

DUMPS ARENA

Cisco Certified Network Associate

Cisco 200-301

Version Demo

Total Demo Questions: 168

Total Premium Questions: 1685

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	204
Topic 2, Network Access	419
Topic 3, IP Connectivity	443
Topic 4, IP Services	174
Topic 5, Security Fundamentals	231
Topic 6, Automation and Programmability	179
Topic 7, Mix Questions	35
Total	1685

QUESTION NO: 1

What are two roles of the Dynamic Host Configuration Protocol (DHCP)? (Choose two.)

- A. The DHCP server assigns IP addresses without requiring the client to renew them.
- B. The DHCP server leases client IP addresses dynamically.
- C. The DHCP client is able to request up to four DNS server addresses.
- D. The DHCP server offers the ability to exclude specific IP addresses from a pool of IP addresses.
- E. The DHCP client maintains a pool of IP addresses it is able to assign.

ANSWER: B D

Explanation:


DHCP's main job is to automatically provide hosts with IPv4 configuration information, most importantly an IP address that is *leased* for a period of time. That makes option B correct: DHCP dynamically leases addresses, and clients must renew the lease periodically (T1/T2 timers) or the address can be reclaimed and reused.

Option D is also correct in the context of how DHCP is commonly implemented on Cisco devices: administrators can prevent certain addresses from being handed out by excluding them from the pool (for example, reserving addresses for servers, printers, or statically addressed network gear). While "exclude" is a configuration feature rather than a wire-protocol message, it is a standard operational role of DHCP services in enterprise networks.

Option A is incorrect because DHCP does require renewal; leases are not permanent by default. Option C is incorrect because DHCP can provide DNS server information, but the "up to four DNS server addresses" statement is not a general DHCP role/requirement and is not universally true across implementations/options. Option E is incorrect because clients do not maintain a pool to assign to others; the server (or relay/server infrastructure) manages address pools.

References: [Cisco DHCP Overview and Configuration](#), [RFC 2131 \(DHCP\)](#)

QUESTION NO: 2



Refer to the exhibit. An engineer must translate the PC1 IP address to 10.199.77.100 and permit PC1 to ping the loopback 0 on router R2. What command set must be used?

- A)

```
○ R1#  
!  
interface Loopback0  
ip address 10.1.1.1 255.255.255.255  
!  
interface FastEthernet0/0  
ip address 10.139.91.1 255.255.255.252  
ip nat inside  
ip virtual-reassembly in  
!  
interface FastEthernet1/0  
ip address 172.16.29.1 255.255.255.0  
ip nat outside  
ip virtual-reassembly in  
!  
router eigrp 100  
network 10.1.1.1 0.0.0.0  
network 10.139.91.0 0.0.0.3  
!  
ip nat inside source static 10.199.77.100 172.16.29.78  
  
R2#  
ip route 10.199.77.100 255.255.255.255 10.139.91.1
```

B)

```
○ R1#  
!  
interface Loopback0  
ip address 10.1.1.1 255.255.255.255  
!  
interface FastEthernet0/0  
ip address 10.139.91.1 255.255.255.252  
ip nat outside  
ip virtual-reassembly in  
!  
interface FastEthernet1/0  
ip address 172.16.29.1 255.255.255.0  
ip nat inside  
ip virtual-reassembly in  
!  
router eigrp 100  
network 10.1.1.1 0.0.0.0  
network 10.139.91.0 0.0.0.3  
!  
ip nat inside source static 172.16.29.78 10.199.77.100  
R2#  
ip route 10.199.77.100 255.255.255.255 10.139.91.1
```

C)

```
R1#  
!  
interface Loopback0  
ip address 10.1.1.1 255.255.255.255  
!  
interface FastEthernet0/0  
ip address 10.139.91.1 255.255.255.252  
ip nat outside  
ip virtual-reassembly in  
!  
interface FastEthernet1/0  
ip address 172.16.29.1 255.255.255.0  
ip nat inside  
ip virtual-reassembly in  
!  
router eigrp 100  
network 10.1.1.1 0.0.0.0  
network 10.139.91.0 0.0.0.3  
!  
ip nat inside source static 172.16.29.78 10.199.77.100  
  
R2#  
ip route 172.16.29.78 255.255.255.255 10.139.91.1
```

D)

```
R1#
!
interface Loopback0
ip address 10.1.1.1 255.255.255.255
!
interface FastEthernet0/0
ip address 10.139.91.1 255.255.255.252
ip nat outside
ip virtual-reassembly in
!
interface FastEthernet1/0
ip address 172.16.29.1 255.255.255.0
ip nat inside
ip virtual-reassembly in
!
router eigrp 100
network 10.1.1.1 0.0.0.0
network 10.139.91.0 0.0.0.3
!
ip nat inside source static 172.16.29.78 10.199.77.100
R2#
ip route 172.16.29.78 255.255.255.255 10.139.91.1
```

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

ANSWER: B


Explanation:

Option **B** is the correct command set because it combines (1) a *static one-to-one NAT* translation for PC1 to the required inside-global address 10.199.77.100 and (2) an ACL that permits only the desired ICMP echo traffic from the translated source toward R2's Loopback0. To make this work, the router performing NAT must mark the LAN-facing interface as `ip nat inside` and the WAN-facing interface as `ip nat outside`, then apply a static mapping such as `ip nat inside source static <PC1-inside-local> 10.199.77.100`. With that in place, when PC1 pings R2's Lo0, the packet's source is rewritten to 10.199.77.100, matching the requirement. The accompanying ACL in option B correctly permits ICMP (echo) from that translated address to the Lo0 destination, while other options either use the wrong NAT type (PAT/dynamic), translate the wrong direction/address, or place the ACL/NAT on incorrect interfaces, which would prevent the ping from succeeding as specified.

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

QUESTION NO: 3

Refer to the exhibit.



Switch A is newly configured. All VLANs are present in the VLAN database. The IP phone and PC A on Gi0/1 must be configured for the appropriate VLANs to establish connectivity between the PCs. Which command set fulfills the requirement?

A)

```
SwitchA(config-if)#switchport mode access
SwitchA(config-if)#switchport access vlan 50
SwitchA(config-if)#switchport voice vlan 51
```

B)

```
SwitchA(config-if)#switchport mode access
SwitchA(config-if)#switchport access vlan 50
SwitchA(config-if)#switchport voice vlan untagged
```

C)

```
SwitchA(config-if)#switchport mode trunk
SwitchA(config-if)#switchport trunk allowed vlan add 50, 51
SwitchA(config-if)#switchport voice vlan dot1p
```

D)

```
SwitchA(config-if)#switchport mode trunk
SwitchA(config-if)#switchport trunk allowed vlan 50, 51
SwitchA(config-if)#mls qos trust cos
```

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

ANSWER: A

Explanation:

The correct configuration for a switchport that connects to an IP phone with a downstream PC is to use an access VLAN for the PC's untagged traffic and a separate voice VLAN for the phone's tagged voice traffic. On Cisco switches, that is done with `switchport mode access`, `switchport access vlan X` (data VLAN), and `switchport voice vlan Y` (voice VLAN). This allows the phone to tag voice frames with the voice VLAN while the attached PC sends untagged frames that the switch places into the access VLAN. That is exactly what Option A represents, so it fulfills the requirement to place PC A and the phone into the appropriate VLANs and restore end-to-end connectivity between PCs in the same data VLAN.

The other options are incorrect because they typically misuse trunking (e.g., forcing the port into trunk mode), omit the voice VLAN command, or attempt to configure multiple data VLANs on a single access port. A phone+PC port should not be configured as a trunk in normal CCNA campus designs unless you have a specific requirement and matching endpoint configuration; the standard best practice is access + voice VLAN.

References: [Cisco – Voice VLAN configuration overview](#), [Cisco – Understanding and configuring Voice VLAN](#).

QUESTION NO: 4

What are two recommendations for protecting network ports from being exploited when located in an office space outside of an IT closet? (Choose two.)

- A. enable the PortFast feature on ports
- B. implement port-based authentication
- C. configure static ARP entries
- D. configure ports to a fixed speed
- E. shut down unused ports

ANSWER: B E

Explanation:

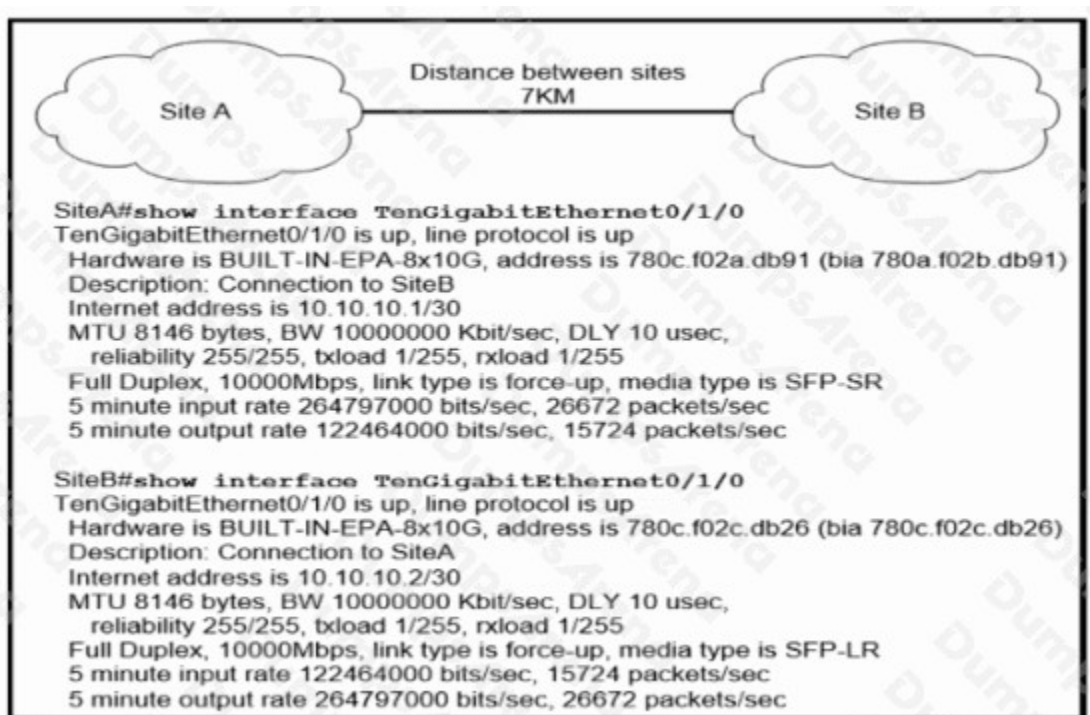
For switchports that are physically accessible in an office (not secured in an IT closet), the best practice is to reduce the chance that an unauthorized device can connect and gain network access. Two common recommendations are (1) using port-based authentication such as IEEE 802.1X, and (2) administratively disabling any ports that are not in use. Port-based authentication (often 802.1X, sometimes combined with MAB) forces a device/user to authenticate before the port is placed into an authorized VLAN, which directly mitigates “walk-up” attacks where someone plugs in a laptop or rogue device. Shutting down unused ports removes the attack surface entirely for ports that shouldn’t be active, preventing casual or opportunistic connections.

The other options don’t primarily address physical port exploitation. PortFast is a spanning-tree convergence feature for edge ports; it can actually increase risk if misused (e.g., without BPDU Guard) because it doesn’t authenticate or restrict access. Static ARP entries are not a practical, scalable control for office access ports and don’t stop someone from plugging in. Forcing a fixed speed/duplex is a troubleshooting/compatibility setting, not a security control.

References: [Cisco – 802.1X Authentication \(overview/tech notes\)](#), [Cisco – IEEE 802.1X on Catalyst switches](#)

QUESTION NO: 5

Refer to the exhibit.



Site A was recently connected to site B over a new single-mode fiber path. Users at site A report Intermittent connectivity Issues with applications hosted at site B. What is the reason for the problem?

- A. Heavy usage is causing high latency.
- B. An incorrect type of transceiver has been inserted into a device on the link.
- C. physical network errors are being transmitted between the two sites.
- D. The wrong cable type was used to make the connection.

ANSWER: B

Explanation:

The most likely cause is a transceiver mismatch on a single-mode fiber link. Single-mode fiber requires optics designed for SMF (for example, 1000BASE-LX/LH, 10GBASE-LR, etc.). If a multimode-only optic (such as many 1000BASE-SX or 10GBASE-SR modules) is inserted on one end, the link may come up in some cases (or appear to work at short distances) but will often show intermittent loss, flapping, or high error rates as the optical budget and modal characteristics don't match the fiber plant. That kind of intermittent application connectivity after a new SMF path is a classic symptom of using the wrong optic type for the medium/distance.

Option D ("wrong cable type") is less precise because the prompt already states the path is single-mode fiber; if the installed plant were actually multimode, it would typically be a consistent failure or a clearly out-of-spec design rather than intermittent behavior framed around a new SMF path. Option A (heavy usage/latency) doesn't align with a newly installed physical path issue and wouldn't usually be described as intermittent connectivity. Option C is vague; physical errors are symptoms, not the root cause, and the question asks for the reason.

References: [Cisco optics overview \(SX/LX/LH and fiber types\)](#), [10GbE optical modules \(SR vs LR and fiber types\)](#).

QUESTION NO: 6

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:

In eBGP (external BGP), the defining characteristic is that the BGP peers are in *different* autonomous systems, so option A is true. Also, BGP does not form adjacencies "automatically" the way some IGP's can; you must explicitly configure the neighbor on each router (neighbor IP and remote-as at minimum). That makes option B true: both sides must be configured to peer with each other for the session to come up.

Option C is false because the `network` statement in BGP controls which prefixes are originated/advertised into BGP; it does not create neighbors dynamically. (Dynamic neighbors exist in some designs like BGP listen ranges, but that's not the default behavior implied here and is not triggered by `network`.) Option D is false because "same autonomous system" describes iBGP, not eBGP. Option E is false because BGP timers do not have to match exactly to form a session; each side proposes timers in the OPEN message and the negotiated values are used (commonly the lower values). Matching timers is a best-practice for predictability, but not a strict requirement.

References: [Cisco BGP Timers Overview](#), [Cisco BGP Configuration Basics](#)

QUESTION NO: 7

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) is responsible for handling user traffic at line rate: receiving frames/packets, applying forwarding decisions already built by the control plane, rewriting headers as needed, and sending traffic out the correct interface. Adding or removing an 802.1Q tag is a classic data-plane function because the switch is modifying the Layer 2 header as it forwards traffic across access/trunk boundaries. Likewise, matching a destination MAC address to the MAC address table (CAM table) is part of Layer 2 forwarding and is performed in the data plane to decide the outbound interface.

By contrast, making configuration changes from an incoming NETCONF RPC is a management-plane activity (device configuration/automation interface), not packet forwarding. Running routing protocols is control-plane work because protocols like OSPF/EIGRP/BGP compute routes and populate the routing table/FIB used later by the data plane. Replying to an ICMP echo request is typically handled by the control plane/CPU (the device is the destination and generates a response), not by the forwarding ASIC as transit traffic. For more background on Cisco's separation of planes, see [Cisco: Internetworking Terms and Concepts \(planes overview\)](#) and [Cisco IOS XE Programmability Overview \(management interfaces like NETCONF\)](#).

QUESTION NO: 8

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

```
D 10.47.114.119/29 [90/6451] via F0/2
D 10.47.114.119/29 [90/52201] via F0/20
R 10.47.114.119/29 [120/9] via F0/12
R 10.47.114.119/29 [120/10] via F0/10
```

- A. F0/2
- B. F0/20
- C. F0/12
- D. F0/10

ANSWER: A

Explanation:

In Cisco routing, the interface used to send traffic toward a destination network is the *outgoing (egress) interface* associated with the best matching route in the routing table (longest-prefix match). In the diagram, the destination network is reached via the path that exits the router on interface **F0/2**, so that is the interface the router will use to forward packets toward that network.

Option A is correct because it represents the router's egress interface toward the next hop/destination network shown in the topology. The other interfaces (F0/20, F0/12, and F0/10) are either connected to different segments or are not on the forwarding path to the destination network in the provided diagram, so they would not be selected as the outgoing interface for that destination.

This behavior aligns with standard Cisco forwarding logic: the router consults the routing table, selects the best route (longest prefix; then administrative distance/metric as needed), and forwards the packet out the route's associated exit interface (or toward the next-hop IP reachable via that interface). See: [Cisco routing table and route selection overview](#) and [Cisco IOS route lookup \(longest match\) behavior](#).

QUESTION NO: 9

Which two options are the best reasons to use an IPv4 private IP space? (Choose two.)

- A. to enable intra-enterprise communication
- B. to implement NAT
- C. to connect applications
- D. to conserve global address space
- E. to manage routing overhead

ANSWER: A D

Explanation:

IPv4 private address space (RFC 1918) exists primarily to let organizations number internal hosts without consuming globally unique public IPv4 addresses. That directly supports **intra-enterprise communication** (A): devices inside the same company can communicate end-to-end using private addressing, and those addresses can be reused across many different organizations because they are not routed on the public Internet. The other key driver is **conserving global address space** (D): by using private addresses internally, an enterprise needs far fewer public IPv4 addresses, which helps mitigate IPv4 exhaustion.

Option (B) "to implement NAT" is commonly associated with private addressing, but it's not the fundamental reason to use private space; rather, NAT is a technique often used *because* private addresses are not Internet-routable and must be translated when accessing the public Internet. (C) "to connect applications" is too generic and not specific to private IPv4 space. (E) "to manage routing overhead" is not a primary reason; private addressing doesn't inherently reduce routing overhead (that's more about summarization and design), and private routes still exist within an enterprise.

References: [RFC 1918 - Address Allocation for Private Internets](#), [Cisco NAT Overview](#).

QUESTION NO: 10 - (SIMULATION)

SIMULATION

-

Guidelines

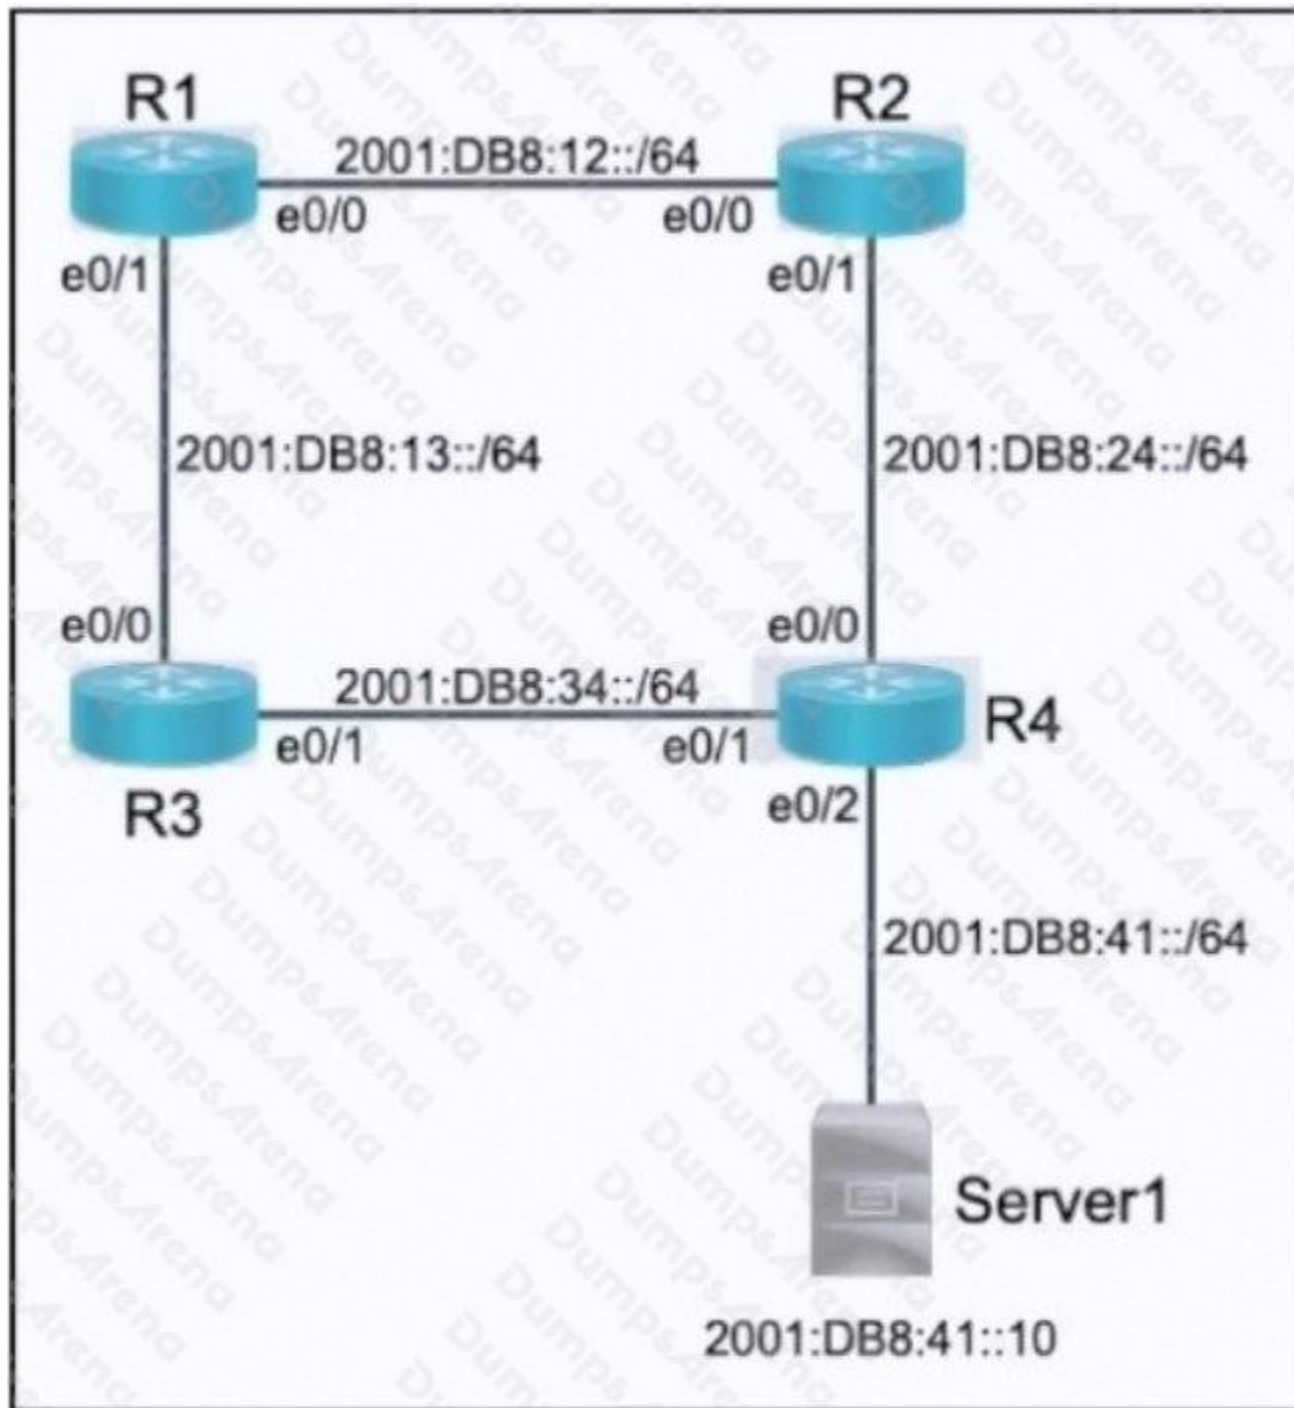
-

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 preconfigurations have been applied.
- 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



Tasks

All physical cabling is in place. Configurations should ensure that connectivity is established end-to-end.

1. Configure a route on R1 to ensure that R1 prefers R2 to reach the 2001:db8:41::/64 network
2. Configure a floating route on R1, and ensure that R1 uses R3 to reach the 2001 :db8:41::/64 network if the connection between R1 and R2 is down
3. Ping and traceroute should be working

ANSWER: See the explanation for answer

Explanation:

```
R1# config t
R1(config)# ipv6 route 2001:db8:41::/64 2001:db8:12::2
R1(config)# ipv6 route 2001:db8:41::/64 2001:db8:13::3 2
R1(config)#end
R1#wr
```

```
R2:
R2> en
R2# config t
R2(config)# ipv6 route 2001:db8:41::/64 2001:db8:24::4
R2(config)#end
R2#wr
```

```
R3:
R3> en
R3# config t
R3(config)# ipv6 route 2001:db8:41::/64 2001:db8:34::4
R3(config)#end
R3#wr
```

```
R4:
R4>en
R4# config t
R4(config)# ipv6 route 2001:db8:12::/64 2001:db8:24::2
R4(config)# ipv6 route 2001:db8:13::/64 2001:db8:34::3
R4(config)#end
4#wr
```

QUESTION NO: 11

Which two command sequences must be configured on a switch to establish a Layer 3 EtherChannel with an open-standard protocol? (Choose two.)

- A. interface GigabitEthernet0/0/1 channel-group 10 mode auto
- B. interface GigabitEthernet0/0/1 channel-group 10 mode on
- C. interface port-channel 10 no switchport
ip address 172.16.0.1 255.255.255.0
- D. interface GigabitEthernet0/0/1 channel-group 10 mode active
- E. interface port-channel 10 switchport switchport mode trunk

ANSWER: C D

Explanation:

To build a **Layer 3 EtherChannel** using an **open-standard protocol**, you use **LACP** (IEEE 802.3ad/802.1AX), not PAgP (Cisco-proprietary) and not static “on”. On the member interfaces, you must enable LACP negotiation with `channel-group ... mode active` (or passive on one side, active on the other). That’s why option **D** is required.

Next, because the question specifically says **Layer 3** EtherChannel, the logical Port-Channel interface must be converted to a routed interface with `no switchport` and then given an IP address. That’s exactly what option **C** does. In contrast, option **E** configures a Layer 2 trunk on the Port-Channel (it remains a switchport), which contradicts the Layer 3 requirement. Option **A** uses `mode auto`, which is PAgP (not open standard). Option **B** uses `mode on`, which is a static EtherChannel mode (no negotiation protocol), so it also doesn’t meet the “open-standard protocol” requirement.

References: [Cisco EtherChannel Configuration Notes](#), [Cisco LACP/PAgP EtherChannel Overview](#).

QUESTION NO: 12

Which two actions influence the EIGRP route selection process? (Choose two)

- A. The router calculates the reported distance by multiplying the delay on the exiting Interface by 256.
- B. The router calculates the best backup path to the destination route and assigns it as the feasible successor.
- C. The router calculates the feasible distance of all paths to the destination route
- D. The advertised distance is calculated by a downstream neighbor to inform the local router of the bandwidth on the link
- E. The router must use the advertised distance as the metric for any given route

ANSWER: B C

Explanation:

EIGRP selects routes based on the lowest *feasible distance (FD)*, which is the local router’s total computed metric to reach a destination via a given neighbor. To do that, the router evaluates all received routes by combining the neighbor’s *reported/advertised distance (RD/AD)* with the local cost to reach that neighbor (bandwidth/delay-based metric by default). This is why calculating the feasible distance for each candidate path directly influences which path becomes the successor (best path). In addition, EIGRP’s DUAL algorithm identifies loop-free backup paths called *feasible successors* by applying the feasibility condition (a neighbor’s RD must be less than the current successor’s FD). Having (or not having) a feasible successor affects route selection behavior during failures and convergence because EIGRP can immediately switch to that backup without recomputation.

Option A is wrong because EIGRP delay is summed in tens of microseconds and the composite metric uses scaling (including 256), but it’s not “delay on the exiting interface multiplied by 256” to form RD. Option D is wrong because RD is the

neighbor's metric to the destination, not merely "bandwidth on the link." Option E is wrong because EIGRP uses FD (not RD alone) as the route metric for selection.

References: [Cisco EIGRP Technical Overview](#), [Cisco DUAL and Feasible Successors](#)

QUESTION NO: 13

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:

When you configure IPv6 on an interface, the interface automatically joins certain well-known multicast groups that are required for basic IPv6 operation on the local link. First, it joins the **all-nodes** multicast group **FF02::1**, which allows a host to receive packets destined to "all IPv6 nodes on this link" (for example, some Neighbor Discovery and general link-local messaging). Second, it joins a **solicited-node** multicast group of the form **FF02::1:FFxx:xxxx**, which is derived from the last 24 bits of the interface's unicast/anycast IPv6 address and is used by Neighbor Discovery for address resolution (the IPv6 equivalent role to ARP, but via multicast rather than broadcast).

FF02::2 is the **all-routers** multicast group; it is joined by IPv6 routers (and by hosts only in special cases, e.g., if they are acting as routers). Since the question is about configuring IPv6 on an interface in general, FF02::2 is not guaranteed. Options A, B, and C are not multicast groups at all (they are unicast/global or ULA ranges), so they don't apply.

References: [Cisco IOS XE IPv6 Multicast Configuration Guide](#), [RFC 4291 \(IPv6 Addressing Architecture\)](#)

QUESTION NO: 14

Refer to the exhibit.



For security reasons, automatic neighbor discovery must be disabled on the R5 Gi0/1 interface. These tasks must be completed:

- Disable all neighbor discovery methods on R5 interface Gi0/1.
- Permit neighbor discovery on R5 interface Gi0/2.
- Verify there are no dynamically learned neighbors on R5 interface Gi0/1.
- Display the IP address of R6's interface Gi0/2.

Which configuration must be used?

- R5(config)#int Gi0/1
R5(config-if)#no cdp run
R5(config-if)#exit
R5(config)#lldp run
R5(config)#cdp enable
R5#sh cdp neighbor
R5#sh lldp neighbor
- R5(config)#int Gi0/1
R5(config-if)#no cdp enable
R5(config-if)#exit
R5(config)#no lldp run
R5(config)#cdp run
R5#sh cdp neighbor
R5#sh lldp neighbor
- R5(config)#int Gi0/1
R5(config-if)#no cdp enable
R5(config-if)#exit
R5(config)#no lldp run
R5(config)#cdp run
R5#sh cdp neighbor detail
R5#sh lldp neighbor
- R5(config)#int Gi0/1
R5(config-if)#no cdp enable
R5(config-if)#exit
R5(config)#lldp run
R5(config)#no cdp run
R5#sh cdp neighbor detail
R5#sh lldp neighbor

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

ANSWER: C

Explanation:

Option C is correct because it is the only choice that satisfies all four requirements using the right Cisco IOS controls and verification commands. To disable *all* automatic neighbor discovery on R5 Gi0/1, you must turn off both CDP and LLDP on that interface (for example, `no cdp enable` and `no lldp transmit/no lldp receive` under the interface). That prevents R5 from advertising itself and from learning neighbors via those protocols on Gi0/1. At the same time, neighbor discovery must remain permitted on Gi0/2, which means CDP/LLDP must remain enabled there (either by default or explicitly with `cdp enable` and `lldp transmit/receive`).

For verification, the correct operational checks are to ensure no neighbors are learned on Gi0/1 (for example, `show cdp neighbors interface gi0/1` and/or `show lldp neighbors interface gi0/1`). Finally, displaying the IP address of R6's Gi0/2 is done with a command like `show ip interface brief` (or `show interfaces gi0/2`) on R6. The other options typically miss disabling both protocols on Gi0/1, incorrectly disable discovery globally (impacting Gi0/2), or use the wrong show commands for the required verification.

QUESTION NO: 15

A network engineer is replacing the switches that belong to a managed-services client with new Cisco Catalyst switches. The new switches will be configured for updated security standards, including replacing Telnet services with encrypted connections and doubling the modulus size from 1024. Which two commands must the engineer configure on the new switches? (Choose two.)

- A. `crypto key generate rsa general-keys modulus 1024`
- B. `transport input all`
- C. `crypto key generate rsa usage-keys`
- D. `crypto key generate rsa modulus 2048`
- E. `transport Input ssh`

ANSWER: D E

Explanation:

To replace Telnet with an encrypted remote-access method on Cisco Catalyst switches, you must enable SSH access on the VTY lines and have RSA keys generated for the SSH server. The command **`transport input ssh`** under the VTY line configuration restricts inbound remote access to SSH only, effectively disabling Telnet on those lines. To “double the modulus size from 1024,” you must generate RSA keys with a 2048-bit modulus. That is done with **`crypto key generate rsa modulus 2048`** (platforms may also accept the “general-keys” keyword, but the key point is the 2048 modulus).

Option A is wrong because it explicitly generates 1024-bit keys, which does not meet the updated standard. Option B is wrong because **`transport input all`** allows both Telnet and SSH (and possibly other protocols), so it does not enforce encrypted-only access. Option C is not the right requirement here; “usage-keys” is not the typical command used to meet the stated goal in CCNA contexts, and it doesn’t specify the required 2048-bit modulus.

References: [Cisco Secure Shell \(SSH\) Configuration Example](#), [Cisco IOS SSH Configuration Guide](#).

QUESTION NO: 16

What is a function of a Next-Generation IPS?

- A. It analyzes and mitigates observed vulnerabilities in a network.
- B. It serves as a controller within a controller-based network
- C. It integrates with a RADIUS server to enforce Layer 2 device authentication rules
- D. It makes forwarding decisions based on learned MAC addresses

ANSWER: A

Explanation:

A Next-Generation Intrusion Prevention System (NGIPS) is designed to inspect traffic and actively prevent malicious activity, using techniques like deep packet inspection, signature- and behavior-based detection, and context (applications/users) to identify and stop attacks. In practice, that means it can detect exploit attempts and other threats and then take action (drop/reset/quarantine) to mitigate the risk to the network. Option A best matches this core purpose: analyzing traffic for threats/vulnerabilities and mitigating them in-line.

Option B is incorrect because a “controller within a controller-based network” describes SDN/controller roles (e.g., Cisco DNA Center, SD-WAN controllers), not an IPS function. Option C is incorrect because integrating with RADIUS for Layer 2 authentication is characteristic of Network Access Control/802.1X (e.g., Cisco ISE), not IPS. Option D is incorrect because making forwarding decisions based on learned MAC addresses is a Layer 2 switch/bridge function, unrelated to IPS.

References: [Cisco Next-Generation IPS \(NGIPS\) overview](#), [Cisco Firepower NGIPS](#).

QUESTION NO: 17

A corporate office uses four floors in a building.

- ☞ Floor 1 has 24 users.
- ☞ Floor 2 has 29 users.
- ☞ Floor 3 has 28 users.
- ☞ Floor 4 has 22 users.

Which subnet summarizes and gives the most efficient distribution of IP addresses for the router configuration?

- A. 192.168.0.0/24 as summary and 192.168.0.0/28 for each floor
- B. 192.168.0.0/23 as summary and 192.168.0.0/25 for each floor
- C. 192.168.0.0/25 as summary and 192.168.0.0/27 for each floor
- D. 192.168.0.0/26 as summary and 192.168.0.0/29 for each floor

ANSWER: C

Explanation:

Each floor needs its own subnet sized to the number of hosts. The largest floor has 29 users, so you need a subnet that supports at least 29 usable IPs. A /27 provides 32 total addresses (30 usable), which fits 29, 28, 24, and 22 users efficiently with minimal waste. Since there are four floors, you need four /27 networks; four contiguous /27s aggregate neatly into a single /25 summary (because $4 \times 32 = 128$ addresses). Therefore, using 192.168.0.0/25 as the summary and allocating /27s per floor is the most efficient and correctly summarizes the block for routing.

Option A is wrong because /28 only provides 14 usable IPs, far too small. Option B is wrong because /25 per floor is massively wasteful (126 usable per floor) and the /23 summary is unnecessary. Option D is wrong because /29 provides only 6 usable IPs per floor, also far too small, and the /26 summary wouldn't even cover four /29s in a meaningful “per floor” design. For subnet sizing and usable host counts, see [Cisco IP Subnetting Reference](#) and [CIDR overview](#).

QUESTION NO: 18

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 (MFA) means a login requires evidence from at least two different factor categories: something you know (knowledge), something you have (possession), and/or something you are (inherence/biometrics). A classic MFA example is using a password (know) plus a one-time code generated by an authenticator app (have). In the options given, **soft tokens** are a possession factor because the user must have access to the device/app that generates or receives the one-time passcode. **Unique user knowledge** maps to the knowledge factor (for example, a password, PIN, or answers to challenge questions), and when combined with another factor it forms MFA.

Single sign-on (SSO) is not an authentication factor; it's an access/authentication architecture that reduces the number of times a user authenticates, but it doesn't inherently add a second factor. **Passwords that expire** are still just passwords (one knowledge factor) with a policy applied, not MFA. A **shared password repository** (password vault) is a storage/management approach and does not itself provide a second factor; it may even reduce security if "shared" implies non-unique credentials.

References: [Cisco – Multi-Factor Authentication \(MFA\)](#), [NIST – Multi-Factor Authentication](#).

QUESTION NO: 19

What is a similarity between 1000BASE-LX and 1000BASE-T standards?

- A. Both use the same data-link header and trailer formats
- B. Both cable types support LP connectors
- C. Both cable types support Rj-45 connectors
- D. Both support up to 550 meters between nodes

ANSWER: A

Explanation:

The key similarity between 1000BASE-LX and 1000BASE-T is that they are both Gigabit Ethernet (IEEE 802.3) PHY standards and therefore carry the same Ethernet frames at Layer 2. That means the MAC header/trailer (e.g., source/destination MAC, EtherType/length, FCS) is identical regardless of whether the physical medium is fiber (1000BASE-LX) or copper twisted pair (1000BASE-T). The PHY encoding and cabling differ, but the data-link framing does not, which is why switches/hosts can interoperate at Layer 2 as long as the interfaces match the medium.

Option B is wrong because "LP connectors" is not a standard connector type for these media; 1000BASE-LX commonly uses LC/SC fiber connectors, while 1000BASE-T uses RJ-45. Option C is wrong because only 1000BASE-T uses RJ-45; 1000BASE-LX is fiber and does not use RJ-45. Option D is wrong because the supported distances differ: 1000BASE-T is up to 100 m over Cat5e/6, while 1000BASE-LX typically supports much longer distances over single-mode fiber (and shorter over multimode depending on conditions), not a shared 550 m limit.

References: [Cisco Ethernet and Gigabit Ethernet Overview](#), [Gigabit Ethernet \(media types and distances\)](#)

QUESTION NO: 20

Using direct sequence spread spectrum, which three 2.4-GHz channels are used to limit collisions?

- A. 1,6,11
- B. 1,5,10
- C. 1,2,3
- D. 5,6,7

ANSWER: A

Explanation:

In the 2.4 GHz band (802.11b/g/n), channels are spaced 5 MHz apart but each channel is ~22 MHz wide (for 802.11b DSSS/CCK), so most adjacent channels overlap and cause co-channel/adjacent-channel interference. To minimize overlap (and therefore reduce contention and retransmissions that people often loosely call “collisions”), you choose the three non-overlapping channels in the common regulatory domains: channels 1, 6, and 11. These center frequencies are far enough apart that their occupied bandwidths do not significantly overlap, making them the standard design choice for multi-AP deployments in 2.4 GHz.

Option A is correct because 1/6/11 is the canonical non-overlapping set used in North America and widely taught in CCNA. Options B, C, and D are incorrect because those channel groupings overlap heavily; deploying APs on those channels in the same area increases interference and reduces throughput. Note that some regions allow channel 13, and you may see 1/5/9/13 patterns discussed for certain channel widths, but the classic DSSS non-overlapping trio remains 1/6/11.

References: [Cisco WLAN Channel Planning and Deployment](#), [List of WLAN channels \(2.4 GHz channel overlap\)](#)

QUESTION NO: 21

Which REST method updates an object in the Cisco DNA Center Intent API?

- A. CHANGE
- B. UPDATE
- C. POST
- D. PUT

ANSWER: D

Explanation:

In RESTful APIs (including Cisco DNA Center Intent APIs), the HTTP method used to update an existing resource is typically **PUT**. PUT is defined as a full update (replace) of a resource at a known URI, and it is idempotent—sending the same PUT request multiple times should result in the same final state. In many Cisco DNA Center endpoints, PUT is used for “update” operations (for example, updating configuration objects where the resource already exists), while POST is commonly used to create new resources or to trigger actions/operations.

