Comptia

(CAS-004)

CompTIA Advanced Security Practitioner (CASP+) CAS-004

Total: **642 Questions**

Link:

Question: 1

An organization is referencing NIST best practices for BCP creation while reviewing current internal organizational processes for mission-essential items.
Which of the following phases establishes the identification and prioritization of critical systems and functions?

A. Review a recent gap analysis.
B. Perform a cost-benefit analysis.
C. Conduct a business impact analysis.
D. Develop an exposure factor matrix.

Answer: C

Explanation:

The correct answer is C: Conduct a business impact analysis (BIA). A BIA is a crucial phase in business continuity planning (BCP) that focuses on identifying and prioritizing an organization's mission-essential functions and critical systems. It aims to understand the potential impact of disruptions to these functions, considering financial, operational, and reputational consequences. The BIA helps determine the recovery time objectives (RTOs) and recovery point objectives (RPOs) for each system, guiding resource allocation and recovery strategies.

Option A, reviewing a recent gap analysis, while valuable, comes into play after identifying the business-critical components. Gap analysis highlights discrepancies between the current state and the desired state (often defined based on BIA findings), showing areas needing improvement. Option B, performing a cost-benefit analysis, assesses the cost of implementing security controls or recovery strategies versus the potential benefits of mitigating risks. This is done after establishing priorities through the BIA. Option D, developing an exposure factor matrix, helps quantify the potential loss from specific threats or vulnerabilities, but again, the 'what' to assess comes from understanding the critical functions identified in the BIA. The BIA provides the foundation for all subsequent steps in the BCP process by clearly defining what needs to be protected and in what order.

The NIST Special Publication 800-34, "Contingency Planning Guide for Federal Information Systems," details the steps involved in BCP, highlighting the BIA as a fundamental activity. This analysis is performed to identify and assess the potential impact resulting from disruptions in business processes, allowing organizations to effectively determine the systems and functions that are most important to restore first. The BIA outlines potential losses, including financial, tangible property, intangible property, and competitive advantages, from these disruptions.

For further reading:

NIST SP 800-34, Contingency Planning Guide for Federal Information Systems:
https://csrc.nist.gov/publications/detail/sp/800-34/rev-1/archive/2010-05-06
Ready.gov - Business Impact Analysis:https://www.ready.gov/business/assessment/impact

Question: 2

An organization is preparing to migrate its production environment systems from an on-premises environment to a cloud service. The lead security architect is concerned that the organization's current methods for addressing risk may not be possible in the cloud environment.
Which of the following BEST describes the reason why traditional methods of addressing risk may not be possible in the cloud?

A. Migrating operations assumes the acceptance of all risk.
B. Cloud providers are unable to avoid risk.

C. Specific risks cannot be transferred to the cloud provider.

D. Risks to data in the cloud cannot be mitigated.

Answer: C

Explanation:

The correct answer is **C. Specific risks cannot be transferred to the cloud provider.**

Here's why:

Cloud computing operates under a shared responsibility model. While the cloud provider assumes responsibility for the security of the cloud (physical infrastructure, network, etc.), the customer retains responsibility for security in the cloud (data, applications, access management, etc.). This means that while you can transfer some operational burdens and associated risks to the provider (e.g., physical security, hardware maintenance), you cannot transfer all risks. Risks related to your data, application vulnerabilities, configuration errors, and compliance remain your responsibility. For example, the cloud provider is responsible for ensuring their servers are physically secure and available, but you are responsible for encrypting your data stored on those servers.

Therefore, traditional risk management methods relying solely on transferring risk (e.g., via insurance or outsourcing everything) won't be fully effective. You still need to assess, mitigate, and manage the risks that remain within your sphere of control within the cloud environment. You cannot simply assume the cloud provider will handle everything. Data breaches, insecure APIs, and compliance violations are just a few examples of risks that the organization retains responsibility for, even when using cloud services. The organisation has to actively manage data in the cloud which leads to risk which cannot be transferred to a cloud provider.

Options A, B, and D are incorrect:

A. Migrating operations assumes the acceptance of all risk. This is false. While some risk acceptance might occur, a comprehensive risk assessment should drive mitigation and transfer strategies before accepting residual risk.

B. Cloud providers are unable to avoid risk. Cloud providers actively work to avoid and mitigate risks within their area of responsibility. They implement security measures, redundancy, and disaster recovery plans to ensure service availability and data protection.

D. Risks to data in the cloud cannot be mitigated. This is also incorrect. Many mitigation strategies exist for data in the cloud, including encryption, access control lists (ACLs), data loss prevention (DLP) tools, and regular backups.

Supporting Resources:

NIST Special Publication 800-146, Cloud Computing Synopsis and Recommendations:
https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-146.pdf
Cloud Security Alliance (CSA):https://cloudsecurityalliance.org/ - Offers guidance on cloud security best practices, including risk management.

ENISA (European Union Agency for Cybersecurity) Cloud Security Guide:
https://www.enisa.europa.eu/topics/cloud-and-big-data/cloud-security/security-guidance

Question: 3

A company created an external application for its customers. A security researcher now reports that the application has a serious LDAP injection vulnerability that could be leveraged to bypass authentication and authorization.
Which of the following actions would BEST resolve the issue? (Choose two.)

A. Conduct input sanitization.
B. Deploy a SIEM.
C. Use containers.
D. Patch the OS
E. Deploy a WAF.
F. Deploy a reverse proxy
G. Deploy an IDS.

Answer: EF

Explanation:

The correct answer is **AE**. Let's break down why:

A. Conduct Input Sanitization: LDAP injection vulnerabilities occur when unsanitized user input is directly incorporated into LDAP queries. By sanitizing input, you neutralize malicious code injected by attackers. This involves validating and encoding user inputs before including them in LDAP queries, ensuring they are treated as data, not executable commands. This prevents attackers from manipulating queries to bypass authentication or access unauthorized data. This is a foundational step to mitigate injection attacks.

E. Deploy a WAF (Web Application Firewall): A WAF acts as a shield between the external application and users, inspecting HTTP traffic for malicious payloads indicative of LDAP injection attempts. WAFs can be configured with rulesets to identify and block suspicious patterns, effectively preventing the exploitation of the LDAP injection vulnerability. Deploying a WAF adds an additional layer of security by proactively preventing malicious requests from reaching the vulnerable application.

Why other options are less effective:

B. Deploy a SIEM (Security Information and Event Management): SIEMs are valuable for log analysis and incident detection, but they are reactive measures. They detect attacks after they've occurred, not necessarily preventing them.

C. Use Containers: While containers offer isolation, they don't directly address the underlying code vulnerability. Containers might limit the impact of a successful attack, but they won't prevent the LDAP injection from occurring if the vulnerable code remains unpatched.

D. Patch the OS: OS patching is crucial for general security, but it won't fix application-level vulnerabilities like LDAP injection. LDAP injection is application level vulnerability that needs application level fix.

F. Deploy a Reverse Proxy: While reverse proxies offer some security benefits like hiding the backend infrastructure, they don't specifically protect against LDAP injection vulnerabilities.

G. Deploy an IDS (Intrusion Detection System): An IDS detects malicious activity, but like a SIEM, it's a reactive measure. It won't prevent the LDAP injection from occurring.

Summary: Input sanitization is a direct preventative measure that ensures user input is safe to use in queries. A WAF provides an extra layer of defence. These two actions are the best way to resolve this specific LDAP injection vulnerability.

Supporting Links:

OWASP LDAP Injection: https://owasp.org/www-project-top-ten/
Web Application Firewall (WAF): https://www.cloudflare.com/learning/cloud/what-is-a-web-application-firewall-waf/

Question: 4

In preparation for the holiday season, a company redesigned the system that manages retail sales and moved it to a cloud service provider. The new infrastructure did not meet the company's availability requirements. During a postmortem analysis, the following issues were highlighted:
1. International users reported latency when images on the web page were initially loading.
2. During times of report processing, users reported issues with inventory when attempting to place orders.
3. Despite the fact that ten new API servers were added, the load across servers was heavy at peak times. Which of the following infrastructure design changes would be BEST for the organization to implement to avoid these issues in the future?

A. Serve static content via distributed CDNs, create a read replica of the central database and pull reports from there, and auto-scale API servers based on performance.

B. Increase the bandwidth for the server that delivers images, use a CDN, change the database to a non-relational database, and split the ten API servers across two load balancers.

C. Serve images from an object storage bucket with infrequent read times, replicate the database across different regions, and dynamically create API servers based on load.

D. Serve static-content object storage across different regions, increase the instance size on the managed relational database, and distribute the ten API servers across multiple regions.

Answer: A

Explanation:

The best solution to address the issues highlighted in the postmortem analysis is **A. Serve static content via distributed CDNs, create a read replica of the central database and pull reports from there, and auto-scale API servers based on performance.**

Here's why:

1. **Latency for International Users (Images):** Content Delivery Networks (CDNs) address latency issues by caching static content (like images) on servers geographically closer to users worldwide. This significantly reduces the time it takes for users to download images, improving their experience.

   Reference:https://www.cloudflare.com/learning/cdn/what-is-a-cdn/
   Reference:https://aws.amazon.com/cloudfront/

2. **Inventory Issues During Report Processing:** Running reports directly on the main database can strain resources and slow down transactional operations (like placing orders). Creating a read replica allows report generation to occur without impacting the primary database's performance.

   Reference:https://aws.amazon.com/rds/features/read-replicas/
   Reference:https://learn.microsoft.com/en-us/azure/azure-sql/database/read-scale-out?view=azuresql

3. **Heavy Load on API Servers:** Despite adding new API servers, the load remains heavy, suggesting a lack of efficient load distribution and scaling. Auto-scaling automatically adjusts the number of API servers based on real-time demand, ensuring optimal performance during peak times and reducing costs during off-peak times.

   Reference:https://azure.microsoft.com/en-us/solutions/architecture/autoscaling-web-apps/
   Reference:https://aws.amazon.com/autoscaling/

Why other options are less suitable:

B: While a CDN helps, changing to a non-relational database might be a significant and potentially unnecessary architectural change if the current relational database adequately handles transactional data.

Splitting API servers across load balancers is a basic step, but auto-scaling provides more dynamic and efficient resource allocation.

C: Object storage with infrequent access is unsuitable for serving frequently accessed images on a retail website. Replicating the database across regions is more focused on disaster recovery than performance issues related to reporting. Dynamic API creation is similar to auto-scaling, but the phrasing lacks precision.

D: Serving static content from object storage across different regions is unclear. Simply increasing the instance size on the managed relational database might not fully resolve contention issues caused by reporting. Distributing API servers across regions can introduce latency if users and servers are geographically distant. Auto-scaling is still a more efficient solution.

Question: 5

During a remodel, a company's computer equipment was moved to a secure storage room with cameras positioned on both sides of the door. The door is locked using a card reader issued by the security team, and only the security team and department managers have access to the room. The company wants to be able to identify any unauthorized individuals who enter the storage room by following an authorized employee.
Which of the following processes would BEST satisfy this requirement?

A. Monitor camera footage corresponding to a valid access request.
B. Require both security and management to open the door.
C. Require department managers to review denied-access requests.
D. Issue new entry badges on a weekly basis.

Answer: A

Explanation:

The best approach to identify unauthorized individuals tailgating into the storage room alongside authorized employees is to **monitor camera footage corresponding to valid access requests (A)**. Here's why:

Tailgating Detection: This directly addresses the specific threat scenario presented in the prompt –unauthorized individuals entering the secure area by following authorized personnel. By correlating camera footage with the card reader logs (valid access requests), security personnel can review the video to identify if anyone followed the authorized employee into the room.

Real-Time Monitoring Capability: Ideally, the monitoring should happen as close to real-time as possible, allowing immediate action if tailgating is detected. This enables security to investigate and potentially intercept unauthorized individuals before they can access the equipment.

Audit Trail: Combining access logs with corresponding video provides a robust audit trail, invaluable for incident response and future security improvements. This trail can be used to analyze the effectiveness of security measures and identify patterns of unauthorized access attempts.

Non-Disruptive: This process is relatively non-disruptive to authorized employees. They simply use their access cards as normal, and the security team monitors the video feed in the background.

Option B (Require both security and management to open the door): While increasing security, this method is highly disruptive and inefficient for normal access, hindering legitimate work.

Option C (Require department managers to review denied-access requests): This addresses failed access attempts but doesn't directly address tailgating following a successful access.

Option D (Issue new entry badges on a weekly basis): This mainly combats lost or stolen badges but doesn't prevent or detect tailgating.

Therefore, monitoring camera footage synchronized with access logs provides the most effective and practical method for detecting and preventing unauthorized entry through tailgating in this scenario. It provides a way to verify that only the authorized employee entered the room during a specific access event.

Supporting Concepts:

Access Control Systems: Card readers and access logs are key components of physical access control systems. Integrating them with video surveillance enhances security by providing a visual verification layer. **Surveillance and Monitoring:** Video surveillance plays a critical role in security, particularly in high-risk areas. The effectiveness is increased when coupled with other security measures.

Audit Trails: Maintaining a comprehensive audit trail (access logs and video recordings) is essential for security investigations, incident response, and compliance.

Authoritative Links:

NIST Special Publication 800-53, Security and Privacy Controls for Information Systems and Organizations: https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final (Focus on access control and audit logging controls)
PCI DSS (Payment Card Industry Data Security Standard): https://www.pcisecuritystandards.org/ (Requirement for physical security controls, including access control)

Question: 6

A company is preparing to deploy a global service.
Which of the following must the company do to ensure GDPR compliance? (Choose two.)

A. Inform users regarding what data is stored.
B. Provide opt-in/out for marketing messages.
C. Provide data deletion capabilities.
D. Provide optional data encryption.
E. Grant data access to third parties.
F. Provide alternative authentication techniques.

Answer: AC

Explanation:

The correct answer is AC. Here's why:

A. Inform users regarding what data is stored: GDPR (General Data Protection Regulation) mandates transparency. Organizations must clearly inform users about what personal data they collect, how it's used, and for what purposes. This is a fundamental principle of GDPR, ensuring individuals have control over their data. Failure to do so violates the core tenets of the regulation.

C. Provide data deletion capabilities: The "right to be forgotten" is a key aspect of GDPR. Individuals have the right to request the deletion of their personal data, and organizations must have mechanisms in place to comply with these requests. This ensures users have the ability to remove their data from systems,
empowering them to control their digital footprint.

Options B, D, E, and F are not core GDPR requirements in the same direct way as options A and C, although some could contribute to compliance depending on the specific scenario.

B. Provide opt-in/out for marketing messages: While important for privacy and marketing best practices (and sometimes required under ePrivacy Directive, which is related to GDPR), it is not a direct, overarching requirement of GDPR related to ALL data processing. It is mainly concerned about direct marketing. **D. Provide optional data encryption:** Encryption is a good security practice and can contribute to GDPR compliance by protecting data, it's not a mandatory requirement under GDPR. Data protection can be achieved through other means. GDPR mandates that organizations implement appropriate technical and

organizational measures to ensure a level of security appropriate to the risk.

E. Grant data access to third parties: GDPR restricts data sharing with third parties. Sharing is only permitted under certain conditions, such as with user consent or a legitimate business need covered by a data processing agreement. This option is essentially the opposite of what GDPR promotes.

F. Provide alternative authentication techniques: While strong authentication is beneficial for security, GDPR doesn't mandate specific authentication methods. It requires organizations to ensure data security, and authentication is just one aspect of it.

Therefore, choices A and C are the most directly and explicitly related to mandatory requirements within the GDPR framework for a global service deployment.

Authoritative Links for Further Research:

Official GDPR Website:https://gdpr.eu/
ICO (Information Commissioner's Office - UK GDPR Regulator) Guide to GDPR:https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/

Question: 7

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: B

Explanation:

Here's a detailed justification for why an expired certificate on a WAF (Web Application Firewall) is the most likely cause of missing logs and lack of visibility into malicious traffic on an external, exposed web server:

Explanation:

A WAF typically sits in front of a web server to inspect incoming HTTP/HTTPS traffic for malicious patterns. If the WAF's SSL/TLS certificate has expired, several issues arise that can prevent it from properly functioning:

1. **SSL/TLS Handshake Failure:** When a client (e.g., a web browser or malicious bot) attempts to connect to the web server via HTTPS, the WAF presents its SSL/TLS certificate. If the certificate is expired, the client's browser will likely refuse the connection or display a warning to the user. The client won't complete the SSL/TLS handshake.

2. **Inability to Decrypt HTTPS Traffic:** If the SSL/TLS handshake fails or is interrupted, the WAF cannot decrypt the HTTPS traffic. Since WAFs rely on decrypting the traffic to inspect its content for malicious payloads, attack patterns, and other suspicious activity, an expired certificate renders the decryption process impossible.

3. **Logging Failure:** Because the WAF cannot inspect the traffic, it cannot generate logs about the traffic, which leads to a blind spot in security monitoring.

4. **Bypass of WAF Protection:** Malicious actors could exploit this by intentionally sending traffic that triggers an error related to the expired certificate, effectively bypassing the WAF.

5. **Impact on visibility:** The visibility into the WAF will be absent because it cannot decrypt the traffic and analyze it.

Why other options are less likely:

A. The user agent client is not compatible with the WAF: While some WAFs might have compatibility issues with certain user agents, it's unlikely that all malicious traffic would originate from incompatible clients. This is usually a more isolated issue. And it would likely be an issue that is addressed by a more modern WAF.

C. HTTP traffic is not forwarding to HTTPS to decrypt: While not forwarding HTTP to HTTPS is a bad security practice, it doesn't directly explain the lack of visibility from the WAF. The WAF should still be able to log and inspect HTTP traffic (although ideally, all traffic should be HTTPS).

D. Old, vulnerable cipher suites are still being used: Using old cipher suites is a security vulnerability. However, it wouldn't necessarily prevent all traffic from being logged. The WAF could still inspect the encrypted traffic if the cipher suite is supported but weak. It is possible that some clients would not support the suites, but it is unlikely to be the cause of all traffic not being logged.

Authoritative Links:

OWASP (Open Web Application Security Project):https://owasp.org/www-project-web-application-firewall/ (General information on WAFs)
SSL/TLS Handshake Explained:https://www.cloudflare.com/learning/ssl/what-happens-in-a-tls-handshake/ (Explanation of the SSL/TLS handshake)
Certificate Authority (CA) Best Practices:https://security.stackexchange.com/questions/12006/best-practices-for-managing-ssl-certificates (Certificate Management best practices)

In summary, the expired certificate directly prevents the WAF from doing its job of decrypting, inspecting, and logging HTTPS traffic, making it the most plausible cause of the described symptoms.

Question: 8

A security analyst is reviewing the following output:

```
Request URL: http://www.largeworldwidebank.org/../../../etc/password
Request Method: GET
Status Code: 200 OK
Remote Address: 107.240.1.127:443
Content-Length: 1245
Content-Type: text/html
Date: Tue, 03 Nov 2020 19:47:14 GMT
Server: Microsoft-IIS/10.0
X-Powered-By: ASP.NET
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,image/webp,image/apng,*/*;q=0.8
Accept-Encoding: gzip, deflate
Accept-Language: en-US,en;q=0.9
Cache-Control: max-age=0
Connection: keep-alive
Host: www.largeworldwidebank.org/
User-Agent: Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/67.0.3396.87 Safari/537.36
```

Which of the following would BEST mitigate this type of attack?

A. Installing a network firewall
B. Placing a WAF inline
C. Implementing an IDS
D. Deploying a honeypot

Answer: B

Explanation:

Placing a WAF inline

Question: 9

Which of the following terms refers to the delivery of encryption keys to a CASB or a third-party entity?

A. Key sharing
B. Key distribution
C. Key recovery
D. Key escrow

Answer: B

Explanation:

The correct answer is **B. Key distribution**. Here's why:

Key distribution, in the context of encryption and cloud security, refers to the process of securely delivering encryption keys to authorized entities. When dealing with Cloud Access Security Brokers (CASBs) or other third-party entities that need access to encrypted data, key distribution becomes crucial. These entities often require the keys to perform functions such as data loss prevention (DLP), threat detection, or compliance monitoring on data stored in the cloud. Secure key distribution mechanisms are essential to prevent unauthorized access to the encryption keys themselves, which would compromise the entire encryption scheme.

Key sharing, while conceptually related, typically implies a more generalized sharing of keys between parties.
Key recovery refers to the process of retrieving lost or corrupted encryption keys. Key escrow is a specific arrangement where encryption keys are held in safekeeping by a trusted third party, allowing access under certain conditions. While a CASB or third-party entity could be part of a key escrow system, the core concept described in the question is the delivery of keys, which falls under key distribution.

