DUMPSBOSS.

CompTIA Network+ Certification Exam

CompTIA N10-009

Version Demo

Total Demo Questions: 20

Total Premium Questions: 431

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

Which of the following source control features allows an administrator to test a new configuration without changing the primary configuration?

A. Central repository

B. Conflict identification

C. Branching

D. Version control

ANSWER: C

Explanation:

The feature you're looking for is **branching**. A branch is basically a separate "copy" of the main configuration (often called main/master) where you can make changes, try new settings, and run tests without touching what's currently in production.

That's why branching is so useful for network and system admins: if your experiment goes sideways, you haven't broken the primary config. And if everything works, you can merge the branch back into the main line later.

The other choices don't fit as well. A **central repository** is just where the code/config is stored. **Conflict identification** helps spot overlapping edits, but it doesn't create an isolated test space. And **version control** is the overall system (like Git), not the specific feature that gives you a safe sandbox.

References: https://git-scm.com/book/en/v2/Git-Branching-Branches-in-a-Nutshell and https://www.atlassian.com/git/tutorials/using-branches

QUESTION NO: 2

Which of the following best describes the amount of time between a disruptive event and the point that affected resources need to be back to fully functional status?

A. RTO

B. MTBF

C. RPO

D. MTTR

ANSWER: A

Explanation:

The metric that fits this description is **RTO (Recovery Time Objective)**. RTO is basically the "how fast do we need to be back up?" number. It's the maximum amount of time the business can tolerate systems being down after an outage or other disruptive event before services must be restored to a fully working state.

The other choices sound similar but measure different things. **RPO** is about data loss—how far back in time you can go (like "we can only afford to lose 15 minutes of data"). **MTBF** is a reliability metric that estimates the average time between failures, and **MTTR** is the average time it takes to fix something once it breaks. Those are useful, but they don't define the required recovery window the way RTO does.

If you want a solid reference, check out NIST's contingency planning guidance, which discusses recovery objectives like RTO: https://csrc.nist.gov/publications/detail/sp/800-34/rev-1/final

QUESTION NO: 3

A network engineer is setting up a new VoIP network for a customer. The current network is segmented only for computers and servers. No additional switchports can be used in the new network. Which of the following does the engineer need to do to configure the network correctly? (Select two).

A. Change network translation definitions

B. Enable 802.1Q

C. Implement a routing protocol

D. Set up voice VLANs

E. Reconfigure the DNS

F. Place devices in the perimeter network

ANSWER: B D

Explanation:

If you can't spare extra switchports, the usual VoIP setup is to let the IP phone and the user's PC share the same physical port (the PC plugs into the phone). To keep that traffic cleanly separated, you configure a dedicated voice VLAN. That way, phone traffic doesn't mix with regular data traffic, and it's much easier to apply QoS and security rules just to voice.

To make multiple VLANs work over that single port/link, you also need VLAN tagging. That's where 802.1Q comes in—it tags frames so the switch knows what belongs to the voice VLAN versus the data VLAN, even though it's all riding the same wire. Without 802.1Q, the switch can't reliably separate voice and data on the shared port.

The other options (DNS changes, routing protocols, NAT translations, or moving devices to a perimeter network) don't solve the core problem here, which is carrying both voice and data over the same switchport while keeping them logically separated.

References: https://en.wikipedia.org/wiki/IEEE_802.1Q and https://www.cisco.com/c/en/us/support/docs/lan-switching/voice-vlan/116364-configure-voice-vlan-00.html

QUESTION NO: 4

A detective is investigating an identity theft case in which the target had an RFID-protected payment card issued and compromised in the same day. The only place the target claims to have used the card was at a local convenience store. The detective notices a video camera at the store is placed in such a way that customers' credentials can be seen when they pay. Which of the following best explains this social engineering technique?

A. Shoulder surfing

B. Impersonation

C. Vishing

D. Tailgating

ANSWER: A

Explanation:

The giveaway here is the camera angle: it's positioned so it can capture what customers type or what's visible on the card/terminal during checkout. That's classic **shoulder surfing**—someone is watching (directly or with a camera) to steal sensitive info like a PIN, card details, or other credentials while the victim is using them.

The other choices don't really fit. **Impersonation** is when the attacker pretends to be someone else (like a bank employee). **Vishing** is voice phishing over the phone. **Tailgating** is physically following someone into a restricted area. None of those involve capturing payment credentials by observation.

If you want a solid reference, see the shoulder surfing description in OWASP's social engineering overview: https://owasp.org/www-community/attacks/Social_Engineering

QUESTION NO: 5

A network engineer is now in charge of all SNMP management in the organization. The engineer must use a SNMP version that does not utilize plaintext data. Which of the following is the minimum version of SNMP that supports this requirement?

A. v1

B. v2c

C. v2u

D. v3

ANSWER: D

Explanation:

The key requirement here is "does not utilize plaintext data." SNMPv1 and SNMPv2c both rely on community strings, which are basically shared passwords sent in clear text. That means anyone who can sniff the traffic could potentially read the community string and pull info from devices.

SNMPv3 is the first version that adds real security features, including authentication (so you know who's talking) and encryption (so the SNMP data isn't readable on the wire). Because encryption is what prevents plaintext exposure, SNMPv3 is the minimum version that meets the requirement.

Option v2u shows up sometimes in old discussions as "SNMPv2u" (user-based security), but it was never widely adopted as a standard replacement in real environments. For practical and exam purposes, SNMPv3 is the clear answer when you need encrypted SNMP.

References: https://www.rfc-editor.org/rfc/rfc3414 and https://www.cisco.com/c/en/us/support/docs/ip/simple-network-management-protocol-snmp/7282-12.html

QUESTION NO: 6

A company has observed increased user traffic to gambling websites and wants to limit this behavior on work computers. Which of the following should the company most likely implement?

A. ACLs

B. Content filter

C. Port security

D. Screened subnet

ANSWER: B

Explanation:

To cut down on gambling site visits from work PCs, the most direct fix is a **content filter**. A content filter (often part of a secure web gateway, DNS filter, or next-gen firewall) can block sites by category—like “gambling”—so users simply can’t load those pages.

