DUMPSARENA

Systems Security Certified Practitioner

ISC2 SSCP

Version Demo

Total Demo Questions: 111

Total Premium Questions: 1114

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Security Operations and Administration	174
Topic 2, Access Controls	229
Topic 3, Risk Identification, Monitoring, and Analysis	68
Topic 4, Incident Response and Recovery	111
Topic 5, Cryptography	158
Topic 6, Network Security	283
Topic 7, Systems and Application Security	90
Topic 8, Mix Questions	1
Total	1114

QUESTION NO: 1

Which of the following statements pertaining to the security kernel is incorrect?

A. The security kernel is made up of mechanisms that fall under the TCB and implements and enforces the reference monitor concept.

B. The security kernel must provide isolation for the processes carrying out the reference monitor concept and they must be tamperproof.

C. The security kernel must be small enough to be able to be tested and verified in a complete and comprehensive manner.

D. The security kernel is an access control concept, not an actual physical component.

ANSWER: D

Explanation:

The statement "The security kernel is an access control concept, not an actual physical component." is incorrect. The access-control concept is the *reference monitor*: an abstract model requiring that every subject's attempted access to an object be checked against the applicable security policy. The security kernel, by contrast, is the actual trusted hardware, firmware, and software mechanism that implements this concept within a system. It is the portion of the trusted computing base responsible for mediating security-relevant accesses and enforcing the system's access-control rules.

A properly designed security kernel embodies the reference-monitor properties: it must be invoked for every relevant access, protected from tampering, and sufficiently small and simple to permit thorough testing and assurance. Although it may not correspond to a separate physical device, it is not merely an abstract access-control idea; it is an implementation component of the trusted computing environment. This distinction is important in SSCP security architecture because the reference monitor describes the policy-enforcement model, whereas the security kernel is the trusted mechanism that realizes that model. NIST defines a security kernel as the hardware, firmware, and software elements of a trusted computing base that implement the reference monitor concept. See NIST's Security Kernel glossary entry and NIST's Reference Monitor glossary entry.

QUESTION NO: 2

What can be defined as a momentary low voltage?

A. Spike

B. Sag

C. Fault

D. Brownout

ANSWER: B

Explanation:

Sag is correct because it is the standard power-quality term for a short-duration reduction in RMS supply voltage. A sag occurs when voltage falls below its normal level for a limited time and then returns to normal. In infrastructure and physical-security operations, sags matter because sensitive electronic equipment—including servers, network devices, controllers, and monitoring systems—may malfunction, reset, or experience service disruption even though electrical power was not completely lost. The duration and depth of the voltage reduction determine its practical effect, so organizations use appropriate power conditioning, uninterruptible power supplies, and monitoring to protect essential systems. IEEE power-quality terminology identifies a voltage sag as a decrease in RMS voltage or current lasting from 0.5 cycle through one minute. This definition distinguishes the event by both its reduced-voltage condition and its momentary duration, which directly matches the question. See the IEEE overview of IEEE 1159 power-quality monitoring terminology and Fluke's discussion of voltage sags, swells, and interruptions.

QUESTION NO: 3

What does the Clark-Wilson security model focus on?

A. Confidentiality

B. Integrity

C. Accountability

D. Availability

ANSWER: B

Explanation:

The Clark-Wilson security model focuses on **integrity**, particularly protecting the correctness and consistency of data used in commercial and business systems. Rather than merely restricting who can view information, the model is designed to ensure that data can be changed only through authorized, well-formed transactions. It distinguishes between constrained data items, which require integrity protection, and transformation procedures, which are the approved processes allowed to modify that data.

A central concept is that users do not manipulate protected data directly. Instead, they use controlled transformation procedures that preserve integrity rules. The model also supports separation of duties: important business actions can be divided among different authorized roles so that no individual can complete a sensitive process alone. Integrity verification procedures provide an additional mechanism for checking that constrained data remains in a valid state. These controls make the model especially applicable to financial, accounting, and transaction-processing environments, where preventing unauthorized or invalid changes is essential. ISC2's SSCP outline includes applying security models and principles in architecture and design, including protection of information assets' integrity. See the ISC2 SSCP Exam Outline and the NIST definition of integrity.

QUESTION NO: 4

What can be defined as: It confirms that users' needs have been met by the supplied solution ?

A. Accreditation

B. Certification

C. Assurance

D. Acceptance

ANSWER: D

Explanation:

Acceptance is correct because it is the formal determination that a delivered solution satisfies the agreed user needs and acceptance criteria. It occurs after the solution has been supplied and evaluated against requirements that are observable and testable, such as required functions, performance measures, interfaces, usability needs, and operational constraints. Acceptance gives the customer or authorized business representative the basis to acknowledge that the delivered system is suitable for its intended use and can be placed into service. User acceptance testing is a common activity supporting this decision: representative users execute realistic business scenarios to confirm that the solution supports required processes in the target environment. Acceptance criteria should be defined early, be traceable to requirements, and provide objective evidence for the acceptance decision. This framing is consistent with ISO/IEC/IEEE 29148, which addresses stakeholder needs, requirements, and verification criteria across the system life cycle, and with NIST guidance describing acceptance testing as testing performed to determine whether a system meets acceptance criteria and is ready for operational use. See ISO/IEC/IEEE 29148:2018 and NIST CSRC: acceptance testing.

QUESTION NO: 5

Which of the following would be LESS likely to prevent an employee from reporting an incident?

A. They are afraid of being pulled into something they don't want to be involved with.

B. The process of reporting incidents is centralized.

C. They are afraid of being accused of something they didn't do.

D. They are unaware of the company's security policies and procedures.

ANSWER: B

Explanation:

The process of reporting incidents is centralized is the condition least likely to discourage reporting. A centralized reporting process gives personnel a defined, consistent route for escalating suspected security events, such as a security operations center, incident-response team, service desk, or designated reporting portal. When employees know exactly where and how to report, they do not need to decide which department owns the issue or try to locate an appropriate individual during a potentially stressful event. This clarity reduces delay and supports timely triage, documentation, and coordination by the incident-response function.

Effective incident-response programs establish clear reporting procedures and communication paths so that events can be reported, assessed, and handled promptly. Centralization also enables the organization to apply a uniform intake process, ensure reports reach trained responders, and maintain appropriate records for follow-up. For this benefit to be realized, the centralized channel should be well communicated, easy to use, and available to all personnel. NIST's incident-response guidance emphasizes defined reporting and coordination mechanisms as part of an effective capability; see NIST SP 800-61 Revision 3 and CISA information-sharing guidance.

QUESTION NO: 6

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 network authentication protocol that uses a trusted third party, called a Key Distribution Center (KDC), to authenticate users and services over an untrusted network. The KDC consists of an Authentication Server and a Ticket-Granting Server. Rather than transmitting a user's password to each requested service, Kerberos uses time-limited cryptographic tickets. After initial authentication, a client obtains a ticket-granting ticket and can then request service tickets for authorized resources. This design supports single sign-on while providing mutual authentication: the client can verify the service's identity and the service can verify the client's identity. Kerberos was originally developed at MIT and is widely implemented in enterprise identity systems, including Microsoft Active Directory environments. Its reliance on shared secret cryptography, tickets, and a centrally trusted KDC is why "A trusted third-party authentication protocol" accurately describes it. For technical protocol details, see the IETF's RFC 4120: The Kerberos Network Authentication Service (V5). MIT also provides an overview of Kerberos and its role in secure network authentication.

QUESTION NO: 7

The Information Technology Security Evaluation Criteria (ITSEC) was written to address which of the following that the Orange Book did not address?

A. integrity and confidentiality.

B. confidentiality and availability.

C. integrity and availability.

D. none of the above.

ANSWER: C

Explanation:

ITSEC was designed as a broader set of security evaluation criteria than the U.S. Trusted Computer System Evaluation Criteria, commonly called the Orange Book. The Orange Book's foundational security objective was protecting classified information through mandatory access controls and related mechanisms that preserve confidentiality. ITSEC retained confidentiality as an important security property but explicitly expanded the evaluation scope to encompass integrity and availability as well. Integrity addresses whether information and system functions remain accurate, complete, and protected from unauthorized modification. Availability addresses whether authorized users can access systems, services, and information when required. Consequently, the correct response is *integrity and availability.*

This difference reflects ITSEC's applicability beyond the Orange Book's original military classification-oriented context. ITSEC allowed a security target to state the particular security objectives required for a product or system, including objectives focused on preventing unauthorized alteration and ensuring dependable service. This broader, objective-based approach helped inform later international evaluation work, including the Common Criteria framework. NIST's overview of historical evaluation criteria describes the Orange Book's emphasis on confidentiality and identifies ITSEC as including confidentiality, integrity, and availability objectives. See NIST ITL Bulletin: An Introduction to the Common Criteria and Common Criteria Portal.

QUESTION NO: 8

Single Sign-on (SSO) is characterized by which of the following advantages?

A. Convenience

B. Convenience and centralized administration

C. Convenience and centralized data administration

D. Convenience and centralized network administration

ANSWER: B

Explanation:

Convenience and centralized administration is correct. Single sign-on enables an authenticated user to access multiple related applications or services without separately entering credentials for each one. This improves the user experience by reducing repetitive password entry and the number of credentials users must remember. It can also reduce password-reset demand and the unsafe practice of reusing passwords across systems. SSO relies on a trusted identity provider or authentication authority that establishes the user's identity and provides assertions or tokens to relying applications. Because authentication and account-access policy can be managed through a central identity service, administrators can more consistently apply controls such as account provisioning and deprovisioning, multifactor-authentication requirements, session policies, and access decisions across integrated services. Centralizing identity administration does not mean centralizing the organization's underlying business data or network administration; the benefit is centralized management of authentication and identity-related access. NIST describes federated identity and SSO as mechanisms that allow users to authenticate once and access multiple applications, while its digital identity guidance addresses centralized identity lifecycle and authenticator management. See NIST SP 800-63C, Federation and Assertions and NIST SP 800-63A, Identity Proofing and Enrollment.

QUESTION NO: 9

Which of the following is a symmetric encryption algorithm?

A. RSA

B. Elliptic Curve

C. RC5

D. El Gamal

ANSWER: C

Explanation:

RC5 is a symmetric-key block cipher designed by Ronald Rivest. Symmetric encryption uses the same shared secret key for both encryption and decryption, so authorized parties must securely establish and protect that key before using the cipher. RC5 is parameterized by its word size, number of rounds, and secret-key length, allowing implementations to select configurations suitable for their operational and security requirements. Its round operations combine modular integer addition, exclusive OR, and data-dependent rotations; this simple structure is commonly summarized as ARX-style design. As a block cipher, RC5 transforms fixed-size blocks of plaintext into ciphertext under the selected secret key, and secure practical use also requires an appropriate mode of operation and key-management process. The original RC5 specification identifies it as a "symmetric block cipher," directly establishing that it belongs to the symmetric-encryption category tested in the question. See Ronald Rivest's RC5 Encryption Algorithm specification and NIST's definition of symmetric key cryptography.

QUESTION NO: 10

A Security Kernel is defined as a strict implementation of a reference monitor mechanism responsible for enforcing a security policy. To be secure, the kernel must meet three basic conditions, what are they?

A. Confidentiality, Integrity, and Availability

B. Policy, mechanism, and assurance

C. Isolation, layering, and abstraction

D. Completeness, Isolation, and Verifiability

ANSWER: D

Explanation:

Completeness, Isolation, and Verifiability are the three conditions required of a security kernel because they are the implementation-oriented properties of the reference monitor concept. Completeness means that every access request to a protected object must be mediated by the security kernel; no process, user, or interface can bypass its enforcement of the security policy. This is also commonly called complete mediation. Isolation means that the kernel and its enforcement functions are protected against unauthorized modification or interference. The kernel must operate within a protected domain so that an attacker or untrusted process cannot alter the rules, access-control decisions, or security-relevant state it relies on. Verifiability means the kernel must be sufficiently small, simple, and well specified that its correct implementation and conformance to the security policy can be analyzed, tested, or otherwise demonstrated. These properties establish that security decisions are consistently applied, protected from tampering, and capable of assurance. The NIST glossary defines a reference monitor as an access-control concept that mediates accesses and is tamperproof and verifiable, characteristics implemented by a security kernel. See the NIST reference monitor glossary entry and NIST's security architecture and assurance resources.

QUESTION NO: 11

What is the primary purpose of a vulnerability scan?

A. To identify known weaknesses in systems and applications

B. To establish legal ownership of digital evidence

C. To restore services after an outage

D. To encrypt traffic between two networks

ANSWER: A

Explanation:

To identify known weaknesses in systems and applications is correct. A vulnerability scan uses automated tools to examine hosts, services, operating systems, and applications for known vulnerabilities, missing patches, insecure configurations, and other conditions that may require remediation. The results help an organization prioritize corrective action according to the severity and exposure of identified issues.

A vulnerability scan differs from a penetration test because it normally identifies potential weaknesses without attempting to exploit them to demonstrate impact. Scanning should be authorized and appropriately configured because active scans can generate substantial network traffic or affect fragile systems. NIST discusses vulnerability scanning as a control assessment activity in NIST SP 800-115, Technical Guide to Information Security Testing and Assessment.

QUESTION NO: 12

The IP header contains a protocol field. If this field contains the value of 2, what type of data is contained within the IP datagram?

A. TCP.

B. ICMP.

C. UDP.

D. IGMP.

ANSWER: D

Explanation:

The correct answer is **IGMP.** In an IPv4 header, the Protocol field is an 8-bit value that identifies the next encapsulated protocol—the type of payload carried by the IP datagram—so that the receiving host can deliver the payload to the appropriate network-layer or transport-layer handler. The Internet Assigned Numbers Authority (IANA) maintains the authoritative protocol-number registry, where decimal value 2 is assigned to the Internet Group Management Protocol (IGMP). IGMP is used by IPv4 hosts and adjacent multicast routers to communicate multicast group-membership information. For example, a host uses IGMP to report that it wishes to receive traffic sent to a particular IPv4 multicast group, while routers use the membership information to determine where multicast traffic should be forwarded. Therefore, an IPv4 packet whose Protocol field contains 2 carries an IGMP message as its payload. The protocol-number assignment is standardized in the IANA Protocol Numbers registry and is also reflected in the IPv4 specification's discussion of the Protocol field. See the IANA Protocol Numbers registry and RFC 791: Internet Protocol.

QUESTION NO: 13

How many bits is the effective length of the key of the Data Encryption Standard algorithm?

A. 168

B. 128

C. 56

D. 64

ANSWER: C

Explanation:

The effective key length of the Data Encryption Standard algorithm is **56** bits. DES accepts a 64-bit key value, but eight of those bits are reserved for parity checking: one parity bit is present in each 8-bit byte of the supplied key. Because these parity bits do not contribute independent secret key material to the cryptographic operation, only 56 bits determine the encryption transformation. Therefore, DES has a 56-bit effective key space, which contains 2^{56} possible keys.

This distinction between the represented key size and effective cryptographic key size is important in security assessments. A 56-bit symmetric key is now far too small to resist exhaustive key-search attacks with modern computing capabilities; accordingly, DES has long been withdrawn from use in U.S. federal systems. NIST formally withdrew the DES standard and identifies its 56-bit key length in the original specification. For modern applications, organizations should use an approved contemporary algorithm, such as AES, with an appropriate key size and implementation mode.

References: NIST FIPS 46-3: Data Encryption Standard and NIST CSRC: FIPS 46-3 publication record.

QUESTION NO: 14

Which of the following statements pertaining to Kerberos is false?

A. The Key Distribution Center represents a single point of failure.

B. Kerberos manages access permissions.

C. Kerberos uses a database to keep a copy of all users' public keys.

D. Kerberos uses symmetric key cryptography.

ANSWER: C

Explanation:

“Kerberos uses a database to keep a copy of all users' public keys.” is false because standard Kerberos is based on shared symmetric secret keys, not a directory of each user’s public keys. A Kerberos Key Distribution Center (KDC) maintains principal records and long-term secret-key material for principals, enabling it to issue tickets and session keys. A client proves knowledge of its secret key, while services use their own secret keys to process tickets issued for them. This trusted-third-party design lets the client and service establish authenticated, time-limited sessions without transmitting the user’s password across the network.

Public-key cryptography can be used by optional Kerberos extensions, such as Public Key Cryptography for Initial Authentication in Kerberos (PKINIT), to perform initial authentication in particular deployments. That optional capability does not change the core Kerberos key database into a repository of every user’s public key, nor does it make public keys the standard credential material described by the protocol. RFC 4120 specifies Kerberos’s use of secret keys and tickets, while RFC 4556 defines the optional PKINIT extension. See RFC 4120: The Kerberos Network Authentication Service (V5) and RFC 4556: PKINIT.

QUESTION NO: 15

What is called the probability that a threat to an information system will materialize?

A. Threat

B. Risk

C. Vulnerability

D. Hole

ANSWER: B

Explanation:

Risk is the correct term because it expresses the possibility that a threat event will occur and adversely affect an information system or its assets. In information-security risk management, risk is assessed by considering both the likelihood of a threat exploiting a vulnerability and the resulting impact if that event occurs. Therefore, when the question focuses on the probability that a threat will materialize, it is identifying the likelihood component of risk and, in common security terminology, is asking for risk.

