DUMPSBOSS.

CompTIA Security+ Exam

CompTIA SY0-701

Version Demo

Total Demo Questions: 119

Total Premium Questions: 1192

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, General Security Concepts	124
Topic 2, Threats, Vulnerabilities, and Mitigations	244
Topic 3, Security Architecture	290
Topic 4, Security Operations	274
Topic 5, Security Program Management and Oversight	260
Total	1192

QUESTION NO: 1

Which of the following must be considered when designing a high-availability network? (Choose two).

A. Ease of recovery

B. Ability to patch

C. Physical isolation

D. Responsiveness

E. Attack surface

F. Extensible authentication

ANSWER: A D

Explanation:

Ease of recovery is a fundamental high-availability consideration because availability depends on restoring normal service quickly after a component, link, configuration, or site failure. A resilient network design should support well-documented recovery procedures, tested backups, redundant paths and devices, automated or straightforward failover, and manageable replacement of failed equipment. These measures reduce mean time to recover and therefore minimize the business impact of an outage.

Responsiveness must also be considered because a network cannot provide meaningful availability if it takes too long to detect a failure, converge on an alternate route, fail over to a standby component, or accommodate changing demand. Health monitoring, routing convergence behavior, load balancing, capacity planning, and alerting all influence how promptly the network reacts while maintaining usable service for clients. High availability is therefore concerned both with continued operation and with timely recovery or adaptation when conditions change.

This aligns with resiliency guidance that emphasizes designing for failure, detecting faults, responding effectively, and recovering services within defined objectives. See the Microsoft Azure Well-Architected Framework Reliability guidance and the AWS Well-Architected Reliability Pillar.

QUESTION NO: 2

A company is developing a critical system for the government and storing project information on a fileshare. Which of the following describes how this data will most likely be classified? (Choose two.)

A. Private

B. Confidential

C. Public

D. Operational

E. Urgent

F. Restricted

ANSWER: B F

Explanation:

Confidential is appropriate because project information for a critical government system is intended for authorized personnel with a legitimate need to know. Design artifacts, technical requirements, implementation details, and test information can reveal capabilities, security architecture, or operational dependencies. A disclosure could adversely affect the government customer, the contractor, or the security of the system. Confidential handling therefore supports access controls, approved sharing, and protection while the information is stored on the file share.

Restricted is also appropriate because the information requires a particularly limited distribution and stronger handling controls. Within an organizational classification scheme, restricted data is typically accessible only to specifically authorized project members or roles. This classification supports enforcing least privilege, using access-control lists or groups, maintaining audit records, and preventing unauthorized copying or disclosure. NIST describes controlled unclassified information as information that requires safeguarding or dissemination controls pursuant to applicable requirements, which aligns with the need to tightly control sensitive government-project material. See NIST's definition of Controlled Unclassified Information and NIST SP 800-171 for safeguarding requirements for sensitive information in nonfederal systems.

QUESTION NO: 3

A new corporate policy requires all staff to use multifactor authentication to access company resources. Which of the following can be utilized to set up this form of identity and access management? (Choose two.)

A. Authentication tokens

B. Least privilege

C. Biometrics

D. LDAP

E. Password vaulting

F. SAML

ANSWER: A C

Explanation:

Authentication tokens and biometrics can be used to implement multifactor authentication because they represent distinct authentication factor categories. An authentication token is an ownership factor: it is something the user has. Examples include a hardware security key, a smart card, a one-time-password token, or an authenticator application on a registered device that generates or approves a sign-in code. This factor can be combined with a password or PIN to require more than one type of evidence before access is granted.

Biometrics provide an inherence factor: something the user is. A fingerprint, facial-recognition scan, or iris scan can verify the identity of the person attempting to authenticate. When biometric verification is used with a separate knowledge or possession factor, such as a password and a security key, the result is MFA. Effective MFA deployments use factors from different categories rather than simply requesting multiple passwords or multiple possession-based credentials. NIST describes MFA as authentication that uses more than one distinct factor and identifies cryptographic devices and biometrics as common mechanisms used in authentication systems. See NIST SP 800-63B: Digital Identity Guidelines and NIST guidance on multi-factor authentication.

QUESTION NO: 4

A security administrator is reissuing a former employee's laptop. Which of the following is the best combination of data handling activities for the administrator to perform? (Choose two.)

A. Data retention

B. Certification

C. Destruction

D. Classification

E. Sanitization

F. Enumeration

ANSWER: B E

Explanation:

Sanitization and **Certification** are the appropriate combination when a laptop is being reassigned to another employee. Sanitization is the process of rendering data on the laptop's storage media inaccessible and unrecoverable at the required protection level. Depending on the device and organizational policy, this can involve cryptographic erase, secure overwrite, or another approved media-sanitization technique. The administrator must ensure that former-user files, browser data, locally cached credentials, encryption keys, application data, and operating-system artifacts are not exposed to the new recipient.

Certification provides the documented confirmation that the required sanitization process was completed successfully. This creates accountability and an auditable record that identifies the asset or media, the method used, the responsible party, and the result. Maintaining this evidence supports organizational data-handling policies, asset lifecycle controls, compliance requirements, and incident investigations. NIST describes sanitization as a process that includes verification and validation activities, with documentation supporting assurance that the selected sanitization method achieved the intended outcome. See NIST SP 800-88 Rev. 2, Guidelines for Media Sanitization and the CompTIA Security+ exam objectives.

QUESTION NO: 5

Visitors to a secured facility are required to check in with a photo ID and enter the facility through an access control vestibule Which of the following but describes this form of security control?

A. Physical

B. Managerial

C. Technical

D. Operational

ANSWER: A

Explanation:

Physical is correct because the described safeguards directly protect entry to a real-world location through barriers and in-person identity verification. An access control vestibule, commonly called a mantrap, is a physical access-control mechanism with interlocking doors. It creates a controlled space in which one door must close before the next door can open, helping ensure that each person entering is individually authorized. This supports prevention of unauthorized entry through tailgating or piggybacking and permits security staff or automated systems to validate access before allowing a visitor farther into the facility.

The photo-identification check also supports physical security because it is performed as part of controlling a person's access to a building. Together, the identity check and vestibule establish a layered facility-entry process: identify the visitor, validate that entry is permitted, and physically restrict movement until the access decision is complete. NIST addresses this type of protection under physical and environmental safeguards, including controlling physical access to systems and facilities. See NIST SP 800-53, Physical and Environmental Protection controls and CompTIA's Security+ exam objectives resources for the control-category context.

QUESTION NO: 6

Which of the following prevents unauthorized modifications to internal processes, assets, and security controls?

A. Change management

B. Playbooks

C. Incident response

D. Acceptable use policy

ANSWER: A

Explanation:

Change management is correct because it is the structured governance process used to request, assess, authorize, implement, document, and review changes to an organization's systems and operational environment. Requiring formal approval and maintaining records of proposed and completed changes prevents personnel from making unreviewed modifications to internal processes, technology assets, configurations, and security controls. This preserves the integrity of the environment, supports accountability, and helps ensure that security implications, dependencies, testing requirements, rollback procedures, and business impact are evaluated before a change is deployed. Effective change management also establishes separation of duties and an audit trail, allowing an organization to verify who requested, approved, and performed a change. These practices reduce the chance that an unauthorized or poorly planned modification introduces vulnerabilities, disrupts critical services, or weakens an established control. Change management is therefore a core security governance and operational control that supports controlled, secure evolution of organizational systems. See the NIST Cybersecurity Framework crosswalk for governance and change-related control context, and NIST SP 800-128, Guide for Security-Focused Configuration Management for guidance on managing authorized changes to information systems.

QUESTION NO: 7

An organization experiences a cybersecurity incident involving a command-and-control server. Which of the following logs should be analyzed to identify the impacted host? (Choose two.)

A. Application

B. Authentication

C. DHCP

D. Network

E. Firewall

F. Database

ANSWER: C E

Explanation:

DHCP and **Firewall** logs should be analyzed together to identify the impacted host associated with command-and-control activity. Firewall logs record connection attempts and permitted or denied sessions at the network boundary. When investigators identify the command-and-control server's IP address, domain, port, or protocol, firewall entries can reveal the internal source IP address that communicated with that infrastructure, along with useful timing and connection details. This establishes the network identity involved in the suspected beaconing or outbound command-and-control session.

DHCP logs provide the essential attribution step for environments where endpoints receive addresses dynamically. They correlate the internal IP address observed in the firewall events to a leased MAC address, hostname, and lease time. By matching the firewall event timestamp to the relevant DHCP lease, responders can determine the actual workstation, server, or other endpoint assigned that address at the time of the activity. This correlation is particularly important because an IP address may have been reassigned after the incident. NIST guidance recognizes firewall and DHCP records as valuable sources for incident investigation and event correlation. See NIST SP 800-61r2 and NIST SP 800-92.

QUESTION NO: 8

Which of the following definitions best describes the concept of log correlation?

A. Combining relevant logs from multiple sources into one location

B. Searching and processing data to identify patterns of malicious activity

C. Making a record of the events that occur in the system

D. Analyzing the log files of the system components

ANSWER: B

Explanation:

Searching and processing data to identify patterns of malicious activity best describes log correlation. Correlation applies logic, rules, analytics, or queries to events so that related activity can be evaluated together rather than as isolated entries. For example, an analyst or SIEM may associate authentication events, endpoint alerts, network connections, and account changes by shared attributes such as timestamps, usernames, hostnames, IP addresses, process identifiers, or session information. This combined context can reveal a security-relevant sequence, such as repeated failed authentications followed by a successful sign-in and unusual access to sensitive resources.

In Security+ terms, correlation is an important capability of security monitoring and SIEM technologies because it turns large volumes of raw event data into actionable detections and alerts. The objective is to recognize relationships and patterns that can indicate malicious activity, policy violations, or other incidents. Effective correlation also helps security teams prioritize investigation by supplying the context needed to understand the scope and timeline of an event. CISA discusses SIEM use for collecting and analyzing security-relevant event information in its SIEM Use Case Compendium. For further background on log management and analysis practices, see the NIST Guide to Computer Security Log Management.

QUESTION NO: 9 - (SIMULATION)

SIMULATION

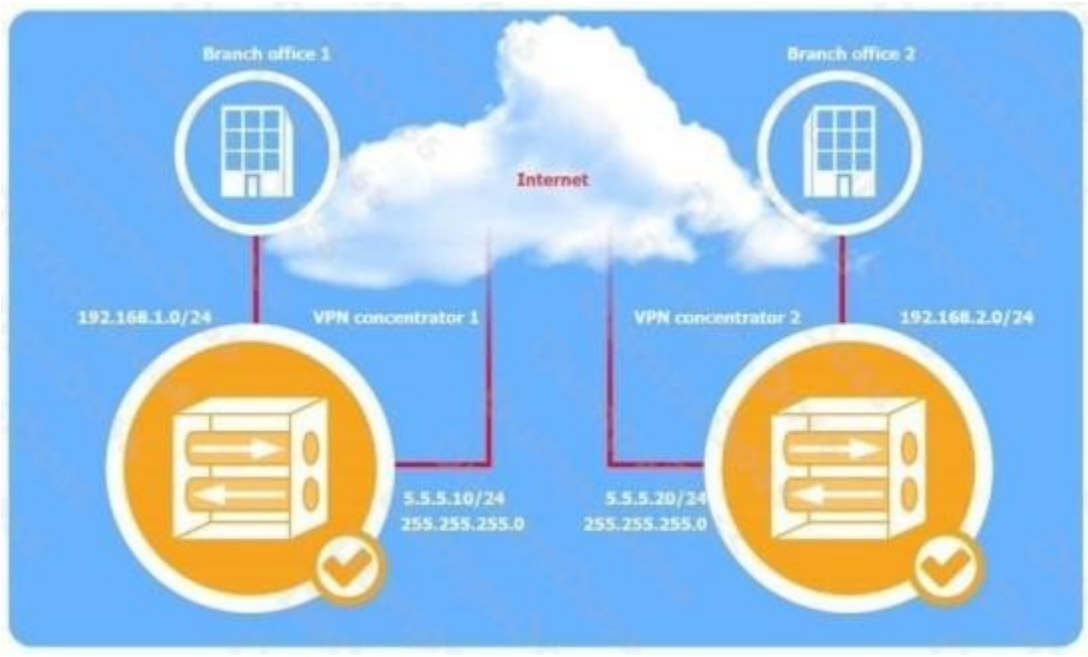
A systems administrator is configuring a site-to-site VPN between two branch offices. Some of the settings have already been configured correctly. The systems administrator has been provided the following requirements as part of completing the configuration:

• Most secure algorithms should be selected

• All traffic should be encrypted over the VPN

• A secret password will be used to authenticate the two VPN concentrators INSTRUCTIONS

Click on the two VPN Concentrators to configure the appropriate settings.



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


VPN Concentrator 1
Phase 1
Phase 2
Peer IP address:
Auth method:
Select
PKI
PSK
RADIUS
Negotiation mode:
MAIN
Encryption algorithm:
Select
AES256
ECC secp160r1
3DES
Hash algorithm:
Select
SHA256
MD5
SHA1
DH key group:
14
Reset to Default
Save
Close


VPN Concentrator 1
Phase 1
Phase 2
Mode:
Tunnel
Protocol:
Select
ESP
AH
Encryption algorithm:
Select
3DES
AES256
BLOWFISH
Hash algorithm:
Select
SHA256
MD5
SHA1
Local network/mask:
Remote network/mask:
Reset to Default
Save
Close


VPN Concentrator 2
Phase 1
Phase 2
Peer IP address:
Auth method:
Select
PKI
RADIUS
PSK
Negotiation mode:
MAIN
Encryption algorithm:
Select
3DES
AES256
ECC secp160r1
Hash algorithm:
Select
SHA256
SHA1
MD5
DH key group:
14
Reset to Default
Save
Close



A. See Explanation section for answer.

ANSWER: See the explanation for the answer

Explanation:

Configure both concentrators with matching Phase 1 and Phase 2 cryptographic settings. Use **PSK** because the concentrators will authenticate using a shared secret password.

VPN Concentrator 1

Phase 1
Peer IP address: `5.5.5.20`
Auth method: `PSK`
Negotiation mode: `MAIN`
Encryption algorithm: `AES256`
Hash algorithm: `SHA256`
DH key group: `14`
Phase 2
Mode: `Tunnel`
Protocol: `ESP`
Encryption algorithm: `AES256`
Hash algorithm: `SHA256`

Local network/mask: `192.168.1.0/24`
Remote network/mask: `192.168.2.0/24`

VPN Concentrator 2

Phase 1
Peer IP address: `5.5.5.10`
Auth method: `PSK`
Negotiation mode: `MAIN`
Encryption algorithm: `AES256`
Hash algorithm: `SHA256`
DH key group: `14`
Phase 2
Mode: `Tunnel`
Protocol: `ESP`
Encryption algorithm: `AES256`
Hash algorithm: `SHA256`
Local network/mask: `192.168.2.0/24`
Remote network/mask: `192.168.1.0/24`

Why: AES-256 and SHA-256 are the strongest available listed choices. ESP provides encryption for VPN traffic; AH provides integrity/authentication only and does not encrypt. Tunnel mode encrypts traffic between the two private branch networks.

QUESTION NO: 10

During the onboarding process, an employee needs to create a password for an intranet account. The password must include ten characters, numbers, and letters, and two special characters. Oncethe password is created, the company will grant the employee access to other company-owned websites based on the intranet profile. Which of the following access management concepts is the company most likely using to safeguard intranet accounts and grant access to multiple sites based on a user's intranet account? (Select two).

A. Federation

B. Identity proofing

C. Password complexity

D. Default password changes

E. Password manager

F. Open authentication

ANSWER: A C

Explanation:

Federation is correct because the organization is using the employee's established intranet identity to provide access to multiple company-owned web resources. Federated identity management enables separate applications or services to rely on a trusted identity provider's authentication and identity assertions, rather than maintaining distinct credentials and authentication processes for every site. This supports a streamlined access experience while allowing the organization to apply access decisions based on the user's intranet profile, such as assigned role, department, or authorization attributes. Microsoft describes federation as establishing trust between identity systems so that applications can use authentication performed by a trusted identity provider: Microsoft Learn: Federated identity.

Password complexity is also correct because the stated password rule explicitly requires a minimum length and a combination of character categories, including numbers, letters, and special characters. Those composition requirements are a password-complexity control intended to make passwords more resistant to guessing and automated password-cracking attempts. Password controls are part of account security and should be implemented with additional protections such as multi-factor authentication and secure password storage. NIST provides digital identity guidance for authentication and memorized secrets in NIST SP 800-63B.

QUESTION NO: 11

During a SQL update of a database, a temporary field that was created was replaced by an attacker in order to allow access to the system. Which of the following best describes this type of vulnerability?

A. Race condition

B. Memory injection

C. Malicious update

D. Side loading

ANSWER: A

Explanation:

Race condition is correct because the described attack relies on an attacker altering a temporary resource during a narrow timing window in an operation. This is commonly associated with a time-of-check to time-of-use (TOCTOU) weakness: an application checks, creates, or otherwise prepares a temporary object, but before the application completes its intended use of that object, an attacker replaces or modifies it. If the application does not verify that the object remains the expected one at the point of use, the substituted resource may be processed with the application's permissions, potentially enabling unauthorized system access.

Race conditions occur when the security or correctness of an operation depends on the order or timing of concurrent events. Temporary files, fields, objects, symbolic links, and shared database resources can be vulnerable when they are predictable, accessible to untrusted users, or not protected by appropriate locking, atomic operations, and ownership or permission checks. Secure designs reduce exposure by using securely generated temporary resources, restrictive permissions, atomic creation and replacement operations, and validation immediately before use. MITRE classifies this issue as CWE-362, Concurrent Execution using Shared Resource with Improper Synchronization ("Race Condition"): MITRE CWE-362. OWASP also describes TOCTOU as a timing-based attack in which a resource changes between validation and use: OWASP TOCTOU.

QUESTION NO: 12

Which of the following security threats aims to compromise a website that multiple employees frequently visit?

A. Supply chain

B. Typosquatting

C. Watering hole

D. Impersonation

ANSWER: C

Explanation:

Watering hole is correct because this attack targets a group by compromising a legitimate website that members of that group are known to visit regularly. Rather than trying to reach every employee directly, an attacker identifies an online location trusted by the intended victims—such as an industry news site, professional forum, partner portal, or other frequently used web resource—and injects malicious code or redirects visitors to attacker-controlled content. When employees visit the compromised site as part of their normal activity, the attacker may attempt browser exploitation, malware delivery, credential theft, or other unauthorized access. This approach is especially effective in targeted campaigns because the attacker selects the site based on the browsing habits of a particular organization, role, or community. The defining characteristic in the question is the compromise of an existing, commonly visited website in order to reach its regular visitors, which directly describes a watering hole attack. CompTIA Security+ includes watering hole attacks among common threat techniques that security professionals should recognize when assessing web-based risks and user exposure. See the CISA guidance on social engineering and phishing attacks and the MITRE ATT&CK Drive-by Compromise technique for related technical context.

QUESTION NO: 13

Which of the following are common VoIP-associated vulnerabilities? (Choose two.)

A. SPIM

B. Vishing

C. Hopping

D. Phishing

E. Credential harvesting

F. Tailgating

ANSWER: A B

Explanation:

SPIM and vishing are common threats associated with Voice over IP environments. SPIM, or spam over Internet messaging, involves sending unsolicited bulk calls, instant messages, or other communications through Internet-based telephony and messaging services. Beyond being disruptive, these campaigns can deliver fraudulent messages, direct recipients to malicious websites, or establish contact for later social-engineering activity. Because VoIP platforms use Internet protocols and can support automated, high-volume communications, they can be abused to distribute SPIM at scale.

Vishing is voice phishing: an attacker uses telephone calls, including calls delivered through VoIP services, to persuade a target to reveal sensitive information or perform an unsafe action. VoIP makes large-scale vishing campaigns more accessible by enabling automated dialing and supporting caller-ID spoofing techniques. A caller may impersonate a bank, help desk, executive, or government agency to request credentials, one-time passcodes, payment information, or remote access. The FCC describes caller-ID spoofing and related scam activity at https://www.fcc.gov/spoofing, while CISA provides guidance on recognizing phishing and social-engineering tactics at https://www.cisa.gov/news-events/news/avoiding-social-engineering-and-phishing-attacks.

QUESTION NO: 14

Which of the following activities are associated with vulnerability management? (Select two).

A. Reporting

B. Prioritization

C. Exploiting

D. Correlation

E. Containment

F. Tabletop exercise

ANSWER: A B

Explanation:

Reporting and **Prioritization** are core activities in a vulnerability-management lifecycle. Reporting communicates discovered vulnerabilities, affected assets, remediation status, trends, and outstanding risk to technical teams and stakeholders. Useful reports establish accountability, document risk acceptance or exceptions, and show whether remediation efforts are reducing organizational exposure over time.

Prioritization determines the order in which identified vulnerabilities should be addressed. Security teams should not rely solely on a severity score; they should also consider asset criticality, business impact, whether the vulnerability is exposed to attack, the availability of public exploits, evidence of active exploitation, and compensating controls. This risk-based approach directs limited remediation resources toward the weaknesses most likely to cause meaningful harm. After

prioritized remediation occurs, teams typically validate that the issue was resolved and continue monitoring for newly discovered vulnerabilities. NIST describes enterprise patch and vulnerability response as a risk-based process involving identification, prioritization, mitigation, and ongoing measurement. CompTIA's Security+ objectives likewise include vulnerability-management activities such as scanning, prioritization, remediation, validation, and reporting. See NIST SP 800-40 Rev. 4 and CompTIA exam objectives.

QUESTION NO: 15

A security manager is implementing MFA and patch management. Which of the following would best describe the control type and category? (Select two).

A. Physical

B. Managerial

C. Detective

D. Administrator

E. Preventative

F. Technical

ANSWER: E F

Explanation:

"Preventative" and "Technical" best describe the intended purpose and implementation of these safeguards. MFA is a technical access-control mechanism: an authentication system enforces a requirement for multiple factors before granting access. Its security purpose is preventive because it stops or reduces the likelihood of unauthorized access before an account, service, or resource is compromised.

Patch management likewise has a preventive objective. Applying vendor security updates removes or mitigates known vulnerabilities before an attacker can exploit them, reducing the attack surface and likelihood of a security incident. In the context of this question, the patching mechanism is treated as a technical control because software, firmware, and system-management technologies implement and enforce the updates on affected assets. Together, MFA and patch management are examples of security controls used proactively to protect systems and data rather than merely identifying an event after it has occurred.

CompTIA's Security+ objectives include control categories and control types as core security concepts; see CompTIA exam objectives. NIST also describes multifactor authentication and flaw remediation as control activities supporting protection against unauthorized access and exploitable weaknesses: NIST SP 800-53 Rev. 5.

QUESTION NO: 16

Which of the following is the first step to take when creating an anomaly detection process?

A. Selecting events

B. Building a baseline

C. Selecting logging options

D. Creating an event log

ANSWER: B

Explanation:

Building a baseline is the first step because anomaly detection depends on having an established, measurable understanding of normal activity. A baseline captures expected patterns for a particular environment, such as typical

authentication times, network throughput, system resource use, DNS requests, application behavior, and the normal volume of security events. Once this expected behavior has been documented over a representative period, security personnel and monitoring tools can identify meaningful deviations from it. For example, a sudden increase in outbound traffic, logons at an unusual time, or an account accessing systems it does not normally use can be compared with the baseline and investigated as a potential anomaly. Baselines must also be reviewed and adjusted when legitimate business operations, infrastructure, applications, or user populations change; otherwise, detection systems may create excessive alerts or fail to recognize newly abnormal behavior. This approach aligns with the Security+ security-operations objective of using monitoring and analysis to identify indicators of malicious activity. NIST guidance similarly describes establishing normal system and network behavior as essential context for detecting security events. See the CompTIA exam objectives resource and NIST SP 800-94, Guide to Intrusion Detection and Prevention Systems.

QUESTION NO: 17

A security analyst reviews firewall configurations and finds that firewalls are configured to fail-open mode in the event of a crash. Which of the following describes the security risk associated with this configuration?

A. There may be increased latency during failover.

B. Authentication tokens may be invalidated during an outage.

C. Traffic will bypass inspection during a failure.

D. All encrypted traffic will be blocked during an outage.

ANSWER: C

Explanation:

Traffic will bypass inspection during a failure is correct because fail-open behavior prioritizes availability when a firewall, inline security appliance, or its inspection function crashes. Rather than interrupting network connectivity, the device permits traffic to continue traversing the network path. The key security consequence is that the failed firewall can no longer reliably enforce its policy, such as packet filtering, stateful inspection, application-layer controls, intrusion-prevention inspection, or blocking rules. Consequently, traffic that would normally be denied or scrutinized can reach internal resources while the security control is unavailable. This creates a temporary gap in the organization's defensive layers and may allow unauthorized connections, malicious payloads, command-and-control communications, or prohibited protocols to pass without the expected inspection. The configuration therefore represents an availability-versus-security tradeoff: fail-open minimizes disruption to business communications, but it reduces protection during the failure condition. Security+ objectives cover the importance of resilient architectures and the impact of availability decisions on security controls. NIST's contingency-planning guidance also emphasizes assessing how system failures affect security safeguards and maintaining appropriate alternate protections. See NIST SP 800-34 Rev. 1 and CompTIA's Security+ certification page.

QUESTION NO: 18 - (SIMULATION)

A recent black-box penetration test of http://example.com discovered that external

website vulnerabilities exist, such as directory traversals, cross-site scripting, cross-site forgery, and insecure protocols.

You are tasked with reducing the attack space and enabling secure protocols.

