Question: 1

Refer to the exhibit.

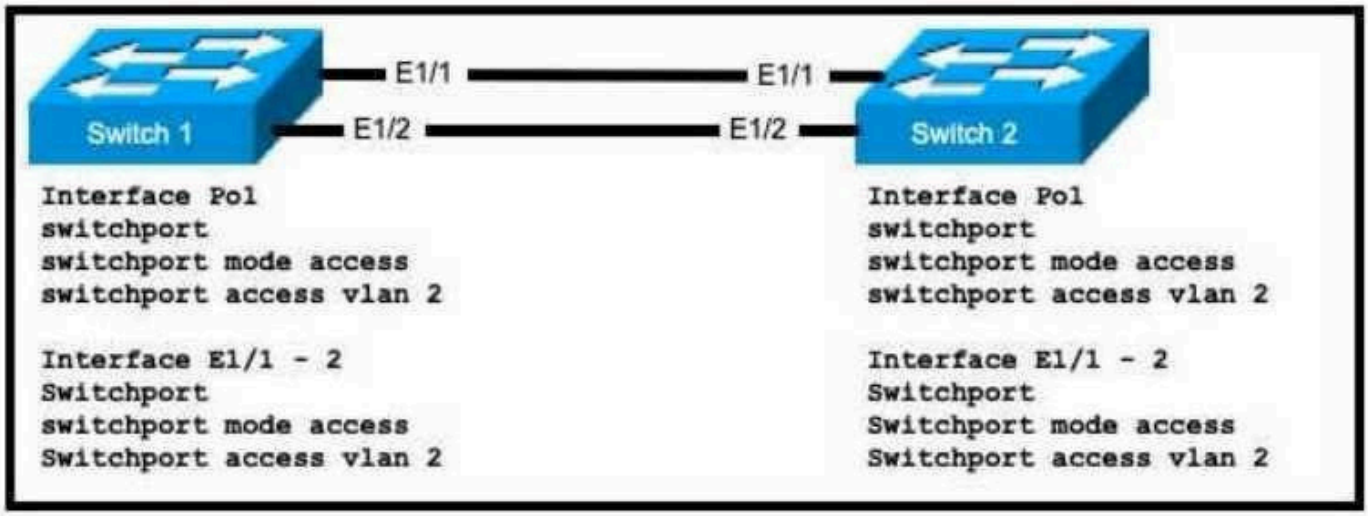

An engineer is configuring an EtherChannel using LACP between Switches 1 and 2 Which configuration must be applied so that only Switch 1 sends LACP initiation packets?

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

B. Switch1(config-if)#channel-group 1 mode paCsesrtiEvmepiSrewitch2(config-if)#channel-group 1 mode active

C. Switch1config-if)channel-group 1 mode active Switch2(config-if)#channel-group 1 mode passive

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

Answer:

C

Explanation:

For an EtherChannel to be formed using LACP (Link Aggregation Control Protocol), ports can be configured in either active or passive mode. Active mode: The port actively sends LACPDU (LACP Data Unit) packets to negotiate the EtherChannel. Passive mode: The port responds to LACPDU packets it receives but does not initiate the LACP negotiation. To ensure that only Switch 1 sends LACP initiation packets, Switch 1 must be configured in active mode, and Switch 2 must be configured in passive mode. This allows Switch 1 to initiate the LACP negotiation, and Switch 2 will respond to form the EtherChannel.

Why Incorrect Options are Wrong:

A: mode on on Switch 1 forces the channel without LACP negotiation, which is incompatible with LACP's passive mode on Switch 2. B: Switch 1 in passive mode will not initiate LACP; Switch 2 in active mode will, which contradicts the requirement. D: mode on on Switch 1 forces the channel without LACP, which is incompatible with LACP's active mode on Switch 2.

References:

Cisco IOS Interface and Hardware Component Configuration Guide, Release 15.0SY, "Configuring EtherChannels". (Search for "LACP Modes" or "channel-group mode active passive").
Specifically, the description of LACP modes:
active: "Places a port into an active negotiating state, in which the port initiates negotiations with other ports by sending LACP packets."
passive: "Places a port into a passive negotiating state, in which the port responds to LACP packets that it receives but does not initiate LACP packet negotiation."
Direct URL (example, actual URL may vary based on specific IOS version but content is consistent):
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750x3560x/software/release/15-01se/configuration/guide/scg3750x/swethchl.html#wp1114950 (Refer to the section on LACP (802.3ad) EtherChannel Configuration and LACP Modes).

IEEE Std 802.3-2018, "IEEE Standard for Ethernet," Clause 43: Link Aggregation. This standard defines LACP behavior.
Section 43.4.2 "LACP functional requirements" describes the active and passive modes.

Question: 2

Refer to the exhibit.

