WGU Managing Cloud Security (JY02, GZO1)

WGU Managing-Cloud-Security

Version Demo

Total Demo Questions: 22

Total Premium Questions: 222

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	50
Topic 2, Cloud Data Security	31
Topic 3, Cloud Platform & Infrastructure Security	31
Topic 4, Cloud Application Security	19
Topic 5, Cloud Security Operations	34
Topic 6, Legal, Risk and Compliance	57
Total	222

QUESTION NO: 1

Which tool provides a dedicated environment to contain and analyze malware?

A. Encryption

B. Gateway

C. Sandbox

D. Controller

ANSWER: C

Explanation:

Sandbox is correct because a sandbox is an isolated, controlled environment designed to execute and observe suspicious files, programs, or URLs without exposing production systems or the broader network to potential compromise. Security teams use sandboxing to determine whether an artifact is malicious and to document its behavior, such as process creation, attempted network connections, file modifications, persistence activity, and other indicators of compromise. This controlled execution is particularly valuable for investigating unknown files and detecting malware that may evade signature-based tools, including previously unseen or obfuscated threats. Effective malware-analysis sandboxes use isolation boundaries and monitoring to prevent the sample's actions from affecting real endpoints, identities, data, or services while still providing analysts with actionable telemetry. In cloud security operations, sandboxing may be delivered as an automated analysis capability within email, endpoint, or threat-detection workflows, allowing suspicious content to be detonated safely at scale. NIST describes dynamic analysis as executing code in a controlled environment to observe its behavior, a core purpose of a malware sandbox. MITRE also documents sandbox evasion as a malware technique, underscoring that sandboxes are widely used defensive analysis environments. See NIST SP 800-115 and MITRE ATT&CK: Virtualization/Sandbox Evasion.

QUESTION NO: 2

Which role ensures that third parties take adequate technical and organizational security measures to safeguard data?

A. Cloud user

B. Data controller

C. Cloud provider

D. Data subject

ANSWER: B

Explanation:

The **Data controller** is correct because it determines the purposes and essential means of processing personal data and remains accountable for ensuring that processing is performed with appropriate protection, including when an external organization handles data on its behalf. In cloud arrangements, this accountability requires the controller to select suitable service providers, establish a data-processing agreement, define documented processing instructions, and confirm that the provider offers sufficient technical and organizational safeguards. Those safeguards can include access control, encryption, incident-response processes, data-location and deletion provisions, audit rights, and mechanisms for managing subcontractors.

Under the GDPR, a controller may use only processors that provide "sufficient guarantees" that appropriate technical and organizational measures will meet regulatory requirements and protect individuals' rights. The controller must also have a binding contract with the processor that addresses required processing and security obligations. This makes third-party due diligence and contractual security oversight a core controller responsibility rather than something that can be fully delegated. See GDPR Article 28 and the European Data Protection Board guidance on controller–processor concepts.

QUESTION NO: 3

As part of an e-discovery process, an employee needs to identify all documents that contain a specific phrase. Which type of discovery method should the employee use to identify these documents?

A. Location-based

B. Content-based

C. Label-based

D. Metadata-based

ANSWER: B

Explanation:

Content-based discovery is the appropriate method because the requirement is to locate documents according to words that appear within the documents themselves. A search for a specific phrase evaluates indexed document text and, where supported, the extracted text of common file types. This allows an e-discovery reviewer to identify items containing the requested phrase regardless of where those items are stored or who created them. Phrase searching is particularly important in legal, compliance, and investigative workflows because it provides a targeted and repeatable way to collect potentially relevant electronically stored information. Modern e-discovery platforms commonly support keyword queries, exact-phrase queries, and additional search conditions so organizations can scope a collection appropriately and preserve results for review. For example, Microsoft Purview eDiscovery supports using keywords and query syntax to search the content of custodial data sources, including exact phrases enclosed in quotation marks. This directly aligns the search method with the stated need: finding the phrase in document content. See Microsoft Purview eDiscovery keyword search and Microsoft Purview eDiscovery search for content.

