PassQueen
a.k.a dumpsqueen.com

Certified Cloud Security Professional (CCSP)

ISC2 CCSP

Version Demo

Total Demo Questions: 89

Total Premium Questions: 891

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Concepts, Architecture and Design	170
Topic 2, Cloud Data Security	172
Topic 3, Cloud Platform & Infrastructure Security	171
Topic 4, Cloud Application Security	101
Topic 5, Cloud Security Operations	100
Topic 6, Legal, Risk and Compliance	176
Topic 7, Mix Questions	1
Total	891

QUESTION NO: 1

Security best practices in a virtualized network environment would include which of the following? Response:

A. Using distinct ports and port groups for various VLANs on a virtual switch rather than running them through thesame port

B. Running iSCSI traffic unencrypted in order to have it observed and monitored by NIDS

C. Adding HIDS to all virtual guests

D. Hardening all outward-facing firewalls in order to make them resistant to attack

ANSWER: A

Explanation:

Using distinct ports and port groups for various VLANs on a virtual switch rather than running them through thesame port is the best answer because virtual-network security depends on enforcing clear logical boundaries between workloads with different trust levels. A virtual-switch port group provides a policy attachment point for virtual NICs, allowing the administrator to apply the appropriate VLAN assignment and consistent network controls to every connected guest. Separating VLANs into distinct port groups supports segmentation, reduces opportunities for unintended cross-zone communication, and makes policy review and traffic troubleshooting more reliable. It also helps ensure that security-sensitive virtual-switch settings can be applied in a controlled and repeatable manner to the relevant group of workloads. This reflects the cloud-security principle of establishing explicit trust boundaries and applying least privilege at the virtualization layer, rather than treating all guest traffic as though it belongs to one shared network context. NIST guidance for secure virtualization specifically identifies virtual networking and virtual-switch configuration as important security considerations and calls for appropriate isolation between virtual machines and networks. See NIST SP 800-125, Guide to Security for Full Virtualization Technologies and NIST SP 800-207, Zero Trust Architecture.

QUESTION NO: 2

Data centers have enormous power resources that are distributed and consumed throughout the entire facility.

Which of the following standards pertains to the proper fire safety standards within that scope?

A. IDCA

B. BICSI

C. NFPA

D. Uptime Institute

ANSWER: C

Explanation:

NFPA is correct because the National Fire Protection Association develops the core consensus codes and standards used to manage fire and life-safety risks in facilities, including data centers with substantial electrical generation, distribution, and consumption systems. Its publications address the risks arising from electrical equipment and associated fire-protection needs, supporting safe facility design, installation, operation, inspection, and maintenance. Particularly relevant examples include NFPA 70, the National Electrical Code, which establishes electrical-installation requirements intended to reduce fire and shock hazards; NFPA 72, which covers fire alarm and signaling systems; and NFPA 75, which addresses fire protection of information technology equipment. Data-center owners and operators use these standards, as adopted or referenced by the applicable authority having jurisdiction, to establish controls for fire detection, notification, suppression, emergency power-related hazards, and life safety. This makes NFPA the authoritative source in the choices for proper fire-safety standards in an environment with extensive power infrastructure. See the NFPA Codes and Standards catalog and NFPA's overview of fire safety research and resources.

QUESTION NO: 3

SOX was enacted because of which of the following? Response:

A. Poor BOD oversight

B. Lack of independent audits

C. Poor financial controls

D. All of the above

ANSWER: D

Explanation:

All of the above is correct. The Sarbanes-Oxley Act of 2002 was enacted following major accounting and corporate-governance failures, including those associated with Enron and WorldCom. These events revealed interconnected deficiencies in board-level oversight, the effectiveness of internal controls supporting financial reporting, and the independence and quality of external audit work. Congress designed SOX as a broad reform statute addressing these related conditions rather than as a remedy for just one isolated control weakness.

SOX strengthened accountability for financial reporting by requiring senior executives to certify reports, increasing expectations for audit committees and board oversight, and requiring management to assess internal control over financial reporting. For applicable issuers, Section 404 also requires an external auditor's attestation on management's assessment. The law additionally established the Public Company Accounting Oversight Board to oversee audits of public companies, reinforcing audit quality, professional standards, and auditor independence. Collectively, these measures directly reflect the governance, audit, and financial-control failures exposed by the scandals that prompted the legislation. The SEC provides a statutory and regulatory overview at SEC: Sarbanes-Oxley Act, and the PCAOB explains its SOX-created oversight role at PCAOB: About.

QUESTION NO: 4

Which of the following standards primarily pertains to cabling designs and setups in a data center?

A. IDCA

B. BICSI

C. NFPA

D. Uptime Institute

ANSWER: B

Explanation:

BICSI is correct because it develops standards and best-practice guidance for information and communications technology infrastructure, with a central emphasis on structured cabling systems and the spaces, pathways, administration, and installation practices that support them. For a data center, this includes designing copper and fiber cabling infrastructure to be scalable, manageable, and resilient; planning routes and distribution areas; maintaining separation and coordination with other building systems; and documenting and labeling connections so that moves, additions, and changes can be performed safely and reliably. These concerns directly affect cloud-service availability and operational maintainability because poor physical-layer design can increase the likelihood and impact of faults or human error. BICSI 002, *Data Center Design and Implementation Best Practices*, specifically addresses data-center infrastructure design and implementation and is therefore the closest fit for a question focused on cabling designs and setups. BICSI's broader standards portfolio also covers telecommunications distribution methods and associated infrastructure requirements used by designers and installers. See the BICSI Standards overview and BICSI's Data Center education and best-practices resources.

QUESTION NO: 5

The cloud deployment model that features organizational ownership of the hardware and infrastructure, and usage only by members of that organization, is known as:

Response:

A. Private

B. Public

C. Hybrid

D. Motive

ANSWER: A

Explanation:

Private is the correct cloud deployment model because it provides cloud infrastructure for the exclusive use of a single organization. Its defining characteristic is that the organization's users, business units, or approved entities are the only consumers of that environment. This exclusivity enables the organization to apply governance, security controls, architecture standards, and operational processes that are tailored to its own business, risk, compliance, and data-protection requirements.

The scenario's references to organizational ownership of hardware and infrastructure and access limited to organizational members are consistent with a private-cloud implementation. In practice, a private cloud may be owned, managed, and operated by the organization, by a third party, or through some combination of the two. It may also reside on premises or off premises. The essential point is that the infrastructure is provisioned for the exclusive use of one organization, rather than being made broadly available to unrelated customers.

NIST defines private cloud as cloud infrastructure provisioned for exclusive use by a single organization, including multiple consumers within that organization. This deployment model is commonly chosen where strong control, isolation, and compliance alignment are required. See NIST SP 800-145 and the ISC2 CCSP Exam Outline.

QUESTION NO: 6

Data labels could include all the following, except: Response:

A. Confidentiality level

B. Distribution limitations

C. Access restrictions

D. Multifactor authentication

ANSWER: D

Explanation:

Multifactor authentication is the correct exception because a data label is metadata that identifies a data asset's sensitivity and communicates the rules that govern its handling. Labels can express a confidentiality level, such as public, internal, or confidential; distribution limitations, such as restrictions on external sharing; and access restrictions that indicate who is authorized to use the information. This metadata supports consistent information governance and enables tools to apply or guide protective actions based on the classification assigned to the data.

Multifactor authentication, by contrast, is an identity and access-management authentication control. It verifies a user's identity by requiring more than one authentication factor before access is granted. An organization may use an access policy that requires multifactor authentication before a user can access information with a particular label, but MFA is enforced by an identity provider or access-control system rather than being a characteristic recorded in the data label itself. Microsoft describes sensitivity labels as persistent metadata that can classify and protect content, including through encryption and content-marking settings: Microsoft Purview sensitivity labels. NIST's digital identity guidance separately addresses authentication and authenticator assurance: NIST SP 800-63B.

QUESTION NO: 7

Which of the following requirements should be addressed when a cloud provider processes personal data on behalf of a customer acting as controller?

(Choose two.)

A. Documented processing instructions

B. Subprocessor controls

C. Unlimited reuse of data for provider marketing

D. Removal of all controller audit rights

ANSWER: A B

Explanation:

Documented processing instructions and **subprocessor controls** are correct. A processor should process personal data only according to documented instructions from the controller, except where applicable law requires otherwise. These instructions establish the permitted purpose and boundaries for processing.

Where the provider uses subprocessors, the controller needs appropriate visibility and control over those parties. The agreement should address authorization, notice of intended changes where appropriate, and flow-down obligations so the subprocessor provides equivalent protection for the personal data. These requirements support accountability across the full cloud service-delivery chain. Article 28 of the GDPR addresses processor contracts and the engagement of other processors; see GDPR, Article 28.

