PassQueen
a.k.a dumpsqueen.com

CompTIA Advanced Security Practitioner (CASP+) Exam

CompTIA CAS-004

Version Demo

Total Demo Questions: 72

Total Premium Questions: 720

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Security Architecture	182
Topic 2, Security Operations	242
Topic 3, Security Engineering and Cryptography	190
Topic 4, Governance, Risk, and Compliance	106
Total	720

QUESTION NO: 1

A municipal department receives telemetry data from a third-party provider The server collecting telemetry sits in the municipal departments screened network and accepts connections from the third party over HTTPS. The daemon has a code execution vulnerability from a lack of input sanitization of out-of-bound messages, and therefore, the cybersecurity engineers would like to Implement nsk mitigations. Which of the following actions, if combined, would BEST prevent exploitation of this vulnerability? (Select TWO).

A. Implementing a TLS inspection proxy on-path to enable monitoring and policy enforcement

B. Creating a Linux namespace on the telemetry server and adding to it the servicing HTTP daemon

C. Installing and configuring filesystem integrity monitoring service on the telemetry server

D. Implementing an EDR and alert on Identified privilege escalation attempts to the SIEM

E. Subscribing to a UTM service that enforces privacy controls between the internal network and the screened subnet

F. Using the published data schema to monitor and block off-nominal telemetry messages

ANSWER: A F

Explanation:

Implementing a TLS inspection proxy on-path to enable monitoring and policy enforcement and using the published data schema to monitor and block off-nominal telemetry messages provide complementary preventative controls at the trust boundary. Because the telemetry stream is transported over HTTPS, an inspection proxy can terminate or decrypt the approved connection as appropriate, inspect the application payload, and enforce rules before data reaches the vulnerable daemon. This provides visibility into content that would otherwise be encrypted in transit and supports policy enforcement for the third-party connection.

The published telemetry schema supplies the specific allowlist needed to recognize valid message structures, fields, lengths, value ranges, and formats. Enforcing that schema enables the organization to reject malformed or out-of-bound messages before they are processed by the daemon. This directly addresses the stated exploitation path: specially crafted input reaching code that does not safely sanitize it. Together, encrypted-traffic inspection and schema-aware validation reduce the likelihood that malicious payloads can traverse the screened-network boundary and trigger the code-execution flaw. This aligns with network boundary protection principles in NIST SP 800-53, SC-7 and secure input-validation guidance from the OWASP Input Validation Cheat Sheet.

QUESTION NO: 2

A SOC suspects that malware on several endpoints is using DNS tunneling to exfiltrate data. The organization wants to improve detection while minimizing disruption to legitimate DNS resolution.

Which of the following actions would BEST support this objective? (Choose two.)

A. Centralize DNS query logging

B. Alert on anomalous DNS query patterns

C. Disable DNS caching on all endpoints

D. Allow outbound DNS only over UDP port 53

E. Replace internal DNS servers with public resolvers

ANSWER: A B

Explanation:

Centralize DNS query logging provides the telemetry needed to identify unusual query volume, uncommon domains, excessive NXDOMAIN responses, long subdomain labels, high-entropy labels, and endpoints communicating with suspicious resolvers. Centralized records also allow analysts to correlate DNS activity with endpoint and identity events.

Alert on anomalous DNS query patterns converts that telemetry into detections for behavior commonly associated with tunneling. Useful analytics include repeated queries to one domain, unusually large query payloads, high query rates, encoded-looking labels, and periodic beaconing behavior. These detective measures preserve normal DNS availability because they do not require blocking legitimate traffic solely on an initial suspicion.

NIST identifies DNS logging and monitoring as important sources of security-relevant network information. See NIST SP 800-81 Rev. 3, Secure Domain Name System Deployment Guide.

QUESTION NO: 3

A systems administrator is preparing to run a vulnerability scan on a set of information systems in the organization. The systems administrator wants to ensure that the targeted systems produce accurate information especially regarding configuration settings.

Which of the following scan types will provide the systems administrator with the MOST accurate information?

A. A passive, credentialed scan

B. A passive, non-credentialed scan

C. An active, non-credentialed scan

D. An active, credentialed scan

ANSWER: D

Explanation:

An active, credentialed scan provides the most accurate information about a target system's configuration because the scanner actively connects to each host and authenticates with authorized credentials. This gives the assessment tool direct visibility into local configuration settings, installed software, patch levels, running services, registry or policy values, file permissions, and other host-based evidence. Rather than inferring a system's state solely from externally observable network behavior, authenticated scanning can inspect the configuration from within the operating system and correlate findings with the actual local state.

This approach substantially improves vulnerability-detection accuracy and supports validation of configuration compliance against organizational baselines. Appropriate credentials should be provisioned according to least-privilege principles and the scan should be authorized, scoped, and scheduled to minimize operational impact. NIST describes authenticated vulnerability scanning as a method that can collect more detailed configuration and patch information than unauthenticated techniques. See NIST SP 800-115, Technical Guide to Information Security Testing and Assessment and Tenable's overview of authenticated scanning.

QUESTION NO: 4

A security analyst has been tasked with providing key information in the risk register. Which of the following outputs or results would be used to BEST provide the information needed to determine the

security posture for a risk decision? (Select TWO).

A. Password cracker

B. SCAP scanner

C. Network traffic analyzer

D. Vulnerability scanner

E. Port scanner

F. Protocol analyzer

ANSWER: B D

Explanation:

SCAP scanner and Vulnerability scanner provide directly actionable assessment evidence for recording and evaluating technical risk in a risk register. A SCAP scanner assesses systems against standardized security automation content, including configuration baselines, known platform and vulnerability identifiers, and compliance-related checks. Its results help establish whether systems conform to the organization's required secure configuration state and identify control gaps that affect exposure.

A Vulnerability scanner identifies known weaknesses in hosts, applications, and services, typically providing affected assets, vulnerability identifiers, severity information, and remediation guidance. This information supports core risk-register fields such as the affected asset, description of the condition, likelihood and impact inputs, existing control status, risk rating, risk owner, and recommended treatment. Combining configuration-compliance findings with vulnerability findings gives decision-makers a defensible view of the organization's technical security posture and a basis for prioritizing remediation or accepting, transferring, or mitigating risk.

NIST describes SCAP as a suite of interoperable specifications for automating vulnerability management, measurement, and policy compliance: NIST SCAP. Guidance on using vulnerability-scanning results as part of ongoing vulnerability management is available from CISA Vulnerability Scanning.

QUESTION NO: 5

A security analyst sees that a hacker has discovered some keys and they are being made available on a public website. The security analyst is then able to successfully decrypt the data using the keys from the website. Which of the following should the security analyst recommend to protect the affected data?

A. Key rotation

B. Key revocation

C. Key escrow

D. Zeroization

E. Cryptographic obfuscation

ANSWER: B

Explanation:

Key revocation is the appropriate recommendation because the disclosed keys must immediately be treated as compromised. Revocation removes the keys' authorization for continued cryptographic use, preventing systems and users from relying on the exposed key material for future encryption, decryption, signing, or verification operations as applicable. The organization should identify all data, systems, certificates, and cryptographic operations associated with the compromised keys; revoke the keys or associated certificates through the relevant key-management or public key infrastructure processes; generate replacement keys in a trusted environment; and re-encrypt affected data where feasible. Revocation status must also be distributed to relying systems—for example, through certificate revocation lists or the Online Certificate Status Protocol when certificates are involved—so that the compromised key is no longer accepted. NIST key-management guidance emphasizes that compromise requires prompt action to prevent further unauthorized use and to replace affected cryptographic keys. See NIST SP 800-57 Part 1 Rev. 5 and NIST's publication record for SP 800-57.

QUESTION NO: 6

An organization wants to perform a scan of all its systems against best practice security configurations.

Which of the following SCAP standards, when combined, will enable the organization to view each of the configuration checks in a machine-readable checklist format for fill automation? (Choose two.)

A. ARF

B. XCCDF

C. CPE

D. CVE

E. CVSS

F. OVAL

ANSWER: B F

Explanation:

XCCDF and **OVAL** are the SCAP components used together to express security-configuration guidance as automatable, machine-readable checks. XCCDF, the Extensible Configuration Checklist Description Format, provides the checklist and policy structure. It defines benchmarks, profiles, rules, remediation information, scoring, and the presentation-oriented details that let an organization organize best-practice configuration requirements into a standardized checklist. This makes it possible to select an applicable profile and view the individual configuration rules being assessed.

OVAL, the Open Vulnerability and Assessment Language, supplies the technical assessment logic that can be associated with those checklist rules. It represents tests, objects, states, and definitions used by an automated assessment tool to collect endpoint configuration information and determine whether a specified condition is met. In practice, an XCCDF rule can reference an OVAL definition, combining human-readable policy/checklist content with precise machine-executable verification logic. This pairing supports consistent configuration-compliance scanning across diverse systems and enables tools to produce standardized assessment results. NIST describes XCCDF as a language for specifying security checklists and OVAL as a language for encoding system-state assessment logic in its XCCDF specification resources and OVAL specification resources.

QUESTION NO: 7

An organization decided to begin issuing corporate mobile device users microSD HSMs that must be installed in the mobile devices in order to access corporate resources remotely. Which of the following features of these devices MOST likely led to this decision? (Select TWO.)

A. Software-backed keystore

B. Embedded cryptoprocessor

C. Hardware-backed public key storage

D. Support for stream ciphers

E. Decentralized key management

F. TPM 2.0 attestation services

ANSWER: B C

Explanation:

An **embedded cryptoprocessor** is a core HSM capability because it performs cryptographic operations within dedicated, tamper-resistant hardware rather than exposing sensitive operations to the mobile device's general-purpose operating system. This lets the microSD HSM generate keys and perform operations such as digital signing, authentication, and decryption while keeping private-key material within the secure module. As a removable factor that must be present to access remote corporate resources, it can also strengthen device and user authentication.

Hardware-backed public key storage is equally important because it provides protected storage for the device or user's public-key credentials and, critically, the associated private keys. Keys can be made non-exportable, so malware or unauthorized applications on the mobile device cannot simply copy them for later impersonation. The organization can therefore use certificate-based authentication and cryptographic proof of possession while retaining stronger control over credential protection. NIST describes cryptographic modules as mechanisms for protecting keys and performing

cryptographic functions, and its key-management guidance emphasizes protecting private keys throughout their lifecycle. See NIST FIPS 140-3 and NIST SP 800-57 Part 1.

QUESTION NO: 8

A company hosts several customer-facing applications in a Kubernetes cluster. A compromised web application pod must not be able to modify cluster resources or communicate directly with database pods unless explicitly required.

Which of the following should the security architect implement? (Choose two.)

A. Least-privilege Kubernetes RBAC

B. Default-deny network policies

C. A privileged container security context

D. A shared cluster-admin service account

E. Host networking for all application pods

ANSWER: A B

Explanation:

Least-privilege Kubernetes RBAC limits what the compromised pod's service account can do through the Kubernetes API. The service account should be granted only the specific verbs and resources required by the application, rather than broad permissions to create, modify, or delete cluster resources. This limits the impact if the pod's credentials are accessed by an attacker.

Default-deny network policies restrict pod-to-pod communications unless an explicit policy permits the traffic. The architect can permit only the required application-to-database flow and deny direct communications from other workloads. This reduces lateral movement and helps prevent a compromised web pod from reaching services that it does not need.

Kubernetes documents RBAC authorization and NetworkPolicy controls in its RBAC documentation and Network Policies documentation.

QUESTION NO: 9 - (SIMULATION)

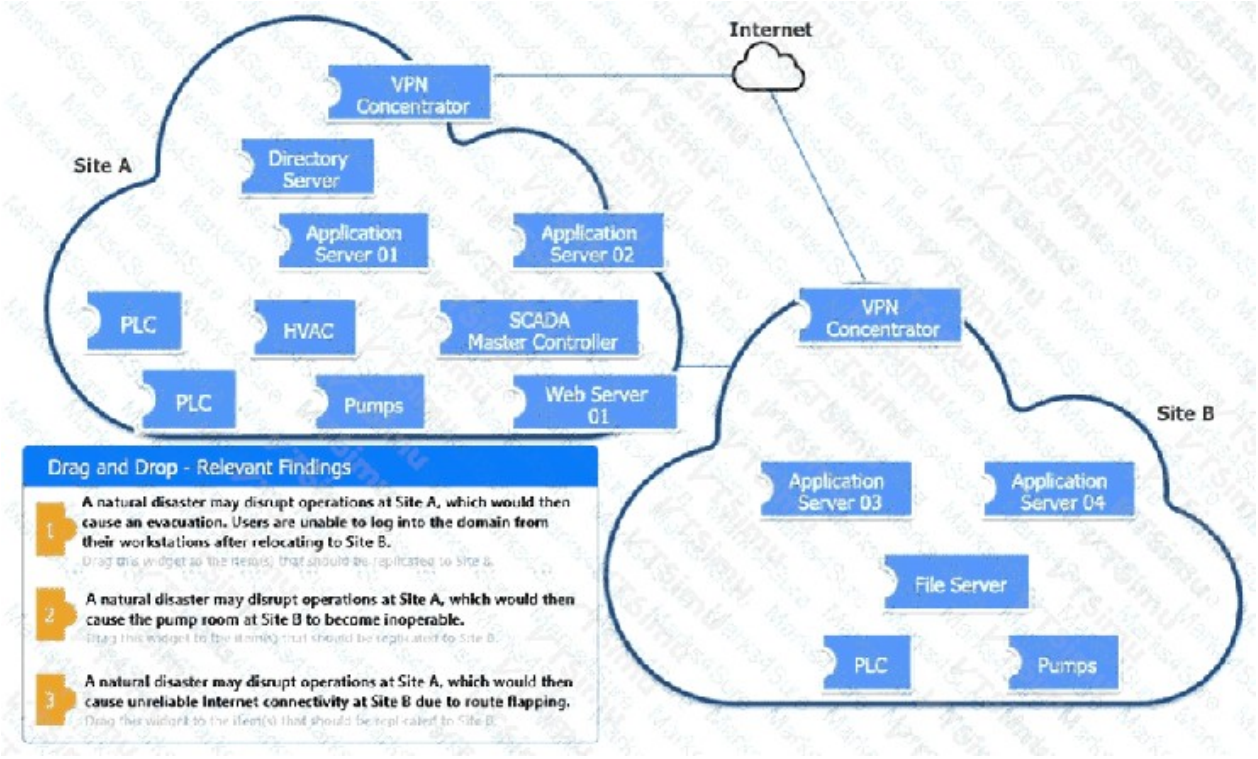

An organization is planning for disaster recovery and continuity of operations.

