol (RSTP), standardized as IEEE 802.1w and incorporated into IEEE 802.1D. RSTP defines three port states: discarding, learning, and forwarding. When a switch port initializes, it begins in the discarding state while the switch exchanges BPDUs, determines the root bridge and port roles, and safely establishes the loop-free spanning-tree topology. In this state, the port neither forwards user data traffic nor learns source MAC addresses into the switching table. This protects the network from transient

Layer 2 loops during convergence. The discarding state consolidates the legacy STP blocking and listening behaviors into one RSTP state. Once the required spanning-tree negotiation has completed, a port that is selected to participate in forwarding can advance through learning and then forwarding. RSTP can make this transition rapidly on eligible links, such as point-to-point links using the proposal/agreement process, and edge ports configured with PortFast can move to forwarding quickly when they connect to endpoint devices. Cisco documents Rapid PVST+ port states and the initial discarding behavior in its spanning-tree guidance: Cisco STP/RSTP Overview and Port States. The IEEE RSTP state model is also summarized by Cisco Rapid Spanning Tree Protocol documentation.

QUESTION NO: 189

Which channel-group mode must be configured when multiple distribution interfaces connected to a WLC are bundled?

A. Channel-group mode passive.

B. Channel-group mode on.

C. Channel-group mode desirable.

D. Channel-group mode active.

ANSWER: B

Explanation:

Channel-group mode on. is correct because a Cisco Wireless LAN Controller configured for link aggregation uses a static EtherChannel (also called Link Aggregation Group, or LAG) across its distribution-system ports. The switch ports facing the controller must therefore be placed into the same static port channel with the `channel-group <number> mode on` command. This setting forces EtherChannel formation without exchanging a link-aggregation negotiation protocol.

On the controller, enabling LAG combines the physical distribution interfaces into one logical interface and provides increased bandwidth plus link redundancy. The corresponding switch configuration must present one compatible Layer 2 EtherChannel to that logical controller interface. Cisco's WLC deployment guidance specifically states that the channel on the switch should be configured statically, using "mode on," when connecting WLC distribution ports configured for LAG. Consistent configuration of the port-channel number, trunking characteristics, VLAN allowances, and native-VLAN settings across member interfaces is also required for proper operation.

See Cisco's Wireless LAN Controller EtherChannel configuration guidance and its EtherChannel configuration and troubleshooting documentation.

QUESTION NO: 190

Which two pieces of information can you learn by viewing the routing table? (Choose two.)

A. whether an ACL was applied inbound or outbound to an interface

B. the EIGRP or BGP autonomous system

C. whether the administrative distance was manually or dynamically configured

D. which neighbor adjacencies are established

E. the length of time that a route has been known

ANSWER: C E

Explanation:

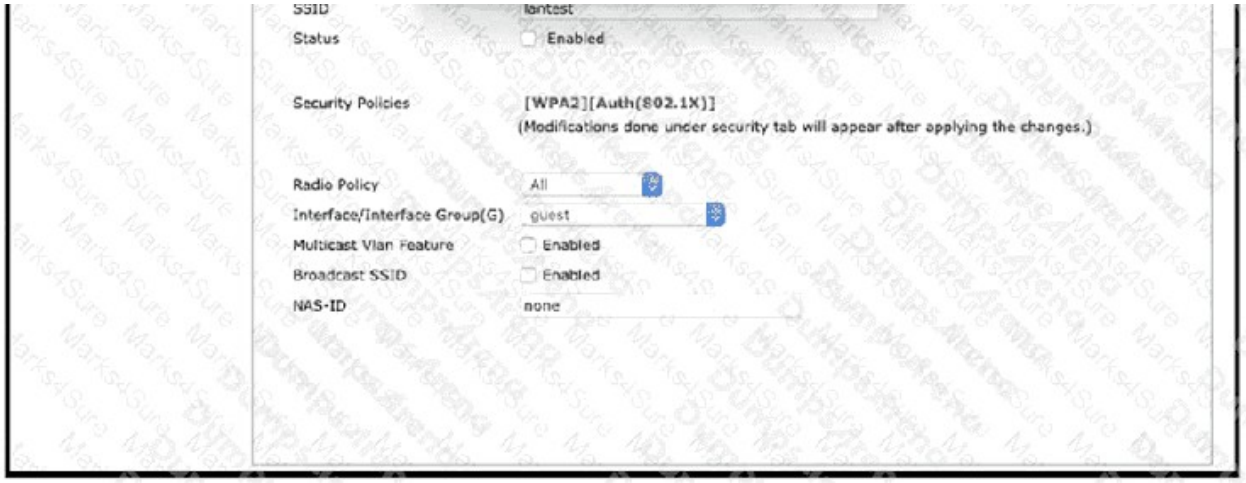
The routing table can reveal whether the administrative distance was manually or dynamically configured because it displays the administrative distance currently assigned to each installed route. In Cisco IOS output, the first value in brackets is the route's administrative distance, followed by its metric; for example, `[110/20]` represents an administrative distance of 110

and a metric of 20. An administrator can compare that displayed distance with the normal default administrative distance for the route's source to recognize a nondefault value and thereby identify that the preference has been configured rather than left at its protocol default. This is useful when determining why one route was selected over another route to the same destination.

The routing table also displays the length of time that a route has been known. The route entry includes an age field after the next-hop information, such as `00:12:33`. This timer shows how long the route has been present or known since its most recent update, helping an administrator assess route stability and recent routing changes. Cisco documents the route age, administrative distance, metric, next hop, and exit interface as fields available in route displays. See Cisco's guide to interpreting the routing table and the Cisco IOS show ip route command reference.

QUESTION NO: 191

Refer to the exhibit.



A Cisco engineer creates a new WLAN called lantest. Which two actions must be performed so that only high-speed 2.4-GHz clients connect? (Choose two.)

A. Enable the Broadcast SSID option

B. Enable the Status option.

C. Set the Radio Policy option to 802.11g Only.

D. Set the Radio Policy option to 802.11a Only.

E. Set the Interface/Interface Group(G) to an interface other than guest

ANSWER: B C

Explanation:

"Enable the Status option." is required because a WLAN must be administratively enabled on the wireless LAN controller before clients can associate with it. Creating and naming the WLAN alone does not make it available for client connections; enabling its status activates the configured WLAN service.

"Set the Radio Policy option to 802.11g Only." restricts the WLAN to the 2.4-GHz 802.11g PHY. IEEE 802.11g uses the 2.4-GHz band and provides PHY data rates up to 54 Mb/s, meeting the requirement for high-speed 2.4-GHz connectivity while excluding legacy 802.11b-only associations. The radio policy is the WLAN-level setting that determines which radio technologies may advertise and service that WLAN, so setting it to 802.11g only enforces the intended client capability requirement. Together, enabling the WLAN and applying the 802.11g-only radio policy make lantest operational and limit it to the specified 2.4-GHz high-speed client type. Cisco's WLAN configuration guidance describes enabling WLANs and selecting radio policies, and Cisco's wireless technology overview summarizes 802.11g operation in the 2.4-GHz spectrum. See Cisco Wireless Controller WLAN Configuration Guide and Cisco wireless standards overview.

QUESTION NO: 192

What is an appropriate use for private IPv4 addressing?

A. on the public-facing interface of a firewall

B. to allow hosts inside to communicate in both directions with hosts outside the organization

C. on internal hosts that stream data solely to external resources

D. on hosts that communicates only with other internal hosts

ANSWER: D

Explanation:

Private IPv4 addressing is appropriate on hosts that communicates only with other internal hosts. The private IPv4 ranges defined by RFC 1918 are intended for use within private networks: 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. These addresses are not globally unique and are not advertised or routed across the public Internet. This makes them well suited to systems whose communication remains inside an organization, such as user endpoints, internal servers, management networks, and other private infrastructure.

An organization can use these ranges repeatedly in isolated environments because Internet routers are expected to filter private-address space rather than forward it as public traffic. If internal hosts later need Internet connectivity, a network edge device can translate private source addresses to a public address by using Network Address Translation. That potential need does not change the core use case: private addresses provide internal addressing while conserving globally routable IPv4 addresses. Cisco networking designs commonly use private addressing for inside networks and apply NAT at the boundary where traffic must reach external networks. See RFC 1918: Address Allocation for Private Internets and Cisco's NAT configuration overview.

QUESTION NO: 193

OSPF must be configured between routers R1 and R2. Which OSPF configuration must be applied to router R1 to avoid a DR/BDR election?

A. router ospf 1
network 192.168.1.1 0.0.0.0 area 0
interface e1/1
ip address 192.168.1.1 255.255.255.252
ip ospf network broadcast

B. router ospf 1
network 192.168.1.1 0.0.0.0 area 0
interface e1/1
ip address 192.168.1.1 255.255.255.252
ip ospf network point-to-point

C. router ospf 1
network 192.168.1.1 0.0.0.0 area 0
interface e1/1
ip address 192.168.1.1 255.255.255.252
ip ospf cost 0

D. router ospf 1
network 192.168.1.1 0.0.0.0 area 0
hello interval 15
interface e1/1
Ip address 192.168.1.1 255.255.255.252

ANSWER: B

Explanation:

The configuration containing `ip ospf network point-to-point` is correct because the point-to-point OSPF network type is intended for a link with two endpoints, such as the R1-to-R2 connection shown. On a point-to-point network, OSPF creates a direct adjacency between the two routers and does not use a designated router or backup designated router. DR and BDR roles exist to reduce adjacency and LSA-flooding overhead on multiaccess segments, where more than two OSPF routers can share the same Layer 2 network. That optimization is unnecessary on a two-router point-to-point link.

The command is applied in interface configuration mode, so placing it under `interface e1/1` correctly changes the OSPF network type for that interface. The interface is then enabled for OSPF in area 0 by the matching `network` statement. For a working neighbor relationship, R2 must use a compatible OSPF network type and matching relevant OSPF interface parameters as well. Cisco documents that point-to-point links have no DR/BDR election and form a full adjacency between the connected routers. See Cisco's OSPF network types documentation and the Cisco IOS OSPF Configuration Guide.

QUESTION NO: 194

What is represented by the word "firewall" within this JSON schema?

```
1 [
2 {"router": "R_pittsburgh", "port":"e0/36"},
3 {"firewall": "FW37", "port":"te23"},
4 {"switch": "SW_toronto", "port":"ge5/28"},
5 ]
```

A. value

B. key

C. object

D. array

ANSWER: B

Explanation:

The correct answer is **key**. JSON represents an object as a set of named members enclosed in curly braces. Each member follows the form `"name": value`: the quoted name before the colon identifies the data item, while the content after the colon supplies that item's value. Therefore, when `"firewall"` appears on the left side of a colon in the displayed JSON schema, it is the key (also formally called a member name). It labels the associated property so that a program, API client, or network-automation workflow can retrieve the corresponding data by name.

This distinction is essential when reading JSON returned by controller APIs or when working with Cisco automation tools. For example, a JSON object may use a key such as `"firewall"` to identify a nested configuration structure or a particular setting; the associated value can itself be a string, number, Boolean value, array, or another object. RFC 8259 specifies that an object is an unordered collection of zero or more name/value pairs and that names are strings, which establishes why `"firewall"` is a key in this context. See RFC 8259: The JavaScript Object Notation (JSON) Data Interchange Format and MDN Web Docs: Working with JSON.

QUESTION NO: 195

What are two benefits of network automation? (Choose two.)

A. reduced hardware footprint

B. reduced operational costs

C. faster changes with more reliable results

D. fewer network failures

E. increased network security

ANSWER: B C

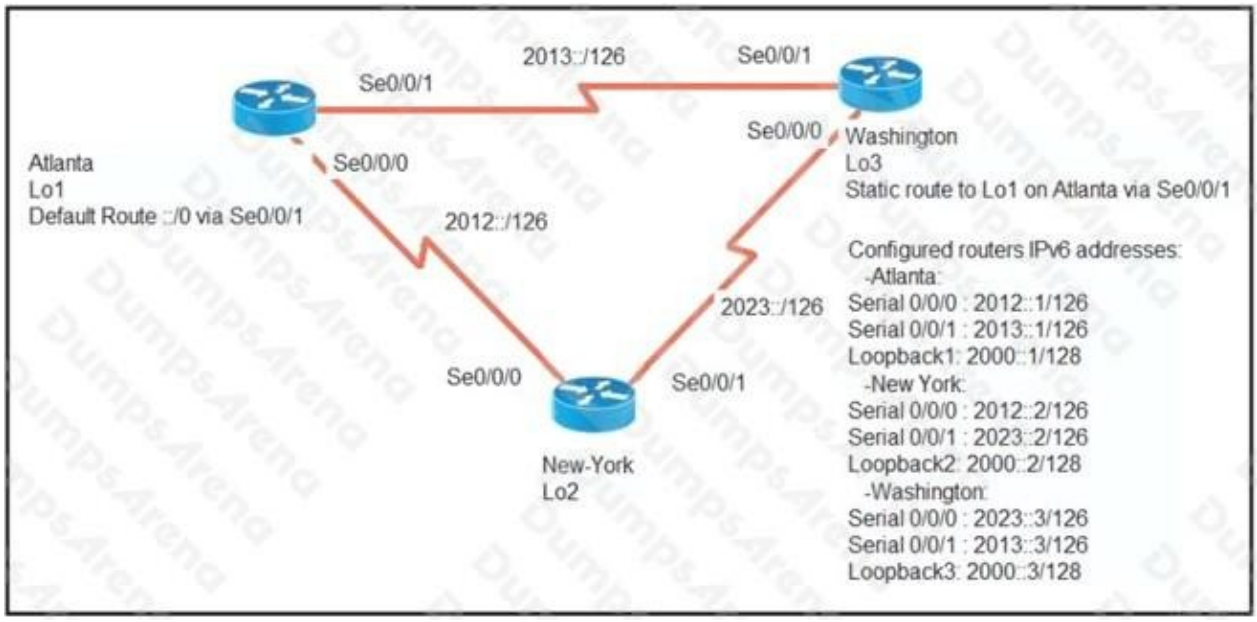
Explanation:

Network automation provides the greatest operational value by making routine network tasks repeatable at scale. **reduced operational costs** is a direct benefit because automation reduces the time engineers spend performing repetitive configuration, validation, inventory, and troubleshooting activities manually. Teams can operate larger environments more efficiently and redirect skilled staff toward design, optimization, and higher-value engineering work. Cisco describes automation as a way to simplify operations and increase IT efficiency.

faster changes with more reliable results is also a core benefit. Automation uses defined workflows, templates, APIs, and source-controlled configuration data to apply intended changes consistently across devices. This shortens deployment time while reducing variability between manually configured devices. When automation is paired with testing, review, and validation processes, teams can make changes more predictably and repeatedly, which is particularly valuable in environments with many network devices or frequent policy updates.

