

DUMPS ARENA

Certified Cloud Security Professional (CCSP)

ISC2 CCSP

Version Demo

Total Demo Questions: 44

Total Premium Questions: 841

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Concepts, Architecture and Design	156
Topic 2, Cloud Data Security	188
Topic 3, Cloud Platform & Infrastructure Security	122
Topic 4, Cloud Application Security	100
Topic 5, Cloud Security Operations	105
Topic 6, Legal, Risk and Compliance	170
Total	841

QUESTION NO: 1

Data masking can be used to provide all of the following functionality, except:

- A. Secure remote access
- B. test data in sandboxed environments
- C. Authentication of privileged users
- D. Enforcing least privilege

ANSWER: C

Explanation:

Data masking is a data protection technique that obscures sensitive values (for example, PII, PHI, PAN) while preserving enough format and realism to keep the data usable for its intended purpose. In CCSP terms, it is primarily a confidentiality control applied to data at rest and data in use, commonly implemented as static data masking (creating a masked copy for non-production use) or dynamic data masking (masking results at query/presentation time based on policy). Because masking changes what data is revealed, it is well-suited to scenarios like providing realistic but non-sensitive test data in sandboxed environments, and it can support least-privilege goals by limiting what fields/values a given user or role can see without changing the underlying dataset. However, data masking is not an identity and access management function: it does not verify identity, establish trust, or perform credential validation. Authentication of privileged users is handled by IAM controls such as MFA, strong authentication protocols, and privileged access management—not by masking. For more on dynamic data masking concepts and use cases, see [Microsoft Learn: Dynamic Data Masking](#) and for broader data masking definitions and approaches see [NIST CSRC Glossary: Data Masking](#).

QUESTION NO: 2

All of the following are activities that should be performed when capturing and maintaining an accurate, secure system baseline, except

Response:

- A. Audit the baseline to ensure that all configuration items have been included and applied correctly
- B. Impose the baseline throughout the environment
- C. Capture an image of the baseline system for future reference/versioning/rollback purposes
- D. Document all baseline configuration elements and versioning data

ANSWER: B

Explanation:

“Impose the baseline throughout the environment” is the exception because it describes a deployment/enforcement activity rather than an activity specifically focused on capturing and maintaining an accurate, secure system baseline. In baseline management, the core work is to define the approved configuration state, record it in a controlled way (including versioning), and verify it remains accurate over time through auditing and change control. Capturing a “gold image” (or equivalent infrastructure-as-code template) and documenting configuration elements and version history are classic baseline-capture and baseline-maintenance practices. Auditing the baseline to confirm all configuration items are included and correctly applied is also part of maintaining baseline integrity and ensuring the baseline remains trustworthy as a reference point for compliance and drift detection. By contrast, imposing/enforcing that baseline across an environment is an operational rollout step (often handled via configuration management, policy-as-code, or orchestration tooling) that comes after the baseline has been established; it is not inherently required to “capture and maintain” the baseline itself. This aligns with common security configuration baseline practices such as those described by NIST guidance on configuration management and security-focused baselines.

References: [NIST SP 800-128 \(Guide for Security-Focused Configuration Management\)](#), [NIST SP 800-53 Rev. 5 \(Configuration Management family\)](#)

QUESTION NO: 3

You are the security manager for a small application development company. Your company is considering the use of the cloud for software testing purposes. Which cloud service model is most likely to suit your needs?

Response:

- A. IaaS
- B. PaaS
- C. SaaS
- D. LaaS

ANSWER: B

Explanation:

PaaS is the cloud service model most likely to fit a small application development company's needs for software testing because it provides an application runtime and development/testing environment without requiring the customer to build and manage the underlying infrastructure. With PaaS, the cloud provider manages the servers, storage, networking, virtualization, and typically the operating system and middleware, while the customer focuses on deploying code, configuring the platform, and running test workloads. This aligns well with software testing use cases where teams need rapid provisioning of environments, consistent build/test pipelines, and easy scaling up/down for short-lived test runs. From a security management perspective, PaaS also helps reduce operational burden and risk associated with patching and hardening lower layers, allowing the organization to concentrate controls on application security, identity and access management, secure configuration, and data handling within the platform. This division of responsibilities matches the shared responsibility model commonly emphasized in CCSP, where the provider secures the platform layers and the customer secures what they deploy and configure. See NIST's cloud service model definitions in [NIST SP 800-145](#) and an overview of shared responsibility in [Microsoft's shared responsibility model](#).

QUESTION NO: 4

During the course of an audit, which of the following would NOT be an input into the control requirements used as part of a gap analysis.

- A. Contractual requirements
- B. Regulations
- C. Vendor recommendations
- D. Corporate policy

ANSWER: C

Explanation:

Vendor recommendations is the correct choice because a gap analysis in an audit context compares the organization's current control environment against defined, authoritative control requirements. Those requirements come from sources that are mandatory or formally adopted by the organization, such as applicable laws and regulations, binding contractual obligations (including customer contracts, SLAs, and flow-down requirements), and internal corporate policies/standards that management has approved and expects to be followed. In contrast, vendor recommendations (for example, hardening guides, reference architectures, or "best practice" configuration suggestions) can be useful implementation guidance, but they are not inherently required unless the organization has explicitly incorporated them into policy/standards or they are referenced in a contract or regulatory framework. Therefore, they are not a primary input into the control requirements baseline used for audit gap analysis; they may inform remediation options after gaps are identified, but they do not define the required state by themselves. This aligns with common audit practice of auditing against criteria that are objective, agreed upon, and enforceable rather than optional guidance. See ISO's overview of management system auditing concepts and

criteria at <https://www.iso.org/iso-19011-guidelines-for-auditing-management-systems.html> and NIST's discussion of security control baselines and tailoring at <https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final>.

QUESTION NO: 5

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

- A. Breach alert
- B. Confidentiality
- C. Communication redundancy
- D. Line conditioning

ANSWER: D

Explanation:

“Line conditioning

” is correct because many uninterruptible power supplies do more than provide short-term battery power during an outage. A common additional function is power conditioning (often called line conditioning), where the UPS helps deliver cleaner, more stable electrical power to connected equipment. This can include smoothing sags and surges, reducing electrical noise, and regulating voltage to keep it within acceptable tolerances. In practice, this capability helps protect sensitive IT hardware from power quality problems that can cause unexpected reboots, data corruption, premature component wear, or intermittent faults that are difficult to troubleshoot. From an availability and resilience perspective—topics emphasized in CCSP operational and infrastructure security—line conditioning is a key benefit because it reduces downtime and prevents damage even when utility power is present but unstable. Many UPS designs (especially line-interactive and online/double-conversion models) explicitly incorporate voltage regulation and conditioning features as part of their normal operation, not just during battery use.