QUESTION NO: 8

You are the security manager of a small firm that has just purchased a DLP solution to implement in your cloud-based production environment.

What should you not expect the tool to address? Response:

A. Sensitive data sent inadvertently in user emails

B. Sensitive data captured by screen shots

C. Sensitive data moved to external devices

D. Sensitive data in the contents of files sent via FTP

ANSWER: B

Explanation:

Sensitive data captured by screen shots is the correct answer because conventional data loss prevention capabilities are principally designed to identify sensitive content and apply policy while that content is stored, transmitted, or handled through supported applications and channels. DLP commonly uses content inspection, classification rules, labels, and contextual controls to detect protected information and then alert, block, quarantine, encrypt, or otherwise govern its use. A screenshot converts the displayed information into an image or set of screen pixels. This creates an “analog hole” problem: the information may be visible to an authorized user but subsequently captured outside the normal content-aware channels on which a DLP policy depends. The risk is even broader if a user photographs the display with an unmanaged device, because no software-based endpoint control can reliably inspect or prevent that capture. Some integrated endpoint, rights-management, or application-specific controls can reduce screen-capture risk in particular supported scenarios, but preventing screenshots is not a core or dependable expectation for a typical DLP deployment. Organizations should treat this as a residual risk and complement DLP with least privilege, data minimization, managed endpoints, application controls, monitoring, and user awareness. NIST describes DLP as controls for detecting and preventing unauthorized disclosure of sensitive information; Microsoft similarly describes DLP policies as identifying and protecting sensitive items across supported locations and activities. See NIST SP 800-53 Rev. 5 and Microsoft Purview data loss prevention documentation.

QUESTION NO: 9

A virtual network interface card (NIC) exists at layer of the OSI model. Response:

A. 2

B. 4

C. 6

D. 8

ANSWER: A

Explanation:

2 is correct because a virtual network interface card functions primarily at the OSI Data Link layer. A vNIC is the software-defined equivalent of an Ethernet network adapter presented to a virtual machine or cloud workload. It has a Media Access Control (MAC) address and provides the guest operating system with the capability to transmit and receive Ethernet frames. These are core Layer 2 responsibilities: frame handling, local-link delivery using MAC addressing, and connection to Layer 2 constructs such as virtual switches and VLAN-backed network segments.

Although a virtual NIC relies on an underlying physical adapter and virtualization platform to move signals and packets, its logical interface to the workload is an Ethernet link-layer interface. Internet Protocol addresses, routing, and subnet decisions are handled above the vNIC by Layer 3 components, while transport protocols operate still higher in the stack. This distinction is especially relevant in cloud environments, where a cloud provider may implement switching and segmentation virtually but still exposes a conventional Layer 2-style network adapter to an instance. Cisco's OSI model overview identifies the Data Link layer as the layer responsible for framing and physical addressing: Cisco: What Is a Network Switch?. Microsoft also describes virtual NICs as virtual network adapters connected to virtual switches: Microsoft Learn: Plan for Hyper-V networking.

QUESTION NO: 10

Most APIs will support a variety of different data formats or structures.

However, the SOAP API will only support which one of the following data formats?

A. XML

B. XSLT

C. JSON

D. SAML

ANSWER: A

Explanation:

XML is correct because SOAP defines its messages as XML documents. A SOAP message has a prescribed XML structure consisting of an *Envelope* element, with optional *Header* and required *Body* elements. These elements provide the standardized message framework used to exchange request, response, fault, and metadata information between a SOAP client and service. SOAP may be carried by several underlying transport protocols, including HTTP and SMTP, but changing the transport does not change the SOAP message format: the message itself remains XML.

This XML requirement also enables SOAP's formal extensibility model. Features such as security headers, transaction context, routing details, and reliable messaging can be represented in the SOAP header while retaining a structured, interoperable XML message. SOAP services are commonly described through WSDL, another XML-based standard, which defines the service operations, message structures, bindings, and endpoints. In the context of the question, the intended distinction is that SOAP uses XML-formatted messages, unlike many modern API styles that can use multiple representations. The W3C SOAP specification explicitly describes SOAP as an XML-based messaging framework and

defines the XML information set for SOAP envelopes. See the W3C SOAP Version 1.2 Messaging Framework and the W3C SOAP specification.

QUESTION NO: 11

Which phase of the cloud data lifecycle represents the first instance where security controls can be implemented?

A. Use

B. Share

C. Store

D. Create

ANSWER: D

Explanation:

Create is the first cloud data lifecycle phase in which security controls can be implemented because it is the point at which data originates or is first captured, generated, received, or materially modified. Security should begin with the earliest possible handling of information so that the data is correctly classified, assigned an owner, and associated with handling requirements from its inception. At creation, organizations can implement controls such as data-classification labels, sensitivity tags, approved data schemas, input validation, data-minimization requirements, metadata requirements, and policies governing the collection of personal or regulated information. These decisions establish the protection requirements that follow the data throughout its lifecycle, including requirements for encryption, access restrictions, retention, monitoring, residency, and secure disposal. Applying controls at creation supports a “security by design” approach and reduces the risk that sensitive information enters cloud processes without an appropriate protection designation. The CCSP examination outline includes cloud data lifecycle responsibilities within cloud data security, emphasizing that candidates must understand data security controls across the lifecycle. Guidance on data-centric security likewise emphasizes identifying and classifying information as early as possible so protection can consistently follow the data. See the ISC2 CCSP Certification Exam Outline and the Cloud Security Alliance Security Guidance.

QUESTION NO: 12

Which of the following are benefits of centralized cloud logging?

(Choose two.)

A. Correlation of events across multiple cloud services

B. Retention of evidence independent of individual workloads

C. Eliminating the need to review security alerts

D. Allowing unrestricted access to log repositories

ANSWER: A B

Explanation:

Correlation of events across multiple cloud services and **retention of evidence independent of individual workloads** are benefits of centralized cloud logging. Aggregating audit, identity, network, application, and platform logs enables analysts to identify related activity across services and accounts that might not be evident from a single workload's records.

Centralized retention also reduces the risk that evidence is lost when a virtual machine, container, or other ephemeral resource is terminated. Access to the logging platform must still be restricted, and log integrity should be protected through appropriate access controls, monitoring, and retention policies. Centralized logging does not eliminate the need for alert review or permit unrestricted administrative access. See NIST SP 800-92, Guide to Computer Security Log Management.

QUESTION NO: 13

Which of the following is NOT a commonly used communications method within cloud environments to secure data in transit?

A. IPSec

B. HTTPS

C. VPN

D. DNSSEC

ANSWER: D

Explanation:

DNSSEC is the correct answer because its purpose is to protect the integrity and authenticity of DNS information, rather than to provide confidentiality for application data transmitted between cloud clients and services. DNSSEC adds digital signatures to DNS records, enabling validating resolvers to establish that received DNS data originates from the authoritative zone and was not modified in transit. This helps mitigate threats such as forged DNS responses and cache-poisoning attacks, which can otherwise redirect users to malicious endpoints.

Although reliable DNS resolution is an important supporting security control in a cloud architecture, DNSSEC does not encrypt the connection that follows a DNS lookup. Protection of actual data in transit requires a transport or network-layer encryption mechanism, such as TLS used by HTTPS or IPsec-based protected networking. DNSSEC can therefore complement in-transit protections by helping a client identify the legitimate service endpoint, but it is not itself a communications method that secures the confidentiality of application payloads while they traverse a network. The IETF's DNSSEC specifications describe it as a system for DNS data origin authentication and integrity; Cloudflare likewise explains that DNSSEC protects DNS records from tampering rather than encrypting traffic.

References: IETF RFC 4033: DNS Security Introduction and Requirements; Cloudflare: What is DNSSEC?

QUESTION NO: 14

For optimal security, trust zones are used for network segmentation and isolation. They allow for the separation of various systems and tiers, each with its own security level.

Which of the following is typically used to allow administrative personnel access to trust zones?

A. IPSec

B. SSH

C. VPN

D. TLS

ANSWER: C

Explanation:

VPN is typically used to provide administrative personnel with controlled access to resources within protected trust zones. Network segmentation separates systems by function and sensitivity, so management interfaces in higher-trust zones should not be exposed directly to untrusted networks or the public Internet. A virtual private network establishes an authenticated, encrypted connection from an authorized administrator to the private network, allowing access to designated management subnets while preserving the isolation boundaries of those zones.

In a well-designed cloud or hybrid environment, the VPN endpoint is commonly integrated with centralized identity controls, multifactor authentication, authorization policies, logging, and restricted routing. After connecting, an administrator may be permitted to reach a bastion host or other specifically authorized management resources, rather than receiving unrestricted access across every network segment. This supports least privilege, creates an auditable administrative entry point, and

protects management traffic in transit. AWS describes Client VPN as a managed client-based VPN service for securely accessing AWS and on-premises resources, while Azure VPN Gateway provides encrypted connectivity to Azure virtual networks. See AWS Client VPN documentation and Azure VPN Gateway documentation.