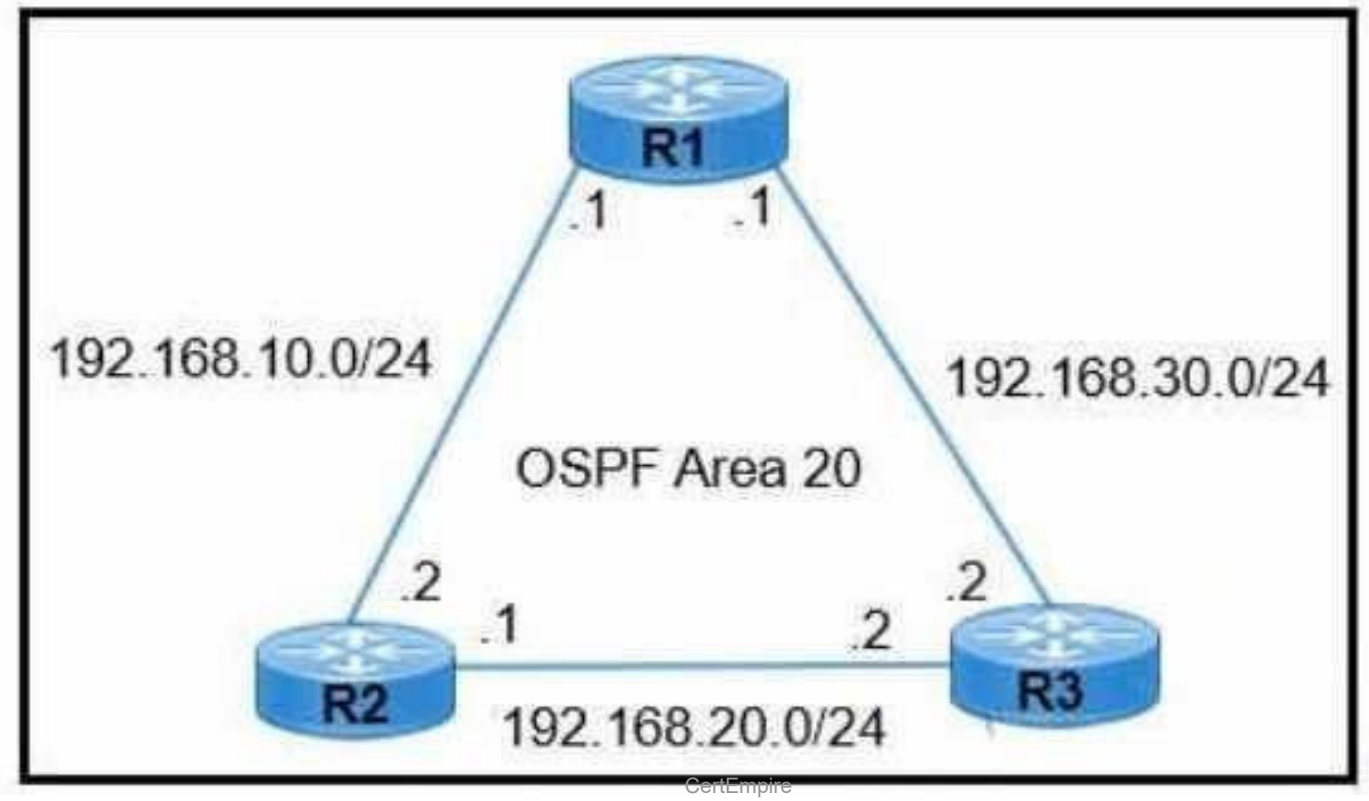

R1 learns all routes via OSPF Which command configures a backup static route on R1 to reach the 192 168.20.0/24 network via R3?

A. R1(config)#ip route 192.168.20.0 255.255.0.0 192.168.30.2

B. R1(config)#ip route 192.168.20.0 255.255.255.0 192.168.30.2 90

C. R1(config)#ip route 192.168.20.0 255.255.255.0 192.168.30.2 111

D. R1(config)#ip route 192.168.20.0 255.255.255.0 192.168.30.2

Answer:

C

Explanation:

To configure a backup static route that is used only when a dynamically learned route (like OSPF) fails, a floating static route is used. This is achieved by setting an administrative distance (AD) for the static route that is higher than the AD of the dynamic routing protocol. OSPF has a default administrative distance of 110. Therefore, the backup static route must have an AD greater than 110. Option C, R1(config)#ip route 192.168.20.0 255.255.255.0 192.168.30.2 111, correctly configures a static route to the 192.168.20.0/24 network via the next-hop 192.168.30.2 (R3's interface) with an administrative distance of 111. Since 111 is greater than OSPF's AD of 110,

this route will only be installed in the routing table if the OSPF route to 192.168.20.0/24 is lost.

Why Incorrect Options are Wrong:

/561 A: The subnet mask 255.255.0.0 is incorrect for a /24 network. Also, no AD is specified, defaulting to 1, which would override OSPF. B: The administrative distance of 90 is lower than OSPF's AD (110). This static route would be preferred over the OSPF route, not a backup. D: No administrative distance is specified, so it defaults to 1. This static route would be preferred over the OSPF route, not a backup.

References:

Cisco IOS IP Routing: Protocol-Independent Configuration Guide - "Configuring IP Routing Protocol-Independent Features" - "Administrative Distance" and "Floating Static Routes".
(A general concept covered in Cisco documentation).
"Administrative distance is the feature that routers use in order to select the best path when there are two or more different routes to the same destination from two different routing protocols. Administrative distance defines the reliability of a routing protocol. Each routing protocol is prioritized in order of most to least reliable (believable) with an administrative distance value."
"A floating static route is a static route that is used to back up a dynamically learned route... A floating static route is configured with a higher administrative distance than the dynamic routing protocol it is backing up."

Default Administrative Distances: OSPF is 110, Static route is 1. (Cisco Press, CCNA 200-301 Official Cert Guide, Volume 1, Chapter 15: Static IPv4 Routing, "Floating Static Routes"; Chapter 18: OSPFv2 Concepts, "OSPFv2 Administrative Distance").
Cisco Command Reference: ip route command.
ip route prefix mask ip-address interface-type interface-number ip-address distance ...
The distance parameter specifies the administrative distance. (Cisco IOS IP Routing: Protocol-Independent Command Reference - ip route).
/561

Question: 3

What is the difference between IPv6 unicast and anycast addressing? A:

A. IPv6 anycast nodes must be explicitly configured to recognize the anycast address, but IPv6 unicast nodes require no special configuration

B. IPv6 unicast nodes must be explicitly configured to recognize the unicast address, but IPv6 anycast

C. An individual IPv6 unicast address is supported on a single interface on one node but an IPv6 anycast address is assigned to a group of interfaces on multiple nodes.

D. Unlike an IPv6 anycast address, an IPv6 unicast address is assigned to a group of interfaces on multiple nodes

Answer:

C

Explanation:

An IPv6 unicast address uniquely identifies a single interface on a single node. Packets sent to a unicast address are delivered to that specific interface. In contrast, an IPv6 anycast address is assigned to a group of interfaces, typically on different nodes. Packets destined for an anycast address are routed to the nearest interface (in terms of routing protocol metrics) that shares that anycast address.

Why Incorrect Options are Wrong:

A: IPv6 anycast addresses are syntactically indistinguishable from unicast addresses; the "anycast" nature comes from configuring the same unicast address on multiple nodes and routing advertisements. B: IPv6 unicast addresses are standard configurations for interfaces; "explicitly configured to recognize" is not a distinguishing feature compared to anycast in this context. D: This statement incorrectly describes IPv6 unicast addresses; unicast addresses are for a single interface, not a group on multiple nodes.

References:

Cisco Press. (2020). CCNA 200-301 Official Cert Guide, Volume 1. Cisco Press.
Chapter 20, "Fundamentals of IP Version 6," section "IPv6 Addressing Formats and Conventions," subsection "Anycast Addresses": "An anycast address is a unicast address that is assigned to more than one interface, typically on different hosts. The routers then determine which host is closest and deliver the packet to that host."
Chapter 20, "Fundamentals of IP Version 6," section "IPv6 Addressing Formats and Conventions," subsection "Unicast Addresses": "A unicast address identifies a single interface on an IPv6 device."

Hinden, R., & Deering, S. (2006). RFC 4291: IP Version 6 Addressing Architecture. IETF.
Section 2.4, "Unicast Addresses": "An identifier for a single interface. A packet sent to a
/561
unicast address is delivered to the interface identified by that address."
Section 2.6, "Anycast Addresses": "An Anycast address is an address that is assigned to more than one interface (typically belonging to different nodes), with the property that a packet sent to an Anycast address is routed to the "nearest" interface having that address..." (Available at: https://www.rfc-editor.org/rfc/rfc4291.html#section-2.6)
Cisco. (n.d.). IPv6 Addressing and Basic Connectivity Configuration Guide, Cisco IOS XE Release 3S. Cisco Systems, Inc.
"IPv6 Unicast Addressing" section: "A unicast address is an identifier for a single interface, on a single node."
"IPv6 Anycast Addresses" section: "An anycast address is an address that is assigned to a set of interfaces that typically belong to different nodes." (Available at: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/xe-3s/ip6-xe-3sbook/ip6-addrg-basic-conn.html - specific page varies by version, search for "IPv6 Anycast Addresses").

Question: 4

Refer to the exhibit.