Therefore, the scenario described – delivering encryption keys to a CASB or third-party – is most accurately defined as key distribution because it highlights the action of transferring and making keys available to the specified entities. The emphasis is on the delivery mechanism itself.

Further research can be done on the following topics:

NIST Special Publication 800-57 Part 1 Revision 5, Recommendation for Key Management: Specifically, Section 5.3 discusses key distribution. (https://csrc.nist.gov/publications/detail/sp/800-57-part-1/rev-5/final) **Cloud Security Alliance (CSA) materials on CASBs and key management:** Search for "CASB key management" on the CSA website. (https://cloudsecurityalliance.org/)

Question: 10

An organization is implementing a new identity and access management architecture with the following objectives:☞ Supporting MFA against on-premises infrastructure
☞ Improving the user experience by integrating with SaaS applications
☞ Applying risk-based policies based on location
☞ Performing just-in-time provisioning
Which of the following authentication protocols should the organization implement to support these requirements?

A. Kerberos and TACACS
B. SAML and RADIUS

C. OAuth and OpenID

D. OTP and 802.1X

Answer: B

Explanation:

The correct answer is **B. SAML and RADIUS**. Here's why:

SAML (Security Assertion Markup Language): SAML is a standard for exchanging authentication and authorization data between security domains, making it ideal for integrating with SaaS applications. It facilitates single sign-on (SSO), improving the user experience by allowing users to access multiple applications with one set of credentials.

RADIUS (Remote Authentication Dial-In User Service): RADIUS is a networking protocol used for centralized authentication, authorization, and accounting (AAA). It supports MFA against on-premises infrastructure through integration with MFA servers.

Risk-based policies based on location: Both protocols are easily integrated with geo-location tools and policies via the AAA server, and the central IAM platform.

Just-in-time provisioning: SAML can also support just-in-time provisioning by provisioning user accounts in a SaaS application when a user first attempts to access it through SSO. RADIUS can be configured to provision access based on a user's attributes during the initial authentication request.

Other options are not suitable because:

Kerberos and TACACS: Kerberos is primarily for authenticating users within a local network and does not easily integrate with SaaS applications. TACACS+ is a Cisco proprietary protocol used for network device administration, not user authentication for applications.

OAuth and OpenID: OAuth is primarily for authorization, allowing users to grant limited access to their resources without sharing their credentials. OpenID Connect is an authentication layer on top of OAuth 2.0, but it's primarily focused on web-based applications and not well-suited for on-premises infrastructure. **OTP and 802.1X:** OTP (One-Time Password) is a form of MFA but doesn't handle the broader authentication and authorization requirements. 802.1X is a network access control protocol, primarily used for wired and wireless network access control.

Authoritative Links:

SAML:https://saml-info.com/
RADIUS:https://datatracker.ietf.org/doc/html/rfc2865

Question: 11

Which of the following allows computation and analysis of data within a ciphertext without knowledge of the plaintext?

A. Lattice-based cryptography

B. Quantum computing

C. Asymmetric cryptography

D. Homomorphic encryption

Answer: D

Explanation:

The correct answer is D. Homomorphic encryption allows computation on ciphertext without decryption. This

means data can be processed while encrypted, preserving privacy and security. The computation yields an encrypted result which, when decrypted, matches the result of the computation performed on the plaintext.

Lattice-based cryptography (A) is a family of cryptographic primitives based on lattices and is not directly related to computation on encrypted data. Quantum computing (B) is a type of computing using quantum mechanics, and while potentially breaking some encryption algorithms, it does not inherently allow for computation on encrypted data. Asymmetric cryptography (C), using key pairs for encryption and decryption, doesn't enable computations directly on encrypted data. Instead, one key encrypts while the other decrypts. Homomorphic encryption's unique property is incredibly valuable in cloud computing where a third party may process data; it allows this processing to occur without exposing the sensitive data itself. This is particularly useful for regulated industries like healthcare and finance where data privacy is paramount. Different types of homomorphic encryption (fully, somewhat, partially) offer varying degrees of computational flexibility.

Further research:

Homomorphic Encryption Explained Simply: https://www.trendmicro.com/vinfo/us/security-news/data-privacy/homomorphic-encryption-explained
Introduction to Homomorphic Encryption: https://www.ibm.com/blogs/research/introduction-homomorphic-encryption/

Question: 12

A company is looking to fortify its cybersecurity defenses and is focusing on its network infrastructure. The solution cannot affect the availability of the company's services to ensure false positives do not drop legitimate traffic.
Which of the following would satisfy the requirement?

A. NIDS
B. NIPS
C. WAF
D. Reverse proxy

Answer: A

Explanation:

The correct answer is A (NIDS - Network Intrusion Detection System). The key requirement is to fortify cybersecurity defenses without affecting availability and avoiding the dropping of legitimate traffic. An NIDS passively monitors network traffic for suspicious activity. It analyzes traffic patterns, signatures, and anomalies and generates alerts when potential threats are detected. Crucially, an NIDS does not actively block or interfere with network traffic. This aligns perfectly with the requirement of not impacting availability or dropping legitimate traffic. Its passive nature allows for threat identification and alerting without interrupting normal network operations.

In contrast, a Network Intrusion Prevention System (NIPS) actively blocks or modifies malicious traffic. While it offers more immediate protection, it carries the risk of false positives, potentially blocking legitimate traffic and impacting availability. This violates the company's requirements. A Web Application Firewall (WAF) protects web applications specifically and is not a general network defense mechanism. A reverse proxy acts as an intermediary between clients and servers, primarily used for load balancing, caching, and security enhancements like hiding the origin servers. While beneficial for security, a reverse proxy is not specifically designed for the real-time detection of network intrusions in the way an NIDS is. Therefore, an NIDS is the most suitable solution that satisfies the constraints of enhancing security monitoring without impacting the availability of services or causing false positives that disrupt legitimate user traffic.

Further reading:

NIDS vs NIPS:https://www.cloudflare.com/learning/security/threats/intrusion-detection-system-ids/ **Network Intrusion Detection Systems (NIDS):**https://www.varonis.com/blog/intrusion-detection-system

Question: 13

A disaster recovery team learned of several mistakes that were made during the last disaster recovery parallel test. Computational resources ran out at 70% of restoration of critical services.
Which of the following should be modified to prevent the issue from reoccurring?

A. Recovery point objective
B. Recovery time objective
C. Mission-essential functions
D. Recovery service level

Answer: D

Explanation:

The correct answer is **D. Recovery service level**.

The scenario highlights a failure in the parallel disaster recovery test where computational resources were exhausted before all critical services were restored. This indicates a discrepancy between the planned and actual capacity needed to meet the desired service levels during recovery.

A **Recovery Service Level (RSL)** defines the acceptable performance and availability of services after a disaster. It outlines the resources, capacity, and functionality that must be restored to meet business requirements. The fact that computational resources ran out at 70% restoration means the RSL was not properly defined or provisioned for in the disaster recovery plan. The DR plan assumed sufficient computing power, which turned out not to be the case in the parallel test. Adjusting the RSL would involve re-evaluating the resource requirements for critical services during recovery and ensuring sufficient capacity is available.

This might involve increasing the allocated computational resources, optimizing restoration processes, or prioritizing services differently.

Recovery Point Objective (RPO) is the acceptable amount of data loss measured in time.
Recovery Time Objective (RTO) is the acceptable amount of time to restore a service.
Mission-essential functions (MEF) identifies the core processes necessary for the organization's survival.

While RPO, RTO, and MEF are all crucial components of disaster recovery, they do not directly address the issue of insufficient computational resources. The RPO defines how far back the data should be recoverable to. The RTO specifies how long it takes to get a service back up and running. Mission essential functions are the critical services that must be available. These parameters do not detail the exact infrastructure
requirements necessary to reach the RTO, therefore focusing directly on the performance and resource needs of the restored services is key, which the RSL represents. Modifying the RSL to account for the resource shortfall would be a direct solution to the problem.

Therefore, the RSL should be modified to accurately reflect the necessary resources required for successful disaster recovery.

Further Reading:

NIST Special Publication 800-34 - Contingency Planning Guide for Federal Information Systems:
https://csrc.nist.gov/publications/detail/sp/800-34/rev-1/archive/2010-05-06
ISO 22301 - Business Continuity Management Systems:https://www.iso.org/iso-22301-business-continuity.html (Paywalled, but generally accepted as a BCMS standard)

Question: 14

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 send traffic to those sites.
The technician will define this threat as:

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

B. a zero-day attack.

C. an advanced persistent threat.

D. an on-path attack.

Answer: C

Explanation:

The correct answer is C, an advanced persistent threat (APT). Here's why:

Prolonged Activity: The file transfer occurring over three months suggests a sustained, long-term campaign, a hallmark of APTs. APTs don't typically execute and vanish; they aim to maintain access for extended periods.

Stealth and Evasion: The use of TLS-protected HTTP (HTTPS) makes detection more difficult. APTs often camouflage their activities using standard protocols to blend in with normal network traffic.

Unusual Communication Patterns: Systems not normally communicating with the remote sites sending data there is suspicious. APTs often compromise internal systems and use them as launchpads for further attacks or data exfiltration.

Data Exfiltration: The large number of files transferred points to data theft as a primary objective. This is a common goal of APTs: to steal sensitive information.

Persistence: The activity ceasing doesn't negate an APT. It could indicate the attackers achieved their initial objective and are lying dormant, preparing for a future phase, or that the compromised host was cleaned. APTs often involve multiple stages.

Why other options are less likely:

A. Decrypting RSA using obsolete and weakened encryption attack: While possible, this describes a specific vulnerability and attack technique. The scenario describes activity, which could stem from many root causes beyond this one specific vulnerability. Furthermore, this attack primarily deals with eavesdropping and decryption, not the sustained exfiltration of data.

B. Zero-day attack: While APTs might utilize zero-day vulnerabilities (exploits for which no patch exists), a zero-day attack itself is just the initial exploit. The scenario describes prolonged activity after a potential initial compromise. A zero-day explains initial access, not the sustained exfiltration.

D. On-path attack (Man-in-the-Middle): This involves intercepting and potentially altering communication between two parties. While possible, the use of TLS (HTTPS) makes this attack significantly more difficult. Also, the scenario emphasizes the source and destination of traffic as the main indicators, not interception en route.

Authoritative Links for Further Research:

NIST (National Institute of Standards and Technology):
(https://csrc.nist.gov/glossary/term/advanced_persistent_threat)
MITRE ATT&CK Framework: (https://attack.mitre.org/) (A knowledge base of adversary tactics and techniques based on real-world observations)
SANS Institute: (https://www.sans.org/) (Provides cybersecurity training and resources)

Question: 15

A security engineer thinks the development team has been hard-coding sensitive environment variables in its code. Which of the following would BEST secure the company's CI/CD pipeline?

A. Utilizing a trusted secrets manager
B. Performing DAST on a weekly basis
C. Introducing the use of container orchestration
D. Deploying instance tagging

Answer: A

Explanation:

The best solution to avoid hard-coding sensitive environment variables is to use a trusted secrets manager (Option A). Secrets managers are designed to securely store, manage, and access sensitive information like API keys, passwords, database connection strings, and other credentials required by applications. By integrating a secrets manager into the CI/CD pipeline, the development team can retrieve these variables at runtime without embedding them directly in the code or configuration files. This significantly reduces the risk of exposing sensitive data through version control systems, deployment artifacts, or during a security breach.

Option B, performing Dynamic Application Security Testing (DAST), is beneficial for identifying vulnerabilities during runtime, but it primarily focuses on external attack vectors and may not detect hard-coded secrets reliably. It's a valuable tool for overall security but not the primary solution for this problem. Option C, container orchestration, helps manage and deploy containerized applications but doesn't inherently address the issue of secret management. While some orchestration platforms offer secret management features, using a dedicated secrets manager is generally more secure and flexible. Option D, deploying instance tagging, is useful for resource management and identification but has no bearing on the security of sensitive data within the code.

Therefore, a secrets manager provides a centralized, secure, and auditable way to handle sensitive information within the CI/CD pipeline. Examples include HashiCorp Vault, AWS Secrets Manager, Azure Key Vault, and Google Cloud Secret Manager. These solutions offer features like encryption at rest and in transit, access control policies, rotation of secrets, and audit logging. Integrating a secrets manager is a best practice for modern application development and deployment, ensuring a more secure and robust CI/CD pipeline.

Further Reading:

OWASP Secrets Management Cheat Sheet:
https://cheatsheetseries.owasp.org/cheatsheets/Secrets_Management_Cheat_Sheet.html
CNCF (Cloud Native Computing Foundation) Security Best Practices: Secrets Management: (Search for relevant resources on the CNCF website regarding secrets management in cloud-native environments)

Question: 16

A small company recently developed prototype technology for a military program. The company's security engineer is concerned about potential theft of the newly developed, proprietary information.
Which of the following should the security engineer do to BEST manage the threats proactively?

A. Join an information-sharing community that is relevant to the company.
B. Leverage the MITRE ATT&CK framework to map the TTP.
C. Use OSINT techniques to evaluate and analyze the threats.
D. Update security awareness training to address new threats, such as best practices for data security.

Answer: B

Explanation:

The best proactive measure to manage potential theft of proprietary information related to a military program is to leverage the MITRE ATT&CK framework to map threat actor tactics, techniques, and procedures (TTPs). Understanding how adversaries operate allows the security engineer to anticipate and prepare for potential attacks targeting the company's sensitive data.

Option B is the best approach because it emphasizes a structured, proactive defense strategy. The MITRE ATT&CK framework is a knowledge base of adversary tactics and techniques based on real-world observations. Mapping potential threats to the framework helps identify vulnerabilities and develop specific countermeasures. This proactive approach allows the company to anticipate and prepare for threats rather than reacting to them.

Option A, joining an information-sharing community, can be helpful but is reactive. The information received is dependent on what others share, which may or may not be timely or relevant.

Option C, using OSINT, provides valuable context on potential threat actors. However, OSINT alone doesn't provide a structured framework for understanding how those threat actors might operate or how to defend against them. It's more useful for identifying potential threats than understanding their methods.

Option D, updating security awareness training, is important, but it's more general and less targeted than mapping TTPs. While training on data security best practices is valuable, it's not as effective as understanding specific adversary techniques. It is a general safeguard that may or may not be directly relevant to the specific threats this company faces.

The MITRE ATT&CK framework (https://attack.mitre.org/) provides a comprehensive understanding of adversary behavior, enabling a proactive and targeted security strategy. By mapping TTPs, the security engineer can identify specific vulnerabilities and implement corresponding security controls, significantly reducing the risk of data theft. Understanding TTPs enables implementing proactive security measures to block and detect malicious activity.

Question: 17

A security engineer has been asked to close all non-secure connections from the corporate network. The engineer is attempting to understand why the corporate
UTM will not allow users to download email via IMAPS. The engineer formulates a theory and begins testing by creating the firewall ID 58, and users are able to download emails correctly by using IMAP instead. The network comprises three VLANs:

The security engineer looks at the UTM firewall rules and finds the following:

- VLAN 30	Guest networks	192.168.20.0/25
- VLAN 20	Corporate user network	192.168.0.0/28
- VLAN 110	Corporate server network	192.168.0.16/29

Rule active	Firewall ID	Source	Destination	Ports	Action	TLS decryption
Yes	58	VLAN 20	15.22.33.45	143	Allow and log	Enabled
Yes	33	VLAN 30	Any	80, 443,	Allow and log	Disabled
Yes	22	VLAN 110	VLAN 20	Any	Allow and log	Disabled
No	21	VLAN 20	15.22.33.45	990	Allow and log	Disabled
Yes	20	VLAN 20	VLAN 110	Any	Allow and log	Enabled
Yes	19	VLAN 20	Any	993, 587	Allow and log	Enabled

Which of the following should the security engineer do to ensure IMAPS functions properly on the corporate user network?

A. Contact the email service provider and ask if the company IP is blocked. B. Confirm the email server certificate is installed on the corporate computers.

C. Make sure the UTM certificate is imported on the corporate computers. D. Create an IMAPS firewall rule to ensure email is allowed.

Answer: B

Explanation:

Confirm the email server certificate is installed on the corporate computers.

Question: 18

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:

The best command to view only active Internet TCP connections on a Linux workstation is sudo netstat -pnut | grep -P ^tcp. Here's why:

netstat: This is a command-line tool for displaying network connections, routing tables, interface statistics, masquerade connections, and multicast memberships.

-p: This option shows the PID (Process ID) and name of the program to which each socket belongs. This helps in identifying which application is using the connection.

-n: This option displays numerical addresses instead of trying to determine symbolic host names, which is faster.

-u: This option shows UDP connections, which is not relevant as the question specifically asks for TCP connections. Including this option causes the command to include UDP in the results.

-t: This option shows TCP connections. This is crucial for answering the question.

grep -P ^tcp: This filters the output to show only lines that begin with tcp. The -P option enables Perl-compatible regular expressions, where ^ signifies the beginning of the line. This is a precise way to isolate TCP connections from other types of connections netstat may display. This effectively filters the output to show only established TCP connections, as opposed to those in other states like LISTEN. Other states may include TCP connection that are not fully established or are in the process of closing, which the question doesn't need.

Option A is incorrect because it focuses on connections in the LISTEN state, which are passively waiting for incoming connections rather than active Internet connections. It also lacks the specificity to filter by protocol effectively.

Option B is incorrect because it uses -l, which only shows listening sockets. While it includes established connections via the ESTABLISHED filter, it doesn't focus exclusively on active Internet TCP connections.

Option C is incorrect because using grep -v Foreign Address does not accurately filter for active internet connections. This attempts to exclude local addresses which may not be completely accurate.

Option D is incorrect because using -w option displays raw socket information and may not properly address the need to isolate internet connections. Using column -t -s

#39;\w' attempts to format the output, but it does not correctly filter for the right type of connections.

In summary, the chosen command provides a clear and concise way to identify established TCP connections by showing the process IDs, numerical addresses, and filtering the results to exclusively target TCP protocols.

Refer to these resources for more information:

netstat manual: man netstat in a Linux terminal.
grep manual: man grep in a Linux terminal.
TCP/IP Protocol Suite: https://www.rfc-editor.org/rfc/rfc793

Question: 19

A shipping company that is trying to eliminate entire classes of threats is developing an SELinux policy to ensure its custom Android devices are used exclusively for package tracking.
After compiling and implementing the policy, in which of the following modes must the company ensure the devices are configured to run?

A. Protecting
B. Permissive
C. Enforcing
D. Mandatory

Answer: B

Explanation:

The correct answer is **B. Permissive**. Here's why:

The shipping company is developing an SELinux policy to lock down their Android devices for package tracking. Before fully deploying such a restrictive policy, it's crucial to understand its impact on the system. The permissive mode is designed for exactly this purpose. In permissive mode, SELinux logs policy violations but does not block actions that would be denied by the policy. This allows administrators to see what the

impact of the policy would be without actually breaking the system or hindering essential operations. Logs are essential for assessing whether a policy is too restrictive or contains errors.

Enforcing mode would be the desired final state, but jumping directly to enforcing without testing could lead to device instability, application failures, or rendering the devices unusable if the policy is too aggressive or contains mistakes. The goal of the company is to eliminate entire classes of threats, so an eventual move to enforcing mode is inevitable. However, rushing into it is not advised, as it could disrupt core business functions.

Protecting and Mandatory are not standard SELinux modes.

Therefore, the permissive mode allows the shipping company to test the SELinux policy, identify potential issues, and refine it before deploying it in enforcing mode for actual threat mitigation. The iterative process of policy development, permissive testing, policy refinement, and eventual enforcement ensures the system is both secure and functional.

Further Reading:

SELinux Documentation: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/using_selinux/index (General SELinux documentation)
Android SELinux: https://source.android.com/docs/security/selinux (SELinux implementation on Android)

Question: 20

A security analyst receives an alert from the SIEM regarding unusual activity on an authorized public SSH jump server. To further investigate, the analyst pulls the event logs directly from /var/log/auth.log:
graphic.ssh_auth_log.
Which of the following actions would BEST address the potential risks posed by the activity in the logs?

A. Altering the misconfigured service account password
B. Modifying the AllowUsers configuration directive
C. Restricting external port 22 access
D. Implementing host-key preferences

Answer: C

Explanation:

The best response is **C. Restricting external port 22 access.** Here's a detailed justification:

The auth.log entries indicate multiple failed SSH login attempts originating from various external IP addresses targeting valid users. This pattern strongly suggests a brute-force attack targeting SSH, a common tactic for gaining unauthorized access. While addressing the other options could provide some benefit, they don't directly mitigate the immediate threat.