QUESTION NO: 4

An organization negotiates a new contract with a cloud provider and wants to ensure that its critical business data is protected if the cloud provider goes out of business. For this reason, the organization wants the cloud provider to store a copy of the organization's data with a neutral third party, which would release the data in case the provider is unable to meet its obligations. Which type of provision should be included in the contract to ensure this?

A. Indemnification

B. Escrow

C. Offboarding

D. Encryption

ANSWER: B

Explanation:

Escrow is the appropriate contractual provision because it establishes an arrangement in which a neutral, independent third party securely holds designated customer assets and releases them when predefined conditions occur. In this case, the escrowed asset is a usable copy of the organization’s critical business data. The agreement should precisely define the data to be deposited, the deposit frequency, security and encryption requirements, verification or restoration testing, the release conditions, and the organization’s right to obtain the data if the provider becomes insolvent, ceases operations, or otherwise cannot meet its contractual obligations.

For cloud services, an effective data-escrow provision supports business continuity by ensuring the customer retains a path to recover essential information independently of the provider’s continued viability. The organization should also ensure the escrow copy is current, in a documented exportable format, and accompanied by the metadata, credentials, and technical documentation required to restore or migrate the data to another environment. This approach aligns with NIST guidance to address cloud-provider exit and portability risks as part of service agreements and contingency planning. See NIST SP 800-146, Cloud Computing Synopsis and Recommendations and NIST Privacy Framework and Cybersecurity Framework resources.

QUESTION NO: 5

Which country lacks a national law assuring individual personal privacy?

A. New Zealand

B. Canada

C. United States

D. Israel

ANSWER: C

Explanation:

United States is correct because it does not have a single, comprehensive national privacy law that broadly assures personal-data privacy for individuals across the private sector. Instead, U.S. privacy protections are distributed across federal sector-specific statutes, such as laws governing health information, financial information, children's data, and consumer reporting, along with state privacy laws. This produces a fragmented framework in which the protections and obligations depend on the type of information, the organization handling it, and the applicable jurisdiction. The U.S. Privacy Act of 1974 does establish privacy requirements for records maintained by federal agencies, but it is not a general national law governing private-sector processing of all individuals' personal data. The Congressional Research Service describes this patchwork approach and notes the absence of a comprehensive federal consumer data privacy statute. For cloud-security governance, this distinction matters because organizations operating in the United States must identify and comply with the particular federal and state requirements relevant to their data, services, and customers rather than relying on one overarching national privacy regime. See the Congressional Research Service overview of data protection law and the U.S. Department of Justice's Privacy Act guidance.

QUESTION NO: 6

What is at risk in the cloud environment when the management console is breached?

A. Breakout of a guest OS that can access other hypervisors

B. Amount of content on the image store

C. Entire infrastructure that the control plane administers

D. Sensitive information that images and snapshots can contain

ANSWER: C

Explanation:

Entire infrastructure that the control plane administers is at risk because a cloud management console is an interface to the cloud provider's control plane. Authorized administrators commonly use it to manage identity and access permissions, provision and alter compute instances, configure virtual networks and security controls, manage storage, deploy services, and view operational information. If an attacker compromises a console account, especially one with broad administrative permissions, that attacker can exercise the permissions assigned to the account across the resources and services it administers.

The resulting impact can extend throughout the environment: an attacker may change configurations, create new identities or credentials for persistence, access workloads and data through management APIs, disrupt availability by stopping or deleting resources, or weaken logging and monitoring. This is why strong authentication, least-privilege authorization, session protection, and audit logging for console access are foundational cloud-security controls. AWS describes the management console as a web application for accessing and managing AWS services, while NIST emphasizes that cloud customers must manage privileged access and administrative interfaces carefully because those interfaces can affect broad portions of the cloud deployment. See AWS Management Console documentation and NIST SP 800-144.

QUESTION NO: 7

Which U.S. standard is used by federal government agencies to manage enterprise risk?