The other options don’t really match the goal. An **ACL** is great for allowing or denying traffic based on IPs, ports, or protocols, but it’s not designed for “block this type of website” unless you’re doing clunky, hard-to-maintain rules. **Port security** is about stopping unknown devices from plugging into switch ports, not controlling where users browse. A **screened subnet (DMZ)** is used to protect public-facing servers by isolating them, and it won’t stop internal users from visiting certain websites.

References: https://www.cisco.com/c/en/us/products/security/web-security-appliance/index.html and https://en.wikipedia.org/wiki/Content-control_software

QUESTION NO: 7

A network engineer is setting up a new VoIP network for a customer. The current network is segmented only for computers and servers. No additional switch ports can be used in the new network. Which of the following does the engineer need to do to configure the network correctly? (Select TWO).

A. Change network translation definitions

B. Enable 802.1Q

C. Implement a routing protocol

D. Set up voice VLANs

E. Reconfigure the DNS

F. Place devices in the perimeter network

ANSWER: B D

Explanation:

Since you can't use any extra switch ports, the phone and the PC have to share the same physical switchport. The normal way to make that work cleanly is to use VLAN tagging so the switch can tell voice traffic from regular data traffic, even though it's coming in on the same port.

That's where **802.1Q** comes in. Enabling 802.1Q lets the switch handle tagged VLAN traffic, which is how IP phones typically separate voice from data when they're daisy-chained with a computer. Then you also need to **set up a voice VLAN** so the voice traffic is placed into its own VLAN automatically. This keeps voice traffic isolated and makes it much easier to apply QoS settings so calls don't sound choppy.

The other choices don't solve the "same port" requirement. NAT and DNS aren't switchport/VLAN features, routing protocols aren't required just to add a voice VLAN, and a perimeter/DMZ network is for public-facing services, not internal VoIP segmentation.

References: https://www.cisco.com/c/en/us/support/docs/smb/switches/cisco-small-business-100-series-unmanaged-switches/smb5541-voice-vlan.html and https://www.cisco.com/c/en/us/support/docs/lan-switching/8021q/17056-741-00.html

QUESTION NO: 8 - (SIMULATION)

SIMULATION

You are tasked with verifying the following requirements are met in order to ensure network security. Requirements:

Datacenter

Ensure network is sub netted to allow all devices to communicate properly while minimizing address space usage

Provide a dedicated server to resolve IP addresses and hostnames correctly and handle port 53 traffic Building A

Ensure network is sub netted to allow all devices to communicate properly while minimizing address space usage

Provide devices to support 5 additional different office users

Add an additional mobile user

Replace the Telnet server with a more secure solution

Screened subnet

Ensure network is sub netted to allow all devices to communicate properly while minimizing address space usage

Provide a server to handle external 80 traffic

Provide a server to handle port 20 traffic

INSTRUCTIONS

Drag and drop objects onto the appropriate locations. Objects can be used multiple times and not all placeholders need to be filled.

Available objects are located in both the Servers and Devices tabs of the Drag & Drop menu.

If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

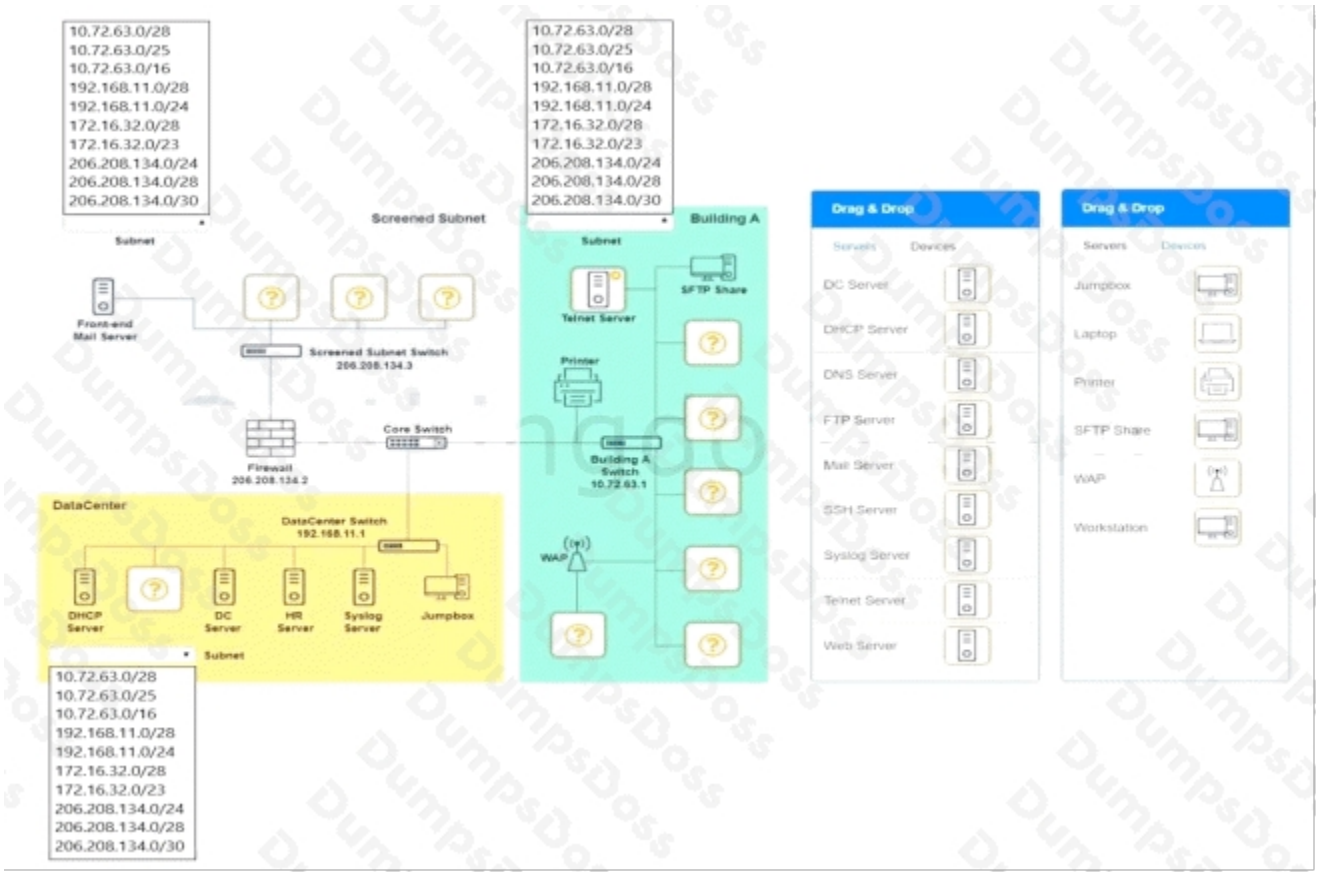

ANSWER: See the explanation for detailed informationon this simulation.

