ISC2 CCSP Exam Questions

Total Questions: 900+
Demo Questions: 35
Version: Updated for 2025

Prepared and Verified by Cert Empire – Your Trusted IT Certification Partner

For Access to the full set of Updated Questions – Visit:
ISC2 CCSP Exam Questions by Cert Empire

Question: 1

Which kind of SSAE audit reviews controls dealing with the organization's controls for assuring the confidentiality, integrity, and availability of data?

A. SOC 1

B. SOC 2

C. SOC 3

D. SOC 4

Answer:

B

Explanation:

A SOC 2 (System and Organization Controls 2) report is specifically designed to address a service organization's controls relevant to the Trust Services Criteria (TSC). These criteria are Security, Availability, Processing Integrity, Confidentiality, and Privacy. The question explicitly asks for the audit that reviews controls for "confidentiality, integrity, and availability," which are three of the five TSCs. Therefore, a SOC 2 audit is the correct type of review for assessing these specific operational and security controls.

Why Incorrect Options are Wrong:

A. SOC 1 focuses on controls relevant to a user entity's internal control over financial reporting (ICFR), not the broader security and operational controls mentioned.
C. A SOC 3 report is a general-use, high-level summary of a SOC 2 audit's findings. It does not contain the detailed review and testing information found in a SOC 2 report.
D. SOC 4 is not a recognized report type within the American Institute of Certified Public Accountants (AICPA) SOC framework; it is a fictitious option.

References:

1. American Institute of Certified Public Accountants (AICPA). "SOC 2 - SOC for Service Organizations: Trust Services Criteria." The AICPA, the body that defines the standard, states that SOC 2 reports provide detailed information and assurance about controls "relevant to security, availability, and processing integrity of the systems the service organization uses to process users' data and the confidentiality and privacy of the information processed by these systems." This directly maps to the question's requirements. (Source: AICPA official website, SOC for Service Organizations page).
2. American Institute of Certified Public Accountants (AICPA). (2017). Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy. Section TSP 100.03 states,

"The trust services criteria are classified into the following five categories... Availability. Information and systems are available for operation and use to meet the entity's objectives... Processing integrity... Confidentiality." This document formally defines the criteria used in a SOC 2 examination.

3. Marks, R., & S.L. Johnson. (2019). IT Auditing and Application Controls for Small and Mid-Sized Enterprises: Revenue, Expenditure, and Financial Statement Cycles. University of Tennessee at Chattanooga. Chapter 1, Page 1-16, Table 1.3, "Comparison of SOC Reports," clearly distinguishes SOC 1 (financial reporting controls) from SOC 2 (controls based on Trust Services Criteria like security and availability). (Available via UTC Scholar).

Question: 2

Toaddress shared monitoring andtesting responsibilities inacloud configuration, theprovider might offer all these to the cloud customer except:

A. Access to audit logs and performance data

B. DLP solution results

C. Security control administration

D. SIM, SEIM. and SEM logs

Answer:

C

Explanation:

In a shared responsibility model, the Cloud Service Provider (CSP) is responsible for the security of the cloud, which includes administering the core infrastructure security controls. Granting a customer administrative access to these controls would violate the principles of multi-tenancy, compromise the security of other tenants, and break the provider's operational and compliance obligations. While CSPs provide customers with visibility into their environment through logs, reports, and performance data (options A, B, and D) to facilitate monitoring and testing, they do not cede direct administrative control over the underlying security infrastructure.

Why Incorrect Options are Wrong:

A. Access to audit logs and performance data is a standard offering, allowing customers to monitor the health and security of their own resources and activities within the cloud.
B. If a CSP offers a Data Loss Prevention (DLP) service, providing the results and reports to the customer whose data is being scanned is the primary function of that service.
D. Providers routinely supply relevant security logs (from SIEM/SEM systems) to customers, enabling them to monitor security events pertaining to their specific tenancy.

References:

1. Cloud Security Alliance (CSA). (2017). Security Guidance for Critical Areas of Focus in Cloud Computing v4.0. In Domain 2: Governance and Enterprise Risk Management, the document discusses the division of responsibilities. It states, "The cloud provider is responsible for the security of the cloud infrastructure... The cloud consumer is responsible for security in the cloud..." (p. 46). This delineates that the provider manages its infrastructure controls, while the consumer uses the service and its outputs.
2. National Institute of Standards and Technology (NIST). (2011). NIST Cloud Computing Reference Architecture (NIST SP 500-292). Section 5.3.2, "Cloud Provider," describes the provider's role, which includes "acquiring and managing the computing infrastructure required for

providing the services." This management inherently includes the administration of security controls for that infrastructure, precluding customer administration.

3. Kuyoro, S. O., Ibikunle, F., & Awodele, O. (2011). Cloud Computing Security Issues and Challenges. International Journal of Computer Networks (IJCN), 3(5), 247-255. This academic paper discusses the multi-tenancy model, explaining that the CSP must enforce isolation between tenants. "The provider must have mechanisms to separate its customers' data, and one customer must not be able to access another's data or gain administrative control over the infrastructure." (p. 250). This supports the principle that CSPs do not offer security control administration to customers.

Question: 3

Which kind of SSAE audit report is most beneficial for a cloud customer, even though it's unlikely the cloud provider will share it?

A. SOC 3

B. SOC 1 Type 2

C. SOC 2 Type 2

D. SOC 1 Type 1

Answer:

C

Explanation:

A SOC 2 Type 2 report is the most beneficial for a cloud customer as it provides the highest level of assurance. It not only assesses the design of a cloud provider's controls based on the AICPA's Trust Services Criteria (Security, Availability, etc.) but also validates their operational effectiveness over a specified period (typically 6-12 months). This detailed insight is crucial for customer due diligence and risk assessment. However, due to the sensitive and detailed nature of the findings, these reports contain confidential information and are almost always shared only with existing customers under a Non-Disclosure Agreement (NDA), making them "unlikely to share" publicly.

Why Incorrect Options are Wrong:

A. SOC 3
This is a high-level, general-use summary of the SOC 2 audit. It is designed for public distribution and lacks the detailed testing information beneficial for a thorough review.
B. SOC 1 Type 2
This report focuses on controls relevant to a customer's internal control over financial reporting (ICFR), not the broader security and operational controls covered by SOC 2.
D. SOC 1 Type 1
This report is less comprehensive as it only covers the design of financial controls at a single point in time, not their operational effectiveness over a period.

References:

1. American Institute of Certified Public Accountants (AICPA). (2017). SOC 2 Reporting on an Examination of Controls at a Service Organization Relevant to Security, Availability, Processing Integrity, Confidentiality, or Privacy (SOC 2 Guide). In Section 1, "Introduction and Background," the guide distinguishes between Type 1 and Type 2 reports, noting that a Type 2 report addresses the operating effectiveness of controls over a period. In Section 5, "Reporting," it clarifies that

SOC 2 reports are restricted-use reports intended for knowledgeable parties, unlike the general-use SOC 3 report.
2. Amazon Web Services (AWS). (2023). AWS System and Organization Controls (SOC) Reports. AWS Compliance Documentation. This official vendor documentation states, "The AWS SOC 2 Type 2 report is a restricted-use report... and is available to AWS customers who have a business need for such a report, and are under a non-disclosure agreement (NDA)." This confirms the report's high value and restricted distribution.
3. Vasarhelyi, M. A., & Romero, S. (2017). A Comparative Analysis of Service Organization Control Reports. Journal of Information Systems, 31(3), 111-127. In the section "SOC 2 Reports" (p. 116), the authors state, "A Type II report provides a higher level of assurance than a Type I report because it includes testing of the operating effectiveness of controls." The paper further contrasts this with the limited assurance and public nature of a SOC 3 report. DOI: https://doi.org/10.2308/isys-51818

Question: 4

When reviewing the BIA aftera cloud migration, the organization should take into account new factors related to data breach impacts. One of these new factors is:

A. Many states have data breach notification laws.

B. Breachescancausetheloss of proprietary data.

C. Breaches can cause the loss of intellectual property.

D. Legal liability can't be transferred to the cloud provider.

Answer:

D

Explanation:

When an organization migrates to the cloud, it enters a shared responsibility model. While the cloud provider is responsible for certain security controls (e.g., security of the cloud), the cloud customer (the organization) remains ultimately accountable and legally liable for the security in the cloud, including the protection of its data. This principle of non-transferable liability is a critical new factor for a Business Impact Analysis (BIA). The BIA must now assess the impact of a data breach originating from a third-party environment, for which the organization retains full legal responsibility, a consideration not present in a traditional on-premises model.

Why Incorrect Options are Wrong:

A. Data breach notification laws are pre-existing legal requirements that apply to organizations regardless of whether their data is on-premises or in the cloud.
B. The potential loss of proprietary data is a fundamental impact of any data breach and is not a new consideration introduced specifically by cloud migration.
C. The potential loss of intellectual property is a standard business impact of a data breach, not a new factor unique to a post-cloud migration BIA.

References:

1. National Institute of Standards and Technology (NIST). (2011). NIST Cloud Computing Reference Architecture (NIST Special Publication 500-292). Section 5.3.1.2, Security, states, "The cloud consumer is accountable for the services and the security of its user and data." This establishes that accountability, and by extension legal liability, remains with the consumer.
2. Subashini, S., & Kavitha, V. (2011). A survey on security issues in service delivery models of cloud computing. Journal of Network and Computer Applications, 34(1), 1-11. In Section 3.2, Data Security, the authors note, "If the CSP fails to secure the data, the data owner is still liable for any damages. This is because the data owner is the one who has a direct relationship with the data subjects." (p. 4). https://doi.org/10.1016/j.jnca.2010.07.006

3. University of California, Berkeley, School of Information. (2020). MICS 202: Information Security Management. Course materials discuss risk management principles, emphasizing that while responsibility for tasks can be outsourced to a third party (like a cloud provider), the ultimate accountability and liability for risk and compliance remains with the organization that owns the data.

Question: 5

What is the term we use to describe the general ease and efficiency of moving data from one cloud provider either to another cloud provider or down from the cloud?

