PassQueen
a.k.a dumpsqueen.com

CompTIA SecurityX Certification Exam

CompTIA CAS-005

Version Demo

Total Demo Questions: 67

Total Premium Questions: 674

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Governance, Risk, and Compliance	82
Topic 2, Security Architecture	197
Topic 3, Security Engineering	140
Topic 4, Security Operations	254
Topic 5, Mix Questions	1
Total	674

QUESTION NO: 1

Audit findings indicate several user endpoints are not utilizing full disk encryption During me remediation process, a compliance analyst reviews the testing details for the endpoints and notes the endpoint device configuration does not support full disk encryption Which of the following is the most likely reason me device must be replaced '

A. The HSM is outdated and no longer supported by the manufacturer

The HSM is outdated: While HSM (Hardware Security Module) is important for security, it is not typically used for full disk encryption.

B. The vTPM was not properly initialized and is corrupt.

The vTPM was not properly initialized: vTPM (virtual TPM) is less common and not typically a reason for requiring hardware replacement.

C. The HSM is vulnerable to common exploits and a firmware upgrade is needed

The HSM is vulnerable to common exploits: This would require a firmware upgrade, not replacement of the device.

D. The motherboard was not configured with a TPM from the OEM supplier.

E. The HSM does not support sealing storage

The HSM does not support sealing storage: Sealing storage is relevant but not the primary reason for requiring TPM for full disk encryption.

ANSWER: D

Explanation:

The motherboard was not configured with a TPM from the OEM supplier. A Trusted Platform Module is a hardware-based root of trust that can securely protect the keys used by endpoint full-disk encryption and validate the device's measured boot state before releasing those keys. In an enterprise deployment, requiring TPM-backed encryption supports stronger assurance that encryption secrets are tied to the physical endpoint and are released only when expected platform-integrity conditions are met.

When a device was supplied without a compatible TPM, the deficiency is inherent to the endpoint hardware platform rather than an operating-system configuration issue. Although some encryption products can operate with alternate startup authenticators, an organization's compliance baseline may specifically require TPM-backed full-disk encryption. If the motherboard cannot provide the required TPM capability, replacing the endpoint is an appropriate remediation to bring it into compliance. Microsoft identifies a compatible TPM as the preferred hardware configuration for BitLocker and explains that TPM technology provides hardware-based security functions used to protect cryptographic material and platform integrity. See Microsoft BitLocker documentation and Microsoft's TPM overview.

QUESTION NO: 2

A developer wants to maintain integrity to each module of a program and ensure the code cannot be altered by malicious users.

Which of the following would be BEST for the developer to perform? (Choose two.)

A. Utilize code signing by a trusted third party.

B. Implement certificate-based authentication.

C. Verify MD5 hashes.

D. Compress the program with a password.

E. Encrypt with 3DES.

F. Make the DACL read-only.

ANSWER: A F

Explanation:

Utilize code signing by a trusted third party and Make the DACL read-only provide complementary protections for software-module integrity. Code signing applies a digital signature to each release artifact, such as an executable, library, driver, or script. The signature binds the signed content to the publisher's certificate and enables a verifier to detect any modification made after signing. A trusted certificate chain also lets systems and users establish that the signer is an approved software publisher. In practice, build and release pipelines should sign finalized modules, protect signing keys, and enforce signature validation before deployment or execution. Microsoft describes code signing as a mechanism that validates both the publisher identity and whether software has changed since it was signed: Microsoft Code Signing.

Make the DACL read-only supports prevention by restricting permissions on deployed program modules. A properly configured discretionary access control list grants ordinary users and service accounts only the read and execute access needed to run the code, while limiting write, modify, and full-control access to tightly controlled administrative or deployment identities. This reduces the opportunity for a malicious or compromised nonprivileged account to overwrite a module. DACLs and their access control entries are the Windows mechanism used to make these authorization decisions: Microsoft DACLs and ACEs.

QUESTION NO: 3

An organization is implementing Kubernetes to host internet-facing microservices. The security architect wants to prevent unapproved or altered container images from being deployed to production clusters.

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

A. Sign container images.

B. Enable Kubernetes audit logging.

C. Enforce signature validation through an admission controller.

D. Deploy a network load balancer.

E. Use a shared cluster administrator account.

F. Disable image registry authentication.

ANSWER: A C

Explanation:

Sign container images and **enforce signature validation through an admission controller** provide complementary supply-chain protections. Image signing establishes cryptographic evidence that an image was produced or approved by a trusted publisher and that the signed artifact has not been modified after signing. The organization should protect the signing keys and define trusted publishers and repositories.

An admission controller enforces the policy at deployment time. It can reject workload manifests that reference unsigned images, images with invalid signatures, or images from untrusted registries before the workload is scheduled. This prevents a developer or attacker from bypassing the intended control by directly submitting an unapproved image to the cluster. Kubernetes documents admission control in its Admission Controllers guidance. Sigstore provides container-signing guidance through cosign documentation.

QUESTION NO: 4

Company A is establishing a contractual with Company B. The terms of the agreement are formalized in a document covering the payment terms, limitation of liability, and intellectual property rights. Which of the following documents will MOST likely contain these elements

A. Company A-B SLA v2.docx

B. Company A OLA v1b.docx

C. Company A MSA v3.docx

D. Company A MOU v1.docx

E. Company A-B NDA v03.docx

ANSWER: C

Explanation:

Company A MSA v3.docx is correct because a Master Services Agreement establishes the overarching legal, commercial, and risk-allocation terms that govern an ongoing relationship between two organizations. An MSA normally defines the foundational provisions the parties do not want to negotiate repeatedly for each future engagement, including payment and invoicing requirements, ownership and licensing of intellectual property, confidentiality, warranties, indemnification, dispute resolution, and limitations of liability. These clauses determine each party's contractual responsibilities and how financial and legal risk is allocated if a dispute, loss, or claim arises.

After the master agreement is executed, specific work can be authorized through statements of work, order forms, or similar documents that identify the scope, deliverables, schedule, and project-specific pricing while incorporating the MSA's baseline terms. This structure supports consistent vendor governance and makes contractual obligations enforceable across multiple engagements. The International Association for Contract & Commercial Management describes master agreements as frameworks for recurring transactions at WorldCC. A practical overview of the provisions commonly included in this type of agreement is also available from DocuSign.

QUESTION NO: 5

A company lined an email service provider called my-email.com to deliver company emails. The company stalled having several issues during the migration. A security engineer is troubleshooting and observes the following configuration snippet:

Which of the following should the security engineer modify to fix the issue? (Select two).

A. The email CNAME record must be changed to a type A record pointing to 192.168.111

B. The TXT record must be changed to "v=spf1 ip4:192.168.1.10 include:my-email.com -all"

C. The srvo1 A record must be changed to a type CNAME record pointing to the email server

D. The email CNAME record must be changed to a type A record pointing to 192.168.1.10

E. The TXT record must be changed to " v=dkim ip4:l92.168.1.11 include my-email.com -ell "

F. The TXT record must be Changed to " v=dkim ip4:192.168.1.10 include:email-all "

G. The srv01 A record must be changed to a type CNAME record pointing to the web01 server

ANSWER: B D

Explanation:

The email CNAME record must be changed to a type A record pointing to 192.168.1.10, because a CNAME record is an alias for another DNS hostname, not an IP address. When the intended destination is the mail system's IPv4 address, an A record is the appropriate DNS record type and provides a direct hostname-to-address mapping. This allows clients and related mail-service records to resolve the intended server address correctly.

The TXT record must be changed to "v=spf1 ip4:192.168.1.10 include:my-email.com -all", because this is an SPF policy, not a DMARC policy. SPF is published as a DNS TXT record and identifies which IP addresses and included domains are authorized to send mail for the organization's domain. The `ip4` mechanism authorizes the specified sending address, `include:my-email.com` evaluates the provider's SPF policy, and `-all` directs receivers to fail sources that are not authorized. Correct DNS addressing and a valid SPF authorization record are both important for successful mail delivery and sender authentication during a provider migration. See RFC 1035 for DNS record behavior and RFC 7208 for SPF syntax and processing.

QUESTION NO: 6

An endpoint security engineer finds that a newly acquired company has a variety of non-standard applications running and no defined ownership for those applications. The engineer needs to find a solution thatrestricts malicious programs and software from running in that environment, while allowing the non-standard applications to function without interruption. Which of the following application control configurations should the engineer apply?

A. Deny list

B. Allow list

C. Audit mode

D. MAC list

ANSWER: A

Explanation:

Deny list is the appropriate application-control configuration because it prevents execution of software that has been identified as malicious, unauthorized, or otherwise unacceptable, while permitting applications that have not been explicitly blocked to continue operating. This is particularly suitable during an acquisition, where the environment contains numerous non-standard applications and there is no reliable application owner or complete inventory available. The organization can immediately block known harmful executables using indicators such as file hashes, publisher certificates, file paths, application reputation, and threat-intelligence data, without first having to identify, validate, and approve every inherited application.

This approach directly balances the two stated requirements: reducing exposure to malicious software while avoiding interruption to legitimate but unfamiliar business applications. A deny-list policy can be deployed as an initial protective control and refined over time as the acquired organization's software estate is discovered, assigned ownership, and brought under governance. Microsoft's application-control guidance describes using policies to control executable code and prevent unwanted or untrusted software from running. NIST SP 800-167 also provides guidance on application whitelisting and related application-control concepts for endpoint protection. See Microsoft App Control for Business and NIST SP 800-167.

QUESTION NO: 7

A company wants to implement hardware security key authentication for accessing sensitive information systems The goal is to prevent unauthorized users from gaining access with a stolen password Which of the following models should the company implement to b«st solve this issue?

A. Rule based

B. Time-based

C. Role based

D. Context-based

ANSWER: B

Explanation:

Time-based is correct because a hardware token can provide a time-based one-time password (TOTP) as an additional possession factor during authentication. The user must supply both the memorized password and a current code generated by the physical device. A stolen password alone is therefore insufficient to access the sensitive system, because the attacker would also need control of the registered hardware token and its current, short-lived code.

TOTP derives authentication codes from a shared secret and the current time interval. The authentication service independently calculates the expected value for that interval and accepts the code only within a narrow validity window. This limits the usefulness of a captured code and prevents reuse after it expires. Hardware tokens are appropriate for this use case because the authentication secret can remain protected in the device rather than being stored as a reusable value on

the endpoint. Using this model implements multifactor authentication by combining something the user knows with something the user possesses, directly addressing password compromise as the stated threat.

The technical specification for TOTP is available in RFC 6238. NIST guidance on OTP authenticators and multifactor authentication is available in NIST SP 800-63B.

QUESTION NO: 8

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:

Embedded cryptoprocessor and **Hardware-backed public key storage** are the features that make a microSD hardware security module suitable for protecting remote access from corporate mobile devices. An embedded cryptoprocessor performs sensitive cryptographic functions within dedicated hardware. This enables operations such as private-key signing, certificate-based authentication, and key establishment to occur in an isolated security boundary rather than exposing private keys to the mobile operating system or applications.

Hardware-backed public key storage is equally important because remote corporate access commonly relies on PKI identities, including client certificates and their associated private keys. A hardware-backed module can generate or import keys and protect them against straightforward extraction, while permitting approved cryptographic operations using those keys. Requiring the microSD HSM to be present also provides a possession factor: the user must have the authorized hardware token installed before the protected credential can be used for remote authentication. This combination improves assurance for enterprise VPN, zero-trust access, and certificate-based authentication workflows, particularly where the organization needs stronger protection than software-only credential storage can provide. NIST defines an HSM as a device that safeguards and manages cryptographic keys and performs cryptographic functions; see the NIST HSM glossary entry. For a related hardware-backed key protection model, see Android Keystore security documentation.

QUESTION NO: 9

Which of the following are risks associated with vendor lock-in? (Select 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 creates dependency on a provider's proprietary services, data formats, tools, contractual terms, and operational processes. "The vendor can change product offerings." is a risk because a dependent client may have limited practical ability to respond if the provider discontinues a service, changes functionality, alters licensing or pricing, or modifies support terms. Migration to an alternative may require substantial time, cost, redesign, data conversion, and retraining, making the client especially exposed to those provider-driven changes.

"The client experiences decreased quality of service." is also a risk because lock-in weakens the client's negotiating leverage and ability to quickly move critical workloads when service performance, support responsiveness, availability, or other service-level outcomes decline. A viable exit strategy and portable architectures help reduce this exposure by preserving the ability to transition services or use alternatives. NIST identifies vendor lock-in as a cloud-computing concern tied to nonportable data and applications, while its cloud security guidance emphasizes evaluating provider dependencies and maintaining appropriate controls for service delivery and transition planning. See NIST SP 800-146, Cloud Computing Synopsis and Recommendations and NIST SP 800-144, Guidelines on Security and Privacy in Public Cloud Computing.

QUESTION NO: 10

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 constrains the number of authentication guesses that can be made against an account. The listed requirements establish a minimum password length, basic character composition, and password-reuse prevention, all of which improve password quality. However, none of them limits repeated failed sign-in attempts against a valid username. An attacker conducting an online brute-force or password-spraying attack can continue submitting guesses unless the authentication service applies a response control.

An account lockout policy defines a failed-attempt threshold and locks the account when that threshold is reached, either for a configured duration or until an authorized administrator unlocks it. This substantially reduces the practical number of guesses available to an attacker, increases the time and effort needed to attack the account, and creates security events that can be monitored and correlated for incident response. The threshold and lockout duration should be tuned to the organization's risk tolerance and operational needs, with monitoring to detect malicious lockout attempts. Microsoft defines the account lockout threshold as the number of failed sign-in attempts that causes an account to be locked: Microsoft Learn. OWASP also recommends login throttling and account lockout protections against automated credential attacks: OWASP Authentication Cheat Sheet.

QUESTION NO: 11 - (SIMULATION)

SIMULATION

A security engineer needs to review the configurations of several devices on the network to meet the following requirements:

• The PostgreSQL server must only allow connectivity in the 10.1.2.0/24

subnet.

• The SSH daemon on the database server must be configured to listen

to port 4022.

• The SSH daemon must only accept connections from a Single

workstation.

• All host-based firewalls must be disabled on all workstations.

• All devices must have the latest updates from within the past eight

days.

•All HDDs must be configured to secure data at rest.

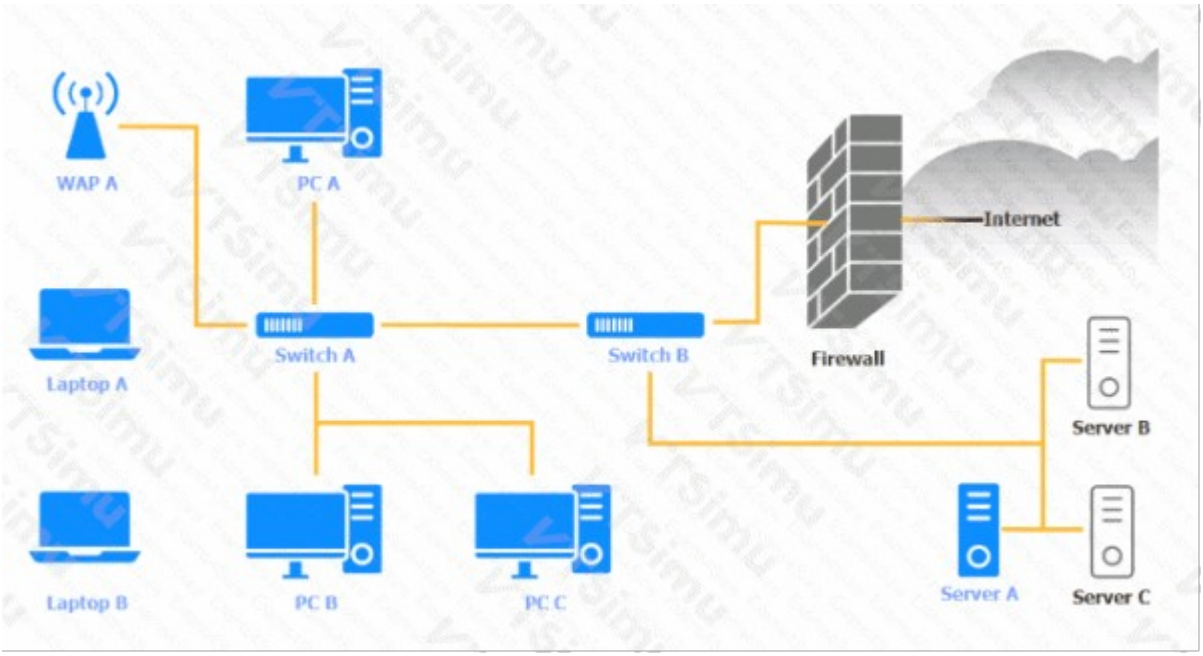
• Cleartext services are not allowed.