This distinction supports practical security decisions: organizations identify threats and vulnerabilities, estimate the likelihood that relevant threat events may occur, evaluate potential business impact, and then prioritize treatment of the resulting risks. NIST describes risk as a measure of the extent to which an entity is threatened by a potential circumstance or event, typically considering likelihood and impact. ISC2's SSCP certification is centered on applying risk-management principles to protect information systems and operations. See the NIST definition of risk and the ISC2 SSCP Exam Outline.

QUESTION NO: 16

Which of the following is not a disadvantage of symmetric cryptography when compared with Asymmetric Ciphers?

A. Provides Limited security services

B. Has no built in Key distribution

C. Speed

D. Large number of keys are needed

ANSWER: C

Explanation:

Speed is correct because it is a principal operational advantage of symmetric cryptography rather than a disadvantage. Symmetric algorithms use one shared secret key for encryption and decryption and are designed to process substantial volumes of data efficiently. Their computational operations are comparatively inexpensive, making them well suited for bulk encryption, such as protecting files, database contents, backups, storage volumes, and network-session traffic.

In practical cryptographic systems, asymmetric cryptography is commonly used only for limited operations such as authentication, establishing a shared secret, or securely transporting a symmetric session key. Once that session key is established, a symmetric algorithm performs the actual data encryption because it delivers much better throughput and lower processing overhead. This hybrid approach is used by protocols such as TLS and reflects the complementary roles of the two cryptographic approaches.

NIST notes that symmetric-key cryptography is appropriate for efficiently protecting data and distinguishes it from public-key techniques used for key establishment and related services. Further guidance on selecting and using cryptographic mechanisms is available in NIST SP 800-57 Part 1 and the NIST definition of symmetric-key cryptography.

QUESTION NO: 17

The MOST common threat that impacts a business's ability to function normally is:

A. Power Outage

B. Water Damage

C. Severe Weather

D. Labor Strike

ANSWER: A

Explanation:

Power Outage is correct because loss of electrical utility service is a highly prevalent operational disruption and can immediately affect nearly every business dependency: lighting, heating and cooling, network connectivity,

telecommunications, physical access systems, point-of-sale equipment, servers, and user endpoints. Even organizations with resilient applications can be unable to deliver normal services when the facilities, communications paths, or supporting infrastructure on which those services depend lose power. For this reason, business continuity and contingency planning commonly treat utility failure as a baseline scenario that must be evaluated for likelihood, duration, operational impact, and recovery requirements.

Appropriate resilience measures include uninterruptible power supplies to bridge short interruptions, standby generators for longer outages, tested fuel and maintenance arrangements, graceful shutdown procedures, redundant utility feeds where feasible, and alternate work or processing capabilities. The U.S. Cybersecurity and Infrastructure Security Agency identifies power outages as hazards requiring preparedness and continuity planning, while Ready.gov provides business-focused guidance for assessing risks and maintaining essential functions during disruptions. These sources support treating power interruption as a routine, high-impact threat scenario in organizational risk management and continuity planning. See CISA: Power Outages and Ready.gov: Business Preparedness.

QUESTION NO: 18

Which of the following is not one of the three goals of Integrity addressed by the Clark-Wilson model?

A. Prevention of the modification of information by unauthorized users.

B. Prevention of the unauthorized or unintentional modification of information by authorized users.

C. Preservation of the internal and external consistency.

D. Prevention of the modification of information by authorized users.

ANSWER: D

Explanation:

Prevention of the modification of information by authorized users. is not one of the three integrity goals addressed by the Clark-Wilson model. Clark-Wilson is a commercial integrity model designed to ensure that data remains valid and that business processes handle it only through controlled, authorized transactions. It does not seek to categorically prevent authorized users from modifying information, because legitimate business operations necessarily require authorized personnel and processes to make changes.

Instead, the model's integrity objective is to ensure that authorized changes occur only in an appropriate, controlled manner. This is accomplished through well-formed transactions, implemented as transformation procedures, and through separation of duties and authorization controls. These mechanisms help prevent an authorized individual from making an unauthorized change or from accidentally making an improper change, while still allowing necessary and approved modifications. Clark-Wilson also emphasizes maintaining consistency by ensuring that certified transformation procedures preserve integrity constraints and that access to constrained data items is mediated through those procedures.

The original Clark-Wilson paper describes integrity in terms of enforcing well-formed transactions and separation of duty in commercial systems: A Comparison of Commercial and Military Computer Security Policies. For additional security-model terminology, see the NIST Computer Security Resource Center glossary.

QUESTION NO: 19

Which communication method is characterized by very high speed transmission rates that are governed by electronic clock timing signals?

A. Asynchronous Communication.

B. Synchronous Communication.

C. Automatic Communication.

D. Full duplex Communication.

ANSWER: B

Explanation:

Synchronous Communication is correct because the communicating devices use a shared or coordinated timing reference—an electronic clock—to determine precisely when each bit or group of bits is sent and received. This timing synchronization lets the receiver identify data boundaries without requiring separate start and stop indicators for every character. Reducing that per-character framing overhead supports efficient, continuous data transfer and makes synchronous transmission well suited to high-speed communications links. In security and networking contexts, understanding the timing model matters because it distinguishes a transmission method based on coordinated clocking from communication descriptions based on direction of traffic flow or unrelated operational terms. The clock signal may be provided over a dedicated timing path, embedded in the transmitted signal, or recovered by the receiving equipment, but its purpose remains to keep the sender and receiver aligned throughout the exchange. The International Telecommunication Union describes synchronous transmission as transmission in which the time interval between characters is fixed, while Cisco's serial-interface documentation similarly identifies clocking as fundamental to synchronous serial links. See ITU-T Recommendation V.250 and Cisco: Understanding Serial Interfaces.

QUESTION NO: 20

Within the OSI model, at what layer are some of the SLIP, CSLIP, PPP control functions provided?

A. Data Link

B. Transport

C. Presentation

D. Application

ANSWER: A

Explanation:

Data Link is correct because SLIP, CSLIP, and PPP operate over point-to-point links and provide functions associated with managing the link that carries network-layer packets. In OSI terms, these are Layer 2 responsibilities: framing or encapsulating packets for transmission over a specific link, establishing link parameters, monitoring the connection, and terminating or testing the link as appropriate.

PPP makes this mapping especially clear. RFC 1661 defines PPP as a method for transporting multiprotocol datagrams over point-to-point links and identifies its Link Control Protocol as responsible for establishing, configuring, and testing the data-link connection. PPP also uses Network Control Protocols to configure the network-layer protocols carried across that established link. Thus, although PPP transports Layer 3 traffic and includes mechanisms that support particular network protocols, its core link-management and encapsulation functions are categorized at the Data Link layer. SLIP and its compressed-header variant, CSLIP, likewise provide serial-link encapsulation mechanisms used to carry IP traffic across point-to-point serial connections. See RFC 1661: The Point-to-Point Protocol and RFC 1055: A Nonstandard for Transmission of IP Datagrams over Serial Lines.

QUESTION NO: 21

Which of the following server contingency solutions offers the highest availability?

A. System backups

B. Electronic vaulting/remote journaling

C. Redundant arrays of independent disks (RAID)

D. Load balancing/disk replication

ANSWER: D

Explanation:

Load balancing/disk replication offers the highest availability among the listed server-contingency solutions because it combines two complementary resilience capabilities. Load balancing distributes service requests across multiple servers or instances, allowing processing to continue when an individual server becomes unavailable and reducing a single-server outage as a point of failure. Disk replication maintains a current or near-current copy of system data on separate storage or at another location, which substantially reduces data loss and supports rapid restoration or failover after a hardware, system, or site disruption. Together, these controls can keep the application service accessible while preserving recoverable data, minimizing both downtime and recovery-point exposure. NIST contingency-planning guidance treats alternate processing and replicated data arrangements as measures that strengthen recovery capability, and its earlier server-contingency comparison identifies this combined approach as providing the greatest availability of the listed technologies. Availability is achieved not merely by retaining a copy of data, but by maintaining sufficient redundant processing and synchronized data to continue or quickly resume the business service. See NIST SP 800-34 Rev. 1, Contingency Planning Guide for Federal Information Systems and the NIST SP 800-34 publication record.

QUESTION NO: 22

In the days before CIDR (Classless Internet Domain Routing), networks were commonly organized by classes. Which of the following would have been true of a Class B network?

A. The first bit of the IP address would be set to zero.

B. The first bit of the IP address would be set to one and the second bit set to zero.

C. The first two bits of the IP address would be set to one, and the third bit set to zero.

D. The first three bits of the IP address would be set to one.

ANSWER: B

Explanation:

The first bit of the IP address would be set to one and the second bit set to zero. This is the defining leading-bit pattern for a classful IPv4 Class B address: binary `10`. In the original classful addressing model, those two leading bits identified the address as Class B, leaving the next 14 bits for the network number and the remaining 16 bits for host addressing. Consequently, Class B addresses occupied the decimal first-octet range 128 through 191, and their traditional default subnet mask was 255.255.0.0, or `/16`.

For example, the first Class B address begins with the binary octet `10000000`, which is decimal 128, as in 128.0.0.0. The final possible first octet with the Class B prefix is binary `10111111`, or decimal 191. Thus, recognizing the initial `10` bit sequence is the key property that distinguished a Class B network under the historical classful scheme. CIDR replaced fixed class boundaries with variable-length prefixes for route aggregation and more flexible allocation, but knowledge of the former class ranges remains useful when interpreting legacy networking material. The original Internet Protocol specification documents the historical address-class interpretation in RFC 791; IANA also maintains the authoritative IPv4 special-purpose address registry.

QUESTION NO: 23

How many bits of a MAC address uniquely identify a vendor, as provided by the IEEE?

A. 6 bits

B. 12 bits

C. 16 bits

D. 24 bits

ANSWER: D

Explanation:

24 bits is the expected answer for the conventional 48-bit MAC address format tested in foundational networking and SSCP material. The first 24 bits are the Organizationally Unique Identifier (OUI), also called a MAC Address Block Large (MA-L) prefix. IEEE assigns this prefix to an organization, thereby identifying the manufacturer or other assignee. The organization then administers the remaining 24 bits to create individual interface addresses within that assigned block. This arrangement supports more than 16 million distinct device addresses for each 24-bit OUI prefix.

The distinction is important because the OUI identifies the organization associated with the address assignment, while the complete 48-bit MAC address identifies a particular network interface address. In hexadecimal notation, the OUI is normally represented by the first three octets, such as `00:1A:2B` in a MAC address written as `00:1A:2B:3C:4D:5E`. IEEE's public registry describes MA-L assignments as 24-bit prefixes and provides the authoritative registry for these assignments. See the IEEE Registration Authority and the IEEE OUI/MA-L registry.

QUESTION NO: 24

You are running a packet sniffer on a network and see a packet containing a long string of "0x90 0x90 0x90 0x90...." in the middle of it traveling to an x86-based machine as a target. This could be indicative of what activity being attempted?

A. Over-subscription of the traffic on a backbone.

B. A source quench packet.

C. A FIN scan.

D. A buffer overflow attack.

ANSWER: D

Explanation:

A buffer overflow attack is the best answer because hexadecimal `0x90` is the x86 instruction opcode for *NOP*, meaning "no operation." A long consecutive sequence of these bytes is commonly called a NOP sled. In classic stack-based buffer-overflow exploitation, an attacker sends oversized input containing a NOP sled followed by injected payload code. The sled creates a broad range of addresses that can safely receive redirected execution: if a vulnerable program's control flow is overwritten to point anywhere within the repeated `0x90` sequence, the processor executes NOP instructions until it reaches the payload. Observing this recognizable pattern in traffic destined for an x86 host can therefore indicate an attempt to exploit a memory-safety flaw by overflowing a buffer and gaining code execution. This signature is especially associated with older or unmitigated native-code applications, although modern protections such as non-executable memory, address-space layout randomization, stack canaries, and control-flow protections make successful exploitation more difficult. NIST describes buffer-overflow vulnerabilities as weaknesses caused by inadequate bounds checking that attackers may use to execute arbitrary code. See MITRE CWE-120: Buffer Copy without Checking Size of Input and NIST SP 800-123, Guide to General Server Security.

QUESTION NO: 25

Which of the following statements pertaining to ethical hacking is incorrect?

A. An organization should use ethical hackers who do not sell auditing, hardware, software, firewall, hosting, and/or networking services.

B. Testing should be done remotely to simulate external threats.

C. Ethical hacking should not involve writing to or modifying the target systems negatively.

D. Ethical hackers never use tools that have the potential of affecting servers or services.

ANSWER: D

Explanation:

“Ethical hackers never use tools that have the potential of affecting servers or services.” is the incorrect statement because authorized security testing can require tools and techniques that, by their nature, may consume resources, trigger defensive controls, alter test data, or create a risk of service disruption. Examples can include vulnerability scanners, password-auditing tools, web-application testing tools, and controlled exploit validation. Ethical practice does not depend on an unrealistic promise that a tool can never affect a target. Instead, it depends on explicit authorization, a defined scope, rules of engagement, careful selection and configuration of tools, appropriate maintenance windows, monitoring, and clear stop conditions.

Before testing begins, the tester and system owner should agree on what systems may be tested, which techniques are permitted, how sensitive findings and data will be handled, escalation contacts, and the acceptable level of operational risk. These controls allow an organization to obtain meaningful assurance while managing the possibility that testing activity could affect availability or system behavior. NIST’s technical guide for security testing describes planning, authorization, and rules of engagement as essential components of a testing program. ISC2’s ethical standards likewise emphasize acting responsibly and protecting society and infrastructure. See NIST SP 800-115 and the ISC2 Code of Ethics.

QUESTION NO: 26

Encapsulating Security Payload (ESP) provides some of the services of Authentication Headers (AH), but it is primarily designed to provide:

A. Confidentiality

B. Cryptography

C. Digital signatures

D. Access Control

E. & KRAUSE, MICKI, Information Security Management Handbook, 4th Edition, Volume 2, 2001, CRC Press, NY, page 164.

ANSWER: A

Explanation:

Confidentiality is correct because IPsec Encapsulating Security Payload was principally defined to protect the secrecy of IP packets by encrypting the payload being transported. ESP can operate in transport mode, protecting the upper-layer payload of an IP packet, or tunnel mode, encapsulating and protecting an entire original IP packet. This encryption prevents an unauthorized party who intercepts protected traffic from reading the protected data.

ESP may also provide integrity and data-origin authentication when the applicable integrity mechanism is selected. Those services overlap with services associated with the Authentication Header protocol, which is why the question notes that ESP provides some Authentication Header services. Nevertheless, encryption and resulting confidentiality are ESP's central purpose. Modern IPsec specifications define ESP as providing confidentiality, and optionally integrity, data-origin authentication, and anti-replay protection; its use of encryption is what distinguishes its primary security objective in this context.

For the IPsec architecture's description of ESP security services, see RFC 4301, Security Architecture for IP. The ESP protocol specification and its confidentiality service are detailed in RFC 4303, IP Encapsulating Security Payload.

QUESTION NO: 27

While using IPsec, the ESP and AH protocols both provides integrity services. However when using AH, some special attention needs to be paid if one of the peers uses NAT for address translation service. Which of the items below would affects the use of AH and it´s Integrity Check Value (ICV) the most?

A. Key session exchange

B. Packet Header Source or Destination address

C. VPN cryptographic key size

D. Cryptographic algorithm used

ANSWER: B

Explanation:

Packet Header Source or Destination address is correct because IPsec Authentication Header (AH) integrity protection covers applicable immutable fields in the IP header, in addition to the protected payload and upper-layer protocol information. The source and destination IP addresses are among the header fields whose values are expected to remain unchanged while an AH-protected packet travels from sender to receiver. AH's Integrity Check Value (ICV) is calculated by the sender over this protected information; the receiver performs the corresponding integrity verification using the same security association and key material.

Network Address Translation changes an IP packet's source address, destination address, or both as the packet passes through the NAT device. That modification changes data covered by AH's authentication calculation. Consequently, the receiving peer calculates a different ICV and integrity verification fails, causing the packet to be rejected. This behavior is fundamental to AH's design: it provides origin authentication and connectionless integrity for the protected IP packet, including selected IP-header fields, rather than allowing intermediary address rewriting.

RFC 4302 specifies that AH authenticates immutable IP-header fields and notes the incompatibility with address translation. NAT traversal designs therefore generally use ESP with NAT Traversal mechanisms, rather than AH, when address translation is present. See RFC 4302, IP Authentication Header and RFC 3948, UDP Encapsulation of IPsec ESP Packets.

QUESTION NO: 28

Which of the following biometric characteristics cannot be used to uniquely authenticate an individual's identity?