A. International Organization for Standardization (ISO) 37500

B. The Statement on Standards for Attestation Engagements 18 (SSAE 18)

C. The Committee of Sponsoring Organizations (COSO) framework

D. The National Institute of Standards and Technology (NIST) SP 800-37

ANSWER: D

Explanation:

The National Institute of Standards and Technology (NIST) SP 800-37 is correct because it defines the Risk Management Framework (RMF) used by U.S. federal information systems and organizations. The RMF provides a structured, repeatable lifecycle for managing security and privacy risk: preparing the organization and system, categorizing systems and information, selecting and implementing controls, assessing their effectiveness, authorizing operation based on risk, and continuously monitoring the security posture. This framework supports the risk-management requirements established under the Federal Information Security Modernization Act (FISMA) and is intended to integrate information-security and privacy activities into an organization's broader enterprise risk management processes. In cloud-security contexts, agencies can apply RMF to evaluate cloud services, select appropriate safeguards, document risk decisions, and maintain ongoing awareness of changes to systems and threats. NIST SP 800-37 Rev. 2 explicitly describes the RMF as the standard federal approach for managing security and privacy risk at the system level while linking those decisions to organizational risk. See the NIST SP 800-37 Rev. 2 publication and NIST's Risk Management Framework project page for the authoritative guidance.

QUESTION NO: 8

Which phase of the software development life cycle includes creating user stories?

A. Developing

B. Designing

C. Defining

D. Planning

ANSWER: D

Explanation:

Planning is the SDLC phase that includes creating user stories. User stories are brief, user-centered descriptions of needed functionality, commonly expressed in a form such as "As a user, I want [capability] so that [benefit]." They help stakeholders, product owners, and development teams identify and communicate requirements from the customer's perspective before implementation begins.

During planning, the team establishes the product vision, scope, priorities, expected outcomes, and high-level requirements. User stories provide an approachable way to document those requirements, support backlog creation and prioritization, and start conversations about acceptance criteria. This work gives the team a shared understanding of what value the application should deliver and informs later design and development activities. In iterative and Agile-oriented cloud development, user stories are regularly refined as planning continues and new information becomes available.

The Agile Alliance describes user stories as short, simple descriptions of a feature told from the perspective of the person who desires the capability: Agile Alliance: User Stories. Additional guidance on organizing and prioritizing user stories in a product backlog is available from Atlassian Agile: User Stories.

QUESTION NO: 9

Which security control could be implemented as part of a layered physical defense at a cloud hosting site?

A. Access control enforcement

B. Background checks

C. Video surveillance capability

D. Multifactor authentication

ANSWER: C

Explanation:

Video surveillance capability is a physical security control that supports defense in depth at a cloud hosting site. Layered physical defense uses multiple, overlapping safeguards around and within a facility to deter unauthorized entry, detect suspicious activity, and support an appropriate response. Cameras can monitor the site perimeter, parking and loading areas, entrances, visitor reception areas, equipment rooms, and other sensitive locations. This visibility helps security personnel identify potential incidents as they occur and provides recorded evidence for incident investigation, audit activities, and forensic review.

In a cloud data center, surveillance is particularly valuable because it complements other facility protections across distinct zones. Monitoring may be continuous, event-triggered, or integrated with a security operations center so that personnel can investigate alerts promptly. The retained footage also establishes an evidentiary record of physical activity near critical infrastructure. This makes video surveillance a practical detection and response component of a layered physical-security program. The NIST physical and environmental protection control family includes monitoring physical access as an important security capability; see NIST SP 800-53 Rev. 5. Cloud providers likewise describe camera monitoring as part of data-center physical security, such as in AWS Data Center Controls.

QUESTION NO: 10

A business process requires that a critical cloud application be restored within one hour of a regional outage. The application can tolerate the loss of no more than five minutes of recently entered transaction data. Which recovery design best meets these requirements?

A. Performing daily backups and rebuilding the environment after an outage

B. Maintaining offline archive copies for monthly restoration testing

C. Using a cold recovery site with manually provisioned infrastructure

D. Using continuous replication to a standby environment with automated failover

ANSWER: D

Explanation:

A continuously replicated standby environment with automated failover is the best choice because it is designed to provide a low recovery point objective (RPO) and a short recovery time objective (RTO). Replication at intervals of five minutes or less limits potential data loss, while a ready standby environment reduces the time required to resume service after an outage.

Daily backups do not satisfy the five-minute RPO. Rebuilding infrastructure after an outage may be less expensive, but it is unlikely to meet the one-hour RTO. A manually restored backup also introduces delay and risks exceeding both recovery objectives.