INSTRUCTIONS

Review the following scenarios and instructions. Match each relevant finding to the affected host.

After associating scenario 3 with the appropriate host(s), click the host to select the appropriate corrective action for that finding.

Each finding may be used more than once.

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

ANSWER: See the explanation for the answer

Explanation:

Match the findings to the Site A hosts that must also be available/replicated at Site B:

Scenario 3 corrective action: Select the action that provides a **redundant/secondary Internet or WAN path (diverse ISP connection)** for the VPN concentrator, with routing configured for stable failover (for example, route damping/controlled failover). This removes Site B's dependency on the Site A path and prevents an unstable route from continually flapping.

Do not select application servers, PLCs, pumps, HVAC, or the web server for scenario 3; routing flapping is a network-edge/VPN-WAN resiliency issue.

QUESTION NO: 10

A healthcare system recently suffered from a ransomware incident As a result the board of directors decided to hire a security consultant to improve existing network security. The security consultant found that the healthcare network was completely flat, had no privileged access limits and had open RDP access to servers with personal health information. As the consultant builds the remediation plan, which of the following solutions would BEST solve these challenges? (Select THREE).

A. SD-WAN

B. PAM

C. Remote access VPN

D. MFA

E. Network segmentation

F. BGP

G. NAC

ANSWER: B C E

Explanation:

PAM, **Remote access VPN**, and **Network segmentation** directly address the three identified security gaps. Privileged access management establishes centralized controls for administrative identities, including least-privilege authorization, credential vaulting, approval workflows, and monitored privileged sessions. These controls substantially reduce the opportunity for compromised administrator credentials to be used to access sensitive healthcare systems.

Network segmentation separates users, workloads, and sensitive PHI-hosting servers into distinct security zones. Enforcing policy-controlled traffic between zones limits lateral movement, which is especially important in ransomware scenarios because a compromise of one endpoint should not provide unrestricted access to clinical or server networks.

A remote-access VPN provides an authenticated, encrypted access path for authorized remote users rather than permitting directly exposed Remote Desktop Protocol services. RDP access to PHI systems can then be constrained to the VPN-connected population and governed by access policies. CISA recommends securing remote access services and reducing exposure of RDP, while NIST's zero-trust guidance emphasizes explicitly controlled access and segmentation principles. See CISA's guidance on securing remote access and NIST SP 800-207, Zero Trust Architecture.

QUESTION NO: 11

A small business would like to provide guests who are using mobile devices encrypted WPA3 access without first distributing PSKs or other credentials. Which of the following features will enable the business to meet this objective?

A. Simultaneous Authentication of Equals

B. Enhanced open

C. Perfect forward secrecy

D. Extensible Authentication Protocol

ANSWER: B

Explanation:

Enhanced open is the appropriate feature because it enables encrypted wireless access for users joining a guest network without a shared password, pre-distributed key, username, certificate, or other onboarding credential. In the WPA3 ecosystem, Enhanced Open uses Opportunistic Wireless Encryption to establish an individual encrypted connection between each client and the access point. This protects traffic over the radio link from passive eavesdropping, even though the network is openly discoverable and guests are not required to authenticate before joining.

This directly matches a small-business guest-access scenario: visitors can connect easily from supported mobile devices, while their over-the-air traffic receives encryption rather than being transmitted as plaintext. Enhanced Open is particularly useful for public or guest wireless deployments where distributing and rotating a common password would create administrative effort and reduce usability. The network operator can still apply normal guest-network controls, such as client isolation, segmentation, captive portals, and firewall policies, independently of the wireless encryption method. The Wi-Fi Alliance describes Enhanced Open as providing data encryption for open networks without requiring users to enter a password. See the Wi-Fi Alliance wireless security overview and Wi-Fi Alliance Enhanced Open FAQ.

QUESTION NO: 12

A software developer must choose encryption algorithms to secure two parts of a mobile application. Given the following part descriptions and requirements:

•The first part of the application is used to transfer large files and must support file parts with transfer start/stop/resume. This part requires strong file encryption.

•The second part of the application uses a bit stream to continuously authenticate both ends of the connection. This part must implement confidentiality for the stream.

Which of the following encryption algorithms should the developer implement in the code to support both parts of the application? (Select two).

A. P384

B. ECDSA

C. RC5

D. ChaCha20

E. bcrypt

F. RIPEMD

ANSWER: C D

Explanation:

RC5 is the applicable selection for the file-transfer component because it is a symmetric block cipher. Block ciphers encrypt data in discrete blocks and can be applied independently to appropriately designed file chunks, supporting encrypted large-file transfer workflows in which chunks are paused, resumed, and processed without requiring the entire file to be held in memory. RC5 is parameterized by word size, number of rounds, and key length; its specification describes its block-oriented design and encryption processing. In a real implementation, each transferred chunk should also have unique nonce/IV handling where the selected mode requires it, plus integrity protection.

ChaCha20 is the applicable selection for the continuous bit-stream component because it is a modern symmetric stream cipher designed to produce a keystream that encrypts data as it is transmitted. This makes it well suited to confidentiality for continuously flowing connection data. For an authenticated connection, ChaCha20 is commonly deployed as the ChaCha20-Poly1305 authenticated-encryption construction, which supplies both encryption and message authentication. The ChaCha20 and Poly1305 construction is standardized in RFC 8439; RC5's block-cipher construction is defined in RFC 2040.

QUESTION NO: 13

An incident response team completed recovery from offline backup for several workstations. The workstations were subjected to a ransomware attack after users fell victim to a spear-phishing campaign, despite a robust training program. Which of the following questions should be considered during the lessons-learned phase to most likely reduce the risk of reoccurrence? (Select two).

A. Are there opportunities for legal recourse against the originators of the spear-phishing campaign?

B. What internal and external stakeholders need to be notified of the breach?

C. Which methods can be implemented to increase speed of offline backup recovery?

D. What measurable user behaviors were exhibited that contributed to the compromise?

E. Which technical controls, if implemented, would provide defense when user training fails?

F. Which user roles are most often targeted by spear phishing attacks?

ANSWER: D E

Explanation:

“What measurable user behaviors were exhibited that contributed to the compromise?” is a lessons-learned question because it turns the incident into evidence that can be used to improve prevention. The team can identify observable patterns such as interaction with malicious links, use of personal credentials on fraudulent sites, reporting delays, or the specific business context that made the message convincing. Measuring these behaviors enables focused improvements to phishing simulations, reporting processes, and security-awareness content rather than relying on generalized retraining.

“Which technical controls, if implemented, would provide defense when user training fails?” is equally important because resilient security architecture assumes that users can occasionally be deceived. Lessons learned should identify practical layered safeguards, such as phishing-resistant multifactor authentication, email authentication and filtering, URL and attachment sandboxing, endpoint detection and response, least privilege, and network segmentation. These controls can prevent credential misuse, block malware execution, or contain ransomware even after a successful phishing interaction. NIST’s incident-handling guidance emphasizes using post-incident lessons learned to improve controls and reduce the likelihood and impact of future incidents. CISA likewise recommends layered phishing defenses in addition to awareness training. See NIST SP 800-61 Rev. 2 and CISA Phishing Guidance.

QUESTION NO: 14

An organization developed a social media application that is used by customers in multiple remote geographic locations around the world. The organization’s headquarters and only datacenter are located in New York City. The Chief Information Security Officer wants to ensure the following requirements are met for the social media application:

Low latency for all mobile users to improve the users’ experience

SSL offloading to improve web server performance

Protection against DoS and DDoS attacks

High availability

Which of the following should the organization implement to BEST ensure all requirements are met?

A. A cache server farm in its datacenter

B. A load-balanced group of reverse proxy servers with SSL acceleration

C. A CDN with the origin set to its datacenter

D. Dual gigabit-speed Internet connections with managed DDoS prevention

ANSWER: C

Explanation:

A CDN with the origin set to its datacenter is the best fit because it places distributed edge servers near mobile users worldwide while retaining the New York datacenter as the authoritative origin. Requests can be served from geographically close edge locations, substantially reducing round-trip time and improving responsiveness for cacheable social-media content. A CDN can also accelerate non-cacheable or dynamic requests by accepting them at a nearby edge and using optimized paths to the origin.

The CDN edge can terminate TLS connections, performing SSL/TLS offloading before requests reach the organization’s web servers. This reduces cryptographic processing demand at the origin and allows the CDN to manage certificates and modern TLS settings at a globally distributed layer. CDNs also provide high availability through redundant points of presence and can absorb or filter volumetric denial-of-service traffic at their large, distributed network edge, rather than allowing that traffic to concentrate on the single datacenter connection. For example, Amazon CloudFront documents global edge locations and HTTPS support at the distribution, while Cloudflare describes its globally distributed DDoS mitigation network. See AWS CloudFront introduction and Cloudflare DDoS protection.

QUESTION NO: 15

Company A acquired Company B. During an initial assessment, the companies discover they are using the same SSO system. To help users with the transition, Company A is requiring the following:

• Before the merger is complete, users from both companies should use a single set of usernames and passwords.

• Users in the same departments should have the same set of rights and privileges, but they should have different sets of rights and privileges if they have different IPs.

• Users from Company B should be able to access Company A ' s available resources.

Which of the following are the BEST solutions? (Select TWO).

A. Installing new Group Policy Object policies

B. Establishing one-way trust from Company B to Company A

C. Enabling multifactor authentication

D. Implementing attribute-based access control

E. Installing Company A ' s Kerberos systems in Company B ' s network

F. Updating login scripts

ANSWER: B D

Explanation:

Establishing one-way trust from Company B to Company A supports the transitional access requirement by allowing identities authenticated in Company B's environment to be recognized when accessing resources owned by Company A. In an Active Directory trust design, the resource-holding environment is configured to trust the account-holding environment, enabling Company B users to use their established credentials for authorized access to Company A resources. This minimizes credential disruption while the organizations are being integrated and works with the shared single sign-on approach. Microsoft explains that trusts provide a mechanism for users in one domain or forest to be authenticated for access to resources in another: Microsoft Learn: Active Directory trusts.

Implementing attribute-based access control is also appropriate because authorization decisions can be based on multiple identity, environmental, and resource attributes. Department can be used to assign a common entitlement baseline for users performing comparable business functions, while the source IP address can be evaluated as an environmental attribute to apply differentiated access rules. This provides the required context-aware authorization model without requiring separate static permission sets for every combination of users and network locations. NIST describes ABAC as an access-control method that evaluates attributes of subjects, objects, operations, and the environment: NIST Attribute Based Access Control.

QUESTION NO: 16

A company created an external application for its customers. A security researcher now reports that the application has a serious LDAP injection vulnerability that could be leveraged to bypass authentication and authorization.

Which of the following actions would BEST resolve the issue? (Choose two.)

A. Conduct input sanitization.

B. Deploy a SIEM.

C. Use containers.

D. Patch the OS

E. Deploy a WAF.

F. Deploy a reverse proxy

G. Deploy an IDS.

ANSWER: A E

Explanation:

Conduct input sanitization is a direct application-layer remediation for LDAP injection. The application must treat all client-supplied values as untrusted, validate them against a strict allowlist appropriate to the expected field, and correctly escape LDAP filter or distinguished-name metacharacters before incorporating values into LDAP operations. This prevents an attacker from altering the intended LDAP query structure to create a filter that bypasses authentication or authorization checks. Where the LDAP library supports it, parameterized or safely constructed LDAP queries should also be used as part of the code fix. OWASP describes LDAP injection as occurring when untrusted input changes an LDAP query and recommends validation and escaping as core defenses: OWASP Injection Prevention Cheat Sheet.

Deploy a WAF provides an important compensating control for this externally accessible application while the code correction is developed, tested, and released. A WAF can inspect inbound HTTP/S requests and apply rules that detect or block common malicious LDAP-injection payload patterns before they reach the application. This reduces exposure to active exploitation and supplies a protective layer at the web boundary, particularly for known attack signatures and anomalous request content. WAF capabilities are designed to filter web requests based on configured rules, as documented by AWS WAF Developer Guide.

QUESTION NO: 17

A networking team asked a security administrator to enable Flash on its web browser. The networking team explained that an important legacy embedded system gathers SNMP information from various devices. The system can only be managed through a web browser running Flash. The embedded system will be replaced within the year but is still critical at the moment.

Which of the following should the security administrator do to mitigate the risk?

A. Explain to the networking team the reason Flash is no longer available and insist the team move up the timetable for replacement.

B. Air gap the legacy system from the network and dedicate a laptop with an end-of-life OS on it to connect to the system via crossover cable for management.

C. Suggest that the networking team contact the original embedded systemâ€™s vendor to get an update to the system that does not require Flash.

D. Isolate the management interface to a private VLAN where a legacy browser in a VM can be used as needed to manage the system.

ANSWER: D

Explanation:

Isolate the management interface to a private VLAN where a legacy browser in a VM can be used as needed to manage the system. This is an appropriate compensating-control strategy for a business-critical legacy system that cannot yet be replaced. Adobe Flash reached end of life on December 31, 2020, and Adobe recommends uninstalling it because it is no longer supported or patched. Running the required legacy browser only inside a dedicated virtual machine limits the exposure to a tightly controlled, purpose-specific environment rather than enabling the obsolete component on ordinary administrative workstations.

Placing only the embedded system’s management interface in a private VLAN provides network segmentation. Access can be restricted to explicitly authorized administrators and management hosts, while firewall or ACL rules can permit only the minimum required management protocols and block connectivity from user networks and the internet. The VM can also be snapshotted, monitored, and rebuilt after use, reducing persistence and simplifying recovery. This design preserves the short-term operational requirement while applying least privilege, isolation, and attack-surface reduction until the planned replacement is complete. Adobe’s Flash end-of-life guidance is available at Adobe Flash Player End of Life, and CISA’s guidance on network segmentation is available at CISA network segmentation guidance.

QUESTION NO: 18

A security analyst wants to keep track of alt outbound web connections from workstations. The analyst's company uses an on-premises web filtering solution that forwards the outbound traffic to a perimeter firewall. When the security analyst gets

the connection events from the firewall, the source IP of the outbound web traffic is the translated IP of the web filtering solution. Considering this scenario involving source NAT. which of the following would be the BEST option to inject in the HTTP header to include the real source IP from workstations?

