

DUMPSBOSS.

**Certified Information Systems Security
Professional (CISSP)**

ISC2 CISSP

Version Demo

Total Demo Questions: 120

Total Premium Questions: 1205

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Security and Risk Management	196
Topic 2, Asset Security	90
Topic 3, Security Architecture and Engineering	162
Topic 4, Communication and Network Security	170
Topic 5, Identity and Access Management (IAM)	163
Topic 6, Security Assessment and Testing	99
Topic 7, Security Operations	178
Topic 8, Software Development Security	147
Total	1205

QUESTION NO: 1

An authentication system that uses challenge and response was recently implemented on an organization's network, because the organization conducted an annual penetration test showing that testers were able to move laterally using authenticated credentials. Which attack method was MOST likely used to achieve this?

- A. Hash collision.
- B. Pass the ticket.
- C. Brute force.
- D. Cross-Site Scripting (XSS).

ANSWER: B

Explanation:

The most likely method is pass the ticket. In many enterprise environments, lateral movement with “authenticated credentials” commonly occurs by stealing and replaying existing authentication artifacts rather than cracking passwords. With Kerberos-based single sign-on, an attacker who obtains a valid Ticket Granting Ticket (TGT) or service ticket from memory (for example, from a compromised workstation) can reuse that ticket to authenticate to other services and systems as that user, enabling lateral movement without needing the user’s plaintext password. This is known as a pass-the-ticket technique and is a classic credential replay scenario.

Challenge-response authentication mechanisms are often introduced specifically to reduce the value of captured password hashes and to mitigate replay-style attacks where an attacker reuses captured authentication material. While challenge-response does not eliminate all Kerberos ticket theft risks, the scenario’s emphasis on lateral movement using authenticated credentials aligns strongly with the reuse of existing Kerberos tickets rather than password guessing or web attacks.

References: [MITRE ATT&CK: Use Alternate Authentication Material \(Pass the Ticket\)](#), [Microsoft: Kerberos authentication overview](#).

QUESTION NO: 2

What part of an organization's strategic risk assessment MOST likely includes information on items affecting the success of the organization?

- A. Threat analysis.
- B. Vulnerability analysis.
- C. Key Performance Indicator (KPI).
- D. Key Risk Indicator (KRI).

ANSWER: D

Explanation:

Key Risk Indicator (KRI).

is the part of a strategic risk assessment that most directly captures information about factors that could affect organizational success. KRIs are measurable metrics used to provide early warning signals of increasing risk exposure and potential deviation from the organization’s risk appetite and strategic objectives. Because strategic risk assessment is concerned with risks that can materially impact mission achievement, market position, regulatory standing, financial viability, or reputation, KRIs are designed to monitor those risk drivers over time (for example, customer churn rate, concentration risk, liquidity

ratios, critical supplier failure rates, or regulatory findings). In contrast to operational security analyses, KRIs are typically reported to senior management and the board to support governance, risk oversight, and timely decision-making. This aligns with common enterprise risk management practices where KRIs complement objectives and performance measures by focusing on uncertainty and downside exposure that could prevent success. A practical way to think about it is that KRIs track “how close we are to unacceptable risk,” which is exactly the kind of information strategic risk assessment aims to surface for leadership.

References: [COSO Enterprise Risk Management](#), [ISACA Journal: Key Risk Indicators](#)

QUESTION NO: 3

Which of the following is NOT part of user provisioning?

- A. Creation and deactivation of user accounts.
- B. Business process implementation.
- C. Maintenance and deactivation of user objects and attributes.
- D. Delegating user administration.

ANSWER: B

Explanation:

Business process implementation is not part of user provisioning. In CISSP terms, user provisioning is an identity and access management (IAM) lifecycle activity focused on creating, modifying, and removing identities and their access based on authorization decisions (often driven by HR events such as hire, transfer, and termination). Provisioning typically includes creating and disabling accounts, maintaining identity records (user objects) and their attributes (such as department, manager, group memberships, and role assignments), and establishing administrative workflows—often including delegated administration—so that appropriate personnel can perform approved changes under defined controls. These activities are about ensuring the right identities exist, have the right entitlements, and are promptly updated or revoked as business needs change. By contrast, implementing business processes is a broader organizational and operational concern (process design, execution, and optimization) and may inform access requirements, but it is not itself an IAM provisioning function. Provisioning supports business processes by enforcing access decisions; it does not implement the processes.

References: [NIST SP 800-63 Digital Identity Guidelines](#), [NIST SP 800-53 Rev. 5 \(Access Control family\)](#)

QUESTION NO: 4

An organization is formulating a strategy to provide access to third-party partners. The information technology (IT) department has been tasked

with providing access by utilizing cloud services. Which of the following technologies is MOST commonly employed for completing the task?

- A. Identity as a Service (IDaaS)
- B. Firewall as a service
- C. Infrastructure as a Service (IaaS)
- D. Software as a Service (SaaS)

ANSWER: A

Explanation:

Identity as a Service (IDaaS) is most commonly used to provide controlled access to third-party partners when leveraging cloud services because it directly addresses identity, authentication, authorization, and federation needs across

organizational boundaries. In partner-access scenarios, the core challenge is typically establishing trusted identities and enforcing consistent access policies (for example, single sign-on, multi-factor authentication, lifecycle provisioning/deprovisioning, and role-based access) without having to create and manage local accounts for every external user. IDaaS offerings commonly support standards such as SAML, OAuth 2.0, and OpenID Connect to enable federated authentication and delegated authorization, which are foundational for business-to-business (B2B) access. This aligns with CISSP domain expectations around identity and access management: the “technology” most directly employed to grant and govern partner access is an identity/federation service rather than general compute infrastructure, application delivery, or perimeter filtering. In practice, organizations use IDaaS to integrate external partners into a centralized IAM control plane, improving security (strong authentication, conditional access) and governance (auditing, access reviews) while reducing administrative overhead.

References: [Cloudflare: What is IDaaS?](#), [Microsoft Learn: Entra External ID \(B2B collaboration\)](#)

QUESTION NO: 5

Which combination of cryptographic algorithms are compliant with Federal Information Processing Standard (FIPS) Publication 140-2 for non-legacy systems?

- A. Diffie-hellman (DH) key exchange: DH (≥ 2048 bits) Symmetric Key: Advanced Encryption Standard (AES) > 128 bits Digital Signature: Digital Signature Algorithm (DSA) (≥ 2048 bits).
- B. Diffie-hellman (DH) key exchange: DH (≥ 2048 bits) Symmetric Key: Advanced Encryption Standard (AES) > 128 bits Digital Signature: Rivest-Shamir-Adleman (RSA) (1024 bits).
- C. Diffie-hellman (DH) key exchange: DH (≤ 1024 bits) Symmetric Key: Blowfish Digital Signature: Rivest-Shamir-Adleman (RSA) (≥ 2048 bits).
- D. Diffie-hellman (DH) key exchange: DH (≥ 2048 bits) Symmetric Key: Advanced Encryption Standard (AES) < 128 bits Digital Signature: Elliptic Curve Digital Signature Algorithm (ECDSA) (≥ 256 bits).

ANSWER: A

Explanation:

The combination that is compliant with FIPS 140-2 expectations for non-legacy systems is the one using Diffie-hellman (DH) key exchange with at least 2048-bit parameters, Advanced Encryption Standard (AES) with a key size greater than 128 bits (that is, AES-192 or AES-256), and Digital Signature Algorithm (DSA) with at least 2048-bit parameters. FIPS 140-2 is a module validation standard, but in practice it relies on using FIPS-approved security functions as specified in associated FIPS and NIST publications. For symmetric encryption, AES is the FIPS-approved block cipher (FIPS 197) and supports 128/192/256-bit keys; using more than 128 bits is acceptable and commonly required by modern policy baselines. For public-key techniques, NIST guidance has long moved away from 1024-bit discrete log/RSA for non-legacy use; 2048-bit finite-field DH and 2048-bit DSA align with current minimum-strength expectations in NIST SP 800-131A transition guidance. Therefore, this set of algorithms and key sizes fits the “non-legacy” intent and aligns with FIPS-approved primitives used in validated modules.

References: [NIST FIPS 197 \(AES\)](#), [NIST SP 800-131A Rev. 2 \(transition guidance\)](#)

QUESTION NO: 6

The Chief Information Security Officer (CISO) of a small organization is making a case for building a security operations center (SOC). While

debating between an in-house, fully outsourced, or a hybrid capability, which of the following would be the MAIN consideration, regardless of the model?

- A. Headcount and capacity
- B. Scope and service catalog

C. Skill set and training

D. Tools and technologies

ANSWER: B

Explanation:

Scope and service catalog is the main consideration because it defines what the SOC is actually responsible for delivering—its mission, boundaries, and measurable outcomes—before any resourcing, tooling, or sourcing decision can be made. In practice, the SOC model (in-house, outsourced, hybrid) is an implementation choice that should follow from clearly defined services such as monitoring use cases, alert triage, incident response coverage, threat hunting, vulnerability coordination, reporting, and escalation paths. Establishing the scope and service catalog also forces clarity on critical constraints like hours of operation (24x7 vs. business hours), supported environments (cloud, endpoints, OT), required SLAs, regulatory obligations, and integration points with IT and incident management. Once these services are defined, the organization can rationally determine what must remain internal (e.g., business context, decision authority) versus what can be delegated, and then size headcount, training needs, and technology accordingly. This aligns with SOC best practices and IT service management thinking: define services and requirements first, then design the operating model to meet them. See guidance on SOC functions and operating models in [NIST SP 800-61](#) and SOC capability considerations in [CISA SOC resources](#).

QUESTION NO: 7

Which is the BEST control to meet the Statement on Standards for Attestation Engagements 18 (SSAE-18) confidentiality category?

A. File hashing.

B. Storage encryption.

C. Data retention policy.

D. Data processing.

ANSWER: B

Explanation:

Storage encryption is the best control to meet the SSAE-18 confidentiality category because confidentiality in SOC reporting (Trust Services Criteria) focuses on protecting information designated as confidential from unauthorized disclosure. Encrypting data at rest directly reduces the likelihood and impact of unauthorized access to stored information by ensuring that, without the appropriate cryptographic keys, the data remains unreadable. This maps cleanly to common confidentiality expectations in SOC 2 examinations, where organizations must demonstrate that confidential data is protected throughout its lifecycle, including while stored on disks, databases, backups, and snapshots. Storage encryption is also a broadly accepted, auditable technical control: auditors can verify encryption configurations, key management practices, and coverage (systems, volumes, databases) as evidence that confidentiality requirements are being met. In practice, encryption at rest is typically paired with strong access controls and key management, but as a single “best” control aligned to confidentiality, storage encryption is the most direct and defensible choice.

References: [AICPA SOC Suite of Services \(SOC and Trust Services\)](#), [Google Cloud: Encryption at rest overview](#).

QUESTION NO: 8

The security architect is designing and implementing an internal certification authority to generate digital certificates for all employees. Which of

the following is the

BEST solution to securely store the private keys?

- A. Physically secured storage device
- B. Trusted Platform Module (TPM)
- C. Encrypted flash drive
- D. Public key infrastructure (PKI)

ANSWER: B

Explanation:

Trusted Platform Module (TPM) is the best solution because it provides hardware-backed protection for cryptographic keys, designed specifically to generate, store, and use private keys in a way that resists extraction. A TPM can keep private keys non-exportable, perform cryptographic operations inside the protected hardware boundary, and leverage tamper-resistant features that significantly reduce the risk of key theft via malware, disk compromise, or casual physical access. In an enterprise internal CA context, protecting private keys is critical because compromise enables certificate forgery and impersonation. While other storage approaches can add physical security or encryption, they generally still allow keys to be copied if the storage medium or host is compromised. A TPM aligns with CISSP best practices for strong key management by using dedicated hardware security features to protect key material at rest and during use. For higher assurance CA key protection, organizations may also use HSMs, but among the provided choices, TPM is the strongest hardware-rooted option for securely storing private keys.

References: [Microsoft TPM overview](#), [NIST SP 800-57 Part 1 \(Key Management\)](#)

QUESTION NO: 9

In access control terms, the word "dominate" refers to which of the following?

- A. Higher or equal to access class.
- B. Rights are superceded.
- C. Valid need-to-know with read privileges.
- D. A higher clearance level than other users.

ANSWER: A

Explanation:

In mandatory access control (MAC) models—especially lattice-based models used to describe multilevel security—the term “dominate” describes an ordering relationship between security labels (or access classes). A subject’s label is said to dominate an object’s label when the subject’s classification level is greater than or equal to the object’s classification level and, in compartmented systems, the subject’s set of categories includes (covers) the object’s categories. This dominance relationship is the formal basis for decisions like whether a subject can read an object in Bell–LaPadula (the “simple security property” relies on the subject dominating the object’s label). In other words, “dominate” is not about having special privileges or overriding rights; it is specifically about the label comparison rule used in MAC to determine whether access is permitted based on the relative ordering of access classes. This is why the best match is the statement that captures the “higher or equal” relationship of access class/label ordering.

References: [NIST CSRC Glossary: Mandatory Access Control](#), [NIST CSRC Glossary: Bell-LaPadula Model](#)

QUESTION NO: 10

Which of the following is the BEST way to protect an organization's data assets?

- A. Encrypt data in transit and at rest using up-to-date cryptographic algorithms.
- B. Monitor and enforce adherence to security policies.

C. Require Multi-Factor Authentication (MFA) and Separation of Duties (SoD).

D. Create the Demilitarized Zone (DMZ) with proxies, firewalls and hardened bastion hosts.

ANSWER: A

Explanation:

Encrypting data in transit and at rest using up-to-date cryptographic algorithms is the best choice because it directly protects the confidentiality (and, when properly designed, integrity) of the data itself, regardless of where the data travels or where it is stored. In CISSP terms, this is a strong example of “data-centric” protection: even if perimeter controls fail, a device is lost, backups are stolen, or traffic is intercepted, properly implemented encryption keeps the data unreadable without the appropriate keys. Modern cryptography also supports authenticated encryption and robust key management practices, which are essential to making encryption effective in real environments. This approach aligns with widely accepted security baselines that emphasize protecting sensitive data both at rest (databases, filesystems, backups) and in transit (TLS-protected sessions, VPN tunnels) using current, vetted algorithms and configurations. While governance, monitoring, access control, and network segmentation are all important, encryption is uniquely positioned as a direct control over the data asset itself and is broadly applicable across systems, users, and networks.

References: [NIST SP 800-53 Rev. 5 \(SC-13, SC-28\)](#), [NIST SP 800-52 Rev. 2 \(TLS Guidelines\)](#)

QUESTION NO: 11

An application is used for funds transfers between an organization and a third-party. During a security audit, an auditor has found an issue with the business continuity disaster recovery policy and procedures for this application. Which of the following reports should the auditor file with the organization?

A. Statement on Auditing Standards (SAS) 70-1.

B. Statement on Auditing Standards (SAS) 70.

C. Service Organization Control (SOC) 1.

D. Service Organization Control (SOC) 2.

ANSWER: C

Explanation:

Service Organization Control (SOC) 1 is the appropriate report when the audited service (here, a third-party involved in funds transfers) can affect the user organization’s financial reporting. SOC 1 reports (under SSAE 18 / AT-C 320) are specifically designed to provide assurance over controls at a service organization that are relevant to internal control over financial reporting (ICFR). Business continuity and disaster recovery controls can be in-scope for SOC 1 when they support the availability and processing integrity needed for complete and accurate transaction processing that ultimately impacts financial statements. In contrast, SOC 2 is oriented to the Trust Services Criteria (security, availability, processing integrity, confidentiality, privacy) for broader assurance needs and is not specifically targeted at financial reporting. Since the scenario is funds transfers—an inherently financial process—the auditor would file a SOC 1 report to communicate control issues relevant to financial reporting reliance on the third-party service. This aligns with the modern replacement of the older SAS 70 framework, which has been superseded by SOC reporting.

References: [AICPA: SOC suite of services](#), [AICPA: Service Organization Reporting \(SOC 1 / SOC 2\)](#)

QUESTION NO: 12

Which of the following activities should a forensic examiner perform FIRST when determining the priority of digital evidence collection at a crime scene?

A. Gather physical evidence.

B. Assign responsibilities to personnel on the scene.

C. Establish a list of files to examine.

D. Establish order of volatility.

ANSWER: D

Explanation:

Establishing an order of volatility is the first step when prioritizing digital evidence collection because it directly determines what evidence is most likely to be lost or altered if not captured immediately. In digital forensics, “volatile” data (such as CPU registers/cache, running processes, active network connections, and contents of RAM) can disappear instantly when a system is powered down, crashes, or is simply left running as applications and memory contents change. By defining the order of volatility up front, the examiner can make defensible, repeatable decisions about what to collect first (typically volatile memory and live system state before less volatile artifacts like disk images and backups), minimizing the risk of evidence spoliation and preserving evidentiary value. This approach is consistent with widely accepted forensic practice and incident handling guidance, where collection sequencing is driven by volatility and the need to preserve perishable data before it changes. Once volatility is addressed, the examiner can proceed with broader scene management tasks and deeper analysis planning without sacrificing transient evidence.

References: [RFC 3227: Guidelines for Evidence Collection and Archiving](#), [NIST SP 800-86: Guide to Integrating Forensic Techniques into Incident Response](#)

QUESTION NO: 13

When telephones in a city are connected by a single exchange, the caller can only connect with the switchboard operator. The operator then

manually connects the call. This is an example of which type of network topology?

A. Point-to-Point Protocol (PPP)

B. Bus

C. Star

D. Tree

ANSWER: C

Explanation:

The correct topology is **Star**. In a star topology, all endpoints connect to a single central device (such as a hub, switch, PBX, or—in the historical telephone example—a central exchange/switchboard). Communication between any two endpoints is mediated through that central point: the caller reaches the exchange first, and then the exchange establishes the connection to the destination line. This matches the described behavior where the caller can only reach the switchboard operator, who then completes the circuit by connecting the call through the exchange. In CISSP terms, this illustrates the defining characteristic of star topologies: centralized connectivity and control, with each node having an individual link to the central system. While modern networks implement this electronically via switches, the manual switchboard is simply an earlier implementation of the same logical topology. For additional background on star topology concepts, see https://en.wikipedia.org/wiki/Network_topology and a practical overview of star-wired networks at <https://www.techtarget.com/searchnetworking/definition/star-topology>.

QUESTION NO: 14

To ensure proper governance of information throughout the lifecycle, which of the following should be assigned FIRST?

A. Owner

- B. Classification
- C. Custodian
- D. Retention

ANSWER: A

Explanation:

Owner is the item that should be assigned first because information governance starts with clear accountability. The data (or information) owner is the role responsible for defining how information is to be handled across its lifecycle, including setting protection requirements, approving access, and ensuring appropriate classification and retention requirements are established. Without an assigned owner, there is no authoritative decision-maker to determine the information's value, sensitivity, and business requirements, which means downstream activities like classification, retention scheduling, and operational handling by custodians lack direction and may be inconsistent. In common governance models referenced in CISSP, the owner drives the "what and why" (business requirements and risk decisions), while custodians implement the "how" (day-to-day safeguards), and classification/retention are outcomes of the owner's decisions and organizational policy. Assigning ownership first therefore enables consistent lifecycle governance from creation/collection through use, storage, sharing, archival, and disposal.

References: [NIST Privacy Framework \(governance and accountability concepts\)](#), [ISO/IEC 27001 overview \(information security roles, responsibilities, and governance\)](#).

QUESTION NO: 15

What term is commonly used to describe hardware and software assets that are stored in a configuration management database (CMDB)?

- A. Configuration item.
- B. Configuration element.
- C. Ledger item.
- D. Asset register.