Option A (Altering the misconfigured service account password) is beneficial in general to improve security hygiene but does not address the brute force attempt on other valid users. Service accounts are not explicitly mentioned in the event logs. Option B (Modifying the AllowUsers configuration directive) could help by limiting which users can log in via SSH, adding a layer of defense. However, it doesn't stop the attacks targeting the allowed users. Option D (Implementing host-key preferences) is not mentioned in the given logs and would be more relevant for preventing man-in-the-middle attacks, which isn't the primary concern here.

Restricting external port 22 access (Option C) directly addresses the root cause by limiting the attack surface. This can be done by allowing only authorized IP addresses or networks to connect to the SSH server (jump server) on port 22, effectively blocking brute-force attempts from unwanted sources. This reduces the

possibility of a successful breach by significantly reducing the available vectors for attack. Limiting SSH access to a specific IP range using a firewall rule, for example, makes it significantly more difficult for attackers to reach the SSH service in the first place. This approach aligns with the principle of least privilege and defense in depth.

Further research on hardening SSH servers and best practices can be found at:

NIST Special Publication 800-123, Guide to Securing Apple OS X 10.9 Systems for IT Professionals: (Although focused on macOS, principles of SSH hardening apply broadly):
https://csrc.nist.gov/publications/detail/sp/800-123/final
SANS Institute: Search for "hardening SSH" on their website (https://www.sans.org/) for various articles and resources.

Question: 21

A high-severity vulnerability was found on a web application and introduced to the enterprise. The vulnerability could allow an unauthorized user to utilize an open- source library to view privileged user information. The enterprise is unwilling to accept the risk, but the developers cannot fix the issue right away.
Which of the following should be implemented to reduce the risk to an acceptable level until the issue can be fixed?

A. Scan the code with a static code analyzer, change privileged user passwords, and provide security training.

B. Change privileged usernames, review the OS logs, and deploy hardware tokens.

C. Implement MFA, review the application logs, and deploy a WAF.

D. Deploy a VPN, configure an official open-source library repository, and perform a full application review for vulnerabilities.

Answer: C

Explanation:

The correct answer is C because it directly addresses the immediate risk posed by the vulnerability. Here's why:

Implement MFA (Multi-Factor Authentication): MFA adds an extra layer of security, making it harder for unauthorized users to exploit the vulnerability, even if they gain access through the open-source library. This mitigates the risk of unauthorized access to privileged user accounts. (Source: https://www.cisa.gov/mfa)

Review Application Logs: Application logs can provide valuable insights into whether the vulnerability is being exploited. Monitoring logs for suspicious activity related to the open-source library can help detect and respond to attacks in real-time. This is a crucial incident response step. (Source: https://owasp.org/www-project-top-ten/ - A10:2021-Security Logging Failure)

Deploy a WAF (Web Application Firewall): A WAF acts as a shield between the web application and the internet. It can be configured with rules to detect and block attempts to exploit the known vulnerability in the open-source library. This provides an immediate protective measure without requiring code changes. (Source: https://www.cloudflare.com/learning/cloud/what-is-a-web-application-firewall-waf/)

Other options are less effective as immediate risk mitigation:

A: Static code analysis is good for finding vulnerabilities but doesn't provide immediate protection. Changing passwords alone isn't sufficient if the application is vulnerable.

B: Changing usernames and deploying hardware tokens are useful security measures, but they don't specifically address the vulnerability in the open-source library.

D: Deploying a VPN and configuring an open-source repository don't directly address the vulnerability's

exploitability. A full application review takes time and doesn't offer immediate protection.

Question: 22

A security analyst discovered that the company's WAF was not properly configured. The main web server was breached, and the following payload was found in one of the malicious requests:

```
<!DOCTYPE doc [
<!ELEMENT doc ANY>
<!ENTITY xxe SYSTEM "file:///etc/password" >]>
<doc>&xxe;</doc>
```

Which of the following would BEST mitigate this vulnerability?

A. CAPTCHA
B. Input validation
C. Data encoding
D. Network intrusion prevention

Answer: B

Explanation:

Reference:
https://hdivsecurity.com/owasp-xml-external-entities-xxe
" target="_blank" style="word-break: break-all;">

Example #1: The attacker attempts to extract data from the server

```
<?xml version="1.0" encoding="ISO-8859-1"?> <!DOCTYPE foo [
<!ELEMENT foo ANY >
<!ENTITY xxe SYSTEM "file:///etc/passwd" >]> <foo>&xxe;</foo>
```

Example #2: An attacker probes the server's private network by changing the above ENTITY line to

```
<!ENTITY xxe SYSTEM "https://192.168.1.1/private" >]>
```

Question: 23

A university issues badges through a homegrown identity management system to all staff and students. Each week during the summer, temporary summer school students arrive and need to be issued a badge to access minimal campus resources. The security team received a report from an outside auditor indicating the homegrown system is not consistent with best practices in the security field and leaves the institution vulnerable.
Which of the following should the security team recommend FIRST?

A. Investigating a potential threat identified in logs related to the identity management system
B. Updating the identity management system to use discretionary access control
C. Beginning research on two-factor authentication to later introduce into the identity management system D. Working with procurement and creating a requirements document to select a new IAM system/vendor

Answer: B

Explanation:

The correct answer is **B. Updating the identity management system to use discretionary access control.**

Here's the justification:

The core issue is that the current homegrown identity management (IdM) system is vulnerable and doesn't align with security best practices. While options A, C, and D address aspects of security, they don't directly tackle the immediate vulnerability identified by the auditor. Option A is reactive and addresses a specific incident, not the underlying flaw. Option C is proactive but deals with a longer-term enhancement. Option D is a long-term solution but takes time, potentially leaving the institution vulnerable in the interim.

Discretionary Access Control (DAC) gives resource owners the ability to control access to their resources. This directly improves the security posture by ensuring users only get access to the resources they need. Implementing DAC allows for granular control over access rights, mitigating potential misuse and reducing the attack surface. It enables the university to define specific permissions for temporary summer students, limiting their access to minimal campus resources as needed.

Addressing the security concerns in the existing system before implementing longer-term improvements is crucial. By applying access control, you are improving your security posture and reducing the risk of unauthorized access by students or external threat actors using the system.

Therefore, updating the identity management system to implement DAC is the most immediate and impactful step to address the identified vulnerability and align with security best practices.

Further Research:

NIST Special Publication 800-162 - Guide to Attribute-Based Access Control (ABAC) Definition and Considerations: While not DAC, this is relevant as it discusses access control principles in depth.
https://csrc.nist.gov/publications/detail/sp/800-162/final
OWASP Access Control:https://owasp.org/www-project-top-ten/

Question: 24

A customer reports being unable to connect to a website at www.test.com to consume services. The customer notices the web application has the following published cipher suite:

```
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
Signature hash algorithm:
sha256
Public key:
RSA (2048 Bits)
.htaccess config:
<VirtualHost> *:80>
ServerName www.test.com
Redirect / https://www.test.com
</VirtualHost>
<VirtualHost _default_:443>
ServerName www.test.com
DocumnetRoot /usr/local/apache2/htdocs
SSLEngine On
...
</VirtualHost>
```

Which of the following is the MOST likely cause of the customer's inability to connect?

A. Weak ciphers are being used.
B. The public key should be using ECDSA.
C. The default should be on port 80.
D. The server name should be test.com.

Answer: B

Explanation:

Reference:
https://security.stackexchange.com/questions/23383/ssh-key-type-rsa-dsa-ecdsa-are-there-easy-answers-f or-which-to-choose-when

Question: 25

An IT administrator is reviewing all the servers in an organization and notices that a server is missing crucial practice against a recent exploit that could gain root access.
Which of the following describes the administrator's discovery?

A. A vulnerability
B. A threat
C. A breach
D. A risk

Answer: A

Explanation:

The correct answer is A. A vulnerability.

Here's why:

A **vulnerability** is a weakness or flaw in a system's hardware, software, or procedures that can be exploited by a threat actor. In this scenario, the missing patch exposes a weakness (the unpatched server) that can be leveraged by the exploit to gain root access. The description directly indicates a potential point of weakness.

A **threat** is a potential danger that might exploit a vulnerability. While the exploit represents a threat, the missing patch itself isn't the threat; it's the vulnerability that the threat seeks to exploit. The question describes something that allows a threat to succeed.

A **breach** is an actual security incident where a system or data is compromised. The description only mentions a potential exploit and a missing patch; there's no indication that the server has already been compromised. It's a potential problem, not an existing one.

A **risk** is the potential for loss or damage resulting from a vulnerability and a threat. While a risk is present in this scenario, the question asks about the specific discovery of the IT administrator, which is the missing patch, representing a vulnerability. Risk combines vulnerability and threat to estimate impact.

The key is that the administrator discovered a missing security patch. This means that server lacks the proper defense against a known exploit, which corresponds precisely to a vulnerability. The missing patch is a specific, identifiable weakness.

Here are some authoritative links for further research:

NIST Definition of Vulnerability:https://csrc.nist.gov/glossary/term/vulnerability

SANS Institute on Vulnerability Management:https://www.sans.org/information-security/glossary/vulnerability/

Question: 26

A security analyst is performing a vulnerability assessment on behalf of a client. The analyst must define what constitutes a risk to the organization.
Which of the following should be the analyst's FIRST action?

A. Create a full inventory of information and data assets.
B. Ascertain the impact of an attack on the availability of crucial resources.
C. Determine which security compliance standards should be followed.
D. Perform a full system penetration test to determine the vulnerabilities.

Answer: A

Explanation:

The correct first step for a security analyst performing a vulnerability assessment and defining risk is **A. Create a full inventory of information and data assets.**

Before any meaningful risk assessment can be conducted, it is paramount to understand what assets are at risk. An asset inventory provides a comprehensive list of all information and data assets, including servers, databases, applications, network devices, endpoints, and sensitive data. This inventory serves as the foundation for identifying potential vulnerabilities and assessing the impact of their exploitation. Without a clear understanding of what needs to be protected, it is impossible to determine the likelihood and impact of threats accurately.

Option B, ascertaining the impact of an attack on availability, is important but dependent on knowing what crucial resources exist. Option C, determining compliance standards, is also crucial for overall security posture but does not precede the fundamental need to know what assets must comply. Option D, performing a penetration test, is a later step in the process, meant to exploit discovered vulnerabilities once they are identified. A penetration test without a proper asset inventory and vulnerability assessment is undirected and inefficient.

The asset inventory allows the analyst to understand the value, criticality, and sensitivity of each asset. This information is essential for prioritizing vulnerability remediation efforts and allocating security resources effectively. Once the assets are known, then the security analyst can proceed to identify vulnerabilities, assess risks, and develop appropriate mitigation strategies. For example, a database containing customer Personally Identifiable Information (PII) will be considered a higher priority asset than a public-facing web server displaying marketing information. Understanding the assets at stake sets the stage for the vulnerability assessment, including what to scan, what to test, and what standards apply.

In a cloud environment, asset discovery can be more challenging due to the dynamic nature of cloud resources. Automated discovery tools and cloud-native asset management solutions can help ensure a complete and up-to-date inventory.

Resources for further research include:

NIST Special Publication 800-30, Guide for Conducting Risk Assessments:
https://csrc.nist.gov/publications/detail/sp/800-30/rev-1/final
SANS Institute Reading Room - Performing a Risk Assessment:https://www.sans.org/reading-room/whitepapers/auditing/performing-risk-assessment-576

Question: 27

While investigating a security event, an analyst finds evidence that a user opened an email attachment from an unknown source. Shortly after the user opened the attachment, a group of servers experienced a large amount of network and resource activity. Upon investigating the servers, the analyst discovers the servers were encrypted by ransomware that is demanding payment within 48 hours or all data will be destroyed. The company has no response plans for ransomware. Which of the following is the NEXT step the analyst should take after reporting the incident to the management team?

A. Pay the ransom within 48 hours.
B. Isolate the servers to prevent the spread.
C. Notify law enforcement.
D. Request that the affected servers be restored immediately.

Answer: B

Explanation:

Isolating the affected servers is the next crucial step after reporting the ransomware incident to management. This action is paramount in containing the spread of the ransomware to other parts of the network. Ransomware thrives on lateral movement, infecting systems connected to the initially compromised host. By disconnecting the infected servers, the analyst effectively creates a barrier, preventing further encryption and damage. This isolation can involve disconnecting network cables, disabling network adapters, or using network segmentation tools to quarantine the compromised servers. Paying the ransom (option A) is generally discouraged, as it doesn't guarantee data recovery and may encourage further attacks. Notifying law enforcement (option C) is important but can be done concurrently or shortly after isolation. Restoring servers immediately (option D) is a viable long-term solution, but requires careful planning to avoid reinfection if the ransomware is still active on the network. Prioritizing isolation provides immediate containment, buying valuable time to assess the situation, develop a recovery plan, and potentially identify the source of the attack. This aligns with incident response best practices focused on minimizing the impact and preventing further damage.

For further research, consider these resources:

SANS Institute: Incident Handler's Handbook:https://www.sans.org/reading-room/whitepapers/incident/incident-handlers-handbook-33901
NIST Special Publication 800-61 Revision 2, Computer Security Incident Handling Guide: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf

Question: 28

A company plans to build an entirely remote workforce that utilizes a cloud-based infrastructure. The Chief Information Security Officer asks the security engineer to design connectivity to meet the following requirements:☞ Only users with corporate-owned devices can directly access servers hosted by the cloud provider.
☞ The company can control what SaaS applications each individual user can access.
☞ User browser activity can be monitored.
Which of the following solutions would BEST meet these requirements?

A. IAM gateway, MDM, and reverse proxy
B. VPN, CASB, and secure web gateway
C. SSL tunnel, DLP, and host-based firewall
D. API gateway, UEM, and forward proxy

Answer: B

Explanation:

Here's a detailed justification for why option B (VPN, CASB, and secure web gateway) is the best solution, along with explanations of why the other options are less suitable:

Why Option B is the Best Choice:

VPN (Virtual Private Network): A VPN ensures that only users connecting from corporate-owned devices can directly access servers hosted by the cloud provider. It establishes an encrypted tunnel, verifying the device's identity and authorizing access to the cloud infrastructure. This addresses the first requirement, restricting server access to corporate assets.

CASB (Cloud Access Security Broker): A CASB is a crucial component for controlling which SaaS applications individual users can access. It sits between users and cloud services, providing visibility and control over cloud application usage. CASBs offer granular control, allowing administrators to specify which SaaS applications each user or group can access, satisfying the second requirement. (Source: https://www.gartner.com/en/information-technology/glossary/cloud-access-security-brokers-casbs)

Secure Web Gateway (SWG): A secure web gateway monitors user browser activity. It filters malicious content, enforces acceptable use policies, and provides visibility into user browsing habits. SWGs inspect web traffic, block access to risky websites, and record browsing activity, meeting the third requirement of monitoring browser activity. (Source: https://www.forcepoint.com/cyber-edu/secure-web-gateway)

Why Other Options Are Less Suitable:

A. IAM gateway, MDM, and reverse proxy: IAM (Identity and Access Management) gateways primarily manage user authentication and authorization but don't directly enforce device restrictions or control SaaS application access as effectively as a CASB. MDM (Mobile Device Management) focuses on managing mobile devices but doesn't inherently monitor browser activity in the same way a SWG does. Reverse proxies are typically used for load balancing and protecting web servers, not for controlling SaaS application access.

C. SSL tunnel, DLP, and host-based firewall: SSL tunnels provide encryption, but they don't offer granular control over SaaS applications or comprehensive browser activity monitoring. DLP (Data Loss Prevention) focuses on preventing data exfiltration but doesn't provide the same level of SaaS application control as a CASB. Host-based firewalls protect individual systems but don't provide centralized control over user access to cloud services.

D. API gateway, UEM, and forward proxy: API gateways manage access to APIs but don't directly address the device access restriction requirement or the SaaS application control requirement. UEM (Unified Endpoint Management) combines MDM and traditional PC management but still lacks the SaaS application control capabilities of a CASB. Forward proxies route user traffic through a central server, but they don't inherently provide the same level of security and monitoring as a SWG.

In summary, the VPN ensures secure and controlled access from corporate devices, the CASB provides granular control over SaaS application access, and the secure web gateway offers comprehensive monitoring of user browser activity. This combination of solutions most effectively meets the requirements outlined in the scenario.

Question: 29

During a system penetration test, a security engineer successfully gained access to a shell on a Linux host as a standard user and wants to elevate the privilege levels.
Which of the following is a valid Linux post-exploitation method to use to accomplish this goal?

A. Spawn a shell using sudo and an escape string such as sudo vim -c '!sh'.

B. Perform ASIC password cracking on the host.
C. Read the /etc/passwd file to extract the usernames.
D. Initiate unquoted service path exploits.
E. Use the UNION operator to extract the database schema.

Answer: A

Explanation:

The correct answer is A: "Spawn a shell using sudo and an escape string such as sudo vim -c '!sh'." This method directly addresses privilege escalation on a Linux system. The sudo command allows a permitted user to execute a command as the superuser (root) or another user, as specified by the security policy. By exploiting a misconfiguration or vulnerability in the sudo configuration or the program being executed with sudo, a standard user can gain root privileges.

Specifically, sudo vim -c '!sh' uses the vim text editor, often available on Linux systems, and leverages its command-line mode (-c) to execute the shell command !sh. If vim is configured in sudoers to allow execution without a password for the current user (which is a common misconfiguration or might be exploited via other means), the shell spawned will inherit the elevated privileges of sudo, thus granting the user root access. This exploit takes advantage of the interaction between sudo's trust in the given program (vim in this case) and the program's ability to execute shell commands.

Option B is incorrect because ASIC password cracking refers to cracking passwords using specialized hardware, which wouldn't be performed directly on the host. Option C, while relevant for user enumeration, doesn't directly lead to privilege escalation as /etc/passwd only contains usernames and user IDs, not password hashes. Option D (unquoted service path exploits) is primarily a Windows vulnerability and isn't applicable to Linux. Option E (UNION operator for database schema extraction) is a technique used in SQL injection attacks and doesn't directly escalate local privileges on a Linux system. Therefore, leveraging sudo and escape strings is a known and valid Linux post-exploitation technique for privilege escalation.

Further Research:

Sudo Manual:man sudo on any Linux terminal
GTFOBins (for vim/sudo abuse):https://gtfobins.github.io/gtfobins/vim/
Privilege Escalation Techniques: Search for "Linux privilege escalation techniques" to find various tutorials and write-ups.

Question: 30

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:

Here's a detailed justification of why UEFI/BIOS is the best answer for boot loader protection, along with relevant concepts and links.

The best choice is **D. UEFI/BIOS**. Modern Unified Extensible Firmware Interface (UEFI) and BIOS (its predecessor, though less secure) provide various mechanisms to protect the boot loader. Secure Boot, a feature of UEFI, is specifically designed to prevent unauthorized software from loading during the boot process. Secure Boot uses cryptographic signatures to verify the integrity of the boot loader, operating system kernel, and other critical system components before they are allowed to execute. If any component's signature is invalid or missing, the system will refuse to boot, preventing the execution of potentially malicious code. This directly addresses the need to secure the boot process before the operating system fully loads. Furthermore, UEFI can implement password protection to prevent unauthorized modifications to boot settings.

Let's analyze why the other options are less suitable in this specific context:

A. TPM (Trusted Platform Module): A TPM is a hardware security module that provides secure storage for cryptographic keys and can perform cryptographic operations. While TPMs can enhance overall system security, they don't directly prevent unauthorized boot loaders. They're more focused on disk encryption key protection and integrity measurements of the boot process, which is useful, but not a direct preventive measure.

B. HSM (Hardware Security Module): An HSM is a dedicated hardware device used to securely manage, process, and store cryptographic keys. HSMs are typically used in high-security environments and primarily handle cryptographic functions. They do not directly manage or protect the boot loader itself.

C. PKI (Public Key Infrastructure): PKI provides a framework for managing digital certificates and public keys. While Secure Boot uses PKI to verify signatures, PKI itself is the infrastructure, not the direct mechanism protecting the boot loader. UEFI Secure Boot is the feature that implements PKI principles for boot loader protection.

In summary, UEFI/BIOS with features like Secure Boot directly provides a security layer for the boot loader by verifying its integrity and preventing the execution of unauthorized code before the operating system loads. The other options are valuable security components but don't offer the same direct boot loader protection as UEFI Secure Boot.

Authoritative Links:

UEFI Forum:https://uefi.org/ (Official UEFI specifications and information)
Microsoft - Secure Boot:https://learn.microsoft.com/en-us/windows-hardware/design/device-experiences/oem-secure-boot (Microsoft's explanation of Secure Boot)
NIST Special Publication 800-147B:https://csrc.nist.gov/publications/detail/sp/800-147b/final (BIOS Integrity Measurement Guidelines)

Question: 31

A developer is creating a new mobile application for a company. The application uses REST API and TLS 1.2 to communicate securely with the external back-end server. Due to this configuration, the company is concerned about HTTPS interception attacks.
Which of the following would be the BEST solution against this type of attack?