A. X-Forwarded-Proto

B. X-Forwarded-For

C. Cache-Control

D. Strict-Transport-Security

E. Content-Security-Policy

ANSWER: B

Explanation:

X-Forwarded-For is the appropriate HTTP request header for retaining the originating workstation's IP address when a web-filtering proxy or other intermediary forwards traffic. The web-filtering solution can add the client address to this header before forwarding the request. Although source NAT causes the perimeter firewall's network-layer connection log to show the filtering solution's translated address as the source, an HTTP-aware firewall, proxy log, web server log, or SIEM parser can inspect and record the value carried in `X-Forwarded-For`. This enables the analyst to correlate outbound web activity with the actual workstation rather than only with the intermediary's NAT address.

The header convention supports a comma-separated chain of addresses when requests traverse multiple proxies. The first value is generally the original client address, while later values represent forwarding intermediaries. Because HTTP headers can be supplied or manipulated by untrusted clients, the organization should configure the web-filtering solution as a trusted proxy, remove any inbound client-provided instance of the header, and insert or append the authoritative value itself. This preserves reliable attribution for internal monitoring and incident investigation. MDN documents `X-Forwarded-For` as the de facto standard header for identifying a client's originating IP address through a proxy: MDN: X-Forwarded-For. OWASP also describes the security considerations for trusting forwarded IP headers: OWASP Cheat Sheet Series.

QUESTION NO: 19

A company is losing hundreds of mobile devices each year due to insider theft. The company wants to prevent these devices from functioning off-site to deter theft, but does not want to prevent the reuse of a device the next day if a device was accidentally taken off-site. Which of the following would best solve this issue?

A. Remote wipe any device taken off-site.

B. Implement full device encryption.

C. Create a geofence around the warehouse.

D. Enable location services to monitor the mobile devices.

ANSWER: C

Explanation:

Create a geofence around the warehouse is the best solution because geofencing establishes a virtual boundary around an approved physical location. When integrated with a mobile-device-management policy, the organization can apply location-based restrictions when a managed device leaves that boundary, such as blocking access to corporate resources or placing the device in a restricted state. This provides an immediate deterrent to removal of devices from the warehouse without requiring destruction of the device's data or permanent deprovisioning.

The key requirement is that an accidentally removed device must be reusable when it is returned the following day. A geofence supports that operational need because the policy can be evaluated again after the device returns to the authorized location, allowing normal access and functionality to be restored according to the organization's configured management policy. This combines a physical-location control with centralized device enforcement, helping reduce insider theft while preserving availability for legitimate business use. NIST recommends that enterprise mobile-device security

programs centrally manage devices and apply organization-defined security policies: NIST SP 800-124 Rev. 2. CompTIA's exam-objectives resources also identify enterprise mobility and security management as relevant advanced-practitioner subject areas: CompTIA Exam Objectives.

QUESTION NO: 20

An organization developed a containerized application. The organization wants to run the application in the cloud and automatically scale it based on demand. The security operations team would like to use container orchestration but does not want to assume patching responsibilities. Which of the following service models best meets these requirements?

A. PaaS

B. SaaS

C. IaaS

D. MaaS

E. CaaS

ANSWER: E

Explanation:

Containers as a Service (CaaS) best meets these requirements because it delivers a managed container-orchestration environment, typically built around Kubernetes or a similar orchestration platform. The cloud provider operates and maintains the orchestration control plane and the underlying managed infrastructure, reducing the organization's responsibility for platform maintenance and patching. This allows the security operations team to focus on securing workloads, identities, configurations, secrets, network policies, and container images rather than administering the orchestration platform itself.

CaaS is designed for organizations that already have containerized applications and need deployment, scheduling, service discovery, health monitoring, and automated scaling capabilities. In a managed Kubernetes offering, workload replicas can scale in response to demand, while the provider can manage node and platform updates according to the selected managed-service configuration. For example, Azure Kubernetes Service supports automatic node-image upgrades to apply Linux security patches, and its cluster autoscaler adjusts node capacity as workloads require it. These managed orchestration capabilities align directly with the requirements for cloud-hosted containers, demand-based scaling, and minimized patch-management responsibility. See Azure Kubernetes Service automatic node OS image upgrades and Azure Kubernetes Service cluster autoscaler.

QUESTION NO: 21

The Chief information Officer (CIO) of a large bank, which uses multiple third-party organizations to deliver a service, is concerned about the handling and security of customer data by the parties. Which of the following should be implemented to BEST manage the risk?

A. Establish a review committee that assesses the importance of suppliers and ranks them according to contract renewals. At the time of contract renewal, incorporate designs and operational controls into the contracts and a right-to-audit clause. Regularly assess the supplier's post-contract renewal with a dedicated risk management team.

B. Establish a team using members from first line risk, the business unit, and vendor management to assess only design security controls of all suppliers. Store findings from the reviews in a database for all other business units and risk teams to reference.

C. Establish an audit program that regularly reviews all suppliers regardless of the data they access, how they access the data, and the type of data, Review all design and operational controls based on best practice standard and report the finding back to upper management.

D. Establish a governance program that rates suppliers based on their access to data, the type of data, and how they access the data Assign key controls that are reviewed and managed based on the supplier's rating. Report finding units that rely on the suppliers and the various risk teams.

ANSWER: D

Explanation:

Establish a governance program that rates suppliers based on their access to data, the type of data, and how they access the data Assign key controls that are reviewed and managed based on the supplier's rating. Report finding units that rely on the suppliers and the various risk teams. This is the best approach because it implements a risk-based third-party risk management model. The bank can identify each supplier's inherent risk by considering the sensitivity of customer data, the extent and method of access, and the supplier's role in delivering critical services. Those factors support tiering suppliers so that due diligence, contractual requirements, security-control validation, monitoring frequency, escalation, and remediation effort are proportionate to the risk posed.

Governance also makes ownership and oversight explicit. Business units that use the supplier, vendor-management personnel, security and risk teams can receive relevant findings and coordinate remediation throughout the relationship lifecycle. This is especially important for a bank, where outsourced activities do not transfer the bank's accountability for protecting customer information. NIST recommends establishing supply-chain risk-management processes that identify, assess, respond to, and monitor supplier-related risks. Similarly, U.S. banking guidance calls for risk-based due diligence, ongoing monitoring, and governance over third-party relationships. See NIST SP 800-161 Rev. 1 and the OCC third-party relationship risk-management guidance.

QUESTION NO: 22

A company with only U S -based customers wants to allow developers from another country to work on the company's website However, the company plans to block normal internet traffic from the other country Which of the following strategies should the company use to accomplish this objective? (Select two).

A. Block foreign IP addresses from accessing the website

B. Have the developers use the company's VPN

C. Implement a WAP for the website

D. Give the developers access to a jump box on the network

E. Employ a reverse proxy for the developers

F. Use NAT to enable access for the developers

ANSWER: B D

Explanation:

Having the developers use the company's VPN provides an authenticated, encrypted, and explicitly authorized remote-access path for the overseas development team. The organization can apply strong identity controls, multifactor authentication, device posture requirements, and narrowly scoped network authorization to this connection. This permits the developers to reach only the necessary development resources while the public-facing controls continue to deny ordinary traffic originating from the relevant country. NIST remote-access guidance emphasizes securing remote connections and controlling the resources that remote users may access.

Giving the developers access to a jump box on the network adds a controlled intermediary administration point. Developers first establish the approved remote connection, then use the hardened jump box to perform authorized work. A jump box can be isolated in a management segment, centrally monitored, restricted to approved accounts, and configured to limit onward connections to the website's development or administrative environment. This design supports least privilege, session logging, and segmentation while avoiding a broad geographic exception for all internet traffic. Together, the company VPN and a jump box create a specific, auditable access route for the developers without opening normal public website access from their country. See NIST SP 800-46 Rev. 2 and NIST SP 800-53 Rev. 5.

QUESTION NO: 23

A company is conducting a quantitative risk assessment for an internet-facing payment application. The application has an asset value of $800,000. A successful compromise is estimated to cause a 25% loss of the asset value.

Which of the following is the single loss expectancy for this scenario?

A. $25,000

B. $200,000

C. $600,000

D. $800,000

ANSWER: B

Explanation:

$200,000 is the single loss expectancy (SLE). SLE represents the expected financial loss from one occurrence of a risk event and is calculated by multiplying the asset value by the exposure factor. In this scenario, the calculation is $800,000 multiplied by 0.25, resulting in an SLE of $200,000.

The annualized loss expectancy would require an annual rate of occurrence in addition to the SLE. The asset value alone does not account for the expected percentage of loss, while the exposure factor is a percentage rather than a dollar amount. Quantitative risk analysis uses these values to support decisions about the appropriate cost of controls and risk treatment. See NIST SP 800-30 Rev. 1.

QUESTION NO: 24

The Chief Information Security Officer is concerned about the possibility of employees downloading 'malicious files from the internet and 'opening them on corporate workstations. Which of the following solutions would be BEST to reduce this risk?

A. Integrate the web proxy with threat intelligence feeds.

B. Scan all downloads using an antivirus engine on the web proxy.

C. Block known malware sites on the web proxy.

D. Execute the files in the sandbox on the web proxy.

ANSWER: D

Explanation:

Execute the files in the sandbox on the web proxy is the best solution because it evaluates downloaded content in an isolated environment before that content reaches a corporate endpoint. A sandbox can detonate a suspicious file and observe its runtime behavior, such as attempts to launch child processes, modify system settings, contact command-and-control infrastructure, exploit vulnerabilities, or encrypt files. This behavior-based analysis is especially valuable for detecting new, obfuscated, or previously unseen malware whose indicators may not yet be present in signatures, blocklists, or external threat-intelligence feeds. When the proxy identifies malicious behavior, it can block the download before an employee can open the file on a workstation, directly addressing both parts of the stated risk: acquisition from the internet and endpoint execution. This is a defense-in-depth control that places inspection at the web egress/ingress point and reduces reliance on users making correct security decisions. CISA describes sandboxing as an isolation technique that can contain potentially malicious code, while NIST identifies malware analysis and controlled execution as useful detection and investigation practices. See CISA Sandboxing and NIST SP 800-83 Rev. 1.

QUESTION NO: 25

A company's SOC has received threat intelligence about an active campaign utilizing a specific vulnerability. The company would like to determine whether it is vulnerable to this active campaign.

Which of the following should the company use to make this determination?

A. Threat hunting

B. A system penetration test

C. Log analysis within the SIEM tool

D. The Cyber Kill Chain

ANSWER: A

Explanation:

Threat hunting is the appropriate activity because the SOC has intelligence tied to an active, real-world campaign and needs to determine its organization's exposure to that campaign. Hunters can convert the intelligence into focused hypotheses and searches based on the campaign's known indicators of compromise, attacker behaviors, affected technologies, exploit activity, persistence mechanisms, and command-and-control patterns. They can then examine endpoint, network, identity, cloud, and application telemetry to identify systems that contain the vulnerable software or configuration and to establish whether exploitation activity has already occurred. This provides actionable context: which assets are relevant, whether the campaign's techniques have been observed, and where containment, remediation, or additional monitoring should be prioritized. Threat hunting is proactive and hypothesis driven, making it well suited to validating intelligence against the organization's actual environment rather than waiting for an alert. NIST incident-response guidance emphasizes analyzing available evidence and threat information to identify and respond to relevant malicious activity, while MITRE ATT&CK provides a structured knowledge base of adversary tactics and techniques that can be used to operationalize campaign intelligence into hunts. See NIST SP 800-61 Rev. 3 and MITRE ATT&CK.

QUESTION NO: 26

A company that uses AD is migrating services from LDAP to secure LDAP. During the pilot phase, services are not connecting properly to secure LDAP. Block is an except of output from the troubleshooting session:

```
openssl s_client -host ldap1.comptia.com -port 636

CONNECTED(00000003)
...
-----BEGIN CERTIFICATE-----
...
-----END CERTIFICATE-----
Subject=/CN=*.comptia.com
Issuer=/DC=com/DC=danville/CN=chicago
```

Which of the following BEST explains why secure LDAP is not working? (Select TWO.)

A. The clients may not trust idapt by default.

B. The secure LDAP service is not started, so no connections can be made.

C. Danvills.com is under a DDoS-inator attack and cannot respond to OCSP requests.

D. Secure LDAP should be running on UDP rather than TCP.

E. The company is using the wrong port. It should be using port 389 for secure LDAP.

F. Secure LDAP does not support wildcard certificates.

G. The clients may not trust Chicago by default.

ANSWER: A F

Explanation:

The clients may not trust idapt by default because LDAPS depends on the client successfully validating the certificate presented by the domain controller. If idapt is the issuing certification authority shown in the troubleshooting output and its root certificate is not installed in each client's trusted root certification authorities store, the TLS handshake cannot establish the trusted channel required for LDAPS. Microsoft's LDAPS guidance requires clients to trust the issuing CA and requires the domain controller to present an appropriate server certificate.

Secure LDAP does not support wildcard certificates in this Active Directory deployment because the certificate used by a domain controller must identify that specific domain controller. Microsoft's requirements state that the domain controller's fully qualified domain name must be present in the certificate subject or Subject Alternative Name, enabling clients to validate the identity of the exact server to which they connected. A certificate using a wildcard name does not provide the required explicit domain-controller FQDN identity for this purpose. The certificate-trust failure and the server-name certificate requirement together explain why the pilot clients cannot create secure LDAP sessions. See Microsoft guidance for enabling LDAP over SSL and Microsoft LDAP over SSL documentation.

QUESTION NO: 27

A developer is creating a new mobile application for a company. The application usesREST APIandTLS 1.2to communicate securely with the external back-end server. Due to this configuration, the company is concerned aboutHTTPS interception attacks. Which of the following would be thebestsolution against this type of attack?

A. Cookies

B. Wildcard certificates

C. HSTS

D. Certificate pinning

ANSWER: D

Explanation:

Certificate pinning is the best solution because the mobile application can be configured to trust only a specific server certificate, public key, or defined set of public keys for its back-end API. During the TLS handshake, the application validates the presented certificate against the pinned identity in addition to normal certificate-chain validation. This substantially limits the ability of an HTTPS-interception proxy to impersonate the API server, even if an attacker has caused a device to trust an additional root certificate or is using a locally installed enterprise inspection certificate. The intercepted connection will fail because the proxy cannot present the expected pinned certificate or key. For a production mobile application, pinning should be implemented with operational resilience in mind: pin public keys rather than a single short-lived leaf certificate where appropriate, support backup pins for planned key rotation, and establish a safe certificate-update process so that legitimate server changes do not cause an outage. OWASP identifies certificate/public-key pinning as a mobile control that helps defend against TLS man-in-the-middle attacks, while Android's Network Security Configuration provides mechanisms for declaring certificate trust anchors and pin sets. See the OWASP Mobile Application Security Testing Guide and Android Network Security Configuration documentation.