• All devices must be hardened when possible.

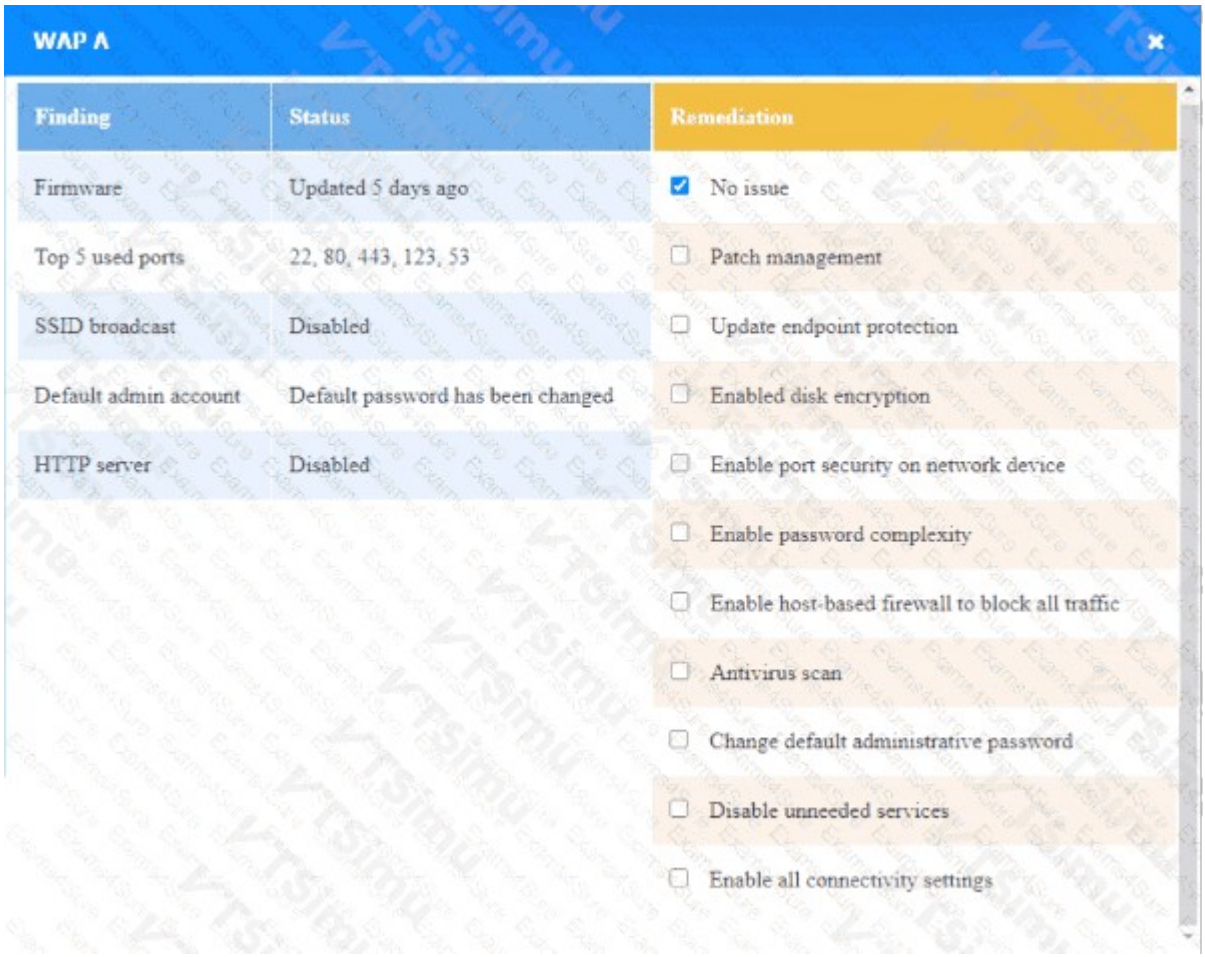
Instructions:

Click on the various workstations and network devices to review the posture assessment results. Remediate any possible issues or indicate that no issue is found.

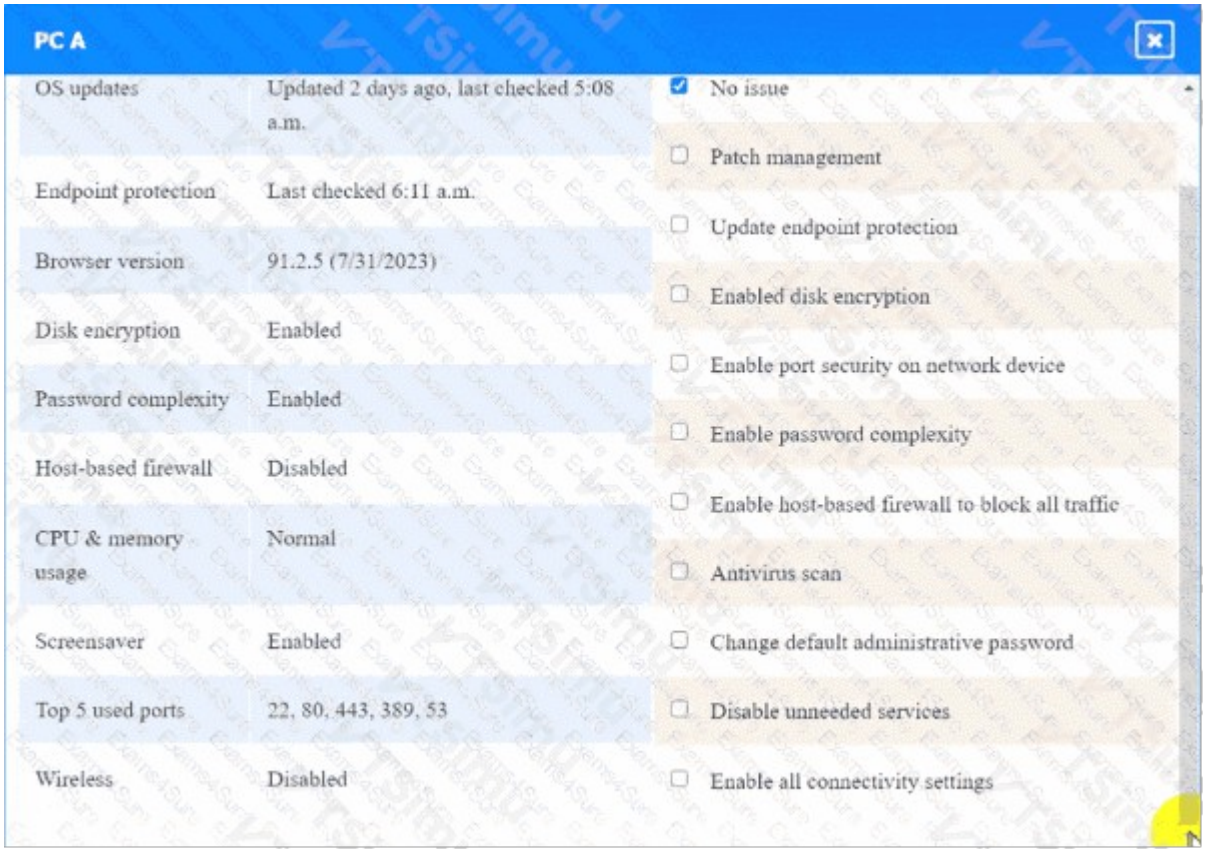
Click on Server A to review output data. Select commands in the appropriate tab to remediate connectivity problems to the pOSTGREsql DATABASE VIA ssh



WAP A


WAP A
Finding
Status
Remediation
Firmware
Updated 5 days ago
Top 5 used ports
22, 80, 443, 123, 53
SSID broadcast
Disabled
Default admin account
Default password has been changed
HTTP server
Disabled
No issue
Patch management
Update endpoint protection
Enabled disk encryption
Enable port security on network device
Enable password complexity
Enable host-based firewall to block all traffic
Antivirus scan
Change default administrative password
Disable unneeded services
Enable all connectivity settings

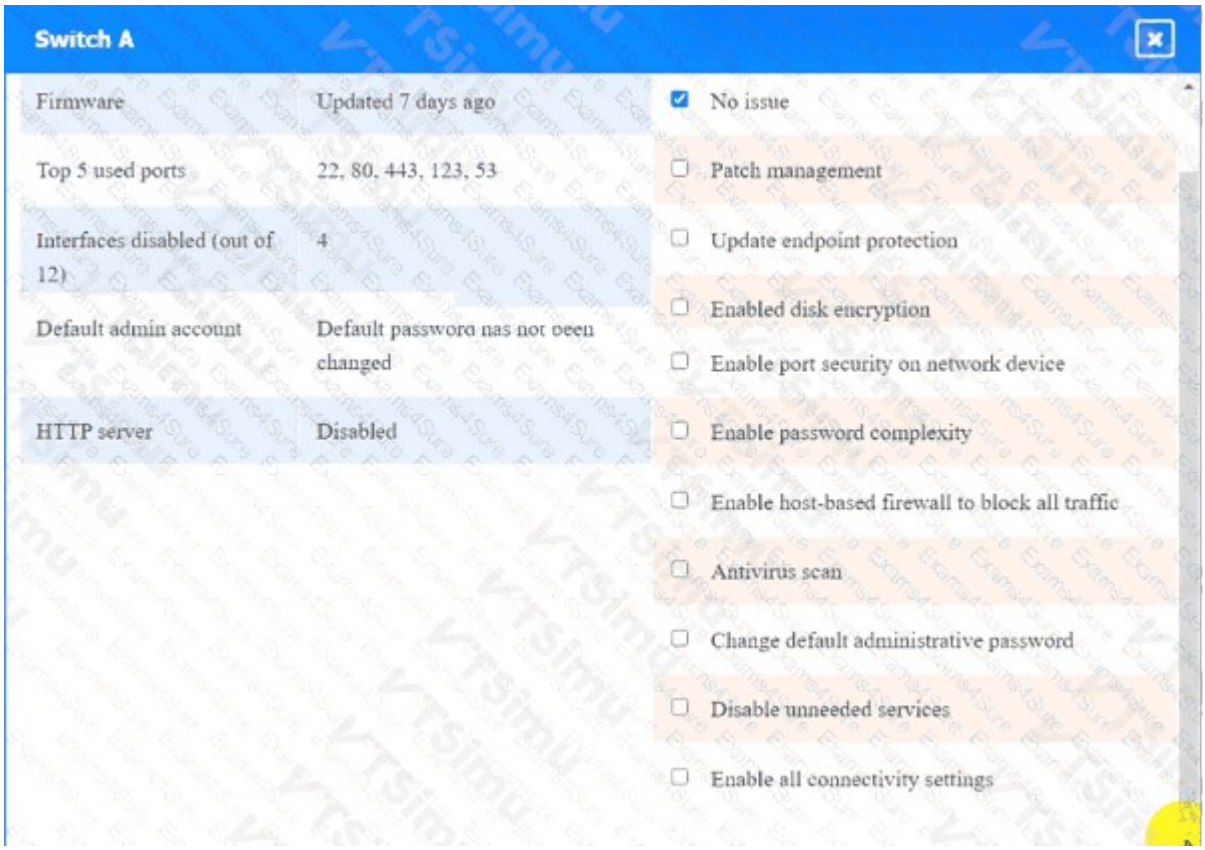
PC A


PC A
OS updates
Updated 2 days ago, last checked 5:08 a.m.
Endpoint protection
Last checked 6:11 a.m.
Browser version
91.2.5 (7/31/2023)
Disk encryption
Enabled
Password complexity
Enabled
Host-based firewall
Disabled
CPU & memory usage
Normal
Screensaver
Enabled
Top 5 used ports
22, 80, 443, 389, 53
Wireless
Disabled
No issue
Patch management
Update endpoint protection
Enabled disk encryption
Enable port security on network device
Enable password complexity
Enable host-based firewall to block all traffic
Antivirus scan
Change default administrative password
Disable unneeded services
Enable all connectivity settings

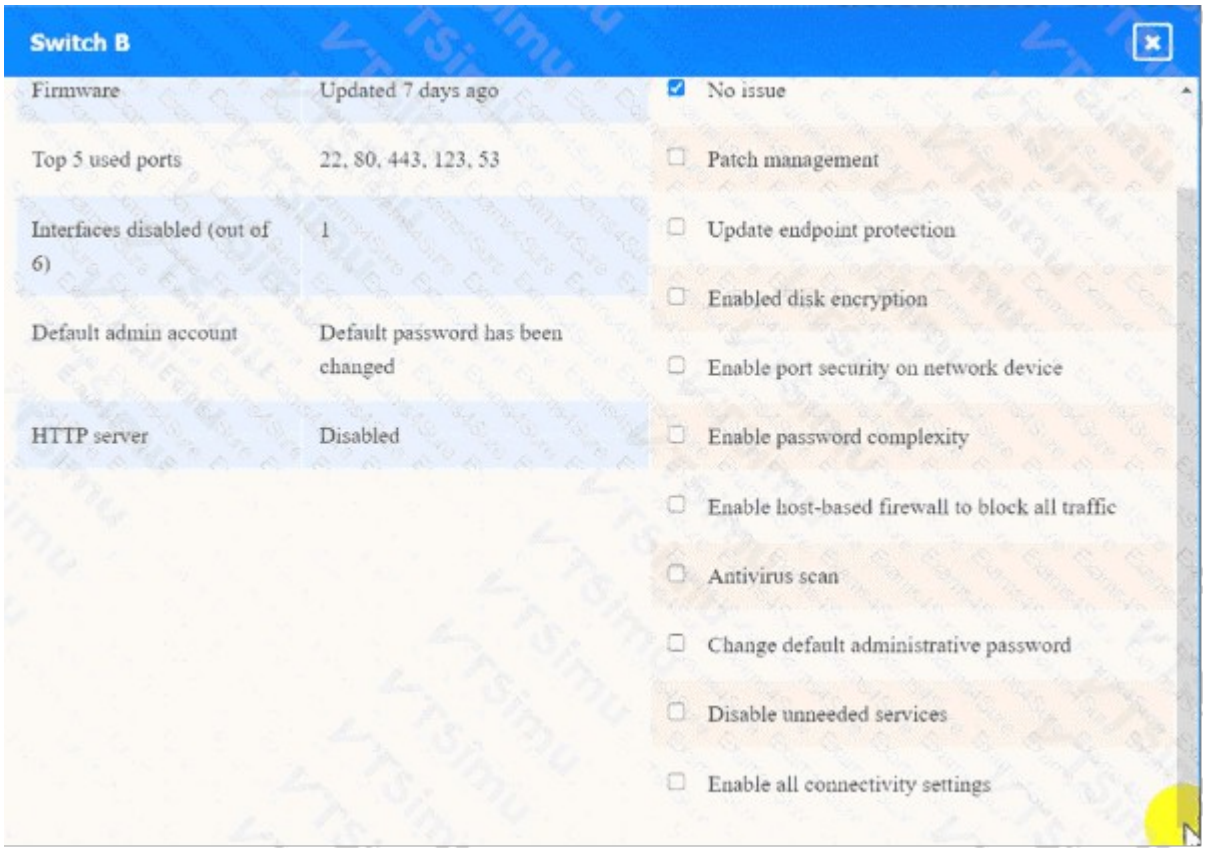
Laptop A



Switch A



Switch B:


Switch B
Firmware
Updated 7 days ago
Top 5 used ports
22, 80, 443, 123, 53
Interfaces disabled (out of 6)
1
Default admin account
Default password has been changed
HTTP server
Disabled
No issue
Patch management
Update endpoint protection
Enabled disk encryption
Enable port security on network device
Enable password complexity
Enable host-based firewall to block all traffic
Antivirus scan
Change default administrative password
Disable unneeded services
Enable all connectivity settings