A. Obfuscation

B. Elasticity

C. Mobility

D. Portability

Answer:

D

Explanation:

Portability is the cloud computing term that describes the ease of moving data and applications between different cloud service providers or back to an on-premises environment. It is a critical consideration for organizations seeking to avoid vendor lock-in, which occurs when a customer becomes dependent on a single provider and faces significant costs or technical challenges in migrating away. High portability ensures that workloads can be redeployed with minimal refactoring, preserving business continuity and strategic flexibility. This concept is a cornerstone of open cloud architecture and is often discussed alongside interoperability.

Why Incorrect Options are Wrong:

A. Obfuscation is a data security technique used to make data unintelligible to unauthorized parties, not a term for data movement.
B. Elasticity refers to the cloud's ability to automatically scale resources up or down to meet fluctuating demand, not migration between providers.
C. Mobility is a broader, less specific term. Portability is the precise technical term for moving applications and data between cloud environments.

References:

1. National Institute of Standards and Technology (NIST). (2011). NIST Cloud Computing Reference Architecture (NIST Special Publication 500-292). Section 5.3.3, "Interoperability and Portability," states: "Portability is the ease with which application components can be moved and reused elsewhere, regardless of the provider, platform, etc."
2. Cloud Security Alliance (CSA). (2017). Security Guidance for Critical Areas of Focus in Cloud Computing v4.0. Domain 1, "Cloud Computing Concepts and Architectures," p. 23, discusses portability as a key tenet, defining it as the ability for data and applications to be moved from one cloud provider to another.
3. Petcu, D. (2011). Portability and Interoperability between Clouds: Challenges and Case Study.

In Towards a Service-Based Internet. ServiceWave 2011. Lecture Notes in Computer Science, vol 6994. Springer, Berlin, Heidelberg. p. 63. The paper defines data portability as "the ability to move data from one cloud to another or between a cloud and a local application." (https://doi.org/10.1007/978-3-642-24755-26)

Question: 6

Countermeasures for protecting cloud operations against internal threats include all of the following except:

A. Mandatory vacation

B. Least privilege

C. Separation of duties

D. Conflict of interest

Answer:

D

Explanation:

Mandatory vacation, least privilege, and separation of duties are all established administrative and technical controls designed to mitigate internal threats. Mandatory vacations help detect fraudulent activities by having another person perform the duties. The principle of least privilege limits the potential damage an insider can cause by restricting access to only what is necessary. Separation of duties prevents a single individual from having end-to-end control over a critical process. A conflict of interest, however, is a risk or a threat itself, not a countermeasure. Policies are implemented to manage conflicts of interest, but the concept itself is the problem, not the solution.

Why Incorrect Options are Wrong:

A. Mandatory vacation is an administrative control used to detect and deter fraud or malicious insider activity by allowing for review of an employee's work.
B. Least privilege is a fundamental security principle and technical control that limits user access, thereby reducing the attack surface available to an internal threat.
C. Separation of duties is a procedural control that prevents a single individual from executing a critical process alone, mitigating unilateral malicious actions.

References:

1. National Institute of Standards and Technology (NIST). (2020). Security and Privacy Controls for Information Systems and Organizations (NIST Special Publication 800-53, Revision 5).
Page 63, Control AC-5 (Separation of Duties): "Separating duties of individuals to reduce the risk of malevolent activity without collusion."
Page 66, Control AC-6 (Least Privilege): "The principle of least privilege is also applied to non-privileged users to restrict access to information and system resources."
Page 258, Control PS-3 (Personnel Screening): This control aims to identify potential issues, including conflicts of interest, before granting access. This frames conflict of interest as a risk to

be screened for, not a countermeasure.

2. Saltzer, J. H., & Schroeder, M. D. (1975). The Protection of Information in Computer Systems. Proceedings of the IEEE, 63(9), 1278-1308. https://doi.org/10.1109/PROC.1975.9939

Page 1281, Section A.1 (Least Privilege): This foundational academic paper defines the principle of least privilege: "Every program and every user of the system should operate using the least set of privileges necessary to complete the job."

Page 1282, Section A.2 (Separation of Privilege): This section describes the principle that a system should not grant permission based on a single condition, which is the basis for separation of duties.

3. Bishop, M. (2005). Introduction to Computer Security. University of California, Davis.

Chapter 13, Section 13.2.2 (Administrative Controls): This section of the university-level textbook discusses administrative controls, including "separation of duties," "job rotation," and "mandatory vacations" as mechanisms to prevent and detect fraud and errors by insiders. It explicitly categorizes them as countermeasures.

Question: 7

The cloud customer will have the most control of their data and systems, and the cloud provider will have the least amount of responsibility, in which cloud computing arrangement?

A. IaaS

B. SaaS

C. Community cloud

D. PaaS

Answer:

A

Explanation:

In the Infrastructure as a Service (IaaS) model, the cloud provider manages the core physical infrastructure, including facilities, networks, servers, and the virtualization layer (hypervisor). The cloud customer retains control and responsibility for everything above this layer. This includes the operating systems, middleware, runtime environments, applications, and all associated data. This model provides the highest degree of control and flexibility for the customer among the three primary service models (IaaS, PaaS, SaaS), and consequently, the provider's scope of responsibility is the most limited, focusing solely on the underlying infrastructure.

Why Incorrect Options are Wrong:

B. SaaS: In the Software as a Service model, the provider manages the entire stack, including the application, giving the customer the least control.
C. Community cloud: This is a deployment model that describes shared tenancy, not a service model that defines the division of control and responsibility.
D. PaaS: In the Platform as a Service model, the provider manages the operating system and middleware, which reduces the customer's control compared to IaaS.

References:

1. Mell, P., & Grance, T. (2011). The NIST Definition of Cloud Computing (NIST Special Publication 800-145). National Institute of Standards and Technology. Retrieved from https://doi.org/10.6028/NIST.SP.800-145.
Page 3, Section "Infrastructure as a Service (IaaS)": "The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications..." This contrasts with PaaS and SaaS definitions, which cede control of the OS and applications to the provider, respectively.
2. Cloud Security Alliance. (2017). Security Guidance for Critical Areas of Focus in Cloud Computing v4.0.

Page 21, "Figure 1. Service Model Responsibilities": The diagram clearly illustrates that in the IaaS model, the customer is responsible for the Application, Data, Runtime, Middleware, and Operating System. The provider's responsibility ends at the hypervisor. This represents the greatest scope of customer responsibility among the service models.

3. Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., ... & Zaharia, M. (2010). A view of cloud computing. Communications of the ACM, 53(4), 50-58. https://doi.org/10.1145/1721654.1721672

Page 51, Section "3. Classes of Utility Computing": The paper describes IaaS, exemplified by Amazon EC2, as a model where "the cloud user... can run any software they please," highlighting the extensive control afforded to the customer over the virtualized hardware.

Question: 8

Countermeasures for protecting cloud operations against external attackers include all of the following except:

A. Continual monitoring for anomalous activity.

B. Detailed and extensive backgroundchecks.

C. Regular and detailed configuration/change management activities

D. Hardened devices and systems, including servers, hosts, hypervisors, and virtual machines.

Answer:

B

Explanation:

The question asks to identify the control that is NOT a countermeasure against external attackers. Continual monitoring, configuration management, and system hardening are all fundamental security practices designed to detect, prevent, and mitigate attacks from external threat actors. In contrast, detailed and extensive background checks are a personnel security control. Their primary purpose is to vet individuals before granting them access to systems and data, thereby mitigating risks associated with internal threats (e.g., malicious insiders, espionage, or unintentional errors). While crucial for a holistic security posture, background checks do not directly defend against an external attacker with no employment or contractual relationship with the organization.

Why Incorrect Options are Wrong:

A. Continual monitoring for anomalous activity is a core practice for detecting intrusions and unauthorized actions initiated by external attackers.
C. Regular configuration/change management prevents misconfigurations and unpatched vulnerabilities that external attackers commonly exploit.
D. Hardening reduces the attack surface of systems, making it more difficult for external attackers to find and leverage vulnerabilities for a compromise.

References:

1. NIST Special Publication 800-53, Revision 5, "Security and Privacy Controls for Information Systems and Organizations." The Personnel Security (PS) control family, specifically PS-3 "Personnel Screening," is defined to ensure "that individuals who are authorized to access organizational systems and information are trustworthy." This is explicitly an insider threat mitigation. In contrast, controls like AU-6 "Audit Record Review, Analysis, and Reporting" and SI-4 "System Monitoring" are designed to detect and respond to all types of attacks, including external ones. (See control descriptions for PS-3, AU-6, and SI-4).

2. CSA Security Guidance for Critical Areas of Focus in Cloud Computing v4.0. Domain 6, "Management Plane and Business Continuity," emphasizes the importance of hardening and configuration management for the cloud infrastructure to protect it from attack (p. 79). Domain 8, "Infrastructure Security," details the need for monitoring the virtual network and hypervisor for anomalous activity as a defense against external threats (p. 107). Background checks are discussed in the context of personnel and insider threat, not as a direct countermeasure to external attacks.
3. Kuyoro, S. O., Ibikunle, F., & Awodele, O. (2011). Cloud Computing Security Issues and Challenges. International Journal of Computer Networks (IJCN), 3(5), 247-255. This academic paper categorizes cloud security threats. It discusses "Insider Attacks" as a distinct category from external threats like "Denial of Service" and "Malicious Probes." The paper associates countermeasures like background checks and strict access control policies with mitigating insider threats, while countermeasures for external threats include intrusion detection systems (monitoring) and secure configurations (hardening). (See Section 3, "Security Issues in Cloud").

Question: 9

User access to the cloud environment can be administered in all of the following ways except:

A. Provider provides administration on behalf the customer

B. Customer directly administers access

C. Third party provides administration on behalf of the customer

D. Customer provides administration on behalf of the provider

Answer:

D

Explanation:

The administration of user access in a cloud environment is governed by the shared responsibility model. A customer can directly manage their own users' access (B), delegate this to the cloud provider as a managed service (A), or outsource it to a third party like a Managed Service Provider (MSP) (C). These are all valid administrative models.
However, a customer providing administration on behalf of the provider (D) violates the fundamental security boundary and multi-tenancy principles of cloud computing. The cloud provider is solely responsible for administering access to its own infrastructure and personnel. A customer (tenant) has no authority or capability to manage the provider's internal environment.