References: [CISA – Uninterruptible Power Supply \(UPS\)](#), [APC/Schneider Electric FAQ – UPS and power conditioning concepts](#)

QUESTION NO: 6

Which concept pertains to cloud customers paying only for the resources they use and consume, and only for the duration they are using them?

Response:

- A. Measured service
- B. Auto-scaling
- C. Portability
- D. Elasticity

ANSWER: A

Explanation:

Measured service is the cloud concept that directly describes pay-per-use and pay-for-what-you-consume billing. In a measured service model, the cloud provider automatically meters resource usage (such as CPU time, storage capacity, network bandwidth, or number of active user accounts) and reports it in a way that supports transparency for both the provider and the customer. This metering enables customers to be charged only for the amount of service actually consumed and typically aligns with time-based consumption (for example, per second/minute/hour for compute instances) and usage-based units (for example, GB-month for storage). This is one of the essential characteristics of cloud computing

defined by NIST and is foundational to the economic value proposition of cloud: shifting from fixed, upfront capacity planning to variable operational expense tied to real consumption. While other cloud concepts can influence cost (for example, scaling features can change how much you consume), the defining characteristic that enables “only pay for what you use, for the duration you use it” is measured service and its associated metering and billing mechanisms.

References: [NIST SP 800-145 \(The NIST Definition of Cloud Computing\)](#), [Google Cloud: What is cloud computing?](#)

QUESTION NO: 7

DLP solutions typically involve all of the following aspects except . Response:

- A. Data discovery
- B. Tokenization
- C. Monitoring
- D. Enforcement

ANSWER: B

Explanation:

Tokenization is the best “except” choice because it is not a typical core aspect of Data Loss Prevention (DLP) solutions. DLP programs and tools are generally built around finding where sensitive data exists (data discovery/classification), observing how it is used and moved (monitoring of data in use, in motion, and at rest), and then taking action based on policy (enforcement such as blocking, quarantining, alerting, encrypting, or applying rights controls). These capabilities align with the common DLP lifecycle: identify sensitive data, detect policy violations, and prevent or respond to exfiltration or misuse.

Tokenization, by contrast, is a data protection technique that replaces sensitive values with non-sensitive tokens, typically used to reduce exposure of regulated data (for example, payment card data) within applications and databases. While tokenization can be used alongside DLP as part of a broader data protection strategy, it is not itself a standard DLP functional pillar in the way discovery, monitoring, and enforcement are.

References: [NIST SP 800-53 Rev. 5 \(Data Loss Prevention control family\)](#), [Microsoft Purview: Learn about data loss prevention](#).

QUESTION NO: 8

The cloud deployment model that features joint ownership of assets among an affinity group is known as: Response:

- A. Private
- B. Public
- C. Hybrid
- D. Community

ANSWER: D

Explanation:

The correct answer is

Community

because a community cloud is specifically defined as a deployment model where the cloud infrastructure is provisioned for exclusive use by a specific community of consumers (an “affinity group”) that share common concerns. Those shared concerns commonly include mission objectives, security requirements, policy/compliance obligations, and governance needs. In a community cloud, the infrastructure and assets can be jointly owned, managed, and governed by one or more

members of the community, a third party, or some combination—making “joint ownership among an affinity group” a defining characteristic. This model is often used when multiple organizations need to collaborate and standardize controls while still keeping the environment restricted to that community, such as government agencies, healthcare consortiums, or financial sector groups with aligned regulatory requirements. This aligns with the widely accepted NIST cloud deployment model definitions used throughout CCSP-aligned materials and industry practice. For the canonical definition of community cloud and its ownership/management possibilities, see NIST SP 800-145 and related NIST cloud computing resources.

References: [NIST SP 800-145 \(The NIST Definition of Cloud Computing\)](#), [NIST SP 800-146 \(Cloud Computing Synopsis and Recommendations\)](#)

QUESTION NO: 9

Which of the following is a valid risk management metric?

- A. KPI
- B. KRI
- C. SOC
- D. SLA

ANSWER: B

Explanation:

“KRI” is a valid risk management metric because it stands for Key Risk Indicator, a commonly used measurement that provides early warning signals about changes in an organization’s risk exposure. In practical risk programs, KRIs are defined, tracked, and trended to help stakeholders detect when risk is increasing (or controls are degrading) so that timely investigation and response can occur. KRIs are typically tied to specific risk scenarios and risk appetite/tolerance thresholds, and they are monitored over time (often with trigger levels) to support governance reporting and decision-making. In cloud security and broader enterprise risk management, examples include rates of privileged access exceptions, percentage of critical vulnerabilities past SLA, frequency of misconfiguration findings, or third-party incidents—each serving as an indicator that risk conditions are shifting. This aligns with standard risk management practices where metrics are used to communicate risk posture and prompt action before losses occur. For additional context on risk indicators and their role in risk monitoring, see [ISACA: Key Risk Indicators](#) and [NIST SP 800-55 Rev.1 \(Performance Measurement Guide for Information Security\)](#).

QUESTION NO: 10

You work for a government research facility. Your organization often shares data with other government research organizations.

You would like to create a single sign-on experience across the organizations, where users at each organization can sign in with the user ID/authentication issued by that organization, then access research data in all the other organizations.

Instead of replicating the data stores of each organization at every other organization (which is one way of accomplishing this goal), you instead want every user to have access to each organization’s specific storage resources.

In order to pass the user IDs and authenticating credentials of each user among the organizations, what protocol/language/motif will you most likely utilize?

Response:

- A. Representational State Transfer (REST)
- B. Security Assertion Markup Language (SAML)
- C. Simple Object Access Protocol (SOAP)
- D. Hypertext Markup Language (HTML)

ANSWER: B

Explanation:

Security Assertion Markup Language (SAML) is the most likely choice because the scenario describes federated single sign-on across separate organizations (separate security domains) where each organization remains the authoritative identity provider for its own users. In a federation model, the user authenticates to their home organization, which then issues a signed assertion containing identity and authentication statements that a partner organization (the service provider) can trust and consume. This is exactly what SAML was designed for: exchanging authentication and authorization data between an identity provider and a service provider, enabling cross-organization SSO without replicating user directories or credentials. In government and other large enterprise environments, SAML is widely used for browser-based SSO and inter-organization trust relationships because it supports strong signing, optional encryption, and well-established metadata/trust configuration patterns. This aligns with the requirement to “pass the user IDs and authenticating credentials” in a secure, standardized way while keeping authentication at the issuing organization. See the OASIS SAML overview at <https://www.oasis-open.org/standards/saml/> and a practical federation explanation in Microsoft's identity platform documentation at <https://learn.microsoft.com/en-us/entra/identity-platform/saml-protocol-reference>.