The other options are not valid HTTP methods used by REST: **CHANGE** and **UPDATE** are not standard HTTP verbs, so they are incorrect. **POST** is a standard REST method, but it is generally associated with creating resources (or invoking non-idempotent operations) rather than updating an existing object in a RESTful sense. While some APIs may use POST for updates in specific cases, the canonical REST method for updating an object is PUT, which is what this question is testing.

References: [Cisco DNA Center Platform Documentation](#), [MDN Web Docs: HTTP PUT](#)

QUESTION NO: 22

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

ANSWER: B C**Explanation:**

On Cisco Wireless LAN Controllers, when you configure WPA2 with a Pre-Shared Key (PSK), the controller lets you enter the PSK in one of two common representations: an ASCII passphrase or a hexadecimal key. The ASCII option is the familiar “password” style entry (typically 8–63 characters for WPA/WPA2-Personal). The hexadecimal option is for entering the raw PSK material as a 64-character hex string (256-bit key) rather than a passphrase that must be expanded into a key.

That’s why **ASCII** and **hexadecimal** are the two valid formats in the WLC GUI for WPA2-PSK. Options like **decimal** and **binary** are not used for WPA2 PSK entry on Cisco WLCs, and **base64** is also not a supported PSK input format in the standard WPA2-PSK WLAN configuration workflow.

References: Cisco WLC configuration guides and WLAN security configuration examples describe PSK entry as ASCII passphrase or 64-hex key. See [Cisco WLAN Security \(WPA/WPA2\) overview](#) and [Cisco 9800 Multi-PSK configuration guide](#).

QUESTION NO: 23

Refer to the exhibit.

The screenshot displays the configuration page for a WLAN on a Cisco Wireless LAN Controller. The 'Security' tab is selected, and the 'Layer 3' sub-tab is active. The 'Layer 2 Security' dropdown is set to 'WPA+WPA2'. The 'MAC Filtering' checkbox is unchecked. Under 'Fast Transition', the 'Fast Transition' checkbox is unchecked. The 'Protected Management Frame' section shows 'PMF' set to 'Required'. The 'WPA+WPA2 Parameters' section includes checkboxes for 'WPA Policy', 'WPA2 Policy', 'WPA2 Encryption' (with 'AES' and 'TKIP' options), and 'OSEN Policy'. The 'Authentication Key Management' section has checkboxes for '802.1X', 'CCKM', 'PSK', 'FT 802.1X', and 'FT PSK', all of which are currently unchecked. The 'PSK Format' dropdown is set to 'ASCII', and there is an empty text input field below it. At the bottom, the 'WPA gtk-randomize State' dropdown is set to 'Disable'.

The network engineer is configuring a new WLAN and is told to use a setup password for authentication instead of the RADIUS servers. Which additional set of tasks must the engineer perform to complete the configuration?

- A. Disable PMF Enable PSK Enable 802.1x
- B. Select WPA Policy Enable CCKM Enable PSK
- C. Select WPA Policy Select WPA2 Policy Enable FT PSK
- D. Select WPA2 Policy Disable PMF Enable PSK

ANSWER: D

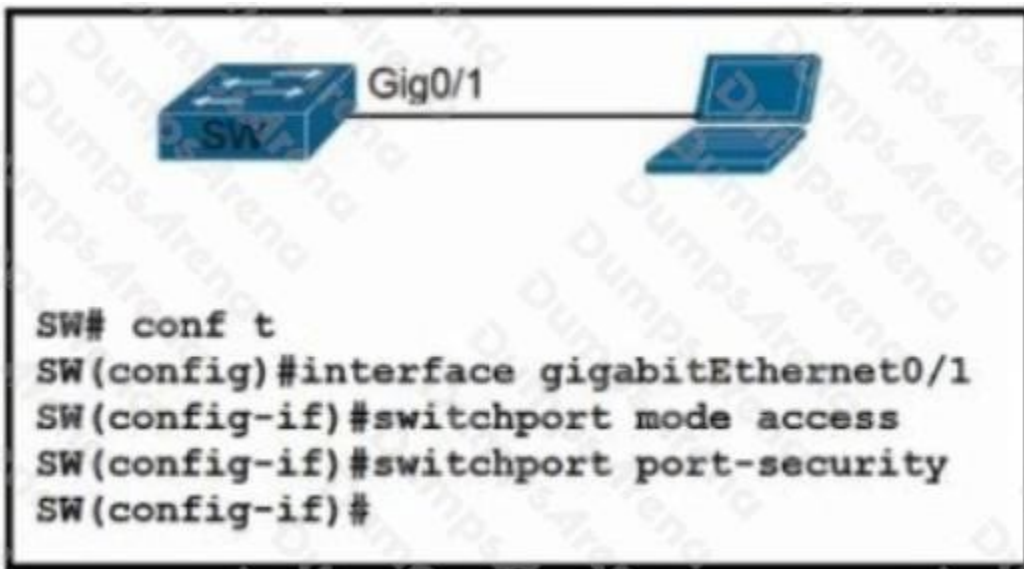
Explanation:

Using a “setup password” for WLAN authentication means the WLAN must use a pre-shared key (PSK) method (personal mode) rather than 802.1X/EAP with a RADIUS server (enterprise mode). On Cisco WLCs, that requires selecting an appropriate WPA/WPA2 security policy and enabling PSK so the controller will accept a shared passphrase instead of attempting 802.1X authentication. In addition, PMF (Protected Management Frames / 802.11w) settings must be compatible with the chosen security mode and client capabilities; many deployments disable PMF (or set it to optional) when the requirement is simply “use a password” and broad client compatibility is desired. Therefore, selecting the WPA2 policy and enabling PSK (with PMF disabled per the option) is the correct set of additional tasks.

Option A is wrong because enabling 802.1X contradicts the requirement to avoid RADIUS-based authentication. Option B is wrong because CCKM is a legacy fast-roaming feature tied to WPA/TKIP-era designs and is not the required step to use PSK. Option C is wrong because FT PSK (802.11r) is an optional roaming enhancement, not a required task to switch from RADIUS to a password-based WLAN.

References: [Cisco WLAN Security \(WPA/WPA2\) overview](#), [Cisco 802.11w / PMF configuration and behavior](#).

QUESTION NO: 24



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 configured to complete this task? (Choose two.)

- A. SW(config-if)#switchport port-security violation restrict
- B. SW(config-if)#switchport port-security mac-address 0010.7B84.45E6

- C. SW(config-if)#switchport port-security maximum 2
- D. SW(config-if)#switchport port-security violation shutdown
- E. SW(config-if)#switchport port-security mac-address sticky

ANSWER: A E

Explanation:

To meet the requirements, the switch must learn MAC addresses dynamically and must generate log messages while keeping the interface up when a violation occurs. For dynamic learning in a way that “remembers” learned addresses, Cisco port security uses *sticky* MAC learning. The command `switchport port-security mac-address sticky` causes the switch to dynamically learn MAC addresses seen on the port and add them to the running configuration (and to the startup configuration if you save). That satisfies the “learned dynamically” requirement without manually specifying a MAC address.

For the violation behavior, `switchport port-security violation restrict` is the correct mode. In restrict mode, frames from unknown source MAC addresses are dropped, the port remains in the up state (not err-disabled), and the switch increments the security violation counter and generates syslog/SNMP notifications—matching the requirement to log without disabling the interface. By contrast, `violation shutdown` err-disables the port, which violates the requirement. A manually configured MAC address does not meet “learn dynamically,” and setting a maximum is not explicitly required by the prompt (sticky can work with the default maximum of 1 unless otherwise specified).

References: [Cisco Support: Configuring Port Security](#), [Cisco Catalyst 2960 Software Configuration Guide \(Port Security\)](#)

QUESTION NO: 25

Which capability does TFTP provide?

- A. loads configuration files on systems without data storage devices
- B. provides authentication for data communications over a private data network
- C. provides encryption mechanisms for file transfer across a WAN
- D. provides secure file access within the LAN

ANSWER: A

Explanation:

TFTP (Trivial File Transfer Protocol) is a very lightweight file transfer protocol that runs over UDP (port 69) and is commonly used for simple, unauthenticated file transfers such as bootstrapping devices and moving configuration or image files. A classic use case is a device that has little or no local storage (for example, diskless workstations or network devices during bootstrap) downloading a configuration or boot file from a TFTP server. That aligns with option A.

The key limitation to remember for CCNA is that TFTP provides no built-in security features: no authentication, no encryption, and no secure access controls. It’s “trivial” by design, which is why it’s often used only in trusted environments or replaced by SCP/SFTP/HTTPS when security is required. Therefore, options B, C, and D are incorrect because they describe authentication, encryption, or secure access—capabilities that TFTP simply does not provide.

References: [RFC 1350 - The TFTP Protocol \(Revision 2\)](#), [Cisco: Trivial File Transfer Protocol \(TFTP\) Overview](#).

QUESTION NO: 26

Which two host addresses are reserved for private use within an enterprise network? (Choose two.)

- A. 10.172.76.200
- B. 12.17.1.20
- C. 172.15.2.250
- D. 172.31.255.100
- E. 192.169.32.10

ANSWER: A D

Explanation:

Private IPv4 addresses are defined by RFC 1918 and are reserved for use inside private enterprise networks (not routable on the public Internet). The three private ranges are: 10.0.0.0/8, 172.16.0.0/12 (172.16.0.0–172.31.255.255), and 192.168.0.0/16.

Option A (10.172.76.200) is within 10.0.0.0/8, so it is a valid private host address. Option D (172.31.255.100) falls within 172.16.0.0/12, so it is also a valid private host address.

Option B (12.17.1.20) is in a public range (12.0.0.0/8 is not RFC1918), so it is not reserved for private use. Option C (172.15.2.250) is just outside the private 172.16.0.0/12 block (it's in 172.0.0.0/8 but not in 172.16–172.31), so it's public. Option E (192.169.32.10) is not in 192.168.0.0/16 (note 169 vs 168), so it's also public. These distinctions are important for NAT design and internal addressing plans.

References: [RFC 1918 - Address Allocation for Private Internets](#), [Cisco: IP Addressing and Subnetting for New Users](#)

QUESTION NO: 27

Which statement correctly compares traditional networks and controller-based networks?

- A. Only traditional networks offer a centralized control plane
- B. Only traditional networks natively support centralized management
- C. Traditional and controller-based networks abstract policies from device configurations
- D. Only controller-based networks decouple the control plane and the data plane

ANSWER: D

Explanation:

The correct comparison is that **controller-based (SDN) networks decouple the control plane from the data plane**, whereas traditional networks generally keep both planes on the same device. In a traditional network, each router/switch runs its own control-plane functions (for example, routing protocols, STP, etc.) and uses them to program its local forwarding (data) plane. That's a distributed control model: intelligence is spread across devices and coordination happens via protocol messages between peers.

In a controller-based architecture, the control logic is centralized (logically centralized, even if implemented as a cluster) in a controller, and network devices focus primarily on forwarding. The controller uses APIs/protocols to program device behavior, which is the essence of decoupling control and forwarding.

Option A is wrong because traditional networks are not characterized by a centralized control plane; they are typically distributed. Option B is wrong because while traditional networks can be centrally *managed* with tools, they don't natively centralize the control plane the way controller-based designs do. Option C is wrong because policy abstraction from device-by-device configuration is a key benefit of controller-based/intent-based networking; traditional networks usually require per-device configuration.

References: [Cisco SDN overview](#), [Cisco DNA Center \(controller-based management\)](#)

QUESTION NO: 28

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 selects routes based on the lowest *feasible distance (FD)*, which is the total metric from the local router to the destination. To do that, the router computes the FD for each candidate path (successor candidates) and installs the path with the lowest FD as the successor. That's why calculating the feasible distance of all paths directly influences route selection (Option B).

EIGRP also tries to precompute loop-free backup routes called *feasible successors*. A path can be a feasible successor only if it meets the feasibility condition: the neighbor's *reported/advertised distance (RD/AD)* to the destination is less than the current successor's FD. When such a path exists, EIGRP keeps it as a backup and can switch to it immediately if the successor fails—this behavior influences which alternate paths are considered usable (Option D).

Option C is wrong because EIGRP does not use advertised distance alone as the route metric; it uses the composite metric (bandwidth/delay by default) to compute FD. Option E is wrong because EIGRP's metric uses delay (in tens of microseconds) and bandwidth, and the overall metric is scaled (commonly by 256), but RD is not simply "delay × 256." Option A is largely explanatory text and contains incorrect phrasing (AD is not "bandwidth on the link"); it's the neighbor's distance to the destination.

References: [Cisco EIGRP Concepts](#), [EIGRP overview \(metric/FD/RD\)](#)

QUESTION NO: 29

Router R1 must send all traffic without a matching routing-table entry to 192.168.1.1. Which configuration accomplishes this task?

- A.** R1#config t
R1(config)#ip routing
R1(config)#ip route default-route 192.168.1.1

```
B. R1#config t
R1(config)#ip routing
R1(config)#ip route 192.168.1.1 0.0.0.0 0.0.0.0
```

```
C. R1#config t
R1(config)#ip routing
R1(config)#ip route 0.0.0.0 0.0.0.0 192.168.1.1
```

```
D. R1#config t
R1(config)#ip routing
R1(config)#ip default-gateway 192.168.1.1
```

ANSWER: C

Explanation:

To forward all packets that do not match a more specific route, R1 needs a default route (also called a gateway of last resort). In Cisco IOS, the standard way to configure this is a static default route using destination `0.0.0.0` with mask `0.0.0.0`, pointing to the next-hop IP address. Option C does exactly that: `ip route 0.0.0.0 0.0.0.0 192.168.1.1`. With this in the routing table, any traffic without a matching entry will be sent to `192.168.1.1`.

Option A is invalid syntax: there is no `ip route default-route` command in IOS for creating a default route. Option B is also incorrect because it attempts to create a route to destination `192.168.1.1` with an invalid mask/format and does not represent a default route. Option D (`ip default-gateway`) is used on devices that are not routing IP (for example, a Layer 2 switch or a router with IP routing disabled). Since the question is about a router sending unmatched traffic, the correct mechanism is a default static route, not `ip default-gateway`.

References: [Cisco: Configuring a Default Route](#), [Cisco IOS XE Static Routes Configuration Guide](#).

QUESTION NO: 30

What are two benefits of private IPv4 addressing? (Choose two.)

- A. propagates routing information to WAN links
- B. provides unlimited address ranges
- C. reuses addresses at multiple sites
- D. conserves globally unique address space
- E. provides external internet network connectivity

ANSWER: C D

Explanation:

Private IPv4 addressing (RFC 1918) provides address space that is not routable on the public Internet and is intended for use inside organizations. Two key benefits follow from that design. First, private space can be **reused at multiple sites** (option C). Because these addresses are not globally unique and are not advertised on the Internet, many different companies can use the same RFC1918 ranges internally without conflict; when connectivity between sites or to the Internet is needed, NAT or proper routing/translation is used to avoid overlap issues. Second, using private addresses **conserves globally unique address space** (option D). By numbering internal hosts with RFC1918 space and translating only the necessary traffic to public addresses, organizations reduce consumption of scarce public IPv4 addresses.

Option A is incorrect because private addressing does not inherently “propagate routing information to WAN links”; routing propagation is a function of routing protocols and design, not a benefit of using RFC1918 space. Option B is incorrect because private IPv4 does not provide unlimited ranges—RFC1918 defines three specific blocks. Option E is incorrect

because private addresses alone do not provide Internet connectivity; they require NAT and/or proxying with public addressing to reach the Internet.

References: [RFC 1918 - Address Allocation for Private Internets](#), [Cisco NAT Overview and Configuration](#)

QUESTION NO: 31

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:

REST-based APIs commonly use standard HTTP verbs to perform CRUD-style actions on resources. Two of the most fundamental and widely used methods are **GET** and **POST**. **GET** is used to retrieve a representation of a resource (read-only). It should be safe and idempotent, meaning repeated GETs do not change server state. **POST** is used to submit data to the server, often to create a new resource or trigger server-side processing; it is not inherently idempotent, which fits many “create” or “action” use cases.

The other options are not valid HTTP methods used in REST APIs. **REMOVE** is not an HTTP verb (the correct RESTful verb for deletion is **DELETE**). **REDIRECT** is not an HTTP method; redirection is handled via HTTP status codes (for example, 301/302) and the `Location` header. **POP** is not an HTTP method (it’s associated with email retrieval protocols like POP3, not HTTP). For reference, the standard HTTP methods are defined in the HTTP Semantics specification, and REST APIs typically map these verbs to resource operations.

References: [RFC 9110: HTTP Semantics](#), [MDN Web Docs: HTTP request methods](#)

QUESTION NO: 32

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:

An Ethernet frame that fails the Frame Check Sequence (FCS) has a bad CRC value when the receiving interface recalculates the checksum and compares it to the FCS field. On Cisco switches/routers, this condition increments the `CRC` counter because the frame is specifically identified as having a CRC/FCS error. At the same time, it also increments the

broader *input errors* counter, which is a roll-up of several receive-side error conditions (including CRC, frame, runts, giants, etc.). In other words, CRC errors are a subset of input errors, so you typically see both counters rise together when FCS failures occur.

The *frame* counter is generally associated with non-CRC framing problems (for example, alignment errors on Ethernet), so it is not the best match for a pure FCS/CRC failure. *Giants* and *runts* relate to invalid frame sizes (too large or too small) and may increment if the bad frame also violates size rules, but an FCS failure alone does not imply an over/undersized frame.

References: [Cisco Support: Understanding Ethernet Interface Counters](#), [Cisco Support: Troubleshooting CRC/Input Errors](#)

QUESTION NO: 33

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 can be added with redundant connections between them.
- B. A spine switch can be added with at least 40 GB uplinks.
- C. A leaf switch can be added with connections to every spine switch.
- D. A leaf switch can be added with a single connection to a core spine switch.

ANSWER: C

Explanation:

Spine-and-leaf scales access ports primarily by adding more leaf switches. In this design, endpoints (servers, services, edge devices) connect to leaf switches, so when you need additional access ports you add another leaf. To preserve predictable performance (consistent latency and bandwidth) and the “fabric” behavior, the new leaf is connected upstream to *every* spine switch. This keeps the topology symmetric and provides multiple equal-cost paths (ECMP) between any two leaves, which is a key benefit of spine-and-leaf designs.

Option C is correct because it describes exactly this scaling method: add a leaf and connect it to all spines. Option A is misleading because adding both a spine and a leaf is not required just to gain access ports; you typically add only leaves until spine capacity becomes the bottleneck. Option B is incorrect because uplink speed (40G, 100G, etc.) is a design choice, not a defining scalability requirement. Option D is wrong because connecting a leaf to only a single “core spine” breaks the full-mesh leaf-to-spine principle and reduces redundancy and path diversity.

References: [Cisco Nexus 9000 Series NX-OS Verified Scalability Guide \(Spine-and-Leaf overview\)](#), [Cisco spine-and-leaf architecture overview](#).

QUESTION NO: 34

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: C F

Explanation:

To ensure that only high-speed 2.4-GHz clients can connect to the WLAN, you must (1) restrict the WLAN to the 2.4-GHz band and (2) restrict the allowed PHY types on that band to the higher-speed standard. On Cisco WLCs, the WLAN "Radio Policy" controls which band(s) the SSID is advertised on. Setting it to **802.11g Only** limits the WLAN to 2.4 GHz (and excludes 5 GHz/802.11a). Then, to prevent slower 802.11b clients from associating, you must disable 802.11b data rates (or require OFDM-only). In practice this is done by disabling 1, 2, 5.5, and 11 Mbps rates and setting an OFDM rate (for example 12 Mbps) as mandatory, which forces clients to use 802.11g (or better) to join.

Options A and B (Broadcast SSID and Status) affect SSID visibility and whether the WLAN is enabled, but they do not enforce "high-speed 2.4 GHz only." Option D (802.11a Only) would force 5 GHz, the opposite of the requirement. Option E only maps the WLAN to a dynamic interface/VLAN and is unrelated to client PHY speed/band selection.

References: [Cisco Wireless LAN Controller Configuration Examples](#), [Cisco Support: WLAN Data Rates and 802.11b Support](#)

QUESTION NO: 35

Which two server types support domain name to IP address resolution? (Choose two >

- A. ESX host
- B. resolver
- C. web
- D. file transfer
- E. authentication

ANSWER: B E

Explanation:

Name-to-IP address resolution is provided by the Domain Name System (DNS). In DNS, a *resolver* (often called a DNS client) is the component that initiates DNS queries on behalf of an end host or application, and a DNS server (name server) answers those queries with records such as A/AAAA (name-to-IP) or CNAME. From the options given, **resolver** is clearly one of the correct choices because it directly participates in translating domain names into IP addresses by querying DNS infrastructure.

The other correct server type is not listed explicitly as "DNS server/name server," so the best match among the remaining choices is **authentication**, because in many enterprise designs DNS is commonly integrated with identity services (for example, Microsoft Active Directory Domain Services typically provides DNS alongside authentication). While "authentication" is not a DNS role by itself, it is the only option that plausibly maps to a server platform that commonly hosts DNS services. The other options (ESX host, web, file transfer) do not inherently provide name-to-IP resolution; they may use DNS but are not responsible for resolving names for clients.

QUESTION NO: 36

How does IPsec provide secure networking for applications within an organization?

- A. It takes advantage of FTP to secure file transfers between nodes on the network.
- B. It provides GRE tunnels to transmit traffic securely between network nodes.
- C. It enables sets of security associations between peers.
- D. It leverages TFTP providing secure file transfers among peers on the network.

ANSWER: C

Explanation:

IPsec secures IP communications by establishing a set of negotiated parameters between peers called Security Associations (SAs). An SA defines how traffic is protected (for example, which encryption and integrity algorithms are used, keying material, lifetimes, and whether ESP or AH is used). In practice, IPsec uses IKE (Internet Key Exchange) to authenticate peers and dynamically negotiate these SAs, then applies the resulting policies to protect packets at Layer 3. This is what enables applications to communicate securely without each application needing to implement its own encryption—IPsec provides confidentiality, integrity, and anti-replay protection for IP traffic based on the agreed SAs.

Option A and D are incorrect because FTP and TFTP are application-layer file transfer protocols and are not what IPsec “takes advantage of” to provide security; in fact, FTP/TFTP are not inherently secure. Option B is incorrect because GRE is a tunneling/encapsulation protocol and does not provide encryption by itself; while GRE can be combined with IPsec, IPsec does not “provide GRE tunnels.” The core IPsec mechanism is the creation and use of SAs between peers.

References: [RFC 4301 - Security Architecture for the Internet Protocol](#), [Cisco: Understanding IKE and IPsec](#)

QUESTION NO: 37

Which management security process is invoked when a user logs in to a network device using their username and password?

- A. authentication
- B. auditing
- C. accounting
- D. authorization

ANSWER: A

Explanation:

The correct process is **authentication**. In AAA (Authentication, Authorization, and Accounting), authentication is the step where the device verifies a user’s identity—typically by checking the supplied username and password against a local database or an external server such as RADIUS or TACACS+. This is exactly what happens at login: the system is answering “Who are you, and can you prove it?”

Authorization happens *after* successful authentication and determines what the authenticated user is allowed to do (for example, which commands they can run or which privilege level they receive). **Accounting** records what the user did (command logging, session start/stop, resource usage) for tracking and reporting. **Auditing** is a broader security concept

related to reviewing logs and events; while it may use accounting data, it is not the AAA process invoked by entering a username and password.

References: Cisco AAA overview and concepts are covered in Cisco documentation such as [Cisco RADIUS/AAA overview](#) and AAA configuration guidance like [Cisco IOS XE AAA configuration guide](#).

QUESTION NO: 38

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

```
SW1#show spanning-tree vlan 30
VLAN0030
Spanning tree enabled protocol rstp
Root ID      Priority      32798
Address      0025.63e9.c800
Cost         19
Port         1 (FastEthernet 2/1)
Hello Time   2 sec
Max Age      30 sec
Forward Delay 20 sec

[Output suppressed]
```

- 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:

From the exhibit output (typical `show spanning-tree / interface` STP detail), two key facts can be concluded. First, the switch is running Rapid PVST+ when the STP mode line indicates `spanning-tree mode rapid-pvst` (or the output explicitly shows Rapid-PVST). That makes option E correct and option C incorrect (PVST+ is the classic 802.1D per-VLAN implementation, while Rapid PVST+ is Cisco's per-VLAN 802.1w implementation).

Second, the root port on a non-root switch is the port that provides the lowest-cost path toward the root bridge; in STP output it is explicitly labeled as the Root Port (RP) for that VLAN/instance. In the exhibit, FastEthernet 2/1 is shown as the RP, so option A is correct. If a port is the root port, it cannot simultaneously be the designated port for that same VLAN/instance on that switch, so option B is incorrect.

Finally, option D ("This is a root bridge") would only be true if the switch's Bridge ID matched the Root ID and there were no root ports (all ports would be designated). Since a root port is present, the switch is not the root bridge.

References: [Cisco STP Concepts and Operation](#), [Cisco Rapid Spanning Tree \(802.1w\) Overview](#)

QUESTION NO: 39

Which two components comprise part of a PKI? (Choose two.)

- A. preshared key that authenticates connections
- B. one or more CRLs
- C. RSA token
- D. CA that grants certificates
- E. clear-text password that authenticates connections

ANSWER: B D

Explanation:

A Public Key Infrastructure (PKI) is the collection of roles, policies, and services used to create, manage, distribute, and revoke digital certificates. Two core PKI components are the Certificate Authority (CA) and certificate status/revocation information. The CA is the trusted entity that issues (signs) certificates binding an identity to a public key, so option D is correct. PKI also needs a way to communicate that a certificate should no longer be trusted before it expires; one common mechanism is a Certificate Revocation List (CRL), which is why option B is also correct.

The other options describe authentication methods that are not PKI components. A preshared key (A) and a clear-text password (E) are symmetric/credential-based authentication approaches and do not provide certificate issuance/revocation services. An RSA token (C) is a form of multi-factor authentication hardware/software token; it can be used alongside PKI, but it is not a required PKI building block. In short, PKI is centered on certificate issuance (CA) and certificate validation/revocation (CRL/OCSP), not on shared secrets or tokens.

References: [RFC 5280 - Internet X.509 Public Key Infrastructure Certificate and CRL Profile](#), [Cisco: PKI Basics](#)

QUESTION NO: 40

Which three statements are typical characteristics of VLAN arrangements? (Choose three.)

- A. A new switch has no VLANs configured.
- B. Connectivity between VLANs requires a Layer 3 device.
- C. VLANs typically decrease the number of collision domains.
- D. Each VLAN uses a separate address space.
- E. A switch maintains a separate bridging table for each VLAN.
- F. VLANs cannot span multiple switches.

ANSWER: B D E

Explanation:

VLANs logically segment a switched network into multiple Layer 2 broadcast domains. Because each VLAN is its own broadcast domain, traffic between VLANs is not switched at Layer 2; it must be routed by a Layer 3 device (a router or multilayer switch using SVIs), so statement B is correct. In practice, each VLAN is normally mapped to a different IP subnet/address space to keep Layer 3 boundaries aligned with the Layer 2 segmentation, making D a typical design characteristic (even though it's a design convention rather than a hard technical requirement). Switches also learn MAC addresses per VLAN; effectively, the MAC address table is maintained with VLAN context (often described as separate bridging tables per VLAN), so E is correct.

A is incorrect because Cisco switches come with default VLANs (notably VLAN 1) and default port membership, so it's not true that a new switch has no VLANs configured. C is incorrect because VLANs increase the number of broadcast domains,

but they do not “decrease the number of collision domains”; on modern switched Ethernet, each switchport is already its own collision domain regardless of VLANs. F is incorrect because VLANs can span multiple switches using trunk links (e.g., IEEE 802.1Q).

References: [Cisco VLAN basics and benefits](#), [Cisco 802.1Q trunking overview](#).

QUESTION NO: 41

Which two values or settings must be entered when configuring a new WLAN in the Cisco Wireless LAN Controller GUI? (Choose two)

- A. management interface settings
- B. QoS settings
- C. Ip address of one or more access points
- D. SSID
- E. Profile name

ANSWER: D E

Explanation:

When you create a new WLAN on a Cisco Wireless LAN Controller (WLC) using the GUI, the controller’s WLAN creation workflow requires you to define two key identifiers up front: the **Profile Name** and the **SSID**. The Profile Name is an internal label used by the WLC to reference the WLAN configuration (it doesn’t have to match what clients see). The SSID is the network name that is advertised (or optionally hidden) and is what wireless clients select when connecting. These are mandatory fields in the initial “New WLAN”/“WLANs > Create New” step, along with a WLAN ID, but “WLAN ID” is not offered as an option here.

Other items listed are not strictly required to be entered to create the WLAN object. **QoS settings** can be left at defaults and adjusted later. **Management interface settings** are part of controller interface configuration (and/or WLAN interface mapping), not something you must enter as part of creating every new WLAN. The **IP address of one or more access points** is not used when defining a WLAN; APs join the controller and then broadcast WLANs based on AP group/site tags policy, not per-AP IP entries.

References: [Cisco WLC Configuration Guides](#), [Cisco WLAN Configuration on WLC \(concepts and steps\)](#)

QUESTION NO: 42

A switch is forwarding a frame out of an interface except the interface that received the frame. What is the technical term for this process?

- A. ARP
- B. CDP
- C. flooding
- D. multicast

ANSWER: C

Explanation:

The technical term for a switch sending a frame out of all ports in the same VLAN except the port it was received on is **flooding**. This behavior is most commonly seen with *unknown unicast* frames (destination MAC not yet in the MAC address table) and with broadcast frames (FF:FF:FF:FF:FF:FF). In both cases, the switch replicates the frame to all other ports in the

VLAN to try to reach the destination (or all hosts, for broadcast), while preventing it from being sent back out the ingress interface.

Option A (ARP) is a protocol used to resolve IPv4 addresses to MAC addresses; while ARP requests are broadcast and therefore get flooded by switches, ARP itself is not the switching process being described. Option B (CDP) is a Cisco Layer 2 discovery protocol and unrelated to frame forwarding behavior. Option D (multicast) refers to a specific destination MAC/IP group delivery method; switches may flood multicast if they lack IGMP snooping state, but “multicast” is not the general term for forwarding out all ports except the source.

References: [Cisco LAN Switching: MAC Address Learning and Flooding](#), [Cisco ARP Overview](#)

QUESTION NO: 43

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:

To form an OSPF neighbor adjacency, both routers must actually run OSPF on the link and they must agree on key OSPF parameters. First, the interfaces connecting R1 and R3 must be made OSPF-active (for example, by using a matching `network` statement under `router ospf` or by enabling OSPF directly on the interface with `ip ospf <pid> area <area>`). If OSPF isn't enabled on the interface, no Hellos are sent/received and no neighbor relationship can form.

Second, the two routers must be in the same OSPF area on that shared link. Area mismatch prevents neighbors from reaching FULL state because the Hello packets include the area ID and must match.

Hello/dead timers must match as well, but they are defaulted to the same values on typical broadcast/point-to-point Ethernet links, so they are not necessarily “next steps” unless they were changed. The OSPF process ID is locally significant and does not need to match between routers. Router IDs must be unique within the OSPF domain; configuring the same router ID would actually create problems rather than fix adjacency.

References: [Cisco OSPF Neighbor Relationship \(support doc\)](#), [Cisco IOS OSPF Command Reference](#)

QUESTION NO: 44

Which two components are needed to create an Ansible script that configures a VLAN on a switch? (Choose two.)

- A. playbook
- B. recipe
- C. model
- D. cookbook
- E. task

ANSWER: A E

Explanation:

To configure a VLAN on a switch using Ansible, you typically write a **playbook** (a YAML file that defines what hosts to target and what automation to run) and within that playbook you define one or more **tasks** (individual actions, usually invoking a module such as a Cisco IOS module to create/configure the VLAN). In other words, the “script” in Ansible terms is the playbook, and the operational steps inside it are tasks.

The other options are terms from different automation/configuration tools, not Ansible. A **cookbook** and **recipe** are Chef concepts, where cookbooks contain recipes that describe desired configuration. A **model** is not a required Ansible component for writing automation; while network automation can involve data models (for example YANG) or structured data, Ansible’s core building blocks for this question are playbooks and tasks.

References: [Ansible documentation: Playbooks](#), [Ansible documentation: Tasks](#).

QUESTION NO: 45

Which action is taken by the data plane within a network device?

- A. forwards traffic to the next hop
- B. constructs a routing table based on a routing protocol
- C. provides CLI access to the network device
- D. looks up an egress interface in the forwarding information base

ANSWER: A

Explanation:

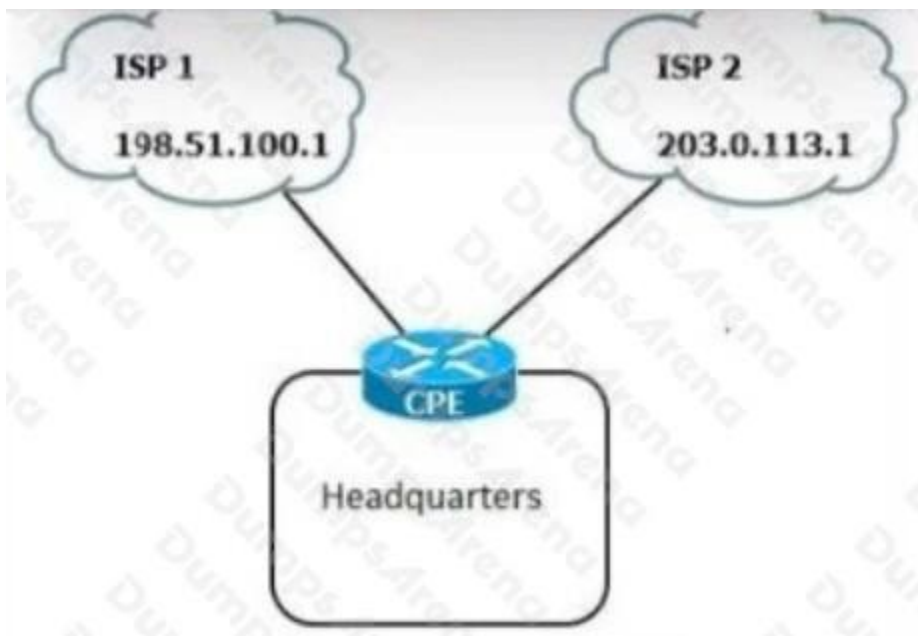
The **data plane** (also called the forwarding plane) is responsible for the real-time handling of user traffic through the device. Its core job is to take an incoming frame/packet and **forward it out the correct egress interface toward the next hop** based on already-built forwarding information. That makes option A the best single answer.

Option B is incorrect because **building the routing table** from routing protocols (OSPF, EIGRP, etc.) is a **control plane** function; the control plane computes routes and installs them into the RIB/FIB. Option C is incorrect because providing **CLI access** is a **management plane** function (SSH, console, AAA, etc.).

Option D describes a mechanism the data plane uses (consulting the **FIB** to determine the egress interface/next hop), but the question is single-choice and asks for an “action taken” by the data plane. The most generally correct, high-level action is that the data plane **forwards traffic**; the FIB lookup is part of how it accomplishes that forwarding.

References: [Cisco IOS Fundamentals \(planes: data/control/management\)](#), [Cisco overview of routing vs forwarding concepts](#).

QUESTION NO: 46



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:

To load-balance Internet-bound traffic across two ISPs using static routing on a Cisco router, you typically configure equal-cost static default routes (ECMP). When two `ip route 0.0.0.0 0.0.0.0` statements point to different next hops and have the same administrative distance (default AD = 1), the router installs both routes in the routing table and performs per-destination (CEF) load sharing while also providing redundancy if one path fails. That is exactly what option C does.

Option D is also valid for load sharing, but it uses “floating” split default routes: one route for 0.0.0.0/1 and one for 128.0.0.0/1, effectively dividing the Internet into two halves and sending each half to a different ISP. This achieves deterministic load distribution and redundancy without relying on ECMP hashing.

Option E is not load balancing: the second default route has a higher AD (2), making it a backup (floating) route only. Option A is inconsistent/misconfigured (mixes high AD defaults plus only one /1 route). Option B combines split routes with defaults, which is unnecessary and can lead to unexpected path selection.

References: [Cisco ECMP load sharing overview](#), [Cisco IOS static route configuration guide](#).

QUESTION NO: 47

What is the benefit of using private IPv4 addressing?

- A. to enable secure connectivity over the Internet
- B. to shield internal network devices from external access
- C. to provide reliable connectivity between like devices
- D. to be routable over an external network

ANSWER: B

Explanation:

Private IPv4 addresses (RFC 1918 ranges) are not globally routable on the public Internet. That characteristic is beneficial because it inherently prevents direct inbound reachability from external networks to hosts using private addresses. In practice, organizations combine private addressing with NAT/PAT at the edge so internal devices can initiate outbound connections while remaining unreachable directly from the Internet unless explicit translations and firewall rules are configured. This “not Internet-routable by default” behavior helps reduce exposure of internal hosts and supports address conservation by allowing many networks to reuse the same private ranges.

Option B is correct because using private addressing helps shield internal devices from direct external access (they can't be routed to across the Internet without NAT and policy). Option A is incorrect because private addressing alone does not provide secure connectivity; security requires mechanisms like VPNs (IPsec/SSL) and proper firewalling. Option C is unrelated to addressing. Option D is the opposite of true: private addresses are specifically *not* routable over external/public networks.

References: [RFC 1918 - Address Allocation for Private Internets](#), [Cisco NAT Overview](#).

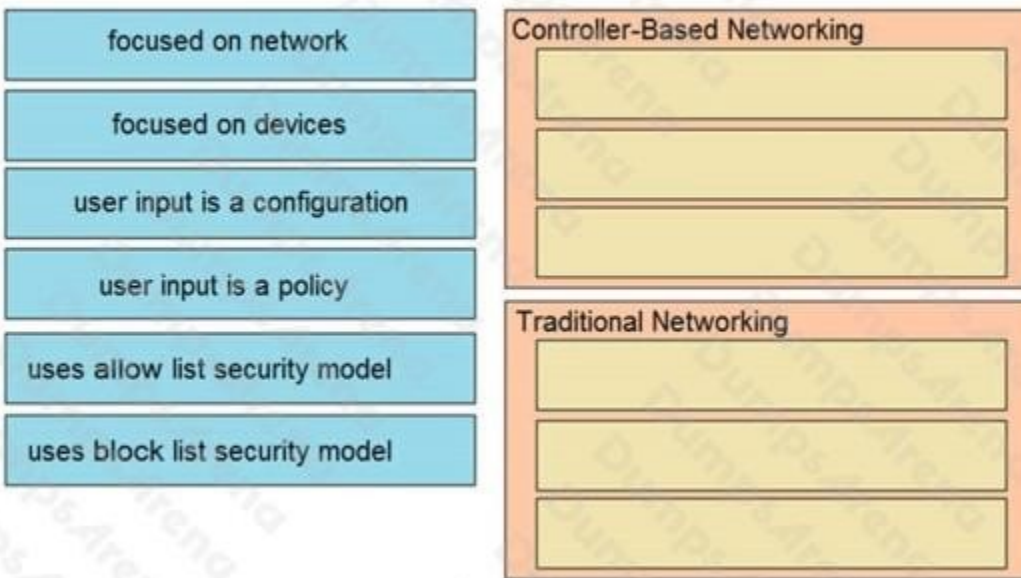
QUESTION NO: 48 - (DRAG DROP)

DRAG DROP

Drag and drop the characteristics of networking from the left onto the correct networking types on the right.