QUESTION NO: 11

Which risk is unable to be highlighted from the outset in a cloud services contract?

A. Introduction of new technology

B. Sunsetting of aging technology

C. Result of an unforeseen event

D. Changes resulting from contract renewals

ANSWER: C

Explanation:

Result of an unforeseen event is correct because a cloud-services contract can identify known, reasonably foreseeable risks at the time it is negotiated, but it cannot fully specify the effects of events that are unknown or unpredictable at that point. These may include major natural disasters, abrupt geopolitical disruptions, unexpected failures in interconnected infrastructure, or novel security threats. Contracts commonly address this uncertainty through force-majeure provisions, incident-management processes, business-continuity commitments, notification requirements, and allocation of responsibilities after a disruptive event occurs. Those provisions provide a framework for responding to the event; they do not make its particular timing, scope, impact, or resulting obligations fully knowable from the outset. Cloud risk management is therefore an ongoing activity rather than a one-time contract exercise. The National Institute of Standards and Technology describes risk management as continuous and responsive to changes in systems, environments, threats, and organizational conditions. Likewise, cloud-security guidance emphasizes planning for resilience and continuity despite uncertain disruptions. A well-written agreement should set expectations for recovery and communication, while recognizing that the precise consequences of an unforeseen event cannot be highlighted in advance. See NIST SP 800-37 Rev. 2 and the Cloud Security Alliance Cloud Controls Matrix.

QUESTION NO: 12

Which document, commonly existing in an IT enterprise, can be used to speed up the process of identifying a potential cloud service provider (CSP)?

A. Entity relationship and data flow diagrams

B. Physical plant blueprint

C. Business continuity and disaster recovery plan

D. Egress safety design

ANSWER: A

Explanation:

Entity relationship and data flow diagrams are the most useful existing enterprise documents for accelerating identification of a potential cloud service provider. These diagrams provide a practical view of how data moves among applications, databases, users, integrations, and external services. That visibility helps an organization identify the workloads in scope, the types and classifications of data involved, dependencies that must be supported, and trust boundaries that require protection. With those requirements identified early, the organization can more efficiently determine whether it needs capabilities such as particular geographic hosting locations, encryption and key-management features, identity federation, network connectivity, logging, or compliance attestations.

The diagrams also support a structured comparison of providers because the enterprise can map its actual technical and data-handling requirements to a prospective provider’s service offerings and shared-responsibility model. Cloud Security Alliance guidance emphasizes understanding data flows and system relationships as part of defining cloud security requirements and assessing cloud services. NIST similarly identifies data flow understanding as a key activity when applying security and privacy controls to information systems. See the Cloud Security Alliance Security Guidance and NIST SP 800-53 Rev. 5.

QUESTION NO: 13

An organization that primarily uses a remote work model is reviewing the documentation of various insurance providers to become eligible for cybersecurity insurance. Competitive insurance providers require the organization to implement security controls to ensure only authorized personnel can access the network, data, emails, and other administrative information. Which commonly required control should the organization implement before applying for cybersecurity insurance from these competitive insurance providers?

A. Network segmentation

B. Application whitelisting

C. Multifactor authentication (MFA)

D. Trusted platform module (TPM)

ANSWER: C

Explanation:

Multifactor authentication (MFA) is the commonly required control because cyber insurers routinely treat it as a baseline safeguard against identity-based compromise. MFA requires a user to verify their identity with more than one factor, such as a password plus an authenticator application, hardware security key, or biometric verification. This added verification makes a stolen, guessed, reused, or phished password substantially less useful to an attacker.

For a remote-work organization, MFA is especially important because staff access email, cloud applications, VPNs, collaboration platforms, and administrative portals through internet-facing authentication services. Applying MFA to these access points helps ensure that access to the organization's network, data, email, and privileged functions is limited to authorized personnel. Insurers focus on this control because business email compromise, credential theft, ransomware, and unauthorized remote access are frequent and costly sources of cyber claims. The U.S. Cybersecurity and Infrastructure Security Agency identifies MFA as a critical protective measure and recommends it for remote access and email. The Cyber Insurance Academy also identifies MFA as a common underwriting requirement. See CISA: Turn On MFA and Cyber Insurance Academy: Cyber Insurance Requirements.