These benefits align with Cisco's automation and intent-based networking approach, which emphasizes simplifying operations, improving consistency, and accelerating deployment through programmable infrastructure. See Cisco's Network Automation overview and its Intent-Based Networking overview.

QUESTION NO: 196



Refer to the exhibit. An engineer is configuring the New York router to reach the Lo1 interface of the Atlanta router using interface Se0/0/0 as the primary path. Which two commands must be configured on the New York router so that it reaches the Lo1 interface of the Atlanta router via Washington when the link between New York and Atlanta goes down? (Choose two.)

A. Ipv6 route 2000::1/128 2012::1

B. Ipv6 route 2000::1/128 2012::1 5

C. Ipv6 route 2000::1/128 2012::2

D. Ipv6 route 2000::1/128 2023::2 5

E. Ipv6 route 2000::1/128 2023::3 5

ANSWER: A D

Explanation:

The correct configuration uses a primary IPv6 static route and a floating IPv6 static route for the same destination prefix, `2000::1/128`. `Ipv6 route 2000::1/128 2012::1` supplies the normal route through the next hop on the direct New York-to-Atlanta path. Because no administrative distance is explicitly configured, this static route uses the default administrative distance of 1 and is preferred while the direct path remains available.

`Ipv6 route 2000::1/128 2023::2 5` provides the alternate route through Washington. Its administrative distance of 5 makes it a floating static route: it is less preferred than the primary static route, so it is not installed as the active route while the primary route is valid. When the direct New York-to-Atlanta link fails and the primary path is no longer usable, the router selects the higher-distance route through Washington, preserving reachability to Atlanta's loopback interface. This design provides deterministic primary-path forwarding with automatic static-route failover, without requiring a dynamic routing protocol.

Cisco documents that static routes have a default administrative distance of 1 and that specifying a greater value creates a floating static route. See Cisco Default Administrative Distances and Cisco IPv6 Static Route Configuration Guide.

QUESTION NO: 197

Which two actions influence the EIGRP route selection process? (Choose two.)

A. The advertised distance is calculated by a downstream neighbor to inform the local router of the bandwidth on the link. Router A considers the first metric (50) as the Advertised distance. The second metric (90), which is from NEVADA to IOWA (through IDAHO), is called the Feasible distance.

The reported distance is calculated in the same way of calculating the metric. By default (K1 = 1, K2 = 0, K3 = 1, K4 = 0, K5 = 0), the metric is calculated as follows:

Feasible successor is the backup route. To be a feasible successor, the route must have an Advertised distance (AD) less than the Feasible distance (FD) of the current successor route.
Feasible distance (FD): The sum of the AD plus the cost between the local router and the next-hop router. The router must calculate the FD of all paths to choose the best path to put into the routing table.
Note: Although the new CCNA exam does not have EIGRP topic but you should learn the basic knowledge of this routing protocol.

B. The router calculates the feasible distance of all paths to the destination route.

C. The router must use the advertised distance as the metric for any given route.

D. The router calculates the best backup path to the destination route and assigns it as the feasible successor.

E. The router calculates the reported distance by multiplying the delay on the exiting interface by 256.

ANSWER: B D

Explanation:

EIGRP uses the Diffusing Update Algorithm (DUAL) to evaluate paths held in its topology table. The router calculates the feasible distance of all paths to the destination route because feasible distance is EIGRP's total composite metric from the local router to that destination through a particular next hop. With the default metric weights, this composite metric is derived primarily from minimum bandwidth and cumulative delay. EIGRP selects the path or paths with the lowest feasible distance as successor routes and installs them in the routing table; equal-cost successors can support load balancing.

The router also calculates the best backup path to the destination route and assigns it as the feasible successor when an alternate path meets EIGRP's feasibility condition. Specifically, the neighbor's reported distance for that alternate path must be lower than the feasible distance of the current successor. This condition ensures the alternate is loop-free. A feasible successor is retained in the topology table rather than installed as the active forwarding route. If the successor becomes unavailable, DUAL can promote the feasible successor immediately without sending queries or waiting for a new network-wide route computation. Therefore, calculating feasible distances and identifying qualifying feasible successors directly determine EIGRP's preferred and immediately usable alternate routes. Cisco documents these DUAL route-selection

concepts in its EIGRP Feasible Successor and Feasibility Condition documentation and Cisco IOS XE EIGRP configuration guide.

QUESTION NO: 198

Refer to the exhibit.

```
Switch#show etherchannel summary
[output omitted]

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------
10     Po10(SU)         LACP      Gi0/0(P)    Gi0/1(P)
20     Po20(SU)         LACP      Gi0/2(P)    Gi0/3(P)
```

Which two commands were used to create port channel 10? (Choose two)

- int range g0/0-1
  channel-group 10 mode active
- int range g0/0-1
  channel-group 10 mode desirable
- int range g0/0-1
  channel-group 10 mode passive
- int range g0/0-1
  channel-group 10 mode auto
- int range g0/0-1
  channel-group 10 mode on

A. Option A

B. Option B

C. Option C

D. Option D

E. Option E

ANSWER: A C

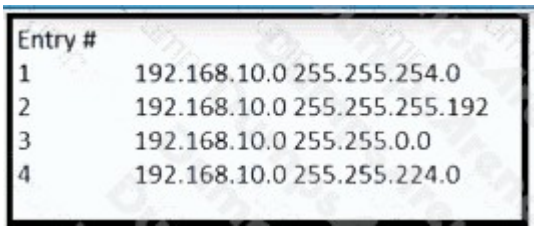
Explanation:

The supplied item marks *Option A* and *Option C* as the two commands used to create Port-Channel 10. EtherChannel configuration is built by assigning physical member interfaces to a common channel-group number; the matching logical Port-Channel interface is then created or used for the bundle. The selected command pair therefore represents the member-link association and the EtherChannel negotiation/configuration needed for channel-group 10. A port channel is a logical interface composed of multiple compatible physical Ethernet links, allowing the switch to treat the links as one bundled

connection for forwarding and configuration purposes. Cisco documents that EtherChannel member interfaces are added with the `channel-group channel-number mode` interface command, and that the resulting Port-Channel interface is identified by that channel number. For LACP-based bundles, the negotiation mode is configured on the member ports and must be compatible on the connected device. This behavior is described in Cisco's EtherChannel configuration guide and in the Cisco IOS EtherChannel documentation.

QUESTION NO: 199

Refer to the exhibit.



Which entry is the longest prefix match for host IP address 192.168.10.5?

A. 1

B. 2

C. 3

D. 4

ANSWER: B

Explanation:

2 is the longest-prefix match for destination host 192.168.10.5. A router evaluates all routing-table prefixes that contain the destination address and selects the route with the greatest prefix length, meaning the route that fixes the largest number of leading address bits. This produces the most specific route to the destination. The entry labeled 2 covers 192.168.10.5 and has the most-specific matching network prefix among the entries in the exhibit, so it is selected for forwarding. This decision is based on prefix specificity before considerations such as administrative distance or route metric; those values matter only when candidate routes have the same destination prefix length. The selected route therefore identifies the appropriate next-hop or exit-interface information for packets addressed to 192.168.10.5. Cisco documents that the routing table uses longest-prefix matching to select the best route when multiple routes can reach a destination. See Cisco's routing table selection guidance and the Cisco IOS IP routing command reference for routing-table behavior.

QUESTION NO: 200

Aside from discarding, which two states does the switch port transition through while using RSTP (802.1w)? (Choose two.)

A. blocking

B. speaking

C. listening

D. learning

E. forwarding

ANSWER: D E

Explanation:

RSTP, specified by IEEE 802.1w and incorporated into IEEE 802.1D, uses three port states: discarding, learning, and forwarding. Therefore, aside from discarding, the two states a switch port can transition through are **learning** and **forwarding**. In the learning state, the port processes received BPDUs and learns source MAC addresses to populate the MAC address table, but it does not forward normal user data frames. This prepares the switch to make efficient forwarding decisions once the port can safely participate in the topology. In the forwarding state, the port continues to learn MAC addresses and forwards user traffic normally. RSTP combines several legacy 802.1D conditions into the discarding state, allowing faster convergence while maintaining loop-prevention behavior. Cisco documentation maps the RSTP port states directly as discarding, learning, and forwarding and describes forwarding as the state in which normal traffic is sent and received. See Cisco's STP and RSTP port-state documentation and the Cisco Spanning Tree Protocol overview.

QUESTION NO: 201

An engineer is configuring remote access to a router from IP subnet 10.139.58.0/28. The domain name, crypto keys, and SSH have been configured. Which configuration enables the traffic on the destination router?

A. line vty 0 15 access-class 120 in ! ip access-list extended 120 permit tcp 10.139.58.0 0.0.0.15 any eq 22

B. interface FastEthernet0/0 ip address 10.122.49.1 255.255.255.252 ip access-group 10 in ! ip access-list standard 10 permit udp 10.139.58.0 0.0.0.7 host 10.122.49.1 eq 22

C. interface FastEthernet0/0 ip address 10.122.49.1 255.255.255.252 ip access-group 110 in ! ip access-list standard 110 permit tcp 10.139.58.0 0.0.0.15 eq 22 host 10.122.49.1

D. line vty 0 15 access-group 120 in ! ip access-list extended 120 permit tcp 10.139.58.0 0.0.0.15 any eq 22

ANSWER: A

Explanation:

`line vty 0 15 access-class 120 in ! ip access-list extended 120 permit tcp 10.139.58.0 0.0.0.15 any eq 22` is the appropriate configuration because SSH access to the router itself is management-plane traffic handled through the VTY lines. Applying an inbound `access-class` on those VTY lines restricts which remote hosts may establish VTY-based management sessions, including SSH. The wildcard mask `0.0.0.15` precisely matches the specified `10.139.58.0/28` source network, covering addresses from 10.139.58.0 through 10.139.58.15. The extended ACL entry permits TCP traffic from that authorized subnet to TCP port 22, which is the well-known SSH port. Because the ACL is associated with the VTY lines in the inbound direction, the router evaluates an incoming SSH connection attempt against the management-access policy before allowing the session. This approach directly limits remote administrative access without relying on the routing path of ordinary transit traffic. Cisco documents VTY access restriction with the `access-class` command and identifies TCP port 22 as the SSH service port. See Cisco's SSH configuration example and the Cisco IOS XE SSH configuration guide.

QUESTION NO: 202

How does CAPWAP communicate between an access point in local mode and a WLC?

A. The access point must directly connect to the WLC using a copper cable

B. The access point must not be connected to the wired network, as it would create a loop

C. The access point must be connected to the same switch as the WLC

D. The access point has the ability to link to any switch in the network, assuming connectivity to the WLC

ANSWER: D

Explanation:

The access point has the ability to link to any switch in the network, assuming connectivity to the WLC. In local mode, an access point joins and communicates with its wireless LAN controller through CAPWAP across an IP network; it does not require a physical, point-to-point connection or attachment to the controller's local switch. The essential requirement is Layer

3 IP reachability between the AP management interface and the WLC management interface. The AP can therefore be cabled to an appropriate access-layer switch at its physical location, receive an IP address, discover a controller, and establish its CAPWAP session through the routed enterprise network.

CAPWAP separates controller communication into a control channel and a data channel. Cisco documents that CAPWAP uses UDP port 5246 for control traffic and UDP port 5247 for data traffic. Network routing and any intervening ACLs or firewalls must permit this traffic so the AP can join and remain connected to the controller. AP discovery can occur through methods such as DHCP option 43, DNS, or local-subnet broadcast, after which the AP selects and joins a reachable WLC. This centralized controller relationship is a fundamental characteristic of Cisco lightweight AP deployments. See Cisco's Wireless Controller Configuration Guide and RFC 5415.

QUESTION NO: 203

Which type of organization should use a collapsed-core architecture?