INSTRUCTIONS

Part 1

Use the drop-down menus to select the appropriate technologies for each location to implement a secure and resilient web architecture. Not all technologies will be used, and technologies may be used multiple times.

Part 2

Use the drop-down menus to select the appropriate command snippets from the drop-down menus. Each command section must be filled.

ANSWER: See the explanation for the answer

Explanation:

Part 1 — secure/resilient placement:

Place an **edge firewall** at the Internet-facing boundary. Permit only required public web traffic (TCP 443; TCP 80 only if needed to redirect clients to HTTPS).
Place a **WAF** in front of the web application/web-server tier (commonly behind the edge firewall or integrated with the reverse proxy/load balancer). The WAF is the control that addresses HTTP-layer attacks, including directory traversal and XSS, and can enforce/request-check CSRF protections.
Place a **load balancer/reverse proxy** in front of multiple web servers. This reduces direct exposure of the web hosts and provides availability/resilience.
Use **multiple web servers** in the application tier; they should not be directly Internet accessible.
Install/use a trusted **PKI/TLS certificate** on the public TLS termination point (the WAF, reverse proxy, or load balancer; or the web server if TLS passes through). Configure it to use HTTPS/TLS rather than insecure HTTP or deprecated SSL/TLS versions.
Keep the database/backend on a separate internal network segment and allow access only from the application/web tier. Do not expose it through the public-facing firewall.

Expected traffic path:

```
Internet → Firewall → WAF → Load balancer/reverse proxy (TLS certificate/HTTPS) → Web-server pool → Internal database
```

Part 2 — select command snippets with these effects:

Redirect HTTP to HTTPS: select the snippet equivalent to `RewriteRule ^ https://%{HTTP_HOST}%{REQUEST_URI} [L,R=301]` (or an HTTP virtual-host redirect such as `Redirect permanent / https://example.com/`).
Disable directory browsing/listing: select `Options -Indexes`.
Use TLS and disable insecure SSL/TLS: select the snippet equivalent to `SSLProtocol -all +TLSv1.2 +TLSv1.3`. Do not select SSLv2, SSLv3, TLS 1.0, or TLS 1.1.
Enable the site certificate/key: select directives equivalent to `SSLEngine on`, `SSLCertificateFile <certificate>`, and `SSLCertificateKeyFile <private-key>`, if those are among the command choices.

The supplied text does not include the actual topology image or the available drop-down values, so exact position labels and literal command-choice wording cannot be verified. Choose the options matching the functions above; the key security answer is **firewall → WAF → load balancer/reverse proxy → isolated web tier**, with a **trusted TLS certificate/HTTPS** at the public entry point and commands that **redirect HTTP, disable indexing, and permit only TLS 1.2/1.3**.

QUESTION NO: 19

A network team segmented a critical, end-of-life server to a VLAN that can only be reached by specific devices but cannot be reached by the perimeter network. Which of the following best describe the controls the team implemented? (Choose two.)

A. Managerial

B. Physical

C. Corrective

D. Detective

E. Compensating

F. Technical

G. Deterrent

ANSWER: E F

Explanation:

Compensating and **Technical** correctly describe this implementation. A compensating control is an alternative safeguard used to reduce risk when the preferred or primary control cannot reasonably be implemented. Because the server is end-of-life, it may no longer receive vendor patches or support, preventing the organization from fully remediating its inherent vulnerabilities. Restricting the server to a dedicated VLAN, allowing access only from specifically authorized devices, and preventing perimeter-network access reduces its reachable attack surface and exposure. This provides risk reduction that compensates for the system's unsupported condition.

The control is technical because it is implemented and enforced through technology. VLAN configuration, switch port assignments, router access-control lists, firewall rules, and network segmentation mechanisms make the access restriction operationally enforceable. These mechanisms constrain network communications based on device location, identity, address, or policy rather than relying on personnel to follow a procedure. NIST describes network segmentation and boundary protections as technical mechanisms for controlling information flows, while VLANs logically separate network traffic into distinct broadcast domains. See NIST SP 800-53 Rev. 5 and Cisco's VLAN overview.

QUESTION NO: 20

Which of the following vulnerabilities would likely be mitigated by setting up an MDM platform?

A. TPM

B. Buffer overflow

C. Jailbreaking

D. SQL injection

ANSWER: C

Explanation:

Jailbreaking is the correct answer because a mobile device management (MDM) platform centrally applies and monitors an organization's security requirements for managed phones and tablets. Jailbreaking removes or bypasses vendor-enforced operating-system restrictions, which can undermine controls such as application sandboxing, code-signing requirements, and limitations on privileged access. This increases the likelihood that untrusted software can run, sensitive organizational data can be exposed, or device configurations can be altered outside approved policy.

An MDM platform can evaluate device compliance and identify indicators that a device has been jailbroken or rooted. The organization can then enforce conditional-access rules that deny corporate email, applications, VPN connectivity, and other protected resources to noncompliant devices. MDM can also require device encryption, approved OS versions, screen-lock policies, and application controls; depending on the deployment and ownership model, it may support selective removal of organizational data or a remote wipe. These capabilities reduce organizational exposure from compromised mobile-device security posture while allowing administrators to consistently enforce mobile security policy at scale. Microsoft's overview of device compliance policies describes using compliance status to protect access to organizational resources: Microsoft Intune device compliance documentation. Apple also documents that device management can query device information and apply configuration and security policies: Apple Platform Deployment: Intro to MDM.

QUESTION NO: 21

Which of the following are the first steps an analyst should perform when developing a heat map? (Choose two.)

A. Methodically walk around the office noting Wi-Fi signal strength.

B. Log in to each access point and check the settings.

C. Create or obtain a layout of the office.

D. Measure cable lengths between access points.

E. Review access logs to determine the most active devices.

F. Remove possible impediments to radio transmissions.

ANSWER: A C

Explanation:

Developing a wireless heat map begins with establishing an accurate representation of the physical environment and gathering RF measurements throughout that environment. “Create or obtain a layout of the office” provides the floor-plan baseline on which survey software can associate measurements with specific locations. The layout should reflect relevant physical details, such as room boundaries, walls, and other structural features, because these affect signal propagation and make the completed visualization actionable for identifying weak-coverage areas.

“Methodically walk around the office noting Wi-Fi signal strength” is the measurement-collection phase of a wireless site survey. The analyst takes readings at defined points across the coverage area, allowing the survey tool to plot received signal levels spatially. These measurements form the data behind a heat map and help reveal where coverage, capacity, or roaming may need further investigation. A methodical walking path and consistent sampling approach improve the reliability of the resulting map.

These activities align with the Security+ expectation that analysts understand wireless site-survey concepts and wireless signal characteristics. Cisco describes a site survey as collecting RF information across the intended deployment area, while NetSpot describes using a map and measurements to produce a Wi-Fi heat map. See Cisco Wireless Site Survey and NetSpot Wi-Fi Heatmap.

QUESTION NO: 22

An employee used a company’s billing system to issue fraudulent checks. The administrator is looking for evidence of other occurrences of this activity. Which of the following should the administrator examine?

A. Application logs

B. Vulnerability scanner logs

C. IDS/IPS logs

D. Firewall logs

ANSWER: A

Explanation:

Application logs are the appropriate evidence source because the suspected fraud occurred through a business application: the company billing system. Well-configured application audit logs can record the identity of the authenticated user, the time of the activity, the action performed, the relevant transaction or check identifier, approval workflow events, changes to payee or payment details, and potentially the source system or client session. These records allow the administrator to search for additional check-issuance events associated with the employee’s account and identify patterns such as unusual payment amounts, repeated vendors, changes made immediately before payment, or activity occurring outside normal work hours.

The administrator should preserve the original logs and relevant transaction data before conducting detailed analysis, maintaining integrity and an appropriate chain of custody. They can then correlate billing-system events with identity-management records and time information to establish a reliable sequence of activity. OWASP recommends application logging of security-relevant events, including authentication, authorization, data changes, and high-value transactions, all of which are directly relevant to investigating fraudulent check issuance. See the OWASP Logging Cheat Sheet and NIST’s guidance on computer security log management.

QUESTION NO: 23

A new employee accessed an unauthorized website. An investigation found that the employee violated the company's rules. Which of the following did the employee violate?

A. MOU

B. AUP

C. NDA

D. MOA

ANSWER: B

Explanation:

AUP is correct because an Acceptable Use Policy defines the permitted and prohibited ways employees may use an organization's information systems, network connections, internet access, applications, and other technology resources. Organizations use this policy to establish clear boundaries for activities such as visiting websites, downloading content, using cloud services, connecting personal devices, and accessing company data. The policy commonly identifies prohibited website categories or states that users may access internet resources only for authorized business purposes. By accessing a website the company did not authorize, the employee violated the organization's established rules for acceptable use of its technology resources. This is also why employees are typically required to acknowledge the policy during onboarding: the acknowledgment documents that they understand the conditions governing access to company systems. In a Security+ context, an Acceptable Use Policy is an administrative control that supports secure and appropriate user behavior and gives the organization a basis for investigating and responding to misuse. Guidance on developing and maintaining security policies, including acceptable-use requirements, is available from SANS Information Security Policy Templates and NIST SP 800-12 Rev. 1.

QUESTION NO: 24

A systems administrator is changing the password policy within an enterprise environment and wants this update implemented on all systems as quickly as possible. Which of the following operating system security measures will the administrator most likely use?

A. Deploying PowerShell scripts

B. Pushing GPO update

C. Enabling PAP

D. Updating EDR profiles

ANSWER: B

Explanation:

Pushing a GPO update is the appropriate way to distribute a revised password policy rapidly and consistently throughout a Windows Active Directory enterprise. Group Policy provides centralized administration: an administrator can configure domain password settings, including password length, complexity requirements, password history, maximum and minimum password age, and account-lockout behavior, in the applicable domain-level Group Policy Object. Domain-joined systems then retrieve and apply that policy automatically during normal Group Policy processing.

When the change must take effect as soon as practical, the administrator can trigger a Group Policy refresh rather than waiting for the standard background-refresh interval. For example, `gpupdate /force` refreshes policy on an individual system, while enterprise management methods can initiate the refresh remotely across many systems. This approach ensures the intended policy is centrally governed, repeatable, and uniformly applied instead of requiring each endpoint to be configured independently.

Microsoft describes Group Policy as infrastructure for centrally managing users and computers in an Active Directory environment, including security-related settings. See the Microsoft Group Policy overview and Group Policy processing documentation.

QUESTION NO: 25

A company wants to reduce the risk of privilege escalation on Windows workstations. Which of the following should the company implement? (Choose two.)

A. Removing local administrator rights from standard users

B. Using shared administrator accounts for routine work

C. Using separate administrative accounts

D. Adding all users to the local Administrators group

E. Disabling endpoint logging

F. Allowing unrestricted software installation

ANSWER: A C

Explanation:

Removing local administrator rights from standard users reduces the risk of privilege escalation by limiting the ability of users to install software, modify protected system settings, and make changes requiring elevated permissions. Users should receive only the permissions needed for their assigned responsibilities.

Using separate administrative accounts further reduces risk by preventing administrators from performing daily email, web browsing, and office tasks with privileged credentials. Administrative accounts should be used only when elevated access is necessary, which limits exposure of those credentials to phishing, malware, and browser-based attacks. These practices support the principle of least privilege. See the NIST definition of least privilege and CIS guidance on account management.

QUESTION NO: 26

A security analyst is investigating a suspected account compromise. Which of the following log sources would provide the most useful information about successful and failed user sign-in attempts? (Choose two.)

A. Authentication

B. Printer

C. VPN

D. Environmental sensor

E. Database schema

F. Building maintenance

ANSWER: A C

Explanation:

Authentication logs provide records of successful and failed sign-in events, including usernames, timestamps, source addresses, authentication methods, and failure reasons. These records help an analyst identify password-spraying activity, impossible travel, repeated failures, and unexpected successful authentications.

VPN logs are also valuable when remote access is involved. They can show which accounts established remote sessions, the assigned internal address, the external source address, connection duration, and the time of access. Correlating VPN and authentication events can help determine whether an account was used from an unauthorized location. NIST recommends collecting and analyzing authentication and remote-access events as part of security log management; see NIST SP 800-92.

QUESTION NO: 27

An organization is retiring several multifunction printers that stored scanned documents on internal drives. Which of the following should the organization perform before transferring the printers to a third party? (Choose two.)

A. Sanitize the internal storage media.

B. Remove network and administrative credentials.

C. Perform a vulnerability scan after transfer.

D. Enable remote printing for the third party.

E. Increase the printer paper capacity.

F. Apply a legal hold to all print jobs.

ANSWER: A B

Explanation:

Sanitize the internal storage media because multifunction printers can retain scanned documents, print jobs, credentials, and configuration information. An approved media-sanitization method reduces the risk that this information can be recovered by the next owner.

Remove network and administrative credentials because stored wireless keys, directory-service accounts, SMTP settings, and administrator passwords could be used to access organizational resources if left on the device. Configuration reset and credential removal should be verified before transfer. NIST provides guidance for selecting and documenting sanitization methods in NIST SP 800-88 Rev. 2, Guidelines for Media Sanitization.

QUESTION NO: 28

A security analyst is performing a packet capture on a series of SOAP HTTP requests for a security assessment. The analyst redirects the output to a file After the capture is complete, the analyst needs to review the first transactions quickly and then search the entire series of requests for a particular string Which of the following would be BEST to use to accomplish the task? (Select TWO).

A. head

B. Tcpdump

C. grep

D. rail

E. curl

F. openssi

G. dd

ANSWER: A C

Explanation:

head and **grep** are the appropriate command-line utilities for this workflow. After output has been redirected to a file, **head** provides an immediate view of the start of that file. By default, it displays the first ten lines, and its line-count option can be used to display a larger or smaller initial segment as needed. This lets the analyst rapidly inspect the earliest captured SOAP HTTP transactions without loading or manually scrolling through the complete output.

grep searches input text for lines that match a supplied pattern or literal string. The analyst can run it against the saved capture output to locate a SOAP element, HTTP header, URI fragment, identifier, error text, or other relevant string across the entire series of requests. Useful operational variants include case-insensitive matching and line-number output, which make findings easier to locate and review. Together, these commands directly meet the two stated needs: quickly examining the beginning of a saved text output and searching all of its contents. The GNU documentation for head and grep describes these functions and supported usage options.

QUESTION NO: 29

Several employees received a fraudulent text message from someone claiming to be the Chief Executive Officer (CEO). The message stated:

“I’m in an airport right now with no access to email. I need you to buy gift cards for employee recognition awards. Please send the gift cards to following email address.”

Which of the following are the best responses to this situation? (Choose two).

A. Cancel current employee recognition gift cards.

B. Add a smishing exercise to the annual company training.

C. Issue a general email warning to the company.

D. Have the CEO change phone numbers.

E. Conduct a forensic investigation on the CEO’s phone.

F. Implement mobile device management.

ANSWER: B C

Explanation:

Add a smishing exercise to the annual company training and **Issue a general email warning to the company** are the best responses. This event is a social-engineering attack using SMS, commonly called smishing, combined with executive impersonation and an urgent gift-card request. A prompt organization-wide warning helps contain the incident by informing employees that the request is fraudulent, directing them not to purchase or send gift cards, and reminding them to report similar messages through established security channels. Rapid internal communication is important because the attacker may target multiple employees and may send follow-up messages from other numbers or addresses.

Including a smishing exercise in awareness training strengthens the preventive control needed for future attempts. Training should reinforce recognizing urgency, unusual payment requests, requests to bypass normal processes, and executive impersonation. It should also teach employees to independently verify sensitive requests through a known, trusted contact method rather than replying to the suspicious text. These actions support a security-awareness program and improve users’ ability to identify and report phishing-based social engineering. CISA provides practical guidance on recognizing and reporting phishing and social-engineering attempts, while NIST documents awareness and training controls for organizations: CISA phishing guidance and NIST SP 800-53 Rev. 5.

QUESTION NO: 30

A vendor needs to remotely and securely transfer files from one server to another using the command line. Which of the following protocols should be Implemented to allow for this type of access? (Select two).

A. SSH

B. SNMP

C. RDP

D. S/MIME

E. SMTP

F. SFTP

ANSWER: A F

Explanation:

SSH and SFTP provide the secure, command-line-capable access required for this scenario. SSH establishes an encrypted, authenticated remote session between systems. It protects the confidentiality and integrity of administrative traffic and can be used from a command-line client to securely access the destination server. SSH also provides the secure transport foundation used by several remote administration and file-transfer capabilities.

SFTP is specifically designed for securely transferring and managing files over an SSH connection. An administrator or vendor can invoke an SFTP client from a command line, authenticate to the remote SSH service, and upload, download, list, rename, or remove files without exposing file contents or credentials in cleartext. Because SFTP operates through the SSH transport, it inherits SSH's encryption and server-authentication protections while providing file-oriented commands rather than only an interactive shell.

These technologies are commonly deployed together: the server runs an SSH service, and SFTP is enabled as an SSH file-transfer subsystem. This combination satisfies both parts of the requirement: remote command-line access and protected server-to-server file transfer. See the SSH protocol architecture in RFC 4251 and the OpenSSH sftp manual for command-line SFTP operation over SSH.

QUESTION NO: 31

A security administrator needs a method to secure data in an environment that includes some form of checks so that the administrator can track any changes. Which of the following should the administrator set up to achieve this goal?

A. SPF

B. GPO

C. NAC

D. FIM

ANSWER: D

Explanation:

File Integrity Monitoring (FIM) is the appropriate control because it establishes a known-good baseline for important files, directories, configurations, and other protected system objects. The baseline commonly includes cryptographic hashes along with attributes such as ownership, permissions, size, and modification time. FIM then periodically scans or continuously monitors those objects and compares their current state with the baseline. When it detects that an item was created, modified, deleted, or had its permissions changed, it records the event and can alert the security administrator for investigation. This provides the requested integrity checks as well as an auditable record of changes. FIM is particularly valuable for security-sensitive assets such as operating-system files, application binaries, configuration files, logs, and web content, where unauthorized modification can indicate malware, persistence, configuration drift, or data tampering. To be effective, the initial baseline should be created from a trusted state, protected from unauthorized alteration, and updated through approved change-management processes. NIST describes file integrity checking as a means of detecting unauthorized changes to critical files; see NIST SP 800-92, Guide to Computer Security Log Management. Additional implementation guidance is available in the OSSEC file integrity monitoring documentation.

QUESTION NO: 32

Which of the following explains how to determine the global regulations that data is subject to regardless of the country where the data is stored?

A. Geographic dispersion

B. Data sovereignty

C. Geographic restrictions

D. Data segmentation

ANSWER: B

Explanation:

Data sovereignty is correct because it describes the principle that data is subject to the laws, regulatory requirements, and governmental authority of the applicable jurisdiction. Organizations must identify which jurisdictions have legal authority over their data and processing activities when determining obligations such as privacy protections, disclosure requirements, retention periods, lawful-access demands, and cross-border transfer conditions. This assessment can involve the organization's location, the locations of its customers or data subjects, contractual commitments, and the jurisdictions in which processing occurs; it is not limited to the physical location of a particular storage system.

For cloud services, data sovereignty is particularly important because workloads may be hosted, replicated, backed up, or accessed across national boundaries. An organization must understand which legal regimes apply before selecting cloud regions, defining data-handling policies, or establishing controls for international transfers. This enables the organization to align its technical and administrative safeguards with the legal obligations governing the information. The concept is closely related to data residency, but sovereignty focuses on the jurisdictional authority and laws applicable to data. See the Cloudflare overview of data sovereignty and the NIST Privacy Framework for supporting privacy and governance context.

QUESTION NO: 33

During the onboarding process, an employee needs to create a password for an intranet account. The password must include ten characters, numbers, and letters, and two special characters. Once the password is created, the company will grant the employee access to other company-owned websites based on the intranet profile. Which of the following access management concepts is the company most likely using to safeguard intranet accounts and grant access to multiple sites based on a user's intranet account? (Choose two.)

A. Federation

B. Identity proofing

C. Password complexity

D. Default password changes

E. Password manager

F. Open authentication

ANSWER: A C

Explanation:

Password complexity is being used to safeguard the intranet account because the organization explicitly requires a password with a minimum length and multiple character types, including letters, numbers, and special characters. These composition requirements are a password-policy control intended to increase resistance to straightforward guessing and automated password-cracking attempts. A strong password policy is part of the authentication controls that protect the account serving as the employee's initial corporate identity.

Federation applies because the employee's intranet identity is used as the trusted basis for access to other company-owned websites. Federated identity management establishes a trust relationship in which separate applications or services accept authentication assertions from an identity provider rather than requiring users to establish and manage independent credentials at every site. This commonly supports a streamlined single sign-on experience while allowing the organization to centrally manage the employee's identity, access status, and profile attributes. The scenario therefore combines a password-complexity policy for protecting the initial credential with federation for extending that trusted identity across multiple sites. See Cloudflare's federated identity management overview and NIST SP 800-63B guidance on authentication and memorized secrets.

QUESTION NO: 34

Which of the following is an example of implementing Zero Trust architecture?

A. Building strong network boundaries to prevent intrusion

B. Verifying user identity once at the start of the session

C. Granting resource access after continuous validation

D. Prioritizing perimeter defense to block external threats

ANSWER: C

Explanation:

Granting resource access after continuous validation is an implementation of Zero Trust architecture because Zero Trust removes the assumption that a user, device, application, or network location is inherently trusted. Instead, access decisions are evaluated using explicit, ongoing verification of relevant contextual signals, such as user identity, device security posture, authentication strength, location, behavior, and the sensitivity of the requested resource. Access is then granted according to least-privilege principles, limiting the user or workload to only the specific resources and permissions required. Continuous validation is especially important because risk can change after an initial sign-in: a device can become noncompliant, a session can exhibit suspicious behavior, or the requested resource can require stronger assurance. Zero Trust therefore treats authorization as a dynamic process rather than a one-time event. This approach aligns with the Zero Trust principles described by the National Institute of Standards and Technology, including continuously evaluating access requests and enforcing policy before allowing access to enterprise resources. Microsoft similarly describes Zero Trust with the core principles of explicitly verifying access, using least-privilege access, and assuming breach. See NIST SP 800-207: Zero Trust Architecture and Microsoft Zero Trust guidance.

QUESTION NO: 35

A security engineer is working to address the growing risks that shadow IT services are introducing to the organization. The organization has taken a cloud-first approach end does not have an on-premises IT infrastructure. Which of the following would best secure the organization?

A. Upgrading to a next-generation firewall

B. Deploying an appropriate in-line CASB solution

C. Conducting user training on software policies

D. Configuring double key encryption in SaaS platforms

ANSWER: B

Explanation:

Deploying an appropriate in-line CASB solution is the best fit because a cloud access security broker provides visibility and policy enforcement for cloud-service use, including unsanctioned services that constitute shadow IT. An in-line CASB operates in the path of user traffic, commonly using forward-proxy controls or endpoint traffic steering, so it can identify cloud applications as they are accessed and apply organization-defined rules in real time. This enables the organization to discover risky or unapproved SaaS use, restrict access to prohibited applications, enforce acceptable-use policies, and apply controls such as data loss prevention and upload/download restrictions where appropriate.

This approach is particularly suitable for a cloud-first organization because the primary security boundary is the interaction among users, devices, identities, data, and cloud services rather than an on-premises network perimeter. CASB capabilities can help maintain consistent governance across many SaaS providers, including services that users may adopt without formal IT approval. Effective deployment should be integrated with the organization's identity provider, device-management posture, and data-classification policies so access decisions and data controls are based on authenticated users and contextual risk. NIST describes cloud access security broker functions and cloud security considerations in its Guidelines on Security and Privacy in Public Cloud Computing; Microsoft also outlines how a CASB provides cloud-app discovery and control in its CASB overview.

QUESTION NO: 36

A company must ensure that log searches are conducted in the shortest time frame. Which of the following should the company do to maintain logs in live storage for 90 days?

A. Conduct deduplication.

B. Conduct archiving.

C. Apply aggregation.

D. Apply compression.

ANSWER: C

Explanation:

Apply aggregation is correct because aggregating logs centralizes event data from endpoints, servers, network devices, applications, and security controls into a common logging platform, such as a SIEM. Rather than requiring an analyst to access and search each source independently, the organization can search a single live repository. This supports faster investigation, particularly when the platform normalizes records, maintains searchable indexes, and enables correlation across sources. Keeping the aggregated data in the platform's live, searchable storage for the required 90-day period ensures that recent operational and security events are immediately available for incident response, threat hunting, and routine monitoring. Effective log aggregation also enables consistent retention handling and makes it practical to apply a common search interface, timestamps, and query capabilities to the retained records. CompTIA Security+ includes log aggregation and analysis as core security monitoring concepts, emphasizing centralized collection and use of log data to support detection and investigation. Guidance from NIST likewise identifies centralized log management as a way to improve the efficiency and effectiveness of reviewing and analyzing organizational log information. See NIST SP 800-92, Guide to Computer Security Log Management and the CompTIA exam objectives resource.

QUESTION NO: 37

During the onboarding process, an employee needs to create a password for an intranet account. The password must include ten characters, numbers, and letters, and two special characters. Once the password is created, the company will grant the employee access to other company-owned websites based on the intranet profile. Which of the following access management concepts is the company most likely using to safeguard intranet accounts and grant access to multiple sites based on a user's intranet account? (Select two).