ANSWER: A

Explanation:

The correct term is "Configuration item." In IT service management (commonly referenced alongside CISSP operations concepts), a configuration management database (CMDB) is used to store information about components that need to be managed to deliver and support IT services. Those components can include hardware, software, systems, network devices, virtual resources, documentation, and even services—along with their attributes and relationships. The standard term for any such managed component recorded in the CMDB is a configuration item (often abbreviated as CI). This terminology is widely used in ITIL/ITSM practices and aligns with the idea of maintaining accurate inventories and dependencies to support change management, incident response, and impact analysis. In practice, treating managed assets as configuration items enables traceability (what changed, when, and by whom), relationship mapping (what depends on what), and better risk-informed operational decisions—topics that map well to CISSP domain expectations around configuration/change control and operational resilience.

References: [AXELOS ITIL Service Management](#), [Microsoft Cloud Adoption Framework - Configuration management](#)

QUESTION NO: 16

Which of the following statements pertaining to packet filtering NOT true?

- A. It is based on ACLs.
- B. It is not application dependent.
- C. It operates at the network layer.
- D. It keeps track of the state of a connection.

ANSWER: D

Explanation:

The statement “It keeps track of the state of a connection.” is the one that is not true for classic packet filtering. Traditional (stateless) packet-filtering firewalls evaluate each packet independently against a rule set (often implemented as access control lists) using header information such as source/destination IP addresses, protocol, and source/destination ports. Because each packet is assessed on its own, a basic packet filter does not maintain session context (for example, whether a TCP handshake has been completed) and therefore cannot inherently distinguish legitimate return traffic from unsolicited inbound traffic beyond what the static rules allow.

Tracking the state of a connection is a capability of stateful inspection (stateful packet filtering) firewalls, which maintain a state table and make decisions based on connection context. In CISSP terms, this is an important distinction: packet filtering is generally associated with network/transport-layer header checks and ACL-style rules, while stateful inspection adds session awareness. For additional background on packet filtering and stateful inspection concepts, see [Cisco: What is a firewall?](#) and [Microsoft: Windows Firewall with Advanced Security](#).

QUESTION NO: 17

A security engineer is assigned to work with the patch and vulnerability management group. The deployment of a new patch has been approved

and needs to be applied. The research is complete, and the security engineer has provided recommendations. Where should the patch be applied FIRST?

- A. Lower environment
- B. Desktop environment
- C. Server environment
- D. Production environment

ANSWER: A

Explanation:

The patch should be applied first in a lower environment (a non-production environment such as development, test, or staging) to validate that it installs correctly and does not introduce unacceptable side effects before impacting business operations. In mature patch and vulnerability management programs, changes are introduced progressively: first to controlled test systems that mirror production as closely as practical, then to broader pilot groups, and only then to production. This approach supports change management goals by reducing risk, enabling rollback planning, confirming compatibility with applications and configurations, and allowing monitoring for performance or stability issues. Even when a patch is approved and research is complete, testing in a lower environment is a key control to prevent outages and to verify that compensating controls or deployment prerequisites (dependencies, reboots, configuration changes) are understood. This aligns with common security and IT service management practices for controlled change introduction and validation prior to production rollout. See NIST guidance on patching and configuration management in [NIST SP 800-40 Rev. 4](#) and change control concepts in [NIST SP 800-128](#).

QUESTION NO: 18

Which part of an operating system (OS) is responsible for providing security interfaces among the hardware, OS, and other parts of the computing

system?

- A. Reference monitor
- B. Trusted Computing Base (TCB)
- C. Time separation
- D. Security kernel

ANSWER: D

Explanation:

The **security kernel** is the part of an operating system that provides the fundamental security interfaces between hardware, the OS, and other components of the computing system. In CISSP terms, it is the core set of mechanisms that enforce the system's security policy by mediating access to resources (such as memory, CPU modes, I/O, and files) and by providing the low-level primitives that higher-level security functions rely on. Conceptually, the security kernel is the "implementation" portion that must be small, isolated, and robust because it underpins trusted enforcement. It is closely related to the reference monitor concept (which describes the properties of always-invoked, tamperproof, and verifiable mediation), but the question asks specifically for the OS part that provides the security interfaces among hardware and the rest of the system—this aligns with the security kernel as the OS's core security mechanism layer. The security kernel is also a key component within the broader Trusted Computing Base (TCB), which includes all hardware/firmware/software elements critical to enforcing security, not just the kernel mechanisms themselves.

References: [NIST CSRC Glossary – Security Kernel](#), [NIST CSRC Glossary – Trusted Computing Base](#)

QUESTION NO: 19

What can be defined as a list of subjects along with their access rights that are authorized to access a specific object?

- A. A capability table.
- B. An access control list.
- C. An access control matrix.
- D. A role-based matrix.

ANSWER: B

Explanation:

An access control list is the correct concept because it is defined as an object-focused list that enumerates which subjects (users, groups, or processes) are permitted to access that specific object and what operations they are allowed to perform (such as read, write, execute, delete). In other words, the list is attached to (or maintained for) the object and answers the question: "Who can access this object, and how?" This is a foundational discretionary access control mechanism used widely in operating systems and file systems, where each file, directory, or resource can have an ACL specifying allowed (and sometimes denied) permissions for multiple subjects. This matches the question's phrasing precisely: a list of subjects along with their access rights for a specific object. Capability tables, by contrast, are typically subject-centric (what objects a subject can access), while an access control matrix is a conceptual model representing subjects, objects, and rights in a grid form rather than the per-object list described here. For practical examples of ACLs in common platforms, see Microsoft's overview of access control lists and permissions: <https://learn.microsoft.com/en-us/windows/win32/secauthz/access-control-lists> and NIST's access control guidance: <https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final>.

QUESTION NO: 20

Which of the following describes the BEST method of maintaining the inventory of software and hardware within the organization?

- A. Maintaining the inventory through a combination of asset owner interviews, open-source system management, and open-source
- B. management tools
- C. Maintaining the inventory through a combination of desktop configuration, administration management, and procurement management tools
- D. Maintaining the inventory through a combination of on premise storage configuration, cloud management, and partner management tools
- E. Maintaining the inventory through a combination of system configuration, network management, and license management tools

ANSWER: E

Explanation:

The best method of maintaining an accurate hardware and software inventory is to use a combination of system configuration management, network management, and license management tools. In practice, this aligns with how mature asset management programs work: configuration management (often via endpoint/agent-based discovery and CMDB processes) captures detailed device and software attributes; network management (including network discovery and monitoring) helps identify unmanaged or rogue devices and validates what is actually connected; and license management (SAM tooling) reconciles installed software against entitlements to maintain compliance and reduce legal/financial risk. Using these tool categories together provides continuous discovery, normalization, reconciliation, and reporting—far more reliable than periodic interviews or procurement records alone. This approach also supports common CISSP expectations around maintaining inventories as a foundational control for vulnerability management, change control, and incident response. For additional guidance, see NIST's asset management discussion in the Cybersecurity Framework (<https://www.nist.gov/cyberframework>) and NIST SP 800-53 control CM-8 (System Component Inventory) (<https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final>).

QUESTION NO: 21

What level of Redundant Array of Independent Disks (RAID) is configured PRIMARILY for high-performance data reads and writes?

- A. RAID-0.
- B. RAID-1.
- C. RAID-5.
- D. RAID-6.

ANSWER: A

Explanation:

RAID-0 is the RAID level configured primarily for high-performance reads and writes because it uses disk striping with no parity or mirroring overhead. Data is split into blocks and written across multiple disks in parallel, which increases throughput for both sequential and many random I/O patterns. Since there is no need to calculate parity (as in parity-based RAID) and no need to duplicate writes to a mirror (as in mirroring), RAID-0 typically delivers the best raw performance and the most usable capacity from the set of drives. The tradeoff—important in CISSP context—is that RAID-0 provides no fault tolerance: the failure of any single disk results in loss of the entire array's data. Therefore, RAID-0 is chosen when performance is the primary objective and data can be recreated or is otherwise protected (for example, via backups, replication, or noncritical

workloads). For a concise overview of RAID levels and their performance characteristics, see https://en.wikipedia.org/wiki/Standard_RAID_levels and a vendor summary such as <https://www.techtarget.com/searchstorage/definition/RAID>.

QUESTION NO: 22

What documentation is produced FIRST when performing an effective physical loss control process?

- A. Deterrent controls list
- B. Security standards list
- C. Asset valuation list
- D. Inventory list

ANSWER: D

Explanation:

The first documentation produced in an effective physical loss control process is an inventory list. Physical loss control starts by identifying and documenting what you have (assets), where they are, and who is responsible for them. Without an accurate inventory, you cannot reliably perform later steps such as valuing assets, selecting appropriate safeguards, or measuring whether controls are adequate and cost-effective. In CISSP terms, this aligns with foundational asset management practices: you establish an asset inventory as an input to risk management, classification, and control selection. Once the inventory exists, you can then proceed to asset valuation (to understand impact), determine required protection levels, and choose deterrent/detective/physical controls accordingly. This sequencing is also consistent with common security governance approaches where asset identification precedes risk analysis and treatment decisions. For additional context on the centrality of asset inventories to security and risk programs, see NIST guidance on asset management and inventories in security programs and controls: [NIST SP 800-53 Rev. 5](#) and NIST's overview of cybersecurity supply chain/asset management concepts: [NIST Cybersecurity Framework](#).

QUESTION NO: 23

Which of the following processes is BEST used to determine the extent to which modifications to an information system affect the security

posture of the system?

- A. Patch management
- B. Continuous monitoring
- C. Configuration change control
- D. Security impact analysis

ANSWER: D

Explanation:

Security impact analysis is the best process for determining how proposed or implemented modifications (such as patches, configuration changes, new components, or code updates) affect an information system's security posture. In mature change management, a security impact analysis is performed as part of evaluating a change request to identify security-relevant effects: shifts in attack surface, changes to trust boundaries, impacts to confidentiality/integrity/availability, new vulnerabilities introduced, control gaps created, and any required updates to security documentation (e.g., SSPs, baselines, and risk assessments). This process directly answers the question's focus on "extent" of impact by explicitly assessing security consequences and required compensating controls before approval and deployment. It is also closely aligned with formal risk-based decision-making and authorization practices used in governance frameworks, where changes are evaluated for security significance and may trigger additional testing, re-authorization activities, or enhanced monitoring. In contrast,

operational processes like patch management or continuous monitoring support security over time, but they are not the primary evaluative method for assessing the security implications of a specific modification.

References: [NIST SP 800-128 \(Security-Focused Configuration Management\)](#), [NIST SP 800-37 Rev. 2 \(Risk Management Framework\)](#)

QUESTION NO: 24

Which of the following is the PRIMARY benefit of implementing an Information Security Management System (ISMS)?

- A. Correlates system events to monitor and demonstrate system health
- B. Improves customer confidence by demonstrating adherence to best practices
- C. Increases employee education and awareness of security policies
- D. Ensures user compliance with computing standards

ANSWER: B

Explanation:

The primary benefit of implementing an Information Security Management System (ISMS) is that it provides a structured, risk-based management framework to protect information assets through governance, defined processes, and continual improvement. In practice, an ISMS (as described by ISO/IEC 27001) formalizes how an organization identifies information security risks, selects and operates controls, assigns accountability, measures performance, and improves over time using a continual improvement cycle. A key business outcome of having this recognized, auditable system in place is increased trust from customers and other stakeholders because the organization can demonstrate that it follows established best practices and manages security systematically rather than ad hoc. This “demonstrability” is often realized through internal assurance and, when applicable, third-party certification, which can reduce due-diligence friction and improve market confidence. While training, monitoring, and compliance activities can be components within an ISMS, the overarching primary benefit is the ability to consistently manage information security in a way that is provable and aligned to business risk and stakeholder expectations.

References: [ISO/IEC 27001 overview \(ISO\)](#), [ISO/IEC 27001 standard page \(ISO\)](#)

QUESTION NO: 25

In supervisory control and data acquisition (SCADA) systems, which of the following controls can be used to reduce device exposure to malware?

- A. Disallow untested code in the execution space of the SCADA device.
- B. Disable all command line interfaces.
- C. Disable Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) port 138 and 139 on the SCADA device.
- D. Prohibit the use of unsecure scripting languages.

ANSWER: A

Explanation:

Disallow untested code in the execution space of the SCADA device is the best control listed to reduce malware exposure because it aligns with application allowlisting (also called execution control). In industrial control/SCADA environments, where systems are often fragile, hard to patch, and expected to run a stable set of known software, preventing unknown or unapproved code from executing is a primary defensive measure. This approach reduces the attack surface for common malware delivery methods (infected removable media, compromised updates, lateral movement) by blocking execution of

binaries, scripts, and libraries that are not explicitly approved and validated for the device. It is also consistent with ICS security guidance that emphasizes strict control of software changes and the use of allowlisting to prevent unauthorized code execution on critical assets. This control is broadly applicable across vendors and protocols and does not depend on specific ports or interfaces; instead, it directly targets the malware's ability to run, which is the key exposure being reduced.

References: [NIST SP 800-82 Rev. 2 \(Guide to Industrial Control Systems Security\)](#), [CISA ICS Recommended Practices](#)

QUESTION NO: 26

What are facets of trustworthy software in supply chain operations?

- A. Functionality, safety, reliability, integrity, and accuracy.
- B. Confidentiality, integrity, availability, authenticity, and possession.
- C. Safety, reliability, availability, resilience, and security.
- D. Reparability, security, upgradability, functionality, and accuracy.

ANSWER: C

Explanation:

Safety, reliability, availability, resilience, and security are commonly cited facets of trustworthy software, especially when discussing software assurance and supply chain risk. In a supply chain context, "trustworthy" software must behave as intended (reliability), remain usable when needed (availability), withstand and recover from faults or attacks (resilience), avoid causing harm (safety), and resist compromise across its lifecycle (security). These facets align with how industry and government guidance frames software trustworthiness: not as a single security property, but as a set of quality and assurance characteristics that collectively reduce operational and systemic risk from third-party components, build pipelines, and delivered artifacts. This framing is consistent with software assurance perspectives used in supply chain security programs, where the goal is to reduce the likelihood and impact of vulnerabilities, malicious code insertion, and operational failures introduced anywhere from development through distribution and deployment. For additional context on software supply chain security practices and the broader trustworthiness goals they support, see NIST's Secure Software Development Framework guidance (<https://csrc.nist.gov/publications/detail/sp/800-218/final>) and CISA's software supply chain resources (<https://www.cisa.gov/supply-chain>).

QUESTION NO: 27

Which of the following is not classified as "Security and Audit Frameworks and Methodologies"?

- A. Bell LaPadula.
- B. Committee of Sponsoring Organizations of the Treadway Commission (COSO).
- C. IT Infrastructure Library (ITIL).
- D. Control Objectives for Information and related Technology (COBIT).

ANSWER: C

Explanation:

IT Infrastructure Library (ITIL) is the best answer because it is primarily an IT service management (ITSM) framework focused on designing, delivering, operating, and continually improving IT services. While ITIL can support security objectives indirectly (for example, through change management, incident management, and service continuity practices), it is not a security or audit framework/methodology in the way governance/control frameworks are. In contrast, commonly recognized security and audit frameworks and methodologies emphasize internal control structures, governance, assurance, and

auditability of IT and business processes. ITIL's core purpose is service lifecycle/value stream management rather than defining control objectives for audits or providing a security evaluation methodology. This distinction is consistent with how ISC2 typically categorizes frameworks: ITIL aligns to service management, whereas audit/security frameworks align to governance, controls, and assurance. For reference, see the official ITIL overview describing its IT service management focus at [AXELOS ITIL](#) and a general description of ITSM/ITIL positioning at [IBM: What is ITIL?](#).

QUESTION NO: 28

During a Disaster Recovery (DR) simulation, it is discovered that the shared recovery site lacks adequate data restoration capabilities to support the implementation of multiple plans simultaneously.

What would be impacted by this fact if left unchanged?

- A. Recovery Point Objective (RPO).
- B. Recovery Time Objective (RTO).
- C. Business Impact Analysis (BIA).
- D. Return on Investment (ROI).

ANSWER: B

Explanation:

Recovery Time Objective (RTO) would be impacted because inadequate data restoration capability at a shared recovery site directly affects how quickly systems and services can be brought back into operation after a disruption. In a shared (often "hot/warm/cold" or subscription-based) recovery environment, multiple organizations may need to restore concurrently during a regional event. If the site cannot support multiple restorations at the same time—due to limited restore infrastructure, bandwidth, storage throughput, staffing, or tooling—then restoration activities will queue or be throttled. That delay extends the elapsed time to recover applications and data to an operational state, which is exactly what RTO measures: the maximum tolerable downtime for a business process or system. While data restoration limitations can indirectly influence other continuity metrics, the most immediate and direct consequence of insufficient restoration capacity is that recovery will take longer than planned, causing RTO commitments to be missed unless capacity, prioritization, or recovery sequencing is redesigned. This is a classic outcome of DR testing: validating whether recovery capabilities can meet stated time-based objectives under realistic load and contention conditions.

References: [NIST SP 800-34 Rev. 1 \(Contingency Planning Guide\)](#), [ISC2 CISSP Exam Outline \(Domain 8: Business Continuity/DR concepts\)](#)

QUESTION NO: 29

What is the MOST common component of a vulnerability management framework?

- A. Risk analysis.
- B. Patch management.
- C. Threat analysis.
- D. Backup management.

ANSWER: B

Explanation:

Patch management is the most common component of a vulnerability management framework because vulnerability management is fundamentally about finding weaknesses and then reducing exposure by remediating them, and the most frequent, repeatable remediation action in most environments is applying vendor fixes. In practice, vulnerability scanning and assessment activities typically feed directly into a patching workflow: identify missing patches or vulnerable versions, prioritize based on severity and asset criticality, test changes, deploy updates, and verify closure. This “detect → remediate → verify” cycle is central to mature vulnerability management programs and is widely reflected in industry guidance and operational tooling (ticketing, endpoint management, configuration management, and continuous monitoring). While broader activities like risk analysis may influence prioritization, patch management is the day-to-day mechanism organizations most commonly use to eliminate known vulnerabilities at scale across operating systems, applications, and firmware. This aligns with established security practice and common control frameworks that emphasize timely remediation of identified vulnerabilities through patching and update processes. See NIST guidance on vulnerability remediation and patching practices in [NIST SP 800-40 Rev. 4](#) and NIST’s vulnerability management lifecycle concepts in [NIST SP 800-53 Rev. 5](#) (e.g., flaw remediation controls).

QUESTION NO: 30

What is the PRIMARY consideration when testing industrial control systems (ICS) for security weaknesses?

- A. ICS often run on UNIX operating systems.
- B. ICS often do not have availability requirements.
- C. ICS are often sensitive to unexpected traffic.
- D. ICS are often isolated and difficult to access.

ANSWER: C

Explanation:

ICS are often sensitive to unexpected traffic is the primary consideration because many control environments are engineered for deterministic, real-time operation and can be fragile when exposed to scanning, fuzzing, aggressive vulnerability assessment, or other test traffic patterns common in IT security testing. Unlike typical enterprise systems, ICS components (PLCs, RTUs, HMIs, historians, and proprietary fieldbus/industrial Ethernet networks) may have limited resources, legacy protocol stacks, and strict timing constraints; even “benign” discovery probes can cause device faults, communication timeouts, process upsets, or safety impacts. As a result, the key security-testing principle in ICS is to avoid disrupting operations and to tightly control test methods, scope, and timing—often favoring passive monitoring, vendor-approved procedures, lab replicas/digital twins, maintenance windows, and strong coordination with operations and safety teams. This aligns with widely accepted ICS guidance that availability and safety are paramount and that active testing must be carefully managed to prevent unintended consequences in physical processes. See CISA’s ICS security guidance for recommended approaches and cautions around testing and scanning, and NIST’s ICS security guide for considerations on assessment methods and operational impact.