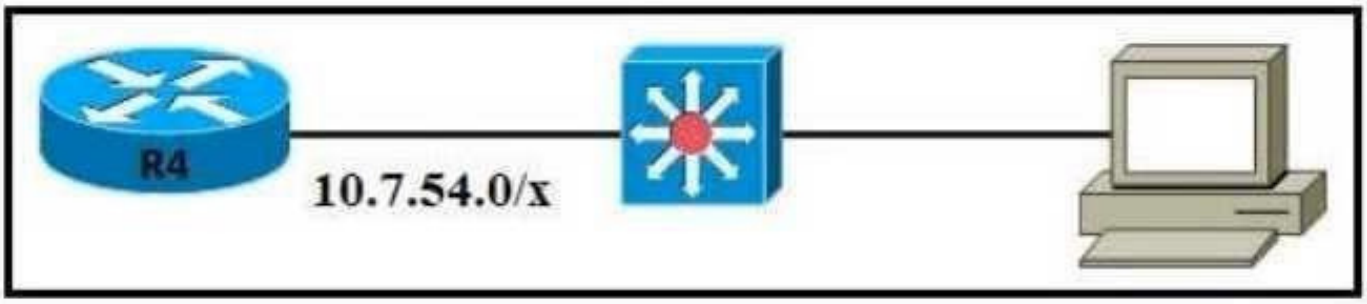

The router has been configured with a supernet to accommodate the requirement for 380 users on a subnet The requirement already considers 30% future growth. Which configuration verifies the IP subnet on router R4? A)

Subnet: 10.7.54.0
Subnet mask: 255.255.254.0
Broadcast address: 10.7.54.255
Usable IP address range: 10.7.54.1 - 10.7.55.254

B)

Subnet: 10.7.54.0
Subnet mask: 255.255.254.0
Broadcast address: 10.7.55.255
Usable IP address range: 10.7.54.1 - 10.7.55.254

C)

Subnet: 10.7.54.0
Subnet mask: 255.255.128.0
Broadcast address: 10.7.55.255
Usable IP address range: 10.7.54.1 - 10.7.55.254

D)

Subnet: 10.7.54.0
Subnet mask: 255.255.255.0
Broadcast address: 10.7.54.255
Usable IP address range: 10.7.54.1 - 10.7.55.254

A: Option A B: Option B C: Option C D: Option D

Answer:

B

Explanation:

The question asks to identify the configuration command that verifies the IP subnet on router R4. An "IP subnet" is defined by an IP address and its corresponding subnet mask. The requirement for 380 users, with 30% future growth, means the network must support at least 380 / (1 - 0.30) 543 hosts. This requires 10 host bits (210 - 2 = 1022 hosts), leading to a /22 prefix (32 - 10 = 22 network bits), or a subnet mask of 255.255.252.0. Let's evaluate the options: A) show ip interface brief: This command displays a summary of IP interface information, including the IP address and interface status, but it does not display the subnet mask. Therefore, it cannot fully verify the IP subnet. B) show ip interface GigabitEthernet0/0: This command displays detailed IP information for the specified interface, including its IP address, subnet mask (often shown as a prefix length, e.g., /22), broadcast address, and other IP-specific parameters. The output Internet address is 192.168.0.1/22 directly verifies the IP address and subnet mask. This is a primary command for IP interface verification. C) show running-config interface GigabitEthernet0/0: This command displays the configuration commands currently active for the specified interface, including the ip address command. The output ip address 192.168.0.1 255.255.252.0 verifies the configured IP subnet. D) show interfaces GigabitEthernet0/0:CTerhtEismpciroemmand provides extensive statistics for interface, covering Layer 1 and Layer 2 details (like MAC address, errors, duplex, speed) as well as Layer 3 information, including the IP address and subnet mask (Internet address is 192.168.0.1/22). While options B, C, and D all provide the necessary information to verify the IP subnet, the command show ip interface GigabitEthernet0/0 (Option B) is the most precise and directly applicable for verifying the IP-specific operational parameters of an interface. It focuses on IP details without the extensive L1/L2 data of show interfaces or showing the configuration text like show running-config.

Why Incorrect Options are Wrong:

A) show ip interface brief: This command does not display the subnet mask, which is essential for verifying the complete IP subnet. C) show running-config interface GigabitEthernet0/0: While it shows the configured IP subnet, show ip interface is generally preferred for verifying operational IP parameters. D) show interfaces GigabitEthernet0/0: This command is very verbose, providing much L1/L2 information; show ip interface is more focused on IP details.

References:

Cisco IOS IP Addressing Services Command Reference - show ip interface command: "To display the usability status of interfaces configured for IP, use the show ip interface command in privileged EXEC mode. This command displays the IP address, broadcast

address, and subnet mask among other IP-related information." (Specific Cisco documentation for this command would confirm its usage for displaying IP address and mask). Example: Cisco IOS Interface and Hardware Component Command Reference, show ip interface section. (A general reference as direct URLs to specific internal docs can be transient. The functionality is standard.) Odom, W. (2019). CCNA 200-301 Official Cert Guide, Volume 1. Cisco Press. Chapter 12, "Verifying IPv4 Addressing and Routing," discusses the use of show ip interface type number as a key command for viewing detailed IPv4 settings, including the IP address and mask. (e.g., page 300). Cisco IOS Master Command List, Release 15M&T - show ip interface: Describes the command as displaying "IP interface status and configuration." (Available via Cisco's website).

Question: 5

Refer to the exhibit.