QUESTION NO: 15

What is a cloud storage architecture that manages the data in caches of copied content close to locations of high demand? Response:

A. Object-based storage

B. File-based storage

C. Database

D. CDN

ANSWER: D

Explanation:

CDN is correct because a content delivery network maintains cached replicas of content at geographically distributed edge locations, placing those copies nearer to users and regions where requests are concentrated. Rather than retrieving every request from the origin server or primary cloud storage location, the CDN can serve eligible content from an edge cache. This reduces network distance and latency, improves delivery performance, and reduces demand on the origin infrastructure. In cloud architectures, a CDN is commonly positioned in front of object storage, web servers, or applications to distribute static content such as images, style sheets, scripts, documents, media, and software downloads. CDN caching behavior is controlled through mechanisms such as cache-control directives, time-to-live settings, cache keys, and invalidations, allowing organizations to balance performance with the need to provide current content. The defining characteristic in the question—copied content managed in caches close to high-demand locations—is therefore the core architectural purpose of a CDN. Cloudflare describes a CDN as a geographically distributed group of servers that caches content close to end users, while AWS documents that CloudFront uses a global network of edge locations to deliver content with low latency. See Cloudflare: What is a CDN? and AWS CloudFront Developer Guide.

QUESTION NO: 16

Which concept of cloud computing pertains to the ability to reuse components and services of an application for other purposes?

A. Portability

B. Interoperability

C. Resource pooling

D. Elasticity

ANSWER: B

Explanation:

Interoperability is correct because it enables independently developed systems, applications, components, and services to interact through compatible interfaces, protocols, and data formats. In a cloud environment, application capabilities are commonly exposed as reusable services through APIs, service-oriented architectures, or microservices. Interoperability allows those capabilities to be discovered, invoked, composed, and consumed by other applications or business processes without being limited to their original implementation context. This makes practical reuse possible: for example, an identity, payment, logging, or notification service can support multiple applications when each can communicate with it using agreed technical interfaces and shared information conventions.

NIST describes interoperability as the ability of systems or components to exchange information and use the information exchanged. That definition directly supports reuse across applications because a component has value beyond its original purpose when other systems can integrate with and use it. Cloud interoperability also supports service composition, where multiple compatible services are combined to deliver new functionality. See NIST SP 500-292, Cloud Computing Reference Architecture and the ISO/IEC 17788 cloud-computing vocabulary.

QUESTION NO: 17

Which of the following practices support an effective cloud data-classification program?

(Choose two.)

A. Assigning data owners

B. Applying handling requirements

C. Allowing each user to determine classification

D. Removing retention requirements

ANSWER: A B

Explanation:

Assigning data owners and **applying handling requirements** are correct. A data owner is accountable for determining the business value, sensitivity, classification, and authorized use of data. Assigning ownership prevents cloud data stores from becoming unmanaged and ensures that classification decisions have clear accountability.

Classification must also lead to defined handling requirements. These requirements can specify access restrictions, encryption needs, approved storage locations, retention periods, sharing limitations, and destruction procedures. Classification without handling rules does not provide an enforceable basis for protecting sensitive information throughout its lifecycle. NIST describes data classification and handling requirements within its security and privacy control framework; see NIST SP 800-53 Rev. 5.

QUESTION NO: 18

Which of the following security measures done at the network layer in a traditional data center are also applicable to a cloud environment?

A. Dedicated switches

B. Trust zones

C. Redundant network circuits

D. Direct connections

ANSWER: B

Explanation:

Trust zones are applicable because the fundamental security goal of separating systems with differing sensitivity, roles, and levels of trust remains the same in both traditional and cloud data centers. A trust zone establishes a logical boundary around workloads, such as public-facing web services, application services, and databases, and defines the permitted communication paths across that boundary. This supports least privilege by allowing only the required north-south and east-west traffic, reducing opportunities for unauthorized access and lateral movement.

In cloud environments, trust zones are normally implemented through software-defined networking rather than dedicated physical network hardware. Cloud-native constructs such as virtual networks, subnets, route tables, security groups or network security groups, firewall policies, and private connectivity provide the boundary and policy-enforcement mechanisms. The implementation may differ from a traditional data center, but the architecture principle is unchanged:

segment the environment, apply controls appropriate to each zone's risk level, and explicitly govern traffic between zones. This is a core component of secure cloud network design and helps limit the blast radius of a compromised workload. AWS describes how a VPC provides a logically isolated virtual network at AWS VPC documentation, while Microsoft explains Azure virtual-network segmentation and connectivity at Azure Virtual Network documentation.

QUESTION NO: 19

You just hired an outside developer to modernize some applications with new web services and functionality. In order to implement a comprehensive test platform for validation, the developer needs a data set that resembles a production data set in both size and composition.

In order to accomplish this, what type of masking would you use?

A. Development

B. Replicated

C. Static

D. Dynamic

ANSWER: C

Explanation:

Static masking is appropriate because the developer needs a persistent, non-production dataset that retains the production environment's approximate volume, structure, relationships, and statistical characteristics. Static data masking creates a separate copy or extract of production data and permanently transforms sensitive values before the dataset is supplied to development or test personnel. This permits realistic functional, integration, and performance validation without exposing actual personal, financial, proprietary, or otherwise sensitive production values to an outside developer.

The transformation is performed before the data reaches the test platform and can use methods such as substitution, shuffling, redaction, or tokenization while preserving required formats and referential integrity. As a result, applications and web services can be tested against data that behaves like production data in scale and composition, while the masked dataset is safe to retain and use in the non-production environment. This supports CCSP data-lifecycle principles of minimizing sensitive-data exposure and applying appropriate protection to data used outside production. IBM describes static data masking as masking data in a copy of a production database for use in non-production environments. See IBM Documentation: Static data masking and Microsoft Learn: Dynamic Data Masking.

QUESTION NO: 20

Which of the following threat types involves leveraging a user's browser to send untrusted data to be executed with legitimate access via the user's valid credentials?

A. Injection

B. Missing function-level access control

C. Cross-site scripting

D. Cross-site request forgery

ANSWER: D

Explanation:

Cross-site request forgery is the threat type described because it abuses the trust that an application places in requests sent by an authenticated user's browser. When a user has an active session with a site, the browser may automatically attach valid session cookies or other ambient authentication credentials to requests for that site. An attacker can cause the browser to submit a crafted request—such as through a malicious web page, embedded content, or a deceptive link—without the user intending to perform that action. Because the request carries the user's valid credentials, the target application can treat

it as an authorized request and execute the action with the user's legitimate access rights. This makes CSRF especially relevant to state-changing operations, including modifying account details, transferring funds, or changing security settings. Effective defenses establish that the request originated from an intended application workflow, commonly through unpredictable anti-CSRF tokens, appropriate SameSite cookie settings, and additional confirmation or reauthentication for sensitive operations. OWASP describes this browser-based use of authenticated credentials in its Cross-Site Request Forgery overview and provides implementation guidance in its CSRF Prevention Cheat Sheet.

QUESTION NO: 21

Which of the following aspects of the BC/DR process poses a risk to the organization? Response:

A. Threat intelligence gathering

B. Preplacement of response assets

C. Budgeting for disaster

D. Full testing of the plan

ANSWER: D

Explanation:

Full testing of the plan poses a risk because a comprehensive BC/DR exercise can affect live business operations while validating recovery capabilities. A full interruption or failover test may involve moving workloads to an alternate environment, changing network routing or DNS records, restoring data, operating with reduced capacity, and subsequently failing back to the primary environment. Each activity can create a possibility of service disruption, data-integrity problems, configuration drift, unexpected dependency failures, or incomplete rollback if the process is not tightly governed.

This does not mean that testing should be avoided. Regular, appropriately scoped testing is essential to demonstrate that recovery objectives, personnel procedures, technical controls, and supplier dependencies will work during an actual disruption. The risk must instead be managed through documented test plans, authorization and change management, stakeholder communications, defined success criteria, backup and rollback procedures, monitoring, and a progression from tabletop exercises to simulations and controlled technical tests before a full-scale exercise. NIST contingency-planning guidance specifically treats testing, training, and exercises as key plan-maintenance activities and emphasizes selecting methods that validate capabilities while controlling operational impact. See NIST SP 800-34 Rev. 1, Contingency Planning Guide and NIST contingency-planning guidance.

QUESTION NO: 22

Which of the following controls support secure use of infrastructure as code in a cloud environment?

(Choose two.)