References: <https://www.cisa.gov/topics/industrial-control-systems>, <https://csrc.nist.gov/publications/detail/sp/800-82/rev-2/final>

QUESTION NO: 31

In the course of responding to and handling an incident, you work on determining the root cause of the incident. In which step are you in?

- A. Recovery.
- B. Containment.
- C. Triage.
- D. Analysis and tracking.

ANSWER: D

Explanation:

Determining the root cause of an incident is part of the incident *analysis* work performed during incident response. In this phase, responders examine evidence (logs, alerts, host artifacts, network captures, timelines, and indicators of compromise) to understand what happened, how it happened, what systems/accounts were affected, and—critically—why it happened (the underlying weakness, misconfiguration, vulnerability, or control failure). Root cause determination supports accurate scoping, eradication planning, and prevention of recurrence, and it is typically documented and tracked as the investigation progresses. This aligns with common incident handling lifecycles (such as NIST’s incident response guidance), where analysis is the step focused on validating the incident, characterizing it, and identifying contributing factors and entry vectors before moving into full remediation and longer-term improvements. While containment and recovery are essential operational steps, they focus on limiting damage and restoring services; root cause analysis is fundamentally an investigative activity that belongs in the analysis/tracking portion of the response process.

References: [NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide](#); [CISA Incident Response resources](#).

QUESTION NO: 32

A security practitioner needs to implement a solution to verify endpoint security protections and operating system (OS) versions. Which of the

following is the

BEST solution to implement?

- A. An intrusion prevention system (IPS)
- B. Network Access Control (NAC)
- C. Active Directory (AD) authentication
- D. A firewall

ANSWER: B

Explanation:

Network Access Control (NAC) is the best solution because it is specifically designed to assess endpoint posture before granting (or while maintaining) network access. NAC solutions commonly perform posture checks such as verifying OS version/patch level, presence and status of endpoint protection (for example, antivirus/EDR), host firewall state, disk encryption, and other compliance attributes. Based on the results, NAC can enforce policy decisions like allowing full access, placing the device into a restricted/quarantine VLAN, requiring remediation, or denying access entirely. This directly matches the requirement to verify endpoint security protections and OS versions in a centralized, enforceable way at the point of network admission. In CISSP terms, NAC is a preventive and detective control that supports strong access control by tying network connectivity to device compliance, reducing the risk of unmanaged or noncompliant endpoints connecting to the environment. NAC is also commonly integrated with directory services and endpoint management tools to improve accuracy and automate remediation workflows.

References: [NIST CSRC Glossary – Network Access Control](#), [Cisco Identity Services Engine \(ISE\) overview \(NAC posture assessment\)](#)

QUESTION NO: 33

Which of the following is a unique feature of attribute-based access control (ABAC)?

- A. A user is granted access to a system at a particular time of day.
- B. A user is granted access to a system based on username and password.
- C. A user is granted access to a system based on group affinity.

D. A user is granted access to a system with biometric authentication.

ANSWER: A

Explanation:

“A user is granted access to a system at a particular time of day.” best reflects a unique feature of attribute-based access control (ABAC) because ABAC makes authorization decisions by evaluating attributes about the subject (user), object (resource), action, and environment (context). Time of day is a classic example of an environmental attribute used in ABAC policies (for example, allowing access only during business hours, or denying access outside a maintenance window). This attribute-driven, policy-based evaluation is what distinguishes ABAC from models that primarily rely on static identity, group/role membership, or a single authentication factor. In ABAC, the same user may be permitted or denied depending on contextual attributes like time, location, device posture, data classification, and requested operation, enabling fine-grained and dynamic access control decisions. This aligns with common ABAC descriptions in standards and guidance where policies combine multiple attributes and conditions to reach an authorization decision.

References: [NIST CSRC: Attribute Based Access Control \(ABAC\)](#), [NIST SP 800-162 \(ABAC Definition and Considerations\)](#)

QUESTION NO: 34

Contracts and agreements are often times unenforceable or hard to enforce in which of the following alternate facility recovery agreement?

- A. hot site.
- B. warm site.
- C. cold site.
- D. reciprocal agreement.

ANSWER: D

Explanation:

“reciprocal agreement.” is the correct choice because reciprocal (or mutual aid) recovery arrangements rely on another organization providing space, systems, or support during a disaster, typically based on a promise to do the same in return. In real incidents, these agreements are frequently difficult to enforce because both parties may be impacted by the same regional event, may have competing priorities, or may lack the spare capacity they assumed they would have. Even when a contract exists, the practical ability to deliver the promised resources under stress can be limited, and the agreement may not contain sufficiently strong service-level commitments, testing requirements, or penalties to ensure performance. CISSP disaster recovery guidance generally treats reciprocal agreements as higher-risk than commercial hot/warm/cold site contracts because they depend heavily on goodwill, availability, and aligned business conditions at the time of need. As a result, organizations are expected to validate feasibility through due diligence and regular testing, and to prefer more reliable alternate-site strategies when recovery time objectives are tight.

References: [Ready.gov – Business Continuity Planning](#), [NIST SP 800-34 Rev. 1 \(Contingency Planning Guide\)](#)

QUESTION NO: 35

A security evaluation report and an accreditation statement are produced in which of the following phases of the system development life cycle?

- A. project initiation and planning phase.
- B. system design specification phase

C. development & documentation phase

D. acceptance phase.

ANSWER: D

Explanation:

The acceptance phase is where a system is formally evaluated against its security requirements and then either approved for operation or sent back for remediation. In classic SDLC security integration (and in the older certification-and-accreditation approach), this is the point at which security testing, assessment, and review activities culminate in formal deliverables: a security evaluation report (documenting the results of the security assessment, residual risk, and any findings) and an accreditation statement (the management/authorizing official's decision to accept the residual risk and authorize the system to operate). This aligns with the idea that accreditation/authorization is a management decision made after the technical evaluation is complete and before the system is placed into production. In modern terminology (NIST RMF), this maps closely to the "Assess" step producing the security assessment report and the "Authorize" step producing the authorization decision, which occur just before ongoing operations/monitoring. For CISSP exam purposes, these artifacts are associated with the SDLC acceptance activities that precede operational deployment.

References: [NIST SP 800-37 Rev. 2 \(RMF: Assess/Authorize\)](#), [NIST SP 800-53A Rev. 5 \(Security assessment reporting\)](#)

QUESTION NO: 36

An employee's home address should be categorized according to which of the following references?

- A. The consent form terms and conditions signed by employees.
- B. An organization security plan for human resources.
- C. Existing employee data classifications.
- D. The organization's data classification model.

ANSWER: D

Explanation:

The organization's data classification model is the correct reference because classification is a formal, enterprise-wide scheme that defines how information types (including personally identifiable information such as an employee's home address) are labeled and handled. In CISSP terms, classification is driven by business requirements and risk, and it must be applied consistently across the organization so that appropriate protection requirements (handling, storage, transmission, retention, and disposal) follow from the assigned class. An employee home address is typically treated as sensitive personal data/PII and should be categorized using the organization's established classification categories (for example, Public, Internal, Confidential, Restricted) and associated handling standards. While consent forms and HR security plans may influence lawful processing and procedural controls, they are not the authoritative mechanism for assigning the organization's information classification label. Using the formal data classification model ensures consistent application of safeguards and aligns with governance expectations for information security programs. For additional context on data classification concepts and handling requirements, see NIST's overview of protecting controlled unclassified information ([NIST SP 800-171](#)) and NIST's guidance on protecting PII ([NIST SP 800-122](#)).

QUESTION NO: 37

An organization is looking to include mobile devices in its asset management system for better tracking. In which system tier of the reference

architecture would mobile devices be tracked?

- A.

- A. 0
- B.
- B. 1
- C.
- C. 2
- D.
- D. 3
- E. Endpoint / end-user device tier (client devices).

ANSWER: E

Explanation:

Mobile devices are considered endpoint assets (client devices) and are therefore tracked in the endpoint/user device tier of a typical enterprise reference architecture used for asset management. In CISSP terms, asset management inventories should include all hardware that stores, processes, or transmits organizational data, and mobile phones/tablets fall squarely into the “end-user device/endpoint” category alongside laptops and desktops. This tier is where organizations maintain device identity (serial/IMEI), ownership, configuration baseline, installed software, compliance state (encryption, screen lock, OS version), and lifecycle status (issued, in repair, retired). In practice, this is commonly implemented through endpoint management/MDM/UEM tooling that feeds the authoritative asset inventory (CMDB/ITAM). Treating mobile devices as endpoints also aligns with security operations expectations: they are managed, monitored, and controlled at the edge where users interact with corporate services, rather than being part of server, network, or application tiers. For additional context on how mobile devices are managed as endpoints via MDM/UEM, see [NIST SP 800-124 Rev. 2](#) and Microsoft’s overview of mobile device management at [Microsoft Intune documentation](#).

QUESTION NO: 38

Which of the following is the MAIN difference between a network-based firewall and a host-based firewall?

- A. A network-based firewall is stateful, while a host-based firewall is stateless.
- B. A network-based firewall blocks network intrusions, while a host-based firewall blocks malware.
- C. A network-based firewall controls traffic passing through the device, while a host-based firewall controls traffic destined for the device.
- D. A network-based firewall verifies network traffic, while a host-based firewall verifies processes and applications.

ANSWER: C

Explanation:

The main difference is scope and enforcement point: a network-based firewall enforces policy at a network boundary (or between network segments) and therefore controls traffic that is passing through the firewall between networks, while a host-based firewall runs on an individual endpoint and controls traffic to and from that specific device. In other words, the network-based firewall is typically positioned inline to filter transit traffic for many systems, whereas the host-based firewall is local to a single host and applies rules based on that host’s interfaces, services, and exposure. This distinction is central in CISSP terms because it affects what each control can see and protect: network firewalls provide centralized, perimeter/segmentation control, while host firewalls provide per-endpoint protection (including when the device is off-network or on untrusted networks) and can be managed as part of endpoint security hardening. Many modern implementations of both types can be stateful, and both can contribute to intrusion and malware defense, but their defining difference is where they sit and which traffic they are responsible for filtering. See [NIST CSRC glossary: Firewall](#) and [Microsoft documentation on host-based firewall concepts](#).

QUESTION NO: 39

Using the cipher text and resultant cleartext message to derive the monoalphabetic cipher key is an example of which method of cryptanalytic

attack?

- A. Known-plaintext attack
- B. Ciphertext-only attack
- C. Frequency analysis
- D. Probable-plaintext attack

ANSWER: A

Explanation:

The correct method is a known-plaintext attack. In a known-plaintext attack, the attacker has access to one or more pairs of corresponding ciphertext and its resulting plaintext (cleartext). With those pairs, the attacker attempts to derive the key or otherwise recover the ability to decrypt additional ciphertexts. For a monoalphabetic substitution cipher, having matching ciphertext and plaintext segments can directly reveal the substitution mapping (the “key”) because each plaintext letter consistently maps to a single ciphertext letter. As more plaintext/ciphertext pairs are observed, the attacker can reconstruct more of the substitution table until the full key is recovered, enabling decryption of other messages encrypted with the same substitution. This aligns with standard cryptanalytic taxonomy used in CISSP: attacks are categorized by what information the attacker possesses (ciphertext-only, known-plaintext, chosen-plaintext, etc.), and the presence of both ciphertext and the resultant cleartext is the defining characteristic of known-plaintext. See NIST’s overview of cryptographic concepts and attack models in [NIST SP 800-175B](#) and a concise description of known-plaintext attacks in [Known-plaintext attack](#).

QUESTION NO: 40

What is the

BEST answer pertaining to the difference between the Session and Transport layers of the OSI model?

- A. The Session layer sets up communication between protocols, while the Transport layer sets up connections between computer systems.
- B. The Transport layer sets up communication between computer systems, while the Session layer sets up connections between applications.
- C. The Session layer sets up communication between computer systems, while the Transport layer sets up connections between protocols.
- D. The Transport layer sets up communication between applications, while the Session layer sets up connections between computer systems.

ANSWER: B

Explanation:

The best answer is: “The Transport layer sets up communication between computer systems, while the Session layer sets up connections between applications.” In the OSI model, the Transport layer (Layer 4) is responsible for end-to-end transport services between hosts, including establishing and managing transport connections, segmentation/reassembly, reliability (when provided, as with TCP), and flow control. In other words, it provides logical communication between two systems across the network, abstracting the underlying network path.

The Session layer (Layer 5) sits above Transport and focuses on managing dialogs (sessions) between applications. It helps establish, manage, and terminate sessions, and can provide dialog control (who can transmit when), synchronization/checkpointing, and session recovery. While many modern protocol stacks don’t implement a distinct

“Session layer” as a separate protocol in the same way the OSI model describes, the conceptual distinction remains important for CISSP: Transport is host-to-host delivery services, whereas Session is application-to-application session management and control.

References: [Cloudflare OSI Model overview](#), [GeeksforGeeks OSI model layers](#).

QUESTION NO: 41

An organization is developing employee training content to increase awareness of Payment Card Industry (PCI) standards. What are the three types of awareness roles applicable to the organization?

- A. All personnel, specialized, management.
- B. Standard, privileged, administrator.
- C. Basic, intermediate, advanced.
- D. Technical, operational, administrative.

ANSWER: A

Explanation:

The correct set of awareness roles is “All personnel, specialized, management.” PCI DSS requires organizations to implement a formal security awareness program and to ensure personnel are trained appropriately for their job responsibilities. In practice, PCI-aligned awareness is commonly structured into tiers: training for all personnel (baseline awareness for anyone who could impact cardholder data security), specialized training (role-based training for staff with specific security or operational responsibilities such as developers, system administrators, help desk, or those handling payment processes), and management training (to ensure leaders understand governance expectations, risk, compliance obligations, and how to sponsor/enforce security practices). This tiering aligns with the intent of PCI DSS Requirement 12.6, which emphasizes security awareness for all personnel and role-appropriate training, and it matches how organizations typically segment content to be effective and auditable. Using these three categories helps ensure broad coverage while still meeting the “appropriate to job function” expectation that PCI assessors look for during compliance validation.

References: [PCI SSC Document Library \(PCI DSS and guidance\)](#), [PCI SSC – Maintaining Payment Security](#)

QUESTION NO: 42

A security practitioner has been asked to model best practices for disaster recovery (DR) and business continuity. The practitioner has decided

that a formal committee is needed to establish a business continuity policy. Which of the following BEST describes this stage of business continuity development?

- A. Developing and Implementing business continuity plans (BCP)
- B. Project Initiation and Management
- C. Risk Evaluation and Control
- D. Business impact analysis (BIA)

ANSWER: B

Explanation:

The stage described is *Project Initiation and Management*. In ISC2-aligned business continuity lifecycle models, the earliest phase focuses on obtaining executive sponsorship, defining scope and objectives, assigning roles and responsibilities, and establishing governance to drive the effort. Creating a formal committee to establish a business continuity policy is a

governance and program-management activity: it sets direction, ensures cross-functional representation (IT, facilities, HR, legal, operations), and provides authority for subsequent steps such as conducting the business impact analysis, selecting recovery strategies, and developing/maintaining plans. This “getting organized” work is foundational because without a chartered team and policy, later BCP/DR activities lack prioritization, funding, and enforcement. Once initiation and management are in place, the organization can proceed to analysis activities (including the BIA) and then strategy and plan development. For additional context on continuity program lifecycle concepts and governance expectations, see NIST guidance on contingency planning and program management: [NIST SP 800-34 Rev. 1](#) and ISO’s overview of business continuity management systems emphasizing leadership and planning: [ISO 22301 overview](#).

QUESTION NO: 43

Which of the following statements pertaining to PPTP (Point-to-Point Tunneling Protocol) is NOT true?

- A. PPTP allows the tunneling of any protocols that can be carried within PPP.
- B. PPTP does not provide strong encryption.
- C. PPTP does not support any token-based authentication method for users.
- D. PPTP is derived from L2TP.

ANSWER: D

Explanation:

The statement “PPTP is derived from L2TP.” is the one that is NOT true. Historically, PPTP and L2TP are separate Layer 2 tunneling approaches that emerged from different efforts and were later positioned as alternatives for VPN tunneling over IP networks. PPTP was developed earlier (led by a consortium including Microsoft) to tunnel PPP over IP using GRE, while L2TP resulted from combining ideas from Cisco’s L2F and Microsoft’s PPTP and was standardized by the IETF. In other words, L2TP is better understood as incorporating concepts from PPTP/L2F rather than PPTP being derived from L2TP. This distinction matters for CISSP-level understanding because it ties to protocol lineage, standardization, and security posture: PPTP is widely considered legacy and weak by modern standards, whereas L2TP is typically paired with IPsec for confidentiality and integrity. For background on L2TP’s origins and standardization, see [RFC 2661 \(Layer Two Tunneling Protocol “L2TP”\)](#). For Microsoft’s current guidance discouraging PPTP due to security weaknesses, see [Microsoft documentation on VPN technologies](#).

QUESTION NO: 44

What is the maximum key size for the RC5 algorithm?

- A. 128 bits.
- B. 256 bits
- C. 1024 bits
- D. 2040 bits.

ANSWER: D

Explanation:

RC5 is a parameterized symmetric block cipher designed by Ronald Rivest. Its specification allows variable block sizes, a variable number of rounds, and—most relevant here—a variable-length secret key. The key length for RC5 is defined as b bytes, where b can range from 0 to 255 bytes. That means the maximum key size is $255 \text{ bytes} \times 8 \text{ bits/byte} = 2040 \text{ bits}$. This is a distinctive property of RC5 compared to many other common symmetric algorithms that have fixed key sizes (for example, AES at 128/192/256 bits). For CISSP-style questions, the key point is remembering that RC5 supports very large, variable key sizes up to 2040 bits, even though such large keys are not typically used in practice. This maximum comes directly from the original RC5 design parameters and is widely cited in cryptography references and standards discussions.

See the RC5 overview and parameter details in <https://en.wikipedia.org/wiki/RC5> and a general description of RC5's variable key length in <https://www.cryptopp.com/wiki/RC5>.

QUESTION NO: 45

Which of the following statements do apply to a hot site?

- A. It is expensive.
- B. There are cases of common overselling of processing capabilities by the service provider.
- C. It provides a false sense of security.
- D. It is accessible on a first come first serve basis. In case of large disaster it might Be accessible.

ANSWER: A

Explanation:

A hot site is a fully equipped alternate processing facility intended to support rapid recovery, typically with preinstalled hardware, network connectivity, environmental controls, and often near-real-time data replication or frequent backups so operations can resume with minimal downtime. Because it is maintained in a ready-to-run state (space, power, cooling, servers, networking, licensing, and ongoing synchronization/testing), it is generally the most costly type of recovery site compared to warm and cold sites. This higher cost is a direct consequence of paying for capacity that is continuously available and periodically validated through exercises and maintenance. In business continuity planning terms, hot sites are selected when the organization's recovery time objective is very short and the business impact of downtime is high, justifying the expense. For additional background on hot sites and alternate site types, see [Ready.gov – Business Continuity Planning](#) and [TechTarget – hot site definition](#).

QUESTION NO: 46

Which of the following value comparisons MOST accurately reflects the agile development approach?

- A. Processes and tools over individuals and interactions
- B. Contract negotiation over customer collaboration
- C. Following a plan over responding to change
- D. Working software over comprehensive documentation.