A. Cookies
B. Wildcard certificates
C. HSTS
D. Certificate pinning

Answer: D

Explanation:

The best defense against HTTPS interception attacks, given the scenario of a mobile application using REST API and TLS 1.2, is **D. Certificate Pinning**. Here's why:

HTTPS interception attacks, also known as man-in-the-middle (MITM) attacks, occur when an attacker intercepts the secure communication between the mobile app and the back-end server. The attacker substitutes the real server's certificate with their own, often achieved by installing a rogue certificate authority (CA) on the user's device or network. This allows the attacker to decrypt and inspect the traffic.

Certificate pinning mitigates this risk by embedding the expected server certificate (or its public key hash) directly into the mobile application code. The app then validates that the server's certificate presented during the TLS handshake exactly matches the pinned certificate. If the certificates don't match, the app refuses the connection, preventing the MITM attack. This removes the reliance on trusting all CAs trusted by the operating system.

HSTS (HTTP Strict Transport Security), option C, primarily focuses on preventing protocol downgrade attacks by instructing the browser to always use HTTPS for a given domain. While HSTS enhances security, it's a mechanism primarily intended for web browsers and does not provide a direct defense against an application being tricked into accepting a rogue certificate, which is the core concern in HTTPS interception attacks. It requires the client to have previously visited the site over HTTPS to establish the HSTS policy. Therefore, it does not provide initial protection.

Cookies, option A, are small pieces of data stored on the user's device to track user activity and maintain session state. They are not directly related to certificate validation or preventing MITM attacks.

Wildcard certificates, option B, cover multiple subdomains with a single certificate (e.g., *.example.com). While useful for simplifying certificate management, they do not inherently prevent HTTPS interception attacks. An attacker could still intercept the connection and present a valid wildcard certificate issued by a rogue CA if the application doesn't perform proper certificate validation.

Certificate pinning provides a more robust and direct defense against HTTPS interception by enforcing strict validation of the server's identity, regardless of the CAs trusted by the device.

Authoritative resources:

OWASP Certificate and Public Key Pinning: https://owasp.org/www-project-top-ten/2017/A7_2017-Insufficient_Attack_Protection
Android Developers - Trusting Certificate Authorities: https://developer.android.com/training/articles/security-config#TrustingCertificates

Question: 32

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 simulation, please click the Reset All button.
Select and Place:

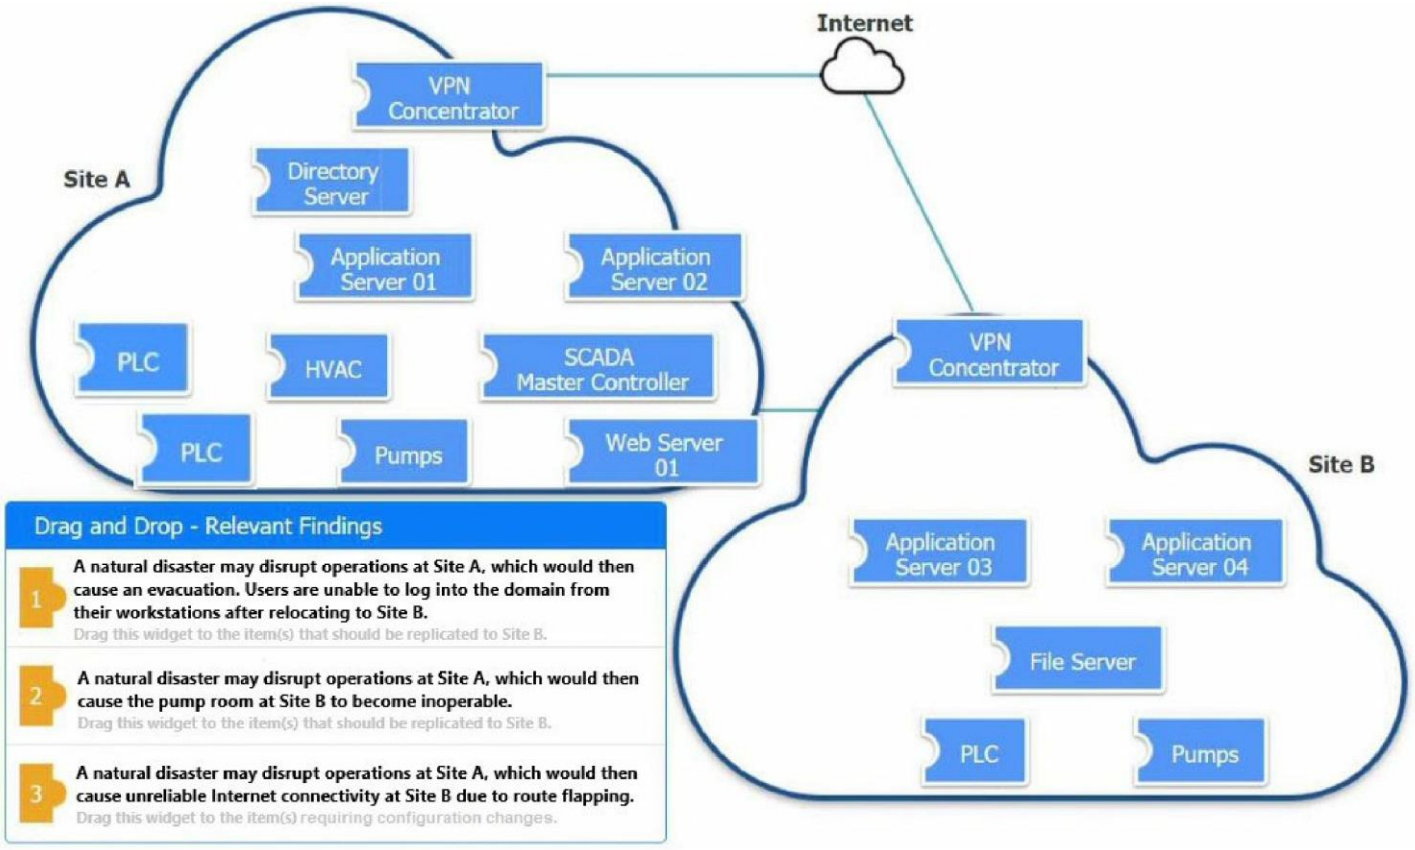

Answer:

1.Directory Server

2-SCADA Master Controller

3-Modify the BGP Configuration

Explanation:

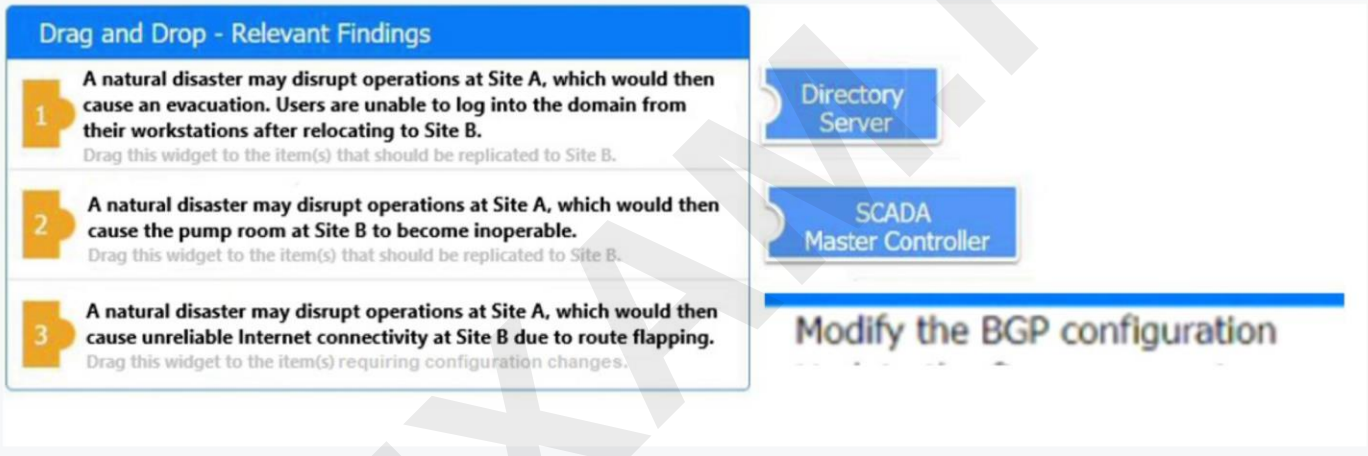

Question: 33

A threat hunting team receives a report about possible APT activity in the network. Which of the following threat management frameworks should the team implement?

A. NIST SP 800-53
B. MITRE ATT&CK
C. The Cyber Kill Chain
D. The Diamond Model of Intrusion Analysis

Answer: B

Explanation:

The MITRE ATT&CK framework is the most suitable choice for a threat hunting team responding to a potential APT (Advanced Persistent Threat) activity report. MITRE ATT&CK provides a comprehensive knowledge base of adversary tactics and techniques based on real-world observations. It maps specific attacker behaviors to stages of an attack lifecycle, enabling the threat hunting team to understand the adversary's potential actions within the network. This allows for proactive identification of indicators of compromise (IOCs) and targeted hunting based on specific APT behaviors. The Cyber Kill Chain, while useful, is more linear and less detailed than MITRE ATT&CK, and is best suited for understanding how the attack unfolds in phases. The Diamond Model of Intrusion Analysis is useful to describe intrusion events, but does not offer a
comprehensive framework for detecting and mitigating threats in the same way as MITRE ATT&CK. NIST SP 800-53, although vital for security controls, is primarily focused on establishing security requirements and controls and does not focus directly on understanding attacker behaviours to hunt for threats. The
framework's extensive knowledge base enables security analysts to effectively prepare for, respond to, and mitigate a range of cyber threats using a standardized approach.

Authoritative Links:

MITRE ATT&CK:https://attack.mitre.org/
Cyber Kill Chain:https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html **NIST SP 800-53:**https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final
The Diamond Model of Intrusion Analysis:https://apps.dtic.mil/sti/pdfs/ADA581475.pdf

Question: 34

Device event logs sourced from MDM software as follows:

Device	Date/Time	Location	Event	Description
ANDROID_1022	01JAN21 0255	38.9072N,77.0369W	PUSH	APPLICATION 1220 INSTALL QUEUED
ANDROID_1022	01JAN21 0301	38.9072N,77.0369W	INVENTORY	APPLICATION 1220 ADDED
ANDROID_1022	01JAN21 0701	25.2854N,51.5310E	CHECK-IN	NORMAL
ANDROID_1022	01JAN21 0900	39.0067N,77.4291W	CHECK-IN	NORMAL
ANDROID_1022	01JAN21 1030	39.0067N,77.4291W	STATUS	LOCAL STORAGE REPORTING 85% FULL

Which of the following security concerns and response actions would BEST address the risks posed by the device in the logs?

A. Malicious installation of an application; change the MDM configuration to remove application ID 1220.
B. Resource leak; recover the device for analysis and clean up the local storage.
C. Impossible travel; disable the device's account and access while investigating.
D. Falsified status reporting; remotely wipe the device.

Answer: C

Explanation:

Due to line 4, a GPS spoofing could be in use either by the newly install app, or before the app was installed.

Question: 35

An energy company is required to report the average pressure of natural gas used over the past quarter. A PLC sends data to a historian server that creates the required reports.

Which of the following historian server locations will allow the business to get the required reports in an ||¢ and IT environment?

A. In the ||¢ environment, use a VPN from the IT environment into the ||¢ environment.

B. In the ||¢ environment, allow IT traffic into the ||¢ environment.

C. In the IT environment, allow PLCs to send data from the ||¢ environment to the IT environment.

D. Use a screened subnet between the ||¢ and IT environments.

Answer: C

Explanation:

The correct answer is C: "In the IT environment, allow PLCs to send data from the ||¢ environment to the IT environment."

Here's why: The goal is to generate reports based on data from the ||¢ (Industrial Control Systems) environment. To achieve this, the historian server, which is responsible for collecting, storing, and analyzing the data, needs access to the PLC data. Placing the historian server in the IT environment and allowing PLCs to send data there enables report generation within a controlled and potentially more secure environment. This approach minimizes the need to extend the IT network's attack surface into the ||¢ environment.

Option A, placing the historian in the ||¢ and using a VPN, creates a VPN connection directly into the ||¢network. While providing access, it potentially increases the risk to the ||¢ environment if the IT network is compromised.

Option B, allowing IT traffic into the ||¢ environment, is generally discouraged. Directly exposing the ||¢environment to broader IT traffic increases the risk of cyberattacks targeting critical infrastructure. ICS networks are often isolated to prevent unauthorized access and modification of control systems.

Option D, using a screened subnet (DMZ) between the ||¢ and IT environments, could be an option, but it's often an intermediate solution. Placing the historian in the DMZ still necessitates data transfer from the ||¢. While better than directly opening the ||¢ to the IT network, it doesn't necessarily leverage the inherent security and management capabilities available within the IT environment for report generation. Moreover, DMZs can become overly complex to manage.

Moving the data from the ||¢ to the IT environment for processing aligns with the principle of defense in depth. The ||¢ environment remains more isolated, and the IT environment, which usually has better monitoring, security controls, and reporting tools, can handle the data analysis. Crucially, this direction of data flow is often preferred because the IT environment is designed to handle larger data volumes and more complex processing. This allows for better scalability and integration with other business reporting tools.

Here are authoritative links for further research:

NIST Special Publication 800-82: Guide to Industrial Control Systems (ICS) Security. This provides comprehensive guidance on securing ICS environments.https://csrc.nist.gov/publications/detail/sp/800-82/r2/final
ISA/IEC 62443 Standards: These standards address security for industrial automation and control systems.https://www.isa.org/standards-and-publications/isa-standards/isa-iec-62443-series-of-standards

Question: 36

Which of the following is a benefit of using steganalysis techniques in forensic response?

A. Breaking a symmetric cipher used in secure voice communications

B. Determining the frequency of unique attacks against DRM-protected media

C. Maintaining chain of custody for acquired evidence

D. Identifying least significant bit encoding of data in a .wav file

Answer: D

Explanation:

Here's a detailed justification for why option D is the most suitable answer:

Steganalysis focuses on detecting the presence of hidden information within seemingly innocuous data. In a forensic investigation, steganalysis is crucial for uncovering covert communication channels or concealed malicious payloads.

Option D, "Identifying least significant bit encoding of data in a .wav file," directly aligns with the core function of steganalysis. Least significant bit (LSB) encoding is a common steganographic technique where secret data is embedded within the least significant bits of a media file (like a .wav audio file). The changes made are subtle enough not to affect the perceived quality of the media file, yet they allow an attacker to hide data in plain sight. Steganalysis techniques specifically aim to detect these subtle alterations indicative of LSB steganography. By successfully identifying LSB encoding, investigators can extract the hidden data.

Options A, B, and C are not related to steganalysis. Breaking a symmetric cipher (A) falls under cryptography or cryptanalysis, not steganalysis. Determining the frequency of attacks against DRM-protected media (B) is related to security monitoring and threat intelligence, and not the detection of hidden data within media.

Maintaining the chain of custody (C) is a general forensic process, not a specific analysis technique like steganalysis.

Therefore, steganalysis directly contributes to the forensic response by uncovering hidden information via techniques such as LSB encoding in media files. It helps reveal covert communications, malware, or other sensitive data that might otherwise be missed. It is an invaluable tool in digital forensics and incident response, where identifying all data, even concealed data, is paramount.

Further Research:

SANS Institute: Provides extensive resources on digital forensics and incident response, often covering steganalysis techniques.

https://www.sans.org/

National Institute of Standards and Technology (NIST): Publishes guidelines and standards related to digital forensics and security.

https://www.nist.gov/

Open Source Steganography Tools: Experimenting with open-source steganalysis tools can provide hands-on understanding.

Search for "open source steganalysis tools" on search engines like DuckDuckGo or Google to find a list of utilities.

Question: 37

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

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

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

A. TLS_AES_128_CCM_8_SHA256
B. TLS_DHE_DSS_WITH_RC4_128_SHA
C. TLS_CHACHA20_POLY1305_SHA256
D. TLS_AES_128_GCM_SHA256

Answer: C

Explanation:

TLS_CHACHA20_POLY1305_SHA256

Question: 38

A security analyst notices a number of SIEM events that show the following activity:

```
10/30/2020 - 8:01 UTC - 192.168.1.1 - sc stop WinDefend
10/30/2020 - 8:05 UTC - 192.168.1.2 - c:\program files\games\comptiacasp.exe
10/30/2020 - 8:07 UTC - 192.168.1.1 - c:\windows\system32\cmd.exe /c powershell https://content.comptia.com/content.exam.ps1
10/30/2020 - 8:07 UTC - 192.168.1.1 - powershell -> 40.90.23.154:443
```

Which of the following response actions should the analyst take FIRST?

A. Disable powershell.exe on all Microsoft Windows endpoints. B.
Restart Microsoft Windows Defender.

C. Configure the forward proxy to block 40.90.23.154.
D. Disable local administrator privileges on the endpoints.

Answer: A

Explanation:

Disable powershell.exe on all Microsoft Windows endpoints.

Question: 39

A company has hired a third party to develop software as part of its strategy to be quicker to market. The company's policy outlines the following requirements:

☞ The credentials used to publish production software to the container registry should be stored in a secure location.
☞ Access should be restricted to the pipeline service account, without the ability for the third-party developer to read the credentials directly.
Which of the following would be the BEST recommendation for storing and monitoring access to these shared credentials?

A. TPM
B. Local secure password file
C. MFA
D. Key vault

Answer: D

Explanation:

The best solution for securely storing and monitoring access to shared credentials for publishing production software to a container registry, while restricting direct access to third-party developers, is a key vault.

A key vault is a secure, centralized repository for storing secrets, such as passwords, API keys, and certificates. It provides strong access control mechanisms to restrict access to these secrets based on identity. In this scenario, the pipeline service account would be granted access to retrieve the credentials from the key vault, allowing it to publish the software. The third-party developer, however, would not have direct access to the credentials within the key vault, satisfying the company policy.

TPM (Trusted Platform Module) is a hardware security module that primarily secures hardware components. While TPM is used to secure the boot process and enable disk encryption, it is not typically used for storing and managing application secrets like container registry credentials.

A local secure password file, while offering some level of security, is not ideal because it is difficult to manage and audit access effectively. Sharing files with third-party developers increases the attack surface.
Monitoring would also be harder to implement.

MFA (Multi-Factor Authentication) enhances authentication by requiring multiple verification factors but does not address the need for secure storage of credentials themselves. MFA would be good for authenticating the pipeline itself, but does not help the pipeline store its credentials to access the container registry.

Key vaults like AWS Key Management Service (KMS), Azure Key Vault, and Google Cloud Key Management offer several advantages. They provide encryption at rest and in transit, granular access control policies (e.g., RBAC), audit logging for monitoring access to secrets, and centralized secret management. They also offer versioning and rotation capabilities, which are essential for maintaining security best practices. The pipeline can then use the identity management services of that same cloud provider to prove that it is authorized to retrieve the credentials without revealing the credentials to the third-party developers.

Here are some authoritative links for further research:

Azure Key Vault:https://learn.microsoft.com/en-us/azure/key-vault/general/overview **AWS KMS:**https://aws.amazon.com/kms/
Google Cloud KMS:https://cloud.google.com/kms

Question: 40

A business stores personal client data of individuals residing in the EU in order to process requests for mortgage loan approvals.
Which of the following does the business's IT manager need to consider?

A. The availability of personal data
B. The right to personal data erasure
C. The company's annual revenue
D. The language of the web application

Answer: B

Explanation:

The correct answer is **B. The right to personal data erasure**.

The scenario describes a business handling personal data of EU residents for mortgage loan approvals. This immediately brings the General Data Protection Regulation (GDPR) into consideration. GDPR grants individuals residing in the EU specific rights regarding their personal data. Among these rights is the "right to be forgotten," formally known as the right to erasure (Article 17 of GDPR). This right entitles individuals to request the deletion of their personal data from the business's systems if certain conditions are met, such as the data no longer being necessary for the purpose for which it was collected, or the individual withdrawing their consent.

Therefore, the IT manager must consider how the business will comply with requests for data erasure. This includes implementing procedures to locate, verify, and permanently delete personal data from all storage locations (including backups) in a timely and secure manner. Failure to comply with data erasure requests can lead to significant fines under GDPR.