QUESTION NO: 11

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 the cloud-related concept that best matches the idea of being able to reuse application components and services for other purposes, because it focuses on the ability of systems, platforms, and services to work together through well-defined interfaces and standards. In cloud environments, reusable components are typically exposed as services (for example, via APIs, microservices, or service-oriented architectures). When those services are interoperable, they can be consumed by different applications, across different environments, and potentially across different cloud providers without requiring extensive rework. This enables practical reuse: a service built for one application can be integrated into another workflow, combined with other services, or orchestrated into new business processes. Interoperability is therefore closely tied to standard protocols, API compatibility, and integration patterns that allow components to be leveraged beyond their original context. This aligns with common CCSP expectations around cloud architecture and design, where service integration and API-driven composition are key enablers of reuse in distributed cloud systems.

References: [NIST SP 800-145 \(Cloud Computing Definition\)](#), [ISO/IEC 17788:2014 Cloud computing — Overview and vocabulary](#)

QUESTION NO: 12

Which of the following best describes data masking? Response:

- A. A method where the last few numbers in a dataset are not obscured. These are often used for authentication.
- B. A method for creating similar but inauthentic datasets used for software testing and user training.
- C. A method used to protect prying eyes from data such as social security numbers and credit card data.
- D. Data masking involves stripping out all similar digits in a string of numbers so as to obscure the original number.

ANSWER: B

Explanation:

Data masking is a technique used to hide sensitive information by replacing it with a non-sensitive substitute while preserving the data's general format and usefulness. The key idea is that the masked data should look realistic enough to be used safely in non-production contexts (such as development, QA, analytics, demonstrations, and training) without exposing real personally identifiable information or regulated data. This is why the best description is "A method for creating similar but inauthentic datasets used for software testing and user training." It captures both the purpose (protecting sensitive values) and the practical outcome (a realistic but fake dataset that still supports application behavior and testing). In cloud environments, data masking is commonly part of broader data security controls to reduce exposure when copying production data into lower environments or when granting limited access to datasets. Masking can be static (creating a masked copy) or dynamic (masking on-the-fly at query time), but in both cases the goal is to prevent disclosure while maintaining utility.

References: [Microsoft Learn: Dynamic Data Masking](#), [Google Cloud DLP: Data masking concepts](#)

QUESTION NO: 13

Which if the following is NOT one of the three components of a federated identity system transaction?

- A. Relying party
- B. Identity provider
- C. User
- D. Proxy relay

ANSWER: D

Explanation:

In a standard federated identity transaction, the core actors are the user (often called the principal), the identity provider (IdP), and the relying party (RP), also commonly referred to as the service provider. The user attempts to access a protected resource at the relying party, the relying party redirects or otherwise initiates an authentication request to the identity provider, and the identity provider authenticates the user and issues an assertion or token (for example, a SAML assertion or an OpenID Connect ID token). The relying party then consumes that assertion/token to establish a session and authorize access. These three roles are the essential components that must exist for federation to function, regardless of the specific protocol used (SAML, OIDC, WS-Federation). A "proxy relay" is not one of the canonical three components; while intermediaries like federation gateways, brokers, or reverse proxies can exist in some architectures, they are optional implementation details rather than a required component of the federated identity transaction model.

References: [NIST SP 800-63C \(Federation and Assertions\)](#), [Microsoft OpenID Connect documentation](#).

QUESTION NO: 14

Your organization is considering a move to a cloud environment and is looking for certifications or audit reports from cloud providers to ensure adequate security controls and processes.

Which of the following is NOT a security certification or audit report that would be pertinent? Response:

- A. FedRAMP
- B. PCI DSS
- C. FIPS 140-2
- D. SOC Type 2

ANSWER: C

Explanation:

FIPS 140-2 is the item that is NOT a security certification or audit report focused on a cloud provider's overall security controls and operational processes. FIPS 140-2 is a U.S./Canadian government standard for validating cryptographic modules (for example, a specific HSM, software crypto library, or appliance) against defined security requirements. While it can be relevant when you need assurance that encryption components used by a cloud service meet government-grade validation, it does not provide an organization-wide assessment of the provider's governance, risk management, control environment, change management, incident response, or other operational control domains. In contrast, cloud assurance artifacts typically sought during provider due diligence are broad control attestations or authorization frameworks that evaluate the provider's security program and processes over time. Therefore, when the goal is to obtain certifications or audit reports demonstrating adequate security controls and processes at the provider level, FIPS 140-2 is not the most pertinent type of report/certification in that category.

References: [NIST Cryptographic Module Validation Program \(CMVP\)](#), [NIST FIPS 140-2 publication](#)

QUESTION NO: 15

DAST checks software functionality in . Response:

- A. The production environment
- B. A runtime state
- C. The cloud
- D. An IaaS configuration

ANSWER: B

Explanation:

DAST checks software functionality in a runtime state because Dynamic Application Security Testing evaluates an application while it is executing, observing its behavior from the outside (often by sending requests to a running web application and analyzing responses). This aligns with how DAST is commonly used to find issues such as injection flaws, authentication/session weaknesses, and misconfigurations that only become apparent when the application is deployed and processing real inputs. Unlike static approaches that analyze source code or binaries without execution, DAST requires the application to be running in an environment where it can be exercised (for example, a test/staging environment or even production in some mature programs), but the defining characteristic is that the assessment occurs during execution. In CCSP terms, this maps to security testing practices within the SDLC and operational assurance: DAST provides evidence of how the application behaves under real runtime conditions, which is essential for validating controls and identifying exploitable vulnerabilities. This is also consistent with industry definitions of DAST as "black-box" or "outside-in" testing against a live application endpoint.

References: [OWASP Source Code Analysis Tools \(includes DAST context\)](#), [MITRE CWE Terminology \(static vs dynamic concepts\)](#)

QUESTION NO: 16

Which of the following would NOT be included as input into the requirements gathering for an application or system?
Response:

- A. Users
- B. Management
- C. Regulators
- D. Auditors

ANSWER: D

Explanation:

Auditors would NOT typically be included as a direct input source during requirements gathering for an application or system. Requirements gathering is primarily driven by stakeholders who define business objectives, operational needs, user workflows, and external obligations that the system must satisfy. Users and management are core stakeholders because they articulate functional needs, priorities, constraints, and success criteria. Regulators (or regulatory requirements) are also a key input because they impose mandatory compliance obligations (for example, privacy, financial, or industry-specific rules) that must be translated into security, retention, logging, and control requirements.

Auditors, by contrast, generally validate and assess whether implemented controls and processes meet stated requirements, policies, and applicable standards. Their role is more commonly associated with assurance activities after requirements have been defined and controls have been designed/implemented, rather than being a primary source of initial requirements. While audit findings can influence future requirement updates, auditors are not usually part of the initial requirements elicitation process for a new system.

References: [ISC2 CCSP Certification Overview](#), [NIST SP 800-160 Vol. 1 \(Systems Security Engineering\)](#)

QUESTION NO: 17

Which of the following is a risk that stems from a virtualized environment? Response:

- A. Live virtual machines in the production environment are moved from one host to another in the clear.
- B. Cloud data centers can become a single point of failure.
- C. It is difficult to find and contract with multiple utility providers of the same type (electric, water, etc.).
- D. Modern SLA demands are stringent and very hard to meet.