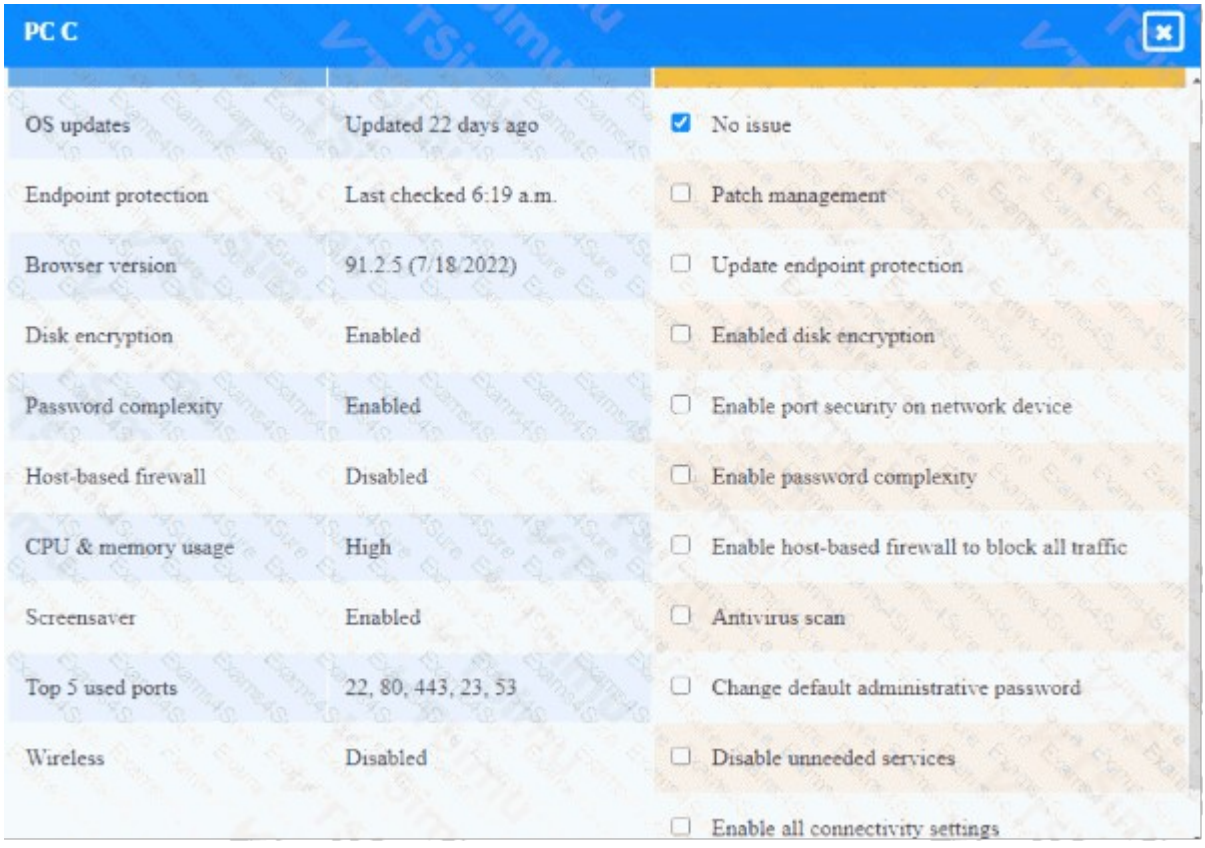
Laptop B


Laptop B
OS updates
Updated 3 days ago, last checked 8:08 a.m.
Endpoint protection
Last checked in 8:11 a.m.
Browser version
81.2.5 (7/31/2023)
Disk encryption
Disabled
Password Complexity
Enabled
Host-based firewall
Disabled
CPU & memory usage
Normal
Screensaver
Enabled
Top 5 used ports
22, 80, 443, 8080, 53
Wireless
Enabled
No issue
Patch management
Update endpoint protection
Enabled disk encryption
Enable port security on network device
Enable password complexity
Enable host-based firewall to block all traffic
Antivirus scan
Change default administrative password
Disable unneeded services
Enable all connectivity settings

PC B


PC B
OS updates
Updated 2 days ago, last checked 5:10 a.m.
Endpoint protection
Last checked in 6:13 a.m.
Browser version
91.2.5 (7/31/2023)
Disk encryption
Enabled
Password complexity
Enabled
Host-based firewall
Disabled
CPU & memory usage
Medium
Screensaver
Enabled
Top 5 used ports
22, 80, 443, 389, 53
Wireless
Disabled
No issue
Patch management
Update endpoint protection
Enabled disk encryption
Enable port security on network device
Enable password complexity
Enable host-based firewall to block all traffic
Antivirus scan
Change default administrative password
Disable unneeded services
Enable all connectivity settings

PC C


PC C
OS updates
Updated 22 days ago
Endpoint protection
Last checked 6:19 a.m.
Browser version
91.2.5 (7/18/2022)
Disk encryption
Enabled
Password complexity
Enabled
Host-based firewall
Disabled
CPU & memory usage
High
Screensaver
Enabled
Top 5 used ports
22, 80, 443, 23, 53
Wireless
Disabled
No issue
Patch management
Update endpoint protection
Enabled disk encryption
Enable port security on network device
Enable password complexity
Enable host-based firewall to block all traffic
Antivirus scan
Change default administrative password
Disable unneeded services
Enable all connectivity settings

Server A


Server A
Nmap
IP Tables
Nmap scan report for psql-srvr.acme.com
Host is up, received arp-response (0.00040s latency).
...
PORT STATE SERVICE VERSION
22/tcp open ssh OpenSSH 8.4
80/tcp closed http
443/tcp closed ssl/http
1433/tcp closed mssql
5432/tcp closed postgresql
...
1 2 3 4
iptables -R INPUT 1 -p tcp -s 10.1.2.25/32 --sport 4022 -j ACCEPT
iptables -D OUTPUT 1
iptables -A OUTPUT -p udp -d 0/0 -s 10.1.2.0/24 --sport 5432 -m state --state ESTABLISHED -j ACCEPT
iptables -A INPUT -p tcp -d 0/0 -s 10.1.2.0/24 --dport 5432 -m state --state NEW,ESTABLISHED -j ACCEPT
1 2 3 4
iptables -R INPUT 1 -p tcp -s 10.1.2.0/24 --dport 4022 -j ACCEPT
iptables -D OUTPUT 2
iptables -A OUTPUT -p tcp -d 0/0 -s 10.1.2.0/24 --sport 5432 -m state --state ESTABLISHED -j ACCEPT
iptables -A INPUT -p tcp -d 0/0 -s 10.1.2.0/24 --dport 5432 -m state --state NEW,ESTABLISHED -j ACCEPT
1 2 3 4
iptables -R OUTPUT 1 -p tcp -s 10.1.2.25/32 --sport 4022 -j ACCEPT
iptables -F OUTPUT
iptables -A OUTPUT -p tcp -d 0/0 -s 10.1.2.0/24 --sport 5432 -m state --state ESTABLISHED -j ACCEPT
iptables -A INPUT -p tcp -d 0/0 -s 10.1.2.0/24 --dport 5432 -m state --state NEW,ESTABLISHED -j ACCEPT
1 2 3 4
iptables -R INPUT 1 -p tcp -s 10.1.2.25/32 --dport 4022 -j ACCEPT
iptables -D OUTPUT 1
iptables -A OUTPUT -p tcp -d 0/0 -s 10.1.2.0/24 --sport 5432 -m state --state ESTABLISHED -j ACCEPT
iptables -A INPUT -p tcp -d 0/0 -s 10.1.2.0/24 --dport 5432 -m state --state NEW,ESTABLISHED -j ACCEPT
Nmap
IP Tables
#iptables --list --verbose
Chain INPUT (policy DROP 5 packets, 341 bytes)
pkts bytes target prot opt in out source destination
0 0 ACCEPT tcp -- any any anywhere anywhere tcp spts:login:65535 dpt:ssh state NEW,ESTABLISHED
1 28 DROP all -- any any anywhere anywhere
Chain FORWARD (policy DROP 0 packets, 0 bytes)

ANSWER: See the explanation for the answer

Explanation:

Select the following posture remediations:

WAP A: *No issue found*. Firmware is current, SSID broadcast and HTTP are disabled, and the default password was changed.
PC A: *No issue found*. Its host-based firewall is disabled as explicitly required; updates and disk encryption are current/enabled.
Laptop A: *No issue found*. It is updated within eight days, encrypted, and its host firewall is disabled as required.
Switch A: Select *Change default administrative password* and *Enable port security on network device*. The default administrative password has not been changed and no interfaces are disabled, so switch hardening is needed.
Switch B: *No issue found*. Firmware is current, HTTP is disabled, the default password is changed, and an unused interface is disabled.
Laptop B: Select *Enable disk encryption* and *Disable unneeded services*. Disk encryption is disabled, and TCP 8080 is an unnecessary cleartext web-service exposure.
PC B: *No issue found*. It is updated, encrypted, has password complexity enabled, and the host firewall is disabled as required.
PC C: Select *Patch management* and *Disable unneeded services*. Its OS was updated 22 days ago, exceeding the eight-day requirement, and TCP 23 is Telnet, a prohibited cleartext service.

Server A — IP Tables tab: Choose **option 3** (the command set that permits SSH on TCP 4022 only from `10.1.2.25/32` and permits PostgreSQL TCP 5432 only for `10.1.2.0/24`).

```
iptables -R INPUT 1 -p tcp -s 10.1.2.25/32 --dport 4022 -j ACCEPT
iptables -D OUTPUT 1
iptables -A OUTPUT -p tcp -d 0/0 -s 10.1.2.0/24 --sport 5432 -m state --state ESTABLISHED -
j ACCEPT
iptables -A INPUT -p tcp -d 0/0 -s 10.1.2.0/24 --dport 5432 -m state --state
NEW,ESTABLISHED -j ACCEPT
```

This selection changes the existing broadly available SSH rule to port `4022` and limits it to the single authorized workstation, while allowing PostgreSQL connections and return traffic only for the `10.1.2.0/24` subnet. Do not choose the options using UDP for PostgreSQL or allowing the entire subnet to SSH.

QUESTION NO: 12

A company implemented a NIDS and a NIPS on the most critical environments. Since this implementation, the company has been experiencing network connectivity issues. Which of the following should the security architect recommend for a new NIDS/NIPS implementation?

A. Implementing the NIDS with a port mirror in the core switch and the NIPS in the main firewall

B. Implementing the NIDS and the NIPS together with the main firewall

C. Implementing a NIDS without a NIPS to increase the detection capability

D. Implementing the NIDS in the bastion host and the NIPS in the branch network router

ANSWER: A

Explanation:

Implementing the NIDS with a port mirror in the core switch and the NIPS in the main firewall is the appropriate design because it separates passive monitoring from active, inline prevention. A network intrusion detection system is normally deployed out of band, receiving a copy of traffic from a network TAP or a switch port-mirroring/SPAN session. It can analyze traffic broadly, including traffic crossing the core, without becoming part of the production forwarding path. Therefore, a sensor outage, overload, or analysis delay will not itself interrupt network connectivity. Cisco documents SPAN as a mechanism that copies traffic from selected ports or VLANs to a destination port for monitoring and analysis: Cisco SPAN configuration guidance.

A network intrusion prevention system must be inline to block malicious packets or sessions. Positioning it at the main firewall creates a managed enforcement point where traffic policies, inspection capacity, high availability, bypass/fail-open behavior, and alerting can be centrally designed and tested. This placement preserves the prevention function while avoiding unnecessary inline deployment throughout critical internal network paths, which can contribute to the connectivity issues described. NIST's guidance on intrusion detection and prevention technologies discusses the operational impact of inline IPS deployment and the importance of appropriate sensor placement: NIST SP 800-94.

QUESTION NO: 13

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 remediate each of the material conditions identified during the assessment. PAM establishes governance over elevated accounts through controls such as least privilege, credential vaulting, just-in-time access, approval workflows, and privileged-session auditing. These controls substantially reduce the opportunity for compromised administrator credentials to be used broadly or persistently.

Remote access VPN replaces direct exposure of Remote Desktop Protocol with a controlled, authenticated remote-entry path. The organization can require authorized users to establish the VPN connection before they can reach internal resources, then apply access policies and logging appropriate to systems holding protected health information. Network segmentation divides the currently flat environment into security zones and enforces policy-controlled traffic between them. Isolating PHI servers and administrative management networks limits lateral movement and reduces ransomware's potential blast radius.

Combined, these measures provide controlled privileged administration, protected remote connectivity, and containment boundaries. This approach is consistent with CISA's Ransomware Guide, which recommends restricting remote services and limiting lateral movement, and with the least-privilege access-control principles in NIST SP 800-53 Rev. 5.

QUESTION NO: 14

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 WAF.

B. A certificate on the WAF is expired.

C. HTTP traffic is not forwarding to HTTPS to decrypt.

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

ANSWER: C

Explanation:

HTTP traffic is not forwarding to HTTPS to decrypt is the most likely cause because the WAF must be deployed in the request path and able to inspect the relevant application traffic in order to generate meaningful security events. In a common protected-web architecture, inbound HTTP requests are redirected to HTTPS, and the WAF or its associated reverse proxy terminates TLS so it can evaluate the request URI, headers, cookies, parameters, and body before passing the request to the origin server. If a portion of traffic reaches the exposed web server through an HTTP path that is not forwarded into the HTTPS inspection flow, that traffic can bypass the WAF policy enforcement and its associated logging. The SOC analyst could then identify malicious activity in web-server logs without finding a corresponding WAF event. Ensuring that all public entry points route through the WAF, that TLS termination and inspection are configured as required, and that direct access to the origin is restricted provides complete coverage and consistent logging. AWS describes that AWS WAF evaluates web requests associated with protected resources in How AWS WAF works. OWASP also explains the role of a WAF in inspecting and filtering HTTP traffic in its Web Application Firewall guidance.

QUESTION NO: 15

A company has moved its sensitive workloads lo the cloud and needs to ensure high availability and resiliency of its web-based application. The cloud architecture team was given the following requirements

• The application must run at 70% capacity at all times

• The application must sustain DoS and DDoS attacks.

• Services must recover automatically.

Which of the following should the cloud architecture team implement? (Select THREE).

A. Read-only replicas

B. BCP

C. Autoscaling

D. WAF

E. CDN

F. Encryption

G. Continuous snapshots

H. Containerization

ANSWER: C D E

Explanation:

Autoscaling, WAF, and CDN together directly support the required availability, resiliency, and attack-resistance goals for a cloud-hosted web application. Autoscaling can use target-tracking policies based on a utilization metric, such as CPU utilization or request count per target, to add or remove healthy application instances while maintaining a target value of 70%. When paired with load-balancer health checks, an autoscaling group can remove unhealthy instances and launch replacements automatically, restoring service capacity without requiring manual intervention. AWS describes this behavior in its target-tracking scaling policy documentation.

WAF provides inspection and filtering of HTTP/S requests before they reach the web application. It can enforce rate-based rules and block malicious web requests, helping reduce the effect of application-layer denial-of-service activity. A CDN adds globally distributed edge locations that cache cacheable content and absorb or distribute large volumes of inbound traffic before it reaches the origin. This reduces origin load, improves availability for legitimate users, and strengthens resilience

during volumetric DDoS events. AWS explains the edge-layer role of content delivery services in its DDoS resiliency guidance.

QUESTION NO: 16

Anorganization has noticed an increase in phishing campaigns utilizingtyposquatting. A security analyst needs to enrich the data for commonly used domains against the domains used in phishing campaigns. The analyst uses a log forwarder to forward network logs to the SIEM. Which of the following would allow the security analyst to perform this analysis?

A. Use acron jobto regularly update and compare domains.

B. Create aparserthat matches domains.

C. Develop aquerythat filters out all matching domain names.

D. Implement adashboardon the SIEM that shows the percentage of traffic by domain.

ANSWER: D

Explanation:

Implement a dashboard on the SIEM that shows the percentage of traffic by domain. A SIEM dashboard can aggregate the forwarded network-log data by domain name and present each domain's relative volume of observed traffic. This establishes a useful baseline for domains that are commonly accessed in the organization and makes unusual or newly observed domains easier to identify. The analyst can then compare domains associated with known phishing activity against this baseline, including visually similar typosquatted domains that may otherwise blend into a large volume of DNS, proxy, or web traffic. Dashboard panels can be backed by searches that normalize domain fields, group events by domain, calculate counts and percentages, and constrain results to relevant periods or data sources. This turns raw log events into actionable security visibility for triage and investigation. SIEM dashboards are specifically intended to visualize search and reporting results in a consolidated view, supporting monitoring and analysis of security-relevant trends. CompTIA's SecurityX certification also emphasizes using data analysis, monitoring, and security controls to investigate and respond to enterprise threats. See Splunk Dashboard Studio documentation and CompTIA SecurityX certification information.

QUESTION NO: 17

A security architect works for a manufacturing organization that has many different branch offices. The architect is looking for a way to reduce traffic and ensure the branch offices receive the latest copy of revoked certificates issued by the CA at the organization's headquarters location. The solution must also have the lowest power requirement on the CA.

Which of the following is the BEST solution?