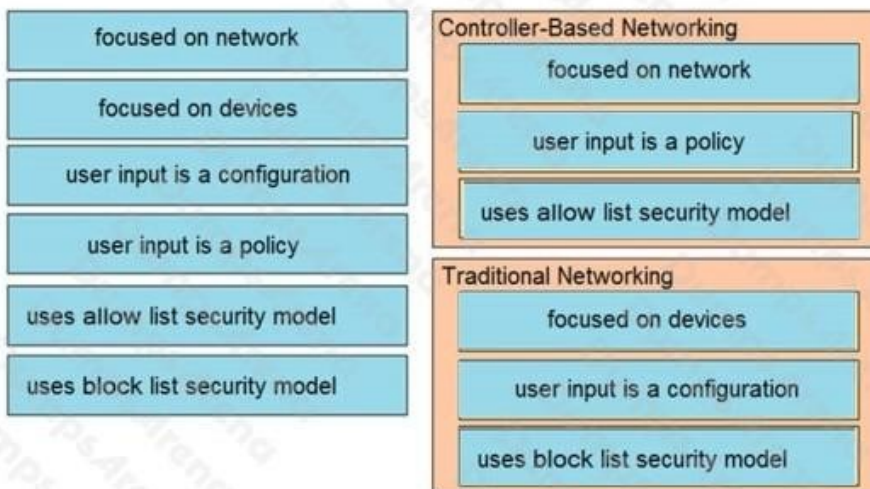
Select and Place:

Answer Area



ANSWER:

Answer Area



Explanation:

In the answered image, the items are grouped into two categories: Controller-Based Networking (SDN-style) and Traditional Networking. This split is mainly about *how* the network is operated and what the administrator provides as “input.” In controller-based networking, you typically express **intent** (policy) and let the controller translate that intent into device-level configuration. That’s why it’s correctly matched with “**focused on network**” and “**user input is a policy.**” The controller provides centralized automation and abstraction, so you manage the network as a system rather than touching each device individually.

Traditional networking is the opposite operational model: you configure each router/switch directly (often via CLI) and troubleshoot per device. That makes “**focused on devices**” and “**user input is a configuration**” the correct matches for Traditional Networking. This aligns with the CCNA view of legacy operations where changes are implemented box-by-box and the desired outcome is achieved by manually crafting configurations.

The security-model pair also matches correctly. Controller-based/intent-based designs commonly align with an **allow-list** (**whitelist**) mindset: only explicitly permitted communications are allowed, which is consistent with modern

policy/segmentation approaches (for example, defining what is allowed between groups). Traditional networks frequently implement security using **block-list (blacklist)** controls such as ACLs that deny specific traffic while allowing the rest by default, which is a classic “block what you don’t want” approach. Cisco’s SDN/controller concepts and traditional device-centric operations are discussed in Cisco’s SDN/controller overviews and enterprise networking resources, such as [Cisco Software-Defined Networking Overview](#) and general Cisco enterprise networking documentation like [What is SDN \(Cisco\)](#).

Because each of the six characteristics is placed under the correct networking type, the provided drag-and-drop answer is correct.

QUESTION NO: 49

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:

REST APIs commonly use standard HTTP verbs to perform CRUD-style actions on resources. **GET** is used to retrieve a representation of a resource (read-only) and is safe/idempotent, which makes it a foundational REST method. **POST** is used to create subordinate resources or to trigger server-side processing when the operation doesn’t fit cleanly into other verbs; it is not idempotent and is also a core REST method.

Options like **REMOVE**, **REDIRECT**, **OPOST**, and **UPOP** are not valid, standardized HTTP methods. In particular, “OPOST” and “UPOP” are not defined HTTP verbs, so they would not be suitable for REST-based API actions. While REST APIs also frequently use other valid methods such as PUT, PATCH, and DELETE, those are not correctly presented among the options here. Therefore, the only suitable choices from the list are GET and POST (with POST added as a corrected option).

References: [MDN Web Docs - HTTP request methods](#), [RFC 9110: HTTP Semantics](#).

QUESTION NO: 50

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:

A public cloud is a cloud deployment model where the cloud infrastructure is owned and operated by a third-party cloud provider (for example, AWS, Microsoft Azure, Google Cloud) and delivered to customers over the Internet. The key idea is multi-tenancy: the provider's underlying physical resources are shared across multiple customers/organizations, while each customer's workloads remain logically isolated. That makes option A correct because it describes provider ownership/maintenance with shared use across multiple organizations.

Option C is also correct because public cloud services are typically consumed over the Internet using provider portals/APIs and public endpoints (even if private connectivity options like VPN/Direct Connect/ExpressRoute can also exist).

Option B is more aligned with private cloud/on-premises control: in public cloud you can configure many services, but you do not "fully customize" the underlying infrastructure the way you can when you own it. Option D describes a single-tenant environment for one company, which is characteristic of a private cloud (or dedicated hosting), not public cloud. Option E describes a hybrid cloud (mix of third-party public resources plus privately owned resources), not a pure public cloud implementation.

References: [NIST Definition of Cloud Computing \(SP 800-145\)](#), [Cisco: What is cloud computing?](#)

QUESTION NO: 51 - (DRAG DROP)

DRAG DROP

-

Drag and drop the characteristic from the left onto the cable type on the right.

Answer Area

uses a single wavelength on light

becomes distorted at longer lengths

transmits data up to 100 Gbits which degrades over distance

has minimal light reflection as it travels down the core

single-mode fiber

multimode fiber

ANSWER:

Answer Area

uses a single wavelength on light	single-mode fiber
becomes distorted at longer lengths	uses a single wavelength on light
transmits data up to 100 Gbits which degrades over distance	has minimal light reflection as it travels down the core
has minimal light reflection as it travels down the core	multimode fiber
	becomes distorted at longer lengths
	transmits data up to 100 Gbits which degrades over distance

Explanation:

To solve this drag-and-drop, you match each description to the fiber type based on how light propagates through the core. **Single-mode fiber (SMF)** is designed so only one mode (one primary light path) propagates down the core. In practice, SMF is typically driven by a laser source and is commonly described in exam questions as using a *single wavelength/light path*. Because there is only one mode, SMF has very low modal dispersion, which is why it's used for long-distance links and high-capacity backbones. That also aligns with the idea that the signal experiences *minimal internal reflections/paths* compared to multimode, where many rays bounce at different angles. So the statements "uses a single wavelength on light" and "has minimal light reflection as it travels down the core" best fit **single-mode fiber**.

Multimode fiber (MMF) has a larger core and allows multiple modes (multiple light paths) to travel simultaneously. Those different paths arrive at slightly different times, causing **modal dispersion**, which shows up as pulse spreading and distortion as distance increases. That's why MMF is often described as becoming *distorted at longer lengths* and why its effective throughput/distance capability *degrades over distance*. While MMF can support very high speeds (including 10G/40G/100G depending on optics and fiber grade), the key CCNA takeaway is that it is more distance-limited than SMF, and performance is constrained by dispersion as the run gets longer.

References: Cisco fiber overview and media characteristics can be cross-checked in Cisco documentation such as [Cisco Fiber-Optic Technology Basics](#) and general fiber mode/dispersion behavior described in [Optical fiber \(modes and dispersion\)](#).

QUESTION NO: 52

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

- A. The administrator can make configuration updates from the CLI
- B. It uses northbound and southbound APIs to communicate between architectural layers
- C. It moves the control plane to a central point.
- D. It decentralizes the control plane, which allows each device to make its own forwarding decisions
- E. It uses Telnet to report system issues.

ANSWER: B C

Explanation:

In a controller-based (SDN) network, the key idea is separating the control plane from the data plane and logically centralizing control. That's why **C** is correct: the controller provides a central (logically centralized) point where control-plane decisions and policy are made, while devices focus on forwarding. Another hallmark is the use of APIs between layers, making **B** correct: *southbound* APIs connect the controller to network devices (for example, to program forwarding behavior), and *northbound* APIs connect the controller to applications/automation systems that express intent and policy.

A is not a defining characteristic of controller-based networking. While you can still use CLI on devices, the architectural characteristic is centralized policy/automation via the controller, not “updates from the CLI.” **D** describes traditional distributed networking where each device runs its own control plane (routing protocols, STP, etc.), which is the opposite of controller-based control-plane centralization. **E** is incorrect because Telnet is not a characteristic mechanism for reporting issues (and is generally discouraged due to lack of encryption); controller-based networks typically use secure management/telemetry methods (e.g., SSH, NETCONF/RESTCONF, streaming telemetry).

References: [Cisco Software-Defined Networking \(SDN\) overview](#), [Cisco APIC and controller-based networking concepts](#)

QUESTION NO: 53

What are two reasons to configure PortFast on a switch port attached to an end host? (Choose two.)

- A.** to block another switch or host from communicating through the port
- B.** to enable the port to enter the forwarding state immediately when the host boots up
- C.** to prevent the port from participating in Spanning Tree Protocol operations
- D.** to protect the operation of the port from topology change processes
- E.** to limit the number of MAC addresses learned on the port to 1

ANSWER: B D

Explanation:

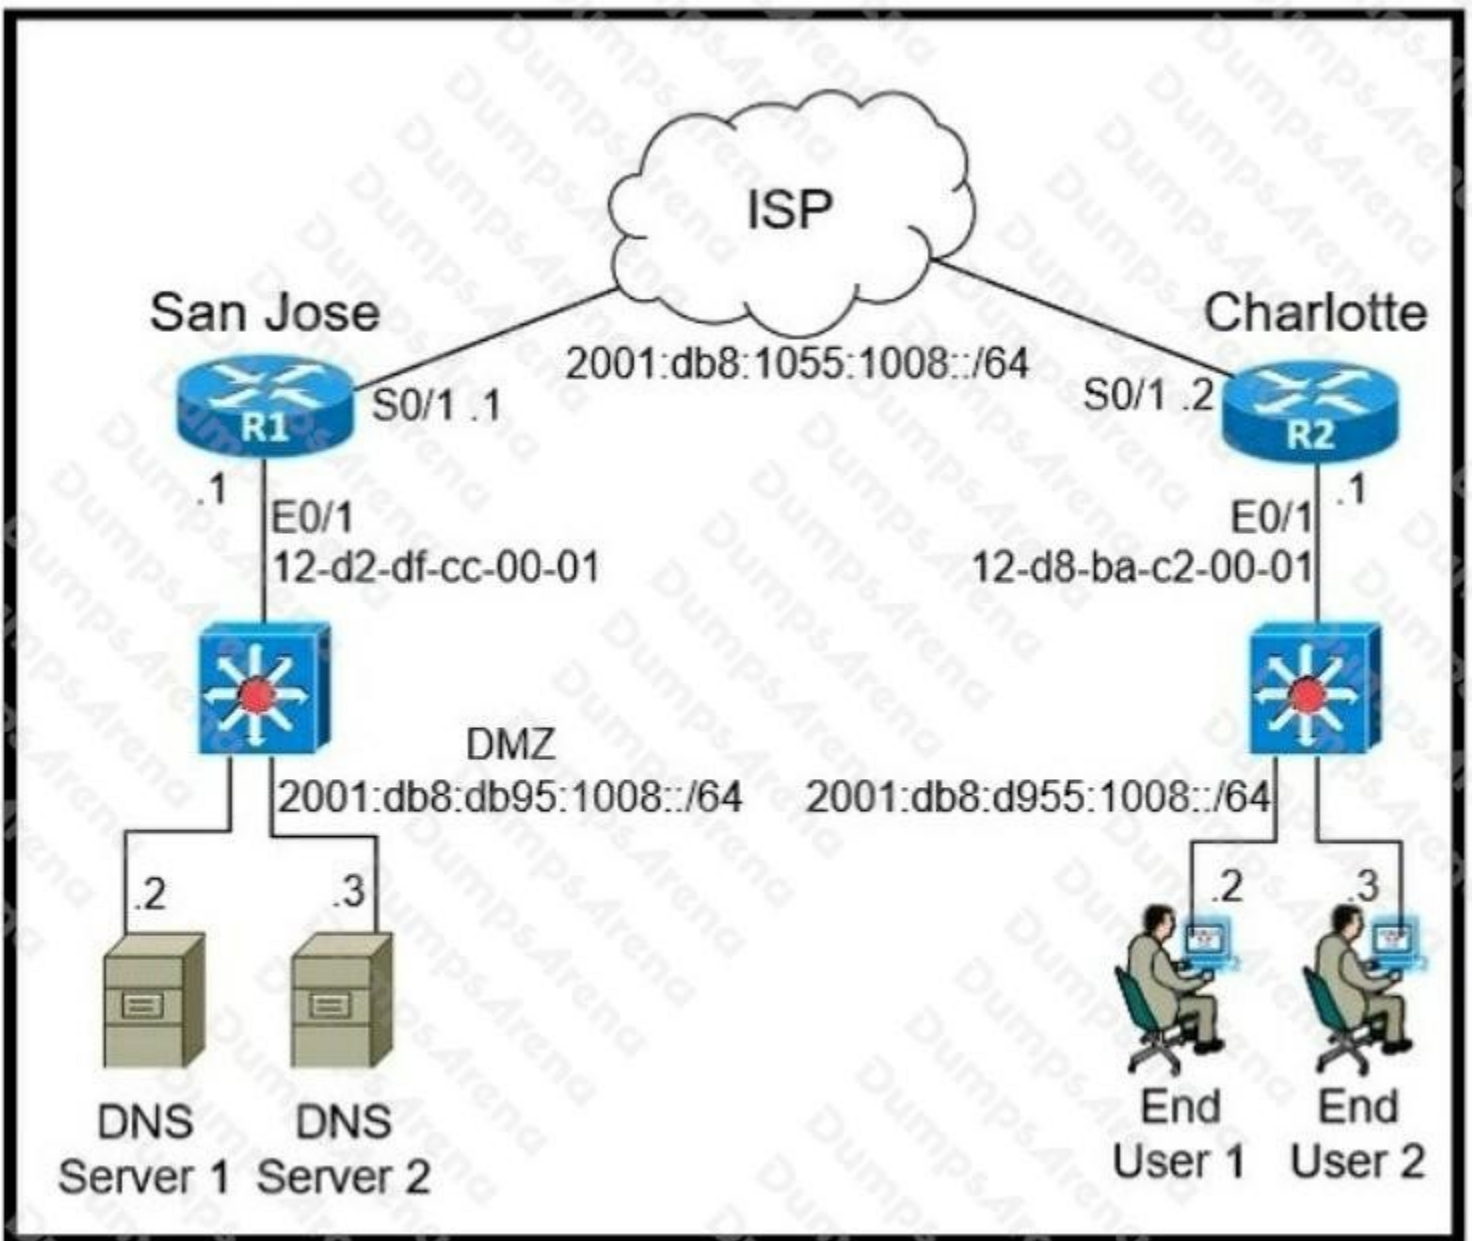
PortFast is designed for access ports that connect to end devices (PCs, printers, servers) so they can start sending/receiving traffic immediately. Normally, STP transitions a port through listening/learning before forwarding, which can delay DHCP and other boot-time traffic. With PortFast enabled, the port skips those transitional states and goes straight to forwarding, which is why option B is correct.

PortFast also changes how STP treats that edge port with respect to topology changes. When an edge (PortFast) port goes up/down due to an end host rebooting or disconnecting, it should not cause the switch to generate topology change notifications (TCNs) that can trigger MAC table aging and unnecessary reconvergence across the LAN. This “insulates” the network from frequent host link flaps, matching option D.

Option C is wrong because PortFast does not disable STP; the port still participates (it will still process BPDUs and can be forced out of PortFast if BPDUs are received, depending on features like BPDU Guard). Option A describes port security or BPDU Guard behavior, not PortFast. Option E is classic port-security (maximum MAC addresses), unrelated to PortFast.

References: [Cisco STP PortFast Configuration and Troubleshooting](#), [Cisco IOS XE Spanning Tree Protocol Configuration Guide](#)

QUESTION NO: 54



Refer to the exhibit. The IPv6 address for the LAN segment on router R2 must be configured using the EUI-64 format. When configured which ipv6 address is produced by the router?

- A. 1:db8:d955:1008:12D8:BAFE:FF01:1
- B. 2001:db8:d955:1008:4598:785F:FE25:1
- C. 2001:db8:d955:1008:1030:ABFF:FECC:1
- D. 2001:db8:d955:1008:10D8:BAFF:FEC2:1

ANSWER: D

Explanation:

With IPv6 EUI-64, the router builds the 64-bit interface ID from the interface's 48-bit MAC address. It does this by splitting the MAC in half, inserting **FFFE** in the middle, and then flipping the **U/L bit** (the 7th bit) of the first byte. In the exhibit, R2's LAN interface MAC is **12:D8:BA:C2:00:01** (as shown). After inserting FFFE you get **12D8:BAFF:FEC2:0001**. Flipping the U/L bit changes the first byte from **0x12** to **0x10**, producing the interface ID **10D8:BAFF:FEC2:0001**. Appending that to the given /64 prefix **2001:db8:d955:1008::/64** yields **2001:db8:d955:1008:10D8:BAFF:FEC2:1** (leading zeros can be compressed, so :0001 becomes :1).

Option A is wrong because it's missing "2001" and uses the wrong "FE" pattern. Option B doesn't match the EUI-64 construction at all. Option C has the right FF:FE pattern but the wrong bytes (it doesn't reflect the MAC shown). Option D matches the correct EUI-64 result.

References: [RFC 4291 \(IPv6 Addressing Architecture\)](#), [Cisco: IPv6 EUI-64 Addressing](#)

QUESTION NO: 55

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 framework that standardizes how network communication functions are described and separated into layers. Option A is accurate because the OSI model defines (at a high level) what kinds of functions belong in each layer (for example, addressing at Layer 3, framing at Layer 2, and physical signaling at Layer 1). Option B is also accurate because the layered model helps people understand and troubleshoot how data is encapsulated and moves through a network from one endpoint to another.

Option C is not accurate as written. While layering promotes modularity, changes in one layer can still affect others (for example, an MTU change at Layer 2 impacts fragmentation/PMTUD behavior at Layer 3/4). The more correct idea is that layers provide well-defined interfaces so implementations can change without necessarily requiring changes in other layers, but "do not impact" is too absolute. Option D is incorrect because the OSI model itself does not "ensure" reliable delivery; reliability is provided by specific protocols and mechanisms (for example, TCP at the transport layer), and many network communications are intentionally unreliable (UDP).

References: [Cisco: Internetworking Basics \(OSI reference\)](#), [OSI model overview](#).

QUESTION NO: 56

Refer to the exhibit. When running EIGRP, what is required for RouterA to exchange routing updates with RouterC?



- A. AS numbers must be changed to match on all the routers
- B. Loopback interfaces must be configured so a DR is elected
- C. The no auto-summary command is needed on Router A and Router C
- D. Router B needs to have two network statements, one for each connected network

ANSWER: A

Explanation:

For Router A to exchange EIGRP routing updates with Router C, the routers must be able to form EIGRP neighbor adjacencies across the path (A–B and B–C). A fundamental requirement for EIGRP neighbors is that they run EIGRP under the same autonomous system (AS) number on the shared link; if the AS numbers don't match, the routers will not become neighbors and no updates will be exchanged end-to-end. Therefore, the AS numbers must be changed to match on all routers participating in that EIGRP domain.

The other options don't address EIGRP adjacency requirements. EIGRP does not elect a DR/BDR (that's OSPF), so loopbacks for DR election are irrelevant. The `no auto-summary` command is not required to form neighbors; it only affects classful route summarization behavior (and in modern IOS, auto-summary is typically disabled by default anyway). Finally, while correct `network` statements are necessary to enable EIGRP on interfaces, the statement "Router B needs two network statements" is not inherently required—one statement can cover multiple interfaces depending on the wildcard, and the key issue for update exchange is matching AS/neighborship.

References: [Cisco EIGRP Technical Overview](#), [Cisco EIGRP Configuration/Support Docs](#)

QUESTION NO: 57

On a corporate network, hosts on the same VLAN can communicate with each other, but they are unable to communicate with hosts on different VLANs. What is needed to allow communication between the VLANs?

- A. a router with subinterfaces configured on the physical interface that is connected to the switch
- B. a router with an IP address on the physical interface connected to the switch
- C. a switch with an access link that is configured between the switches
- D. a switch with a trunk link that is configured between the switches

ANSWER: A

Explanation:

Hosts in different VLANs are in different Layer 2 broadcast domains, so a Layer 2 switch will not forward traffic between them. To enable inter-VLAN communication, you need a Layer 3 device to perform routing between the VLAN subnets. One common CCNA solution is "router-on-a-stick": a router connects to the switch over a single physical link configured as an 802.1Q trunk, and the router uses one subinterface per VLAN. Each subinterface is tagged for its VLAN (encapsulation dot1q X) and has an IP address that serves as the default gateway for that VLAN. That is exactly what option A describes.

Option B is insufficient because a single IP address on the router's physical interface would only provide Layer 3 connectivity for one VLAN (typically the native/untagged VLAN) unless subinterfaces (or SVIs on a multilayer switch) are used to handle multiple VLANs over the trunk. Options C and D describe switch-to-switch links; while trunks are used to carry multiple VLANs between switches, they do not provide Layer 3 routing between VLANs by themselves.

References: [Cisco Support: Inter-VLAN Routing \(Router-on-a-Stick\)](#), [Cisco Support: IEEE 802.1Q Trunking](#)

QUESTION NO: 58

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

ANSWER: B D

Explanation:

In SNMPv2, the key enhancement for efficiently retrieving large amounts of management data is the **GetBulk** operation. GetBulk was introduced specifically to reduce the number of request/response exchanges needed to walk large portions of the MIB (for example, pulling many rows from a table). It does this by allowing the manager to request multiple lexicographically successive variables in a single PDU, effectively “batching” what would otherwise require many GetNext operations.

While **GetNext** is not new to SNMPv2 (it exists in SNMPv1 as well), it is still a mechanism used to retrieve sequences of objects (such as walking a table) and, in practice, is associated with retrieving lots of data—just less efficiently than GetBulk. By contrast, **Get** retrieves specific OIDs (not optimized for bulk retrieval), and **Set** is for writing values, not reading. **Inform** is an SNMPv2 notification type (manager-to-manager) that expects an acknowledgment; it’s unrelated to bulk data retrieval.

References: [RFC 1905 - SNMPv2 Protocol Operations \(GetBulk\)](#), [RFC 3416 - SNMPv2 Operations \(Get, GetNext, GetBulk\)](#)

QUESTION NO: 59



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 requirements imply two different IPv6 address assignment methods on two different interfaces. The uplink to the ISP must be a manual (static) IPv6 address, which on Cisco IOS is done by configuring an explicit IPv6 address/prefix on the WAN-facing interface (no EUI-64 keyword). The LAN interface must be "self-provisioned," which in this context means using EUI-64 so the router derives the interface ID from the MAC address while still using the correct /64 LAN prefix.

Option B correctly applies EUI-64 on Gi0/0 with a /64 prefix, which matches the typical LAN requirement for SLAAC/EUI-64-based addressing. Options A, C, and E attempt to configure a /127 with a full 128-bit-looking address but include an extra colon before "/127" (invalid syntax as written). Also, a /127 is commonly used on point-to-point links, but the address must be syntactically correct and match the exact WAN subnet shown in the exhibit. Option D uses EUI-64 but with a different /64 prefix than option B, so it would not use the applicable LAN network. Option F is not a configuration at all.

References: [Cisco IPv6 Configuration Guide \(IOS\)](#), [RFC 6164 \(/127 for IPv6 point-to-point links\)](#)

QUESTION NO: 60

```
{
  "Interfaces": ["ethernet0/3", "ethernet0/4", "ethernet0/5"]
}
```

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

- A. Boolean
- B. string
- C. object

D. sequence

ANSWER: C

Explanation:

The exhibit is showing a JSON **object**. In JSON, an object is represented by curly braces { } and contains one or more *name/value* pairs (members). Each member name is a string in double quotes, followed by a colon, followed by a value (which can be a string, number, boolean, null, array, or another object). This structure is the most common way APIs return structured data such as device attributes (for example, interface details, hostname, or configuration elements) in Cisco automation contexts.

Option A (Boolean) is incorrect because a boolean value in JSON is only `true` or `false`, not a collection of key/value pairs. Option B (string) is incorrect because a JSON string is a single quoted value like `"GigabitEthernet0/1"`, not a brace-delimited structure. Option D (sequence) is not a standard JSON data type name; JSON uses the term **array** (written with square brackets []) for ordered lists. Since the exhibit uses { } and key/value members, it is clearly an object.

References: [RFC 8259 \(The JavaScript Object Notation \(JSON\) Data Interchange Format\)](#), [MDN JSON documentation](#).

QUESTION NO: 61

Router A learns the same route from two different neighbors, one of the neighbor routers is an OSPF neighbor and the other is an EIGRP neighbor. What is the administrative distance of the route that will be installed in the routing table?

- A. 20
- B. 90
- C. 110
- D. 115

ANSWER: B

Explanation:

When a router learns the same destination prefix from multiple routing sources (different routing protocols, static routes, etc.), it chooses which route to install in the routing table based first on **administrative distance (AD)**. The route with the **lowest AD** is considered the most trustworthy and is installed (assuming the routes are otherwise valid and the prefix length is the same).

In this scenario, Router A learns the same route via OSPF and via EIGRP. By default, Cisco routers use an AD of **90** for **internal EIGRP** routes and an AD of **110** for **OSPF** routes. Since 90 is lower than 110, the EIGRP-learned route wins and the route installed in the routing table will therefore have an administrative distance of **90**.

Why the other options are wrong: **20** is the default AD for eBGP, not OSPF/EIGRP. **110** would apply only if OSPF were the best (lower AD) source, which it is not here. **115** is not the default AD for either OSPF or EIGRP (EIGRP external is 170).

References: [Cisco: Administrative Distance \(Routing Protocol Preference\)](#), [Cisco IOS XE IP Routing: Route Overview \(AD selection\)](#)

QUESTION NO: 62

An office has 8 floors with approximately 30-40 users per floor. One subnet must be used. Which command must be configured on the router Switched Virtual Interface to use address space efficiently?

- A. ip address 192.168.0.0 255.255.0.0
- B. ip address 192.168.0.0 255.255.254.0

C. ip address 192.168.0.0 255.255.255.128

D. ip address 192.168.0.0 255.255.255.224

ANSWER: B

Explanation:

You need a single subnet that can accommodate all hosts across 8 floors. With ~30–40 users per floor, that’s roughly 240–320 endpoints total (and in practice you also need room for printers, phones, APs, and the SVI default gateway). So the subnet must support well over 254 usable IPs.

A /23 mask (255.255.254.0) provides 512 total addresses and 510 usable host addresses, which comfortably covers 240–320+ devices while still being more efficient than a /16. A /24 (255.255.255.0) would only provide 254 usable addresses and would likely be too small; /25 (255.255.255.128) provides 126 usable; /27 (255.255.255.224) provides 30 usable—both far too small.

Note: while the options show 192.168.0.0, in real Cisco IOS you would typically assign a usable host address to the SVI (for example 192.168.0.1 255.255.254.0), not the network address itself. But among the given masks, /23 is the only one that meets the host requirement and uses address space reasonably.

References: [Cisco IP Subnetting Reference](#), [CIDR and subnet sizing overview](#)

QUESTION NO: 63

SIP-based Call Admission Control must be configured 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 interlace of the WLC.
- E. Configure two different QoS rates for data and voice traffic.

ANSWER: C E

Explanation:

For SIP-based CAC on a Cisco WLC, after defining the SIP call-snooping ports, you must enable the features that actually inspect SIP signaling and track the associated RTP media flows per WLAN. Specifically, you enable SIP call snooping (signaling inspection) and Media Session Snooping (media/RTP session tracking) on the WLAN where voice clients connect. These two WLAN-level settings allow the controller to learn active calls and enforce admission control based on bandwidth/call limits.

QoS settings (silver/platinum) are important for voice quality, but they are not the required “next steps” to make SIP CAC function; CAC can operate independently of choosing a particular QoS profile, and “silver or greater” is not a defined CAC prerequisite. “Traffic shaping for the LAN interface” is not a standard required step for SIP CAC on WLC; CAC is enforced by call counting/bandwidth accounting rather than enabling generic shaping on the wired uplink. Likewise, “configure two different QoS rates for data and voice traffic” is not a typical WLC SIP CAC requirement and is ambiguously worded.

References: Cisco WLC QoS and WLAN QoS profiles overview ([Cisco QoS on WLC](#)), and Cisco Wireless Controller configuration guides landing page for feature configuration context ([Cisco WLC Configuration Guides](#)).

QUESTION NO: 64

Which two QoS tools provide congestion management? (Choose two.)

- A. CBWFQ
- B. FRTS
- C. CAR
- D. PBR
- E. PQ

ANSWER: A E

Explanation:

Congestion management refers to QoS mechanisms that decide how packets are queued and scheduled for transmission when an interface is congested (output is oversubscribed). In Cisco QoS, this is primarily done with queuing/scheduling tools.

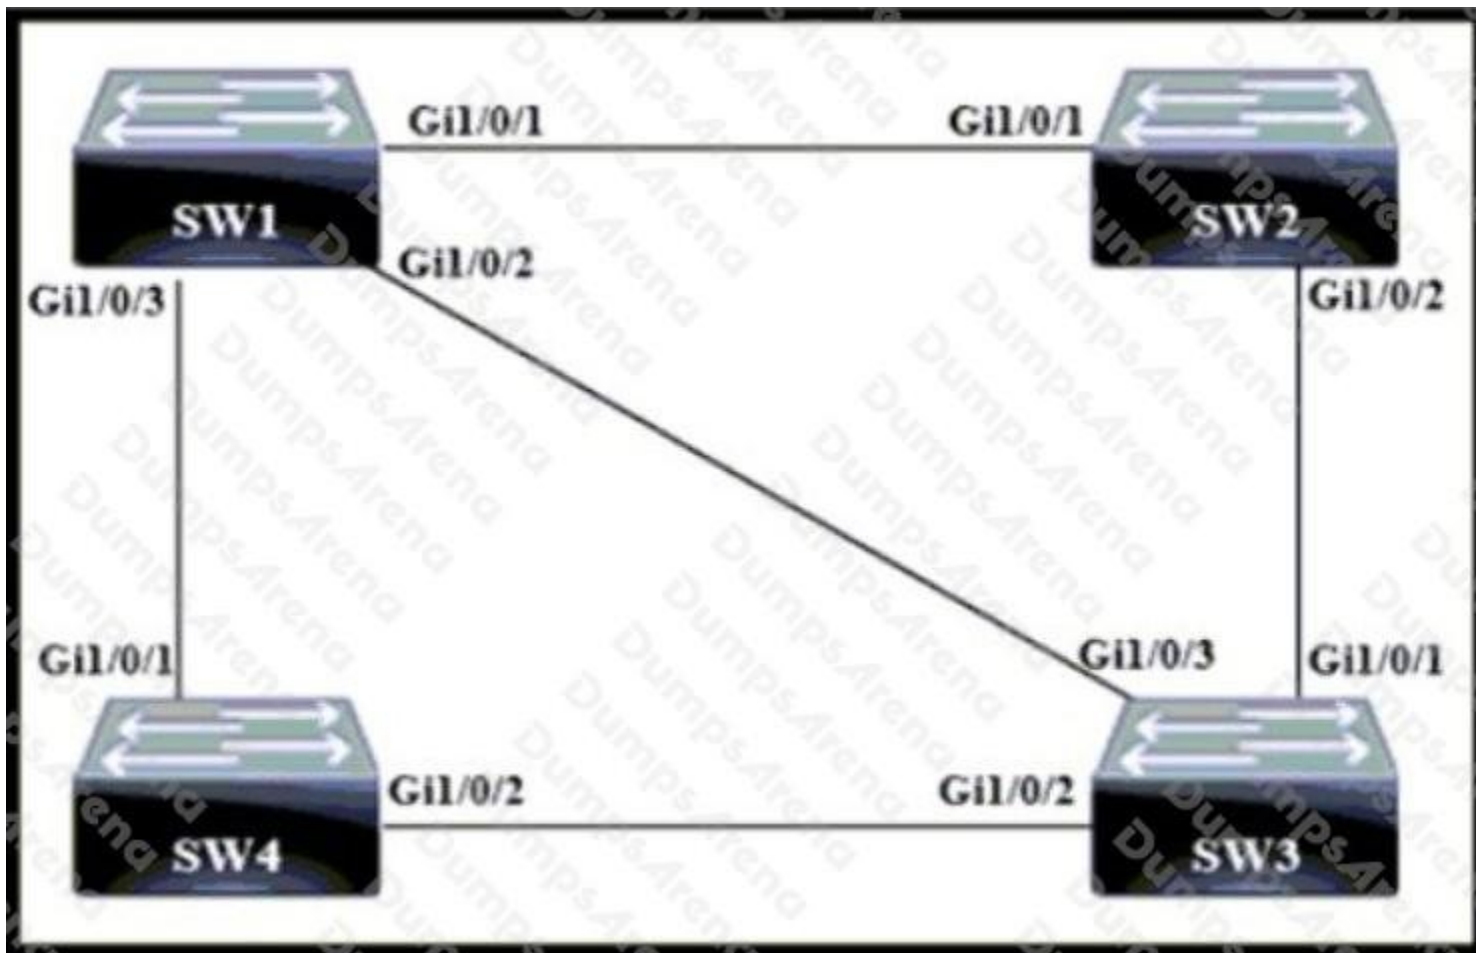
CBWFQ (Class-Based Weighted Fair Queuing) is a congestion-management tool because it creates separate queues per class and services them according to configured bandwidth/weights. It's used to manage congestion by ensuring each traffic class gets a predictable share of the link during contention.

PQ (Priority Queuing) is also congestion management. It places traffic into priority levels and always services higher-priority queues first, which directly controls how congestion is handled (though it can starve lower queues if overused).

FRTS (Frame Relay Traffic Shaping) is congestion *avoidance/conditioning* via shaping to a CIR, not a queuing scheduler for general interface congestion management. **CAR** (Committed Access Rate) is policing (rate limiting), which is traffic conditioning rather than congestion management. **PBR** (Policy-Based Routing) influences path selection/forwarding decisions, not queuing or scheduling under congestion.

References: [Cisco QoS Tools Overview](#), [Cisco Queuing Overview \(Congestion Management\)](#)

QUESTION NO: 65



Refer to the exhibit. Which switch becomes the root bridge?

A. SW 4 -

Bridge Priority - 49152 -
mac-address 06:8e:bc:7e:5b:85

B. SW 3 -

Bridge Priority - 49152 -
mac-address 0d:e4:96:da:ee:95

C. SW 1 -

Bridge Priority - 36864 -
mac-address 05:a7:23:5b:52:25

D. SW 2 -

Bridge Priority - 36864 -
mac-address 04:1e:c4:bf:02:55

ANSWER: D

Explanation:

In STP/RSTP, the root bridge is the switch with the lowest Bridge ID (BID). The BID is compared first by bridge priority (including the extended system ID/VLAN component), and if there's a tie, by the lowest MAC address.

From the options, SW1 and SW2 both have the lowest bridge priority value (36864). SW3 and SW4 have a higher priority (49152), so they cannot be elected root as long as a lower-priority switch exists. With the priority tie between SW1 and SW2, STP breaks the tie using the MAC address: the lower MAC wins. SW2's MAC (04:1e:c4:bf:02:55) is lower than SW1's MAC (05:a7:23:5b:52:25), so SW2 has the lowest BID overall and becomes the root bridge.

Why the others are wrong: SW1 loses only because its MAC is higher than SW2's under the same priority. SW3 and SW4 lose immediately because their bridge priority is higher, regardless of MAC.

References: [Cisco STP Root Bridge Election \(bridge ID/priority/MAC\)](#), [Spanning Tree Protocol \(Bridge ID election rules\)](#).

QUESTION NO: 66

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:

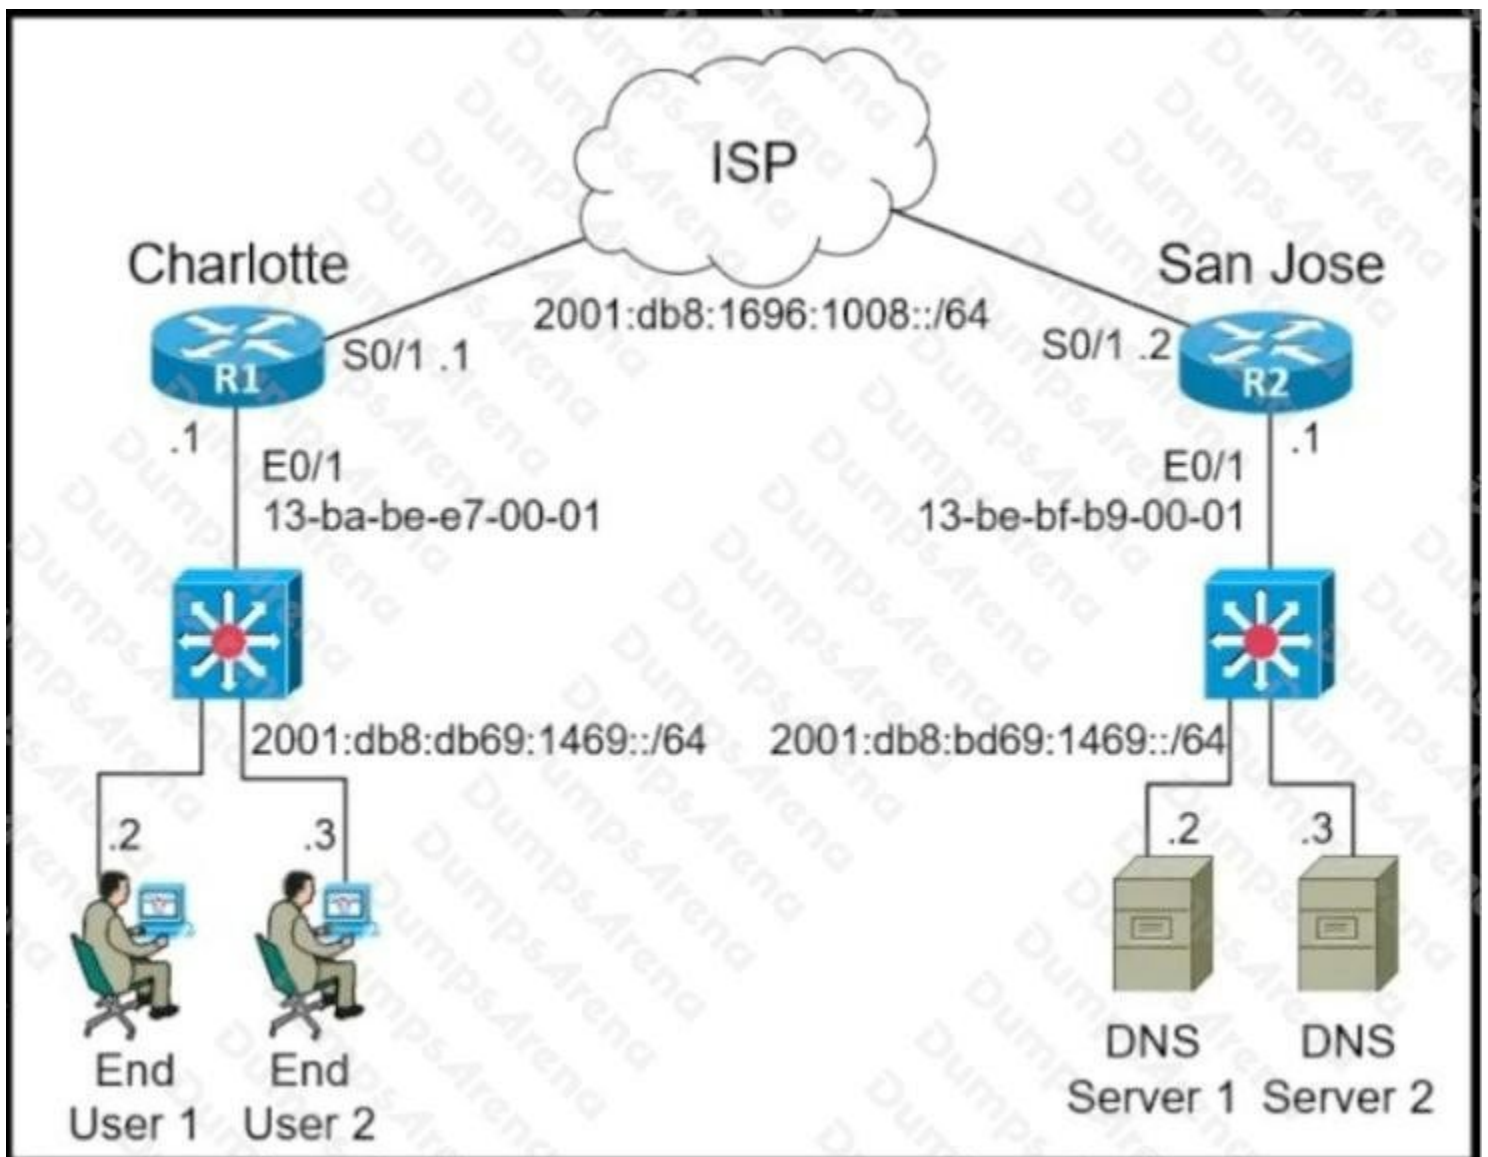
From a router's routing table (for example, via `show ip route`), you can learn details about each installed route such as its administrative distance/metric and how long the route has been in the table. Cisco IOS displays routes in a format like `D 10.1.1.0/24 [90/2172416] via 10.0.0.2, 00:12:33, GigabitEthernet0/0`. The bracketed values include the administrative distance (AD) and metric, and the timestamp (e.g., `00:12:33`) indicates the age—how long the route has been known/installed—so option E is correct.