While option A, "The availability of personal data," is important for business continuity, it is not the primary GDPR consideration in this specific scenario focusing on EU residents' rights. Option C, "The company's annual revenue," is relevant to calculating potential GDPR fines, but not the immediate, fundamental consideration for the IT manager in complying with individual rights. Option D, "The language of the web application," relates more to accessibility and user experience, but isn't directly related to the core data protection rights granted by GDPR. The core consideration revolves around the rights granted to EU citizens under GDPR, in this instance, the right to have their data erased.

GDPR Article 17 (Right to Erasure - "Right to be Forgotten"): https://gdpr-info.eu/art-17-gdpr/
Information Commissioner's Office (ICO) - Right to erasure: https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/individual-rights/right-to-erasure/

Question: 41

A company publishes several APIs for customers and is required to use keys to segregate customer data sets. Which of the following would be BEST to use to store customer keys?

A. A trusted platform module
B. A hardware security module
C. A localized key store
D. A public key infrastructure

Answer: B

Explanation:

The best option for storing customer API keys in a secure and scalable manner is a Hardware Security Module (HSM). Here's why:

HSMs are purpose-built for key management: They are dedicated hardware devices specifically designed to

securely store and manage cryptographic keys.

Enhanced Security: HSMs offer tamper-resistant protection, meaning keys are protected from unauthorized access, extraction, or modification. This is crucial for safeguarding sensitive customer data and maintaining the integrity of the APIs.

Compliance and Auditability: HSMs often meet stringent security certifications (like FIPS 140-2 Level 3 or higher), which can aid in compliance with industry regulations and facilitate audits.

Scalability: Many HSM solutions can handle a large number of keys, making them suitable for companies with numerous customers and APIs.

Dedicated Key Lifecycle Management: HSMs provide tools for key generation, storage, rotation, and destruction, which are essential for robust key management practices.

Let's consider the other options:

Trusted Platform Module (TPM): TPMs are generally used for device integrity and boot-time security, rather than managing a large number of API keys.

Localized Key Store: A localized key store (e.g., a file on a server) is less secure than an HSM because it's more vulnerable to compromise.

Public Key Infrastructure (PKI): PKI is used for managing digital certificates and enabling secure communication, but it's not the most efficient solution for storing and managing API keys.

In summary, using an HSM provides the highest level of security, manageability, and compliance for storing customer API keys, especially when dealing with sensitive data and regulatory requirements.

Authoritative Links:

NIST Special Publication 800-57 Part 2:https://csrc.nist.gov/publications/detail/sp/800-57-part-2/rev-1/final (Guidelines for Key Management)
Hardware Security Module (HSM) - Thales:https://cpl.thalesgroup.com/encryption/hardware-security-modules-hsms (Example HSM provider and description)

Question: 42

An organization wants to perform a scan of all its systems against best practice security configurations. Which of the following SCAP standards, when combined, will enable the organization to view each of the configuration checks in a machine-readable checklist format for full automation? (Choose two.)

A. ARF
B. XCCDF
C. CPE
D. CVE
E. CVSS
F. OVAL

Answer: BF

Explanation:

The correct answer is B. XCCDF and F. OVAL. Here's why:

XCCDF (Extensible Configuration Checklist Description Format): XCCDF is a language used to write security checklists and benchmarks. It provides a structured, machine-readable format to describe security policies and configuration rules. This allows organizations to define the desired security state of their systems in a way that can be automatically validated. Essentially, it contains the rules and checks.

OVAL (Open Vulnerability and Assessment Language): OVAL is a language used to express low-level tests to check for the presence of vulnerabilities or misconfigurations. It defines the actual tests that are run to evaluate whether a system conforms to the standards defined in an XCCDF document. OVAL provides the specific instructions for assessing compliance.

Combining XCCDF and OVAL allows for fully automated configuration assessment. XCCDF defines what configurations need to be checked, while OVAL defines how those checks are performed. Together, they provide a comprehensive framework for automating security configuration compliance scanning. ARF (Asset Reporting Format) is used for reporting after the assessment has been completed. CPE (Common Platform Enumeration), CVE (Common Vulnerabilities and Exposures) and CVSS (Common Vulnerability Scoring System) address software or vulnerabilities and how to identify them, but don't focus on configuration standards and checklists.

Authoritative Links:

NIST's Security Content Automation Protocol (SCAP): https://scap.nist.gov/
XCCDF at NIST: https://csrc.nist.gov/projects/security-content-automation-protocol/specifications/xccdf OVAL at MITRE: https://oval.mitre.org/

Question: 43

A company is migrating from company-owned phones to a BYOD strategy for mobile devices. The pilot program will start with the executive management team and be rolled out to the rest of the staff in phases. The company's Chief Financial Officer loses a phone multiple times a year.
Which of the following will MOST likely secure the data on the lost device?

A. Require a VPN to be active to access company data.
B. Set up different profiles based on the person's risk.
C. Remotely wipe the device.
D. Require MFA to access company applications.

Answer: C

Explanation:

The best approach to secure data on a lost device in a BYOD environment, especially considering the CFO's history of device loss, is to remotely wipe the device (Option C). This action completely erases all data from the phone, preventing unauthorized access to sensitive company information. While VPNs (Option A) and MFA (Option D) enhance security during device use, they are ineffective when the device is lost and potentially in malicious hands. VPNs only protect data in transit, not data at rest on the device. MFA requires the user to be present and responsive. Setting up different profiles based on risk (Option B) is a good practice in general for security management and device compliance, but doesn't address the immediate threat of a lost device.

Remote wipe offers the most immediate and comprehensive protection in this scenario. Mobile Device Management (MDM) solutions often provide the remote wipe feature. BYOD policies should clearly outline the company's right to remotely wipe a device in case of loss or theft. This is essential to mitigate data breaches and maintain compliance.

Further Reading:

NIST Guidelines on BYOD: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-124r2.pdf SANS Institute on Mobile Device Security: https://www.sans.org/reading-room/whitepapers/mobile/mobile-device-security-challenges-solutions-34105

Question: 44

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:

The best solution is **C. Configure clients to use OCSP.**

Here's a detailed justification:

OCSP (Online Certificate Status Protocol) offers a real-time alternative to CRLs (Certificate Revocation Lists) for checking the revocation status of digital certificates. Instead of downloading and processing entire CRLs, which can be large and bandwidth-intensive, clients using OCSP send a targeted request to an OCSP responder to check the status of a specific certificate. This greatly reduces network traffic, especially in a scenario with many branch offices. Since only specific requests are handled, the load on the CA (Certificate Authority) is reduced compared to generating and distributing full CRLs frequently.

Delta CRLs (B) are smaller, incremental updates to a full CRL. While they reduce bandwidth compared to full CRLs, clients still need to download and process them regularly. In a large, geographically distributed organization, this still generates considerable traffic.

Deploying an RA (Registration Authority) (A) in each branch office would primarily address certificate enrollment, not revocation checking. It also adds significant administrative overhead and complexity. RAs do not directly solve the problem of CRL distribution and bandwidth consumption.

Using GPO (Group Policy Object) (D) to distribute CRLs is a common method, but it doesn't address the fundamental issue of CRL size and the resulting network traffic. Every time a new CRL is issued, it must be distributed to all clients via GPO, consuming bandwidth. Also, GPO is primarily a Windows-centric solution, so it might not be suitable for environments with diverse operating systems.

OCSP provides a more efficient and scalable approach for real-time certificate status checking. It significantly reduces traffic because clients only query the status of certificates they're actively using. This lowers the processing burden on the CA and minimizes power consumption on the CA, addressing the stated requirements. OCSP stapling further improves performance and security by allowing the server to provide the OCSP response along with the certificate during the TLS handshake, eliminating the need for the client to contact the OCSP responder directly in most cases. This also enhances privacy because the OCSP responder doesn't see which websites a client is visiting.

Supporting resources:

Microsoft - How Certificate Revocation Works:https://learn.microsoft.com/en-us/windows-server/networking/core-network-guide/cncg/security/how-certificate-revocation-works
Cloudflare - What is OCSP stapling?:https://www.cloudflare.com/learning/ssl/what-is-ocsp-stapling/

Question: 45

After a security incident, a network security engineer discovers that a portion of the company's sensitive external traffic has been redirected through a secondary
ISP that is not normally used.
Which of the following would BEST secure the routes while allowing the network to function in the event of a single provider failure?

A. Disable BGP and implement a single static route for each internal network.
B. Implement a BGP route reflector.
C. Implement an inbound BGP prefix list.
D. Disable BGP and implement OSPF.

Answer: C

Explanation:

The best solution to secure routes and allow for provider redundancy in this scenario is to implement an inbound BGP prefix list (Option C). Here's why:

The incident indicates a potential BGP route hijacking or manipulation where unauthorized routes are being advertised. Disabling BGP altogether (Options A and D) would eliminate the dynamic routing benefits crucial for internet connectivity and failover. Static routes (Option A) are not scalable or adaptable to network changes, especially in a dynamic environment with multiple ISPs. OSPF (Option D) is an interior gateway protocol (IGP) and is typically used within an organization's network, not between the organization and external ISPs. Therefore, it is not suitable for controlling traffic routing with external providers.

A BGP route reflector (Option B) primarily addresses scalability issues within an autonomous system (AS) and doesn't directly prevent route hijacking. Route reflectors help distribute routing information without requiring full mesh connectivity between BGP routers within the same AS. While helpful in large networks, it doesn't inherently secure against malicious or incorrect route advertisements from external entities.

An inbound BGP prefix list (Option C) is the most appropriate solution because it allows administrators to define which prefixes (network addresses) the company will accept from its ISPs. By creating a whitelist of legitimate prefixes, the company can filter out any unauthorized or malicious route advertisements, preventing traffic redirection through unauthorized ISPs. This filter acts as a security measure, ensuring only trusted routes are accepted. This approach allows continued use of BGP for its dynamic routing benefits and failover capabilities while securing the routing process against unwanted changes. In essence, it enables the network to function normally, utilizing both ISPs, while still only accepting routes that the organization trusts, even if a single provider fails.

Further Reading:

BGP Prefix Filtering:https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/15986-prefix-list.html
Understanding BGP Route Filtering:
https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/bgp-route-filtering.html

Question: 46

A company's SOC has received threat intelligence about an active campaign utilizing a specific vulnerability. The company would like to determine whether it is vulnerable to this active campaign.
Which of the following should the company use to make this determination?

A. Threat hunting

B. A system penetration test
C. Log analysis within the SIEM tool
D. The Cyber Kill Chain

Answer: B

Explanation:

The correct answer is **B. A system penetration test**.

Here's why:

Penetration testing directly assesses vulnerability: A penetration test is a simulated attack against a system or network to identify vulnerabilities that an attacker could exploit. In this scenario, the company wants to determine if it is vulnerable. A pentest is designed to do exactly that – prove whether the vulnerability is exploitable in the company's environment.

Threat hunting (A) proactively seeks out threats: Threat hunting is a proactive security search through networks, endpoints, and datasets to hunt malicious activities that have evaded existing security tools. While important, it aims to find threats, not necessarily confirm the exploitability of a specific vulnerability mentioned in threat intelligence.

Log analysis (C) retrospectively investigates events: Log analysis within a SIEM tool is used to correlate and analyze security events captured in logs. It can help detect if an exploitation attempt has occurred in the past, but it doesn't proactively confirm whether the system is currently vulnerable. Although log analysis will be an important part of assessing whether the vulnerability has been exploited, it can't alone be used to test whether a system is vulnerable.

Cyber Kill Chain (D) is a conceptual framework: The Cyber Kill Chain is a framework for understanding the stages of a cyberattack. While it's valuable for incident response and threat modeling, it doesn't directly determine if a system is vulnerable to a specific exploit.

Therefore, a penetration test is the most direct and effective method for determining whether the company is vulnerable to the active campaign described in the threat intelligence. It validates the existence and exploitability of the specific vulnerability, providing the definitive answer the company needs. It actively tests the system to determine if an attacker could successfully exploit the reported vulnerability.

Supporting Links:

NIST Special Publication 800-115: Technical Guide to Information Security Testing and Assessment: https://csrc.nist.gov/publications/detail/sp/800-115/rev-1/final (Explains the principles and practices of security testing, including penetration testing.)
OWASP Penetration Testing:https://owasp.org/www-project-top-ten/ (Provides guidance and resources on penetration testing methodologies.)

Question: 47

A security engineer needs to recommend a solution that will meet the following requirements:
➪ Identify sensitive data in the provider's network
➪ Maintain compliance with company and regulatory guidelines
➪ Detect and respond to insider threats, privileged user threats, and compromised accounts
➪ Enforce datacentric security, such as encryption, tokenization, and access control
Which of the following solutions should the security engineer recommend to address these requirements?

A. WAF
B. CASB

C. SWG
D. DLP

Answer: C

Explanation:

The correct answer is **C. SWG (Secure Web Gateway)**, but it's important to understand why, and why the other options are less suitable. While CASB is closer, SWG aligns better with the comprehensive requirements. Let's break it down:

Identifying sensitive data: While a CASB can identify sensitive data in cloud applications, SWGs inspect web traffic going in and out, making them effective at identifying sensitive data leaving the network, including traffic to cloud providers.

Compliance: SWGs can enforce policies that ensure compliance with company and regulatory guidelines by filtering traffic, blocking access to prohibited sites or applications, and inspecting content for sensitive information that violates policies.

Insider threats and compromised accounts: SWGs monitor user activity and can detect unusual or suspicious behavior, such as a user accessing sensitive data they normally wouldn't or attempting to exfiltrate large amounts of data. They can be configured to block access based on user identity, location, and other factors.

Datacentric security: Although SWGs are not primarily designed for data-at-rest protection, they can enforce datacentric security measures. For example, they can enforce the use of encryption for web traffic and prevent the transmission of unencrypted sensitive data. Some advanced SWGs can also integrate with tokenization solutions.

Why the other options are less suitable:

A. WAF (Web Application Firewall): WAFs protect web applications from attacks like SQL injection and cross-site scripting. They don't address insider threats, data identification across a network, or datacentric security beyond protecting the application itself.

B. CASB (Cloud Access Security Broker): CASBs focus on securing cloud applications. While they can identify sensitive data within cloud applications and enforce some access controls, they don't provide as broad network-wide visibility and control over all web traffic as an SWG. They are less focused on privileged user/insider threat monitoring across the entire network to all resources. CASBs are better at governance within sanctioned cloud applications, less good at preventing data leaving through unsanctioned channels or detecting general insider abuse.

D. DLP (Data Loss Prevention): While DLP is crucial for data protection, its primary focus is on preventing data exfiltration. It typically operates by scanning data at rest and in transit, but it doesn't provide the same level of real-time monitoring and control over web traffic as an SWG. Also, it's less effective at insider threat detection, especially with privileged accounts, compared to SWG's behavioral analysis capabilities.

In summary: The requirements of identifying sensitive data in a network, maintaining compliance, detecting insider threats, and enforcing datacentric security are best met by an SWG, which provides a comprehensive layer of security for web traffic, including traffic to and from cloud providers.

Supporting Links:

Secure Web Gateway:https://www.paloaltonetworks.com/cyberpedia/what-is-a-secure-web-gateway
CASB:https://www.netskope.com/what-is-casb

While a CASB complements an SWG strategy, the combination is more costly than a complete SWG deployment that has similar capabilities.

Question: 48

A security engineer estimates the company's popular web application experiences 100 attempted breaches per day. In the past four years, the company's data has been breached two times.
Which of the following should the engineer report as the ARO for successful breaches?

A. 0.5
B. 8
C. 50
D. 36,500

Answer: A

Explanation:

Here's a detailed justification for why A (0.5) is the correct Annualized Rate of Occurrence (ARO) in this scenario:

The Annualized Rate of Occurrence (ARO) represents the estimated frequency with which a threat is expected to be successful in a given year. In risk management, ARO is a key component when calculating the Annualized Loss Expectancy (ALE).

The question provides that the company experienced two successful data breaches over the past four years.
To determine the ARO, we need to calculate the average number of breaches occurring per year. This is achieved by dividing the total number of breaches (2) by the number of years (4). Therefore, 2 breaches / 4 years = 0.5 breaches per year.

This result indicates that, on average, the company can expect a successful data breach to occur 0.5 times per year. The other options are incorrect because they don't accurately reflect the historical breach data provided. Option B (8) and C (50) and D (36,500) are not supported by any given calculation from the question.

Therefore, the security engineer should report an ARO of 0.5 for successful breaches.

Further Research:

NIST Special Publication 800-30, Guide for Conducting Risk Assessments: https://csrc.nist.gov/publications/detail/sp/800-30/rev-1/final (This document discusses risk assessment methodologies and includes guidance on calculating ARO.)
SANS Institute: Search for risk management resources on the SANS website (https://www.sans.org/). They often have articles and courses that explain ARO and ALE in detail.

Question: 49

A network architect is designing a new SD-WAN architecture to connect all local sites to a central hub site. The hub is then responsible for redirecting traffic to public cloud and datacenter applications. The SD-WAN routers are managed through a SaaS, and the same security policy is applied to staff whether working in the office or at a remote location. The main requirements are the following:
1. The network supports core applications that have 99.99% uptime.
2. Configuration updates to the SD-WAN routers can only be initiated from the management service.
3. Documents downloaded from websites must be scanned for malware.
Which of the following solutions should the network architect implement to meet the requirements?

A. Reverse proxy, stateful firewalls, and VPNs at the local sites
B. IDSs, WAFs, and forward proxy IDS
C. DoS protection at the hub site, mutual certificate authentication, and cloud proxy

D. IPSs at the hub, Layer 4 firewalls, and DLP

Answer: B

Explanation:

The correct answer is **B. IDSs, WAFs, and forward proxy IDS**. Here's why:

Requirement 1 (99.99% Uptime): While no single component guarantees 99.99% uptime, a Web Application Firewall (WAF) and Intrusion Detection System (IDS) contribute significantly to overall availability by preventing attacks that could lead to downtime. A forward proxy provides load balancing and caching, improving performance and availability.

Requirement 2 (Configuration Updates Initiated from Management Service): This requirement points to the importance of secure control plane communication and access control. A forward proxy and IDS does not directly fulfill this, but indirectly supports the SD-WAN's ability to manage configurations by improving security and visibility into network traffic, which could help identify and prevent unauthorized modifications to the router configurations.

Requirement 3 (Malware Scanning of Downloads): A forward proxy can intercept web traffic and integrate with malware scanning engines. This ensures that downloaded documents are inspected before reaching the user, mitigating the risk of infection. Some forward proxies even offer integrated data loss prevention (DLP) capabilities.

Why other options are less suitable:

A: Reverse proxies are typically used to protect web servers and don't directly address malware scanning. VPNs are useful for secure remote access but are not the primary solution for malware scanning.

C: DoS protection is essential but doesn't directly address malware scanning. Mutual certificate authentication is for secure authentication but not for content inspection. Cloud proxies can scan for malware but are best used in conjunction with other security elements.

D: IPSs are valuable for intrusion prevention but don't always handle malware scanning of downloads. Layer 4 firewalls control traffic based on IP addresses and ports, not content. DLP can prevent data leaks but doesn't directly scan for malware in downloaded files.

In summary, option B provides the most comprehensive solution by addressing the key requirements of high availability through WAF/IDS, support for centralized SD-WAN management via improved network monitoring, and malware scanning of downloads with a forward proxy, making it the most suitable choice for the network architect.

Supporting Links:

Forward Proxy:https://www.cloudflare.com/learning/security/what-is-a-forward-proxy/
Web Application Firewall (WAF):https://www.cloudflare.com/learning/cloud/what-is-a-web-application-firewall-waf/
Intrusion Detection System (IDS):https://www.cloudflare.com/learning/security/threats/intrusion-detection-system/

Question: 50

A security engineer needs to implement a solution to increase the security posture of user endpoints by providing more visibility and control over local administrator accounts. The endpoint security team is overwhelmed with alerts and wants a solution that has minimal operational burdens. Additionally, the solution must maintain a positive user experience after implementation.
Which of the following is the BEST solution to meet these objectives?

A. Implement Privileged Access Management (PAM), keep users in the local administrators group, and enable local administrator account monitoring.

B. Implement PAM, remove users from the local administrators group, and prompt users for explicit approval when elevated privileges are required.

C. Implement EDR, remove users from the local administrators group, and enable privilege escalation monitoring.

D. Implement EDR, keep users in the local administrators group, and enable user behavior analytics.

Answer: A

Explanation:

The best solution is A: Implement Privileged Access Management (PAM), keep users in the local administrators group, and enable local administrator account monitoring.

Here's why:

Visibility and Control: PAM solutions specifically address the need for increased visibility and control over privileged accounts, including local administrator accounts. Monitoring those accounts provides insights into their usage.

Minimal Operational Burden: Keeping users in the local administrators group (while implementing PAM monitoring) minimizes disruption to existing workflows, thus reducing the operational burden on the endpoint security team. Removing users from the local admin group (as in options B and C) creates significant helpdesk tickets and user frustration, adding to the operational burden.