ANSWER: D

Explanation:

“Working software over comprehensive documentation.” most accurately reflects the agile development approach because it is one of the four core values explicitly stated in the Agile Manifesto. Agile does not reject documentation; rather, it prioritizes delivering usable, tested increments of software as the primary measure of progress and the most reliable way to validate requirements and reduce risk. In practice, agile teams aim to keep documentation “just enough” to support shared understanding, onboarding, operations, and compliance needs, while avoiding excessive up-front documentation that can become outdated as requirements evolve. This value aligns well with iterative delivery, frequent feedback, and continuous improvement—key characteristics of agile methods such as Scrum and XP. From a CISSP perspective, this also maps to secure SDLC expectations: security requirements and controls should be integrated into each iteration, with evidence and artifacts produced as needed, but the emphasis remains on delivering secure, functioning software increments that can be assessed and improved continuously. The authoritative source for this value comparison is the Agile Manifesto itself.

References: [Agile Manifesto](#), [The Scrum Guide](#)

QUESTION NO: 47

A breach investigation found a website was exploited through an open source component. What is the FIRST step in the process that could have prevented this breach?

- A. Application whitelisting
- B. Vulnerability remediation
- C. Web application firewall (WAF)
- D. Software inventory

ANSWER: D

Explanation:

Software inventory is the first step because you can't effectively manage or secure what you don't know you have. In the context of an open source component being exploited, the preventive process starts with identifying and tracking all software components and dependencies in use (often implemented as an asset inventory plus a software bill of materials, SBOM). Once the organization has an accurate inventory, it can then perform vulnerability identification (e.g., mapping components/versions to known CVEs), prioritize risk, and execute vulnerability remediation/patching or compensating controls. Without inventory, remediation and monitoring efforts are incomplete and may miss vulnerable libraries embedded in applications. This aligns with widely accepted security practice and control frameworks that place asset/software inventory at the foundation of vulnerability management and secure software supply chain practices. Maintaining an inventory also supports ongoing governance activities like version control, end-of-life tracking, and rapid response when new vulnerabilities (such as in common open source libraries) are disclosed.

References: [NIST SP 800-53 Rev. 5 \(CM-8 System Component Inventory\)](#), [OWASP Software Component Verification Standard \(SCVS\)](#)

QUESTION NO: 48

Which of the following was designed as a more fault-tolerant topology than Ethernet, and very resilient when properly implemented?

- A. Token Link.
- B. Token system.
- C. Token Ring.
- D. Duplicate ring.

ANSWER: C

Explanation:

Token Ring is the topology/protocol family commonly cited in CISSP study materials as having been designed with stronger determinism and resilience characteristics than classic shared-media Ethernet. Token Ring uses token passing (only the station holding the token transmits), which avoids collisions and provides predictable access behavior. In addition, Token Ring deployments were often implemented with physical/logical ring management features (for example, MAUs that can bypass a failed station) and could be built with redundant ring paths in some implementations, improving availability when properly engineered. While modern switched Ethernet is highly reliable, the question is framed historically and conceptually around "designed as more fault-tolerant than Ethernet," which aligns with Token Ring's design goals and operational characteristics in enterprise LANs. This is why Token Ring is the best match for a fault-tolerant, resilient topology in the provided choices.

References: [IETF RFC 1042 \(IP over IEEE 802 networks, including Token Ring\)](#), [IBM Docs: Token-Ring network overview](#)

QUESTION NO: 49

Which of the following is NOT an example of a detective control?

- A. System Monitor.
- B. IDS.
- C. Motion detector.
- D. Backup data restore.

ANSWER: D

Explanation:

“Backup data restore.” is not a detective control; it is a corrective (and often also recovery) control. Detective controls are designed to discover or identify that an event has occurred or is occurring—typically by monitoring, alerting, or logging—so that an organization can respond. In contrast, restoring from backup is an action taken after an incident or failure to return systems and data to a known-good state. That makes it part of correcting the impact of an event and enabling recovery, rather than detecting the event in the first place. In CISSP terms, backups themselves are commonly discussed as recovery/corrective measures, and the act of restoring is clearly a corrective response activity that helps re-establish availability and integrity after loss, corruption, ransomware, or operational errors. This aligns with widely used security control taxonomies where “detective” focuses on identification (for example, monitoring and intrusion detection), while “corrective/recovery” focuses on remediation and restoration of service. For additional context on control types and how they map to functions like detection and recovery, see NIST SP 800-53 control families and discussion of security control functions in [NIST SP 800-53 Rev. 5](#) and the NIST Cybersecurity Framework functions (Detect vs. Recover) in [NIST CSF](#).

QUESTION NO: 50

Why are mobile devices sometimes difficult to investigate in a forensic examination?

- A. There are no forensics tools available for examination.
- B. They may contain cryptographic protection.
- C. They have password-based security at logon.
- D. They may have proprietary software installed to protect them.

ANSWER: B

Explanation:

They may contain cryptographic protection is the best answer because modern mobile devices commonly use strong, hardware-backed encryption that can prevent forensic acquisition and analysis without the correct unlock credentials. Full-disk/file-based encryption (often tied to a device’s secure element/TPM-like component such as Secure Enclave or Titan M) means that even if an examiner can physically access storage, the data at rest is not intelligible without the keys. Those keys are typically derived from the user’s passcode/biometrics and protected by rate limiting, secure key storage, and anti-hammering controls, which significantly increases the time and complexity of lawful access attempts. In addition, encryption is frequently integrated with secure boot and application sandboxing, limiting the ability to bypass the operating system to extract artifacts. From a CISSP perspective, this maps to the core forensic challenge of acquiring a forensically sound image when confidentiality controls (encryption) are correctly implemented, making evidence collection dependent on key recovery, device state (locked/unlocked), or vendor-assisted access under legal process.

References: [Apple Platform Security – Data Protection](#), [Android Open Source Project – Encryption](#)

QUESTION NO: 51

Which of the following criteria ensures information is protected relative to its importance to the organization?

- A. Legal requirements, value, criticality, and sensitivity to unauthorized disclosure or modification.
- B. The value of the data to the organization's senior management.
- C. Organizational stakeholders, with classification approved by the management board.
- D. Legal requirements determined by the organization headquarters' location.

ANSWER: A

Explanation:

The correct criterion is “Legal requirements, value, criticality, and sensitivity to unauthorized disclosure or modification.” In CISSP terms, protecting information relative to its importance is the purpose of an information classification scheme. A sound classification approach considers multiple drivers: external obligations (laws, regulations, contracts), business value and mission impact (criticality), and the potential harm from loss of confidentiality and integrity (sensitivity to unauthorized disclosure or modification). These factors collectively determine the appropriate classification level and, in turn, the baseline handling requirements (labeling, access controls, encryption, retention, transmission rules, and disposal). Using only a single dimension—such as what senior management values most—doesn’t reliably capture regulatory exposure or operational impact. Likewise, focusing only on headquarters location misses multi-jurisdictional and contractual requirements. ISC2 emphasizes that classification should be risk-based and aligned to business needs and compliance obligations, ensuring controls are commensurate with the information’s importance and potential impact if compromised. For additional background on classification concepts and how sensitivity/impact drive protection requirements, see [NIST FIPS 199](#) and [NIST SP 800-60](#).

QUESTION NO: 52

Which is the PRIMARY mechanism for providing the workforce with the information needed to protect an agency's vital information resources?

- A. Implementation of access provisioning process for coordinating the creation of user accounts
- B. Incorporating security awareness and training as part of the overall information security program
- C. An information technology (IT) security policy to preserve the confidentiality, integrity, and availability of systems
- D. Execution of periodic security and privacy assessments to the organization

ANSWER: B

Explanation:

Incorporating security awareness and training as part of the overall information security program is the primary mechanism because it is the direct, organization-wide method for ensuring personnel actually receive, understand, and can apply the guidance needed to protect information resources. In CISSP terms, policies set direction, but awareness and training operationalize that direction by communicating expectations, teaching secure behaviors, and building role-appropriate competence across the workforce. This includes baseline security awareness for all users (e.g., phishing recognition, data handling, incident reporting) and targeted training for privileged users and specialized roles. In many regulatory and standards frameworks, awareness and training are explicitly required as a foundational control to reduce human-related risk and to ensure consistent protection practices across the enterprise. For example, NIST SP 800-53 includes the Awareness and Training (AT) control family, emphasizing that organizations must provide security awareness and role-based training to personnel. Similarly, NIST SP 800-50 focuses on building an effective IT security awareness and training program as a core element of an agency’s security posture. These references align with the question’s focus on providing the workforce with the information they need, which is best achieved through a structured awareness and training program.

QUESTION NO: 53

Which of the following is the MOST common method of memory protection?

- A. Error correction.
- B. Virtual local area network (VLAN) tagging
- C. Segmentation.
- D. Compartmentalization.

ANSWER: C

Explanation:

The MOST common method of memory protection is **Segmentation**. In operating systems and CPU memory-management designs, segmentation is a foundational technique used to separate a process's memory into logical regions (segments) such as code, data, heap, and stack. Each segment can have its own base address, limit, and access permissions, which enables the system to enforce boundaries and prevent one segment (or one process) from reading or writing outside its authorized range. This is a core goal of memory protection: isolating execution contexts and preventing unintended or malicious memory access that could lead to corruption, privilege escalation, or information disclosure.

Segmentation is also commonly discussed alongside paging and virtual memory; many real-world architectures combine these concepts, but segmentation remains a classic and widely taught mechanism for implementing protection domains and bounds checking at the memory-management unit (MMU) level. From a CISSP perspective, segmentation is a standard, broadly applicable control that maps directly to enforcing least privilege and isolation within system memory.

References: https://en.wikipedia.org/wiki/Memory_segmentation, <https://learn.microsoft.com/en-us/windows/win32/memory/virtual-address-space>

QUESTION NO: 54

Which of the following does the security design process ensure within the System Development Life Cycle (SDLC)?

- A. Proper security controls, security objectives, and security goals are properly initiated.
- B. Security objectives, security goals, and system test are properly conducted.
- C. Proper security controls, security goals, and fault mitigation are properly conducted.
- D. Security goals, proper security controls, and validation are properly initiated.

ANSWER: A

Explanation:

Within the SDLC, the security design process ensures that security requirements are translated into an actionable architecture and set of safeguards for the system. In practice, this means selecting and specifying appropriate security controls (technical, administrative, and physical) that meet defined security objectives and support the organization's broader security goals. This is the point in the lifecycle where security is "built in" by design—controls are chosen, designed, and integrated into the system's architecture and detailed design artifacts so they can be implemented consistently and verified later. This aligns with the CISSP view of integrating security throughout the SDLC: requirements drive design decisions, and the design phase produces the security control specifications that will be implemented and then validated/verified during later testing and acceptance activities. While validation activities occur later, the design process ensures the system is set up so that validation can be performed against the stated goals and objectives, with controls defined clearly enough to test. See

NIST's SDLC guidance on incorporating security into system design and development and NIST's control selection concepts for how objectives/goals map to controls: <https://csrc.nist.gov/publications/detail/sp/800-64/rev-2/final> and <https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final>.

QUESTION NO: 55

Which of the following statements BEST describes least privilege principle in a cloud environment?

- A. A single cloud administrator is configured to access core functions.
- B. Internet traffic is inspected for all incoming and outgoing packets.
- C. Routing configurations are regularly updated with the latest routes.
- D. Network segments remain private if unneeded to access the internet.

ANSWER: A

Explanation:

The least privilege principle in a cloud environment means identities (users, admins, services, and workloads) are granted only the minimum permissions necessary to perform their required tasks, for only as long as needed, and ideally scoped as narrowly as possible (specific resources, actions, and conditions). In practice, this is implemented through fine-grained IAM policies/roles, separation of duties, and avoiding broad, standing administrative access. The statement "A single cloud administrator is configured to access core functions." most closely aligns with the idea of restricting access to only what is needed (core functions) rather than granting wide access to many people or systems. While real-world least privilege would typically avoid concentrating power in a single account and would use role-based access with just-in-time elevation, among the provided choices this option is the best match to the concept of limiting permissions. Least privilege is a foundational control for reducing blast radius in the event of credential compromise and for meeting governance requirements in shared-responsibility cloud models.

References: [NIST Glossary – Least Privilege](#), [Microsoft – Secure least-privileged access](#)

QUESTION NO: 56

Why is data classification control important to an organization?

- A. To enable data discovery
- B. To ensure security controls align with organizational risk appetite
- C. To ensure its integrity, confidentiality and availability
- D. To control data retention in alignment with organizational policies and regulation

ANSWER: B

Explanation:

To ensure security controls align with organizational risk appetite is correct because data classification is fundamentally a risk-based governance mechanism. By labeling information according to sensitivity and business impact (for example, public, internal, confidential, restricted), an organization can apply proportionate safeguards—such as encryption requirements, access control strength, monitoring, handling procedures, and sharing restrictions—where they are most needed. This prevents both under-protection of high-value data (increasing likelihood and impact of compromise) and over-protection of low-value data (wasting resources and impeding business operations). In CISSP terms, classification drives the selection and rigor of administrative, technical, and physical controls based on the organization's risk tolerance and compliance obligations, and it provides a consistent basis for data handling throughout its lifecycle. Classification also supports communicating protection expectations to users and system owners so that security investments and operational friction are

aligned with what the organization is willing to accept as residual risk. See NIST guidance on categorizing information and systems: <https://csrc.nist.gov/publications/detail/fips/199/final> and NIST SP 800-60 (security categorization process): <https://csrc.nist.gov/publications/detail/sp/800-60/vol-1-rev-1/final>.

QUESTION NO: 57

Which of the following methods is MOST effective in mitigating Cross-Site Scripting (XSS) vulnerabilities within HyperText Markup Language

(HTML) websites?

- A. Use antivirus and endpoint protection on the server to secure the web-based application
- B. Place the web-based system in a defined Demilitarized Zone (DMZ)
- C. Use .NET framework with .aspx extension to provide a higher level of security to the web application so that the web server display can be locked down
- D. Not returning any HTML tags to the browser client

ANSWER: D

Explanation:

The most effective way to mitigate Cross-Site Scripting is to ensure that untrusted input is never interpreted by the browser as executable markup or script. The option “Not returning any HTML tags to the browser client” best aligns with this goal because it implies rendering user-supplied content as plain text rather than as HTML. In practice, this is achieved through context-appropriate output encoding/escaping (for example, HTML-encoding characters like <, >, and & so the browser displays them instead of treating them as tags), combined with strong input validation and safe templating. This approach directly addresses the root cause of XSS: injection of attacker-controlled content into a page in a way the browser can execute. While real-world applications often must return HTML, the security principle remains the same: never allow untrusted data to be emitted into an HTML/JS context without proper encoding, and prefer frameworks that auto-escape by default. Additional defense-in-depth controls like Content Security Policy can further reduce impact, but correct output handling is the primary mitigation. See OWASP’s guidance on XSS prevention and output encoding: [OWASP XSS](#) and [OWASP XSS Prevention Cheat Sheet](#).

QUESTION NO: 58

Suppose you are a domain administrator and are choosing an employee to carry out backups. Which access control method would be BEST for this scenario?

- A. RBAC - Role-Based Access Control
- B. MAC - Mandatory Access Control
- C. DAC - Discretionary Access Control
- D. RBAC - Rule-Based Access Control.

ANSWER: A

Explanation:

RBAC - Role-Based Access Control is the best fit when you want to assign permissions based on a job function such as “Backup Operator” rather than granting broad administrative rights or managing permissions object-by-object. In this scenario, the goal is to delegate a specific operational task (performing backups) to an employee while keeping privileges tightly scoped and consistent across systems. RBAC supports least privilege by bundling the exact permissions required for backup activities into a role and then assigning that role to the selected employee (and removing it just as easily when duties change). This approach also improves administrative efficiency and auditability: access reviews can focus on role membership, and logs can be correlated to role-based entitlements. In Windows domain environments, this maps cleanly to

built-in or custom groups/roles (for example, Backup Operators) and aligns with common enterprise governance practices for separation of duties. Overall, RBAC is the standard access control model for delegating repeatable business tasks in a controlled, scalable way.

References: [NIST RBAC Project](#), [Microsoft documentation on Active Directory security groups](#)

QUESTION NO: 59

A Virtual Machine (VM) environment has five guest Operating Systems (OS) and provides strong isolation. What MUST an administrator review to audit a user's access to data files?

- A. Host VM monitor audit logs.
- B. Guest OS access controls.
- C. Host VM access controls.
- D. Guest OS audit logs.

ANSWER: D

Explanation:

Guest OS audit logs must be reviewed because file access (read/write/modify/delete) is enforced and recorded at the operating system and filesystem layer where the data files actually reside. In a strongly isolated VM environment, each guest OS is its own security boundary for identity, authorization, and object access; therefore, the authoritative evidence of a specific user's access to specific files is found in that guest's auditing subsystem (for example, Windows Security Event Log with Object Access auditing, or Linux auditd records). Hypervisor/host-level logs can show VM lifecycle events, console access, and some management actions, but they typically do not provide per-file, per-user access events inside the guest without additional in-guest agents or specialized introspection tooling. For an audit focused on "a user's access to data files," the required source is the guest OS audit trail that records user identities, timestamps, object names/paths, and access types. This aligns with the principle that auditing should be performed at the layer that mediates the access decision and can attribute actions to a subject identity within that security context.

References: [Microsoft Learn: Audit object access](#), [Red Hat: Auditing the system \(auditd\)](#)

QUESTION NO: 60

A large international organization that collects information from its consumers has contracted with a Software as a Service (SaaS) cloud provider to process this data. The SaaS cloud provider uses additional data processing to demonstrate other capabilities it wishes to offer to the data owner. This vendor believes additional data processing activity is allowed since they are not disclosing to other organizations. Which of the following BEST supports this rationale?

- A. The data was encrypted at all times and only a few cloud provider employees had access.
- B. As the data owner, the cloud provider has the authority to direct how the data will be processed.
- C. As the data processor, the cloud provider has the authority to direct how the data will be processed.
- D. The agreement between the two parties is vague and does not detail how the data can be used.

ANSWER: D

Explanation:

The best support for the vendor's rationale is that the agreement between the two parties is vague and does not detail how the data can be used. In practice, a cloud provider may try to justify "extra" internal processing (beyond the explicit service

purpose) by pointing to ambiguous contract language, missing purpose limitation clauses, or the absence of explicit prohibitions on secondary use. While this is a weak position from a privacy and compliance standpoint, it is the only choice that directly aligns with the vendor's stated reasoning: they believe it is allowed because nothing clearly forbids it and they are not sharing the data externally.

From an ISC2/CISSP perspective, this scenario highlights the importance of clear contractual controls for cloud services: defining roles (controller/processor), permitted processing purposes, restrictions on secondary use, and audit/monitoring rights. Many privacy frameworks (including GDPR concepts) require purpose limitation and that processors act only on documented instructions from the controller; therefore, contracts should explicitly constrain processing activities to avoid "implied permission" arguments based on vagueness. See [GDPR Article 28 \(Processor\)](#) and NIST guidance on cloud governance and contracts in [NIST SP 800-146](#).

QUESTION NO: 61

A system developer has a requirement for an application to check for a secure digital signature before the application is accessed on a user's laptop. Which security mechanism addresses this requirement?

- A. Trusted Platform Module (TPM).
- B. Certificate revocation list (CRL) policy.
- C. Key exchange.
- D. Hardware encryption.

ANSWER: A

Explanation:

Trusted Platform Module (TPM) is the security mechanism that best addresses the requirement to verify a secure digital signature before an application is accessed on a user's laptop. A TPM is a dedicated hardware root of trust that can securely store cryptographic keys and perform cryptographic operations such as signing and verification. In endpoint and laptop scenarios, TPM-backed measurements and attestation can be used to validate the integrity and authenticity of software components (for example, verifying that code or boot/application-related artifacts are signed by a trusted issuer) before allowing access or execution. This aligns with the requirement's intent: ensuring the application is trusted (authentic and untampered) by checking a digital signature in a way that is resistant to software-only attacks. TPMs are commonly used to support secure/measured boot, device health attestation, and protection of signing/verification keys so that signature checks can be enforced reliably on the client device. For background on TPM capabilities and how it provides hardware-based trust and cryptographic functions, see [Microsoft's TPM overview](#) and the [Trusted Computing Group TPM summary](#).