QUESTION NO: 28

Users are claiming that a web server is not accessible. A security engineer logs for the site. The engineer connects to the server and runs netstat -an and receives the following output:

TCP	192.168.5.107:54585	64.78.243.12:443	ESTABLISHED
TCP	192.168.5.107:54587	54.164.78.234:80	ESTABLISHED
TCP	192.168.5.107:54636	104.16.33.27:5228	ESTABLISHED
TCP	192.168.5.107:54676	69.65.64.94:443	ESTABLISHED
TCP	192.168.5.107:54689	91.190.130.171:443	TIME_WAIT
TCP	192.168.5.107:54775	91.190.130.171:443	FIN_WAIT_2
TCP	192.168.5.107:54789	91.190.130.171:443	ESTABLISHED
TCP	192.168.5.107:55983	79.136.88.109:31802	ESTABLISHED
TCP	192.168.5.107:56234	50.112.252.181:443	TIME_WAIT
TCP	192.168.5.107:56874	40.117.100.83:443	ESTABLISHED
TCP	192.168.5.107:00	213.37.55.67:600873	TIME_WAIT
TCP	192.168.5.107:00	213.37.55.67:600874	TIME_WAIT
TCP	192.168.5.107:00	213.37.55.67:600875	TIME_WAIT
TCP	192.168.5.107:00	213.37.55.67:600876	TIME_WAIT
TCP	192.168.5.107:00	213.37.55.67:600877	TIME_WAIT
TCP	192.168.5.107:00	213.37.55.67:600878	TIME_WAIT
TCP	192.168.5.107:00	213.37.55.67:600879	TIME_WAIT
TCP	192.168.5.107:00	213.37.55.67:600880	TIME_WAIT

Which of the following is MOST likely happening to the server?

A. Port scanning

B. ARP spoofing

C. Buffer overflow

D. Denial of service

ANSWER: D

Explanation:

Denial of service is the most likely condition because the netstat output indicates an unusually high volume of connections associated with the same remote IP address, using numerous ephemeral source ports. This pattern can consume a server's finite network and operating-system resources, including connection-tracking entries, socket buffers, CPU time, and available worker capacity. Once those resources are heavily consumed, legitimate clients can experience slow responses, connection failures, or complete inability to access the website.

The displayed connection states are also significant evidence for a service-availability event. A large accumulation of short-lived or incomplete connection activity is consistent with connection-flooding behavior, where an attacker generates enough traffic or connection setup activity to degrade normal service. The `netstat -a -n` command displays active connections and listening ports using numerical addresses, making it useful for identifying abnormal concentrations of connections during incident triage. Microsoft documents these netstat capabilities at Netstat. NIST describes denial-of-service attacks as attempts to prevent or impair authorized use of systems or networks in its denial-of-service glossary entry.

QUESTION NO: 29

A company created an external application for its customers. A security researcher now reports that the application has a serious LDAP injection vulnerability that could be leveraged to bypass authentication and authorization.

Which of the following actions would BEST resolve the issue? (Choose two.)

A. Conduct input sanitization.

B. Deploy a SIEM.

C. Use containers.

D. Patch the OS

E. Deploy a WA

F. Deploy a reverse proxy

G. Deploy an IDS.

ANSWER: A E

Explanation:

Conduct input sanitization is a primary corrective action because LDAP injection occurs when application-controlled LDAP search or authentication filters are constructed with untrusted input. The application should validate input against an allowlist appropriate to the expected value and safely escape LDAP special characters before incorporating input into a directory query. This ensures characters such as parentheses, asterisks, backslashes, and null characters cannot alter the intended LDAP filter structure. Input handling must be implemented consistently in the application's authentication and authorization paths, rather than relying on client-side checks.

Deploy a WA is also correct when this refers to a web application firewall (WAF). A WAF can provide compensating protection at the HTTP/S boundary by identifying and blocking requests that contain recognizable LDAP-injection payloads. This is especially valuable while the application fix is being developed, tested, and deployed, and it provides an additional defense-in-depth control after secure input handling is in place. OWASP describes LDAP injection risks and recommends escaping untrusted input used in LDAP queries; its LDAP Injection Prevention Cheat Sheet provides encoding guidance. See OWASP LDAP Injection and OWASP LDAP Injection Prevention Cheat Sheet.

QUESTION NO: 30

Due to internal resource constraints, the management team has asked the principal security architect to recommend a solution that shifts most of the responsibility for application-level controls to the cloud provider. In the shared responsibility model, which of the following levels of service meets this requirement?

A. IaaS

B. SaaS

C. Faas

D. PaaS

ANSWER: B

Explanation:

SaaS is the appropriate service model because it places the greatest share of responsibility for operating and securing the application with the cloud provider. In a Software as a Service offering, the provider delivers a complete, ready-to-use application and is responsible for the application software, its supporting platform, runtime, operating systems, and underlying infrastructure. This substantially reduces the organization's internal operational burden for application-level activities such as application deployment, patching, maintenance, and operation of the service. The customer's responsibilities remain important but are generally focused on configuring the service securely, managing identities and access permissions, protecting and governing its data, and ensuring the service is used in accordance with organizational policy and regulatory obligations. NIST defines SaaS as a model in which the consumer uses the provider's applications and does not manage or control the underlying cloud infrastructure, including the application capabilities, except for limited user-specific configuration settings. This allocation directly meets the requirement to shift most application-control responsibility to the provider. See the NIST cloud-computing service-model definition and Microsoft's shared responsibility guidance.

QUESTION NO: 31

A vulnerability analyst identified a zero-day vulnerability in a company's internally developed software. Since the current vulnerability management system does not have any checks for this vulnerability, an engineer has been asked to create one.

Which of the following would be BEST suited to meet these requirements?

A. ARF

B. ISACs

C. Node.js

D. OVAL

ANSWER: D

Explanation:

OVAL is best suited because it is an open standard specifically designed to express and assess detailed endpoint configuration, inventory, and vulnerability conditions. An engineer can create an OVAL definition describing the precise indicators associated with the newly identified zero-day, such as a vulnerable software version, a particular file version or hash, a registry value, an installed package, or a configuration state. The vulnerability management platform can then evaluate that definition against managed assets and report which systems meet the defined vulnerable condition.

This approach is particularly useful when a commercial scanner has not yet released a signature or plug-in for a newly discovered issue. OVAL separates the definition of the condition to evaluate from the system characteristics collected from endpoints, allowing the organization to build a targeted, repeatable detection check and update it as the vulnerability research develops. OVAL was created to standardize machine-readable security assessment content and is commonly associated with automated vulnerability and compliance assessment workflows. See the MITRE OVAL project and NIST's overview of Security Content Automation Protocol, which incorporates OVAL for automated assessment content.

QUESTION NO: 32

A company wants to implement a new website that will be accessible via browsers with no mobile applications available. The new website will allow customers to submit sensitive medical information securely and receive online medical advice. The company already has multiple other websites where it provides various public health data and information. The new website must implement the following:

• The highest form Of web identity validation

• Encryption of all web transactions

• The strongest encryption in-transit

• Logical separation based on data sensitivity

Other things that should be considered include:

• The company operates multiple other websites that use encryption.

• The company wants to minimize total expenditure.

• The company wants to minimize complexity

Which of the following should the company implement on its new website? (Select TWO).

A. Wildcard certificate

B. EV certificate

C. Mutual authentication

D. Certificate pinning

E. SSO

F. HSTS

ANSWER: B F

Explanation:

An EV certificate and HSTS satisfy the stated browser-based security requirements. An EV certificate represents the highest validation tier for publicly trusted TLS certificates. Issuance requires the certificate authority to validate the requesting organization's legal identity, operational existence, authority to request the certificate, and control of the domain. This provides strong organizational identity assurance for a site that handles sensitive medical submissions. The CA/Browser Forum's EV Guidelines define these validation requirements: CA/Browser Forum EV Guidelines.

HSTS instructs compatible browsers to use HTTPS for the designated host for a specified period. Once received, the policy prevents ordinary HTTP access and causes browsers to automatically upgrade attempted HTTP connections to HTTPS, helping ensure that transactions remain encrypted in transit. HSTS can also be configured with the `includeSubDomains` directive, allowing the organization to apply a consistent HTTPS-only policy to related subdomains when appropriate. Deploying HSTS on the dedicated sensitive-data site supports logical separation from the company's public-information sites while using standard browser and web-server capabilities. Implementation and header directives are specified in RFC 6797.

QUESTION NO: 33

A company is migrating its data center to the cloud. Some hosts had been previously isolated, but a risk assessment convinced the engineering team to reintegrate the systems. Because the systems were isolated, the risk associated with vulnerabilities was low. Which of the following should the security team recommend be performed before migrating these servers to the cloud?

A. Performing patching and hardening

B. Deploying host and network IDS

C. Implementing least functionality and time-based access

D. Creating a honeypot and adding decoy files

ANSWER: A

Explanation:

Performing patching and hardening is the appropriate action before these formerly isolated hosts are connected to a cloud environment. Isolation can reduce the likelihood that an exploitable vulnerability will be reached, but it does not eliminate the vulnerability itself. Once the hosts are reintegrated and migrated, their network exposure, connectivity paths, and potential attack surface can increase substantially. The security team should first bring operating systems, applications, firmware, and supporting components to an approved patch baseline, prioritizing known exploitable and critical vulnerabilities. Hardening should then establish a secure configuration baseline for the intended cloud deployment, including removal or disabling of unnecessary software and services, secure configuration of accounts and authentication, restriction of administrative interfaces, host firewall configuration, and validation that only required ports and protocols are enabled. This work reduces the probability that legacy weaknesses accumulated during isolation will be exposed during or after migration. It also provides a controlled, documented baseline that can be assessed continuously in the cloud. NIST emphasizes enterprise patch-management processes for identifying, prioritizing, acquiring, testing, deploying, and verifying patches, while its server-security guidance describes applying secure configurations and minimizing unnecessary functionality. See NIST SP 800-40 Rev. 4 and NIST SP 800-123.

QUESTION NO: 34

A company based in the United States holds insurance details of EU citizens. Which of the following must be adhered to when processing EU citizens' personal, private, and confidential data?

A. The principle of lawful, fair, and transparent processing

B. The right to be forgotten principle of personal data erasure requests

C. The non-repudiation and deniability principle

D. The principle of encryption, obfuscation, and data masking

ANSWER: A

Explanation:

The principle of lawful, fair, and transparent processing is correct because it is a foundational requirement of the EU General Data Protection Regulation (GDPR) for the processing of personal data. A US-based organization can fall within the GDPR's territorial scope when its processing relates to offering goods or services to individuals in the EU or monitoring their behavior. Insurance details are personal data and may contain sensitive information, so the organization must establish and document an appropriate lawful basis before processing, such as performance of a contract, compliance with a legal obligation, or explicit consent where required. Fair processing means using the data in ways individuals would reasonably expect and avoiding unjustified adverse effects. Transparency requires clear privacy information describing what data is processed, why it is used, the lawful basis, retention periods, recipients, and the individual's applicable rights. These obligations apply throughout the data lifecycle, including collection, use, storage, sharing, and deletion. Article 5(1)(a) of the GDPR expressly establishes lawfulness, fairness, and transparency as a core processing principle. The European Data Protection Board also explains how GDPR obligations can apply to controllers and processors located outside the EU. See GDPR, Article 5 and EDPB Guidelines on territorial scope.

QUESTION NO: 35

A security administrator is setting up a virtualization solution that needs to run services from a single host. Each service should be the only one running in its environment. Each environment needs to have its own operating system as a base but share the kernel version and properties of the running host. Which of the following technologies would best meet these requirements?

A. Containers

B. Type 1 hypervisor

C. Type 2 hypervisor

D. Virtual desktop infrastructure

E. Emulation

ANSWER: A

Explanation:

Containers best meet these requirements because they isolate each service into a separate execution environment while using the host operating system's kernel. A container packages an application or service with its required user-space libraries, binaries, configuration, and dependencies. This gives each service a consistent, self-contained operating environment without requiring a separate guest kernel or a full virtual machine for every workload. Kernel sharing makes containers substantially more resource-efficient than machine virtualization, allowing multiple isolated services to run on one host with lower startup time and overhead. This design is particularly appropriate for microservices, where independently deployed services need process, filesystem, network, and resource isolation while remaining portable across compatible hosts. In security architecture, the shared-kernel model also means administrators should apply timely host-kernel patching, enforce least privilege, and use container runtime controls, since the host kernel is part of the trusted computing base for all containers. NIST describes application containers as technologies that provide operating-system-level virtualization and isolate applications while sharing the host OS kernel. Docker likewise explains that containers package an application and its dependencies into isolated processes running on a shared host kernel. See NIST SP 800-190, Application Container Security Guide and Docker: What is a container?.

QUESTION NO: 36

A company's Chief Information Security Officer is concerned that the company's proposed move to the cloud could lead to a lack of visibility into network traffic flow logs within the VP

C.

Which of the following compensating controls would be BEST to implement in this situation?

A. EDR

B. SIEM

C. HIDS

D. UEBA

ANSWER: B

Explanation:

SIEM is the best compensating control because it provides a centralized platform to collect, retain, normalize, correlate, and alert on security telemetry generated in cloud environments. The organization can configure its cloud platform's virtual network flow logs to be exported to a logging destination and then ingested into the SIEM alongside identity, workload, DNS, firewall, and application logs. This restores operational visibility even when the infrastructure is provider-managed and security teams do not have the same direct access to network devices that they had on premises.

Once flow-log data is available, the SIEM can correlate connections with authentication activity, threat-intelligence indicators, and events from other environments. That correlation supports detection and investigation of suspicious east-west traffic, unexpected egress, policy violations, and potentially compromised workloads. It also provides durable search, reporting, and alerting capabilities that help meet monitoring and audit requirements across hybrid or multicloud deployments. For example, AWS describes VPC Flow Logs as records of IP traffic to and from network interfaces, which can be published to supported logging destinations for analysis. See AWS VPC Flow Logs documentation and NIST's guidance on centralized log management in NIST SP 800-92.

QUESTION NO: 37

A software development company is implementing a SaaS-based password vault for customers to use. The requirements for the password vault include:

Vault encryption using a variable block and key size

Resistance to brute-force attacks

Which of the following should be implemented to meet these requirements? (Select two.)