Explanation:

Screened Subnet devices " Web server, FTP server

Building A devices " SSH server top left, workstations on all 5 on the right, laptop on bottom left Datacenter devices " DNS server.

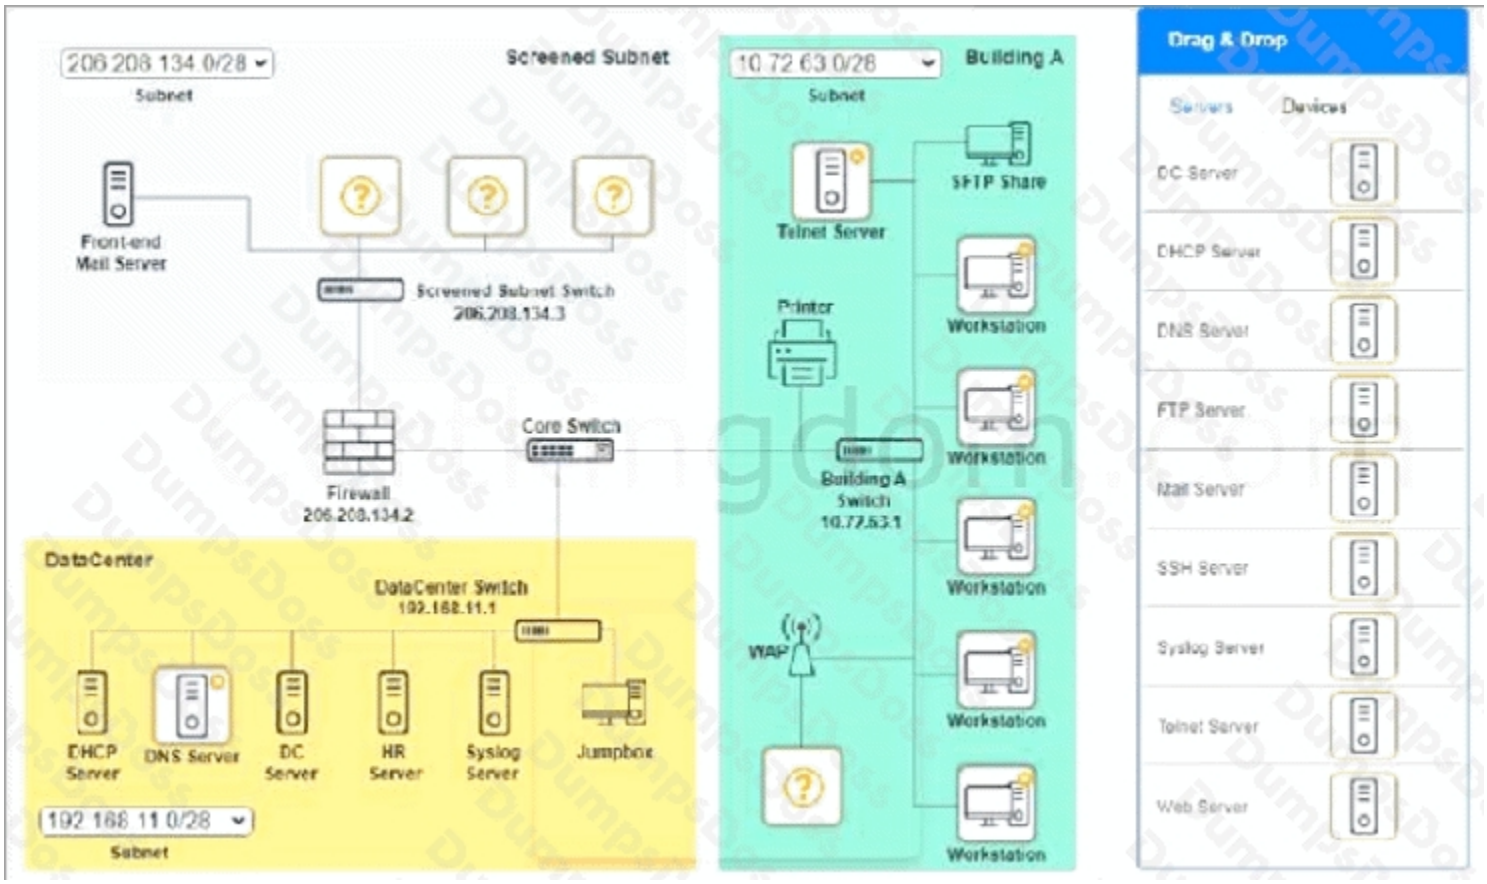

A screenshot of a computer AI-generated content may be incorrect.

A screenshot of a computer AI-generated content may be incorrect. A screenshot of a computer AI-generated content may be incorrect.

QUESTION NO: 9

An employee in a corporate office clicks on a link in an email that was forwarded to them. The employee is redirected to a splash page that says the page is restricted. Which of the following security solutions is most likely in place?

A. DLP

B. Captive portal

C. Content filtering

D. DNS sink holing

ANSWER: C

Explanation:

A “restricted page” splash screen after clicking a link is classic web/content filtering behavior. The company’s filter (often on a secure web gateway or firewall) checks the URL against categories and policies—like phishing, malware, gambling, or “uncategorized”—and if it breaks the rules, it blocks the request and shows a block page instead of loading the site.

DLP doesn’t usually stop you from visiting a website; it’s more about preventing sensitive data from leaving the company (uploads, email attachments, copy/paste of protected data, etc.). A captive portal is that login/accept-terms page you see on guest Wi-Fi before you get internet access, not a targeted “this site is restricted” message. DNS sinkholing can redirect known bad domains to a safe IP, but the user experience is often a DNS failure or a security landing page for malicious domains specifically—not general policy-based “restricted content” blocking.