QUESTION NO: 62 - (DRAG DROP)

DRAG DROP -

Match the following generic software testing methods with their major focus and objective. Drag each testing method next to its corresponding set of testing objectives.

Testing Method

Nonfunctional testing



Functional testing



Structural testing



Regression testing



Testing Objectives

Tests functionality related to changes in software or the environment

Tests suitability, accuracy, interoperability, and security characteristics

Tests control flow, call hierarchies, menu, component, and integration characteristics

Tests reliability, usability, efficiency, and compatibility characteristics

ANSWER:

Testing Method

Nonfunctional testing

Regression testing

Functional testing

Functional testing

Structural testing

Structural testing

Regression testing

Nonfunctional testing

Testing Objectives

Tests functionality related to changes in software or the environment

Tests suitability, accuracy, interoperability, and security characteristics

Tests control flow, call hierarchies, menu, component, and integration characteristics

Tests reliability, usability, efficiency, and compatibility characteristics

Explanation:

This drag-and-drop is asking you to connect common testing method categories to the kind of assurance each one is meant to provide. The key is to remember that each method name hints at what it examines: behavior, internal structure, quality attributes, or the impact of change.

Regression testing is specifically about change impact. After a patch, new feature, configuration update, or platform/environment change, regression testing re-runs relevant tests to confirm previously working functionality still works. That's why it aligns to the objective that mentions "changes in software or the environment." This is a core practice in secure SDLC because fixes can unintentionally reintroduce defects or weaken controls.

Structural testing focuses on the internal construction of the software rather than just external behavior. It's commonly associated with white-box techniques that examine control flow, call hierarchies, and how modules/components integrate. That directly matches the objective referencing control flow, call hierarchies, menus/components, and integration characteristics.

Nonfunctional testing targets quality attributes (often called “-ilities”) such as reliability, usability, performance/efficiency, and compatibility. These are not about whether a specific feature works, but how well the system operates under expected conditions and constraints. Therefore, it maps to the objective listing reliability, usability, efficiency, and compatibility characteristics.

Functional testing validates that the system’s functions meet requirements—what the system does and whether it does it correctly. In many security and assurance contexts, functional requirements can include correct behavior for interoperability and security-related functions (for example, authentication flows, authorization decisions, and correct handling of security-relevant inputs). That’s why functional testing aligns to the objective mentioning suitability, accuracy, interoperability, and security characteristics.

References: [NIST CSRC Glossary – Regression Testing](#), [ISO/IEC 25010 \(software product quality model\)](#), [ISTQB Foundation Level \(testing concepts and categories\)](#).

QUESTION NO: 63

What BEST describes the confidentiality, integrity, availability triad?

- A. A vulnerability assessment to see how well the organization's data is protected.
- B. The three-step approach to determine the risk level of an organization.
- C. The implementation of security systems to protect the organization's data.
- D. A tool used to assist in understanding how to protect the organization's data.

ANSWER: D

Explanation:

The confidentiality, integrity, availability (CIA) triad is best described as a foundational model used to guide information security objectives and decision-making. It helps security professionals and organizations think systematically about what it means to “secure” information: confidentiality focuses on preventing unauthorized disclosure of information; integrity focuses on preventing unauthorized or improper modification and ensuring accuracy and trustworthiness; availability focuses on ensuring timely and reliable access to information and systems for authorized users. In practice, the CIA triad is used as a conceptual tool to evaluate requirements, design controls, and balance trade-offs (for example, stronger confidentiality controls like encryption and strict access control can affect availability or usability). Because it is a model for understanding and framing protection goals rather than a specific assessment method, risk calculation process, or implementation activity, the best description is the one that characterizes it as a tool to assist in understanding how to protect an organization’s data.

References: [NIST CSRC Glossary – Confidentiality](#), [NIST CSRC Glossary – Integrity](#)

QUESTION NO: 64

A project manager for a large software firm has acquired a government contract that generates large amounts of Controlled Unclassified

Information (CUI). The organization's information security manager had received a request to transfer project-related CUI between systems of differing security classifications. What role provides the authoritative guidance for this transfer?

- A. PM
- B. Information owner
- C. Data Custodian
- D. Mission/Business Owner

ANSWER: B

Explanation:

The role that provides authoritative guidance for transferring Controlled Unclassified Information (CUI) between systems of differing security classifications is the information owner. In CISSP terms, the information owner (often aligned with the data owner) is the party with formal accountability for the data and the authority to define how it may be used, shared, labeled, and protected. That includes approving or directing cross-boundary transfers and specifying required controls (for example, approved methods, encryption requirements, and any need for sanitization, downgrading, or release approvals) when information moves between environments with different protection levels. By contrast, operational roles typically implement and maintain controls but do not set the binding rules for data movement. This aligns with common governance models where ownership establishes policy and handling requirements, while custodians execute those requirements. For CUI specifically, organizations must follow defined handling and dissemination rules, and the accountable owner is the appropriate authority to provide direction on whether and how such transfers are permitted within the organization's governance framework.

References: [NIST SP 800-171 Rev. 3](#), [NARA Controlled Unclassified Information \(CUI\)](#).

QUESTION NO: 65

When designing a Cyber-Physical System (CPS), which of the following should be a security practitioner's first consideration?

- A. Detection of sophisticated attackers
- B. Topology of the network used for the system
- C. Risk assessment of the system.
- D. Resiliency of the system.

ANSWER: C**Explanation:**

Risk assessment of the system is the first consideration because CPS security design must start by understanding what can go wrong, how likely it is, and what the impact would be to safety, mission, operations, and the physical environment. In CISSP terms, this aligns with the foundational principle that security controls and architecture are selected based on risk, not on technology preferences or isolated goals like "detect sophisticated attackers." For CPS/ICS environments, risk assessment also explicitly incorporates safety and operational constraints (availability and integrity often dominate), identifies critical assets and trust boundaries, and drives requirements such as segmentation, fail-safe behavior, monitoring, and recovery objectives. Only after risk is characterized can you rationally prioritize resiliency targets, choose an appropriate network topology, and determine detection capabilities that are proportionate and feasible for the system's real-time and safety requirements. This approach is consistent with widely used security frameworks that begin with risk and requirements definition before control selection and implementation. See NIST's guidance on risk management (<https://csrc.nist.gov/projects/risk-management>) and NIST SP 800-82 guidance for ICS security programs that are risk-informed (<https://csrc.nist.gov/publications/detail/sp/800-82/rev-2/final>).

QUESTION NO: 66

A large law firm would like to enable employees to participate in a bring your own device (BYOD) program. Only devices with up-to-date antivirus and operating system (OS) patches will be allowed on the network. Which solution will BEST enforce the security requirements?

- A. Endpoint Detection and Response.
- B. Next-Generation Firewall.
- C. Intrusion detection and prevention system (IDPS).

D. Network Access Control (NAC).

ANSWER: D

Explanation:

Network Access Control (NAC) is the best fit because it is specifically designed to enforce device posture requirements at the point of network admission. In a BYOD scenario, the organization needs a control that can evaluate whether a connecting endpoint meets defined compliance checks—such as current antivirus status and up-to-date OS patch levels—before granting access. NAC solutions commonly integrate with authentication (for example, 802.1X), endpoint posture assessment agents or agentless checks, and policy engines to either allow, deny, or quarantine devices into a remediation network until they meet requirements. This directly matches the requirement of “only devices with up-to-date antivirus and OS patches will be allowed on the network,” because enforcement occurs prior to or at the moment of access, not merely through detection after the fact. NAC also supports differentiated access (role-based or risk-based), which is valuable for a law firm handling sensitive client data and needing strong control over unmanaged personal devices. For additional background on NAC concepts and posture-based access control, see [Cisco Identity Services Engine \(NAC\) overview](#) and an overview of 802.1X network access control at [Microsoft NPS and 802.1X](#).

QUESTION NO: 67

The acquisition of personal data being obtained by a lawful and fair means is an example of what principle?

- A. Collection Limitation Principle.
- B. Openness Principle.
- C. Purpose Specification Principle.
- D. Data Quality Principle.

ANSWER: A

Explanation:

The correct principle is the Collection Limitation Principle. In widely cited privacy frameworks such as the OECD Privacy Guidelines (which strongly influence modern privacy laws and CISSP privacy concepts), collection limitation requires that there be limits to the collection of personal data and that any such data should be obtained by lawful and fair means. It also commonly includes the idea that, where appropriate, collection should occur with the knowledge or consent of the data subject. The key phrase in the question—personal data being acquired via lawful and fair means—maps directly to this principle because it focuses on how data is collected (the acquisition step), not on later stages like disclosure, transparency notices, or ensuring accuracy. In CISSP terms, this is a foundational privacy principle governing the legitimacy and fairness of data acquisition and minimizing unnecessary collection, aligning with privacy-by-design expectations and regulatory requirements for lawful processing/collection. For the canonical wording, see the OECD privacy principles and their description of collection limitation.

References: [OECD Privacy Framework \(Collection Limitation Principle\)](#), [European Commission: Data protection overview](#)

QUESTION NO: 68

Which of the following protection is provided when using a Virtual Private Network (VPN) with Authentication Header (AH)?

- A. Sender non-repudiation.
- B. Multi-factor authentication (MFA).
- C. Payload encryption.

D. Sender confidentiality.

E. Data origin authentication and integrity (with anti-replay protection).

ANSWER: E

Explanation:

Authentication Header (AH) is an IPsec protocol that provides connectionless integrity, data origin authentication, and anti-replay protection for IP packets. In practical terms, AH uses a keyed hash (HMAC) over selected parts of the IP header and the payload to ensure the packet was not altered in transit and that it was generated by a party that possesses the shared secret key. This is the key protection AH adds to a VPN: it authenticates the sender at the packet level and detects tampering, while also helping prevent replay attacks via sequence numbers.

AH does not provide encryption (confidentiality) of the payload; that function is provided by IPsec Encapsulating Security Payload (ESP). AH also does not inherently provide multi-factor authentication; MFA is an access control/authentication mechanism typically implemented at the VPN gateway or identity provider, not by AH itself. Finally, “non-repudiation” is generally not a property delivered by AH because AH commonly relies on symmetric keys; non-repudiation typically requires asymmetric cryptography and digital signatures to prevent a sender from plausibly denying having sent the data.

References: [RFC 4302: IP Authentication Header](#), [RFC 4301: Security Architecture for the Internet Protocol](#)

QUESTION NO: 69

Which of the following BEST ensures the integrity of transactions to intended recipients?

- A. Public key infrastructure (PKI).
- B. Blockchain technology.
- C. Pre-shared key (PSK).
- D. Web of trust.

ANSWER: A

Explanation:

Public key infrastructure (PKI) best ensures the integrity of transactions to intended recipients because it enables strong, scalable use of digital signatures bound to validated identities via certificates. In practice, a sender signs a transaction with their private key, and the recipient (or any verifier) uses the sender’s public key from a trusted certificate to verify that the message has not been altered (integrity) and that it was produced by the holder of the corresponding private key (authentication and nonrepudiation). The “intended recipient” aspect is supported by PKI’s trust model: certificate chains (root/intermediate CAs) allow recipients to validate that the public key they are using truly belongs to the claimed entity, reducing the risk of man-in-the-middle substitution of keys. This is the common foundation for secure email (S/MIME), TLS, code signing, and many enterprise transaction systems where integrity and identity binding are required at scale. While other mechanisms can provide integrity in narrower contexts, PKI is the broad, standard approach for ensuring transaction integrity with verifiable identity binding across untrusted networks.

References: [NIST SP 800-63-3 \(Digital Identity Guidelines\)](#), [RFC 5280 \(X.509 PKI Certificate and CRL Profile\)](#)

QUESTION NO: 70

What is the benefit of an operating system (OS) feature that is designed to prevent an application from executing code from a non-executable memory region?

- A. Identifies which security patches still need to be installed on the system.
- B. Reduces the risk of polymorphic viruses from encrypting their payload.
- C. Stops memory resident viruses from propagating their payload.
- D. Helps prevent certain exploits that store code in buffers.

ANSWER: D

Explanation:

The described OS feature is commonly known as Data Execution Prevention (DEP) or the NX (No-eXecute) bit. Its benefit is that it marks certain memory regions—typically those used for data such as the stack and heap—as non-executable, so the CPU will block attempts to run instructions from those areas. This directly mitigates a large class of memory-corruption attacks (for example, classic stack-based buffer overflows) where an attacker injects shellcode into a buffer and then diverts control flow to execute that injected code. By preventing execution from data pages, DEP/NX forces attackers to use more complex techniques (such as return-oriented programming) rather than straightforward “inject-and-execute” payloads. In CISSP terms, this is a preventive technical control that reduces exploitability of vulnerabilities related to improper memory handling. It does not identify missing patches, and it is not primarily an anti-virus mechanism; instead, it is an OS/hardware-enforced memory protection that helps prevent certain exploits that store code in buffers. See Microsoft's overview of DEP at <https://learn.microsoft.com/en-us/windows/win32/memory/data-execution-prevention> and a general description of the NX bit at https://en.wikipedia.org/wiki/NX_bit.

QUESTION NO: 71

An organization has approved deployment of a virtual environment for the development servers and has established controls for restricting access to resources. In order to implement best security practices for the virtual environment, the security team MUST also implement which of the following steps?

- A. Implement a dedicated management network for the hypervisor.
- B. Deploy Terminal Access Controller Access Control System Plus (TACACS+) for authentication.
- C. Implement complex passwords using Privileged Access Management (PAM).
- D. Capture network traffic for the network interface.

ANSWER: A

Explanation:

Implementing a dedicated management network for the hypervisor is a core virtualization security best practice because it isolates the most sensitive control plane traffic (hypervisor/host management, APIs, consoles, live migration controls, and administrative services) from general user and workload networks. This reduces the attack surface by preventing ordinary VM or developer network segments from directly reaching management interfaces, limiting opportunities for credential theft, management-plane scanning, and lateral movement into the hypervisor layer. It also supports stronger monitoring and tighter firewalling/ACLs around a small, well-defined set of administrative endpoints, and enables more reliable enforcement of least privilege and segmentation (often paired with jump hosts/bastions and MFA). In CISSP terms, this is a practical application of network segmentation and separation of duties applied to virtual infrastructure, helping protect the hypervisor as a high-value target and reducing the blast radius of a compromised VM. This control is broadly recommended across virtualization platforms and aligns with general guidance to separate management, storage, and VM traffic where feasible. See VMware guidance on isolating management traffic: [VMware vSphere Security](#) and Microsoft guidance on securing virtualization hosts and management: [Microsoft Learn: Secure Hyper-V](#).

QUESTION NO: 72

What is the PRIMARY objective of an application security assessment?

- A. Obtain information security management approval.
- B. Maintain the integrity of the application.
- C. Obtain feedback before implementation.
- D. Identify vulnerabilities.

ANSWER: D

Explanation:

The primary objective of an application security assessment is to identify vulnerabilities in the application and its supporting components so they can be prioritized and remediated before they are exploited. In CISSP terms, an application security assessment (which may include activities like secure code review, static/dynamic analysis, configuration review, and penetration testing) is fundamentally a risk-reduction activity: it discovers weaknesses such as injection flaws, broken access control, insecure deserialization, misconfigurations, and other defects that could compromise confidentiality, integrity, or availability. While outcomes like improved integrity, governance approvals, or implementation feedback can be secondary benefits, the core purpose of the assessment itself is vulnerability discovery and characterization (often including severity, exploitability, and business impact) to drive corrective action and continuous improvement in the SDLC. This aligns with widely accepted industry guidance that frames application security testing as a means to find and fix security weaknesses early and repeatedly across development and deployment. See OWASP's overview of application security testing and common vulnerability classes for context: [OWASP Web Security Testing Guide](#) and [OWASP Top 10](#).

QUESTION NO: 73

To minimize the vulnerabilities of a web-based application, which of the following FIRST actions will lock down the system and minimize the risk of an attack?

- A. Apply the latest vendor patches and updates.
- B. Run a vulnerability scanner.
- C. Review access controls.
- D. Install an antivirus on the server.

ANSWER: A

Explanation:

"Apply the latest vendor patches and updates." is the best first action because patching directly reduces the system's exposed attack surface by removing known, exploitable weaknesses in the operating system, web server, application frameworks, and supporting libraries. In CISSP terms, this is a foundational hardening step: before you measure (scanning) or fine-tune controls (access control review), you should eliminate widely known vulnerabilities that attackers routinely target with automated tooling and public exploit code. Patching is also one of the most time-effective risk-reduction activities because it addresses entire classes of vulnerabilities at once (for example, remote code execution flaws in web servers or critical library CVEs). Once current patches are applied, subsequent actions like vulnerability scanning become more meaningful (fewer "known fixed" findings) and access control reviews can focus on configuration and authorization design rather than urgent, known defects. This aligns with industry best practice guidance that emphasizes timely remediation/patch management as a primary control for reducing exploitability and improving security posture.

References: [NIST SP 800-40 Rev. 4 \(Guide to Enterprise Patch Management Planning\)](#), [OWASP Top 10 \(web application risk context\)](#)

QUESTION NO: 74

Which of the following statements pertaining to disaster recovery planning is incorrect?

- A. Every organization must have a disaster recovery plan.
- B. A disaster recovery plan contains actions to be taken before, during and after a disruptive event.
- C. The major goal of disaster recovery planning is to provide an organized way to make decisions if a disruptive event occurs.
- D. A disaster recovery plan should cover return from alternate facilities to primary facilities.

ANSWER: C

Explanation:

The incorrect statement is “The major goal of disaster recovery planning is to provide an organized way to make decisions if a disruptive event occurs.” In ISC2 terms, disaster recovery planning is primarily about restoring IT services and capabilities after a disruption within business-defined objectives (such as RTO and RPO). While clear decision-making and governance are important elements of any response, the central purpose of disaster recovery is operational recovery: re-establishing systems, applications, data, and supporting infrastructure so the organization can resume required levels of service. A plan that only emphasizes decision organization without focusing on recovery strategies, restoration steps, roles, communications, and validation/testing would not meet the intent of DR planning. Effective DR planning also includes preparation activities (e.g., backups, alternate processing, documentation), execution activities during the event (e.g., failover, restoration), and post-event activities (e.g., reconstitution and lessons learned), all aligned to business continuity needs and risk appetite.

References: [NIST SP 800-34 \(Contingency Planning Guide for Federal Information Systems\)](#), [Ready.gov Business Continuity/Continuity Planning](#)

QUESTION NO: 75

One of Canada’s leading pharmaceutical firms recently hired a Chief Data Officer (CDO) to oversee its data privacy program. The CDO has discovered the firm’s marketing department has been collecting information from individuals without their knowledge and consent via the company website. Which of the following privacy regulations should concern the CDO regarding this practice?

- A. The Health Insurance Portability and Accountability Act (HIPAA).
- B. The Privacy Act of 1974.
- C. The Fair Information Practice Principles (FIPPs).
- D. The Personal Information Protection and Electronic Documents Act (PIPEDA).