A. Deploy an RA on each branch office.

B. Use Delta CRLs at the branches.

C. Configure clients to use OCSP.

D. Send the new CRLs by using GPO.

ANSWER: C

Explanation:

Configure clients to use OCSP is the best solution because the Online Certificate Status Protocol enables a client to request the revocation status of an individual certificate when validation is needed. Instead of distributing complete certificate revocation lists to every branch and having each client download and process those lists, an OCSP client receives a small signed status response indicating whether the specified certificate is good, revoked, or unknown. This substantially reduces WAN traffic in an organization with many geographically distributed offices.

For the lowest workload on the certification authority, the organization can deploy an OCSP responder, also called an Online Responder, rather than having the CA directly process status requests. The responder handles client queries and returns

cached or pre-signed revocation-status responses, while the CA continues its core responsibility of issuing certificates and publishing revocation information. This architecture provides timely certificate-status validation without requiring the CA to serve large revocation-list downloads to each branch. OCSP is specified by the IETF in RFC 6960. Microsoft also describes the Online Responder role as a scalable means for clients to determine certificate revocation status in an AD CS PKI: Online Responder overview.

QUESTION NO: 18

An organization uses infrastructure as code to provision cloud environments. A security review found that developers can commit cloud access keys to repositories and deploy public storage buckets without approval.

Which of the following controls should the organization implement to BEST address these findings? (Choose two.)

A. Secrets scanning in repositories and CI/CD pipelines.

B. Policy-as-code checks before deployment.

C. Manual approval of all DNS queries.

D. Network packet capture on developer workstations.

E. Static IP addressing for cloud resources.

F. Disable version control for infrastructure files.

ANSWER: A B

Explanation:

Secrets scanning in repositories and CI/CD pipelines and **policy-as-code checks before deployment** directly address the identified issues. Secrets scanning detects exposed access keys, tokens, and other credentials in source code, commit history, and build pipelines. When integrated into the CI/CD workflow, the control can block releases, alert security personnel, and trigger credential rotation when an exposed secret is identified.

Policy-as-code evaluates proposed infrastructure configurations against defined requirements before resources are created. The organization can use this control to prevent public storage buckets, overly permissive identity policies, unencrypted resources, or deployments that do not meet approved baseline requirements. These preventive checks shift security validation earlier in the delivery lifecycle. NIST's Secure Software Development Framework recommends protecting code and related artifacts, while OWASP provides guidance for secrets management.

QUESTION NO: 19

A security engineer needs to remediate a SWEET32 vulnerability in an OpenSSH-based application and review existing configurations. Which of the following should the security engineer do? (Select two.)

A. Disable Twofish algorithms

B. cat /etc/ashd/ash_config | grep "HMAC"

C. Disable RSA algorithms

D. cat /etc/sshd/ssh_config | grep "PermitRootLogin"

E. Disable 3DES algorithms

F. cat /etc/ssh/sshd_config | grep "Ciphers"

ANSWER: E F

Explanation:

SWEET32 is a practical collision attack against long-lived use of ciphers with 64-bit block sizes. In SSH deployments, the relevant legacy cipher is Triple DES (3DES), commonly represented in OpenSSH cipher policy as `3des-cbc`. Disabling 3DES algorithms removes the affected 64-bit-block cipher from the SSH service's permitted encryption algorithms, preventing clients from negotiating it for new sessions. This is the appropriate remediation step for the specific SWEET32 exposure.

Reviewing the `Ciphers` directive in the OpenSSH server configuration is also required to validate that 3DES is not explicitly enabled and that the remaining allowed cipher suite matches organizational cryptographic policy. On typical OpenSSH installations, the server configuration file is `/etc/ssh/sshd_config`; filtering it for `Ciphers` provides a quick review of configured cipher restrictions. OpenSSH documents the server configuration directives, including `Ciphers`, in its sshd_config manual. Background on the 64-bit block-cipher collision issue addressed by SWEET32 is available from the SWEET32 research site.

QUESTION NO: 20

Which of the following key management practices ensures that an encryption key is maintained within the organization?

A. Encrypting using a key stored in an on-premises hardware security module

B. Encrypting using server-side encryption capabilities provided by the cloud provider

C. Encrypting using encryption and key storage systems provided by the cloud provider

D. Encrypting using a key escrow process for storage of the encryption key

ANSWER: A

Explanation:

Encrypting using a key stored in an on-premises hardware security module is the appropriate practice because it keeps the encryption key under the organization's physical and administrative control. A hardware security module is a dedicated cryptographic device designed to generate, store, protect, and use cryptographic keys without exposing key material to general-purpose systems. When the HSM is deployed on premises and managed by the organization, the organization controls access policies, administrator roles, lifecycle activities such as rotation and destruction, and the facility in which the key is held. This model supports strong key custody and separation of duties while allowing cryptographic operations to occur in a protected boundary. NIST key-management guidance emphasizes protecting cryptographic keys throughout their lifecycle and establishing controls for key storage, access, and management. Hardware security modules may be validated against FIPS 140 requirements, which define security requirements for cryptographic modules and provide a recognized basis for evaluating the protected boundary used for key operations. See NIST SP 800-57 Part 1 Rev. 5 and NIST Cryptographic Module Validation Program.

QUESTION NO: 21

An engineer wants to automate several tasks by running commands daily on a UNIX server. The engineer has only built-in, default tools available. Which of the following should the engineer use to best assist with this effort? (Select Two).

A. Python

B. Cron

C. Ansible

D. PowerShell

E. Bash

F. Task Scheduler

ANSWER: B E

Explanation:

Cron and **Bash** are the appropriate built-in UNIX tools for this automation requirement. Cron is the standard time-based job scheduler on UNIX-like systems. An administrator can create a crontab entry that invokes a command or script at a specific time every day, making it suitable for recurring operational work such as backups, maintenance, reporting, and log-processing tasks. A daily schedule can be defined directly in a user or system crontab, with output redirected or logged as needed for operational visibility.

Bash provides the scripting environment to place several commands into one repeatable executable script. The script can sequence commands, use variables and conditional logic, check exit statuses, write logs, and handle simple errors before Cron invokes it on the defined daily schedule. This division of responsibilities is a conventional UNIX administration pattern: Bash implements the task workflow, while Cron supplies dependable scheduled execution. The crontab format and scheduling behavior are documented in the crontab manual page; shell scripting capabilities are documented in the GNU Bash Reference Manual.

QUESTION NO: 22

Emails that the marketing department is sending to customers are pomp to the customers ' spam folders. The security team is investigating the issue and discovers that the certificates used by the email server were reissued, but DNS records had not been updated. Which of the following should the security team update in order to fix this issue? (Select three.)

A. DMARC

B. SPF

C. DKIM

D. DNSSEC

E. SASC

F. SAN

G. SOA

ANSWER: A B C

Explanation:

DMARC, SPF, and DKIM are the DNS-based email-authentication controls that should be reviewed and updated so recipient mail systems can validate messages sent by the marketing department. SPF authorizes the infrastructure permitted to send mail for the organization’s domain. If the outbound service, sending addresses, or mail-routing configuration changed during the certificate reissue or associated server change, the published SPF policy must accurately authorize that sender. DKIM depends on a public key published in DNS under a selector. When signing keys or related mail-server configuration are replaced, the matching public key record must be available so receivers can verify the cryptographic signature on outgoing messages. DMARC publishes the domain’s policy for handling messages based on aligned SPF and/or DKIM authentication results, along with reporting destinations. Maintaining an appropriate DMARC record ensures receiving providers evaluate authenticated mail consistently and lets the organization obtain reports that identify authentication or alignment problems. Together, these records establish sender authorization, message-signature verification, and domain-level handling policy, which are core controls for improving legitimate-mail delivery and reducing spam-folder placement. See RFC 7208: Sender Policy Framework and RFC 7489: DMARC.

QUESTION NO: 23

A security engineer needs to remediate a SWEET32 vulnerability in an OpenSSH-based application and review existing configurations. Which of the following should the security engineer do? (Select two.)

A. Disable Twofish algorithms

B. cat /etc/ashd/ash_config | grep "HMAC"

C. Disable RSA algorithms

D. cat /etc/sshd/ssh_config | grep "PermitRootLogin"

E. Disable 3DES algorithms

F. cat /etc/ssh/sshd_config | grep "Ciphers"

ANSWER: E F

Explanation:

SWEET32 is a birthday-bound cryptographic attack against encryption algorithms that use 64-bit block sizes when enough encrypted traffic is processed under the same key. Triple DES is the principal relevant OpenSSH cipher because its 64-bit block size makes it susceptible to this class of attack. Therefore, disabling 3DES algorithms removes the vulnerable `3des-cbc` cipher from the SSH service's permitted cipher set. OpenSSH supports specifying permitted symmetric encryption algorithms through the `Ciphers` directive in the server configuration. Reviewing that directive in `/etc/ssh/sshd_config` allows the engineer to identify whether 3DES remains enabled and to verify that modern ciphers are configured instead. Typical current OpenSSH choices include AES-CTR, AES-GCM, and ChaCha20-Poly1305, which do not use a 64-bit block cipher construction. The effective configuration can also be validated after changes with `sshd -T`, then the SSH daemon should be safely reloaded or restarted according to the organization's change process. OpenSSH documents the server-side `Ciphers` configuration directive at OpenBSD sshd_config: Ciphers. Technical details of the collision risk underlying SWEET32 are available from the SWEET32 research site.

QUESTION NO: 24

A company receives several complaints from customers regarding its website. An engineer implements a parser for the web server logs that generates the following output:

Browser	User location	Load time	HTTP response
Mozilla 5.0	United States	190ms	302
Chrome 110	France	1.2s	302
Microsoft Edge	India	3.7s	307
Microsoft Edge	Australia	6.4s	200

which of the following should the company implement to best resolve the issue?

A. IDS

B. CDN

C. WAF

D. NAC

ANSWER: B

Explanation:

CDN is the appropriate implementation because the log output indicates that website response times differ substantially by the customer's geographic location. When users who are farther from the origin web server have consistently longer load times, network latency and the distance that content must traverse are likely contributing to the poor experience. A content delivery network distributes cached copies of website assets, such as images, scripts, style sheets, and downloadable content, to edge locations around the world. Requests can then be served from an edge server close to the user rather than always requiring a round trip to the company's centralized origin infrastructure.

Using a CDN reduces latency, can improve throughput through localized delivery, and offloads repeated requests for cacheable content from the origin server. This directly addresses customer complaints that correlate with geography while supporting availability and performance at scale. The organization should configure appropriate cache-control behavior, ensure dynamic or sensitive content is handled according to its requirements, and monitor CDN and origin metrics after deployment to validate the improvement. Cloudflare provides an overview of how CDN edge caching works at What is a CDN?, and AWS describes the use of edge locations for low-latency content delivery at Amazon CloudFront Developer Guide.

QUESTION NO: 25

A company is migrating from a Windows Server to Linux-based servers. A security engineer must deploy a configuration management solution that maintains security software across all the Linux servers. Which of the following configuration file snippets is the most appropriate to use?

A. ---
- name: deployment
hosts: linux_servers
remote_user: root
tasks:
- name: Install security software
ansible.builtin.apt:
name: security-software
state: present

B. < hosts > linux_servers < /hosts > < os_type > Linux 3.1 < /os_type > < SElinux > true < /SElinux > < source > com.canonical.io < /source >

C. { " name " : " deployment " , " hosts " : " linux_servers " , " remote_user " : " Administrator " , " tasks " :{ " name " : " Install security software " , " com.microsoft.store.latest " } }

D. { " task " : " install " , " hosts " : " linux_servers " , " remote_user " : " root " , " se_linux " : " false " , " application " : " AppX " }

ANSWER: A

Explanation:

The Ansible YAML playbook that targets `linux_servers` and uses `ansible.builtin.apt` is the appropriate configuration-management snippet. Ansible is designed to automate consistent configuration and software deployment across an inventory group of managed hosts. A playbook identifies the target host group, defines the tasks to run, and declares the intended software state. This allows a security engineer to repeatedly install and maintain a security package across the Linux servers rather than configuring each server individually. The package declaration and `state: present` setting express desired state: the selected security software must be installed. Reapplying the playbook is idempotent, meaning the automation can be run repeatedly to verify and enforce that state without unnecessarily reinstalling software that is already present. This supports scalable, repeatable administration during a platform migration and provides a configuration artifact that can be version controlled and reviewed. The `ansible.builtin.apt` module is specifically appropriate for Debian-family Linux hosts, including Debian and Ubuntu, because it manages packages through the APT package manager. Ansible documents the structure and execution model of YAML playbooks at Ansible Playbooks and documents package installation and state management at ansible.builtin.apt module.

QUESTION NO: 26

A company isolated its OT systems from other areas of the corporate network These systems are required to report usage information over the internet to the vendor Which oi the following b*st reduces the risk of compromise or sabotage ' (Select two).

A. Implementing allow lists

B. Monitoring network behavior

C. Encrypting data at rest

D. Performing boot integrity checks

E. Executing daily health checks

F. Implementing a site-to-site IPSec VPN

ANSWER: A F

Explanation:

Implementing allow lists is appropriate because OT reporting traffic is normally highly predictable: a limited set of devices needs to communicate with a known vendor endpoint, using specific ports, protocols, and services. A default-deny policy that explicitly permits only those required flows minimizes the externally reachable attack surface and helps prevent unauthorized connections, command traffic, or data transfers. This preserves the OT network's isolation while allowing the required business function.

Implementing a site-to-site IPSec VPN provides an authenticated and encrypted communications path between the company and the vendor over the public internet. IPSec protects the confidentiality and integrity of the usage information in transit and validates the communicating peers, reducing opportunities for interception, alteration, or spoofing of the vendor connection. Used with restrictive firewall and allow-list rules, the VPN can limit external connectivity to the specific vendor relationship rather than exposing OT assets directly to the internet. NIST guidance for operational technology emphasizes segmentation, tightly controlled conduits between zones, and protection of remote or external communications. CISA likewise recommends securing industrial-control-system communications and limiting network access to necessary services. See NIST SP 800-82 Rev. 3 and CISA Industrial Control Systems.

QUESTION NO: 27

A systems administrator is in the process of hardening the host systems before connecting to the network. The administrator wants to add protection to the boot loader to ensure the hosts are secure before the OS fully boots.

Which of the following would provide the BEST boot loader protection?

A. TPM

B. HSM

C. PKI

D. UEFI/BIOS

ANSWER: D

Explanation:

UEFI/BIOS is correct because it provides the firmware-level controls used to protect the pre-operating-system boot process, most importantly UEFI Secure Boot. Secure Boot establishes a chain of trust before the operating system loads: firmware verifies the signature of the boot manager, and the boot manager subsequently verifies allowed boot components. This prevents an untrusted, modified, or malicious boot loader from executing, helping defend against bootkits and other malware that attempts to gain control before operating-system security services are active. During host hardening, the administrator should enable Secure Boot in the UEFI firmware configuration and manage the trusted certificates and signing keys appropriately for the organization's approved operating systems and boot loaders. This control directly addresses the requirement to verify boot software before the OS fully boots, rather than relying on protections that primarily operate after startup. Microsoft describes Secure Boot as a security standard that helps ensure a device boots using software trusted by the original equipment manufacturer; see Microsoft Secure Boot documentation. The underlying firmware specification and Secure Boot architecture are maintained by the UEFI Forum.

QUESTION NO: 28

A cloud engineer wants to configure mail security protocols to support email authenticity and enable the flow of email security information to a third-party platform for further analysis. Which of the following must be configured to achieve these requirements? (Select two).

A. DMARC

B. DKIM

C. TLS

D. SPF

E. DNSSEC

F. MX

ANSWER: A B

Explanation:

DMARC and **DKIM** together meet the stated requirements. DKIM provides message-level authenticity and integrity: the sending domain cryptographically signs selected email headers and content, and the receiving system retrieves the public key from DNS to validate that signature. This allows recipients and downstream security tooling to determine whether a message was authorized by the signing domain and whether signed content was altered in transit.

DMARC applies domain-level policy and alignment to the underlying email-authentication results. Crucially for the requirement to share security information with another platform, DMARC supports aggregate and forensic/failure reporting through the `rua` and `ruf` reporting URI tags. A domain can direct these reports to an authorized external reporting destination, enabling a third-party analysis platform to collect authentication outcomes, identify spoofing attempts, and monitor mail sources at scale.

Configuring DKIM establishes a cryptographically verifiable authentication signal, while configuring DMARC supplies the policy, alignment evaluation, and standardized reporting flow needed for external analysis. DMARC reporting receivers should follow the external-destination authorization process to accept reports for domains they analyze. See RFC 6376: DomainKeys Identified Mail and RFC 7489: DMARC.