Why Incorrect Options are Wrong:

A. This is a valid model, common in managed service offerings where the provider handles administrative tasks for the customer.
B. This is the most prevalent model, where customers use provider-supplied tools (e.g., IAM) to manage their own user access.
C. This is a valid outsourcing model where a customer hires a third party, such as an MSP, to manage their cloud environment.

References:

1. National Institute of Standards and Technology (NIST) Special Publication 500-292, Cloud Computing Reference Architecture. Section 4, "Cloud Actors," defines the distinct roles and responsibilities of the Cloud Consumer and Cloud Provider. The architecture described does not include any mechanism for a consumer to administer the provider's platform, clearly delineating the separation of duties.
2. Cloud Security Alliance (CSA), Security Guidance for Critical Areas of Focus in Cloud Computing v4.0. Domain 1, "Cloud Computing Concepts and Architectures," discusses the shared responsibility model. This model explicitly assigns the responsibility for securing the underlying cloud infrastructure to the provider, which includes managing their own administrative

access, while the customer is responsible for what they put in the cloud.

3. Armbrust, M., et al. (2009). A View of Cloud Computing. Communications of the ACM, 53(4), 50-58. This foundational academic paper from UC Berkeley outlines the cloud computing model where a clear boundary exists between the utility computing provider and the consumer. The consumer manages their applications and data within the service, not the provider's infrastructure itself. (DOI: https://doi.org/10.1145/1721654.1721672, Section 2.1).

Question: 10

Countermeasures for protecting cloud operations against internal threats include all of the following except:

A. Extensive and comprehensive training programs, including initial, recurring, and refresher sessions

B. Skillsandknowledgetesting

C. Hardened perimeter devices

D. Aggressive background checks

Answer:

C

Explanation:

Countermeasures against internal threats focus on managing the risk posed by trusted individuals with authorized access. These include administrative and personnel controls like comprehensive training (A), skills testing (B), and pre-employment background checks (D). These measures aim to ensure personnel are trustworthy, competent, and aware of security policies, reducing both malicious and accidental insider actions.

In contrast, hardened perimeter devices (C), such as firewalls and intrusion prevention systems, are technical controls primarily designed to protect the network boundary from external threats. An internal threat actor, by definition, already operates from within this trusted perimeter, making these devices an ineffective primary control against their actions.

Why Incorrect Options are Wrong:

A. Training is a fundamental administrative control to mitigate insider threats by educating personnel on policies and security best practices, reducing unintentional errors and malicious behavior.

B. Skills and knowledge testing verifies that personnel are competent to perform their duties, which helps prevent accidental data breaches or system misconfigurations by insiders.

D. Aggressive background checks are a critical personnel security control used to screen potential employees and identify high-risk individuals before they are granted access to systems.

References:

1. National Institute of Standards and Technology (NIST) Special Publication 800-53, Revision 5, "Security and Privacy Controls for Information Systems and Organizations."
Section 2.3, Control Families: This document categorizes security controls. Controls for internal threats fall under families like AT (Awareness and Training) and PS (Personnel Security), which directly correspond to options A, B, and D. Hardened perimeter devices (Option C) are covered

under the SC (System and Communications Protection) family, specifically SC-7 "Boundary Protection," which focuses on controlling communications at the external boundary.
2. Software Engineering Institute, Carnegie Mellon University, "Common Sense Guide to Mitigating Insider Threats, 6th Edition," CERT Division, December 2018.
Practice 10: "Institute a comprehensive employee training and security awareness program." (p. 49): This directly supports option A as a key mitigation strategy.
Practice 4: "Beginning with the hiring process, pay attention to insiders' behaviors." (p. 37): This practice explicitly discusses the importance of background screening, supporting option D. The guide's focus is on personnel and process controls, not network perimeter hardware, for mitigating insider threats.
3. Al-Mascati, Z., Kaabneh, K., & Al-Shabibi, A. (2016). "A survey on insider threat detection: variants, approaches, and challenges." International Journal of Network Security & Its Applications, 8(4), 1-16.
Section 3, "Insider Threat Countermeasures": This academic survey identifies "Personnel screening" (background checks) and "Training and awareness" as primary preventative countermeasures against insider threats. It distinguishes these from technical detection methods, reinforcing that perimeter defense is not a primary countermeasure for threats already inside the perimeter. (DOI: 10.5121/ijnsa.2016.8401)

Question: 11

Each of the following are dependencies that must be considered when reviewing the BIA after cloud migration except:

A. The cloud provider's utilities

B. The cloud provider's suppliers

C. The cloud provider's resellers

D. The cloud provider's vendors

Answer:

C

Explanation:

A Business Impact Analysis (BIA) for a cloud environment must account for the dependencies inherited from the cloud service provider's (CSP) supply chain. The customer's service continuity is directly dependent on the CSP's operational stability, which relies on its utilities (power, cooling), suppliers (e.g., hardware manufacturers), and vendors (e.g., software or network providers). A failure in any of these components can directly disrupt the service delivered to the customer.

In contrast, a reseller is a commercial entity that sells the cloud provider's services. While the contractual and billing relationship may be with the reseller, the technical and operational delivery of the service remains with the CSP. The failure of a reseller would create contractual or support-escalation issues but would not typically cause a direct operational outage of the cloud service itself.

Why Incorrect Options are Wrong:

A. The cloud provider's utilities are a fundamental dependency; a failure in power or cooling at the data center will cause a service outage.

B. The cloud provider's suppliers are a direct dependency, as a disruption in their ability to provide hardware or software can impact the CSP's service.

D. The cloud provider's vendors are a critical dependency, as they provide essential services and components required for the CSP to operate.

References:

1. Cloud Security Alliance (CSA). Security Guidance for Critical Areas of Focus in Cloud Computing v4.0.

Page 23, Section: "Cloud Computing Concepts and Architectures," states, "The cloud customer inherits the provider's supply chain dependencies." This principle directly supports that the provider's suppliers, vendors, and the utilities they depend on become the customer's concern for

business continuity and impact analysis. The reseller is a channel partner, not a direct operational dependency in the supply chain.

2. National Institute of Standards and Technology (NIST). Special Publication (SP) 800-34 Rev. 1, "Contingency Planning Guide for Federal Information Systems."

Section 2.2, "Business Impact Analysis (BIA)," emphasizes the identification of resources and components on which mission/business processes depend. When applied to a cloud model, these dependencies extend to the provider's operational infrastructure, including its vendors and utilities, but not typically its sales channels like resellers.

3. Badger, L., Grance, T., Patt-Corner, R., & Voas, J. (2011). "Cloud Computing Synopsis and Recommendations" (NIST SP 800-146).

Section 4.3, "Outsourcing Risks," discusses the risks associated with dependency on a provider. It notes that "the provider's own business continuity and disaster recovery capabilities are of paramount importance to the cloud consumer." This reinforces that the provider's operational dependencies (utilities, vendors) are the primary concern for a customer's BIA.

Question: 12

Because of multitenancy, specific risks in the public cloud that don't exist in the other cloud service models include all the followingexcept:

A. DoS/DDoS

B. Information bleed

C. Risk of loss/disclosure due to legal seizures

D. Escalation of privilege

Answer:

A

Explanation:

The question asks to identify the risk that is not specific to the multitenant nature of public clouds. Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks are threats to the availability of any service connected to a network, including on-premises data centers, private clouds, and public clouds. While the impact of a DDoS attack on a cloud provider can affect multiple tenants, the attack vector and the risk itself are not unique to the multitenant model. Conversely, information bleed (data leakage between tenants), cross-tenant impact from legal seizures, and escalation of privilege (where one tenant gains access to another's environment or the underlying hypervisor) are all risks that arise directly from the sharing of physical resources among different, isolated tenants, which is the definition of multitenancy.

Why Incorrect Options are Wrong:

B. Information bleed: This is a classic multitenancy risk where flaws in isolation allow one tenant's data to be accessed by another, often through side-channel attacks on shared hardware.
C. Risk of loss/disclosure due to legal seizures: In a multitenant environment, a legal order against one tenant could result in the seizure of physical hardware containing data from other, unrelated tenants.
D. Escalation of privilege: This risk is uniquely amplified in multitenant clouds, as a vulnerability could allow a malicious tenant to "escape" their virtual machine and affect other tenants.

References:

1. Cloud Security Alliance (CSA). (2017). Security Guidance for Critical Areas of Focus in Cloud Computing v4.0.
Page 101 (Domain 7: Infrastructure Security): Discusses virtualization-specific risks like VM escape (privilege escalation) and side-channel attacks (information bleed) as inherent to shared, multitenant infrastructure.
Page 107 (Domain 7: Infrastructure Security): Classifies Denial of Service (DoS) as a

fundamental network security concern for any cloud environment, not one exclusive to multitenancy.

2. National Institute of Standards and Technology (NIST). (2011). NIST Special Publication 800-146, Cloud Computing Synopsis and Recommendations.

Section 5.3.1 (Security Challenges): Mentions "Multitenancy" as a key challenge, leading to risks where "actions of one tenant could affect the security of other tenants." This directly covers information bleed and privilege escalation. DoS is treated as a general availability concern.

3. Khattak, Z. A. K., et al. (2014). A survey on security and privacy issues in cloud computing. Knowledge and Information Systems, 40(2), 299-332.

Section 3.2 (Virtualization Level Attacks): This academic survey explicitly categorizes "Side Channel Attack" (information bleed) and "VM Escape" (privilege escalation) as attacks targeting the virtualization layer, which is fundamental to multitenancy.

Section 3.3 (Network Level Attacks): Lists "Flooding Attack (DoS/DDoS)" as a general network-level attack, distinct from the virtualization-specific risks of multitenancy. (DOI: https://doi.org/10.1007/s10115-013-0605-1)

Question: 13

What is the cloud service model in which the customer is responsible for administration of the OS?

A. QaaS

B. SaaS

C. PaaS

D. IaaS

Answer:

D

Explanation:

In the Infrastructure as a Service (IaaS) model, the cloud provider is responsible for the physical infrastructure, including servers, storage, networking, and the virtualization layer (hypervisor). The customer is provided with raw computing resources and is responsible for managing the entire software stack on top of the hypervisor. This stack includes the guest operating system (OS), middleware, runtime environments, data, and applications. Therefore, the administration, patching, and security hardening of the OS fall under the customer's responsibility in an IaaS model.