Positive User Experience: Maintaining users' local administrator rights ensures a smooth transition and reduces the likelihood of negative impacts on their work. Prompting users for explicit approval every time elevation is required (as in option B) is disruptive and frustrating.

EDR Limitations: While Endpoint Detection and Response (EDR) provides threat detection and response capabilities, it doesn't primarily focus on managing privileged access. EDR without PAM lacks specific features for controlling and monitoring local administrator accounts effectively.

User Behavior Analytics: While user behavior analytics can detect anomalous activity, it doesn't directly provide control over local administrator rights or offer immediate intervention.

Monitoring Existing Privileges: Option A directly addresses the issue of existing local admin rights with oversight. While risky, the organization might have valid reasons (legacy applications, specific user roles) that justify retaining those rights. PAM monitoring provides security while those conditions remain.

In contrast, removing users from local administrator groups and enabling privilege escalation (options B and C) can lead to many user support requests and is generally not a smooth experience. User behavior analytics alone (option D) does not provide adequate control over local administrator rights.

Therefore, the combination of PAM, retaining current administrator rights, and monitoring those rights provides the best balance between security, operational burden, and user experience in this scenario.

Privileged Access Management (PAM) Best Practices

Endpoint Detection and Response (EDR)

Question: 51

An organization's hunt team thinks a persistent threats exists and already has a foothold in the enterprise network. Which of the following techniques would be BEST for the hunt team to use to entice the adversary to uncover malicious activity?

A. Deploy a SOAR tool.

B. Modify user password history and length requirements.

C. Apply new isolation and segmentation schemes.

D. Implement decoy files on adjacent hosts.

Answer: D

Explanation:

The correct answer is D: Implement decoy files on adjacent hosts. This is because the hunt team suspects an attacker already has access and is moving within the network (a foothold). The goal is to detect the attacker's activity without disrupting legitimate users or alerting the attacker prematurely.

Decoy files (honeypots in file form) are designed to be enticing to an attacker. Placed on hosts near compromised systems, they can bait the adversary into interacting with them. Any access to these files would raise a red flag, providing concrete evidence of malicious activity and helping the hunt team understand the attacker's tactics, techniques, and procedures (TTPs). This approach is proactive and focuses on attracting the attacker rather than passively waiting for suspicious network traffic or endpoint events, which might be missed by other security controls.

Option A (Deploy a SOAR tool) is helpful for automation but won't directly entice an adversary. Option B (Modify password requirements) primarily impacts legitimate users and, while improving overall security posture, doesn't actively draw out an attacker who is already inside. Option C (Apply new isolation and segmentation schemes) is a good security practice to contain breaches, but it also may alert the adversary to the fact that they have been detected and cause them to cease operations and cover their tracks before being identified.

Decoy files offer a lower-risk, targeted approach for confirming a persistent threat and gathering valuable intelligence about the attacker.

Supporting Links:

Honeypots:https://www.sans.org/reading-room/whitepapers/honeypots/introduction-honeypots-387 **MITRE ATT&CK Framework - Deception Techniques:**https://attack.mitre.org/techniques/T1036/ (Note: T1036 covers masquerading but illustrates the attacker's need to appear legitimate, making deception techniques effective)

Question: 52

A junior developer is informed about the impact of new malware on an Advanced RISC Machine (ARM) CPU, and the code must be fixed accordingly. Based on the debug, the malware is able to insert itself in another process memory location. Which of the following technologies can the developer enable on the ARM architecture to prevent this type of malware?

A. Execute never

B. No-execute

C. Total memory encryption

D. Virtual memory protection

Answer: A

Explanation:

The correct answer is A. Execute Never (XN).

The scenario describes malware injecting itself into another process's memory space on an ARM CPU. The

goal is to prevent code execution from data memory regions. Execute Never (XN), also known as Data Execution Prevention (DEP), is a hardware-level security feature that marks certain memory regions as non-executable. This prevents the CPU from executing code residing in these marked memory areas. Therefore, if the malware attempts to execute code it has injected into a data segment, the CPU will refuse, effectively blocking the attack.

Option B, No-Execute, is essentially the same concept as Execute Never; however, "Execute Never" is the more common and technically accurate term used in the context of ARM architectures. While both terms refer to preventing code execution from data regions, Execute Never specifically relates to ARM's implementation.

Option C, Total Memory Encryption (TME), encrypts the entire memory space. While it protects data confidentiality, it does not directly prevent code execution from data segments. Even with memory encryption, if the injected code lands in a decrypted data segment, it could still potentially execute if the XN bit is not set.

Option D, Virtual Memory Protection, is a broader concept related to memory management and isolation of processes. While it provides a level of protection, it does not specifically prevent code execution from data memory. Virtual memory protection is more focused on isolating address spaces to prevent one process from directly accessing another's memory. Implementing XN is a more targeted approach.

Therefore, to specifically prevent malware from executing injected code in data memory regions on an ARM CPU, enabling Execute Never is the most effective solution.

Supporting Links:

ARM Security Features: https://developer.arm.com/documentation/102418/0100/armv8-m-security-extensions/security-extensions/memory-protection
Data Execution Prevention (DEP): https://en.wikipedia.org/wiki/Data_Execution_Prevention

Question: 53

A company is implementing SSL inspection. During the next six months, multiple web applications that will be separated out with subdomains will be deployed.
Which of the following will allow the inspection of the data without multiple certificate deployments?

A. Include all available cipher suites.
B. Create a wildcard certificate.
C. Use a third-party CA.
D. Implement certificate pinning.

Answer: B

Explanation:

The correct answer is B: Create a wildcard certificate. Here's why:

The scenario involves SSL inspection across multiple web applications residing on subdomains within the same domain. The primary goal is to avoid deploying numerous certificates as new applications are rolled out over the next six months. Wildcard certificates are designed precisely for this purpose.

A wildcard certificate secures a domain and all its first-level subdomains with a single certificate. For example, a wildcard certificate issued for *.example.com would secure app1.example.com, app2.example.com, and any other subdomain directly under example.com. This eliminates the need to purchase and manage individual certificates for each subdomain, simplifying certificate management significantly.

In the context of SSL inspection, the inspection device needs to present a valid certificate to the client when it

intercepts and decrypts the SSL/TLS traffic. With a wildcard certificate, the SSL inspection appliance can use the same certificate for all subdomains under the primary domain, streamlining the process.

Option A (Include all available cipher suites) relates to the encryption algorithms supported during SSL/TLS handshakes and doesn't address certificate management for multiple subdomains. Option C (Use a third-party CA) is about the certificate authority used to issue the certificate, not the type of certificate, and does not help reduce the number of certificate deployments. Option D (Implement certificate pinning) is a security mechanism that ties an application to a specific certificate or public key, making it harder for attackers to intercept traffic with fraudulent certificates. However, it's typically used to prevent SSL inspection, not to facilitate it across multiple subdomains without certificate proliferation. Moreover, certificate pinning would be extremely difficult to manage with a large number of subdomains undergoing frequent change. Pinning a certificate issued by the SSL inspection device would make the application trust the MitM proxy, but would require updates to the pinned certificates frequently.

Therefore, a wildcard certificate offers the most efficient solution by allowing a single certificate to cover all the subdomains, thereby minimizing certificate deployments and management overhead during the rollout of new web applications.

Further Reading:

Wildcard Certificates: https://www.digicert.com/wildcard-ssl-certificates
SSL Inspection: https://www.fortinet.com/resources/cyberglossary/ssl-tls-inspection

Question: 54

A small business requires a low-cost approach to theft detection for the audio recordings it produces and sells. Which of the following techniques will MOST likely meet the business's needs?

A. Performing deep-packet inspection of all digital audio files
B. Adding identifying filesystem metadata to the digital audio files
C. Implementing steganography
D. Purchasing and installing a DRM suite

Answer: C

Explanation:

The correct answer is **C. Implementing steganography.** Here's why:

Steganography, the practice of concealing a message within another, non-secret, message or medium, is a relatively low-cost approach to theft detection in audio recordings. In this scenario, the business can embed a unique, virtually imperceptible digital watermark or identifier within each audio file. If an unauthorized copy surfaces, the presence of the watermark proves ownership and illegal distribution. It does not rely on expensive infrastructure.

Deep packet inspection (DPI), while effective for network traffic analysis (option A), is unsuitable and resource-intensive for analyzing individual audio files to detect theft. DPI is normally applied when inspecting network traffic to identify the protocol in use, malicious code and attacks and not for verifying the data in a file against metadata.

Adding filesystem metadata (option B), like author or copyright information, is easily stripped or modified by someone intending to steal the audio. Metadata is part of the file but it is easily modified and can be easily removed by any modern operating system.

Purchasing and installing a Digital Rights Management (DRM) suite (option D) is generally a high-cost

approach, as DRM systems often involve licensing fees, complex implementation, and ongoing maintenance. For a small business seeking a low-cost solution, this is not a feasible first step.

Steganography is a more suitable option because the embedding process can be automated and does not necessarily require purchasing a expensive third-party tool.**Further reading:**

Steganography:https://en.wikipedia.org/wiki/Steganography

Question: 55

Clients are reporting slowness when attempting to access a series of load-balanced APIs that do not require authentication. The servers that host the APIs are showing heavy CPU utilization. No alerts are found on the WAFs sitting in front of the APIs.
Which of the following should a security engineer recommend to BEST remedy the performance issues in a timely manner?

A. Implement rate limiting on the API.
B. Implement geoblocking on the WAF.
C. Implement OAuth 2.0 on the API.
D. Implement input validation on the API.

Answer: C

Explanation:

The question states that the APIs do not require authentication. Implementing OAuth 2.0 (Option C), before rate limiting (Option A), is unlikely to directly address the root cause of high CPU utilization and performance issues, especially considering no authentication is currently required. The increased overhead from authentication processing could even exacerbate the problem if the issue is a denial-of-service or similar attack. Therefore, **Option C is incorrect.**

The problem description suggests a potential overload of the API servers, leading to high CPU utilization. Implementing rate limiting (Option A) will restrict the number of requests a client can make within a specific timeframe. This can effectively prevent a single client (malicious or otherwise) from overwhelming the API servers, thereby reducing CPU load and improving performance for all users. This aligns directly with addressing the symptoms described.

Geoblocking (Option B) might help if the traffic originates from a specific geographic region known for malicious activity, but it's a blunt instrument and could block legitimate users. It is unlikely to remedy the immediate performance issue as rate limiting would.

Input validation (Option D) is important for security but is unlikely to be the immediate cause of high CPU utilization, especially if the APIs aren't currently susceptible to input-based attacks.

Therefore, the most effective and timely solution is to implement rate limiting (Option A) to control the traffic volume hitting the API servers, thereby alleviating CPU load and improving responsiveness. Rate limiting is a common technique to mitigate denial-of-service attacks and prevent resource exhaustion.

https://cloud.google.com/api-gateway/docs/rate-limitinghttps://aws.amazon.com/blogs/architecture/rate-limiting-strategies-for-api-gateways/https://learn.microsoft.com/en-us/azure/architecture/patterns/rate-limiting-pattern

Question: 56

An organization is considering a BYOD standard to support remote working. The first iteration of the solution will utilize only approved collaboration applications and the ability to move corporate data between those applications.
The security team has concerns about the following:
- Unstructured data being exfiltrated after an employee leaves the organization
- Data being exfiltrated as a result of compromised credentials
- Sensitive information in emails being exfiltrated

Which of the following solutions should the security team implement to mitigate the risk of data loss?

A. Mobile device management, remote wipe, and data loss detection
B. Conditional access, DoH, and full disk encryption
C. Mobile application management, MFA, and DRM
D. Certificates, DLP, and geofencing

Answer: C

Explanation:

The correct answer is **C. Mobile application management, MFA, and DRM**. This combination directly addresses the stated security concerns in the BYOD scenario.

Mobile Application Management (MAM) focuses on controlling the applications used on personal devices, enabling the organization to manage and secure the approved collaboration applications. This limits the risk of unstructured data exfiltration by containing corporate data within managed apps.

Multi-Factor Authentication (MFA) significantly reduces the risk of data exfiltration resulting from compromised credentials. Even if a password is stolen, the attacker needs a second factor (e.g., a code from a mobile app) to gain access.

Digital Rights Management (DRM) helps prevent sensitive information in emails and other documents from being copied, forwarded, or otherwise exfiltrated. It ensures that access and usage restrictions are enforced.

Option A (Mobile device management, remote wipe, and data loss detection) is less focused on the application level. Remote wipe is drastic and data loss detection is reactive.

Option B (Conditional access, DoH, and full disk encryption) focuses on device access and general security but doesn't directly manage data movement between applications.

Option D (Certificates, DLP, and geofencing) provides security at the device and network level, and for data loss prevention in transit, but does not specifically protect data from exfiltration within or between approved applications on the device itself. Certificates secure communication, DLP prevents outbound data leaks, and Geofencing restricts access based on location.

Therefore, the solution most effectively addresses the concerns of unstructured data exfiltration after employee departure, data exfiltration due to compromised credentials, and sensitive email leakage. MAM, MFA, and DRM work together to control the usage and security of corporate data within approved applications on personal devices.

Authoritative Links:

Mobile Application Management (MAM):https://www.vmware.com/topics/glossary/content/mobile-application-management.html
Multi-Factor Authentication (MFA):https://www.nist.gov/itl/applied-cybersecurity/nist-cybersecurity-framework/online-learning-module-multi-factor
Digital Rights Management (DRM):https://www.w3.org/TR/drm-overview/

Question: 57

A Chief Information Officer is considering migrating all company data to the cloud to save money on expensive

SAN storage.
Which of the following is a security concern that will MOST likely need to be addressed during migration?

A. Latency
B. Data exposure
C. Data loss
D. Data dispersion

Answer: B

Explanation:

The correct answer is **B. Data exposure**. Here's a detailed justification:

When migrating data to the cloud, especially sensitive company data, a primary concern is **data exposure**.
This risk arises from various factors intrinsic to cloud environments. Firstly, data is now managed by a third-party provider, introducing a new set of trust relationships and potential vulnerabilities. Improperly configured cloud storage can lead to data being unintentionally made public or accessible to unauthorized users. For example, simple misconfigurations in access control lists (ACLs) or storage bucket policies can expose sensitive information.

Secondly, cloud environments are often targets for cyberattacks. Attackers constantly probe for weaknesses in cloud infrastructure, seeking to exploit vulnerabilities and gain access to stored data. A successful breach can result in significant data leakage.

Furthermore, compliance regulations (like HIPAA, GDPR, or PCI DSS) mandate stringent data protection measures. Moving data to the cloud necessitates ensuring the cloud provider meets these compliance requirements and that the data is protected in transit and at rest through encryption and other security controls.

Latency (A), while a concern for application performance, is not primarily a security issue. **Data loss (C)** is a general risk but is often mitigated by cloud providers' built-in redundancy and backup mechanisms. **Data dispersion (D)**, the distribution of data across multiple locations, can present challenges but is also a feature leveraged for redundancy and resilience, and not inherently a security vulnerability.

The overriding concern is the potential exposure of sensitive data to unauthorized entities or vulnerabilities introduced by the cloud environment itself. Appropriate security measures, such as encryption, strong access controls, vigilant monitoring, and regular audits, are critical to mitigating this risk. Selecting a reputable cloud provider with robust security certifications is also paramount.

Further Research:

Cloud Security Alliance (CSA):https://cloudsecurityalliance.org/
NIST Cloud Computing Program:https://www.nist.gov/itl/applied-cybersecurity/nccoe/what-cloud-computing

Question: 58

Due to locality and budget constraints, an organization's satellite office has a lower bandwidth allocation than other offices in the organization. As a result, the local security infrastructure staff is assessing architectural options that will help preserve network bandwidth and increase speed to both internal and external resources while not sacrificing threat visibility.
Which of the following would be the BEST option to implement?

A. Distributed connection allocation
B. Local caching

C. Content delivery network

D. SD-WAN vertical heterogeneity

Answer: B

Explanation:

The best option is local caching (B). Given the bandwidth constraints and need for increased speed, caching frequently accessed resources locally reduces the need to repeatedly download them from remote servers, thus conserving bandwidth. A local caching server stores copies of web pages, files, and other content, serving them directly to users within the satellite office. This minimizes latency and improves response times for frequently requested data.

Distributed connection allocation (A) might help balance connections but doesn't fundamentally reduce bandwidth usage. A content delivery network (CDN) (C) would improve external resource access speed but might be too costly for a satellite office with budget constraints, and its primary benefit is for external content, not internal. SD-WAN vertical heterogeneity (D), while offering potential benefits like traffic prioritization, doesn't directly address the core issue of bandwidth conservation for frequently accessed content as efficiently as local caching does. Local caching directly tackles the problem by reducing redundant data transfers, making it the most impactful and budget-friendly solution.

Therefore, local caching provides the most efficient and cost-effective solution for addressing the bandwidth limitations and improving speed for frequently accessed resources at the satellite office.

References:

Caching:https://www.cloudflare.com/learning/cdn/what-is-caching/
Bandwidth Management:https://www.cisco.com/c/en/us/solutions/small-business/resource-center/networking/bandwidth-management.html

Question: 59

A security analyst is concerned that a malicious piece of code was downloaded on a Linux system. After some research, the analyst determines that the suspected piece of code is performing a lot of input/output (I/O) on the disk drive.

```
procs ---------memory--------swap---io--  --system-- -----cpu------
r b swpd free  buff   cache  si so bi     bo         in  cs   us sy id wa st
1 0 0    44408 110052 623096 0  0  300    200003     88  1446 31 4  65 0  0
0 0 0    44524 110052 623096 0  0  400020 20         84  872  11 2  87 0  0
0 2 0    44516 110052 623096 0  0  10     0          149 142  18 5  77 0  0
0 0 0    44524 110052 623096 0  0  0      0          60  431  14 1  85 0  0
```

Based on the output above, from which of the following process IDs can the analyst begin an investigation?

A. 65

B. 77

C. 83

D. 87

Answer: D

Explanation:

87 is a correct answer.

Question: 60

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

A. The client can seamlessly move data.
B. The vendor can change product offerings.
C. The client receives a sufficient level of service.
D. The client experiences decreased quality of service.
E. The client can leverage a multicloud approach.
F. The client experiences increased interoperability.

Answer: BD

Explanation:

The correct answer is B and D. Vendor lock-in refers to a situation where a customer becomes dependent on a specific vendor's products or services, making it difficult to switch to another vendor without significant costs, disruption, or compatibility issues.

Option B, "The vendor can change product offerings," is a risk because once locked in, the client becomes vulnerable to the vendor's decisions. The vendor may alter features, discontinue products, or change pricing models, potentially negatively impacting the client's operations and budget. The client has limited negotiation power due to the difficulty of switching to a different provider.

Option D, "The client experiences decreased quality of service," is also a risk. A locked-in vendor may become complacent, knowing the client faces barriers to leaving. This can lead to a decline in service levels, delayed support, or reduced innovation, as the vendor feels less pressure to maintain competitiveness.

Options A, E, and F are incorrect because they represent scenarios that mitigate vendor lock-in. Option A (seamless data movement) and E (multicloud approach) provide flexibility and alternative options. Option F (increased interoperability) allows for easier integration with other services, reducing dependency on a single vendor. Option C (sufficient level of service) is generally a positive outcome, not a risk.

Therefore, the two main risks of vendor lock-in are the vendor's ability to unilaterally change product offerings and the potential for decreased quality of service due to the reduced incentive for the vendor to remain competitive. Clients should carefully assess the potential for lock-in when choosing cloud providers or other technology solutions and consider strategies like open standards, multicloud environments, and well-defined exit strategies to mitigate these risks.

Supporting Resources:

NIST Special Publication 800-146:https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-146.pdf (Discusses vendor lock-in in the context of cloud computing.)
Cloud Security Alliance:https://cloudsecurityalliance.org/ (Provides resources and best practices for cloud security, including vendor lock-in mitigation.)

Question: 61

An organization recently experienced a ransomware attack. The security team leader is concerned about the attack reoccurring. However, no further security measures have been implemented.
Which of the following processes can be used to identify potential prevention recommendations?

A. Detection

B. Remediation
C. Preparation
D. Recovery

Answer: C

Explanation:

The correct answer is **C. Preparation**.

Preparation, in the context of cybersecurity, is a proactive process that involves implementing measures to prevent attacks and mitigate their impact should they occur. Given that the security team leader is concerned about a repeat ransomware attack and no further security measures have been implemented, focusing on preparation is paramount. This involves identifying vulnerabilities, hardening systems, educating users, and developing incident response plans before another attack happens. It is a vital step in creating a resilient defense.