A. PBKDF2

B. RC5

C. AES

D. P256

E. ECDSA

F. RIPEMD

ANSWER: A B

Explanation:

PBKDF2 is appropriate for resistance to brute-force password attacks because it is a password-based key-derivation function designed to make each password guess computationally expensive. It derives cryptographic keying material from a password using a salt and an adjustable iteration count. A unique, cryptographically random salt per vault and a sufficiently high, periodically reviewed work factor help prevent efficient precomputation and increase the cost of offline guessing attempts against stolen password-verifier data. RFC 8018 defines PBKDF2 and its use in password-based cryptography: RFC 8018.

RC5 meets the stated variable block-size and key-size encryption requirement. RC5 is a parameterized symmetric block cipher: its word size determines the block size, and it supports variable-length secret keys. Its parameters also include a configurable number of rounds, allowing the algorithm configuration to be expressed according to the required security and

performance characteristics. The RC5 specification documents its variable word-size/block-size and key-length design: RFC 2040. Therefore, PBKDF2 supplies password-guessing resistance, while RC5 supplies the specifically requested variable block and key sizing for vault encryption.

QUESTION NO: 38

The results of an internal audit indicate several employees reused passwords that were previously included in a published list of compromised passwords. The company has the following employee password policy:

Which of the following should be implemented to best address the password reuse issue? (Select two).

A. Increase the minimum age to two days.

B. Increase the history to 20.

C. Increase the character length to 12.

D. Add case-sensitive requirements to character class.

E. Decrease the maximum age to 30 days.

F. Remove the complexity requirements

G. Implement a compromised-password blocklist that rejects passwords found in published breach data.

ANSWER: G

Explanation:

Implementing a compromised-password blocklist directly addresses the audit finding because it prevents users from selecting passwords known to have been exposed in public breaches or credential dumps. During password creation or change, the identity system should compare a proposed password against a continuously updated denylist of commonly used, expected, and compromised values. If there is a match, the system rejects the password and requires the employee to choose a different one. This control is effective even when a password has never been used previously within the company, because the risk arises from its exposure elsewhere and the likelihood of credential-stuffing attacks. NIST Digital Identity Guidelines specifically require verifiers to compare prospective passwords against blocklists containing passwords obtained from previous breach corpuses and other commonly used or compromised values. This approach focuses the control on known high-risk credentials rather than relying on arbitrary expiration intervals or composition rules. The blocklist should be integrated with the organization's password-management or identity-provider workflow and maintained from reputable breach-intelligence sources. See NIST SP 800-63B password verifier requirements and Have I Been Pwned Pwned Passwords.

QUESTION NO: 39

A company is implementing centralized logging for cloud workloads that process regulated data. The security architect needs to ensure security events are available for investigations and cannot be altered by the administrators who manage the workloads.

Which of the following should the architect implement? (Choose two.)

A. Send logs to a separate security account.

B. Apply write-once retention to the centralized log repository.

C. Store logs only on the local disks of workload servers.

D. Reduce retention to 24 hours for all security events.

E. Grant application developers permission to delete centralized logs.

F. Disable logging during high-volume processing periods.

ANSWER: A B

Explanation:

Send logs to a separate security account provides administrative separation between workload operators and the personnel or systems responsible for security monitoring. If a workload administrator account is compromised, the attacker should not also be able to modify, disable, or delete the central logging destination. This separation improves the reliability of evidence available for incident response and compliance reviews.

Apply write-once retention to the centralized log repository protects retained records against modification and deletion for the defined retention period. Write-once retention, often implemented through object-lock or immutable-storage features, helps preserve log integrity even when an attacker gains elevated access to the source environment or attempts to cover activity by deleting evidence.

Storing logs only on local workload disks makes them vulnerable to tampering and loss during an incident. Reducing log retention and granting developers permission to delete logs undermine the stated investigative and integrity requirements. NIST log-management guidance recommends centralized collection, protection, and retention of logs appropriate to organizational requirements. See NIST SP 800-92, Guide to Computer Security Log Management.

QUESTION NO: 40

A company's software developers have indicated that the security team takes too long to perform application security tasks. A security analyst plans to improve the situation by implementing security into the SDLC. The developers have the following requirements:

1. The solution must be able to initiate SQL injection and reflected XSS attacks.

2. The solution must ensure the application is not susceptible to memory leaks.

Which of the following should be implemented to meet these requirements? (Select two).

A. Side-channel analysis

B. Protocol scanner

C. HTTP interceptor

D. DAST

E. Fuzz testing

F. SAST

G. SCAP

ANSWER: D F

Explanation:

DAST and SAST should be implemented together as automated security-testing controls within the SDLC. DAST evaluates a running web application from the outside, sending crafted HTTP requests and observing the application's responses. This makes it well suited to actively exercising input-handling flaws, including SQL injection and reflected cross-site scripting, in an environment where the application is deployed and reachable. It can be incorporated into CI/CD pipelines or pre-release testing so that developers receive repeatable findings sooner. OWASP describes dynamic testing as testing an application while it is running: OWASP source-code analysis tools.

SAST examines application source code, bytecode, or compiled artifacts without executing the application. It identifies insecure coding patterns and data/control-flow conditions early in development. For software written in memory-unsafe languages, static analysis can detect issues associated with incorrect allocation, release, and lifetime management that can result in memory leaks. Running SAST during code review or builds provides feedback before deployment and supports the goal of shifting application-security tasks left. The NIST Secure Software Development Framework specifically recommends using automated static analysis to identify code vulnerabilities as part of secure development practices: NIST SP 800-218, Secure Software Development Framework.

QUESTION NO: 41

A SOC analyst is reviewing malicious activity on an external, exposed web server. During the investigation, the analyst determines specific traffic is not being logged, and there is no visibility from the WAF for the web application.

Which of the following is the MOST likely cause?

A. The user agent client is not compatible with the WA

B. A certificate on the WAF is expired.

C. HTTPS traffic is not being forwarded to the WAF for decryption.

D. Old, vulnerable cipher suites are still being used.

ANSWER: C

Explanation:

HTTPS traffic is not being forwarded to the WAF for decryption is the most likely cause because a WAF requires visibility into the HTTP request contents—such as the URI, headers, parameters, cookies, and request body—to inspect, enforce rules, and generate useful security logs. When TLS-encrypted sessions terminate directly at the web server, or when encrypted traffic is routed around the WAF, the WAF cannot inspect the protected application-layer data. This creates a monitoring gap in which the web server may receive malicious requests but the SOC has no corresponding WAF telemetry or event records for those sessions. Correct deployment places the WAF inline or as a reverse proxy and terminates TLS there, or otherwise provides the WAF with access to decrypted traffic before it reaches the protected application. This also ensures that detection signatures and policy controls can operate on the actual web requests rather than opaque encrypted packets. NIST describes TLS as protecting application data in transit and highlights the importance of correctly managing TLS termination and inspection points; see NIST SP 800-52 Rev. 2. OWASP also identifies WAFs as controls that inspect HTTP(S) traffic and must be deployed where they can observe the relevant requests: OWASP Web Application Firewall.

QUESTION NO: 42

A cloud security engineer is setting up a cloud-hosted WAF. The engineer needs to implement a solution to protect the multiple websites the organization hosts. The organization websites are:

*

*

* campus.mycompany.com

* wiki. mycompany.org

The solution must save costs and be able to protect all websites. Users should be able to notify the cloud security engineer of any on-path attacks. Which of the following is the BEST solution?

A. Purchase one SAN certificate.

B. Implement self-signed certificates.

C. Purchase one certificate for each website.

D. Purchase one wildcard certificate.

ANSWER: A

Explanation:

Purchase one SAN certificate. A subject alternative name (SAN) certificate can contain multiple fully qualified domain names in one publicly trusted TLS certificate, including names in different parent domains, such as `campus.mycompany.com` and `wiki.mycompany.org`. This makes it suitable for a cloud-hosted WAF that terminates TLS for several organizational

websites while reducing the cost and administrative overhead of obtaining, deploying, renewing, and tracking separate certificates for each hostname. The WAF can present the same SAN certificate when it receives HTTPS requests for any hostname listed in the certificate's SAN extension.

Using a certificate issued by a trusted certificate authority also lets browsers validate the WAF's presented identity. If an attacker performs an on-path interception and presents an untrusted, expired, or hostname-mismatched certificate, the browser displays a certificate warning. Users can then report that warning to the cloud security engineer, supporting detection and notification of suspected on-path attacks. RFC 6125 specifies that service identity matching uses the subjectAltName extension and requires the presented DNS name to match the requested service name: RFC 6125. AWS likewise documents that an ACM certificate can include multiple domain names as SANs: AWS Certificate Manager documentation.

QUESTION NO: 43

An IT director is working on a solution to meet the challenge of remotely managing laptop devices and securely locking them down. The solution must meet the following requirements:

• Cut down on patch management.

• Make use of standard configurations.

• Allow for custom resource configurations.

• Provide access to the enterprise system from multiple types of devices.

Which of the following would meet these requirements?

A. MDM

B. Emulator

C. Hosted hypervisor

D. VDI

ANSWER: D

Explanation:

VDI is the appropriate solution because it centrally hosts user desktop operating systems and applications in virtual machines in the enterprise environment, rather than requiring each endpoint to be managed as a full, independently configured workstation. Administrators can maintain approved, standardized desktop images and apply operating-system and application updates to the central images, significantly reducing repetitive endpoint patching work. Desktop pools and image-based provisioning also support secure lockdown through centrally enforced policies and controlled configurations.

At the same time, VDI can provide customized virtual hardware and application resources for particular users, groups, or workloads. For example, a resource-intensive user can be assigned a desktop with additional processor, memory, graphics, or application capacity while other users receive a standard desktop configuration. Users connect to their assigned virtual desktop through supported client software or a browser from different endpoint types, allowing enterprise access from managed laptops, thin clients, tablets, or other authorized devices without placing the full enterprise desktop environment on each device. This centralized control and flexible remote delivery model directly addresses the stated management, standardization, customization, and multi-device access requirements. See Citrix: What is VDI? and Microsoft: Azure Virtual Desktop overview.

QUESTION NO: 44

A security engineer is implementing a server-side TLS configuration that provides forward secrecy and authenticated encryption with associated data. Which of the following algorithms, when combined into a cipher suite, will meet these requirements? (Choose three.)

A. EDE

B. CBC

C. GCM

D. AES

E. RSA

F. RC4

G. ECDSA

H. ECDHE

ANSWER: C D G H

Explanation:

A secure TLS configuration meeting these goals uses ephemeral elliptic-curve Diffie-Hellman key exchange, an authenticated server signature, and an AEAD encryption construction. **ECDHE** supplies forward secrecy because it creates an ephemeral session key agreement; compromise of the server's long-term private key does not retroactively expose captured sessions. **ECDSA** can authenticate the server's ephemeral key-exchange parameters using its certificate private key. **AES** is the symmetric block cipher used to encrypt application data, and **GCM** is the AES mode that provides authenticated encryption with associated data, protecting ciphertext confidentiality and integrity while authenticating unencrypted associated data. Together, these components describe the structure of a TLS suite such as ECDHE-ECDSA-AES-GCM. The original choices omit ECDHE, so it must be added to accurately satisfy the forward-secrecy requirement. NIST specifies GCM as an AEAD mode in SP 800-38D, while TLS 1.3 requires (EC)DHE-based key establishment and AEAD cipher suites, as described in RFC 8446.

QUESTION NO: 45

A new web server must comply with new secure-by-design principles and PCI DSS. This includes mitigating the risk of an on-path attack. A security analyst is reviewing the following web server configuration:

```
TLS_AES_256_GCM_SHA384
TLS_CHACHA20_POLY1305_SHA256
TLS_AES_128_GCM_SHA256
TLS_AES_128_CCM_8_SHA256
TLS_RSA_WITH_AES_128_CBC_SHA256
TLS_DHE_DSS_WITH_RC4_128_SHA
RSA_WITH_AES_128_CCM
```

Which of the following ciphers should the security analyst remove to support the business requirements?

A. TLS_AES_128_CCM_8_SHA256

B. TLS_DHE_DSS_WITH_RC4_128_SHA

C. TLS_CHACHA20_POLY1305_SHA256

D. TLS_AES_128_GCM_SHA256

ANSWER: B

Explanation:

TLS_DHE_DSS_WITH_RC4_128_SHA should be removed because it uses RC4, a legacy stream cipher with well-documented statistical biases that can expose plaintext when an attacker captures sufficient encrypted traffic. An on-path attacker is particularly relevant because such an attacker can observe, replay, and potentially influence TLS traffic; retaining known-weak cryptography unnecessarily increases the opportunity for compromise. The presence of ephemeral Diffie-Hellman in the suite does not make RC4 suitable, because every cryptographic component of a negotiated suite must meet the organization's security requirements. RFC 7465 formally prohibits RC4 cipher suites in TLS and states that TLS clients and servers must not negotiate RC4. Removing this suite supports secure-by-design by eliminating an obsolete, known-vulnerable algorithm rather than relying on compensating controls. It also aligns with PCI DSS expectations to use strong cryptography and secure protocols to protect cardholder data in transit. The server should be configured to negotiate only current TLS versions and strong authenticated-encryption cipher suites. See RFC 7465: Prohibiting RC4 Cipher Suites and the PCI DSS document library.

QUESTION NO: 46

An auditor needs to scan documents at rest for sensitive text. These documents contain both text and Images. Which of the following software functionalities must be enabled in the DLP solution for the auditor to be able to fully read these documents? (Select TWO).

A. Document interpolation

B. Regular expression pattern matching

C. Optical character recognition functionality

D. Baseline image matching

E. Advanced rasterization

F. Watermarking

ANSWER: A C

Explanation:

Document interpolation and optical character recognition functionality must be enabled to inspect all content in documents that combine native text with image-based text. Document interpolation enables the DLP platform to process and interpret the document's embedded or native textual content, allowing the inspection engine to access text across supported document formats. This provides the foundation for content discovery scans of data at rest.

Optical character recognition functionality is necessary for text that exists only as pixels, such as text in scanned documents, screenshots, photographs, or images embedded in office files and PDFs. OCR converts the visual representation of characters into machine-readable text that the DLP engine can inspect alongside the document's native text. With both capabilities enabled, the auditor can apply the organization's sensitive-data detection policies consistently to text regardless of whether it was stored directly as text or presented within an image.