QUESTION NO: 14

Which data retention policy addresses how long data must be retained to meet regulatory requirements?

A. Formats

B. Classification

C. Retrieval

D. Retention periods

ANSWER: D

Explanation:

Retention periods is correct because this policy element establishes the length of time an organization must preserve a given category of data before it can be dispositioned according to approved procedures. Regulatory, legal, contractual, operational, and evidentiary obligations frequently impose minimum retention durations. For example, requirements may apply to financial records, healthcare information, security logs, or records subject to litigation holds. A well-defined retention period connects each data type to its required preservation timeframe, accountable owner, storage location, and final disposition action.

In cloud environments, retention periods must be implemented through both governance and technical controls. These can include retention labels, immutable or write-once storage settings, lifecycle policies, backup-retention configurations, and audited deletion workflows. The organization should also periodically review the schedule as regulations and business obligations change. Retaining information for the required period supports compliance and availability for audits, investigations, and recovery needs, while disposing of data at the end of its authorized lifecycle reduces unnecessary storage, exposure, and privacy risk. NIST describes records-management controls, including retention and disposal

requirements, in NIST SP 800-53 Rev. 5. Microsoft also documents how retention policies preserve and delete information over defined time periods in Microsoft Purview retention policies.

QUESTION NO: 15

Which factor is a primary consideration when analyzing the legal and privacy implications of using cloud technologies?

A. Jurisdiction of the cloud provider and users

B. Level of encryption that the cloud solution provides

C. Configuration details specified in the contract

D. Penalties specified in the service level agreement

ANSWER: A

Explanation:

Jurisdiction of the cloud provider and users is a primary consideration because applicable laws can depend on where cloud data is stored, processed, transmitted, and accessed, as well as where the customer and provider are established. Those jurisdictions may determine privacy obligations, requirements for lawful cross-border transfers, data-retention duties, breach-notification rules, and the circumstances in which governmental authorities can seek access to information. Cloud services commonly use distributed infrastructure, so an organization must understand both the provider's data-location practices and the legal locations relevant to its own users and regulated data. This assessment helps the organization select appropriate regions, establish compliant contractual and technical controls, and evaluate whether use of the service supports obligations under laws such as the GDPR or sector-specific requirements. The European Commission explains that transfers of personal data outside the European Economic Area require a valid transfer mechanism and appropriate safeguards, making the location and legal destination of processing directly relevant. NIST also identifies legal and jurisdictional issues as important cloud-computing considerations. See European Commission: International data transfers and NIST SP 800-146: Cloud Computing Synopsis and Recommendations.

QUESTION NO: 16

A network administrator is concerned about the loss of physical control when moving data to the cloud. Which countermeasure should be implemented to avoid this threat?

A. Multi-layer control

B. Tertiary control

C. Ancillary control

D. Compensating control

ANSWER: D

Explanation:

Compensating control is correct because moving systems and data to a cloud provider means the customer no longer operates the physical facility, controls physical entry to the data center, or directly manages infrastructure protections such as building security and hardware handling. A compensating control provides an alternative safeguard that achieves the security intent when the original or direct control cannot practically be implemented by the organization. In this situation, customer-managed technical and administrative safeguards compensate for the loss of direct physical control.

Appropriate compensating safeguards include encrypting data both in transit and at rest, retaining strong control of encryption keys, enforcing least-privilege identity and access management, enabling detailed audit logging, monitoring for anomalous access, and evaluating the provider's physical-security assurances and compliance reports. Together, these controls reduce the chance that loss of facility access translates into unauthorized disclosure or alteration of cloud-hosted information. This aligns with the cloud shared-responsibility model: the provider protects the underlying facilities and

infrastructure, while the customer remains responsible for configuring and operating protections for its data, identities, and workloads. NIST discusses the use of compensating controls when specified controls cannot be implemented, and AWS explains the allocation of responsibilities in cloud environments. See NIST SP 800-53 Rev. 5 and AWS Shared Responsibility Model.

QUESTION NO: 17

Which description accurately characterizes the movement of applications to the cloud?