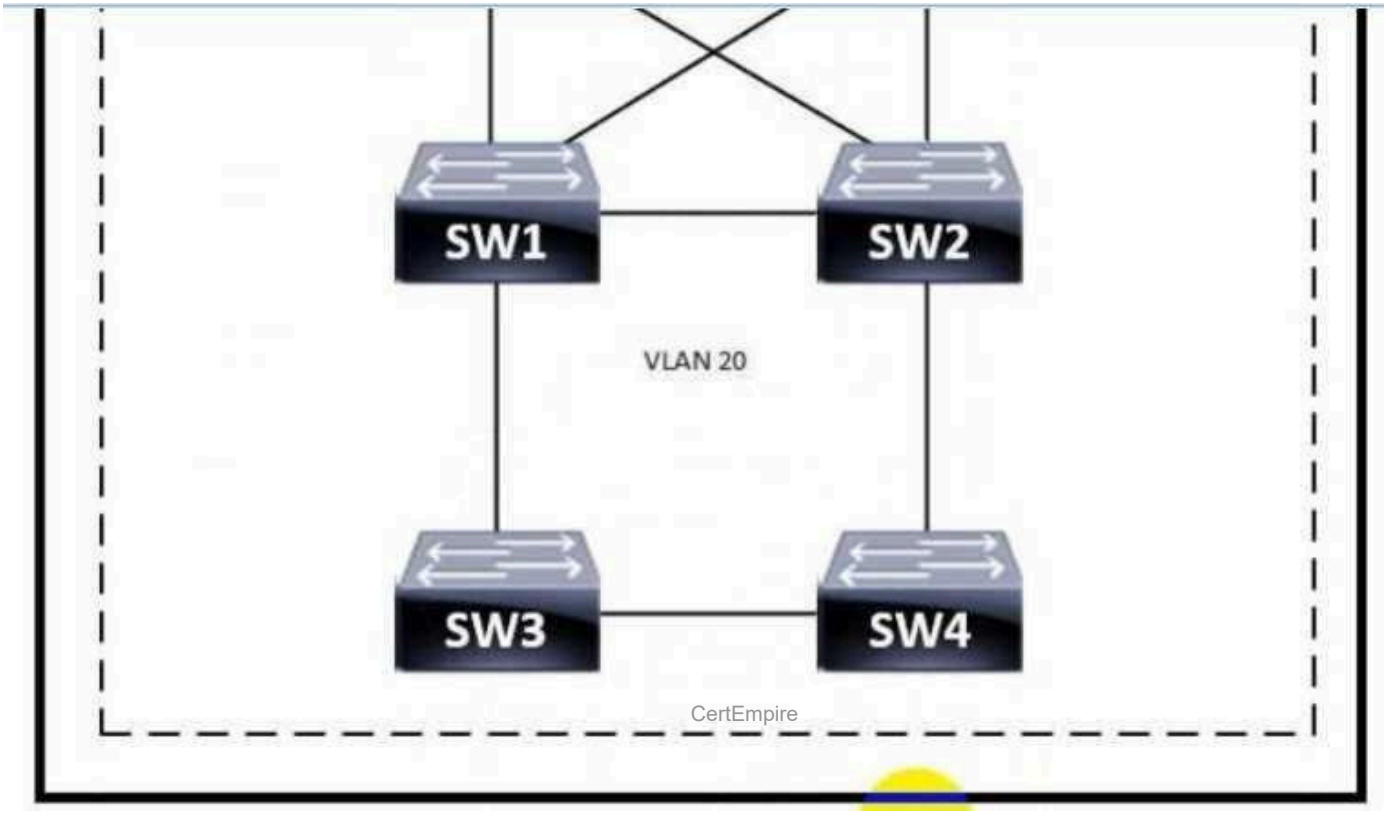

SW1 = 24596 0018.184e.3c00
SW2 = 28692 004a.14e5.4077
SW3 = 32788 0022.55cf.dd00
SW4 = 64000 0041.454d.407f

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

A. SW1

B. SW2

C. SW3

D. SW4

Answer:

C

Explanation:

The root bridge in a Spanning Tree Protocol (STP) topology is elected based on the lowest Bridge ID (BID). The BID consists of a 2-byte Bridge Priority and a 6-byte MAC address. For Per-VLAN Spanning Tree Plus (PVST+), the Bridge Priority field is a combination of a configurable priority value (a multiple of 4096) and the VLAN ID (Extended System ID). For VLAN 20, the STP priorities are calculated as follows: /561 SW1: Configurable Priority 32768 + VLAN ID 20 = 32788. BID: 32788:0000.0C11.1111 SW2: Configurable Priority 32768 + VLAN ID 20 = 32788. BID: 32788:0000.0C22.2222 SW3: Configurable Priority 28672 + VLAN ID 20 = 28692. BID: 28692:0000.0C33.3333 SW4: Configurable Priority 32768 + VLAN ID 20 = 32788. BID: 32788:0000.0C44.4444 Comparing the BIDs, SW3 has the lowest priority value (28692). Therefore, SW3 becomes the root bridge for VLAN 20. The equal speed of links affects path cost calculations for non- root bridges but not the root bridge election itself.

Why Incorrect Options are Wrong:

A (SW1): SW1 has a priority of 32788 for VLAN 20, which is higher than SW3's priority of 28692. B (SW2): SW2 has a priority of 32788 for VLAN 20, which is higher than SW3's priority of 28692. D (SW4): SW4 has a priority of 32788 for VLAN 20, which is higher than SW3's priority of 28692.

References:

Cisco Systems, "Understanding and Configuring Spanning Tree Protocol (STP) on Catalyst Switches" - Search on Cisco.com for this document or similar STP configuration guides. Specifically, sections covering Root Bridge Election and Bridge ID (BID) components including priority and extended system ID (VLAN ID).
Example content: The BID is an 8-byte value (2-byte priority + 6-byte MAC). The switch with the numerically lowest BID becomes the root. The priority field includes the extended system ID, which is the VLAN ID for PVST+.
Odom, W. (2019). CCNA 200-301 Official Cert Guide, Volume 1. Cisco Press. Chapter 10: "Spanning Tree Protocol Concepts", Section: "Electing a Root Bridge".
Specifically, the explanation of the Bridge ID (Priority and MAC Address) and how the Extended System ID (VLAN ID) is part of the priority field in PVST+. The default priority is 32768. The switch with the lowest priority value is chosen; if priorities are equal, the lowest MAC address is the tiebreaker. (Approx. pages 268-271 in some editions).
IEEE Std 802.1D-2004, "IEEE Standard for Local and metropolitan area networks--Media Access Control (MAC) Bridges." Section 17.3 "Bridge identifiers" and Section 17.4 "Root Bridge election". (While CCNA focuses on Cisco's implementation, the underlying principles are based on this standard).

Question: 6

Which protocol uses the SSL?

A. HTTP

B. SSH

C. HTTPS

D. Telnet

Answer:

C

Explanation:

HTTPS (Hypertext Transfer Protocol Secure) is the protocol that explicitly uses SSL (Secure Sockets Layer), or its successor TLS (Transport Layer Security), to encrypt the communication channel. This secures data exchanged between a web browser (client) and a web server, ensuring confidentiality and integrity.

Why Incorrect Options are Wrong:

A: HTTP: HTTP (Hypertext Transfer Protocol) is the standard protocol for transmitting hypermedia documents, but it is unencrypted and does not inherently use SSL. B: SSH: SSH (Secure Shell) is a cryptographic network protocol for secure remote login and other secure network services; it uses its own distinct security mechanisms, not SSL/TLS, for its core operations. D: Telnet: Telnet is an older, unsecure protocol for remote terminal access that transmits data, including credentials, in clear text and does not use SSL.

References:

1. Cisco. (n.d.). What Is HTTPS?. Cisco. Retrieved from https://www.cisco.com/c/en/us/products/security/what-is-https.html (This page states: "HTTPS (Hypertext Transfer Protocol Secure) is a secure version of the HTTP protocol that uses the SSL/TLS protocol for encryption and authentication.")
2. Kurose, J. F., & Ross, K. W. (2021). Computer Networking: A Top-Down Approach (8th ed.). Pearson. (Chapter 8, Section 8.6 "Securing TCP Connections: TLS" explains: "HTTP running on top of TLS is often referred to as HTTPS.")
3. Cisco. (n.d.). Transport Layer Security (TLS) Overview. Cisco. Retrieved from https://www.cisco.com/c/en/us/support/docs/security-vpn/public-key-infrastructurepki/212800-transport-layer-security-tls-overview.html (This document states: "It TLS is most familiar to users through its use in secure web browsing with HTTPS, which is HTTP

secured by TLS.")

Question: 7

Which two spanning-tree states are bypassed on an interface running PortFast? (Choose two.)