A. large and requires a flexible, scalable network design

B. large and must minimize downtime when hardware fails

C. small and needs to reduce networking costs currently

D. small but is expected to grow dramatically in the near future

ANSWER: C

Explanation:

A collapsed-core architecture is appropriate for an organization that is small and needs to reduce networking costs currently. In the traditional three-tier campus model, the access, distribution, and core layers have distinct roles. A collapsed-core design combines the distribution and core functions, usually into a resilient pair of multilayer switches. This reduces the number of devices, interconnections, and operational elements required in the campus, which lowers both capital expense and management complexity.

Cisco positions this approach for small campus deployments in which network scale, traffic volume, and service requirements do not justify building and maintaining a dedicated core layer. The organization can still use hierarchical design principles, including redundant uplinks and Layer 3 routing at the collapsed core, while avoiding the expense of an additional core infrastructure. The design can therefore provide an effective balance of availability, performance, and cost for a smaller environment with current budget constraints.

As the campus expands, the organization can reassess whether its traffic and resiliency requirements warrant evolving to a dedicated core. Cisco campus design guidance describes the hierarchical campus model and the role of collapsed-core deployments: Cisco Enterprise Campus Design Zone and Cisco Campus Network Overview.

QUESTION NO: 204

Which two primary drivers support the need for network automation? (Choose two.)

A. Increasing reliance on self-diagnostic and self-healing

B. Eliminating training needs

C. Policy-driven provisioning of resources

D. Reducing hardware footprint

E. Providing a single entry point for resource provisioning

ANSWER: C E

Explanation:

Policy-driven provisioning of resources is a primary automation driver because modern networks must apply business and security intent consistently across many devices, users, and locations. Rather than relying on repeated device-by-device command-line changes, an automation platform can translate defined policies into configurations and deploy them in a repeatable manner. This supports consistent implementation, reduces configuration variation, and makes it practical to manage changes at scale. Cisco describes intent-based networking as using business intent and automation to simplify network operations and continuously align the network with the intended policy.

Providing a single entry point for resource provisioning is also a key driver because centralized controllers and orchestration tools provide a common operational interface and APIs for provisioning network resources. Centralized provisioning enables administrators and integrated systems to request, validate, deploy, and track services through consistent workflows rather than treating each device as an independent management target. This model improves operational speed, repeatability, governance, and visibility while supporting scalable automation across the infrastructure. Cisco Catalyst Center is an example of a centralized network-management platform that provides automation and programmable workflows for enterprise networks.

References: Cisco Intent-Based Networking; Cisco Catalyst Center.

QUESTION NO: 205

What is represented by the word "fe5/42" within this JSON schema?

```
1 [
2 {"load balancer": "LB_milwaukee", "port":"fe5/42"},
3 {"VPN concentrator": "VPNadmin", "port":"e1/39"},
4 {"firewall": "FW_chicago", "port":"te3/42"},
5 ]
```

A. array

B. object

C. value

D. key

ANSWER: C

Explanation:

The correct answer is **value**. In JSON, an object is composed of member pairs in the form `"name": value`. The quoted text `"fe5/42"` is a JSON string and, when it occurs to the right of a colon in the displayed object, it supplies the data associated with that member's property name. Therefore, its role is the value of the key/value pair. A JSON value may be a string, number, object, array, Boolean, or null; here, `"fe5/42"` specifically is a string value. The slash is simply a permitted character within the string and does not change the JSON data type or its role in the object. JSON syntax requires member names to be strings and separates each name from its associated value with a colon, which makes the position of this text in the shown structure the key detail for identifying it correctly. RFC 8259 defines JSON object members and the permitted forms of a JSON value: RFC 8259: The JavaScript Object Notation Data Interchange Format. Additional practical JSON structure guidance is available from MDN Web Docs: Working with JSON.

QUESTION NO: 206

What are two characteristics of a public cloud implementation? (Choose two.)

A. It is owned and maintained by one party, but it is shared among multiple organizations

B. It enables an organization to fully customize how it deploys network resources

C. It provides services that are accessed over the Internet

D. It is a data center on the public Internet that maintains cloud services for only one company

E. It supports network resources from a centralized third-party provider and privately-owned virtual resources

ANSWER: A C

Explanation:

“It is owned and maintained by one party, but it is shared among multiple organizations” describes the multitenant nature of a public-cloud deployment. A cloud service provider owns, operates, and maintains the underlying physical infrastructure, such as data centers, servers, storage, and networking. Multiple independent customer organizations consume services from that provider’s shared resource pool, while virtualization, access controls, and tenant isolation keep each customer’s workloads and data logically separated. This provider-operated, shared-consumption model enables broad scalability and elastic capacity without each organization having to purchase and operate all underlying infrastructure itself.

“It provides services that are accessed over the Internet” is also a defining public-cloud characteristic. Public-cloud services are commonly made available through internet-accessible provider portals, web interfaces, and programmatic service endpoints. Customers can provision and manage resources remotely and consume compute, storage, networking, and application services on demand. Organizations can additionally use private connectivity for particular workloads, but the provider’s services remain public-cloud offerings because they are delivered from provider-owned, multitenant infrastructure. NIST identifies broad network access and resource pooling as essential cloud-computing characteristics, and Cisco describes public cloud as services delivered by a third-party provider to multiple customers. See NIST SP 800-145 and Cisco: What Is Cloud Computing?.

QUESTION NO: 207

What are two behaviors of a point-to-point WAN topology? (Choose two.)

A. It leverages a dedicated connection.

B. It provides direct connections between each router in the topology.

C. It delivers redundancy between the central office and branch offices.

D. It uses a single router to route traffic between sites.

E. It connects remote networks through a single line.

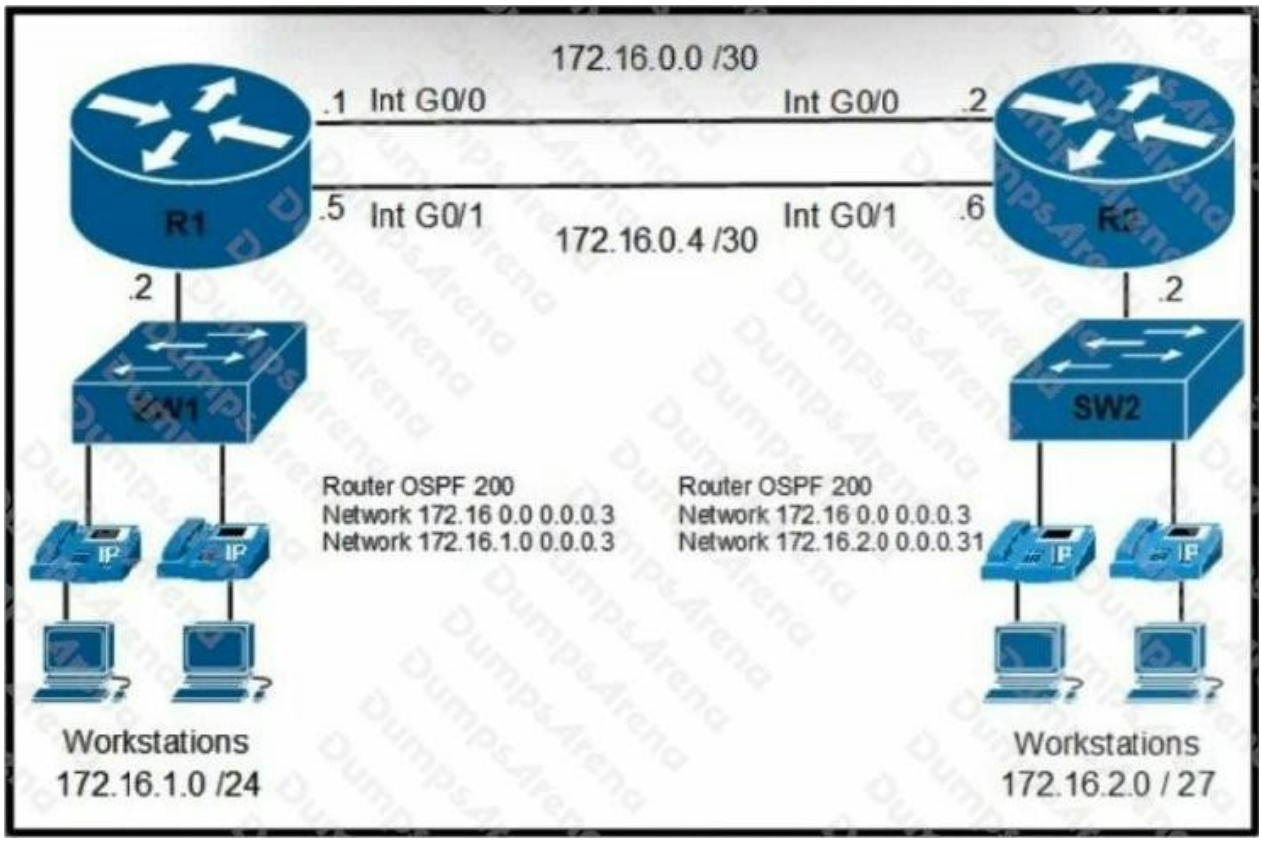
ANSWER: A E

Explanation:

A point-to-point WAN topology uses a connection that has exactly two endpoints, such as two routers at separate sites. “It leverages a dedicated connection” is correct because this topology is commonly implemented with a dedicated physical or logical circuit provisioned between those endpoints. Classic examples include leased lines and provider Ethernet virtual circuits. The circuit’s capacity and forwarding path are allocated for communication between the two connected locations rather than being a multipoint access medium shared by several customer sites.

“It connects remote networks through a single line” is also correct because the fundamental topology is one WAN link between one site and one other site. Each endpoint uses that single point-to-point circuit to exchange traffic with the remote network. This simple, direct site-to-site arrangement makes addressing, routing, and fault isolation straightforward. Point-to-point is a topology description; the connection can carry different Layer 2 encapsulations or services, including PPP and carrier-provided Ethernet services, provided that the service is presented as a two-endpoint circuit. Cisco describes PPP as a protocol for transporting multiprotocol datagrams over point-to-point links, reflecting this two-node WAN model. See Cisco Point-to-Point Protocol technology support and Cisco: What is a WAN?.

QUESTION NO: 208



Refer to the exhibit. The primary route across Gi0/0 is configured on both routers. A secondary route must be configured to establish connectivity between the workstation networks. Which command set must be configured to complete this task?

A. R1 -
ip route 172.16.2.0 255.255.255.248 172.16.0.5 110

R2 -
ip route 172.16.1.0 255.255.255.0 172.16.0.6 110

B. R1 -
ip route 172.16.2.0 255.255.255.240 172.16.0.2 113

R2 -
ip route 172.16.1.0 255.255.255.0 172.16.0.1 114

C. R1 -
ip route 172.16.2.0 255.255.255.224 172.16.0.6 111

R2 -
ip route 172.16.1.0 255.255.255.0 172.16.0.5 112

D. R1 -
ip route 172.16.2.0 255.255.255.240 172.16.0.5 89
R2 -
ip route 172.16.1.0 255.255.255.0 172.16.0.6 89

ANSWER: C

Explanation:

The command set that configures `ip route 172.16.2.0 255.255.255.224 172.16.0.6 111` on R1 and `ip route 172.16.1.0 255.255.255.0 172.16.0.5 112` on R2 correctly creates a secondary, or floating, static route. Each route identifies the remote workstation LAN with its proper network mask and uses the next-hop address on the alternate inter-router path: R1 forwards through 172.16.0.6, while R2 forwards through 172.16.0.5. This makes the backup path distinct from the existing Gi0/0 primary path.

The final numeric value in each static-route command is the administrative distance. Static routes normally have an administrative distance of 1. Values of 111 and 112 are higher than that primary static-route distance, so IOS does not install these backup routes into the active routing table while the primary route remains available. If the primary route is withdrawn because its path fails, the corresponding floating static route becomes the preferred available route and restores traffic between the workstation networks. Cisco describes this design as a floating static route: a backup static route configured with a higher administrative distance than the preferred route. See Cisco's floating static route documentation and the Cisco IOS XE static routing guide.

QUESTION NO: 209

What are two functions of DHCP servers? (Choose two.)

A. issue DHCPDISCOVER messages when added to the network

B. respond to client DHCPOFFER requests by Issuing an IP address

C. support centralized IP management

D. assign dynamic IP configurations to hosts in a network

E. prevent users from assigning their own IP addresses to hosts

ANSWER: C D

Explanation:

DHCP servers **support centralized IP management** by maintaining address pools (scopes), reservations, exclusions, lease durations, and records of addresses currently allocated to clients. Rather than requiring an administrator to configure addressing parameters separately on every endpoint, the network's addressing policy can be administered at the DHCP server. This centralization makes it practical to control address utilization and consistently distribute network settings across a LAN.

DHCP servers also **assign dynamic IP configurations to hosts in a network**. When a host joins a network and requests configuration, the DHCP process enables the server to lease an available IPv4 address and provide associated parameters. These commonly include the subnet mask, default gateway, DNS server addresses, domain name, and lease time, although the precise parameters depend on the configured DHCP options. A lease-based assignment lets addresses be reused when clients leave or no longer renew their leases, which is particularly useful for user devices that connect intermittently.

Cisco describes DHCP as a protocol that dynamically assigns IP addresses and related configuration information to hosts. The DHCP message and lease behavior are also defined by the IETF in RFC 2131. See Cisco's DHCP overview and message exchange for the operational flow.

QUESTION NO: 210

What are two benefits of controller-based networking compared to traditional networking? (Choose two.)