Detection (A) is essential for recognizing an ongoing attack, but it is not the primary process for identifying potential prevention recommendations. While detection mechanisms could reveal attack patterns, they are more reactive than proactive. Remediation (B) focuses on fixing the damage after an attack has occurred.

While it's related to recovery, it doesn't directly assist in discovering preventative controls. Recovery (D) is about restoring systems and data to their pre-attack state. Like remediation, it's reactive rather than proactive in identifying prevention strategies.

Effective preparation includes vulnerability assessments, penetration testing, security awareness training for employees (particularly phishing awareness, since ransomware often enters via phishing), creating and regularly testing backups, patching systems promptly, implementing strong access controls, and having a documented incident response plan that is practiced via tabletop exercises or simulations. The incident response plan will detail the steps to be taken to minimize damage and downtime in the event of another attack. The entire security lifecycle must be continuously improved based on lessons learned from past incidents and threat intelligence. These preventative measures and processes significantly reduce the likelihood and potential damage of future attacks.

Authoritative Links for Further Research:

NIST Cybersecurity Framework (CSF):https://www.nist.gov/cyberframework (Focus on the "Identify" and "Protect" functions)
SANS Institute:https://www.sans.org/ (Offers a wealth of resources on cybersecurity best practices, including incident response and prevention)
CISA (Cybersecurity and Infrastructure Security Agency):https://www.cisa.gov/ (Provides guidance and resources on cybersecurity threats and mitigation strategies).

Question: 62

A security architect is implementing a web application that uses a database back end. Prior to the production, the architect is concerned about the possibility of
XSS attacks and wants to identify security controls that could be put in place to prevent these attacks. Which of the following sources could the architect consult to address this security concern?

A. SDLC
B. OVAL
C. IEEE
D. OWASP

Answer: D

Explanation:

The correct answer is D, OWASP (Open Web Application Security Project). Here's why:

OWASP is a non-profit organization dedicated to improving the security of software. Its primary focus is on web application security. They provide a wealth of resources, tools, and documentation specifically designed to help developers and security professionals understand and mitigate web application vulnerabilities, including Cross-Site Scripting (XSS).

OWASP provides resources such as:

OWASP Top Ten: A list of the most critical web application security risks, XSS is usually on the list. **OWASP Testing Guide:** A comprehensive guide on how to test for web application vulnerabilities. **OWASP Cheat Sheet Series:** Provides concise advice on how to prevent specific vulnerabilities.
OWASP Zed Attack Proxy (ZAP): A free, open-source web application security scanner.

These resources provide the architect with specific guidance on identifying, preventing, and mitigating XSS vulnerabilities in the web application. The resources offer concrete coding examples and mitigation techniques tailored to XSS attacks, covering various types of XSS vulnerabilities and the contexts where they can occur.

Here's why the other options are less suitable:

A. SDLC (Software Development Life Cycle): While the SDLC is a crucial framework for software development, it provides a broader perspective on the entire development process. While security should be integrated into the SDLC, it doesn't specifically offer the in-depth technical guidance on XSS prevention that OWASP does.

B. OVAL (Open Vulnerability and Assessment Language): OVAL is primarily used for assessing the presence of known vulnerabilities on systems. It's more focused on system-level security configurations and patch management than on specific web application vulnerabilities like XSS.

C. IEEE (Institute of Electrical and Electronics Engineers): IEEE is a professional organization that publishes standards related to a wide range of electrical and electronic engineering topics. While IEEE does have standards related to software engineering, they are typically more general and do not provide specific guidance on web application security or XSS prevention.

Therefore, OWASP is the most appropriate resource for the security architect to consult for addressing XSS concerns in the web application.

Authoritative Links:

OWASP: https://owasp.org/
OWASP Top Ten: https://owasp.org/www-project-top-ten/
OWASP Cheat Sheet Series: https://cheatsheetseries.owasp.org/

Question: 63

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:

The correct answer is **C. Track the library versions and monitor the CVE website for related vulnerabilities.**

Here's why:

Open-source libraries, while beneficial for rapid development, introduce security risks if not managed properly. Simply performing SAST/DAST on these libraries (option A) isn't sufficient. SAST/DAST are designed to find vulnerabilities in the application code, not necessarily vulnerabilities inherent in the libraries themselves. While they might indirectly identify some library-related issues, they are not a comprehensive solution.

Option B, implementing SDLC security guidelines, is a general practice that should be in place, but it doesn't specifically address the unique challenges posed by open-source libraries. It's a good broader practice, but it doesn't provide the targeted protection that is needed for the dependencies.

Option D, unit testing the open-source libraries, is impractical. The organization did not develop the library; therefore, the libraries will have their own testing suites and the organization will be unlikely to be able to perform unit tests on a package that they are simply using as a dependency.

Tracking the library versions and monitoring the CVE (Common Vulnerabilities and Exposures) website is the most effective approach (option C). This allows the organization to be proactively notified when vulnerabilities are discovered in the specific versions of the libraries they are using. Once a CVE is discovered, the organization can then take swift action to remediate by patching or upgrading the library in their codebase.

This is a vital component of dependency management. This process ensures that known vulnerabilities are quickly addressed.

Dependency management is vital, and the following will enable one to learn more:

NIST National Vulnerability Database (NVD):https://nvd.nist.gov/
OWASP Dependency Check:https://owasp.org/www-project-dependency-check/

These processes combined help secure the SDLC and minimize risks associated with open-source components.

Question: 64

A security analyst is investigating a possible buffer overflow attack. The following output was found on a user's workstation: graphic.linux_randomization.prg
Which of the following technologies would mitigate the manipulation of memory segments?

A. NX bit

B. ASLR

C. DEP

D. HSM

Answer: B

Explanation:

The correct answer is **B. ASLR (Address Space Layout Randomization)**.

Here's a detailed justification:

The question describes a scenario involving a potential buffer overflow attack targeting the program graphic.linux_randomization.prg. Buffer overflow attacks exploit vulnerabilities where a program writes data beyond the allocated memory buffer, potentially overwriting adjacent memory regions. This can be used to inject and execute malicious code.

ASLR is a security technique that randomizes the memory locations used by a program's key data areas. These areas include the base address of the executable, the location of libraries, the heap, and the stack. By randomizing these locations, ASLR makes it much more difficult for an attacker to reliably predict the target address of injected code.

In a successful buffer overflow attack without ASLR, an attacker can calculate the exact memory address where they want to overwrite the return address on the stack with the address of their malicious code. ASLR thwarts this by changing the base address of the program and its libraries each time it is executed, making the hardcoded addresses in an exploit unreliable.

The linux_randomization part of the filename hints at the presence or intended use of memory address randomization. If ASLR is correctly implemented and enabled, even if an attacker manages to overflow a buffer, they cannot easily redirect the execution flow to malicious code because they don't know where that code is located in memory.

While the NX bit (A) and DEP (C) also contribute to memory protection, they primarily prevent the execution of code in data segments (like the stack or heap). They help against attacks that inject code into these areas, but they don't specifically address the issue of predictability in memory addresses that ASLR does. The NX bit/DEP would prevent the execution of code injected into an overflowing buffer, while ASLR prevents the predictable location where such code would attempt to execute from in the first place. An HSM (D) is a hardware security module that manages cryptographic keys and is irrelevant to this type of memory-based attack. ASLR directly mitigates the effectiveness of buffer overflow attacks by making memory addresses unpredictable.

Therefore, ASLR is the most appropriate technology to mitigate the manipulation of memory segments in a buffer overflow attack context, specifically in the case described where the target program's name even hints at an attempt for randomization, and its potential defeat by exploitation.

Further reading:

Address Space Layout Randomization (ASLR):
https://en.wikipedia.org/wiki/Address_space_layout_randomization
Buffer Overflow:https://owasp.org/www-community/vulnerabilities/Buffer_Overflow

Question: 65

An e-commerce company is running a web server on premises, and the resource utilization is usually less than 30%. During the last two holiday seasons, the server experienced performance issues because of too many connections, and several customers were not able to finalize purchase orders. The company is looking to change the server configuration to avoid this kind of performance issue.
Which of the following is the MOST cost-effective solution?

A.Move the server to a cloud provider.
B.Change the operating system.
C.Buy a new server and create an active-active cluster.
D.Upgrade the server with a new one.

Answer: A

Explanation:

The most cost-effective solution is to move the server to a cloud provider (Option A). Here's why:

Scalability: Cloud providers offer on-demand scalability. During peak seasons, the e-commerce company can automatically scale up its resources (CPU, memory, bandwidth) to handle the increased traffic. This elasticity ensures that the server can handle a large number of connections without performance issues. Once the holiday season ends, the resources can be scaled down, reducing costs.

Cost Efficiency: Compared to purchasing and maintaining a new server (Options C and D), cloud solutions can be more cost-effective, especially considering the server utilization is typically low. Cloud providers offer various pricing models (pay-as-you-go, reserved instances) that allow businesses to pay only for the resources they consume. Buying a new server might mean it sits idle for most of the year.

High Availability: Cloud providers offer built-in redundancy and high availability features. Moving the server to a cloud environment can improve the server's uptime and reduce the risk of downtime during peak seasons. Many cloud providers offer automated failover mechanisms to ensure that the application remains available even if one server fails.

Active-Active Cluster Complexity: Creating an active-active cluster (Option C) involves purchasing and configuring multiple servers, setting up load balancing, and managing data synchronization. This can be complex and expensive.

Operating System Change (Option B): Changing the operating system might not directly address the scalability issues. While a more optimized OS may offer some performance improvements, it is unlikely to solve the problem of handling a sudden surge in connections.

Vendor Management: With cloud solutions, vendor management of underlying hardware becomes a non-issue.

Example: Amazon Web Services (AWS) Auto Scaling can automatically adjust the number of EC2 instances based on traffic demands. https://aws.amazon.com/autoscaling/

Example: Azure Virtual Machine Scale Sets allow you to create and manage a group of identical, load-balanced VMs. https://azure.microsoft.com/en-us/products/virtual-machine-scale-sets

Therefore, migrating to the cloud provides the most cost-effective way to handle peak season traffic without significant capital expenditure or ongoing maintenance overhead, and it offers scalability and high availability that on-premises solutions might struggle to match in a cost-efficient manner.

Question: 66

A company has decided to purchase a license for software that is used to operate a mission-critical process. The third-party developer is new to the industry but is delivering what the company needs at this time.
Which of the following BEST describes the reason why utilizing a source code escrow will reduce the operational risk to the company if the third party stops supporting the application?

A. The company will have access to the latest version to continue development.
B. The company will be able to force the third-party developer to continue support.
C. The company will be able to manage the third-party developer's development process.
D. The company will be paid by the third-party developer to hire a new development team.

Answer: A

Explanation:

Here's a detailed justification for why option A is the best answer, along with supporting concepts and resources:

Justification for Option A: The company will have access to the latest version to continue development.

Source code escrow is a vital risk mitigation strategy, especially when dealing with smaller or newer software vendors for mission-critical applications. Its core purpose is to ensure business continuity if the vendor becomes unable or unwilling to support the software. This commonly happens due to bankruptcy, acquisition, or a change in business strategy by the third-party developer.

The essence of a source code escrow arrangement involves depositing the software's source code, along with necessary build instructions and documentation, with a trusted third-party escrow agent. The escrow agent holds this material under specific contractual terms. These terms define the "release conditions," or events that trigger the release of the source code to the licensee (in this case, the company using the software).

If the trigger events described in the contract occur (such as the vendor ceasing support or going out of business), the company gains access to the escrowed source code. This enables the company to either continue development internally, hire another vendor to provide support and enhancement, or adapt the software to their changing needs. This greatly reduces the risk of business disruption due to the vendor's inability to maintain the software.

Options B, C, and D are not the primary purposes of source code escrow. It cannot force support (B), manage development (C), or ensure payment for a new development team (D). The primary focus is ensuring access to the source code in contingency situations.

Therefore, source code escrow is a contingency plan that grants the company access to the latest version of the software's code. This access provides the ability to continue development, maintenance, and adaptation, thus reducing the operational risk associated with the third-party developer's potential inability to provide future support.

Authoritative Resources:

Source Code Escrow: A Guide for Businesses:
https://www.ironmountain.com/resources/whitepapers/s/source-code-escrow-guide
Understanding Source Code Escrow: https://www.lexology.com/library/detail.aspx?g=3131e084-105f-457b-8208-2b15c12f16a2

Question: 67

A security analyst is researching containerization concepts for an organization. The analyst is concerned about potential resource exhaustion scenarios on the
Docker host due to a single application that is overconsuming available resources.
Which of the following core Linux concepts BEST reflects the ability to limit resource allocation to containers?

A. Union filesystem overlay
B. Cgroups
C. Linux namespaces
D. Device mapper

Answer: B

Explanation:

The correct answer is B, Cgroups (Control Groups). Cgroups are a core Linux kernel feature designed to limit,

account for, and isolate the resource usage (CPU, memory, disk I/O, network) of a set of processes. They are fundamental to containerization as they provide the mechanism to prevent a single container or application from monopolizing system resources and potentially starving other applications running on the same host. By using Cgroups, the security analyst can effectively limit the resource allocation to containers, mitigating the risk of resource exhaustion on the Docker host. This ensures that even if one container tries to consume excessive resources, it will be constrained by the limits defined in its Cgroup configuration, preventing it from impacting other containers or the host system's stability.

Union filesystem overlays (A) provide a layered approach to file systems but do not directly manage resource allocation. Linux namespaces (C) offer isolation by providing separate views of the system (process IDs, network interfaces, etc.) but do not directly control resource limits. Device mapper (D) is a framework for managing block devices but is not directly related to resource allocation limits for processes. Cgroups are specifically designed for resource management and are essential for stable and secure container environments.

Authoritative Links:

Red Hat Documentation on Cgroups:https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/resource_management_guide/ch01
Docker Documentation on Resource Constraints:
https://docs.docker.com/config/containers/resource_constraints/

Question: 68

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: AB

Explanation:

The correct answer is A and B. Here's why:

A. Utilize code signing by a trusted third party: Code signing is a digital signature added to executable files and scripts that verifies the software's origin and guarantees that the code hasn't been altered or corrupted since it was signed by the developer. Using a trusted third-party Certificate Authority (CA) for signing provides a higher level of assurance, as the CA has validated the developer's identity. This assures users that the code comes from a reputable source and hasn't been tampered with. This mechanism ensures integrity by detecting any unauthorized modifications. Learn more about code signing.

B. Implement certificate-based authentication: While primarily used for authentication (verifying identity), certificate-based authentication enhances integrity in this context. By requiring cryptographic certificates for access or modification, the program's modules are protected from unauthorized changes. It ensures that only individuals or processes with the correct certificate can access and modify the code. If the code is altered outside the approved and certified process, the alteration won't have a valid certificate to be accepted.
Certificate authentication strengthens access controls which preserves integrity. Further reading on

certificate-based authentication.

C. Verify MD5 hashes: While MD5 hashes can detect changes, they are considered cryptographically broken and are vulnerable to collisions. This means that a malicious user could potentially create a different piece of code with the same MD5 hash, defeating the integrity check.

D. Compress the program with a password: Compressing a program with a password provides a basic level of protection against casual users, but it's easily bypassed by a determined attacker. It doesn't offer strong integrity protection.

E. Encrypt with 3DES: While encryption protects confidentiality, it doesn't directly address the need to verify the code's integrity. The code could be decrypted and modified, and there would be no mechanism to detect that change without additional integrity checks. Also, 3DES is considered outdated and less secure than newer encryption algorithms.

F. Make the DACL read-only: Making the DACL (Discretionary Access Control List) read-only can prevent accidental changes to the code, but it doesn't prevent a user with sufficient privileges from modifying the DACL and then altering the code. It is more for data availability protection than integrity.

Question: 69

A company is moving most of its customer-facing production systems to the cloud-facing production systems to the cloud. IaaS is the service model being used.
The Chief Executive Officer is concerned about the type of encryption available and requires the solution must have the highest level of security.
Which of the following encryption methods should the cloud security engineer select during the implementation phase?

A. Instance-based
B. Storage-based
C. Proxy-based
D. Array controller-based

Answer: B

Explanation:

The correct answer is **B. Storage-based**. Here's a detailed justification:

When a company moves its production systems to the cloud using IaaS, it essentially gains control over the infrastructure elements, including storage. The CEO's concern regarding the highest level of security necessitates choosing an encryption method that provides comprehensive data protection at rest. Storage-based encryption directly encrypts the data as it's stored within the cloud storage service. This means that regardless of who gains access to the physical storage or the underlying infrastructure, the data remains encrypted and unusable without the proper decryption keys.

Instance-based encryption (A) focuses on encrypting the virtual machine instances themselves. While this adds a layer of security, it might not protect all data at rest, particularly data stored outside the instance's local storage.

Proxy-based encryption (C) typically involves encrypting data in transit, not at rest. While important for secure communication, it doesn't address the CEO's concern about protecting data residing in the cloud storage.

Array controller-based encryption (D) is more relevant in on-premise storage solutions and less applicable in IaaS cloud environments where the storage infrastructure is managed by the cloud provider.

Storage-based encryption provides the strongest level of security for data at rest in an IaaS environment because it encrypts the data directly within the storage system, irrespective of the underlying infrastructure vulnerabilities. This fulfills the CEO's requirement for the "highest level of security." The cloud provider will typically provide the tools and services needed to manage storage-based encryption, allowing the cloud security engineer to implement and maintain the encryption keys and policies.

Further Reading:

AWS Storage Encryption:https://aws.amazon.com/security/server-side-encryption/
Azure Storage Encryption:https://learn.microsoft.com/en-us/azure/storage/common/storage-service-encryption
Google Cloud Storage Encryption:https://cloud.google.com/storage/docs/encryption/customer-managed-keys

Question: 70

A vulnerability analyst identified a zero-day vulnerability in a company's internally developed software. Since the current vulnerability management system does not have any checks for this vulnerability, an engineer has been asked to create one.
Which of the following would be BEST suited to meet these requirements?

A. ARF
B. ISACs
C. Node.js
D. OVAL

Answer: D

Explanation:

The best solution for creating a vulnerability check for a zero-day vulnerability in internally developed software, where the existing vulnerability management system lacks the necessary checks, is OVAL (Open Vulnerability and Assessment Language). Here's why:

OVAL's Purpose: OVAL is a language designed to represent system configuration information and assess machine state for the presence of vulnerabilities or compliance issues. It is specifically built to automate the process of vulnerability assessment and compliance auditing.

Flexibility and Customization: OVAL provides a standardized way to define vulnerability tests, allowing security engineers to create custom checks tailored to specific vulnerabilities, even zero-days. Since the existing system doesn't have checks, OVAL allows you to build one specific to the zero-day threat.

Machine-Readability: OVAL definitions are machine-readable, meaning they can be processed and executed by various vulnerability scanners and assessment tools. This contrasts with static documents or alerts.

Platform Support: OVAL is designed to be platform-independent, meaning that a single OVAL definition can be used to assess vulnerabilities on multiple operating systems.

Community and Resources: While specifically defining the zero-day requires internal creation, OVAL benefits from a community and associated resources which can aid in the development and integration of the custom check.

Why other options are less suitable:

ARF (Asset Reporting Format): ARF is a reporting format for security assessments. It describes how the results of vulnerability scans and other security evaluations should be structured. It is for reporting not

definition of checks.

ISACs (Information Sharing and Analysis Centers): ISACs are industry-specific groups that share threat intelligence and security best practices. While helpful for awareness, they don't provide a technical means to define a vulnerability check.

Node.js: Node.js is a JavaScript runtime environment. While it could potentially be used to create a vulnerability scanner from scratch, it's a more complex and time-consuming approach compared to using a dedicated vulnerability definition language like OVAL. Furthermore, it doesn't provide a standardized format for defining vulnerability checks.

In summary, OVAL's design for defining system states, assessing vulnerabilities, and its flexibility in creating custom checks make it the ideal choice for addressing a zero-day vulnerability scenario where existing systems lack the required checks.

Authoritative Links:

OVAL Website:https://oval.cisecurity.org/
NIST - National Vulnerability Database (NVD):https://nvd.nist.gov/ (Often references OVAL definitions).
CIS Security:https://www.cisecurity.org/

Question: 71

An organization recently started processing, transmitting, and storing its customers' credit card information. Within a week of doing so, the organization suffered a massive breach that resulted in the exposure of the customers' information.

Which of the following provides the BEST guidance for protecting such information while it is at rest and in transit?

A. NIST
B. GDPR
C. PCI DSS
D. ISO

Answer: C

Explanation:

The correct answer is C, PCI DSS (Payment Card Industry Data Security Standard). Here's a detailed justification:

The scenario describes an organization handling credit card data and experiencing a data breach. This directly falls under the purview of PCI DSS. PCI DSS is a set of security standards designed to protect cardholder data and reduce credit card fraud. It provides specific requirements for securing credit card information both at rest (storage) and in transit (transmission).