ANSWER: A

Explanation:

“Live virtual machines in the production environment are moved from one host to another in the clear.” is a virtualization-specific risk because it describes exposure introduced by hypervisor features such as live migration (e.g., vMotion or similar mechanisms). In a virtualized environment, workloads can be moved between physical hosts to support maintenance, load balancing, and high availability. If the live migration channel is not properly isolated and protected (for example, using a dedicated migration network, strong segmentation, and encryption), sensitive in-memory state and VM execution context can be intercepted or tampered with while in transit. This risk is directly tied to virtualization management planes and east-west traffic patterns that do not exist in the same way for non-virtualized, single-host systems. CCSP emphasizes understanding hypervisor/virtualization threats, including inter-host migration exposure and management network security, as part of securing cloud infrastructure and virtualized compute. Practical mitigations include encrypting vMotion/migration traffic, restricting who can initiate migrations, and isolating management and migration networks from tenant/user networks.

References: [Microsoft Hyper-V Live Migration overview](#), [VMware vSphere Security \(vMotion and migration traffic protection\)](#)

QUESTION NO: 18

Which aspect of cloud computing pertains to cloud customers only paying for the resources and services they actually use?

- A. Metered service
- B. Measured billing
- C. Metered billing
- D. Measured service

ANSWER: D

Explanation:

The correct aspect is measured service. In the NIST essential characteristics of cloud computing, measured service means the cloud provider automatically controls and optimizes resource use by leveraging a metering capability appropriate to the type of service (for example, storage, processing, bandwidth, and active user accounts). This metering enables transparency for both the provider and consumer, and it directly supports the pay-as-you-go model where customers are charged based on actual consumption rather than fixed capacity. In practice, measured service underpins usage-based billing, chargeback/showback, and cost allocation, because consumption is monitored, reported, and often billed per unit (such as per vCPU-hour, GB-month, or requests). This is distinct from general “metered” wording; the recognized cloud-computing term in the canonical definition is specifically “measured service.” For CCSP purposes, tying “only paying for what you use” to the NIST characteristic is the key: metering and reporting of usage is what enables consumption-based pricing and accountability in shared cloud environments.

References: [NIST SP 800-145 \(The NIST Definition of Cloud Computing\)](#), [NIST CSRC publication page for SP 800-145](#).

QUESTION NO: 19

SOC 2 reports were intended to be . Response:

- A. Released to the public
- B. Only technical assessments
- C. Retained for internal use
- D. Nonbinding

ANSWER: C

Explanation:

SOC 2 reports were intended to be retained for internal use in the sense that they are not general-public reports; they are “restricted use” reports meant for a defined set of knowledgeable parties (typically the service organization’s management, existing customers, prospective customers, and their auditors/business partners) who need assurance over controls aligned to the Trust Services Criteria. Unlike SOC 3, which is designed for broad distribution and marketing-style public sharing, SOC 2 contains detailed descriptions of the system, the auditor’s tests, and results that can create security and confidentiality concerns if widely published. In cloud due diligence, this restricted distribution model supports risk-based vendor assessment while limiting unnecessary exposure of control details. Practically, organizations often treat SOC 2 as confidential and share it under NDA or controlled portals, aligning with the intent that it supports internal governance and customer assurance rather than public release. This is consistent with AICPA guidance distinguishing SOC 2 (restricted use) from SOC 3 (general use). See the AICPA overview of SOC reporting and intended users: <https://www.aicpa-cima.com/resources/landing/system-and-organization-controls-soc-suite-of-services> and the SOC 2 description and use context: <https://www.aicpa-cima.com/topic/audit-assurance/service-organization-controls>.

QUESTION NO: 20

You have been tasked by management to offload processing and validation of incoming encoded data from your application servers and their associated APIs. Which of the following would be the most appropriate device or software to consider?

Response:

- A. XML accelerator
- B. XML firewall
- C. Web application firewall
- D. Firewall

ANSWER: A

Explanation:

An **XML accelerator** is specifically designed to offload CPU-intensive XML and related encoded-message processing from application servers and API tiers. In many enterprise and cloud-integrated architectures, XML (and similar structured payloads) requires parsing, schema validation, canonicalization, transformation (for example XSLT), and sometimes cryptographic operations tied to message-level security. These tasks can be computationally expensive and can become a bottleneck on application servers, reducing throughput and increasing latency. By placing an XML accelerator in front of the application/API layer, you can move that parsing and validation workload to specialized software or hardware optimized for these operations, improving performance and scalability while keeping the application servers focused on business logic. This aligns with CCSP architectural best practices of using purpose-built components to handle protocol/message processing at the edge of the application tier and to reduce load on core services. For additional context on XML validation concepts and schema-based validation (which accelerators commonly offload), see [W3C XML Schema](#) and Microsoft's overview of XML schema validation in application stacks at [Microsoft Learn](#).

QUESTION NO: 21

Access should be based on . Response:

- A. Regulatory mandates
- B. Business needs and acceptable risk
- C. User requirements and management requests
- D. Optimum performance and security provision

ANSWER: B

Explanation:

“Business needs and acceptable risk” is the best basis for access decisions because access control is fundamentally a risk-management activity aligned to organizational objectives. In CCSP/ISC2 terms, access should be granted according to the principle of least privilege and need-to-know, which requires understanding what the business is trying to achieve (business need) and what level of exposure the organization is willing to tolerate (acceptable risk). This framing also supports consistent governance: access is not granted simply because someone asks, nor solely because a regulation exists, but because the organization has determined that enabling a specific activity is necessary and that the residual risk after controls (authentication strength, authorization model, monitoring, segmentation, etc.) is within risk appetite/tolerance. In cloud environments, this approach maps directly to designing IAM roles, policies, and conditional access based on workload criticality, data classification, and threat models, then validating that the resulting risk is acceptable to stakeholders. This is consistent with ISC2's emphasis on aligning security controls to business requirements and risk decisions. See NIST's access control guidance in [NIST SP 800-53 Rev. 5](#) (AC family) and the overview of least privilege/need-to-know concepts in [NIST SP 800-12 Rev. 1](#).

QUESTION NO: 22

All of the following might be used as data discovery characteristics in a content-analysis-based data discovery effort except .