A. controller-based increases network bandwidth usage, while traditional lightens the load on the network

B. controller-based reduces network configuration complexity, while traditional increases the potential for errors

C. controller-based allows for fewer network failures, while traditional increases failure rates

D. controller-based provides centralization of key IT functions, while traditional requires distributed management functions

E. controller-based inflates software costs, while traditional decreases individual licensing costs

ANSWER: B D

Explanation:

Controller-based networking reduces network configuration complexity by abstracting device-level configuration into centrally defined policies, intent, and reusable automation workflows. Rather than requiring an administrator to configure each router, switch, or wireless device separately, a controller can provision and maintain consistent configurations across the infrastructure. This centralized, model-driven approach improves operational efficiency and helps reduce configuration errors caused by repetitive manual changes. Cisco describes its controller-based enterprise solutions as providing automated provisioning and policy-based automation across the network.

Controller-based networking also provides centralization of key IT functions. A controller gives network teams a centralized point for management, automation, policy definition, telemetry collection, monitoring, and assurance. This allows administrators to obtain broader network visibility and apply business or security policies consistently, instead of relying exclusively on distributed, device-by-device management. In Cisco enterprise architectures, Cisco Catalyst Center provides centralized management and automation capabilities that support these controller-based functions.

These benefits are fundamental to software-defined networking concepts covered at the CCNA level: separation of control logic from individual devices, centralized policy and management, and automation of operational tasks. See Cisco's Cisco Catalyst Center overview and its Software-Defined Access overview.

QUESTION NO: 211

Refer to Exhibit.

How does SW2 interact with other switches in this VTP domain?

A. It processes VTP updates from any VTP clients on the network on its access ports.

B. It receives updates from all VTP servers and forwards all locally configured VLANs out all trunk ports

C. It forwards only the VTP advertisements that it receives on its trunk ports.

D. It transmits and processes VTP updates from any VTP Clients on the network on its trunk ports

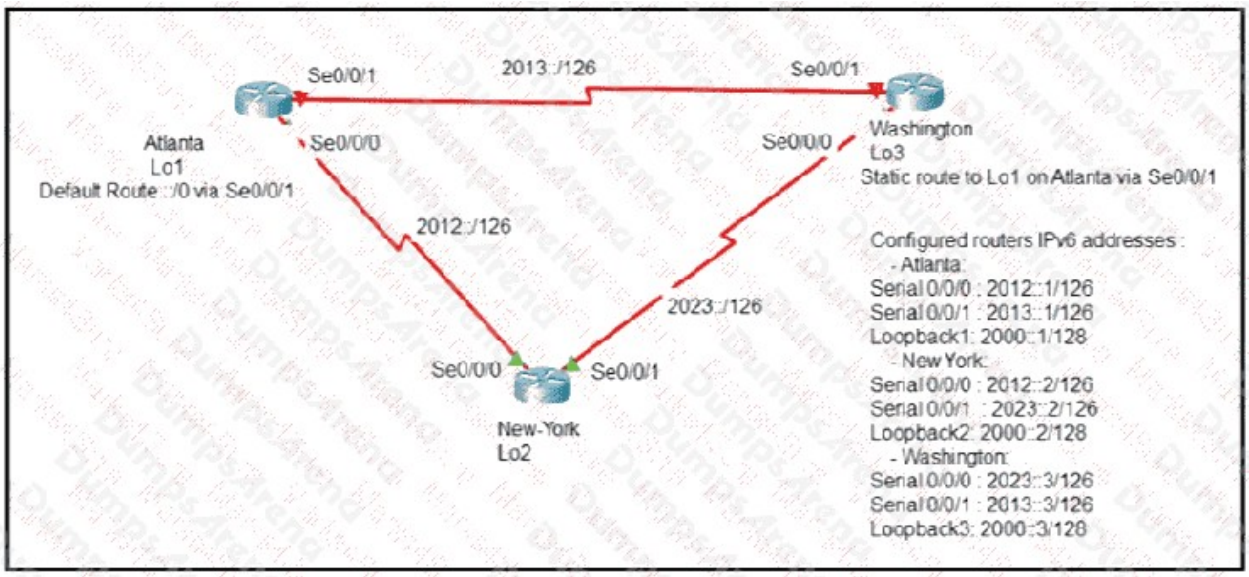
ANSWER: C

Explanation:

It forwards only the VTP advertisements that it receives on its trunk ports. is correct because SW2 is operating in VTP transparent mode. A transparent-mode switch does not synchronize its VLAN database with VTP advertisements received from VTP servers, and it does not advertise its locally created VLANs as VTP database updates. Instead, it acts as a transit device for VTP control-plane information: advertisements received over a trunk link are forwarded through its other trunk links so that participating VTP switches farther along the Layer 2 topology can receive them. This behavior allows a transparent switch to sit between switches that actively participate in the same VTP domain without itself accepting or applying the advertised VLAN database changes. Locally configured VLAN information remains local to the transparent switch and is stored independently from the VTP database learned by server and client modes. VTP advertisements use trunk connections because VTP is a Layer 2 switching protocol used to distribute VLAN configuration information between interconnected switches. Cisco documents that transparent mode forwards received VTP advertisements while not synchronizing its VLAN database with them. See Cisco's Understanding VTP and the Cisco IOS XE VTP configuration guide.

QUESTION NO: 212

Refer to Exhibit.



An engineer is configuring the NEW York router to reach the Lo1 interface of the Atlanta router using interface Se0/0/0 as the primary path. Which two commands must be configured on the New York router so that it can reach the Lo1 interface of the Atlanta router via Washington when the link between New York and Atlanta goes down? (Choose two)

A. ipv6 route 2000::1/128 2012::1

B. ipv6 route 2000::1/128 2012::1 5

C. ipv6 route 2000::1/128 2012::2

D. ipv6 route 2000::1/128 2023::2 5

E. ipv6 route 2000::1/128 2023::3 5

ANSWER: A E

Explanation:

The required configuration uses two IPv6 static routes to Atlanta's loopback prefix, `2000::1/128`. The command `ipv6 route 2000::1/128 2012::1` installs the normal primary static route through the directly connected Atlanta next-hop address on the New York-to-Atlanta serial path. A static route without an explicitly configured administrative distance uses the default static-route administrative distance of 1, so it is preferred while that direct path is available.

The command `ipv6 route 2000::1/128 2023::3 5` provides the alternate route through Washington. Its next-hop address is Washington's address on the New York-to-Washington link. The administrative distance of 5 is higher than the primary route's default distance of 1, making this a floating static route. Therefore, it remains in the routing table only as a less-preferred candidate and is selected when the primary route is removed because the New York-to-Atlanta link fails. Cisco IOS uses the `ipv6 route` global configuration command for IPv6 static routes and permits an administrative distance to control route preference. See Cisco's IPv6 static-route configuration guide and the administrative-distance documentation.

QUESTION NO: 213

What are two advantages of implementing a controller-based architecture instead of a traditional network architecture? (Choose two.)

A. It allows for seamless connectivity to virtual machines.

B. It supports complex and high-scale IP addressing schemes.

C. It enables configuration task automation.

D. It provides increased scalability and management options.

E. It increases security against denial-of-service attacks.

ANSWER: C D

Explanation:

Controller-based architecture separates network policy and management decisions from individual device-by-device configuration. This centralized approach lets administrators define intended network behavior once and have the controller consistently deploy the required configurations to the appropriate infrastructure. Therefore, **It enables configuration task automation.** Controllers expose programmatic interfaces and integrate with automation and orchestration workflows, allowing repeatable provisioning, policy deployment, and configuration changes. This reduces manual effort and improves consistency across the environment.

It provides increased scalability and management options. is also an advantage because the controller maintains a broader view of the network and offers centralized tools for monitoring, policy management, inventory, assurance, and lifecycle operations. Rather than managing each network device independently, operations teams can apply policy and intent across groups of devices, sites, or fabric domains. The centralized management model supports growth by making large numbers of network elements easier to operate consistently while retaining device-level forwarding functions in the infrastructure. Cisco describes software-defined networking as an architecture that centralizes control and enables programmability, and Cisco Catalyst Center provides centralized automation and assurance capabilities for enterprise networks. See Cisco Software-Defined Networking overview and Cisco Catalyst Center.

QUESTION NO: 214

What occurs when overlapping Wi-Fi channels are implemented?

A. Users experience poor wireless network performance.

B. Wireless devices are unable to distinguish between different SSIDs.

C. The wireless network becomes vulnerable to unauthorized access.

D. Network communications are open to eavesdropping.

ANSWER: A

Explanation:

Users experience poor wireless network performance because partially overlapping Wi-Fi channels create adjacent-channel interference. In the 2.4 GHz band, for example, nearby access points should generally use nonoverlapping 20 MHz channels such as 1, 6, and 11 where permitted by the regulatory domain. Channels that overlap transmit energy into each other's spectrum. That energy is treated as interference by receivers, increasing errors and causing corrupted frames. The affected devices must contend for airtime and retransmit frames, reducing usable capacity and increasing delay. The result is often lower throughput, unstable connectivity, and a poor experience for applications that need consistent performance, including voice, video, and interactive services. Sound RF design therefore uses a channel plan that minimizes adjacent-channel interference, while also selecting an appropriate channel width, transmit-power level, and access-point density. Cisco describes the importance of channel planning and nonoverlapping channel use in its Wireless LAN Channel Planning Best Practices. For an overview of the interference mechanisms involved, see Cisco's WLAN Interference overview.

QUESTION NO: 215

What are two reasons for an engineer to configure a floating state route? (Choose two)

A. to automatically route traffic on a secondary path when the primary path goes down

B. to route traffic differently based on the source IP of the packet

C. to enable fallback static routing when the dynamic routing protocol fails

D. to support load balancing via static routing

E. to control the return path of traffic that is sent from the router

ANSWER: A C

Explanation:

A floating static route provides a standby route to a destination. It is configured with an administrative distance that is higher than that of the preferred route source, such as a normal static route or a dynamically learned route. Since Cisco routers select routes with lower administrative distance first, the floating route is not installed as the active route while the preferred route remains available. If that preferred route is withdrawn or becomes unavailable, the floating static route can enter the routing table and forward traffic through the backup next hop.

Therefore, *to automatically route traffic on a secondary path when the primary path goes down* is a valid use case. This supplies deterministic failover without requiring manual intervention. Likewise, *to enable fallback static routing when the dynamic routing protocol fails* is correct because loss of a routing-protocol neighbor, prefix withdrawal, or another protocol failure removes the preferred learned route, allowing the higher-distance static route to become active. Engineers commonly use this design to retain basic reachability for critical prefixes when dynamic route learning is interrupted. Cisco describes floating static routes as backup routes that take effect after routes with lower administrative distance are removed. See Cisco's floating static route documentation and the Cisco IOS XE static-route configuration guide.

QUESTION NO: 216

Refer to the exhibit.

An engineer deploys a topology in which R1 obtains its IP configuration from DHCP. If

the switch and DHCP server configurations are complete and correct. Which two sets of commands must be configured on R1 and R2 to complete the task? (Choose two)

A. R1(config-if)# ip helper-address 198.51.100.100

B. R2(config-if)# ip helper-address 198.51.100.100

C. R1(config-if)# ip address dhcp

D. R2(config-if)# ip address dhcp

E. R1(config-if)# ip helper-address 192.0.2.2

ANSWER: B C

Explanation:

`R1(config-if)# ip address dhcp` is required on R1's interface because it makes that interface a DHCP client. After the interface is enabled, R1 sends DHCP discovery messages to obtain its IPv4 address and associated parameters, such as the subnet mask, default gateway, and possibly DNS-server information, from the DHCP infrastructure. Cisco IOS XE documents the `ip address dhcp` interface command as the mechanism for dynamically acquiring an IPv4 address through DHCP.

`R2(config-if)# ip helper-address 198.51.100.100` is required on R2's interface that receives R1's client broadcast. DHCP discovery traffic is initially broadcast and cannot cross a routed boundary by itself. The helper address makes R2 operate as a DHCP relay agent: it receives the local DHCP broadcast and forwards it as a unicast request to the DHCP server at `198.51.100.100`. This lets a DHCP server located on a different IP network provide an address lease to R1 while retaining the client's originating subnet information for correct pool selection. Cisco describes the relay-agent behavior and helper-address configuration in its DHCP documentation.

References: Cisco IOS DHCP Client Configuration Guide; Cisco IOS DHCP Relay Agent Configuration Guide.

QUESTION NO: 217

An engineer must configure an OSPF neighbor relationship between router R1 and R3. The authentication configuration has been configured and the connecting interfaces are in the same 192.168.1.0/30 subnet. What are the next two steps to complete the configuration? (Choose two.)