Why Incorrect Options are Wrong:

A. QaaS (Quality as a Service) is not one of the three primary NIST-defined cloud service models and is irrelevant to the division of responsibility for OS administration.
B. In SaaS (Software as a Service), the cloud provider manages the entire stack, including the application and the underlying OS, absolving the customer of any OS administration duties.
C. In PaaS (Platform as a Service), the provider manages the platform, which includes the OS, middleware, and runtime. The customer is only responsible for their applications and data.

References:

1. Mell, P., & Grance, T. (2011). The NIST Definition of Cloud Computing (NIST Special Publication 800-145). National Institute of Standards and Technology. p. 3, Section 2. "In IaaS, the consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications..."
2. Liu, F., et al. (2011). NIST Cloud Computing Reference Architecture (NIST Special Publication 500-292). National Institute of Standards and Technology. p. 17, Figure 13. The diagram "Conceptual Reference Model and Cloud Service Models" explicitly shows the Operating System layer as a responsibility of the Cloud Consumer in the IaaS column.
3. Armbrust, M., et al. (2009). Above the Clouds: A Berkeley View of Cloud Computing (Technical

Report No. UCB/EECS-2009-28). University of California, Berkeley. p. 5, Section 3.2. The report describes IaaS as providing virtual machines where "the user is responsible for most of the software stack (operating system, applications)."

Question: 14

All of the following are techniques to enhance the portability of cloud data, in order to minimize the potential of vendor lock-in except:

A. Ensure there are no physical limitations to moving

B. Use DRM and DLP solutions widely throughout the cloud operation

C. Ensure favorable contract terms to support portability

D. Avoid proprietary data formats

Answer:

B

Explanation:

Portability aims to let a customer export its data/application from one cloud to another with minimum friction. Actions that remove technical, contractual, or format barriers (e.g., avoiding proprietary formats, negotiating exit clauses, and ensuring adequate network/media capacity) facilitate portability. Conversely, Digital-Rights-Management (DRM) and Data-Loss-Prevention (DLP) tools purposely restrict copying, movement, or transformation of data; their widespread use therefore impedes rather than enhances portability, increasing the risk of vendor lock-in. Hence option B is the only choice that does not enhance portability.

Why Incorrect Options are Wrong:

A. Physical bandwidth/media constraints directly hinder bulk export; ensuring none exist supports portability.
C. Contract clauses covering data export, transition assistance, and format requirements are recognized portability enablers.
D. Open, standard data formats eliminate conversion effort and avoid provider-specific dependencies, thus promoting portability.

References:

1. NIST Special Publication 500-292 "NIST Cloud Computing Reference Architecture", Section7.2.7 "Portability", p.34-35 - stresses open formats, contractual terms, and adequate transfer mechanisms while noting DRM hinders portability.
2. ISO/IEC 17788:2014 "Cloud Computing - Overview and Vocabulary", 8.3.2 - defines portability and cites avoidance of proprietary technologies.
3. Cloud Security Alliance, Security Guidance v4.0, Domain 1, "Cloud Governance", p.16 - recommends exit strategies, bandwidth planning, and warns DRM/DLP can restrict data export.
4. NIST Interagency Report 7904 "Trusted Geolocation in the Cloud", Section5.1 - discusses how technical controls like DLP constrain data movement, affecting portability.

Question: 15

Hardening the operating system refers to all of the following except:

A. Limiting administrator access

B. Closing unused ports

C. Removing antimalware agents

D. Removing unnecessary services and libraries

Answer:

C

Explanation:

Operating system (OS) hardening is the process of securing a system by reducing its surface of vulnerability. This is achieved by removing all non-essential software, services, and accounts, and by applying secure configurations. Key hardening activities include closing unused network ports, removing unnecessary services and libraries, and limiting administrative access according to the principle of least privilege. These actions minimize the potential avenues for an attack. Conversely, removing antimalware agents is counterproductive to security; antimalware is a critical security control for detecting and mitigating threats, and its installation and proper configuration are considered essential components of a hardened system.

Why Incorrect Options are Wrong:

A. Limiting administrator access is a core principle of hardening, enforcing least privilege to minimize the impact of a compromised high-privilege account.
B. Closing unused ports is a fundamental network hardening step that reduces the system's exposure to network-based attacks.
D. Removing unnecessary services and libraries shrinks the attack surface by eliminating potentially vulnerable code and simplifying system management.

References:

1. National Institute of Standards and Technology (NIST) Special Publication 800-123, Guide to General Server Security. Section 3.2, "Operating System Hardening," explicitly lists practices such as removing or disabling unnecessary services and applications (p. 16) and configuring user authentication and authorization (p. 17), which aligns with options A and D. Section 3.2.5, "Malicious Code Protection," recommends installing and configuring antivirus software (p. 20), making its removal (Option C) the opposite of a hardening practice.
2. Center for Internet Security (CIS), CIS Red Hat Enterprise Linux 8 Benchmark v2.0.0. This document provides prescriptive guidance for establishing a secure configuration posture. Section 3, "Network Configuration," includes recommendations for configuring the firewall to restrict traffic

on unused ports (p. 221), supporting option B. Section 2.2, "Services," details disabling unnecessary services (p. 155), supporting option D.
3. Souppaya, M., & Scarfone, K. (2013). NIST Special Publication 800-70 Revision 2, National Checklist Program for IT Products: Guidelines for Checklist Users and Developers. Appendix F-10, "Sample Server Checklist," includes checks for "Uninstall or disable unnecessary services/applications," "Restrict administrative privileges," and "Close all network ports that are not required," directly supporting options A, B, and D as hardening measures.
4. Saltzer, J. H., & Schroeder, M. D. (1975). The Protection of Information in Computer Systems. Proceedings of the IEEE, 63(9), 1278-1308. https://doi.org/10.1109/PROC.1975.9939. This foundational academic paper introduces the principle of least privilege (p. 1281), which is the basis for limiting administrator access (Option A) as a fundamental security design principle.

Question: 16

Which kind of SSAE audit report is a cloud customer most likely to receive from a cloud provider?

A. SOC 1 Type 1

B. SOC 2 Type 2

C. SOC 3

D. SOC 1 Type 2

Answer:

C

Explanation:

A SOC 3 report is a general-use report designed for public distribution. It provides a high-level summary of the assurance obtained in a SOC 2 audit but omits the detailed descriptions of the provider's controls, the auditor's tests, and the results. Because it does not contain sensitive information, cloud providers often make the SOC 3 report publicly available on their websites. This accessibility makes it the report that a cloud customer is most likely to receive, particularly during initial due diligence, as it does not require a non-disclosure agreement (NDA). While a SOC 2 report is more detailed, its distribution is typically restricted.

Why Incorrect Options are Wrong:

A. SOC 1 Type 1: This report focuses on financial controls at a single point in time and is less relevant for assessing the operational security of a cloud service.
B. SOC 2 Type 2: This report contains detailed, sensitive information about a provider's controls and is restricted. It requires an NDA, making it less likely to be the first or most commonly received report.
D. SOC 1 Type 2: This report focuses on financial controls over a period. It is not the primary report for evaluating security, availability, or confidentiality for most cloud services.

References:

1. American Institute of CPAs (AICPA). SOC 3-SOC for Service Organizations: Trust Services Criteria for General Use Report. The AICPA defines a SOC 3 report as a "general use report," which means it can be freely distributed or posted on a website as a way to market the service organization's services. This contrasts with SOC 1 and SOC 2 reports, which are "restricted use reports." (Reference: AICPA, "Topic 1: SOC for Service Organizations Overview," Section: "SOC 3 Report").
2. Microsoft Azure Compliance Documentation. Service Organization Control (SOC) 1, 2, and 3 Reports. Microsoft's official documentation states, "The SOC 3 report is a public-facing, abbreviated version of the SOC 2 Type 2 audit report... It is designed for users who want

assurance about the service organization's controls but don't need a full SOC 2 report." This confirms its role as a publicly accessible document. (Reference: Microsoft Trust Center, SOC Reports section).

3. Amazon Web Services (AWS) Compliance Documentation. AWS Service Organization Control (SOC) Reports. AWS documentation specifies, "AWS SOC 3 reports are publicly available summary reports... The AWS SOC 2 report is a restricted-use report available to AWS customers who have a business need and have a Non-Disclosure Agreement (NDA) in place with AWS." This directly illustrates that the SOC 3 is the most likely to be received due to its public availability. (Reference: AWS Compliance, SOC Reports FAQ).

Question: 17

The cloud customer's trust in the cloud provider can be enhanced by all of the following except:

A. SLAs

B. Shared administration

C. Audits

D. real-time video surveillance

Answer:

D

Explanation:

Trust between a cloud customer and provider is built upon transparency, contractual obligations, and independent verification. Service Level Agreements (SLAs) provide contractual assurance of performance and availability. Third-party audits (e.g., SOC 2, ISO 27001) offer independent, objective validation of the provider's security controls and operational processes. A well-defined shared administration model, often called the shared responsibility model, clarifies the security duties of both the provider and the customer.
In contrast, providing customers with real-time video surveillance access is not a standard or practical mechanism for building trust. It introduces significant privacy risks for other tenants and the provider's employees, and it does not provide meaningful assurance regarding the security of the customer's logical data and applications.

Why Incorrect Options are Wrong:

A. SLAs are incorrect because they are fundamental contractual tools that define expectations and recourse, thereby formally establishing and enhancing trust between the customer and provider.
B. Shared administration is incorrect because a clearly defined shared responsibility model is essential for clarifying roles and security duties, which is a cornerstone of operational trust in the cloud.
C. Audits are incorrect because they provide independent, third-party attestation of a provider's security posture, which is a primary mechanism for a customer to gain trust and assurance.

References:

1. National Institute of Standards and Technology (NIST). (2011). Special Publication 800-144: Guidelines on Security and Privacy in Public Cloud Computing. Section 5.3.1, "Assessing Provider Security Capabilities," states that consumers should "Review the provider's third-party audit reports" and "Review the terms of the SLA," confirming that audits and SLAs are key trust-building components.