A. Federation

B. Identity proofing

C. Password complexity

D. Default password changes

E. Password manager

F. Open authentication

ANSWER: A C

Explanation:

Password complexity is correct because the stated requirement establishes password-composition rules: a minimum length plus required character types, including letters, numbers, and special characters. Such requirements are part of an organization's authentication policy and are intended to make passwords more resistant to guessing and automated password-cracking attempts. Password policy requirements should be implemented alongside broader controls, such as secure password storage, multifactor authentication where appropriate, and protections against online guessing attempts. NIST discusses the role of password requirements within digital identity and authenticator management guidance at NIST SP 800-63B.

Federation is correct because the employee's intranet identity is used as a trusted basis for obtaining access to other company-owned websites. In a federated identity arrangement, an identity provider authenticates the user and conveys identity attributes or claims to relying applications. Those applications can then make access decisions based on the centrally maintained intranet profile, reducing the need for separate credentials at every site. Federation commonly supports single sign-on experiences across related services while preserving centralized identity and access management. Microsoft's

identity platform documentation describes federation as an arrangement that enables trust between identity systems and applications: Microsoft Entra federation.

QUESTION NO: 38

Which of the following most securely protects data at rest?

A. TLS 1.2

B. AES-256

C. Masking

D. Salting

ANSWER: B

Explanation:

AES-256 most securely protects data at rest because it is a strong symmetric encryption standard that converts stored plaintext into ciphertext. Data stored on devices, database volumes, removable media, backups, and file systems remains unreadable to an unauthorized party unless that party can obtain the appropriate decryption key. The “256” designation refers to the key length, which provides a very large keyspace and is commonly selected for high-assurance protection of sensitive information. AES is a NIST-approved block cipher and is widely implemented in full-disk encryption, database encryption, encrypted backups, and cloud storage encryption services. Effective AES-256 protection also depends on secure key generation, storage, access control, rotation, and recovery procedures; encryption keys must be protected separately from the data they encrypt. For Security+ purposes, encryption is the core control for preserving confidentiality of data at rest, while AES-256 is a recognized strong algorithm for performing that encryption. NIST specifies AES in Federal Information Processing Standard 197, and NIST guidance on storage encryption is available in Special Publication 800-111.

QUESTION NO: 39

A network administrator deploys an FDE solution on all end user workstations. Which of the following data protection strategies does this describe?

A. Masking

B. Data in transit

C. Obfuscation

D. Data at rest

E. Data sovereignty

ANSWER: D

Explanation:

Data at rest is correct because full-disk encryption (FDE) protects information stored on a workstation’s internal drive or other persistent storage. FDE encrypts the disk contents automatically, generally including user files, the operating system, application data, and temporary files. The information remains unreadable without successful authentication and access to the appropriate cryptographic key. This protection is especially important for endpoint devices because a lost or stolen workstation, or a drive removed from its original system, could otherwise expose locally stored organizational and personal information.

In Security+ terms, data at rest is data residing in a storage location rather than actively moving across a network or being processed in memory. FDE is therefore a confidentiality control for stored data. Platforms commonly implement this capability through features such as BitLocker, which encrypts volumes and can use a Trusted Platform Module to help protect encryption keys during the boot process. NIST likewise identifies storage encryption as a safeguard for data held on

endpoint devices and other storage media. See Microsoft BitLocker documentation and NIST SP 800-111, Guide to Storage Encryption Technologies for End User Devices.

QUESTION NO: 40

A security administrator needs to reduce the attack surface in the company ' s data centers. Which of the following should the security administrator do to complete this task?

A. Implement a honeynet.

B. Define Group Policy on the servers.

C. Configure the servers for high availability.

D. Upgrade end-of-support operating systems.

ANSWER: D

Explanation:

Upgrade end-of-support operating systems is the best action because unsupported operating systems expand an organization's attack surface. Once a vendor ends support, the operating system no longer receives routine security updates for newly discovered vulnerabilities. Consequently, any unpatched flaws remain available to attackers, who can use public exploit code, vulnerability scans, or targeted techniques to compromise affected servers. Replacing or upgrading these systems with supported versions restores access to vendor security patches and modern platform security capabilities, such as improved encryption support, stronger authentication protections, and current endpoint-security compatibility. This directly reduces the number of known and unremediated weaknesses exposed in the data center environment. A sound lifecycle-management process should inventory operating systems, identify end-of-support dates, test an upgrade or migration plan, and retire the legacy systems after workloads have been safely moved. Where an immediate upgrade is not feasible, the organization should apply documented compensating controls while prioritizing remediation, but those temporary measures do not eliminate the underlying exposure. This approach aligns with NIST guidance that organizations should identify and remediate vulnerabilities through patch and configuration management. See NIST SP 800-40 Rev. 4, Guide to Enterprise Patch Management Planning and CISA's Known Exploited Vulnerabilities Catalog.

QUESTION NO: 41

A technician needs to apply a high-priority patch to a production system. Which of the following steps should be taken first?

A. Air gap the system.

B. Move the system to a different network segment.

C. Create a change control request.

D. Apply the patch to the system.

ANSWER: C

Explanation:

Create a change control request. Even when a patch is high priority, production changes should begin through the organization's established change-management process. A change control request documents the purpose and scope of the patch, identifies affected systems and business services, records risk and impact considerations, obtains the required authorization, and establishes an implementation plan. Critically, it also captures validation and rollback requirements so the organization can restore service quickly if the patch introduces an operational issue.

Security+ emphasizes that secure operational practices must balance timely remediation with availability, accountability, and controlled risk. A high-priority patch may qualify for an emergency-change procedure, which can accelerate review and approval, but it does not eliminate the need to record, authorize, and manage the production change. Beginning with a

change request ensures relevant stakeholders are notified, an appropriate maintenance window or emergency process is selected, backups and recovery preparations are considered, and the patch can be audited after implementation.

This approach supports repeatable, accountable patch management while reducing the chance that an urgent remediation causes an unplanned outage. See Atlassian's change-management overview and NIST's guide to enterprise patch management.

QUESTION NO: 42

A company's website is www.company.com. Attackers purchased the domain www.c0mpany.com. Which of the following types of attacks describes this example?

A. Typosquatting

B. Brand impersonation

C. On-path

D. Watering-hole

ANSWER: A

Explanation:

Typosquatting is correct because the attackers registered a domain name designed to be visually confused with the company's legitimate domain. The fraudulent address replaces the letter "o" in "company" with the numeral "0," a small substitution that can be easily overlooked in a browser address bar, email hyperlink, advertisement, or manually entered URL. This technique relies on users mistyping an address or failing to recognize a subtle difference between a trusted domain and an attacker-controlled look-alike.

Typosquatting domains are commonly used to support phishing and credential theft. For example, an attacker can build a page that closely resembles the genuine company website, prompt visitors to sign in, and capture submitted credentials or payment information. They may also redirect users to malicious downloads or use the domain in convincing business-email and social-engineering campaigns. The defining element is the registration and use of a deceptively similar domain name, rather than compromise of the legitimate domain itself.

CompTIA Security+ treats typosquatting as a domain-based social-engineering and impersonation technique. ICANN also identifies confusingly similar domain names as a common mechanism in phishing activity. See Cloudflare: What is typosquatting? and ICANN: Phishing.

QUESTION NO: 43

An organization is implementing a COPE mobile device management policy. Which of the following should the organization include in the COPE policy? (Choose two.)

A. Remote wiping of the device

B. Data encryption

C. Requiring passwords with eight characters

D. Data usage caps

E. Employee data ownership

F. Personal application store access

ANSWER: A B

Explanation:

COPE, meaning corporate-owned, personally enabled, describes devices that are purchased, issued, and managed by the organization while allowing limited personal use by the employee. Because the organization retains ownership and must protect business information throughout the device life cycle, **Remote wiping of the device** is an essential policy control. The policy should authorize IT to erase the device when it is lost, stolen, compromised, or being reassigned or retired. A wipe capability limits exposure of corporate accounts, locally stored files, application data, credentials, and security configurations when the device is no longer under authorized control.

Data encryption is also fundamental for a COPE deployment. Encryption protects data at rest on the device so that storage contents are not readable merely because an unauthorized person obtains physical possession of the phone or tablet. The MDM policy should require supported device encryption, define compliance expectations, and ensure that access to encrypted data is tied to appropriate authentication. Together, encryption reduces the risk from physical loss or theft, while remote wiping provides a management response when the organization needs to remove information from the device. These controls align with mobile-device security practices for organization-managed endpoints. See NIST SP 800-124 Rev. 2 and Apple Platform Deployment: Remote wipe.

QUESTION NO: 44

An organization recently updated its security policy to include the following statement:

Regular expressions are included in source code to remove special characters such as $, |, ;. &, `, and ? from variables set by forms in a web application.

Which of the following best explains the security technique the organization adopted by making this addition to

the policy?

A. Identify embedded keys

B. Code debugging

C. Input validation

D. Static code analysis

ANSWER: C

Explanation:

Input validation is the security technique described because the web application evaluates data submitted through forms and applies defined rules before allowing that data to be processed. In this case, regular expressions are used to identify and remove characters that can have special meaning to interpreters, databases, operating-system shells, and other downstream components. This reduces the chance that untrusted input will be interpreted as part of a command, query, or script rather than as ordinary application data.

Effective input validation should be performed server-side and as close as possible to the point where the application receives the data. A robust implementation normally uses an allowlist approach: it defines the expected type, length, format, and permitted character set for each field. For example, a phone-number field should accept only the characters needed for a phone number, while an account identifier should accept only its documented format. Regular expressions can enforce these expected patterns. When data is later used in SQL, OS commands, or similar contexts, parameterized queries and context-appropriate encoding should also be used as complementary protections.

This policy statement therefore establishes a control for handling untrusted form input and helping prevent injection-related attacks. OWASP describes input validation practices at OWASP Input Validation Cheat Sheet and injection defenses at OWASP Injection Prevention Cheat Sheet.

QUESTION NO: 45

A security analyst needs to improve the company's authentication policy following a password audit. Which of the following should be included in the policy? (Choose two.)

A. Length

B. Complexity

C. Least privilege

D. Something you have

E. Security keys

F. Biometrics

ANSWER: A B

Explanation:

Length and **Complexity** are appropriate password-policy controls to address after a password audit. A minimum password length increases the number of possible password combinations and raises the work factor for offline cracking and brute-force attempts. Current guidance strongly favors longer, memorable passwords or passphrases because length provides substantial resistance to guessing when users select unique values. A policy should therefore establish a minimum length appropriate to the organization's risk level and authentication system capabilities.

Complexity addresses password predictability by requiring passwords to avoid easily guessed patterns and by establishing rules for acceptable password construction. In an exam policy context, complexity commonly includes using varied character types and preventing obvious or common passwords. These controls help ensure that credentials identified during the audit are less susceptible to dictionary-based guessing and common password attacks. Password-policy requirements should also be supported by technical enforcement in the identity provider or directory service, as well as password screening against known-compromised values. NIST's digital identity guidance discusses memorized-secret length, screening, and related verifier requirements, while CISA provides practical guidance for choosing and protecting passwords. See NIST SP 800-63B and CISA Secure Our World.

QUESTION NO: 46

Which of the following would best allow a company to prevent access to systems from the Internet?

A. Containerization

B. Virtualization

C. SD-WAN

D. Air-gapped

ANSWER: D

Explanation:

Air-gapped is correct because an air-gapped system or network is deliberately isolated from external networks, including the public Internet. This separation removes the normal network route an Internet-based user or attacker would need in order to reach the protected systems. For environments with especially sensitive workloads, such as critical operational systems, classified data, or highly regulated assets, air gapping provides a strong architectural control by ensuring the environment has no direct network connectivity to Internet-facing infrastructure. The isolation must be implemented comprehensively: network interfaces, wireless connectivity, remote administration paths, and any bridging devices must be controlled so they do not create an unintended connection to an untrusted network. Data exchange with an air-gapped environment is typically handled through tightly governed, limited processes, such as scanned removable media or purpose-built transfer mechanisms, rather than ordinary Internet communications. NIST defines an air gap as a physical separation between a system or network and another network, supporting the principle that physical and logical isolation can prevent external network access. See the NIST Computer Security Resource Center definition of air gap and NIST guidance on security and privacy controls.

QUESTION NO: 47

An organization is reviewing its identity and access management processes after discovering that several former employees still have access to cloud applications. Which of the following actions would best reduce the likelihood of this issue occurring again? (Choose two.)

A. Implement an automated deprovisioning process.

B. Conduct periodic access reviews.

C. Increase the password expiration frequency.

D. Require users to change their usernames annually.

E. Enable guest access for all cloud applications.

ANSWER: A B

Explanation:

Implementing an automated deprovisioning process is appropriate because it removes or disables accounts and application access when an employee leaves the organization. Integrating human resources termination events with the identity provider helps ensure access is revoked promptly and consistently.

Conducting periodic access reviews is also appropriate because reviews allow application owners and managers to validate that assigned access remains necessary. These reviews can identify orphaned accounts, excessive permissions, and accounts that were missed by an offboarding workflow. Together, automated deprovisioning and recurring reviews support account lifecycle management and least privilege. NIST access-control guidance addresses account management and periodic review requirements in NIST SP 800-53 Rev. 5.

QUESTION NO: 48

Which of the following are the best security controls for controlling on-premises access? (Select two.)

A. Swipe card

B. Picture ID

C. Phone authentication application

D. Biometric scanner

E. Camera

F. Memorable

ANSWER: A D

Explanation:

Swipe card and **Biometric scanner** are appropriate physical access controls for managing entry to an on-premises facility. A swipe card, commonly implemented as an access badge or proximity credential, allows an organization to grant, deny, modify, and revoke access centrally. Access can be limited according to the person's role, permitted doors, and authorized times, while access-control systems can also retain entry records for auditing and incident investigations.

A biometric scanner strengthens physical access decisions by verifying a characteristic associated with the individual requesting entry, such as a fingerprint, face, or iris. This helps establish that the person presenting for access is the authorized person and is commonly used where stronger assurance is needed, including sensitive areas and data centers. These controls can be used together: a badge supplies a credential and the biometric check verifies the credential holder, providing multifactor physical access control. CompTIA includes badges and biometrics among physical security controls, and NIST describes physical access control systems as mechanisms that control and monitor access to facilities and protected areas. See the CompTIA exam objectives resource and NIST SP 800-53 physical and environmental protection controls.

QUESTION NO: 49

A company is changing its mobile device policy. The company has the following requirements:

• Company-owned devices

• Ability to harden the devices

• Reduced security risk

• Compatibility with company resources

Which of the following would best meet these requirements?

A. BYOD

B. CYOD

C. COPE

D. COBO

ANSWER: D

Explanation:

COBO (company-owned, business-only) best satisfies all of the stated requirements. Under this ownership model, the organization purchases, issues, and retains control of each mobile device. Centralized management through a mobile device management platform lets administrators apply a secure baseline consistently, including enforced screen-lock settings, device encryption, supported operating-system versions, approved applications, certificate deployment, remote wipe capability, and restrictions on high-risk device features. This level of administrative control makes device hardening practical and auditable.

The business-only use restriction also reduces exposure by preventing personal applications, personal accounts, and unmanaged consumer services from sharing the same device environment as organizational information. The organization can limit access to only authorized business applications and services, reducing opportunities for data leakage and helping maintain a predictable security posture.

COBO also supports reliable access to company resources because IT can standardize device models, operating-system releases, network settings, identity certificates, VPN configurations, and required business apps. That standardization improves interoperability, policy enforcement, and supportability for corporate systems. CompTIA includes mobile-device deployment and management considerations within Security+ security architecture concepts; see CompTIA Security+. Guidance for securing enterprise mobile devices is also available from NIST SP 800-124 Rev. 2.

QUESTION NO: 50

A company is publishing an API for a mobile application. The company wants to limit each application to only the API operations it requires and reduce the impact of a stolen access token. Which of the following should the company implement? (Choose two.)

A. OAuth scopes

B. Short token expiration periods

C. Shared API credentials for all mobile applications

D. Disable TLS certificate validation

E. Use a static password in each API request

ANSWER: A B

Explanation:

OAuth scopes are appropriate because scopes limit the actions that an access token authorizes. For example, a token issued to a mobile application can be restricted to reading a user's profile instead of permitting administrative changes or access to unrelated APIs. This supports least privilege for API access.

Short token expiration periods reduce the window in which a stolen token can be used. After the token expires, the client must obtain a new token through the authorized authentication process. Token lifetimes should be balanced with operational needs and can be combined with revocation and monitoring capabilities. OAuth 2.0 defines scopes as a mechanism for limiting access-token permissions. See RFC 6749: The OAuth 2.0 Authorization Framework.

QUESTION NO: 51

Which of the following job roles would sponsor data quality and data entry initiatives that ensure business and regulatory requirements are met?

A. The data owner

B. The data processor

C. The data steward

D. The data privacy officer.

ANSWER: C

Explanation:

The data steward is the appropriate role because data stewardship focuses on the operational governance of organizational data. A data steward helps establish and maintain data definitions, quality standards, validation requirements, and processes for accurate data capture. This includes coordinating data-entry practices so records are complete, consistent, reliable, and suitable for their intended business use.

In a governance program, the data steward commonly acts as the business-facing champion for initiatives that improve the quality of a specific data domain, such as customer, employee, financial, or inventory data. The steward works with users, process owners, and technical teams to apply business rules to data collection and maintenance. By promoting consistent entry standards and monitoring data quality, the role helps the organization satisfy internal governance expectations as well as external regulatory obligations that depend on trustworthy records.

This responsibility aligns with established data-governance practice: data stewardship assigns accountability for the day-to-day definition, quality, and proper use of data assets. See IBM’s overview of data stewardship and the DAMA Body of Knowledge for data-management and governance concepts.

QUESTION NO: 52

A security engineer is implementing full-disk encryption (FDE) for all laptops in an organization. Which of the following are the most important for the engineer to consider as part of the planning process? (Choose two.)

A. Key escrow

B. TPM presence

C. Digital signatures

D. Data tokenization

E. Public key management

F. Certificate authority linking

ANSWER: A B

Explanation:

Key escrow and TPM presence are essential considerations when planning an organization-wide full-disk encryption deployment. Key escrow provides a controlled, centrally managed method to retain and retrieve recovery keys. This is critical for maintaining access to organizational data when a user forgets a preboot PIN, a device's boot configuration changes, a laptop is reassigned, or personnel leave the organization. Recovery keys must be protected using strong access controls, auditing, and documented recovery procedures so that recovery remains available without creating an unnecessary key-exposure risk. NIST defines key escrow as the secure storage of cryptographic keys for authorized recovery: NIST Key Escrow Glossary.

TPM presence must also be assessed across the laptop fleet because a Trusted Platform Module can protect disk-encryption key material in dedicated hardware and release it only when measured boot conditions match expected values. This helps bind decryption to the intended device and trusted startup state, strengthening protection against offline theft and tampering. During planning, the engineer should inventory TPM availability and version, verify that TPMs are enabled and provisioned in UEFI/BIOS, and establish a process for devices that lack suitable TPM support. The Trusted Computing Group provides an overview of TPM capabilities at TPM Summary.

QUESTION NO: 53

Which of the following would be the best ways to ensure only authorized personnel can access a secure facility? (Select two).

A. Fencing

B. Video surveillance

C. Badge access

D. Access control vestibule

E. Sign-in sheet

F. Sensor

ANSWER: C D

Explanation:

Badge access and an access control vestibule are the best combination for restricting entry to authorized personnel because they provide both credential validation and protection against unauthorized entry alongside an authorized person. Badge access is a preventive physical access control: a door controller checks the presented credential against its authorization rules before releasing the lock. This allows access to be granted according to an individual's identity, assigned permissions, time of day, and location. Organizations can also revoke a badge promptly when it is lost, stolen, or no longer needed, while maintaining entry logs for accountability and investigations.

An access control vestibule, commonly called a mantrap, strengthens that credential-based control at the doorway. Its interlocking doors allow the first door to close before the second door can open, creating a controlled space in which an entrant can be validated before proceeding. This design materially reduces the opportunity for tailgating or piggybacking and can support an additional identity check where the facility's risk level requires it. Together, badge access verifies permission to enter, while the vestibule enforces a one-person-at-a-time entry process. These are core examples of physical security and access-control mechanisms addressed in the CompTIA exam objectives and CISA's physical security guidance.

QUESTION NO: 54

Which of the following physical controls can be used to both detect and deter? (Choose two.)

A. Lighting

B. Fencing

C. Signage

D. Sensor

E. Bollard

F. Lock

ANSWER: A D

Explanation:

Lighting can serve both deterrent and detection purposes when deployed as security lighting. Well-lit entrances, parking areas, perimeters, and access paths make unauthorized activity more conspicuous and increase the perceived likelihood that an intruder will be seen. This visible exposure discourages attempted intrusion while also enabling guards, occupants, and surveillance systems to observe suspicious activity more effectively. In physical-security design, lighting is commonly used to support monitoring and assessment of activity around protected facilities.

Sensor is also both detective and deterrent. Door contacts, motion sensors, glass-break detectors, and similar devices detect physical conditions or unauthorized movement and can initiate alarms, notifications, or an automated response. When sensors, alarm placards, or protected-entry equipment are visible, they also discourage attempts by signaling that intrusion is likely to be identified quickly. NIST describes physical-access monitoring as a control for detecting and responding to physical-access events, while CISA identifies lighting and intrusion-detection measures as complementary elements of facility security. See NIST SP 800-53 Rev. 5 and CISA Physical Security.

QUESTION NO: 55

A company is planning to set up a SIEM system and assign an analyst to review the logs on a weekly basis. Which of the following types of controls is the company setting up?

A. Corrective

B. Preventive

C. Detective

D. Deterrent

ANSWER: C

Explanation:

Detective is correct because the SIEM and the analyst's scheduled log reviews are intended to identify indicators of malicious activity, policy violations, anomalous behavior, and potential security incidents. A security information and event management platform centralizes and correlates event data from sources such as endpoints, servers, network devices, applications, and identity systems. This visibility enables the organization to discover suspicious patterns, such as repeated failed authentication attempts, unexpected privilege changes, unusual network connections, or malware-related events.

The weekly review cadence does not change the control's classification. Although more frequent monitoring would generally improve the organization's ability to respond quickly, the stated activity still exists to discover events that have occurred or may be occurring. Detective controls provide the evidence and alerting needed for incident investigation, escalation, and response; they do not themselves primarily block an action or restore a system.

NIST describes audit and accountability controls that create, retain, review, and analyze audit records to support detection and investigation. SIEM capabilities operationalize these practices by aggregating and analyzing log data. See NIST SP 800-53 Rev. 5, Update 1 and the CISA information-sharing resources.

QUESTION NO: 56

Which of the following best practices gives administrators a set period to perform changes to an operational system to ensure availability and minimize business impacts?

A. Impact analysis

B. Scheduled downtime

C. Backout plan

D. Change management boards

ANSWER: B

Explanation:

Scheduled downtime is correct because it establishes a predefined maintenance window in which administrators can apply patches, upgrades, configuration changes, or other maintenance to a production system. Setting aside this specific period allows the organization to plan the work around business requirements, notify affected users and stakeholders, and avoid performing disruptive changes during peak operating hours. This makes service availability more predictable and reduces the operational effect of necessary system maintenance.

A scheduled maintenance window also supports disciplined change management by ensuring the responsible team has time to execute approved procedures, validate that the system is functioning as intended after the change, and restore normal service in a controlled manner. When changes are coordinated with system owners and business units, administrators can select an appropriate time based on service-level commitments, critical business processes, and resource availability. This directly addresses the requirement for a set period to make changes while minimizing business impact.

CompTIA includes change management processes and operational procedures among Security+ administrative and security-operations concepts. See the CompTIA exam objectives resource page and NIST guidance on maintaining controlled configuration and change processes in NIST SP 800-128.

QUESTION NO: 57

An engineer moved to another team and is unable to access the new team's shared folders while still being able to access the shared folders from the former team. After opening a ticket, the engineer discovers that the account was never moved to the new group. Which of the following access controls is most likely causing the lack of access?

A. Role-based

B. Discretionary

C. Time of day

D. Least privilege

ANSWER: A

Explanation:

Role-based access control (RBAC) is correct because permissions are assigned according to job functions or organizational roles, commonly implemented through directory-service groups. A team’s shared-folder permissions can be granted to its associated group rather than individually to every user. When the engineer changed teams, the account needed to be added to the new team’s group to inherit that group’s access rights. Because that group-membership update did not occur, the engineer did not receive authorization for the new team’s shared folders.

The continued ability to access the former team’s folders reinforces this conclusion: the account remains associated with the former group and therefore continues to inherit the permissions assigned to that role. This is a common identity and access management lifecycle issue during a transfer, where access should be provisioned for the new role and adjusted or removed for the previous role. NIST defines RBAC as controlling access based on roles and associated permissions, which directly matches access changing with team or group assignment. See the NIST RBAC glossary definition and NIST’s RBAC model publication.

QUESTION NO: 58

A security administrator is developing a backup strategy for a database that supports a critical business application. Which of the following should be included to improve the organization's ability to recover from ransomware? (Choose two.)

A. Maintaining an offline backup copy

B. Granting all users write access to backup repositories

C. Testing restoration procedures

D. Storing backup credentials in a shared spreadsheet

E. Disabling backup encryption

F. Using production administrator accounts for all backup operations

ANSWER: A C

Explanation:

Maintaining an offline backup copy is important because ransomware can encrypt or delete backups that remain continuously accessible from the production environment. An offline or otherwise isolated copy provides a recovery source that is more difficult for an attacker to alter.

Testing restoration procedures is also essential. A backup is useful only if the organization can restore the required data and systems within its recovery objectives. Testing validates backup integrity, recovery documentation, system dependencies, required credentials, and the time needed to return the application to service. CISA recommends maintaining offline, encrypted backups and regularly testing restoration procedures in its Ransomware Guide. See also NIST guidance on contingency planning.