A. Retina scans

B. Iris scans

C. Palm scans

D. Skin scans

ANSWER: D

Explanation:

Skin scans is the correct answer in the conventional biometric-authentication taxonomy used for SSCP-style access-control questions. A biometric authenticator must measure sufficiently distinctive, repeatable physiological or behavioral features and compare a newly captured sample with a previously enrolled reference. Although skin may be examined as part of a specialized medical, forensic, or liveness-detection process, "skin scans" is not a standard, uniquely identifying biometric modality as stated. The term does not identify a stable, established feature set or matching method comparable to the named biometric modalities.

In contrast, established biometric systems use characteristics such as iris or retina patterns and palm features, including palm prints or vein patterns, because they can produce templates with enough discriminatory information for matching. Biometric systems also require attention to presentation-attack detection, error rates, enrollment quality, and secure protection of biometric reference data; a biometric match supports identity verification but should be implemented within a broader authentication design. NIST describes biometric comparison and its role in identity proofing and authentication in NIST SP 800-63B. Additional technical guidance on biometric performance evaluation is available from NIST's Biometric Technology Evaluation program.

QUESTION NO: 29

In the context of access control, locks, gates, guards are examples of which of the following?

A. Administrative controls

B. Technical controls

C. Physical controls

D. Logical controls

ANSWER: C

Explanation:

Locks, gates, and guards are examples of **Physical controls**. Physical access controls protect facilities, equipment, personnel, and other tangible assets by limiting or regulating access to physical spaces. A lock restricts entry through a door, gate, cabinet, or other barrier; a gate establishes a controlled boundary around a property or protected area; and a guard provides a human security presence to verify authorization, monitor activity, and respond to incidents. These safeguards are commonly layered together as part of defense in depth. For example, a secured facility may use perimeter gates, locked exterior and interior doors, badge-controlled entry points, surveillance, and security personnel to establish increasingly protected zones around valuable assets. The SSCP access-control domain distinguishes physical measures from controls implemented through policy and procedures or through information-system technologies. Physical controls are essential because technical protections alone cannot prevent an unauthorized individual from entering a server room, stealing equipment, or tampering with infrastructure. NIST's physical and environmental protection control family similarly addresses safeguards for physical access to systems, facilities, and components. See NIST SP 800-53 Rev. 5, Physical and Environmental Protection (PE) controls, and ISC2's overview of the SSCP Certification Exam Outline.

QUESTION NO: 30

If an employee's computer has been used by a fraudulent employee to commit a crime, the hard disk may be seized as evidence and once the investigation is complete it would follow the normal steps of the Evidence Life Cycle. In such case, the Evidence life cycle would not include which of the following steps listed below?

A. Acquisition collection and identification

B. Analysis

C. Storage, preservation, and transportation

D. Destruction

ANSWER: D

Explanation:

Destruction is the step not normally included for the seized hard disk in this scenario. Digital-forensic evidence handling is intended to preserve the item and its evidentiary value from acquisition through examination, analysis, reporting or presentation, and final disposition. Preservation includes maintaining documented custody and protecting the original media from alteration, damage, or loss while it is stored or transported. Once the legal need for the hard disk has ended and the court or investigative authority authorizes release, an ordinary piece of property such as an employee computer's disk is normally returned to its lawful owner or otherwise disposed of under the applicable legal process.

Destruction can occur under exceptional legal circumstances—for example, where property is contraband, subject to a destruction order, or cannot lawfully be returned—but those circumstances are not present in the stated case. A hard disk used as evidence in a fraud investigation is not inherently contraband. NIST's digital-forensics guidance describes the core forensic process as collection, examination, analysis, and reporting, with preservation and integrity maintained throughout. The U.S. Department of Justice also stresses preserving electronic evidence and documenting its handling to support admissibility and reliable findings. See NIST SP 800-86, Guide to Integrating Forensic Techniques into Incident Response and the U.S. DOJ Electronic Crime Scene Investigation guide.

QUESTION NO: 31

Which of the following exemplifies proper separation of duties?

A. Operators are not permitted modify the system time.

B. Programmers are permitted to use the system console.

C. Console operators are permitted to mount tapes and disks.

D. Tape operators are permitted to use the system console.

ANSWER: A

Explanation:

“Operators are not permitted modify the system time.” exemplifies separation of duties because it prevents personnel who operate production systems from also controlling a security-sensitive setting that affects the reliability of audit evidence. System timestamps establish the sequence of events in logs, alerts, transactions, backups, and other records used for monitoring, incident response, and investigations. Restricting time changes to a separately authorized administrative function helps ensure that an operator cannot conceal or alter the apparent timing of activity by changing the system clock. This division creates an independent control over a critical action and reduces the opportunity for error, misuse, or fraud by a single individual. Separation of duties is intended to divide responsibilities so that no one role can perform all parts of a sensitive process without oversight. NIST’s Access Control guidance specifically calls for separating duties among different individuals to reduce the risk of malevolent activity occurring without detection. See NIST SP 800-53 Rev. 5, control AC-5. Maintaining trustworthy time information is also a recognized security control because consistent, accurate timestamps support audit-log correlation and analysis; see NIST SP 800-92, Guide to Computer Security Log Management.

QUESTION NO: 32

Which of the following would be MOST important to guarantee that the computer evidence will be admissible in court?

A. It must prove a fact that is immaterial to the case.

B. Its reliability must be proven.

C. The process for producing it must be documented and repeatable.

D. The chain of custody of the evidence must show who collected, secured, controlled, handled, transported the evidence, and that it was not tampered with.

ANSWER: D

Explanation:

The chain of custody of the evidence must show who collected, secured, controlled, handled, transported the evidence, and that it was not tampered with. This is the most important choice because digital evidence is particularly susceptible to alteration, duplication, and questions about its origin. A complete, contemporaneous chain-of-custody record establishes the evidence’s history from acquisition through storage, analysis, and presentation. It identifies each person who possessed the item, the dates and times of transfers, the purpose of handling, and the controls used to preserve integrity.

Maintaining this documented history enables an examiner or investigator to demonstrate that the presented data is the same data that was originally collected and that it was protected against unauthorized modification. This supports the evidence’s authenticity and reliability, foundational considerations for courtroom admissibility. In practice, evidence handling should include secure storage, controlled access, integrity verification such as cryptographic hashes for forensic images, and records of every transfer or examination. These practices provide a defensible basis for testifying that the evidence remained intact. NIST’s forensic guidance emphasizes preserving evidence and documenting actions throughout the investigative process; see NIST SP 800-86. Additional evidence-handling guidance is available from the U.S. National Institute of Justice.

QUESTION NO: 33

The IP header contains a protocol field. If this field contains the value of 1, what type of data is contained within the IP datagram?

A. TCP.

B. ICMP.

C. UDP.

D. IGMP.

ANSWER: B

Explanation:

ICMP. is correct. In an IPv4 packet header, the Protocol field is an 8-bit value that identifies the next-header protocol encapsulated in the IP payload. A value of 1 is assigned to the Internet Control Message Protocol (ICMP). Consequently, a datagram whose IPv4 Protocol field contains 1 carries an ICMP message, such as a diagnostic, error-reporting, or network-control message. ICMP is used by IP hosts and routers to communicate conditions related to packet delivery and network reachability. Common tools, including ping and traceroute implementations, rely on ICMP messages to test reachability and help identify forwarding-path behavior. The authoritative protocol-number registry maintained by IANA lists decimal protocol number 1 as ICMP, and RFC 791 defines the IPv4 header Protocol field as the field used to identify the protocol carried in the IP datagram's data portion. See the IANA Protocol Numbers registry and RFC 791: Internet Protocol.

QUESTION NO: 34

What is the main characteristic of a bastion host?

A. It is located on the internal network.

B. It is a hardened computer implementation

C. It is a firewall.

D. It does packet filtering.

ANSWER: B

Explanation:

A bastion host is distinguished primarily by being a hardened computer implementation. It is deliberately configured to resist attack because it is exposed to a higher-risk network location, commonly at a network boundary or within a demilitarized zone. Hardening reduces the host's attack surface through measures such as disabling unnecessary services, removing unneeded software and accounts, applying timely security updates, enforcing strong authentication, tightly controlling administrative access, and monitoring system activity. A bastion host commonly provides a narrowly defined function, such as a jump server, proxy, or externally reachable application service. Its defining characteristic is not a particular network device role or packet-processing capability; rather, it is the elevated security posture and restricted configuration appropriate to a system that may be accessed from less-trusted networks. NIST guidance describes system hardening as securing systems by minimizing unnecessary functionality and applying configuration controls, practices central to operating a bastion host securely. ISC2 SSCP candidates should recognize that this hardened design is what enables the host to serve as a controlled, defensible point of access across trust boundaries. See NIST SP 800-123, Guide to General Server Security and CISA guidance on securing network infrastructure devices.

QUESTION NO: 35

What works as an E-mail message transfer agent?

A. SMTP

B. SNMP

C. S-RPC

D. S/MIME

ANSWER: A

Explanation:

SMTP is correct because the Simple Mail Transfer Protocol is the standard application-layer protocol used to transfer email messages between mail servers. When a user submits an email, the sending mail system uses SMTP to relay the message toward the recipient's mail server, potentially through one or more intermediary mail transfer agents. SMTP therefore provides the message submission and server-to-server transfer function central to email delivery. Modern SMTP operation is defined in RFC 5321, which describes how an SMTP client establishes a connection to an SMTP server and transmits mail using commands such as MAIL, RCPT, and DATA. In practical enterprise environments, mail transfer agents implement SMTP for outbound relay and inbound receipt of email, commonly applying controls such as authentication, TLS encryption, anti-spam filtering, and routing policies as part of their processing. The distinction is important: an MTA is the software role that accepts, relays, and delivers messages, while SMTP is the protocol that enables those message-transfer exchanges. See RFC 5321: Simple Mail Transfer Protocol and the IANA SMTP service registration.

QUESTION NO: 36

Which type of password provides maximum security because a new password is required for each new log-on?

A. One-time or dynamic password

B. Congnitive password

C. Static password

D. Passphrase

ANSWER: A

Explanation:

One-time or dynamic password is correct because its authentication value is valid for only one login or transaction and is replaced with a newly generated value at the next authentication attempt. Consequently, intercepting, observing, or obtaining a previously used value does not ordinarily enable an attacker to replay it successfully for a later logon. One-time passwords can be generated by hardware or software tokens, mobile authenticator applications, or challenge-response mechanisms. They are commonly used as a factor in multifactor authentication, where their short validity period and single-use property materially reduce the risk associated with credential replay. NIST describes one-time passwords as secrets that are valid for only a single authentication event and notes their role in authentication protocols. The CISSP/SSCP authentication concept tested here is specifically the use of a fresh password for every logon, which directly defines a one-time password rather than a reusable password format. See NIST SP 800-63B: Digital Identity Guidelines—Authentication and Lifecycle Management and CISA guidance on implementing phishing-resistant MFA.

QUESTION NO: 37

Which of the following is best at defeating frequency analysis?

A. Substitution cipher

B. Polyalphabetic cipher

C. Transposition cipher

D. Ceasar Cipher

ANSWER: B

Explanation:

Polyalphabetic cipher is correct because it encrypts plaintext characters using more than one substitution alphabet. Rather than preserving one fixed ciphertext equivalent for each plaintext letter throughout the message, the applicable alphabet

changes according to the key or key position. Consequently, a frequently occurring plaintext letter can be represented by several different ciphertext letters. This distributes and obscures the single-letter frequency patterns on which straightforward frequency analysis depends, making it substantially more resistant than a fixed monoalphabetic substitution approach.

In the classic Vigenère implementation, for example, the repeated key selects different Caesar-style substitution alphabets for successive characters. An analyst may still attack a weak or repeating-key polyalphabetic system using sufficient ciphertext and techniques such as identifying key length, so this is not equivalent to modern cryptographic security. Nevertheless, among the listed classical cipher categories, a polyalphabetic cipher is the best answer to defeating simple frequency analysis because its changing substitutions conceal the direct plaintext-to-ciphertext frequency relationship. Modern systems avoid relying on classical substitution designs and instead use vetted cryptographic algorithms with appropriately managed keys. See the historical and technical description of the Vigenère/polyalphabetic method from Encyclopaedia Britannica and NIST guidance on using approved cryptography in NIST's Cryptographic Module Validation Program.

QUESTION NO: 38

Each data packet is assigned the IP address of the sender and the IP address of the:

A. recipient.

B. host.

C. node.

D. network.

ANSWER: A

Explanation:

The correct completion is **recipient.** At the Internet Protocol layer, an IP packet header contains a source IP address and a destination IP address. The source address identifies the host or interface that originated the packet, while the destination address identifies the intended recipient host or interface. Routers examine the destination IP address to decide how to forward the packet toward its final destination. This source-and-destination addressing model enables hosts on separate networks to communicate and permits return traffic or responses to be directed back to the originating address.

Although an IP address can be interpreted using network and host portions—depending on the addressing architecture and subnet mask—the packet itself is addressed end-to-end using its source and destination IP address fields. In current IPv4 terminology, these fields are explicitly named “Source Address” and “Destination Address”; IPv6 likewise includes “Source Address” and “Destination Address” fields in its base header. Therefore, “recipient” accurately describes the entity identified by the packet’s destination IP address. See the IPv4 specification in RFC 791 and the IPv6 specification in RFC 8200.

QUESTION NO: 39

Once evidence is seized, a law enforcement officer should emphasize which of the following?

A. Chain of command

B. Chain of custody

C. Chain of control

D. Chain of communications

ANSWER: B

Explanation:

Chain of custody is correct because it provides the documented, chronological record showing the seizure, custody, control, transfer, analysis, storage, and final disposition of evidence. From the moment an officer collects an item, every person who handles it and every transfer or change in location should be recorded. The record commonly includes

identifying information for the item, the date and time of each action, the individual releasing and receiving it, the purpose of the transfer, and secure storage details.

Maintaining this record establishes that the evidence presented in an investigation or legal proceeding is the same evidence originally seized and that it has remained protected against loss, substitution, alteration, or unauthorized access. This supports the evidence's integrity and authenticity, which are essential for establishing confidence in forensic findings and for demonstrating that proper evidence-handling procedures were followed. The U.S. National Institute of Justice describes chain-of-custody documentation as a core element of preserving evidence integrity. Guidance from NIST likewise stresses documentation and controlled handling to preserve the integrity of digital evidence. See the National Institute of Justice evidence-collection guidance and NIST SP 800-86.

QUESTION NO: 40

Network-based Intrusion Detection systems:

A. Commonly reside on a discrete network segment and monitor the traffic on that network segment.

B. Commonly will not reside on a discrete network segment and monitor the traffic on that network segment.

C. Commonly reside on a discrete network segment and does not monitor the traffic on that network segment.

D. Commonly reside on a host and and monitor the traffic on that specific host.

ANSWER: A

Explanation:

Network-based intrusion detection systems commonly reside at a strategic point on a network segment and monitor traffic traversing that segment. A NIDS sensor obtains packet and flow visibility from a network tap, a switch SPAN/mirror port, or an inline deployment. It then analyzes observed network activity for known attack signatures, protocol anomalies, policy violations, or other indicators of malicious behavior. Because its scope is the traffic visible at its monitoring location, placing sensors near key network boundaries or high-value segments provides useful detection coverage without requiring an agent on every endpoint.

This reflects the distinction between a network-based sensor and a host-based intrusion detection system: the relevant characteristic here is monitoring communications on a network segment rather than activity local to an individual host. NIST describes network-based IDPS technologies as monitoring network traffic for particular network segments or devices, with sensors deployed at appropriate points in the network. ISC2's SSCP domain coverage similarly expects practitioners to understand the placement and monitoring role of network security controls, including intrusion detection capabilities. See NIST SP 800-94, Guide to Intrusion Detection and Prevention Systems and ISC2 SSCP Certification Exam Outline.

QUESTION NO: 41

What is the goal of the Maintenance phase in a common development process of a security policy?

A. to review the document on the specified review date

B. publication within the organization

C. to write a proposal to management that states the objectives of the policy

D. to present the document to an approving body

E. & KRAUSE, MICKI, Information Security Management Handbook, 4th Edition, Volume 3, 2002, Auerbach Publications. Also: KRUTZ, Ronald L. & VINES, Russel D., The CISSP Prep Guide: Mastering the Ten Domains of Computer Security, John Wiley & Sons, 2001, Chapter 8: Business Continuity Planning and Disaster Recovery Planning (page 286).

ANSWER: A

Explanation:

The goal of the Maintenance phase is **to review the document on the specified review date**. A security policy is not a static document: business priorities, technologies, threats, legal obligations, organizational roles, and risk tolerances change over time. Maintenance establishes the recurring process for reassessing whether the approved policy remains accurate, relevant, enforceable, and aligned with the organization's current security objectives. A defined review date makes this responsibility deliberate and auditable rather than relying on an informal decision to revisit the policy only after an incident or major organizational change.

During a scheduled review, policy owners and relevant stakeholders can identify necessary updates, confirm that referenced standards and procedures remain current, and route substantive revisions through the organization's established governance and approval process. This lifecycle approach supports continual policy effectiveness after the policy has been created, approved, and distributed. NIST describes policy as a foundational element of an information security program and emphasizes periodic review and updating as program conditions evolve. See NIST SP 800-12 Rev. 1 and the ISC2 SSCP Exam Outline for the governance and operational-security context in which policy lifecycle management is applied.