Response:

- A. Keywords
- B. Pattern-matching
- C. Frequency
- D. Inheritance

ANSWER: D

Explanation:

“Inheritance

” is the correct exception because content-analysis-based data discovery focuses on inspecting the actual content of files/objects to identify sensitive data, rather than relying on access control or classification properties that may be inherited from a parent container. In practice, content analysis techniques include searching for keywords (e.g., “confidential”, “SSN”), using pattern matching/regular expressions (e.g., credit card formats, national identifiers), and applying statistical or heuristic methods such as frequency analysis to detect likely sensitive content (for example, repeated occurrences of specific token types or data elements). Inheritance, by contrast, is a metadata/permission behavior (e.g., inherited labels, inherited ACLs, inherited folder permissions) and is not a characteristic derived from analyzing the content itself. CCSP-aligned data discovery discussions commonly distinguish content inspection methods from metadata- or context-based approaches; inheritance belongs to the latter category and therefore does not fit as a content-analysis characteristic. For additional context on content inspection approaches used in data discovery/DLP, see [Microsoft Purview DLP policies](#) and [Google Cloud Sensitive Data Protection inspection concepts](#).

QUESTION NO: 23

Which of the following is an example of useful and sufficient data masking of the string “CCSP”? Response:

- A. XCSP
- B. PSCC
- C. TtLp
- D. 3X91

ANSWER: C**Explanation:**

Data masking is intended to de-identify sensitive values while still preserving enough characteristics to keep the data useful for non-production purposes (such as testing, analytics, or troubleshooting). “Useful and sufficient” masking typically means the original value cannot be feasibly reconstructed from the masked output, while the masked output still looks like the original data in terms of format and length. For the string “CCSP”, a good masked value should preserve the same character count and general data type (alphabetic characters), but replace the original characters so that the original is not revealed. The response “TtLp” is a good example because it maintains a four-character alphabetic string, which supports application validation and downstream processing, while not exposing the original “CCSP”. This aligns with common masking/tokenization goals described in cloud security guidance: reduce exposure of sensitive data while maintaining functional utility for permitted use cases. In CCSP terms, this supports data security controls that minimize data disclosure and reduce risk when using lower-trust environments or broader access groups.

References: [NIST Glossary – Data Masking](#), [Microsoft Learn – Dynamic Data Masking](#)

QUESTION NO: 24

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:**

The objectives of change management (often aligned with ITIL change enablement and widely adopted in cloud/IT operations) focus on ensuring that changes to services and infrastructure are handled in a way that delivers business value while minimizing risk. A core objective is to respond to changing business requirements while maximizing value and reducing incidents, disruption, and rework—this captures the balance between agility and stability that change management is meant to provide. Another key objective is ensuring that changes are recorded and evaluated so there is traceability, accountability, and an informed decision process (for example, assessing risk, impact, and authorization needs). Finally, change management aims to ensure changes are prioritized, planned, tested, implemented, documented, and reviewed in a controlled manner. This end-to-end control supports governance, reduces the likelihood of outages or security regressions, and enables continuous improvement through post-implementation review. These objectives are especially important in cloud environments where frequent releases and infrastructure-as-code can increase change velocity, making disciplined recording, evaluation, and controlled execution essential to maintaining availability and security. See ITIL-aligned guidance on change enablement objectives and practices in [AXELOS ITIL](#) and an example of structured change control in [Microsoft Cloud Adoption Framework change management](#).

QUESTION NO: 25

With an application hosted in a cloud environment, who could be the recipient of an eDiscovery order?

- A. Users
- B. Both the cloud provider and cloud customer
- C. The cloud customer
- D. The cloud provider

ANSWER: B

Explanation:

Both the cloud provider and cloud customer is correct because eDiscovery orders (or similar legal process such as subpoenas, preservation orders, or litigation holds) can be served on any entity that has custody, control, or possession of potentially relevant electronically stored information. In cloud deployments, the cloud customer typically owns the data and determines its processing purposes, but the cloud provider often has physical/administrative custody of the underlying infrastructure and may operate the platforms and tooling needed to preserve, search, and export data. As a result, depending on jurisdiction, contract terms, and where the data resides, either party may be directly compelled by a court or regulator, and practical compliance frequently requires coordinated action (e.g., provider-assisted collection, chain of custody, and preservation while the customer identifies scope and authorizes access). This aligns with CCSP legal and compliance expectations that responsibilities are shared and must be clearly addressed in contracts, incident/legal response procedures, and governance processes to ensure timely, defensible eDiscovery.

References: [Microsoft Purview eDiscovery overview](#), [Google Cloud eDiscovery and legal process guidance](#).

QUESTION NO: 26

The Cloud Security Alliance (CSA) publishes the Notorious Nine, a list of common threats to organizations participating in cloud computing.

According to the CSA, what is one reason the threat of insecure interfaces and APIs is so prevalent in cloud computing?
Response:

- A. Most of the cloud customer's interaction with resources will be performed through APIs.
- B. APIs are inherently insecure.
- C. Attackers have already published vulnerabilities for all known APIs.
- D. APIs are known carcinogens.

ANSWER: A

Explanation:

The statement “Most of the cloud customer’s interaction with resources will be performed through APIs.” is correct because cloud services are fundamentally operated and managed through software-defined control planes exposed via web consoles and, more importantly, APIs. Provisioning compute, configuring storage, managing identity and access, deploying applications, and automating operations are typically performed through these interfaces. As a result, APIs become a high-value and high-frequency attack surface: if authentication, authorization, input validation, logging, or transport protections are weak, an attacker can directly manipulate cloud resources or extract sensitive data. The CSA highlights that the security and availability of cloud services depend heavily on the security of these interfaces, since customers and third parties rely on them for orchestration and integration. This heavy dependence makes insecure interfaces and APIs a prevalent threat in cloud computing environments, especially where organizations expose management endpoints, use overly permissive tokens/keys, or integrate multiple services via API calls.

References: [Cloud Security Alliance – Top Threats to Cloud Computing](#), [OWASP API Security Project](#)

QUESTION NO: 27

Over time, what is a primary concern for data archiving?

- A. Size of archives
- B. Format of archives
- C. Recoverability
- D. Regulatory changes

ANSWER: C

Explanation:

Recoverability is a primary long-term concern for data archiving because archives only provide value if the organization can reliably restore and read the data throughout the entire retention period. As storage media ages, bit rot and hardware failures become more likely; as platforms and applications evolve, the tools, keys, and dependencies needed to decrypt, decompress, interpret, and restore archived data can become unavailable. Effective archiving therefore requires ongoing assurance that data remains intact and usable, including periodic restore testing, integrity verification (e.g., checksums), key management/escrow for encrypted archives, and migration/refresh strategies to newer storage systems before obsolescence. In cloud contexts, this also includes ensuring the organization can retrieve data from long-term storage tiers within required time objectives and that the archived content remains accessible despite changes in services, APIs, or account structures. This focus aligns with core security and resilience principles emphasized in CCSP: maintaining availability and ensuring data can be retrieved when needed for legal hold, eDiscovery, audits, or business continuity. Practical guidance on long-term preservation and ensuring continued accessibility is reflected in standards and industry recommendations for digital preservation and records retention. See [ISO 14721 \(OAIS\) overview](#) and NIST guidance on storage/media reliability and integrity considerations in [NIST SP 800-209](#).