ANSWER: D

Explanation:

The Personal Information Protection and Electronic Documents Act (PIPEDA) should be the primary concern because it is Canada’s federal private-sector privacy law governing how organizations collect, use, and disclose personal information in the course of commercial activities. A core requirement under PIPEDA is meaningful consent: individuals must be informed about what personal information is being collected and for what purposes, and consent must be obtained (with limited exceptions). Collecting information via a company website “without their knowledge and consent” directly conflicts with these principles, especially in a marketing context where transparency, purpose specification, and consent are central expectations. In practice, this means the organization should provide clear privacy notices, identify purposes at or before collection, and implement consent mechanisms appropriate to the sensitivity of the data and the reasonable expectations of the individual. For a pharmaceutical firm, even marketing-related data collection can become sensitive depending on context

(for example, if it relates to health interests), increasing the need for explicit, well-documented consent and strong governance controls.

References: [PIPEDA \(Justice Laws Website\)](#); [Office of the Privacy Commissioner of Canada – PIPEDA overview](#).

QUESTION NO: 76

Which of the following MUST be scalable to address security concerns raised by the integration of third-party identity services?

- A. Mandatory Access Controls (MAC).
- B. Enterprise security architecture.
- C. Enterprise security procedures.
- D. Role Based Access Controls (RBAC).

ANSWER: B

Explanation:

Enterprise security architecture is the element that must be scalable when integrating third-party identity services (for example, federated identity/SSO using SAML or OpenID Connect, external IdPs, and cloud directory integrations). As organizations add more applications, partners, and identity providers, the security model must scale in a consistent, governed way: trust relationships, authentication and authorization flows, token handling, session management, identity lifecycle processes (provisioning/deprovisioning), logging/monitoring, and resilience requirements all need to be designed so they can expand without creating gaps or one-off exceptions. A scalable architecture provides standardized patterns and controls (e.g., centralized policy decision points, consistent MFA requirements, strong assurance levels, and well-defined boundaries for third-party trust) that can be applied repeatedly as integrations grow. This aligns with CISSP domain expectations around security architecture and engineering: architecture sets the blueprint that enables secure, repeatable, and manageable implementations as complexity increases. For background on federated identity concepts and how third-party identity integration works at scale, see [OpenID Foundation: How OpenID Connect works](#) and [OASIS SAML standard overview](#).

QUESTION NO: 77

A systems engineer is designing a wide area network (WAN) environment for a new organization. The WAN will connect sites holding information at various levels of sensitivity, from publicly available to highly confidential. The organization requires a high degree of interconnectedness to support existing business processes.

What is the BEST design approach to securing this environment?

- A. Use reverse proxies to create a secondary "shadow" environment for critical systems.
- B. Place firewalls around critical devices, isolating them from the rest of the environment.
- C. Layer multiple detective and preventative technologies at the environment perimeter.
- D. Align risk across all interconnected elements to ensure critical threats are detected and handled.

ANSWER: C

Explanation:

Layer multiple detective and preventative technologies at the environment perimeter is the best approach because it reflects a defense-in-depth strategy applied at key control points in a highly interconnected WAN. When business requirements

demand broad connectivity between sites and data of mixed sensitivity, you typically cannot rely on strict isolation alone without breaking workflows. Instead, you design a layered security architecture that combines preventative controls (for example, stateful filtering, next-generation firewall policy, strong authentication, encryption/VPN, and segmentation enforcement) with detective controls (for example, IDS/IPS, centralized logging/SIEM, and continuous monitoring) so that failures or bypasses of one control are compensated by others. In CISSP terms, this is a security architecture approach that reduces the likelihood and impact of compromise while preserving required connectivity, and it supports consistent policy enforcement and visibility across the WAN. Perimeter layering is also a practical place to concentrate controls in a WAN because it creates manageable choke points for inspection, access control, and monitoring between trusted and less-trusted networks and between sites.

References: [NIST Glossary – Defense in Depth](#), [NIST SP 800-41r1 – Guidelines on Firewalls and Firewall Policy](#)

QUESTION NO: 78

Which software defined networking (SDN) architectural component is responsible for translating network requirements?

- A. SDN Controller
- B. SDN Datapath
- C. SDN Northbound Interfaces
- D. SDN Application

ANSWER: C

Explanation:

SDN Northbound Interfaces is the SDN architectural component responsible for translating network requirements because it provides the API boundary where applications express intent (business or application requirements) to the SDN control plane in a consumable form. In SDN, applications define desired outcomes such as segmentation, QoS, or access policy; those high-level requirements must be communicated to the controller using northbound APIs. The northbound interface is therefore the “translation” point between application intent and controller-consumable requests, enabling the controller to compute and enforce the necessary network behavior via its southbound mechanisms. This aligns with the common SDN architecture model (application plane → northbound APIs → control plane), where northbound APIs abstract the underlying network details and allow requirements to be expressed programmatically and consistently. In practice, these interfaces are often RESTful APIs or intent-based interfaces that map application needs into controller policies and service definitions, which the controller then turns into device-level instructions.

References: [Open Networking Foundation \(ONF\) SDN definition](#), [RFC 7426: Software-Defined Networking \(SDN\): Layers and Architecture Terminology](#)

QUESTION NO: 79

What is Kerberos?

- A. A three-headed dog from the Egyptian mythology.
- B. A trusted third-party authentication protocol.
- C. A security model.
- D. A remote authentication dial-in user server.

ANSWER: B

Explanation:

Kerberos is a trusted third-party authentication protocol designed to provide strong, centralized authentication over an untrusted network. It relies on a Key Distribution Center (KDC), which acts as the trusted third party and is typically

composed of an Authentication Server (AS) and a Ticket Granting Server (TGS). Rather than sending passwords across the network, Kerberos uses symmetric-key cryptography and “tickets” to prove identity. A user first authenticates to the KDC to obtain a Ticket Granting Ticket (TGT). The user then presents the TGT to request service tickets for specific resources (such as file servers or applications). Those service tickets are presented to the target service to gain access, enabling single sign-on (SSO) while reducing exposure of long-term secrets. Kerberos also supports mutual authentication, meaning the client can verify the server’s identity and the server can verify the client’s identity, which helps mitigate certain spoofing and replay risks when properly configured (including time synchronization and appropriate ticket lifetimes). This aligns with CISSP expectations: Kerberos is an authentication protocol using a trusted third party, not a general security model and not a dial-in remote authentication server like RADIUS.

References: [RFC 4120: The Kerberos Network Authentication Service \(V5\)](#), [Microsoft: Kerberos authentication overview](#)

QUESTION NO: 80

While performing a security review for a new product, an information security professional discovers that the organization's product development

team is proposing to collect government-issued identification (ID) numbers from customers to use as unique customer identifiers. Which of the following recommendations should be made to the product development team?

- A. Customer identifiers should be a variant of the user's government-issued ID number.
- B. Customer identifiers should be a cryptographic hash of the user's government-issued ID number.
- C. Customer identifiers that do not resemble the user's government-issued ID number should be used.
- D. Customer identifiers should be a variant of the user's name, for example, "jdoe" or "john.doe."

ANSWER: C

Explanation:

Customer identifiers that do not resemble the user's government-issued ID number should be used. Government-issued ID numbers (for example, SSNs, national IDs, driver's license numbers) are highly sensitive personal data and are frequently used for identity verification and fraud. Using them as internal identifiers increases privacy risk, expands breach impact, and can violate data minimization expectations. Even if the identifier is transformed (such as hashing), it can still be linkable and vulnerable to guessing or brute-force attacks because ID numbers often come from constrained formats and limited ranges; this can enable re-identification. CISSP best practice aligns with privacy-by-design and data minimization: collect only what is necessary for the stated purpose, and avoid using sensitive attributes as primary keys or identifiers when a random, non-derivable surrogate identifier will meet the business need. A better design is to generate a unique customer ID (for example, a random UUID) that has no semantic meaning and cannot be used to infer or reconstruct the underlying government ID. This reduces exposure, simplifies compliance, and limits harm if the identifier is leaked.

References: [NIST SP 800-122 \(Protecting the Confidentiality of PII\)](#); [ISO/IEC 27701 \(privacy information management\)](#)

QUESTION NO: 81

Which of the following is NOT a two-factor authentication mechanism?

- A. Something you have and something you know.
- B. Something you do and a password.
- C. A smartcard and something you are.
- D. Something you know and a password.

ANSWER: D

Explanation:

“Something you know and a password.” is not a two-factor authentication mechanism because it uses two credentials from the same authentication factor category: knowledge. In CISSP terms, multi-factor authentication requires combining at least two different factor types, typically drawn from knowledge (something you know, like a PIN/password), possession (something you have, like a smartcard or token), and inherence (something you are, like a fingerprint). Using two knowledge-based secrets (for example, a password plus another password or PIN) may increase complexity, but it remains single-factor authentication because an attacker who compromises the user’s knowledge (through phishing, credential stuffing, keylogging, or reuse) can potentially obtain both. Two-factor authentication specifically aims to reduce this risk by requiring an additional, independent factor type—most commonly a possession factor (hardware token, authenticator app, smartcard) or an inherence factor (biometric). This distinction is central to access control design and is consistent with industry guidance on MFA factor categories and independence. For additional reference on factor types and MFA concepts, see [NIST SP 800-63B \(Digital Identity Guidelines: Authentication and Lifecycle Management\)](#) and [CISA guidance on implementing phishing-resistant MFA](#).

QUESTION NO: 82

Which of the following are the three MAIN categories of security controls?

- A. Preventative, corrective, detective.
- B. Administrative, technical, physical.
- C. Corrective, detective, recovery.
- D. Confidentiality, integrity, availability.

ANSWER: B

Explanation:

The three main categories of security controls in CISSP (and in common security governance practice) are administrative (also called managerial), technical (also called logical), and physical controls. This classification groups controls by their nature and how they are implemented: administrative controls are policies, standards, procedures, risk management, training, and governance mechanisms that direct people and processes; technical controls are implemented through technology such as authentication, access control lists, encryption, and system hardening; and physical controls protect facilities and hardware through measures like locks, guards, mantraps, and environmental protections. While other groupings like preventive/detective/corrective describe control *function*, the question asks for the three main *categories*, which in ISC2-aligned terminology refers to administrative/technical/physical. This taxonomy is widely used in security frameworks and maps cleanly to how organizations design layered defenses across people, process, and technology.

References: [NIST Glossary – Security Control](#), [NIST SP 800-53 Rev. 5](#)

QUESTION NO: 83

What are the three key benefits that application developers should derive from the northbound application programming interface (API) of

software defined networking (SDN)?

- A. Network syntax, abstraction of network flow, and abstraction of network protocols
- B. Network syntax, abstraction of network commands, and abstraction of network protocols
- C. Familiar syntax, abstraction of network topology, and definition of network protocols
- D. Familiar syntax, abstraction of network topology, and abstraction of network protocols

ANSWER: D

Explanation:

The northbound API in SDN is intended to make it easier for application developers to request network behavior without needing to deal with device-specific details. The key benefits are that developers can use a familiar, developer-friendly interface (often RESTful APIs and common data models), work with an abstracted view of the network topology (so they can express intent like connectivity, paths, and policies without manually handling per-switch configuration), and rely on abstraction of network protocols (so the controller translates application intent into underlying southbound mechanisms such as OpenFlow, NETCONF/YANG, gNMI, or vendor-specific protocols). This separation of concerns is central to SDN: applications talk “up” to the controller using higher-level constructs, and the controller handles the complexity of programming the network consistently. In practice, this improves agility, reduces coupling to specific network hardware, and enables automation/orchestration use cases such as traffic engineering and microsegmentation. See the ONF SDN architecture overview for the northbound/southbound separation and API intent, and OpenDaylight documentation for examples of northbound interfaces and abstractions: <https://opennetworking.org/sdn-definition/> and <https://docs.opendaylight.org/>.

QUESTION NO: 84

Which of the following is MOST important to follow when developing information security controls for an organization?

- A. Use industry standard best practices for security controls in the organization.
- B. Exercise due diligence with regard to all risk management information to tailor appropriate controls.
- C. Review all local and international standards and choose the most stringent based on location.
- D. Perform a risk assessment and choose a standard that addresses existing gaps.

ANSWER: B

Explanation:

The most important principle to follow when developing information security controls is to ensure they are driven by risk management and implemented with due diligence. In CISSP terms, controls should be selected and tailored based on the organization’s risk posture, business objectives, threat environment, and tolerance for risk—not simply copied from “best practices” or chosen because they are the most stringent. Exercising due diligence over risk management information means leadership and security professionals are making informed, reasonable decisions using current risk data (e.g., risk assessments, threat intelligence, asset criticality, and compliance obligations) and then translating that into appropriate administrative, technical, and physical controls. This aligns with the core governance expectation that security is risk-based and business-aligned, and it supports defensible decision-making (including demonstrating that the organization took reasonable steps to protect information). A risk-based approach also enables prioritization and cost-effective control selection, ensuring controls reduce risk to an acceptable level rather than maximizing control strength without regard to impact or feasibility.

References: [NIST SP 800-53 Rev. 5 \(control selection within a risk management context\)](#), [NIST SP 800-37 Rev. 2 \(Risk Management Framework\)](#).

QUESTION NO: 85

Which of the following is the MOST effective way to ensure the endpoint devices used by remote users are compliant with an organization's

approved policies before being allowed on the network?

- A. Network Access Control (NAC)
- B. Privileged Access Management (PAM)
- C. Group Policy Object (GPO)
- D. Mobile Device Management (MDM)

ANSWER: A

Explanation:

Network Access Control (NAC) is the most effective choice because it is specifically designed to enforce device posture and policy compliance *before* granting network access. NAC solutions can evaluate endpoints against defined requirements (for example: OS version/patch level, endpoint protection status, disk encryption, configuration baselines, or presence of prohibited software) at the time a device attempts to connect. Based on the assessment, NAC can allow, deny, or quarantine the device into a remediation network segment until it meets policy. This “pre-admission” control is exactly what the question asks for: ensuring remote endpoints are compliant prior to being allowed onto the network. In modern deployments, NAC commonly integrates with identity providers, EDR/AV tools, vulnerability scanners, and MDM/UEM platforms to make access decisions dynamically and continuously, which aligns well with zero trust principles and remote access realities. This makes NAC the most direct and comprehensive control for enforcing compliance gating at the network edge.

References: [NIST SP 800-46 Rev. 2 \(Telework, Remote Access, and BYOD Security\)](#), [NIST SP 800-207 \(Zero Trust Architecture\)](#)

QUESTION NO: 86

A new internal auditor is tasked with auditing the supply chain. The system owner stated that the last internal auditor was terminated because the

auditor discovered too many deficient controls. The auditor reports this conversation to their manager. Which of the following audit integrity principles BEST applies to this situation?

- A. Demonstrate competence while performing professional duties.
- B. Perform professional duties with honesty, diligence, and responsibility.
- C. Perform professional duties in accordance with company policy.
- D. Be aware of any influences that may be exerted on professional judgement.

ANSWER: D

Explanation:

The best-fitting audit integrity principle is to be aware of any influences that may be exerted on professional judgement. In this scenario, the system owner’s comment about the prior auditor being terminated for finding “too many deficient controls” is a clear attempt to create pressure and bias the new auditor’s conclusions. This is a classic independence/objectivity threat: intimidation and undue influence that could cause the auditor to soften findings, narrow scope, or avoid reporting deficiencies. By escalating the conversation to their manager, the auditor is appropriately recognizing and disclosing a potential impairment to objective judgement so that safeguards can be applied (for example, management support, reassignment, additional oversight, or formal documentation of the attempted influence). This aligns with widely accepted internal auditing ethics and standards emphasizing objectivity and the need to manage threats to impartial judgement during audit work. Maintaining awareness of such influences is essential to preserve the credibility of audit results and the integrity of the audit function, especially when auditing sensitive areas like supply chain controls where stakeholders may have strong incentives to avoid negative findings.

References: [IIA Standards and Code of Ethics](#), [ISC2 Code of Ethics](#)

QUESTION NO: 87

What type of investigation applies when malicious behavior is suspected between two organizations?

- A. Regulatory.
- B. Operational.
- C. Civil.

D. Criminal.

ANSWER: C

Explanation:

“Civil.” is the correct choice because disputes involving suspected malicious actions between two organizations are typically handled as civil matters (for example, breach of contract, tort claims, or other business-to-business disputes). In a CISSP context, when the parties are organizations and the issue is framed as one entity harming another (rather than the state prosecuting a violation of criminal law), the investigation and resulting legal process commonly align with civil litigation. Civil investigations focus on establishing facts to support claims such as damages, injunctions, or other remedies, and they often involve internal investigations, outside counsel, eDiscovery, and preservation of evidence to meet civil procedure requirements. While criminal investigations can occur if laws are broken, the question’s emphasis on “between two organizations” points to a civil dispute framework rather than a government-led prosecution. This aligns with standard incident response and legal considerations where organizations pursue civil remedies against other entities for harmful acts impacting business operations or assets.

References: [Cornell Law School – Civil action](#), [U.S. Department of Justice – Types of cases](#)

QUESTION NO: 88

What is the benefit of using Network Admission Control (NAC)?

- A. NAC only supports Windows operating systems (OS).
- B. NAC supports validation of the endpoint's security posture prior to allowing the session to go into an authorized state.
- C. NAC can require the use of certificates, passwords, or a combination of both before allowing network admission.
- D. Operating system (OS) versions can be validated prior to allowing network access.

ANSWER: B

Explanation:

“NAC supports validation of the endpoint's security posture prior to allowing the session to go into an authorized state.” is the key benefit of Network Admission Control. NAC is designed to enforce policy at the moment a device attempts to connect, commonly by checking posture attributes such as patch level, endpoint protection status, configuration baselines, and other compliance signals before granting access. Based on the result, NAC can permit, deny, or place the device into a restricted/quarantine network where remediation can occur. This is a core control for reducing the risk of unmanaged or noncompliant endpoints gaining access to internal resources, and it supports Zero Trust-style access decisions by making access conditional on device health and policy compliance. While NAC deployments often integrate authentication methods (for example, 802.1X with certificates and/or passwords), authentication is not the primary “benefit” being tested here; the distinguishing value is posture assessment and enforcement prior to authorization. This aligns with common NAC/Network Access Control definitions and implementations used in enterprise networks.

References: [Cisco Identity Services Engine \(NAC posture and policy enforcement\)](#), [Microsoft Network Access Protection overview \(conceptually similar posture-based admission control\)](#)

QUESTION NO: 89

Which of the following goals represents a modern shift in risk management according to National Institute of Standards and Technology (NIST)?

- A. Provide an improved mission accomplishment approach.
- B. Focus on operating environments that are changing, evolving, and full of emerging threats.

- C. Enable management to make well-informed risk-based decisions justifying security expenditure.
- D. Secure information technology (IT) systems that store, mass, or transmit organizational information.

ANSWER: B

Explanation:

The goal that best represents NIST's modern shift in risk management is focusing on operating environments that are changing, evolving, and full of emerging threats. In NIST's view (especially as expressed in the Risk Management Framework and related guidance), risk management is no longer a static, compliance-driven activity performed at fixed points in time. Instead, it is intended to be continuous and responsive to rapidly changing technologies, missions, dependencies (including supply chain), and adversary tactics. This "ongoing" posture emphasizes continuous monitoring, frequent reassessment of risk, and adapting controls and risk responses as the environment evolves. The modern shift is therefore characterized by recognizing dynamic threat landscapes and operational contexts, and integrating risk management into day-to-day organizational decision-making rather than treating it as a one-time authorization event. This aligns with NIST's broader messaging that effective security and privacy risk management must keep pace with change and support resilient operations in the face of emerging threats and vulnerabilities.

References: [NIST SP 800-37 Rev. 2 \(Risk Management Framework\)](#), [NIST SP 800-39 \(Managing Information Security Risk\)](#)

QUESTION NO: 90

An organization recently suffered from a web-application attack that resulted in stolen user session cookie information. The attacker was able to obtain the information when a user's browser executed a script upon visiting a compromised website. What type of attack MOST likely occurred?