This supports the CASP+ expectation that practitioners select appropriate technical controls for data discovery and protection across varied data formats. OCR is a standard mechanism for extracting text from images; see Microsoft Learn: Optical character recognition. CompTIA's CASP+ exam information is available at CompTIA Advanced Security Practitioner (CASP+).

QUESTION NO: 47

A company created an external, PHP-based web application for its customers. A security researcher reports that the application has the Heartbleed vulnerability. Which of the following would BEST resolve and mitigate the issue? (Select TWO).

A. Deploying a WAF signature

B. Fixing the PHP code

C. Changing the web server from HTTPS to HTTP

D. UsingSSLv3

E. Changing the code from PHP to ColdFusion

F. Updating the OpenSSL library

ANSWER: A F

Explanation:

Deploying a WAF signature and updating the OpenSSL library are the best paired actions for resolving and mitigating Heartbleed. Heartbleed was an implementation flaw in vulnerable versions of OpenSSL's TLS heartbeat extension, rather than a weakness in PHP application logic. Updating the OpenSSL library to a version containing the vendor fix removes the vulnerable heartbeat processing and is the essential remediation. The OpenSSL project's advisory identifies the affected versions and states that the flaw was corrected in OpenSSL 1.0.1g; systems should use a supported, patched version supplied by their platform or application vendor.

Deploying a WAF signature provides a compensating control while patching is planned, tested, and deployed. A suitably configured signature can recognize malformed TLS heartbeat traffic associated with exploitation attempts and block it before the request reaches the vulnerable service. This is particularly useful for an externally accessible customer application, where the exposure may be continuously scanned or attacked. Following remediation, the organization should also replace the service's TLS private key and certificates where compromise is plausible, because Heartbleed could disclose process memory containing sensitive cryptographic material or session data. See the OpenSSL security advisory and the OWASP Heartbleed Bug overview.

QUESTION NO: 48

A major broadcasting company that requires continuous availability to streaming content needs to be resilient against DDoS attacks Which of the following is the MOST important infrastructure security design element to prevent an outage7

A. Supporting heterogeneous architecture

B. Leveraging content delivery network across multiple regions

C. Ensuring cloud autoscaling is in place

D. Scaling horizontally to handle increases in traffic

ANSWER: B

Explanation:

Leveraging content delivery network across multiple regions is the most important design element because a CDN places distributed edge servers between viewers and the streaming provider's origin infrastructure. Viewers retrieve cacheable streaming assets and other content from geographically close edge locations, which distributes normal demand and reduces the amount of traffic that must reach the origin. During a distributed denial-of-service event, a globally distributed CDN also provides a large, geographically dispersed traffic-absorption layer. Its edge network can identify and filter malicious requests, apply rate limiting and web-application protections, and route legitimate users to available regional points of presence before attack traffic can exhaust origin bandwidth, connection capacity, or compute resources. This architecture directly supports the availability objective: even when one region or edge location is under severe load, requests can be served through other locations and the origin remains insulated from much of the attack volume. For a broadcaster whose service depends on uninterrupted public streaming access, placing DDoS mitigation at the CDN edge is therefore a foundational availability control rather than merely a capacity enhancement. Cloudflare describes how its anycast CDN network distributes traffic and mitigates DDoS attacks at the network edge: Cloudflare DDoS mitigation. AWS likewise explains that CloudFront can help protect applications by absorbing and filtering traffic before it reaches origins: AWS DDoS resiliency guidance.

QUESTION NO: 49

Immediately following the report of a potential breach, a security engineer creates a forensic image of the server in question as part of the organization incident response procedure. Which of the must occur to ensure the integrity of the image?

A. The image must be password protected against changes.

B. A hash value of the image must be computed.

C. The disk containing the image must be placed in a seated container.

D. A duplicate copy of the image must be maintained

ANSWER: B

Explanation:

A hash value of the image must be computed. A cryptographic hash establishes a repeatable integrity baseline for the forensic image at the time it is acquired. The investigator records the hash value, typically using a strong algorithm such as SHA-256, and can recalculate it whenever the image is transferred, stored, examined, or presented as evidence. Matching values demonstrate that the bit-level contents of the image have not changed since the baseline was created. This validation is central to forensic soundness and supports the organization's ability to document and defend the evidence-handling process. Integrity validation should be documented with acquisition details, timestamps, the personnel involved, the tool used, and the resulting hash value as part of the evidence record and chain of custody. NIST's incident-handling guidance identifies preserving evidence and maintaining its integrity as key incident-response activities, while its forensic guidance discusses using hashes to verify that copied forensic media remains unchanged. See NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide and NIST SP 800-86, Guide to Integrating Forensic Techniques into Incident Response.

QUESTION NO: 50

A security administrator wants to detect a potential forged sender claim in tt-e envelope of an email. Which of the following should the security administrator implement? (Select TWO).

A. MX record

B. DMARC

C. SPF

D. DNSSEC

E. S/MIME

F. TLS

ANSWER: B C

Explanation:

SPF and **DMARC** provide complementary controls for identifying forged email sender claims. SPF lets the owner of a sending domain publish a DNS policy identifying the IP addresses and mail servers authorized to send messages using that domain in the SMTP envelope sender (the MAIL FROM/Return-Path identity). A receiving mail system evaluates the connecting sender against that policy; a failure can indicate that an unauthorized server is attempting to send mail for the claimed domain.

DMARC builds on email authentication results and applies domain alignment to the user-visible From domain. It evaluates whether SPF authentication aligns with the message's From identity and enables the domain owner to publish an enforcement policy, such as monitoring, quarantine, or rejection. DMARC also supplies reporting capabilities, allowing administrators to identify unauthorized sources and investigate spoofing attempts at scale. Used together, SPF validates authorization for the envelope-sender domain, while DMARC makes that validation actionable in relation to the sender identity presented to the recipient. This combination is a core best practice for reducing domain spoofing and strengthening inbound email authentication. See RFC 7208: Sender Policy Framework and RFC 7489: DMARC.

QUESTION NO: 51

A security researcher detonated some malware in a lab environment and identified the following commands running from the EDR tool:

```
netsh advfirewall set allprofiles firewall policy blockinbound, blockoutbound
netsh advfirewall set allprofiles state on
init.ps1 -win32_shadow copy
```

With which of the following MITRE ATT & CK TTPs is the command associated? (Select TWO).

A. Indirect command execution

B. OS credential dumping

C. Inhibit system recovery

D. External remote services

E. System information discovery

F. Network denial of service

ANSWER: B E

Explanation:

OS credential dumping is associated with use of Mimikatz commands that access credential material held by Windows, commonly including credentials or password hashes from LSASS memory. Adversaries use this technique to obtain authentication material that can enable account compromise, privilege escalation, and movement to other systems. MITRE ATT&CK identifies this behavior as OS Credential Dumping and notes that adversaries may dump credentials from operating-system components such as LSASS. See the MITRE ATT&CK OS Credential Dumping technique.

System information discovery is associated with WMIC system-information queries. Windows Management Instrumentation command-line activity can collect details about the affected endpoint, including its operating-system configuration, hardware, installed components, hostname, and related environmental attributes. Malware commonly performs this reconnaissance after execution to profile the host, determine whether it meets targeting criteria, and guide later actions. MITRE ATT&CK categorizes collection of this type of host and operating-system detail as System Information Discovery. See the MITRE ATT&CK System Information Discovery technique.

QUESTION NO: 52

A company wants to quantify and communicate the effectiveness of its security controls but must establish measures. Which of the following is MOST likely to be included in an effective assessment roadmap for these controls?

A. Create a change management process.

B. Establish key performance indicators.

C. Create an integrated master schedule.

D. Develop a communication plan.

E. Perform a security control assessment.

ANSWER: B

Explanation:

Establish key performance indicators is correct because KPIs provide defined, repeatable measures for showing whether security controls are operating as intended and contributing to stated security objectives. An assessment roadmap needs measurable criteria so stakeholders can track control performance over time, identify trends, set targets and thresholds, prioritize remediation, and communicate the security program's value in business-relevant terms. Useful security-control KPIs may measure such outcomes as remediation time for critical vulnerabilities, patch compliance, percentage of privileged

accounts reviewed, control-coverage rates, or the percentage of systems meeting a defined configuration baseline. For a KPI to be effective, the organization should document its purpose, data source, calculation method, collection frequency, owner, target, and reporting audience. This makes results consistent and auditable rather than dependent on subjective assessments. NIST's security performance measurement guidance describes using measures to facilitate decision-making, improve accountability, and demonstrate how information security activities support organizational goals. NIST's control-assessment guidance also emphasizes determining the effectiveness of implemented controls through defined assessment procedures and evidence. Together, these practices support an assessment roadmap that produces meaningful, communicable evidence of control effectiveness. See NIST SP 800-55 Rev. 1 and NIST SP 800-53A Rev. 5.

QUESTION NO: 53

A security operations team is investigating a suspected compromise of a Linux application server. The team believes an attacker may have used a fileless payload and established an encrypted command-and-control session.

Which of the following evidence sources should the team collect FIRST to preserve the most volatile evidence? (Choose two.)

A. A memory image

B. Active network connection information

C. A full filesystem backup

D. Archived vulnerability scan reports

E. A current asset inventory export

ANSWER: A B

Explanation:

A memory image is essential because it can preserve volatile artifacts that may not exist on disk, including running processes, injected code, loaded modules, decrypted payloads, command history, credentials, and cryptographic material. A fileless payload may be lost when the system is powered down or restarted.

Active network connection information is also highly volatile and can identify current remote endpoints, local processes, ports, connection states, and potentially the encrypted command-and-control channel. Collecting this information promptly helps responders scope containment actions and correlate network activity with processes found in memory.

NIST recommends collecting volatile data before less volatile evidence during incident handling and digital forensic activities. See NIST SP 800-86, Guide to Integrating Forensic Techniques into Incident Response.

QUESTION NO: 54

A company's employees are not permitted to access company systems while traveling internationally. The company email system is configured to block logins based on geographic location, but some employees report their mobile phones continue to sync email traveling . Which of the following is the MOST likely explanation? (Select TWO.)

A. Outdated escalation attack

B. Privilege escalation attack

C. VPN on the mobile device

D. Unrestricted email administrator accounts

E. Chief use of UDP protocols

F. Disabled GPS on mobile devices

ANSWER: C F

Explanation:

VPN on the mobile device is correct because a VPN can tunnel the phone's traffic to an egress point in the company's home country or another permitted region. The email service then evaluates the source public IP address of the VPN gateway rather than the phone's actual international network location. As a result, an IP-based geographic access policy can regard the connection as originating from an approved location and permit synchronization. Geographic and network-location controls are an important part of the enterprise security architecture and risk-management focus covered by CompTIA's advanced practitioner certification materials: CompTIA Advanced Security Practitioner.

Disabled GPS on mobile devices is also correct where the organization's mobile access controls rely on device-reported location information as part of location validation. If GPS location services are unavailable or disabled, the device cannot provide reliable geographic evidence to support enforcement or detection of an international location. Mobile-device location reporting depends on enabled location services and available location sources, as described in Apple Platform Security: Location Services. A robust implementation should use multiple trustworthy signals, such as authenticated network egress, device compliance, and conditional-access telemetry, rather than assuming a single device-location signal always proves physical location.

QUESTION NO: 55

Which of the following are risks associated with vendor lock-in? (Choose two.)

A. The client can seamlessly move data.

B. The vendor can change product offerings.

C. The client receives a sufficient level of service.

D. The client experiences decreased quality of service.

E. The client can leverage a multicloud approach.

F. The client experiences increased interoperability.

ANSWER: B D

Explanation:

Vendor lock-in occurs when a client becomes substantially dependent on a provider's proprietary technologies, formats, interfaces, contractual terms, or integrated services, making migration costly, complex, or operationally disruptive. **The vendor can change product offerings.** is a risk because the client's limited ability to move workloads or data gives the provider greater leverage to modify features, retire services, alter licensing or packaging, and change commercial terms. The client may have little practical choice but to accommodate the change while a migration is difficult.

The client experiences decreased quality of service. is also a risk. Once switching providers is prohibitively expensive or technically challenging, the customer may be exposed to degraded support, performance, reliability, or service responsiveness without an easy alternative. Vendor lock-in therefore weakens the customer's negotiating position and reduces the flexibility to select a provider that better meets evolving technical and business requirements. Cloudflare describes lock-in as dependency that can make moving data, applications, and services difficult: What is vendor lock-in?. NIST also identifies portability and interoperability concerns as important considerations in cloud computing: NIST SP 800-146.

QUESTION NO: 56

A security analyst is designing a touch screen device so users can gain entry into a locked room by touching buttons numbered zero through nine in a specific numerical sequence. The analyst designs thekeypad so that the numbers are randomly presented to the user each time the device is used. Which of the following best describes the design trade-offs? (Select two.)

A. The risk of someone overseeing a pattern as a user enters the numbers is decreased.

B. The routines to generate the random sequences are trivial to implement.

C. This design makes entering numbers more difficult for users.

D. The device needs to have additional power to compute the numbers.

E. End users will have a more difficult time remembering the access numbers.

F. Weak or easily guessed access numbers are more likely.

ANSWER: A C

Explanation:

“The risk of someone overseeing a pattern as a user enters the numbers is decreased.” is correct because a randomly arranged keypad prevents an observer from learning the secret merely by watching the same screen locations being touched. With a fixed arrangement, observing a sequence of positions can reveal the associated digits directly. When the digit locations change for each authentication attempt, the observed touch pattern is not reusable as the same numerical sequence, improving resistance to casual observation and shoulder-surfing attacks. NIST recognizes observation resistance as an important consideration for authentication mechanisms used in environments where an attacker may watch the claimant enter a secret; see NIST SP 800-63B.

“This design makes entering numbers more difficult for users.” is also correct because users cannot rely on learned spatial or muscle-memory patterns. They must visually locate each digit before every touch, which adds time and attention to the entry process. This is a genuine usability cost associated with the security benefit: the interface reduces predictability for observers but increases the cognitive effort and interaction time required of authorized users. Usable security guidance emphasizes that authentication designs must balance security controls with the burden imposed on legitimate users; see NIST Usable Cybersecurity Primer.

QUESTION NO: 57

A company just released a new video card. Due to limited supply and nigh demand, attackers are employing automated systems to purchase the device through the company ' s web store so they can resell it on the secondary market. The company ' s Intended customers are frustrated. A security engineer suggests implementing a CAPTCHA system on the web store to help reduce the number of video cards purchased through automated systems. Which of the following now describes the level of risk?