QUESTION NO: 59

After a company was compromised, customers initiated a lawsuit. The company's attorneys have requested that the security team initiate a legal hold in response to the lawsuit. Which of the following describes the action the security team will most likely be required to take?

A. Retain the emails between the security team and affected customers for 30 days.

B. Retain any communications related to the security breach until further notice.

C. Retain any communications between security members during the breach response.

D. Retain all emails from the company to affected customers for an indefinite period of time.

ANSWER: B

Explanation:

Retain any communications related to the security breach until further notice. A legal hold, also called a litigation hold, is a directive to preserve potentially relevant electronically stored information and records when litigation is pending or reasonably anticipated. In this scenario, the customer lawsuit makes breach-related communications potentially discoverable evidence. The security team must therefore stop normal deletion, alteration, or rotation processes for relevant data and preserve it until the legal department releases or modifies the hold. Relevant communications can include email, chat messages, incident tickets, investigation notes, alerts, reports, and correspondence involving internal teams, counsel, vendors, or affected parties. Preservation should be broad enough to protect information that may relate to the incident, response activities, scope, timelines, decisions, and customer impact. The hold remains in effect based on legal direction rather than a preselected retention period, because disposing of relevant records after receiving notice could create legal and evidentiary risk. The U.S. Federal Rules of Civil Procedure address the preservation of electronically stored information in anticipation of or during litigation; see Federal Rule of Civil Procedure 37(e). For practical legal-hold guidance, see the EDRM framework.

QUESTION NO: 60

Which of the following factors are the most important to address when formulating a training curriculum plan for a security awareness program? (Choose two.)

A. Channels by which the organization communicates with customers

B. The reporting mechanisms for ethics violations

C. Threat vectors based on the industry in which the organization operates

D. Secure software development training for all personnel

E. Cadence and duration of training events

F. Retraining requirements for individuals who fail phishing simulations

ANSWER: C E

Explanation:

“Threat vectors based on the industry in which the organization operates” is essential because an effective awareness curriculum begins with an assessment of the organization’s risks, operating environment, workforce roles, and likely attacks. Training should prepare personnel to recognize and respond to the threats that are most relevant to the organization, such as targeted phishing, business email compromise, ransomware, removable-media risks, or sector-specific social-engineering techniques. This risk-based approach makes learning objectives meaningful and supports measurable reductions in human-related security risk.

“Cadence and duration of training events” is also a core curriculum-planning consideration. Awareness is an ongoing behavior-change program rather than a one-time presentation. Defining how often training occurs, how long each activity should take, and when refresher material or reinforcement is delivered helps ensure personnel can complete the program without excessive operational disruption while retaining and applying the content over time. A planned cadence also enables the organization to align training with emerging threats, policy updates, simulations, and metrics gathered from prior events. NIST’s awareness, training, and education guidance emphasizes establishing a structured, role- and risk-informed program with recurring activities and evaluation. See NIST SP 800-50 Rev. 1 and CISA’s Cybersecurity Awareness Program guidance.

QUESTION NO: 61

Which of the following are cases in which an engineer should recommend the decommissioning of a network device? (Select two).

A. The device has been moved from a production environment to a test environment.

B. The device is configured to use cleartext passwords.

C. The device is moved to an isolated segment on the enterprise network.

D. The device is moved to a different location in the enterprise.

E. The device's encryption level cannot meet organizational standards.

F. The device is unable to receive authorized updates.

ANSWER: E F

Explanation:

The device's encryption level cannot meet organizational standards is a valid reason for decommissioning because encryption capabilities are often constrained by the device’s hardware, operating system, and supported firmware. If a device cannot support the organization’s approved cryptographic protocols, cipher suites, or key strengths, it may be unable to protect management access and network traffic in accordance with security policy, regulatory obligations, or risk-management requirements. Replacing or retiring the device removes a persistent exposure rather than relying on an unsupported workaround.

The device is unable to receive authorized updates is also a clear decommissioning trigger. Network devices require vendor-authorized firmware and security updates to correct vulnerabilities and maintain secure operation. Once a device reaches end of support, newly discovered flaws may remain permanently unpatched, leaving the organization without a sustainable

way to manage the associated risk. A supported replacement enables ongoing patch management, vulnerability remediation, and lifecycle governance. CISA recommends maintaining network infrastructure through secure configuration and timely updates, while NIST emphasizes applying software updates to address known security weaknesses. See CISA guidance on securing network infrastructure devices and NIST guidance on software updates.

QUESTION NO: 62

A staff member finds a USB drive in the office ' s parking lot. Which of the following should the staff member do?

A. Notify the file owner after reviewing the contents of the drive.

B. Use an air-gapped system to open the files without exposing the network.

C. Wipe the drive immediately using a secure method.

D. Submit the device to the security team without connecting it.

ANSWER: D

Explanation:

Submit the device to the security team without connecting it. An unattended USB drive from an unknown source must be treated as potentially malicious. This situation is a well-known example of removable-media baiting, in which an attacker deliberately leaves a device in a location where a curious employee is likely to find and connect it. The device could contain malware, exploit firmware or USB interface behavior, or present deceptive files intended to obtain credentials or cause an employee to bypass security procedures. Connecting it to a workstation—even briefly—can expose the endpoint and potentially the organization to compromise.

The appropriate employee action is to avoid inserting, browsing, modifying, or destroying the device and to promptly transfer it through the organization’s established incident-reporting and evidence-handling process. The security team can determine whether the device requires forensic examination, isolation, documentation, or other response actions using approved tools and procedures. This approach supports the Security+ objectives of recognizing social-engineering indicators, following security policies, and escalating suspected incidents to the proper personnel. CISA specifically advises organizations to prevent use of unknown removable media and to establish controls for removable devices. See CISA guidance on USB drive safety and CompTIA exam objectives resources.

QUESTION NO: 63

Which of the following are cases in which an engineer should recommend the decommissioning of a network device? (Choose two.)

A. The device has been moved from a production environment to a test environment.

B. The device is configured to use cleartext passwords.

C. The device is moved to an isolated segment on the enterprise network.

D. The device is moved to a different location in the enterprise.

E. The device's encryption level cannot meet organizational standards.

F. The device is unable to receive authorized updates.

ANSWER: E F

Explanation:

The device's encryption level cannot meet organizational standards and the device is unable to receive authorized updates are appropriate reasons to recommend decommissioning. A network device that cannot support the organization’s required cryptographic protocols, cipher suites, key lengths, or certificate capabilities creates a persistent security and compliance concern. Modern cryptographic requirements protect the confidentiality and integrity of management traffic, user traffic, and

device-to-device communications. When a device's hardware or supported firmware cannot meet those requirements, replacement is the sustainable remediation.

A device that can no longer receive authorized vendor updates should also be retired. Security patches and signed firmware updates address newly discovered vulnerabilities, including flaws that may permit unauthorized access, disruption, or compromise of the network infrastructure. Once a product reaches end of support, known vulnerabilities can remain permanently uncorrected. Although compensating controls may temporarily reduce exposure while a replacement is planned, they do not restore the device to a supported, patchable state. Asset lifecycle management should therefore include removing unsupported infrastructure and securely sanitizing its configuration and stored credentials before disposal or reassignment. NIST identifies supported, secure configuration and maintenance practices as important elements of system security; see NIST SP 800-53 Rev. 5. CISA also provides practical lifecycle and patching guidance at CISA Patching.

QUESTION NO: 64

Which of the following are the best for hardening end-user devices? (Selecttwo)

A. Full disk encryption

B. Group-level permissions

C. Account lockout

D. Endpoint protection

E. Proxy server

F. Segmentation

ANSWER: A D

Explanation:

Full disk encryption and **Endpoint protection** are core, complementary controls for hardening end-user devices such as laptops, desktops, and mobile endpoints. Full disk encryption protects data at rest across the device's storage. If an endpoint is lost, stolen, or otherwise accessed without authorization while powered off, encryption helps prevent an attacker from reading the files and operating-system data simply by removing or booting from the drive. Proper deployment includes secure key management and recovery procedures so that protection remains practical for the organization.

Endpoint protection provides host-level defensive capabilities, commonly including antimalware detection, behavioral analysis, exploit prevention, endpoint detection and response functions, and host firewall enforcement. These capabilities reduce the chance that malicious code, ransomware, or suspicious activity can compromise the endpoint, while also providing security teams with visibility and response capabilities. Together, protecting stored information through full disk encryption and actively defending the operating endpoint through endpoint protection address major risks to end-user devices and align with Security+ device-hardening and organizational risk-mitigation concepts. CompTIA's current Security+ exam information is available at CompTIA Security+; NIST guidance on storage encryption is available in NIST SP 800-111.

QUESTION NO: 65

Which of the following is most important to maintain when it comes to evidence during a digital forensic investigation?

A. Change management

B. Availability

C. Chain of custody

D. Reporting

ANSWER: C

Explanation:

Chain of custody is most important to maintain because it establishes a documented, chronological record of evidence handling from collection through presentation or disposition. For each item of digital evidence, this record identifies who collected it, the date and time of collection, the source or location, each person who subsequently possessed or accessed it, transfers between custodians, and relevant storage or handling details. This documentation supports the ability to demonstrate that the evidence presented is the same evidence originally acquired and that it has remained under controlled handling. Maintaining this record is central to evidentiary integrity, accountability, and legal admissibility. Digital evidence can be copied, modified, or contaminated easily, so forensic personnel also use practices such as cryptographic hashes and write blockers to support integrity; chain-of-custody documentation connects those technical safeguards to accountable human processes. A complete record enables an investigator or court to trace the evidence lifecycle and assess whether handling procedures preserved its reliability. NIST's guidance on integrating forensic techniques emphasizes documenting the chain of custody and preserving evidence integrity throughout the investigation. See NIST SP 800-86, Guide to Integrating Forensic Techniques into Incident Response and the NIST Computer Forensics Tool Testing program.

QUESTION NO: 66

An organization experiences a cybersecurity incident involving a command-and-control server. Which of the following logs should be analyzed to identify the impacted host? (Select two).

A. Application

B. Authentication

C. DHCP

D. Network

E. Firewall

F. Database

ANSWER: C E

Explanation:

DHCP and **Firewall** logs should be correlated to identify the affected endpoint in a command-and-control incident. Firewall logs record connections crossing the network boundary and can reveal an internal source IP address communicating with the known or suspected command-and-control infrastructure, along with relevant timestamps, ports, protocols, connection outcomes, and traffic volumes. This provides the key network evidence linking activity to an internal address.

DHCP logs provide the attribution needed when internal addresses are assigned dynamically. By matching the source IP address and time observed in the firewall records with DHCP lease history, investigators can determine the device identity associated with that address, commonly through its hostname, MAC address, client identifier, and assigned lease period. This correlation is especially important because the same IP address can be reassigned to different devices over time. Together, these records establish a defensible connection between suspicious external command-and-control traffic and the specific impacted host. NIST guidance emphasizes collecting and correlating relevant event-log data during incident handling; see NIST SP 800-92, Guide to Computer Security Log Management and CISA Malware guidance.

QUESTION NO: 67

A company needs to keep the fewest records possible, meet compliance needs, and ensure destruction of records that are no longer needed. Which of the following best describes the policy that meets these requirements?

A. Security policy

B. Classification policy

C. Retention policy

D. Access control policy

ANSWER: C

Explanation:

Retention policy is correct because it establishes the organization's approved records life cycle: which records must be retained, the minimum and maximum periods for keeping them, and the approved process for disposition once they are no longer required. By defining retention periods according to legal, regulatory, contractual, operational, and evidentiary requirements, the organization can preserve required information without maintaining unnecessary records indefinitely. This supports data minimization by limiting stored records to information with a continuing business or compliance purpose.

A sound retention policy also defines how expired records are disposed of securely and consistently. Depending on the medium and data sensitivity, this may include secure deletion, cryptographic erasure, physical destruction, or shredding. Documented disposition procedures help ensure that records are not retained beyond their authorized period and that their destruction can be demonstrated during an audit. NIST's control framework addresses records retention and disposal requirements under media protection, while the GDPR's storage-limitation principle requires personal data to be kept no longer than necessary for its processing purpose. See NIST SP 800-53 Rev. 5 and GDPR Article 5.

QUESTION NO: 68

Which of the following has been implemented when a host-based firewall on a legacy Linux system allows connections from only specific internal IP addresses?

A. Compensating control

B. Network segmentation

C. Transfer of risk

D. SNMP traps

ANSWER: A

Explanation:

Compensating control is correct because the host-based firewall restriction provides an alternate safeguard for a legacy Linux system whose preferred security remediation may not be feasible. Legacy systems can have technical or operational limitations that prevent timely upgrades, patches, agent deployment, or replacement. In that situation, allowing network connections only from explicitly approved internal IP addresses reduces the system's exposure without requiring changes that the older platform may not support.

This is a risk-based control choice: the firewall enforces a narrow communication path and limits the number of hosts that can reach the legacy asset. The control does not remove the underlying legacy-system risk, but it meaningfully reduces the likelihood that unauthorized or unneeded systems can access vulnerable services. Its purpose is therefore to provide protection that compensates for a security capability or remediation that cannot currently be implemented. Such restrictions should be documented, monitored, and reviewed as part of the organization's risk-management process until the system can be modernized or retired.

NIST describes security controls and their role in reducing organizational risk in NIST SP 800-53 Rev. 5. NIST's definition of a compensating security control also supports using an alternate control when a primary control cannot be implemented.

QUESTION NO: 69

A company requires hard drives to be securely wiped before sending decommissioned systems to recycling. Which of the following best describes this policy?

A. Enumeration

B. Sanitization

C. Destruction

D. Inventory

ANSWER: B

Explanation:

Sanitization is the correct term for a policy requiring hard drives to be securely wiped before systems are recycled. Media sanitization is the process of applying an approved technique to render data on storage media inaccessible and unrecoverable at a level appropriate to the organization's risk, the sensitivity of the information, and the planned disposition of the device. For drives that will leave organizational control for recycling or reuse, the organization should define and document the approved sanitization method, verify that the method completed successfully, and retain appropriate records of the action.

NIST describes sanitization as a set of processes that can include clearing, purging, or physical destruction, selected according to the media type and confidentiality requirements. Secure wiping is therefore a sanitization control: it addresses residual data that could otherwise remain on a decommissioned drive and be exposed after disposal. A sound policy also assigns responsibility for the process and requires validation before the device enters the recycling stream. See NIST SP 800-88 Rev. 1, Guidelines for Media Sanitization and NIST's publication record for SP 800-88 Rev. 1.

QUESTION NO: 70

Which of the following is an effective tool to stop or prevent the exfiltration of data from a network?

A. DLP

B. NIDS

C. TPM

D. FDE

ANSWER: A

Explanation:

DLP is the appropriate control because Data Loss Prevention technologies are specifically designed to identify, monitor, and enforce policies governing sensitive data as it is used, stored, and transmitted. To prevent network exfiltration, a DLP solution can inspect outbound email, web uploads, cloud-service transfers, and other egress channels for defined sensitive-data patterns, such as personally identifiable information, payment-card data, health information, intellectual property, or classified document labels. When a policy match occurs, DLP can block the transfer, quarantine the content, require justification or encryption, and alert security personnel. These capabilities make it a preventative control for unauthorized movement of data outside an organization rather than merely a means of detecting anomalous network activity. Effective DLP implementation begins with data classification and clearly defined handling policies, then applies those policies across endpoint, network, email, and cloud environments as appropriate. This directly supports Security+ objectives concerning data protection and the prevention of data leakage or exfiltration. The Cybersecurity and Infrastructure Security Agency describes DLP as a strategy and set of tools for detecting and preventing unauthorized transmission of sensitive information. See CISA's Data Loss Prevention guidance and Microsoft Purview Data Loss Prevention documentation.

QUESTION NO: 71

A security analyst is collecting a laptop as evidence during an internal investigation. Which of the following actions are the most important to preserve the integrity of the evidence? (Choose two.)

A. Document the chain of custody.

B. Calculate a cryptographic hash of the forensic image.

C. Install the latest operating system updates.

D. Delete temporary files from the laptop.

E. Allow the employee to review the collected evidence.

ANSWER: A B

Explanation:

Documenting the chain of custody is essential because it records who collected, transferred, accessed, stored, and analyzed the laptop. This documentation establishes accountability and supports the ability to demonstrate that evidence was controlled throughout the investigation.

Calculating a cryptographic hash of the forensic image is also essential because the hash provides a value that can be compared later to verify that the acquired image has not changed. Investigators should perform analysis on a verified forensic copy rather than modifying the original evidence. NIST describes preserving evidence integrity and documenting evidence-handling activities in NIST SP 800-86.

QUESTION NO: 72

A Chief Information Officer wants to ensure that network devices cannot connect to the public internet or the local network to directly perform firmware updates. The IT team must manually perform the update process by using a portable device. Which of the following architecture types best fits this description?

A. Microservices

B. Air-gapped

C. Software-defined networking

D. Serverless

ANSWER: B

Explanation:

Air-gapped is the correct architecture because it isolates a system or network from external connectivity, including the public internet and the organization's normal local network. In this scenario, the network devices cannot retrieve firmware directly from a vendor site, management server, or other connected resource. Instead, administrators must obtain, validate, and transport the firmware to the isolated environment using a portable device or removable media. That manual transfer requirement is a key operational characteristic of an air-gapped environment.

Air gaps are commonly used where minimizing exposure to remote attacks is especially important, such as industrial control systems, critical infrastructure, classified systems, and highly sensitive operational technology. Although isolation reduces the remote attack surface, the update process still requires disciplined controls: firmware should be obtained from an authorized source, its integrity and authenticity should be verified using vendor-provided hashes or digital signatures, and portable media should be controlled and scanned before use. These practices preserve the security benefits of isolation while allowing necessary maintenance. NIST describes air-gapped systems as isolated from external networks and recognizes the importance of secure media handling in protected environments. See NIST's air gap definition and CISA guidance on industrial control systems.

QUESTION NO: 73

A business is expanding to a new country and must protect customers from accidental disclosure of specific national identity information. Which of the following should the security engineer update to best meet business requirements?

A. SIEM

B. SCAP

C. DLP

D. WAF

ANSWER: C

Explanation:

DLP is the appropriate technology to update because it is designed to identify and prevent unauthorized or accidental exposure of sensitive information, including national identity numbers and other personally identifiable information. A security engineer can configure DLP policies with detection rules tailored to the identity-data format and requirements of the new country, such as pattern matching, checksums, keywords, data labels, and exact-data matching. These policies can monitor sensitive information in email, endpoint activity, cloud collaboration services, and network traffic, then alert, quarantine, encrypt, redact, or block a transfer when policy conditions are met. This directly supports the business requirement to protect customers against accidental disclosure rather than merely recording an event after it occurs. DLP also provides reporting that helps demonstrate that controls are being applied consistently as the organization expands into a new regulatory jurisdiction. Microsoft's overview of data loss prevention describes how DLP policies detect sensitive information and help prevent inappropriate sharing across organizational locations: Microsoft Learn: Learn about data loss prevention. For broader context, the NIST Privacy Framework identifies data processing controls as an important means of managing privacy risk: NIST Privacy Framework.

QUESTION NO: 74

A security team is addressing a risk associated with the attack surface of the organization's web application over port 443. Currently, no advanced network security capabilities are in place. Which of the following would be best to set up? (Choose two.)

A. NIDS

B. Honeypot

C. Certificate revocation list

D. HIPS

E. WAF

F. SIEM

ANSWER: D E

Explanation:

WAF and HIPS provide complementary preventive controls for an internet-facing web application using HTTPS on port 443. A WAF is specifically designed to sit in front of a web application and evaluate HTTP/S traffic at the application layer. It can enforce rules that identify and block common web attacks, including injection attempts, cross-site scripting payloads, malicious requests, and abnormal request patterns. Because it understands web protocols, a WAF directly reduces the exposed application attack surface while allowing legitimate HTTPS traffic to reach the application.

HIPS adds protection on the web server itself, where HTTPS traffic has already been decrypted and application activity can be evaluated. A host-based intrusion prevention system can detect and actively stop malicious behavior through host telemetry, such as suspicious processes, exploit behavior, unauthorized changes, or malicious activity directed at the web-server environment. Deploying HIPS therefore provides a prevention layer even if a harmful request bypasses edge filtering or exploits an application weakness. Together, WAF and HIPS implement defense in depth at the web edge and on the protected host. See the OWASP Web Application Firewall guidance and CISA guidance on intrusion detection and prevention systems.

QUESTION NO: 75

During a security incident, the security operations team identified sustained network traffic from a malicious IP address:

10.1.4.9. A security analyst is creating an inbound firewall rule to block the IP address from accessing the organization?s network. Which of the following fulfills this request?

A. access-list inbound deny ip source 0.0.0.0/0 destination 10.1.4.9

B. access-list inbound deny ip source 10.1.4.9 destination 0.0.0.0/0

C. access-list inbound permit ip source 10.1.4.9 destination 0.0.0.0/0

D. access-list inbound permit ip source 0.0.0.0/0 destination 10.1.4.9

ANSWER: B

Explanation:

access-list inbound deny ip source 10.1.4.9 destination 0.0.0.0/0 fulfills the stated goal because it creates a deny rule for traffic whose source is the identified malicious host, 10.1.4.9. In an inbound filtering context, the source address represents the remote system sending traffic toward the organization. Matching that source directly is therefore the appropriate way to stop packets originating from the known malicious address before they reach protected internal resources.

The destination prefix `0.0.0.0/0` represents any IPv4 destination. Using it ensures the block applies regardless of which organizational host, subnet, or service the malicious system attempts to contact. This is appropriate when the incident response requirement is to block the source IP from accessing the organization's network generally, rather than only restricting a particular destination or port. The `deny` action causes matching packets to be discarded. In practice, the rule should be placed where it is evaluated before a broader allow rule that could otherwise permit the traffic, and responders should document and monitor the block as part of incident containment. Cisco's access-control-list guidance describes matching traffic using source and destination addresses and applying permit or deny actions: Cisco ACL configuration guidance. NIST also identifies containment as a core incident-response activity: NIST SP 800-61 Rev. 2.

QUESTION NO: 76

A company plans to secure its systems by:

• Preventing users from sending sensitive data over corporate email

• Restricting access to potentially harmful websites

Which of the following features should the company set up? (Choose two.)

A. DLP software

B. DNS filtering

C. File integrity monitoring

D. Stateful firewall

E. Guardrails

F. Antivirus signatures

ANSWER: A B

Explanation:

DLP software is appropriate because data loss prevention policies inspect data in motion, including outbound corporate email messages and attachments. Administrators can define sensitive-data patterns and classifications—such as personally identifiable information, payment-card data, intellectual property, or confidential document labels—and configure the system to warn users, block delivery, quarantine messages, or require additional justification before a message leaves the organization. This directly supports prevention of unintended or unauthorized disclosure through email.

DNS filtering is appropriate for restricting access to potentially harmful websites. When a user or endpoint attempts to resolve a domain name, the filtering service can compare that request with threat-intelligence and category policies. Domains associated with phishing, malware distribution, command-and-control activity, or prohibited web categories can be blocked before the user establishes a connection to the destination. DNS filtering is therefore a practical preventive control for reducing exposure to malicious and inappropriate websites across managed devices and networks.

These controls address the two stated goals at complementary layers: DLP software governs sensitive information leaving the organization through email, while DNS filtering controls access to unsafe internet destinations. See CISA's Data Loss Prevention resource and Cloudflare's DNS filtering overview.

QUESTION NO: 77

Which of the following methods can be used to detect attackers who have successfully infiltrated a network? (Choose two.)

A. Tokenization

B. CI/CD

C. Honeypots

D. Threat modeling

E. DNS sinkhole

F. Data obfuscation

ANSWER: C E

Explanation:

Honeypots and a **DNS sinkhole** are methods that can reveal malicious activity after an adversary has gained access to, or compromised a device within, a network. A honeypot is a deliberately deployed decoy asset, such as a host, service, credential, or share, designed to appear valuable to an attacker. Because it should have no normal business use, connections, authentication attempts, scans, or changes involving the honeypot create a high-confidence alert for investigation. It can also provide security teams with useful telemetry about attacker tools, techniques, and lateral-movement behavior.

A DNS sinkhole identifies potentially compromised internal hosts by intercepting DNS requests for known malicious domains, particularly command-and-control infrastructure. Rather than allowing the client to reach the requested destination, the sinkhole directs the request to a controlled address and records the originating system. These records enable defenders to identify endpoints attempting to communicate with malicious infrastructure and to begin containment and remediation. Together, deception telemetry and DNS monitoring provide meaningful detection coverage for intrusions that have already progressed into the environment. See CISA's overview of honeypots and Infoblox's DNS sinkhole definition.

QUESTION NO: 78

An organization wants to implement a secure solution for remote users. The users handle sensitive PHI on a regular basis and need to access an internally developed corporate application. Which of the following best meet the organization's security requirements? (Choose two.)

A. Local administrative password

B. Perimeter network

C. Jump server

D. WAF

E. MFA

F. VPN

ANSWER: E F

Explanation:

MFA and VPN together best support secure remote access to an internal application used to handle protected health information. A VPN establishes an encrypted connection across an untrusted network, protecting the confidentiality and integrity of traffic between the remote endpoint and organizational resources. This reduces exposure of sensitive application sessions and data while personnel work outside the corporate network. VPN access can also be configured to require authenticated users and managed devices before permitting access to the internal environment.