- A. SQL injection (SQLi).
- B. Extensible Markup Language (XML) external entities.
- C. Cross-Site Scripting (XSS).
- D. Cross-Site Request Forgery (CSRF).

ANSWER: C

Explanation:

Cross-Site Scripting (XSS) is the most likely attack because the scenario describes a malicious script executing in the user's browser when the user visits a compromised website, followed by theft of session cookie information. That is the hallmark of XSS: attacker-supplied script content runs in the victim's browser within the context of a trusted site or page, enabling access to browser-accessible data such as session identifiers (especially if cookies are not protected with flags like HttpOnly) and allowing exfiltration to an attacker-controlled endpoint. In CISSP terms, this is a client-side injection that abuses the browser's trust model and same-origin context to perform actions or read data as the user. While modern defenses (HttpOnly, SameSite, CSP, output encoding) can reduce impact, the described mechanism—script execution on page visit leading to cookie theft—maps directly to XSS, commonly reflected or stored depending on how the compromised content was introduced. For authoritative background, see OWASP's XSS overview and prevention guidance.

References: [OWASP: Cross Site Scripting \(XSS\)](#), [OWASP XSS Prevention Cheat Sheet](#)

QUESTION NO: 91

A code, as is pertains to cryptography:

- A. is a generic term for encryption.

- B. is specific to substitution ciphers.
- C. deals with linguistic units.
- D. is specific to transposition ciphers.

ANSWER: C

Explanation:

In cryptography, a code deals with linguistic units (semantic meaning) rather than operating directly on the individual characters or bits of a message. Practically, this means codes replace words, phrases, or larger message elements with other symbols, numbers, or groups (often via a codebook). For example, a code might map “attack at dawn” or “rendezvous point” to a short numeric group. This is distinct from a cipher, which transforms plaintext according to an algorithm at the level of characters/bits (such as substitution or transposition at the symbol level, or modern block/stream ciphers). CISSP commonly tests this distinction: codes work on meaning-bearing units and typically require secure distribution and protection of the codebook, while ciphers rely on an algorithm and key. Understanding this difference helps in evaluating historical systems and in recognizing operational risks (like codebook compromise) versus cryptanalytic risks against ciphers. For additional background on the code vs. cipher distinction, see [https://en.wikipedia.org/wiki/Code_\(cryptography\)](https://en.wikipedia.org/wiki/Code_(cryptography)) and <https://en.wikipedia.org/wiki/Cipher>.

QUESTION NO: 92

Which access control model would a lattice-based access control model be an example of?

- A. Mandatory access control.
- B. Discretionary access control.
- C. Non-discretionary access control.
- D. Rule-based access control.

ANSWER: A

Explanation:

Mandatory access control is the correct choice because lattice-based access control (LBAC) is a formal model typically used to implement mandatory policies where access decisions are derived from system-enforced security labels and a defined dominance relationship. In a lattice, subjects and objects are assigned classifications (for example, Top Secret, Secret, Confidential) and possibly compartments/categories, and the system evaluates whether a subject’s label dominates an object’s label before allowing access. This is characteristic of mandatory access control: end users (including data owners) do not get to arbitrarily grant permissions; instead, the operating environment enforces the policy consistently based on labels and the lattice rules. This is the classic foundation for multilevel security systems and aligns with well-known MAC models such as Bell–LaPadula (confidentiality-focused) and Biba (integrity-focused), which can be represented using lattice structures. For additional background on MAC and label-based controls, see [NIST CSRC glossary: Mandatory Access Control](#) and [NIST CSRC glossary: Access Control](#).

QUESTION NO: 93

Which of the following is the MOST challenging issue in apprehending cyber criminals?

- A. They often use sophisticated method to commit a crime.
- B. It is often hard to collect and maintain integrity of digital evidence.
- C. The crime is often committed from a different jurisdiction.

D. There is often no physical evidence involved.

ANSWER: C

Explanation:

The crime is often committed from a different jurisdiction is the most challenging issue because cybercrime investigations frequently cross city, state, and national boundaries, creating major legal and procedural barriers to identification, evidence collection, and arrest. Even when investigators can attribute activity to an IP address, hosting provider, or suspect, taking action typically requires cooperation from multiple organizations and governments. Differences in laws (what is criminalized, required evidentiary standards, privacy rules, data retention requirements), as well as the need for formal processes such as mutual legal assistance treaties (MLATs) or other cross-border requests, can introduce significant delays or prevent access to key evidence entirely. This jurisdictional complexity also affects extradition, prosecution venue, and chain-of-custody handling across agencies, all of which can slow or block apprehension. CISSP-aligned best practice recognizes that while technical sophistication and digital evidence challenges are real, cross-jurisdictional legal authority and international cooperation are often the dominant constraint in actually apprehending offenders.

References: [INTERPOL – Cybercrime](#), [U.S. DOJ – Mutual Legal Assistance Treaties \(MLATs\)](#)

QUESTION NO: 94

Which of the following is the BEST way to determine the success of a patch management process?

- A. Change management
- B. Configuration management (CM)
- C. Analysis and impact assessment
- D. Auditing and assessment

ANSWER: D

Explanation:

Auditing and assessment is the best way to determine whether a patch management process is successful because it provides objective, verifiable evidence that patches are actually deployed, effective, and meeting policy/SLAs. A mature patch program defines measurable requirements (coverage, timeliness, exception handling, rollback, and verification) and then validates them through compliance checks, vulnerability scans, and periodic internal/external audits. This closes the loop between “we planned/approved patches” and “systems are demonstrably at the required patch level,” including confirming that compensating controls or documented exceptions exist where patching cannot occur. In practice, auditing and assessment uses sources such as endpoint configuration/patch status reports, authenticated vulnerability scanning results, and sampling/testing to confirm that the environment’s real state matches the intended state. This aligns with CISSP expectations that effectiveness is proven through monitoring, metrics, and independent verification rather than process activities alone. For example, NIST emphasizes assessing control effectiveness and verifying remediation results as part of continuous monitoring and security assessment activities. See [NIST SP 800-40 Rev. 4 \(Enterprise Patch Management Planning\)](#) and [NIST SP 800-53 Rev. 5 \(Assessment, Audit, and Continuous Monitoring concepts\)](#).

QUESTION NO: 95

Which of the following can BEST eliminate dial-up access through a Remote Access Server as a hacking vector?

- A. Using a TACACS+ server.
- B. Installing the Remote Access Server outside the firewall and forcing legitimate users to authenticate to the firewall.
- C. Setting modem ring count to at least 5.

D. Only attaching modems to non-networked hosts.

ANSWER: D

Explanation:

Only attaching modems to non-networked hosts best eliminates dial-up access through a Remote Access Server as a hacking vector because it removes the pathway from a dial-in connection into the organization's internal network. The core risk with dial-up remote access is that an attacker who discovers and successfully authenticates (or exploits) a dial-in line can land on a system that has routed or bridged connectivity into the enterprise environment, bypassing many perimeter controls. By ensuring that any dial-up-capable system is not connected to the network (no routing, no bridging, no dual-homing into internal segments), the dial-up line cannot be used as an entry point to compromise internal resources. This is a classic "eliminate the attack surface" control: rather than trying to harden dial-up authentication or place the remote access server in a different zone, it prevents dial-up from becoming a network ingress vector at all. This aligns with CISSP guidance to reduce exposure by removing unnecessary services and limiting connectivity paths, especially legacy remote access methods like dial-up. See NIST's general guidance on reducing attack surface and controlling remote access in [NIST SP 800-46 Rev. 2](#) and broader security control concepts in [NIST SP 800-53 Rev. 5](#).

QUESTION NO: 96

Which of the following is BEST suited for exchanging authentication and authorization messages in a multi-party decentralized environment?

- A. Lightweight Directory Access Protocol (LDAP).
- B. Security Assertion Markup Language (SAML).
- C. Internet Mail Access Protocol
- D. Transport Layer Security (TLS).

ANSWER: B

Explanation:

Security Assertion Markup Language (SAML) is best suited for exchanging authentication and authorization information across multiple independent parties because it is a federation standard designed specifically for passing identity assertions between an identity provider and one or more service providers. In a multi-party, decentralized environment (common in B2B, SaaS, and cross-domain single sign-on), SAML enables organizations to keep identity management local while still allowing trusted partners to consume standardized authentication statements and authorization-related attributes. This supports scalable trust relationships, reduces the need for each service to maintain separate credential stores, and enables centralized authentication with distributed consumption of assertions. SAML's core value is its structured, signed assertions (typically XML) that can be validated by relying parties, providing integrity and supporting strong federation workflows such as web browser SSO. This aligns with CISSP expectations around federated identity and cross-domain authentication/authorization message exchange, where SAML is a canonical protocol for assertion-based identity federation.

References: [OASIS SAML 2.0 Core Specification](#), [Cloudflare: What is SAML?](#)

QUESTION NO: 97

Which of the following techniques evaluates the secure design principles of network or software architectures?

- A. Risk modeling.
- B. Waterfall method.
- C. Threat modeling.

D. Fuzzing.

ANSWER: C

Explanation:

Threat modeling is the technique used to evaluate whether a network or software architecture aligns with secure design principles by systematically identifying what can go wrong in the design and how to mitigate it. It is performed early (and iteratively) in the system development life cycle, focusing on architecture, trust boundaries, data flows, assets, and attacker goals. By enumerating threats and mapping them to controls (for example, using STRIDE or similar frameworks), threat modeling helps validate that principles such as least privilege, defense in depth, secure defaults, and separation of duties are actually reflected in the proposed design. This makes it a design-centric activity: it assesses the architecture before (or alongside) implementation and testing, and it produces actionable design changes and security requirements. In CISSP terms, it supports secure architecture and engineering by ensuring the design addresses likely threat scenarios and reduces attack surface in a structured, repeatable way. For additional context, see OWASP's overview of threat modeling and Microsoft's guidance on threat modeling practices.

References: [OWASP Threat Modeling](#), [Microsoft Threat Modeling](#)

QUESTION NO: 98

What is the PRIMARY benefit of relying on Security Content Automation Protocol (SCAP)?

- A. Standardize specifications between software security products.
- B. Achieve organizational compliance with international standards.
- C. Improve vulnerability assessment capabilities.
- D. Save security costs for the organization.

ANSWER: A

Explanation:

The primary benefit of relying on Security Content Automation Protocol (SCAP) is that it standardizes how security configuration, vulnerability, and compliance-related information is expressed and exchanged between tools. SCAP is a suite of open specifications (for example, CVE, CCE, CPE, CVSS, XCCDF, and OVAL) designed to enable automation and interoperability across different vendors' security products. By using common, machine-readable formats and identifiers, organizations can more consistently perform tasks like configuration assessment, vulnerability checking, and reporting across heterogeneous environments, and they can more easily compare results from different scanners and management platforms. This standardization is the foundational value that enables the other downstream benefits (such as better assessments and more efficient reporting), but the core advantage is the shared language and specifications that allow security content to be reused and tools to work together reliably.

References: [NIST CSRC – Security Content Automation Protocol \(SCAP\)](#), [NIST NVD – SCAP](#)

QUESTION NO: 99

Which of the following is BEST practice to employ in order to reduce the risk of collusion?

- A. Least Privilege.
- B. Job Rotation
- C. Separation of Duties
- D. Mandatory Vacations.

ANSWER: C

Explanation:

Separation of Duties is the best practice to reduce the risk of collusion because it is a foundational internal control designed to prevent any single person (or small group) from having end-to-end control over a critical process. By splitting key steps (for example, request, approval, execution, and reconciliation) across different roles, an organization reduces the opportunity for two or more individuals to coordinate actions that bypass oversight. In other words, the control objective is to ensure that completing a sensitive transaction requires independent participation and review, which makes coordinated fraud harder to execute and easier to detect. This is why Separation of Duties is emphasized in governance and audit frameworks and is a core concept in CISSP security operations and risk management: it directly addresses the “two-person rule” problem space and limits the ability to conceal unauthorized activity through shared access or shared responsibilities. While other administrative controls can help detect fraud, Separation of Duties is the primary preventive control specifically aimed at reducing collusion risk in business processes.

References: [NIST CSRC Glossary – Separation of Duties](#), [ISACA COBIT resources \(governance and control concepts\)](#)

QUESTION NO: 100

A security architect is responsible for the protection of a new home banking system. Which of the following solutions can BEST improve the confidentiality and integrity of this external system?

- A. Intrusion Prevention System (IPS).
- B. Denial of Service (DoS) protection solution.
- C. One-time Password (OTP) token.
- D. Web Application Firewall (WAF).

ANSWER: D

Explanation:

Web Application Firewall (WAF) is the best choice to improve confidentiality and integrity for an external (internet-facing) home banking system because it directly protects the application layer where most banking-specific attacks occur. A WAF inspects HTTP/HTTPS traffic and can detect and block common web application threats such as SQL injection, cross-site scripting (XSS), request smuggling, and other OWASP Top 10-style attacks. Preventing these attacks helps preserve confidentiality by reducing the risk of unauthorized data disclosure and helps preserve integrity by preventing unauthorized modification of transactions, session data, or backend database content. In CISSP terms, a WAF is a compensating and preventive technical control placed in front of a web application to enforce security policy at Layer 7, often with positive/negative security models, signatures, and anomaly detection. For a home banking system exposed to untrusted networks, this targeted control typically provides more direct protection of application data and transaction integrity than network-only controls. WAFs are also commonly deployed as reverse proxies and can add protections like request normalization, protocol enforcement, and virtual patching for known vulnerabilities while application fixes are developed and tested.

References: [OWASP Top 10](#), [Cloudflare: What is a WAF?](#)

QUESTION NO: 101

Which of the following features is MOST effective in mitigating against theft of data on a corporate mobile device which has been stolen?

- A. Mobile Device Management (MDM) with device wipe
- B. Mobile device tracking with geolocation
- C. Virtual private network (VPN) with traffic encryption

D. Whole device encryption with key escrow

ANSWER: D

Explanation:

Whole device encryption with key escrow is the most effective mitigation for data theft after a corporate mobile device has already been stolen because it directly protects data at rest on the device. With strong full-disk (or file-based) encryption, an attacker who gains physical possession cannot read stored corporate data without the decryption key (typically derived from a user credential and/or hardware-backed key storage). This remains true even if the attacker removes storage media, boots alternate operating systems, or attempts offline analysis. Key escrow supports enterprise recovery and continuity by allowing the organization to regain access to encrypted data when a user forgets credentials or when lawful/authorized recovery is required, without weakening the fundamental protection provided by encryption. In contrast, controls like remote wipe depend on the device being powered on and able to receive management commands, and geolocation or VPN encryption do not inherently prevent offline access to locally stored data. Full-device encryption is a core best practice for mobile endpoints handling sensitive information and is widely recommended in security guidance for protecting lost or stolen devices.

References: [NIST SP 800-124 Rev. 2 \(Guidelines for Managing the Security of Mobile Devices\)](#), [Apple Platform Security – Data Protection](#)

QUESTION NO: 102

What is the MOST effective method to enhance security of a single sign-on (SSO) solution that interfaces with critical systems?

- A. Two-factor authentication.
- B. Reusable tokens for application level authentication.
- C. High performance encryption algorithms.
- D. Secure Sockets Layer (SSL) for all communications.

ANSWER: A

Explanation:

Two-factor authentication is the most effective enhancement for an SSO solution protecting access to critical systems because SSO centralizes authentication: if an attacker compromises the single set of credentials, they can potentially pivot into many connected applications. Adding a second factor (something you have/are) materially reduces the likelihood that stolen passwords, phishing, credential stuffing, or password reuse will result in account takeover. This is especially important for high-impact targets where the threat model includes sophisticated social engineering and malware. In modern enterprise SSO (for example, SAML/OIDC-based identity providers), enforcing MFA at the identity provider creates a consistent, centralized control point that can be applied conditionally (risk-based/step-up) for privileged actions or sensitive applications, improving security without requiring each downstream application to implement its own strong authentication. This aligns with CISSP best practices around strong authentication, centralized access control, and reducing single points of failure introduced by SSO. For further reading on MFA concepts and why it mitigates password compromise, see [NIST SP 800-63B \(Digital Identity Guidelines\)](#) and an overview of MFA in identity systems at [Microsoft Entra MFA: how it works](#).

QUESTION NO: 103

Which of the following is NOT a valid reason to use external penetration service firms rather than corporate resources?

- A. They are more cost-effective.
- B. They offer a lack of corporate bias.

- C. They use highly talented ex-hackers.
- D. They ensure a more complete reporting.

ANSWER: C

Explanation:

“They use highly talented ex-hackers.” is not a valid reason in the CISSP sense because it relies on an informal, non-governance-based assumption about who performs the work rather than on measurable business and security drivers. ISC2-aligned best practice focuses on objective reasons to outsource penetration testing such as independence (reducing organizational bias and conflicts of interest), access to specialized skills and tooling as a capability (not a particular background), and improved credibility and defensibility of results for stakeholders and auditors. The value proposition is the firm’s demonstrated competence, methodology, and adherence to rules of engagement, legal constraints, and professional standards—not employing “ex-hackers,” which is neither necessary nor inherently indicative of quality, ethics, or reliability. In professional penetration testing, qualifications are typically evidenced through experience, repeatable processes, and recognized certifications/standards, along with clear scoping and reporting requirements. This aligns with the CISSP emphasis on governance, risk management, and due diligence when selecting third-party security service providers.

References: [ISC2 CISSP Certification Overview](#), [OWASP Web Security Testing Guide \(WSTG\)](#)

QUESTION NO: 104

Which of the following should NOT normally be allowed through a firewall?

- A. SNMP
- B. SMTP
- C. HTTP
- D. SSH.

ANSWER: A

Explanation:

SNMP should not normally be allowed through a firewall, especially from untrusted networks into an internal environment. SNMP is a management protocol intended for monitoring and administering network devices (routers, switches, servers, printers). Exposing it beyond a tightly controlled management network increases risk because SNMP can reveal sensitive configuration and inventory data (via “read” access) and, if misconfigured or using weak community strings/credentials, can enable unauthorized changes (via “write” access). Even with SNMPv3’s stronger security features (authentication and encryption), best practice is to restrict SNMP to dedicated management segments, limit it to specific management hosts, and filter it at network boundaries rather than broadly permitting it through perimeter firewalls. In contrast, protocols like HTTP and SMTP are commonly business-required and are often intentionally allowed to specific destinations/ports with additional controls (reverse proxies, mail gateways). SSH may be allowed for tightly scoped administrative access (e.g., from a bastion host/VPN), but it is still typically restricted. The key point is that SNMP is primarily an internal management-plane protocol and is generally not appropriate to expose across security boundaries unless there is a well-justified, tightly controlled management use case.

References: [NIST SP 800-41 Rev. 1 \(Guidelines on Firewalls and Firewall Policy\)](#), [RFC 3411 \(SNMP Framework\)](#)

QUESTION NO: 105

Which RAID level concept is considered more expensive and is applied to servers to create what is commonly known as server fault tolerance?

- A. RAID level 0.

- B. RAID level 1.
- C. RAID level 2.
- D. RAID level 5.

Answer: B

Explanation:

- A. RAID level 1.
- B. RAID level 2.
- C. RAID level 5.

ANSWER: A

Explanation:

RAID level 1 is the RAID concept commonly associated with “server fault tolerance” in many CISSP-style questions because it provides redundancy through disk mirroring. With mirroring, every write is duplicated to a second disk (or set of disks), so if one drive fails, the system can continue operating using the mirror without data loss and typically with minimal interruption. This approach is considered more expensive than parity-based RAID levels because it requires at least 2 drives and effectively sacrifices 50% of raw disk capacity to maintain the duplicate copy. In other words, you pay for additional disks to gain higher availability and simpler recovery behavior. While other RAID levels can also improve availability, RAID 1 is the straightforward “duplicate everything” design that maps well to the idea of fault tolerance at the storage layer for servers.