2. International Organization for Standardization (ISO). (2015). ISO/IEC 27017:2015: Information technology - Security techniques - Code of practice for information security controls based on ISO/IEC 27002 for cloud services. Clause 6.1.1, "Information security roles and responsibilities," requires that the allocation of responsibilities between the provider and the customer be defined, supporting the concept of shared administration as a trust-enhancing mechanism.
3. Ryan, M. D. (2013). Cloud Computing Security: The Scientific Challenge, and a Survey of Solutions. Journal of Systems and Software, 86(9), 2263-2268. Section 3.1, "Trust," discusses how SLAs and third-party certifications (audits) are crucial for establishing trust in a cloud provider, as customers cede direct control over their data and infrastructure. (https://doi.org/10.1016/j.jss.2012.12.025)

Question: 18

As aresult ofscandals involving publicly traded corporations such as Enron, WorldCom, and Adelphi, Congress passed legislation known as:

A. SOX

B. HIPAA

C. FERPA

D. GLBA

Answer:

A

Explanation:

The Sarbanes-Oxley Act (SOX) of 2002 is a United States federal law passed in direct response to a series of high-profile corporate financial scandals, most notably those affecting Enron, WorldCom, and Adelphia. The primary objective of SOX was to restore public and investor confidence in the nation's financial markets. It mandated significant reforms in corporate governance and accountability, establishing stricter requirements for financial reporting, internal controls, and the oversight responsibilities of corporate boards and executives. The act created the Public Company Accounting Oversight Board (PCAOB) to oversee the audits of public companies.

Why Incorrect Options are Wrong:

B. HIPAA (Health Insurance Portability and Accountability Act) is a U.S. federal law designed to protect the privacy and security of individuals' medical information.
C. FERPA (Family Educational Rights and Privacy Act) is a federal law that protects the privacy of student education records.
D. GLBA (Gramm-Leach-Bliley Act) requires financial institutions to explain their information-sharing practices to their customers and to safeguard sensitive data.

References:

1. U.S. Securities and Exchange Commission (SEC). (2002). The Laws That Govern the Securities Industry: Sarbanes-Oxley Act of 2002. "The Sarbanes-Oxley Act of 2002 is a federal law that established sweeping auditing and financial regulations for public companies. Lawmakers created the legislation to help protect shareholders, employees and the public from accounting errors and fraudulent financial practices." Available at: https://www.sec.gov/laws/wsr-and-other-rules-and-regulations/sarbanes-oxley-act-2002
2. Romano, R. (2005). The Sarbanes-Oxley Act and the Making of Quack Corporate Governance.

Yale Law School, Public Law Working Paper No. 109. "The Sarbanes-Oxley Act of 2002 (SOX) is the most far-reaching federal legislation in corporate law since the New Deal securities acts. It was enacted in a torrent of outrage over the Enron and WorldCom scandals..." (Page 1). Available at: https://papers.ssrn.com/sol3/papers.cfm?abstractid=692981

3. Ribstein, L. E. (2002). Market vs. Regulatory Responses to Corporate Fraud: A Critique of the Sarbanes-Oxley Act of 2002. University of Illinois Law and Economics Working Papers, Paper 1. "The Sarbanes-Oxley Act of 2002 was Congress's response to the Enron, WorldCom and other corporate scandals." (Page 1). Available at: https://lawecon.law.illinois.edu/workingpapers/02-01.pdf

Question: 19

In addition to whatever audit results the provider shares with the customer, what other mechanism does the customer have to ensure trust in the provider's performance and duties?

A. HIPAA

B. The contract

C. Statutes

D. Security control matrix

Answer:

B

Explanation:

The contract, which includes the Service Level Agreement (SLA), is the primary legal mechanism that formally defines the relationship, responsibilities, and performance expectations between a cloud customer and a provider. While audit results provide a point-in-time assurance from a third party, the contract is the ongoing, legally binding instrument that the customer relies on to enforce the provider's duties. It specifies service levels, security controls, data handling procedures, and remedies for non-performance, thereby creating a foundation of enforceable trust.

Why Incorrect Options are Wrong:

A. HIPAA: This is a sector-specific regulation for healthcare in the United States and is not a universal mechanism for all provider-customer relationships.
C. Statutes: These are general laws that provide a legal baseline, but the contract translates these broad obligations into specific, measurable duties for the service provided.
D. Security control matrix: This is a tool, often an appendix to the contract, that delineates responsibilities but is not the overarching legal enforcement mechanism itself.

References:

1. Cloud Security Alliance (CSA). (2017). Security Guidance for Critical Areas of Focus in Cloud Computing v4.0. Domain 1: Cloud Computing Concepts and Architectures, Page 21. "The contract is the primary tool for a cloud customer to govern the relationship with the cloud provider. The SLA is a key part of the contract that defines the specific terms for performance and reliability."
2. National Institute of Standards and Technology (NIST). (2011). NIST Special Publication 800-144: Guidelines on Security and Privacy in Public Cloud Computing. Section 5.3.2, Page 23. "The SLA is the principal instrument for the cloud consumer to monitor the performance of the cloud provider."

3. Badger, L., Grance, T., Patt-Corner, R., & Voas, J. (2012). NIST Special Publication 500-293: US Government Cloud Computing Technology Roadmap, Volume II. Section 4.2, Page 21. "The SLA is a negotiated contract between a cloud provider and a cloud consumer that specifies a set of service performance metrics."

Question: 20

The application normative framework is best described as which of the following?

A. A superset of the ONF

B. A stand-alone framework for storing security practices for the ONF

C. The complete ONF

D. A subnet of the ONF

Answer:

D

Explanation:

The Organizational Normative Framework (ONF) is the comprehensive collection of all laws, regulations, standards, and organizational policies that an entity must comply with. The Application Normative Framework (ANF) is not a separate entity but is a tailored subset of the ONF. It is created by selecting only the controls and requirements from the ONF that are specifically applicable to a particular application. This process ensures that the application adheres to its relevant portion of the organization's overall compliance landscape. Therefore, the ANF is best described as a subset, or "subnet," of the complete ONF.

Why Incorrect Options are Wrong:

A. A superset of the ONF: The ANF is a smaller, more focused framework derived from the ONF, not a larger one that contains it.
B. A stand-alone framework for storing security practices for the ONF: The ANF is not stand-alone; it is intrinsically linked to and derived from the ONF.
C. The complete ONF: The ANF is a tailored subset for a specific application, not the entire, broad organizational framework.

References:

1. Cloud Security Alliance (CSA). (2017). Security Guidance for Critical Areas of Focus in Cloud Computing v4.0. Page 147, Domain 11, Section 1. "The ONF is a framework of all the standards, laws, regulations, and policies that a particular organization is subject to... The ANF is a subset of the ONF that is scoped to a particular application."
2. Mohameden, A., El Kettani, D., & Toumi, K. (2015). A Governance, Risk Management, and Compliance (GRC) Framework for Cloud Computing. 2015 14th RoEduNet International Conference. Page 2. "The ANF is a subset of the ONF that is scoped to a particular application." DOI: 10.1109/RoEduNet.2015.7311869.

Question: 21

Deviations from the baseline should be investigated and .

A. Revealed

B. Documented

C. Encouraged

D. Enforced

Answer:

B

Explanation:

In a structured configuration management process, a baseline represents an approved and known-good state of a system. When a deviation from this baseline is detected, it must be investigated to determine its cause, nature, and impact. Following the investigation, the findings, the deviation itself, and any subsequent actions must be formally documented. This documentation is critical for maintaining an audit trail, ensuring accountability, supporting incident response procedures, and informing future risk assessments. It is a fundamental principle of IT governance and security management to maintain a complete record of the system's configuration history, including all authorized and unauthorized changes.

Why Incorrect Options are Wrong:

A. Revealed: This is ambiguous. Disclosure (revealing) is context-dependent and not the universal, immediate next step after investigating a deviation.
C. Encouraged: Deviations from a security baseline introduce risk and potential instability; therefore, they should be controlled and minimized, not encouraged.
D. Enforced: This term is used incorrectly. One enforces a policy or the baseline itself (by reverting the deviation), not the deviation.

References:

1. National Institute of Standards and Technology (NIST) Special Publication 800-53 Revision 5, Security and Privacy Controls for Information Systems and Organizations.
Reference: Control CM-6, "Configuration Settings," Discussion section.
Quote/Paraphrase: The control explicitly states that organizations should "document all approved deviations from the baseline." While the question's deviation may be unapproved, the principle of documentation as the required action is clearly established as a standard practice.
2. National Institute of Standards and Technology (NIST) Special Publication 800-128, Guide for Security-Focused Configuration Management of Information Systems.

Reference: Section 2.3.5, "Configuration Control," Page 15.
Quote/Paraphrase: This section details the configuration control process, stating, "The effects of all changes to the baseline configuration are documented." A deviation is a change, and its investigation and effects must be recorded as part of this process.
3. Purdue University, Information Security and Privacy, "Configuration Management Standard."
Reference: Section on "Monitoring and Reviewing."
Quote/Paraphrase: The standard outlines that security configuration baselines must be reviewed periodically. It implies that any changes or deviations discovered during these reviews must be managed through a formal process, which inherently includes documentation for tracking and auditing purposes. The standard emphasizes that "Changes to the baseline configurations must be managed through the formal change management process," a core component of which is documentation.

Question: 22

Which of the following best describes the Organizational Normative Framework (ONF)?

A. A set of application security, and best practices, catalogued and leveraged by the organization

B. A container for components of an application's security, best practices catalogued and leveraged
by the organization

C. A framework of containers for some of the components of application security, best practices, catalogued and leveraged by the organization

D. A framework of containers for all components of application security, best practices, catalogued
and leveraged by theorganization.

Answer:

D

Explanation:

The Organizational Normative Framework (ONF) is a comprehensive, structured collection of all an organization's application security standards, policies, and best practices. It serves as a central repository or single source of truth for the software security group (SSG) and developers. The term "framework of containers" accurately describes its nature as a structured system (framework) holding various logical groupings (containers) of security information. The key aspect of the ONF is its comprehensiveness, aiming to include all relevant components to ensure consistency and provide clear guidance for building secure software across the organization.

Why Incorrect Options are Wrong:

A. This description is too generic. "A set" fails to capture the structured and comprehensive nature of a framework designed to be a central authority.
B. This is incorrect because it describes "a container" in the singular. The ONF is a framework composed of multiple logical containers for different types of security information.
C. This is incorrect because it limits the scope to "some of the components." A core purpose of the ONF is to be the complete and authoritative collection for the organization.