QUESTION NO: 28

The most pragmatic option for data disposal in the cloud is which of the following?

- A. Cryptoshredding
- B. Overwriting
- C. Cold fusion
- D. Melting

ANSWER: A

Explanation:

Cryptoshredding is the most pragmatic option for data disposal in the cloud because cloud customers typically cannot reliably control or verify physical media handling, block-level overwrites, or the full lifecycle of replicas, snapshots, and backups across distributed storage systems. In many cloud architectures, data may be replicated across multiple devices and locations for resilience, and storage virtualization/abstraction makes it difficult (or impossible) for a tenant to target every physical block that ever held the data. Cryptoshredding addresses these realities by rendering data unreadable through secure destruction of the encryption keys protecting it. When strong encryption is used and keys are properly managed (e.g., centralized KMS/HSM controls, separation of duties, rotation, and auditable key deletion), destroying the keys effectively makes all encrypted copies—primary, replicated, cached, and backed up—computationally infeasible to recover. This aligns with common cloud security guidance that emphasizes encryption and key management as practical compensating controls for media sanitization challenges in multi-tenant environments. As a result, cryptoshredding is widely recognized as a feasible, scalable approach for cloud data disposal when implemented with robust cryptographic and key management practices.

References: [NIST SP 800-88 Rev.1 \(Media Sanitization\)](#), [Google Cloud KMS key deletion \(crypto-erasure concept\)](#)

QUESTION NO: 29

What are third-party providers of IAM functions for the cloud environment?

- A. AESs
- B. SIEMs
- C. DLPs
- D. CASBs

ANSWER: D

Explanation:

CASBs are commonly used as third-party providers of identity and access management-related functions in cloud environments because they sit between cloud consumers and cloud service providers to enforce enterprise security policies. In practice, CASBs can integrate with identity providers and cloud applications to provide visibility into cloud usage, apply access controls (including conditional access and session controls), enforce authentication/authorization policy decisions, and support governance requirements such as monitoring user activity and detecting anomalous behavior. While CASBs are broader than pure IAM (they also cover data security and threat protection), they are a well-established third-party control point for cloud access and policy enforcement, which is why they are frequently referenced in CCSP materials as a key cloud security control for managing access to SaaS and other cloud services. This aligns with industry guidance that positions CASBs as a security policy enforcement layer for cloud services, often complementing or extending native IAM capabilities of cloud providers and enterprise IdPs.

References: [NIST SP 800-210 \(Cloud Security\)](#), [Microsoft Defender for Cloud Apps \(CASB\) overview](#).

QUESTION NO: 30

With the rapid emergence of cloud computing, very few regulations were in place that pertained to it specifically, and organizations often had to resort to using a collection of regulations that were not specific to cloud in order to drive audits and policies.

Which standard from the ISO/IEC was designed specifically for cloud computing?

- A. ISO/IEC 27001
- B. ISO/IEC 19889

C. ISO/IEC 27001:2015

D. ISO/IEC 27018

ANSWER: D

Explanation:

ISO/IEC 27018 is the ISO/IEC standard designed specifically for cloud computing, focusing on the protection of personally identifiable information (PII) in public cloud environments acting as PII processors. In CCSP terms, it provides cloud-specific privacy controls and implementation guidance that extend a general information security management system by addressing risks unique to multi-tenant, outsourced processing—such as limitations on secondary use of customer data, transparency around data location and subcontractors, breach notification expectations, and customer ability to access, correct, and delete PII where applicable. This makes it directly applicable when an organization is selecting cloud providers, defining contractual privacy requirements, and preparing for cloud-focused audits and assurance activities. While ISO/IEC 27001 (including its 2015 edition) is foundational for an ISMS, it is not cloud-specific; ISO/IEC 27018 fills that gap by tailoring privacy controls to public cloud service providers and their customers. See ISO's overview of ISO/IEC 27018 here: <https://www.iso.org/standard/76559.html> and an accessible summary of its cloud privacy focus here: <https://www.itgovernance.co.uk/iso27018>.

QUESTION NO: 31

Federation allows across organizations. Response:

- A. Role replication
- B. Encryption
- C. Policy
- D. Access

ANSWER: D

Explanation:

Federation allows access across organizations by establishing trust relationships between separate identity domains. In a federated identity model, one organization (the identity provider) authenticates the user and issues an assertion or token (for example, SAML assertions or OpenID Connect/JWT tokens) that another organization (the service provider/relying party) accepts to grant access to resources. This enables single sign-on (SSO) and cross-domain access without requiring each organization to create and manage duplicate user accounts locally. In CCSP terms, federation is a core identity and access management (IAM) capability used to extend authentication and authorization beyond a single administrative boundary while maintaining centralized control over identity proofing, authentication strength, and attribute release. It is commonly implemented with standards-based protocols such as SAML 2.0 and OpenID Connect, which are specifically designed to convey identity and authorization information between organizations in a secure, interoperable way. This is why “access” is the best completion of the statement: federation's primary purpose is to enable users to access services in another organization using their home credentials and established trust.

References: [OASIS SAML 2.0 Core Specification](#), [OpenID Connect Core 1.0](#)

QUESTION NO: 32

All of the following are activities that should be performed when capturing and maintaining an accurate, secure system baseline except .

Response:

- A. Remove all nonessential programs from the baseline image

- B. Exclude the target system you intend to baseline from any scheduled updates/patching used in production systems
- C. Include the baseline image in the asset inventory/configuration management database
- D. Configure the host OS according to the baseline requirements

ANSWER: B

Explanation:

Exclude the target system you intend to baseline from any scheduled updates/patching used in production systems is the exception because maintaining an accurate and secure baseline depends on keeping the baseline build aligned with the organization's approved patch levels and configuration standards over time. A baseline is not a "frozen forever" image; it is a controlled reference configuration that must be updated through change management as patches, security hardening guidance, and approved software versions evolve. If the system used to create or validate the baseline is excluded from scheduled patching, it will quickly drift behind current security requirements, increasing the risk that newly deployed systems inherit known vulnerabilities and fail compliance checks. In practice, baseline images are updated regularly (or rebuilt) to incorporate approved patches, and any changes are documented and version-controlled as part of configuration management. This aligns with standard secure configuration and continuous maintenance practices described in widely used security baselines and hardening guidance. See [NIST SP 800-128 \(Configuration Management\)](#) and [Microsoft Security Baselines](#).

QUESTION NO: 33

What is the data encapsulation used with the SOAP protocol referred to?