QUESTION NO: 42

What does it mean to say that sensitivity labels are "incomparable"?

A. The number of classification in the two labels is different.

B. Neither label contains all the classifications of the other.

C. the number of categories in the two labels are different.

D. Neither label contains all the categories of the other.

ANSWER: D

Explanation:

"Neither label contains all the categories of the other." correctly describes incomparable sensitivity labels in a multilevel security system. An MLS label combines a hierarchical classification, such as Secret or Top Secret, with a set of categories (also called compartments), such as project, country, or mission identifiers. Labels are compared using the dominance relationship: for one label to dominate another, its classification must be at least as high and its category set must include every category in the other label's set.

Labels are incomparable when neither category set is a subset of the other. For example, a label with categories {VENUS} and another with {ALPHA} are incomparable because each has a category absent from the other. This produces a partial ordering rather than a simple linear ranking: some labels can be ordered by dominance, while others cannot. The lattice-based model uses this relationship to enforce mandatory access control decisions, ensuring subjects access only information for which their clearances dominate the relevant labels. NIST describes mandatory access control and sensitivity labels in its Mandatory Access Control glossary entry; the underlying formal lattice model is specified in the DoD Trusted Computer System Evaluation Criteria.

QUESTION NO: 43

What refers to legitimate users accessing networked services that would normally be restricted to them?

A. Spoofing

B. Piggybacking

C. Eavesdropping

D. Logon abuse

ANSWER: D

Explanation:

Logon abuse is the correct term for a situation in which a legitimate, authenticated user accesses networked services, systems, or resources beyond the scope of permissions normally assigned to that user. The user has a valid identity and may be permitted to use the network generally, but abuses that legitimate access by obtaining or using access to restricted services. This can occur through misuse of credentials, inappropriate privilege use, exploitation of a weak access-control configuration, or use of an authenticated session to reach resources outside the user's authorized role.

This concept directly concerns enforcement of authorization after authentication. Security controls must not merely establish who a user is; they must also enforce which objects, services, and actions that identity is allowed to access. NIST describes access enforcement as enforcing approved authorizations for logical access to information and system resources, a principle that addresses the risk represented by logon abuse. This is aligned with the access-controls and security-operations knowledge assessed in the SSCP curriculum. See NIST SP 800-53, Access Control (AC) family and the ISC2 SSCP Examination Outline.

QUESTION NO: 44

Secure Shell (SSH-2) provides all the following services except:

A. secure remote login

B. command execution

C. port forwarding

D. user authentication

ANSWER: D

Explanation:

user authentication is the correct exception when the question is referring to the SSH-2 Transport Layer Protocol. The transport layer establishes the secure foundation for an SSH connection: it performs key exchange, authenticates the SSH server (host) to the client, protects data confidentiality through encryption, and protects packet integrity through message authentication. It deliberately does not authenticate the human user or client account at this layer. RFC 4253 explicitly states that the transport protocol "does not perform user authentication"; instead, a higher-level protocol is designed to run over the protected transport connection for that purpose.

This layered design is central to SSH-2 architecture. After the secure transport channel is established, the SSH User Authentication Protocol requests and verifies the client user's identity using methods such as public-key authentication, passwords, or keyboard-interactive authentication. SSH connection services can then support interactive sessions, remote command execution, and TCP/IP forwarding over that authenticated, encrypted channel. Thus, user authentication is not a service of the SSH transport layer itself, even though it is available within the broader SSH protocol suite through a separate protocol layer. See RFC 4253, The Secure Shell Transport Layer Protocol and RFC 4252, The Secure Shell Authentication Protocol.

QUESTION NO: 45

Which of the following steps should be one of the first step performed in a Business Impact Analysis (BIA)?

A. Identify all CRITICAL business units within the organization.

B. Evaluate the impact of disruptive events.

C. Estimate the Recovery Time Objectives (RTO).

D. Identify and Prioritize Critical Organization Functions

ANSWER: D

Explanation:

Identify and Prioritize Critical Organization Functions is correct because a Business Impact Analysis begins by establishing which business processes, services, and supporting functions are most important to the organization's mission. This initial identification creates the baseline for the rest of the analysis: stakeholders can then determine the consequences if each function is unavailable, define acceptable outage periods, identify dependencies, and establish recovery priorities. A function's priority is determined by the operational, financial, legal, regulatory, contractual, and reputational effects of disruption. This means the BIA is focused first on the business need and the importance of the function, rather than beginning with a technical recovery target. Once critical functions have been identified and ranked, the organization can use the impact information to support recovery time objectives and continuity strategies that reflect actual business requirements. NIST's contingency-planning guidance describes the BIA as identifying and prioritizing system and business-process requirements based on the impact of a disruption. This approach also aligns with ISC2 continuity and resilience principles: recovery planning must be driven by prioritized organizational functions and their dependencies. See NIST SP 800-34 Rev. 1 and ISC2 SSCP Certification Exam Outline.

QUESTION NO: 46

Controls to keep password sniffing attacks from compromising computer systems include which of the following?

A. static and recurring passwords.

B. encryption and recurring passwords.

C. one-time passwords and encryption.

D. static and one-time passwords.

ANSWER: C

Explanation:

One-time passwords and encryption are controls that directly reduce the value and exposure of credentials to a password-sniffing attack. Encryption protects authentication traffic while it crosses a network by preventing a passive listener from reading the password in cleartext. In practice, this means using properly configured cryptographic protocols for remote access, web authentication, and administrative connections, such as TLS and SSH, rather than protocols that transmit credentials without encryption. Encryption must include certificate validation and current secure protocol configurations so that the connection is actually protected against interception.

One-time passwords add a separate and important safeguard: each authentication value is valid for only one use or for a very short period. Therefore, even if an attacker captures that value, replaying it later should not authenticate the attacker. This principle is commonly implemented through time-based codes, hardware tokens, challenge-response mechanisms, or other multifactor authentication methods. Combining encrypted transport with one-time passwords provides defense in depth by reducing the likelihood that credentials can be read and ensuring that captured authentication values are not reusable. NIST discusses replay resistance and phishing-resistant authentication properties in NIST SP 800-63B; CISA also recommends multifactor authentication as a key protection for accounts in its password security guidance.

QUESTION NO: 47

Which xDSL flavour, appropriate for home or small offices, delivers more bandwidth downstream than upstream and over longer distance?

A. VDSL

B. SDSL

C. ADSL

D. HDSL

ANSWER: C

Explanation:

ADSL is correct because it was designed for the typical residential and small-office access pattern, where users generally receive substantially more data than they transmit. Its asymmetric design allocates more of the available copper-loop capacity to downstream traffic while retaining a smaller upstream channel for requests, email, uploads, and other outbound activity. This made ADSL well suited to Internet access over existing telephone subscriber loops without requiring the symmetric capacity commonly needed for business-grade services.

The distance element also fits ADSL in the traditional xDSL comparison: attainable rate declines as the copper loop to the telephone exchange or DSL access multiplexer becomes longer, but ADSL was intended to operate over materially longer loops than very-high-speed DSL technologies. Actual throughput and reach depend on loop length, copper quality, noise, line configuration, and the particular ADSL standard deployed. The ITU's ADSL recommendation defines an asymmetric digital-subscriber-line transceiver system, including distinct downstream and upstream transmission arrangements. Broadband Forum technical material likewise describes ADSL as an asymmetric DSL access technology for delivering broadband over copper telephone lines. See ITU-T Recommendation G.992.1 and the Broadband Forum DSL reference architecture.

QUESTION NO: 48

What would be considered the biggest drawback of Host-based Intrusion Detection systems (HIDS)?

A. It can be very invasive to the host operating system

B. Monitors all processes and activities on the host system only

C. Virtually eliminates limits associated with encryption

D. They have an increased level of visibility and control compared to NIDS

ANSWER: A

Explanation:

It can be very invasive to the host operating system is the best answer because a host-based intrusion detection system operates directly on the endpoint it protects. To detect suspicious activity, a HIDS agent needs broad visibility into host events such as process execution, file and registry changes, local logs, user activity, and configuration changes. This depth of monitoring commonly requires privileged access and ongoing use of system resources, which can increase CPU, memory, disk-I/O, administrative, and compatibility overhead. Deployment also requires installing, configuring, updating, and maintaining an agent on every monitored endpoint. Consequently, organizations must assess the operational impact carefully, especially for performance-sensitive, legacy, or highly regulated systems. The same endpoint-level perspective that enables a HIDS to identify activity unavailable to purely network-based monitoring is therefore also its principal implementation drawback: it interacts closely with the operating system and its workloads. NIST describes host-based IDS technology as monitoring characteristics and events occurring within hosts, including audit logs and operating-system activity; this necessarily depends on software and data collection at the host level. See NIST SP 800-94, Guide to Intrusion Detection and Prevention Systems and the ISC2 SSCP Exam Outline.

QUESTION NO: 49

What is a hot-site facility?

A. A site with pre-installed computers, raised flooring, air conditioning, telecommunications and networking equipment, and UPS.

B. A site in which space is reserved with pre-installed wiring and raised floors.

C. A site with raised flooring, air conditioning, telecommunications, and networking equipment, and UPS.

D. A site with ready made work space with telecommunications equipment, LANs, PCs, and terminals for work groups.

ANSWER: A

Explanation:

A hot-site facility is a fully equipped alternate processing location designed to support rapid recovery of critical business operations after a disruption. It has the necessary environmental and infrastructure components—such as raised flooring, climate control, telecommunications, network connectivity, and uninterruptible power supplies—already in place. Crucially, it also includes pre-installed computer systems or processing equipment, enabling an organization to restore its applications and data and resume processing with minimal delay once current backups and configurations are loaded. This level of readiness distinguishes a hot site as an important business continuity and disaster recovery capability for systems with stringent recovery-time objectives. The facility may be operated by the organization or contracted from a disaster-recovery provider, but its defining characteristic is that the technology environment is operationally ready rather than merely being reserved space or basic infrastructure. NIST describes alternate processing sites as recovery facilities whose capabilities should be selected according to system recovery requirements, including the time needed to become operational. A hot site represents the highest-readiness category in the traditional hot/warm/cold site model. See NIST SP 800-34 Rev. 1, Contingency Planning Guide and the ISC2 SSCP Exam Outline.

QUESTION NO: 50

Which layer of the DoD TCP/IP Model ensures error-free delivery and packet sequencing?

A. Internet layer

B. Network access layer

C. Host-to-host

D. Application layer

ANSWER: C

Explanation:

The **Host-to-host** layer is correct because it supplies end-to-end transport services between communicating hosts in the TCP/IP (DoD) model. Its reliability functions include sequencing transmitted data so that the receiving system can restore it to the proper order, detecting transmission errors, acknowledging successfully received data, and retransmitting data when necessary. These capabilities are principally associated with the Transmission Control Protocol (TCP), which operates at this layer. TCP uses sequence numbers and acknowledgments to provide reliable, ordered delivery of a byte stream, while checksum processing helps detect corruption. Flow-control mechanisms also help ensure that a sender does not overwhelm a receiver. In terminology that maps TCP/IP to the OSI model, the Host-to-host layer corresponds to the OSI Transport layer. Although the TCP/IP suite also supports UDP at this layer, UDP does not provide guaranteed delivery or sequencing; the layer is nevertheless the correct model layer because it is where TCP's reliable end-to-end delivery service resides. See the Internet Engineering Task Force's TCP specification in RFC 9293 and the TCP/IP model overview in RFC 1122.

QUESTION NO: 51

All hosts on an IP network have a logical ID called a(n):

A. IP address.

B. MAC address.

C. TCP address.

D. Datagram address.

ANSWER: A

Explanation:

An IP address is the logical identifier assigned to a network interface on an Internet Protocol network. It provides the network-layer addressing information used to identify the interface and to route packets between networks. In IPv4, this identifier is a 32-bit value, while IPv6 uses a 128-bit value; both support logical addressing that can be configured, assigned dynamically, or changed when a device moves to another network. This differs from a physical link-layer identifier because

IP addressing is designed for internetwork communication and routing. The Internet Engineering Task Force's Internet Protocol specification describes an IP address as an identifier for a source or destination interface and explains its role in delivering datagrams across interconnected networks. For SSCP purposes, recognize the distinction between logical network-layer addressing and hardware/link-layer addressing: an IP address identifies a host or interface logically within IP communications. See the IETF's Internet Protocol (IPv4) specification and Internet Protocol, Version 6 (IPv6) specification.

QUESTION NO: 52

Which of the following backup methods makes a complete backup of every file on the server every time it is run?

A. full backup method.

B. incremental backup method.

C. differential backup method.

D. tape backup method.

ANSWER: A

Explanation:

The correct answer is **full backup method.** A full backup copies all selected data in the backup scope each time the job runs, regardless of whether individual files have changed since an earlier backup. For a server-wide backup scope, this means every included file is backed up during every execution. Full backups provide a self-contained recovery point: restoring data generally requires only the relevant full-backup set, which simplifies restoration and supports dependable recovery planning. Because each run recopies the complete selected dataset, full backups usually require more storage capacity and longer backup windows than change-based approaches. Organizations commonly schedule them periodically and use them as the baseline recovery set within a broader backup strategy. The National Institute of Standards and Technology describes full backups as copying all selected files, and its contingency-planning guidance emphasizes maintaining recoverable backups as part of system restoration capabilities. See NIST's definition of full backup and NIST SP 800-34 Rev. 1, Contingency Planning Guide.

QUESTION NO: 53

A DMZ is located:

A. right behind your first Internet facing firewall

B. right in front of your first Internet facing firewall

C. right behind your first network active firewall

D. right behind your first network passive Internet http firewall

ANSWER: A

Explanation:

The correct placement is **right behind your first Internet facing firewall**. A demilitarized zone is a separate perimeter network segment positioned between an untrusted external network, such as the Internet, and the organization's trusted internal network. In a screened-subnet design, traffic from the Internet must first traverse the Internet-facing (perimeter) firewall before reaching publicly accessible services hosted in the DMZ. A second control boundary commonly separates the DMZ from the internal network.

This placement limits direct exposure of the internal network while allowing carefully controlled access to services that must be reachable externally, such as public web, DNS, mail relay, VPN gateway, or reverse-proxy services. Firewall rules should permit only the specific protocols, ports, source and destination paths required for each service. The DMZ therefore provides network segmentation and containment: compromising a public-facing system should not automatically grant an attacker unrestricted access to internal systems or data.

This architecture aligns with NIST guidance on using screened subnets and boundary protections to isolate public-facing components from internal networks. See NIST SP 800-41 Rev. 1, Guidelines on Firewalls and Firewall Policy and CISA guidance on network security architecture.

QUESTION NO: 54

Which of the following would be true about Static password tokens?

A. The owner identity is authenticated by the token

B. The owner will never be authenticated by the token.

C. The owner will authenticate himself to the system.

D. The token does not authenticates the token owner but the system.

ANSWER: A

Explanation:

The owner identity is authenticated by the token is the intended description of a static password-token system. A static token, such as a card, smart card, or other device assigned to a particular user, serves as a possession-based authenticator: its presentation supplies a fixed credential or otherwise demonstrates that the user possesses the issued token. The organization establishes the association between the token and its owner when it enrolls and issues the token. When the token is then presented to a reader, lock, workstation, or authentication service, the system can use that established binding to authenticate the claimed identity. In practical terms, this is the familiar “something you have” factor—for example, using an assigned access card to obtain entry through a controlled door. A static token differs from a one-time-password token because it does not generate a newly changing value for each authentication event; its credential remains fixed unless it is replaced, revoked, or reconfigured. NIST describes authenticators as evidence controlled by a claimant and recognizes possession as an authentication factor; the relying system verifies the presented authenticator and its binding to the account. See NIST SP 800-63B, Digital Identity Guidelines and CISA guidance on access control systems.

QUESTION NO: 55

Which of the following backup sites is the most effective for disaster recovery?

A. Time brokers

B. Hot sites

C. Cold sites

D. Reciprocal Agreement

ANSWER: B

Explanation:

Hot sites are the most effective choice for disaster recovery because they are pre-equipped, operational alternate processing facilities designed to support rapid restoration of critical business services. A hot site typically includes compatible computing infrastructure, network connectivity, telecommunications, environmental controls, and the software and data arrangements needed to resume processing with minimal delay. Depending on the organization’s architecture and replication design, it can support recovery in hours or less, producing a substantially lower recovery time objective than facilities that must first be equipped or activated. The site must be matched to documented business continuity requirements, including the recovery time objective and recovery point objective for each critical system. A hot-site strategy also supports realistic, repeatable recovery exercises: personnel can validate technical procedures, communications, access, data restoration, and application functionality before an actual disruption occurs. ISC2’s SSCP examination outline identifies business continuity and disaster recovery planning, including recovery strategies and testing, as relevant security-practitioner knowledge. NIST similarly describes alternate processing sites as a contingency-planning recovery strategy and emphasizes selecting solutions that

meet system recovery requirements. See the ISC2 SSCP Certification Exam Outline and NIST SP 800-34 Rev. 1, Contingency Planning Guide.