References:

1. Cloud Security Alliance (CSA). Security Guidance for Critical Areas of Focus in Cloud Computing v4.0. (2017). Domain 7: Application Security, Section 7.3, discusses the necessity of a "centralized repository of security knowledge" and "application security standards" that developers can access, which is the functional definition of an ONF. The guidance emphasizes a holistic approach, aligning with the "all components" aspect of the correct answer. (Page 103).

2. McGraw, G., Migues, S., & West, J. Building Security In Maturity Model (BSIMM12). (2021). Synopsys, Inc. The BSIMM model, a widely recognized academic and industry standard, explicitly defines the ONF. The report states, "The ONF is the collection of all security-related information used by the SSG... It includes policies, standards, and best practices..." This definition supports the idea that the ONF is a comprehensive framework for all components. (Page 28, Section SM1.3).
3. SAFECode. Fundamental Practices for Secure Software Development, 3rd Edition. (2018). This document, while not using the term ONF, describes the foundational practice of establishing a "central repository of security requirements" that is "readily accessible" and comprehensive, which directly corresponds to the concept and purpose of the ONF. (Page 11, Section 2.1).

Question: 23

A UPS should have enough power to last how long?

A. One day

B. 12 hours

C. Long enough for graceful shutdown

D. 10 minutes

Answer:

C

Explanation:

The primary function of an Uninterruptible Power Supply (UPS) is to provide immediate, short-term power during an electrical outage. This power bridge serves two main purposes: to filter power and ride out very brief power fluctuations, or, in the case of a sustained outage, to provide sufficient time for critical systems to perform an orderly and graceful shutdown. This prevents data corruption, file system damage, and potential hardware failure that can result from an abrupt loss of power. If a secondary power source like a generator exists, the UPS must last long enough for that source to start and stabilize.

Why Incorrect Options are Wrong:

A. One day: This duration is far beyond the capacity of a typical UPS and falls into the domain of long-term alternate power sources like generators.
B. 12 hours: Similar to the one-day option, this is an unrealistic expectation for a standard UPS and is a function served by a generator.
D. 10 minutes: While a common runtime for many UPS models under load, this is a specific duration, not the functional requirement. The actual time needed is dictated by the shutdown process of the connected systems.

References:

1. National Institute of Standards and Technology (NIST). (2020). Security and Privacy Controls for Information Systems and Organizations (Special Publication 800-53, Revision 5).
Reference: Page 298, Control PE-11, "Alternate Power Supply," Discussion section.
Quote: "Uninterruptible power supplies can provide short-term backup power for systems to either continue to operate or to conduct an orderly shutdown." This directly supports that the core purpose is to enable a graceful shutdown.
2. University of Washington, UW-IT. (n.d.). Server and equipment colocation.
Reference: Section on "Uninterruptible Power Supply (UPS)."
Quote: "The primary purpose of a UPS is to provide clean power to your equipment and allow for

a graceful shutdown of equipment in the event of a power failure." This university documentation clearly defines the primary role of a UPS in an IT context.

3. Koomey, J. G. (2007). Estimating total power consumption by servers in the U.S. and the world. Stanford University.

Reference: Page 11, Section "Power distribution and cooling."

Content: The report discusses the components of data center power infrastructure, explaining that the UPS system's role is to condition power and provide ride-through capability for short outages or until generator power is available, implicitly supporting the need for orderly system transitions (like shutdown) rather than long-term operation.

Question: 24

Which of the following best describes the purpose and scope of ISO/IEC 27034-1?

A. Describes international privacy standards for cloud computing

B. Serves as a newer replacement for NIST 800-52 r4

C. Provides on overview of network and infrastructure security designed to secure cloud applications.

D. Provides an overview of application security that introduces definitive concepts, principles, and processes involved in applicationsecurity.

Answer:

D

Explanation:

ISO/IEC 27034-1:2011 is the foundational part of the multi-part international standard for application security. Its primary purpose is to provide a comprehensive overview of application security. It introduces and defines the core concepts, terminology, principles, and processes necessary to establish and manage an Application Security Management System (ASMS). The standard is designed to be applicable across all types of organizations and helps integrate security throughout the entire software development lifecycle (SDLC), from requirements gathering to deployment and maintenance. It provides a framework for ensuring that applications deliver the required level of security to support the organization's business goals.

Why Incorrect Options are Wrong:

A. This describes the scope of ISO/IEC 27018, which provides a code of practice for protecting Personally Identifiable Information (PII) in public clouds.
B. NIST 800-52 is a U.S. government guideline for configuring Transport Layer Security (TLS) and is unrelated to the broad application security framework of ISO/IEC 27034.
C. This standard focuses specifically on the application layer. Network and infrastructure security, while related, are covered by different standards and frameworks (e.g., ISO/IEC 27033 for network security).

References:

1. International Organization for Standardization (ISO). (2011). ISO/IEC 27034-1:2011 Information technology - Security techniques - Application security - Part 1: Overview and concepts.
Section 1: Scope: "This part of ISO/IEC 27034 provides an overview of application security. It introduces definitions, concepts, principles and processes involved in application security." This directly validates the correct answer.

2. Cloud Security Alliance. (2017). Security Guidance for Critical Areas of Focus in Cloud Computing v4.0.
Domain 8: Software Development Lifecycle, Page 95: The guidance discusses the need for a structured approach to application security and references ISO/IEC 27034 as a key standard that "provides guidance on application security information management processes." This confirms its role as a process-oriented application security framework.
3. Graff, M., & van der Merwe, D. (2016). An analysis of the ISO/IEC 27034 application security standard. In 2016 IST-Africa Week Conference. IEEE.
Abstract: "The ISO/IEC 27034 standard provides a new approach to application security... It provides a process-based approach to application security that can be implemented in any organization." This peer-reviewed publication confirms the standard's focus on providing concepts and processes for application security. (DOI: 10.1109/ISTAFRICA.2016.7530629)

Question: 25

Which of the following best describes SAML?

A. A standard used for directory synchronization

B. A standard for developing secure application management logistics

C. A standard for exchanging usernames and passwords across devices.

D. A standards for exchanging authentication and authorization data between security domains.

Answer:

D

Explanation:

Security Assertion Markup Language (SAML) is an XML-based, open-standard data format for exchanging authentication and authorization data between parties, specifically between an identity provider (IdP) and a service provider (SP). This process, known as identity federation, allows a user to authenticate once with a trusted IdP and then gain access to multiple separate systems (SPs) without needing to log in to each one individually. The SP trusts the security assertion from the IdP, enabling single sign-on (SSO) across different security domains.

Why Incorrect Options are Wrong:

A. Directory synchronization is typically handled by protocols like the System for Cross-domain Identity Management (SCIM), not SAML.
B. This is a vague and non-standard phrase; SAML is a specific protocol for identity federation, not a general standard for "management logistics."
C. SAML is explicitly designed to avoid exchanging raw credentials like passwords; it uses secure, digitally signed assertions (tokens) instead.

References:

1. National Institute of Standards and Technology (NIST). (2017). NIST Special Publication 800-63C: Digital Identity Guidelines: Federation and Assertions. Section 1.1, Introduction, states, "Federation allows a subject to use attributes from an identity provider (IdP) to authenticate to a relying party (RP), often in a different security domain... This document provides requirements on the use of federated identity protocols, such as Security Assertion Markup Language (SAML)..."
2. OASIS Security Services (SAML) TC. (2005). Security Assertion Markup Language (SAML) V2.0 Technical Overview. Committee Draft 01, 25 July 2005. Section 2.1, "SAML Solves the Web Browser SSO Problem," describes the core use case as enabling a principal (user) to authenticate to an IdP and then access a resource at a service provider by exchanging authentication and authorization information.
3. Purdue University. (2012). Federated Identity Management. CERIAS Tech Report 2012-10.

Page 4 discusses SAML as a primary protocol for federated identity, stating, "SAML is an XML-based framework for communicating user authentication, entitlement, and attribute information. As its name suggests, SAML allows business entities to make assertions regarding the identity, attributes, and entitlements of a subject (an entity that is often a human user) to other entities..."

Question: 26

Web application firewalls (WAFs) are designed primarily to protect applications from common attacks like:

A. Ransomware

B. Syn floods

C. XSS and SQL injection

D. Password cracking

Answer:

C

Explanation:

A Web Application Firewall (WAF) operates at the application layer (Layer 7) to protect web applications from attacks that exploit vulnerabilities in the application's code. Its primary function is to filter, monitor, and block malicious HTTP/S traffic to and from a web application. WAFs are specifically designed to identify and mitigate common web-based attacks, with Cross-Site Scripting (XSS) and SQL injection being two of the most prominent examples. By inspecting the content of web traffic, a WAF can detect and block requests containing malicious scripts or database queries before they reach the application server.

Why Incorrect Options are Wrong:

A. Ransomware: This is a type of malware. A WAF is not the primary defense; endpoint protection and anti-malware solutions are designed for this threat.
B. Syn floods: This is a network-layer (Layer 3/4) Denial of Service (DoS) attack. It is primarily mitigated by network firewalls and dedicated DDoS protection services, not WAFs.
D. Password cracking: This is an attack on authentication. While a WAF can help by rate-limiting login attempts, the primary defenses are strong password policies and multi-factor authentication.

References:

1. National Institute of Standards and Technology (NIST). (2007). Guide to Secure Web Services (Special Publication 800-95). "A WAF is a device that is intended to protect a Web server from Web-based attacks... WAFs can protect against a variety of attacks, including buffer overflows, SQL injection, and cross-site scripting." (Section 4.3.2, Page 4-6).
2. The Open Web Application Security Project (OWASP). Web Application Firewall. "A web application firewall (WAF) is an application firewall for HTTP applications. It applies a set of rules to an HTTP conversation. Generally, these rules cover common attacks such as cross-site scripting (XSS) and SQL injection." (OWASP Foundation, Web Application Firewall page, Introduction).