A. Reviewing code changes before deployment

B. Scanning templates for security misconfigurations

C. Storing credentials in source-code repositories

D. Allowing direct production changes outside the deployment workflow

ANSWER: A B

Explanation:

Reviewing code changes before deployment and **scanning templates for security misconfigurations** support secure infrastructure as code. Code review helps identify unauthorized, erroneous, or risky changes before infrastructure is provisioned. Automated scanning can identify conditions such as unrestricted network access, missing encryption settings, overly permissive identity policies, and publicly exposed storage resources.

Infrastructure as code should be version controlled, tested, approved, and deployed through controlled pipelines. Storing credentials in source-code repositories creates unnecessary exposure, and allowing direct production changes outside the controlled workflow produces configuration drift and weakens accountability. NIST describes secure software-development practices in SP 800-218, Secure Software Development Framework.

QUESTION NO: 23

A variety of security systems can be integrated within a network--some that just monitor for threats and issue alerts, and others that take action based on signatures, behavior, and other types of rules to actively stop potential threats.

Which of the following types of technologies is best described here?

A. IDS

B. IPS

C. Proxy

D. Firewall

ANSWER: B

Explanation:

IPS is correct because an intrusion prevention system is designed to inspect network traffic for malicious activity or policy violations and then actively intervene to prevent the activity from succeeding. Its detection methods can include known attack signatures, protocol and rule analysis, behavioral indicators, and anomaly-based techniques. Crucially, an IPS has the capability to enforce an automated response in or alongside the traffic path. Depending on its deployment and configuration, that response can include dropping malicious packets, terminating a session, blocking a source address, or applying another defined security policy. This combines threat detection with preventive enforcement, matching the scenario's emphasis on systems that take action to stop potential threats rather than only produce alerts. NIST describes intrusion detection and prevention systems as technologies that monitor events and analyze them for signs of security incidents; prevention capabilities enable the system to attempt to stop detected incidents. See NIST SP 800-94, Guide to Intrusion Detection and Prevention Systems. Cisco also summarizes the operational distinction: an IPS can block or reject malicious traffic as it is detected, which is the defining capability described here: Cisco IDS and IPS overview.

QUESTION NO: 24

Which United States law is focused on accounting and financial practices of organizations?

A. Safe Harbor

B. GLBA

C. SOX

D. HIPAA

ANSWER: C

Explanation:

SOX is correct because the Sarbanes-Oxley Act of 2002 is a U.S. federal law designed to strengthen the accuracy, reliability, and integrity of public-company financial reporting. Enacted following major accounting scandals, it established significant requirements for corporate governance, audit oversight, executive responsibility for financial statements, and internal control over financial reporting. In particular, Section 404 requires management to assess the effectiveness of internal controls related to financial reporting, with external-auditor attestation requirements applying in relevant circumstances.

For cloud security professionals, SOX obligations commonly affect the systems that create, process, transmit, or retain financial-reporting information. Organizations need demonstrable controls over access, changes, logging, record retention,

backup and recovery, and the integrity of records. When cloud services support an in-scope business process, the customer must obtain sufficient assurance that relevant controls operate effectively, often through contractual commitments, audit reports, evidence collection, and shared-responsibility governance. The U.S. Securities and Exchange Commission provides an overview of the law and related implementation materials at SEC: Sarbanes-Oxley Act. The enacted legislation and its congressional history are available at Congress.gov: H.R. 3763, Sarbanes-Oxley Act of 2002.

QUESTION NO: 25

If you're using iSCSI in a cloud environment, what must come from an external protocol or application?

A. Kerberos support

B. CHAP support

C. Authentication

D. Encryption

ANSWER: D

Explanation:

Encryption is correct because iSCSI defines how SCSI block-storage commands are transported over TCP/IP, but it does not itself provide encryption for the data payload in transit. Consequently, when confidentiality is required for iSCSI communications in a cloud environment, encryption must be supplied by a separate security mechanism or application layer, such as IPsec, a VPN tunnel, or an appropriately implemented encrypted network overlay. This is particularly important in cloud and virtualized environments, where storage traffic may traverse shared physical infrastructure, virtual switches, or networks outside the direct control of the workload owner. Applying external encryption protects the confidentiality and integrity of transmitted storage data against interception or unauthorized modification while it travels between the iSCSI initiator and target. The iSCSI standard specifies security negotiation and authentication capabilities, including CHAP, but states that IPsec may be used to provide confidentiality and integrity services. This separation illustrates a key cloud-security principle: do not assume that a storage transport protocol inherently supplies all required data-protection controls; implement compensating controls for data in transit according to the organization's risk and compliance requirements. See the iSCSI protocol specification (RFC 7143) and the IPsec security architecture (RFC 4301).

QUESTION NO: 26

A federated identity system is composed of three main components. Which of the following is NOT one of the three main components?

Response:

A. Identity provider

B. User

C. Relying party

D. API

ANSWER: D

Explanation:

API is correct because it is not one of the foundational participants in a federated identity system. Federation establishes a trust relationship through which a user's identity and authentication information can be accepted by a separate application or service without requiring that service to authenticate the user independently. The core participants are the user, who requests access; the identity provider, which authenticates the user and produces a signed assertion or token; and the relying party, which trusts and consumes that assertion or token to make an access decision.

An API may be the resource ultimately accessed after federation occurs, and it can validate an access token issued through an OAuth 2.0 or OpenID Connect flow. That role, however, describes a protected interface or resource server, not a required principal component of the basic federation trust model. In OpenID Connect terminology, the identity provider function is performed by an OpenID Provider, while the application consuming identity information is the relying party. This separation of responsibilities is central to federated identity architecture. See the OpenID Connect Core specification and NIST's Digital Identity Guidelines: Federation and Assertions.

QUESTION NO: 27

What is the primary security benefit of tokenization when used to protect payment card data in a cloud application?

Response:

A. Reducing exposure of the original cardholder data

B. Eliminating the need for access controls

C. Providing nonrepudiation for transactions

D. Compressing cardholder data for storage

ANSWER: A

Explanation:

Reducing exposure of the original cardholder data is correct because tokenization substitutes a non-sensitive token for a sensitive value such as a primary account number. Applications can store and process the token for authorized business functions without routinely retaining the original card value. The relationship between the token and original data is maintained separately in a protected tokenization system.

Unlike encryption, tokenization does not rely on reversible cryptographic transformation of the original value by the applications that use the token. Properly implemented tokenization can reduce the number of systems that store or process cardholder data, although the organization must still assess the actual scope of its environment and protect the tokenization service. The PCI Security Standards Council discusses tokenization and its effect on cardholder-data environments in its Tokenization Guidelines Information Supplement.

QUESTION NO: 28

What are the objectives of change management? (Choose all that apply.) Response:

A. Respond to a customer's changing business requirements while maximizing value and reducing incidents, disruption, and rework

B. Ensure that changes are recorded and evaluated

C. Respond to business and IT requests for change that will disassociate services with business needs

D. Ensure that all changes are prioritized, planned, tested, implemented, documented, and reviewed in a controlled manner

ANSWER: A B D

Explanation:

Change management, called "change enablement" in ITIL 4, is intended to make beneficial changes practical and timely while protecting service quality and business operations. "Respond to a customer's changing business requirements while maximizing value and reducing incidents, disruption, and rework" reflects the central objective: balancing responsiveness to evolving needs with control of operational risk. This balance is particularly important for cloud environments, where automated deployment pipelines and infrastructure as code can substantially increase the rate of change.

"Ensure that changes are recorded and evaluated" is also an objective because a reliable change record provides traceability and supports informed assessment of a proposed change's value, risks, impacts, dependencies, authorization

requirements, and implementation approach. Recorded evaluation also establishes accountability and enables organizational learning.

“Ensure that all changes are prioritized, planned, tested, implemented, documented, and reviewed in a controlled manner” describes the lifecycle discipline required to ensure changes are authorized, appropriately sequenced, validated, and reviewed after deployment. These practices support availability, security, compliance evidence, rollback readiness, and continual improvement. ITIL describes change enablement as maximizing successful service and product changes by ensuring risks have been properly assessed, authorizing changes, and managing the change schedule. See AXELOS ITIL service management and Microsoft Cloud Adoption Framework guidance on change management.

QUESTION NO: 29

Which of the following provides assurance, to a predetermined acceptable level of certainty, that an entity is indeed who they claim to be?

A. Authentication

B. Identification

C. Proofing

D. Authorization

ANSWER: A

Explanation:

Authentication is the process that establishes confidence that an entity is genuinely the identity it claims to be. An entity first asserts an identity, such as by supplying a username, account identifier, or certificate subject. Authentication then tests that assertion by requiring evidence bound to the identity. This evidence may include a memorized secret, a cryptographic authenticator or possession factor, or a biometric characteristic. When the supplied evidence is successfully validated, the relying system has assurance—at a defined level of confidence—that the claimant is the asserted entity.