MFA provides stronger assurance that the individual requesting access is the authorized user. It requires an additional authentication factor beyond a password, such as a hardware security key, authenticator application approval, or biometric factor. This is particularly important for PHI because compromised passwords alone should not enable access to sensitive systems. Combining MFA with VPN access implements defense in depth: the VPN protects the communication path, while MFA strengthens identity verification before the internal corporate application can be reached. CISA recommends phishing-resistant MFA as a key protection for accounts and sensitive resources, and NIST describes VPNs as protected communications over public networks. See CISA guidance on phishing-resistant MFA and NIST guidance on secure remote access.

QUESTION NO: 79

Which of the following organizational documents is most often used to establish and communicate expectations associated with integrity and ethical behavior within an organization?

A. AUP

B. SLA

C. EULA

D. MOA

E. Code of ethics

ANSWER: E

Explanation:

A code of ethics is the organizational document used to formally establish and communicate expectations for integrity and ethical behavior. It defines the values and principles that guide personnel decisions and conduct, such as acting honestly, avoiding conflicts of interest, protecting confidential information, treating stakeholders fairly, reporting suspected misconduct, and complying with applicable laws and organizational requirements. By documenting these standards, leadership provides a consistent baseline for employee behavior and supports a culture of accountability. A code of ethics is also commonly incorporated into security-awareness, onboarding, governance, and compliance activities so that personnel understand their responsibility to act responsibly when handling organizational systems, information, and assets. In Security+ governance concepts, ethical behavior is an important administrative and organizational control because security depends on trustworthy actions by people as well as technical safeguards. CompTIA's exam-objectives resources cover governance, risk, compliance, and security-program concepts relevant to organizational policies and expectations: CompTIA Exam Objectives. The U.S. Office of Government Ethics likewise describes ethics standards as principles intended to promote integrity in organizational service: Standards of Ethical Conduct.

QUESTION NO: 80

An organization's web servers host an online ordering system. The organization discovers that the servers are vulnerable to a malicious JavaScript injection, which could allow attackers to access customer payment information. Which of the following mitigation strategies would be most effective for preventing an attack on the organization's web servers? (Choose two.)

A. Regularly updating server software and patches

B. Implementing strong password policies

C. Encrypting sensitive data at rest and in transit

D. Utilizing a web-application firewall

E. Performing regular vulnerability scans

F. Removing payment information from the servers

ANSWER: A D

Explanation:

Regularly updating server software and patches is a key preventive control because patch management remediates known vulnerabilities in the web server, application framework, runtime environment, content-management components, and dependencies that may permit JavaScript injection. Maintaining a defined process to identify, test, and promptly deploy security updates reduces the attack surface and removes exploitable flaws before an attacker can use them against the ordering system.

Utilizing a web-application firewall provides a compensating and defense-in-depth control at the HTTP/S layer. A WAF can inspect requests for patterns associated with web attacks, including malicious script payloads and attempts to inject untrusted content into application inputs. It can block or alert on suspicious requests before they reach the web application. This is particularly valuable while remediation is being developed, tested, and deployed, and it continues to add protection against known attack signatures and abnormal request behavior.

Together, patching addresses the underlying vulnerable components, while a WAF supplies an enforcement layer in front of the application. This aligns with OWASP guidance on reducing web-application risk through secure maintenance and protective controls. See the OWASP Top 10 and the OWASP Web Application Firewall project.

QUESTION NO: 81

Which of the following techniques can be used to sanitize the data contained on a hard drive while allowing for the hard drive to be repurposed?

A. Degaussing

B. Drive shredder

C. Retention platform

D. Wipe tool

ANSWER: D

Explanation:

Wipe tool
is correct because it sanitizes a hard drive by logically clearing data while preserving the physical device for continued use. For a traditional magnetic hard disk, an approved overwrite-based sanitization process writes specified patterns across the addressable storage area, preventing practical recovery of the previous contents. Once the wipe completes successfully and is verified, the organization can reimage, redeploy, donate, or otherwise repurpose the drive according to its asset-management policy.

The exact sanitization procedure should match the media type, data sensitivity, organizational policy, and reuse destination. NIST describes clearing as a sanitization method intended to protect data from simple, noninvasive recovery techniques while allowing the media to remain usable. For hard disk drives that will stay under organizational control or be repurposed at an appropriate risk level, a properly implemented and validated wipe is therefore an appropriate approach. The process should include verification and documentation of the sanitization result. NIST's media-sanitization guidance is available in NIST SP 800-88 Rev. 1, and CISA provides complementary disposal and sanitization guidance in its Media Destruction Guidance.

QUESTION NO: 82

Which of the following elements of digital forensics should a company use if it needs to ensure the integrity of evidence?

A. Preservation

B. E-discovery

C. Acquisition

D. Containment

ANSWER: A

Explanation:

Preservation is the correct digital-forensics element because it protects evidence from alteration and maintains its evidentiary integrity from the time it is identified through examination and possible legal proceedings. Effective preservation practices include isolating the original media, restricting access, maintaining a complete chain of custody, and creating forensic copies rather than conducting analysis on the original evidence. Investigators also calculate and record cryptographic hash values, such as SHA-256 hashes, for collected files or forensic images. Recalculating and comparing those values later provides demonstrable assurance that the evidence has remained unchanged.

Preservation must be planned and documented carefully: the organization should record who handled the evidence, when and where transfers occurred, the tools and methods used, and the hash results associated with each item. These records allow the organization to establish that the evidence presented for analysis is the same evidence originally collected. NIST describes preserving data integrity and documenting handling procedures as essential components of sound forensic processes. For additional guidance, see NIST SP 800-86, Guide to Integrating Forensic Techniques into Incident Response and NIST's publication record for the guide.

QUESTION NO: 83

An organization ' s web servers host an online ordering system. The organization discovers that the servers are vulnerable to a malicious JavaScript injection, which could allow attackers to access customer payment information. Which of the following mitigation strategies would be most effective for preventing an attack on the organization ' s web servers? (Select two).

A. Regularly updating server software and patches

B. Implementing strong password policies

C. Encrypting sensitive data at rest and in transit

D. Utilizing a web-application firewall

E. Performing regular vulnerability scans

F. Removing payment information from the servers

ANSWER: A D

Explanation:

"Regularly updating server software and patches" is correct because prompt remediation removes known weaknesses in the web-server platform, application framework, runtime components, and supporting libraries that could be exploited to deliver or enable script-injection attacks. A disciplined patch-management process includes maintaining an asset inventory, monitoring vendor security advisories, testing updates appropriately, and deploying security fixes on a defined schedule. This reduces the attack surface before an attacker can take advantage of a disclosed flaw.

"Utilizing a web-application firewall" is also correct because a WAF provides a protective control specifically at the HTTP/S application layer. It can inspect requests and responses and apply rules designed to detect and block common web-application attack patterns, including suspicious script payloads associated with cross-site scripting. A WAF is particularly valuable as a compensating control while code fixes are being developed and deployed, and as defense in depth after remediation. Effective WAF operation requires tuning rules to the organization's application behavior and reviewing alerts so that malicious requests are blocked without unnecessarily disrupting legitimate customer orders. These controls align with the Security+ emphasis on secure application and infrastructure operations. See OWASP's Cross Site Scripting guidance and CompTIA exam objectives resources.

QUESTION NO: 84

An employee decides to collect PII data from the company ' s system for personal use. The employee compresses the data into a single encrypted file before sending the file to their personal email. The security department becomes aware of the attempted misuse and blocks the attachment from leaving the corporate environment. Which of the following types of employee training would most likely reduce the occurrence of this type of issue?

(Select two).

A. Privacy legislation

B. Social engineering

C. Risk management

D. Company compliance

E. Phishing

F. Remote work

ANSWER: A D

Explanation:

Privacy legislation and company compliance are the most relevant training areas because the incident is an intentional insider attempt to remove personally identifiable information from authorized corporate systems for a nonbusiness purpose. Privacy-legislation training establishes employees' responsibilities for collecting, accessing, using, retaining, and disclosing personal data. It helps personnel understand that personal data is subject to legal and regulatory protections, that access must be limited to a legitimate business need, and that unauthorized disclosure can create serious consequences for affected individuals and the organization. The Federal Trade Commission describes privacy and data-security expectations for businesses at FTC Privacy and Security.

Company-compliance training translates those external and internal requirements into practical employee obligations: acceptable-use rules, data-classification and handling procedures, restrictions on personal email or unauthorized storage, reporting responsibilities, and disciplinary consequences for misuse. This training directly addresses the behavior shown by reinforcing that an employee's ability to access a system does not grant permission to extract or use its data personally. These topics align with Security+ security-program and governance outcomes, including compliance, privacy, and user awareness, as reflected in CompTIA's current Security+ exam objectives: CompTIA Exam Objectives.

QUESTION NO: 85

A security analyst sees an increase of vulnerabilities on workstations after a deployment of a company group policy. Which of the following vulnerability types will the analyst most likely find on the workstations?

A. Misconfiguration

B. Zero-day

C. Malicious update

D. Supply chain

ANSWER: A

Explanation:

Misconfiguration is the most likely vulnerability type because Group Policy centrally applies security and system settings across many workstations. If a policy is deployed with insecure values, missing restrictions, excessive permissions, disabled security controls, or unintended configuration changes, every affected endpoint can inherit the same weakness. This can quickly produce a measurable increase in findings from vulnerability scans, configuration-compliance tools, or endpoint-security monitoring. Examples include a policy that weakens password or lockout settings, enables unnecessary services,

changes firewall rules, grants local administrative rights, or disables protective features. The key indicator is the timing and scope: vulnerabilities increased after a company Group Policy deployment, which points to an erroneous or insecure configuration being propagated at scale. Microsoft describes Group Policy as infrastructure for centrally managing configuration settings for users and computers in an Active Directory environment, making it especially important to validate policies before broad deployment and monitor their effects afterward. See Microsoft's Group Policy overview and Microsoft's Security Compliance Toolkit documentation for information on centrally managed settings and security baselines.

QUESTION NO: 86

A security engineer is implementing FDE for all laptops in an organization. Which of the following are the most important for the engineer to consider as part of the planning process? (Choose two.)

A. Key escrow

B. TPM presence

C. Digital signatures

D. Data tokenization

E. Public key management

F. Certificate authority linking

ANSWER: A B

Explanation:

Key escrow and TPM presence are essential planning considerations for an organization-wide full-disk encryption deployment. Key escrow provides a securely managed recovery path for encryption recovery keys. This is critical when a user forgets a preboot authentication factor, a device has a hardware or boot configuration change, or IT must access a device during support, reassignment, or incident response. The organization should define who can retrieve recovery material, require appropriate authorization and auditing, and protect escrowed keys from unauthorized disclosure.

TPM presence must be assessed across the laptop fleet because a Trusted Platform Module can protect disk-encryption key material in hardware and can bind key release to measured boot conditions. This supports an effective, manageable deployment model for products such as BitLocker, including the option to require an additional PIN for stronger preboot authentication. Inventorying TPM availability, version, ownership, activation status, and firmware posture allows the engineer to identify exceptions and choose suitable controls for systems that cannot use the intended TPM-backed configuration. Microsoft describes TPM-based BitLocker protection and recovery planning in its BitLocker documentation. NIST also emphasizes managing cryptographic keys throughout their lifecycle in NIST SP 800-57 Part 1.

QUESTION NO: 87

Which of the following is best used to detect fraud by assigning employees to different roles?

A. Least privilege

B. Mandatory vacation

C. Separation of duties

D. Job rotation

ANSWER: D

Explanation:

Job rotation is correct because it periodically moves employees among roles or responsibilities, allowing a different person to perform and review work that may previously have been controlled by one individual. This is especially valuable for fraud detection because concealed irregularities, unauthorized changes, or undocumented workarounds are more likely to surface

when another employee assumes the function and must reconcile records, follow established procedures, or complete the same operational tasks. Rotation reduces the likelihood that an employee can retain exclusive familiarity with a sensitive process long enough to hide fraudulent activity. It also supports cross-training and operational resilience: the organization gains more than one person who understands essential duties, which makes absence coverage and independent review more practical. Within a security program, personnel controls such as job rotation help manage insider-risk exposure alongside technical and administrative safeguards. CompTIA includes personnel-related security controls and governance concepts in the Security+ certification objectives; see the CompTIA Security+ SY0-701 exam objectives. The Association of Certified Fraud Examiners also identifies job rotation as an anti-fraud control that can increase the opportunity to detect misconduct through fresh oversight; see its anti-fraud controls guidance.

QUESTION NO: 88

A company plans to secure its systems by:

Preventing users from sending sensitive data over corporate email

Restricting access to potentially harmful websites

Which of the following features should the company set up? (Select two).

A. DLP software

B. DNS filtering

C. File integrity monitoring

D. Stateful firewall

E. Guardralls

F. Antivirus signatures

ANSWER: A B

Explanation:

DLP software is appropriate for preventing sensitive information from leaving the organization through corporate email. Data loss prevention policies can inspect email messages and attachments for defined sensitive-data patterns, such as personally identifiable information, payment-card data, or confidential document classifications. When a policy match occurs, the organization can block the message, quarantine it, encrypt it, or require a business justification before delivery. This directly supports control of data exfiltration through approved communication channels.

DNS filtering supports the goal of restricting access to harmful websites by evaluating requested domain names against security and category policies. It can block domains known to distribute malware, conduct phishing campaigns, or fall into prohibited content categories before the endpoint establishes a connection to the site. DNS-layer enforcement is useful for applying consistent web-access controls across managed users and devices, including users working away from the corporate network when the protective DNS service is configured accordingly.

These controls address distinct layers of the stated requirements: DLP governs sensitive data leaving through email, while DNS filtering governs access to risky internet destinations. CompTIA Security+ includes data loss prevention and DNS filtering among relevant security-control concepts. See the CompTIA exam objectives resources and Microsoft documentation on DLP.

QUESTION NO: 89

An administrator assists the legal and compliance team with ensuring information about customer transactions is archived for the proper time period. Which of the following data policies is the administrator carrying out?

A. Compromise

B. Retention

C. Analysis

D. Transfer

E. Inventory

ANSWER: B

Explanation:

Retention is correct because a data-retention policy defines how long an organization must preserve information before it can be securely deleted or otherwise disposed of. Customer transaction data is commonly subject to legal, regulatory, contractual, tax, audit, and operational recordkeeping requirements. In this scenario, the administrator is working with legal and compliance personnel to ensure the transaction information is archived for the required duration. That activity is the direct implementation of a retention schedule: identifying a record category, applying its required preservation period, and maintaining the records in an accessible and protected archive until the period expires.

Retention requirements should be documented, consistently enforced, and aligned with applicable laws and organizational obligations. They also support defensible records management by showing that data is not deleted prematurely and is not kept indefinitely without a legitimate business or compliance need. At the end of the defined period, the organization's lifecycle processes can authorize appropriate disposition, subject to any litigation hold or other preservation requirement. NIST describes the importance of managing personally identifiable information throughout its lifecycle, including retention and disposal, in NIST SP 800-122. The U.S. National Archives also explains that records schedules establish how long records are retained and their eventual disposition: NARA Records Scheduling.

QUESTION NO: 90

Which of the following metrics are used to calculate the risk rating in a matrix format? Select two.

A. Likelihood

B. Quantitative

C. SLE

D. Impact

E. ALE

F. ARO

ANSWER: A D

Explanation:

Likelihood and Impact are the metrics used in a risk matrix to determine a risk rating. Likelihood represents the estimated probability or frequency that a threat event will occur. It may be expressed with qualitative values, such as rare, unlikely, possible, likely, and almost certain, or with organization-defined numerical scales. Impact represents the anticipated magnitude of harm if the event occurs, such as financial loss, operational disruption, regulatory consequences, harm to personnel, or reputational damage.

A matrix places likelihood on one axis and impact on the other. The intersection of their assigned values produces a rating, commonly expressed as low, medium, high, or critical. This provides a consistent, repeatable method for prioritizing risks and deciding where treatment resources should be focused. NIST describes risk as a function of the likelihood that a threat will exploit a vulnerability and the resulting impact, which is the foundational principle represented by this type of matrix. See NIST SP 800-30 Rev. 1 and the NIST Cybersecurity Framework.

QUESTION NO: 91

Various stakeholders are meeting to discuss their hypothetical roles and responsibilities in a specific situation, such as a security incident or major disaster. Which of the following best describes this meeting?

A. Penetration test

B. Continuity of operations planning

C. Tabletop exercise

D. Simulation

ANSWER: C

Explanation:

Tabletop exercise is correct because it is a discussion-based method for evaluating preparedness, plans, decision-making, communications, and assigned responsibilities during a hypothetical event. Stakeholders such as security personnel, incident responders, business leaders, legal staff, communications teams, and operational owners meet in a facilitated setting to talk through an incident or disaster scenario. They describe the actions they would take, the information they would need, escalation paths, dependencies, and coordination requirements. The goal is to validate that documented procedures are usable and that participants understand their roles before a real incident occurs. Tabletop exercises are particularly valuable because they can expose planning gaps, unclear ownership, missing contact information, and communication problems without disrupting production systems or requiring a full operational deployment. CompTIA Security+ includes exercising incident response and business continuity capabilities as an important part of security operations and resilience. NIST describes tabletop exercises as discussion-based exercises in which personnel discuss roles, responsibilities, coordination, and responses to a scenario. See the NIST Computer Security Resource Center definition of a tabletop exercise and NIST's Guide to Test, Training, and Exercise Programs for IT Plans and Capabilities.

QUESTION NO: 92

A security administrator is addressing an issue with a legacy system that communicates data using an unencrypted protocol to transfer sensitive data to a third party. No software updates that use an encrypted protocol are available, so a compensating control is needed. Which of the following are the most appropriate for the administrator to suggest? (Select two.)

A. Tokenization

B. Cryptographic downgrade

C. SSH tunneling

D. Segmentation

E. Patch installation

F. Data masking

ANSWER: C D

Explanation:

SSH tunneling is an appropriate compensating control because it encapsulates the legacy application's otherwise unencrypted traffic within an encrypted SSH connection. This protects the confidentiality and integrity of sensitive data while it traverses untrusted or shared networks, without requiring a change to the legacy application's native protocol. The tunnel endpoints can be implemented through a managed SSH client, bastion host, gateway, or compatible proxy arrangement, allowing the third-party transfer path to use strong cryptographic protection.

Segmentation is also appropriate because it restricts the legacy system's network exposure and limits which systems, users, and network paths can communicate with it. Administrators can place the system in a dedicated VLAN, subnet, or isolated security zone and enforce tightly scoped firewall or access-control rules that permit only the required communication to the SSH tunnel endpoint and authorized third party. This reduces the attack surface and helps contain potential compromise of a system that cannot be modernized. These layered controls align with the principle of applying compensating safeguards when a technical weakness cannot be directly remediated. See CISA Secure by Design and NIST SP 800-53 security controls.

QUESTION NO: 93

Which of the following best describe the benefits of a microservices architecture when compared to a monolithic architecture? (Choose two.)

A. Easier debugging of the system

B. Reduced cost of ownership of the system

C. Improved scalability of the system

D. Increased compartmentalization of the system

E. Stronger authentication of the system

F. Reduced complexity of the system

ANSWER: C D

Explanation:

“Improved scalability of the system” is a key benefit because microservices divide an application into independently deployable services. An organization can assign additional compute capacity to the specific service experiencing increased demand, such as a payment, search, or ordering service, without having to scale every component of the application. This supports more efficient, targeted performance scaling as workload requirements change.

“Increased compartmentalization of the system” is also a core benefit. Microservices establish clearer boundaries between application functions, allowing teams to develop, deploy, maintain, and secure services separately. These boundaries can reduce the operational and security blast radius of a service-level failure by isolating functionality and enabling more targeted controls, updates, access permissions, and monitoring. Effective compartmentalization also aligns with security design practices such as segmentation and least privilege, because each service can have narrowly scoped responsibilities and interfaces.

These benefits depend on sound architecture, including well-defined service APIs, identity controls between services, centralized observability, and resilient deployment practices. When implemented appropriately, independent scaling and service isolation are fundamental characteristics that distinguish microservices from a single, tightly coupled monolithic deployment. See Microservices.io and IBM: What are microservices?.

QUESTION NO: 94

Which of the following activities would involve members of the incident response team and other stakeholders simul-ating an event?

A. Lessons learned

B. Digital forensics

C. Tabletop exercise

D. Root cause analysis

ANSWER: C

Explanation:

Tabletop exercise is correct because it is a facilitated, discussion-based simulation in which incident response personnel, business leaders, communications staff, legal representatives, and other relevant stakeholders work through a realistic hypothetical incident. Participants review a scenario and discuss the actions they would take, including escalation paths, decision authority, internal and external communications, evidence-handling procedures, recovery priorities, and coordination with third parties. The purpose is to validate that the organization’s incident response plan is practical, that everyone understands their assigned responsibilities, and that procedural or communication gaps can be identified before a real event occurs. A tabletop exercise generally does not require changes to production systems or the execution of live

technical response actions, making it an effective and low-risk preparedness activity. NIST describes tabletop exercises as discussion-based exercises that allow personnel to examine plans, policies, and procedures in a simulated setting. Regularly conducting these exercises supports incident-response readiness and continuous improvement. See NIST SP 800-84, Guide to Test, Training, and Exercise Programs and CISA Incident Response resources.

QUESTION NO: 95

A security engineer is building a file transfer solution to send files to a business partner over the internet, and it needs to be secure. Users will drop files into a specific directory and the server will send them to the partner. Which of the following can be used?

A. S/MIME

B. LDAPS

C. SSH

D. SRTP

E. SFTP (SSH File Transfer Protocol)

ANSWER: C E

Explanation:

SSH and SFTP (SSH File Transfer Protocol) can be used to implement this secure partner file-transfer workflow. SSH provides an encrypted, integrity-protected remote-communication channel across an untrusted network such as the internet. A server can authenticate to the partner's system using managed credentials, preferably public-key authentication, and transmit files without exposing their contents or authentication exchanges in cleartext. SSH is therefore the underlying secure transport and security architecture used for the transfer.

SFTP is the file-transfer protocol designed to operate over an SSH connection. It supports the practical operations needed for an automated drop-directory process, including uploading files, creating or managing remote directories when authorized, checking transfer status, and resuming or handling file operations through a single secured SSH session. An automation process can monitor the local directory, establish an SSH-authenticated session to the business partner, and upload each file through SFTP. This design provides confidentiality and integrity in transit while also enabling strong partner-system authentication and operational controls such as restricted service accounts and key rotation.

For protocol details, see the IETF's SSH Protocol Architecture (RFC 4251) and the OpenSSH documentation.

QUESTION NO: 96

The Cruel Information Security Officer (CISO) asks a security analyst to install an OS update to a production VM that has a 99% uptime SLA. The CISO tells me analyst the installation must be done as quickly as possible. Which of the following courses of action should the security analyst take first?

A. Log in to the server and perform a health check on the VM.

B. Install the patch Immediately.

C. Confirm that the backup service is running.

D. Take a snapshot of the VM.

ANSWER: D

Explanation:

Taking a snapshot of the VM is the appropriate first action because it creates a rapid rollback point immediately before the production change. An OS update can introduce boot failures, driver conflicts, application incompatibilities, or unexpected configuration changes. For a VM governed by a 99% uptime SLA, the ability to revert promptly to its known-good pre-update

state reduces the duration and operational impact of an unsuccessful patch deployment. A snapshot is particularly well suited to this scenario because it can be created quickly and preserves the virtual machine's disk state; depending on the selected settings, it may also preserve memory and device state. This supports a controlled change process even when the requested maintenance must occur quickly. The analyst should ensure the snapshot completes successfully before beginning the update and should follow the organization's approved change-management and maintenance procedures. CompTIA Security+ includes change management, backups, and recovery considerations within its published exam objectives: CompTIA Security+ SY0-701 Exam Objectives. VMware also describes snapshots as preserving a VM state that can later be reverted: VMware vSphere snapshot documentation.

QUESTION NO: 97

Which of the following must be considered when designing a high-availability network? (Select two).

A. Ease of recovery

B. Ability to patch

C. Physical isolation

D. Responsiveness

E. Attack surface

F. Extensible authentication

ANSWER: A D

Explanation:

High availability is the capability of a network and its supporting services to remain accessible and operational with minimal interruption. **Ease of recovery** is essential because failures will still occur despite resilient design. Administrators need documented, practical recovery procedures and technologies that restore connectivity and service quickly, such as redundant paths, automated failover, backups, and tested recovery processes. Recovery time directly affects the duration and impact of an outage.

Responsiveness must also be considered because a service that is technically online but too slow to use does not meet a meaningful availability objective. High-availability designs must accommodate expected traffic, avoid single points of congestion, provide sufficient capacity, and continue to meet performance requirements during a component failure or failover event. This includes planning for load balancing, bandwidth, latency, and the reduced capacity that may exist while redundant components are carrying production traffic.

Together, rapid restoration and sustained performance help ensure that network services meet organizational uptime and usability requirements. NIST describes availability as ensuring timely and reliable access to and use of information, while its contingency-planning guidance addresses recovery planning and restoration of systems and services. See NIST's definition of availability and NIST SP 800-34 contingency planning guidance.

QUESTION NO: 98

A company is decommissioning its physical servers and replacing them with an architecture that will reduce the number of individual operating systems. Which of the following strategies should the company use to achieve this security requirement?

A. Microservices

B. Containerization

C. Virtualization

D. Infrastructure as code

ANSWER: B

Explanation:

Containerization is the appropriate strategy because containers isolate applications while sharing the host system's operating-system kernel. Rather than deploying a separate full guest operating system for every workload, the organization can run multiple containerized applications on a single appropriately secured and maintained host OS. This directly reduces the number of individual operating-system instances that must be deployed, hardened, monitored, patched, and retired. Reducing those instances can simplify vulnerability and configuration management, provided the shared host and container runtime are securely configured and kept current. Containers package the application and its required libraries and dependencies into portable units, enabling consistent deployment without the overhead of a separate operating system per application. Docker describes containers as isolated processes that share the host kernel, and NIST explains that container technologies provide application isolation through OS-level virtualization. Therefore, containerization best fulfills the stated goal of reducing individual operating systems while retaining separation between workloads. See Docker's container overview and NIST SP 800-190, Application Container Security Guide.