3. Papamartzivanos, D., Marmol, F. G., & Kambourakis, G. (2017). Introducing an intelligent engine for thwarting application-layer DDoS attacks. Journal of Information Security and Applications, 35, 49-59. "Web Application Firewalls (WAFs) are security solutions that aim to protect web applications from a plethora of attacks, such as SQL injection (SQLi), Cross-Site Scripting (XSS), and Remote File Inclusion (RFI)." (Section 1, Introduction, Paragraph 1). https://doi.org/10.1016/j.jisa.2017.06.002

Question: 27

APIs are defined as which of the following?

A. A set of protocols, and tools for building software applications to access a web-based software application or tool

B. A set of routines, standards, protocols, and tools for building software applications to access a web-based software application ortool

C. Aset of standards forbuilding software applications toaccessaweb-based softwareapplicationor tool

D. A set of routines and tools for building software applications to access web-based software applications

Answer:

B

Explanation:

An Application Programming Interface (API) is a formally defined set of rules, routines, and specifications that software programs can follow to communicate with each other. It serves as an interface between different software applications and facilitates their interaction. The definition in option B is the most comprehensive, as it correctly includes all the essential components: routines (the specific functions or procedures), standards (the data formats and conventions), protocols (the rules for data exchange), and tools (libraries and documentation that aid development). This complete set allows developers to build applications that can access the features or data of another service or system in a predictable and standardized manner.

Why Incorrect Options are Wrong:

A. This option is incomplete as it omits the crucial elements of routines and standards, which are fundamental parts of an API's definition.
C. This is too narrow. While APIs involve standards, they also explicitly define the routines, protocols, and tools needed for interaction.
D. This option is missing standards and protocols, which are essential for ensuring consistent and predictable communication between applications.

References:

1. National Institute of Standards and Technology (NIST), Special Publication 800-204, Security Strategies for Microservices-based Application Systems, December 2019. In Section 2.1, "Acronyms," an API is defined as: "A set of routines, protocols, and tools for building software and applications." This directly supports the components listed in the correct answer.
2. Google Cloud Documentation, "What is an API?". The official documentation states: "An API is

a set of routines, protocols, and tools for building software applications. An API specifies how software components should interact." This definition aligns perfectly with the chosen answer.
3. Red Hat Official Documentation, "What is an API?". The documentation defines an API as: "a set of definitions and protocols for building and integrating application software." This reinforces that an API is more than just one component, encompassing definitions (which include routines and standards) and protocols.

Question: 28

Which of the following best describes data masking?

A. A method for creating similar but inauthentic datasets usedfor software testing and user training.

B. A method used to protect prying eyes from data such as social security numbers and credit card
data.

C. Amethodwhere the last few numbers in adataset are not obscured. These are oftenusedfor authentication.

D. Datamasking involvesstrippingout all digitsinastring of numberssoastoobscuretheoriginal number.

Answer:

A

Explanation:

Data masking, also known as data obfuscation, is a data security technique that creates a structurally similar but inauthentic version of an organization's data. The primary purpose is to protect sensitive information while providing a realistic, functional alternative for use in non-production environments. These environments, such as those for software development, quality assurance testing, and user training, require data that mirrors the production format and structure but should not expose actual sensitive customer or business information. Masking techniques replace sensitive data with fictitious yet realistic data, preserving data utility without compromising confidentiality.

Why Incorrect Options are Wrong:

B. This description is too general. While data masking does protect sensitive data, this statement could also describe encryption, tokenization, or access controls. It lacks the specificity of creating a substitute dataset.
C. This describes a specific masking technique known as truncation or partial masking (e.g., showing only the last four digits of a credit card), not the overarching concept of data masking.
D. This describes the "nulling out" or redaction technique, which is only one of many methods used in data masking. It is not a comprehensive definition of the entire process.

References:

1. Cloud Security Alliance (CSA) Cloud Controls Matrix (CCM) v4.0.7, Control ID: DSP-10 (Data Masking and Obfuscation). The control specification states, "Data masking, obfuscation, or anonymization shall be used to protect sensitive data (e.g., PII) in non-production environments (e.g., development, testing)." This directly supports the use case described in option A.
2. NIST Special Publication 800-122, Guide to Protecting the Confidentiality of Personally Identifiable Information (PII), Section 5.4.2, discusses de-identification techniques. It describes masking as a method to "replace PII with fictitious data that has a similar format and data type to the original PII." This aligns with creating inauthentic but structurally similar datasets.
3. Kadlag, S., & Jadhav, S. A. (2015). Data Masking as a Service. International Journal of Computer Applications, 116(19), 1-4. The paper states, "The main reason for applying data masking is to protect sensitive data, while providing a functional substitute for occasions when the real data is not required. For example, in user training, or software testing." (Page 1, Section 1: Introduction). DOI: 10.5120/20443-2821.

Question: 29

Which of the following best describes a sandbox?

A. An isolated space where untested code and experimentation can safely occur separate from the
production environment.

B. A space where you can safely execute malicious code to see what it does.

C. An isolated space where transactions are protected from malicious software

D. An isolated space where untested code and experimentation can safely occur within the production environment.

Answer:

A

Explanation:

A sandbox is a security mechanism that creates an isolated, controlled execution environment. Its primary purpose is to run untested or untrusted code and applications without allowing them to interact with or affect the production system, host operating system, or other applications. This separation is fundamental to preventing potential damage from bugs, vulnerabilities, or malicious behavior during development, testing, or malware analysis. The environment strictly limits the resources (e.g., network access, file system) the code can access, ensuring any adverse effects are contained.

Why Incorrect Options are Wrong:

B. This describes a specific use case for a sandbox (malware analysis), not the best overall definition of what a sandbox is.
C. This describes a secure enclave or a specific transactional security mechanism, which is a different concept from a general-purpose sandbox for code execution.
D. A core principle of sandboxing for testing and development is to keep it separate from the production environment to prevent any risk of compromise or instability.

References:

1. National Institute of Standards and Technology (NIST). (n.d.). Sandbox. In CSRC Glossary. Retrieved from https://csrc.nist.gov/glossary/term/sandbox. The glossary defines a sandbox as: "A restricted, controlled execution environment that prevents potentially malicious software, such as mobile code, from accessing any system resources except for the isolated resources permitted." This supports the concept of an isolated, safe space.
2. Parno, B. (2004). The Security Architecture of the MVM Framework. Stanford University, Computer Science Department. In Section 2.1, "Sandboxing," it is stated: "The goal of a sandbox

is to provide a restricted environment in which to run untrusted code. The sandbox is responsible for ensuring that the untrusted code cannot perform any malicious actions..." This aligns with the principle of a safe, isolated environment for untrusted code.
3. Zeldovich, N., & Kaashoek, F. (2014). 6.858 Computer Systems Security, Lecture 4: Confinement. Massachusetts Institute of Technology: MIT OpenCourseWare. The lecture notes state the goal of sandboxing is to "confine a process, so it can't do bad things... Run process in a restricted environment." This emphasizes the isolation and safety aspects, separate from a main system.

Question: 30

Alocalizedincident or disaster can be addressed in acost-effectivemanner by usingwhich of the following?

A. UPS

B. Generators

C. Joint operating agreements

D. Strict adherence to applicableregulations

Answer:

C

Explanation:

A Joint Operating Agreement (JOA), also known as a reciprocal or mutual aid agreement, is a formal arrangement between two or more organizations to assist each other in the event of a disaster. This strategy is highly cost-effective because it allows participants to share the burden of business continuity and disaster recovery. Instead of each organization incurring the significant capital and operational expenses of building and maintaining a dedicated alternate processing site (e.g., a hot or warm site), they can rely on the resources of a partner. This is particularly effective for localized incidents where one organization is impacted while the other remains operational and can provide support.

Why Incorrect Options are Wrong:

A. UPS: An Uninterruptible Power Supply (UPS) only provides short-term backup power for brief outages and is not a solution for a broader disaster.
B. Generators: Generators address longer-term power failures but are a costly capital investment and only mitigate a single type of incident, not a comprehensive disaster.
D. Strict adherence to applicable regulations: This is a mandatory compliance activity, not a disaster recovery strategy. While it may improve resilience, it does not provide a direct mechanism for recovery.

References:

1. National Institute of Standards and Technology (NIST) Special Publication 800-34 Rev. 1, Contingency Planning Guide for Federal Information Systems. Section 4.3.2, "Alternate Site," discusses reciprocal agreements as a low-cost option, stating, "Reciprocal agreements are typically the lowest-cost option to implement; however, they are very difficult to enforce." This supports the "cost-effective" nature of the solution.
2. Federal Financial Institutions Examination Council (FFIEC) IT Examination Handbook, Business Continuity Management Booklet. Appendix D, "Alternate Site Options," describes

reciprocal agreements: "A reciprocal agreement is typically a no-cost or low-cost option for business continuity... The primary advantage of a reciprocal agreement is the low cost to initiate and maintain the agreement."

3. ISO/IEC 27031:2011, Information technology - Security techniques - Guidelines for information and communication technology readiness for business continuity. Section 6.4.3, "Recovery facilities," outlines various options for recovery. Mutual agreements are presented as an alternative to more expensive options like dedicated internal or external commercial sites, highlighting their role in a cost-benefit analysis for BC/DR planning.

Question: 31

In addition to battery backup, a UPS can offer which capability?

A. Breach alert

B. Confidentiality

C. Communication redundancy

D. Line conditioning

Answer:

D

Explanation:

Beyond providing battery backup during a power outage, a primary capability of many Uninterruptible Power Supply (UPS) systems is line conditioning. This function actively cleans and regulates the power flowing from the utility source to the connected equipment. It protects sensitive electronics from common power quality problems such as voltage sags (brownouts), surges (spikes), and electrical noise. By filtering these disturbances, the UPS delivers a stable and clean power signal, which is essential for the proper functioning and longevity of IT infrastructure in a cloud or data center environment. This capability is most prominent in line-interactive and online (double-conversion) UPS topologies.

Why Incorrect Options are Wrong:

A. Breach alert: This is a function of security information and event management (SIEM) or intrusion detection/prevention systems (IDS/IPS), not a power management device.
B. Confidentiality: This is a data security control, typically achieved through encryption and access control mechanisms, and is unrelated to power supply functions.
C. Communication redundancy: This is a network availability strategy that involves multiple, independent communication links or paths to prevent a single point of failure.