Option C is also correct in the practical sense that the routing table reveals the AD value being used for that route (the first number in brackets). While the table doesn't explicitly say "manually configured" vs "default," you can determine whether AD has been altered by comparing the displayed AD to the protocol's default (e.g., EIGRP internal 90, OSPF 110).

Options A and D are not shown in the routing table; ACL direction and neighbor adjacencies are verified with `interface/ACL` and `routing-protocol neighbor` commands (e.g., `show ip interface`, `show ip ospf neighbor`). Option B is not learned from the routing table; EIGRP/BGP AS information is seen in protocol-specific outputs (e.g., `show ip protocols`, `show ip bgp summary`).

References: [Cisco - Understanding and Interpreting the Routing Table](#), [Cisco IOS Command Reference - show ip route](#)

QUESTION NO: 67



Refer to the exhibit The IPv6 address for the LAN segment on router R2 must be configured using the EUI-64 format. When configured which ipv6 address is produced by the router?

- A. 1:db8:bd69:1469:12D8:BAFE:FF01:1
- B. 2001:db8:bd69:1469:1130:ABFF:FECC:1
- C. 2001:db8:bd69:1469:4628:255F:FE32:1
- D. 2001:db8:bd69:1469:11BE:BFFF:FEB9:1

ANSWER: D

Explanation:

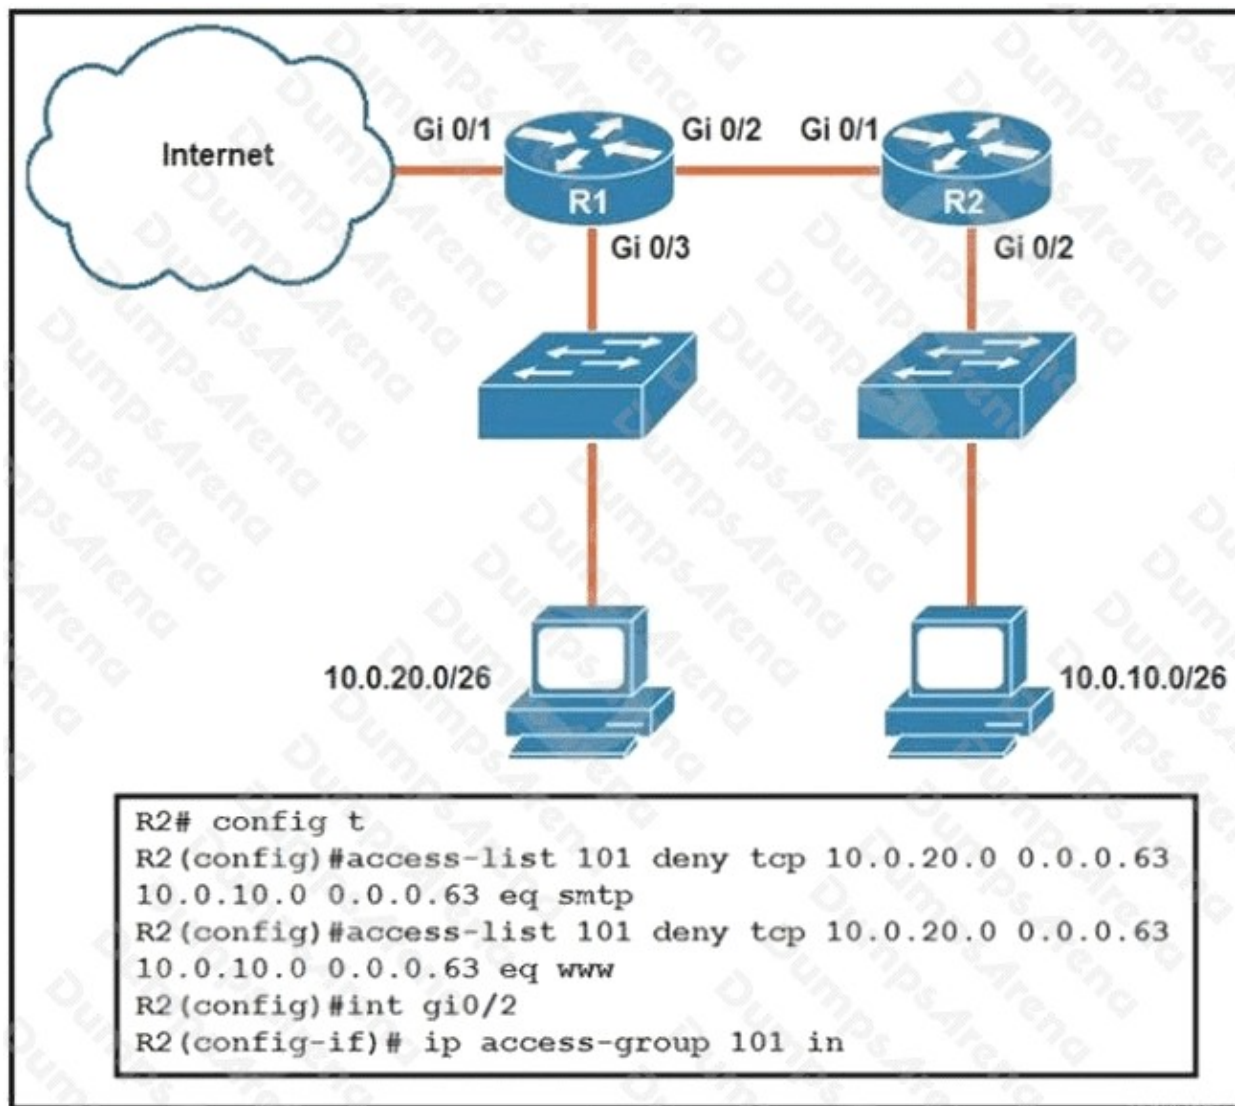
With IPv6 EUI-64, the router builds the 64-bit interface ID from the interface MAC address. It takes the 48-bit MAC, inserts **FFFE** in the middle (between the OUI and NIC portions), and flips the **U/L bit** (the 7th bit) of the first byte. That “bit flip” is what changes the first byte value and is the key detail many wrong answers miss.

From the exhibit, R2’s LAN interface MAC is **13:BE:BF:B9:00:01**. Inserting FFFE yields **13BE:BFFF:FEB9:0001**. Flipping the U/L bit in the first byte (0x13) changes it to **0x11**, producing the interface ID **11BE:BFFF:FEB9:0001**. Combined with the given /64 prefix **2001:DB8:BD69:1469::/64**, the resulting IPv6 address is **2001:db8:bd69:1469:11BE:BFFF:FEB9:1**, which matches option D.

Options A, B, and C are incorrect because their interface IDs don't correctly reflect the required EUI-64 steps (FFFE insertion and/or U/L bit inversion) based on the MAC shown in the exhibit.

References: [RFC 4291 - IPv6 Addressing Architecture](#), [Cisco: IPv6 EUI-64 Addressing](#)

QUESTION NO: 68



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 at the end of ACL 101 for allowed traffic.
- B. Add a "permit ip any any" statement to the beginning of ACL 101 for allowed traffic.
- C. The ACL must be moved to the Gi0/1 interface outbound on R2.
- D. The source and destination IPs must be swapped in ACL 101.
- E. The ACL must be configured the Gi0/2 interface inbound on R1.

ANSWER: A D

Explanation:

To block only TCP ports 25 (SMTP) and 80 (HTTP) from 10.0.10.0/26 to 10.0.20.0/26 while allowing everything else, two ACL fundamentals must be satisfied. First, the ACEs must match the correct traffic direction: in an extended ACL, the *source* should be the 10.0.10.0/26 subnet and the *destination* should be 10.0.20.0/26. If these are reversed, the deny statements won't match the intended flows, so swapping source/destination fixes the logic. Second, Cisco ACLs have an implicit "deny ip any any" at the end. If you only add deny statements and don't explicitly permit the rest, you will unintentionally block all other traffic. Adding "permit ip any any" at the end ensures only the specified TCP ports are denied and all other IP traffic is allowed.

Placing "permit ip any any" at the beginning would override the denies (everything would be permitted). Moving the ACL to a different interface is not required to meet the stated goal if it's already applied in the correct path; the described failure is more consistent with incorrect ACE matching and missing explicit permit. See Cisco ACL behavior and implicit deny details in [Cisco ACL configuration notes](#) and general ACL concepts in [Cisco IP ACL configuration guide](#).

QUESTION NO: 69

Why choose Cisco DNA Center for automated lifecycle management?

- A. to allow SSH access to all nodes in the network
- B. to provide software redundancy in the network
- C. to perform upgrades without service interruption
- D. to provide fast and accurate deployment of patches and updates

ANSWER: D

Explanation:

Cisco DNA Center is commonly chosen for automated lifecycle management because it centralizes and automates ongoing operational tasks such as image management, software maintenance, and compliance. In particular, DNA Center can orchestrate distribution and activation of network device software images and help push updates consistently across many devices, reducing manual effort and human error. That aligns best with "fast and accurate deployment of patches and updates," which is a core lifecycle-management value proposition.

Option A is incorrect because DNA Center is not primarily about enabling SSH access to all nodes; SSH is a device access method, while DNA Center focuses on intent-based automation and assurance. Option B is incorrect because "software redundancy" is not a lifecycle-management driver or a primary DNA Center feature; redundancy is typically addressed via network design and HA mechanisms, not via DNA Center's lifecycle workflows. Option C is incorrect because DNA Center can help schedule and automate upgrades, but it cannot universally guarantee "upgrades without service interruption" (that depends on platform capabilities like ISSU, topology, and maintenance windows). DNA Center improves the process, not the physics of downtime.

References: [Cisco DNA Center overview](#), [Cisco DNA Center Platform documentation](#)

QUESTION NO: 70

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:

On Cisco switches/routers, an interface state of **down/down** means the physical layer is down and, as a result, the line protocol is also down. Two common administrative/operational causes on a switch are: (1) the port is **administratively shut down** (shown as *administratively down/down* in many outputs), and (2) the port has been placed into an **err-disabled** state due to a detected violation (for example, port-security, BPDU Guard, UDLD, etc.). In both cases, the interface is not forwarding traffic and will present as down at Layer 1/2 until it is re-enabled (e.g., `shutdown/no shutdown`) or the err-disable condition is cleared and recovery occurs.

By contrast, a **duplex mismatch** or **speed mismatch** typically still allows the physical link to come up (often resulting in poor performance, errors, or flapping), so it's less associated with a steady down/down condition. A generic "protocol mismatch" is not a typical Ethernet cause for down/down; protocol issues more commonly show as **up/down** (physical up, line protocol down) in other contexts.

References: [Cisco: Troubleshooting Ethernet Interfaces](#), [Cisco: Understanding and Configuring Errdisable Recovery](#)

QUESTION NO: 71

A network engineer is replacing the switches that belong to a managed-services client with new Cisco Catalyst switches. The new switches will be configured for updated security standards including replacing.

Telnet services with encrypted connections and doubling the modulus size from 1024. Which two commands must the engineer configure on the new switches?

(Choose two.)

- A. `transport input ssh`
- B. `transport input all`
- C. `crypto key generate rsa modulus 2048`
- D. `crypto key generate rsa general-keys modulus 1024`
- E. `crypto key generate rsa usage-keys`

ANSWER: A C

Explanation:

To replace Telnet with encrypted remote access on Cisco Catalyst switches, you must allow SSH on the VTY lines and (typically) restrict inbound protocols so Telnet is not accepted. The command **transport input ssh** under **line vty** ensures only SSH is permitted for inbound remote sessions, effectively disabling Telnet access on those lines. To meet the requirement of "doubling the modulus size from 1024," you must generate RSA keys with a 2048-bit modulus. The command **crypto key generate rsa modulus 2048** creates the RSA key pair used by the SSH server process.

Option B (**transport input all**) is the opposite of the requirement because it allows multiple protocols (commonly including Telnet and SSH), so it does not enforce encrypted-only access. Option D generates 1024-bit keys, which fails the "doubling" requirement. Option E is incomplete as written because it omits the modulus size; also, "usage-keys" is a different key type selection and still requires proper parameters to meet the 2048-bit requirement.

References: [Cisco Secure Shell \(SSH\) overview/config guidance](#), [Cisco IOS XE Secure Shell Configuration Guide](#).

QUESTION NO: 72

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:

CCMP (Counter Mode with Cipher Block Chaining Message Authentication Code Protocol) is the AES-based security protocol introduced with IEEE 802.11i and is the required data confidentiality/integrity mechanism for WPA2. In other words, when you see "WPA2 (AES)", that's CCMP providing both encryption (AES in counter mode) and integrity (CBC-MAC). WPA3 continues to use AES-CCMP as the baseline data protection method (with WPA3-Personal using SAE for authentication and WPA3-Enterprise adding stronger options like 192-bit mode). So both WPA2 and WPA3 are associated with CCMP for encryption and data integrity.

Wi-Fi 6 (802.11ax) is a PHY/MAC amendment (performance/efficiency), not a security standard; it can run with WPA2 or WPA3, but it doesn't itself define CCMP. WEP is the legacy RC4-based scheme with weak integrity (CRC-32) and does not use CCMP. WPA (original WPA) primarily used TKIP (also RC4-based) rather than CCMP, although some transitional "WPA with AES" existed, it's not the standard association tested at CCNA level.

References: [https://en.wikipedia.org/wiki/CCMP_\(cryptography\)](https://en.wikipedia.org/wiki/CCMP_(cryptography)), <https://www.wi-fi.org/discover-wi-fi/security>

QUESTION NO: 73

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

- A. int range g0/0-1
channel-group 10 mode active
- B. int range g0/0-1 chanm.l-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:

To create an EtherChannel Port-Channel 10, you place the member interfaces into channel-group 10 with a valid negotiation mode that matches the protocol you intend to use. In this question, the two correct choices are the LACP modes **active** and **passive**. LACP forms an EtherChannel when at least one side is **active** (initiates negotiation); the other side can be either **active** or **passive**. Therefore, using `channel-group 10 mode active` on one switch and `channel-group 10 mode passive` on the other (on the same member ports) will successfully create Port-Channel 10.

Option B (**desirable**) and option D (**auto**) are PAgP modes, not LACP. They can form an EtherChannel only with compatible PAgP settings (desirable/auto), and they won't interoperate with LACP modes. Option E (**on**) forces a static EtherChannel with no negotiation; while it can create a port-channel, it must be configured as **on** on both ends and is not the negotiated setup implied by the exhibit-style question. Cisco generally tests LACP/PAgP mode matching explicitly.

References: [Cisco EtherChannel and LACP/PAgP configuration overview](#), [Cisco LACP configuration guide](#)

QUESTION NO: 74

Which two events occur automatically when a device is added to Cisco DNA Center? (Choose two.)

- A. The device is placed into the Managed state.
- B. The device is placed into the Unmanaged state.
- C. The device is assigned to the Local site.
- D. The device is assigned to the Global site.
- E. The device is placed into the Provisioned state.

ANSWER: A D

Explanation:

When you add (discover/import) a network device into Cisco DNA Center, two things happen by default: the device is brought under management by Cisco DNA Center and it is placed into the default site hierarchy. In practice, newly added devices show up as *Managed* (DNA Center starts monitoring/collecting inventory and can perform operations using the provided credentials) and they are associated to the *Global* site unless you explicitly assign them to another site. That "Global" site is the root container in the site hierarchy and is the automatic landing spot for devices that haven't been placed into a specific building/floor/site.

By contrast, *Unmanaged* is not the normal automatic outcome of adding a device; devices become unmanaged when you intentionally change their management state or remove credentials/access. Likewise, *Provisioned* is a later lifecycle step that requires explicit provisioning actions (assigning to a site, pushing settings/images, etc.)—it does not occur merely by adding the device. "Local site" is not the default; it's a specific site you would manually assign.

References: [Cisco DNA Center overview](#), [Cisco DNA Center Platform documentation](#).

QUESTION NO: 75

What are two reasons a switch experiences frame flooding? (Choose two.)

- A. A defective patch cable is connected to the switch port
- B. Topology changes are occurring within spanning-tree
- C. An aged MAC (able entry is causing excessive updates

- D. Port-security is configured globally
- E. The forwarding table has overflowed

ANSWER: B E

Explanation:

A switch floods frames when it cannot make a specific forwarding decision based on its MAC address table (CAM table). The most common case is an unknown unicast: the destination MAC isn't in the table yet, so the switch forwards the frame out all ports in the VLAN except the ingress port. This can happen if the MAC table is being cleared or entries are being removed frequently. Spanning Tree topology changes can contribute to this behavior because a topology change triggers MAC address table aging/flush behavior (implementation-dependent), which temporarily increases unknown unicasts and therefore flooding until the table is relearned.

Another classic cause is a MAC address table overflow (CAM table exhaustion). If the table is full, the switch can't learn additional source MAC addresses, which again increases unknown unicast forwarding and flooding. A defective patch cable is not a typical root cause of flooding (it more often causes errors, link flaps, or loss), and "aged MAC entry causing excessive updates" is not a standard flooding cause—aging out simply leads to normal relearning. Port-security being configured globally does not inherently cause flooding; it restricts/controls MAC learning and may err-disable ports, but flooding isn't the expected outcome.

References: [Cisco LAN Switching: Understanding MAC Address Tables](#), [Cisco Spanning Tree Protocol and Topology Changes](#)

QUESTION NO: 76

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:

After deploying a new WLAN controller (WLC), two common follow-up tasks are ensuring the wired infrastructure supports the APs and ensuring the network segmentation exists for wireless traffic. First, you should consider deploying PoE-capable access-layer switches because most enterprise APs require Power over Ethernet to operate; without PoE (or injectors), AP rollout and placement become difficult and may fail entirely. Second, you should consider configuring additional VLANs (and the associated trunking/routing) because WLANs/SSIDs are typically mapped to VLANs for user, guest, voice, or IoT segmentation. Even if the WLC is configured correctly, missing VLANs (or missing allowed VLANs on trunks) will prevent clients from getting DHCP, reaching gateways, or being properly isolated.

"Deploy load balancers" is not a typical requirement for WLC deployments; redundancy is usually handled with WLC HA/SSO, N+1 designs, or AP/controller discovery mechanisms rather than generic load balancers. "Configure multiple VRRP groups" is also not inherently tied to WLC deployment; VRRP is a first-hop redundancy feature for default gateways and may already exist, but it's not a standard "must-do" WLC follow-up. "Configure additional security policies" can be relevant (ACLs, firewall rules, AAA), but the question asks for two tasks most directly associated with bringing a new WLC/WLAN online; VLANs and PoE are the most universally applicable.

References: [Cisco Enterprise Wireless](#), [Cisco PoE Support Documentation](#)

QUESTION NO: 77

```
GigabitEthernet1 is up, line protocol is up
Hardware is CSR vNIC, address is 5000.0004.0000 (bia 5000.0004.0000)
Internet address is 192.168.1.1/24
MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Full Duplex, 1000Mbps, link type is auto, media type is RJ45
```

Refer to the exhibit. Which format matches the Modified EUI-64 IPv6 interface address for the network 2001:db8::/64?

- A. 2001:db8::5000:00ff:fe04:0000/64
- B. 2001:db8::4332:5800:41ff:fe06:/64
- C. 2001:db8::5000:0004:5678:0090/64
- D. 2001:db8::5200:00ff:fe04:0000/64

ANSWER: D

Explanation:

With Modified EUI-64, the 64-bit interface ID is derived from the 48-bit MAC address by (1) splitting the MAC into two 24-bit halves, (2) inserting `FFFF` in the middle, and (3) flipping the Universal/Local (U/L) bit (the 7th bit) of the first byte. In the exhibit, the MAC begins with `50` (hex). Flipping the U/L bit means XORing that first byte with $0x02$: $0x50 \oplus 0x02 = 0x52$. That produces an interface ID that starts with `52` rather than `50`, and it must include the inserted `FFFF` in the middle. Therefore the correct IPv6 address format for prefix `2001:db8::/64` is the one that contains `::5200:00ff:fe04:0000` as the interface ID.

Option D matches this exact Modified EUI-64 construction. Option A is wrong because it does not flip the U/L bit (it starts with `50`, not `52`). Options B and C do not follow the required `FFFF` insertion and/or do not match the MAC-derived pattern.

References: [RFC 4291 \(IPv6 Addressing Architecture\)](#), [Cisco: IPv6 EUI-64 Addressing](#)

QUESTION NO: 78

Which two IPv6 addresses are used to provide connectivity between two routers on a shared link? (Choose two)

- A. ::ffif 1014 1011/96
- B. 2001 7011046:1111:1/64
- C. ;jff06bb43cd4dd111bbff02 4545234d
- D. 2002 5121204b 1111:1/64
- E. FF02::0WIFF00:0I)00/104
- F. FE80::1

ANSWER: B

Explanation:

On a shared IPv6 link between two routers, the addresses that provide basic connectivity are typically (1) a global unicast address from a /64 assigned to that link and (2) each interface's link-local address (FE80::/10). Global unicast addresses are routable and are used for end-to-end reachability, while link-local addresses are always present on IPv6 interfaces and are used for on-link operations and routing protocol neighbor adjacencies (for example, OSPFv3 and EIGRP for IPv6 commonly form neighbors using link-local addresses). In the provided options, only option B resembles a valid global unicast IPv6 address with a /64 prefix, which is appropriate for a router-to-router shared segment. None of the other options are valid IPv6 representations (they contain invalid characters, malformed hexets, or incorrect formatting), and none provides a valid FE80::/10 link-local address. Therefore, a new option is required to represent the missing correct link-local choice.

References: [RFC 4291 \(IPv6 Addressing Architecture\)](#), [Cisco: IPv6 Link-Local Addresses](#).

QUESTION NO: 79

How does MAC learning function?

- A. overwrites the known source MAC address in the address table
- B. enabled by default on all VLANs and interfaces
- C. protects against denial of service attacks
- D. forwards frames to a neighbor port using CDP

ANSWER: A

Explanation:

MAC learning is the switch behavior of dynamically building the MAC (CAM) address table by examining the *source* MAC address of incoming Ethernet frames. When a frame arrives, the switch records (or refreshes) an entry that maps that source MAC to the ingress port and VLAN. If the same source MAC is later seen on a different port in the same VLAN, the switch updates/overwrites the existing entry to point to the new port (this is normal behavior and is also why MAC flapping can be detected). This learned table is then used to make forwarding decisions for *destination* MAC addresses (known unicast vs unknown unicast flooding).

Option A best matches this core function: the switch updates/overwrites the existing MAC table entry when it learns a source MAC again (especially if it appears on a different port). Option B is misleading: while MAC learning is generally enabled by default, it's not accurately described as something you "enable on all VLANs and interfaces" as a configurable feature; it's inherent to Layer 2 switching behavior. Option C describes a security goal (DoS protection) rather than MAC learning itself (that's more aligned with features like port security/storm control). Option D is incorrect because CDP is a discovery protocol and not used for frame forwarding.

References: [Cisco: Ethernet Switching - How a Switch Works](#), [MAC address table \(overview\)](#)

QUESTION NO: 80

Which two statements distinguish authentication from accounting? (Choose two.)

- A. Only authentication supports user-activity audits.
- B. Only authentication challenges users for their credentials and returns a response.
- C. Only authentication validates "who you are."
- D. Only authentication records the duration of a user's connection.
- E. Only authentication provides supporting information for billing users.

ANSWER: B C

Explanation:


In AAA, **authentication** answers the question “Who are you?” by prompting a user/device for credentials (for example, username/password, certificates, tokens) and validating them against a local database or an external server such as RADIUS/TACACS+. That’s why option **B** is correct: authentication is the step that performs the credential challenge/response (or equivalent exchange) to establish identity. Option **C** is also correct because authentication is specifically the identity-verification function (“who you are”).

Accounting, by contrast, tracks and logs what the user did and for how long: session start/stop, duration, bytes/packets, commands (depending on the system), and other usage data. This information supports audits and billing/chargeback. Therefore, options **A**, **D**, and **E** are incorrect because those are classic *accounting* functions, not authentication. Auditing user activity (A), recording connection duration (D), and providing billing-support data (E) are all outcomes of accounting records, not identity verification.

References: [Cisco AAA overview \(RADIUS/AAA concepts\)](#), [Cisco TACACS+/AAA concepts](#).

QUESTION NO: 81

Refer to the exhibit.



Which command must be executed for Gi1.1 on SW1 to become a trunk port if Gi1/1 on SW2 is configured in desirable or trunk mode?

- A. switchport mode trunk
- B. switchport mode dot1-tunnel
- C. switchport mode dynamic auto
- D. switchport mode dynamic desirable

ANSWER: C

Explanation:

To have Gi1/1 on SW1 become a trunk when the neighbor (SW2 Gi1/1) is set to *dynamic desirable* or forced *trunk*, the required configuration on SW1 is **switchport mode dynamic auto**. In DTP (Dynamic Trunking Protocol), *dynamic auto* is passive: it does not actively try to form a trunk, but it *will* negotiate and become a trunk if the other side is actively negotiating (*dynamic desirable*) or is already configured as a trunk (assuming DTP negotiation is allowed and trunking is supported on both ends). This matches the question’s wording: SW1 must “become a trunk” given SW2 is in desirable or trunk mode.

Option A (*switchport mode trunk*) would also make SW1 a trunk, but it is not the “must be executed” command for negotiation-based trunking; it forces trunking regardless of DTP. Option D (*dynamic desirable*) would also form a trunk, but it is not required—*dynamic auto* is sufficient. Option B (*dot1-tunnel*) is for 802.1Q tunneling (Q-in-Q) and is unrelated to DTP trunk negotiation.

References: [Cisco DTP and trunking modes overview](#), [Cisco VLAN trunking and negotiation concepts](#)

QUESTION NO: 82

Refer to the exhibit.

```
ip domain-name CNAC.com
!
interface GigabitEthernet0/0/0
 ip address 192.168.1.10 255.255.255.0
 duplex auto
 speed auto
!
line vty 0 15
 login local

R1#show crypto key mypubkey rsa

R1#show ssh
%No SSHv2 server connections running.
%No SSHv1 server connections running.
```

Which two commands must be added to update the configuration of router R1 so that it accepts only encrypted connections? (Choose two)

- A. username CNAC secret R!41!4319115@
- B. ip ssh version 2
- C. line vty 0 4
- D. crypto key generate rsa 1024
- E. transport input ssh

ANSWER: D E

Explanation:

To ensure R1 accepts only encrypted remote-management connections, you must restrict inbound VTY access to SSH only. On Cisco IOS, that is done under the VTY lines with **transport input ssh**, which disables Telnet (clear-text) and permits only SSH (encrypted). In addition, SSH requires RSA keys to be present on the device; without them, the SSH server cannot operate. Therefore, **crypto key generate rsa 1024** (or another supported modulus size) is required to create the key pair used for SSH encryption and key exchange.

Option B (**ip ssh version 2**) is a best practice because SSHv2 is more secure than SSHv1, but it is not strictly required to meet the wording “accepts only encrypted connections” as long as SSH is the only allowed transport. Option A (creating a local username/secret) is also commonly needed when using *login local* for SSH authentication, but the question asks specifically for the two commands that must be added to accept only encrypted connections; restricting transport and enabling SSH via RSA keys are the essentials. Option C (**line vty 0 4**) is just a mode change, not a security-enforcing command by itself.

References: [Cisco SSH configuration example](#), [Cisco IOS Secure Shell \(SSH\) configuration guide](#).

QUESTION NO: 83

What is the primary function of a Layer 3 device?

- A. to analyze traffic and drop unauthorized traffic from the Internet

- B. to transmit wireless traffic between hosts
- C. to pass traffic between different networks
- D. forward traffic within the same broadcast domain

ANSWER: C

Explanation:

A Layer 3 device (most commonly a router, or a multilayer switch performing routing) primarily forwards packets between different IP networks. It makes forwarding decisions based on Layer 3 information—logical addressing (IPv4/IPv6) and the routing table—so it can move traffic from one subnet/VLAN to another and select the best next hop toward a remote destination. This is what distinguishes Layer 3 from Layer 2 switching, which forwards frames within the same broadcast domain using MAC addresses.

Option C is correct because “pass traffic between different networks” describes routing/inter-VLAN routing, the core job of Layer 3. Option D is incorrect because forwarding within the same broadcast domain is a Layer 2 switch function. Option A describes a firewall/UTM security function (traffic inspection and policy enforcement), not the primary role of a generic Layer 3 device. Option B describes wireless access/bridging behavior (Layer 1/2), not Layer 3 routing.

References: [Cisco IP Routing Overview \(routing fundamentals\)](#), [Cisco Networking Basics \(switching vs routing concepts\)](#).

QUESTION NO: 84

Refer to the exhibit.

```
Switch(config)#hostname R1
R1(config)#interface FastEthernet0/1
R1(config-if)#no switchport
R1(config-if)#ip address 10.100.20.42 255.255.255.0
R1(config-if)#line vty 0 4
R1(config-line)#login
```

An engineer booted a new switch and applied this configuration via the console port. Which additional configuration must be applied to allow administrators to authenticate directly to enable privilege mode via Telnet using a local username and password?

- R1(config)#username admin privilege 15 secret p@ss1234
R1(config-if)#line vty 0 4
R1(config-line)#login local
- R1(config)#username admin secret p@ss1234
R1(config-if)#line vty 0 4
R1(config-line)#login local
R1(config)#enable secret p@ss1234
- R1(config)#username admin
R1(config-if)#line vty 0 4
R1(config-line)#password p@ss1234
R1(config-line)#transport input telnet
- R1(config)#username admin
R1(config-if)#line vty 0 4
R1(config-line)#password p@ss1234

A. Option A

B. Option B

C. Option C

D. Option D

E. line vty 0 15
login local
privilege level 15
transport input telnet

ANSWER: E

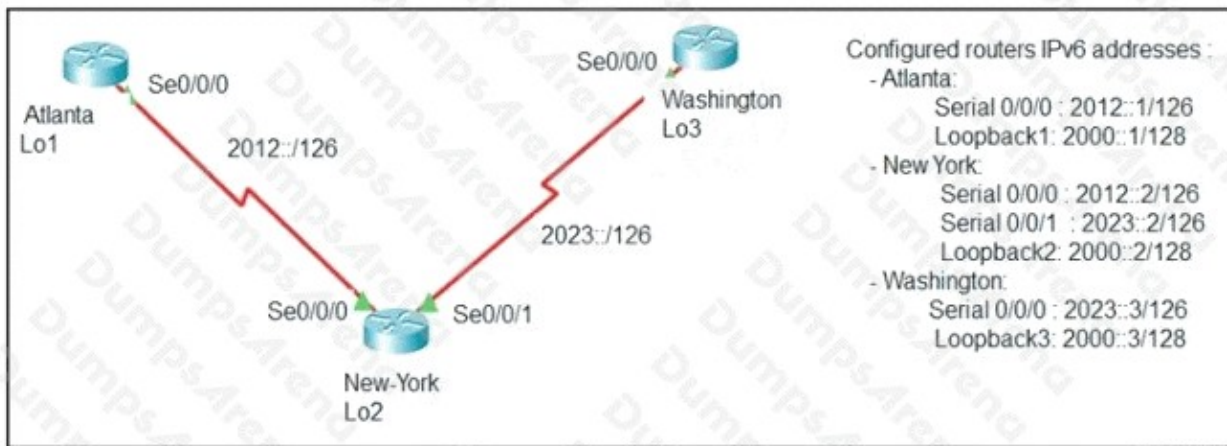
Explanation:

To let administrators Telnet to the switch and authenticate using a *local username/password*, the VTY lines must be configured to use the local user database (that is, `login local`) and Telnet must be permitted on those lines (for example, `transport input telnet` or `transport input all`). In addition, the question requires users to authenticate “directly to enable privilege mode”, which means they should land in privileged EXEC immediately after login. The standard way to do that is to set the VTY lines to privilege level 15 (for example, `privilege level 15` under `line vty`). Without that, users would typically enter user EXEC and then be prompted for the enable secret to reach privileged EXEC.

Options that only set an `enable secret` or only create a local username are insufficient by themselves, because Telnet authentication is controlled under the VTY lines. Options that configure VTY passwords (`password + login`) are also not “local username and password” authentication; that would be line-password authentication, not per-user local database authentication.

References: [Cisco AAA and local authentication overview](#), [Cisco IOS login local / line authentication configuration](#).

QUESTION NO: 85



Refer to the exhibit. An engineer configured the New York router with static routes that point to the Atlanta and Washington sites. Which command must be configured on the Atlanta and Washington routers so that both sites are able to reach the loopback2 interface on the New York router?

A. `ipv6 route::/0 Serial 0/0/0`

B. `ipv6 route::/0 Serial 0/0/1`

C. `ipv6 route:0/0 Serial 0/0/0`

D. `ip route 0.0.0.0 0.0.0.0 Serial 0/0/0`

E. `ipv6 route::/0 2000::2`

ANSWER: A

Explanation:


To allow Atlanta and Washington to reach New York's Loopback2, those two routers must have a route that sends traffic destined for unknown (non-local) IPv6 networks toward New York. Because the question says New York already has static routes pointing to Atlanta and Washington, the missing piece is the return path from Atlanta/Washington back to New York's networks (including the Loopback2 /128). The simplest way—given the single-choice options—is to configure an IPv6 default route on each remote router pointing out the WAN interface that leads to New York.

Option A (`ipv6 route ::/0 Serial 0/0/0`) is the correct syntax for an IPv6 default static route using an exit interface, and it matches the typical CCNA-style hub-and-spoke design where the spoke's Serial0/0/0 faces the hub (New York). With that default route, Atlanta and Washington will forward traffic for New York's Loopback2 toward New York, and New York's existing static routes provide the return reachability.

Option B could be correct only if the New York-facing link were on Serial0/0/1, which is not indicated. Option C is invalid syntax (`ipv6 route:0/0`). Option D is IPv4, not IPv6. Option E is incomplete/incorrect because a next-hop-only default route typically requires a full next-hop address (and often the outgoing interface on multiaccess links); as written it's not the best-practice/expected answer here.

References: [Cisco IOS XE IPv6 Static Routes](#), [Cisco Routing Basics \(default route concepts\)](#)

QUESTION NO: 86



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)# interface fa0/0
R1(config-if)# ip helper-address 198.51.100.100
- B. R2(config)# interface gi0/0
R2(config-if)# ip helper-address 198.51.100.100
- C. R1(config)# interface fa0/0 R1(config-if)# ip address dhcp
R1(config-if)# no shutdown
- D. R2(config)# interface gi0/0 R2(config-if)# ip address dhcp
- E. R1(config)# interface fa0/0
R1(config-if)# ip helper-address 192.0.2.2

ANSWER: B C

Explanation:

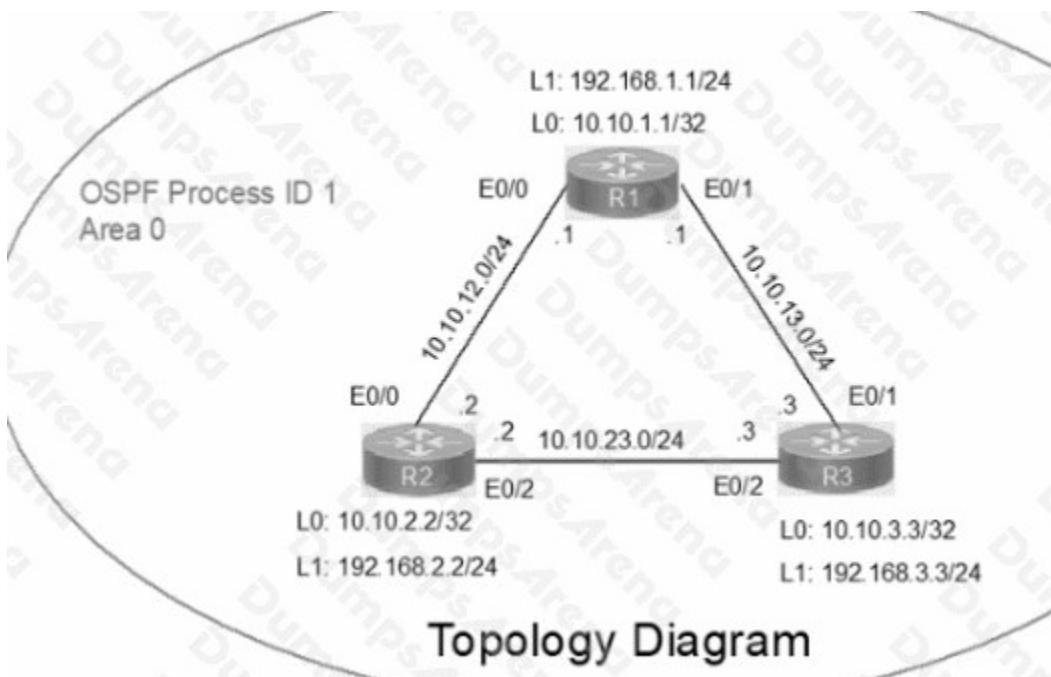
For R1 to obtain an address via DHCP, its client-facing interface must be configured as a DHCP client. That is exactly what option C does with `ip address dhcp` (and `no shutdown` to ensure the interface is operational). However, DHCP uses broadcasts (DISCOVER/REQUEST) that do not cross a router by default. Because the DHCP server is on a different network than R1, the router interface that receives the client broadcast must relay it to the DHCP server using `ip helper-`

address. In this topology, that relay function belongs on R2's interface connected toward R1's subnet (Gi0/0 in the options), pointing to the DHCP server's IP address (198.51.100.100), which is option B.

Option A is incorrect because configuring a helper on R1's client interface is not the needed relay point when R1 itself is the DHCP client; R1 generates the DHCP messages and routes unicast traffic normally once it has an address. Option D is incorrect because R2 is not intended to be a DHCP client here. Option E is incorrect because it relays to the wrong server address.

References: [Cisco Support: DHCP Relay \(ip helper-address\) overview](#), [Cisco IOS XE DHCP Relay Agent Configuration Guide](#).

QUESTION NO: 87 - (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 preconfigurations have been applied.
- 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.

IP connectivity between the three routers is configured. OSPF adjacencies must be established.

1. Configure R1 and R2 Router IDs using the interface IP addresses from the link that is shared between them.
2. Configure the R2 links with a max value facing R1 and R3. R2 must become the DR. R1 and R3 links facing R2 must remain with the default OSPF configuration for DR election. Verify the configuration after clearing the OSPF process.
3. Using a host wildcard mask, configure all three routers to advertise their respective Loopback1 networks.
4. Configure the link between R1 and R3 to disable their ability to add other OSPF routers.

ANSWER: See the explanation for answer

Explanation:

Answer as below configuration:

on R1

conf terminal

interface Loopback0

ip address 10.10.1.1 255.255.255.255

!

interface Loopback1

ip address 192.168.1.1 255.255.255.0

!

interface Ethernet0/0

no shut

ip address 10.10.12.1 255.255.255.0

ip ospf 1 area 0

duplex auto

!

interface Ethernet0/1

no shut

ip address 10.10.13.1 255.255.255.0

ip ospf 1 area 0

duplex auto

!

router ospf 1

router-id 10.10.12.1

network 10.10.1.1 0.0.0.0 area 0

network 192.168.1.0 0.0.0.255 area 0

!

copy run star

On R2

conf terminal

interface Loopback0

ip address 10.10.2.2 255.255.255.255

!

```
interface Loopback1
ip address 192.168.2.2 255.255.255.0
!
interface Ethernet0/0
no shut
ip address 10.10.12.2 255.255.255.0
ip ospf priority 255
ip ospf 1 area 0
duplex auto
!
interface Ethernet0/2
no shut
ip address 10.10.23.2 255.255.255.0
ip ospf priority 255
ip ospf 1 area 0
duplex auto
!
router ospf 1
network 10.10.2.2 0.0.0.0 area 0
network 192.168.2.0 0.0.0.255 area 0
!
```

```
copy runs start
```

```
-----
```

```
On R3
conf ter
interface Loopback0
ip address 10.10.3.3 255.255.255.255
!
interface Loopback1
ip address 192.168.3.3 255.255.255.0
!
interface Ethernet0/1
no shut
ip address 10.10.13.3 255.255.255.0
ip ospf 1 area 0
```

```
duplex auto
!  
interface Ethernet0/2  
no shut  
ip address 10.10.23.3 255.255.255.0  
ip ospf 1 area 0  
duplex auto  
!  
router ospf 1  
network 10.10.3.3 0.0.0.0 area 0  
network 192.168.3.0 0.0.0.255 area 0  
!  
copy run start  
!
```

QUESTION NO: 88

What are two roles of Domain Name Services (DNS)? (Choose two.)