QUESTION NO: 99

A security analyst needs to complete an assessment. The analyst is logged into a server and must use native tools to map services running on it to the server's listening ports. Which of the following tools can BEST accomplish this talk?

A. Netcat

B. Netstat

C. Nmap

D. Nessus

ANSWER: B

Explanation:

Netstat is the best tool because it provides a local view of the server's network sockets, including ports in the listening state and established connections. Since the analyst is already authenticated to the server, a host-based command is appropriate for identifying which TCP and UDP ports are exposed by local processes. On Windows, `netstat -ano` displays connections and listening ports along with the associated process ID. The PID can then be correlated with a running service or executable using native process-management tools. On Linux systems, commonly used netstat parameters such as `-tulpn` display listening TCP and UDP sockets and, when permissions allow, the owning program and process ID.

This directly supports the assessment objective: mapping locally running services to their listening ports without relying on an external scan. The command's output helps an analyst validate the system's active network exposure, identify unexpected listeners, and establish which local process is responsible for each service endpoint. Microsoft documents the Windows `netstat` command and its PID-related output at Microsoft Learn: netstat. Linux netstat behavior and options are documented in the netstat(8) Linux manual page.

QUESTION NO: 100

Easy-to-guess passwords led to an account compromise. The current password policy requires at least 12 alphanumeric characters, one uppercase character, one lowercase character, a password history of two passwords, a minimum password age of one day, and a maximum password age of 90 days. Which of the following would reduce the risk of this incident from happening again? (Choose two.)

A. Increasing the minimum password length to 14 characters.

B. Upgrading the password hashing algorithm from MD5 to SHA-512.

C. Increasing the maximum password age to 120 days.

D. Reducing the minimum password length to ten characters.

E. Reducing the minimum password age to zero days.

F. Including a requirement for at least one special character.

ANSWER: A F

Explanation:

Increasing the minimum password length to 14 characters reduces the likelihood of successful password guessing because password length substantially expands the possible search space. It also gives users more room to create memorable passphrases rather than relying on short, predictable words, names, dates, or common substitutions. Longer passwords are therefore more resistant to both automated guessing and offline cracking attempts when an attacker has obtained password-verifier data.

Including a requirement for at least one special character increases the permitted character set beyond the current alphanumeric-only policy. In the context of this question's existing complexity-based policy, this makes simple and commonly guessed password patterns less likely and raises the effort required for brute-force and rule-based guessing. Together, greater length and a broader character set directly address the stated cause: passwords that were easy to guess.

Password-policy guidance emphasizes password length as a primary strength factor and recommends screening proposed passwords against lists of known compromised or commonly used values. These measures help prevent users from selecting passwords that attackers are likely to try first. See NIST SP 800-63B: Digital Identity Guidelines and CISA guidance on phishing-resistant MFA.

QUESTION NO: 101

A security administrator is addressing an issue with a legacy system that communicates data using an unencrypted protocol to transfer sensitive data to a third party. No software updates that use an encrypted protocol are available, so a compensating control is needed. Which of the following are the most appropriate for the administrator to suggest? (Choose two.)

A. Tokenization

B. Cryptographic downgrade

C. SSH tunneling

D. Segmentation

E. Patch installation

F. Data masking

ANSWER: C D

Explanation:

SSH tunneling is an appropriate compensating control because it creates an encrypted channel around the legacy application's otherwise unencrypted communications. The legacy system can continue using its required protocol locally, while the sensitive data traverses the untrusted network inside the protected SSH connection. This provides confidentiality and integrity protections for data in transit without requiring the unavailable application update. SSH port forwarding is specifically intended to securely forward traffic through an encrypted SSH connection, making it useful when an older service cannot natively support modern transport encryption.

Segmentation is also appropriate because the legacy system should be isolated in a tightly controlled network segment. Firewall and access-control rules can restrict communication to only the required third-party endpoint, approved management systems, and explicitly necessary ports. Reducing the number of systems that can reach or observe the legacy service limits exposure to interception and lateral movement, while supporting the organization's compensating-control strategy. Together, SSH tunneling protects the data path and segmentation minimizes the reachable attack surface associated with the legacy system. See SSH tunneling guidance from SSH and CISA guidance on network segmentation.

QUESTION NO: 102

A security analyst discovers several .jpg photos from a cellular phone during a forensics investigation involving a compromised system. The analyst runs a forensics tool to gather file metadat a. Which of the following would be part of the images if all the metadata is still intact?

A. The GPS location

B. When the file was deleted

C. The total number of print jobs

D. The number of copies made

ANSWER: A

Explanation:

The GPS location is correct because photographs captured by cellular phones commonly contain Exchangeable Image File Format (EXIF) metadata embedded within the JPEG file. EXIF is metadata created by the camera or phone at capture time and can preserve useful forensic attributes such as the device manufacturer and model, image dimensions, capture date and time, camera settings, and geolocation information. When the phone's camera had location tagging enabled and permission to access location data, the image can include GPS latitude, longitude, altitude, and related location fields.

This information is valuable during a forensic investigation because it can help an analyst associate an image with a physical location, establish context for an event, and support timeline analysis using the image's capture-related metadata. "All the metadata is still intact" indicates that the embedded EXIF data has not been removed through image editing, export, social-media processing, or metadata-sanitization tools. A forensic utility can extract and display these fields without altering the original evidence file, preserving the information for analysis and documentation.

For technical details on GPS-related EXIF fields, see ExifTool GPS Tag Names. The Security+ exam objectives also include digital forensics and evidence collection concepts: CompTIA Exam Objectives.

QUESTION NO: 103

An employee receives a text message that appears to have been sent by the payroll department and is asking for credential verification. Which of the following social engineering techniques are being attempted? (Choose two.)

A. Typosquatting

B. Phishing

C. Impersonation

D. Vishing

E. Smishing

F. Misinformation

ANSWER: B C E

Explanation:

Phishing, **Impersonation**, and **Smishing** accurately describe elements of this attack. Phishing is the broad social-engineering practice of using deceptive communications to induce a target to disclose sensitive information, such as account credentials. A fraudulent request to "verify" credentials clearly serves that phishing objective. Smishing is the specific phishing delivery method used when the deceptive communication is sent through SMS or another text-message service. Because the scenario explicitly identifies a text message, smishing applies directly. Impersonation also applies because the attacker falsely presents the message as originating from the payroll department, exploiting the authority and trust associated with an internal business function to make the credential request appear legitimate. These terms are not mutually exclusive: a single malicious text can be a phishing attempt, specifically a smishing attempt, while also impersonating a trusted department. NIST defines phishing as deceptive solicitation intended to obtain sensitive information and identifies SMS-based phishing as smishing. See the NIST phishing glossary entry and the CISA phishing guidance.

QUESTION NO: 104

A company has begun labeling all laptops with asset inventory stickers and associating them with employee IDs. Which of the following security benefits do these actions provide? (Choose two.)

A. If a security incident occurs on the device, the correct employee can be notified.

B. The security team will be able to send user awareness training to the appropriate device.

C. Users can be mapped to their devices when configuring software MFA tokens.

D. User-based firewall policies can be correctly targeted to the appropriate laptops.

E. When conducting penetration testing, the security team will be able to target the desired laptops.

F. Company data can be accounted for when the employee leaves the organization.

ANSWER: A F

Explanation:

“If a security incident occurs on the device, the correct employee can be notified.” is correct because an asset tag establishes a unique identifier for a specific laptop, while the employee-ID association documents its assigned custodian. During incident triage, analysts can use that record to rapidly identify the person responsible for the affected endpoint, contact them for details, and provide containment instructions. This supports accountability and timely incident response by connecting endpoint activity and alerts to an accountable user and known asset.

“Company data can be accounted for when the employee leaves the organization.” is also correct because asset-to-custodian records support offboarding and asset recovery. The organization can determine which laptop was assigned to the departing employee, recover it, validate that it is returned, and apply appropriate handling such as disabling access, preserving evidence when necessary, or securely sanitizing the device before reassignment. This improves control of organizational information that may reside on endpoint storage and reduces the likelihood that an unmanaged device remains in a former employee’s possession.

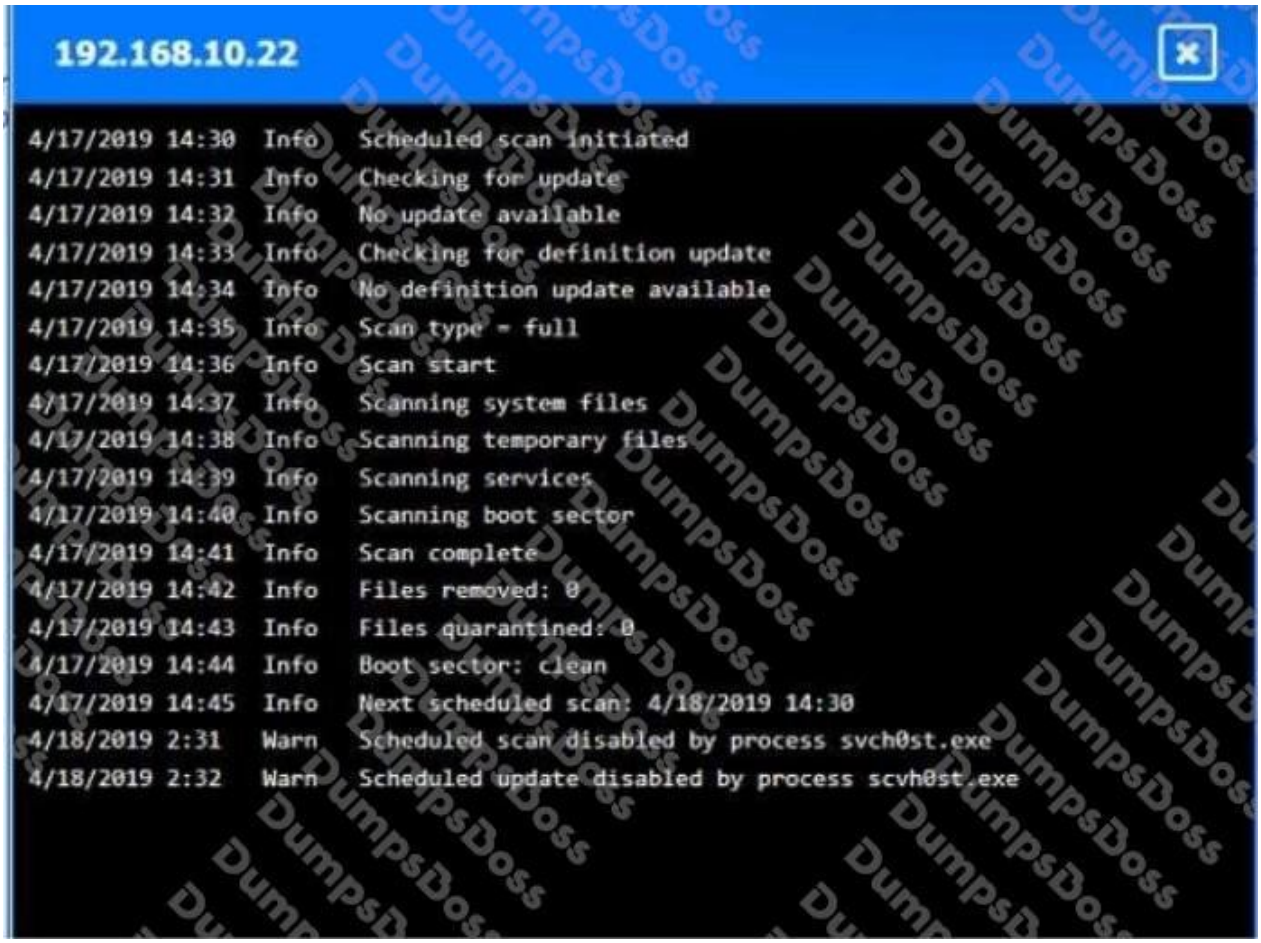
NIST identifies maintaining an inventory of system components as a foundational configuration-management control and requires organizations to identify accountable individuals for assigned components. See NIST SP 800-53 Rev. 5 and NIST SP 800-171 Rev. 2.

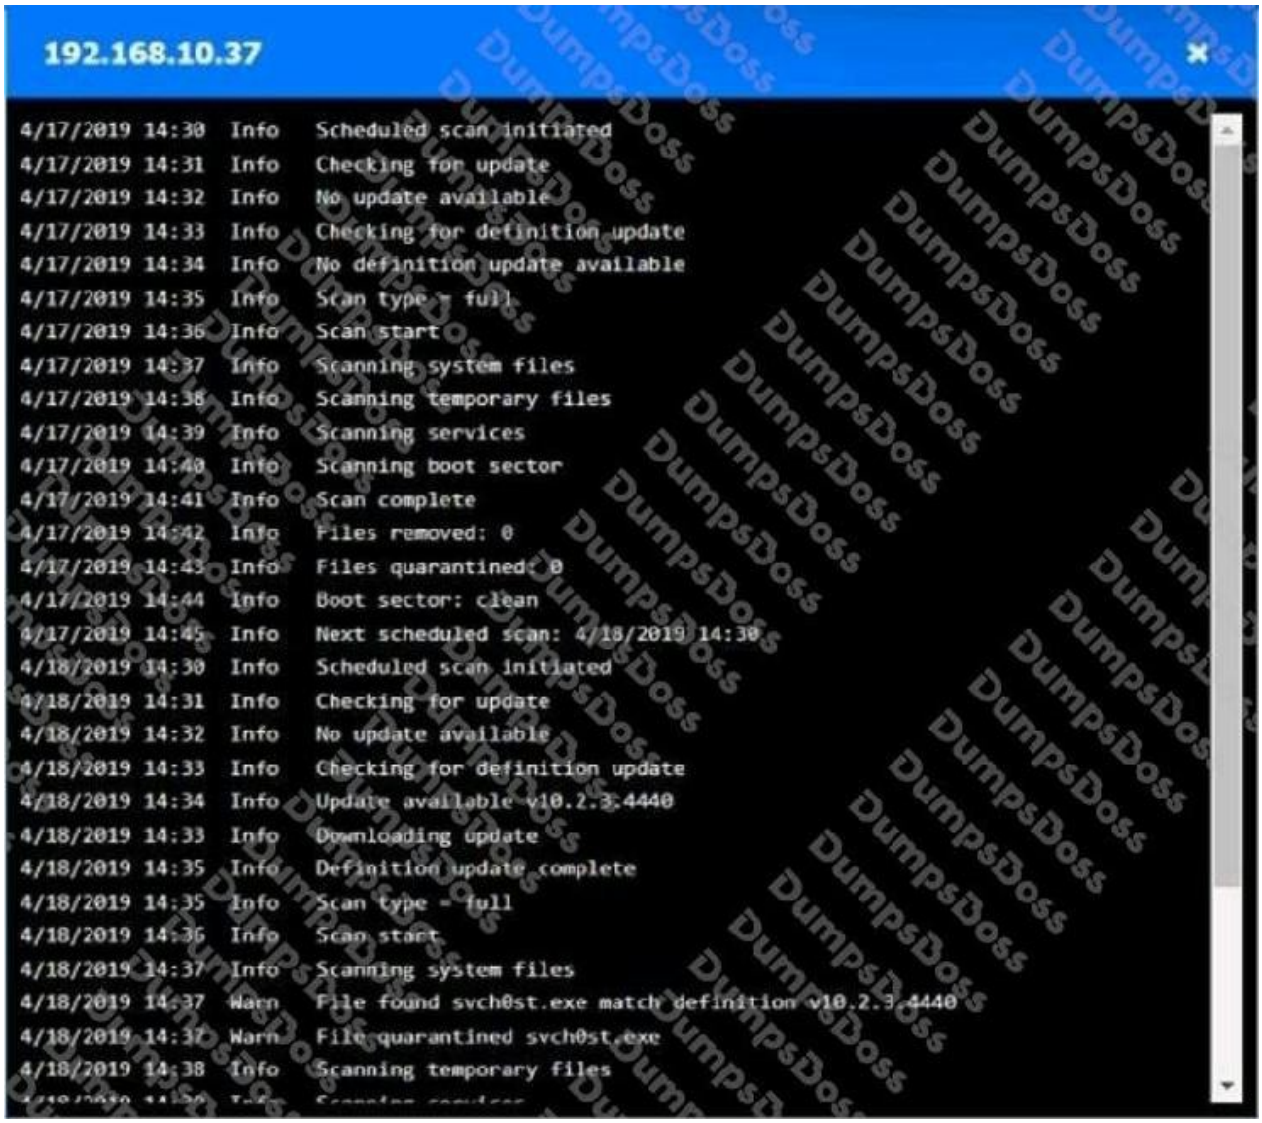
QUESTION NO: 105 - (HOTSPOT)

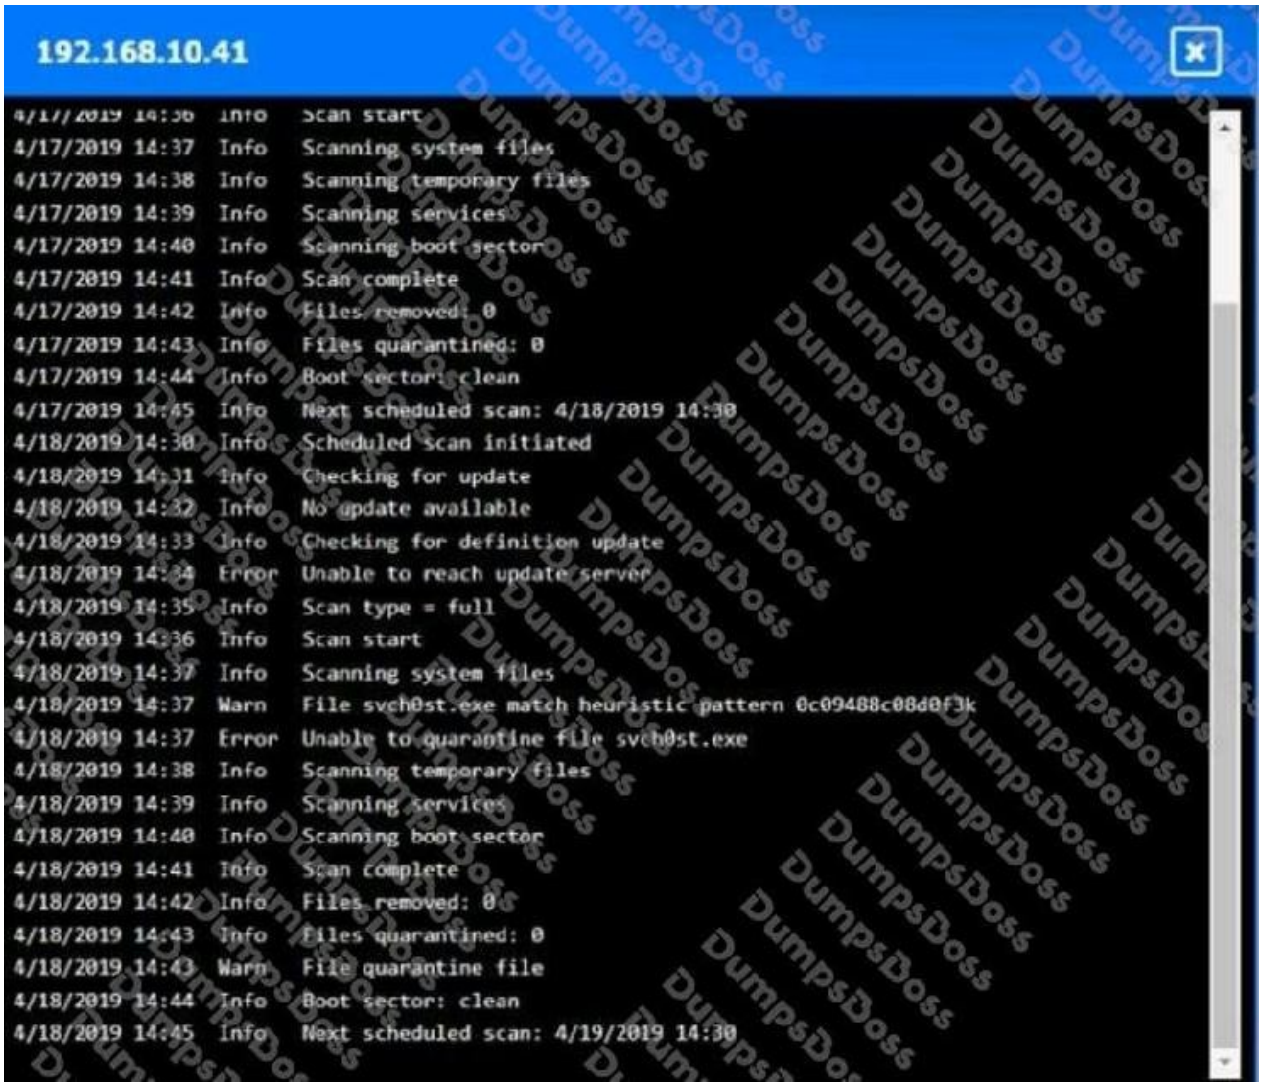
HOTSPOT

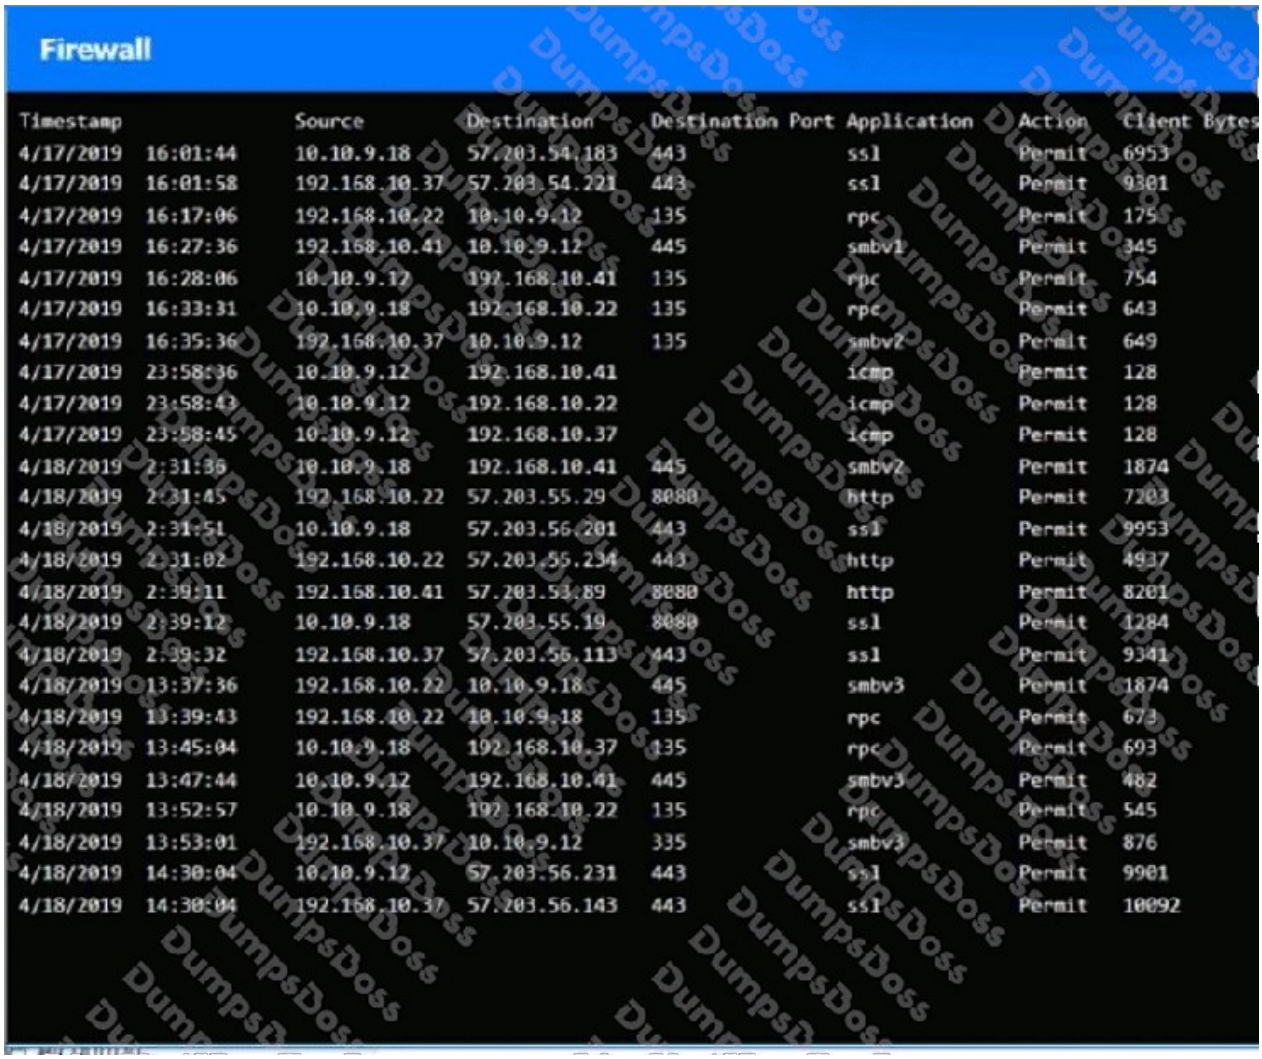
You are security administrator investigating a potential infection on a network.