QUESTION NO: 29

An organization is implementing a zero trust access architecture for contractors who require access to a private application. The organization does not want to expose the application directly to the internet or grant contractors broad network-level access.

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

A. Application-specific access through a zero trust network access gateway.

B. Device-posture validation before access is granted.

C. Direct inbound access to the application from the internet.

D. A shared contractor VPN account.

E. A flat network with unrestricted east-west traffic.

ANSWER: A B

Explanation:

Application-specific access through a zero trust network access gateway and **device-posture validation before access is granted** best meet the requirements. ZTNA brokers access to a specific private application after evaluating the user, device, and contextual policy. This avoids exposing the application directly to the internet and avoids providing the broad internal network access commonly associated with a traditional remote-access VPN.

Device-posture validation adds an important policy signal by confirming that the contractor endpoint satisfies defined requirements, such as approved management status, supported operating-system version, encryption, endpoint protection, or the absence of known compromise indicators. Access can be denied or restricted when the endpoint does not meet the required posture. NIST describes zero trust as an approach that evaluates access requests using identity, device, and contextual information in SP 800-207. CISA also provides Zero Trust Maturity Model guidance.

QUESTION NO: 30

A financial technology firm works collaboratively with business partners in the industry to share threat intelligence within a central platform This collaboration gives partner organizations the ability to obtain and share data associated with emerging threats from a variety of adversaries Which of the following should the organization most likely leverage to facilitate this activity? (Select two).

A. CWPP

B. YAKA

C. ATTACK

D. STIX

E. TAXII

F. JTAG

ANSWER: D E

Explanation:

STIX and TAXII are designed to work together to enable standardized, scalable cyber-threat-intelligence sharing among trusted organizations. STIX, or Structured Threat Information Expression, provides a common, machine-readable model for representing intelligence such as indicators, observed data, threat actors, malware, campaigns, attack patterns, relationships, and confidence information. Using a shared structure lets partners consistently interpret intelligence originating from different security tools and teams, while retaining context about emerging adversary activity.

TAXII, or Trusted Automated Exchange of Intelligence Information, provides the application-layer protocol and API patterns through which that intelligence can be requested, published, discovered, and collected. A central threat-intelligence platform can expose TAXII collections to authorized partner organizations, allowing automated retrieval and distribution of STIX content rather than relying on manual reports or incompatible formats. This combination supports timely collaboration, operationalizes shared intelligence in security tooling, and improves consistency across the partner community. The OASIS specifications describe STIX 2.1 as the core threat-intelligence language and define TAXII 2.1 as the protocol for exchanging that intelligence.

QUESTION NO: 31

An analyst reviews a SIEM and generates the following report:

Host	Rule	Offense Trigger
VM002	Network connection	TCP connection generated to web.corp.local
HOST002	Network connection	Web navigation to comptia.org
HOST002	File download	File download from web browser from web.corp.local
VM002	Network connection	Web navigation to web.corp.local
HOST002	Network connection	Web navigation to comptia.org/files
HOST002	Log-in activity	Log-in successful after two attempts

OnlyHOST002is authorized for internet traffic. Which of the following statements is accurate?

A. The VM002 host is misconfigured and needs to be revised by the network team.

B. The HOST002 host is under attack, and a security incident should be declared.

C. The SIEM platform is reporting multiple false positives on the alerts.

D. The network connection activity is unusual, and a network infection is highly possible.

ANSWER: D

Explanation:

The network connection activity is unusual, and a network infection is highly possible. The key fact is the defined network baseline: only HOST002 is authorized to communicate with the internet. SIEM evidence showing outbound internet connections from other internal systems is therefore a meaningful anomaly that requires security investigation. Unexpected external communications are a common indicator of compromised endpoints, particularly when malware is attempting command-and-control beaconing, retrieving additional payloads, receiving instructions, or transferring collected data outside the organization.

An analyst should validate the alert by correlating the affected hosts' network flows with DNS requests, destination reputation, endpoint process telemetry, proxy logs, and recent authentication activity. This provides the context needed to determine whether the activity represents malicious external communication and to prioritize containment. MITRE identifies Command and Control as the adversary tactic used to communicate with compromised systems, often through common network protocols and external infrastructure: MITRE ATT&CK: Command and Control. CISA's incident-response guidance likewise emphasizes investigating suspicious indicators across network and host data to confirm and respond to a potential compromise: CISA: Incident Detection and Response.

QUESTION NO: 32

A security engineer discovers that some legacy systems are still in use or were not properly decommissioned. After further investigation, the engineer identifies that an unknown and potentially malicious server is also sending emails on behalf of the company. The security engineer extracts the following data for review:

Server IP	DNS	Status	Auth. pass rate
200.100.50.25	marketing.company.com	Authorized	97%
200.50.25.10	29mail.mycrosoft.info	Unauthorized	0%
210.20.20.5	mail.google.com	Authorized	100%

Which of the following actions should the security engineer take next? (Select two).

A. Rotate the DKIM selector to use another key.

B. Change the DMARC policy to reject and remove references to the server.

C. Remove the unnecessary servers from the SPF record.

D. Change the SPF record to enforce the hard fail parameter.

E. Update the MX record to contain only the primary email server.

F. Change the DMARC policy to none and monitor email flow to establish a new baseline.

ANSWER: C D

Explanation:

Removing unnecessary servers from the SPF record is an immediate corrective action because an SPF record identifies the IP addresses, hosts, and third-party services permitted to send mail using the organization's domain. Legacy systems that were retained in DNS after decommissioning unnecessarily expand the set of authorized senders. Eliminating those entries limits authorization to the currently approved mail infrastructure and helps prevent an unmanaged server from successfully representing the company's domain.

Changing the SPF record to enforce the hard fail parameter is also appropriate after validating the legitimate sending services. The `-all` mechanism communicates that mail sent from a source not matched by the preceding authorized mechanisms fails SPF evaluation. This provides a clear authorization boundary for recipient mail systems and strengthens the organization's sender-authentication posture. Maintaining a minimal, accurate SPF record and using a hard-fail policy ensures that outdated infrastructure cannot remain implicitly trusted while legitimate mail continues to be authorized. SPF

evaluation and the meaning of the `-all` mechanism are specified in RFC 7208. Operational guidance for publishing and maintaining SPF records is available from Google Workspace Admin Help.

QUESTION NO: 33

A Chief Information Security Officer is concerned about the operational impact of ransomware. In the event of a ransomware attack, the business requires the integrity of the data to remain intact and an RPO of less than one hour. Which of the following storage strategies best satisfies the business requirements?

A. Full disk encryption

B. Remote journaling

C. Immutable

D. RAID 10

ANSWER: C

Explanation:

Immutable is correct because ransomware recovery depends on retaining recovery data that an attacker cannot encrypt, modify, or delete. Immutable storage applies enforced retention controls to backups, snapshots, or object versions, creating a protected, known-good recovery point even when a ransomware incident affects production systems or privileged accounts. This directly addresses the integrity requirement: the organization can restore from data whose contents and retention state have been protected against unauthorized alteration.

The recovery-point objective can be met by creating immutable snapshots, backup versions, or replicated recovery points on a schedule of less than one hour. For example, frequent protected snapshots can ensure that the maximum acceptable amount of lost data remains below the stated threshold while still preserving a reliable restore source. Immutability is therefore a ransomware-resilience control that combines recoverability with assurance that retained data has not been corrupted or encrypted during the attack. CISA recommends maintaining backups that are protected from modification and deletion as part of ransomware preparedness. Object-lock capabilities are one technical mechanism for implementing such retention protections. See the CISA StopRansomware Guide and AWS documentation for S3 Object Lock.

QUESTION NO: 34

A game developer wants to reach new markets and is advised by legal counsel to include specific age-related sign-up requirements. Which of the following best describes the legal counsel ' s concerns?

A. GDPR

B. LGPD

C. PCI DSS

D. COPPA

ANSWER: D

Explanation:

COPPA is correct because age-related registration controls for an online game directly concern protections for children’s personal information. The Children’s Online Privacy Protection Act and its implementing rule apply to operators of online services, including games and mobile apps, that are directed to children under age 13 or that have actual knowledge they collect personal information from children under 13. A developer subject to COPPA must design its sign-up and data-collection processes to provide clear privacy notices, obtain verifiable parental consent before collecting covered personal information in most circumstances, limit collection to what is reasonably necessary, and provide parents the ability to review or delete their child’s information. Age screening can be an important component of identifying when these obligations apply, particularly when a game is offered to a broad audience that may include children. COPPA also encompasses certain online

identifiers used for persistent tracking, so the concern reaches beyond obvious registration fields such as a name or email address. These requirements make COPPA the legal framework most closely associated with age-gated sign-up flows for a game developer. The FTC's Children's Privacy guidance and the official COPPA Rule in 16 CFR Part 312 describe these obligations.

QUESTION NO: 35

A company is moving several of its systems to a multicloud environment and wants to automate the creation of the new servers using a standard image. Which of the following should the company implement to best support this goal?

A. PowerShell

B. Bash

C. Terraform

D. Ansible

ANSWER: C

Explanation:

Terraform is the best fit because it is an infrastructure-as-code tool designed to define, provision, and manage infrastructure through declarative configuration files. A company can describe its desired server instances, network resources, storage, and standard machine-image references in version-controlled Terraform configurations, then consistently deploy those resources through the relevant cloud-provider integrations. This supports repeatable server builds and reduces configuration drift and manual provisioning errors as systems are moved into more than one cloud environment.

For a multicloud deployment, Terraform is particularly useful because its workflow and HashiCorp Configuration Language remain consistent even when the underlying infrastructure is provided by different platforms. Teams can create reusable modules that encode approved build patterns, including the required standard image and associated security or network settings. Those modules can be reviewed, tested, and applied consistently, supporting governance and operational scalability. Terraform's state management also enables it to track deployed infrastructure and determine the changes needed to reconcile the environment with the declared configuration. HashiCorp describes Terraform as an infrastructure-as-code tool that can provision and manage resources across cloud providers, while its provider ecosystem supplies the interfaces for specific platforms. See HashiCorp Terraform introduction and HashiCorp Terraform providers documentation.

QUESTION NO: 36

An organization is developing an in-house software platform to support capital planning and reporting functions. In addition to role-based access controls and auditing/logging capabilities, the product manager must include requirements associated with archiving data and immutable backups. Which of the following organizational considerations are most likely associated with this requirement? (Select two)

A. Crypto-export management controls

B. Supply chain weaknesses

C. Device attestation

D. Quality assurance

E. Legal hold compliance

F. Ransomware resilience

ANSWER: E F

Explanation:

Legal hold compliance and **Ransomware resilience** are the organizational considerations most directly supported by requirements for data archiving and immutable backups. Legal holds require an organization to preserve relevant electronically stored information when litigation, an investigation, an audit, or a regulatory matter is anticipated. Archived records and protected backup copies help ensure that required information remains available, intact, and defensible throughout the retention period. Immutable storage is especially valuable because it prevents records from being altered or deleted before the applicable hold is released. Microsoft describes legal holds as a means of preserving content relevant to eDiscovery and compliance matters in Microsoft Purview eDiscovery documentation.

Ransomware resilience depends on maintaining recovery data that an attacker cannot encrypt, modify, or erase. Immutable backups provide a known-good restoration point even when a threat actor compromises production systems, privileged accounts, or accessible backup repositories. This capability supports business continuity for critical capital-planning and reporting functions by enabling recovery of trustworthy data after a destructive incident. CISA specifically recommends maintaining offline, encrypted, and immutable backups as a key ransomware preparedness and recovery measure in its StopRansomware Guide.

QUESTION NO: 37

A technician is reviewing the logs and notices a large number of files were transferred to remote sites over the course of three months. This activity then stopped. The files were transferred via TLS-protected HTTP sessions from systems that do not normally send traffic to those sites. The technician will define this threat as:

A. A decrypting RSA using an obsolete and weakened encryption attack.

B. A zero-day attack.

C. An advanced persistent threat.

D. An on-path attack.

ANSWER: C

Explanation:

An advanced persistent threat is the best classification because the observed behavior indicates a deliberate, sustained intrusion involving covert data exfiltration. The transfers continued over three months, which demonstrates persistence rather than a brief or isolated security event. Their origin from systems that do not normally communicate with the remote destinations is a strong anomaly and can indicate that compromised internal hosts were used to collect and send sensitive data. Using HTTP protected by TLS also helps an adversary conceal the content of outbound transfers within legitimate-looking encrypted web traffic, a common technique for evading inspection and detection. The eventual end of activity is consistent with an actor completing an objective, changing infrastructure, or reducing activity to avoid detection. An advanced persistent threat generally involves prolonged unauthorized access, targeted collection of information, and carefully managed exfiltration through channels that blend with ordinary network traffic. MITRE ATT&CK identifies exfiltration over web services and encrypted channels as relevant adversary behaviors: MITRE ATT&CK: Exfiltration. CISA also describes advanced persistent threats as sophisticated, long-term campaigns aimed at gaining and maintaining access to victim environments: CISA: Advanced Persistent Threats.

QUESTION NO: 38

A recent data breach revealed that a company has a number of files containing customer data across its storage environment. These files are individualized for each employee and are used in tracking various customer orders, inquiries, and issues. The files are not encrypted and can be accessed by anyone. The senior management team would like to address these issues without interrupting existing processes.

Which of the following should a security architect recommend?

A. A DLP program to identify which files have customer data and delete them

B. An ERP program to identify which processes need to be tracked

C. A CMDB to report on systems that are not configured to security baselines

D. A CRM application to consolidate the data and provision access based on the process and need

ANSWER: D

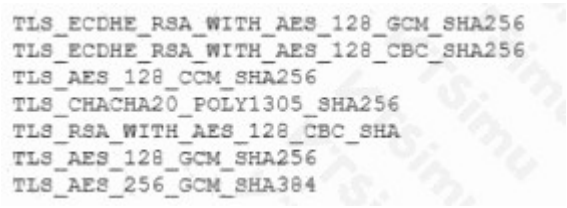
Explanation:

A CRM application to consolidate the data and provision access based on the process and need is the appropriate recommendation. The employee-specific files are being used to record customer orders, inquiries, and issues, which are fundamental customer relationship management functions. A centralized CRM preserves these business activities while replacing dispersed, broadly accessible files with a managed repository designed for customer records and case workflows.

Implementing a CRM enables the organization to define access according to job role and business process, so personnel receive only the customer information necessary to perform assigned duties. This supports least privilege and need-to-know access while reducing exposure caused by unencrypted files that any user can access. The CRM can also apply consistent controls such as authentication, authorization, audit logging, retention policies, secure backups, and workflow ownership. Consolidation gives the organization a practical path to protect and govern customer data without requiring employees to abandon the established order, inquiry, and issue-tracking processes; instead, those processes are migrated into a platform built to support them. NIST identifies least privilege as a key access-control principle in NIST SP 800-53 Rev. 5. For an overview of CRM capabilities and customer-interaction management, see Salesforce: What Is CRM?.

QUESTION NO: 39

A vulnerability scan was performed on a website, and the following encryption suites were found:



Which of the following actions will remediate the vulnerability?

A. Removing any ciphers utilizing cipher block chaining

B. Rearranging the order of the ciphers from strongest to weakest

C. Deploying a WAF to monitor web traffic

D. Reissuing new SSL certificates for the website

ANSWER: A

Explanation:

Removing any ciphers utilizing cipher block chaining is the appropriate remediation because the identified issue is exposure created by TLS cipher suites that use CBC-mode encryption. CBC suites have been associated with practical weaknesses in older TLS deployments, particularly when paired with legacy protocol behavior and vulnerable implementation details. Eliminating these suites from the web server's enabled TLS configuration prevents clients from negotiating the affected cryptographic construction. The server should instead offer modern TLS versions and authenticated-encryption cipher suites, such as AES-GCM or ChaCha20-Poly1305, which provide confidentiality and integrity protections designed for current TLS usage. This is a configuration-level cryptographic remediation: after disabling the affected suites, the service must be rescanned or tested to confirm that no CBC-based suites can be negotiated. NIST recommends configuring TLS servers to support strong, modern cipher suites and to disable weak or deprecated protocol and cipher-suite choices; see NIST SP 800-52 Rev. 2. Additional current TLS implementation guidance, including preferred modern cipher suites, is available from Mozilla SSL Configuration Generator.