A. disabled B. listening C. forwarding D. learning E. blocking

Answer:

B, D

Explanation:

PortFast allows an interface to transition directly from the blocking state to the forwarding state, effectively bypassing the listening and learning states. This feature is intended for ports connected to end stations (e.g., PCs, servers) that are not expected to create Layer 2 loops. By skipping the listening and learning states, PortFast minimizes the time it takes for these devices to start communicating on the network after link-up.

Why Incorrect Options are Wrong:

A: disabled: A disabled port is administratively down or not operational; PortFast doesn't bypass this state, it applies to operational ports. C: forwarding: Forwarding is the final operational state that PortFast aims to reach quickly, not bypass. E: blocking: Blocking is the initial state from which PortFast transitions; it is not bypassed.

References:

Cisco, "Understanding and Configuring Spanning Tree Protocol (STP) on Catalyst Switches," (Search for "PortFast" section). While a specific URL changes, this document type is standard. A general search on Cisco's site for "Spanning Tree PortFast" will yield relevant official documentation. For example, the "Spanning Tree Protocol Configuration Guide, Cisco IOS XE Bengaluru 17.6.x (Catalyst 9300 Switches)" discusses PortFast behavior. Cisco, "Configuring Spanning Tree," Cisco IOS LAN Switching Configuration Guide, Release 12.2. (See section on PortFast). "When you enable PortFast on a switch port, the port immediately changes from the blocking state to the forwarding state, bypassing the usual listening and learning states." (This is a common statement in Cisco STP documentation across various platforms and IOS versions). IEEE Std 802.1D-2004, "IEEE Standard for Local and metropolitan area networks- Media

Access Control (MAC) Bridges," (See section on Port States and PortFast equivalent concepts like edge ports). While not using the term "PortFast" (a Cisco proprietary feature name), the standard discusses mechanisms for rapid transition for edge ports. Clause 17.19 "Port states" describes the standard states. PortFast is a Cisco implementation that accelerates this for specific port types.

Question: 8

How does Rapid PVST+ create a fast loop-free network topology? A:

A. It requires multiple links between core switches

B. It generates one spanning-tree instance for each VLAN

C. It maps multiple VLANs into the same spanning-tree instance

D. It uses multiple active paths between end stations.

Answer:

B

Explanation:

Rapid PVST+ (Rapid Per-VLAN Spanning Tree Plus) creates a fast, loop-free network topology by running an independent instance of Rapid Spanning Tree Protocol (RSTP, IEEE 802.1w) for each VLAN. This per-VLAN operation allows each VLAN to establish its own optimized, loop-free path. The RSTP mechanisms within each instance ensure rapid convergence to this stable, loop-free state, significantly faster than traditional STP.

Why Incorrect Options are Wrong:

A: It requires multiple links between core switches. RPVST+ manages existing links to prevent loops; multiple links provide redundancy but are not a requirement for its loop- prevention mechanism. C: It maps multiple VLANs into the same spanning-tree instance. This describes the behavior of Multiple Spanning Tree Protocol (MSTP), not RPVST+, which uses one instance per VLAN. D: It uses multiple active paths between end stations. RPVST+, like all STP variants, prevents loops by ensuring only a single active path between any two end stations within a VLAN.

References:

1. Cisco. (n.d.). Spanning Tree Protocol Configuration Guide, Cisco IOS XE Amsterdam 17.3.x (Catalyst 9300 Switches). "Configuring Rapid PVST+". Retrieved from https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/173/configurationguide/lyr23/b173lyr23spanningtree9300cg/configuringspanningtree.html#ID108
Relevant text: "Rapid PVST+ is the Cisco implementation of RSTP. It supports PVST+ (one spanning-tree instance for each VLAN)." and "Rapid PVST+ is the IEEE 802.1w (RSTP) standard implemented per VLAN. A single instance of STP runs on each configured VLAN..."
2. Odom, W. (2019). CCNA 200-301 Official Cert Guide, Volume 1. Cisco Press. Chapter 7, "Spanning Tree Protocol Concepts," section "Rapid PVST+ Concepts."
Relevant text: "Rapid PVST+ (RPVST+) is simply Cisco's version of RSTP that provides a

separate instance of 802.1w per VLAN." This highlights that the per-VLAN structure is fundamental to RPVST+, with RSTP providing the fast convergence.

Question: 9

Refer to the exhibit.

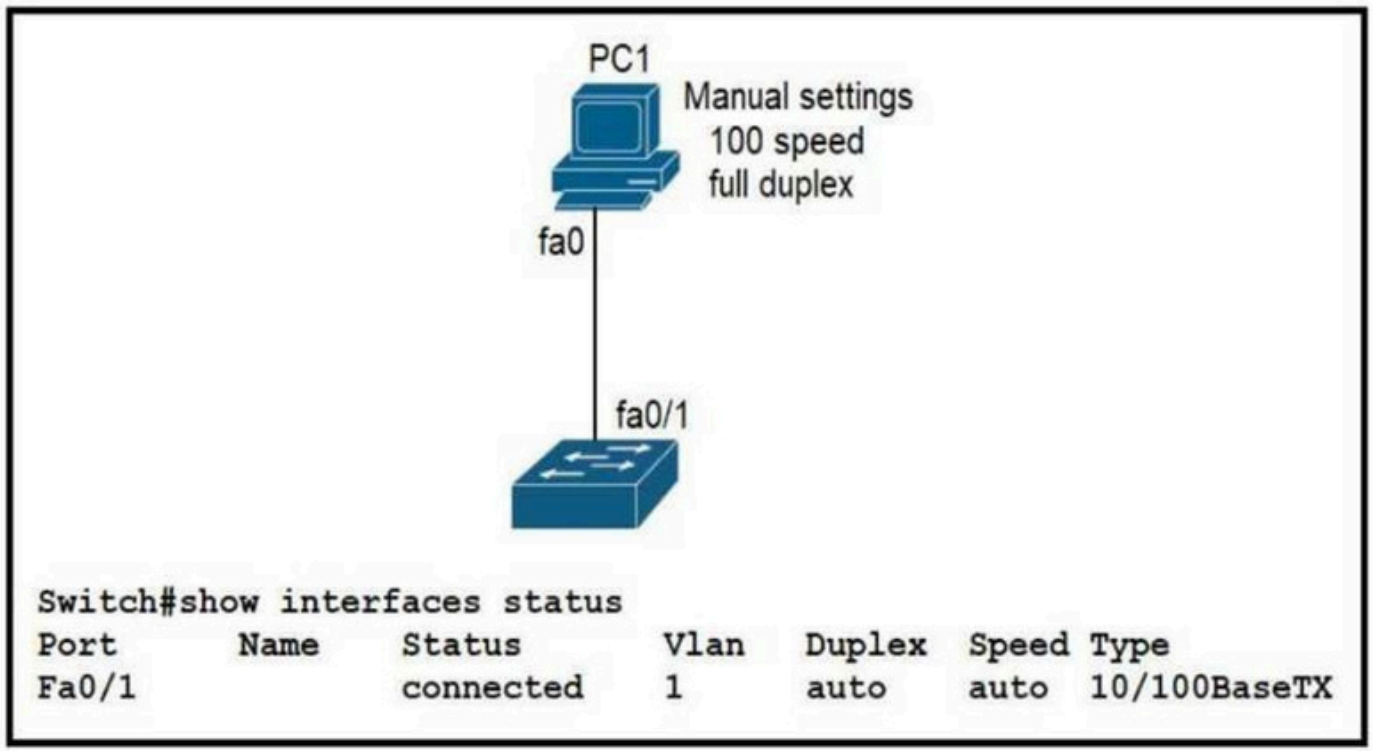

The link between PC1 and the switch is up. but is performing poorly. Which interface condition is causing the performance problem?