- A. Packet
- B. Envelope
- C. Payload
- D. Object

ANSWER: B

Explanation:

In SOAP messaging, the fundamental unit of encapsulation is the SOAP envelope. SOAP is an XML-based messaging protocol used primarily with web services, and every SOAP message is structured as an XML document whose outermost element is the *Envelope*. This envelope defines the message boundaries and provides a standardized container that can include an optional header (for metadata such as security tokens, routing, and transaction context) and a required body (which carries the actual application data or fault information). Because SOAP is designed to be transport-agnostic, the envelope is what makes the message self-describing and consistently parseable regardless of whether it is carried over HTTP, SMTP, or another transport. In CCSP terms, recognizing the SOAP envelope matters for understanding how message-level security, integrity, and policy enforcement can be applied at the header/body level rather than relying only on transport security.

References: [W3C SOAP Version 1.2 Part 1: Messaging Framework](#), [Microsoft Learn: SOAP messages \(WCF\)](#).

QUESTION NO: 34

To address shared monitoring and testing responsibilities in a cloud configuration, the provider might offer all these to the cloud customer except:

- A. Access to audit logs and performance data
- B. DLP solution results
- C. Security control administration

ANSWER: C

Explanation:

Security control administration is the best answer because, in cloud shared responsibility models, the customer and provider may share visibility (logs, telemetry, reports) but they do not typically “share” administration of the provider’s security controls. The provider is responsible for operating and administering the security controls that protect the cloud infrastructure and the managed services layer (for example, hypervisor, physical network, and core platform controls). Customers may configure security settings within their tenant or subscription boundary, but that is customer-side administration of customer controls—not shared administration of the provider’s controls. In practice, providers commonly expose audit logs, performance/health data, and security event feeds (including SIEM-relevant logs) to enable customer monitoring, incident response, and compliance evidence collection. This aligns with CCSP expectations that customers must be able to obtain sufficient logging and reporting to meet governance and assurance needs, while the provider retains operational control over the underlying platform security mechanisms. For more on the shared responsibility concept and how providers retain control of platform security while customers consume logs and configure tenant settings, see <https://learn.microsoft.com/en-us/azure/security/fundamentals/shared-responsibility> and <https://aws.amazon.com/compliance/shared-responsibility-model/>.

QUESTION NO: 35

What does a cloud customer purchase or obtain from a cloud provider?

- A. Services
- B. Hosting
- C. Servers
- D. Customers

ANSWER: A

Explanation:

A cloud customer purchases or obtains services from a cloud provider. In CCSP terms, cloud computing is fundamentally a service-delivery model: the provider offers capabilities (compute, storage, networking, platforms, and applications) that are delivered on-demand, measured, and accessed over the network. Whether the customer is consuming infrastructure, a development/runtime platform, or a complete application, what is being acquired is still a service with defined characteristics (such as elasticity, metering, and self-service provisioning) and governed by contracts and SLAs. This aligns with the NIST cloud definition, which frames cloud computing as a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released—i.e., services delivered by the provider. CCSP also emphasizes that the customer-provider relationship is centered on service consumption and shared responsibility for securing those services. Specific items like hosting environments or servers/virtual machines are examples of service components, but the overarching, correct concept is that the customer is obtaining services from the provider.

References: [NIST SP 800-145 \(The NIST Definition of Cloud Computing\)](#), [ISO/IEC 17788:2014 Cloud computing — Overview and vocabulary](#)

QUESTION NO: 36

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 ASHRAE's data center environmental guidance targets a mid-range humidity band that balances two major risks in IT environments: electrostatic discharge (ESD) at low humidity and corrosion/condensation risk at high humidity. Keeping relative humidity in this moderate range helps reduce static buildup that can damage sensitive electronics while also avoiding excessive moisture that can accelerate corrosion on contacts and circuit boards or create conditions where condensation could occur during temperature transients. In practice, many data centers operationalize this by controlling both relative humidity and dew point (or humidity ratio) to keep conditions stable across seasonal changes and varying heat loads. This aligns with the intent of ASHRAE's thermal guidelines for data processing environments, which emphasize maintaining environmental conditions that support IT reliability and availability while enabling efficient cooling strategies. For additional context on ASHRAE's data center thermal guidance and commonly cited humidity targets, see ASHRAE's overview materials and summaries such as [ASHRAE TC 9.9](#) and a vendor-neutral discussion of ASHRAE humidity/dew point recommendations in data centers like [Cisco Data Center Environmental Guidelines](#).

QUESTION NO: 37

Which cloud storage type requires special consideration on the part of the cloud customer to ensure they do not program themselves into a vendor lock-in situation?

- A. Unstructured
- B. Object
- C. Volume
- D. Structured

ANSWER: D

Explanation:

Structured storage is the cloud storage type that most commonly drives vendor lock-in risk because it is typically consumed as a managed database or other provider-specific data service (often aligned with PaaS). In these services, the provider controls the underlying platform, APIs, scaling mechanisms, backup/restore workflows, and sometimes proprietary query features or extensions. When an application is written to depend on those provider-specific interfaces and behaviors, migrating to another cloud can require significant refactoring, data transformation, and revalidation—far beyond simply copying data. CCSP guidance emphasizes understanding cloud service models and portability risks: higher-level managed services generally increase dependency on the provider's implementation choices. Therefore, customers should design for portability (e.g., abstraction layers, standards-based SQL where possible, minimizing proprietary features, and planning export/replication strategies) when using structured storage services. This is a key part of avoiding “programming in” a dependency that becomes costly to unwind later.

References: [NIST SP 800-146 \(Cloud Computing Synopsis and Recommendations\)](#), [NIST SP 800-145 \(Cloud Computing Definition\)](#)

QUESTION NO: 38

Which of the following statements accurately describes VLANs?

- A. They are not restricted to the same data center or the same racks.

- B. They are not restricted to the name rack but restricted to the same data center.
- C. They are restricted to the same racks and data centers.
- D. They are not restricted to the same rack but restricted to same switches.

ANSWER: A

Explanation:

They are not restricted to the same data center or the same racks is the accurate description. A VLAN (Virtual Local Area Network) is a logical Layer 2 segmentation mechanism that groups switch ports (and/or tagged traffic) into separate broadcast domains independent of physical location within the same Layer 2 switching environment. In practice, this means devices can be placed into the same VLAN even if they are connected to different access switches and located in different racks, as long as the VLAN is carried across the switching infrastructure (typically via 802.1Q trunking) and the network is designed to extend that VLAN. This logical separation is commonly used to isolate workloads, reduce broadcast scope, and enforce security boundaries (often combined with routing/ACLs between VLANs). While VLANs are typically confined to a single administrative Layer 2 domain (and extending them across data centers is possible but requires specific technologies and careful design), the key concept tested here is that VLAN membership is not tied to physical rack placement or a single switch. See IEEE 802.1Q VLAN tagging concepts and practical VLAN/trunking behavior in vendor documentation such as [Cisco VLAN Basics](#) and [Microsoft networking segmentation overview](#).

QUESTION NO: 39

At which phase of the SDLC process should security begin participating?