References: https://en.wikipedia.org/wiki/Content-control_software and https://www.cloudflare.com/learning/access-management/what-is-a-secure-web-gateway/

QUESTION NO: 10

A secure communication link needs to be configured between data centers via the internet. The data centers are located in different regions. Which of the following is the best protocol for the network administrator to use?

A. DCI

B. GRE

C. VXLAN

D. IPSec

ANSWER: D

Explanation:

For a secure link between two data centers over the public internet, you typically want a site-to-site VPN. IPSec is the go-to protocol for that because it can encrypt the traffic, authenticate the endpoints, and protect the data from being changed in transit. In other words, it’s built specifically to make “untrusted network” links (like the internet) safe for private traffic.

GRE is useful when you need to wrap (encapsulate) one type of traffic inside another—often to carry routing protocols or multicast—but by itself it doesn’t encrypt anything. People commonly pair GRE with IPSec (GRE over IPSec) when they need both encapsulation and encryption, but if you have to pick the best single protocol for security, IPSec wins.

VXLAN is more of a data center overlay technology for stretching Layer 2 networks over Layer 3, not a straightforward “secure tunnel over the internet” solution. And “DCI” is a broad concept (data center interconnect), not a specific security protocol.

References: https://www.rfc-editor.org/rfc/rfc4301, https://www.cloudflare.com/learning/network-layer/what-is-ipsec/

QUESTION NO: 11

An ISP provided a company with a pre-configured modem and five public static IP addresses. Which of the following does the company's firewall require to access the internet? (Select TWO).

A. NTP server

B. Default gateway

C. The modem's IP address

D. One static IP address

E. DNS servers

F. DHCP server

ANSWER: B D

Explanation:

To get a firewall online with public static IPs, you need two basic things: an IP address to use on the WAN side, and a way to reach “the rest of the internet.” That’s why **one static IP address** (assigned to the firewall’s outside/WAN interface) is required—without it, the firewall can’t participate on the public network at all.

You also need the **default gateway** from the ISP. The gateway is the next hop the firewall sends traffic to when the destination isn’t on its local subnet (which is basically everything on the internet). No gateway means traffic never leaves your ISP-facing network, even if the IP is correct.

The other choices are nice-to-haves, not must-haves for raw connectivity. DNS helps you browse by names (like example.com), but you can still reach sites by IP without it. NTP is only for time sync. DHCP doesn’t apply because you’re using static addressing. And the modem’s management IP is only relevant if you’re logging into the modem, not for routing traffic through it.

References: https://www.cloudflare.com/learning/network-layer/what-is-a-default-gateway/ and https://en.wikipedia.org/wiki/Static_IP_address

QUESTION NO: 12 - (SIMULATION)

SIMULATION

Users are unable to access files on their department share located on file server 2. The network administrator has been tasked with validating routing between networks hosting workstation A and file server 2.

INSTRUCTIONS

Click on each router to review output, identify any issues, and configure the appropriate solution.

If at any time you would like to bring back the initial state of the simulation, please click the Reset All

button.