The reference to a “predetermined acceptable level of certainty” is important because assurance is not necessarily absolute. The required confidence depends on the system’s risk context and is increased by using stronger authenticators, phishing-resistant mechanisms, or multiple independent factors. In cloud environments, authentication is a core component of identity and access management: it establishes the verified identity context on which subsequent access-control decisions can safely rely. NIST defines authentication as the process of determining whether a claimant is the person they claim to be and explains how authenticator strength supports assurance in its digital identity guidance. See the NIST authentication glossary definition and NIST SP 800-63B, Digital Identity Guidelines: Authentication and Lifecycle Management.

QUESTION NO: 30

What are the six components that make up the STRIDE threat model? Response:

A. Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, and Elevation of Privilege

B. Spoofing, Tampering, Non-Repudiation, Information Disclosure, Denial of Service, and Elevation of Privilege

C. Spoofing, Tampering, Repudiation, Information Disclosure, Distributed Denial of Service, and Elevation of Privilege

D. Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, and Social Engineering

ANSWER: A

Explanation:

Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, and Elevation of Privilege are the six threat categories defined by Microsoft’s STRIDE threat-modeling methodology. STRIDE provides a structured way to identify threats during application and cloud-system design, typically by examining each element and data flow in a data-flow diagram. Spoofing covers impersonation of an identity; tampering concerns unauthorized modification of data or system

components; repudiation addresses the ability to deny an action when adequate evidence, such as audit logs, is absent; information disclosure concerns exposure of protected information; denial of service concerns making a service or resource unavailable; and elevation of privilege concerns obtaining permissions beyond those intended. These categories map directly to fundamental security objectives, including authentication, integrity, accountability, confidentiality, availability, and authorization. Applying STRIDE early in the development life cycle helps cloud architects identify applicable protections before deployment, such as strong identity controls, integrity validation, centralized logging, encryption, resiliency measures, and least-privilege access design. Microsoft describes STRIDE as a core threat-modeling framework and identifies these exact six categories in its guidance. See Microsoft's STRIDE threat guidance and Microsoft Threat Modeling Tool documentation.

QUESTION NO: 31

Federation should be to the users. Response:

A. Hostile

B. Proportional

C. Transparent

D. Expensive

ANSWER: C

Explanation:

Federation should be **transparent** to users. In a federated identity architecture, the user authenticates through an identity provider and then receives access to relying-party applications or cloud services based on trusted assertions or tokens. The federation protocols, trust configuration, metadata exchange, token validation, and attribute release decisions should operate in the background rather than requiring the user to understand or manually manage them.

This transparency is central to single sign-on: after completing the required authentication steps, users can move between authorized services with a consistent, low-friction experience. It also supports stronger cloud IAM operations by centralizing authentication and policy enforcement at the identity provider, reducing password proliferation and avoiding separate credentials for every service. Transparent federation does not mean invisible security; users may still be prompted for multifactor authentication or step-up authentication when risk, policy, or assurance requirements demand it. Instead, it means the underlying cross-domain identity trust is seamless from the user's perspective.

NIST describes federation as using assertions to convey authentication and attribute information from a credential service provider to a relying party: NIST SP 800-63C: Federation and Assertions. ISC2's CCSP domain outline also identifies identity and access management as a core cloud-security area: CCSP Certification Exam Outline.

QUESTION NO: 32

Which of the following roles involves overseeing billing, purchasing, and requesting audit reports for an organization within a cloud environment?

A. Cloud service user

B. Cloud service business manager

C. Cloud service administrator

D. Cloud service integrator

ANSWER: B

Explanation:

The correct role is **Cloud service business manager**. This role is responsible for the commercial, financial, and administrative aspects of an organization's use of cloud services. Its responsibilities commonly include managing cloud-

service accounts and subscriptions, monitoring and allocating costs, overseeing billing, approving or coordinating procurement, and ensuring that service consumption aligns with contractual commitments and organizational budget requirements. The role also coordinates with the cloud provider and internal procurement, finance, risk, and compliance stakeholders.

Requesting audit and assurance reports is a natural responsibility for this business-facing cloud role. Organizations need evidence such as SOC reports, ISO/IEC certification information, and related provider assurance documentation to support vendor governance, due diligence, compliance reviews, and contractual oversight. While technical staff may evaluate the control details in these artifacts, the business manager commonly facilitates obtaining them through the provider relationship and ensures they are available to the appropriate governance functions. ISO/IEC 17788 identifies cloud-service business management among cloud-service customer activities, including business-related responsibilities surrounding cloud-service use. See ISO/IEC 17788:2014 Cloud computing — Overview and vocabulary and the Cloud Security Alliance Cloud Controls Matrix for cloud governance and assurance context.

QUESTION NO: 33

Maintenance mode requires all of these actions except:

A. Remove all active production instances

B. Ensure logging continues

C. Initiate enhanced security controls

D. Prevent new logins

ANSWER: C

Explanation:

Initiate enhanced security controls is the exception because maintenance mode is principally an operational state used to protect system integrity while planned changes, servicing, or recovery activities occur. The essential practices are to stop or limit new user access, remove or quiesce active production workloads as appropriate for the maintenance activity, and retain logging so administrator actions and system events remain accountable and available for review. Security controls must continue to operate during maintenance, particularly access control and audit logging, but maintenance mode does not inherently require adding or elevating controls beyond the organization's established security baseline. Any temporary enhanced measures—such as additional monitoring, approvals, or privileged-access restrictions—are selected according to the scope and risk of a particular change under the organization's change-management process. NIST control AU-2 requires organizations to define auditable events and ensure their logging supports accountability, including during administrative activity. NIST CM-3 also establishes that changes should be controlled and documented, supporting the use of a managed maintenance state rather than normal production operation. See NIST SP 800-53 Rev. 5, Release 1 and NIST Risk Management Framework.

QUESTION NO: 34

Managed cloud services exist because the service is less expensive for each customer than creating the same services for

themselves in a legacy environment.

Using a managed service allows the customer to realize significant cost savings through the reduction of

. Response:

A. Risk

B. Security controls

C. Personnel

D. Data

ANSWER: C

Explanation:

Personnel is correct because managed cloud services reduce the customer's need to employ people to operate an equivalent technology service internally. In a traditional environment, the organization must provide staff and associated operational capabilities for activities such as platform deployment, configuration, maintenance, patching, monitoring, backup administration, capacity planning, troubleshooting, and potentially around-the-clock incident response. Those staffing requirements include not only salaries, but also recruiting, training, benefits, management overhead, specialized tools, and on-call coverage.

A managed-service provider performs much of this operational work as a repeatable service for many customers. This lets the provider spread the cost of specialized engineers, automation, operational processes, and facilities across its customer base. The customer can therefore consume the service without building and maintaining a comparably sized internal operations team. The customer still needs personnel for governance, vendor management, architecture decisions, and responsibilities retained under the shared-responsibility model, but the total staffing burden for day-to-day service operation is substantially reduced.

This reflects the cloud economic principle of economies of scale and the reduction of operational overhead for undifferentiated technical work. See the AWS overview of cloud-computing advantages and the NIST Cloud Computing Standards Roadmap.

QUESTION NO: 35

Which of the following statements best describes a Type 1 hypervisor?

A. The hypervisor software runs within an operating system tied to the hardware.

B. The hypervisor software runs as a client on a server and needs an external service to administer it.

C. The hypervisor software runs on top of an application layer.

D. The hypervisor software runs directly on "bare metal" without an intermediary.

ANSWER: D

Explanation:

The hypervisor software runs directly on "bare metal" without an intermediary. This describes a Type 1 hypervisor, also known as a bare-metal or native hypervisor. Its defining architectural characteristic is that the virtualization layer operates directly on the physical server's hardware rather than being installed as an application within a general-purpose host operating system. The hypervisor allocates and schedules physical resources such as processor time, memory, storage, and network access for guest virtual machines, while maintaining isolation between those workloads. This direct architecture is widely used in enterprise virtualization and cloud infrastructure because it provides a purpose-built platform for hosting multiple virtual machines and avoids placing a conventional host OS beneath the hypervisor. Management functions may be delivered through local or remote administrative tools and a separate management plane, but these operational details do not alter the Type 1 classification. The classification depends on where the hypervisor executes in the system stack: directly over the server hardware. VMware identifies ESXi as a bare-metal hypervisor installed directly on a physical server, and Microsoft documents Hyper-V as a hypervisor-based virtualization technology for Windows Server. See VMware's hypervisor overview and Microsoft Learn: Hyper-V on Windows Server.

QUESTION NO: 36

Which of the following activities are important during cloud incident response?

(Choose two.)