A. configure the interfaces as OSPF active on both sides

B. configure both interfaces with the same area ID

C. configure the hello and dead timers to match on both sides

D. configure the same process ID for the router OSPF process

E. configure the same router ID on both routing processes

ANSWER: A B

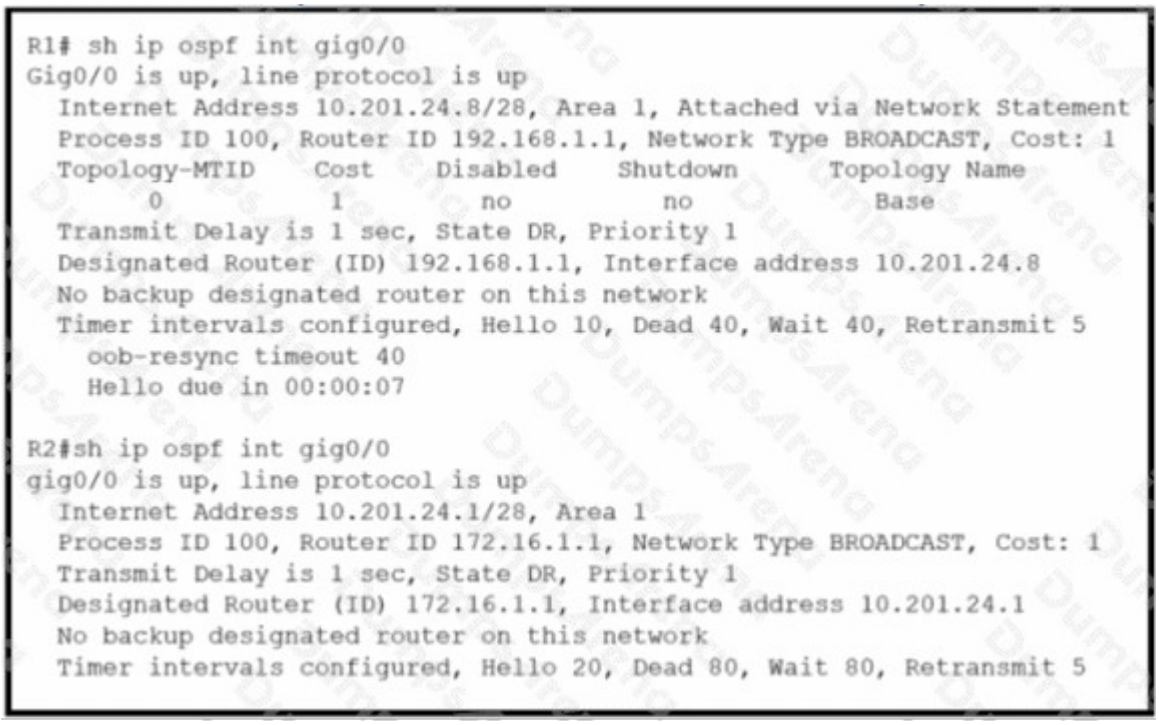
Explanation:

OSPF must be enabled on each interface that connects R1 and R3. This can be done by matching the interface address with a `network` command under the appropriate OSPF routing process, or by enabling OSPF directly under each interface with `ip ospf process-id area area-id`. Once enabled, each router sends and receives OSPF Hello packets over the shared 192.168.1.0/30 link, allowing neighbor discovery to begin.

The shared interfaces must also be assigned to the same OSPF area. The OSPF area ID is carried in Hello packets and is a required adjacency parameter; therefore, both endpoints of this link need a matching area assignment. With IP connectivity already present and authentication configured consistently, enabling OSPF on both endpoints and placing the link in one common area supplies the essential configuration needed to establish the adjacency. OSPF process IDs are locally significant, while each router needs its own unique router ID within the OSPF domain. Cisco's OSPF configuration guide describes enabling OSPF per interface or with network statements, and Cisco's neighbor-adjacency guidance identifies matching area parameters as an adjacency requirement. See Cisco IOS XE OSPF Configuration Guide and Cisco OSPF Neighbor Relationship documentation.

QUESTION NO: 218

Refer to the exhibit.



```
R1# sh ip ospf int gig0/0
Gig0/0 is up, line protocol is up
  Internet Address 10.201.24.8/28, Area 1, Attached via Network Statement
  Process ID 100, Router ID 192.168.1.1, 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) 192.168.1.1, Interface address 10.201.24.8
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:07

R2#sh ip ospf int gig0/0
gig0/0 is up, line protocol is up
  Internet Address 10.201.24.1/28, Area 1
  Process ID 100, Router ID 172.16.1.1, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 172.16.1.1, Interface address 10.201.24.1
  No backup designated router on this network
  Timer intervals configured, Hello 20, Dead 80, Wait 80, Retransmit 5
```

What action establishes the OSPF neighbor relationship without forming an adjacency?

A. modify hello interval

B. modify process ID

C. modify priority

D. modify network type

ANSWER: C

Explanation:

modify priority is correct because OSPF interface priority controls whether a router is eligible to become the designated router or backup designated router on a broadcast or nonbroadcast multiaccess segment. Configuring an interface with an OSPF priority of 0 makes that router ineligible for either election role. On these network types, OSPF does not require every router to form a full database-exchange adjacency with every other router. Instead, routers that are neither the designated router nor backup designated router—often called DROTHER routers—can recognize one another as neighbors and remain in the 2-Way state, while forming full adjacencies with the designated router and backup designated router. Thus, priority manipulation can produce the intended state: hello packets establish bidirectional neighbor visibility, but no full adjacency is formed between the applicable routers. This behavior reduces the number of required adjacencies and link-state database exchanges on a shared multiaccess network. Cisco documents that the DR and BDR establish adjacencies with all routers on the segment, whereas routers that are not DR/BDR do not become adjacent to each other. The priority setting is also the mechanism Cisco IOS uses to determine DR/BDR election eligibility. See Cisco's OSPF neighbor-state documentation and Cisco IOS OSPF configuration guide.

QUESTION NO: 219

An engineer must configure the IPv6 address 2001:0db8:0000:0000:0700:0003:400F:572B on the serial0/0 interface of the HQ router and wants to compress it for easier configuration. Which command must be issued on the router interface?

A. ipv6 address 2001:db8::700:3:400F:572B

B. ipv6 address 2001:db8:0::700:3:4F:572B

C. ipv6 address 2001::db8:0000::700:3:400F:572B

D. ipv6 address 2001:0db8::7:3:4F:572B

ANSWER: A

Explanation:

The correct compressed representation is `ipv6 address 2001:db8::700:3:400F:572B`. IPv6 addresses consist of eight 16-bit hexadecimal fields, called hextets. Textual compression first removes leading zeroes within an individual hextet: `0db8` becomes `db8`, `0700` becomes `700`, and `0003` becomes `3`. It must not remove non-leading zeroes, because doing so would change the numeric value represented by that hextet. The two adjacent all-zero hextets in the original address, `0000:0000`, can then be represented once by `::`. Combining these valid transformations produces `2001:db8::700:3:400F:572B`. Hexadecimal characters are case-insensitive in IPv6 notation, so the uppercase `F` and `B` are valid. RFC 5952 describes the conventional rules for representing IPv6 addresses in compressed text form, including suppressing leading zeroes and using a single double-colon for a contiguous zero sequence. Cisco IOS IPv6 interface configuration uses this IPv6 address notation as part of the `ipv6 address` interface command; an operational configuration also requires the appropriate prefix length for the addressing plan. See RFC 5952: A Recommendation for IPv6 Address Text Representation and Cisco IPv6 Addressing and Configuration.

QUESTION NO: 220

What is a characteristic of cloud-based network topology?

A. wireless connections provide the sole access method to services

B. onsite network services are provided with physical Layer 2 and Layer 3 components

C. services are provided by a public, private, or hybrid deployment

D. physical workstations are configured to share resources

ANSWER: C

Explanation:

Services are provided by a public, private, or hybrid deployment is a defining characteristic of cloud-based networking. Cloud computing delivers shared computing, storage, and network resources as services, with the deployment model indicating ownership, access scope, and the way resources are operated. A public cloud provides services through a provider's infrastructure for broad customer use; a private cloud is provisioned for exclusive use by one organization; and a hybrid cloud combines distinct cloud infrastructures in a way that enables portability or interoperability between them. In a cloud-based topology, connectivity is designed around reaching and integrating these cloud-hosted services, commonly through Internet access, VPNs, or dedicated provider connectivity. This makes the deployment model fundamental to describing where services reside and how an organization consumes them. Cisco identifies public, private, and hybrid cloud as common cloud deployment approaches, while the NIST cloud-computing definition formally defines these deployment models. See Cisco: What Is Cloud Computing? and NIST SP 800-145.

QUESTION NO: 221

What are two requirements for an HSRP group? (Choose two.)

A. exactly one active router

B. one or more standby routers

C. one or more backup virtual routers

D. exactly one standby active router

E. exactly one backup virtual router

F. exactly one standby router

ANSWER: A F

Explanation:

An HSRP group requires exactly one active router and exactly one standby router. The active router owns the HSRP virtual IP and virtual MAC address and forwards packets that hosts send to their configured default gateway. Having a single active router provides a consistent forwarding point for the virtual gateway and avoids duplicate ownership of the same virtual addressing information on the segment.

The standby router is the designated successor for that group. It monitors HSRP hello messages from the active router and assumes the active role when the active router becomes unavailable, allowing hosts to continue using the same virtual default-gateway IP address without requiring host-side changes. HSRP can include additional participating routers, but they are not additional standby routers; they remain in other HSRP states until an election makes one router the single standby router. Cisco's HSRP operation documentation describes the group roles as one active router and one standby router, with other group members able to participate in the election process. See Cisco HSRP Overview and Operation and Cisco IOS XE HSRP Configuration Guide.

QUESTION NO: 222

Which statements describe the routing protocol OSPF? (Choose three.)

A. It supports VLSM.

B. It is used to route between autonomous systems.

C. It confines network instability to one area of the network.

D. It increases routing overhead on the network.

E. It allows extensive control of routing updates.

F. It is simpler to configure than RIP v2.

ANSWER: A C E

Explanation:

OSPF is a classless interior gateway protocol that carries subnet-mask information in its routing information. Therefore, *It supports VLSM.* is correct: routers can advertise and calculate routes for networks using different prefix lengths, allowing efficient address allocation and route design. Cisco identifies OSPF as a link-state IGP and documents its support for classless addressing and hierarchical area-based operation.

It confines network instability to one area of the network. is correct because OSPF uses a hierarchical area design. Link-state advertisements are flooded within their relevant scope, and each area maintains its own link-state database. Area boundaries help contain many topology changes and SPF recalculations, preventing local changes from requiring every router in the OSPF domain to process the same detailed topology information. The backbone area provides the inter-area transit structure.

It allows extensive control of routing updates. is also correct. OSPF provides administrative control through area design, route summarization at area border routers and autonomous system boundary routers, authentication, interface costs, passive interfaces, and specialized area types. OSPF normally sends link-state information when topology changes occur and refreshes LSAs periodically, rather than relying on full periodic routing-table broadcasts. These controls support scalable, policy-aware internal routing. See Cisco's OSPF overview and RFC 2328.

QUESTION NO: 223

What is a requirement when configuring or removing LAG on a WLC?

A. The Incoming and outgoing ports for traffic flow must be specified If LAG Is enabled.

B. The controller must be rebooted after enabling or reconfiguring LAG.

C. The management interface must be reassigned if LAG disabled.

D. Multiple untagged interfaces on the same port must be supported.

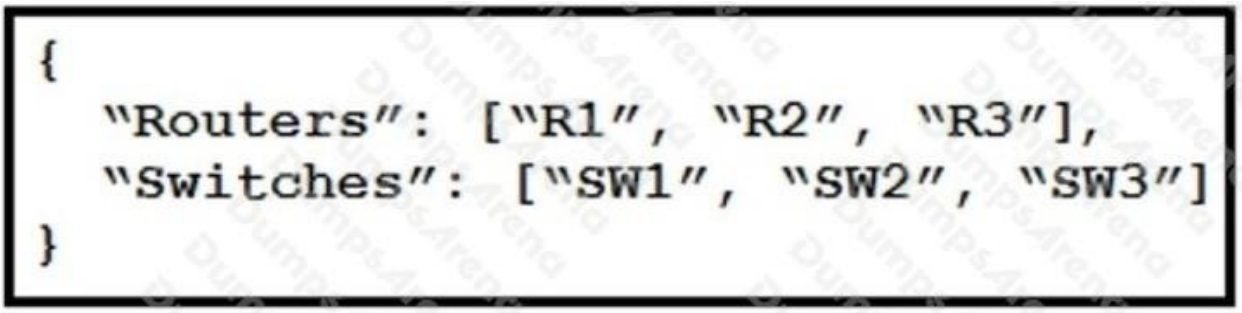
ANSWER: B

Explanation:

The controller must be rebooted after enabling or reconfiguring LAG. On Cisco AireOS wireless LAN controllers, Link Aggregation is a controller-wide setting that changes how the distribution-system Ethernet ports operate: rather than functioning as independently mapped ports, they are treated as members of one logical aggregated connection. Because this setting affects the controller's underlying port operation and interface mapping behavior, it cannot be fully applied dynamically while the controller is running. A reboot is therefore required after enabling LAG, and it is likewise required when disabling it, so that the controller initializes the physical interfaces in the selected mode.

Before making this change, the connected switch should be prepared with a compatible EtherChannel/port-channel configuration and the required VLAN trunking. Planning a maintenance window is important because the reboot interrupts controller operation, including wireless client connectivity and controller management access until the WLC returns to service. Cisco's AireOS configuration documentation identifies LAG as a setting that requires a controller reboot for the configuration change to take effect. See the Cisco Wireless Controller Configuration Guide, Release 8.10 and Cisco's Wireless LAN Controller configuration guide library.

QUESTION NO: 224



Refer to the exhibit. What is represented by `R1` and `SW1` within the JSON output?