A. There is a duplex mismatch on the interface

B. There is an issue with the fiber on the switch interface.

C. There is a speed mismatch on the interface.

D. There is an interface type mismatch

Answer:

A

Explanation:

A link that is "up" but "performing poorly" is a classic symptom of a duplex mismatch on an Ethernet interface. When one side of the connection operates in full-duplex mode and the other in half-duplex mode, the full-duplex side may transmit while the half-duplex side is also transmitting or listening. This results in collisions detected by the half-duplex side and often CRC errors or other input errors on the full-duplex side, leading to retransmissions and severely degraded performance. The link itself, however, remains operational.

Why Incorrect Options are Wrong:

B: There is an issue with the fiber on the switch interface. Fiber issues (like dirty connectors or breaks) typically cause the link to be down or flap, not just perform poorly while remaining up. Also, PCs usually connect via copper. C: There is a speed mismatch on the interface. A speed mismatch (e.g., 100 Mbps vs. 1000 Mbps) usually results in the link not coming up at all if autonegotiation fails to find a common speed. D: There is an interface type mismatch. This is vague; if it implies incompatible physical media (e.g., copper to fiber without a converter), the link would not be up.

References:

Cisco IOS LAN Switching Configuration Guide, Release 12.2, "Configuring Interface Characteristics": "A common cause of performance issues on 10/100-Mb Ethernet links is when one port on the link operates at half-duplex while the other port operates at full-duplex." (This general principle is widely documented in Cisco's Ethernet troubleshooting guides).
Note: While a direct URL to a specific page for the 200-301 exam might be from a newer guide, the principle of duplex mismatch causing poor performance is fundamental and consistently documented. A representative modern guide would be:
Cisco, "Troubleshooting Ethernet" (General Cisco documentation often covers this). For example, in many Cisco troubleshooting guides for switches, duplex mismatch is highlighted. Example from a general troubleshooting context: "Symptoms of a Duplex Mismatch: ...The most common symptom of a duplex mismatch is slow throughput..." (Search for "duplex mismatch symptoms cisco" on cisco.com).
IEEE Std 802.3-2018, "IEEE Standard for Ethernet": Clause 28 (Physical Layer link signaling for 10 Mb/s, 100 Mb/s, and 1000 Mb/s Auto-Negotiation) and Clause 30 (Management) describe the auto-negotiation process. Failure in this process, or manual misconfiguration, can lead to duplex mismatches. The standard implies that if duplex modes are mismatched, collisions will occur on the half-duplex link, degrading performance. (e.g., Section 2, Clause 28.2.3.3 "Negotiated ability").
Wendell Odom, "CCNA 200-301 Official Cert Guide, Volume 1", Cisco Press. Chapter 2, "Fundamentals of Ethernet LANs," discusses Ethernet operation, including duplex settings and the impact of mismatches. (While commercial prep, Official Cert Guides are Cisco Press and generally align with official documentation principles). Specific reference to Cisco documentation is preferred over Cert Guides if available for direct citation.
Cisco Learning Network resources often reiterate these concepts. For instance, discussions on Ethernet troubleshooting.
Cisco Support Community, "Troubleshooting Switch Port Problems": This type of document often details duplex mismatch as a primary cause of poor performance on an active link. (e.g., https://community.cisco.com/t5/switching/troubleshooting-switch-port-problems/tap/3106605

- Section on Duplex Mismatch).

Question: 10

Which PoE mode enables powered-device detection and guarantees power when the device is detected?

A. dynamic

B. static

C. active

D. auto

Answer:

B

Explanation:

The static PoE mode on a Cisco switch pre-allocates and reserves the maximum configured power for a port, regardless of whether a powered device (PD) is currently connected. This mode enables PD detection, and crucially, it guarantees that power will be available and supplied to the port once a compatible PD is detected and connected, as the power has already been set aside from the switch's total power budget.

Why Incorrect Options are Wrong:

A: dynamic: "Dynamic" is not a standard Cisco PoE configuration mode keyword. While the auto mode behaves dynamically with power allocation, "dynamic" itself isn't the specific mode that guarantees power. C: active: While PoE operation involves active detection and negotiation, "active" is not a distinct Cisco configuration mode for PoE ports that specifically guarantees power allocation like static mode does. D: auto: The auto mode enables PD detection and supplies power if a PD is detected and if sufficient power is available in the switch's overall budget. It does not guarantee power if the budget is constrained.

References:

Cisco Systems, "Configuring PoE," Catalyst 9300 Series Switches Software Configuration Guide, Cisco IOS XE Bengaluru 17.6.x. Available: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/176/configurationguide/inthw/b176inthw9300cg/configuringpoe.html (See section "PoE Modes" or similar, describing auto and static functionalities).
Specifically, under "PoE Modes": "When you specify static, the switch reserves power for this port even when no device is connected. The switch guarantees that power will be provided to the port when a powered device is connected."
"When you specify auto, the switch automatically detects if the connected device is a powered device and allocates power from the switch power budget. If a device is connected

to a port, the switch turns on power only if it discovers the device and if power is available."

Question: 11

What is an expected outcome when network management automation is deployed?

A. A distributed management plane must be used.

B. Software upgrades are performed from a central controller

C. Complexity increases when new device configurations are added

D. Custom applications are needed to configure network devices

Answer:

B

Explanation:

Network management automation, particularly through centralized controllers like Cisco DNA Center, aims to simplify and streamline network operations. One significant and expected outcome is the ability to perform software upgrades for network devices from a central point. This reduces manual effort, ensures consistency, and allows for scheduled updates across the network, enhancing efficiency and reliability.

Why Incorrect Options are Wrong:

A: A distributed management plane must be used. Automation can utilize centralized management (e.g., a single controller) or distributed models; a distributed plane is not a mandatory outcome. C: Complexity increases when new device configurations are added. Automation typically aims to reduce complexity by using templates and standardized processes, making adding new devices easier. D: Custom applications are needed to configure network devices. While APIs allow for custom solutions, many automation platforms provide built-in tools and interfaces, so custom applications are not always needed.