QUESTION NO: 56

This is a common security issue that is extremely hard to control in large environments. It occurs when a user has more computer rights, permissions, and access than what is required for the tasks the user needs to fulfill. What best describes this scenario?

A. Excessive Rights

B. Excessive Access

C. Excessive Permissions

D. Excessive Privileges

ANSWER: D

Explanation:

Excessive Privileges best describes a situation in which an account has more rights, permissions, or access capabilities than are necessary to perform its assigned job functions. This condition is the inverse of the principle of least privilege: users and processes should receive only the minimum authorization required to complete authorized tasks. In a large environment, role changes, temporary access grants, inherited group membership, and inconsistent account reviews can cause privileges to accumulate over time. Those unneeded capabilities expand the account's potential impact if it is misused, compromised, or used in error.

Privilege is the broadest and most appropriate term because it encompasses the authority granted through rights, permissions, access controls, and elevated administrative capabilities. Controlling excessive privileges requires a defined authorization process, role-based access management where appropriate, periodic entitlement reviews, and prompt removal or adjustment of access when duties change. NIST Special Publication 800-53 defines least privilege as restricting privileged accounts and functions to authorized users and necessary functions, directly supporting this classification. See NIST SP 800-53 Rev. 5 and the related NIST definition of least privilege.

QUESTION NO: 57

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

A. SNMP

B. SMTP

C. HTTP

D. SSH

ANSWER: A

Explanation:

SNMP is the correct choice because it is a network-management protocol intended to monitor and, depending on configuration and access rights, administer managed devices such as routers, switches, servers, and security appliances. Exposing SNMP across an organizational perimeter is generally inappropriate: it can disclose valuable operational information about device configuration, interfaces, routes, and status, while SNMP write access can enable remote configuration changes. A secure firewall policy follows least privilege, permitting only the explicitly required communications to tightly defined destinations. Therefore, SNMP should normally be blocked at the external firewall and limited to authorized management networks, hosts, and secure administrative paths. Where remote management is required, organizations should use SNMPv3 because it supports authentication and privacy protections, and should still restrict access with firewall rules and access-control lists rather than expose the service broadly. The IETF's SNMP architecture describes the

framework and security model in RFC 3411. NIST also recommends that firewall policies be based on documented business requirements and allow only necessary traffic; see NIST SP 800-41 Rev. 1.

QUESTION NO: 58

Which is the last line of defense in a physical security sense?

A. people

B. interior barriers

C. exterior barriers

D. perimeter barriers

E. (2006). CompTIA Security+ study guide: Y0-101. Indianapolis, IN: Sybex.

ANSWER: A

Explanation:

people is the correct answer because physical security is based on defense in depth: protections are arranged in layers to deter, delay, detect, and respond to unauthorized access. Barriers, locks, access-control devices, alarms, and surveillance provide important protective layers, but they do not independently make an informed security decision or take responsive action. People—such as security officers, employees, reception staff, system owners, and incident responders—provide the final safeguard by recognizing suspicious activity, challenging or reporting unauthorized access, following procedures, and escalating incidents.

This principle also includes every authorized individual's responsibility to protect facilities and assets. For example, an employee may notice tailgating through a controlled door, an unattended visitor, or signs of forced entry and notify the appropriate security personnel. Human awareness and response therefore complete the protective function when earlier physical controls have been bypassed, fail, or require intervention. This aligns with NIST's physical and environmental protection control family, which addresses access enforcement, monitoring, visitor controls, and the personnel-driven operation of those safeguards. See NIST's SP 800-53 Physical and Environmental Protection controls and CISA's physical security guidance.

QUESTION NO: 59

In the statement below, fill in the blank:

Law enforcement agencies must get a warrant to search and seize an individual's property, as stated in the _____ Amendment.

A. Fourth

ANSWER: A

Explanation:

Fourth is the correct completion. The Fourth Amendment to the U.S. Constitution protects people against unreasonable searches and seizures by the government. It provides that warrants may be issued only upon probable cause, supported by oath or affirmation, and must particularly describe the place to be searched and the persons or things to be seized. Thus, the amendment is the constitutional foundation for the warrant requirement in law-enforcement investigations and for the requirements that a warrant be authorized by a neutral judicial officer and limited in scope. In information-security and digital-forensics contexts, this principle is important when government investigators seek access to devices, accounts, records, or other property: the authorized search must stay within the warrant's lawful scope. The warrant requirement is subject to established legal exceptions in specific circumstances, but the amendment identified by the statement is still the Fourth Amendment. The amendment's text is available from the Constitution Annotated, Congress.gov, and its application to searches and seizures is summarized by the Legal Information Institute at Cornell Law School.

QUESTION NO: 60

Which conceptual approach to intrusion detection system is the most common?

A. Behavior-based intrusion detection

B. Knowledge-based intrusion detection

C. Statistical anomaly-based intrusion detection

D. Host-based intrusion detection

ANSWER: B

Explanation:

Knowledge-based intrusion detection is correct. This approach, commonly called signature-based or misuse-based detection in current security terminology, compares observed activity against a maintained set of known attack patterns, indicators, rules, or vulnerability-exploitation signatures. It has historically been—and remains—the most broadly deployed conceptual IDS approach because known malicious activity can be recognized consistently and alerts can be tied to identifiable detection content. Security teams can update the signature or rule set as new threats are identified, allowing the detection capability to recognize those newly documented techniques.

NIST describes signature-based IDS technologies as using signatures of known attacks and notes that signature-based detection is effective at identifying previously known threats when the relevant signatures are present and current. This corresponds directly to the knowledge-based model named in the question. In operational environments, network and host detection products frequently incorporate this model through rules, signatures, and threat-intelligence indicators. Keeping the knowledge base current is therefore central to the approach's effectiveness. See NIST's Guide to Intrusion Detection and Prevention Systems (SP 800-94) and the NIST glossary entry for signature-based detection.

QUESTION NO: 61

Which of the following items is NOT a benefit of cold sites?

A. No resource contention with other organisation

B. Quick Recovery

C. A secondary location is available to reconstruct the environment

D. Low Cost

ANSWER: B

Explanation:

"Quick Recovery" is the item that is not a benefit of a cold site. A cold site generally provides a suitable alternate physical location with basic supporting infrastructure, such as space, power, environmental controls, and telecommunications capability. However, it does not normally have the production systems, applications, current data, or fully configured equipment needed to resume operations immediately. Following a disruption, the organization must obtain, transport, install, configure, test, and restore the required technology and information before business processes can be restarted. Those activities make the recovery time comparatively long.

This trade-off is central to disaster-recovery site selection: cold sites are typically less expensive to maintain because the organization is not paying for continuously operational duplicate systems and staff. They are therefore appropriate when the business can tolerate a longer recovery time objective and has documented procedures and resources to build out the recovery environment after an event. NIST describes cold sites as facilities with adequate space and infrastructure to support recovery activities, distinguishing them from more operationally ready warm and hot sites. See NIST SP 800-34 Rev. 1 and its official PDF publication.

QUESTION NO: 62

SMTP can best be described as:

A. a host-to-host email protocol.

B. an email retrieval protocol.

C. a web-based e-mail reading protocol.

D. a standard defining the format of e-mail messages.

ANSWER: A

Explanation:

SMTP is correctly described as “a host-to-host email protocol.” The Simple Mail Transfer Protocol is the Internet standard protocol for submitting and relaying email messages between mail systems. A mail client commonly uses SMTP to submit an outbound message to its configured mail server, and mail transfer agents use SMTP to relay that message onward to the recipient domain’s mail server. Thus, SMTP supports the sending and server-to-server transport side of electronic mail delivery.

The protocol’s core behavior, including the SMTP command-and-response model and message transfer process, is defined in RFC 5321. SMTP is also commonly used with extensions such as SMTP authentication and transport-layer security to support authenticated, protected mail submission and relay. Email message retrieval by recipients is a separate function handled by other protocols, while message-body formatting and attachments are addressed by MIME standards. The distinction is important operationally: SMTP moves mail toward the receiving system, where it can then be accessed through an appropriate user-facing mail service. See RFC 5321: Simple Mail Transfer Protocol and the IETF RFC standards index.

QUESTION NO: 63

When a station communicates on the network for the first time, which of the following protocol would search for and find the Internet Protocol (IP) address that matches with a known Ethernet address?

A. Address Resolution Protocol (ARP).

B. Reverse Address Resolution Protocol (RARP).

C. Internet Control Message protocol (ICMP).

D. User Datagram Protocol (UDP).

ANSWER: B

Explanation:

Reverse Address Resolution Protocol (RARP) is correct because its purpose is to let a host that knows its own hardware address, such as an Ethernet MAC address, obtain the Internet Protocol address assigned to that address. A newly booted diskless workstation was a classic use case: it could transmit a RARP request containing its MAC address on the local network, and a RARP server with the appropriate address mapping could return the corresponding IP address. This is precisely the reverse of the normal address-resolution direction: starting with a known link-layer address and requesting the associated network-layer address.

RARP is defined in RFC 903, which describes a mechanism for dynamically determining a protocol address from a hardware address. The standard also notes that RARP requests are broadcast and that the service is generally limited to the local network segment. In modern deployments, RARP has largely been replaced by more capable bootstrap and address-configuration technologies, particularly BOOTP and DHCP, because they can provide additional configuration information such as gateway and name-server settings. Nevertheless, for the protocol function stated in this question, Reverse Address Resolution Protocol (RARP) is the accurate answer. See also the DHCP specification (RFC 2131) for the later configuration approach.

QUESTION NO: 64

What is the goal of the Maintenance phase in a common development process of a security policy?

A. to review the document on the specified review date

B. publication within the organization

C. to write a proposal to management that states the objectives of the policy

D. to present the document to an approving body

ANSWER: A

Explanation:

The goal of the Maintenance phase is **"to review the document on the specified review date."** Security policies must remain current as the organization's business objectives, legal and regulatory obligations, information systems, threats, and risk posture change. A defined review date establishes an accountable, repeatable policy-lifecycle activity rather than leaving policy updates to occur only after an incident or an audit finding. During maintenance, policy owners and relevant stakeholders reassess whether the policy still reflects organizational requirements, identify necessary revisions, obtain approval for material changes when required, and communicate the updated version under document-control practices. This ongoing review helps ensure that personnel are operating under guidance that is applicable, authorized, and aligned with the organization's current security governance framework. NIST describes policy as a foundational element of an information-security program and notes that policies should be reviewed and updated to remain effective as organizational needs evolve. Likewise, ISO/IEC 27001 guidance emphasizes maintaining documented information and keeping the information security management system suitable and current. See NIST SP 800-12 Rev. 1 and ISO/IEC 27001.

QUESTION NO: 65

What kind of certificate is used to validate a user identity?

A. Public key certificate

B. Attribute certificate

C. Root certificate

D. Code signing certificate

ANSWER: A

Explanation:

Public key certificate is correct because its core purpose is to bind a subject's verified identity to that subject's public key. In a public key infrastructure, a certificate authority validates identity information at an assurance level appropriate to the certificate's intended use, then digitally signs the certificate. A relying party can validate the certificate authority's signature and certification path, confirm that the certificate is within its validity period and has not been revoked, and then trust that the contained public key is associated with the identified user or entity. This supports identity validation in operations such as secure email, client authentication, document signing, and establishing authenticated encrypted sessions.

The certificate commonly follows the X.509 profile and includes a subject name or subject alternative name, the public key, issuer information, serial number, validity dates, and the issuing authority's digital signature. NIST describes public-key certificates as mechanisms that bind public keys to named entities and support authentication through certificate-path validation. See NIST SP 800-32: Introduction to Public Key Technology and the Federal PKI Infrastructure and the IETF X.509 Public Key Infrastructure Certificate and CRL Profile (RFC 5280).

QUESTION NO: 66

Which of the following is defined as a key establishment protocol based on the Diffie-Hellman algorithm proposed for IPsec but superseded by IKE?

A. Diffie-Hellman Key Exchange Protocol

B. Internet Security Association and Key Management Protocol (ISAKMP)

C. Simple Key-management for Internet Protocols (SKIP)

D. OAKLEY

ANSWER: D

Explanation:

OAKLEY is correct. It is a key-establishment protocol built on Diffie-Hellman key exchange and was designed to provide negotiable security properties for establishing authenticated shared keys. Its defined exchange modes address services such as authentication, identity protection, and perfect forward secrecy. OAKLEY was designed to work with the Internet Security Association and Key Management Protocol framework, which supplies mechanisms for negotiating and managing security associations rather than prescribing one specific key-exchange algorithm. The Internet Key Exchange protocol subsequently combined components and concepts from OAKLEY, SKEME, and ISAKMP into the key-management approach used by IPsec. Therefore, in the historical IPsec terminology described by the question, OAKLEY is the named Diffie-Hellman-based key-establishment protocol that was incorporated into and superseded operationally by IKE. RFC 2409 documents IKE and explicitly identifies its use of the Oakley key-determination protocol together with ISAKMP and SKEME-related methods. The Internet Security Glossary also defines OAKLEY as a key-establishment protocol based on Diffie-Hellman, proposed for IPsec and superseded by IKE. See RFC 2409: The Internet Key Exchange (IKE) and RFC 2828: Internet Security Glossary.

QUESTION NO: 67

Which virus category has the capability of changing its own code, making it harder to detect by anti-virus software?

A. Stealth viruses

B. Polymorphic viruses

C. Trojan horses

D. Logic bombs

ANSWER: B

Explanation:

Polymorphic viruses is the correct answer. A polymorphic virus alters its identifiable code representation as it replicates or infects additional files, commonly by changing encryption keys, using different decryptor routines, or otherwise varying its signature while preserving its malicious function. This variation produces many distinct instances of the same malware, making simple signature-based detection less reliable because a scanner cannot depend on one fixed byte pattern. Security products therefore use additional techniques, such as heuristic analysis, behavioral detection, emulation, and machine-learning-assisted analysis, to identify suspicious activity even when the malware's static appearance changes. The term reflects an important malware-evasion concept: the code's form can change from one generation to the next without changing the underlying payload or objective. MITRE describes polymorphic code as a defense-evasion technique in which malicious code changes form while retaining functionality; see MITRE ATT&CK: Obfuscated Files or Information. For additional context on malware detection and protective technologies, consult the CISA malware guidance.

QUESTION NO: 68

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:

“It operates on 64-bit plaintext blocks and uses a 128 bit key.” is not a property of Rijndael. Rijndael, the algorithm selected by NIST as the basis for the Advanced Encryption Standard, uses a 128-bit block size when implemented as AES. AES supports key sizes of 128, 192, and 256 bits. The broader Rijndael design also permits block and key sizes of 128, 160, 192, 224, or 256 bits, in combinations defined by its design; it is therefore not a 64-bit-block cipher.

The stated 64-bit plaintext block size paired with a 128-bit key identifies the International Data Encryption Algorithm (IDEA), rather than Rijndael. IDEA is a symmetric block cipher that processes 64-bit data blocks and uses a 128-bit key. Distinguishing a cipher’s block size from its key size is important in security architecture and exam questions: these are separate algorithm parameters, and Rijndael/AES is specifically standardized around 128-bit blocks. NIST’s AES standard specifies the AES block size and its permitted key lengths, while the original Rijndael submission describes the wider Rijndael parameter set. See NIST FIPS 197: Advanced Encryption Standard and NIST Block Cipher Techniques.

QUESTION NO: 69

What would BEST define a covert channel?

A. An undocumented backdoor that has been left by a programmer in an operating system

B. An open system port that should be closed.

C. A communication channel that allows transfer of information in a manner that violates the system's security policy.

D. A trojan horse.

ANSWER: C

Explanation:

A communication channel that allows transfer of information in a manner that violates the system's security policy is the best definition of a covert channel. A covert channel is an unintended or unauthorized communication path that can be used to move information in a way not permitted or mediated by the system’s intended security controls. It may exploit shared system resources, such as processor time, memory allocation, file locks, or observable variations in system behavior, to encode and transmit data between subjects that should not be able to communicate under the security policy.

This concept is important in systems enforcing confidentiality policies, particularly multilevel-security environments, because the normal access-control mechanisms may correctly prevent direct reads or writes while an indirect signaling path still permits information flow. The U.S. National Institute of Standards and Technology defines a covert channel as a communications channel that is not intended for information transfer and can be used to violate a security policy. ISC2 SSCP candidates should recognize that the defining feature is the unauthorized policy-violating information flow, rather than a particular malware type, network service, or implementation flaw. See NIST’s covert-channel glossary entry and NIST Trusted Systems Technology Guide.

QUESTION NO: 70

Which of the following ports does NOT normally need to be open for a mail server to operate?