A. object

B. value

C. key

D. array

ANSWER: C

Explanation:

`key` is correct. JSON represents an object as a collection of member name/value pairs enclosed in curly braces. In each pair, the member name appears to the left of a colon and identifies the data that follows. In the displayed JSON, `R1` and `SW1` serve as those member names: they label and provide access to their associated data structures or values. In common JSON terminology, such a member name is called a key; it may also be described formally as the name portion of a name/value pair. Keys are strings in JSON syntax, even if an application uses them to represent meaningful identifiers such as router or switch hostnames. Their placement as labels within the object, rather than as data following a colon, establishes that they are keys. This distinction is important when interpreting API output and automation data used in network programmability: software reads a key to locate the associated value, nested object, or list. RFC 8259 defines a JSON object as an unordered collection of zero or more name/value pairs and specifies the syntax for names and values. See RFC 8259, Section 4: Objects and MDN Web Docs: Working with JSON.

QUESTION NO: 225

Which two HTTP methods are suitable for actions performed by REST-based APIs? (Choose two.)

A. REMOVE

B. REDIRECT

C. POST

D. GET

E. PUT

ANSWER: C D E

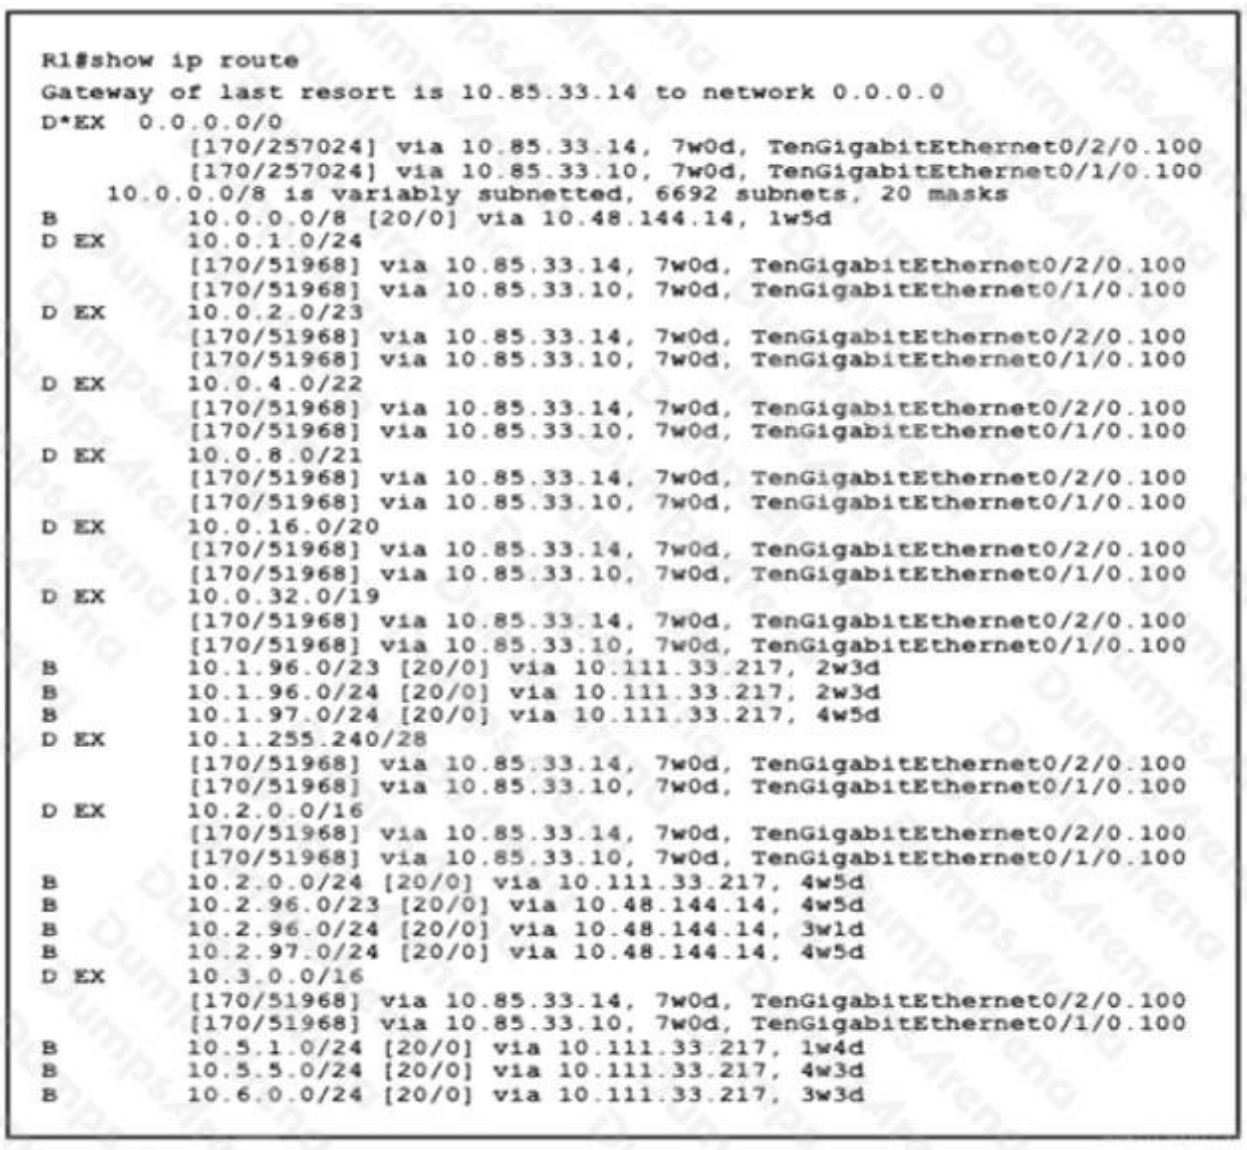
Explanation:

REST-based APIs use standardized HTTP request methods to express the operation a client wants to perform on a resource. **GET** is suitable when the API action is to retrieve a representation of a resource or collection of resources. It is defined as a safe, read-only method, making it the normal choice for API queries. **POST** is suitable when a client submits data for processing, commonly to create a subordinate resource or invoke processing that has server-side effects. **PUT** is also a valid REST API method: it requests that the supplied representation be created or replace the state of the resource identified by the request URI. Consequently, all three listed methods are suitable for REST-based API operations, even though the question wording requests exactly two selections. HTTP method semantics are defined by the IETF HTTP

Semantics specification, while Cisco's CCNA automation objectives include recognizing REST APIs and the HTTP methods used to interact with resources. See RFC 9110, HTTP request methods and Cisco's CCNA exam topics.

QUESTION NO: 226

Refer to the exhibit. Which two statements about the network environment of router R1 must be true? (Choose two.)



```
R1#show ip route
Gateway of last resort is 10.85.33.14 to network 0.0.0.0
D*EX  0.0.0.0/0
          [170/257024] via 10.85.33.14, 7w0d, TenGigabitEthernet0/2/0.100
          [170/257024] via 10.85.33.10, 7w0d, TenGigabitEthernet0/1/0.100
    10.0.0.0/8 is variably subnetted, 6692 subnets, 20 masks
B         10.0.0.0/8 [20/0] via 10.48.144.14, 1w5d
D EX      10.0.1.0/24
          [170/51968] via 10.85.33.14, 7w0d, TenGigabitEthernet0/2/0.100
          [170/51968] via 10.85.33.10, 7w0d, TenGigabitEthernet0/1/0.100
D EX      10.0.2.0/23
          [170/51968] via 10.85.33.14, 7w0d, TenGigabitEthernet0/2/0.100
          [170/51968] via 10.85.33.10, 7w0d, TenGigabitEthernet0/1/0.100
D EX      10.0.4.0/22
          [170/51968] via 10.85.33.14, 7w0d, TenGigabitEthernet0/2/0.100
          [170/51968] via 10.85.33.10, 7w0d, TenGigabitEthernet0/1/0.100
D EX      10.0.8.0/21
          [170/51968] via 10.85.33.14, 7w0d, TenGigabitEthernet0/2/0.100
          [170/51968] via 10.85.33.10, 7w0d, TenGigabitEthernet0/1/0.100
D EX      10.0.16.0/20
          [170/51968] via 10.85.33.14, 7w0d, TenGigabitEthernet0/2/0.100
          [170/51968] via 10.85.33.10, 7w0d, TenGigabitEthernet0/1/0.100
D EX      10.0.32.0/19
          [170/51968] via 10.85.33.14, 7w0d, TenGigabitEthernet0/2/0.100
          [170/51968] via 10.85.33.10, 7w0d, TenGigabitEthernet0/1/0.100
B         10.1.96.0/23 [20/0] via 10.111.33.217, 2w3d
B         10.1.96.0/24 [20/0] via 10.111.33.217, 2w3d
B         10.1.97.0/24 [20/0] via 10.111.33.217, 4w5d
D EX      10.1.255.240/28
          [170/51968] via 10.85.33.14, 7w0d, TenGigabitEthernet0/2/0.100
          [170/51968] via 10.85.33.10, 7w0d, TenGigabitEthernet0/1/0.100
D EX      10.2.0.0/16
          [170/51968] via 10.85.33.14, 7w0d, TenGigabitEthernet0/2/0.100
          [170/51968] via 10.85.33.10, 7w0d, TenGigabitEthernet0/1/0.100
B         10.2.0.0/24 [20/0] via 10.111.33.217, 4w5d
B         10.2.96.0/23 [20/0] via 10.48.144.14, 4w5d
B         10.2.96.0/24 [20/0] via 10.48.144.14, 3w1d
B         10.2.97.0/24 [20/0] via 10.48.144.14, 4w5d
D EX      10.3.0.0/16
          [170/51968] via 10.85.33.14, 7w0d, TenGigabitEthernet0/2/0.100
          [170/51968] via 10.85.33.10, 7w0d, TenGigabitEthernet0/1/0.100
B         10.5.1.0/24 [20/0] via 10.111.33.217, 1w4d
B         10.5.5.0/24 [20/0] via 10.111.33.217, 4w3d
B         10.6.0.0/24 [20/0] via 10.111.33.217, 3w3d
```

A. The EIGRP administrative distance was manually changed from 90 to 170.

B. There are 20 different network masks within the 10.0.0.0/8 network.

C. Ten routes are equally load-balanced between Te0/1/0.100 and Te0/2/0.100.

D. The 10.0.0.0/8 network was learned via external EIGRP.

E. A static default route to 10.85.33.14 was defined.

ANSWER: B C

Explanation:

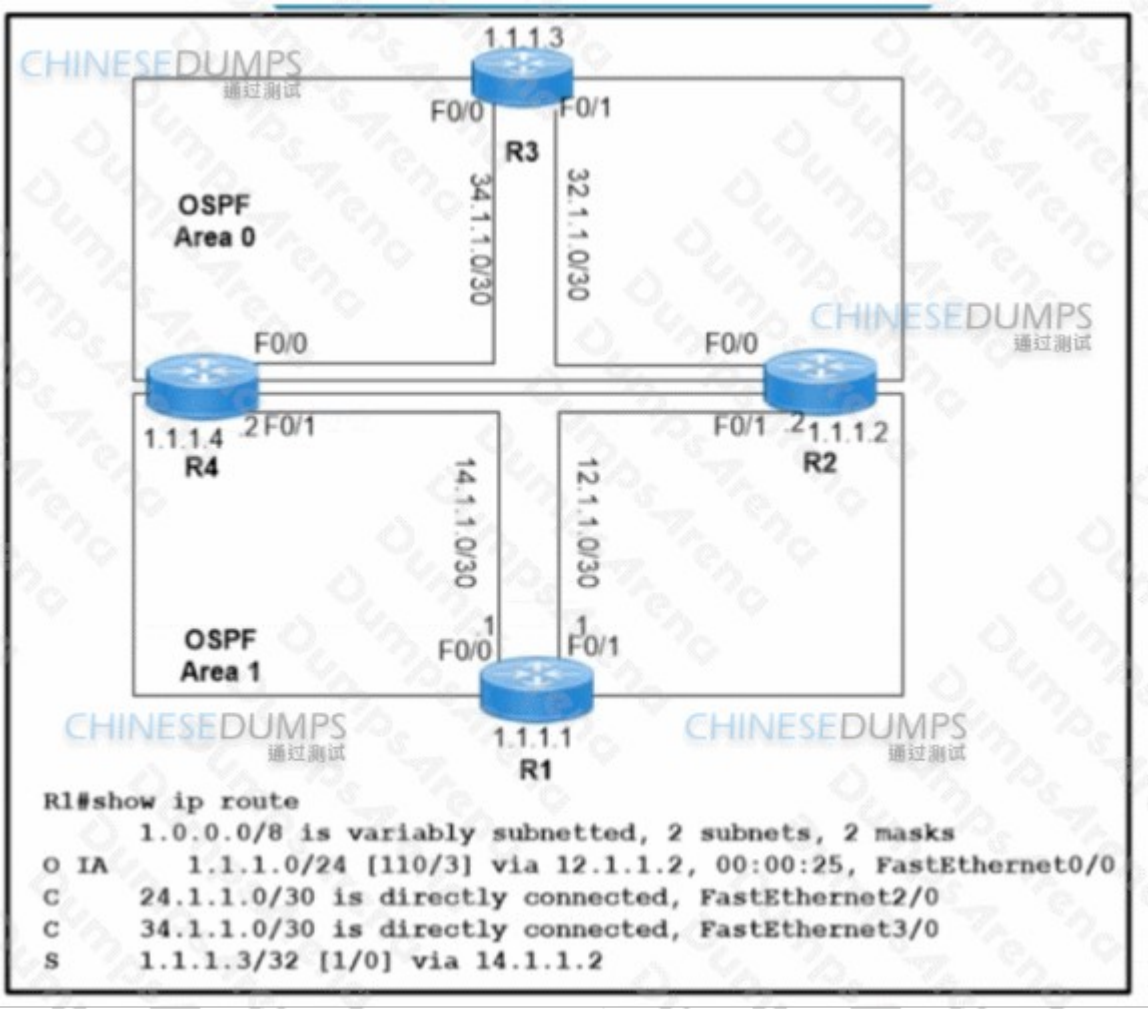
The statement *"There are 20 different network masks within the 10.0.0.0/8 network."* follows directly from the Cisco IOS routing-table summary for the 10.0.0.0 major network. IOS displays a "variably subnetted" line when routes beneath a classful major network use variable-length subnet masks. The values shown on that line identify both the number of subnet

routes and the number of masks represented in that major network. Therefore, a display of “20 subnets, 20 masks” establishes that 20 masks are present under 10.0.0.0/8.