- A. builds a flat structure of DNS names for more efficient IP operations
- B. encrypts network Traffic as it travels across a WAN by default
- C. improves security by protecting IP addresses under Fully Qualified Domain Names (FQDNs)
- D. enables applications to identify resources by name instead of IP address
- E. allows a single host name to be shared across more than one IP address

ANSWER: D E

Explanation:

DNS primarily provides name resolution: it lets users and applications refer to resources using human-friendly names (like hostnames/FQDNs) and then resolves those names to IP addresses (and sometimes the reverse via PTR records). That's why option D is correct—applications can identify resources by name instead of needing to know the IP address.

DNS also supports mapping one name to multiple IP addresses (multiple A/AAAA records). This is commonly used for basic load distribution, redundancy, or multi-site services, so option E is also correct.

Option A is incorrect because DNS is explicitly a hierarchical, distributed database (root, TLDs, authoritative zones), not a "flat structure." Option B is incorrect because DNS does not encrypt WAN traffic by default; standard DNS uses UDP/TCP 53 in cleartext (encryption requires add-ons like DoT/DoH). Option C is misleading/incorrect: DNS doesn't "protect" IP addresses by hiding them under FQDNs—DNS exists to publish mappings, and security improvements come from features like DNSSEC, not from merely using names.

References: [Cloudflare: What is DNS?](#), [IANA: DNS Root Zone](#)

QUESTION NO: 89

Company has decided to require multifactor authentication for all systems. Which set of parameters meets the requirement?

- A. personal 10-digit PIN and RSA certificate
- B. complex password and personal 10-digit PIN
- C. password of 8 to 15 characters and personal 12-digit PIN
- D. fingerprint scanning and facial recognition

ANSWER: A

Explanation:

Multifactor authentication (MFA) requires using credentials from at least two different authentication factor categories: something you know (password/PIN), something you have (token/smart card/certificate stored on a device), and something you are (biometrics). Option A meets MFA because it combines a personal PIN (something you know) with an RSA certificate (something you have, typically stored on a smart card, token, or device keystore). That is two distinct factors, satisfying the requirement.

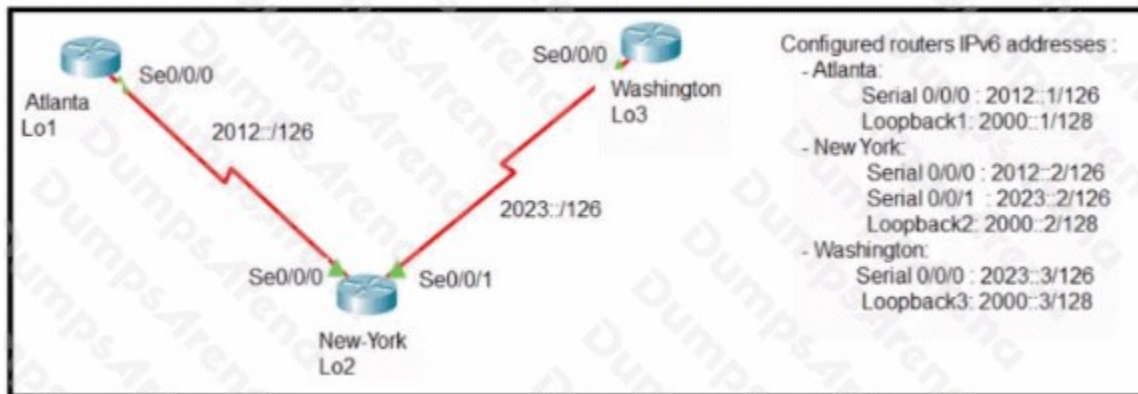
Options B and C do not meet MFA because both items in each pair are “something you know” (a password and a PIN). Using two knowledge-based secrets is still single-factor (knowledge) authentication, even if the secrets differ in complexity/length.

Option D also fails the MFA requirement as written because it uses two biometric methods (fingerprint and face), which are both in the same factor category (“something you are”). That’s multi-method biometric authentication, but not multi-factor.

For background on authentication factors and MFA concepts, see NIST’s Digital Identity Guidelines ([NIST SP 800-63B](#)) and Cisco’s overview of MFA concepts ([Cisco Duo MFA](#)).

QUESTION NO: 90

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 Serial0/0/0 interfaces on the Atlanta and Washington routers can reach one another?

(Choose two.)

- A. Configure the ipv6 route 2012::/126 2023::1 command on the Washington router.
- B. Configure the ipv6 route 2023::/126 2012::1 command on the Atlanta router.
- C. Configure the ipv6 route 2012::/126 s0/0/0 command on the Atlanta router.
- D. Configure the ipv6 route 2023::/126 2012::2 command on the Atlanta router.
- E. Configure the ipv6 route 2012::/126 2023::2 command on the Washington router.

ANSWER: D E

Explanation:

For Atlanta and Washington to reach each other's Serial0/0/0 networks using static routing, each router must have a route to the other router's WAN /126 prefix, pointing to the correct next-hop address on the shared New York link. New York already has static routes to both sites, but that alone does not give Atlanta a route to Washington (and vice versa). Each "spoke" must send traffic for the other spoke toward New York as the next hop.


That's why Atlanta needs a static route to Washington's WAN prefix (2023::/126) with a next hop of New York's address on the Atlanta–New York link (2012::2). Similarly, Washington needs a static route to Atlanta's WAN prefix (2012::/126) with a next hop of New York's address on the Washington–New York link (2023::2). These match options D and E.

Options A and B use next-hop addresses that are not the New York-facing next hop for those routers (they point to the wrong side). Option C is incomplete/incorrect as written because using only an exit interface for an IPv6 static route on a multi-access style scenario can be problematic; best practice is to specify the next-hop IPv6 address (or both exit interface and next hop) to ensure proper neighbor discovery and forwarding.

References: [Cisco – Configuring Static Routes](#), [Cisco IOS IPv6 Static Route Configuration](#)

QUESTION NO: 91

Refer to the exhibit.



Which configuration establishes a Layer 2 LACP EtherChannel when applied to both switches?

- A. Interface range G1/1 – 1/3 switchport mode trunk channel-group 1 mode active no shutdown
- B. Interface range G1/1 – 1/3 switchport mode access channel-group 1 mode passive no shutdown
- C. Interface range G1/1 – 1/3 switchport mode trunk channel-group 1 mode desirable no shutdown
- D. Interface range G1/1 – 1/3 switchport mode access channel-group 1 mode on no shutdown

ANSWER: A

Explanation:

To establish a Layer 2 EtherChannel using LACP, the member interfaces on both switches must be placed into the same channel-group and use LACP negotiation modes (**active** or **passive**). At least one side must be **active** for LACP to form; **active/active** is a common, reliable choice. Option A correctly configures an interface range, sets the ports as Layer 2 switchports in trunk mode (appropriate if the EtherChannel is intended to carry multiple VLANs), and enables LACP with `channel-group 1 mode active`. When applied to both switches, this will negotiate and bundle the links into Port-Channel 1.

Option B uses LACP **passive** on both sides (as stated "applied to both switches"), which will not form an EtherChannel because neither side initiates LACP negotiation. Option C uses **desirable**, which is a PAgP mode (Cisco proprietary), not LACP. Option D uses **mode on**, which forces a static EtherChannel (no negotiation protocol), so it is not LACP.

References: [Cisco EtherChannel and LACP/PAgP configuration overview](#), [Cisco IOS XE LAN EtherChannel Configuration Guide](#)

```
R1#config t
R1(config)# interface gil/1
R1(config-if)# ip address 192.168.0.1 255.255.255.0

R1(config)# router bgp 65000
R1(config-router)# neighbor 192.168.0.2 remote-as 65001
R1(config-router)# network 10.1.1.0 mask 255.255.255.0

R1(config)# router ospf 1
R1(config)# router-id 1.1.1.1
R1(config)# network 192.168.0.1 0.0.0.0 area 0
R1(config)# network 10.1.1.0 0.0.0.255 area 0

R1(config)# router eigrp 1
R1(config)# eigrp router-id 1.1.1.1
R1(config)# network 10.1.1.0 0.0.0.255
R1(config)# network 192.168.0.1 0.0.0.0

R2#config t
R2(config)# interface gil/1
R2(config-if)# ip address 192.168.0.2 255.255.255.0

R2#config t
R2(config)# router bgp 65001
R2(config-router)# neighbor 192.168.0.1 remote-as 65000

R2(config)# router ospf 1
R2(config)# router-id 2.2.2.2
R2(config)# network 192.168.1.2 0.0.0.0 area 0

R2(config)# router eigrp 1
R2(config)# eigrp router-id 1.1.1.1
R2(config)# network 192.168.0.1 0.0.0.0

R2(config)# ip route 10.1.1.0 255.255.255.0 192.168.0.1
```

Refer to the exhibit. Router R2 is configured with multiple routes to reach network 10.1.1.0/24 from router R1. Which path is chosen by router R2 to reach the destination network 10.1.1.0/24?

- A. static
- B. EIGRP
- C. eBGP
- D. OSPF

ANSWER: A

Explanation:

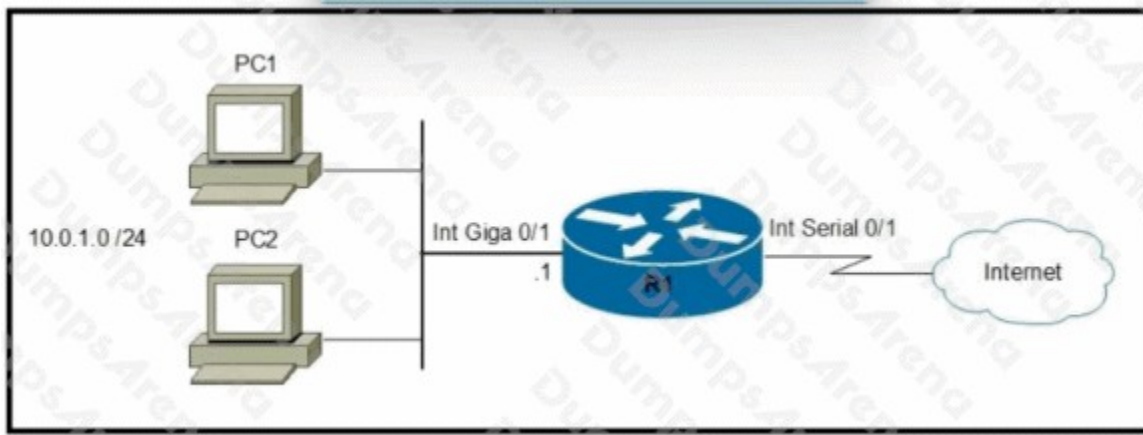
When a router has multiple routes to the same destination prefix learned from different sources (static and dynamic routing protocols), it selects the route with the lowest Administrative Distance (AD). AD is Cisco's measure of how trustworthy a route source is; the lower the AD, the more preferred the route. By default on Cisco IOS, a static route has an AD of 1, eBGP has an AD of 20, EIGRP (internal) has an AD of 90, and OSPF has an AD of 110. Therefore, assuming all routes are for the same prefix length (10.1.1.0/24) and are valid, R2 will install and use the static route to reach 10.1.1.0/24 because it has the lowest AD.

The other options are incorrect because their default ADs are higher than a static route's AD. eBGP (20) would beat EIGRP and OSPF, but it still loses to a static route (1). EIGRP and OSPF are even less preferred by default. Note that only if the static route were configured as a "floating static" (higher AD than the dynamic route) would a dynamic protocol be chosen instead.

References: [Cisco: Administrative Distance \(Routing Protocol Preference\)](#), [Cisco IOS IP Routing: Route Selection Concepts](#)

QUESTION NO: 93

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. transport input telnet
- B. crypto key generate rsa
- C. ip ssh pubkey-chain
- D. login console
- E. username cisco password 0 Cisco

ANSWER: B E

Explanation:

"Secure remote-access connections" on a Cisco router typically refers to SSH (not Telnet). For a router to accept inbound SSH sessions, it must have (1) an RSA key pair and (2) a local user database (or AAA) to authenticate users. The command **crypto key generate rsa** creates the RSA keys required by the SSH server process; without keys, SSH cannot be enabled/used. The command **username cisco password 0 Cisco** creates a local username/password that can be used with VTY lines (for example, with *login local*) to authenticate SSH logins.

Option A (**transport input telnet**) explicitly allows Telnet, which is insecure and does not enable secure access. Option C (**ip ssh pubkey-chain**) is only needed when configuring SSH public key authentication; it is not required for basic SSH with password authentication. Option D (**login console**) is not a valid/typical command for enabling remote access on VTY lines; remote access is controlled under *line vty* with commands like *transport input ssh* and *login local*.

References: [Cisco SSH configuration example](#), [Cisco IOS SSH configuration guide](#).

QUESTION NO: 94

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:

A DHCP server's core job is to automatically provide hosts with IP configuration information (at minimum an IP address, and commonly also subnet mask, default gateway, DNS servers, and lease timers). That makes option D correct: DHCP assigns dynamic IP configurations to hosts using a lease-based process. Another key function/benefit is centralized address management (option C). By keeping scopes/pools, exclusions, and lease records in one place, DHCP simplifies administration and reduces manual configuration errors across the network.

Option A is incorrect because DHCPDISCOVER messages are sent by DHCP clients (broadcast) when they need an address; servers listen and respond. Option B is incorrect because the message flow is reversed: the server sends a DHCPOFFER to the client, and the client sends a DHCPREQUEST to accept an offer; the server then replies with DHCPACK. Option E is not a true DHCP "function": DHCP does not technically prevent a user from statically configuring an IP on their device; it only provides an automated alternative (and admins may use other controls like DHCP snooping/port security to mitigate misuse).

References: [Cisco DHCP overview and message exchange](#), [DHCP protocol summary \(DORA process\)](#).

QUESTION NO: 95

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

- A. Southbound API
- B. REST API
- C. NETCONF
- D. OpenFlow

ANSWER: B

Explanation:

In SDN, applications (often called "SDN apps" or "northbound clients") communicate with the SDN controller using a *northbound interface*, which is commonly implemented as a RESTful API over HTTP/HTTPS. This lets applications request network state, push policy/intent, and automate provisioning through standard web methods (GET/POST/PUT/DELETE) and data formats like JSON. Therefore, a REST API is the appropriate technology for communication between an SDN controller and applications.

Option A (Southbound API) is the opposite direction: southbound interfaces connect the controller to network devices (switches/routers) to program forwarding and collect telemetry. Option D (OpenFlow) is a classic example of a southbound protocol used between the controller and switches, not between the controller and applications. Option C (NETCONF) is primarily a device configuration/management protocol (often used controller-to-device, i.e., southbound/management plane), not the typical controller-to-application interface in SDN architectures.

References: Cisco SDN overview and controller APIs ([Cisco SDN Overview](#)) and REST architectural style background ([What is a REST API](#)).

QUESTION NO: 96

An engineer observes high usage on the 2.4GHz channels and lower usage on the 5GHz channels. What must be configured to allow clients to preferentially use 5GHz access points?

- A. Re-Anchor Roamed Clients
- B. 11ac MU-MIMO
- C. OEAP Split Tunnel
- D. Client Band Select

ANSWER: D

Explanation:

To encourage dual-band clients to use 5 GHz instead of 2.4 GHz, you configure **band selection** (often called *Client Band Select* on Cisco WLC). Band Select works by influencing the client's band choice during association—typically by delaying or ignoring some 2.4 GHz probe/association requests so the client is more likely to discover and join the 5 GHz BSSID when both are available. This helps reduce congestion on 2.4 GHz (which has fewer non-overlapping channels and more interference) and better utilizes 5 GHz capacity.

The other options don't address steering clients between bands. **Re-Anchor Roamed Clients** relates to mobility/anchoring (often for guest WLAN designs), not RF band preference. **11ac MU-MIMO** is a PHY feature that can improve throughput/efficiency on supported 5 GHz clients, but it does not cause clients to choose 5 GHz in the first place. **OEAP Split Tunnel** is for OfficeExtend AP traffic forwarding behavior and is unrelated to RF band utilization.

References: [Cisco Wireless LAN Controller Configuration Examples \(includes Band Select concepts\)](#), [Cisco WLC Configuration Guide \(Wireless Networks / Band Select\)](#)

QUESTION NO: 97

What is a function of a southbound API?

- A. Use orchestration to provision a virtual server configuration from a web server
- B. Automate configuration changes between a server and a switching fabric
- C. Manage flow control between an SDN controller and a switching fabric
- D. Facilitate the information exchange between an SDN controller and application

ANSWER: C

Explanation:

A southbound API is the interface used by an SDN controller to communicate with and program the network devices (the data plane), such as switches and routers. Its core function is to push forwarding/flow rules and other control instructions from the controller down to the switching/forwarding infrastructure. A classic example of a southbound protocol/API is OpenFlow, which allows the controller to manage how traffic is handled by the switches (i.e., flow entries and forwarding behavior).

Option C matches this: it describes managing flow control between an SDN controller and a switching fabric, which is exactly the controller-to-device (southbound) relationship. Option D is incorrect because facilitating information exchange between an SDN controller and applications is a northbound API function (controller-to-app). Options A and B are more aligned with

orchestration/automation use cases and do not specifically describe the SDN controller-to-network-device programming interface that defines southbound APIs.

References: [Cisco SDN overview](#), [Open Networking Foundation \(ONF\) SDN definition](#).

QUESTION NO: 98

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:

From the exhibit (a routing table view), the key clues are the route codes, the bracketed values, and the “variably subnetted” summary line. The “variably subnetted” line for 10.0.0.0/8 indicates how many subnets and masks exist within that major network. If it shows “20 subnets, 20 masks,” then it must be true that there are 20 different masks used inside 10.0.0.0/8, which makes option B correct.

The routing table also indicates equal-cost multipath (ECMP) when the same prefix appears with multiple next-hops/interfaces and identical metrics. If the exhibit shows ten routes/prefixes with two equal next-hops via Te0/1/0.100 and Te0/2/0.100 (same administrative distance/metric), then those routes are being load-balanced across both interfaces, making option C correct.

Option A is not required: EIGRP’s default AD is 90 (internal) and 170 (external), and seeing 170 does not prove it was manually changed. Option D is incorrect if the route is marked internal EIGRP (D) rather than external (D EX). Option E would require an “S* 0.0.0.0/0 via 10.85.33.14” style entry; without that exact default-route notation, it cannot be concluded.

References: [Cisco EIGRP overview and administrative distance](#), [Cisco: Understanding routing table output \(route codes/AD/metrics\)](#)

QUESTION NO: 99

Which Rapid PVST+ port state does a port operate in without receiving BPDUs from neighbors or updating the address database?

- A. listening
- B. forwarding
- C. disabled
- D. blocking

ANSWER: D

Explanation:

In Rapid PVST+ (Cisco’s per-VLAN implementation of RSTP), the port state that does *not* learn MAC addresses (doesn’t update the CAM/address table) and does *not* forward user traffic is the **discarding** behavior. In classic STP terms, that behavior maps to the **blocking** (and listening) states; in RSTP/Rapid PVST+, those are consolidated into the discarding state. Among the provided options, **blocking** is the best match because a blocking/discarding port does not populate the MAC address table and does not forward frames. It can still receive BPDUs (and in RSTP, it processes them), but the key differentiator in exam questions is typically “no learning/no forwarding,” which aligns with blocking/discarding.

Why the others are wrong: **Listening** is a legacy 802.1D STP state (RSTP replaces it with discarding), so it’s not a Rapid PVST+ state. **Forwarding** both forwards frames and learns MAC addresses, so it clearly doesn’t fit. **Disabled** means the interface is administratively down or not participating at all; it’s not the operational STP/RSTP state described for an active port in the topology.

References: [Cisco STP/RSTP Port States and Roles](#), [Rapid Spanning Tree Protocol \(RSTP\) overview](#)

QUESTION NO: 100

Which Rapid PVST+ feature should be configured on a switch port to immediately send traffic to a connected server as soon as it is active?

- A. portfast
- B. uplinkfast
- C. BPDU guard
- D. loop guard

ANSWER: A

Explanation:

The correct feature is **PortFast**. In Rapid PVST+ (RSTP), a port connected to an end host (like a server) can be configured as an “edge” port using PortFast. This allows the port to transition to the forwarding state immediately when the link comes up, instead of waiting through the normal spanning-tree convergence process. That’s exactly what you want for server-facing access ports so the server can start sending/receiving traffic right away (for example, during boot or after a link flap).

UplinkFast is a legacy enhancement for classic 802.1D STP to speed up convergence on access-layer switches with redundant uplinks; it’s not the feature used to make a single server-facing port forward immediately, and it’s not part of RSTP behavior. **BPDU Guard** is a protection mechanism typically used with PortFast to err-disable a port if BPDUs are received—useful for preventing accidental switch connections, but it doesn’t itself make the port forward immediately. **Loop Guard** prevents certain STP failures from creating loops by keeping ports from transitioning to forwarding when BPDUs stop unexpectedly; it also doesn’t provide immediate forwarding.

References: [Cisco STP PortFast and BPDU Guard](#), [Cisco Understanding Rapid Spanning Tree Protocol \(802.1w\)](#)

QUESTION NO: 101

Which interface condition is occurring in this output?

```
R45# show interface fa0/0
FastEthernet0/0 is up, line protocol is up
Hardware is DEC21140, address is ca02.7788.0000 (bia ca02.7788.0000)
Description: atlanta_subnet
Internet address is 10.32.102.2/30
MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
reliability 255/255, txload 255/255, rxload 255/255
Encapsulation ARPA, loopback not set
Keepalive set (60 sec)
Full-duplex, 100 Mb/s, 100BaseTX/FX
ARP type: ARPA, ARP Timeout 04:00:00
Last input 00:00:01, output 00:00:00, output hang never
Last clearing of "show interface" counters never
Input queue: 0/300/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/300 (size/max)
30 second input rate 234712855 bits/sec, 0 packets/sec
30 second output rate 228528957 bits/sec, 0 packets/sec
7331 packets input, 7101162 bytes
Received 267 broadcasts (0 IP multicasts)
0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
0 watchdog
0 input packets with dribble condition detected
3927 packets output, 1440403 bytes, 0 underruns
0 output errors, 0 collisions, 0 interface resets
0 unknown protocol drops
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier
0 output buffer failures, 0 output buffers swapped out
```

- A. broadcast storm
- B. collisions
- C. high throughput
- D. duplex mismatch

ANSWER: C

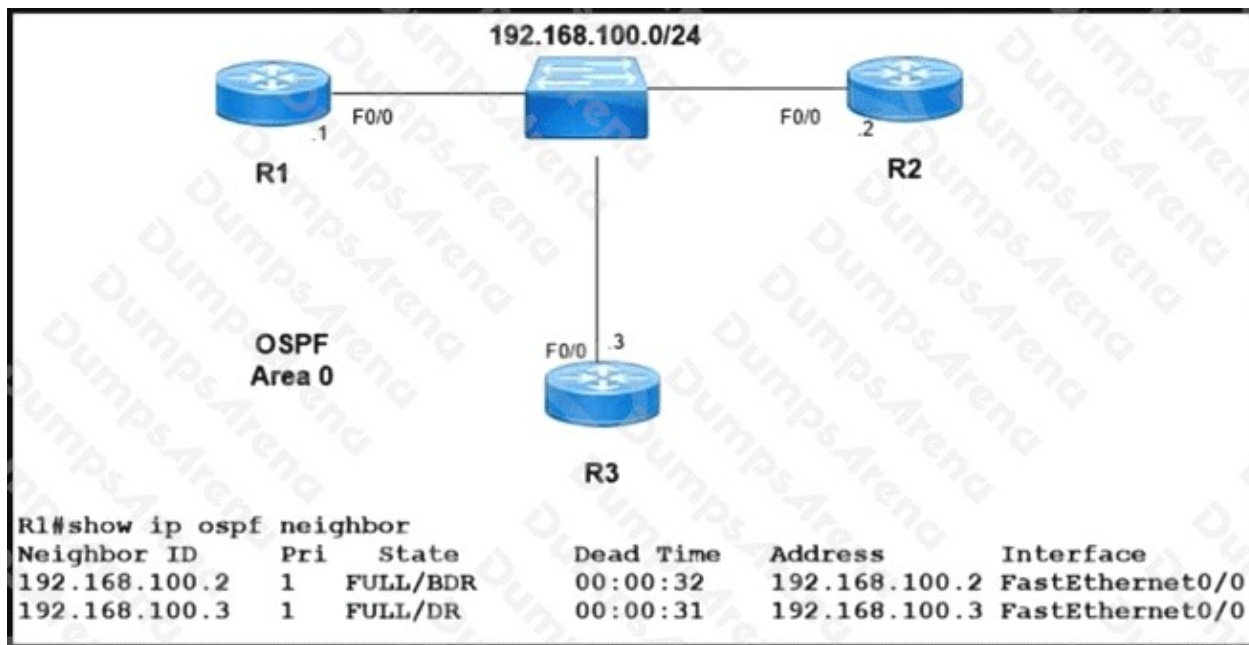
Explanation:

The interface condition indicated by this kind of output is **high throughput**. In Cisco IOS, the `show interfaces` output commonly includes “5 minute input rate” and “5 minute output rate” values (bits/sec and packets/sec). When those rates are consistently high relative to the interface’s capacity (for example, a large percentage of a 100 Mb/s or 1 Gb/s link), it points to the interface being heavily utilized—i.e., high throughput. This is a normal operational condition (though it may indicate congestion risk) and is not, by itself, an error state.

The other options imply error symptoms that would show up as specific counters increasing. A **broadcast storm** would typically be suggested by unusually high broadcast/multicast traffic and potentially rising input drops/CPU impact, not just high overall rates. **Collisions** are a half-duplex Ethernet phenomenon; on modern switched full-duplex links you should see zero collisions, and the output would explicitly show collision counters incrementing. A **duplex mismatch** usually presents as late collisions, CRC/FCS errors, runts/giants, and poor performance; again, you’d expect those error counters to climb rather than simply seeing high 5-minute rates.

References: [Cisco: Interpreting Ethernet Interface Statistics](#), [Cisco: Using show interfaces to Troubleshoot](#)

QUESTION NO: 102



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 a broadcast multiaccess segment (like Ethernet), OSPF elects a DR/BDR based first on the interface OSPF priority (highest wins), and if there’s a tie, on the highest OSPF router ID. A router with interface priority 0 is ineligible to become DR/BDR. Therefore, to ensure R1 becomes the DR, you must (1) make R1’s interface priority higher than the others and/or (2) make competing routers ineligible by setting their priority to 0.

Option C sets R1's FastEthernet0/0 OSPF priority to 200, which strongly biases the election in R1's favor (default is 1). Option A sets R3's priority to 0, removing R3 from DR/BDR eligibility, further ensuring R1 wins even if other routers exist on the segment. Option B (setting router-id) can influence elections only when priorities tie; by itself it does not guarantee DR if another router has a higher priority. Option D makes R1 ineligible (priority 0), the opposite of the goal. Option E increases R3's priority, making it more likely R3 becomes DR, not R1.

Note: after changing priority/router-id, you typically need to clear the OSPF process or bounce the interface to force a new election. References: [Cisco OSPF DR/BDR Election](#), [Cisco IOS OSPF Commands \(ip ospf priority\)](#).

QUESTION NO: 103

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

ANSWER: A

Explanation:

On a Cisco WLC (AireOS-style workflow), a basic WLAN build typically follows this order: define the RADIUS/authentication server, create the client VLAN mapping (dynamic interface), then create the WLAN (SSID) and map/bind it to the appropriate interface so client traffic is placed into the correct VLAN. Without creating the WLAN and binding it to the dynamic interface, clients may see no SSID to join or, if they can join, they won't be placed into the intended VLAN and authentication policies won't be applied as expected. Therefore, the next required step to complete the basic configuration is to create the WLAN and bind it to the dynamic interface.

High availability/redundancy for APs is a design/availability enhancement, not required for basic WLAN functionality. Enabling Telnet is not a best practice (SSH is preferred), and RADIUS for management access is unrelated to client WLAN setup. "Install the management interface and add the management IP" is something you do during initial controller setup; it's not the next step after already creating a dynamic interface and configuring an auth server.

References: [Cisco Wireless LAN Controller Configuration Guides](#), [Cisco WLAN security and RADIUS/802.1X overview](#).

QUESTION NO: 104

What are two purposes of HSRP? (Choose two.)

- A. It provides a mechanism for diskless clients to autoconfigure their IP parameters during boot.
- B. It improves network availability by providing redundant gateways.
- C. It groups two or more routers to operate as one virtual router.
- D. It passes configuration information to hosts in a TCP/IP network.
- E. It helps hosts on the network to reach remote subnets without a default gateway.

ANSWER: B C

Explanation:

HSRP (Hot Standby Router Protocol) is a Cisco first-hop redundancy protocol used to make the default gateway function highly available for end hosts. Its main purpose is to improve network availability by providing gateway redundancy: hosts use a single virtual IP/MAC as their default gateway, and if the active router fails, a standby router takes over with minimal disruption. That directly matches option B.

HSRP achieves this by forming an HSRP group in which two or more routers cooperate to present themselves as one “virtual router” to the LAN. One router is elected active (for forwarding) and another is standby (ready to take over). This behavior is exactly what option C describes.

Option A describes BOOTP/DHCP-style boot-time IP parameter assignment for diskless clients, not HSRP. Option D also aligns with DHCP (passing configuration info like DNS server, default gateway, etc.), not HSRP. Option E is incorrect because hosts generally require a default gateway to reach remote subnets; HSRP doesn’t eliminate that requirement—it provides a resilient default gateway instead.

References: [Cisco HSRP Overview and Configuration](#), [HSRP \(overview\)](#)

QUESTION NO: 105

How do servers connect to the network in a virtual environment?

- A. a cable connected to a physical switch on the network
- B. wireless to an access point that is physically connected to the network
- C. a virtual switch that links to an access point that is physically connected to the network
- D. a software switch on a hypervisor that is physically connected to the network

ANSWER: D

Explanation:

In a virtualized environment, a server (VM) connects to the network through a virtual NIC (vNIC) that plugs into a virtual switch (vSwitch) implemented in software on the hypervisor. The vSwitch then forwards traffic out of the host via one or more physical NICs (uplinks) that connect to the physical network (typically a physical switch). That’s exactly what option D describes: a software switch on the hypervisor with physical connectivity to the network.

Option A is how a physical server connects (a physical cable from the server NIC to a physical switch), not how a VM connects. Option B is generally irrelevant for data-center virtualization: servers/hosts are not typically “wireless to an access point” for production networking, and VMs don’t directly associate to Wi-Fi like clients. Option C is incorrect because it introduces an access point as the physical attachment; virtual switches uplink to physical NICs/switch ports, not to APs as a standard server connectivity model.

References: [Cisco: What is a network switch?](#), [VMware glossary: Virtual switch](#)

QUESTION NO: 106

An engineer is configuring switch SW1 to act an NTP server when all upstream NTP server connectivity fails. Which configuration must be used?

A)

```
SW1# config t
SW1(config)#ntp peer 192.168.1.1
SW1(config)#ntp access-group peer accesslist1
```

B)

```
SW1# config t
SW1(config)#ntp master
SW1(config)#ntp server 192.168.1.1
```

C)

```
SW1# config t
SW1(config)#ntp server 192.168.1.1
SW1(config)#ntp access-group server accesslist1
```

D)

```
SW1# config t
SW1(config)#ntp backup
SW1(config)#ntp server 192.168.1.1
```

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

ANSWER: B

Explanation:

To have SW1 provide time to other devices only as a fallback when it cannot reach its upstream NTP sources, you use the **ntp master** feature. In Cisco IOS, **ntp master [stratum]** makes the device act as an NTP server (authoritative clock source) at the configured stratum (or default stratum 8). This is commonly used as a “local clock backup” so downstream clients can still synchronize if WAN/upstream NTP connectivity fails. Therefore, the correct configuration is the one that enables SW1 as an NTP master (option B).

The other options are incorrect because they do not make the switch serve time as a master clock. Configurations that only define upstream servers (for example, **ntp server x.x.x.x**) make SW1 an NTP client, not a fallback server. Likewise, commands such as **ntp peer** establish symmetric peering rather than explicitly making SW1 a master clock source, and authentication-only or source-interface-only snippets don't create the required fallback time source. The key requirement in the question is “act as an NTP server when upstream fails,” which is exactly what **ntp master** is designed for.

References: [Cisco IOS Time and NTP Configuration Guide](#), [Cisco NTP Best Practices/Overview](#)

QUESTION NO: 107

Which physical component is distributed among multiple virtual machines running on the same hypervisor?

- A. external storage
- B. network interfaces
- C. backplane network
- D. hardware resources

ANSWER: D

Explanation:

On a hypervisor, multiple virtual machines (VMs) share the underlying host's physical compute capacity. The hypervisor abstracts and schedules the host's CPU, memory, and I/O so that each VM receives a portion of those physical resources according to configuration (for example, vCPU count, memory allocation, shares/limits/reservations). In other words, the

physical *hardware resources* of the host are distributed across the VMs running on that same hypervisor, which is a core concept of server virtualization.

Option A (external storage) is not inherently “distributed among VMs on the same hypervisor”; storage may be local or shared (SAN/NAS), but it’s typically presented as datastores/volumes and allocated as virtual disks—sharing depends on the storage architecture, not simply co-residency on a hypervisor. Option B (network interfaces) can be shared via virtual switches and vNICs, but the question asks for the physical component distributed among VMs; the broad, correct concept tested in CCNA is the host’s hardware resources. Option C (backplane network) is not a standard hypervisor physical component in this context.

References: [Cisco: What is virtualization?](#), [VMware glossary: Hypervisor](#)

QUESTION NO: 108

What causes a port to be placed in the err-disabled state?

- A. shutdown command issued on the port
- B. port security violation
- C. nothing plugged into the port
- D. latency

ANSWER: B

Explanation:

An interface goes into the *err-disabled* state when the switch detects a condition it considers serious enough to automatically shut the port down to protect the network. A classic and very common trigger is a **port-security violation** (for example, more MAC addresses are learned than allowed, or an unauthorized MAC appears). When the violation action is set to `shutdown` (the default on many platforms), the switch places the port into err-disabled, and it stays down until manually recovered (`shutdown` or `errdisable recovery`) or until errdisable recovery is configured.

Option A is wrong because an administratively issued `shutdown` puts the interface into an *administratively down* state, not err-disabled. Option C is wrong because an unplugged cable typically results in *down/down* (link down), not err-disabled. Option D is wrong because “latency” is not an errdisable cause; errdisable is tied to specific detected faults (e.g., port-security, BPDU Guard, UDLD, link-flap, etc.).

References: [Cisco Support: Understanding and Troubleshooting Errdisable](#), [Cisco Support: Configuring Port Security](#)

QUESTION NO: 109

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:

In a controller-based (SDN) network, the key idea is that the control plane is logically centralized in a controller, while the data plane remains on the network devices. That makes option E correct: the controller becomes the central point where policy and control decisions are made and then pushed to devices. Another hallmark is the use of APIs between layers: southbound APIs (controller-to-device, e.g., NETCONF/gNMI/OpenFlow depending on implementation) and northbound APIs (controller-to-applications/automation/orchestration). That makes option C correct.

Option A is incorrect because Telnet is not a defining characteristic of controller-based networking and is generally discouraged due to lack of encryption (SSH is preferred). Option B is incorrect because while you can still use CLI on devices, controller-based networks emphasize centralized policy/automation via the controller rather than per-device CLI updates as a characteristic. Option D is incorrect because it describes a traditional distributed control plane (each device runs its own control logic), which is the opposite of controller-based networking's logically centralized control.

References: [Cisco SDN overview](#), [Cisco DNA Center \(controller-based networking\)](#).

QUESTION NO: 110

Refer to the exhibit. A network associate has configured OSPF with the command:

```
City(config-router)# network 192.168.12.64 0.0.0.63 area 0
```

After completing the configuration, the associate discovers that not all the interfaces are participating in OSPF. Which three of the interfaces shown in the exhibit will participate in OSPF according to this configuration statement? (Choose three.)

```
City#show ip interface brief
```

Interface	IP-Address	OK?	Method	Status	Protocol
FastEthernet0/0	192.168.12.48	Yes	manual	up	up
FastEthernet0/1	192.168.12.65	Yes	manual	up	up
Serial0/0	192.168.12.121	Yes	manual	up	up
Seriak0/1	unassigned	Yes	unset	up	up
Serial0/1.102	192.168.12.125	Yes	manual	up	up
Serial0/1.103	192.168.12.129	Yes	manual	up	up
Serial0/1.104	192.168.12.133	Yes	manual	up	up

```
City#
```

- A. FastEthernet0 /0
- B. FastEthernet0 /1
- C. Serial0/0
- D. Serial0/1.102
- E. Serial0/1.103
- F. Serial0/1.104

ANSWER: B C D

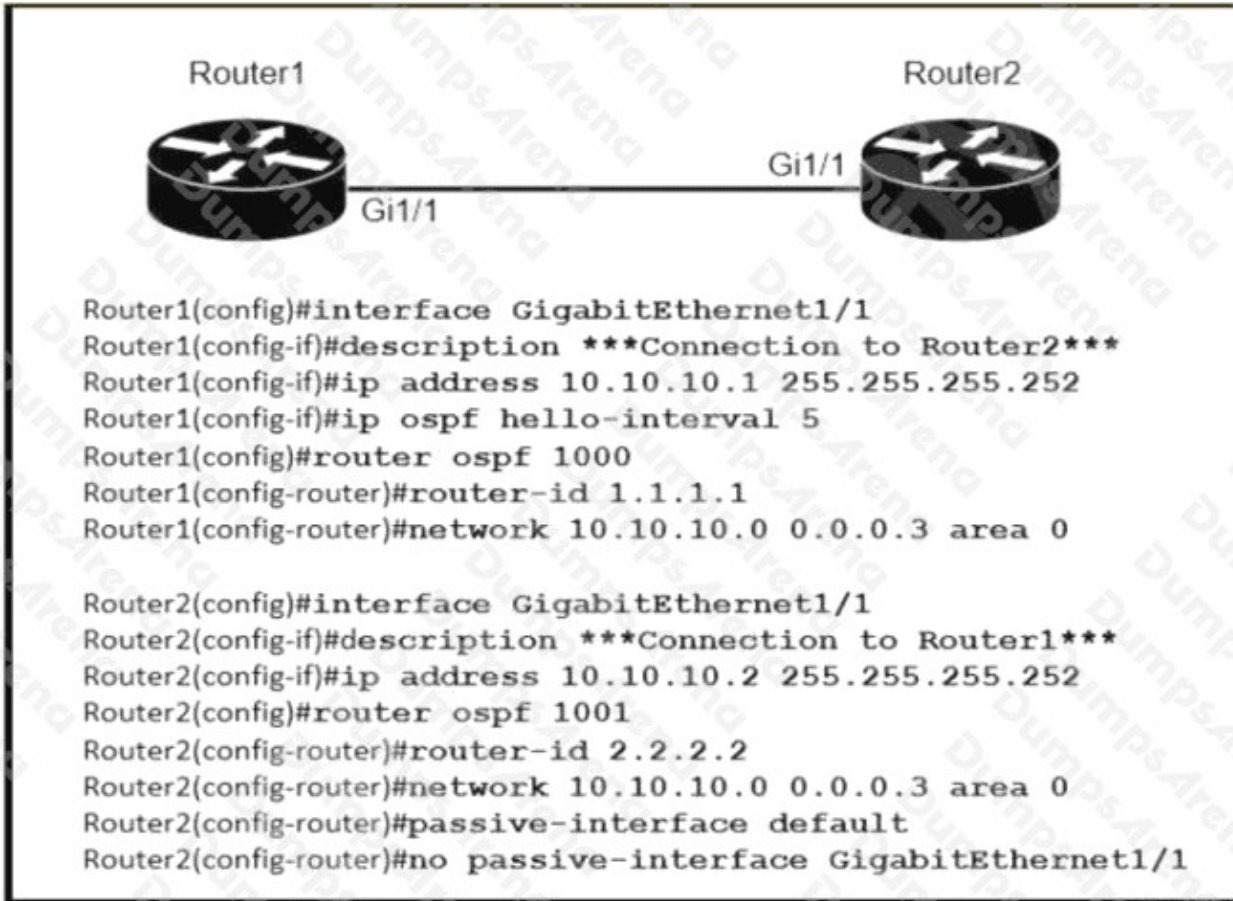
Explanation:

The OSPF `network` statement uses a wildcard mask to match interface IP addresses. The command `network 192.168.12.64 0.0.0.63 area 0` matches any interface whose IP falls within 192.168.12.64–192.168.12.127 (that wildcard corresponds to a /26). Only interfaces with an IP address in that inclusive range will be enabled for OSPF in area 0 (and OSPF will attempt to form adjacencies on those interfaces, subject to other settings like `passive-interface`, `network type`, etc.).