A. InherentLow

B. Mitigated

C. Residual

D. Transferred

ANSWER: C

Explanation:

Residual is correct because it describes the risk that remains after an organization applies a security control. In this scenario, automated purchasing bots create a business and customer-impact risk: inventory is rapidly consumed by resellers, legitimate buyers cannot obtain the product, and the company may experience reputational or revenue consequences. Implementing a CAPTCHA is a mitigating control intended to reduce the likelihood of successful automated transactions by requiring a challenge that is more difficult for unsophisticated bots to complete.

The CAPTCHA does not eliminate the threat entirely. Attackers may use CAPTCHA-solving services, compromised accounts, browser automation that can pass some challenges, or human-assisted purchasing methods. The organization therefore retains some exposure even after the control is deployed. That remaining exposure is residual risk. NIST describes risk assessment and risk response as considering how controls affect likelihood and impact, while recognizing that controls may not fully remove risk. See NIST SP 800-30 Rev. 1. CAPTCHA implementations are also commonly used to help distinguish human activity from automated abuse, as discussed in the OWASP Blocking Brute Force Attacks Cheat Sheet.

QUESTION NO: 58

An organization is concerned with a critical legacy application that is only supported on an end-of-life operating system. The organization would like to limit network communication from this device to only a select number of other devices. Which of the following primary and compensating controls should the organization use to reduce risk? (Select two).

A. Host-based firewalls

B. UEBA

C. HIDS

D. Antivirus

E. EDR

F. SEDs

ANSWER: A C

Explanation:

Host-based firewalls provide the primary preventive control because they can enforce a default-deny, allow-list policy directly on the legacy device. Rules can restrict inbound and outbound traffic by source and destination address, protocol, port, and, where supported, application. This permits the critical application to communicate only with its explicitly authorized peer systems while reducing the exposed network attack surface of an operating system that no longer receives vendor security updates. NIST identifies boundary protection as a control area that includes monitoring and controlling communications at managed interfaces; an endpoint firewall applies this principle at the host interface.

HIDS is an appropriate compensating detective control for the end-of-life system. A host intrusion detection system can monitor host activity such as suspicious processes, unauthorized configuration or file changes, log events, and indicators of compromise. Alerts from the HIDS give security personnel visibility into potential misuse or compromise that may occur despite the restrictive firewall policy, supporting a layered approach where the legacy platform cannot be remediated through normal patching. NIST describes information-system monitoring as a key control family for detecting attacks and indicators of potential compromise. See NIST SP 800-53 Rev. 5 and the NIST Cybersecurity Framework.

QUESTION NO: 59

A global organization's Chief Information Security Officer (CISO) has been asked to analyze the risks involved in a plan to move the organization's current MPLS-based WAN network to use commodity Internet and SD-WAN hardware. The SD-WAN provider is currently highly regarded but Is a regional provider. Which of the following is MOST likely identified as a potential risk by the CISO?

A. The SD-WAN provider would not be able to handle the organization's bandwidth requirements.

B. The operating costs of the MPLS network are too high for the organization.

C. The SD-WAN provider uses a third party for support.

D. Internal IT staff will not be able to properly support remote offices after the migration.

ANSWER: A

Explanation:

The SD-WAN provider would not be able to handle the organization's bandwidth requirements is the most likely risk because the provider’s regional scope is directly relevant to whether it can deliver consistent capacity and service performance for a global enterprise. Moving from MPLS to Internet-based SD-WAN changes the organization’s dependency model: application availability and site connectivity depend not only on local Internet circuits, but also on the provider’s ability to operate sufficient points of presence, backbone or cloud connectivity, provisioning processes, monitoring, and support across every required geography. A provider can be highly regarded within its established region while lacking the network scale, carrier relationships, or operational capacity needed for enterprise traffic volumes worldwide. The CISO should therefore assess

peak and sustained bandwidth requirements, regional coverage, performance commitments, capacity-expansion procedures, and enforceable service-level agreements before approving the migration. This is an external-service supply-chain risk that should be evaluated before relying on the service for critical operations. NIST guidance calls for managing risks associated with externally provided services and supply-chain dependencies; see NIST SP 800-53, Rev. 5 and NIST SP 800-161 Rev. 1.

QUESTION NO: 60

A PKI engineer is defining certificate templates for an organization's CA and would like to ensure at least two of the possible SAN certificate extension fields populate for documentation purposes. Which of the following are explicit options within this extension? (Select two).

A. Type

B. Email

C. OCSP responder

D. Registration authority

E. Common Name

F. DNS name

ANSWER: B F

Explanation:

Email and **DNS name** are explicit GeneralName forms defined for the Subject Alternative Name (SAN) extension in X.509 certificates. A DNS name SAN binds the certificate to a fully qualified domain name, such as `service.example.com`, and is the standard identity form used by TLS clients to validate a server hostname. An email SAN, formally the `rfc822Name` form, binds an email address to the certificate subject and can support certificate use cases involving secure email identities. A SAN extension can contain multiple entries, including several DNS names, several email addresses, or a mixture of supported GeneralName forms. This makes it possible for a certificate template to document and carry more than one identity associated with the certificate holder or service. The Internet X.509 Public Key Infrastructure profile specifies both `rfc822Name` and `dNSName` as choices in the SubjectAltName GeneralName structure. Microsoft's certificate documentation also identifies DNS and email as supported SAN types for certificate requests and enrollment scenarios. See RFC 5280, Subject Alternative Name and Microsoft Learn: Configure certificate templates.

QUESTION NO: 61

A Chief Information Security Officer (CISO) received a call from the Chief Executive Officer (CEO) about a data breach from the SOC lead around 9:00 a.m. At 10:00 a.m. The CEO informs the CISO that a breach of the firm is being reported on national news. Upon investigation, it is determined that a network administrator has reached out to a vendor prior to the breach for information on a security patch that failed to be installed. Which of the following should the CISO do to prevent this from happening again?

A. Properly triage events based on brand imaging and ensure the CEO is on the call roster.

B. Create an effective communication plan and socialize it with all employees.

C. Send out a press release denying the breach until more information can be obtained.

D. Implement a more robust vulnerability identification process.

ANSWER: B

Explanation:

Create an effective communication plan and socialize it with all employees. The scenario identifies a breakdown in communication and escalation: a network administrator had awareness of a potentially material patch issue through contact with the vendor, yet the appropriate security and executive stakeholders were not effectively informed before the incident became publicly reported. A documented incident- and vulnerability-communication plan establishes who must be notified, what events and conditions trigger escalation, required communication channels, notification time frames, decision authorities, and procedures for coordinating with external parties. Socializing the plan ensures that administrators, SOC personnel, leadership, legal staff, and communications teams understand their responsibilities and can act consistently under time pressure. This includes communicating material vulnerability and patch-deployment exceptions to the CISO or incident-response authority early enough to assess risk, apply compensating controls, and prepare accurate internal and external messaging. NIST incident-response guidance emphasizes defining communication and coordination processes as part of incident-response preparation, while CISA's vulnerability-response guidance highlights coordinated handling and communication of vulnerability information. See NIST SP 800-61 Rev. 3 and CISA Vulnerability Management.

QUESTION NO: 62

An organization is preparing to process personal data from customers in multiple jurisdictions through a new cloud-based analytics platform. The organization needs to reduce regulatory exposure before data is ingested into the platform.

Which of the following actions should the security architect recommend? (Choose two.)

A. Classify data and identify applicable handling requirements.

B. Document data flows and processing locations.

C. Deploy a CDN in every cloud region.

D. Increase the analytics cluster CPU allocation.

E. Disable all audit logging for personal-data workloads.

F. Use a shared service account for all data processors.

ANSWER: A B

Explanation:

Classify data and identify applicable handling requirements is necessary because the organization must understand what personal, regulated, confidential, and sensitive data will be processed before it can apply appropriate controls. Classification supports decisions about retention, encryption, access restrictions, transfer limitations, and required legal or contractual safeguards.

Document data flows and processing locations enables the organization to identify where data is collected, stored, processed, replicated, and accessed. This is critical when data may cross jurisdictional boundaries or be handled by cloud providers and subprocessors. A documented data-flow inventory supports privacy impact assessments, residency decisions, processor oversight, and incident-response planning.

Deploying a CDN or increasing compute capacity may improve service performance, but neither determines whether the intended processing complies with applicable privacy obligations. NIST privacy guidance emphasizes understanding data processing activities and managing privacy risks across the data lifecycle. See NIST Privacy Framework and the official GDPR text.

QUESTION NO: 63

A company that uses AD is migrating services from LDAP to secure LDAP. During the pilot phase, services are not connecting properly to secure LDAP. Block is an except of output from the troubleshooting session:

Which of the following BEST explains why secure LDAP is not working? (Select TWO.)

A. The clients may not trust idapt by default.

B. The secure LDAP service is not started, so no connections can be made.

C. Danvills.com is under a DDoS-inator attack and cannot respond to OCSP requests.

D. Secure LDAP should be running on UDP rather than TCP.

E. The company is using the wrong port. It should be using port 389 for secure LDAP.

F. Secure LDAP does not support wildcard certificates.

G. The clients may not trust Chicago by default.

H. The company is using the wrong port; LDAPS normally uses TCP port 636.

ANSWER: B H

Explanation:

The secure LDAP service is not started, so no connections can be made. The troubleshooting output indicates that the domain controller is listening for standard LDAP but does not have an active LDAPS listener. In Active Directory, LDAP and LDAP over SSL/TLS are separate listener endpoints. A functioning standard LDAP service on its usual port does not by itself make the secure LDAP endpoint available. The domain controller must have an appropriate server certificate installed and the AD DS service must successfully initialize the LDAPS listener.

The company is using the wrong port; LDAPS normally uses TCP port 636. Standard LDAP conventionally uses TCP/UDP port 389, while LDAPS uses TCP port 636. Therefore, clients configured for secure LDAP must connect to the LDAPS TCP endpoint rather than treating port 389 as the implicit LDAPS port. Administrators should verify that TCP 636 is listening after installing a valid certificate whose subject/SAN matches the directory server name and whose issuing CA is trusted by clients. Microsoft documents the Active Directory LDAPS port requirements and certificate prerequisites in its LDAP over SSL connection troubleshooting guidance and LDAPS certificate deployment guidance.

QUESTION NO: 64

A security administrator configured the account policies per security implementation guidelines. However, the accounts still appear to be susceptible to brute-force attacks. The following settings meet the existing compliance guidelines:

Must have a minimum of 15 characters

Must use one number

Must use one capital letter

Must not be one of the last 12 passwords used

Which of the following policies should be added to provide additional security?

A. Shared accounts

B. Password complexity

C. Account lockout

D. Password history

E. Time-based logins

ANSWER: C

Explanation:

Account lockout is the appropriate additional policy because it directly limits repeated authentication attempts against an account. Even when passwords meet length, composition, and history requirements, an attacker may continue making online password guesses until a valid credential is found. An account-lockout policy establishes a failed-attempt threshold and, once that threshold is reached, temporarily locks the account or requires administrative action before further authentication can occur. This slows or stops automated guessing activity and gives defenders an opportunity to detect, investigate, and respond to a suspected attack.

The policy should be configured with a threshold, lockout duration, and reset counter that balance resistance to brute-force attempts with operational needs. In particular, the threshold should be low enough to constrain rapid guessing while avoiding unnecessary user disruption or a denial-of-service condition caused by intentional lockouts. Windows documents that the Account lockout threshold policy determines the number of invalid sign-in attempts that trigger a locked account: Microsoft Learn: Account lockout threshold. This control aligns with the account-management recommendations in NIST SP 800-63B, which addresses throttling failed authentication attempts.

QUESTION NO: 65

city government's IT director was notified by the City council that the following cybersecurity requirements must be met to be awarded a large federal grant:

+ Logs for all critical devices must be retained for 365 days to enable monitoring and threat hunting.

+ All privileged user access must be tightly controlled and tracked to mitigate compromised accounts.

+ Ransomware threats and zero-day vulnerabilities must be quickly identified.

Which of the following technologies would BEST satisfy these requirements? (Select THREE).

A. Endpoint protection

B. Log aggregator

C. Zero trust network access

D. PAM

E. Cloud sandbox

F. SIEM

G. NGFW

ANSWER: B D F

Explanation:

Log aggregator is appropriate because it centrally collects event records from critical devices and supports the defined long-term retention period. Centralized collection gives security personnel a consistent historical data set for investigations, monitoring, and threat-hunting queries. NIST's log-management guidance describes the need to establish log retention requirements and maintain logs in centralized infrastructure appropriate to organizational needs.

PAM is the direct control for privileged accounts. A privileged-access-management deployment can enforce least privilege, require strong authentication and approval workflows, vault and rotate credentials, and record or audit administrative sessions. These controls provide both the restriction and accountability needed when a privileged credential is compromised or misused. NIST identifies least privilege as a core access-control principle for limiting authorized access to only what is necessary.

SIEM is appropriate for rapid security identification because it ingests and correlates security events across systems, applies detection analytics and threat intelligence, and generates alerts for suspicious behavior associated with ransomware and emerging exploit activity. Its correlation and alerting capabilities make retained telemetry actionable for timely investigation and response. See NIST SP 800-92, Guide to Computer Security Log Management and NIST SP 800-53, Access Control and Least Privilege.

QUESTION NO: 66

A software company is developing an application in which data must be encrypted with a cipher that requires the following:

* Initialization vector

* Low latency

* Suitable for streaming

Which of the following ciphers should the company use?

A. Cipher feedback

B. Cipher block chaining message authentication code

C. Cipher block chaining

D. Electronic codebook

ANSWER: A

Explanation:

Cipher feedback is the appropriate choice because it is a block-cipher mode designed to operate in a stream-like manner. It begins with an initialization vector, which seeds the feedback process and must be handled securely so that encryption does not reuse the same starting state with the same key. The mode produces output that is combined with plaintext in small units, allowing data to be encrypted as it becomes available rather than requiring an entire message or a padded final block. This characteristic supports low-latency processing and makes the mode suitable for continuous or streaming data.

NIST describes cipher feedback as a confidentiality mode that can use segment sizes smaller than the underlying cipher block, including byte-oriented segments. That capability is the key feature matching the application's streaming requirement. In a real deployment, the company should also provide integrity and authentication through an authenticated-encryption design or a separate, correctly implemented authentication mechanism; confidentiality-only encryption modes do not detect modification of ciphertext. See NIST's Recommendation for Block Cipher Modes of Operation and its full publication PDF for the specification and IV requirements.