References:

1. Cisco Press, "CCNA 200-301 Official Cert Guide, Volume 2" by Wendell Odom.
Chapter 22, "Introduction to Controller-Based Networking," discusses how controllers centralize network management. Specifically, features like Software Image Management (SWIM) in Cisco DNA Center exemplify centralized software upgrades. (e.g., Section: "Cisco DNA Center Assurance for Network Automation").
"Cisco DNA Center automates several key tasks, including... software image management (SWIM) to upgrade device OS images." (Paraphrased from typical descriptions of DNA Center capabilities).
2. Cisco Learning Network, "CCNA Study Material - Understanding Automation and Programmability."
This resource often highlights the benefits of automation, including centralized management

and simplified operations. The ability to push software updates from a central controller is a /561 key example of these benefits. (Specific page/section may vary, but the concept is core to Cisco's automation narrative). URL: (General reference to Cisco's official learning materials for CCNA) e.g., Content on Cisco DNA Center capabilities. 3. Cisco Documentation, "Cisco DNA Center User Guide" or "Cisco DNA Center Solution Overview." These documents detail the features of Cisco DNA Center, explicitly mentioning Software Image Management (SWIM) as a core function that allows administrators to manage and deploy software images to network devices from a central location. Example (conceptual, actual URL path may vary): https://www.cisco.com/c/en/us/products/cloud-systems-management/dna-center/index.html (See features related to automation and device lifecycle management).

Question: 12

Refer to the exhibit.

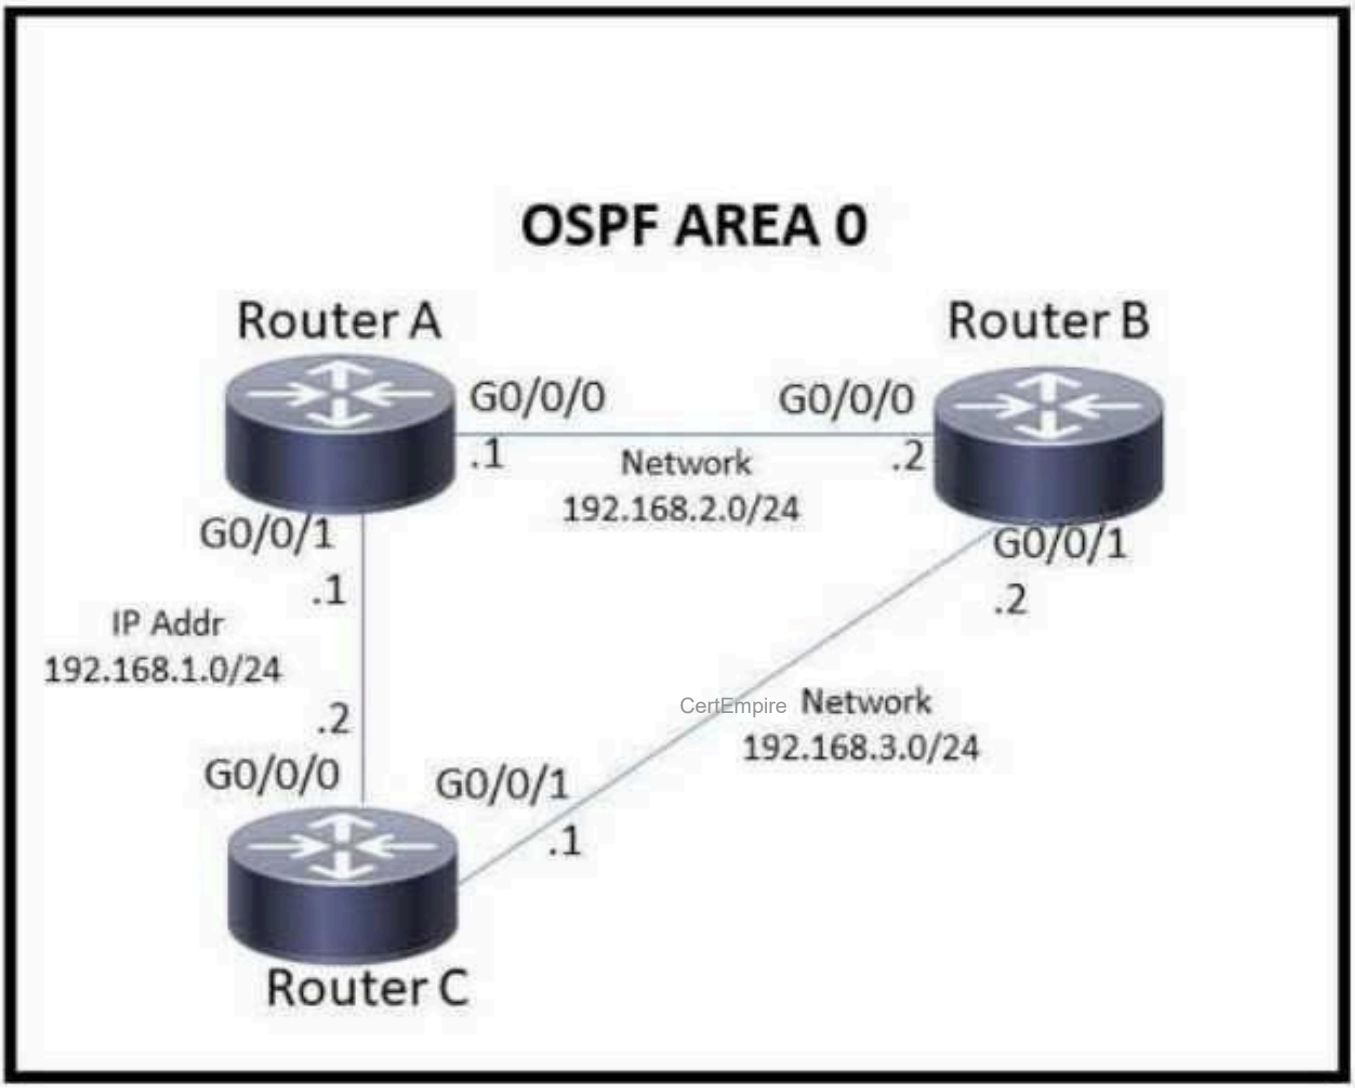

Which action must be taken to ensure that router A is elected as the DR for OSPF area 0? A:

A. Configure the OSPF priority on router A with the lowest value between the three routers.

B. Configure router B and router C as OSPF neighbors of router A.

C. Configure the router A interfaces with the highest OSPF priority value within the area.

D. Configure router A with a fixed OSPF router ID

Answer:

C

Explanation:

In OSPF, the Designated Router (DR) election on a multiaccess network segment (like Ethernet, which is implied by the diagram showing multiple routers connected) is determined primarily by the OSPF interface priority. The router with the highest OSPF priority on its interface connected to that segment will be elected as the DR. If priorities are tied, the router with the highest Router ID (RID) wins. To ensure Router A is elected DR, its interface(s) within Area 0 must have the highest OSPF priority among all routers on that segment.