QUESTION NO: 40

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: A

Explanation:

Key rotation is appropriate because the keys have been publicly exposed and have been demonstrated to decrypt the affected data. The organization must treat those keys as compromised: they can no longer provide confidentiality for data encrypted with them or reliably authenticate cryptographic operations. Rotating the keys means generating replacement keys using an approved, trusted key-management process; securely distributing and activating them in the relevant systems; and retiring the exposed key material. Where the compromised key encrypted stored data, backups, or archived records that must remain confidential, the organization should re-encrypt that data under the replacement key. Key rotation should also include identifying the key's scope and dependencies, such as applications, services, certificates, databases, and key-encryption keys, so all affected cryptographic uses are transitioned safely. NIST defines key compromise as a condition requiring key replacement and describes lifecycle practices for replacing keys that are no longer suitable for use. OWASP likewise recommends designing key-management processes to support rotation and responding promptly when compromise is suspected or confirmed. See NIST SP 800-57 Part 1 Rev. 5 and the OWASP Key Management Cheat Sheet.

QUESTION NO: 41

A security engineer needs 10 implement a CASB to secure employee user web traffic. A Key requirement is mat relevant event data must be collected from existing on-premises infrastructure components and consumed by me CASB to expand traffic visibility. The solution must be nighty resilient to network outages. Which of the following architectural components would BEST meet these requirements?

A. Log collection

B. Reverse proxy

C. AWAF

D. API mode

ANSWER: A

Explanation:

Log collection is the best fit because it enables a CASB to ingest event and traffic records generated by existing on-premises infrastructure, such as firewalls, web proxies, and secure web gateways. These records provide the CASB with the information needed to identify cloud application use, users, source systems, destinations, uploaded data, and other activity indicators. This is a common approach for cloud discovery, in which the CASB analyzes network-device logs to expand visibility into sanctioned and unsanctioned cloud-service use without requiring every employee connection to pass through a new inline enforcement component.

It also directly supports the resilience requirement. Log collectors and forwarding systems can store, queue, and retry log uploads when a connection to the CASB service is temporarily unavailable. Once connectivity returns, accumulated records can be sent for analysis, preserving visibility while avoiding a dependency on continuous real-time connectivity for employee web access. Microsoft documents Cloud Discovery as analyzing traffic logs from supported firewalls and proxies, and it supports log collectors for automated upload: Microsoft Defender for Cloud Apps: Set up Cloud Discovery and Microsoft Defender for Cloud Apps: Automatic log upload using log collectors.

QUESTION NO: 42

After some employees were caught uploading data to online personal storage accounts, a company becomes concerned about data leaks related to sensitive, internal documentation. Which of the following would the company most likely do to decrease this type of risk?

A. Improve firewall rules to avoid access to those platforms.

B. Implement a cloud-access security broker

C. Create SIEM rules to raise alerts for access to those platforms

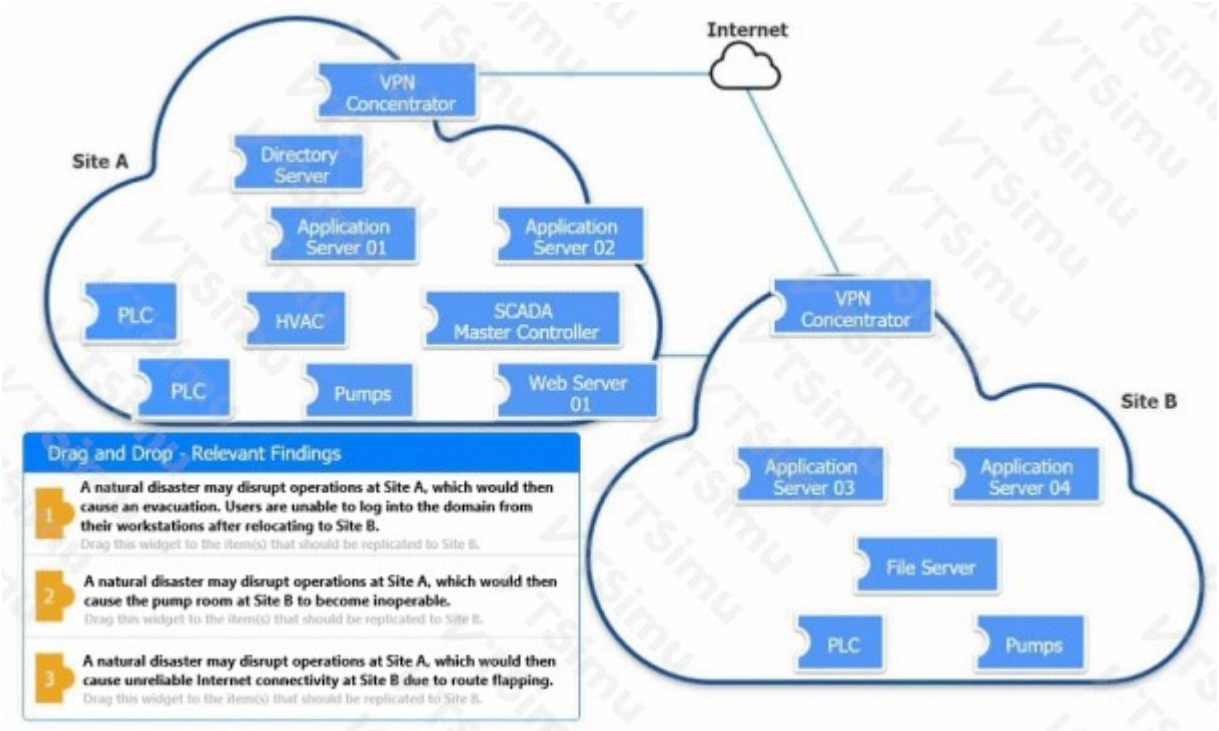
D. Deploy an internet proxy that filters certain domains

ANSWER: B

Explanation:

Implement a cloud-access security broker is the most appropriate measure because a CASB is designed to apply organizational security controls to cloud-service use, including sanctioned and unsanctioned personal storage services. It provides centralized visibility into cloud applications being accessed by users and can enforce policies that prevent sensitive internal documents from being uploaded, shared externally, or accessed inappropriately. In this scenario, the organization needs more than awareness that employees reached a personal-storage site; it needs controls that inspect and govern data movement to cloud applications. A CASB can integrate data loss prevention capabilities to identify sensitive content using classifications, labels, patterns, or contextual rules, then block uploads, quarantine content, require encryption, or create a governed workflow. It can also support shadow-IT discovery, allowing the company to identify unapproved cloud services in use and apply risk-based policy decisions. These functions directly reduce the probability of confidential documentation leaving approved enterprise-controlled repositories. This aligns with SecurityX cloud-security objectives addressing cloud access controls and data-protection mechanisms. See the CompTIA exam objectives resource page and Microsoft's overview of CASB capabilities.

QUESTION NO: 43 - (DRAG DROP)



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

INSTRUCTIONS

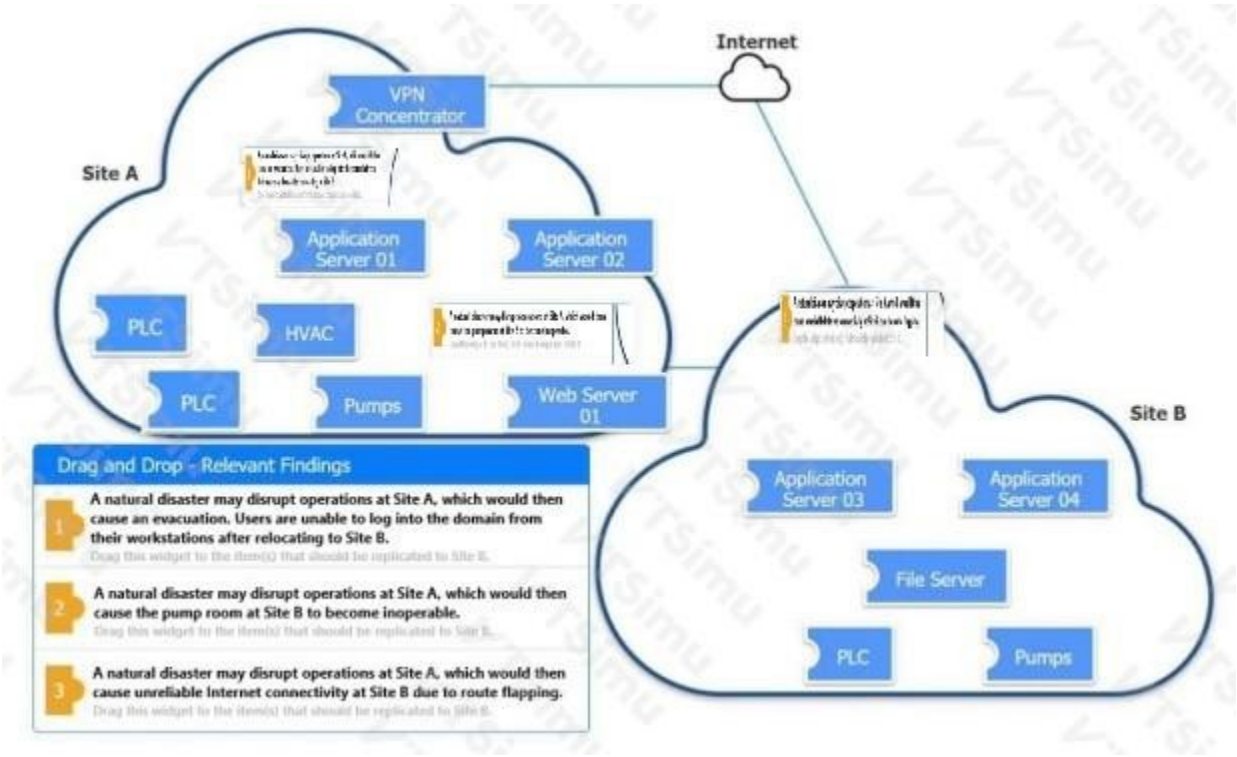
Review the following scenarios and instructions. Match each relevant finding to the affected host.

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:

Continuity planning starts by identifying the services that users and operational systems must still have when the primary site is unavailable. The finding concerning users who cannot log in after relocating to Site B belongs with the Site A Directory Server. Active Directory provides centralized authentication and authorization, and it commonly also provides the directory information needed for domain policy and access to domain resources. Providing a replicated directory-service capability at the recovery site allows relocated users to authenticate without relying on the failed Site A environment. This is consistent with the dependency identification and recovery planning principles in NIST SP 800-34 Rev. 1.

The pump-room continuity finding belongs with the Site A SCADA Master Controller. Although the pumps are physically located at Site B, their continued operation can depend on supervisory control, monitoring, command, and coordination functions provided by the master controller. Replicating or recovering that controller function at Site B preserves the industrial process during a Site A outage. Maintaining these essential industrial-control capabilities reflects the resilience and availability guidance described in NIST SP 800-82 Rev. 3.

The unreliable Internet-connectivity finding belongs with the VPN Concentrator at Site B, because that device is the local network edge affected by unstable route advertisements. The appropriate corrective action is route-flap damping, also called route-flap dampening. This mechanism suppresses repeatedly changing route information for a controlled period, reducing instability while routes converge and preventing continual route changes from disrupting connectivity. The behavior and purpose of route-flap damping are documented in RFC 2439.

QUESTION NO: 44

A security engineer was auditing an organization's current software development practice and discovered that multiple open-source libraries were Integrated into the organization's software. The organization currently performs SAST and DAST on the software it develops.

Which of the following should the organization incorporate into the SDLC to ensure the security of the open-source libraries?

A. Perform additional SAST/DAST on the open-source libraries.

B. Implement the SDLC security guidelines.

C. Track the library versions and monitor the CVE website for related vulnerabilities.

D. Perform unit testing of the open-source libraries.

ANSWER: C

Explanation:

Track the library versions and monitor the CVE website for related vulnerabilities is correct because open-source components are third-party software dependencies whose security posture can change after an application is built and deployed. The organization needs an accurate, continuously maintained inventory of each library and its exact version so it can determine whether a newly disclosed vulnerability affects its software. This is commonly implemented through software composition analysis and a software bill of materials, which identify included components and correlate them with vulnerability intelligence. When a relevant CVE is published, the organization can assess exposure, prioritize remediation according to exploitability and business impact, and upgrade, replace, patch, or otherwise mitigate the affected dependency. This control complements SAST and DAST by addressing known vulnerabilities in externally sourced components throughout the dependency life cycle, including transitive dependencies. NIST's Secure Software Development Framework calls for maintaining provenance and security information for software components and addressing vulnerabilities in third-party components. The CVE Program provides the standardized vulnerability identifiers needed to track disclosures across vendors and tools. See NIST Secure Software Development Framework and CVE Program.

QUESTION NO: 45

A SOC analyst suspects compromised endpoints are using DNS tunneling to exfiltrate data. The organization currently allows endpoints to query any external DNS resolver directly.

Which of the following actions would BEST improve the organization's ability to detect and limit this activity? (Choose two.)

A. Centralize DNS query logging.

B. Restrict endpoints to approved recursive resolvers.

C. Disable DHCP on all user subnets.

D. Replace DNS with static host files.

E. Allow DNS over any outbound port.

F. Configure split tunneling for all remote users.

ANSWER: A B

Explanation:

Centralize DNS query logging provides the telemetry needed to identify tunneling indicators, such as unusually long subdomain labels, high query volumes, elevated NXDOMAIN responses, uncommon record types, high-entropy query names, and communication with suspicious domains. Central collection also enables correlation of DNS events with endpoint, proxy, identity, and network telemetry.

Restrict endpoints to approved recursive resolvers limits opportunities for a compromised system to bypass organizational monitoring by communicating directly with attacker-controlled or public resolvers. Egress controls should permit DNS traffic only to the approved resolvers and block direct outbound DNS requests from ordinary endpoints. The approved resolvers can then apply logging, filtering, reputation controls, and response policies consistently. CISA discusses DNS security and protective DNS services in its Protective DNS guidance.

QUESTION NO: 46

Asecuntv administrator is performing a gap assessment against a specific OS benchmark The benchmark requires the following configurations be applied to endpomts:

• Full disk encryption

* Host-based firewall

• Time synchronization

* Password policies

• Application allow listing

* Zero Trust application access

Which of the following solutions best addresses the requirements? (Select two).

A. CASB

B. SBoM

C. SCAP

D. SASE

E. HIDS

ANSWER: C D

Explanation:

SCAP is appropriate because it is a standardized suite of specifications for expressing security configuration benchmarks and automating assessment of endpoint compliance. SCAP content can represent and evaluate required operating-system settings, including encryption status, host firewall configuration, time-related settings, password-policy parameters, and application-control policies where supported by the platform and associated checks. This makes it particularly useful during a gap assessment: administrators can compare endpoints against a defined benchmark consistently, produce compliance results, and use those results to drive remediation. NIST describes SCAP as a framework for automated vulnerability management, measurement, and policy compliance: NIST SCAP.

SASE addresses the requirement for Zero Trust application access by delivering security services through a cloud-based architecture that commonly incorporates zero trust network access. Rather than trusting a user or device merely because it is connected to a corporate network, this model evaluates identity, device context, and policy before granting access to a specific application. SASE policy enforcement can therefore restrict application access to approved users, devices, and conditions, supporting the benchmark's Zero Trust access objective alongside endpoint-focused configuration compliance. CISA's Zero Trust maturity model provides relevant guidance on application and resource access decisions based on explicit verification: CISA Zero Trust Maturity Model.

QUESTION NO: 47

A threat intelligence company ' s business objective is to allow customers to integrate data directly to different TIPs through an API. The company would like to address as many of the following objectives as possible:

• Reduce compute spend as much as possible.

• Ensure availability for all users.

• Reduce the potential attack surface.

• Ensure the integrity of the data provided.

Which of the following should the company consider to best meet the objectives?

A. Configuring a unique API secret key for accounts

B. Publishing a list of IoCs on a public directory

C. Implementing rate limiting for each registered user

D. Providing a hash of all data that is made available

ANSWER: C

Explanation:

Implementing rate limiting for each registered user best addresses the greatest number of stated API-service objectives. Rate limiting constrains how many requests a consumer can make during a defined period, preventing a single customer, compromised account, or automated process from consuming a disproportionate amount of API capacity. This directly helps control compute consumption and associated operating costs.

It also supports availability by preserving resources for other legitimate users during high-demand periods and by reducing the impact of request floods or application-layer denial-of-service activity. Applying limits to authenticated, registered users allows the provider to enforce fair usage policies, monitor anomalous request behavior, and reduce exposure to abusive API consumption. These controls are especially relevant for an API intended for automated integrations with multiple threat intelligence platforms.