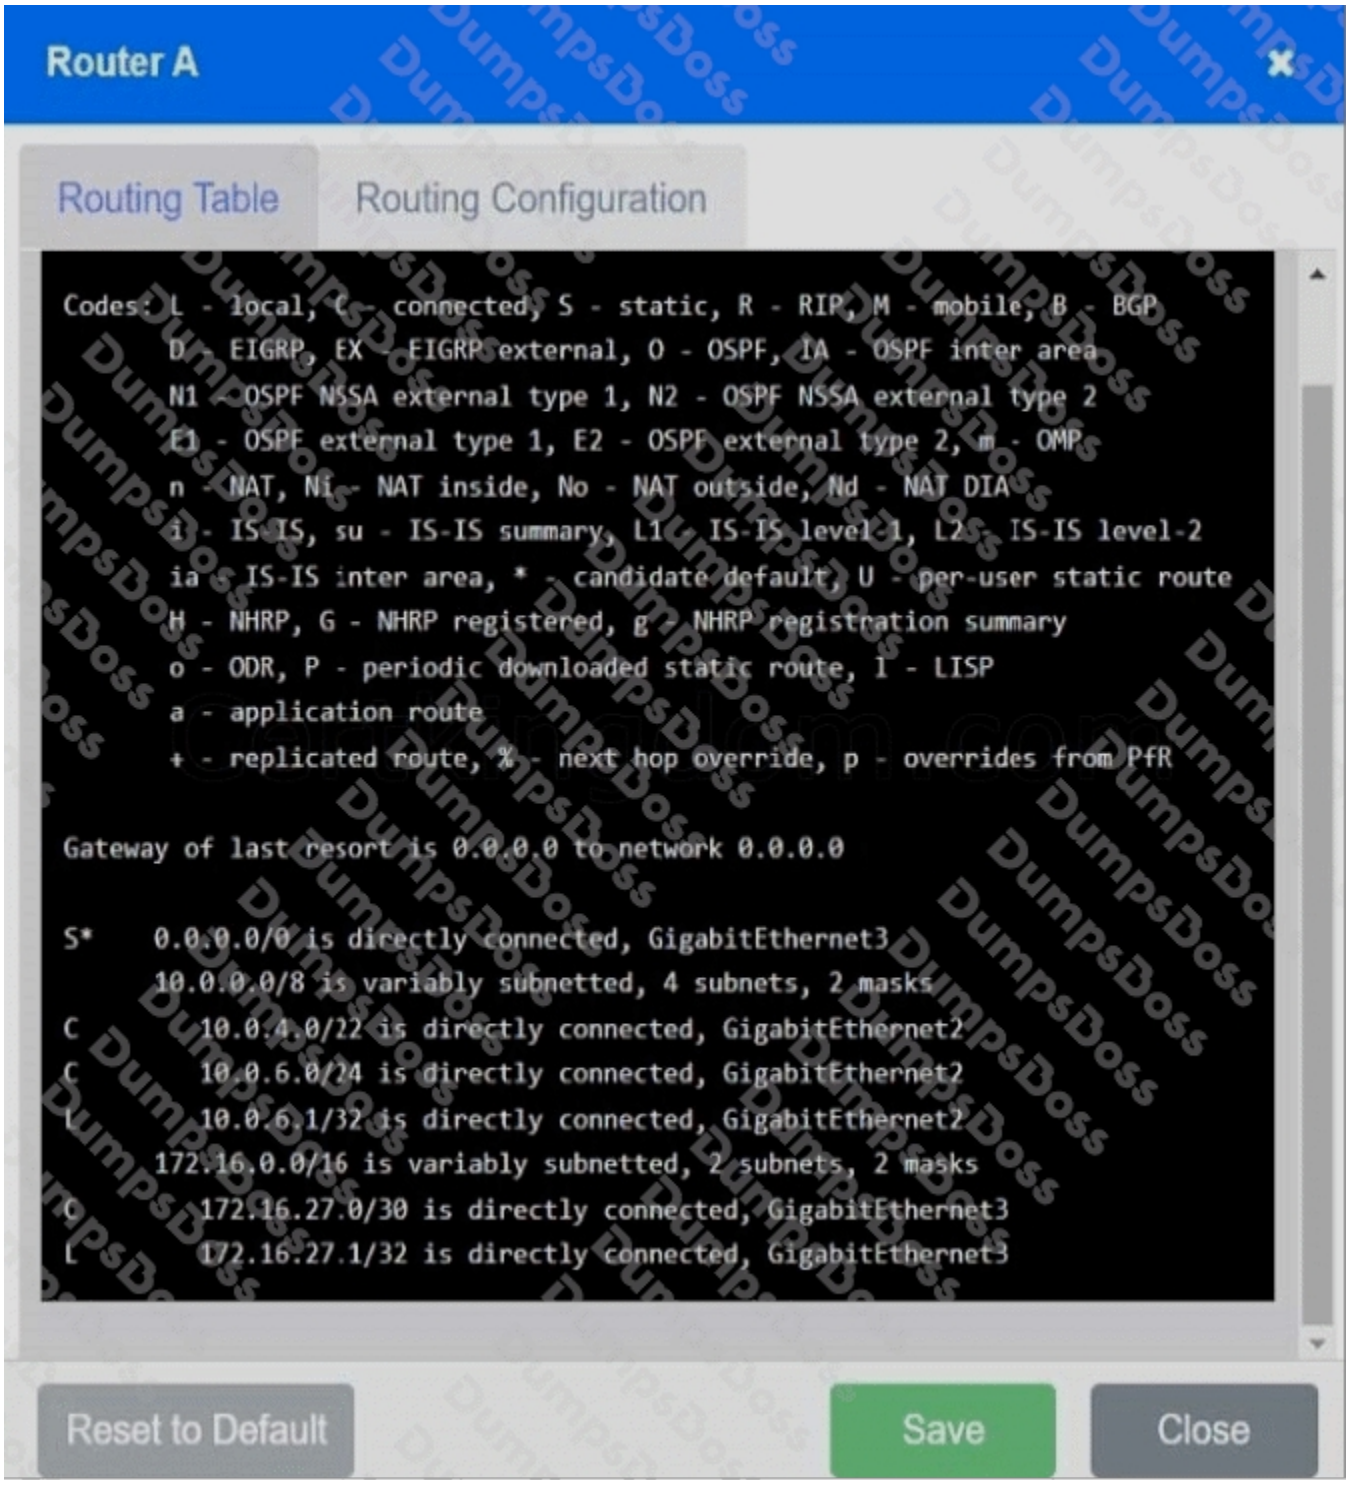
Router A
Routing Table
Routing Configuration
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
H - NHRP, G - NHRP registered, g - NHRP registration summary
o - ODR, P - periodic downloaded static route, l - LISP
a - application route
+ - replicated route, % - next hop override, p - overrides from PfR
Gateway of last resort is 0.0.0.0 to network 0.0.0.0
S* 0.0.0.0/0 is directly connected, GigabitEthernet3
10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C 10.0.4.0/22 is directly connected, GigabitEthernet2
C 10.0.6.0/24 is directly connected, GigabitEthernet2
L 10.0.6.1/32 is directly connected, GigabitEthernet2
172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C 172.16.27.0/30 is directly connected, GigabitEthernet3
L 172.16.27.1/32 is directly connected, GigabitEthernet3
Reset to Default
Save
Close


Router C
Routing Table
Routing Configuration
Router-C# show ip route
Codes: L - local, C - connected, S - static, 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, m - OMP
n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
H - NHRP, G - NHRP registered, g - NHRP registration summary
o - ODR, P - periodic downloaded static route, l - LISP
a - application route
+ - replicated route, % - next hop override, p - overrides from PfR
10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
S 10.0.0.0/22 [1/0] via GigabitEthernet1
S 10.0.4.0/22 [1/0] via GigabitEthernet2
172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C 172.16.27.0/30 is directly connected, GigabitEthernet2
L 172.16.27.2/32 is directly connected, GigabitEthernet2
C 172.16.27.4/30 is directly connected, GigabitEthernet1
L 172.16.27.6/32 is directly connected, GigabitEthernet1
Reset to Default
Save
Close