A. Preserving relevant logs and snapshots

B. Coordinating with the cloud provider under defined procedures

C. Disabling logging to reduce storage costs

D. Altering evidence without documenting the change

ANSWER: A B

Explanation:

Preserving relevant logs and snapshots and **coordinating with the cloud provider under defined procedures** are important cloud incident-response activities. Cloud evidence can be volatile because resources may be automatically scaled, terminated, overwritten, or reconfigured. Preserving relevant logs, snapshots, and configuration information supports investigation while maintaining evidence integrity.

The cloud provider may control infrastructure-level telemetry, physical systems, or service capabilities that the customer cannot access directly. Defined coordination procedures, including escalation contacts and contractual support obligations, help ensure timely containment and evidence collection. Disabling logging reduces visibility, while altering evidence without documentation compromises forensic reliability. See NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide and NIST SP 800-86, Guide to Integrating Forensic Techniques into Incident Response.

QUESTION NO: 37

You are the security policy lead for your organization, which is considering migrating from your on-premises, legacy environment into the cloud. You are reviewing the Cloud Security Alliance Cloud Controls Matrix (CSA CCM) as a tool for your organization.

What is probably the best benefit offered by the CCM? Response:

A. The low cost of the tool

B. Allowing your organization to leverage existing controls across multiple frameworks so as not to duplicate effort

C. Simplicity of control selection from the list of approved choices

D. Ease of implementation by choosing controls from the list of qualified vendors

ANSWER: B

Explanation:

The best benefit is “Allowing your organization to leverage existing controls across multiple frameworks so as not to duplicate effort.” The Cloud Security Alliance Cloud Controls Matrix is a cloud-specific control framework designed to provide a common set of security and privacy control objectives for cloud environments. A central CCM capability is its mapping, or crosswalk, to numerous widely adopted standards, regulations, and frameworks. This enables an organization to compare its existing governance, risk, compliance, and security controls with cloud-focused requirements without starting its control program from scratch.

For a migration from a legacy on-premises environment, these mappings help the security policy lead identify reusable policies, procedures, control evidence, and assessment activities. The organization can rationalize overlapping requirements and establish a unified cloud control baseline that supports multiple assurance obligations. This is especially useful for communicating expectations among internal security teams, auditors, cloud customers, and cloud service providers under the shared-responsibility model. The CCM also serves as the foundational control framework for CSA’s Security, Trust, Assurance, and Risk program, supporting consistent cloud assurance assessments. See the Cloud Security Alliance Cloud Controls Matrix and the CSA STAR program.

QUESTION NO: 38

Which of the following is the optimal humidity level for a data center, per the guidelines established by the America Society of Heating, Refrigeration, and Air Conditioning Engineers (ASHRAE)?

A. 30-50 percent relative humidity

B. 50-75 percent relative humidity

C. 20-40 percent relative humidity

D. 40-60 percent relative humidity

ANSWER: D

Explanation:

40-60 percent relative humidity is the best answer because it represents the traditional, practical operating band commonly used to maintain a stable data-center environment. This moderate range controls the principal humidity-related risks to IT equipment: low relative humidity encourages electrostatic-charge accumulation and electrostatic discharge, while excessive moisture can contribute to corrosion and, under changing thermal conditions, condensation risk. Maintaining humidity within this band therefore supports the reliability and availability of servers, storage, network hardware, and their electrical connections.

ASHRAE's current thermal guidance for data-processing environments expresses recommended and allowable moisture conditions using a combination of dew point, relative humidity, and humidity ratio rather than relying solely on one relative-humidity range. Nevertheless, 40-60 percent relative humidity remains the conventional exam-oriented target because it captures the intended balance between static control and moisture control. In implementation, facilities should monitor dew point as well as relative humidity, since relative humidity varies with temperature even when the actual amount of water vapor in the air does not. ASHRAE Technical Committee 9.9 develops the environmental guidance for data centers; see ASHRAE TC 9.9 and Cisco's summary of data-center environmental guidelines.

QUESTION NO: 39

What is the minimum regularity for testing a BCDR plan to meet best practices?

A. Once a year

B. Once a month

C. Every six months

D. When the budget allows it

ANSWER: A

Explanation:

Once a year is the appropriate minimum frequency under widely accepted business continuity and disaster recovery best practices. An annual exercise provides a recurring, structured opportunity to validate that recovery procedures, assigned roles, contact details, dependencies, backup restoration processes, and recovery objectives remain workable as the organization changes. Testing is not simply a documentation review: the organization should select an exercise appropriate to its maturity and risk, such as a tabletop exercise, technical recovery test, functional exercise, or full simulation, and use its results to update the BCDR plan.

Annual testing should be treated as a baseline rather than a ceiling. The plan should also be reviewed and retested after significant changes, including cloud-service migrations, material architecture changes, new critical applications, supplier changes, major incidents, or changes to recovery time and recovery point requirements. Higher-risk environments and applicable contractual, legal, or regulatory obligations can require more frequent validation. NIST's contingency-planning guidance identifies testing, training, and exercises as necessary to validate plans and maintain organizational readiness; its guidance is available in NIST SP 800-34 Rev. 1. FEMA likewise describes exercises as a means to assess and improve preparedness in its exercise guidance.

QUESTION NO: 40

Which of the following should be included in a cloud-provider exit strategy?

(Choose two.)

A. Data-export requirements

B. Secure data-deletion procedures

C. Increased use of proprietary interfaces

D. Indefinite service-termination periods

ANSWER: A B

Explanation:

Data-export requirements and **secure data-deletion procedures** are essential components of a cloud-provider exit strategy. Data-export requirements establish the format, completeness, timing, and assistance needed to retrieve customer data and associated metadata before service termination. These requirements support continuity, portability, and compliance obligations.

Secure deletion procedures establish how customer data will be returned, retained where legally required, and then sanitized or rendered inaccessible after the contract ends. The strategy should also define responsibilities, transition support, validation activities, and evidence of completion. Increasing proprietary dependencies makes transition more difficult, while an indefinite termination period fails to establish a usable exit process. See NIST SP 800-146, Cloud Computing Synopsis and Recommendations and NIST SP 800-88 Rev. 2, Guidelines for Media Sanitization.

QUESTION NO: 41

Without the extensive funds of a large corporation, a small-sized company could gain considerable and cost-effective services for which of the following concepts by moving to a cloud environment?

A. Regulatory

B. Security

C. Testing

D. Development

ANSWER: B

Explanation:

Security is correct because cloud computing can give a small organization access to security capabilities that would be difficult and expensive to build, staff, and operate independently. Major cloud service providers can spread the costs of highly resilient data centers, physical access controls, redundant power and environmental protections, infrastructure monitoring, specialized security personnel, threat detection, and vulnerability-management processes across a large customer base. This economy of scale enables customers to consume robust security services at a cost aligned more closely with their actual usage and business size.

Cloud security still operates under a shared-responsibility model: the provider protects security *of* the cloud for the services it operates, while the customer remains responsible for configuring and managing protections applicable to its data, identities, workloads, and use of the service. Consequently, moving to cloud does not eliminate the need for governance or security expertise, but it can materially improve a small company's ability to obtain mature physical and infrastructure security capabilities without making the same capital investment required for an equivalent private environment. This aligns with CCSP's emphasis on evaluating cloud-provider controls and clearly assigning security responsibilities. See ISC2's CCSP Exam Outline and NIST's Guidelines on Security and Privacy in Public Cloud Computing.

QUESTION NO: 42

Which of the following design choices improve availability for a cloud-hosted application?

(Choose two.)

A. Deploying across multiple availability zones

B. Eliminating single points of failure

C. Using one production instance for all workloads

D. Disabling health monitoring

ANSWER: A B

Explanation:

Deploying across multiple availability zones and **eliminating single points of failure** are correct. Availability zones are designed as physically separate locations within a cloud region. Distributing application components across multiple zones helps the application continue operating if one zone experiences a localized disruption.

Eliminating single points of failure requires identifying dependencies that could independently cause a service outage, including compute instances, databases, load balancers, network paths, and identity or key-management services. Redundancy, automated failover, health checks, and tested recovery procedures should be aligned with the application's recovery requirements. See AWS Well-Architected Reliability Pillar and NIST SP 800-34 Rev. 1.

QUESTION NO: 43

Which of the following are characteristics of a properly designed cloud backup strategy?

(Choose two.)

A. Regular restoration testing

B. Backup access controls

C. Storing all copies in one availability zone

D. Disabling backup audit logs

ANSWER: A B

Explanation:

Regular restoration testing and **backup access controls** are correct. A backup is useful only if it can be recovered within the required timeframe and produces complete, usable data. Regular restoration testing validates backup integrity, confirms that procedures work, and identifies dependency or configuration issues before an actual recovery event.