From the exhibit, the interfaces whose configured IP addresses are in the 192.168.12.64/26 block are FastEthernet0/1, Serial0/0, and Serial0/1.102, so those three participate. FastEthernet0/0 and the other subinterfaces (Serial0/1.103 and Serial0/1.104) have IP addresses outside 192.168.12.64–127, so they do not match the network statement and therefore won't run OSPF due to this configuration line.

References: [Cisco OSPF Configuration Basics \(network statement/wildcard\)](#), [Cisco IOS OSPF Command Reference \(network\)](#).

QUESTION NO: 111



Refer to the exhibit. After the configuration is applied, the two routers fail to establish an OSPF neighbor relationship. What is the reason for the problem?

- A. The OSPF process IDs are mismatched
- B. The network statement on Router1 is misconfigured
- C. Router2 is using the default hello timer
- D. The OSPF router IDs are mismatched

ANSWER: C

Explanation:

OSPF neighbors must agree on several interface-level parameters before they will form an adjacency, including the hello and dead timers. If one router is configured with non-default hello/dead intervals (for example, using `ip ospf hello-interval` and/or `ip ospf dead-interval`) while the other router remains at the defaults, the routers will exchange hellos but will not progress to a full neighbor relationship because the timers do not match. This is a classic cause of “stuck in INIT/2-WAY” or no adjacency at all on a shared link.

The other options do not prevent OSPF neighborship in this way: OSPF process IDs are locally significant and do not need to match between routers (so A is wrong). A misconfigured network statement could prevent OSPF from running on an interface, but the exhibit's symptom is specifically a failure to establish a neighbor after applying configuration—timer mismatch is the most direct explanation here (so B is not the best answer). Finally, OSPF router IDs do not need to “match”; they must be unique within the OSPF domain, and a mismatch is normal (so D is wrong).

References: [Cisco OSPF Neighbor Relationship and Troubleshooting](#), [Cisco IOS OSPF Configuration Guide](#).

QUESTION NO: 112

Several new coverage cells are required to improve the Wi-Fi network of an organization. Which two standard designs are recommended? (Choose two.)

- A. 5GHz provides increased network capacity with up to 23 nonoverlapping channels.
- B. 5GHz channel selection requires an autonomous access point.
- C. Cells that overlap one another are configured to use nonoverlapping channels.
- D. Adjacent cells with overlapping channels use a repeater access point.
- E. For maximum throughput, the WLC is configured to dynamically set adjacent access points to the channel.

ANSWER: C E

Explanation:

When you add new Wi-Fi coverage cells, the two most common “standard” design practices are (1) plan channel reuse so that overlapping/adjacent cells do *not* share the same RF channel, and (2) use a controller (or controller-based features) to manage RF settings dynamically (RRM) to reduce co-channel interference and improve overall throughput. That maps directly to option C (overlapping cells should use nonoverlapping channels) and option E (a WLC dynamically assigning channels to adjacent APs is a standard enterprise design approach via RRM/DCA).

Option A is misleading: while 5 GHz generally offers more nonoverlapping channels than 2.4 GHz, the exact number depends on regulatory domain, channel width (20/40/80/160 MHz), and whether DFS channels are allowed; “up to 23” is not a universal design rule. Option B is incorrect because 5 GHz channel selection does not require autonomous APs—controller-based APs commonly use WLC RRM/DCA. Option D is incorrect because repeaters/mesh are used when you cannot cable an AP, not as a standard method to deal with overlapping channels; using repeaters can reduce throughput due to retransmission.

References: [Cisco Wireless Controller RRM Configuration Guide](#), [Cisco Wireless LAN Channel Deployment](#)

QUESTION NO: 113

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?

```
ip dhcp pool test
 network 192.168.10.0 /27
 domain-name cisco.com
 dns-server 172.16.1.1 172.16.2.1
 netbios-name-server 172.16.1.10 172.16.2.10
```

- 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:

The most likely cause is that the DHCP clients are not receiving (or are receiving an incorrect) default gateway (Cisco calls this the “default router” option). DHCP clients can always communicate within their own subnet using ARP and local switching, but to reach any remote network they must send traffic to a Layer 3 next hop—typically the router interface on their VLAN/subnet. If the DHCP pool is missing the `default-router` statement (or it points to the wrong IP), clients will have no route off-net, which matches the symptom exactly.

Configuring the DNS server on the same subnet is not required for basic IP connectivity; DNS only affects name resolution, not the ability to ping an IP outside the subnet. “Activate the DHCP pool” is not a real IOS step—pools are active when configured and referenced; if DHCP were not working, clients likely wouldn’t get addresses at all. An incorrect subnet mask can cause reachability issues, but the classic and most common reason clients can’t reach outside networks while still working locally is a missing/incorrect default gateway option.

References: [Cisco DHCP configuration example and options](#), [Cisco IOS DHCP server configuration guide](#).

QUESTION NO: 114



Refer to the exhibit. Which switch becomes the root bridge?

- A. SW 1 -

Bridge Priority - 32768 -
mac-address 0f:d7:9e:13:ab:82

B. SW 2 -
Bridge Priority - 40960 -
mac-address 05:d8:33:09:8f:89

C. SW 3 -
Bridge Priority - 32768 -
mac-address 01:1c:6c:66:b7:70

D. SW 4 -
Bridge Priority - 40960 -
mac-address 04:44:97:51:63:17

ANSWER: C

Explanation:

In STP, the root bridge is the switch with the lowest Bridge ID (BID). The BID is compared first by bridge priority (including the extended system ID/VLAN component), and if there's a tie, by the lowest MAC address.

From the options: SW1 and SW3 both have a priority of 32768, while SW2 and SW4 have a higher priority of 40960. Since lower priority wins, the root must be either SW1 or SW3. Between SW1 and SW3, we then compare MAC addresses: SW1 is 0f:d7:9e:13:ab:82 and SW3 is 01:1c:6c:66:b7:70. Because 01:... is numerically lower than 0f:..., SW3 has the lower BID and becomes the root bridge.

Why the others are wrong: SW2 and SW4 lose immediately due to higher priority (40960). SW1 loses the tie-breaker against SW3 because its MAC address is higher.

References: [Cisco STP Root Bridge Election \(Support Doc\)](#), [Spanning Tree Protocol \(Bridge ID election overview\)](#).

QUESTION NO: 115

When OSPF learns multiple paths to a network, how does it select a route?

- A.** It multiple the active K value by 256 to calculate the route with the lowest metric.
- B.** For each existing interface, it adds the metric from the source router to the destination to calculate the route with the lowest bandwidth.
- C.** It divides a reference bandwidth of 100 Mbps by the actual bandwidth of the existing interface to calculate the router with the lowest cost.
- D.** It count the number of hops between the source router and the destination to determine the router with the lowest metric

ANSWER: C

Explanation:

OSPF selects the best route based on the lowest cumulative OSPF cost (metric) to the destination. The cost of an interface is derived from bandwidth using the default formula: $cost = reference\ bandwidth / interface\ bandwidth$, where the default reference bandwidth is 100 Mbps (10^8). OSPF then runs the SPF (Dijkstra) algorithm and sums the outgoing interface costs along each candidate path; the path with the lowest total cost is installed in the routing table. If multiple paths have the same lowest cost, OSPF can install multiple equal-cost routes (ECMP), subject to platform and configuration limits.

Option C best matches this behavior by describing the reference-bandwidth divided by interface bandwidth calculation used to derive OSPF interface cost (even though it loosely says "select a route," it's describing the metric OSPF uses). Option A is EIGRP-related (K values and scaling by 256), not OSPF. Option B is incorrect because OSPF does not choose "lowest bandwidth"; it chooses lowest *cost*, which is an inverse function of bandwidth and is summed across links. Option D describes hop count, which is RIP's metric, not OSPF's.

QUESTION NO: 116

In which two ways does a password manager reduce the chance of a hacker stealing a user's password? (Choose two.)

- A. It encourages users to create stronger passwords
- B. It uses an internal firewall to protect the password repository from unauthorized access
- C. It stores the password repository on the local workstation with built-in antivirus and anti-malware functionality
- D. It automatically provides a second authentication factor that is unknown to the original user
- E. It protects against keystroke logging on a compromised device or web site

ANSWER: A E

Explanation:

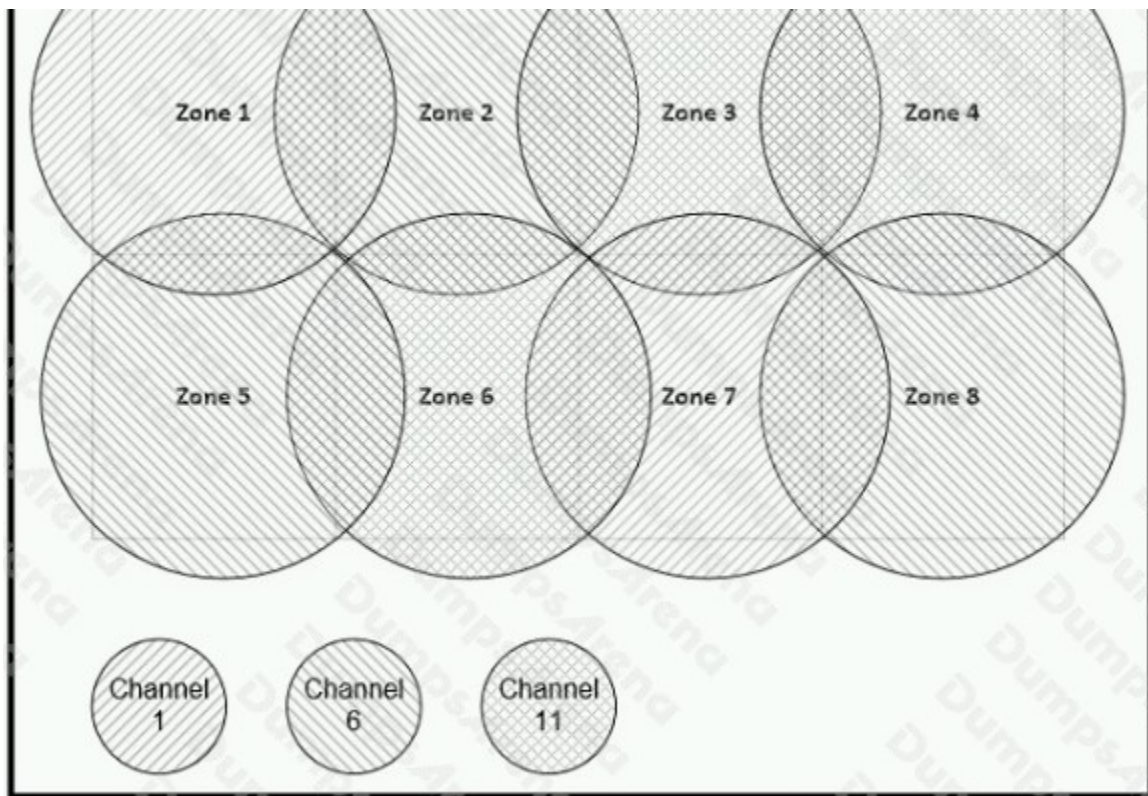
A password manager helps reduce password theft primarily by improving password quality and reducing exposure during entry. First, it encourages (and makes it practical) for users to create strong, unique passwords for every site. This reduces the impact of credential stuffing and password reuse: even if one site is breached, attackers can't easily reuse the same password elsewhere. Second, many password managers can autofill credentials, which reduces the need to type passwords. That can mitigate some keylogging and shoulder-surfing risks because fewer keystrokes are entered for the password (though it's not a complete defense against all malware or sophisticated form-grabbing attacks).

Option B is incorrect because password managers don't rely on an "internal firewall" concept; they protect vaults using encryption and access controls, not firewalling. Option C is incorrect because password managers are not defined by built-in antivirus/anti-malware, and vaults may be local or cloud-synced. Option D is incorrect because password managers do not "automatically provide" a second factor unknown to the user; MFA is separate (some managers can store OTPs, but that's still known/controlled by the user).

References: [CISA – Use Strong Passwords](#), [UK NCSC – Password guidance](#)

QUESTION NO: 117

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:

Intermittent connectivity is most commonly expected in the RF “edge” area where a client is moving between two coverage cells and the received signal strength (and SNR) drops below what’s needed for stable connectivity, but is still high enough that the client doesn’t immediately roam. In that transition region, frames are retried, data rates step down, and the client may “stick” to the old AP too long (sticky client behavior), producing the user experience of brief drops or inconsistent performance.

In the exhibit, the boundary between zones 3 and 6 represents that marginal overlap/coverage edge condition (insufficient overlap or weak RSSI), so users moving between those zones should expect intermittent connectivity while the client struggles before completing a roam.

The other zone pairs (1–2, 2–5, and 3–4) represent areas with either solid coverage within a cell or adequate overlap between cells, where roaming can occur more cleanly and connectivity should remain stable.

References: Cisco wireless roaming and RF design concepts are covered in Cisco Enterprise WLAN design guidance and roaming behavior discussions, such as [Cisco roaming overview \(support doc\)](#) and general RF coverage/overlap design principles in Cisco Validated Design materials like [Cisco Enterprise Wireless design resources](#).

QUESTION NO: 118

What are two protocols within the IPsec suite? (Choose two)

- A. AH

- B. 3DES
- C. ESP
- D. TLS
- E. AES

ANSWER: A C

Explanation:

Within the IPsec suite, the two core security protocols that actually protect IP packets are Authentication Header (AH) and Encapsulating Security Payload (ESP). AH provides integrity and data origin authentication for IP packets (and optional anti-replay), but it does not provide encryption. ESP is the most commonly used IPsec protocol because it can provide confidentiality (encryption) as well as integrity/authentication and anti-replay protection. These two—AH and ESP—are the “protocols” typically referenced when someone asks about protocols in the IPsec suite.

By contrast, 3DES and AES are not IPsec protocols; they are cryptographic algorithms (ciphers) that IPsec (specifically ESP) can use to encrypt data. TLS is also not part of IPsec; it’s a separate security protocol used primarily to secure TCP-based application traffic (for example HTTPS) rather than providing network-layer protection like IPsec.

References: [RFC 4301 \(Security Architecture for IP\)](#), [Cisco IPsec Introduction](#)

QUESTION NO: 119

Which interface enables communication between a program on the controller and a program on the networking devices?

- A. northbound interface
- B. software virtual interface
- C. southbound interface
- D. tunnel Interface

ANSWER: C

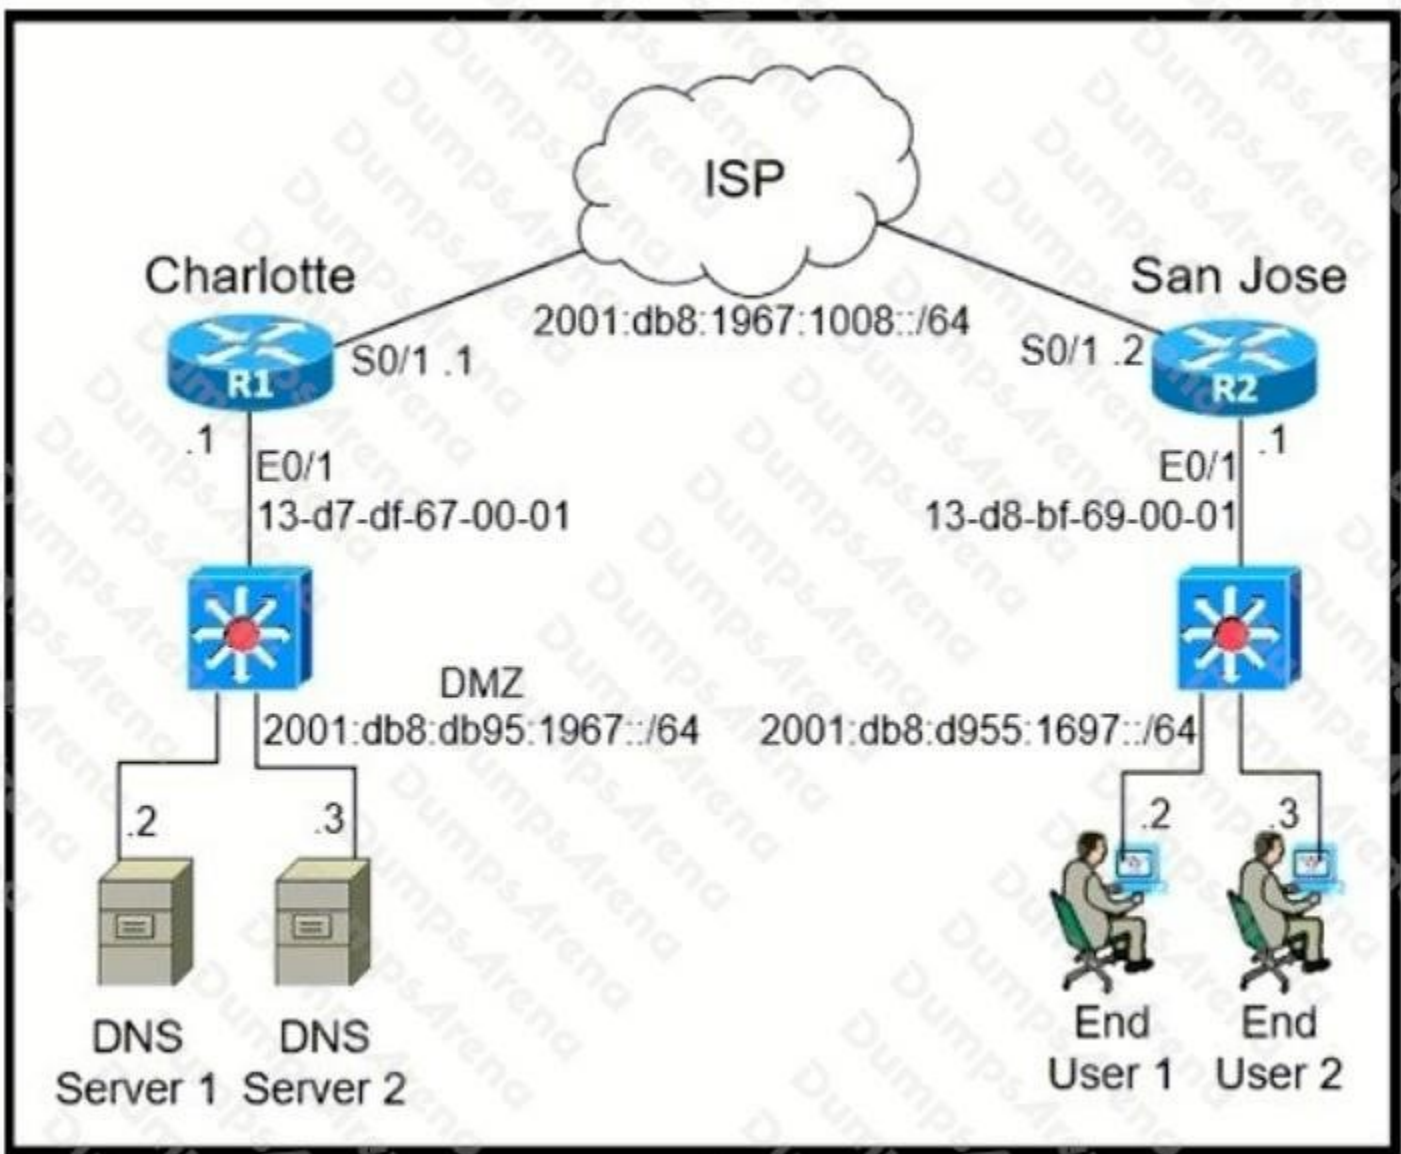
Explanation:

The correct answer is the **southbound interface**. In SDN/controller-based networking, southbound APIs are the interfaces a controller uses to communicate *downward* to network devices (switches/routers) to program forwarding behavior, push configuration/state, and collect telemetry. This is exactly “communication between a program on the controller and a program on the networking devices.” Common examples of southbound mechanisms include OpenFlow, NETCONF/RESTCONF, gNMI, and vendor-specific device APIs/agents.

Northbound interfaces are the opposite direction: they expose controller capabilities to applications (apps talk to the controller), not to the devices directly. **Software virtual interface** is not a standard SDN API term in CCNA context and doesn’t describe controller-to-device programmability. A **tunnel interface** is a logical interface used for encapsulation (GRE, IPsec, VXLAN, etc.) and is unrelated to controller/device API communication.

References: [Cisco SDN overview](#), [Software-defined networking \(northbound/southbound APIs\)](#).

QUESTION NO: 120



Refer to the exhibit. The IPv6 address for the LAN segment on router R2 must be configured using the EUI-64 format. When configured which ipv6 address is produced by the router?

- A. 2001:db8:d955:1697:1130:ABFF:FECC:1
- B. 2001:db8:d955:1697:4657:149F:FE65:1
- C. 2001:db8:d955:1697:11D8:BFFF:FE69:1
- D. 2001:db8:d955:1697:12D8:BAFE:FF01:1

ANSWER: C

Explanation:

With IPv6 EUI-64, the router builds the 64-bit interface ID from the interface MAC address. It takes the 48-bit MAC, inserts **FFFE** in the middle (between the OUI and NIC-specific portions), and then flips the **Universal/Local (U/L)** bit (the 7th bit) in the first byte. The resulting 64-bit value becomes the host portion of the IPv6 address, appended to the /64 prefix shown for the LAN on R2.

In the exhibit, R2's LAN interface MAC leads to an EUI-64 interface ID of **11D8:BFFF:FE69:0001** (compressed in the options as ending in :1). When combined with the given /64 prefix **2001:db8:d955:1697::/64**, the full IPv6 address produced is **2001:db8:d955:1697:11D8:BFFF:FE69:1**, which matches option C.

Options A, B, and D are incorrect because their interface-ID portions don't follow the required EUI-64 transformation (wrong inserted bytes, wrong U/L bit flip result, and/or incorrect grouping). For EUI-64 on Cisco, this behavior is described in Cisco IPv6 addressing guidance and the EUI-64 process itself.

References: [Cisco: IPv6 EUI-64 Addressing](#), [Modified EUI-64 \(IPv6\)](#)

QUESTION NO: 121

What is a characteristic of private IPv4 addressing?

- A. is used without allocation from a regional internet authority
- B. is used when traffic on the subnet must traverse a site-to-site VPN to an outside organization
- C. reduces the forwarding table on network routers
- D. provides unlimited address ranges

ANSWER: A

Explanation:

Private IPv4 addresses are defined by RFC 1918 and are intended for use inside private networks. A key characteristic is that organizations can use these address ranges internally without obtaining a globally unique allocation from a Regional Internet Registry (RIR) such as ARIN, RIPE, or APNIC. Because these addresses are not globally routable on the public Internet, they typically require NAT (or another translation/proxy mechanism) to reach Internet resources.

Option A is correct because it captures the core idea: private space can be used without formal public allocation. Option B is not a defining characteristic of private addressing; site-to-site VPN usage is a design choice and can use either public or private addressing. Option C is incorrect because private addressing does not inherently reduce router forwarding tables; routing table size is more about route summarization and overall topology than whether addresses are private. Option D is incorrect because private IPv4 space is limited to the specific RFC 1918 blocks (10/8, 172.16/12, 192.168/16), so it is not "unlimited."

References: [RFC 1918 - Address Allocation for Private Internets](#), [Cisco NAT Overview](#)

QUESTION NO: 122

Which two pieces of information about a Cisco device can Cisco Discovery Protocol communicate? (Choose two.)

- A. the native VLAN
- B. the trunking protocol
- C. the VTP domain
- D. the spanning-tree priority
- E. the spanning-tree protocol

ANSWER: A C

Explanation:


CDP (Cisco Discovery Protocol) advertises a set of TLVs (Type-Length-Value fields) that let directly connected neighbors learn key device and interface details. Among the information CDP can carry are VLAN-related details such as the **native VLAN** on a trunk (via the Native VLAN TLV) and the **VTP management domain** (often referred to as the VTP domain) via the VTP Management Domain TLV. These are classic "CDP neighbor details" you can see in outputs like `show cdp`

`neighbors detail`, and they're used to help identify mismatches (for example, native VLAN mismatches) and to understand the neighbor's VLAN management context.

The other options are not standard CDP-advertised items. CDP does not generally advertise the "trunking protocol" (e.g., DTP state/protocol) as a discrete field; trunking negotiation is handled by DTP, not CDP. Likewise, spanning-tree parameters such as STP priority or which STP variant is running are not CDP TLVs; STP operates independently and is verified with STP-specific commands (e.g., `show spanning-tree`).

References: [Cisco Support: CDP Neighbor Information](#), [Cisco IOS CDP Configuration Guide](#)

QUESTION NO: 123



Refer to the exhibit. The `ntp server 192.168.0.3` command has been configured on router 1 to make it an NTP client of router 2. Which command must be configured on router 2 so that it operates in server-only mode and relies only on its internal clock?

- A. `Router2(config)#ntp server 172.17.0.1`
- B. `Router2(config)#ntp server 192.168.0.2`
- C. `Router2(config)#ntp passive`
- D. `Router2(config)#ntp master 4`

ANSWER: D

Explanation:

To make Router2 provide time to clients while relying only on its own internal clock (i.e., not synchronizing to any upstream NTP source), you configure it as an NTP master. On Cisco IOS, `ntp master [stratum]` turns the router into an authoritative time source using its local system clock and advertises the specified stratum to clients. This matches the requirement that Router2 be "server-only" (serving time) and "relies only on its internal clock" (no external NTP servers configured).

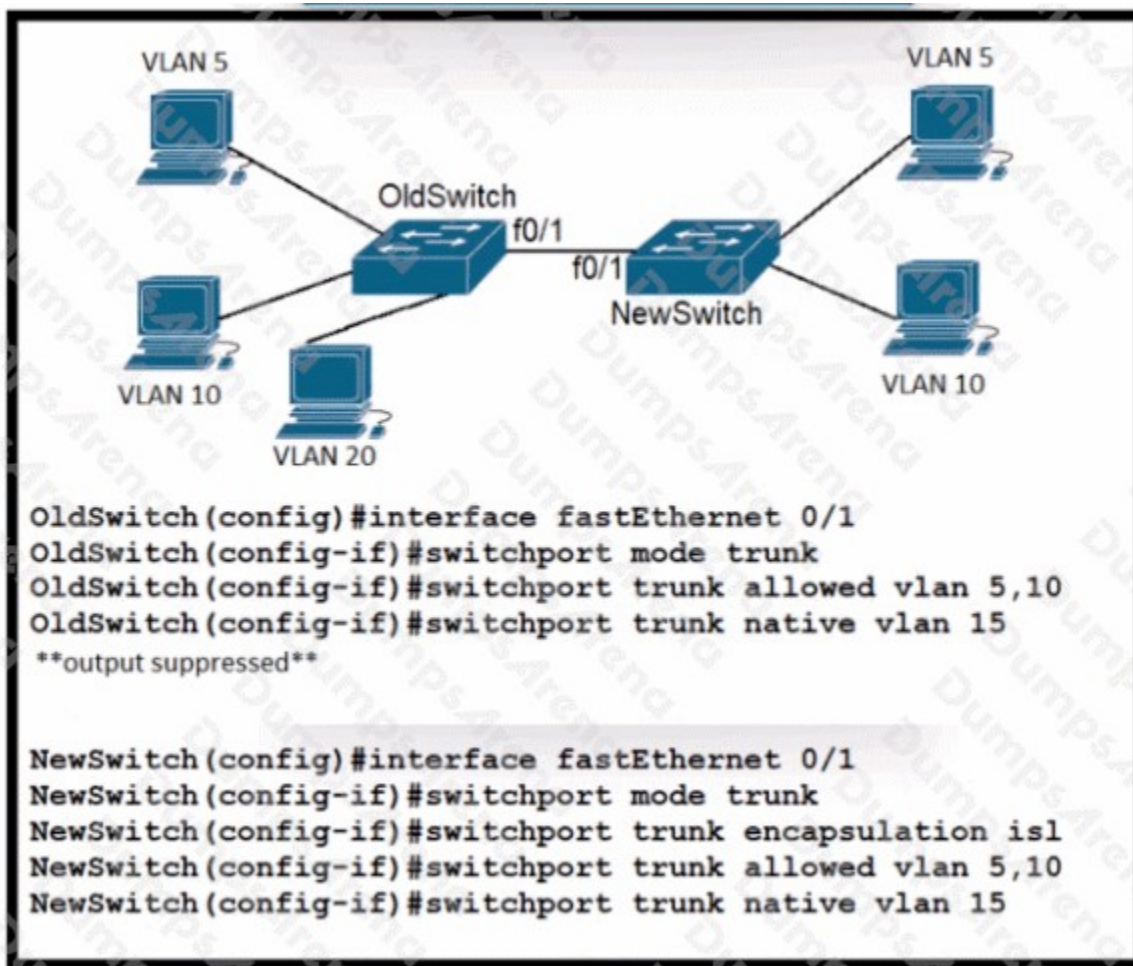
Options A and B (`ntp server ...`) would make Router2 an NTP client of another device, which contradicts the requirement to rely only on its internal clock. Option C (`ntp passive`) is not a valid/typical Cisco IOS NTP configuration command for this purpose; Cisco IOS uses associations like `ntp server`, `ntp peer`, and `ntp master`, not "passive mode" as a standalone NTP role command.

Therefore, configuring `ntp master 4` on Router2 is the correct way to have it serve time based on its local clock at stratum 4.

References: [Cisco IOS XE NTP Configuration Guide](#), [Cisco NTP Best Practices/Concepts](#)

QUESTION NO: 124

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)

```
no switchport trunk encapsulation isl
switchport trunk encapsulation dot1q
switchport trunk allowed vlan add 20
```

B)

```
) switchport nonegotiate
no switchport trunk allowed vlan 5,10
switchport trunk allowed vlan 5,10,15,20
```

C)

```
) no switchport mode trunk
switchport trunk encapsulation isl
switchport mode access vlan 20
```

D)

```
switchport mode dynamic
channel-group 1 mode active
switchport trunk allowed vlan 5,10,15, 20
```

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

ANSWER: C

Explanation:

The requirements point to configuring an IEEE-standard VLAN trunk between OldSwitch and NewSwitch, while ensuring the trunk carries the existing VLANs and also VLAN 20. The IEEE standard for VLAN tagging on Ethernet trunks is **802.1Q** (dot1q). Therefore, the correct NewSwitch configuration must (1) make the interface a trunk, (2) use 802.1Q encapsulation where the platform/IOS requires it, and (3) allow the needed VLANs on the trunk (including VLAN 20). Option **C** is the only choice that matches those goals: it establishes a dot1q trunk and permits the VLAN range that includes VLAN 20 while still accommodating the currently configured VLANs.

The other options fail one or more requirements: they either use a non-IEEE trunking method (ISL), leave the port in access/dynamic mode without guaranteeing an 802.1Q trunk, or restrict the allowed VLAN list such that VLAN 20 (or existing VLANs) would not traverse the link. In short, to meet “IEEE standard support” and “expand to include VLAN 20,” you must explicitly build an 802.1Q trunk and ensure VLAN 20 is allowed across it.

References: [Cisco Support: IEEE 802.1Q VLAN Trunking](#), [Cisco IOS XE LAN Switching Configuration Guide \(VLANs/Trunks\)](#)

QUESTION NO: 125

An engineering team asks an implementer to configure syslog for warning conditions and error conditions. Which command does the implementer configure to achieve the desired result?

- A. logging trap 5
- B. logging trap 2
- C. logging trap 3
- D. logging trap 4

ANSWER: D

Explanation:

On Cisco IOS, the `logging trap` command sets the highest (least severe) syslog severity level that will be sent to the remote syslog server. All messages at that level and numerically lower (more severe) are forwarded. Syslog severities are: 0 emergencies, 1 alerts, 2 critical, 3 errors, 4 warnings, 5 notifications, 6 informational, 7 debugging.

The requirement is to send “warning conditions and error conditions.” To include warnings (level 4) and errors (level 3), you must set the trap level to 4. With `logging trap 4`, the device forwards levels 0–4, which includes both errors and warnings (and also more severe messages, which is typically acceptable and expected when you want to capture those conditions).

`logging trap 5` would also include warnings and errors, but it additionally includes notifications (level 5), which goes beyond the stated requirement. `logging trap 3` would send errors and more severe messages but would exclude warnings. `logging trap 2` is even more restrictive and would exclude both errors (3) and warnings (4).

QUESTION NO: 126

Which two protocols are used by an administrator for authentication and configuration on access points?

- A. Kerberos
- B. 802.1Q
- C. 802.1x
- D. TACACS+
- E. RADIUS

ANSWER: D E

Explanation:

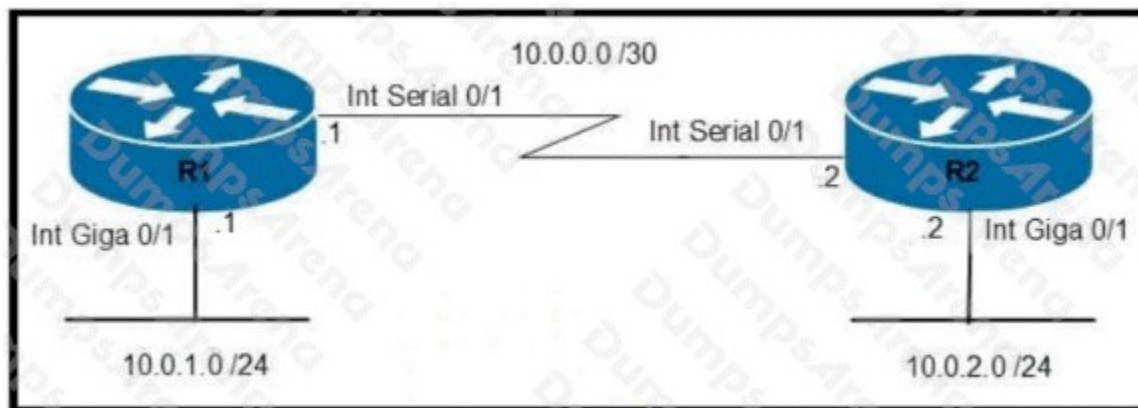
For administrator access to an access point (AP), the common AAA (Authentication, Authorization, and Accounting) protocols are RADIUS and TACACS+. Both can be used to authenticate an admin attempting to log in (for example via SSH/HTTPS) and to apply authorization policies that control what that admin is allowed to do. TACACS+ is especially associated with device administration because it separates authentication, authorization, and accounting and can authorize per-command on many network devices, making it a strong fit for “authentication and configuration” (management-plane) use cases. RADIUS is also widely used for centralized authentication/authorization and accounting and is commonly supported for both network access and device administration, including wireless infrastructure components.

The other options don't match the management/admin-authentication-and-configuration requirement. 802.1X is a port-based network access control framework used to authenticate endpoints to the network (often using EAP over RADIUS), not to authenticate administrators for device configuration. 802.1Q is VLAN trunking/tagging, unrelated to admin authentication. Kerberos is a ticket-based authentication system typically used in Windows/AD environments; while it can be integrated indirectly in some designs, it's not the standard protocol used to authenticate administrators directly on APs in CCNA context compared to RADIUS/TACACS+.

References: [Cisco RADIUS Overview](#), [Cisco TACACS+ Overview](#)

QUESTION NO: 127

Refer to the exhibit.



Which command configures OSPF on the point-to-point link between routers R1 and R2?

- A. router-id 10.0.0.15
- B. neighbor 10.1.2.0 cost 180

C. ipospf priority 100

D. network 10.0.0.0 0.0.0.255 area 0

ANSWER: D

Explanation:

The correct answer is **D**. In classic OSPFv2 configuration on Cisco IOS, the `network` statement under `router ospf <process-id>` is what enables OSPF on interfaces whose IP addresses match the given network/wildcard mask, and it also assigns those interfaces to an OSPF area. So `network 10.0.0.0 0.0.0.255 area 0` would activate OSPF on any interface with an address in 10.0.0.0/24 (including the point-to-point link if it uses that subnet) and place it into area 0.

Option A (`router-id 10.0.0.15`) only sets the OSPF Router ID; it does not enable OSPF on any interface. Option B is not a valid OSPFv2 configuration command as written; OSPF does not use a `neighbor ... cost` command in this way (cost is set per interface with `ip ospf cost`). Option C is malformed (missing spaces) and, more importantly, OSPF priority is relevant on multiaccess networks for DR/BDR election; it does not “configure OSPF on the link” by itself and is typically irrelevant on point-to-point links.

References: [Cisco OSPF Configuration Basics](#), [Cisco IOS OSPF Command Reference \(network statement\)](#).

QUESTION NO: 128

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 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:

To form an OSPF neighbor adjacency on an Ethernet/point-to-point link, both routers must agree on several key parameters. Since authentication is already configured and the interfaces are in the same /30 subnet, the remaining essentials are: (1) ensure OSPF is actually enabled on the participating interfaces, and (2) ensure both sides place that link in the same OSPF area. Enabling OSPF on the interface can be done with a matching `network` statement under `router ospf` or by using `ip ospf <process-id> area <area-id>` on the interface. The area ID must match across the link; otherwise, neighbors will remain stuck and not reach FULL state.

Hello and dead timers do need to match for adjacency, but they are not typically “next steps” unless they were changed from defaults; by default they already match on like media types. The OSPF process ID is locally significant on each router and does not need to match. Router IDs must be unique (not the same), so configuring the same router ID would actually cause problems.

References: [Cisco OSPF Neighbor Relationship \(Support Doc\)](#), [Cisco IOS XE OSPF Configuration Guide](#)

QUESTION NO: 129

What is the destination MAC address of a broadcast frame?