Although data integrity requires complementary controls such as cryptographic signing, hashes, and secure transport, rate limiting is the single listed measure that most broadly supports efficient, resilient, and defensible API operations. CompTIA SecurityX includes availability, secure architecture, and the selection of appropriate security controls among its assessed domains. See the CompTIA SecurityX certification page and OWASP's guidance on rate limiting and request controls.

QUESTION NO: 48 - (SIMULATION)

A security team is responding to malicious activity and needs to determine the scope of impact the malicious activity appears to affect certain version of an application used by the organization Which of the following actions best enables the team to determine the scope of Impact?

A. Performing a port scan

B. Inspecting egress network traffic

C. Reviewing the asset inventory

D. Analyzing user behavior

ANSWER: See the explanation for the answer

Explanation:

Answer: C. Reviewing the asset inventory

Review the organization's hardware/software asset inventory to identify every system running the affected application and version. This establishes the potentially impacted population, supports scoping of the incident, and enables prioritization of investigation, containment, and remediation.

Port scans identify exposed services but generally do not reliably identify installed application versions.
Egress traffic inspection may identify active compromise or exfiltration, but it does not determine every affected instance.
User behavior analysis can identify anomalous activity but does not inventory vulnerable application versions.

QUESTION NO: 49

An organization's existing infrastructure includes site-to-site VPNs between datacenters. In the past year, a sophisticated attacker exploited a zero-day vulnerability on the VPN concentrator. Consequently,

the Chief Information Security Officer (CISO) is making infrastructure changes to mitigate the risk of service loss should another zero-day exploit be used against the VPN solution.

Which of the following designs would be BEST for the CISO to use?

A. Adding a second redundant layer of alternate vendor VPN concentrators

B. Using Base64 encoding within the existing site-to-site VPN connections

C. Distributing security resources across VPN sites

D. Implementing IDS services with each VPN concentrator

E. Transitioning to a container-based architecture for site-based services

ANSWER: A

Explanation:

Adding a second redundant layer of alternate vendor VPN concentrators is the best design because it directly improves the availability and cyber resilience of the inter-datacenter connectivity when a future zero-day affects a VPN platform. Redundancy supplies an independent failover capability, allowing critical site-to-site traffic to continue when a vulnerable concentrator must be isolated, rebuilt, or patched. Selecting an alternate vendor adds diversity: a vulnerability in one vendor's codebase, appliance operating system, management plane, or implementation is less likely to compromise both VPN solutions simultaneously. This reduces the organization's dependence on a single technology monoculture and creates an alternate path for essential services during incident response. The design therefore addresses the stated concern—preventing service loss from an unknown VPN exploit—rather than only improving detection or making unrelated application-architecture changes. NIST identifies redundancy and diversity as cyber-resiliency techniques that help systems continue operating under adverse conditions; see NIST SP 800-160, Volume 2, Revision 1. NIST contingency-planning guidance also supports establishing alternate capabilities to sustain essential functions during disruptions: NIST SP 800-34 Revision 1.

QUESTION NO: 50

A security analyst is reviewing network connectivity on a Linux workstation and examining the active TCP connections using the command line.

Which of the following commands would be the BEST to run to view only active Internet connections?

A. sudo netstat -antu | grep "LISTEN" | awk '{print$5}'

B. sudo netstat -nlt -p | grep "ESTABLISHED"

C. sudo netstat -plntu | grep -v "Foreign Address"

D. sudo netstat -pnut -w | column -t -s $'\w'

E. sudo netstat -pnut | grep -P ^tcp

ANSWER: E

Explanation:

`sudo netstat -pnut | grep -P ^tcp` is the best choice because it produces a focused view of active TCP Internet sockets and associates each socket with its owning process where privileges permit. In the `netstat` command, `-t` selects TCP sockets and `-u` selects UDP sockets, while `-n` preserves numeric IP addresses and port numbers. Numeric output is especially useful during security analysis because it avoids name-resolution delays and makes endpoints immediately identifiable. The `-p` switch displays the PID and program name for a socket, allowing the analyst to correlate a connection with a local process. By default, `netstat` reports connected sockets rather than including listening sockets; the command does not use `-a` or `-l`, which would broaden the output to include listeners. Finally, `grep -P ^tcp` retains rows that begin with `tcp`, including IPv4 `tcp` and IPv6 `tcp6` entries, and removes the UDP rows selected earlier. The resulting display is therefore tailored to active TCP network connections and the responsible local processes. The flag behavior is documented in the netstat manual page; Red Hat also provides context on using netstat and ss for socket statistics.

QUESTION NO: 51

During a gap assessment, an organization notes that OYOD usage is asignificant risk. The organization implemented administrative policies prohibiting BYOD usage However, the organization has not implemented technical controls to prevent the unauthorized use of BYOD assets when accessing the organization's resources. Which of the following solutions should the organization implement to b»« reduce the risk of OYOD devices? (Select two).

A. Cloud 1AM to enforce the use of token based MFA

B. Conditional access, to enforce user-to-device binding

C. NAC, to enforce device configuration requirements

D. PAM. to enforce local password policies

E. SD-WAN. to enforce web content filtering through external proxies

F. DLP, to enforce data protection capabilities

ANSWER: B C

Explanation:

Conditional access, to enforce user-to-device binding and NAC, to enforce device configuration requirements provide complementary technical enforcement for an organization that prohibits unmanaged personal devices. Conditional access evaluates access requests using contextual signals, including the authenticated user and device state. By requiring a device to be registered, joined, managed, or marked compliant before access is granted, it can bind access to an approved device rather than relying solely on an administrative BYOD prohibition. This control is particularly effective for protecting cloud applications and other identity-integrated organizational resources. Network access control (NAC) applies a similar policy decision at the network connection layer. NAC can identify connecting endpoints, assess whether they meet defined configuration and security requirements, and restrict, quarantine, or deny access when they do not. Used together, these controls reduce the opportunity for an unauthorized BYOD endpoint to access resources: conditional access governs resource access based on identity and device context, while NAC controls whether the endpoint can obtain appropriate network access. This layered approach turns the organization’s stated policy into enforceable technical controls. Microsoft describes device-based Conditional Access controls at Microsoft Learn, and Cisco outlines NAC endpoint visibility and access-policy capabilities at Cisco Identity Services Engine.

QUESTION NO: 52

A vulnerability can on a web server identified the following:

Which of the following actions would most likely eliminate on path decryption attacks? (Select two).

A. Disallowing cipher suites that use ephemeral modes of operation for key agreement

B. Removing support for CBC-based cipher suites

C. Adding TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384

D. Implementing HIPS rules to identify and block BEAST attack attempts

E. Restricting cipher suites to only allow TLS_RSA_WITH_AES_128_CBC_SHA

F. Increasing the key length to 256 for TLS_RSA_WITH_AES_128_CBC_SHA

ANSWER: B C

Explanation:

Removing support for CBC-based cipher suites and adding `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384` address the cryptographic conditions that enable attacks such as BEAST against vulnerable TLS implementations. CBC-mode encryption in older TLS versions has construction and initialization-vector handling characteristics that can allow a capable on-path attacker to infer protected plaintext under specific conditions. Eliminating CBC suites prevents the server and client from negotiating that vulnerable encryption mode.

The specified ECDHE/ECDSA AES-GCM suite instead combines ephemeral Elliptic Curve Diffie-Hellman key establishment with AES in Galois/Counter Mode, an authenticated-encryption mode. ECDHE creates fresh session keying material for each connection and provides forward secrecy, while AES-GCM provides confidentiality and integrity without CBC-mode chaining. Together, these settings move TLS negotiation to modern protections rather than attempting to detect attacks after a weak session has already been established. NIST recommends configuring TLS implementations with approved modern cipher suites and avoiding obsolete or vulnerable configurations; its TLS guidance is available in NIST SP 800-52 Rev. 2. The AES-GCM TLS cipher-suite definitions are specified in RFC 5288.

QUESTION NO: 53

A company reduced its staff 60 days ago, and applications are now starting to fail. The security analyst is investigating to determine if there is malicious intent for the application failures. The security analyst reviews the following logs:

22:03:50 sshd[21502]: Success login for user01 from 192.168.2.5

22:10:00 sshd[21502]: Failed login for user10 from 192.168.2.5

22:11:40 sshd[21502]: Success login for user07 from 192.168.2.58

22:12:00 sshd[21502]: Failed login for user10 from 192.168.2.5

22:13:00 sshd[21502]: Failed login for user10 from 192.168.2.5

22:13:00 sshd[21502]: Success login for user03 from 192.168.2.27

22:13:00 sshd[21502]: Failed login for user10 from 192.168.2.5

Which of the following is the most likely reason for the application failures?

A. The user's account was set as a service account.

B. The user's home directory was deleted.

C. The user does not have sudo access.

D. The root password has been changed.

ANSWER: B

Explanation:

The user's home directory was deleted is the most likely reason for the application failures. The repeated SSH failures for user10 originate from the same internal address, which is consistent with an automated application, job, or service trying to authenticate with credentials associated with that account. In the context of a staff reduction, user10 may have been offboarded and the account removed as part of normal deprovisioning. If the account's home directory was also deleted, the application could lose files it needs to operate, such as SSH keys, configuration files, scripts, runtime data, or application-specific environment settings.

The delayed onset is also plausible: a scheduled process may run only periodically, or the application may have continued using cached data until it next required a file or credential stored in that directory. This is an availability issue caused by incomplete dependency discovery during offboarding rather than evidence, by itself, of malicious activity. Secure account-management procedures should identify nonhuman and application dependencies before deleting accounts and their associated data. The Linux `userdel` utility can remove a user's home directory and mail spool when used with its removal option, as documented in the userdel manual page. NIST also emphasizes managing account lifecycle actions, including account removal or disabling, in NIST SP 800-53, AC-2.

QUESTION NO: 54

A compliance officer isfacilitating abusiness impact analysis (BIA)and wantsbusiness unit leadersto collect meaningful data. Several business unit leaders want more information about the types of data the officer needs.

Which of the following data types would be the most beneficial for the compliance officer?(Select two)

A. Inventory details

B. Applicable contract obligations

C. Costs associated with downtime

D. Network diagrams

E. Contingency plans

F. Critical processes

ANSWER: C F

Explanation:

Costs associated with downtime and **Critical processes** are the most beneficial data types because they provide the core inputs needed for a business impact analysis. A BIA identifies the business functions that are essential to the organization and evaluates the consequences if those functions are disrupted. Business unit leaders are best positioned to identify their unit's critical processes, including the services, workflows, resources, upstream dependencies, and downstream activities required to deliver essential outcomes. This lets the organization establish sound recovery priorities based on mission and operational importance.

Downtime-cost data translates a disruption into measurable business impact. It can include lost revenue, lost productivity, missed service-level commitments, contractual penalties, regulatory exposure, customer impact, and escalating operational losses over time. Together, the criticality of a process and its downtime impact support determination of maximum tolerable downtime and recovery-time objectives, which are key continuity and resilience planning outputs. NIST describes the BIA as identifying critical resources, impacts from disruptions, and allowable outage periods in support of contingency planning. Ready.gov similarly notes that BIAs assess disruption consequences and collect financial and operational impact information. See NIST SP 800-34 Rev. 1 and Ready.gov: Business Impact Analysis.

QUESTION NO: 55

Source code snippets for two separate malware samples are shown below:

Sample 1:

```
knockEmDown(String e) {
if(target.isAccessed()) {
target.toShell(e);
System.out.printIn(e.toString());
c2.sendTelemetry(target.hostname.toString + " is " + e.toString());
} else {
target.close();
}
}
```

Sample 2:

```
targetSys(address a) {
if(address.islpv4()) {
address.connect(1337);
address.keepAlive("paranoid");
```

```
String status = knockEmDown(address.current);

remote.sendC2(address.current + " is " + status);

} else {

throw Exception e;

}

}
```

Which of the following describes the most important observation about the two samples?

A. Telemetry is first buffered and then transmitted in paranoid mode.

B. The samples were probably written by the same developer.

C. Both samples use IP connectivity for command and control.

D. Sample 1 is the target agent while Sample 2 is the C2 server.

ANSWER: B

Explanation:

The samples were probably written by the same developer. The strongest observation is the apparent reuse of the distinctive `knockEmDown` routine: it is defined in one sample and called by the other. This is a meaningful shared code artifact, particularly because the function name is unusual and therefore less likely to be an independently chosen generic identifier. The snippets also exhibit parallel implementation patterns for reporting status to command-and-control functionality. Both construct a message from a host or address value, concatenate the literal text `" is "`, and send the resulting status through a C2-related method. Shared helper routines, unique naming conventions, formatting habits, and telemetry-message structures are commonly used to associate malware samples with a common codebase, author, or development team. This conclusion is appropriately probabilistic: code may be copied, inherited, or modified by another actor, so the snippets alone do not establish definitive attribution. Nevertheless, code similarity is a high-value malware-analysis indicator for clustering related samples and tracking adversary software. MITRE documents malware as adversary software that can be tracked and analyzed across campaigns; see MITRE ATT&CK Software. NIST also describes analyzing malicious-code artifacts and indicators as part of malware incident handling in NIST SP 800-83 Rev. 1.

QUESTION NO: 56

A cloud engineer needs to identify appropriate solutions to:

• Provide secure access to internal and external cloud resources.

• Eliminate split-tunnel traffic flows.

•Enable identity and access management capabilities.

Which of the following solutions arc the most appropriate? (Select two).

A. Federation

Federation: Useful for identity management but does not eliminate split-tunnel traffic or provide comprehensive security.

B. Micro segmentation

Microsegmentation: Enhances security within the network but does not directly address secure access to cloud resources or split-tunnel traffic.

C. CASB

D. PAM

PAM (Privileged Access Management): Focuses on managing privileged accounts and does not provide comprehensive access control for internal and external resources.

E. SD-WAN

SD-WAN: Enhances WAN performance but does not inherently provide the identity and access management capabilities or eliminate split-tunnel traffic.

F. SASE

ANSWER: C F

Explanation:

CASB and SASE are the appropriate combined solutions because they provide cloud-focused security enforcement while supporting identity-aware access and full traffic inspection. CASB acts as a policy enforcement point for cloud-service use. It provides visibility into cloud applications, applies access and data-protection policies, and can integrate with an organization's identity provider to make decisions based on user identity, group, device state, location, and session risk. This supports governed access to external cloud services and strengthens identity and access management controls. Microsoft describes cloud access security broker capabilities as visibility, data control, threat protection, and compliance controls for cloud applications in its Defender for Cloud Apps documentation.

SASE delivers networking and security functions from cloud service edges, enabling users to access private/internal applications as well as internet and cloud resources through centrally enforced, identity- and context-aware policies. Its architecture commonly combines secure web gateway, zero trust network access, firewall as a service, and cloud access security broker capabilities. Steering traffic through those security services rather than allowing direct split-tunnel bypasses enables consistent inspection and policy enforcement. Gartner's SASE definition describes this convergence of network and security capabilities.

QUESTION NO: 57

A company updates itscloud-based services by saving infrastructure code in a remote repository. The code is automatically deployed into the development environment every time the code is saved lo the repository The developers express concern that the deployment often fails, citing minor code issues and occasional security control check failures in the development environment Which of the following should a security engineer recommend to reduce the deployment failures? (Select two).

A. Software composition analysis

B. Pre-commit code linting

C. Repository branch protection

D. Automated regression testing

E. Code submit authorization workflow

F. Pipeline compliance scanning

ANSWER: B F

Explanation:

Pre-commit code linting and **Pipeline compliance scanning** are the appropriate recommendations because they move the specific failure checks identified by the developers earlier in the delivery workflow. Pre-commit linting evaluates infrastructure-as-code syntax, formatting, and defined coding conventions before changes enter the shared repository. This gives developers rapid, local feedback on minor defects that could otherwise trigger an automated deployment and fail after consuming pipeline and environment resources. For infrastructure code, linting can also identify common configuration errors before provisioning begins.

Pipeline compliance scanning integrates required security-control and policy checks into the CI/CD process. Since some deployments are failing due to security control checks, running those checks consistently as an explicit pipeline stage provides clear, repeatable feedback before deployment to the development environment. The pipeline can enforce

organization requirements such as approved configurations, prohibited settings, secrets detection, and infrastructure policy rules, allowing remediation before an attempted release. Together, these controls support a shift-left DevSecOps approach: code-quality issues are found before commit, while security and compliance requirements are validated automatically as part of delivery. OWASP describes integrating security practices throughout CI/CD in its DevSecOps Guideline, and GitHub documents how automated workflows can enforce checks before changes are accepted or deployed in its GitHub Actions documentation.