A. Port 110

B. Port 25

C. Port 119

D. Port 143

ANSWER: C

Explanation:

Port 119 is the correct answer because it is the well-known TCP port assigned to the Network News Transfer Protocol (NNTP), a protocol for distributing and retrieving Usenet news articles. NNTP is separate from standard Internet email services and is not a normal functional requirement for a mail server. A mail server's required exposed services depend on its role and configuration, but email transport is typically provided through SMTP, while mailbox retrieval services may be provided through POP3 or IMAP for end users. NNTP neither transfers email between mail systems nor provides normal email mailbox access, so opening it would not be necessary merely to operate a mail server. In a secure deployment, services and firewall rules should follow least functionality: enable and expose only the ports required by the organization's actual mail architecture. Keeping unrelated services such as NNTP closed reduces unnecessary attack surface and simplifies monitoring. The Internet Assigned Numbers Authority lists port 119 as the registered NNTP service port and separately records the standard mail-related assignments. See the IANA service-name and port-number registry for port 119 and the NNTP specification (RFC 3977).

QUESTION NO: 71

A potential problem related to the physical installation of the Iris Scanner in regards to the usage of the iris pattern within a biometric system is:

A. concern that the laser beam may cause eye damage

B. the iris pattern changes as a person grows older.

C. there is a relatively high rate of false accepts.

D. the optical unit must be positioned so that the sun does not shine into the aperture.

ANSWER: D

Explanation:

The optical unit must be positioned so that the sun does not shine into the aperture. Iris-recognition systems capture detailed images of the iris using a camera, typically with near-infrared illumination to obtain consistent iris texture while minimizing effects from visible-light conditions. Direct sunlight entering the camera aperture can create glare, reflections, excessive illumination, and reduced contrast. These conditions can obscure iris features or prevent the device from acquiring an image of sufficient quality for reliable template generation and matching. Therefore, installation location is a practical security and usability consideration: the scanner should be placed and aimed to avoid direct sun exposure, with lighting controlled as appropriate for the device and its capture environment.

This concern is specifically tied to physical installation because it depends on the scanner's orientation, nearby windows or exterior entrances, changing sun angles, and the availability of shielding or controlled illumination. NIST's biometric-capture guidance addresses the need for controlled acquisition conditions and quality assessment so that captured biometric samples are suitable for recognition. The NIST Biometric Specifications for Personal Identity Verification describes iris-image capture requirements, and NIST's Iris Recognition Vendor Test provides further context on iris-recognition performance and image-quality-dependent evaluation.

QUESTION NO: 72

Which of the following is an advantage in using a bottom-up versus a top-down approach to software testing?

A. Interface errors are detected earlier.

B. Errors in critical modules are detected earlier.

C. Confidence in the system is achieved earlier.

D. Major functions and processing are tested earlier.

ANSWER: B

Explanation:

“Errors in critical modules are detected earlier.” is the advantage associated with a bottom-up integration-testing approach. Bottom-up testing starts by testing the lowest-level units or modules, then progressively integrates them into larger subsystems and, ultimately, the complete application. Because foundational modules are exercised first, defects in core processing logic, data handling, calculations, and reusable services can be identified and corrected before higher-level components depend on them. This can reduce the likelihood that a defect in a low-level component will obscure results or cause widespread failures during later integration stages.

In practice, the approach is particularly valuable when lower-level modules contain technically complex or business-critical processing. Test teams can establish confidence in those building blocks before integrating user-facing workflows and broader system behavior. The ISTQB describes integration testing as focusing on interactions among components or systems and recognizes that integration may proceed incrementally in different orders. Bottom-up integration specifically builds upward from lower-level components, making early validation of those components a defining benefit. See the ISTQB glossary entry for bottom-up testing and the ISTQB Certified Tester Foundation Level syllabus resources.

QUESTION NO: 73

What is the Maximum Tolerable Downtime (MTD)?

A. Maximum elapsed time required to complete recovery of application data

B. Minimum elapsed time required to complete recovery of application data

C. Maximum elapsed time required to move back to primary site after a major disruption

D. It is maximum delay businesses can tolerate and still remain viable

ANSWER: D

Explanation:

“It is maximum delay businesses can tolerate and still remain viable” correctly describes Maximum Tolerable Downtime (MTD). MTD is the total maximum period that a mission or business process may be disrupted before the resulting impact becomes unacceptable to the organization. It is a business-driven limit: process owners and leadership determine it by considering operational, financial, legal, regulatory, reputational, safety, and customer-service consequences of an outage. Once that duration is exceeded, the organization may no longer be able to sustain the effects of the interruption or meet its essential obligations.

MTD provides a critical boundary for continuity and recovery planning. Recovery strategies, technical capabilities, staffing, alternate facilities, and recovery procedures must be designed so that the affected service is restored—and any necessary follow-on processing can be completed—within the established tolerable period. In practice, recovery time objectives are set to support meeting this business limit, while recovery point objectives address the amount of data loss that can be accepted. NIST describes MTD as the total amount of time an owner or authorizing official is willing to accept for a mission or business-process disruption, including all relevant impacts. See NIST SP 800-34 Rev. 1 and ISC2’s SSCP Exam Outline.

QUESTION NO: 74

Configuration Management controls what?

A. Auditing of changes to the Trusted Computing Base.

B. Control of changes to the Trusted Computing Base.

C. Changes in the configuration access to the Trusted Computing Base.

D. Auditing and controlling any changes to the Trusted Computing Base.

ANSWER: D

Explanation:

Auditing and controlling any changes to the Trusted Computing Base. is correct because configuration management provides a disciplined process for maintaining the security-relevant configuration of a system throughout its life cycle. In the context of a Trusted Computing Base, this means proposed modifications are identified, evaluated, authorized, implemented in a controlled manner, documented, and subsequently reviewed or audited. These activities preserve the integrity of the security functions that enforce the system's security policy and provide accountability for what changed, when it changed, and who approved it.

NIST describes configuration management as managing and controlling changes to information-system configurations, including processes for configuration control and monitoring. Security-focused configuration management also requires documenting baselines and assessing changes so that the system remains in its intended, secure state. Auditing is therefore an essential companion to change control: it supplies the records and verification needed to confirm that approved changes were made as intended and that the Trusted Computing Base continues to operate with its required protections. This combined control-and-audit approach supports traceability, repeatability, and assurance for security-critical system components. See NIST SP 800-128, Guide for Security-Focused Configuration Management of Information Systems and the NIST configuration management glossary entry.

QUESTION NO: 75

Which of the following concerning the Rijndael block cipher algorithm is false?

A. The design of Rijndael was strongly influenced by the design of the block cipher Square.

B. A total of 25 combinations of key length and block length are possible

C. Both block size and key length can be extended to multiples of 64 bits.

D. The cipher has a variable block length and key length.

ANSWER: C

Explanation:

"Both block size and key length can be extended to multiples of 64 bits." is the false statement. In the Rijndael specification, block length and cipher-key length are independent parameters. Each may be selected in 32-bit increments, from 128 bits through 256 bits. Thus, the supported Rijndael values are 128, 160, 192, 224, and 256 bits, rather than only values that are multiples of 64 bits. This distinction is significant because 160-bit and 224-bit block or key lengths are permitted by Rijndael even though they are not multiples of 64.

The original Rijndael design is broader than the AES standard derived from it. AES standardizes only a 128-bit block size and permits 128-, 192-, and 256-bit keys; however, that AES restriction must not be confused with Rijndael's variable-size design. The Rijndael submission documentation explicitly describes the allowable 32-bit granularity for these parameters, while NIST's AES standard documents the fixed AES profile. See the Rijndael amended submission specification and NIST FIPS 197.

QUESTION NO: 76

What is used to protect programs from all unauthorized modification or executional interference?

A. A protection domain

B. A security perimeter

C. Security labels

D. Abstraction

ANSWER: A

Explanation:

A protection domain is correct because it defines the security boundaries and access rights within which a process or program executes. It associates an executing subject with a permitted set of objects and authorized operations, such as reading, writing, executing, or modifying resources. By constraining a program to only the memory, instructions, files, devices, and system services for which it has authorization, the protection domain helps prevent unauthorized changes to the program and prevents improper interference with its execution. This reflects the foundational security-architecture principle of isolation: separately executing processes should be restricted so that one process cannot arbitrarily alter another process's code, data, or execution state. Protection domains can be implemented through hardware and operating-system mechanisms, including processor privilege levels, virtual-memory protection, process isolation, and access-control enforcement. In an SSCP context, this is an example of applying access-control and least-privilege concepts to operating-system processes and resources. NIST describes access control as limiting system access to authorized users, processes acting on behalf of users, and devices, as well as limiting permitted transactions and functions. ISC2's SSCP outline also includes applying security concepts and controls to protect systems and their resources. See NIST's Access Control glossary and the ISC2 SSCP Exam Outline.

QUESTION NO: 77

Which of the following statements pertaining to key management is incorrect?

A. The more a key is used, the shorter its lifetime should be.

B. When not using the full keyspace, the key should be extremely random.

C. Keys should be backed up or escrowed in case of emergencies.

D. A key's lifetime should correspond with the sensitivity of the data it is protecting.

ANSWER: B

Explanation:

"When not using the full keyspace, the key should be extremely random." is the incorrect statement because a cryptographic key must be generated across the full set of values permitted by its defined keyspace, with each possible key value having an appropriately equal likelihood of selection. Randomness is necessary, but it does not compensate for limiting the range of possible values. Restricting generation to only part of the keyspace reduces the effective entropy of the key and makes exhaustive-search attacks more practical, because an attacker has fewer candidate values to test. Sound key management therefore requires an approved random or deterministic random-bit generator, sufficient entropy, and a generation process that produces keys at the required algorithm strength without introducing bias or predictable restrictions. In practical terms, a key that is randomly selected from a reduced subset is weaker than a key of the same nominal length selected uniformly from the full keyspace. NIST key-management guidance emphasizes using approved methods to generate cryptographic keying material with the security strength required for the protected information and associated cryptographic mechanism. See NIST SP 800-57 Part 1 Rev. 5 and NIST SP 800-133 Rev. 2.

QUESTION NO: 78

During which phase of an IT system life cycle are security requirements developed?

A. Operation

B. Initiation

C. Functional design analysis and Planning

D. Implementation

ANSWER: B

Explanation:

Initiation is correct because this is the life-cycle phase in which the organization establishes the system's purpose, scope, business need, high-level architecture context, stakeholders, and initial risk posture. These inputs enable security requirements to be developed as part of the system requirements baseline. Requirements should cover the protection needs of the information and services the system will provide, including requirements driven by organizational policy, applicable law, regulation, privacy obligations, and the system's required confidentiality, integrity, and availability objectives.

NIST Special Publication 800-64 Rev. 2 explicitly identifies defining security requirements as an initiation-phase security activity. Establishing them at this point ensures security is built into planning, budgeting, acquisition, architecture, and later engineering work rather than added after technical decisions have already been made. The requirements are subsequently refined, analyzed, designed, implemented, tested, and maintained through the remaining life-cycle activities, but their initial development belongs to initiation. This early definition supports a risk-based approach and provides the criteria against which security controls and system acceptance can later be evaluated. See NIST SP 800-64 Rev. 2 and NIST's Risk Management Framework resources.

QUESTION NO: 79

Which of the following is a CHARACTERISTIC of a decision support system (DSS) in regards to Threats and Risks Analysis?

A. DSS is aimed at solving highly structured problems.

B. DSS emphasizes flexibility in the decision making approach of users.

C. DSS supports only structured decision-making tasks.

D. DSS combines the use of models with non-traditional data access and retrieval functions.

ANSWER: B

Explanation:

"DSS emphasizes flexibility in the decision making approach of users." is correct because a decision support system is designed to help people analyze information, apply models, and exercise judgment when making decisions that are not fully routine or predetermined. In threat and risk analysis, decision makers commonly need to evaluate uncertain likelihoods, potential business impacts, alternative treatments, and changing assumptions. A DSS supports this work by allowing users to explore data, compare scenarios, adjust model inputs, and use analytical techniques without forcing every decision into a fixed processing sequence.

This flexibility is particularly valuable for semi-structured decisions: the system can provide relevant data, calculations, visualizations, and model outputs, while the analyst or risk owner supplies professional judgment about risk acceptance, prioritization, and treatment. IBM describes decision support systems as tools that support organizational decision-making through data, documents, knowledge, and models, particularly where human judgment remains important. NIST's risk-management guidance likewise frames risk assessment as an activity involving the evaluation of threats, vulnerabilities, likelihood, and impact, all of which benefit from adaptable analytical support. See IBM: What is a decision support system? and NIST SP 800-30 Rev. 1: Guide for Conducting Risk Assessments.

QUESTION NO: 80

A common way to create fault tolerance with leased lines is to group several T1s together with an inverse multiplexer placed:

A. at one end of the connection.

B. at both ends of the connection.

C. somewhere between both end points.

D. in the middle of the connection.

ANSWER: B

Explanation:

The correct answer is **"at both ends of the connection."** Inverse multiplexing combines the capacity of multiple lower-speed circuits, such as T1 leased lines, into one logical connection. Unlike conventional multiplexing, which aggregates several inputs onto a single higher-capacity path, inverse multiplexing distributes traffic across multiple physical links at the transmitting endpoint. A corresponding inverse multiplexer at the receiving endpoint accepts the traffic from those links, compensates for differences in delay where necessary, and reassembles the data stream in its proper sequence.

Deploying the equipment at both endpoints is essential because both sides must participate in the link-bonding process: one side divides outbound traffic among the available T1s, while the other side recombines the received traffic. This arrangement also provides resilience. If an individual T1 fails, the inverse-multiplexing group can generally continue carrying traffic over the remaining circuits, although with reduced aggregate bandwidth. This makes grouped leased lines a practical historical approach to improving availability without relying on one larger circuit. The IETF describes inverse multiplexing as distributing data over multiple links and reconstructing it at the remote endpoint in RFC 2686. Additional technical background on T1 digital transmission is available from ETSI TS 101 318.

QUESTION NO: 81

Which of the following would NOT violate the Due Diligence concept?

A. Security policy being outdated

B. Data owners not laying out the foundation of data protection

C. Network administrator not taking mandatory two-week vacation as planned

D. Latest security patches for servers being installed as per the Patch Management process

ANSWER: D

Explanation:

"Latest security patches for servers being installed as per the Patch Management process

" would not violate due diligence because it demonstrates that the organization has identified vulnerability and patching risk, established a defined process to manage that risk, and is operating the process in practice. Due diligence requires an organization to make a continuing, reasonable effort to understand relevant threats, assess risk, and ensure that appropriate security safeguards remain effective. A documented patch-management process, coupled with installation of current server patches according to that process, provides evidence of this ongoing attention and governance.

Patch management is not merely a one-time technical action. It includes identifying affected assets, evaluating patches, prioritizing remediation according to risk, deploying updates in a controlled manner, and confirming completion. Performing these activities consistently helps reduce exposure to known vulnerabilities and shows that the organization is taking reasonable steps to maintain the security of its systems. NIST describes enterprise patch management as a process for identifying, acquiring, testing, deploying, and monitoring patches, with the goal of reducing vulnerabilities. See NIST SP 800-40 Rev. 4, Guide to Enterprise Patch Management Planning. This ongoing, risk-informed maintenance aligns with the governance and care expected of an SSCP practitioner; see the ISC2 SSCP Exam Outline.

QUESTION NO: 82

In Discretionary Access Control the subject has authority, within certain limitations,

A. but he is not permitted to specify what objects can be accessible and so we need to get an independent third party to specify what objects can be accessible.

B. to specify what objects can be accessible.

C. to specify on a aggregate basis without understanding what objects can be accessible.

D. to specify in full detail what objects can be accessible.

ANSWER: B

Explanation:

In discretionary access control, the owner of an object, or a subject acting with the owner's delegated authority, can determine which other subjects may access that object and what permissions they receive. Therefore, "to specify what objects can be accessible" correctly completes the statement. This discretion is commonly implemented through object-specific permissions, such as access control lists or file and folder sharing permissions. For example, a file owner may grant another user permission to read, modify, or execute a file, subject to organizational policy and the technical constraints enforced by the operating system or application.

The phrase "within certain limitations" is important: discretionary control does not mean unlimited authority. System administrators, security policies, privilege boundaries, and mandatory protections can constrain an owner's ability to grant access. NIST describes discretionary access control as a means of restricting access based on the identity of subjects and/or the groups to which they belong, with access decisions often governed by permissions associated with objects. ISC2's SSCP domain coverage also includes applying access-control concepts and mechanisms, including discretionary models. See NIST's definition of discretionary access control and the ISC2 SSCP Exam Outline.

QUESTION NO: 83

When we encrypt or decrypt data there is a basic operation involving ones and zeros where they are compared in a process that looks something like this:

0101 0001 Plain text

0111 0011 Key stream 0010 0010 Output

What is this cryptographic operation called?

A. Exclusive-OR

B. Bit Swapping

C. Logical-NOR

D. Decryption

ANSWER: A

Explanation:

Exclusive-OR is correct because the displayed bitwise calculation produces the stated output when each plaintext bit is compared with the corresponding keystream bit. Exclusive-OR outputs zero when its two inputs are identical and one when they differ. Applying that rule to 0101 0001 and 0111 0011 produces 0010 0010: the first pair, 0 and 0, yields 0; the next pair, 1 and 1, yields 0; and differing pairs yield 1. This operation is fundamental to stream-cipher encryption, where plaintext is combined with a generated keystream to form ciphertext.