References:

Cisco IOS IP Routing: OSPF Configuration Guide - OSPF Network Design Solutions: Designated Router Election: "OSPF elects a DR and a BDR on every multiaccess network... The router with the highest OSPF priority on a segment will be elected the DR for that segment. The same process is repeated for the BDR. If there is a tie in priority, the router with the higher router ID will be chosen."
URL: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iprouteospf/configuration/15-mt/iro15-mt-book/iro-ospf-overview.html#GUID-70A29167-363F-47F9-959F-263787A396C0
(Refer to the section on DR/BDR election)
Cisco IOS IP Routing: OSPF Command Reference - ip ospf priority: "To set the router priority, which helps determine the designated router (DR) for a network, use the ip ospf priority command in interface configuration mode. The router with the highest priority is elected as the DR."
URL: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iprouteospf/command/iro-crbook/ospf-i1.html#wp1113609037
RFC 2328 - OSPF Version 2, Section 9.4: Electing the Designated Router: "The Designated Router is elected from the set of routers belonging to the network that are eligible to be Designated Router. .... If two routers on the network have the same Router Priority, the one with the highest Router ID is chosen." (This confirms priority is checked first).
URL: https://datatracker.ietf.org/doc/html/rfc2328#section-9.4

Question: 13

Refer to the exhibit.

```
R1# show ip route
     Codes:
     C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP, D -
     EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area N1 - OSPF NSSA
     external type 1, N2 - OSPF NSSA external type 2, E1 - OSPF external type
     1, E2 - OSPF external type 2, E - EGP
     i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default,
     U - per-user static route, o - ODR
Gateway of last resort is not set
C   10.0.0.0/8 is directly connected, Loopback0
    10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O   10.0.1.3/32 [110/100] via 10.0.1.100, 00:39:08, Serial0
C   10.0.1.0/24 is directly connected, Serial0
O   10.0.1.5/32 [110/5] via 10.0.1.50, 00:39:08, Serial0
O   10.0.10.0/24 [110/10] via 10.0.1.4, 00:39:08, Gigabit Ethernet 0/0
D   10.0.10.0/24 [90/10] via 10.0.1.5, 00:39:08, Gigabit Ethernet 0/1
```

Web traffic is coming in from the WAN interface. Which route takes precedence when the router is processing traffic destined for the LAN network at 10 0.10.0/24?

A. via next-hop 10.0.1.5 B.

via next-hop 10 0 1.4 C.

via next-hop 10.0 1.50 D.

via next-hop 10.0 1 100

Answer:

B

Explanation:

The router selects the best path based on the longest prefix match, then administrative distance (AD), and finally the metric. In this scenario, all listed routes are for the same destination prefix 10.0.10.0/24. Therefore, the selection will be based on the Administrative Distance. The ADs are: Static route (S): 1 EIGRP route (D): 90 OSPF route (O): 110 RIP route (R): 120 The static route S 10.0.10.0/24 1/0 via 10.0.1.4 has the lowest AD (1) and will be chosen.

References:

Cisco Press. CCNA 200-301 Official Cert Guide, Volume 1. Chapter 15: IP Routing. Section: "How Routers Make Forwarding Decisions". (Administrative Distance is a key factor when multiple routes to the same destination prefix exist).
Cisco Systems. IP Routing: Protocol-Independent Configuration Guide, Cisco IOS XE

Release 3S - Route Selection in Cisco IOS Support. (URL: https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocoleigrp/8651-21.html - This document explains route selection, including the role of AD).
"If the prefix length is the same, the router prefers the route with the lower administrative distance."
Cisco Systems. Default Administrative Distances. (URL: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproutepi/configuration/15-mt/iri-15-mtbook/iri-adm-dist.html - Lists default ADs: Static route = 1, EIGRP (internal) = 90, OSPF = 110, RIP = 120).

Question: 14

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

A. preshared key that authenticates connections

B. RSA token

C. CA that grants certificates

D. clear-text password that authenticates connections

E. one or more CRLs

Answer:

C, E

Explanation:

A Public Key Infrastructure (PKI) is a framework for managing digital certificates and public- key encryption. Core components include a Certificate Authority (CA), which issues and manages these certificates, binding public keys to entities. Another vital component is the Certificate Revocation List (CRL), published by the CA or a CRL issuer, which lists certificates that are no longer trusted before their official expiry. These elements ensure the integrity and trustworthiness of digital identities and communications within tCheertEPmKpiIresystem.

Why Incorrect Options are Wrong:

A: preshared key that authenticates connections: Preshared keys are symmetric; PKI uses asymmetric cryptography and certificates, not shared secrets. B: RSA token: RSA tokens are authentication devices (e.g., for OTPs), not structural components of the PKI itself. D: clear-text password that authenticates connections: Clear-text passwords offer no cryptographic security and are unrelated to PKI's certificate-based authentication.

References:

1. Cisco Systems, "Public Key Infrastructure Overview" (Part of Cisco IOS Security Configuration Guide).
Direct URL: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/secconnpki/configuration/15mt/sec-pki-15-mt-book/sec-pki-overview.html
Relevant Section: The document lists "Certificate authority (CA)" and "Certificate revocation list (CRL)" as components of a PKI.
2. IETF RFC 5280, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile".
Direct URL: https://datatracker.ietf.org/doc/html/rfc5280
Relevant Section: Section 3, "PKI Components," explicitly lists "CA: certification authority"

and mentions "CRL issuer" and repositories for CRLs as key parts of a PKI.
3. Odom, W. (2020). CCNA 200-301 Official Cert Guide, Volume 2. Cisco Press.
/561
Relevant Chapter: Chapter 26, "Basic Security Concepts," discusses PKI and its components, including Certificate Authorities and the mechanisms for certificate validation and revocation (like CRLs). (Specific page numbers vary by edition, but the concept is covered in the PKI section).

Question: 15

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

A. They enable automatic failover of the default gateway.

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

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

D. They prevent loops in the Layer 2 network.

E. They allow encrypted traffic.

Answer:

A, B

Explanation:

First Hop Redundancy Protocols (FHRPs) enhance network reliability by providing default gateway redundancy. They enable automatic failover (A) if the primary gateway device fails, ensuring uninterrupted connectivity for end-user devices. FHRPs also allow multiple physical routers to present themselves as a single virtual gateway (B) to hosts on the network, using a shared virtual IP and MAC address. This simplifies host configuration and improves fault tolerance.

Why Incorrect Options are Wrong:

C: Bundling multiple ports to increase bandwidth is a function of EtherChannel (Link Aggregation), not FHRPs. D: Preventing loops in the Layer 2 network is the primary role of Spanning Tree Protocol (STP). E: FHRPs are designed for gateway redundancy and do not inherently provide traffic encryption; other protocols (e.g., IPsec, TLS) handle encryption.

References:

1. Odom, W. (2019). CCNA 200-301 Official Cert Guide, Volume 1. Cisco Press. Chapter 10, "First Hop Redundancy Protocols," states, "The primary benefit of FHRPs is that the end-user devices can continue to send packets to the same default gateway IP address, and the FHRP takes care of the rest, failing over to a working router when the currently active router fails" (supports A) and "FHRPs allow all hosts to use a single default gateway IP address and MAC address, while also allowing the network to have multiple physical routers that can act as that default gateway" (supports B).
2. Cisco. (n.d.). IP Routing: HSRP Configuration Guide, Cisco IOS XE Release 3S - HSRP Overview. Cisco. Retrieved from https://www.cisco.com/c/en/us/td/docs/iosxml/ios/iproutehsrp/configuration/xe-3s/irh-xe-3s-book/irh-hsrp.html. This document explains that HSRP (an FHRP) "provides first-hop routing redundancy" and involves "selecting an