PCI DSS mandates security controls like encryption, access controls, secure coding practices, regular security assessments, and incident response plans specifically tailored to environments handling payment card data. Organizations processing, storing, or transmitting credit card data are legally and contractually obligated to comply with PCI DSS. A failure to do so can result in significant fines, penalties, and reputational damage.

While NIST (National Institute of Standards and Technology) provides broader cybersecurity frameworks and guidance applicable to various types of data and organizations, it isn't specifically tailored to credit card data security like PCI DSS. GDPR (General Data Protection Regulation) focuses on the protection of personal data of EU citizens, which can include credit card information, but its scope is much wider than just payment card data and it doesn't offer as detailed guidance on card data security specifically. ISO (International Organization for Standardization) standards like ISO 27001 provide a framework for information security

management systems, but they lack the detailed, specific requirements for credit card data protection found in PCI DSS.

Therefore, given the scenario's focus on protecting credit card information, PCI DSS is the most relevant and appropriate guidance for ensuring the security of this sensitive data at rest and in transit. It is the industry standard and legal requirement for organizations handling payment card data.

Further Research:

PCI Security Standards Council:https://www.pcisecuritystandards.org/ **NIST Cybersecurity Framework:**https://www.nist.gov/cyberframework
GDPR:https://gdpr.eu/
ISO 27001:https://www.iso.org/isoiec-27001-information-security.html

Question: 72

Which of the following is the MOST important security objective when applying cryptography to control messages that tell an ICS how much electrical power to output?

A. Improving the availability of messages
B. Ensuring non-repudiation of messages
C. Enforcing protocol conformance for messages
D. Assuring the integrity of messages

Answer: D

Explanation:

The correct answer is D: Assuring the integrity of messages.

In the context of Industrial Control Systems (ICS) managing critical infrastructure like electrical power output, message integrity is paramount. These control messages dictate the amount of power generated or distributed, and any alteration, whether accidental or malicious, could have catastrophic consequences.

Availability (A) is important, but a corrupted message that is available is worse than a temporarily unavailable one. Non-repudiation (B) is useful for accountability after an incident, but preventing the incident in the first place is more crucial. Protocol conformance (C) ensures messages are properly formatted but doesn't guarantee their content hasn't been tampered with.

Integrity, achieved through cryptographic techniques such as hashing or digital signatures, ensures that the control messages received by the ICS are exactly as they were sent. If the integrity check fails, the system should reject the message, preventing potentially dangerous or damaging actions based on compromised instructions. For example, imagine a malicious actor altering a message to instruct a power plant to drastically increase output beyond safe limits. This could lead to equipment failure, blackouts, or even physical damage. Cryptographic integrity checks can prevent such scenarios. The focus must be on preventing manipulation to avoid cascading effects.

Ultimately, the objective in this scenario isn't about guaranteeing who sent the message or that the system operates without interruption. The primary goal is to guarantee the message hasn't been altered and is what the sender intended.

For further research:

NIST Special Publication 800-82, Guide to Industrial Control Systems (ICS) Security: This document covers security considerations for ICS, including the importance of message integrity.

https://csrc.nist.gov/publications/detail/sp/800-82/rev-2/final
SANS Institute ICS Security: SANS offers numerous resources and training on ICS security, including topics like cryptographic controls and message integrity. https://www.sans.org/industrial-control-systems-security/

Question: 73

A company wants to protect its intellectual property from theft. The company has already applied ACLs and DACs. Which of the following should the company use to prevent data theft?

A. Watermarking
B. DRM
C. NDA
D. Access logging

Answer: B

Explanation:

The correct answer is B. DRM (Digital Rights Management). Here's why:

DRM technologies are specifically designed to control access to and usage of digital content, thus directly addressing the need to protect intellectual property from theft. They accomplish this by embedding restrictions within the digital files themselves, dictating how the content can be used, copied, shared, or even accessed. This helps prevent unauthorized duplication and distribution.

Watermarking (A) can identify the owner or origin of a document but doesn't technically prevent copying; it just provides evidence of ownership if a copy is found. NDAs (C) are legal agreements that prohibit the disclosure of confidential information; while useful, they don't technologically prevent data theft. Access logging (D) tracks who accessed what resources and when. While useful for auditing and identifying potential breaches after they've occurred, it does not prevent the initial data theft.

DRM employs techniques like encryption, license management, and digital watermarks (as a secondary component) to enforce usage rules. A company wanting to protect its intellectual property would implement DRM solutions that directly control how its digital assets can be accessed, used, and shared. ACLs and DACs control access at the operating system/file system level, while DRM operates at the content level, providing a finer-grained level of control that survives even if the file is copied. It addresses the problem after it gets past ACLs and DACs, helping stop dissemination and usage beyond what the company permits.

Therefore, DRM is the most direct and effective solution to prevent data theft of intellectual property beyond access controls.

Relevant links for further research:

Digital Rights Management (DRM):https://en.wikipedia.org/wiki/Digital_rights_management **W3C - Encrypted Media Extensions:**https://www.w3.org/TR/encrypted-media/

Question: 74

A satellite communications ISP frequently experiences outages and degraded modes of operation over one of its legacy satellite links due to the use of deprecated hardware and software. Three days per week, on average, a contracted company must follow a checklist of 16 different high-latency commands that must be run in serial to restore nominal performance. The ISP wants this process to be automated.
Which of the following techniques would be BEST suited for this requirement?

A. Deploy SOAR utilities and runbooks.

B. Replace the associated hardware.

C. Provide the contractors with direct access to satellite telemetry data.

D. Reduce link latency on the affected ground and satellite segments.

Answer: A

Explanation:

The best approach to automate the restoration process of the satellite link is to deploy SOAR (Security Orchestration, Automation, and Response) utilities and runbooks (Option A). SOAR platforms are designed to automate repetitive tasks, especially those involving multiple steps and disparate systems. In this scenario, a SOAR platform can be configured with a runbook that executes the 16 high-latency commands in the required sequence, eliminating the need for manual intervention by the contracted company. The SOAR platform can integrate with the existing satellite communication infrastructure and execute the commands through APIs or scripts. Automation reduces the time it takes to restore performance and minimizes the impact of outages.

Replacing the hardware (Option B), although a potential long-term solution, addresses the root cause but does not provide an immediate automated solution for the existing process. Providing contractors with telemetry data (Option C) is useful for monitoring, but does not automate the restoration process. Reducing link latency (Option D) might improve overall performance, but does not eliminate the need to execute the 16 commands when the system degrades. Therefore, SOAR is the most suitable approach for immediate automation and improved efficiency.

SOAR platforms leverage automation and orchestration to streamline security operations, incident response, and threat management. Runbooks are pre-defined workflows that automate tasks and procedures. This allows organizations to improve efficiency, reduce response times, and free up security analysts to focus on more complex tasks.

Supporting Links:

NIST Special Publication 800-207, Zero Trust Architecture: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf (While not directly SOAR focused, it highlights the importance of automation in modern security architectures, a key element that SOAR utilizes) **SANS Institute - Security Orchestration, Automation, and Response (SOAR):**https://www.sans.org/white-papers/security-orchestration-automation-response-soar/ (Provides a general overview of SOAR and its benefits)

Question: 75

A company processes data subject to NDAs with partners that define the processing and storage constraints for the covered data. The agreements currently do not permit moving the covered data to the cloud, and the company would like to renegotiate the terms of the agreements.
Which of the following would MOST likely help the company gain consensus to move the data to the cloud?

A. Designing data protection schemes to mitigate the risk of loss due to multitenancy

B. Implementing redundant stores and services across diverse CSPs for high availability

C. Emulating OS and hardware architectures to blur operations from CSP view

D. Purchasing managed FIM services to alert on detected modifications to covered data

Answer: A

Explanation:

The best answer is **A. Designing data protection schemes to mitigate the risk of loss due to multitenancy.**

Here's why: The core issue is convincing partners to allow data migration to the cloud, despite existing NDA constraints. Partners' primary concern is likely the security and confidentiality of their data when stored in a cloud environment.

Addressing Security Concerns: Multitenancy, a common cloud characteristic where multiple tenants share the same infrastructure, introduces potential risks of data leakage or unauthorized access. Designing data protection schemes specifically aimed at mitigating these multitenancy-related risks directly addresses the partners' main concern. This approach demonstrates a proactive commitment to maintaining data confidentiality and integrity in the cloud.

Building Trust and Confidence: By focusing on security, the company is essentially saying, "We understand your concerns about cloud security, and we're taking concrete steps to address them." This will likely build more trust and confidence in the company's ability to protect the data in the cloud.

Relevant Cloud Concepts: Data protection schemes could include:

Encryption: Encrypting data at rest and in transit using strong encryption algorithms ensures that even if unauthorized access occurs, the data remains unreadable.

Data Masking: Masking sensitive data elements replaces actual values with realistic but fake values, protecting sensitive information while still allowing for testing or analysis.

Tokenization: Replacing sensitive data with non-sensitive tokens, which are then stored securely separately from the actual data.

Access Controls: Implementing strict role-based access control (RBAC) policies and multi-factor authentication (MFA) ensures that only authorized individuals can access the data.

Data Loss Prevention (DLP): Implementing DLP measures to detect and prevent sensitive data from leaving the cloud environment.

Let's also consider why the other options are less suitable:

B. Implementing redundant stores and services across diverse CSPs for high availability: While high availability is important, it doesn't directly address the partners' primary concern about data security and confidentiality. It is more related to business continuity and disaster recovery.

C. Emulating OS and hardware architectures to blur operations from CSP view: This is a more complex and potentially risky approach. It might even create more security vulnerabilities. Furthermore, it's unlikely to alleviate partners' concerns about data security.

D. Purchasing managed FIM services to alert on detected modifications to covered data: This approach is reactive, meaning it only alerts to changes after they occur. Partners are likely to want proactive measures to prevent unauthorized modifications in the first place.

Authoritative Links:

NIST Special Publication 800-144: "Guidelines on Security and Privacy in Public Cloud Computing" (https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-144.pdf)
Cloud Security Alliance (CSA): (https://cloudsecurityalliance.org/) Offers a wealth of information on cloud security best practices.

Question: 76

Ransomware encrypted the entire human resources fileshare for a large financial institution. Security operations personnel were unaware of the activity until it was too late to stop it. The restoration will take approximately four hours, and the last backup occurred 48 hours ago. The management team has indicated that the
RPO for a disaster recovery event for this data classification is 24 hours.
Based on RPO requirements, which of the following recommendations should the management team make?

A. Leave the current backup schedule intact and pay the ransom to decrypt the data.

B. Leave the current backup schedule intact and make the human resources fileshare read-only.

C. Increase the frequency of backups and create SIEM alerts for IOCs.

D. Decrease the frequency of backups and pay the ransom to decrypt the data.

Answer: C

Explanation:

The correct answer is C. Here's why:

The scenario highlights a failure to meet the Recovery Point Objective (RPO). The RPO of 24 hours means that the financial institution can tolerate losing a maximum of 24 hours' worth of data. In this incident, the last backup was 48 hours ago, exceeding the acceptable data loss.

Option A is incorrect because paying the ransom is generally discouraged. There's no guarantee the data will be decrypted, and it encourages future attacks. Also, it does not address the root cause: an inadequate backup strategy.

Option B is incorrect because making the fileshare read-only might prevent further encryption, but it doesn't recover the lost data or address the unmet RPO. It's a reactive measure, not a preventative one.

Option D is incorrect because decreasing the backup frequency would only worsen the RPO issue, increasing the potential data loss in future incidents. Furthermore, paying the ransom has the same downsides as option A.

Option C, increasing the backup frequency, directly addresses the RPO failure. By backing up more frequently (e.g., every 12 or 24 hours), the organization reduces the potential data loss to within the acceptable 24-hour window. Creating SIEM alerts for Indicators of Compromise (IOCs) adds a proactive layer of security by helping detect and respond to ransomware attacks earlier. This prevents similar incidents from reaching the point of complete encryption and data loss. A SIEM (Security Information and Event Management) solution helps centralize security logs and event data for analysis.

Therefore, the management team should recommend increasing the frequency of backups to align with the 24-hour RPO and creating SIEM alerts for early detection of malicious activity. This approach addresses the immediate issue of data loss vulnerability and proactively reduces the likelihood of future ransomware attacks.For further research, consult the following links:

NIST Cybersecurity Framework:https://www.nist.gov/cyberframework
SANS Institute on RPO and RTO:https://www.sans.org/ (Search their site for "RPO" and "RTO" for relevant articles and resources)
CISA on Ransomware:https://www.cisa.gov/ransomware

Question: 77

A company undergoing digital transformation is reviewing the resiliency of a CSP and is concerned about meeting SLA requirements in the event of a CSP incident.
Which of the following would be BEST to proceed with the transformation?

A. An on-premises solution as a backup

B. A load balancer with a round-robin configuration

C. A multicloud provider solution

D. An active-active solution within the same tenant

Answer: D

Explanation:

The best approach for a company concerned about CSP incident resiliency during a digital transformation is D. An active-active solution within the same tenant. This is because an active-active setup provides immediate failover capabilities. In an active-active configuration, multiple instances of an application or service are running simultaneously and actively serving requests. If one instance fails due to a CSP incident, the other active instance(s) can seamlessly take over, minimizing downtime and ensuring SLA requirements are met.

While a multicloud provider solution (C) offers geographic diversity and reduces vendor lock-in, it introduces complexity in managing multiple environments and ensuring data consistency. Setting up and maintaining effective failover between different cloud providers can be challenging. An on-premises solution as a backup (A) also adds complexity and cost, potentially requiring significant infrastructure investments and ongoing maintenance.

A load balancer with a round-robin configuration (B) simply distributes traffic; it doesn't inherently provide high availability unless the backend servers themselves are highly available and fault-tolerant. Round-robin can become detrimental if one of the backend servers is unavailable but the load balancer continues to send requests to it. An active-active solution can use a load balancer, but the key aspect is the application's capability to seamlessly switch between running instances.

An active-active setup within the same tenant, offers simpler management compared to a multicloud setup while delivering superior availability compared to a single instance or a traditional active-passive setup. This architecture is specifically designed for rapid failover and high availability within the existing cloud environment. It allows for the application to immediately continue serving requests from another instance running and ready within the same environment. Therefore, addressing the company's main concern about maintaining SLA during a CSP incident.

Authoritative Links:

Active-Active vs. Active-Passive:https://www.nginx.com/blog/high-availability-load-balancing-active-passive-active-active/ (Explains the differences and benefits of active-active setups)
High Availability Architectures:https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/design-for-high-availability.html (AWS documentation on designing highly available systems)

Question: 78

A company has hired a security architect to address several service outages on the endpoints due to new malware. The Chief Executive Officer's laptop was impacted while working from home. The goal is to prevent further endpoint disruption. The edge network is protected by a web proxy.
Which of the following solutions should the security architect recommend?

A. Replace the current antivirus with an EDR solution.
B. Remove the web proxy and install a UTM appliance.
C. Implement a deny list feature on the endpoints.
D. Add a firewall module on the current antivirus solution.

Answer: A

Explanation:

The correct answer is A: Replace the current antivirus with an EDR solution. Here's a detailed justification:

The scenario highlights a critical need for improved endpoint protection due to recent malware outbreaks causing service disruptions, even affecting the CEO's laptop while working remotely. While the company uses a web proxy, it's clearly insufficient to prevent all threats. Traditional antivirus solutions often rely on

signature-based detection, which struggles to identify new or polymorphic malware variants effectively. This makes them reactive rather than proactive.

An Endpoint Detection and Response (EDR) solution provides a much more comprehensive approach to endpoint security. EDR solutions continuously monitor endpoint behavior, collect detailed data about processes, network connections, and file system activity, and use advanced analytics, including machine learning, to detect suspicious patterns and anomalies that might indicate malware infections or malicious activity. They can identify threats that traditional antivirus might miss. This behavioral analysis is crucial for detecting zero-day exploits and advanced persistent threats (APTs).

Furthermore, EDR solutions often include response capabilities that allow security teams to isolate infected endpoints, block malicious processes, and remediate threats remotely. They also integrate with threat intelligence feeds, providing real-time updates on emerging threats and helping to prioritize security efforts. For remote workers, EDR is invaluable as it extends security beyond the corporate network.

Option B is less effective. Removing the web proxy weakens perimeter security. While a UTM appliance consolidates several security functions, its endpoint visibility might be limited compared to a dedicated EDR solution. Option C, a deny list, is insufficient as it only blocks known bad files and is easily bypassed by new malware. Option D, adding a firewall module, is a step in the right direction but doesn't provide the advanced detection and response capabilities of EDR.

In summary, EDR offers proactive threat detection, behavioral analysis, and incident response capabilities, making it the most effective solution to prevent further endpoint disruptions from malware. Its ability to identify and respond to sophisticated threats, even on remote devices, makes it the superior choice.

For further reading:

NIST SP 800-123 Guide to General Server Security:https://csrc.nist.gov/publications/detail/sp/800-123/rev-1/final
SANS Institute Reading Room on Endpoint Security:https://www.sans.org/reading-room/whitepapers/endpoint/

Question: 79

All staff at a company have started working remotely due to a global pandemic. To transition to remote work, the company has migrated to SaaS collaboration tools. The human resources department wants to use these tools to process sensitive information but is concerned the data could be:
☞ Leaked to the media via printing of the documents
☞ Sent to a personal email address
Accessed and viewed by systems administrators

☞ Uploaded to a file storage site
Which of the following would mitigate the department's concerns?

- A. Data loss detection, reverse proxy, EDR, and PGP
  B. VDI, proxy, CASB, and DRM
  C. Watermarking, forward proxy, DLP, and MFA
  D. Proxy, secure VPN, endpoint encryption, and AV

Answer: B

Explanation:

VDI, proxy, CASB, and DRM

Question: 80

A home automation company just purchased and installed tools for its SOC to enable incident identification and response on software the company develops. The company would like to prioritize defenses against the following attack scenarios:

☞ Unauthorized insertions into application development environments

☞ Authorized insiders making unauthorized changes to environment configurations

Which of the following actions will enable the data feeds needed to detect these types of attacks on development environments? (Choose two.)

A. Perform static code analysis of committed code and generate summary reports.

B. Implement an XML gateway and monitor for policy violations.

C. Monitor dependency management tools and report on susceptible third-party libraries.

D. Install an IDS on the development subnet and passively monitor for vulnerable services.

E. Model user behavior and monitor for deviations from normal.

F. Continuously monitor code commits to repositories and generate summary logs.

Answer: CD

Explanation:

The correct answer is **A and F**. Here's why:

A. Perform static code analysis of committed code and generate summary reports: Static code analysis examines the source code before it's executed. It can identify potential vulnerabilities, coding errors, and security flaws that could lead to unauthorized insertions into application development environments. By generating summary reports, the SOC can prioritize and address critical issues proactively. This aligns directly with detecting unauthorized insertions before they become exploitable.

F. Continuously monitor code commits to repositories and generate summary logs: Monitoring code commits provides real-time visibility into changes made to the codebase. By analyzing commit messages, author information, and the code diffs, the SOC can detect unauthorized or suspicious modifications.

Summary logs offer a consolidated view of these activities, allowing for efficient investigation of potential security breaches or insider threats related to unauthorized changes. This is particularly useful in identifying authorized insiders making unauthorized changes.

Why other options are less suitable:

B. Implement an XML gateway and monitor for policy violations: An XML gateway is primarily used to secure and manage XML-based web services and APIs. It's not directly related to monitoring code commits or analyzing code for vulnerabilities within a development environment.

C. Monitor dependency management tools and report on susceptible third-party libraries: While important for overall security posture and vulnerability management, focusing solely on vulnerable third-party libraries doesn't directly address unauthorized insertions or unauthorized insider changes to configurations. It's more focused on external vulnerabilities.

D. Install an IDS on the development subnet and passively monitor for vulnerable services: While an Intrusion Detection System (IDS) can detect attacks in progress, it is reactive. It would only alert after an unauthorized insertion had already occurred. Static code analysis and monitoring code commits are more proactive and can prevent the attack.

E. Model user behavior and monitor for deviations from normal: While helpful in detecting anomalies, modeling user behavior and monitoring for deviations is not a direct approach to detecting unauthorized insertions into application development or unauthorized configuration changes. It can raise red flags but is not the primary solution.

In summary, static code analysis and continuous monitoring of code commits are the most effective actions to detect the specified attack scenarios because they focus on the source code itself and changes made to it, allowing for proactive identification and prevention of unauthorized activities.