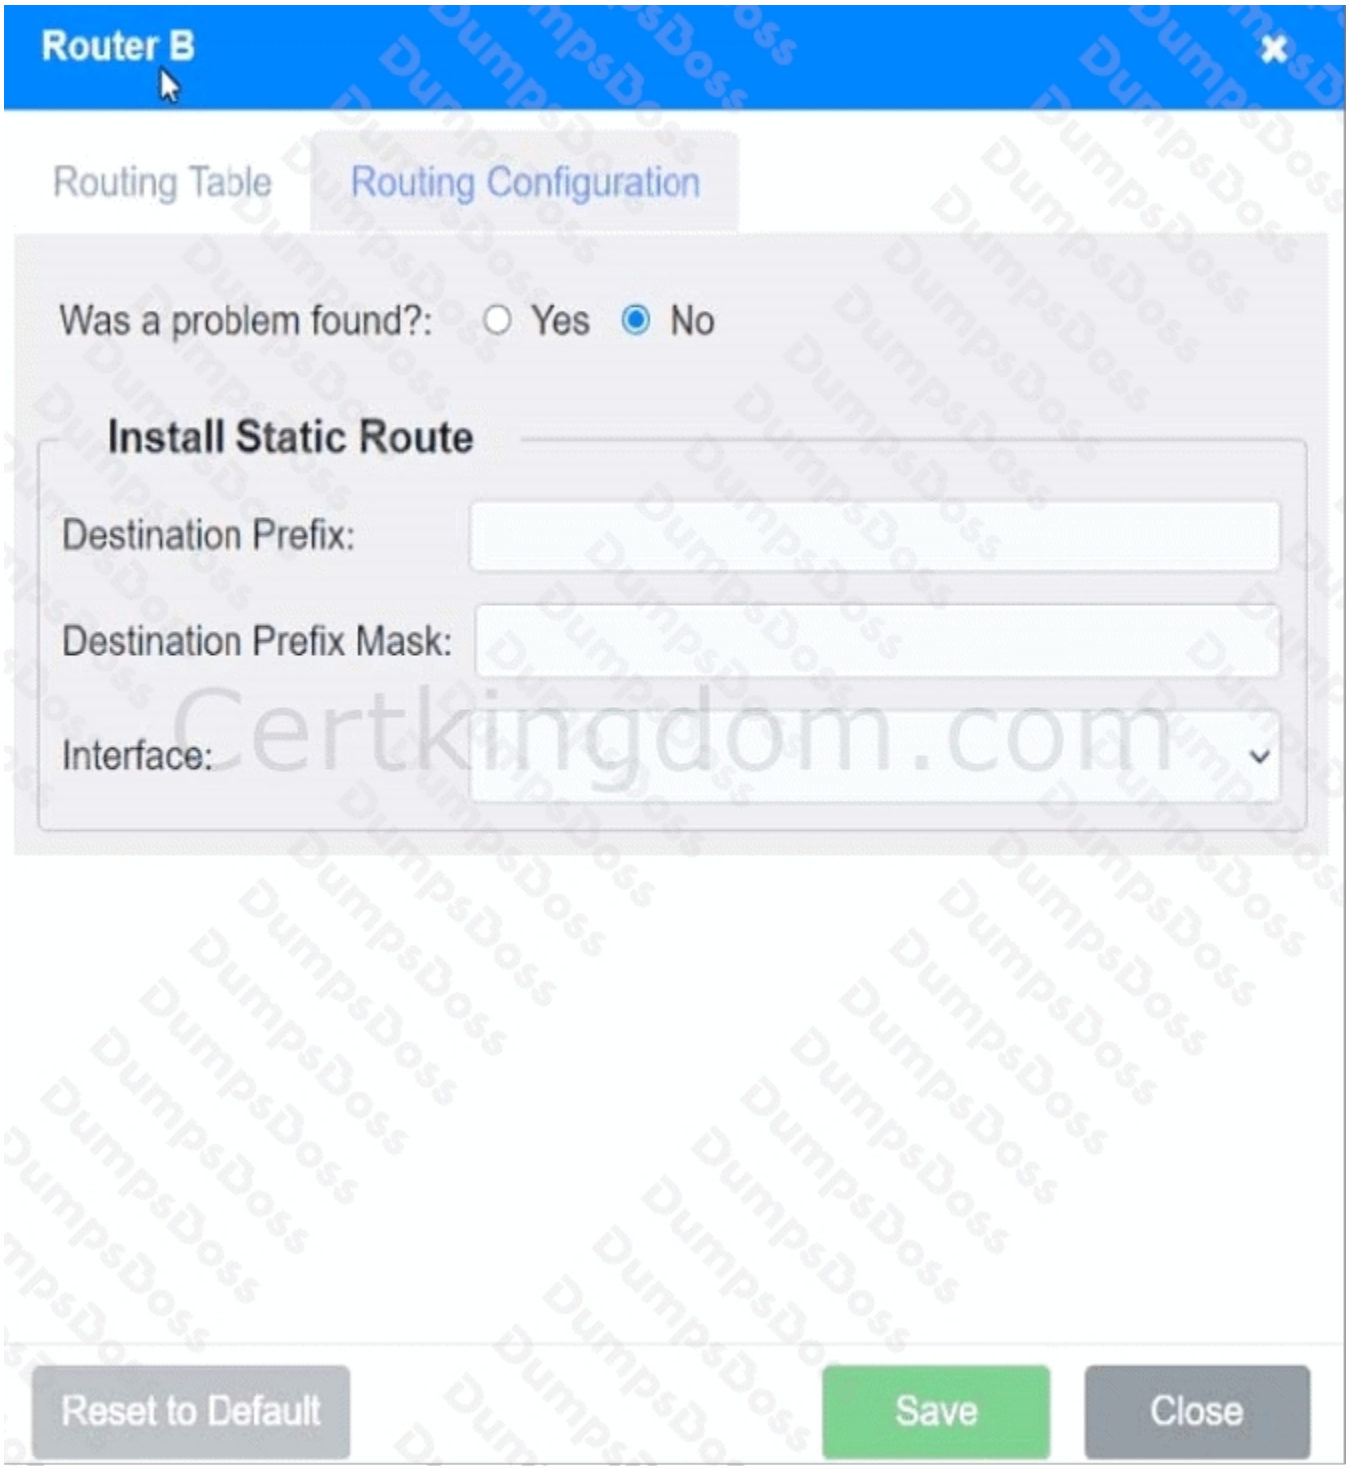

ANSWER: See the explanation for detailed informationon this simulation

Explanation:

Explanation.

To validate routing between networks hosting Workstation A and File Server 2, follow these steps:

Review Routing Tables:

Check the routing tables of Router A, Router B, and Router C to identify any missing routes.

Identify Missing Routes:

Ensure that each router has routes to the networks on which Workstation A and File Server 2 are located.

Add Static Routes:

If a route is missing, add a static route to the relevant destination network via the correct interface.