The statement *“Ten routes are equally load-balanced between Te0/1/0.100 and Te0/2/0.100.”* is also established by the route entries in the exhibit. Cisco IOS installs equal-cost multipath routes when multiple paths to the same destination prefix have the same route source, administrative distance, and metric. The routing table then lists each equal next hop/interface for that prefix. The repeated equal-cost entries through Te0/1/0.100 and Te0/2/0.100 show that R1 has installed two equal paths for each of the ten indicated routes, allowing traffic to be load-balanced across those interfaces. Cisco documents routing-table route selection and equal-cost routing behavior at Cisco EIGRP documentation and IOS routing-table output conventions at Cisco’s routing table guide.

QUESTION NO: 227

Refer to the exhibit.



Which two values does router R1 use to identify valid routes for the R3 loopback address 1.1.1.3/32? (Choose two.)

A. lowest cost to teach the next hop

B. highest metric

C. highest administrative distance

D. lowest metric

E. lowest administrative distance

ANSWER: D E

Explanation:

Router R1 evaluates candidate routes to the same destination prefix by using administrative distance and the routing protocol metric. **Lowest administrative distance** identifies the most trusted routing information source when R1 has routes to 1.1.1.3/32 that originated from different sources, such as a static route and a dynamically learned route. Administrative distance is a locally significant Cisco value that expresses the relative trustworthiness of the source; the route source with the lower value is preferred for installation in the routing table.

After routes are compared within the same routing protocol or route source, **lowest metric** identifies the preferable path. A metric is the protocol-specific measurement of path quality. For example, OSPF uses cost, while EIGRP uses a composite metric. R1 selects the path with the numerically lowest metric among equivalent candidate routes from that protocol, and it can install equal-cost paths when their metrics are equal and equal-cost multipath is supported and enabled. Thus, administrative distance determines the preferred source of route information, and the metric determines the preferred path offered by that source. Cisco documents the default administrative-distance behavior at Routing Protocol Administrative Distances and describes OSPF's cost-based path metric at OSPF Route Selection.

QUESTION NO: 228

Which two principles must be considered when using per-hop behavior in QoS? (Choose two.)

A. Policing is not supported on subinterfaces.

B. Shaping and rate limiting have the same effect.

C. Shaping drops excessive traffic without adding traffic delay.

D. Shaping levels out traffic bursts by delaying excess traffic.

E. Policing is performed in the inbound and outbound directions.

ANSWER: D E

Explanation:

Shaping levels out traffic bursts by delaying excess traffic. A shaper measures traffic against a configured rate and places packets that exceed the current transmission allowance into a queue, transmitting them later as tokens become available. This buffering behavior smooths bursts and produces a more consistent output rate, which is particularly useful when sending traffic into a service with a contractual bandwidth limit. Because queued packets wait before transmission, shaping can introduce latency and, if the queue fills, may ultimately drop packets.

Policing is performed in the inbound and outbound directions. Cisco QoS policy maps can apply policing actions to traffic entering an interface or leaving it. A policer monitors the traffic rate and enforces the configured rate by taking an action on traffic that exceeds the allowance, such as dropping or remarking packets. Applying policing at either traffic direction allows administrators to enforce policy at the appropriate network boundary and classify traffic consistently with the intended QoS design. Cisco's QoS documentation describes both policing and shaping as traffic-conditioning mechanisms and documents the directions in which policing may be configured. See Cisco IOS XE QoS policing and shaping overview and Cisco IOS XE policing configuration guide.

QUESTION NO: 229

Which two minimum parameters must be configured on an active interface to enable OSPFV2 to operate? (Choose two.)

A. OSPF process ID

B. OSPF MD5 authentication key

C. OSPF stub flag

D. IPv6 address

E. OSPF area

ANSWER: A E

Explanation:

OSPF process ID and OSPF area are the two required configuration values for enabling OSPFv2 on an IPv4 interface. The process ID identifies the local OSPF routing process on the router. It is locally significant, so it does not have to match the process identifier used by neighboring routers. The area identifies the OSPF area to which the interface belongs and is essential because OSPF organizes topology information, neighbor relationships, and link-state databases by area. An interface participating in OSPF must therefore be associated with both an OSPF process and an area.

On Cisco IOS, this association can be made globally with a `network` command under the OSPF routing process, which includes matching interfaces in a specified area. It can also be configured directly on the interface with `ip ospf process-id area area-id`. In either configuration method, the process identifier and area identifier provide the minimum OSPF-specific values needed to activate OSPFv2 on an operational IPv4 interface. Cisco documents both the routing-process configuration and assignment of interfaces to OSPF areas in its OSPF Configuration Example and Command Overview and IOS XE OSPF Configuration Guide.

QUESTION NO: 230

Which 802.11 management frame type is sent when a client roams between access points on the same SSID?

A. Reassociation Request

B. Authentication Request

C. Association Request

D. Probe Request

ANSWER: A

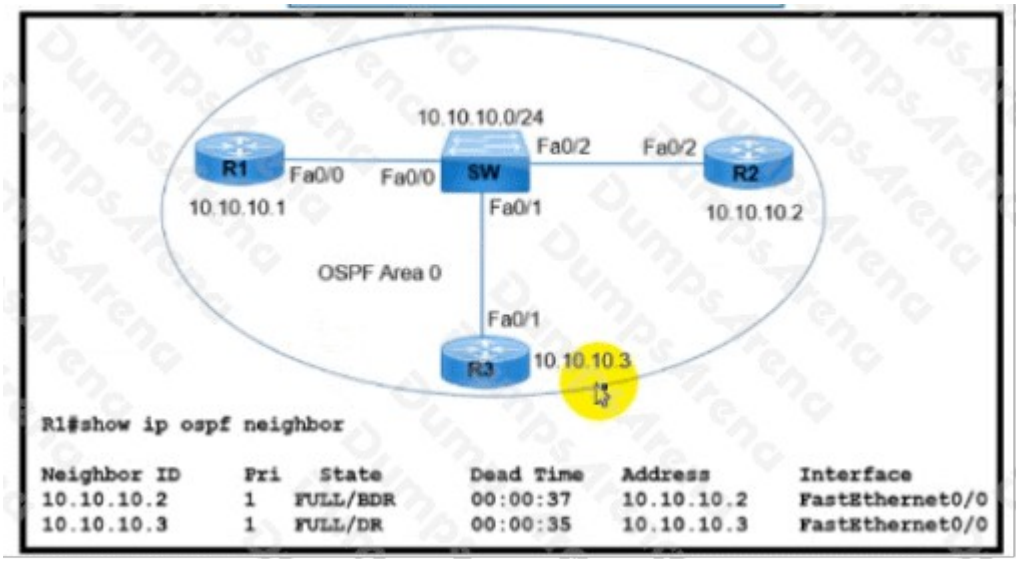
Explanation:

Reassociation Request is the 802.11 management frame a wireless client sends to the target access point when it moves, or roams, within the same extended service set (ESS), such as between access points advertising the same SSID. Reassociation enables the new access point and the distribution system to recognize that the station was previously associated elsewhere and to update the station's attachment point. The request contains information identifying the client's current/previous access point, which supports continuity of service and handling of traffic that may have been buffered during the transition.

This frame is specifically defined for changing an existing association from one access point to another. After receiving the request, the target access point returns a Reassociation Response indicating whether it accepts the client. In enterprise WLANs, the exact authentication and key-management activity around the roam can vary based on the security configuration and roaming enhancements, but reassociation remains the management-frame mechanism that establishes the station's association with the new access point during a roam. Cisco's roaming documentation describes this client movement and reassociation process, while Cisco's wireless fundamentals material identifies reassociation as an IEEE 802.11 management-frame function. See Cisco Wireless Roaming and Cisco 802.11 Design Guide.

QUESTION NO: 231

Refer to the exhibit.



R1 has taken the DROTHER role in the OSPF DR/BDR election process. Which configuration must an engineer implement so that R1 is elected as the DR?



A. Option A

B. Configure a higher OSPF interface priority on R1's multiaccess interface.

C. Option C

D. Option D

ANSWER: B

Explanation:

Configure a higher OSPF interface priority on R1’s interface connected to the shared broadcast or NBMA segment, using the interface-level `ip ospf priority` command. OSPF chooses the designated router by comparing interface priorities first: the eligible router with the highest priority becomes the DR, and the next-highest becomes the BDR. The priority range is 0 through 255; a value of 0 makes a router ineligible for DR or BDR election. If priorities tie, OSPF uses the highest router ID as the tie-breaker. Therefore, assigning R1 a priority higher than every other eligible router on that segment ensures it wins a subsequent election.

DR/BDR elections are not preemptive. If a DR and BDR are already established, increasing R1’s priority alone does not displace the current DR. An engineer must also cause a new election after applying the priority, such as by resetting the OSPF adjacency or cycling the participating interface during an appropriate maintenance window. Cisco documents the election criteria and nonpreemptive behavior in its OSPF Designated Router Election guide. The interface command syntax and priority behavior are documented in the Cisco IOS OSPF command reference.

QUESTION NO: 232

An organization secures its network with multi-factor authentication using an authenticator app on employee smartphone. How is the application secured in the case of a user’s smartphone being lost or stolen?

A. The application requires an administrator password to reactivate after a configured Interval.

B. The application requires the user to enter a PIN before it provides the second factor.

C. The application challenges a user by requiring an administrator password to reactivate when the smartphone is rebooted.

D. The application verifies that the user is in a specific location before it provides the second factor.

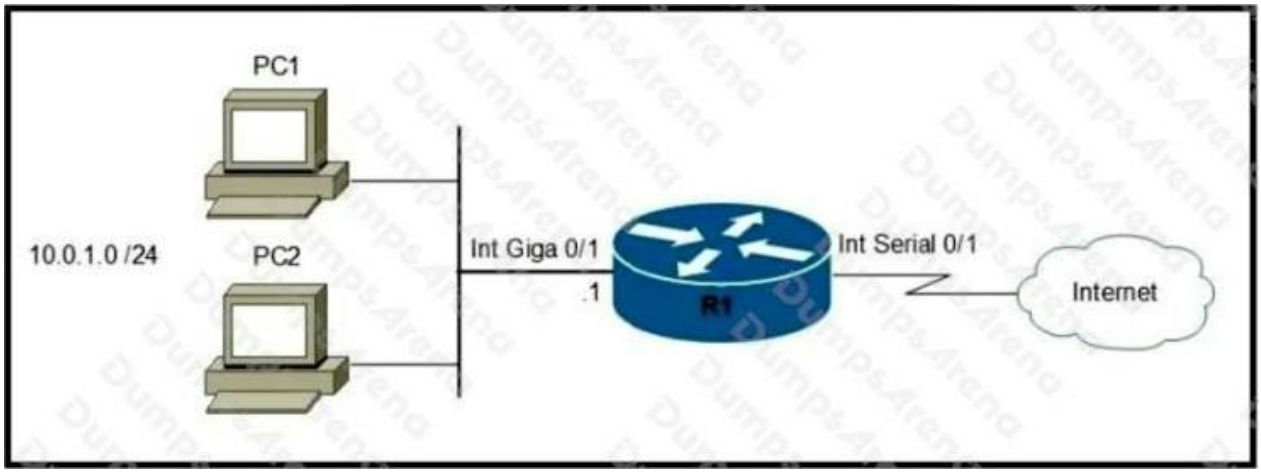
ANSWER: B

Explanation:

The application requires the user to enter a PIN before it provides the second factor. This adds a local verification step to the possession factor represented by the enrolled smartphone. A person who merely obtains a lost or stolen device should not be able to open the authenticator application and immediately retrieve time-based one-time passcodes or approve an authentication prompt. Requiring a PIN, or an equivalent local device/app unlock method such as biometrics, helps ensure that the authenticator remains protected even when the physical phone is no longer under the employee's control.

This approach reflects the core MFA principle that factors should be protected independently. The smartphone provides something the user has, while the PIN provides something the user knows. In practice, the organization should also have procedures to revoke the lost device's registered authenticator and enroll a replacement device promptly. Microsoft documents that Microsoft Authenticator can require app lock before authentication approvals or code access, and NIST guidance emphasizes protecting authentication secrets and authenticators against theft and unauthorized use. See Microsoft Entra documentation for Microsoft Authenticator and NIST SP 800-63B Digital Identity Guidelines.