Click on each host and firewall. Review all logs to determine which host originated the Infecton and then deny each remaining hosts clean or infected.


192.168.10.22
4/17/2019 14:30 Info Scheduled scan initiated
4/17/2019 14:31 Info Checking for update
4/17/2019 14:32 Info No update available
4/17/2019 14:33 Info Checking for definition update
4/17/2019 14:34 Info No definition update available
4/17/2019 14:35 Info Scan type = full
4/17/2019 14:36 Info Scan start
4/17/2019 14:37 Info Scanning system files
4/17/2019 14:38 Info Scanning temporary files
4/17/2019 14:39 Info Scanning services
4/17/2019 14:40 Info Scanning boot sector
4/17/2019 14:41 Info Scan complete
4/17/2019 14:42 Info Files removed: 0
4/17/2019 14:43 Info Files quarantined: 0
4/17/2019 14:44 Info Boot sector: clean
4/17/2019 14:45 Info Next scheduled scan: 4/18/2019 14:30
4/18/2019 2:31 Warn Scheduled scan disabled by process svch0st.exe
4/18/2019 2:32 Warn Scheduled update disabled by process scvh0st.exe

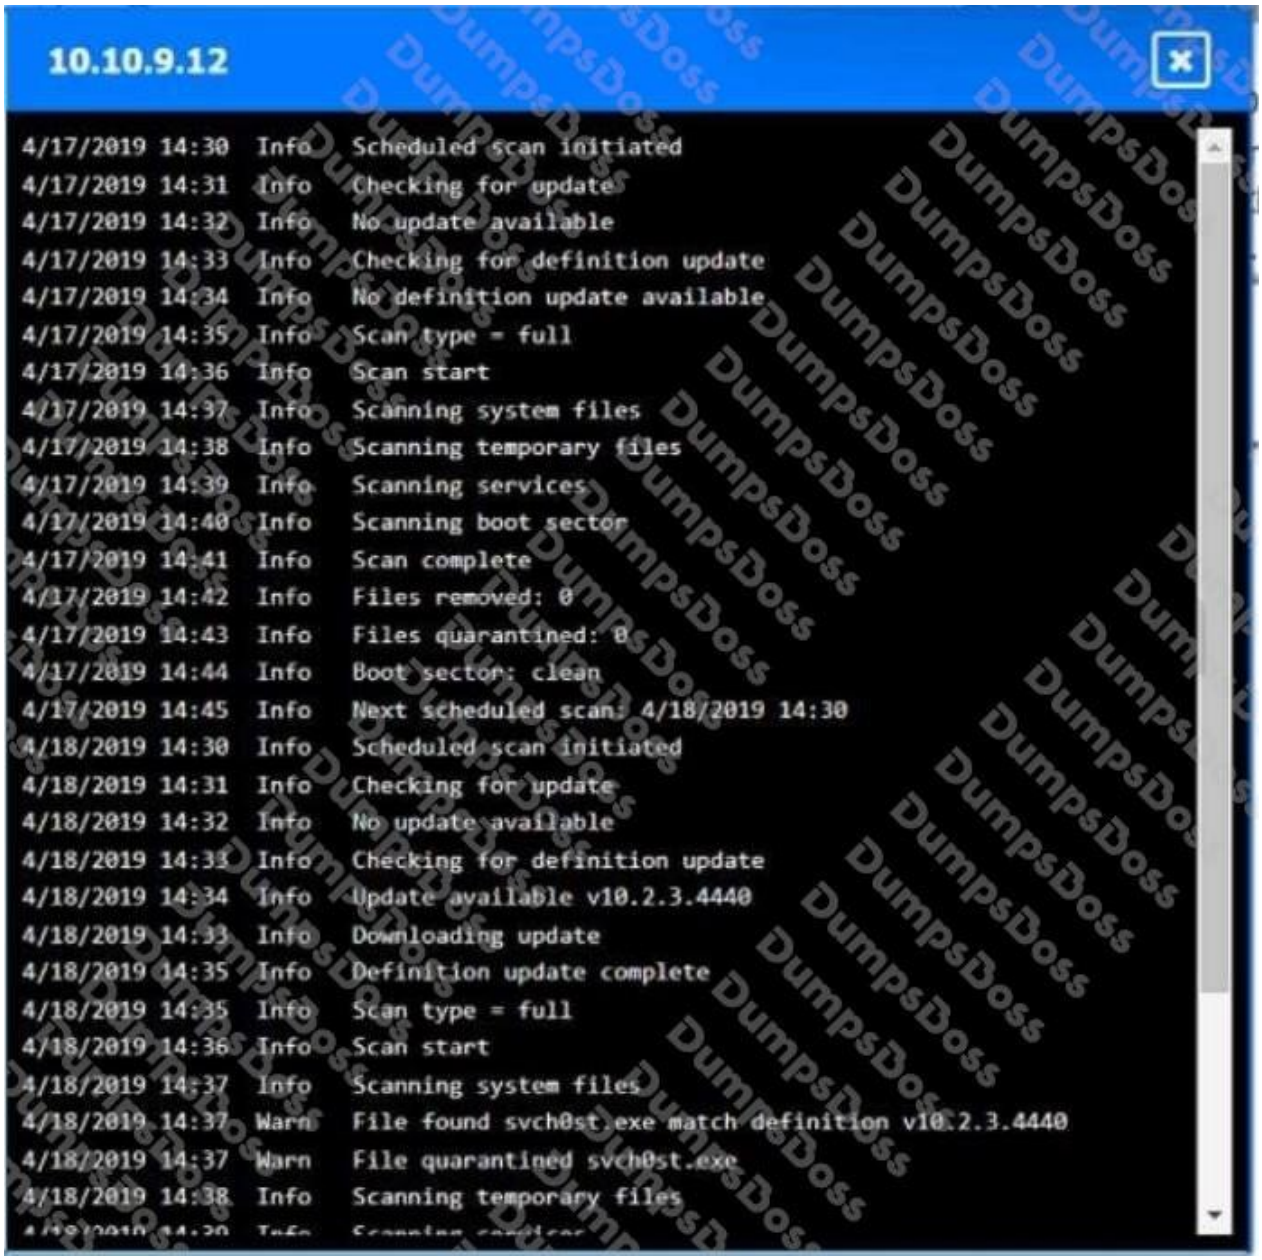

192.168.10.37
4/17/2019 14:30 Info Scheduled scan initiated
4/17/2019 14:31 Info Checking for update
4/17/2019 14:32 Info No update available
4/17/2019 14:33 Info Checking for definition update
4/17/2019 14:34 Info No definition update available
4/17/2019 14:35 Info Scan type = full
4/17/2019 14:36 Info Scan start
4/17/2019 14:37 Info Scanning system files
4/17/2019 14:38 Info Scanning temporary files
4/17/2019 14:39 Info Scanning services
4/17/2019 14:40 Info Scanning boot sector
4/17/2019 14:41 Info Scan complete
4/17/2019 14:42 Info Files removed: 0
4/17/2019 14:43 Info Files quarantined: 0
4/17/2019 14:44 Info Boot sector: clean
4/17/2019 14:45 Info Next scheduled scan: 4/18/2019 14:30
4/18/2019 14:30 Info Scheduled scan initiated
4/18/2019 14:31 Info Checking for update
4/18/2019 14:32 Info No update available
4/18/2019 14:33 Info Checking for definition update
4/18/2019 14:34 Info Update available v10.2.3.4440
4/18/2019 14:33 Info Downloading update
4/18/2019 14:35 Info Definition update complete
4/18/2019 14:35 Info Scan type = full
4/18/2019 14:36 Info Scan start
4/18/2019 14:37 Info Scanning system files
4/18/2019 14:37 Warn File found svch0st.exe match definition v10.2.3.4440
4/18/2019 14:37 Warn File quarantined svch0st.exe
4/18/2019 14:38 Info Scanning temporary files

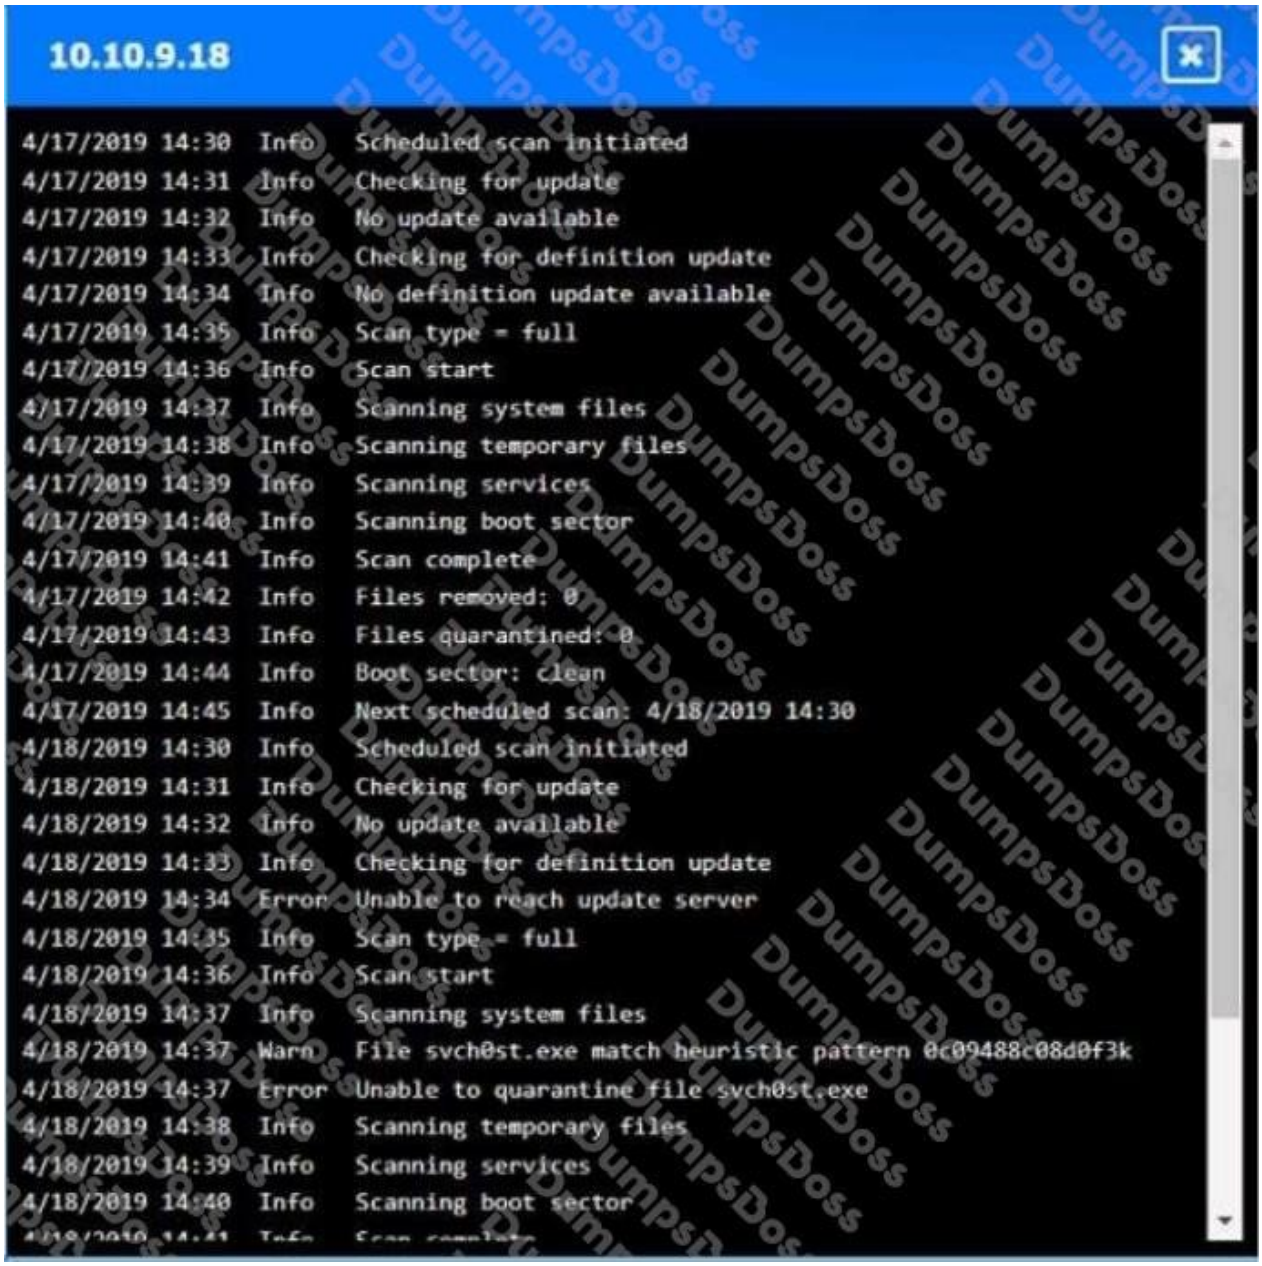

192.168.10.41
4/17/2019 14:36 Info Scan start
4/17/2019 14:37 Info Scanning system files
4/17/2019 14:38 Info Scanning temporary files
4/17/2019 14:39 Info Scanning services
4/17/2019 14:40 Info Scanning boot sector
4/17/2019 14:41 Info Scan complete
4/17/2019 14:42 Info Files removed: 0
4/17/2019 14:43 Info Files quarantined: 0
4/17/2019 14:44 Info Boot sector: clean
4/17/2019 14:45 Info Next scheduled scan: 4/18/2019 14:30
4/18/2019 14:30 Info Scheduled scan initiated
4/18/2019 14:31 Info Checking for update
4/18/2019 14:32 Info No update available
4/18/2019 14:33 Info Checking for definition update
4/18/2019 14:34 Error Unable to reach update server
4/18/2019 14:35 Info Scan type = full
4/18/2019 14:36 Info Scan start
4/18/2019 14:37 Info Scanning system files
4/18/2019 14:37 Warn File svch0st.exe match heuristic pattern 0c09488c08d0f3k
4/18/2019 14:37 Error Unable to quarantine file svch0st.exe
4/18/2019 14:38 Info Scanning temporary files
4/18/2019 14:39 Info Scanning services
4/18/2019 14:40 Info Scanning boot sector
4/18/2019 14:41 Info Scan complete
4/18/2019 14:42 Info Files removed: 0
4/18/2019 14:43 Info Files quarantined: 0
4/18/2019 14:43 Warn File quarantine file
4/18/2019 14:44 Info Boot sector: clean
4/18/2019 14:45 Info Next scheduled scan: 4/19/2019 14:30

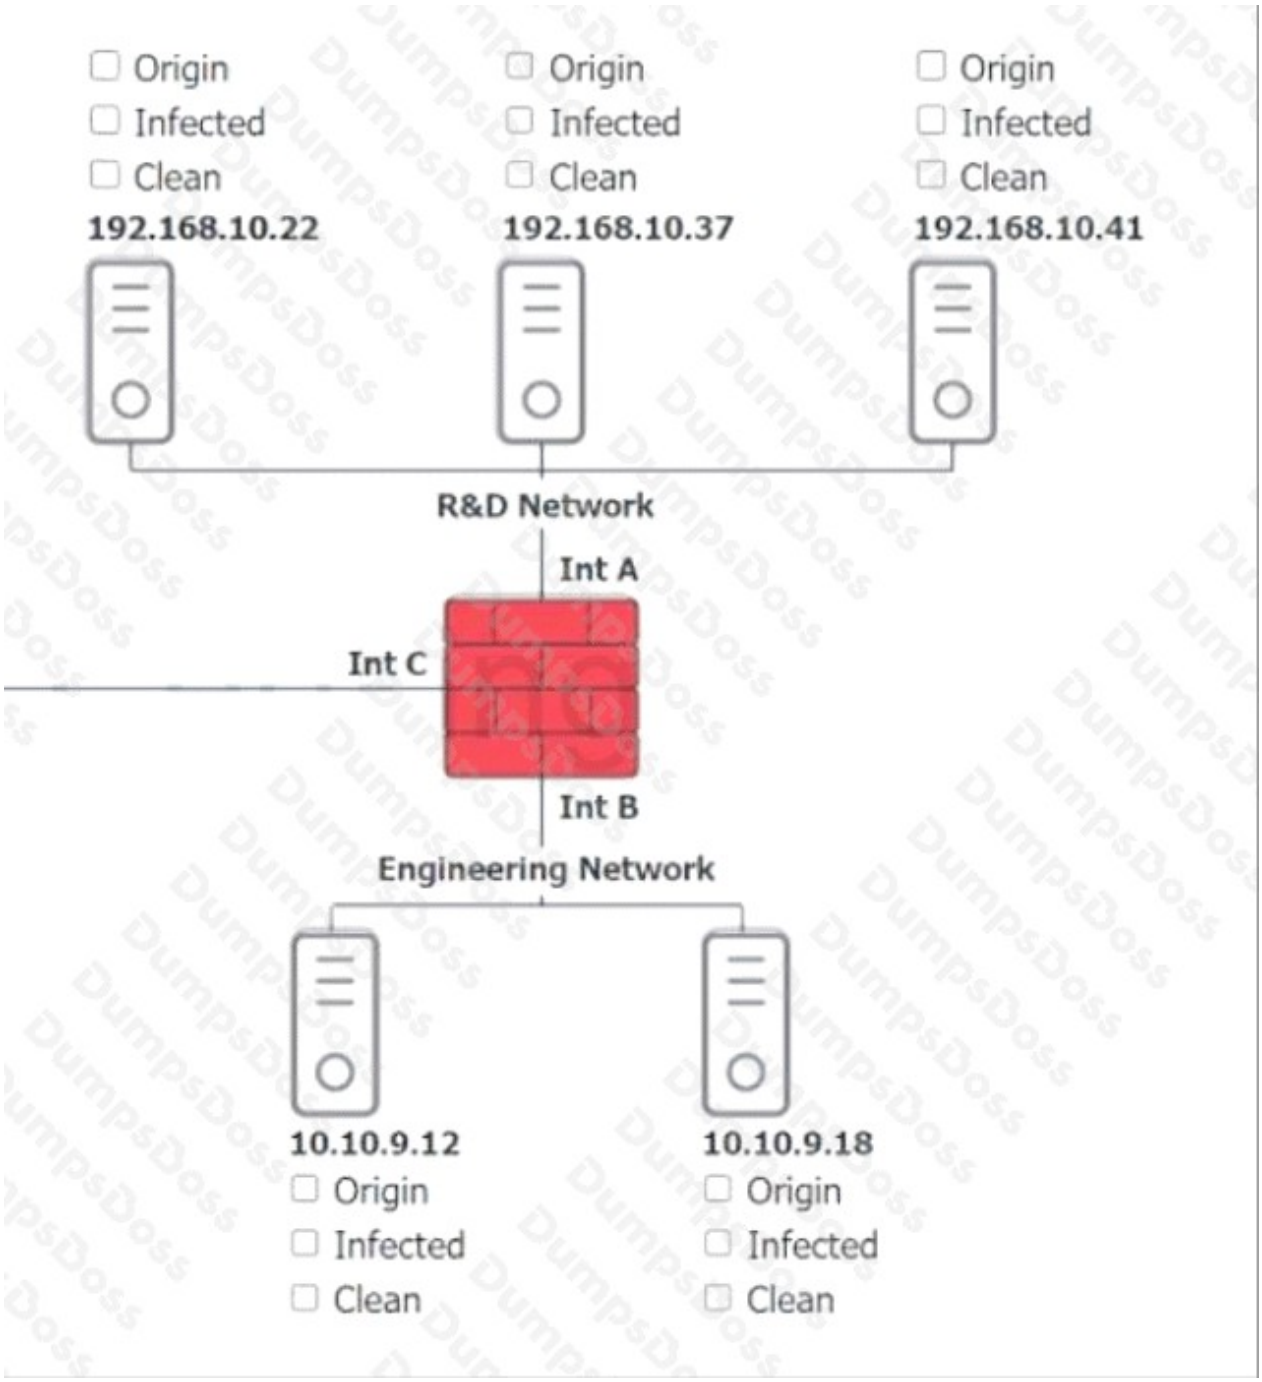


Firewall

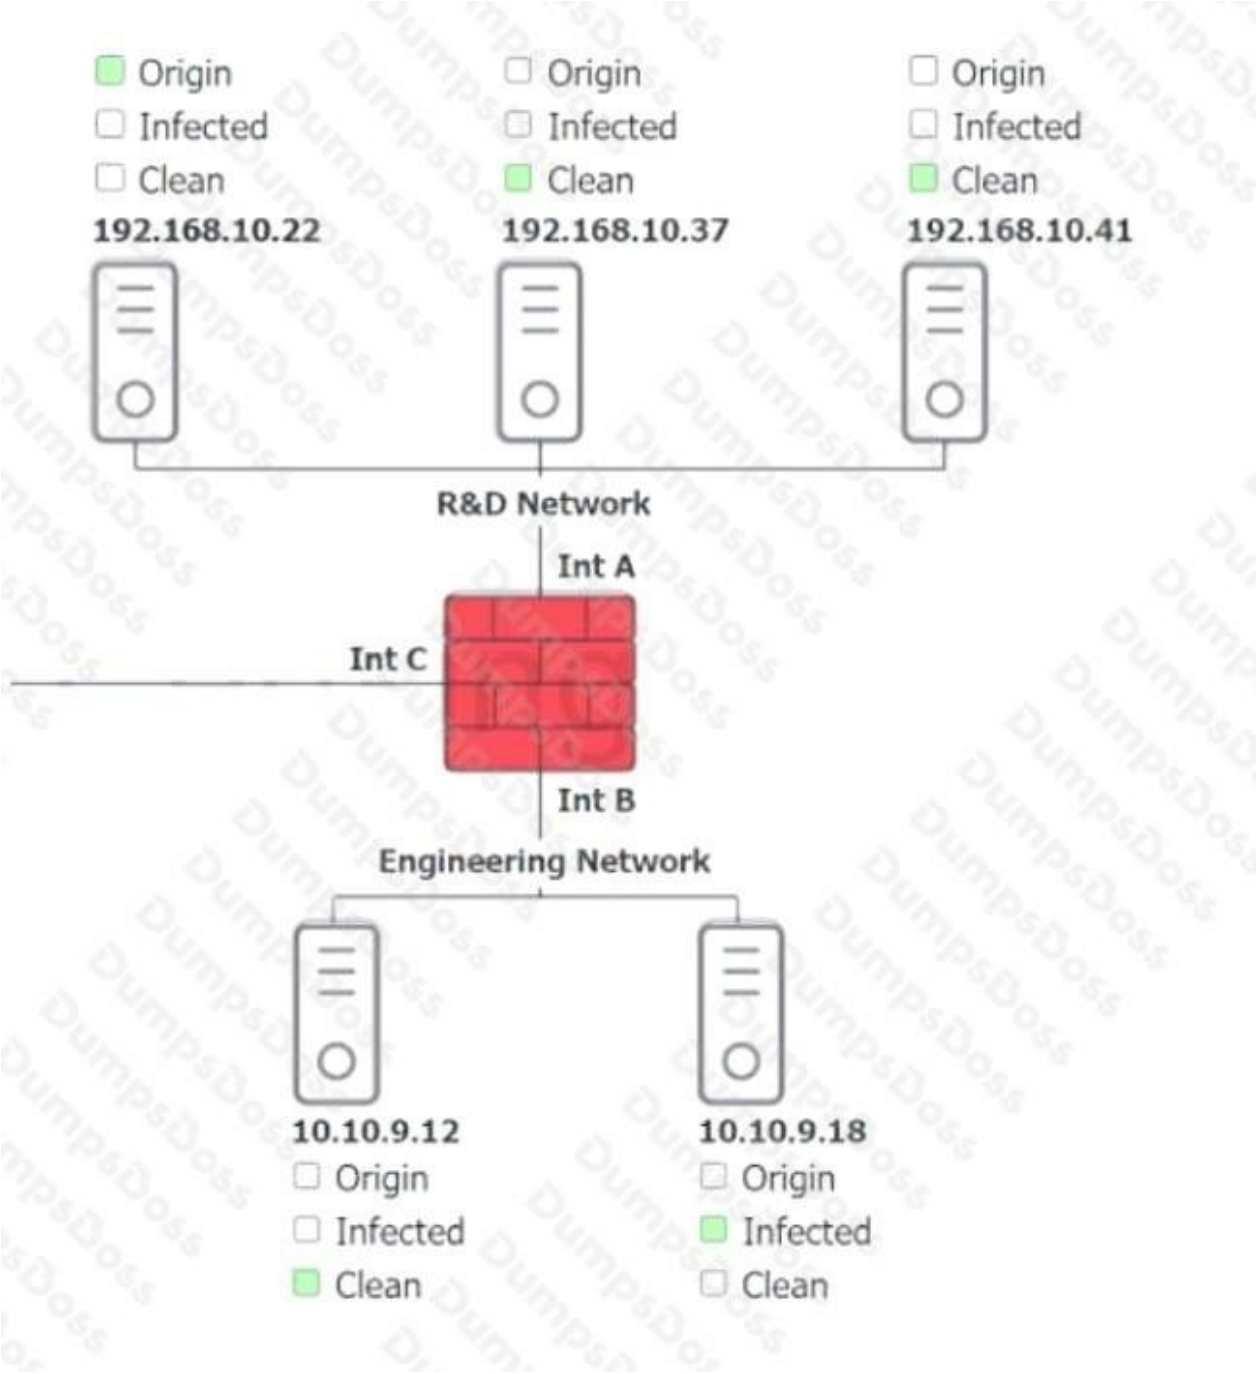
Timestamp	Source	Destination	Destination Port	Application	Action	Client Bytes
4/17/2019 16:01:44	10.10.9.18	57.203.54.183	443	ssl	Permit	6953
4/17/2019 16:01:58	192.168.10.37	57.203.54.221	443	ssl	Permit	9301
4/17/2019 16:17:06	192.168.10.22	10.10.9.12	135	rpc	Permit	175
4/17/2019 16:27:36	192.168.10.41	10.10.9.12	445	smbv1	Permit	345
4/17/2019 16:28:06	10.10.9.12	192.168.10.41	135	rpc	Permit	754
4/17/2019 16:33:31	10.10.9.18	192.168.10.22	135	rpc	Permit	643
4/17/2019 16:35:36	192.168.10.37	10.10.9.12	135	smbv2	Permit	649
4/17/2019 23:58:36	10.10.9.12	192.168.10.41		icmp	Permit	128
4/17/2019 23:58:43	10.10.9.12	192.168.10.22		icmp	Permit	128
4/17/2019 23:58:45	10.10.9.12	192.168.10.37		icmp	Permit	128
4/18/2019 2:31:36	10.10.9.18	192.168.10.41	445	smbv2	Permit	1874
4/18/2019 2:31:45	192.168.10.22	57.203.55.29	8080	http	Permit	7203
4/18/2019 2:31:51	10.10.9.18	57.203.56.201	443	ssl	Permit	9953
4/18/2019 2:31:02	192.168.10.22	57.203.55.234	443	http	Permit	4937
4/18/2019 2:39:11	192.168.10.41	57.203.53.89	8080	http	Permit	8201
4/18/2019 2:39:12	10.10.9.18	57.203.55.19	8080	ssl	Permit	1284
4/18/2019 2:39:32	192.168.10.37	57.203.56.113	443	ssl	Permit	9341
4/18/2019 13:37:36	192.168.10.22	10.10.9.18	445	smbv3	Permit	1874
4/18/2019 13:39:43	192.168.10.22	10.10.9.18	135	rpc	Permit	673
4/18/2019 13:45:04	10.10.9.18	192.168.10.37	135	rpc	Permit	693
4/18/2019 13:47:44	10.10.9.12	192.168.10.41	445	smbv3	Permit	482
4/18/2019 13:52:57	10.10.9.18	192.168.10.22	135	rpc	Permit	545
4/18/2019 13:53:01	192.168.10.37	10.10.9.12	335	smbv3	Permit	876
4/18/2019 14:30:04	10.10.9.12	57.203.56.231	443	ssl	Permit	9901
4/18/2019 14:30:04	192.168.10.37	57.203.56.143	443	ssl	Permit	10092