Routing Table:

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

S* 0.0.0.0/0 is directly connected, GigabitEthernet3

10.0.0.0 is variably subnetted, 4 subnets, 2 masks

C 10.0.4.0 is directly connected, GigabitEthernet2

C 10.0.6.0 is directly connected, GigabitEthernet2

L 10.0.6.1 is directly connected, GigabitEthernet2

172.16.0.0 is variably subnetted, 2 subnets, 2 masks

C 172.16.27.0 is directly connected, GigabitEthernet3 L 172.16.27.1 is directly connected, GigabitEthernet3 Routing Table:

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

S* 0.0.0.0/0 is directly connected, GigabitEthernet1

10.0.0.0 is variably subnetted, 4 subnets, 2 masks

C 10.0.0.0 is directly connected, GigabitEthernet1

L 10.0.0.1 is directly connected, GigabitEthernet1

172.16.0.0 is variably subnetted, 2 subnets, 2 masks

C 172.16.27.4 is directly connected, GigabitEthernet1 L 172.16.27.5 is directly connected, GigabitEthernet1 Routing Table:

10.0.0.0 is variably subnetted, 4 subnets, 2 masks

S 10.0.0.0 [1/0] via GigabitEthernet1

S 10.0.4.0 [1/0] via GigabitEthernet2

172.16.0.0 is variably subnetted, 2 subnets, 2 masks

C 172.16.27.0 is directly connected, GigabitEthernet2 L 172.16.27.2 is directly connected, GigabitEthernet2

C 172.16.27.4 is directly connected, GigabitEthernet1

L 172.16.27.6 is directly connected, GigabitEthernet1

Install Static Route to 10.0.0.0 via 172.16.27.1 (assuming Router C's IP is 172.16.27.1):

Destination Prefix: 10.0.0.0

Destination Prefix Mask: 255.255.252.0

Interface: GigabitEthernet3

Install Static Route to 10.0.4.0 via 172.16.27.5 (assuming Router C's IP is 172.16.27.5):

Destination Prefix: 10.0.4.0

Destination Prefix Mask: 255.255.252.0

Interface: GigabitEthernet1

Install Static Route to 10.0.6.0 via 172.16.27.2 (assuming Router A's IP is 172.16.27.2):

Destination Prefix: 10.0.6.0

Destination Prefix Mask: 255.255.255.0

Interface: GigabitEthernet2

Install Static Route to 10.0.0.0 via 172.16.27.1 (assuming Router B's IP is 172.16.27.1):

Destination Prefix: 10.0.0.0

Destination Prefix Mask: 255.255.252.0 Interface: GigabitEthernet1 Summary of Static Routes:

Router A:

ip route 10.0.0.0 255.255.252.0 GigabitEthernet3 Router B:

ip route 10.0.4.0 255.255.252.0 GigabitEthernet1 Router C:

ip route 10.0.6.0 255.255.255.0 GigabitEthernet2 ip route 10.0.0.0 255.255.252.0 GigabitEthernet1 These configurations ensure that each router knows the correct paths to reach Workstation A and File Server 2, resolving the connectivity issue.

QUESTION NO: 13

Which of the following are the best device-hardening techniques for network security? (Select two).

A. Disabling unused ports

B. Performing regular scanning of unauthorized devices

C. Monitoring system logs for irregularities

D. Enabling logical security such as SSO

E. Changing default passwords

F. Ensuring least privilege concepts are in place

ANSWER: A E

Explanation:

The two best “device hardening” moves here are disabling unused ports and changing default passwords. When you turn off ports and services you’re not using, you shrink the attack surface—there’s simply less for an attacker to probe or exploit. This is a common baseline step on switches, routers, firewalls, and even servers, because open-but-unused interfaces are just unnecessary risk.

Changing default passwords is just as important. Default credentials are easy to guess and often publicly documented, so leaving them in place is basically an open door. Swapping them for strong, unique passwords (and ideally using a password manager) blocks a huge chunk of opportunistic attacks.

The other options are good security practices, but they’re either broader operational controls (like log monitoring and scanning) or more about identity/access management (like SSO and least privilege) than classic “device hardening” basics. For quick, high-impact hardening on the device itself, A and E are the best fit.

References: https://www.cisa.gov/news-events/news/securing-network-infrastructure-devices and https://www.nist.gov/itl/smallbusinesscyber/guidance-topic/passwords

QUESTION NO: 14

A client wants to increase overall security after a recent breach. Which of the following would be best to implement? (Select two.)

A. Least privilege network access

B. Dynamic inventories

C. Central policy management

D. Zero-touch provisioning

E. Configuration drift prevention

F. Subnet range limits

ANSWER: A C

Explanation:

After a breach, two of the biggest wins usually come from tightening who can access what and making sure security settings are applied consistently everywhere. That’s why **least privilege network access** is a great move—it limits users and systems to only the permissions they truly need, which shrinks the “blast radius” if an account gets compromised again. It’s a core security principle and it directly reduces unnecessary exposure.

Central policy management is the other strong choice because it helps you push and enforce security rules (like authentication requirements, access rules, and baseline configs) across the whole environment from one place. That consistency matters a lot—breaches often happen because one system is misconfigured or didn’t get the same protections as the others. Central management also makes auditing and quick response much easier.

The other options can be useful, but they’re not as directly tied to improving overall security posture right after a breach. For example, zero-touch provisioning and dynamic inventories help operations, and drift prevention helps stability, but least privilege plus centralized policy control are the most broadly effective security improvements.

References: https://www.nist.gov/itl/smallbusinesscyber/guidance-topic/access-control and https://www.cisa.gov/resources-tools/resources/identity-and-access-management

QUESTION NO: 15

A user cannot access an external server for a client after connecting to a VPN. Which of the following commands would a support agent most likely use to examine the issue? (Select two).

A. nslookup

B. tcpdump

C. arp

D. dig

E. tracert

F. route print

ANSWER: E F

Explanation:

When someone connects to a VPN and suddenly can't reach an external server, a really common culprit is routing. The VPN may be forcing all traffic through the tunnel (full-tunnel) or the routes might not be set up correctly for split-tunneling. That's why checking the path and the routing table is usually the fastest way to narrow it down.

tracert helps you see where the traffic is actually going. If the trace dies right after the VPN gateway, or it starts taking an unexpected path, you've got a strong clue that the VPN route or upstream firewall rules are involved. It's a quick "where does it break?" tool.

route print shows the local routing table, which is perfect for confirming whether the default route changed when the VPN connected, or whether there's a more specific route that's hijacking traffic meant for that external server. This is often the smoking gun with VPN access problems.

References: https://learn.microsoft.com/en-us/windows-server/administration/windows-commands/tracert and https://learn.microsoft.com/en-us/windows-server/administration/windows-commands/route

QUESTION NO: 16

A company implements a video streaming solution that will play on all computers that have joined a particular group, but router ACLs are blocking the traffic. Which of the following is the most appropriate IP address that will be allowed in the ACL?

A. 127.0.0.1

B. 172.17.1.1

C. 224.0.0.1

D. 240.0.0.1

ANSWER: C

Explanation:

Since the video stream needs to reach “all computers that have joined a particular group,” that’s a classic use case for multicast. Multicast lets one sender transmit a single stream and have the network deliver it to multiple receivers that subscribed to the multicast group, which is way more efficient than sending separate unicast streams to every PC.