References: [IBM Docs – RAID levels overview](#), [TechTarget – RAID 1 \(disk mirroring\)](#)

QUESTION NO: 106

How is protection for hypervisor host and software administration functions BEST achieved?

- A. Enforce network controls using a host-based firewall.
- B. Deploy the management interface in a dedicated virtual network segment.
- C. The management traffic pathway should have separate physical network interface cards (NIC) and network.
- D. Deny permissions to specific virtual machines (VM) groups and objects.

ANSWER: C

Explanation:

The management traffic pathway should have separate physical network interface cards (NIC) and network is the best answer because hypervisor administration is a high-value control plane that should be isolated from tenant/guest and general production traffic. Using dedicated physical NICs and a separate management network reduces the attack surface by preventing untrusted workloads or compromised guest networks from reaching the hypervisor’s management interfaces. This also limits opportunities for lateral movement, sniffing, spoofing, and misrouting of administrative traffic, and it supports stronger monitoring and access control (for example, restricting management network access to jump hosts and privileged admin workstations). In virtualization security best practice, management plane isolation is a core design principle: keep management interfaces on a dedicated network, ideally physically separated, and tightly control who can reach it. While logical segmentation can help, physical separation provides stronger assurance against misconfiguration and certain classes of virtual switch or VLAN hopping issues. This aligns with common guidance from virtualization vendors and cloud security references emphasizing separation of management, storage, and VM data networks.

References: [VMware vSphere Security Configuration Guide](#), [Microsoft Hyper-V best practices](#)

QUESTION NO: 107

Which of the following assertions is NOT true about pattern matching and anomaly detection in intrusion detection?

- A. Anomaly detection tends to produce more data.
- B. A pattern matching IDS can only identify known attacks.
- C. Stateful matching scans for attack signatures by analyzing individual packets instead of traffic streams.
- D. An anomaly-based engine develops baselines of normal traffic activity and throughput, and alerts on deviations from these baselines.

ANSWER: C

Explanation:

The statement “Stateful matching scans for attack signatures by analyzing individual packets instead of traffic streams.” is the assertion that is NOT true. In intrusion detection, stateful inspection (often called stateful pattern matching in IDS/IPS contexts) correlates packets across a session or traffic flow, tracking connection state and reassembling or logically associating traffic streams to detect signatures that may be split across multiple packets. This is specifically done to avoid the limitations of looking at isolated packets, where an attacker can evade detection by fragmenting payloads or distributing an exploit sequence across multiple packets. In other words, stateful matching is about understanding context over time (streams/flows), not just per-packet analysis. This aligns with common IDS/IPS design principles discussed in security engineering and network security references: signature-based detection benefits from stream reassembly and state tracking to improve accuracy and reduce evasion opportunities. For additional background on stateful inspection concepts and how state tracking differs from stateless packet filtering, see [Cisco: What is a firewall?](#) and [Stateful firewall \(overview\)](#).

QUESTION NO: 108

Which of the following processes has the PRIMARY purpose of identifying outdated software versions, missing patches, and lapsed system updates?

- A. Penetration testing.
- B. Vulnerability management.
- C. Software Development Life Cycle (SDLC).
- D. Life cycle management.

ANSWER: B

Explanation:

Vulnerability management is the process primarily focused on continuously identifying, assessing, prioritizing, and remediating security weaknesses across systems and applications. A core activity within vulnerability management is vulnerability scanning and configuration/asset assessment, which commonly detects outdated software versions, missing security patches, and systems that have fallen behind on required updates. This is typically done on a recurring schedule (and after significant changes) to maintain an up-to-date view of exposure and to drive patching and other corrective actions. In CISSP terms, this aligns with ongoing operational security practices: maintaining system hygiene through regular identification of known weaknesses and ensuring timely remediation through patch and update management workflows. While other activities may incidentally reveal missing patches, vulnerability management is the discipline whose primary purpose is to find and track these issues over time and ensure they are addressed according to risk and policy. See NIST guidance on vulnerability management and scanning in [NIST SP 800-40 Rev. 4](#) and related vulnerability scanning concepts in [NIST SP 800-115](#).

QUESTION NO: 109

Which mechanism provides the BEST protection against buffer overflow attacks in memory?

- A. Address Space Layout Randomization (ASLR)
- B. Memory management unit
- C. Stack and heap allocation
- D. Dynamic random access memory (DRAM)

ANSWER: A

Explanation:

Address Space Layout Randomization (ASLR) provides the best protection among the listed choices because it makes exploitation of memory corruption bugs (including many buffer overflows) significantly harder by randomizing the locations of key process memory regions (such as the stack, heap, and loaded libraries) each time a program runs. Buffer overflow attacks often rely on predicting or knowing target addresses for return-oriented programming (ROP), jumping to injected shellcode, or reusing existing code in shared libraries. By introducing unpredictability into where code and data reside, ASLR reduces the attacker's ability to reliably redirect execution flow, increasing the likelihood that an exploit will crash rather than succeed. While ASLR is not a complete fix by itself (it is typically most effective when combined with other controls like non-executable memory and stack canaries), it is a widely deployed OS-level mitigation specifically aimed at making memory corruption exploitation less deterministic and therefore less practical at scale. This aligns with CISSP expectations around defense-in-depth and platform protections against common software exploitation techniques.

References: [Microsoft: Address Space Layout Randomization](#), [OWASP: Buffer Overflow Attack](#)

QUESTION NO: 110

Password management falls into which control category?

- A. Compensating.
- B. Detective.
- C. Preventive.
- D. Technical.

ANSWER: C

Explanation:

Password management is best categorized as a preventive control because its primary purpose is to stop unauthorized access before it occurs. In CISSP terms, preventive controls are designed to avoid or block security incidents (for example, by enforcing authentication requirements, password complexity/length rules, rotation policies where applicable, and preventing password reuse). Effective password management reduces the likelihood of successful guessing, brute force, credential stuffing, and misuse of weak or shared credentials, thereby preventing compromise rather than merely detecting it after the fact.

While password mechanisms are often implemented using technical means (such as directory services, PAM modules, or identity providers), the question asks for the control category (by function). Functionally, password management aligns with prevention because it enforces access control at the point of entry. This maps cleanly to ISC2's common control-function taxonomy (preventive, detective, corrective, deterrent, compensating, recovery) where authentication controls are classic examples of preventive measures.

References: [NIST SP 800-53 Rev. 5 \(Access Control / Identification and Authentication control families\)](#), [NIST SP 800-63B \(Digital Identity Guidelines: Authentication and lifecycle management\)](#)

QUESTION NO: 111

A security architect is reviewing an implemented security framework. After the review, the security architect wants to enhance the security by implementing segregation of duties (SoD) to address protection against fraud. Which security model BEST protects the integrity of data?

- A. The Brewer-Nash model.
- B. The Biba Integrity model.
- C. The Bell-LaPadula model.
- D. The Clark-Wilson model.

ANSWER: D

Explanation:

The Clark-Wilson model best protects the integrity of data in scenarios focused on preventing fraud through segregation of duties. Clark-Wilson is an integrity model designed around real-world commercial and financial processing, where the primary goal is to ensure that data can only be modified in authorized, well-formed ways. It accomplishes this by requiring that users do not directly manipulate critical data items; instead, they invoke certified transformation procedures (controlled programs/transactions) that enforce business rules and validation. A key feature is the explicit support for separation of duties: different roles are required to certify procedures, execute them, and audit the results, reducing the likelihood that a single individual can both perpetrate and conceal fraudulent activity. The model also emphasizes auditing and accountability, typically via append-only logs, so that changes are traceable and tampering is detectable. These characteristics align tightly with integrity protection in enterprise environments where SoD is a primary control objective. For additional background on Clark-Wilson and integrity-focused access control concepts, see [NIST CSRC \(Integrity\)](#) and [NIST CSRC \(Access Control\)](#).

QUESTION NO: 112

What is the MOST effective way to mitigate distributed denial of service (DDoS) attacks?

- A. Deploy a web application firewall (WAF).
- B. Block access to Transmission Control Protocol (TCP) ports under attack.
- C. Detect and block bad Internet Protocol (IP) subnets on the corporate firewall.
- D. Engage an upstream Internet service provider (ISP).

ANSWER: D

Explanation:

Engaging an upstream Internet service provider (ISP) is the most effective mitigation for many DDoS attacks because DDoS traffic volumes commonly exceed the capacity of an organization's internet circuit and on-premises security devices. Once a link is saturated, local controls (such as firewalls, ACLs, or application-layer filtering) may never see legitimate traffic, making "in-house only" responses insufficient. An upstream ISP (or a specialized DDoS scrubbing provider) can apply filtering, rate limiting, blackholing/sinkholing, or traffic diversion to scrubbing centers closer to the attack sources and with far greater bandwidth and infrastructure scale. This approach reduces the attack traffic before it reaches the victim's network edge, preserving availability—one of the core security objectives emphasized in CISSP. In practice, this is implemented via pre-arranged DDoS protection services, BGP-based diversion/RTBH, and coordinated incident response runbooks with the provider, often combined with CDN/Anycast for additional absorption. For background on DDoS concepts and common mitigation patterns, see [CISA: Understanding Denial-of-Service Attacks](#) and [Cloudflare: What is a DDoS attack?](#).

QUESTION NO: 113

A security professional needs to find a secure and efficient method of encrypting data on an endpoint. Which solution includes a root key?

- A. Bitlocker
- B. Trusted Platform Module (TPM)
- C. Virtual storage array network (VSAN)
- D. Hardware security module (HSM)

ANSWER: B

Explanation:

Trusted Platform Module (TPM) is the best answer because a TPM is specifically designed to provide a hardware-based root of trust for an endpoint, including the secure generation, storage, and use of cryptographic keys. In TPM terminology, the “root key” concept maps to the Storage Root Key (SRK) (and, in TPM 2.0, the storage hierarchy root), which anchors key protection: other keys can be wrapped (encrypted) under this root so they can only be unwrapped and used by the same TPM in an expected state. This makes endpoint encryption both secure (keys are protected by tamper-resistant hardware and platform measurements) and efficient (cryptographic operations and key handling are offloaded to dedicated hardware). TPM-backed designs are commonly used to protect full-disk encryption key material and to support secure boot/attestation workflows on endpoints. For example, Windows can use a TPM to protect BitLocker’s volume master key, but the component that “includes a root key” in the sense of a hardware root-of-trust key hierarchy is the TPM itself. See [Microsoft TPM overview](#) and [Trusted Computing Group TPM Library Specification](#).

QUESTION NO: 114

Which of the following roles is responsible for ensuring that important datasets are developed, maintained, and are accessible within their defined specifications?

- A. Data Custodian.
- B. Data Reviewer.
- C. Data User.
- D. Data Owner.

ANSWER: A

Explanation:

Data Custodian is the role responsible for the day-to-day operational handling of data to ensure it is properly maintained and made available according to the requirements set by the organization. In CISSP terms, the data owner (often a business role) defines the data’s classification, intended use, and access requirements, while the data custodian (often IT/operations) implements and administers those requirements. That operational responsibility includes maintaining the datasets, ensuring appropriate storage, backups, integrity controls, and making the data accessible to authorized users in line with defined specifications (such as availability targets, retention rules, and handling procedures). This aligns with common governance models where custodians manage the technical environment and controls that keep data usable and available, while ownership remains accountable for the data’s business purpose and protection requirements. For additional background on these governance roles and how custodians implement controls to meet defined requirements, see [NIST glossary \(data steward concept\)](#) and [ISC2 overview of security roles and responsibilities](#).

QUESTION NO: 115

Which of the following is not a property of the Rijndael block cipher algorithm?

- A. It employs a round transformation that is comprised of three layers of distinct and invertible transformations.
- B. It is suited for high speed chips with no area restrictions.
- C. It operates on 64-bit plaintext blocks and uses a 128 bit key.
- D. It could be used on a smart card.

ANSWER: C

Explanation:

The statement “It operates on 64-bit plaintext blocks and uses a 128 bit key.” is not a property of the Rijndael block cipher (and of AES, which is the standardized subset of Rijndael). Rijndael/AES is designed around a 128-bit block size (the AES standard fixes the block size at 128 bits), and it supports key sizes of 128, 192, or 256 bits. A 64-bit block size is characteristic of older block ciphers such as DES/3DES, not Rijndael/AES. This distinction matters in practice because block size affects security properties like birthday-bound collision risks in modes of operation; moving from 64-bit blocks to 128-bit blocks significantly increases the amount of data that can be safely encrypted under a single key in common modes. For authoritative details, see NIST’s AES specification, which defines AES with a 128-bit block and the three approved key sizes, and NIST’s general AES information page for context on the Rijndael selection and standardization.

References: [NIST FIPS 197 \(AES\)](#), [NIST AES Project](#)

QUESTION NO: 116

Which of the following is the MOST appropriate control for asset data labeling procedures?

- A. Categorizing the types of media being used.
- B. Logging data media to provide a physical inventory control.
- C. Reviewing off-site storage access controls.
- D. Reviewing audit trails of logging records.

ANSWER: A

Explanation:

Categorizing the types of media being used is the most appropriate control for asset data labeling procedures because effective labeling depends on first defining what is being labeled and ensuring the labeling scheme is consistently applied across all relevant asset/media types. In practice, organizations establish a data classification and labeling standard (for example: Public, Internal, Confidential, Restricted) and then map that standard to specific media and information containers (paper records, removable media, endpoints, databases, cloud storage objects, backups, etc.). This categorization step enables clear procedures for how labels are created, displayed, and handled, and it supports downstream controls such as handling requirements, storage, transmission protections, and retention/destruction rules. Within CISSP asset security concepts, labeling is tightly coupled to classification and to the identification of assets and their containers; without categorizing media types, labeling procedures become incomplete and inconsistently implemented. This aligns with widely adopted guidance that classification and labeling must be defined and applied to information in all forms. See [ISO/IEC 27001](#) (information classification/handling expectations) and NIST guidance on categorizing information and systems in [FIPS 199](#).

QUESTION NO: 117

As users switch roles within an organization, their accounts are given additional permissions to perform the duties of their new position. After a recent audit, it was discovered that many of these accounts maintained their old permissions as well.

The obsolete permissions identified by the audit have been remediated and accounts have only the appropriate permissions to complete their jobs.

Which of the following is the BEST way to prevent access privilege creep?

- A. Implementing Identity and Access Management (IAM) solution
- B. Time-based review and certification.
- C. Internet audit.
- D. Trigger-based review and certification.

ANSWER: D

Explanation:

Trigger-based review and certification is the best way to prevent access privilege creep because it ties access revalidation to the event that most commonly causes privilege accumulation: a change in a user's role, department, manager, or job function. When access reviews are automatically initiated by these lifecycle events (often called "joiner/mover/leaver" processes), the organization can promptly remove entitlements that are no longer required and ensure the user's access aligns with least privilege for the new position. This is more preventative than periodic, time-based reviews because it reduces the window in which excessive access can persist after a role change. In mature IAM/GRC programs, trigger-based certifications are commonly implemented via HR-driven identity lifecycle workflows and access recertification campaigns that fire on "mover" events, ensuring old group memberships, application roles, and privileged access are re-evaluated immediately. This directly addresses the scenario described—permissions being added without removing prior ones—by making deprovisioning and re-approval an integral, automated step of the role-change process. See NIST guidance on access control and account management practices in [NIST SP 800-53 Rev. 5](#) and identity lifecycle concepts in [NIST SP 800-63-3](#).

QUESTION NO: 118

Which of the following is MOST appropriate to collect evidence of a zero-day attack?

- A. Honeypot
- B. Antispam
- C. Antivirus
- D. Firewall.

ANSWER: A

Explanation:

A honeypot is most appropriate for collecting evidence of a zero-day attack because it is a deliberately instrumented decoy system designed to attract and observe attacker behavior in a controlled environment. Since a zero-day exploit may not match known signatures or established detection logic, relying on traditional preventive controls alone is less effective for evidence collection. A honeypot (or honeynet) can be configured with extensive logging, full packet capture, file integrity monitoring, and process/audit telemetry to record the attacker's actions, tools, and exploit chain. This makes it particularly valuable for incident response and threat intelligence: you can capture indicators of compromise, exploit artifacts, command-and-control patterns, and post-exploitation techniques without exposing production assets. In CISSP terms, honeypots are classic detective controls that support monitoring, attribution, and forensic readiness by generating high-signal alerts and rich evidence when interacted with. Properly deployed, they also help validate whether an observed technique is novel (potentially zero-day) by providing reproducible traces for analysis and sharing with vendors or internal detection engineering teams. See [NIST Glossary: Honeypot](#) and [OWASP: Honeypots](#).

QUESTION NO: 119

A security architect is implementing an authentication system for a distributed network of servers. This network will be accessed by users on

workstations that cannot trust the identity of the user. Which solution should the security architect use to have the users trust one another?

- A. One-way authentication
- B. Kerberos
- C. Mutual authentication
- D. Single session software tokens

ANSWER: C

Explanation:

Mutual authentication is the correct solution because it ensures both parties in a connection validate each other's identity before any sensitive actions occur. In a distributed server environment where endpoints (workstations) cannot inherently trust the claimed identity presented by a user or a remote system, the key requirement is bidirectional assurance: the user (or client) must be confident they are talking to the legitimate service, and the service must be confident it is communicating with the legitimate user/client. Mutual authentication directly addresses this by requiring each side to prove its identity—commonly using certificates (mTLS), challenge-response mechanisms, or protocols that support two-way verification. This reduces the risk of man-in-the-middle attacks, rogue servers, and credential forwarding, and it is a foundational design principle for secure distributed systems and zero trust-style architectures. In practice, mutual authentication is frequently implemented via TLS with client certificates (mTLS) or via authentication frameworks that explicitly support reciprocal verification. For additional background on mutual TLS as a common mutual-auth pattern, see [Cloudflare: What is mutual TLS?](#) and [Microsoft: Mutual TLS pattern](#).

QUESTION NO: 120

A healthcare insurance organization chose a vendor to develop a software application. Upon review of the draft contract, the information security professional notices that software security is not addressed. What is the BEST approach to address the issue?

- A. Update the contract to require the vendor to perform security code reviews.
- B. Update the service level agreement (SLA) to provide the organization the right to audit the vendor.
- C. Update the contract so that the vendor is obligated to provide security capabilities.
- D. Update the service level agreement (SLA) to require the vendor to provide security capabilities.

ANSWER: C

Explanation:

The best approach is to update the contract so that the vendor is obligated to provide security capabilities. In a third-party software development engagement, security requirements must be established as binding contractual obligations (e.g., secure SDLC expectations, security requirements, acceptance criteria, vulnerability remediation timelines, and compliance obligations such as HIPAA-related safeguards where applicable). A contract is the governing legal instrument that defines deliverables and enforceable responsibilities; it is the right place to ensure security is treated as a required feature of the delivered application rather than an optional activity. While specific practices like code reviews or audit rights can be valuable, they are implementation mechanisms or oversight tools; they do not, by themselves, ensure the delivered product meets defined security outcomes. By contractually obligating security capabilities, the organization can then attach detailed security requirements, testing/verification expectations, and remedies for noncompliance (e.g., rework at vendor expense,

withholding payment, termination). This aligns with CISSP best practices for vendor management and secure acquisition: define security requirements up front and make them enforceable through procurement and contracting.

References: [NIST SP 800-161r1 \(Cybersecurity Supply Chain Risk Management\)](#), [OWASP SAMM \(Software Assurance Maturity Model\)](#)