QUESTION NO: 58

After remote desktop capabilities were deployed in the environment, various vulnerabilities were noticed.

• Exfiltration of intellectual property

• Unencrypted files

• Weak user passwords

Which of the following is the best way to mitigate these vulnerabilities? (Select two).

A. Implementing data loss prevention

B. Deploying file integrity monitoring

C. Restricting access to critical file services only

D. Deploying directory-based group policies

E. Enabling modern authentication that supports MFA

F. Implementing a version control system

G. Implementing a CMDB platform

ANSWER: A D

Explanation:

Implementing data loss prevention is appropriate because DLP controls discover, classify, monitor, and protect sensitive organizational data. For remote desktop use, DLP policies can identify intellectual property and apply controls to reduce unauthorized transfer through channels such as copying data to removable media, uploading it to cloud services, printing, or clipboard use, depending on the deployed platform and policy configuration. This directly mitigates the stated risk of intellectual-property exfiltration.

Deploying directory-based group policies provides centralized, enforceable configuration management for domain-joined endpoints and users. Administrators can use policy to require password complexity, minimum length, password history, and account-lockout settings, addressing weak user passwords. Policies can also centrally deploy and enforce encryption-related settings, such as BitLocker configuration and recovery-key escrow, so files stored on managed endpoints are protected rather than left unencrypted. Combining centrally enforced endpoint and identity policy with DLP addresses all three identified findings: data exfiltration, lack of encryption, and weak passwords. This aligns with SecurityX objectives covering data protection and enterprise security controls. See Microsoft BitLocker documentation and Microsoft Purview DLP documentation.

QUESTION NO: 59

A company's help desk is experiencing a large number of calls from the finance department slating access issues to www bank com The security operations center reviewed the following security logs:

Which of the following is most likely the cause of the issue?

A. Recursive DNS resolution is failing

B. The DNS record has been poisoned.

C. DNS traffic is being sinkholed.

D. The DNS was set up incorrectly.

ANSWER: C

Explanation:

DNS traffic is being sinkholed. DNS sinkholing is a protective control in which a security product or DNS service deliberately responds to queries for a prohibited or known-malicious domain with an address chosen by the defender, rather than allowing the client to reach the domain's legitimate destination. The controlled destination might be a block page, a monitoring host, or a non-routable address. This approach both prevents the outbound connection and provides useful telemetry about affected hosts and users.

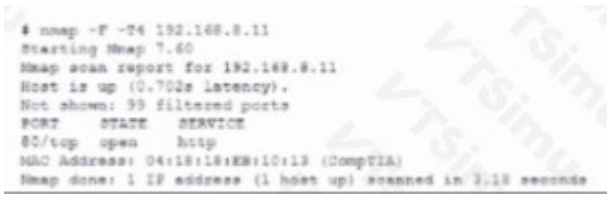
In this scenario, repeated access failures limited to users attempting to reach the same banking-related destination are consistent with an intentional DNS-policy action. If the security logs show the queried name resolving to an internal, security-controlled, or otherwise unexpected address, the resulting HTTPS connection can fail because the server reached after redirection cannot present a certificate valid for the requested banking hostname. A client-certificate or TLS-related HTTP error can therefore appear after the DNS redirection, even though the underlying access-control action occurred at DNS resolution. The concentration of calls from finance also makes this especially important to validate promptly, since financial destinations are commonly subject to phishing and fraud-protection policies. CISA describes protective DNS as a service that blocks requests to known or suspected malicious domains: CISA Protective DNS Services. Additional DNS security guidance is available from the UK NCSC: NCSC DNS security guidance.

QUESTION NO: 60

----------72

A security auditor needs to review the manner in which an entertainment device operates. The auditor is analyzing the output of a port scanning tool to determine the next steps in the security review. Given the following log output.

The best option for the auditor to use NEXT is:



A. A SCAP assessment.

B. Reverse engineering

C. Fuzzing

D. Network interception.

ANSWER: D

Explanation:

Network interception is the appropriate next step because a port scan establishes only that network services are reachable; it does not reveal how the entertainment device behaves in normal operation or what information it exchanges with other systems. Capturing traffic while the device is used allows the auditor to identify the actual protocols in use, destination hosts, DNS queries, authentication exchanges, software-update requests, telemetry, and any unencrypted or weakly protected data flows. This provides the evidence needed to determine which exposed services and communications warrant more focused testing. In a controlled assessment environment, the auditor can collect packets using a network tap, span port, or an appropriately configured proxy, then inspect the traffic with a protocol analyzer. NIST's technical security-testing guidance includes packet capture and analysis among techniques used to gather evidence about networked-system behavior and identify potential vulnerabilities. Wireshark likewise documents packet capture and protocol dissection capabilities that

support inspection of device communications. See NIST SP 800-115, Technical Guide to Information Security Testing and Assessment and the Wireshark User's Guide.

QUESTION NO: 61

During an incident response activity, the response team collected some artifacts from a compromised server, but the following information is missing:

• Source of the malicious files

• Initial attack vector

• Lateral movement activities

The next step in the playbook is to reconstruct a timeline. Which of the following best supports this effort?

A. Executing decompilation of binary files

B. Analyzing all network routes and connections

C. Performing primary memory analysis

D. Collecting operational system logs and storage disk data

ANSWER: D

Explanation:

Collecting operational system logs and storage disk data best supports timeline reconstruction because these evidence sources contain complementary, time-stamped records of activity on the compromised server. Operating system and security logs can establish sequences involving account authentication, remote access, process execution, service changes, scheduled-task creation, and other events relevant to identifying the initial access path and tracing activity that may indicate lateral movement. Disk collection provides persistent forensic artifacts, including file-system timestamps, executable and script metadata, event-log files, persistence artifacts, browser or command-history data, and potentially deleted files. Correlating these artifacts enables responders to place the arrival or creation of malicious files into a broader chronology and associate that activity with user accounts, processes, and system changes. This evidence-driven reconstruction helps establish what occurred, when it occurred, and how the affected host was involved in the incident. NIST incident-response guidance describes analyzing available incident data to determine the nature, extent, and source of an incident, while its forensic guidance discusses using data from systems and storage media to reconstruct events. See NIST SP 800-61 Rev. 2 and NIST SP 800-86.

QUESTION NO: 62

A security team is hardening an Active Directory environment against Kerberoasting attacks. The team identified several legacy service accounts that use weak passwords and have service principal names assigned.

Which of the following actions would BEST reduce the risk? (Choose two.)

A. Migrate eligible services to group Managed Service Accounts.

B. Use long, randomly generated passwords for remaining service accounts.

C. Configure all service accounts as Domain Admins.

D. Disable Kerberos authentication for domain users.

E. Use a shared password for all service accounts.

ANSWER: A B

Explanation:

Migrate eligible services to group Managed Service Accounts and **use long, randomly generated passwords for remaining service accounts** are the best actions. Kerberoasting targets service accounts associated with service principal names by requesting service tickets and attempting offline password cracking against the ticket material. Group Managed Service Accounts provide automatically managed, complex passwords and reduce reliance on manually maintained service-account credentials.

Where a legacy service cannot use a gMSA, assigning a long, randomly generated password substantially increases the effort required to crack a captured service ticket. The organization should also apply least privilege and remove unnecessary SPNs, but disabling Kerberos entirely or broadly removing service accounts would disrupt domain operations. Microsoft documents password management and service identity benefits of group Managed Service Accounts. MITRE ATT&CK describes Kerberoasting under Steal or Forge Kerberos Tickets: Kerberoasting.

QUESTION NO: 63

The material finding from a recent compliance audit indicate a company has an issue with excessive permissions. The findings show that employees changing roles or departments results in privilege creep. Which of the following solutions are the best ways to mitigate this issue? (Select two).

Setting different access controls defined by business area

A. Implementing a role-based access policy

B. Designing a least-needed privilege policy

C. Establishing a mandatory vacation policy

D. Performing periodic access reviews

E. Requiring periodic job rotation

ANSWER: A D

Explanation:

Implementing a role-based access policy and performing periodic access reviews directly address privilege creep arising when employees transfer between roles, teams, or departments. Role-based access control assigns permissions according to defined job functions and business responsibilities rather than maintaining ad hoc, user-specific grants. A well-managed role change process can therefore remove entitlements associated with the employee's former role and provision only the access required by the new role. This operationalizes least privilege while making access administration more consistent and scalable.

Performing periodic access reviews provides the governance control needed to identify and remediate permissions that were not removed during transfers or that are no longer justified by current business responsibilities. Managers and application or data owners can recertify access, confirm an ongoing business need, and revoke stale entitlements. Together, role-based assignment and recurring review support an effective joiner-mover-leaver identity lifecycle, produce evidence for audit remediation, and reduce the risk that accumulated permissions enable unauthorized access. NIST's Role-Based Access Control project describes RBAC as a means of managing permissions through organizational roles. NIST SP 800-53 also identifies account management and periodic review of access authorizations as important access-control practices; see NIST SP 800-53 Rev. 5, Update 1.

QUESTION NO: 64

A development team created a mobile application that contacts a company's back-end APIs housed in a PaaS environment. The APIs have been experiencing high processor utilization due to scraping activities. The security engineer needs to recommend a solution that will prevent and remedy the behavior.

Which of the following would BEST safeguard the APIs? (Choose two.)

A. Bot protection

B. OAuth 2.0

C. Input validation

D. Autoscaling endpoints

E. Rate limiting

F. CSRF protection

ANSWER: A E

Explanation:

Bot protection and **rate limiting** directly address the automated, high-volume nature of scraping while preserving API availability. Bot protection identifies and mitigates unwanted automated clients using signals such as client reputation, behavioral patterns, request characteristics, fingerprinting, and bot-management policies. This allows the organization to distinguish approved mobile-app traffic and legitimate users from scraping tools, then challenge, block, or otherwise restrict malicious automation before it consumes substantial PaaS compute capacity.

Rate limiting enforces defined request-volume thresholds over a time interval. Limits can be applied by API key, access token, authenticated account, client identifier, source address, endpoint, or a combination of these attributes. Applying limits at an API gateway, web application firewall, or PaaS ingress layer constrains excessive request bursts and sustained high-frequency retrieval attempts. It therefore reduces processor exhaustion, gives back-end services a predictable operating boundary, and provides a practical remediation mechanism during an active scraping event.

Used together, these controls provide layered protection: bot protection targets abusive automation specifically, while rate limiting contains request consumption even when traffic has not yet been conclusively classified. OWASP identifies resource-consumption controls, including rate limiting, as essential API protections; see the OWASP API Security Top 10 guidance on unrestricted resource consumption and the OWASP REST Security Cheat Sheet.

QUESTION NO: 65

A security architect is reviewing the following proposed corporate firewall architecture and configuration:

```
DMZ architecture
Internet---------70.54.30.1-[Firewall_A]----192.168.1.0/24----[Firewall_B]----10.0.0.0/16----corporate net

Firewall_A ACL
10 PERMIT FROM 0.0.0.0/0 TO 192.168.1.0/24 TCP 80,443
20 DENY FROM 0.0.0.0/0 TO 0.0.0.0/0 TCP/UDP 0-65535

Firewall_B ACL
10 PERMIT FROM 10.0.0.0/16 TO 192.168.1.0/24 TCP 80,443
20 PERMIT FROM 10.0.0.0/16 TO 0.0.0.0/0 TCP/UDP 0-65535
30 PERMIT FROM 192.168.1.0/24 TO $DB_SERVERS TCP/UDP 3306
40 DENY FROM 192.168.1.0/24 TO 10.0.0.0/16 TCP/UDP 0-65535
```

Both firewalls are stateful and provide Layer 7 filtering and routing. The company has the following requirements:

Web servers must receive all updates via HTTP/S from the corporate network.

Web servers should not initiate communication with the Internet.

Web servers should only connect to preapproved corporate database servers.

Employees' computing devices should only connect to web services over ports 80 and 443.

Which of the following should the architect recommend to ensure all requirements are met in the MOST secure manner? (Choose two.)

A. Add the following to Firewall_A: 15 PERMIT FROM 10.0.0.0/16 TO 0.0.0.0/0 TCP 80,443

B. Add the following to Firewall_A: 15 PERMIT FROM 192.168.1.0/24 TO 0.0.0.0 TCP 80,443

C. Add the following to Firewall_A: 15 PERMIT FROM 10.0.0.0/16 TO 0.0.0.0/0 TCP/UDP 0-65535

D. Add the following to Firewall_B: 15 PERMIT FROM 0.0.0.0/0 TO 10.0.0.0/16 TCP/UDP 0-65535

E. Add the following to Firewall_B: 15 PERMIT FROM 10.0.0.0/16 TO 0.0.0.0 TCP/UDP 0-65535

F. Add the following to Firewall_B: 15 PERMIT FROM 192.168.1.0/24 TO 10.0.2.10/32 TCP 80,443

ANSWER: A F

Explanation:

Adding *15 PERMIT FROM 10.0.0.0/16 TO 0.0.0.0/0 TCP 80,443* to Firewall_A enforces the employee-device requirement with a least-privilege egress rule. It permits corporate endpoints in the specified network to establish only HTTP and HTTPS connections to web services. Because Firewall_A is stateful, return packets for those permitted, established sessions can be allowed automatically without adding a broad inbound rule. Layer 7 inspection can further validate that the permitted traffic is actually web traffic rather than merely traffic using those port numbers.

Adding *15 PERMIT FROM 192.168.1.0/24 TO 10.0.2.10/32 TCP 80,443* to Firewall_B creates a narrowly scoped path from the web-server subnet to one explicitly approved corporate database host. The source subnet, individual destination address, transport protocol, and permitted ports are all constrained. This supports the required server-to-database connectivity while preserving the intended segmentation between the web tier and the corporate network. Combined with the existing policy and stateful filtering, these rules avoid granting the web-server subnet general Internet initiation capability or unrestricted access to internal systems. This is consistent with NIST firewall guidance to implement explicit, restrictive policy rules between network zones and with SecurityX's emphasis on architecture, segmentation, and least privilege. See NIST SP 800-41 Rev. 1 and CompTIA SecurityX.

QUESTION NO: 66

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, **PAM**, and **SIEM** collectively satisfy the required logging, privileged-access, and detection capabilities. A log aggregator centrally collects telemetry from critical systems and retains it in a searchable repository. This provides the historical data needed to meet the 365-day retention requirement and enables analysts to query device activity during monitoring, incident response, and threat-hunting activities. NIST identifies log generation, collection, storage, retention, and analysis as essential log-management functions in NIST SP 800-92, Guide to Computer Security Log Management.

PAM controls and records access to elevated accounts, including administrative credentials and privileged sessions. Capabilities such as credential vaulting, just-in-time access, approval workflows, least-privilege enforcement, session recording, and audit trails directly address the need to tightly control and track privileged user activity.

SIEM ingests and correlates security telemetry at scale, applies detection analytics and threat intelligence, and creates prioritized alerts for suspicious behavior. Those capabilities support rapid identification of ransomware indicators and anomalous activity associated with exploitation of previously unknown vulnerabilities. A SIEM also makes the retained logs operationally useful by linking events across endpoints, identities, applications, and network devices. See Microsoft Sentinel SIEM overview for an example of centralized security analytics, detection, investigation, and response.

QUESTION NO: 67

A cloud engineer needs to identify appropriate solutions to:

• Provide secure access to internal and external cloud resources.

• Eliminate split-tunnel traffic flows.

•Enable identity and access management capabilities.

Which of the following solutions arc the most appropriate? (Select two).

A. Federation

B. Microsegmentation

C. CASB

D. PAM

E. SD-WAN

F. SASE

ANSWER: C F

Explanation:

CASB and SASE together most directly satisfy the stated cloud-access requirements. A CASB provides policy enforcement and visibility for cloud-service use, helping an organization control access to sanctioned and unsanctioned cloud applications, protect data, and apply controls based on user identity, device posture, and session context. In practice, a CASB commonly integrates with an enterprise identity provider so that cloud access is governed through centralized authentication and authorization.

SASE delivers cloud-based networking and security capabilities at the access edge. It combines secure connectivity with security services such as secure web gateways, zero-trust network access, firewall capabilities, and data protection controls. This model supports identity-aware access to both private/internal applications and internet or SaaS resources. By steering user traffic through the SASE security service for inspection and policy enforcement, an organization can avoid a split-tunnel design in which some traffic bypasses enterprise security controls. The result is consistent access policy and inspection regardless of where the user, application, or cloud resource resides.

The Cloud Security Alliance describes CASB capabilities and architecture at CASB Reference Architecture. CISA also outlines the SASE model and its security-service integration in its SASE Reference Architecture.