QUESTION NO: 233



Refer to the exhibit. Which two commands must be configured on router R1 to enable the router to accept secure remote-access connections? (Choose two.)

A. ip ssh pubkey-chain

B. username cisco password 0 cisco

C. crypto key generate rsa

D. transport input telnet

E. login console

ANSWER: B C

Explanation:

Secure remote CLI access to an IOS router is normally provided with SSH. The command `crypto key generate rsa` generates the RSA key pair used by the router's SSH server to identify itself and to establish the cryptographic protection

required for an SSH session. Cisco IOS requires an RSA key pair before SSH can operate; the router hostname and IP domain name must already be configured before this command can successfully generate the key.

The command `username cisco password 0 cisco` creates a local user credential. When the VTY configuration uses local authentication, such as `login local`, this credential is used to authenticate the inbound SSH user. Together, an RSA key pair and a local username provide the essential SSH encryption/identity and user-authentication components for the secure remote-access configuration shown in the exhibit. Cisco's SSH configuration guidance demonstrates generating RSA keys and configuring local credentials for SSH access. See Cisco: Configuring Secure Shell on Routers and Switches and the Cisco IOS XE SSH Configuration Guide.

QUESTION NO: 234

Which two ICMPv6 messages are used by Stateless Address Autoconfiguration? (Choose two.)

A. router solicitation

B. router advertisement

C. echo request

D. redirect

E. destination unreachable

ANSWER: A B

Explanation:

Router solicitation messages allow an IPv6 host to request configuration information from local routers without waiting for the next periodic advertisement. A newly connected host commonly sends a router solicitation to the all-routers multicast address.

Router advertisement messages provide the information required for Stateless Address Autoconfiguration, including the IPv6 prefix, prefix length, default gateway information, and address-configuration flags. A host can combine the advertised prefix with an interface identifier to create its IPv6 address. RFC 4861 defines these Neighbor Discovery messages, and RFC 4862 defines Stateless Address Autoconfiguration. See RFC 4861 and RFC 4862.

QUESTION NO: 235

Which two outcomes are predictable behaviors for HSRP? (Choose two)

A. The two routers share a virtual IP address that is used as the default gateway for devices on the LAN.

B. The two routers negotiate one router as the active router and the other as the standby router

C. D The two routers synchronize configurations to provide consistent packet forwarding

D. The two routed share the same IP address, and default gateway traffic is load-balanced between them

ANSWER: A B

Explanation:

HSRP provides first-hop gateway redundancy by presenting hosts on a LAN with a single virtual default gateway. The routers participating in an HSRP group use a shared virtual IP address, which hosts configure as their default gateway. This lets endpoint configuration remain unchanged even if the currently forwarding router fails. HSRP also elects routers into distinct operational roles: one router becomes the active router and forwards traffic sent to the virtual gateway, while another becomes the standby router and monitors the active router's state. If the active router becomes unavailable, the standby router transitions to active status and takes over forwarding for the virtual IP and virtual MAC address. This predictable failover behavior is the core purpose of HSRP and helps preserve LAN hosts' access to off-subnet destinations without

requiring those hosts to select a new gateway. Cisco documents that the active router performs packet forwarding for the virtual router and that the standby router assumes this responsibility after an active-router failure. See Cisco's HSRP configuration guide and the protocol specification in RFC 2281.

QUESTION NO: 236

What are two reasons lo configure PortFast on a switch port attached to an end host? (Choose two.)

A. to enable the number of MAC addresses learned on the port to 1

B. to protect the operation of the port from topology change processes

C. to enable the port to enter the forwarding state immediately when the host boots up

D. to prevent the port from participating in Spanning Tree Protocol operations

E. to block another switch or host from communicating through the port

ANSWER: B C

Explanation:

PortFast is intended for an access-layer switch port that connects directly to a single end host, such as a workstation, server, printer, or IP phone. Its primary benefit is that it allows the port to transition directly to the forwarding state when the physical link comes up, rather than waiting through the normal Spanning Tree Protocol listening and learning states. This lets the attached host obtain network connectivity promptly during startup, which is especially useful for services that begin network communication immediately after booting.

PortFast also treats the interface as an edge port for topology-change behavior. Link transitions on a PortFast-enabled access port do not trigger the usual STP topology-change notifications, avoiding unnecessary forwarding-table aging and STP reconvergence effects elsewhere in the switched network when an end device connects, disconnects, or reboots. This protects normal port operation from topology-change processes caused by routine endpoint link events. Cisco recommends enabling PortFast only where a switch or bridge will not be connected; BPDU Guard can additionally shut down the port if BPDUs are received. See Cisco: Using PortFast and Other Command Extensions to Enhance STP and Cisco IOS XE STP Configuration Guide.

QUESTION NO: 237

R1 as an NTP server must have:

- NTP authentication enabled
- NTP packets sourced from Interface loopback 0
- NTP stratum 2
- NTP packets only permitted to client IP 209.165.200.225

How should R1 be configured?

A. ntp authenticate ntp authentication-key 2 sha1 CISCO123 ntp source Loopback0 ntp access-group server-only 10 ntp master 2
!
access-list 10 permit udp host 209.165.200.225 any eq 123

B. ntp authenticate ntp authentication-key 2 md5 CISCO123 ntp interface Loopback0 ntp access-group server-only 10 ntp stratum 2
!
access-list 10 permit 209.165.200.225

C. ntp authenticate
ntp authentication-key 2 md5 CISCO123

ntp trusted-key 2
ntp source Loopback0
ntp access-group serve-only 10
ntp master 2
!
access-list 10 permit 209.165.200.225

D. ntp authenticate ntp authentication-key 2 md5 CISCO123 ntp source Loopback0 ntp access-group server-only 10 ntp stratum 2
!
access-list 10 permit udp host 209.165.200.225 any eq 123

ANSWER: C

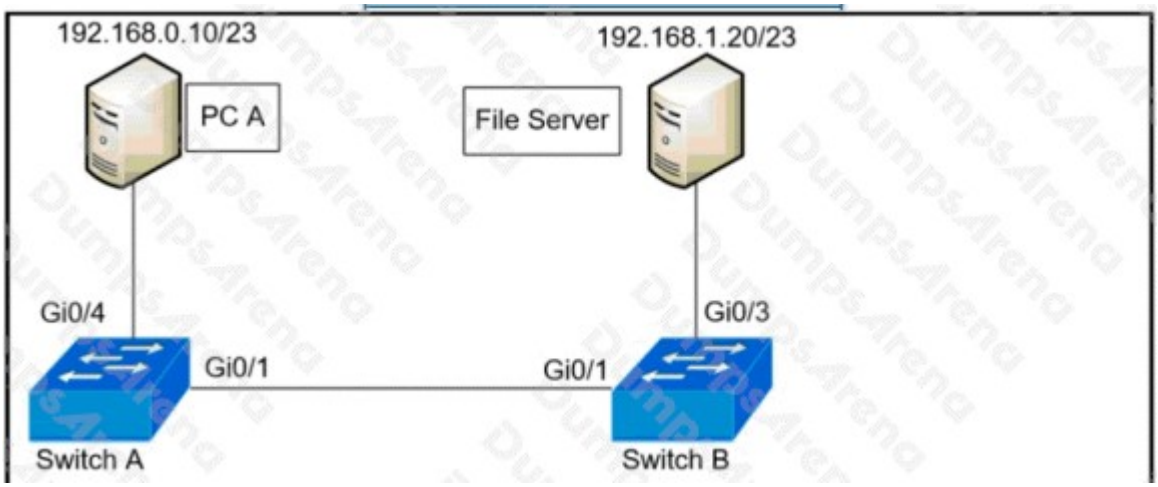
Explanation:

The configuration using `ntp authenticate`, an MD5 authentication key, `ntp source Loopback0`, `ntp master 2`, and `ntp access-group serve-only 10` meets all stated requirements. `ntp authenticate` globally enables NTP authentication, while `ntp authentication-key 2 md5 CISCO123` defines the key material and `ntp trusted-key 2` marks that key as acceptable for NTP authentication. This provides a complete local authentication-key configuration.

`ntp source Loopback0` causes NTP packets generated by R1 to use the Loopback0 address as their source, providing a stable source address independent of physical-interface state. `ntp master 2` makes R1 operate as an NTP master clock at stratum 2 when it is not synchronized to an upstream source. The `serve-only` NTP access group permits matching hosts to receive NTP service without granting peer association access. Standard ACL 10 matches only source address `209.165.200.225`, so only that client is permitted to synchronize with R1. Cisco documents the NTP master, source, authentication, and access-group commands in its IOS NTP Command Reference and provides authentication configuration guidance in its NTP Configuration Example.

QUESTION NO: 238

Refer to the exhibit.



Switch A
Vlan 10,11,12,13

```
interface GigabitEthernet0/1
 switchport mode trunk
 switchport trunk allowed vlan 10-12
!
interface GigabitEthernet0/4
 switchport access vlan 13
 switchport mode access
```

Switch B
Vlan 10,11,12,13

```
interface GigabitEthernet0/1
switchport mode trunk
!
interface GigabitEthernet0/3
 switchport access vlan 13
 switchport mode access
```

A network administrator assumes a task to complete the connectivity between PC A and the File Server. Switch A and Switch B have been partially configured with VLAN 10, 11, 12, and 13. What is the next step in the configuration?

A. Add PC A to VLAN 10 and the File Server to VLAN 11 fa VLAN segmentation

B. Add VLAN 13 to the trunk links on Switch A and Switch B for VLAN propagation

C. Add a router on a stick between Switch A and Switch B allowing for Inter-VLAN routing.

D. Add PC A to the same subnet as the Fie Server allowing for intra-VLAN communication.

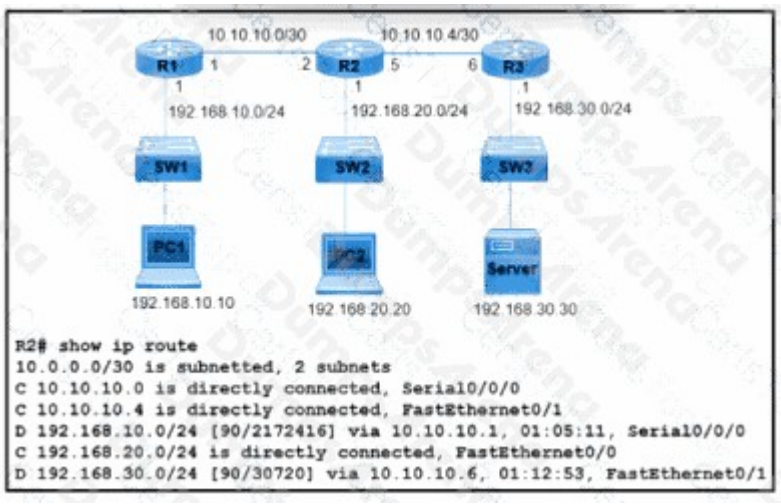
ANSWER: B

Explanation:

Add VLAN 13 to the trunk links on Switch A and Switch B for VLAN propagation. For two endpoints assigned to VLAN 13 but connected through separate switches, the interswitch connection must carry VLAN 13 traffic. An IEEE 802.1Q trunk transports traffic for multiple VLANs by tagging frames with the VLAN identifier. However, a trunk can be restricted to an allowed VLAN list. If VLAN 13 is absent from that list on either end of the connection, VLAN 13 frames cannot cross between the switches, even when VLAN 13 has been created and the endpoint access ports have been assigned correctly.

Adding VLAN 13 to the allowed VLAN set on the relevant trunk interfaces enables Layer 2 forwarding for that VLAN from PC A's switch to the File Server's switch. The VLAN must be permitted consistently on both trunk endpoints, and the trunk must be operational, so that the switches can preserve the VLAN 13 association while forwarding the frames. This completes same-VLAN connectivity without requiring Layer 3 routing, because devices in the same VLAN communicate within one Layer 2 broadcast domain. Cisco documents 802.1Q VLAN tagging and trunk operation at Cisco 802.1Q Trunking and VLAN Tagging and provides VLAN trunk configuration guidance at Cisco VLAN Trunking Configuration.

QUESTION NO: 239



Refer to the exhibit. What is the next-hop IP addressss for R2 so that PC2 reaches the application server via EIGRP?

A. 192.168.20.1

B. 10.10.10.5

C. 192.168.30.1

D. 10.10.10.6

ANSWER: D

Explanation:

10.10.10.6 is the correct next-hop address because it is the IP address of R2's directly connected EIGRP neighbor in the direction of the application-server network. EIGRP routers form neighbor adjacencies on connected interfaces and exchange reachability information with those neighbors. When R2 selects the EIGRP route for the application server's destination prefix, its routing-table entry identifies the adjacent router that should receive packets as the next hop. R2 therefore encapsulates packets for the application server in a Layer 2 frame addressed to the neighbor reachable at 10.10.10.6, while retaining the application server's IP address as the Layer 3 destination. This is the essential distinction between a final destination address and a routing next hop: the destination remains the remote host, but the immediate forwarding target is the neighboring router that provides the selected EIGRP path. Cisco documents that EIGRP learns routes from neighbors and uses its topology and routing information to select forwarding paths. See Cisco's EIGRP Configuration Guide and its EIGRP overview.