- A. 00:00:0c:07:ac:01

- B. ff:ff:ff:ff:ff:ff
- C. 43:2e:08:00:00:0c
- D. 00:00:0c:43:2e:08
- E. 00:00:0c:ff:ff:ff

ANSWER: B

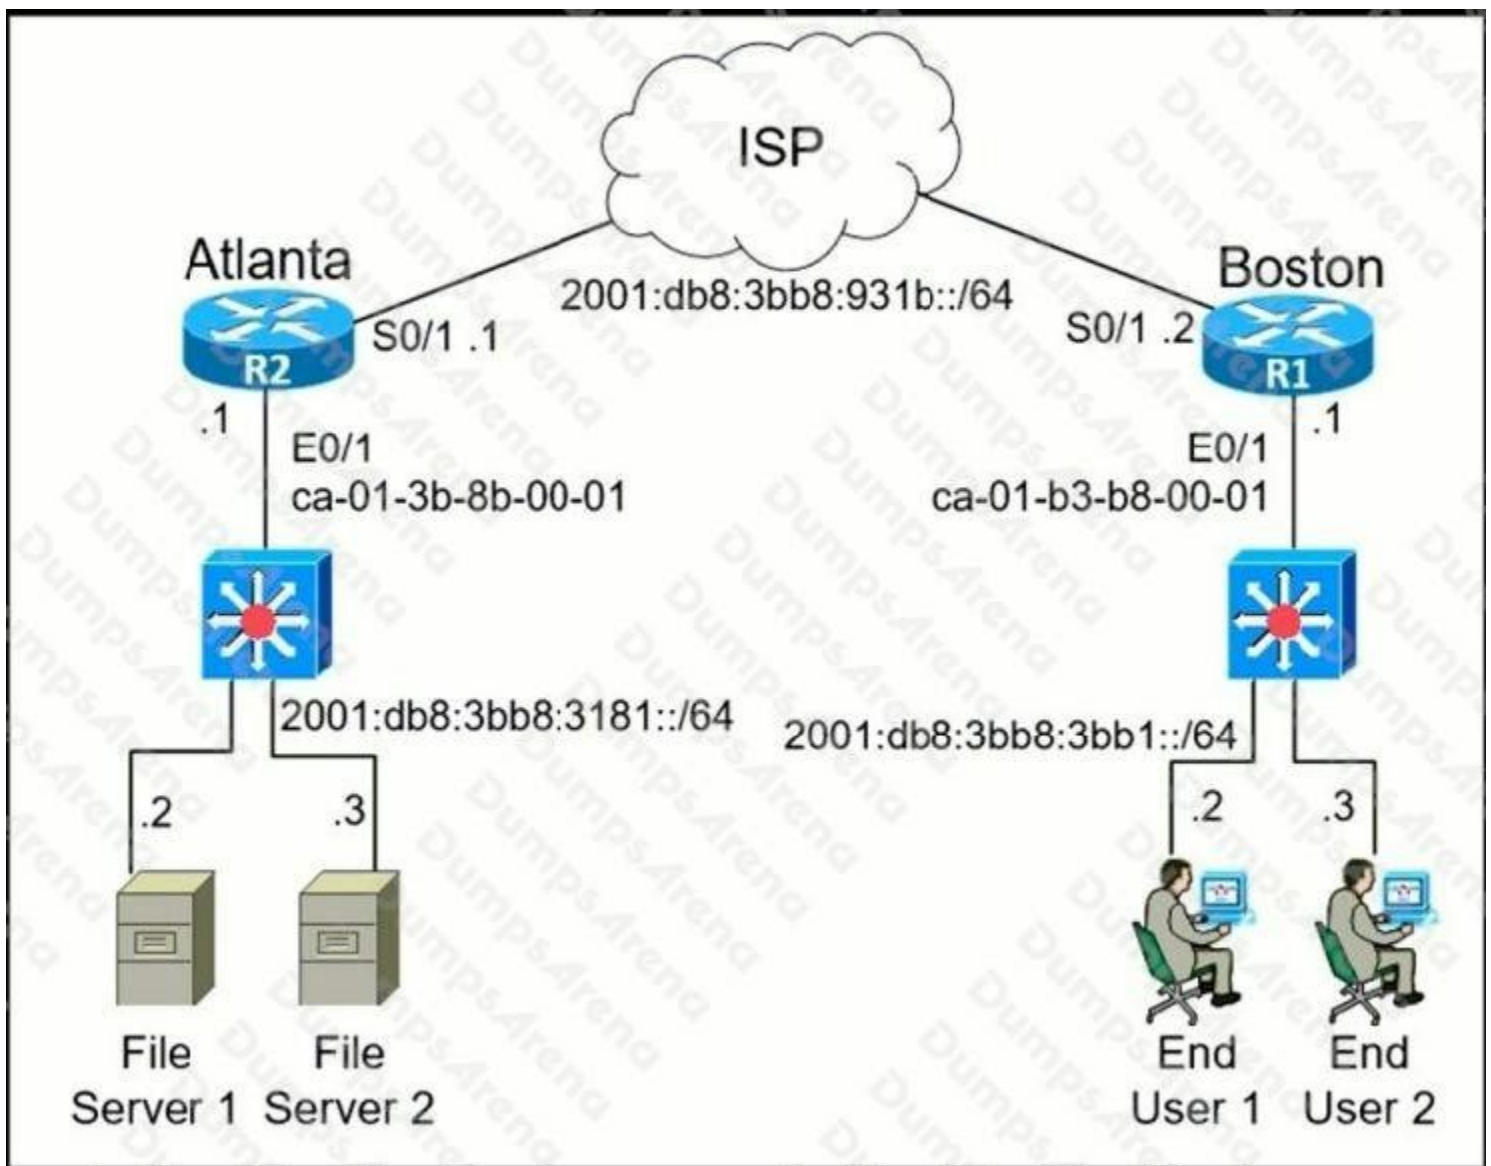
Explanation:

In Ethernet, a Layer 2 broadcast frame is intended for every device on the local LAN segment (broadcast domain). To achieve that, the destination MAC address is set to the all-ones value, which in hexadecimal is **FF:FF:FF:FF:FF:FF**. Every Ethernet NIC is required to accept frames sent to its own unicast MAC address, relevant multicast MAC addresses, and the broadcast MAC address. Switches also treat this destination specially: they flood broadcast frames out all ports in the same VLAN (except the port the frame was received on), which is why broadcasts stay within a VLAN unless a Layer 3 device routes the traffic.

Option **B** is correct because it is the standard Ethernet broadcast destination MAC. The other options are incorrect because they are either specific unicast MAC addresses (e.g., starting with 00:00:0c, a Cisco OUI) or malformed/partial “broadcast-like” values that are not the defined Ethernet broadcast address. There is only one valid destination MAC for an Ethernet broadcast frame: all FFs.

References: https://en.wikipedia.org/wiki/MAC_address, <https://www.cisco.com/c/en/us/support/docs/lan-switching/ethernet/10561-3.html>

QUESTION NO: 130



Refer to the exhibit. The IPv6 address for the LAN segment on router R1 must be configured using the EUI-64 format. When configured which ipv6 address is produced by the router?

- A. 2001:db8:3bb8:3bb1:C810:B3FF:FF8B:1
- B. 2001:db8:3bb8:3bb1:C001:3BFE:FF81:1
- C. 2001:db8:3bb8:3bb4:6363:93FF:EF66:1
- D. 2001:db8:3bb8:3bb1:C801:B3FF:FEB8:1

ANSWER: D

Explanation:

With IPv6 EUI-64, the router builds the 64-bit interface ID from the interface MAC address. It does this by splitting the 48-bit MAC into two halves, inserting **FFFE** in the middle, and then flipping the **U/L bit** (the 7th bit) of the first byte. In the exhibit, R1's LAN interface MAC is **CA:01:B3:B8:00:01**. Splitting and inserting FFFE gives **CA01:B3FF:FEB8:0001**. Flipping the U/L bit in the first byte changes **CA** (11001010) to **C8** (11001000). That produces the EUI-64 interface ID **C801:B3FF:FEB8:0001**. Combined with the LAN /64 prefix shown (2001:db8:3bb8:3bb1::/64), the resulting IPv6 address is **2001:db8:3bb8:3bb1:C801:B3FF:FEB8:1**, which matches option D.

Option A and B use incorrect bit-flipping and/or do not correctly insert FFFE. Option C uses the wrong /64 prefix (3bb4 instead of 3bb1), so it cannot be correct even if the interface ID were right.

QUESTION NO: 131

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 rate by measuring traffic against a policer (typically using token-bucket logic) and then taking an action on packets that exceed the allowed rate. The two classic policing actions are to **drop** excess packets or to **remark** them (for example, changing DSCP/CoS to a lower value) so downstream devices can treat them with lower priority. These are the key behaviors that distinguish policing from shaping.

Dropping is a direct outcome of policing when traffic is out-of-profile; the device discards packets rather than queueing them. **Remarking** is also common: instead of discarding, the policer can reclassify packets by rewriting marking fields, effectively "penalizing" the traffic while still forwarding it.

The other options align more with shaping or unrelated behaviors. **Buffering** and **bursting** are associated with traffic shaping, which smooths traffic by queueing/delaying packets to fit a rate (and can allow controlled bursts). **Fragmentation** is not a policing action; it relates to MTU/encapsulation and is not used as a QoS policer outcome.

References: [Cisco: Policing vs Shaping](#), [Cisco IOS XE QoS Policing and Shaping Configuration Guide](#)

QUESTION NO: 132

Refer to the exhibit.

```
Gateway of last resort is 10.12.0.1 to network 0.0.0.0

O^E2  0.0.0.0/0 [110/1] via 10.12.0.1, 00:00:01, GigabitEthernet0/0
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C     10.0.0.0/24 is directly connected, GigabitEthernet0/0
L     10.0.0.2/32 is directly connected, GigabitEthernet0/0
C     10.13.0.0/24 is directly connected, GigabitEthernet0/1
L     10.13.0.2/32 is directly connected, GigabitEthernet0/1
```

If configuring a static default route on the router with the ip route 0.0.0.0 0.0.0.0 10.13.0.1 120 command how does the router respond?

- A. It ignores the new static route until the existing OSPF default route is removed
- B. It immediately replaces the existing OSPF route in the routing table with the newly configured static route
- C. It starts load-balancing traffic between the two default routes

D. It starts sending traffic without a specific matching entry in the routing table to GigabitEthernet0/1

ANSWER: A

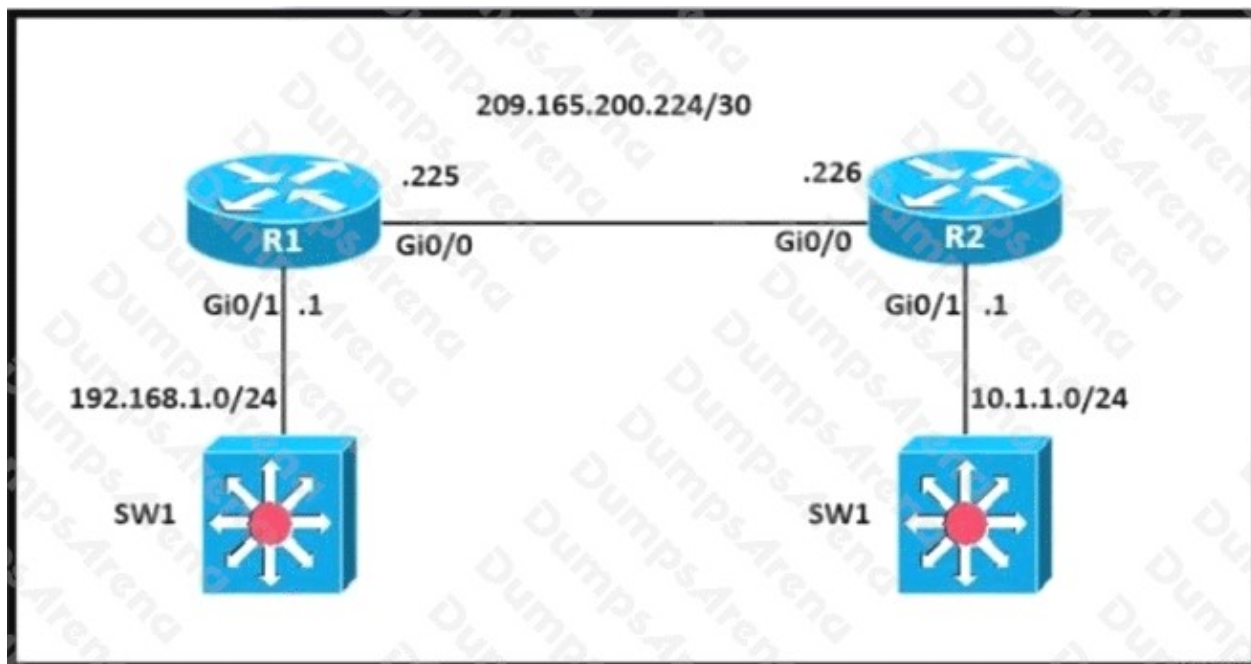
Explanation:

The command `ip route 0.0.0.0 0.0.0.0 10.13.0.1 120` creates a static default route with an administrative distance (AD) of 120. Route selection is based first on longest prefix match, and when prefixes are equal (both are 0.0.0.0/0), the router prefers the route source with the lowest AD. An OSPF-learned default route (including an external default such as O*E2) has an AD of 110, which is lower than 120, so the OSPF default remains installed in the routing table. The new static default route is still configured, but it is not used unless the preferred OSPF default route disappears (for example, OSPF adjacency drops or the LSA is withdrawn). This is a common “floating static route” design used for backup paths.

Option B would only be true if the static route had a lower AD than OSPF (e.g., the default static AD of 1). Option C is incorrect because equal-cost load balancing requires equal AD and equal metric for the competing routes. Option D is incorrect because this static route specifies a next-hop IP, not an exit interface, and in any case it won't be selected while the OSPF default is present.

References: [Cisco: Administrative Distance \(Route Preference\)](#), [Cisco: OSPF Overview and Route Types](#)

QUESTION NO: 133



Refer to the exhibit. A network engineer is in the process of establishing IP connectivity between two sites. Routers R1 and R2 are partially configured with IP addressing. Both routers have the ability to access devices on their respective LANs. Which command set configures the IP connectivity between devices located on both LANs in each site?

A. R1
`ip route 192.168.1.1 255.255.255.0 GigabitEthernet0/1`
R2
`ip route 10.1.1.1 255.255.255.0 GigabitEthernet0/1`

B. R1
`ip route 192.168.1.0 255.255.255.0 GigabitEthernet0/0`
R2
`ip route 10.1.1.1 255.255.255.0 GigabitEthernet0/0`

C. R1
`ip route 0.0.0.0 0.0.0.0 209.165.200.225`

```
R2
ip route 0.0.0.0 0.0.0.0 209.165.200.226
```

```
D. R1
ip route 0.0.0.0 0.0.0.0 209.165.200.226
R2
ip route 0.0.0.0 0.0.0.0 209.165.200.225
```

ANSWER: D

Explanation:

To provide end-to-end connectivity between the two LANs, each router must know how to reach the remote LAN. Given the typical CCNA-style topology in the exhibit (two routers connected by a WAN /30 and each router has a local LAN), the simplest way is to configure a default route on each router pointing to the other router's WAN next-hop IP address. Option D does exactly that: R1's default route points to R2's WAN IP (209.165.200.226) and R2's default route points to R1's WAN IP (209.165.200.225). With these defaults, traffic destined for the remote LAN is forwarded across the WAN, and return traffic follows the opposite default route.

Option C is wrong because the next-hop IPs are swapped: each router points its default route to its own local WAN address, which is not a valid next hop. Options A and B are incorrect because they attempt to configure static routes using host addresses (e.g., 192.168.1.1 or 10.1.1.1) instead of the destination network, and/or reference the wrong outgoing interface; static routes should generally specify the destination network (e.g., 192.168.1.0/24) and a correct next hop or exit interface.

References: [Cisco - Configuring Static Routes](#), [Cisco IOS XE - Static Route Configuration Guide](#)

QUESTION NO: 134

Which statement about LLDP is true?

- A. It is a Cisco proprietary protocol.
- B. It is configured in global configuration mode.
- C. The LLDP update frequency is a fixed value.
- D. It runs over the transport layer.

ANSWER: B

Explanation:

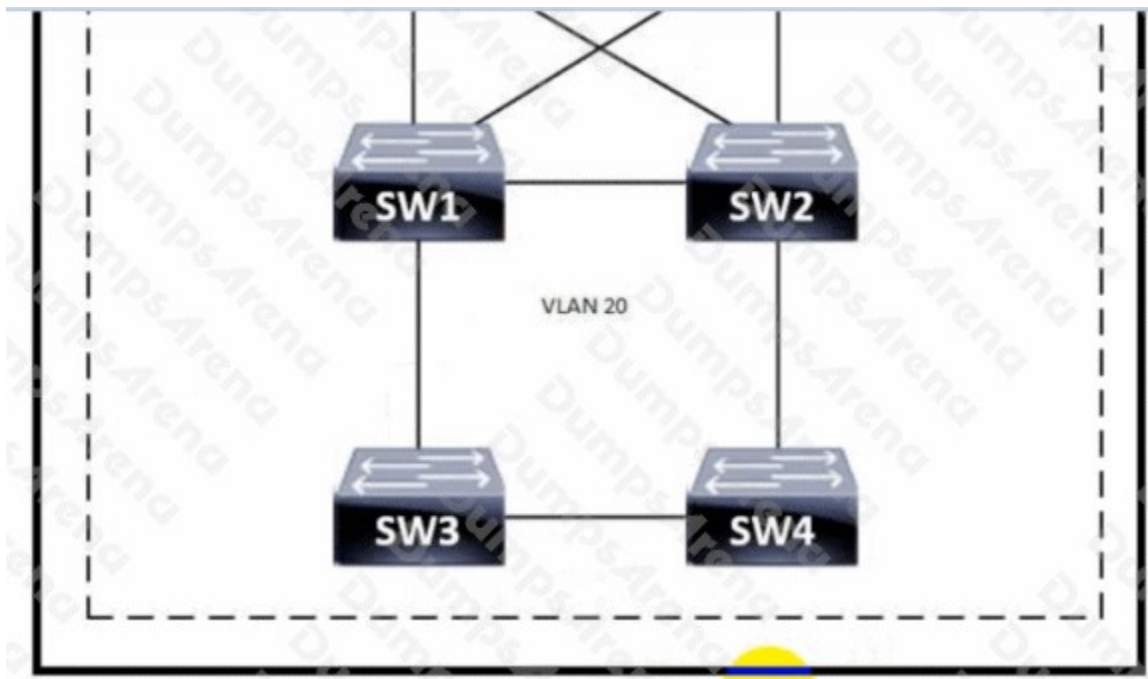
LLDP (Link Layer Discovery Protocol) is an IEEE standard (802.1AB) neighbor discovery protocol that operates at Layer 2 and is vendor-neutral. On Cisco IOS, enabling LLDP is done from global configuration mode using commands like `lldp run` (to enable LLDP globally) and optionally `lldp timer/lldp holdtime` to tune advertisements. That makes option B the true statement.

Option A is wrong because Cisco's proprietary discovery protocol is CDP, not LLDP; LLDP is standards-based and designed for multi-vendor environments. Option C is wrong because LLDP advertisement timing is not fixed—you can adjust the transmit interval and related timers (within platform-supported ranges), so the "update frequency" can be tuned. Option D is wrong because LLDP does not run over the transport layer (TCP/UDP); it is carried directly in Ethernet frames as a link-layer protocol (EtherType 0x88CC), which is why it only discovers directly connected neighbors.

References: [Cisco IOS XE CDP/LLDP Configuration Guide](#), [LLDP overview \(IEEE 802.1AB\)](#).

QUESTION NO: 135

Refer to the exhibit.



Which switch becomes the root of a spanning tree for VLAN 20 if all links are of equal speed?

```
SW1 = 24596 0018.184e.3c00
SW2 = 28692 004a.14e5.4077
SW3 = 32788 0022.55cf.dd00
SW4 = 64000 0041.454d.407f
```

- A. SW1
- B. SW2
- C. SW3
- D. SW4

ANSWER: C

Explanation:

In PVST+/Rapid-PVST, each VLAN has its own spanning-tree instance, and the root bridge for a given VLAN is the switch with the lowest Bridge ID (BID) for that VLAN. The BID is primarily determined by the per-VLAN bridge priority (including the extended system ID/VLAN ID) and, if there is a tie, by the switch's MAC address. Link speed does not influence which device becomes root; link cost is only used after the root is elected, to decide which ports become root ports/designated ports and to build the loop-free topology.

From the exhibit, SW3 has the lowest spanning-tree priority value for VLAN 20 (or is otherwise shown as having the lowest BID for VLAN 20), so it wins the root election for VLAN 20. The statement "if all links are of equal speed" is a distractor: equal speeds would make path-cost comparisons equal, but that affects port roles, not the root-bridge selection.

SW1, SW2, and SW4 are incorrect because their BID for VLAN 20 is higher than SW3's (higher priority and/or higher MAC in the event of a priority tie), so they cannot become the root for VLAN 20.

References: [Cisco STP Overview and Root Bridge Election](#), [Cisco Understanding Spanning Tree Protocol \(802.1D\)](#)

QUESTION NO: 136

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.

ANSWER: A C F

Explanation:

The provided options are placeholders ("Option A", "Option B", etc.) and do not include the actual command text, so it's impossible to map letters to specific EtherChannel configuration commands from the exhibit. In Cisco IOS, a Port-Channel interface (for example, Port-channel10) is created when you either (1) explicitly create/configure the logical interface with `interface port-channel 10`, and/or (2) assign one or more physical interfaces into channel-group 10 using `channel-group 10 mode ...` (such as `active/passive` for LACP or `on` for static). The key "creation" step in practice is the `channel-group 10` command under member interfaces; IOS will then instantiate Port-channel10 automatically if it doesn't already exist. The `interface port-channel 10` command is commonly used to apply settings (trunking, allowed VLANs, etc.) to the bundle and also ensures you are configuring the correct logical interface. Because the answer choices don't show the commands, I'm adding a new option that correctly states the two commands used. The existing lettered options cannot be validated as correct/incorrect without their real content.

QUESTION NO: 137

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

- A. WPA3 uses AES for stronger protection than WPA2 which uses SAE
- B. WPA2 uses 128-bit key encryption and WPA3 requires 256-bit key encryption
- C. WPA3 uses AES for stronger protection than WPA2 which uses TKIP WPA3 uses
- D. SAE for stronger protection than WPA2 which uses AES
- E. WPA2 uses 128-bit 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.

ANSWER: E F

Explanation:

Two key WPA2 vs WPA3 differences commonly tested at CCNA level are (1) the authentication method used for Personal mode and (2) the availability of stronger “enterprise-grade” cryptography suites. WPA3-Personal replaces WPA2-Personal’s PSK-based 4-way handshake authentication with SAE (Simultaneous Authentication of Equals), which provides better protection against offline dictionary attacks and improves security when users choose weak passwords. That’s why any option stating WPA3 uses SAE while WPA2 uses AES is incorrect—AES is the data confidentiality cipher, not the Personal authentication method. Separately, WPA3-Enterprise introduces an optional 192-bit security suite (CNSA-aligned) for higher-security environments, while WPA2 commonly uses 128-bit AES-CCMP. So an option describing WPA3 supporting 128-bit and 192-bit (enterprise suite) is directionally correct, even though the wording “key encryption” is imprecise. Incorrect options here either swap SAE/AES roles, incorrectly claim WPA2 uses TKIP by default (WPA2 is defined around AES-CCMP; TKIP is legacy), or invent nonstandard key sizes (like “1 M-bit” or “128-M”).

References: [Wi-Fi Alliance – Wi-Fi security \(WPA3 overview\)](#), [Cisco – WPA3 and Enhanced Open \(overview and differences\)](#)

QUESTION NO: 138

What are two characteristics of a small office / home office connection environment? (Choose two.)

- A. It requires 10Gb ports on all uplinks.
- B. It supports between 1 and 50 users.
- C. It supports between 50 and 100 users.
- D. A router port connects to a broadband connection.
- E. It requires a core, distribution, and access layer architecture.

ANSWER: B D

Explanation:

A SOHO (small office/home office) environment is typically a small-scale network designed for a limited number of users and simple connectivity needs. A common CCNA characterization is that SOHO networks support a small user base (often cited as roughly 1–50 users), which aligns with option B. Another hallmark is that Internet access is usually provided via a

broadband service (cable/DSL/fiber), and the edge device (often an integrated router/wireless router) connects to that broadband link—matching option D.

Option A is incorrect because SOHO networks do not inherently require 10Gb uplinks; they commonly use 1Gb Ethernet and may even use lower speeds depending on the ISP and equipment. Option C is less accurate for SOHO because 50–100 users is more typical of a small/medium business rather than a home office or very small office. Option E is incorrect because a three-tier core/distribution/access design is an enterprise campus design concept; SOHO networks are usually flat or have minimal hierarchy (often a single router/switch/AP device).

References: [Cisco — What is a SOHO network?](#), [Cisco — Campus network design overview \(hierarchical model\)](#)

QUESTION NO: 139

Refer to the exhibit.

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

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

ANSWER: B C

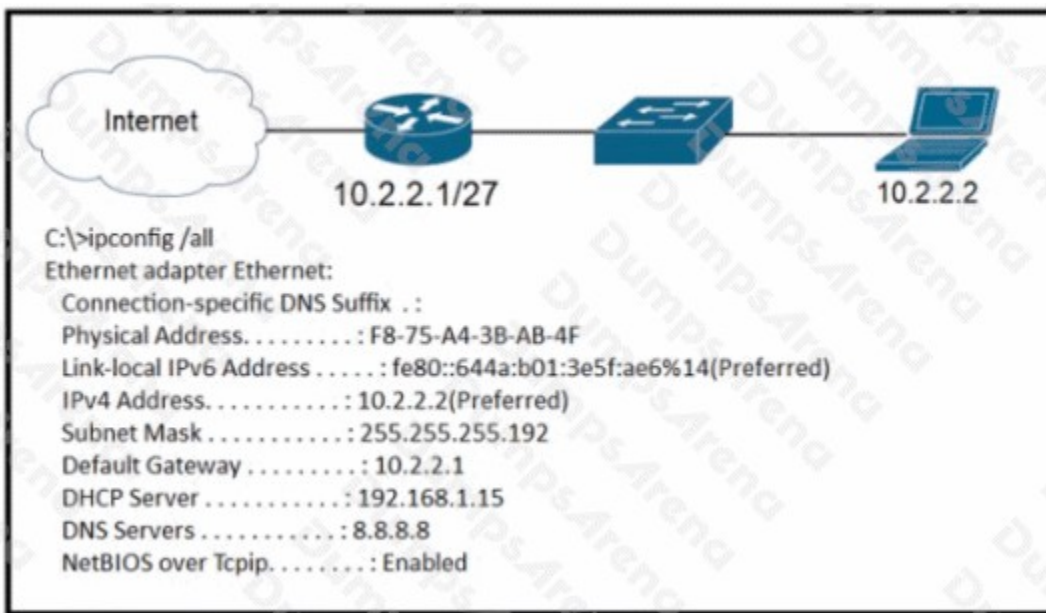
Explanation:

The exhibit's routing table entry represents a summarized network (a single route that covers a contiguous range of IP addresses). From the answer choices, the only pair that fits a single common subnet range is **192.168.1.61** and **192.168.1.64**. These two addresses fall within the same block when the route is a /26 (255.255.255.192) covering **192.168.1.64–192.168.1.127**, or when the entry is a broader summary that still includes both values. In contrast, **192.168.1.17** is in the lower portion of the /24 (it would be in 192.168.1.0/27, /28, /26, etc., depending on mask), and **192.168.1.254** is near the end of the /24; those would not be included if the routing entry is for the mid-range block shown in typical CCNA exhibits (e.g., 192.168.1.64/26). **192.168.1.127** is a boundary/broadcast candidate for a /26 and is commonly excluded as a usable host address, which is why it's not selected here as a "prefix included" choice in many routing-table questions.

For subnet range calculations and how prefixes map to address blocks, see Cisco's IP addressing/subnetting overview and route lookup behavior: [Cisco IP Subnetting Reference](#) and [Cisco Route Lookup \(Longest Match\)](#).

QUESTION NO: 140

Refer to the exhibit.



A newly configured PC fails to connect to the internet using TCP port 80 to www.cisco.com Which setting must be modified for the connection to work?

- A. Subnet Mask
- B. DNS Servers
- C. Default Gateway
- D. DHCP Server

ANSWER: B

Explanation:

The setting that must be modified is the **DNS server** configuration. The PC is trying to reach *www.cisco.com* (a hostname) over TCP port 80 (HTTP). Before the PC can open a TCP connection to that site, it must resolve the hostname to an IP address using DNS. If the DNS server address is missing or incorrect, name resolution fails, and the browser/app cannot initiate the TCP session—even if the PC has a valid IP address, subnet mask, and default gateway.

The other options don't best match the symptom described. A wrong **subnet mask** can cause local routing issues, but the question specifically highlights failure to reach a URL by name, which strongly indicates DNS. A wrong **default gateway** would typically prevent reaching any off-subnet IP destinations (including public IPs), but you could still test connectivity by pinging an Internet IP directly; the question emphasizes the hostname. **DHCP server** is not a host setting you "modify" on the PC to fix browsing; DHCP is just one method to obtain addressing, and a statically configured PC can work fine without DHCP.

References: [Cisco DNS overview and troubleshooting](#), [Cisco explanation of default gateway behavior](#).

QUESTION NO: 141

```
{
  "aaaUser": {
    "attributes": {
      "pwd": "password1",
      "firstName": "Abraham",
      "lastName": "Lincoln",
      "phone": "5555551212",
      "email": "test@cisco.com"
    },
    "children": [
      {
        "aaaUserDomain": {
          "attributes": {
            "name": "ExampleCisco"
          },
          "children": [
            {
              "aaaUserRole": {
                "attributes": {
                  "name": "admin"
                }
              }
            }
          ]
        }
      }
    ]
  }
}
```

Refer to the exhibit. How many objects are present in the given JSON-encoded data?

- A. One
- B. Four
- C. Seven
- D. Nine

ANSWER: D

Explanation:

In JSON, an **object** is a collection of name/value pairs enclosed in curly braces { . . . }. Arrays are enclosed in square brackets [. . .], and they can contain objects, arrays, or primitive values. To answer this question, you count how many distinct { . . . } structures appear in the JSON shown in the exhibit (including nested objects), not how many keys, not how many array elements (unless those elements are objects), and not how many total values.

In the provided JSON-encoded data, there are **nine** separate JSON objects (curly-brace blocks). That includes the top-level object plus all nested objects inside it (for example, objects inside arrays and objects nested as values of other keys). Therefore, option **D** is correct.

Option A (One) would only be true if the JSON contained only a single top-level object with no nested objects, which is not the case. Options B (Four) and C (Seven) undercount because they miss some nested { . . . } blocks present in the structure.

References: [RFC 8259 \(The JavaScript Object Notation \(JSON\) Data Interchange Format\)](#), [MDN: JSON](#)

QUESTION NO: 142

What does physical access control regulate?

- A. access to networking equipment and facilities
- B. access to servers to prevent malicious activity
- C. access to specific networks based on business function
- D. access to computer networks and file systems

ANSWER: A

Explanation:

Physical access control regulates who can physically enter or touch protected spaces and assets—things like buildings, wiring closets, data centers, and the networking equipment inside them. This includes controls such as locks, badges, mantraps, guards, cameras, and biometric readers. In Cisco security fundamentals (and in general security practice), “physical” controls are distinct from “logical” controls: physical controls protect the environment and hardware from theft, tampering, and unauthorized presence, while logical controls govern access to systems, networks, and data using authentication/authorization mechanisms.

Option A is correct because it directly describes controlling access to facilities and networking equipment, which is the core scope of physical access control. Option B is too narrow and mixes in intent (“prevent malicious activity”); while servers are physical assets, the phrasing implies a logical/operational security goal rather than defining what physical access control regulates. Option C describes role-based or business-function-based network access, which is logical access control (authorization). Option D also describes logical access control to networks and file systems (AAA, permissions, ACLs), not physical entry.

References: [Cisco – Physical Security \(overview\)](#), [Physical security \(concept and controls\)](#)

QUESTION NO: 143

What are two facts that differentiate optical-fiber cabling from copper cabling? (Choose two.)

- A. It is less expensive when purchasing patch cables.
- B. It carries electrical current further distances for PoE devices.
- C. It provides greater throughput options.
- D. It has a greater sensitivity to changes in temperature and moisture.

E. It carries signals for longer distances.

ANSWER: C E

Explanation:

Optical fiber differs from copper primarily in performance characteristics tied to how it transmits data: fiber uses light, while copper uses electrical signals. Because light experiences far less attenuation and is immune to electromagnetic interference (EMI/RFI), fiber links can typically run much longer distances than copper at the same data rates. That makes option E correct: fiber is the go-to medium for longer uplinks (for example, building-to-building or long campus runs) where copper Ethernet would exceed its distance limits.

Fiber also supports very high bandwidth/throughput options. Modern Ethernet standards scale to very high speeds over fiber (commonly 10G/40G/100G and beyond), and fiber is frequently chosen when higher throughput and future scalability are required. That makes option C correct.

The incorrect options fail for clear reasons: fiber patch cables and optics are generally more expensive than copper patch cords (A is false). Fiber does not carry electrical current and therefore cannot deliver PoE at all (B is false). Finally, while fiber can be physically fragile and has bend-radius considerations, it is not generally characterized as being more sensitive than copper to temperature/moisture in the way stated (D is false); fiber's key differentiators are distance, bandwidth, and EMI immunity.

References: [Cisco: Fiber-Optic Cabling Overview/Guidance](#), [Cisco: Ethernet \(media and speeds overview\)](#)

QUESTION NO: 144

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 primarily to improve scalability and stability. By splitting the network into areas, OSPF limits how far link-state information must be flooded; routers inside an area only maintain detailed LSDB information for that area. This directly **reduces routing overhead** (CPU, memory, and LSA flooding), making option B correct. Hierarchy also **confines instability**—for example, frequent link flaps—mostly within the local area so that SPF recalculations and LSA churn don't ripple across the entire domain, which makes option E correct. Finally, smaller LSDBs and fewer SPF runs generally **speed up convergence** from an overall network perspective because fewer routers are impacted by a given change and SPF calculations are performed on a smaller topology graph (option A).

Option C is unrelated to OSPF design; replacing routers with switches is an architecture/cost decision, not a reason for OSPF hierarchy. Option D is incorrect because OSPF hierarchy doesn't increase bandwidth; it reduces control-plane traffic. Option F is not a primary driver; multi-area OSPF often adds design and configuration considerations (ABRs, area types), even though it improves operational scalability.

References: [Cisco OSPF Design Guide \(support doc\)](#), [Cisco IOS XE OSPF Configuration Guide](#).

QUESTION NO: 145

Which mechanism allows WPA3 to provide a higher degree of security than its predecessors?

- A. automatic device pairing
- B. SAE password-based key exchange
- C. certificate-based authentication
- D. special-character support in preshared keys

ANSWER: B

Explanation:

WPA3 improves security primarily by replacing WPA2-Personal's PSK/4-way handshake reliance with SAE (Simultaneous Authentication of Equals), also known as Dragonfly. SAE is a password-authenticated key exchange (PAKE) that provides stronger protection against offline dictionary attacks: an attacker who captures the handshake can't simply test password guesses offline at high speed. Instead, each guess requires active interaction with the network, making large-scale guessing far less practical. SAE also provides forward secrecy, meaning that even if the password is later compromised, previously captured traffic is not automatically decryptable.

Option A is incorrect because "automatic device pairing" is not a core WPA3 security mechanism (that's more aligned with Wi-Fi Easy Connect/DPP, which is separate from WPA3 itself). Option C is incorrect because certificate-based authentication is characteristic of WPA2/WPA3-Enterprise (802.1X/EAP), not the key differentiator that makes WPA3 broadly "more secure than predecessors" in the Personal mode context. Option D is incorrect because special characters have long been allowed in passphrases; WPA3's improvement is not about character sets but about the authentication/key exchange method.

References: [Wi-Fi Alliance – Wi-Fi security \(WPA3 overview\)](#), [Simultaneous Authentication of Equals \(SAE\)](#)

QUESTION NO: 146

You have two paths for the 10.10.10.0 network - one that has a feasible distance of 3072 and the other of 6144.

What do you need to do to load balance your EIGRP routes?

- A. Change the maximum paths to 2
- B. Change the configuration so they both have the same feasible distance
- C. Change the variance for the path that has a feasible distance of 3072 to 2
- D. Change the IP addresses so both paths have the same source IP address

ANSWER: A C

Explanation:

EIGRP installs multiple routes to the same destination for load balancing. By default it performs equal-cost load balancing only, meaning the paths must have the same metric (same feasible distance) to both be installed. In your case, the feasible distances are 3072 and 6144, so they are not equal and the higher-metric path will not be used unless you enable unequal-cost load balancing.

Unequal-cost load balancing is enabled with the *variance* command. The variance value is a multiplier applied to the best (successor) feasible distance. Any feasible successor whose FD is less than or equal to (best FD × variance) can be installed for load balancing. Here, best FD is 3072. With variance 2, the threshold is 6144, so the second path (FD 6144) becomes eligible and can be installed alongside the best path.

Changing *maximum-paths* to 2 only helps if there are already multiple eligible paths (equal-cost successors, or unequal-cost paths made eligible via variance). Changing IP addresses is irrelevant. Making both paths have the same feasible distance could allow equal-cost load balancing, but it's not required and often not practical; variance is the intended mechanism.

References: [Cisco Support: EIGRP Load Balancing and Variance](#), [Cisco IOS XE EIGRP Load Balancing \(variance\)](#)

QUESTION NO: 147

What are two fundamentals of virtualization? (Choose two.)

- A. It allows logical network devices to move traffic between virtual machines and the rest of the physical network.
- B. It allows multiple operating systems and applications to run independently on one physical server.
- C. It allows a physical router to directly connect NICs from each virtual machine into the network.
- D. It requires that some servers, virtual machines, and network gear reside on the Internet.
- E. The environment must be configured with one hypervisor that serves solely as a network manager to monitor SNMP traffic.

ANSWER: A B

Explanation:

Two core fundamentals of virtualization are (1) abstracting physical resources so multiple isolated workloads can share the same hardware, and (2) providing virtualized networking so VMs can communicate with each other and with the external physical network. Option B is correct because server virtualization uses a hypervisor to let multiple guest operating systems and their applications run independently on a single physical host, sharing CPU, memory, storage, and NICs while remaining isolated. Option A is also correct because virtualization commonly includes virtual switching (vSwitch/virtual NICs) that forwards traffic between VMs and uplinks to the physical network, enabling connectivity beyond the host.

Option C is incorrect because VMs do not "directly connect NICs" to a physical router; instead, VMs use virtual NICs connected to a virtual switch/port group, and the host's physical NICs provide uplink connectivity. Option D is incorrect because virtualization does not require Internet residency; it can be entirely on-premises. Option E is incorrect because hypervisors are not required to be dedicated SNMP network managers; monitoring can be done many ways, and virtualization does not mandate a single hypervisor for that purpose.

References: [Cisco: What is a virtual switch?](#), [VMware glossary: Hypervisor](#)

QUESTION NO: 148

What are two port types used by a Cisco WLC for out-of-band management? (Choose two.)

- A. service
- B. console
- C. management
- D. distribution system
- E. redundant

ANSWER: A B

Explanation:

On Cisco Wireless LAN Controllers (WLCs), “out-of-band management” refers to managing the controller through interfaces that do not rely on the in-band data network. Two common out-of-band methods are the **console** port (local serial access for initial setup, recovery, and troubleshooting) and the dedicated **service** port (an Ethernet port intended for isolated management access, separate from the distribution/system ports used for AP and client traffic). These two ports let you reach the WLC even when the production network is down or misconfigured.

The **management** interface (option C) is typically an in-band logical interface carried over the controller’s network ports and VLANs; it’s used for GUI/SSH/SNMP, but it’s not inherently out-of-band because it depends on the network path. The **distribution system** ports (option D) are the data/uplink ports that connect to the wired network for AP/client traffic, so they are in-band. The **redundant** port (option E) is used for high availability/redundancy (e.g., SSO/HA links on certain models), not for routine out-of-band management.

References: [Cisco WLC Configuration Guides](#), [Cisco WLC Interface/Port Overview \(Cisco Support\)](#).