A key security property makes Exclusive-OR particularly useful in this setting: it is self-inverting. Applying the identical keystream to the ciphertext restores the original plaintext, because a bit combined with the same bit twice returns to its initial value. Formally, plaintext XOR keystream equals ciphertext, and ciphertext XOR the same keystream equals plaintext. This is the construction used by common stream-cipher designs; for example, the ChaCha20 specification describes encryption as XORing a keystream with plaintext. See the NIST glossary definition of Exclusive-OR and RFC 8439, ChaCha20 encryption operation.

QUESTION NO: 84

Access Control techniques do not include which of the following choices?

A. Relevant Access Controls

B. Discretionary Access Control

C. Mandatory Access Control

D. Lattice Based Access Control

ANSWER: A

Explanation:

“Relevant Access Controls” is correct because it is not a recognized access-control model or technique. In security architecture, access control is the set of mechanisms and policies used to identify subjects, authenticate them where required, authorize permitted actions on objects or resources, and enforce those decisions. Established access-control approaches have defined rules for determining access rights, such as owner-directed permissions, centrally imposed classification rules, or access assignments based on job functions. “Relevant Access Controls,” by contrast, is only a vague descriptive phrase. It does not identify a defined authorization model, a decision method, or an enforcement technique that could be implemented and assessed as part of an access-control system. For an SSCP exam question, the expected distinction is between formally named access-control approaches and terminology that does not represent a standard model. ISC2 describes access controls as measures that protect information systems and their assets by managing and enforcing authorized access. This requires policy-driven mechanisms with identifiable criteria and rules, not merely controls described as relevant. See the ISC2 SSCP Exam Outline and NIST’s Access Control glossary definition.

QUESTION NO: 85

How would nonrepudiation be best classified as?

A. A preventive control

B. A logical control

C. A corrective control

D. A compensating control

ANSWER: A

Explanation:

Nonrepudiation is best classified as **A preventive control**. It provides evidence that binds an action or transaction to its originator and, where applicable, to its recipient. For example, a properly implemented digital signature can demonstrate that a particular private key signed a message, while signed receipts or transaction logs can establish that a recipient received it. This evidence makes it substantially more difficult for a party to credibly deny having performed an action after the fact.

In functional control classifications, a preventive control is intended to stop an undesirable event or policy violation from succeeding. Nonrepudiation serves that purpose by preventing successful repudiation of a transaction: the sender cannot plausibly deny sending the information, and the receiver cannot plausibly deny receiving it when the supporting evidence is retained and protected. This directly supports accountability, since individuals can be held responsible for actions traceable to them. NIST describes nonrepudiation as a security service supporting accountability and explains that it provides assurance that parties cannot successfully deny actions such as sending or receiving information. See NIST SP 800-33 and the NIST CSRC definition of nonrepudiation.

QUESTION NO: 86

Which layer of the TCP/IP protocol stack corresponds to the ISO/OSI Network layer (layer 3)?

A. Host-to-host layer

B. Internet layer

C. Network access layer

D. Session layer

ANSWER: B

Explanation:

The Internet layer is correct because it performs the logical addressing and packet-routing functions that correspond to the ISO/OSI Network layer, which is layer 3. In the TCP/IP model, the Internet layer uses Internet Protocol (IP) to provide a connectionless, best-effort method for delivering packets across interconnected networks. IP addresses identify source and destination hosts or networks, while routers use routing information to forward packets toward their intended destinations. These functions align directly with the OSI Network layer's responsibility for moving packets between networks rather than merely delivering frames over a local link. Common protocols associated with this layer include IPv4, IPv6, ICMP, and IPsec. The TCP/IP architecture does not map every OSI layer one-to-one; instead, it groups related OSI functions into broader layers. Nevertheless, the Internet layer is the established TCP/IP equivalent for OSI layer 3. The Internet Engineering Task Force describes IP as the protocol that enables the interconnection of networks and the transmission of datagrams across them in RFC 8200. A concise TCP/IP model overview is also available from Cloudflare's OSI model reference.

QUESTION NO: 87

In biometric identification systems, the parts of the body conveniently available for identification are:

A. neck and mouth

B. hands, face, and eyes

C. feet and hair

D. voice and neck

ANSWER: B

Explanation:

hands, face, and eyes is correct because these biometric characteristics are readily accessible during ordinary identity-verification interactions, without requiring a person to remove typical clothing or use an intrusive collection process. Hands support fingerprint, palm-print, hand-geometry, and vein-pattern recognition. The face supports facial-recognition systems, which can acquire an image with a camera. Eyes provide highly distinctive iris and retinal characteristics, with iris recognition in particular commonly used where high confidence is required. These modalities are based on physiological traits that can be measured and compared against a previously enrolled biometric reference template. Their practical accessibility is a major reason they are frequently considered in biometric-system design, alongside operational factors such as accuracy, presentation-attack resistance, user acceptance, privacy, and the security of stored biometric data. NIST's digital identity guidance addresses biometric comparison and the need to protect biometric information and associated processes. The UK National Cyber Security Centre also describes fingerprints, facial images, and iris patterns as examples of biometric data used to identify or authenticate individuals. See NIST SP 800-63B and NCSC biometrics guidance.

QUESTION NO: 88

What is called the act of a user professing an identity to a system, usually in the form of a log-on ID?

A. Authentication

B. Identification

C. Authorization

D. Confidentiality

ANSWER: B

Explanation:

Identification is correct because it is the act of asserting or claiming an identity to an information system. In a typical sign-in process, a person supplies an identifier, such as a username, account name, email address, or log-on ID, to tell the system which identity is being claimed. At this stage, the user has not yet demonstrated that the claim is genuine; the identifier simply allows the system to associate the request with a particular account or digital identity.

This distinction is fundamental to access control and identity and access management. After identification, the system can request evidence that validates the asserted identity, such as a password, cryptographic authenticator, hardware token, or biometric factor. NIST's digital identity guidance describes authentication as a process in which a claimant proves control of one or more authenticators associated with an identity, which necessarily follows the user's identity assertion. Maintaining this conceptual separation helps security practitioners design clear sign-in workflows, account-management controls, audit records, and access decisions. See NIST SP 800-63B: Authentication and Authenticator Management and the NIST CSRC glossary entry for identification.

QUESTION NO: 89

Which of the following statements pertaining to Secure Sockets Layer (SSL) is false?

A. The SSL protocol was developed by Netscape to secure Internet client-server transactions.

B. The SSL protocol's primary use is to authenticate the client to the server using public key cryptography and digital certificates.

C. Web pages using the SSL protocol start with HTTPS

D. SSL can be used with applications such as Telnet, FTP and email protocols.

ANSWER: B

Explanation:

The statement "The SSL protocol's primary use is to authenticate the client to the server using public key cryptography and digital certificates." is false. In the usual SSL/TLS deployment model, the server presents an X.509 digital certificate during the handshake. The client validates that certificate, including its trust chain and the server name, to gain assurance that it is communicating with the intended server. This server authentication is especially important for HTTPS because it helps protect users from connecting to an impersonating site.

SSL/TLS can support client-certificate authentication, in which the server requests and validates a certificate from the client. However, this is optional and is not the protocol's primary or most common authentication use. Most public HTTPS services authenticate the server to the client while clients authenticate through application-level methods such as passwords, multifactor authentication, or tokens. SSL itself has been deprecated because of known security weaknesses, and modern secure deployments use TLS; nevertheless, the described authentication model remains the relevant conceptual distinction for this question. The TLS handshake and certificate-based server authentication process are described in RFC 8446, and NIST's guidance on TLS configuration is available in NIST SP 800-52 Rev. 2.

QUESTION NO: 90

Why does fiber optic communication technology have significant security advantage over other transmission technology?

A. Higher data rates can be transmitted.

B. Interception of data traffic is more difficult.

C. Traffic analysis is prevented by multiplexing.

D. Single and double-bit errors are correctable.

ANSWER: B

Explanation:

“Interception of data traffic is more difficult.” is correct because optical fiber carries information as light within a glass or plastic core rather than as electrical signals propagating through a conductive medium. This makes passive interception substantially harder: an attacker generally must physically access the cable and bend, split, or otherwise alter the fiber to extract enough light for monitoring. Such activity requires specialized equipment and skills and can degrade the optical signal, making tampering more likely to be detected through link monitoring or loss-of-signal conditions.

Fiber also provides very limited electromagnetic radiation compared with copper cabling. Consequently, it offers an important protection against electromagnetic eavesdropping, in which an attacker attempts to capture information from emissions without directly connecting to the transmission path. These characteristics make fiber a strong choice where confidentiality of data in transit is important, although it should not be treated as impossible to tap; physical protections, encryption, monitoring, and sound cable-management practices remain necessary.

NIST’s guidance on physical and environmental security recognizes the importance of protecting telecommunications cabling and transmission paths: NIST SP 800-53, PE-4: Access Control for Transmission. The U.S. National Security Agency also describes fiber-optic links as a means of reducing compromising electromagnetic emissions: NSA TEMPEST Fundamentals.

QUESTION NO: 91

Who should direct short-term recovery actions immediately following a disaster?

A. Chief Information Officer.

B. Chief Operating Officer.

C. Disaster Recovery Manager.

D. Chief Executive Officer.

ANSWER: C

Explanation:

Disaster Recovery Manager. is the appropriate role to direct short-term recovery immediately after a disaster. This role is responsible for operationally leading the disaster-recovery effort: initiating the recovery plan, coordinating assigned recovery teams, establishing the immediate recovery priorities, tracking status, escalating material issues, and communicating progress through the organization’s defined governance structure. The person directing these first actions needs detailed familiarity with the disaster recovery plan, its recovery procedures, dependencies, alternate arrangements, and the personnel who perform technical and business recovery tasks.

Effective recovery planning distinguishes executive sponsorship and authorization from hands-on command of recovery activities. Senior executives provide authority, resources, strategic direction, and risk-based business decisions, while the designated recovery manager coordinates the plan’s execution at the time-sensitive operational level. This assignment should be documented in advance, including delegation procedures in case the primary recovery manager is unavailable. NIST contingency-planning guidance emphasizes defined recovery teams, roles, responsibilities, and coordination for restoring systems and services after a disruption. See NIST SP 800-34 Rev. 1, Contingency Planning Guide for Federal Information Systems and CISA Business Resilience.

QUESTION NO: 92

Which of the following remote access authentication systems is the most robust?

A. TACACS+

B. RADIUS

C. PAP

D. TACACS

ANSWER: A

Explanation:

TACACS+ is the best answer because it provides a comprehensive AAA framework—authentication, authorization, and accounting—designed for centralized control of administrative access to network devices. Its architecture separates these functions, allowing organizations to apply granular authorization policies after a user has authenticated. This is especially valuable for privileged remote access because permissions can be tailored to a user, role, device, or even individual command. TACACS+ also encrypts the body of protocol packets, protecting authentication information and authorization details exchanged between the network access device and the AAA server. The protocol supports detailed accounting records, which strengthens auditability and helps security teams investigate administrative activity. These features make TACACS+ particularly well suited to robust, centrally managed access control for routers, switches, firewalls, and similar infrastructure. The IETF's TACACS+ specification describes the protocol's authentication, authorization, accounting, packet protection, and command-authorization capabilities: RFC 8907: The Terminal Access Controller Access-Control System Plus (TACACS+) Protocol. Cisco also documents TACACS+ as a mechanism for centralized device-administration access and policy enforcement: Cisco TACACS+ configuration and deployment guidance.

QUESTION NO: 93

Which port does the Post Office Protocol Version 3 (POP3) make use of?

A. 110

B. 109

C. 139

D. 119

ANSWER: A

Explanation:

POP3 uses TCP port 110 for its standard, unencrypted client-to-server mail-retrieval service. POP3 is an application-layer protocol that lets a mail client connect to a mail server and download messages, commonly for local storage and offline reading. The Internet Assigned Numbers Authority (IANA) registers the service name *pop3* with transport protocol TCP and port number 110. For secure POP3 using implicit TLS/SSL—often called POP3S—the commonly assigned TCP port is 995; however, that secure port is not what this question asks for. Therefore, *110* is the correct port for standard POP3. Port assignments are important operationally because firewall rules, intrusion-detection signatures, service inventories, and network troubleshooting frequently use them to identify and control mail-related traffic. Administrators should also recognize that a service can be configured on a nonstandard port, but the well-known/default POP3 assignment remains TCP 110. See the IANA Service Name and Transport Protocol Port Number Registry at IANA: POP3 service registration and RFC 1939, the POP3 protocol specification, at RFC 1939.

QUESTION NO: 94

Which security model is based on the military classification of data and people with clearances?

A. Brewer-Nash model

B. Clark-Wilson model

C. Bell-LaPadula model

D. Biba model

ANSWER: C

Explanation:

The Bell-LaPadula model is correct because it is a formal confidentiality model designed around the kind of hierarchical security classifications used in military and government environments. It assigns security labels to subjects, such as users or

processes, and to objects, such as files or records. A subject's clearance and an object's classification determine whether access is permitted. Typical labels map to ordered levels such as Unclassified, Confidential, Secret, and Top Secret, potentially combined with compartments or categories.

Its central objective is preventing unauthorized disclosure of classified information. The model expresses this through its simple-security property, commonly summarized as "no read up," and its star property, commonly summarized as "no write down." Together, these rules ensure that a person cannot read information above their authorized clearance and cannot write sensitive information into a lower-classified destination. This makes Bell-LaPadula especially appropriate where protecting confidentiality across classification levels is the primary security requirement. The model is widely recognized in security architecture education as a foundational mandatory access control model. See the NIST Bell-LaPadula Model glossary entry and CISA's Introduction to Computer Security.

QUESTION NO: 95

A business application can tolerate the loss of no more than four hours of transaction data following a disaster. This requirement defines the application's:

A. Recovery point objective

B. Recovery time objective

C. Maximum tolerable downtime

D. Mean time between failures

ANSWER: A

Explanation:

Recovery point objective is correct. The recovery point objective (RPO) identifies the maximum acceptable amount of data loss measured in time. If the RPO is four hours, recovery arrangements must ensure that restored data is no more than four hours behind the point of disruption.

The recovery time objective is different: it specifies the maximum acceptable time required to restore a service after an interruption. Backup frequency, replication design, and log retention are selected in part to meet the established RPO. NIST discusses recovery objectives and contingency-planning considerations in SP 800-34 Rev. 1, Contingency Planning Guide for Federal Information Systems.

QUESTION NO: 96

You work in a police department forensics lab where you examine computers for evidence of crimes. Your work is vital to the success of the prosecution of criminals.

One day you receive a laptop and are part of a two man team responsible for examining it together. However, it is lunch time and after receiving the laptop you leave it on your desk and you both head out to lunch.

What critical step in forensic evidence have you forgotten?

A. Chain of custody

B. Locking the laptop in your desk

C. Making a disk image for examination

D. Cracking the admin password with chntpw

ANSWER: A

Explanation:

Chain of custody is correct because the laptop is potential evidence and must remain under controlled, documented protection immediately upon receipt. Leaving it unattended on a desk creates an unrecorded period in which unauthorized access, alteration, damage, removal, or even an allegation of such activity could occur. The issue is not merely whether anyone actually changed the device; investigators must be able to demonstrate that the evidence was continuously safeguarded and that its integrity can be trusted from collection through examination and eventual presentation in court.

A complete chain of custody records the identity of each person who receives, transfers, transports, stores, or examines the item, along with dates, times, locations, and the purpose of each transfer. It also requires appropriate secure storage whenever the evidence is not actively being handled. For digital evidence, this control is particularly important because electronic data can be modified or destroyed without visible physical signs. By promptly documenting receipt and placing the laptop in approved secured evidence storage before leaving, the team preserves accountability and supports the evidence's authenticity and admissibility.

NIST's guidance on integrating forensic techniques emphasizes maintaining the integrity of collected evidence and documenting handling activities; see NIST SP 800-86. The U.S. National Institute of Justice also addresses evidence handling in its Electronic Crime Scene Investigation guide.

QUESTION NO: 97

Which of the following is NOT a task normally performed by a Computer Incident Response Team (CIRT)?

A. Develop an information security policy.

B. Coordinate the distribution of information pertaining to the incident to the appropriate parties.

C. Mitigate risk to the enterprise.

D. Assemble teams to investigate the potential vulnerabilities.

ANSWER: A

Explanation:

Develop an information security policy is the task not normally performed by a Computer Incident Response Team. Enterprise security policy is a governance document that establishes management's direction, expectations, authorities, and organization-wide requirements for protecting information and systems. Senior leadership and the organization's security governance function normally approve and maintain this policy, because it drives risk decisions, resource allocation, compliance obligations, and accountability across the enterprise.