Backups frequently contain the same sensitive information as production systems and may be retained for longer periods. Access must therefore be restricted through least privilege, strong authentication, encryption, and logging. Otherwise, an attacker who cannot access production data may target backup repositories to steal information or prevent recovery. NIST backup and contingency-planning guidance emphasizes testing, protected storage, and controlled recovery capabilities. See NIST SP 800-34 Rev. 1 and NIST SP 800-53 Rev. 5.

QUESTION NO: 44

What is the biggest benefit to leasing space in a data center versus building or maintain your own?

A. Certification

B. Costs

C. Regulation

D. Control

ANSWER: B

Explanation:

Costs is correct because leasing data-center space, commonly through a colocation arrangement, avoids the substantial capital investment and long-term facilities burden of constructing and operating a private data center. An organization does not need to purchase land, build the facility, install large-scale redundant electrical and cooling systems, or fund the specialized facilities operations required to keep those systems available and safe. Instead, it pays a recurring charge for the space, power, cooling, and physical infrastructure capacity it uses.

This model improves financial flexibility because expenditures are more predictable and capacity can generally be expanded without a new construction project. Colocation providers also spread the cost of resilient utility feeds, backup generation, environmental controls, fire suppression, monitoring, and building security across many customers. That economy of scale can make enterprise-grade availability practical at a lower overall cost than an organization could achieve by building and maintaining equivalent facilities itself. The customer may still own and administer its equipment, but the costly building-level infrastructure is provided as a service. NIST identifies cost reduction and shifting resources toward operational expenses as important cloud-related economic considerations; its discussion is relevant to the shared-infrastructure value proposition underlying outsourced data-center facilities. See NIST SP 800-146, Cloud Computing Synopsis and Recommendations and the NIST publication record.

QUESTION NO: 45

Which type of cloud model typically presents the most challenges to a cloud customer during the "destroy" phase of the cloud data lifecycle?

A. IaaS

B. DaaS

C. SaaS

D. PaaS

ANSWER: C

Explanation:

SaaS typically presents the greatest challenge during the destroy phase because the customer has the least operational control over the application, platform, storage infrastructure, and media-handling processes. The provider operates the complete service stack, so the customer generally cannot directly verify where data resides or execute sanitization procedures on the underlying systems. Customer data can also persist in multiple provider-managed locations, including production systems, replicated storage, logs, search indexes, backups, and disaster-recovery environments. In a multitenant SaaS environment, secure deletion must be performed without disrupting data belonging to other customers, which makes physical media destruction or customer-directed wiping impractical.

Accordingly, a CCSP should ensure that the SaaS agreement defines deletion triggers, retention periods, backup-expiration handling, sanitization methods, and the evidence the provider will supply to demonstrate fulfillment. Provider attestations, audit reports, and documented media-sanitization procedures are especially important because the customer must rely on contractual and assurance mechanisms rather than direct technical control. NIST's media-sanitization guidance emphasizes selecting and documenting appropriate sanitization techniques based on the data and media involved. See NIST SP 800-88 Rev. 2, Guidelines for Media Sanitization and the Cloud Security Alliance Cloud Controls Matrix.

QUESTION NO: 46

Which of the following controls help reduce container-image security risk before deployment to a cloud environment?

(Choose two.)

A. Scanning images for vulnerabilities

B. Using trusted base images

C. Embedding administrator passwords in images

D. Allowing unsigned images from any registry

ANSWER: A B

Explanation:

Scanning images for vulnerabilities and **using trusted base images** are correct. Image vulnerability scanning identifies known vulnerable packages, libraries, and operating-system components before an image is deployed. Findings can then be remediated or evaluated according to the organization's risk process.

Using trusted, approved base images reduces the likelihood that an application inherits unnecessary software, malicious content, unsupported components, or unknown configuration settings. Organizations should obtain images from controlled registries, maintain provenance, and update images when security fixes become available. Container images should not be treated as inherently safe merely because they are obtained from a public registry. See NIST SP 800-190, Application Container Security Guide.

QUESTION NO: 47

Which of the following might make crypto-shredding difficult or useless? Response:

A. Cloud provider also managing the organization's keys

B. Lack of physical access to the environment

C. External attackers

D. Lack of user training and awareness

ANSWER: A

Explanation:

Cloud provider also managing the organization's keys might make crypto-shredding difficult or useless because cryptographic erasure depends on the permanent destruction of every key required to decrypt the protected data. This technique is effective only when the organization can control and verify the encryption-key lifecycle, including key generation, storage, backup, replication, recovery, rotation, and destruction. Where a cloud provider manages or retains access to the organization's keys, the provider may hold copies of key material or possess recovery mechanisms that remain capable of decrypting retained ciphertext. The organization therefore may be unable to demonstrate that decryption is irreversibly impossible after a deletion request. For cloud data deletion, assurance must cover both the encrypted data and all relevant cryptographic material across the service's operational and backup environments. CCSP-aligned practice is to evaluate key ownership, separation of duties, contractual commitments, and auditable key-destruction processes. Customer-controlled keys, including keys maintained in a customer-controlled or external key-management system when appropriate, provide stronger assurance that key destruction can render associated encrypted data inaccessible. NIST recognizes cryptographic erase as a sanitization technique when the media encryption key is effectively destroyed and cannot be recovered. See NIST SP 800-88 Rev. 1, Guidelines for Media Sanitization and Cloud Security Alliance Security Guidance v4.

QUESTION NO: 48

In a federated identity arrangement using a trusted third-party model, who is the identity provider and who is the relying party?

A. The users of the various organizations within the federations within the federation/a CASB

B. Each member organization/a trusted third party

C. Each member organization/each member organization

D. A contracted third party/the various member organizations of the federation

ANSWER: D

Explanation:

In the trusted third-party federation model, **A contracted third party/the various member organizations of the federation** correctly identifies the respective roles. The contracted third party acts as the identity provider because it performs authentication and supplies identity assertions or claims to the federation's participants. Its role is to provide a trusted authentication service that each participating organization can accept, avoiding the need for every organization to independently authenticate the same user or maintain separate credentials for every connected service.

The various member organizations are relying parties because they consume the identity provider's assertions when deciding whether to establish a session and grant a user access to their applications, services, or resources. This separation of duties is fundamental to federation: the identity provider authenticates the subject and communicates trusted identity information, while the relying party uses that information according to its access-control and authorization policies. In common federation protocols, such as SAML and OpenID Connect, the same relationship may be described as an identity provider and service provider/client, respectively. Trust is established in advance through federation agreements, metadata, certificates, and protocol configuration. See the OASIS Security Assertion Markup Language (SAML) standards and OpenID Connect Core.

QUESTION NO: 49

Which of the following are contractual components that the CSP should review and understand fully when contracting with a cloud service provider?

(Choose two.)

A. Concurrently maintainable site infrastructure

B. Use of subcontractors

C. Redundant site infrastructure capacity components

D. Scope of processing

ANSWER: B D

Explanation:

Use of subcontractors and **Scope of processing** are essential cloud-contract terms because they establish who may handle customer data and precisely how that handling is permitted. A provider commonly relies on subcontractors, often called subprocessors, for hosting, support, security operations, or other service functions. The contract should therefore establish authorization requirements, flow-down security and privacy obligations, accountability, and notification or objection processes for subprocessor changes. This gives the customer visibility into the full service-delivery chain and supports supplier-risk management.

The scope of processing must clearly state the subject matter and duration of processing, the nature and purpose of processing, the categories of data involved, and the categories of data subjects. These boundaries ensure that service-provider processing remains limited to documented customer instructions and that the parties can map contractual commitments to applicable privacy, security, retention, location, and compliance requirements. For personal data processing, Article 28 of the GDPR expressly requires a processor contract to contain these details and requires prior specific or general written authorization for engaging another processor. These provisions make both terms central to cloud governance and shared-responsibility oversight. See GDPR, Article 28 and the Cloud Security Alliance Cloud Controls Matrix.

QUESTION NO: 50

Which of the following is NOT considered a type of data loss?

A. Data corruption

B. Stolen by hackers

C. Accidental deletion

D. Lost or destroyed encryption keys

ANSWER: B

Explanation:

Stolen by hackers is the best answer because it primarily describes unauthorized acquisition or disclosure of information—a confidentiality compromise—rather than data loss. In a data-loss event, the organization's authoritative data is deleted, destroyed, corrupted, or otherwise becomes inaccessible or unusable. By contrast, an attacker can copy or exfiltrate data while the organization retains its original, usable data. The resulting incident is therefore principally a data breach, although it may have serious legal, privacy, financial, and operational consequences.