QUESTION NO: 149

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.)

```
Port Security : Enabled
Port Status : Secure-up
Violation Mode : Protect
Aging Time : 0 mins
Aging Type : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses : 4
Total MAC Addresses : 3
Configured MAC Addresses: 1
Sticky MAC Addresses : 2
Last Source Address:Vlan : 0001:0fAA.33BB:1
Security Vioalntion Count : 0
```

- 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:

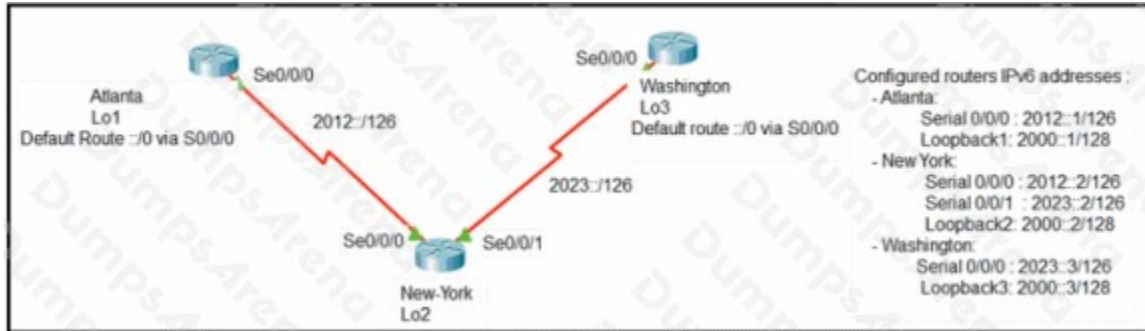
This question is describing a switchport with port security enabled that has already learned its maximum number of secure MAC addresses. When a frame arrives with a source MAC address that is not in the secure MAC table, a port-security violation occurs. The exact behavior depends on the configured violation mode (protect, restrict, or shutdown). In the common CCNA context, when the port is configured for *restrict*, the switch drops frames from unknown source MAC addresses and generates a syslog message (and also increments the violation counter). That matches two of the listed events: traffic from the unknown MAC is dropped, and a syslog message is generated.

Option A is incorrect because in restrict mode the violation counter *does* increment (protect is the mode that drops silently without incrementing or logging). Option C (error-disabled) is only true for the shutdown violation mode, not for restrict/protect. Option B is not a defined port-security behavior; LEDs may change depending on platform, but it's not the key event Cisco tests for.

References: [Cisco Port Security Configuration/Behavior \(violation modes\)](#), [Cisco Catalyst Port Security Configuration Guide](#)

QUESTION NO: 150

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:

Because the requirement is end-to-end reachability between Atlanta Lo1 (2000::1/128) and Washington Lo3 (2000::3/128), the New York router must have host (/128) routes for both loopback addresses pointing toward the correct next-hop routers on each side. In a typical three-router chain (Atlanta—New York—Washington), New York forwards traffic destined to Atlanta's loopback toward Atlanta's adjacent interface address, and traffic destined to Washington's loopback toward Washington's adjacent interface address. Option A installs a host route to 2000::1/128 via next hop 2012::1 (the Atlanta-side next hop). Option B installs a host route to 2000::3/128 via next hop 2023::3 (the Washington-side next hop). With these two routes, New York can forward packets in both directions, enabling Atlanta to reach Washington's loopback (and return traffic to reach Atlanta's loopback).

Option C and E use only an exit interface. For IPv6, Cisco best practice is to specify a next-hop address (or both interface and next hop) to avoid ambiguity and recursive resolution issues on multi-access links; interface-only static routes are generally discouraged unless the link type makes it unambiguous. Option D points 2000::1/128 to 2012::2, which would be New York's own address (wrong next hop), so it would not correctly forward toward Atlanta.

References: [Cisco IP Routing - Static Routes \(concepts apply to IPv4/IPv6\)](#), [Cisco IOS IPv6 Routing Configuration Guide \(ipv6 route\)](#)

QUESTION NO: 151

What is the default port-security behavior on a trunk link?

- A. It causes a network loop when a violation occurs.

- B. It disables the native VLAN configuration as soon as port security is enabled.
- C. It places the port in the err-disabled state if it learns more than one MAC address.
- D. It places the port in the err-disabled state after 10 MAC addresses are statically configured.
- E. Port security is not enabled/operational on a trunk link by default (it is typically only supported on access ports), so there is no default trunk port-security violation behavior.

ANSWER: E

Explanation:

On Cisco switches, port security is primarily intended for access ports. By default, you cannot enable port security on a trunk interface; IOS will reject the configuration unless the port is an access port (or you change the port mode away from trunk). Because of that, there is no “default violation behavior on a trunk link” in the sense of learning MAC addresses and err-disabling—port security simply isn’t operational on a trunk by default.

Option C describes the default violation action (shutdown/err-disabled) and a common default maximum (1 MAC) for an access port with port security enabled, but it’s not the default behavior on a trunk link because port security isn’t enabled/allowed there by default. Option A is unrelated—port security violations do not “cause a network loop.” Option B is incorrect because port security has nothing to do with disabling the native VLAN. Option D is also incorrect: the default maximum is not “10 statically configured MAC addresses,” and again the trunk context is wrong.

References: [Cisco Support: Port Security Configuration and Troubleshooting](#), [Cisco Catalyst 2960 Software Configuration Guide \(Port Security section\)](#)

QUESTION NO: 152

An engineer needs to add an old switch back into a network. To prevent the switch from corrupting the VLAN database, with action must be taken?

- A. Add the switch in the VTP domain with a lower revision number.
- B. Add the switch in the VTP domain with a higher revision number.
- C. Add the switch with DTP set to dynamic desirable.
- D. Add the switch with DTP set to desirable.

ANSWER: A

Explanation:

In VTP (VLAN Trunking Protocol), the key risk when introducing an “old” switch is its VTP configuration revision number. If a switch joins a VTP domain with a *higher* revision number than the current domain, it can advertise its VLAN database and overwrite (corrupt) the VLAN information on other switches in that domain. To prevent this, you must ensure the reintroduced switch has a *lower* (ideally reset to 0) revision number before connecting it to the production VTP domain/trunks. Common best practice is to change the VTP domain name (or set VTP to transparent/off), delete vlan.dat, and/or reload so the revision resets, then reconfigure appropriately.

Option A aligns with this concept: the switch must not have a higher revision than the domain, so it should be introduced with a lower revision number. Option B is exactly the dangerous scenario (higher revision overwrites the domain). Options C and D discuss DTP (Dynamic Trunking Protocol), which negotiates trunking but does not protect the VLAN database from VTP revision issues; DTP settings are unrelated to preventing VTP-based VLAN database corruption.

References: [Cisco Support: Understanding and Configuring VTP](#), [Cisco Support: VTP Configuration Revision Number](#)

QUESTION NO: 153

What is a benefit of VRRP?

- A. It provides the default gateway redundancy on a LAN using two or more routers.
- B. It provides traffic load balancing to destinations that are more than two hops from the source.
- C. It prevents loops in a Layer 2 LAN by forwarding all traffic to a root bridge, which then makes the final forwarding decision.
- D. It allows neighbors to share routing table information between each other.

ANSWER: A

Explanation:

VRRP (Virtual Router Redundancy Protocol) provides default gateway redundancy for hosts on a LAN. Multiple routers participate in a VRRP group and present a single virtual default gateway IP/MAC to end hosts. One router is elected the Master and actively forwards traffic; the others are in Backup state and take over if the Master fails. This improves availability because end devices keep using the same default gateway address without needing reconfiguration or waiting for manual intervention.

Option A is correct because it describes exactly this first-hop (default gateway) redundancy function. Option B is incorrect because VRRP is not a traffic engineering or multi-hop load-balancing mechanism; it's a first-hop redundancy protocol (FHRP). While some FHRPs can be designed for limited gateway load sharing, VRRP's primary benefit is redundancy, not balancing traffic to remote destinations. Option C describes Spanning Tree Protocol behavior (Layer 2 loop prevention), not VRRP. Option D describes dynamic routing protocols (like OSPF/EIGRP) exchanging routing information, which is unrelated to VRRP's purpose.

References: [Cisco VRRP Overview](#), [RFC 5798 \(VRRPv3\)](#)

QUESTION NO: 154

Refer to the exhibit.

```
Switch2# show lldp
Global LLDP Information
  Status: ACTIVE
  LLDP advertisements are sent every 30 seconds
  LLDP hold time advertised is 120 seconds
  LLDP interface reinitialization delay is 2 seconds
```

A network engineer must update the configuration on switch2 so that it sends LLDP packets.

A)

```
Switch2(config)#lldp timer 1
Switch2(config)#lldp tlv-select 3
```

B)

```
Switch2(config)#lldp timer 1
Switch2(config)#lldp holdtime 3
```

C)

```
Switch2(config)#lldp timer 60
Switch2(config)#lldp holdtime 180
```

D)

```
Switch2(config)#lldp timer 60  
Switch2(config)#lldp tlv-select 180
```

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

ANSWER: C

Explanation:

To make switch2 *send* LLDP packets, LLDP must be enabled globally and (if needed) enabled for transmission on the relevant interface(s). On Cisco IOS, LLDP is controlled with global commands `lldp run` (enables LLDP process) and interface-level commands `lldp transmit/lldp receive` (control per-port behavior). The correct configuration is the one that enables LLDP operation so the switch actually advertises itself to neighbors. That aligns with option C in the exhibit.

The other options are incorrect because they either (1) enable only reception and not transmission, (2) disable LLDP globally (for example using `no lldp run`), or (3) configure CDP instead of LLDP (CDP uses `cdp run / cdp enable` and is a different discovery protocol). Since the requirement is specifically to send LLDP packets, the configuration must include LLDP being enabled and transmitting.

References: [Cisco IOS XE CDP/LLDP Configuration Guide](#), [Cisco LLDP Configuration Example](#).

QUESTION NO: 155

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

```
D 10.214.247.237.28 [90/2170] via G0/12  
D 10.214.247.237.28 [90/46985] via G0/19  
O 10.214.247.237.28 [110/665] via G0/9  
O 10.214.247.237.28 [110/3399] via G0/1
```

- A. G0/9
- B. G0/12
- C. G0/19
- D. G0/1

ANSWER: B

Explanation:

To determine which interface is used to send traffic to a destination network, you look up the destination in the routing table and then use the route's "exit interface" (or the interface associated with the next-hop). In the exhibit, the best match for the destination network is the longest-prefix match, and that selected route points out interface **G0/12**. That means packets destined for that network will be forwarded out G0/12 (either directly connected on that interface or toward a next-hop reachable via G0/12).

Option A (G0/9), option C (G0/19), and option D (G0/1) are incorrect because they are not the exit interface referenced by the routing table entry that matches the destination network with the most specific prefix. Even if those interfaces appear in

the routing table for other networks (or as connected interfaces), they won't be used unless their route is the best match for the destination.

This is standard Cisco forwarding behavior: routers select the most specific route (longest prefix) and forward out the associated outgoing interface/next hop. References: [Cisco – How to read the routing table \(route selection concepts\)](#), [Longest prefix match \(concept overview\)](#).

QUESTION NO: 156

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

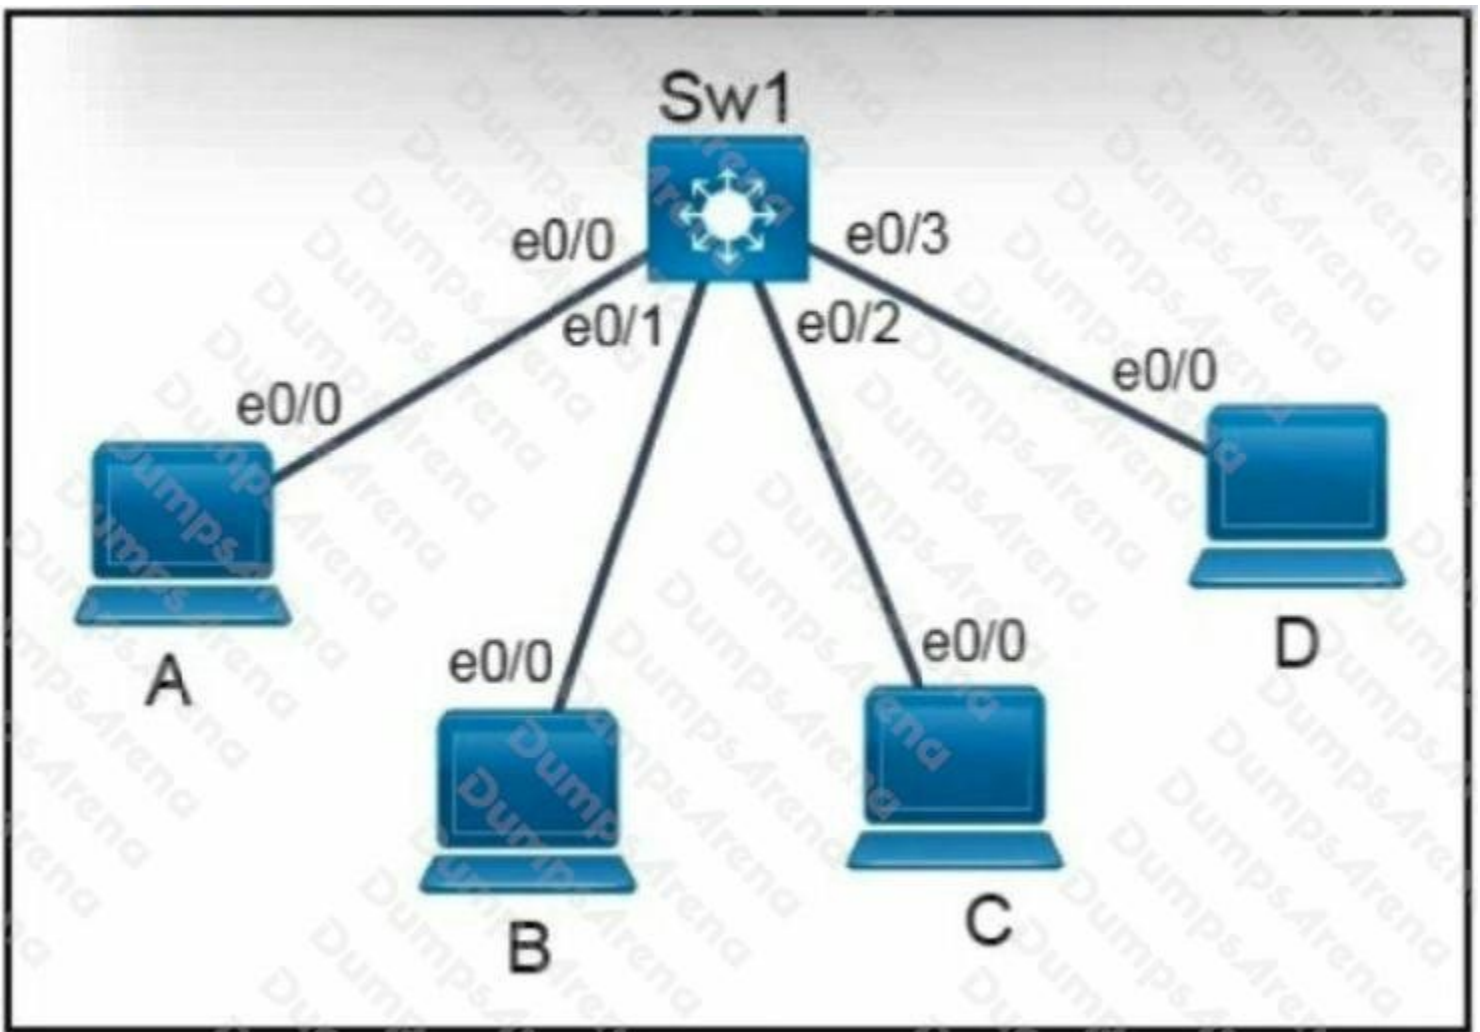
Explanation:

Virtual machines are software-defined compute instances that are abstracted from the physical server by a hypervisor. A key characteristic is consolidation: multiple VMs can run simultaneously on the same physical hardware while remaining logically separated. That makes option A correct. Another defining trait is isolation/independence: each VM has its own virtual hardware (vCPU, vNIC, vDisk) and runs its own OS and applications without being affected by other VMs on the same host (aside from shared-resource contention). That makes option E correct.

Option B is misleading as written: a guest OS does not “depend” on the hypervisor in the sense of being tied to a specific one; it depends on virtual hardware presented to it, and VMs can often be migrated between compatible hypervisors/hosts. Option C is incorrect because VMs are not automatically interconnected; connectivity requires virtual switching/port groups and appropriate network configuration. Option D is incorrect because hypervisors do not guarantee equal sharing of resources—CPU and memory allocation can be weighted, limited, reserved, and overcommitted depending on policy and configuration.

References: [Cisco: What is virtualization?](#), [VMware glossary: Virtual Machine](#)

QUESTION NO: 157



Refer to the exhibit. Host A switch interface is configured in VLAN 2. Host D sends a unicast packet destined for the IP address of host A.

```
Sw1#show mac-address table
Mac Address Table
```

Vlan	Mac Address	Type	Ports
2	000c.859c.bb7b	DYNAMIC	e0/1
3	000c.859c.bb7b	DYNAMIC	e0/1
2	0010.11dc.3e91	DYNAMIC	e0/2
3	0010.11dc.3e91	DYNAMIC	e0/2
2	0043.29d9.c045	DYNAMIC	e0/3

```
Sw1#
```

What does the switch do when it receives the frame from host D?

- A. It floods the frame out of every ports except the source port.
- B. It creates a broadcast storm.
- C. It shuts down the source port and places it in err-disable mode.
- D. It drops the frame from the MAC table of the switch.

ANSWER: A

Explanation:

When the switch receives Host D's Ethernet frame, it makes a forwarding decision based on the *destination MAC address*, not the destination IP address. To send a unicast IP packet to Host A, Host D must encapsulate it in an Ethernet frame addressed to Host A's MAC (learned via ARP) or, if ARP resolution hasn't completed yet, it will send an ARP broadcast first. In the scenario implied by the question, the switch receives a unicast frame but does not have an entry for the destination MAC in its CAM/MAC address table (an "unknown unicast"). In that case, the switch floods the frame out all ports in the same VLAN as the incoming frame, except the port it arrived on. Flooding is limited to the VLAN/broadcast domain; it is not sent to ports in other VLANs.

Option A is the best match to this behavior (though more precisely it floods within the VLAN). Option B is incorrect because a broadcast storm requires a Layer 2 loop and uncontrolled replication, not a single unknown-unicast event. Option C is incorrect because err-disable is triggered by specific features (e.g., port-security violations), not normal forwarding. Option D is incorrect because switches don't "drop the frame from the MAC table"; they either forward, flood, or drop the received frame.

References: [Cisco: Ethernet Switching \(MAC learning and forwarding\)](#), [Cisco Catalyst: Switching and MAC address table behavior](#)

QUESTION NO: 158

Router#					
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge					
S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,					
D - Remote, C - CVTA, M - Two-port Mac Relay					
Device ID	Local Interface	Holdtime	Capability	Platform	Port ID
10.1.1.2	Gig 37/3	176	R I	CPT 600	Gig 36/41
10.1.1.2	Gig 37/1	174	R I	CPT 600	Gig 36/43
10.1.1.2	Gig 36/41	134	R I	CPT 600	Gig 37/3
10.1.1.2	Gig 36/43	134	R I	CPT 600	Gig 37/1
10.1.1.2	Ten 3/2	132	R I	CPT 600	Ten 4/2
10.1.1.2	Ten 4/2	174	R I	CPT 600	Ten 3/2

Refer to the exhibit. Which command provides this output?

- A. show ip route
- B. show cdp neighbor
- C. show ip interface
- D. show interface

ANSWER: B

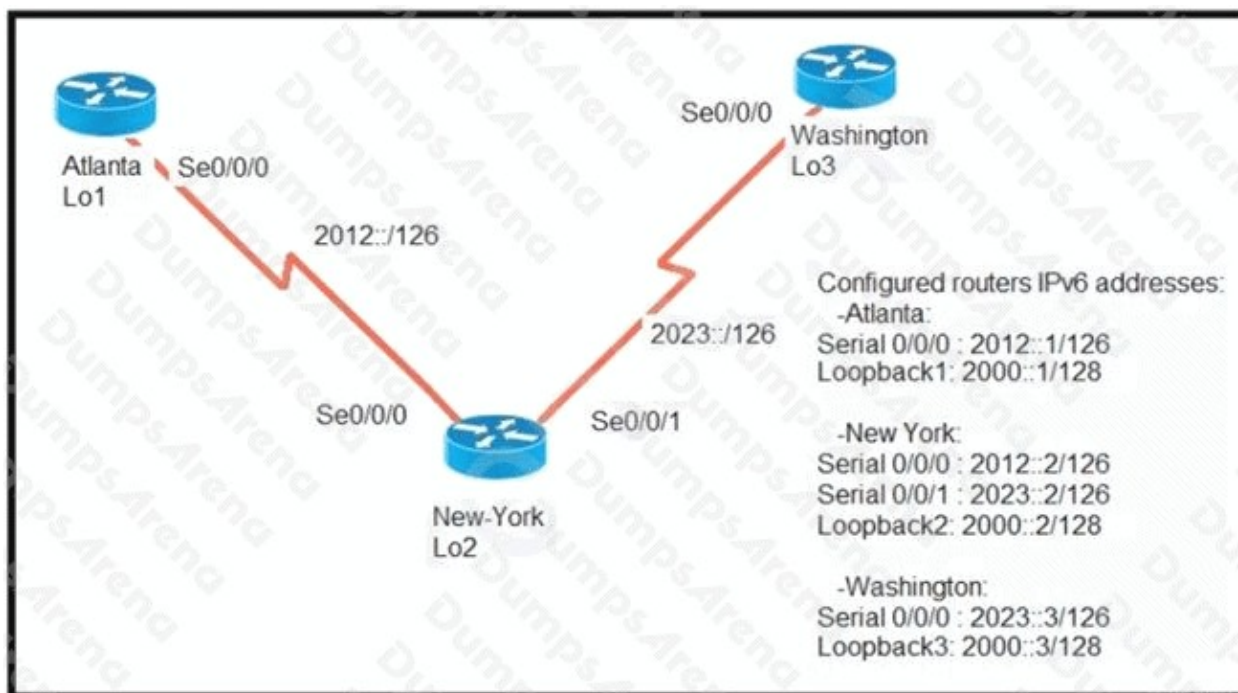
Explanation:

The output shown in the exhibit matches the typical format of Cisco Discovery Protocol (CDP) neighbor information: it lists directly connected Cisco devices and includes fields such as the neighbor *Device ID*, the local interface, the *Holdtime*, the neighbor's *Capability* (for example, R for router, S for switch), the *Platform*, and the neighbor's *Port ID*. That exact style of table is produced by the **show cdp neighbors** command (note: IOS accepts both singular/plural in many cases, but the canonical command is plural).

The other options don't fit the exhibit's structure. **show ip route** displays the routing table (route codes, prefixes, next hops), not neighbor device IDs and platforms. **show ip interface** (and related variants like *show ip interface brief*) focuses on interface IP addressing and interface status, not CDP-discovered neighbors. **show interfaces** provides detailed per-interface statistics (errors, duplex/speed, counters), again not a neighbor summary table.

References: [Cisco CDP Overview and show command examples](#), [Cisco IOS CDP Command Reference \(show cdp neighbors\)](#).

QUESTION NO: 159



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

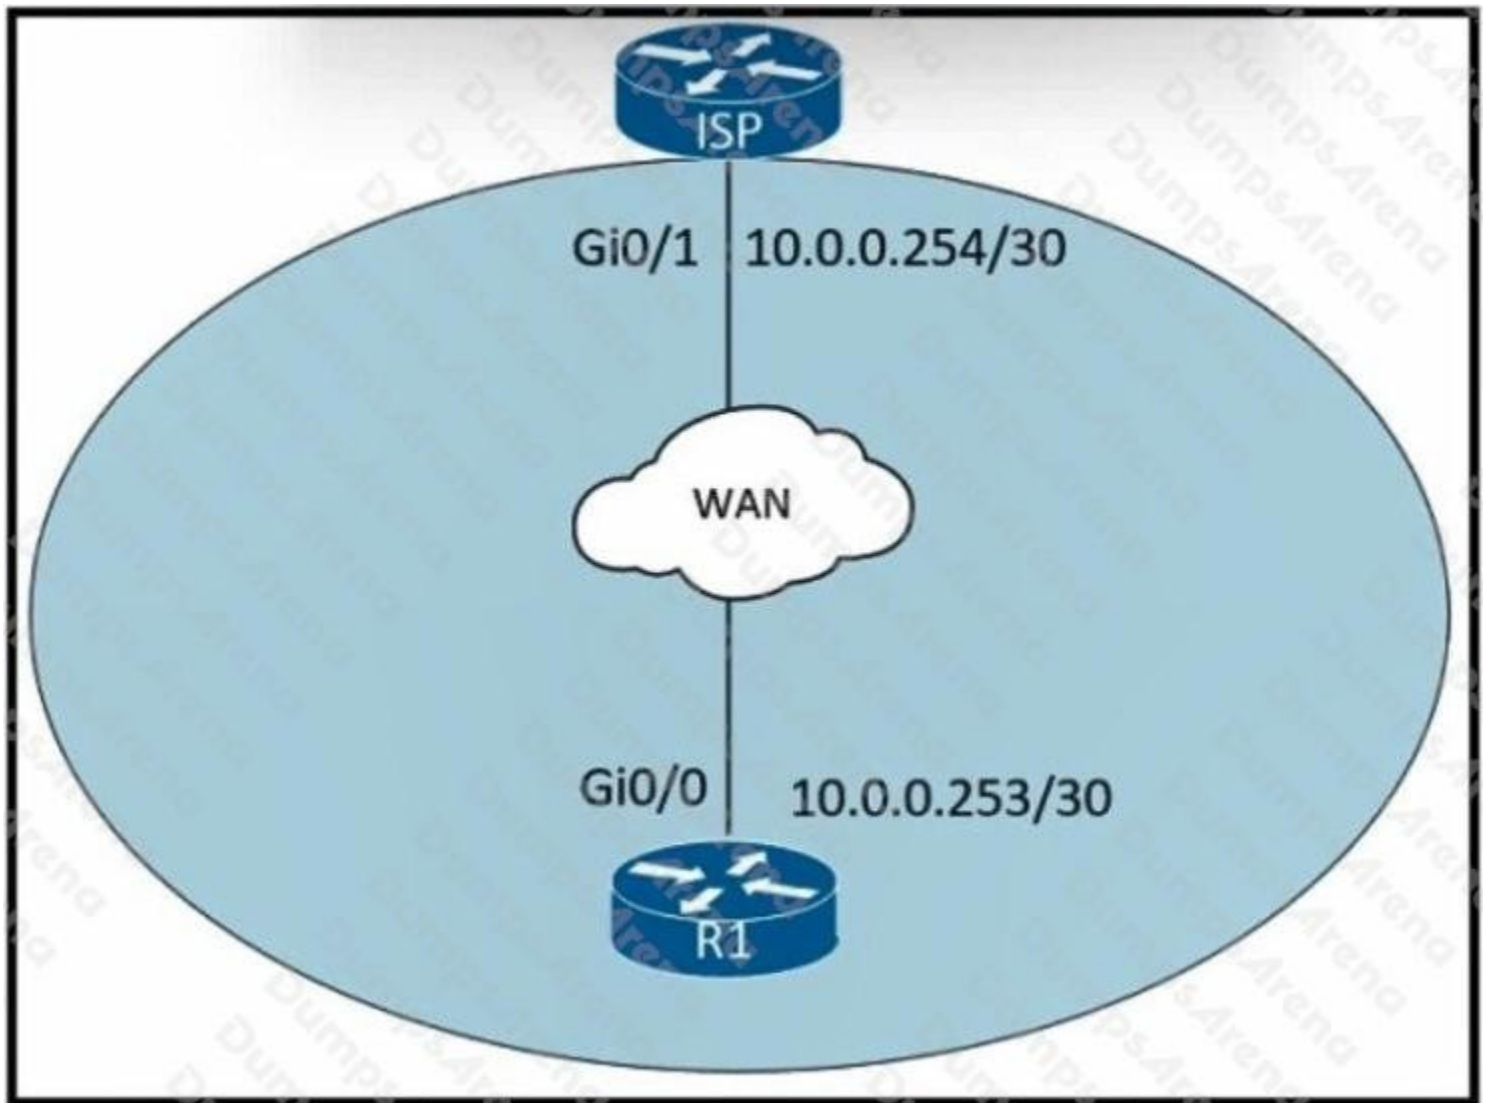
Explanation:

Atlanta and Washington are not directly connected; they must route through New York. New York already has static routes to both remote sites, but for end-to-end reachability each “edge” router also needs a route to the other edge router’s WAN prefix via New York as the next hop. In the diagram, the Atlanta–New York serial link uses the 2012::/126 network and the New York address on that link is 2012::2. The Washington–New York serial link uses the 2023::/126 network and the New York address on that link is 2023::1. Therefore, Atlanta must add a static route to the Washington link (2023::/126) pointing to New York’s next-hop address on the Atlanta link (2012::2). Likewise, Washington must add a static route to the Atlanta link (2012::/126) pointing to New York’s next-hop address on the Washington link (2023::1).

Option D matches the required route on Atlanta. Option C matches the required route on Washington. Option A uses the wrong next-hop (2012::1 is Atlanta itself). Option B uses the wrong next-hop (2023::2 is Washington itself). Option E is not appropriate as written because it specifies an outgoing interface that is not the correct path toward 2012::/126 from Atlanta (and IPv6 static routes typically use a next-hop, or an exit interface only when it makes sense on that router).

References: [Cisco IPv6 Static Routing Configuration](#), [Cisco IOS XE IPv6 Unicast Routing \(Static Routes\)](#)

QUESTION NO: 160



Refer to the exhibit. An administrator must turn off the Cisco Discovery Protocol on the port configured with address last usable address in the 10.0.0.0/30 subnet. Which command set meets the requirement?

- A. interface gi0/1 no cdp enable
- B. interface gi0/0 no cdp run
- C. interface gi0/0 no cdp advertise-v2

D. interface gi0/1 clear cdp table

ANSWER: A

Explanation:

A /30 network (10.0.0.0/30) provides 4 addresses: network 10.0.0.0, usable hosts 10.0.0.1 and 10.0.0.2, and broadcast 10.0.0.3. The “last usable” address is therefore 10.0.0.2. From the exhibit, the interface configured with 10.0.0.2/30 is Gi0/1, so CDP must be disabled on that specific interface.

To disable CDP on a single interface, the correct interface-level command is **no cdp enable** under the interface. That is exactly what option A does on Gi0/1.

Option B uses **no cdp run**, which is a global configuration command (it disables CDP for the entire device) and it’s incorrectly shown under an interface; it also targets Gi0/0, not the interface with 10.0.0.2. Option C (**no cdp advertise-v2**) only affects CDP version 2 advertisements and does not fully disable CDP. Option D clears the CDP neighbor table but does not disable CDP.

References: [Cisco CDP overview and configuration](#), [Cisco IOS CDP Configuration Guide](#).

QUESTION NO: 161

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:

On a network, a server’s core role is to provide services/resources to multiple clients and respond to their requests. Option A is correct because servers are designed to accept and process concurrent requests from many workstations (clients) at the same time (for example, multiple users accessing files, web pages, or authentication services simultaneously). Option E is also correct because servers commonly run network applications/services (such as web, file, DNS, DHCP, email, or database services) that receive client requests and then send back data or perform actions on the clients’ behalf.

Option B is incorrect because redundancy is not achieved “exclusively” through virtual server clustering; high availability can be implemented in many ways (clustering, load balancing, redundant hardware, multiple servers, etc.), and clustering is not a mandatory or exclusive server function. Option C is incorrect because servers can be located in many environments (on-prem, shared data centers, cloud) and are not “solely” housed in a dedicated single-client data center. Option D is incorrect because servers do not need to run the same operating system to communicate; interoperability is provided by standard network protocols (TCP/IP, DNS, HTTP, SMB, etc.), not identical OSs.

References: [Cisco – What is a server?](#), [Cisco – Networking fundamentals \(TCP/IP-based communication\)](#)

QUESTION NO: 162

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 neighbors must agree on several key parameters and be able to exchange OSPF control packets. A mismatch in **Hello and Dead timers** will prevent neighbors from reaching the 2-Way/Full state because each router expects Hellos at a different interval and will declare the other down prematurely; timer matching is a core OSPF neighbor requirement (along with area ID, authentication, stub flags, etc.). Also, **duplicate router IDs** can prevent stable adjacency: OSPF requires unique router IDs within the OSPF domain, and duplicates commonly lead to neighbor resets and failure to maintain adjacency.

Option A is wrong because OSPF is an IGP and does not use autonomous system numbers for neighbor formation (that concept is associated with EIGRP/BGP). Option B is wrong because OSPFv2 uses multicast **224.0.0.5** (AllSPFRouters) and **224.0.0.6** (AllDRouters), not 224.0.0.10 (that address is used by EIGRP). Option C is wrong because the OSPF process ID is locally significant on a router; neighbors do not need to match process IDs to form adjacency.

References: [Cisco OSPF Neighbor Relationship \(support doc\)](#), [Cisco OSPF Design Guide / OSPF operation details](#).

QUESTION NO: 163

Which protocol does Ansible use to push modules to nodes in a network?

- A. Telnet
- B. Kerberos
- C. SNMP
- D. SSH

ANSWER: D

Explanation:

Ansible is an agentless automation tool. For most network and Linux/Unix node management, it connects from the control node to the managed node and “pushes” the required module code over the existing remote management transport. The primary protocol Ansible uses for this push model is **SSH** (Secure Shell). Over SSH, Ansible can authenticate (password or keys), copy temporary module code to the remote system, execute it, and then remove it—without requiring a persistent agent on the managed node.

Option D is therefore correct because SSH is the standard transport Ansible uses for remote execution on network devices (via `network_cli`) and on servers (via the `ssh` connection plugin). Telnet (A) is insecure and not the default/expected Ansible transport. Kerberos (B) is an authentication mechanism that can be used in some environments, but it is not the protocol Ansible uses to push modules; it would typically still ride over SSH or other transports. SNMP (C) is primarily for monitoring/management data retrieval and traps, not for pushing and executing automation modules.

References: [Ansible Network Platform Guide \(connection methods\)](#), [Ansible SSH connection plugin documentation](#).

QUESTION NO: 164

Which two transport layer protocols carry syslog messages? (Choose two.)

- A. IP
- B. RTP
- C. TCP
- D. UDP
- E. ARP

ANSWER: C D

Explanation:

Syslog messages can be transported using both **UDP** and **TCP**, which is why options D and C are correct. Traditionally, syslog uses UDP port 514 for message delivery because it's lightweight and has low overhead—useful for high-volume logging where occasional loss may be acceptable. However, UDP provides no delivery guarantee, ordering, or congestion control.

To address reliability and security needs, syslog is also commonly carried over TCP. TCP provides session-based delivery with acknowledgments and retransmissions, making it a better fit when you want higher assurance that logs arrive (for example, compliance or forensic logging). In modern deployments, TCP is also frequently paired with TLS (often referred to as “syslog over TLS”) to protect log confidentiality and integrity in transit.

The other options are not transport-layer protocols used for syslog: **IP** is a network-layer protocol, **RTP** is an application-layer protocol typically used for real-time media streams, and **ARP** is used for Layer 2 address resolution and is not a transport protocol.

References: [RFC 5424 \(The Syslog Protocol\)](#), [RFC 6587 \(Syslog over TCP\)](#)

QUESTION NO: 165

Which two outcomes are predictable behaviors for HSRP? (Choose two.)

- A. The two routers negotiate one router as the active router and the other as the standby router.
- B. The two routers share the same interface IP address, and default gateway traffic is load-balanced between them.
- C. The two routers synchronize configurations to provide consistent packet forwarding.
- D. Each router has a different IP address, both routers act as the default gateway on the LAN, and traffic is load-balanced between them.
- E. The two routers share a virtual IP address that is used as the default gateway for devices on the LAN.

ANSWER: A E

Explanation:

HSRP (Hot Standby Router Protocol) provides first-hop redundancy by presenting hosts with a single, consistent default gateway while using two (or more) routers behind the scenes. A predictable behavior is that HSRP elects roles: one router becomes the *active* router that forwards traffic sent to the virtual gateway, and another becomes the *standby* router ready to take over if the active fails (Option A). Another core behavior is the use of a *virtual IP address* (and virtual MAC) that end hosts configure as their default gateway; the active router answers ARP for that virtual IP and forwards packets (Option E).

Option B is wrong because HSRP does not “share the same interface IP address” between two physical routers, and HSRP by itself does not inherently load-balance default-gateway traffic (that’s more associated with GLBP, or with multiple HSRP groups configured intentionally). Option C is wrong because HSRP does not synchronize router configurations; it only exchanges hello messages and state information for gateway redundancy. Option D is wrong because hosts do not use two different default gateways simultaneously in basic HSRP, and load-balancing is not a default/predictable HSRP outcome.

QUESTION NO: 166

What is a characteristic of an SSID in wireless networks?

- A. allows easy file sharing between endpoints
- B. provides protection against spyware
- C. associates a name to a wireless network
- D. eliminates network piggybacking

ANSWER: C

Explanation:

An SSID (Service Set Identifier) is essentially the human-readable name that identifies a WLAN. Access points advertise (or can be configured not to advertise) the SSID in beacon frames, and wireless clients use the SSID to select and join the intended wireless network. In other words, the SSID's key characteristic is that it associates a name with a wireless network so users and devices can distinguish one WLAN from another (for example, separating "Guest" from "Corporate").

Option A is incorrect because file sharing is an application/host function and not something an SSID enables. Option B is incorrect because an SSID provides no malware/spyware protection; security is provided by mechanisms like WPA2/WPA3, 802.1X/EAP, and proper endpoint controls. Option D is incorrect because simply having an SSID does not prevent "piggybacking" (unauthorized use). Preventing unauthorized access requires authentication and encryption (again, WPA2/WPA3, strong passphrases, or enterprise authentication). Even hiding an SSID is not a real security control because the SSID can still be discovered via other management frames.

References: [Cisco: Understanding SSIDs](#), [Cisco Wireless LAN Basics \(SSID concept\)](#)

QUESTION NO: 167

What are two reasons a network administrator would use CDP? (Choose two.)

- A. to verify the type of cable interconnecting two devices
- B. to determine the status of network services on a remote device
- C. to obtain VLAN information from directly connected switches
- D. to verify Layer 2 connectivity between two devices when Layer 3 fails
- E. to obtain the IP address of a connected device in order to telnet to the device
- F. to determine the status of the routing protocols between directly connected routers

ANSWER: D E

Explanation:

Cisco Discovery Protocol (CDP) is a Cisco-proprietary Layer 2 neighbor discovery protocol used to learn information about directly connected Cisco devices. Two common administrative uses are (1) validating that a neighbor is reachable at Layer 2 and identifying what it is (device ID, platform, port ID), which helps troubleshoot cases where Layer 3 is failing but the physical/L2 adjacency may still be up; and (2) learning management addressing details (such as the neighbor's IP address) so you can connect to it for management (today typically SSH, historically Telnet). CDP advertisements also include other useful fields like the neighbor's capabilities and (on many platforms) the native VLAN/voice VLAN, which can assist with switchport troubleshooting.

Option D is correct because CDP operates at Layer 2 and can confirm a directly connected neighbor even when IP routing or other Layer 3 functions are broken. Option E is correct because CDP can display a neighbor's management address, enabling you to initiate a remote management session.

Option A is incorrect because CDP does not identify the physical cable type. Option B is incorrect because CDP does not report "network services" status (that's more in the realm of SNMP, NMAP, or specific show commands). Option C is not generally correct as "VLAN information" broadly (VLAN database/trunks) is not obtained via CDP (though native/voice VLAN may appear). Option F is incorrect because CDP does not indicate routing protocol status.

References: [Cisco CDP neighbors and troubleshooting](#), [Cisco IOS CDP Configuration Guide](#)

QUESTION NO: 168

Which two northbound APIs are found in a software-defined network? (Choose two.)

- A. REST
- B. OpenFlow
- C. SOAP
- D. NETCONF
- E. OpFlex

ANSWER: A C

Explanation:

In SDN, the **northbound API** is the interface exposed by the SDN controller to applications and orchestration systems (for example, automation tools, policy apps, and service catalogs). These northbound interfaces are commonly implemented using web-service styles such as **REST** (typically RESTful APIs over HTTP/HTTPS using JSON) and sometimes **SOAP** (XML-based web services). That's why **REST** and **SOAP** are valid northbound API examples.

By contrast, **OpenFlow** is a classic **southbound** protocol used between the controller and the forwarding devices to program flow tables. **NETCONF** is a network configuration protocol used for device configuration/management (often considered a management-plane interface and can be used by controllers, but it's not typically categorized as the SDN controller's northbound API in CCNA SDN context). **OpFlex** is associated with Cisco ACI as a policy protocol between policy elements and devices—again not a typical "northbound API" example for SDN applications.

References: [Cisco SDN overview](#), [Open Networking Foundation \(ONF\) SDN definition](#).