The address 224.0.0.1 is a well-known multicast address (in the 224.0.0.0–239.255.255.255 range). In practice, you’d allow the specific multicast group address your streaming app uses, but among these choices, 224.0.0.1 is the only valid multicast option, so it’s the best ACL match.

The others don’t fit: 127.0.0.1 is loopback (it never leaves the local machine), 172.17.1.1 is just a private unicast host address, and 240.0.0.1 is in a reserved/experimental range that isn’t used for normal multicast streaming.

References: https://www.iana.org/assignments/multicast-addresses/multicast-addresses.xhtml and https://datatracker.ietf.org/doc/html/rfc1112

QUESTION NO: 17

After installing a new wireless access point, an engineer tests the device and sees that it is not performing at the rated speeds. Which of the following should the engineer do to troubleshoot the issue? (Select two.)

A. Ensure a bottleneck is not coming from other devices on the network.

B. Install the latest firmware for the device.

C. Create a new VLAN for the access point.

D. Make sure the SSID is not longer than 16 characters.

E. Configure the AP in autonomous mode.

F. Install a wireless LAN controller.

ANSWER: A B

Explanation:

If a brand-new AP isn’t hitting its advertised speeds, the first thing I’d check is whether the slowdown is actually somewhere else. A fast AP can still feel “slow” if the uplink is limited (like a 100Mb switch port), the cable is bad, the switch is overloaded, or the internet connection is the real bottleneck. So validating there’s no bottleneck on the wired side (switchport speed/duplex, cabling, upstream throughput) is a smart early step.

Next, make sure the AP is on current firmware. Vendors regularly fix performance bugs, driver issues, and radio stability problems in firmware updates. An AP can absolutely underperform if it shipped with older code, especially with newer client devices and newer Wi-Fi features.

The other choices don’t really target “rated speed” issues. VLANs and SSID length won’t change throughput in any meaningful way, and switching modes or adding a controller is more of an architecture decision than a basic performance troubleshooting step.

References: https://www.cisco.com/c/en/us/support/docs/wireless-mobility/wireless-lan-wlan/200132-Wireless-LAN-Controller-and-Access.html and https://www.tp-link.com/us/support/faq/2292/

QUESTION NO: 18

Which of the following can also provide a security feature when implemented?

A. NAT

B. BGP

C. FHRP

D. EIGRP

ANSWER: A

Explanation:

NAT is the best answer here because it can act like a small “shield” for your internal network. With NAT (especially PAT/overload), devices on your LAN use private IP addresses that aren’t directly reachable from the internet. From the outside, it looks like traffic is coming from the NAT device’s public IP, so your internal addressing scheme is hidden.

It’s not a replacement for a firewall, but it does reduce exposure by making unsolicited inbound connections harder unless you intentionally set up port forwarding or static NAT rules. That’s why NAT is often mentioned as providing a security benefit as a side effect of how it works.

The other choices (BGP, FHRP, EIGRP) are routing or high-availability technologies. They’re great for path selection and uptime, but they don’t inherently “hide” hosts or add a security layer the way NAT can.

References: https://en.wikipedia.org/wiki/Network_address_translation and https://www.cisco.com/c/en/us/support/docs/ip/network-address-translation-nat/13772-12.html

QUESTION NO: 19

Following a fire in a data center, the cabling was replaced. Soon after, an administrator notices network issues. Which of the following are the most likely causes of the network issues? (Select two).

A. The switches are not the correct voltage.

B. The HVAC system was not verified as fully functional after the fire.

C. The VLAN database was not deleted before the equipment was brought back online.

D. The RJ45 cables were replaced with unshielded cables.

E. The wrong transceiver type was used for the new termination.

F. The new RJ45 cables are a higher category than the old ones.

ANSWER: D E

Explanation:

After a fire, it’s common to re-pull cabling quickly, and that’s where mistakes creep in. If the RJ45 runs were replaced with unshielded twisted pair (UTP) when the environment really needs shielding, you can end up with lots of EMI/RFI noise from

power gear, HVAC motors, or other equipment. That kind of interference usually shows up as flaky links, retransmits, and random drops—especially in a busy data center.

The other big “gotcha” is using the wrong transceiver for the fiber/copper termination. SFP/SFP+/QSFP modules have to match the speed, wavelength, and fiber type (single-mode vs. multi-mode). If the module doesn’t match the cabling or the port expectations, the link may not come up at all, or it may negotiate incorrectly and perform poorly.

For quick reference on transceivers and matching optics, see https://en.wikipedia.org/wiki/Small_Form-factor_Pluggable_transceiver. For a straightforward overview of shielding and why it matters in noisy environments, see https://en.wikipedia.org/wiki/Shielded_cable.

QUESTION NO: 20

After installing a new 6E wireless router in a small office, a technician notices that some wireless devices are not able to achieve the rated speeds.

Which of the following should the technician check to troubleshoot the issue? (Select two)

A. Client device compatibility

B. Back-end cabling

C. Weather phenomena

D. Voltage source requirements

E. Interference levels

F. Processing power

ANSWER: A E

Explanation:

Two big things usually explain why “Wi-Fi 6E speeds” don’t show up on some devices: the device may not actually support Wi-Fi 6E features, or the wireless environment may be noisy.

First, check **client device compatibility**. Wi-Fi 6E means 6 GHz support, and plenty of laptops/phones are only Wi-Fi 5 or Wi-Fi 6 (2.4/5 GHz). If a device can’t use 6 GHz (or doesn’t support wider channels like 160 MHz), it’ll connect on older bands and top out at lower speeds. A quick look at the client’s wireless adapter specs and what band/channel width it negotiated will usually confirm this. See https://www.wi-fi.org/discover-wi-fi/wi-fi-6e

Second, check **interference levels**. Even with 6 GHz being cleaner, devices might still be connecting on 2.4/5 GHz where neighboring networks, Bluetooth, microwaves, and bad channel choices can crush throughput. A site survey or Wi-Fi analyzer can reveal congestion, overlap, or poor SNR that forces slower rates. Reference: https://www.cisco.com/c/en/us/products/wireless/what-is-wi-fi-6e.html