This distinction matters in CCSP practice because the security objectives and response priorities differ. A breach response focuses on identifying and containing unauthorized access, preserving evidence, assessing exposed data, and meeting notification obligations. Data-loss resilience focuses on maintaining recoverable copies, validating backup and restoration processes, protecting encryption keys, and ensuring data remains available and intact throughout its lifecycle. NIST defines a breach as the loss of control, compromise, unauthorized disclosure, unauthorized acquisition, unauthorized access, or similar term referring to situations in which an unauthorized person accesses or potentially accesses sensitive information. See the NIST CSRC definition of breach and NIST's Computer Security Incident Handling Guide.

QUESTION NO: 51

A user signs on to a cloud-based social media platform. In another browser tab, the user finds an article worth posting to the social media platform. The user clicks on the platform's icon listed on the article's website, and the article is

automatically posted to the user's account on the social media platform. This is an example of what?

Response:

A. Single sign-on

B. Insecure direct identifiers

C. Identity federation

D. Cross-site scripting

ANSWER: C

Explanation:

Identity federation is the best answer because the action relies on an established trust relationship that allows a user's identity and authorization context to be recognized across separate web domains. The user has already authenticated to the social-media service, and the article site can initiate an interaction with that service that results in content being posted under the user's account without requiring the user to enter credentials again. This illustrates the practical goal of federation: enabling interoperable identity-based access among independently operated services while the authoritative service continues to manage the user's session and permissions.

In modern cloud environments, this type of cross-service interaction is commonly supported by browser sessions together with token-based delegated authorization. OAuth 2.0 defines a framework through which a resource owner can authorize a third-party application to access protected resources or perform actions on the owner's behalf, without sharing the owner's password. Federated identity architectures similarly use standardized assertions or tokens to convey trusted identity and authorization information between parties. These mechanisms reduce repeated credential entry and support secure integration between cloud services. See RFC 6749: The OAuth 2.0 Authorization Framework and the OpenID Connect overview.

QUESTION NO: 52

Cloud vendors are held to contractual obligations with specified metrics by: Response:

A. SLAs

B. Regulations

C. Law

D. Discipline

ANSWER: A

Explanation:

SLAs are the contractual mechanism used to hold cloud vendors accountable for delivering agreed service outcomes against defined, measurable targets. A service-level agreement establishes specific commitments such as service availability, performance levels, support response times, maintenance notifications, incident communication requirements, and the methods used to measure and report those commitments. It also establishes accountability by documenting the scope of the service, roles and responsibilities, exclusions, escalation procedures, and the remedies available when the provider fails to meet its commitments, such as service credits or other contractual recourse. For CCSP governance and vendor-management purposes, an SLA translates a customer's business and technical requirements into objective criteria that can be monitored throughout the provider relationship. The SLA may be included directly in the cloud services contract or incorporated by reference, but it is the agreement that makes service metrics contractually enforceable between the customer and cloud service provider. NIST describes cloud service agreements as addressing service-level objectives and related contractual terms; cloud providers likewise publish SLAs that define availability commitments and applicable credits. See NIST SP 800-146, Cloud Computing Synopsis and Recommendations and AWS Compute Service Level Agreements.

QUESTION NO: 53

DLP can be combined with what other security technology to enhance data controls? Response:

A. DRM

B. SIEM

C. Kerberos

D. Hypervisors

ANSWER: A

Explanation:

DRM is the appropriate technology to combine with DLP because it adds persistent, data-centric protections to the content that DLP identifies and governs. Data loss prevention discovers sensitive information, classifies it, and applies policy when a user or process attempts to transmit, share, upload, or otherwise handle that information in a way that creates risk. Digital rights management extends protection after the data is accessed or distributed by applying rights directly to the document or file. These rights can include encryption, restrictions on viewing, editing, copying, printing, or forwarding, identity-based access requirements, and expiration or revocation of access. Together, DLP provides detection and policy enforcement at handling and egress points, while DRM maintains control over authorized use of sensitive data wherever it travels. This pairing is particularly relevant to cloud environments, where organizational data may be shared with external users or accessed from unmanaged locations. Microsoft describes DLP as identifying and protecting sensitive items across services and endpoints, while its rights-management capabilities apply encryption and usage rights that persist with protected content. See the Microsoft Purview DLP overview and Azure Rights Management overview.

QUESTION NO: 54

Which aspect of cloud computing makes data classification even more vital than in a traditional data center?

A. Interoperability

B. Virtualization

C. Multitenancy

D. Portability

ANSWER: C

Explanation:

Multitenancy is correct because cloud services commonly deliver shared computing, storage, network, and management infrastructure to multiple independent customers. Although the cloud provider uses logical isolation to separate tenants, an organization must still know the sensitivity, regulatory status, ownership, and handling requirements of its data before placing it into a cloud service. Data classification is the mechanism that drives those decisions. A classification such as public, internal, confidential, or regulated determines the appropriate safeguards, including identity and access controls, encryption and key-management requirements, monitoring, retention, backup handling, geographic residency restrictions, and approved service or deployment models.

In a multitenant environment, properly classifying data is especially important because the organization relies on provider and customer-configured controls to preserve separation from other tenants on shared resources. Classification enables the cloud customer to select and enforce controls proportionate to the data's impact level and compliance obligations, while maintaining accountability under the shared-responsibility model. NIST identifies multitenancy as a significant public-cloud security consideration and emphasizes that customers must evaluate protections for their data. See NIST SP 800-144, Guidelines on Security and Privacy in Public Cloud Computing and the ISC2 CCSP certification overview.

QUESTION NO: 55

Which of the following practices help protect cloud-management-plane access?

(Choose two.)

A. Requiring multifactor authentication for privileged accounts

B. Enforcing least-privilege roles

C. Using shared administrative accounts

D. Retaining long-lived access keys without rotation

ANSWER: A B

Explanation:

Requiring multifactor authentication for privileged accounts and **enforcing least-privilege roles** help protect cloud-management-plane access. Management-plane interfaces can create, modify, and delete cloud resources and can change security configurations. Compromise of a highly privileged management identity can therefore affect the confidentiality, integrity, and availability of an entire cloud environment.

Multifactor authentication reduces the likelihood that a stolen password alone can be used to obtain administrative access. Least-privilege roles limit identities to the specific permissions needed for approved tasks and reduce the impact of an account compromise or administrative error. Shared administrative accounts undermine accountability, while long-lived access keys increase exposure when they are not rotated or replaced with temporary credentials. See NIST SP 800-63-3, Digital Identity Guidelines and AWS IAM security best practices.

QUESTION NO: 56

When data discovery is undertaken, three main approaches or strategies are commonly used to determine what the type of data, its format, and composition are for the purposes of classification.

Which of the following is NOT one of the three main approaches to data discovery?

A. Content analysis

B. Hashing

C. Labels

D. Metadata

ANSWER: B

Explanation:

Hashing is the correct choice because it is a cryptographic transformation rather than a primary data-discovery strategy. A hash function produces a fixed-length digest from an input, allowing systems to compare a known value with a later value for integrity checking or matching. This can be useful operationally—for example, to identify a file that exactly matches a previously known file—but it does not determine an unknown data set's business meaning, sensitivity, format, or composition for classification purposes.

Data discovery and classification instead rely on examining information that describes or identifies the data. Content analysis inspects the actual values and patterns in data, such as national identification numbers, payment-card numbers, or health-related terms. Metadata supplies contextual characteristics, including file type, database field names, ownership, location, and creation attributes. Labels provide explicit classification or sensitivity markings that can be applied by users or automated systems and then used by security controls. Together, these approaches support identifying and classifying cloud data so that appropriate handling, access, retention, and protection controls can be selected. NIST describes secure hash functions as mechanisms for generating message digests in FIPS 180-4, Secure Hash Standard. For an example of classification based on sensitive-information types, trainable classifiers, and labels, see the Microsoft Purview data classification overview.

QUESTION NO: 57

Which cloud storage type uses an opaque value or descriptor to categorize and organize data? Response:

A. Volume

B. Object

C. Structured

D. Unstructured

ANSWER: B

Explanation:

Object storage is correct because it organizes data as independently addressable objects rather than as blocks in a volume or files in a hierarchical directory. Each object consists of the data, a unique identifier or key used to retrieve it, and metadata that describes the object. The identifier functions as an opaque value from the consumer's perspective: an application uses the object key or identifier to access the content without needing awareness of the underlying physical devices, disk sectors, or placement mechanisms. This abstraction lets cloud providers distribute and replicate objects across large-scale infrastructure while presenting a simple retrieval interface. Metadata also enables categorization, management, lifecycle controls, and policy application at object scale. This storage model is particularly well suited to cloud-native repositories of unstructured content, including backups, images, video, documents, log data, and data-lake assets. AWS describes object storage as storing data in objects that include the data, metadata, and a unique identifier, while Google Cloud Storage describes objects as immutable pieces of data stored in buckets and accessed through object names. See AWS: What is Object Storage? and Google Cloud Storage: Objects.