A CIRT operates primarily as an incident-handling capability. Its core mission is to prepare for, detect, analyze, coordinate, contain, eradicate, and recover from cybersecurity incidents, while ensuring that relevant stakeholders receive timely incident information. The team may provide valuable operational input to policy authors, particularly when lessons learned identify gaps in incident-response requirements, but drafting or owning the overarching corporate information security policy is not its normal responsibility.

NIST's incident-response guidance distinguishes incident-response activities from the broader policy and governance structure that authorizes and supports them. See NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide and the NIST Cybersecurity Framework for the relationship between organizational governance and incident-response operations.

QUESTION NO: 98

Which of the following best describes the purpose of a demilitarized zone (DMZ)?

A. To isolate public-facing services from the internal network

B. To replace multifactor authentication for remote users

C. To store backup media outside the primary facility

D. To provide encryption for wireless traffic

ANSWER: A

Explanation:

To isolate public-facing services from the internal network is correct. A DMZ is a separately controlled network segment used for systems that must communicate with untrusted networks, such as web servers, mail gateways, or public DNS services. Firewall rules restrict traffic between the Internet, the DMZ, and the internal network.

This segmentation limits exposure if a public-facing host is compromised. A compromised DMZ host should not automatically provide unrestricted access to internal resources because additional policy enforcement points separate the zones. NIST describes use of DMZs and firewall architectures in SP 800-41 Rev. 1, Guidelines on Firewalls and Firewall Policy.

QUESTION NO: 99

The communications products and services, which ensure that the various components of a network (such as devices, protocols, and access methods) work together refers to:

A. Netware Architecture.

B. Network Architecture.

C. WAN Architecture.

D. Multiprotocol Architecture.

ANSWER: B

Explanation:

Network Architecture is correct because it is the overarching design framework that defines how a network's components and services operate together. It encompasses the arrangement and interaction of network devices, communication protocols, transmission methods, access technologies, addressing, and the services that enable reliable communication. Rather than describing only one network type or one vendor technology, network architecture provides the coordinated structure that allows heterogeneous components to interoperate and deliver communications services. This aligns with the layered approach used in networking: protocols and functions are organized so that systems can communicate predictably despite differences in underlying hardware or access mechanisms. A well-defined architecture also supports scalable design, consistent implementation, secure operation, and troubleshooting by establishing how network functions relate to one another. The ISO/IEC Open Systems Interconnection reference model is a foundational example of a network architecture, describing layered communication functions and their relationships. See the ISO overview of the OSI Basic Reference Model and the IETF's Requirements for Internet Hosts—Communication Layers for an example of how layered network communication functions are specified.

QUESTION NO: 100

Crackers today are MOST often motivated by their desire to:

A. Help the community in securing their networks.

B. Seeing how far their skills will take them.

C. Getting recognition for their actions.

D. Gaining Money or Financial Gains.

ANSWER: D

Explanation:

Gaining Money or Financial Gains. is the best answer because financially motivated crime is the dominant driver of contemporary malicious cyber activity. Criminal actors use unauthorized access to obtain direct payments, steal payment-

card or banking data, deploy ransomware, commit business-email compromise, sell stolen credentials, or monetize personal and business information through fraud and extortion. These activities make system compromise a means of generating revenue rather than merely a technical challenge.

This focus aligns with current threat reporting. The FBI's Internet Crime Complaint Center reports multibillion-dollar annual victim losses across cyber-enabled fraud schemes, including investment fraud, business email compromise, and ransomware-related crimes. Likewise, the Verizon Data Breach Investigations Report consistently identifies financial gain as the leading motivation in analyzed breaches. For an SSCP practitioner, recognizing this primary motive supports risk-based security decisions: controls should prioritize protecting accounts, payment processes, sensitive data, backups, and recovery capabilities against the techniques used to monetize compromise.

See the FBI IC3 2023 Internet Crime Report and the Verizon Data Breach Investigations Report.

QUESTION NO: 101

Related to information security, integrity is the opposite of which of the following?

A. abstraction

B. alteration

C. accreditation

D. application

ANSWER: B

Explanation:

"alteration" is correct. In information security, integrity is the property that information remains accurate, complete, and protected from unauthorized or improper modification. Therefore, the condition integrity is intended to prevent is alteration: a change to data, whether accidental, unauthorized, or malicious, that compromises its trustworthiness. Maintaining integrity includes using controls such as access restrictions, change-management procedures, hashes, digital signatures, audit logging, and validation mechanisms. These measures provide confidence that data has not changed unexpectedly and that authorized changes can be identified and verified. The NIST definition of integrity expressly describes it as guarding against improper information modification or destruction, which directly establishes alteration as the opposing concept in this question. This meaning applies to data at rest, in processing, and in transit; for example, an integrity check can detect whether a transmitted file differs from the original file. SSCP security practice treats integrity as a core security objective alongside confidentiality and availability, with emphasis on preserving the reliability and correctness of systems and information. See NIST's Integrity glossary definition and FIPS 199.

QUESTION NO: 102

Which of the following packets should NOT be dropped at a firewall protecting an organization's internal network?

A. Inbound packets with Source Routing option set

B. Router information exchange protocols

C. Inbound packets with an internal address as the source IP address

D. Outbound packets with an external destination IP address

ANSWER: D

Explanation:

Outbound packets with an external destination IP address should not be dropped when they are legitimate traffic originating from the organization's internal network and allowed by its egress policy. This is the normal direction for users and systems accessing Internet-hosted services: a packet leaves an internal interface, uses an authorized internal source address, and is addressed to a routable external destination. A firewall should inspect this traffic and enforce the organization's policy—for

example, permitting required web, DNS, VPN, or business-service connections while logging or blocking prohibited destinations and ports—but external destination addressing alone is not a reason to discard it.

This principle reflects the distinction between ingress anti-spoofing and normal outbound forwarding. At a network boundary, filtering rules should validate that traffic is plausible for its direction and interface. The Internet Engineering Task Force's ingress-filtering guidance describes filtering packets whose source addresses would be invalid for the interface on which they arrive; it does not prescribe blocking properly addressed outbound communications merely because their destinations are outside the organization. See RFC 2827: Network Ingress Filtering. For practical firewall policy concepts, including controlling outbound connections according to documented policy, see NIST SP 800-41 Rev. 1, Guidelines on Firewalls and Firewall Policy.

QUESTION NO: 103

In which of the following security models is the subject's clearance compared to the object's classification such that specific rules can be applied to control how the subject-to-object interactions take place?

A. Bell-LaPadula model

B. Biba model

C. Access Matrix model

D. Take-Grant model

ANSWER: A

Explanation:

The Bell-LaPadula model is correct because it is a formal multilevel security model that bases access decisions on the relationship between a subject's security clearance and an object's classification. It was designed to preserve confidentiality in environments where information exists at multiple sensitivity levels, such as Unclassified, Confidential, Secret, and Top Secret. A subject may access an object only when the model's mandatory-access-control rules permit the interaction according to those labels and any applicable categories or compartments.

Bell-LaPadula expresses these protections through well-known properties, including the simple-security property ("no read up"), which prevents a subject from reading data classified above the subject's clearance, and the *-property ("no write down"), which prevents a subject from writing sensitive information to a lower classification level. These rules help prevent unauthorized disclosure as information is accessed and handled across a multilevel system. The model is therefore directly characterized by comparing subject clearances with object classifications and applying formal rules to govern subject-to-object access. ISC2's SSCP examination outline includes access-control models and mandatory access control among the access-controls knowledge areas: ISC2 SSCP Certification Exam Outline. A detailed description of Bell-LaPadula's confidentiality properties is also available from NIST CSRC.

QUESTION NO: 104

A server cluster looks like a:

***A.** single server from the user's point of view*

***B.** dual server from the user's point of view*

***C.** triple server from the user's point of view*

***D.** quardle server from the user's point of view*

ANSWER: A

Explanation:

A server cluster is designed to present services through a unified access point, so it appears to users and client systems as a single server. Although the cluster contains multiple physical or virtual nodes behind the scenes, clustering software

coordinates those nodes to provide a common service identity, often using a virtual hostname or IP address. Clients connect to that shared identity rather than selecting an individual cluster member. This abstraction supports high availability: if the active node fails, another eligible node can take over the workload while clients continue using the same service endpoint. In some architectures, multiple nodes can also process requests concurrently for scalability, but the client-facing service is still managed as one logical system. Therefore, "single server from the user's point of view" accurately describes the usual operational model of a server cluster. Microsoft's overview of failover clustering explains that clustered servers work together to increase availability of applications and services, while Red Hat describes clusters as computers working together as a single system. See Microsoft Learn: Failover Clustering overview and Red Hat: What is high availability?.

QUESTION NO: 105

What is the primary role of cross certification?

***A.** Creating trust between different PKIs*

***B.** Build an overall PKI hierarchy*

***C.** set up direct trust to a second root CA*

***D.** Prevent the nullification of user certificates by CA certificate revocation*

ANSWER: A

Explanation:

Creating trust between different PKIs is correct because cross-certification enables separate certification authorities, and therefore their separate public key infrastructures, to recognize and validate certificates issued under one another's trust domains. In a conventional hierarchy, relying parties build certificate paths to a shared trust anchor. Independent organizations or business units may instead operate distinct root CAs with no common superior CA. Cross-certificates address that interoperability need: one CA issues a certificate for the other CA's public key, establishing a cryptographic trust relationship that can be used during certification-path construction. This permits a relying party in one PKI to obtain and validate an appropriate path for a certificate issued in the partner PKI, subject to the policies, name constraints, key-usage restrictions, and trust decisions configured by the participating organizations. Cross-certification is therefore principally a federation mechanism between otherwise independent PKIs, supporting secure communication and certificate validation across organizational boundaries without requiring both PKIs to be reorganized beneath a single root. RFC 5280 defines the X.509 certificate and path-validation framework used for these relationships, while NIST's PKI guidance discusses cross-certification as a way to establish trust among CA domains. See RFC 5280 and NIST SP 800-32.

QUESTION NO: 106

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 correct because a lattice-based model organizes security labels into an ordered structure and makes authorization decisions according to the relationship between a subject's clearance and an object's classification or sensitivity label. The lattice provides concepts such as dominance, least upper bounds, and greatest lower bounds, allowing the system to consistently evaluate whether a subject may access information at a particular security level and within specified categories. These decisions are centrally enforced by policy rather than being freely changed by the object owner.

This is the classic structure associated with mandatory controls in multilevel-security environments. A subject generally needs sufficient clearance and an appropriate need-to-know category relationship to read an object, while write restrictions can protect against inappropriate information flow between levels. The defining feature is that labels and authorization rules are imposed and administered by the system or security authority. NIST describes mandatory access control as a policy-based approach in which access is controlled using security labels and authorizations: NIST CSRC, Mandatory Access Control. The SSCP examination outline includes access controls as a core knowledge area: ISC2 SSCP Certification Exam Outline.

QUESTION NO: 107

An Architecture where there are more than two execution domains or privilege levels is called:

***A.** Ring Architecture.*

***B.** Ring Layering*

***C.** Network Environment.*

***D.** Security Models*

ANSWER: A

Explanation:

***Ring Architecture.** is correct because it describes a protection architecture in which the processor and operating system enforce multiple hierarchical privilege levels, commonly called protection rings. Each execution domain has a defined authority level: software executing in a more privileged domain can perform sensitive operations and access protected resources, while software in a less privileged domain must use controlled mechanisms to request those services. This separation supports least privilege, limits the impact of faults or compromise in user-mode software, and helps prevent untrusted code from directly manipulating kernel-managed memory, hardware, or security-critical functions.*

In the commonly cited x86 implementation, privilege levels are numbered from 0 through 3, with level 0 being the most privileged. Operating systems commonly place the kernel in the most privileged domain and applications in the least privileged domain, using system calls and other controlled transitions when an application needs an operating-system service. The broader security principle is that distinct execution domains establish boundaries that can be enforced by hardware and the operating system. Intel documents these hierarchical privilege levels and their access-control role in its system-programming guidance: Intel Software Developer's Manual. ISC2's SSCP certification overview also identifies security architecture and access-control concepts as relevant practitioner knowledge: ISC2 SSCP.

QUESTION NO: 108

What can best be defined as the sum of protection mechanisms inside the computer, including hardware, firmware and software?

***A.** Trusted system*

***B.** Security kernel*

***C.** Trusted computing base*

***D.** Security perimeter*

ANSWER: C

Explanation:

***Trusted computing base** is the correct term for the totality of protection mechanisms within a computer system that are responsible for enforcing its security policy. It encompasses the relevant hardware, firmware, software, and procedural components that must operate correctly to provide the system's security assurances. These components form the trusted portion of the system because they are relied upon to correctly implement access-control and other security-relevant rules.*

In a trusted-system architecture, the trusted computing base establishes the foundation on which security depends. Its scope can include processor and memory-protection features, operating-system code, security services, and firmware involved in enforcing policy. A smaller, carefully designed subset of this base may implement the reference-monitor concept, but the trusted computing base is the broader term describing the complete collection of trusted protection mechanisms. This definition is consistent with the Common Criteria description of a TCB as the set of components responsible for enforcing a security policy and with NIST terminology for trusted computing capabilities.

References: NIST CSRC: Trusted Computing Base; Common Criteria, Part 1.

QUESTION NO: 109

Which of the following is BEST defined as a physical control?

***A.** Monitoring of system activity*

***B.** Fencing*

***C.** Identification and authentication methods*

***D.** Logical access control mechanisms*

ANSWER: B

Explanation:

Fencing is a physical control because it is a tangible, on-site safeguard that establishes a perimeter around a facility or protected area. Its primary purpose is to deter, delay, and restrict unauthorized physical entry, helping protect personnel, buildings, equipment, and other assets before an individual can reach information systems or sensitive resources. Physical controls are implemented through environmental or facility-based measures, such as barriers, locks, guards, gates, lighting, cameras, and secure doors. In a layered security program, fencing is commonly part of the outermost layer of defense and may be combined with controlled gates, security staff, signage, and surveillance to support effective perimeter protection. This classification aligns with the NIST concept of physical and environmental protection controls, which address limiting physical access to systems, equipment, and facilities to authorized individuals. ISC2's SSCP domain coverage also includes applying physical security controls as part of protecting organizational assets and operations. See NIST Special Publication 800-53 Rev. 5, Physical and Environmental Protection control family: https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final, and the ISC2 SSCP certification outline: https://www.isc2.org/certifications/sscp/sscp-certification-exam-outline.

QUESTION NO: 110

Which of the following is an Internet IPsec protocol to negotiate, establish, modify, and delete security associations, and to exchange key generation and authentication data, independent of the details of any specific key generation technique, key establishment protocol, encryption algorithm, or authentication mechanism?

***A.** OAKLEY*

***B.** Internet Security Association and Key Management Protocol (ISAKMP)*

***C.** Simple Key-management for Internet Protocols (SKIP)*

***D.** IPsec Key exchange (IKE)*

ANSWER: B

Explanation:

Internet Security Association and Key Management Protocol (ISAKMP) is correct because the wording in the question describes ISAKMP's defined role almost verbatim. ISAKMP supplies a general framework and message format for security-association management: it can negotiate, establish, modify, and delete security associations, while also carrying information needed for key generation and authentication. Its key characteristic is separation of the management framework from the particular cryptographic method used. In other words, ISAKMP defines the mechanics and payload structure through which

peers manage security associations and exchange key-management information, without requiring one specific key-exchange algorithm, cipher, or authentication mechanism.

This architectural separation allowed ISAKMP to support different underlying approaches to key establishment and authentication. Historically, IKE combined ISAKMP's framework with key-exchange concepts drawn from protocols such as Oakley and SKEME to create a practical IPsec key-management protocol. Therefore, although IKE is widely associated with IPsec security-association negotiation in operational deployments, the broad, technique-independent protocol description used by this question specifically identifies ISAKMP. The definitive historical specification is RFC 2408, Internet Security Association and Key Management Protocol. The terminology also appears in the Internet Security Glossary (RFC 2828).

QUESTION NO: 111

The IP header contains a protocol field. If this field contains the value of 6, what type of data is contained within the ip datagram?

***A.** TCP.*

***B.** ICMP.*

***C.** UDP.*

***D.** IGMP.*

ANSWER: A

Explanation:

*The correct answer is **TCP.** In an IPv4 packet, the Protocol field identifies the next-layer protocol carried in the IP payload, allowing the receiving host to deliver the encapsulated data to the appropriate transport or network-layer protocol handler. The assigned decimal value **6** identifies the Transmission Control Protocol (TCP). Thus, an IP datagram with Protocol set to 6 contains a TCP segment as its payload.*

TCP is a connection-oriented transport protocol that provides reliable, ordered delivery of an application data stream. When a system receives such a datagram, it processes the IP header, recognizes protocol number 6, and passes the payload to TCP for such functions as validating the TCP header, tracking sequence and acknowledgment numbers, and delivering data to the relevant application socket. The Internet Assigned Numbers Authority maintains the authoritative protocol-number registry, where decimal protocol number 6 is assigned to TCP. RFC 791 defines the IPv4 Protocol field as identifying the protocol used in the data portion of the internet datagram. See the IANA Protocol Numbers registry and RFC 791.