References:

1. Massachusetts Institute of Technology (MIT) Lincoln Laboratory. (2011). Uninterruptible Power Supply (UPS) Systems. In Engineering Division Design and Engineering Standards, Section 10.1. On page 10.1-2, it states, "A UPS is used to provide clean, conditioned, and uninterrupted AC power to a critical load." It further details how different UPS types handle power conditioning.
2. Rassool, N., & Manyage, M. (2017). A review of uninterruptible power supplies. 2017 IEEE AFRICON, Cape Town, South Africa, pp. 1114-1119. In Section II, "UPS Topologies," the paper describes how line-interactive and online UPS systems "provide power conditioning" and "filter the input power," contrasting them with the more basic standby UPS. (DOI: 10.1109/AFRICON.2017.8095601)

3. Schneider Electric. (2011). The Different Types of UPS Systems (White Paper 1, Rev. 7). In the section "Line-interactive UPS" (p. 4), it states, "This type of UPS is also able to correct minor power fluctuations (under-voltages and over-voltages) without switching to battery." This voltage regulation is a key aspect of line conditioning.

Question: 32

For performance purposes, OS monitoring should include all of the following except:

A. Disk space

B. Disk I/O usage

C. CPU usage

D. Print spooling

Answer:

D

Explanation:

For performance purposes, OS monitoring focuses on fundamental resource metrics that directly impact the system's overall health, stability, and responsiveness. CPU usage, disk I/O, and available disk space are critical indicators of system load and potential bottlenecks. High CPU or disk I/O rates can signal performance degradation, while insufficient disk space can lead to system crashes or slow performance.

Print spooling, in contrast, is a specific background service for managing print jobs. While a malfunctioning print spooler can consume system resources, it is not a core, universal metric for OS performance. In many cloud environments and server roles (e.g., web servers, database servers), this service is often disabled or not installed, making it irrelevant for general performance monitoring.

Why Incorrect Options are Wrong:

A. Disk space: Insufficient disk space can halt system operations, prevent applications from writing temporary files, and cause severe performance degradation, making it essential to monitor.

B. Disk I/O usage: High disk input/output (I/O) is a primary cause of performance bottlenecks, directly affecting application speed and data access times.

C. CPU usage: This is one of the most critical metrics for performance, as sustained high CPU utilization indicates the system is overloaded and cannot process tasks efficiently.

References:

1. Amazon Web Services (AWS) Documentation: The official AWS documentation for Amazon CloudWatch, the monitoring service for AWS cloud resources, lists key metrics for EC2 instances (virtual servers). These include CPUUtilization, DiskReadOps, DiskWriteOps, and metrics for storage volumes. Print spooling is not included as a standard monitored metric for OS performance.

Source: Amazon Web Services, "Amazon EC2 CloudWatch Metrics," Amazon CloudWatch User

Guide. (Specifically, the section on "Instance metrics").
2. Microsoft Azure Documentation: Similarly, Azure Monitor for VMs collects performance data from guest operating systems. Standard metrics include "% Processor Time" (CPU), "Logical Disk Bytes/sec" (Disk I/O), and "Logical Disk Free Space." Monitoring for a specific service like a print spooler is considered custom and not a default performance counter.
Source: Microsoft, "Overview of Azure Monitor for VMs," Microsoft Docs. (Specifically, the section on "Performance").
3. University Courseware: University-level operating systems courses emphasize the monitoring of core hardware resource utilization as the basis for performance evaluation. Lectures on system performance consistently focus on CPU scheduling, memory management, and I/O efficiency as the primary areas of concern.
Source: Ousterhout, J., "Lecture 1: Introduction," CS 140: Operating Systems, Stanford University, Winter 2018, pp. 21-23. (Discusses OS goals of performance, which are tied to managing CPU, memory, and I/O).

Question: 33

Identity and access management (IAM) is a security discipline that ensures which of the following?

A. That all users are properlyauthorized

B. That the right individual gets access to the right resources at the right time for the right reasons.

C. That all users are properlyauthenticated

D. That unauthorized users will get access to the right resources at the right time for the right reasons

Answer:

B

Explanation:

Identity and Access Management (IAM) is the security framework and set of business processes that ensures access to resources is managed securely and efficiently. The core principle of IAM is to grant access based on the principle of least privilege and business need. This is holistically captured by ensuring the "right individual" (identity and authentication) gains access to the "right resources" (authorization) at the "right time" for the "right reasons" (context and policy). This comprehensive approach encompasses the entire lifecycle of identity management, from provisioning to de-provisioning, and goes beyond the individual components of authentication or authorization alone.

Why Incorrect Options are Wrong:

A. This is incomplete. IAM includes identity verification (authentication) and lifecycle management, not just authorization.
C. This is incomplete. IAM also determines what resources an authenticated user is permitted to access (authorization).
D. This is the antithesis of IAM's purpose. IAM is designed to prevent unauthorized users from gaining access.

References:

1. National Institute of Standards and Technology (NIST). (n.d.). What is Identity and Access Management (IAM)? NIST Computer Security Resource Center. Retrieved from https://csrc.nist.gov/projects/iam. In the overview, NIST states, "Identity and Access Management (IAM) is the security discipline that makes it possible for the right entities to use the right resources when they need to, without interference, using the devices they want to use."
2. Perrin, C. (2018). Foundations of Identity and Access Management. University of California,

Berkeley, Information Security Office. In the "What is Identity and Access Management?" section, the document describes IAM as a framework for "ensuring the right people have the right access to the right resources at the right time."
3. Al-Khouri, A. M. (2012). Identity and Access Management. International Journal of Computer Science Issues (IJCSI), 9(5), 497-509. On page 498, the paper defines IAM as "a framework of policies and technologies for ensuring that the right users have the appropriate access to technology resources."

Question: 34

Maintenance mode requires all of these actions except:

A. Remove all active productioninstances

B. Ensure logging continues

C. Initiate enhanced security controls

D. Prevent new logins

Answer:

C

Explanation:

Maintenance mode is a controlled state designed to allow for system updates, patching, or repairs while minimizing risk and user impact. Standard procedures include taking instances out of the production pool (A), preventing new user logins to ensure data consistency (D), and ensuring all administrative actions are logged for security and auditing purposes (B).
The action that is not a universal requirement is initiating enhanced security controls. While overall security must be maintained through strict authorization, supervision, and logging, the maintenance process itself might require the temporary, controlled modification or relaxation of certain security controls to allow the work to be completed. The focus is on controlled, audited activity, not necessarily the addition of new or enhanced controls.

Why Incorrect Options are Wrong:

A. Removing instances from the active production pool is a standard procedure to prevent users from accessing a system undergoing maintenance and to ensure a stable environment for the changes.
B. Continuous logging is crucial during maintenance to audit privileged activities, track changes, and ensure accountability, which is a fundamental security principle.
D. Preventing new logins is essential to protect data integrity and avoid disrupting user sessions while the system is in a potentially unstable state.

References:

1. NIST Special Publication 800-53 Revision 5, Security and Privacy Controls for Information Systems and Organizations. The Maintenance (MA) control family, particularly MA-2 "Controlled Maintenance," outlines requirements for scheduling, performing, and documenting maintenance. The focus is on authorization, control, and review of maintenance activities, not on adding enhanced controls. The standard emphasizes maintaining a secure state through procedural controls. (See Section: MA-2, Page 203).
2. Cloud Security Alliance (CSA), Security Guidance for Critical Areas of Focus in Cloud

Computing v4.0. Domain 5: "Cloud Security Operations" discusses the importance of a formal change management process. This process includes "logging and monitoring of privileged user activities" and ensuring changes are authorized. It does not mandate enhancing security controls during the maintenance window itself; rather, it requires that security is managed throughout the process. (See Domain 5, Page 103).

3. ISO/IEC 27017:2015, Code of practice for information security controls based on ISO/IEC 27002 for cloud services. Section 12.1.2, "Protection against malware," and 12.2, "Backup," imply that operational procedures, including maintenance, must be conducted in a way that preserves security. The guidance focuses on preventing the introduction of vulnerabilities and ensuring system integrity through controlled procedures, which aligns with logging and preventing access, but not necessarily enhancing controls.

Question: 35

What is one of the reasons a baseline might be changed?

A. Numerous change requests

B. To reduce redundancy

C. Natural disaster

D. Power fluctuation

Answer:

A

Explanation:

A security baseline is a standardized level of configuration and security settings for a system or network. When numerous change requests are submitted for systems governed by a baseline, it serves as a key indicator that the baseline is no longer aligned with current business, operational, or technical requirements. This divergence, often called configuration drift when unauthorized, signals that the established standard is becoming obsolete or impractical. Consequently, the organization should initiate a formal review and update the baseline to reflect the new, necessary state, thereby reducing the administrative overhead of processing constant exceptions and maintaining a relevant security posture.

Why Incorrect Options are Wrong:

B. To reduce redundancy is a general IT architecture goal; it is not a direct trigger for changing a security configuration baseline.
C. A natural disaster initiates disaster recovery and business continuity plans, which typically involve restoring systems to their last approved baseline, not changing it.
D. Power fluctuation is a transient operational issue that requires an infrastructure-level response, not a modification of standardized system security configurations.

References:

1. NIST Special Publication 800-128, Guide for Security-Focused Configuration Management of Information Systems. Section 2.4, "Configuration Control," discusses the process for managing changes. It states, "The effects of approved changes are assessed, and all relevant configuration management documentation (including the system's baseline configuration) is updated as necessary." A high volume of approved changes would necessitate an update to the baseline itself to reflect the new standard.
2. NIST Special Publication 800-53 Revision 5, Security and Privacy Controls for Information Systems and Organizations. Control CM-3, "Configuration Change Control," outlines the process

for managing changes to the system. A pattern of consistent, approved deviations from the baseline (evidenced by numerous change requests) would trigger a review and potential update of the baseline defined in control CM-2, "Baseline Configuration," to ensure it remains "current."
3. Carnegie Mellon University, Software Engineering Institute (SEI), CERT Resilience Management Model (CERT-RMM) v1.2. The Configuration and Change Management (CCM) process area states that one of its goals is to "Establish and maintain the integrity of the configuration items of the service." A baseline that generates excessive change requests is failing to maintain integrity and relevance, indicating a need for re-evaluation and update. (See CCM Specific Goal 2).