QUESTION NO: 67

A security engineer needs to select the architecture for a cloud database that will protect an organization's sensitive data. The engineer has a choice between a single-tenant or a multitenant database architecture offered by a cloud vendor. Which of the following best describes the security benefits of the single-tenant option? (Select two).

A. Most cost-effective

B. Ease of backup and restoration

C. High degree of privacy

D. Low resilience to side-channel attacks

E. Full control and ability to customize

F. Increased geographic diversity

ANSWER: C E

Explanation:

High degree of privacy is a key benefit of a single-tenant database architecture because the customer's database environment and associated resources are dedicated to that one organization rather than shared with unrelated tenants. This stronger logical and often physical separation reduces exposure to cross-tenant access paths and supports clearer security boundaries for sensitive data. Dedicated tenancy can be particularly valuable when an organization has stringent privacy, regulatory, data-residency, or workload-isolation requirements.

Full control and ability to customize is also a security benefit because a dedicated environment generally permits the organization to apply security configurations, hardening standards, monitoring integrations, encryption-key approaches, network controls, and change-management processes that align with its own risk appetite. This control makes it easier to tailor the environment to specialized compliance obligations and to maintain a consistent security baseline without needing to accommodate the shared operational model of multiple customers. CompTIA's CASP+ objectives emphasize selecting secure architectures and applying appropriate security controls based on organizational requirements. See the CompTIA CASP+ certification page and the NIST guidance on security and privacy in public cloud computing.

QUESTION NO: 68

A cybersecurity engineer analyst a system for vulnerabilities. The tool created an OVAL. Results document as output. Which of the following would enable the engineer to interpret the results in a human readable form? (Select TWO.)

A. Text editor

B. XML editor

C. Event Viewer

D. XML style sheet

E. SCAP tool

F. Debugging utility

ANSWER: B D

Explanation:

An **XML editor** can open and display the structured XML elements, attributes, tests, and result states contained in an OVAL Results document. This lets the engineer inspect the raw assessment output directly and understand the individual definitions that evaluated to true, false, unknown, or error. Although raw XML is structured rather than report-like, an XML-aware editor makes the hierarchical content readable and navigable for an analyst.

An **XML style sheet**, typically an XSLT stylesheet, can transform the OVAL Results XML into a presentation format such as HTML. The resulting report is much easier for a human to review because it can organize findings, affected systems, definition identifiers, and evaluation outcomes into readable sections or tables. OVAL is defined as an XML-based language, and the OVAL Results schema represents assessment results as XML. Applying a stylesheet is therefore a standard way to convert the machine-readable results into a human-readable report. The OVAL specification and schema resources are available from OVAL Project, and NIST describes OVAL as part of the SCAP ecosystem at NIST SCAP OVAL specifications.

QUESTION NO: 69 - (DRAG DROP)

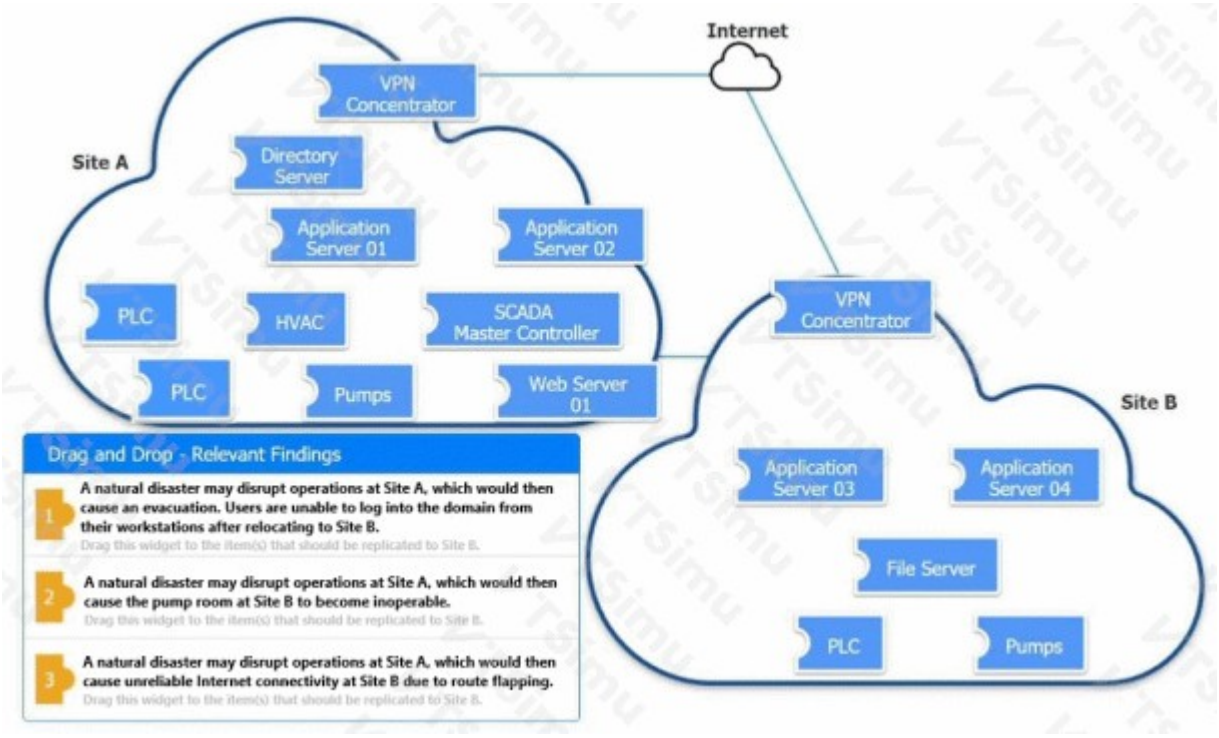

An organization is planning for disaster recovery and continuity of operations.

INSTRUCTIONS

Review the following scenarios and instructions. Match each relevant finding to the affected host.

After associating scenario 3 with the appropriate host(s), click the host to select the appropriate corrective action for that finding.

Each finding may be used more than once.

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

ANSWER:

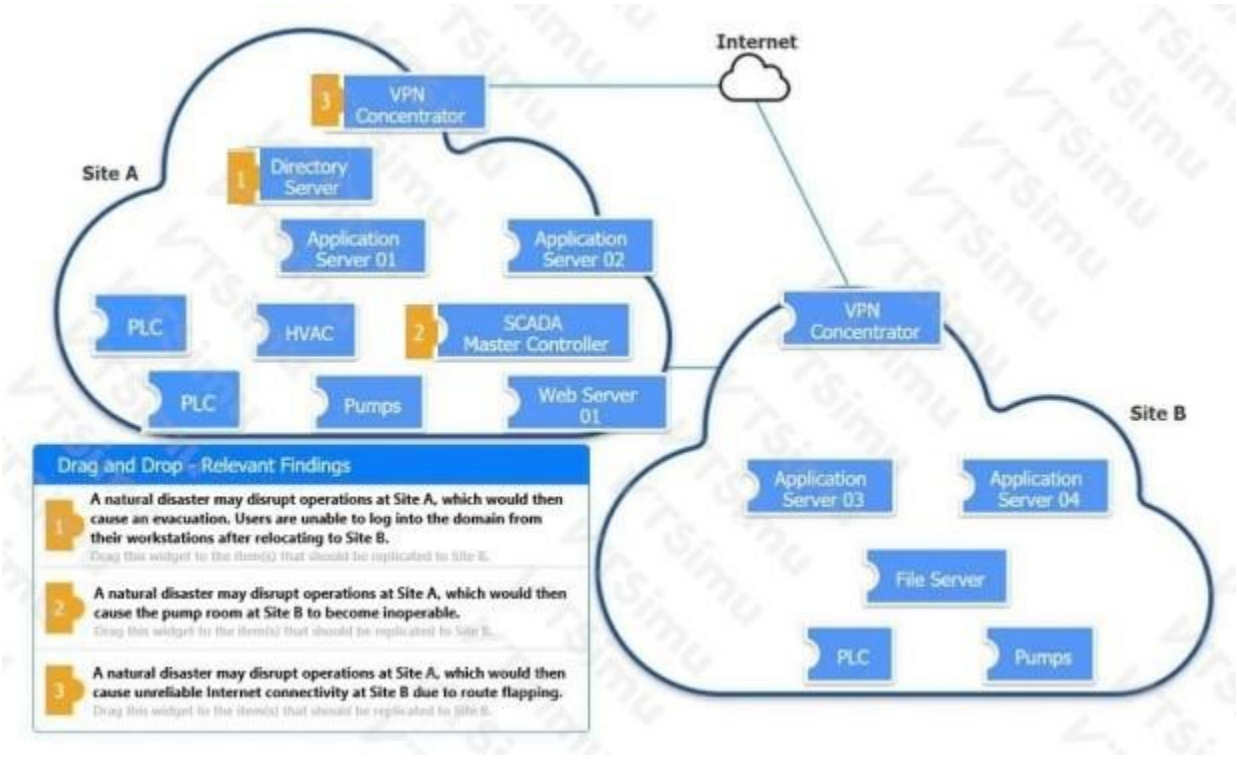

Explanation:

The disaster-recovery design should preserve the services that allow Site B to continue both business access and industrial operations when Site A is unavailable. Replicating the Directory Server to Site B lets users who evacuate from Site A authenticate locally and regain access to domain resources without depending on the impaired site. Active Directory's disaster-recovery guidance emphasizes maintaining recoverable, available directory services because identity, authentication, and authorization underpin dependent services. See Microsoft's Active Directory forest recovery guide.

Replicating the SCADA Master Controller provides the supervisory control capability required to operate the pump environment from the recovery site. In an operational-technology setting, the master control function coordinates field equipment and enables the pumps to be monitored and controlled after a site disruption. This supports continuity of essential physical processes, consistent with NIST's guidance that contingency planning addresses restoration and continuation of mission-essential systems and services; see NIST SP 800-34 Rev. 1.

Replicating the VPN Concentrator supports secure remote connectivity into the recovery environment. For the stated route-flapping condition, the selected corrective action should enable route-flap damping on the VPN concentrator's relevant routing path. Route-flap damping suppresses unstable routes for a period rather than repeatedly propagating route withdrawals and announcements, which stabilizes connectivity during frequent route changes. The behavior and operational considerations for route-flap damping are defined in RFC 2439. Together, these placements and the route-stability action provide a usable and resilient Site B recovery capability.

QUESTION NO: 70

A DNS forward lookup zone named complia.org must:

• Ensure the DNS is protected from on-path attacks.

• Ensure zone transfers use mutual authentication and are authenticated and negotiated.

Which of the following should the security architect configure to meet these requirements? (Select two).

A. Public keys

B. Conditional forwarders

C. Root hints

D. DNSSEC

E. CNAME records

F. SRV records

ANSWER: A D

Explanation:

DNSSEC protects DNS data against on-path manipulation by digitally signing zone data. A validating resolver can use the zone's published DNSKEY records and signature records to verify that a response is authentic and has not been altered in transit. This establishes cryptographic origin authentication and data integrity for signed DNS records, directly addressing attacks in which an intermediary attempts to forge or modify DNS responses. DNSSEC's security model and signed-record validation process are defined by the IETF in RFC 4033.

Public keys support the mutual authentication requirement for protected zone transfers when used through certificate-based authenticated TLS. Each DNS server can validate the peer's certificate and corresponding public key before transferring zone information, so both the primary and secondary authenticate one another as part of establishing the encrypted session. Authenticated TLS also provides integrity protection for the transfer and negotiates the cryptographic session parameters. DNS zone-transfer support over TLS is specified in RFC 9103. Together, DNSSEC protects the authenticity and integrity of DNS answers, while public-key-based mutual authentication protects the server-to-server zone-transfer channel.

QUESTION NO: 71

A software development company makes Its software version available to customers from a web portal. On several occasions, hackers were able to access the software repository to change the package that is automatically published on the website. Which of the following would be the BEST technique to ensure the software the users download is the official software released by the company?

A. Distribute the software via a third-party repository.

B. Close the web repository and deliver the software via email.

C. Email the software link to all customers.

D. Display the SHA checksum on the website.

E. Digitally sign software packages and have users verify the signature using the company's public key.

ANSWER: E

Explanation:

Digitally sign software packages and have users verify the signature using the company's public key is the best technique because a digital signature provides cryptographic assurance of both publisher identity and package integrity. Before release, the company signs the finalized package with a private signing key that is protected separately from the web portal and repository. A customer's operating system, package manager, or verification tool can then validate the signature with the corresponding trusted public key or certificate. If an attacker replaces the published package, the altered file will not validate against the legitimate signature, alerting users that it was not the official release.

This control remains useful even when the distribution repository is compromised because trust is anchored in the signing key and certificate-validation process rather than in the repository's current contents. Effective implementation includes protecting the private key in secure key storage, restricting signing authority, applying timestamping where appropriate, and revoking or replacing certificates if a signing key is suspected of compromise. NIST describes digital signatures as mechanisms that detect unauthorized modification and authenticate the signatory; see NIST FIPS 186-5. Microsoft also describes how Authenticode code signing establishes software publisher identity and verifies integrity: Microsoft Authenticode documentation.

QUESTION NO: 72

A security analyst detected a malicious PowerShell attack on a single server. The malware used the Invoke-Expression function to execute an external malicious script. The security analyst scanned the disk with an antivirus application and did not find any IOCs. The security analyst now needs to deploy a protection solution against this type of malware.

Which of the following BEST describes the type of malware the solution should protect against?

A. Worm

B. Logic bomb

C. Fileless

D. Rootkit

ANSWER: C

Explanation:

Fileless is the best description because PowerShell-based attacks commonly execute malicious code directly in memory rather than installing a conventional malicious executable or leaving persistent artifacts on disk. `Invoke-Expression` can evaluate and run a downloaded or otherwise supplied PowerShell command string, enabling an attacker to use the legitimate PowerShell interpreter as a living-off-the-land technique. This helps explain why an antivirus disk scan may find no indicators of compromise: the malicious payload may have been memory-resident, fetched remotely, encoded, or executed only transiently.

Protection against this threat class should therefore include controls with behavioral and memory-level visibility, such as endpoint detection and response, antimalware scanning interfaces for PowerShell content, PowerShell script block logging,

and constrained application-control policies. Analysts should also collect and inspect PowerShell operational logs to identify suspicious command execution, including encoded commands, remote content retrieval, and use of dynamic execution functions. Microsoft documents that script block logging records PowerShell commands and script blocks for investigative use, while its antimalware integration scans PowerShell script content before execution. See Microsoft PowerShell logging documentation and Microsoft PowerShell antimalware documentation.