- A. Requirements gathering
- B. Requirements analysis
- C. Design
- D. Testing

ANSWER: A

Explanation:

Security should begin participating at requirements gathering because this is where the system's intended business outcomes, constraints, and stakeholder expectations are first captured, and it is the earliest point to define security and privacy requirements as first-class requirements. In secure SDLC practices, "shift left" means building security in from the start: identifying high-level security objectives, compliance needs, data classification, trust boundaries, and misuse/abuse cases before architecture and design decisions lock in risk. Getting security involved during requirements gathering enables activities like initial risk assessment, defining security acceptance criteria, and ensuring nonfunctional requirements (confidentiality, integrity, availability, resiliency, auditability) are explicitly documented and testable later. This reduces rework and cost, because fixing security gaps after design or during testing is typically far more expensive and may require architectural changes. This approach aligns with widely accepted secure development guidance that emphasizes integrating security throughout the SDLC starting at the earliest phases. See NIST's Secure Software Development Framework for integrating security practices across development (<https://csrc.nist.gov/publications/detail/sp/800-218/final>) and OWASP guidance on building security into requirements and design (<https://owasp.org/www-project-samm/>).

QUESTION NO: 40

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

Which of these activities should you perform before deploying the tool? Response:

- A. Survey your company's departments about the data under their control

- B. Reconstruct your firewalls
- C. Harden all your routers
- D. Adjust the hypervisors

ANSWER: A

Explanation:

Survey your company's departments about the data under their control is the key prerequisite activity before deploying a DLP solution because DLP effectiveness depends on knowing what sensitive data exists, where it lives, who owns it, and how it flows. In CCSP terms, you need to establish data governance inputs (data ownership, classification, handling requirements, and lifecycle) so the DLP can be configured with accurate policies (what to detect, what to block/quarantine, what to encrypt, what to alert on) and tuned to reduce false positives/negatives. Without this discovery and stakeholder alignment, DLP rules are typically generic, overly restrictive, or miss critical regulated data (PII/PHI/IP), leading to business disruption or inadequate protection. A departmental survey (often paired with data discovery/scanning and interviews) helps build an inventory of data types and repositories (SaaS, IaaS storage, endpoints), define acceptable use and exception processes, and map requirements to controls before rollout in production. This aligns with common DLP implementation guidance: identify and classify sensitive data first, then implement monitoring/enforcement policies based on that classification.

References: [NIST SP 800-53 Rev. 5 \(data classification/handling and governance-related controls\)](#), [Microsoft Purview DLP overview \(emphasizes identifying sensitive info and defining policies\)](#)

QUESTION NO: 41

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:

Key contractual components to review and fully understand when engaging a cloud service provider include the use of subcontractors and the scope of processing. Subcontractor use is a critical contract element because it directly affects the cloud customer's risk posture, visibility, and accountability chain. Contracts should clearly define whether subcontractors are permitted, how they are vetted, what security and privacy obligations flow down to them, and how the customer is notified of changes. This is especially important for compliance and auditability, since subprocessors may handle customer data or provide material parts of the service.

The scope of processing is equally fundamental because it defines what data is processed, for what purposes, where processing occurs, and under what constraints. Clear scope language supports governance, regulatory compliance (including privacy obligations), and alignment to the shared responsibility model by setting boundaries for permitted processing activities and data handling requirements. These two items are commonly emphasized in cloud contracts, data processing addenda, and supplier agreements as core terms that determine legal, security, and compliance outcomes. See [Cloud Security Alliance CAIQ](#) and [GDPR Article 28 \(Processor\)](#) for widely used guidance on subprocessors and processing scope in cloud/service provider contracting.

QUESTION NO: 42

Which of the following is NOT part of a retention policy?

- A. Format
- B. Costs
- C. Accessibility
- D. Duration

ANSWER: B

Explanation:

“Costs” is not typically considered a core element of a data retention policy. In CCSP terms, a retention policy is primarily a governance control that defines what information must be kept, for how long, and under what conditions it must remain available and protected. That includes retention duration (often driven by legal/regulatory requirements), acceptable storage formats/media, accessibility requirements (who can retrieve it and how quickly), and related handling constraints such as preservation, integrity, and disposition (secure deletion) at end of life. While cost considerations may influence implementation choices (e.g., tiered storage, archiving solutions, or lifecycle rules), they are generally treated as part of budgeting, total cost of ownership analysis, or operational planning rather than being a defining component of the policy itself. In other words, the policy states the “what/when/how long,” and the organization then selects cost-effective mechanisms to meet those requirements without weakening compliance, security, or eDiscovery readiness. This aligns with common records management and information governance practice where retention schedules and policies are driven by business, legal, and risk requirements, not by the accounting line items of the solution used to enforce them.

References: [NIST SP 800-53 Rev. 5 \(controls for retention and disposal\)](#), [Microsoft Purview retention documentation](#)

QUESTION NO: 43

Which of the following service capabilities gives the cloud customer the least amount of control over configurations and deployments?

- A. Platform
- B. Infrastructure
- C. Software
- D. Desktop

ANSWER: C

Explanation:

The service capability that gives the cloud customer the least control over configurations and deployments is software as a service. In a software as a service model, the provider delivers a complete, ready-to-use application stack and retains responsibility for the underlying infrastructure, virtualization, operating system, middleware/runtime, and the application itself (including most configuration and deployment decisions). The customer’s control is typically limited to user-specific settings and application-level configuration exposed by the vendor (for example, tenant settings, roles, and some feature toggles), rather than control over how the application is deployed, patched, scaled, or architected. This aligns with the core cloud shared responsibility concept: as you move from infrastructure to platform to software services, customer control decreases while provider responsibility increases. CCSP exam content commonly tests this gradient of control across service models, where software as a service sits at the “least customer control” end of the spectrum. For a clear depiction of responsibilities by cloud service model, see [Microsoft’s shared responsibility model](#) and the NIST definition of cloud service models in [NIST SP 800-145](#).

QUESTION NO: 44

APIs are defined as which of the following?

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

ANSWER: B

Explanation:

An API (Application Programming Interface) is best defined as a set of routines, standards, protocols, and tools that enables software components to communicate and allows developers to build applications that interact with other software services. In cloud contexts (highly relevant to CCSP), APIs are the primary mechanism customers and administrators use to provision resources, automate operations, integrate services, and manage cloud environments. This definition captures both the “how” (protocols and routines that govern interaction) and the “what” (tools and standards that support building and integrating software), which is why it is the most complete characterization among the choices. Importantly, APIs are not limited to web-based applications, but modern cloud service APIs are commonly exposed over HTTP(S) using REST/JSON or similar patterns; still, the core concept remains an interface contract that defines requests, responses, authentication/authorization expectations, and allowable operations. From a security perspective, understanding APIs as a formalized interface is critical because API endpoints become part of the attack surface and must be governed with strong identity controls, authorization, input validation, logging, and rate limiting. For additional authoritative background, see [Red Hat: What is an API?](#) and [IBM: What is an API?](#).