10.10.9.12
4/17/2019 14:30 Info Scheduled scan initiated
4/17/2019 14:31 Info Checking for update
4/17/2019 14:32 Info No update available
4/17/2019 14:33 Info Checking for definition update
4/17/2019 14:34 Info No definition update available
4/17/2019 14:35 Info Scan type = full
4/17/2019 14:36 Info Scan start
4/17/2019 14:37 Info Scanning system files
4/17/2019 14:38 Info Scanning temporary files
4/17/2019 14:39 Info Scanning services
4/17/2019 14:40 Info Scanning boot sector
4/17/2019 14:41 Info Scan complete
4/17/2019 14:42 Info Files removed: 0
4/17/2019 14:43 Info Files quarantined: 0
4/17/2019 14:44 Info Boot sector: clean
4/17/2019 14:45 Info Next scheduled scan: 4/18/2019 14:30
4/18/2019 14:30 Info Scheduled scan initiated
4/18/2019 14:31 Info Checking for update
4/18/2019 14:32 Info No update available
4/18/2019 14:33 Info Checking for definition update
4/18/2019 14:34 Info Update available v10.2.3.4440
4/18/2019 14:33 Info Downloading update
4/18/2019 14:35 Info Definition update complete
4/18/2019 14:35 Info Scan type = full
4/18/2019 14:36 Info Scan start
4/18/2019 14:37 Info Scanning system files
4/18/2019 14:37 Warn File found svch0st.exe match definition v10.2.3.4440
4/18/2019 14:37 Warn File quarantined svch0st.exe
4/18/2019 14:38 Info Scanning temporary files


10.10.9.18
4/17/2019 14:30 Info Scheduled scan initiated
4/17/2019 14:31 Info Checking for update
4/17/2019 14:32 Info No update available
4/17/2019 14:33 Info Checking for definition update
4/17/2019 14:34 Info No definition update available
4/17/2019 14:35 Info Scan type = full
4/17/2019 14:36 Info Scan start
4/17/2019 14:37 Info Scanning system files
4/17/2019 14:38 Info Scanning temporary files
4/17/2019 14:39 Info Scanning services
4/17/2019 14:40 Info Scanning boot sector
4/17/2019 14:41 Info Scan complete
4/17/2019 14:42 Info Files removed: 0
4/17/2019 14:43 Info Files quarantined: 0
4/17/2019 14:44 Info Boot sector: clean
4/17/2019 14:45 Info Next scheduled scan: 4/18/2019 14:30
4/18/2019 14:30 Info Scheduled scan initiated
4/18/2019 14:31 Info Checking for update
4/18/2019 14:32 Info No update available
4/18/2019 14:33 Info Checking for definition update
4/18/2019 14:34 Error Unable to reach update server
4/18/2019 14:35 Info Scan type = full
4/18/2019 14:36 Info Scan start
4/18/2019 14:37 Info Scanning system files
4/18/2019 14:37 Warn File svch0st.exe match heuristic pattern 0c09488c08d0f3k
4/18/2019 14:37 Error Unable to quarantine file svch0st.exe
4/18/2019 14:38 Info Scanning temporary files
4/18/2019 14:39 Info Scanning services
4/18/2019 14:40 Info Scanning boot sector


Origin
Infected
Clean
192.168.10.22
Origin
Infected
Clean
192.168.10.37
Origin
Infected
Clean
192.168.10.41
R&D Network
Int A
Int C
Int B
Engineering Network
10.10.9.12
Origin
Infected
Clean
10.10.9.18
Origin
Infected
Clean

ANSWER:



Explanation:

The originating host is 192.168.10.22. The incident evidence identifies activity on this R&D host that is consistent with the first compromised system in the observed infection path, including suspicious use of a process named `svchost.exe` and abnormal outbound HTTPS connections. HTTPS normally uses TCP port 443, so a service-like process generating numerous external connections can be significant when it does not match the expected host role, baseline behavior, or legitimate service configuration. In this scenario, the communications observed from 192.168.10.22 precede and explain the infection activity later associated with 10.10.9.18 across the network boundary.

Accordingly, 192.168.10.22 must receive the **Origin** classification. Host 10.10.9.18 must receive the **Infected** classification because it displays corresponding suspicious behavior and communication associated with the incident, but it is not the initial source. The available evidence does not show the same infection indicators on 192.168.10.37, 192.168.10.41, or 10.10.9.12, so each of those systems should be classified as **Clean**. This distinction is operationally important during containment: responders isolate the origin and all infected endpoints, then use the clean-host findings to focus further validation, monitoring, and scope confirmation. CISA's incident-response guidance describes identifying affected assets, containing confirmed compromise, and preserving evidence as core response activities; see CISA Incident Response. NIST likewise outlines the importance of analysis and containment in its Computer Security Incident Handling Guide.

QUESTION NO: 106

Which of the following are the best methods for hardening end user devices? (Select two)

A. Full disk encryption

B. Group-level permissions

C. Account lockout

D. Endpoint protection

E. Proxy server

F. Segmentation

ANSWER: A D

Explanation:

Full disk encryption and **Endpoint protection** are complementary, high-value controls for hardening user workstations, laptops, and similar endpoints. Full disk encryption protects data at rest by cryptographically safeguarding the device's storage. If a device is lost, stolen, or accessed outside the authorized operating environment, an attacker cannot simply remove the drive and read its contents. This directly supports endpoint confidentiality and is particularly important for portable devices that may hold organizational or personally identifiable data.

Endpoint protection hardens the host against active threats. Modern endpoint protection platforms commonly combine anti-malware capabilities with behavioral detection, exploit prevention, host isolation, telemetry, and response functions. These capabilities help prevent malicious code from executing and give security teams visibility into suspicious activity occurring on the device. Used together, encryption reduces the impact of physical device compromise, while endpoint protection reduces exposure to malware and other host-based attacks during normal use.

This layered approach aligns with Security+ endpoint-security concepts and established guidance for protecting end-user computing devices. See the CompTIA exam objectives resource page and NIST's Guide to Storage Encryption Technologies for End User Devices.

QUESTION NO: 107

A service provider wants a cost-effective way to rapidly expand from providing internet links to managing them. Which of the following methods will allow the service provider to best scale its services while maintaining performance consistency?

A. Escalation support

B. Increased workforce

C. Baseline enforcement

D. Technical debt

ANSWER: C

Explanation:

Baseline enforcement is the best method because it establishes a documented, repeatable standard for deploying, configuring, monitoring, and managing the provider's services. As the provider adds managed links, devices, customers, and supporting infrastructure, standardized baselines reduce variation between implementations. This makes the environment easier to automate and operate at scale, while helping ensure that service performance, security settings, and operational processes remain consistent for every customer deployment.

In practice, a baseline can define approved device configurations, required firmware versions, monitoring thresholds, naming conventions, access-control settings, and change-management requirements. Teams can then use templates, orchestration, and configuration-management tooling to apply those requirements quickly instead of designing each new managed service

independently. This reduces manual effort and rework, supporting rapid growth without sacrificing predictable service quality. It also provides a known-good reference point for performance monitoring and troubleshooting as the environment expands.

This aligns with Security+ governance and operational-resilience concepts, including the use of standard processes and configuration baselines to manage secure, scalable operations. See the CompTIA Security+ certification overview and the NIST Cybersecurity Framework resources for related guidance on repeatable, governed cybersecurity practices.

QUESTION NO: 108

Which of the following describes the procedures a penetration tester must follow while conducting a test?

A. Rules of engagement

B. Rules of acceptance

C. Rules of understanding

D. Rules of execution

ANSWER: A

Explanation:

Rules of engagement is correct because it defines the authorized procedures, boundaries, and operational expectations for a penetration test. Before testing starts, the organization and the tester establish these rules to document what assets, networks, applications, locations, and accounts are within scope, as well as any targets or actions that are explicitly excluded. The rules also identify permitted testing techniques, the approved testing period, communication and escalation contacts, evidence-handling expectations, and procedures for pausing work if testing causes an unexpected operational impact.

Clear rules of engagement provide the authorization that distinguishes legitimate security testing from unauthorized activity. They ensure the tester can demonstrate permission, remain within agreed legal and business constraints, and coordinate safely with the organization if a critical finding or service disruption occurs. This planning supports controlled testing while producing results that are meaningful to the organization's security objectives. NIST describes planning and rules of engagement as key elements of security testing methodology in NIST SP 800-115. The concept also aligns with penetration-testing topics included in the CompTIA Security+ certification framework.

QUESTION NO: 109

A security analyst at an organization observed several user logins from outside the organization's network. The analyst determined that these logins were not performed by individuals within the organization. Which of the following recommendations would reduce the likelihood of future attacks? (Choose two.)

A. Disciplinary actions for users

B. Conditional access policies

C. More regular account audits

D. Implementation of additional authentication factors

E. Enforcement of content filtering policies

F. A review of user account permissions

ANSWER: B D

Explanation:

Conditional access policies and **Implementation of additional authentication factors** directly reduce the chance that an attacker can reuse compromised credentials to access organizational resources. Conditional access applies context-aware

access decisions during authentication. For example, an organization can require stronger authentication for sign-ins originating outside trusted network locations, block access from locations that do not meet policy, or require a compliant managed device before granting access. This makes unusual external sign-in activity substantially harder to convert into a successful account compromise.

Additional authentication factors implement multifactor authentication, requiring proof of identity beyond a password, such as a phishing-resistant hardware security key, authenticator application approval, or a time-based one-time passcode. Since passwords can be guessed, stolen, reused, or obtained through phishing, a second independent factor prevents possession of the password alone from being sufficient for account access. Together, these controls provide layered protection: conditional access evaluates the context and risk of the sign-in, while multifactor authentication verifies the user with additional evidence. This aligns with Security+ identity and access management practices and with NIST guidance on using multifactor authentication to strengthen digital identity assurance. See Microsoft Entra Conditional Access overview and NIST SP 800-63B Digital Identity Guidelines.

QUESTION NO: 110

For which of the following reasons would a systems administrator leverage a 3DES hash from an installer file that is posted on a vendor's website?

A. To test the integrity of the file

B. To validate the authenticity of the file

C. To activate the license for the file

D. To calculate the checksum of the file

ANSWER: A

Explanation:

To test the integrity of the file is correct because a vendor-published cryptographic digest gives the administrator a known expected value for the installer. After downloading the installer, the administrator calculates the same type of digest locally and compares the result with the value published by the vendor. Matching values provide strong evidence that the downloaded file is identical to the file for which the vendor published the digest, meaning it was not accidentally corrupted or modified after the digest was created. A mismatch indicates that the installer's contents differ and it should not be executed until the discrepancy is resolved.

The term "3DES hash" is technically inaccurate: Triple DES is a symmetric encryption algorithm rather than a cryptographic hash algorithm. In actual deployment, vendors commonly publish SHA-256 or SHA-512 values for download verification. The security objective described by the question remains integrity verification. A hash comparison alone verifies content consistency with the published digest; assurance of publisher identity is normally provided through code signing and certificate validation. NIST describes approved hash functions and modern hash usage in its Hash Functions project, while Microsoft documents how cryptographic hash functions create fixed-size values used to identify data in its Cryptographic Hash Functions documentation.

QUESTION NO: 111

An organization is implementing a COPE mobile device management policy. Which of the following should the organization include in the COPE policy? (Select two).

A. Remote wiping of the device

B. Data encryption

C. Requiring passwords with eight characters

D. Data usage caps

E. Employee data ownership

F. Personal application store access

ANSWER: A B

Explanation:

COPE (corporate-owned, personally enabled) provides the organization with ownership and administrative control of the mobile device while still allowing limited personal use by the employee. A COPE policy should therefore establish technical safeguards that protect organizational information throughout the device lifecycle. **Remote wiping of the device** is appropriate because the organization must be able to remove corporate data if a device is lost, stolen, compromised, reassigned, or when employment ends. Mobile-device-management tooling can issue this action remotely when the device next connects, reducing the risk that sensitive business information remains accessible on an unmanaged endpoint.

Data encryption is also essential because corporate-owned mobile devices commonly store or access sensitive organizational data, including email, files, authentication tokens, and application data. Requiring device encryption protects data at rest if physical possession of the device is lost. Encryption should be paired with centrally managed device configuration and an appropriate screen-lock mechanism so that the encryption keys are protected. NIST recommends that organizations use enterprise mobility management capabilities to enforce security controls such as encryption and remote-management actions for mobile devices. See NIST SP 800-124 Rev. 2: Guidelines for Managing the Security of Mobile Devices in the Enterprise and Android Enterprise device-management implementation guidance.

QUESTION NO: 112

A company installed cameras and added signs to alert visitors that they are being recorded. Which of the following controls did the company implement? (Choose two.)

A. Directive

B. Deterrent

C. Preventive

D. Detective

E. Corrective

F. Technical

ANSWER: B D

Explanation:

Deterrent and **Detective** are the correct control categories. The prominent notification signs stating that visitors are being recorded provide a deterrent effect: they make surveillance visible and are intended to influence behavior before an incident occurs. A person considering theft, vandalism, or unauthorized activity may decide not to proceed when they know their actions can be observed and linked to them. The cameras reinforce that deterrent value because the surveillance capability is apparent rather than hidden.

The cameras also provide a detective capability. They observe and record activity in the monitored area, allowing security personnel to identify suspicious events, confirm that an incident occurred, establish a timeline, and collect evidence for investigation or response. Reviewing recordings can help determine who accessed an area and what actions took place. A single security measure can serve more than one functional purpose, and visible video surveillance commonly provides both discouragement and event detection. NIST's physical and environmental protection guidance includes monitoring physical access as an important security capability, while its control glossary describes detection-oriented controls as identifying security-relevant events. See NIST's detective control glossary entry and NIST SP 800-53, Physical and Environmental Protection controls.

QUESTION NO: 113

An organization has been experiencing outages during holiday sales and needs to ensure availability of its point-of-sale systems The IT administrator has been asked to improve both server-data fault tolerance and site availability under high consumer load Which of the following are the BEST options to accomplish this objective'? (Select TWO)

A. Load balancing

B. Incremental backups

C. UPS

D. RAID

E. Dual power supply

F. NIC teaming

ANSWER: A D

Explanation:

Load balancing and **RAID** directly address the two availability requirements. Load balancing distributes incoming point-of-sale requests across multiple application or web servers rather than allowing a single server to handle all holiday traffic. This increases capacity during demand spikes and supports continued service when a backend server becomes unavailable, provided health checks and redundant load-balancing infrastructure are properly configured. It is therefore an appropriate control for maintaining site availability under high consumer load.

RAID improves server-side data and storage fault tolerance by combining physical disks into a redundant array. A fault-tolerant configuration such as mirroring or a parity-based RAID level can allow the server's storage to remain operational after an individual disk failure, reducing the likelihood that a disk outage interrupts point-of-sale services. RAID should be implemented with monitoring, prompt failed-disk replacement, and tested backups, because its availability benefit depends on maintaining array health. Together, load balancing protects service delivery and capacity, while RAID protects the availability of server data during storage-device failures. These are complementary controls for a high-availability POS environment. See Cloudflare's load-balancing overview and IBM's RAID overview.

QUESTION NO: 114

A financial institution would like to store its customer data m the cloud but still allow the data to be accessed and manipulated while encrypted. Doing so would prevent the cloud service provider from being able to decipher the data due to its sensitivity. The financial institution Is not concerned about computational overheads and slow speeds. Which of the following cryptographic techniques would best meet the requirement?

A. Asymmetric

B. Symmetric

C. Homomorphic

D. Ephemeral

ANSWER: C

Explanation:

Homomorphic encryption is the appropriate technique because it supports computation on ciphertext without first exposing the underlying plaintext. The financial institution can encrypt customer information before it is stored with the cloud provider, and the provider can perform approved operations—such as calculations, aggregation, or certain query-related processing—against the encrypted values. The resulting encrypted output can then be returned to the institution, which decrypts it using its protected key material. This preserves the confidentiality of sensitive customer data from the cloud service provider while still enabling the outsourced environment to process that data.

The question explicitly states that computational overhead and slow processing are acceptable. That is significant because homomorphic encryption, particularly fully homomorphic encryption, has traditionally imposed substantial performance and implementation costs compared with conventional encryption approaches. Its key benefit in this scenario is therefore not speed, but the ability to maintain encryption throughout processing and reduce the need to disclose plaintext to the cloud environment. NIST describes homomorphic encryption as enabling operations to be performed on encrypted data, and Microsoft provides an implementation-oriented overview of this privacy-preserving model.

References: NIST Cryptographic Standards and Guidelines; Microsoft Learn: Fully Homomorphic Encryption.

QUESTION NO: 115

A security team purchases a tool for cloud security posture management. The team is quickly overwhelmed by the number of misconfigurations that the tool detects. Which of the following should the security team configure to establish workflows for cloud resource security?

A. CASB

B. IAM

C. SOAR

D. XDR

ANSWER: C

Explanation:

SOAR is correct because security orchestration, automation, and response platforms are designed to turn high volumes of security findings into consistent, repeatable workflows, often called playbooks. A cloud security posture management tool identifies configuration weaknesses and policy violations across cloud resources, but its findings can become difficult to manage when the environment is large or changes frequently. Integrating those findings with SOAR lets the team define triage and remediation processes, enrich alerts with contextual data, assign tickets to resource owners, request approvals, send notifications, and automate approved corrective actions through cloud-provider APIs. For example, a playbook can validate that a publicly exposed storage resource violates policy, gather ownership and business context, create a service-management ticket, and initiate a remediation workflow according to the organization's change-control requirements. This reduces manual effort while improving prioritization, documentation, consistency, and response time. CompTIA's Security+ objectives include security orchestration and automation concepts as tools for streamlining operational security activities. CISA also describes automation as an important capability for improving the speed and scalability of cyber defense operations. See the CompTIA exam objectives resources and CISA's Secure Cloud Business Applications project.

QUESTION NO: 116

A company is developing a critical system for the government and storing project information on a fileshare. Which of the following describes how this data will most likely be classified? (Select two).

A. Private

B. Confidential

C. Public

D. Operational

E. Urgent

F. Restricted

ANSWER: B F

Explanation:

Confidential and **Restricted** are appropriate classifications for fileshare information associated with development of a critical government system. Project artifacts can include system designs, technical configurations, implementation details, security architecture, test results, identified weaknesses, and operational timelines. Disclosure of such material could adversely affect the confidentiality, integrity, or availability of a system that supports an important government function.

Confidential identifies information that must be accessed only by authorized personnel because unauthorized disclosure could cause meaningful harm. Applying this label supports controls such as role-based or need-to-know access, encryption, audit logging, secure sharing procedures, and retention requirements.

Restricted communicates an even narrower distribution requirement for especially sensitive project material. It is suitable where access should be limited to a defined group, such as cleared project staff, authorized government stakeholders, and security personnel. Restricting access helps protect details that could enable targeting, compromise, disruption, or misuse of a critical system.

These classifications align with risk-based government information protection practices. NIST's impact framework describes protecting information based on the potential adverse effect of a compromise, while NIST control guidance addresses access enforcement, least privilege, and protection of sensitive system information. See NIST FIPS 199 and NIST SP 800-53.

QUESTION NO: 117

A security analyst is reviewing the following logs:

```
[10:00:00 AM] Login rejected - username administrator - password Spring2023
[10:00:01 AM] Login rejected - username jsmith - password Spring2023
[10:00:01 AM] Login rejected - username guest - password Spring2023
[10:00:02 AM] Login rejected - username cpolk - password Spring2023
[10:00:03 AM] Login rejected - username fmartin - password Spring2023
```

Which of the following attacks is most likely occurring?

A. Password spraying

B. Account forgery

C. Pass-the-hash

D. Brute-force

ANSWER: A

Explanation:

Password spraying is occurring because the log pattern shows repeated authentication attempts using one likely common password, such as `password123`, against multiple distinct user accounts. Rather than concentrating many password guesses on one account, the activity distributes a small number of common-password attempts across a population of accounts. This approach is designed to identify an account that uses a weak or commonly used password while reducing the chance that any individual account reaches a failed-login threshold quickly.

This distinction is operationally important for monitoring and response. Analysts should correlate failed authentication events by source address, user-agent or tooling indicators, password value when it is safely available in controlled telemetry, and the number of unique accounts targeted over a short time window. Useful mitigations include multifactor authentication, blocking known-bad source indicators, smart lockout or password-protection controls, detection rules for one-password-to-many-account behavior, and preventing use of common passwords. Microsoft describes password spray attacks as attempts to use a few commonly used passwords against many accounts and provides detection and investigation guidance in its password spray attack documentation. MITRE ATT&CK also catalogs this behavior as Password Spraying, a credential-attack technique.

QUESTION NO: 118

Which of the following is a key reason to follow data retention policies during asset decommissioning?

A. To ensure data is securely destroyed when no longer needed

B. To make backup copies of all company data before disposing of hardware

C. To allow employees to access old files even after the hardware is recycled

D. To keep all customer data available in case it is required in the future

ANSWER: A

Explanation:

To ensure data is securely destroyed when no longer needed is correct because data retention policies establish how long an organization must preserve information and when it is eligible for disposal. During asset decommissioning, storage devices can contain sensitive business records, customer information, credentials, intellectual property, or regulated data. Applying the retention policy ensures that information is neither discarded before its required retention period nor retained indefinitely without a business, legal, or regulatory need. Once the applicable retention period, legal hold, and operational requirements have ended, the organization should sanitize or destroy the data using a method appropriate to the media and data sensitivity. Secure destruction helps prevent unauthorized recovery of residual data from retired drives, removable media, and other assets, while also supporting privacy and records-management obligations. NIST describes media sanitization as rendering access to target data infeasible for a given level of effort, including when media is reused or disposed of. See NIST SP 800-88 Rev. 1, Guidelines for Media Sanitization. CompTIA's Security+ exam objectives also identify data lifecycle and data-remnant considerations as relevant security concepts; see the CompTIA exam objectives page.

QUESTION NO: 119

A company is implementing a policy to allow employees to use their personal equipment for work. However, the company wants to ensure that only company-approved applications can be installed. Which of the following addresses this concern?

A. MDM

B. Containerization

C. DLP

D. FIM

ANSWER: B

Explanation:

Containerization is correct because it creates a distinct, managed workspace for corporate applications and data on an employee-owned device. This approach supports BYOD by preserving separation between the employee's personal environment and the organization's work environment. Within the managed container, the organization can define and enforce which corporate applications are available, require apps to meet organizational security requirements, and prevent unapproved software from being used in the work context. The organization therefore retains control over its apps and data without needing to administer every aspect of the employee's personal device.

In practice, containerization is commonly implemented through an enterprise mobility management or unified endpoint management platform. The management platform applies application protection and configuration policies to managed work apps, while the container enforces the logical boundary that distinguishes organizational resources from personal resources. This is particularly appropriate when employees use personally owned phones, tablets, or laptops, since the policy objective is to control the company workspace rather than take ownership of the entire endpoint. Microsoft describes app protection policies and managed app controls at Microsoft Learn. IBM also describes how containerization separates and manages corporate content on mobile devices at IBM Documentation.