A. In a desktop as a service (DaaS) environment, the customer is responsible for securing the underlying infrastructure.

B. In an infrastructure as a service (IaaS) environment, the CSP is responsible for securing the platform.

C. In a platform as a service (PaaS) environment, the customer is responsible for securing the underlying infrastructure.

D. In a software as a service (SaaS) environment, the CSP is responsible for securing the platform.

ANSWER: D

Explanation:

In a software as a service (SaaS) environment, the CSP is responsible for securing the platform. SaaS delivers a complete, provider-operated application to the customer, so the cloud service provider manages and secures the service's supporting platform, including the application infrastructure, operating environment, service availability, patching of provider-managed components, and the physical cloud facilities. This model reduces the customer's operational responsibility for the technology stack compared with infrastructure as a service or platform as a service.

The customer still has important security duties in SaaS, particularly configuring the service securely, managing identities and access permissions, protecting the data entered into the service, and ensuring appropriate use of the application. However, those customer responsibilities do not transfer responsibility for securing the provider-operated platform. This division is an application of the cloud shared-responsibility model: the provider secures the cloud service and its underlying managed components, while the customer secures its use of that service and its information. AWS describes this model in its Shared Responsibility Model, and NIST describes SaaS as a provider-hosted application model in SP 800-145.

QUESTION NO: 18

Which device is used to create and manage encryption keys used for data transmission in a cloud-based environment?

A. Hardware security module (HSM)

B. Memory controller

C. RAID controller

D. Trusted platform module (TPM)

ANSWER: A

Explanation:

Hardware security module (HSM) is correct because an HSM is specialized, tamper-resistant hardware that generates, protects, stores, and uses cryptographic keys without exposing their sensitive key material to general-purpose systems. For cloud-based data transmission, HSMs can securely perform cryptographic operations supporting TLS/SSL certificates, encryption, decryption, signing, and key wrapping. This enables organizations to protect keys used to establish secure communications while maintaining strong control over how those keys are accessed and used. Cloud providers commonly offer HSM-backed key-management services as well as dedicated HSM offerings, allowing customers to use validated cryptographic hardware while integrating it with cloud applications and services. HSMs also support security and compliance requirements through controls such as role-based administration, audit logging, secure key backup mechanisms where applicable, and hardware protections against physical tampering. The National Institute of Standards and Technology describes cryptographic modules and their security requirements in FIPS 140-3, a standard frequently used when evaluating

HSM security assurances. AWS also documents CloudHSM as a cloud service that provides dedicated hardware security modules for generating and using encryption keys. See NIST FIPS 140-3 and AWS CloudHSM documentation.

QUESTION NO: 19

An organization is undergoing an ISO 27001 audit that includes a software as a service (SaaS) solution within scope, and the auditor has requested evidence of controls. What evidence should the organization provide the auditor?

A. Network firewall rules

B. Provider compliance attestation

C. Operating system patch logs

D. Physical diagram of the data center

ANSWER: B

Explanation:

Provider compliance attestation is the appropriate evidence because, in a SaaS model, the service provider operates and controls much of the underlying environment used to deliver the application. A current independent assurance artifact—such as the provider's ISO/IEC 27001 certificate and statement of applicability, or a relevant SOC 2 report—gives the organization and its auditor evidence that the provider's applicable security controls have been independently assessed. The organization should obtain an attestation that is current, identifies the covered service and locations, and is relevant to the organization's intended use of the SaaS offering. It should also review the document for any scope limitations, exceptions, complementary customer controls, and corrective-action information. This supports the organization's supplier-relationship and outsourced-service assurance activities while recognizing the cloud shared-responsibility model. ISO describes ISO/IEC 27001 as the standard for establishing and maintaining an information security management system: ISO/IEC 27001. For assurance reports, the AICPA explains the purpose and structure of SOC examinations: AICPA SOC suite of services.

QUESTION NO: 20

Which cloud storage architecture allows the digital rights management (DRM) solutions to associate metadata with the materials in storage?

A. Object-based

B. Volume

C. Relational database

D. File

ANSWER: A

Explanation:

Object-based storage is correct because it stores each item as a self-contained object that combines the content itself, a unique object identifier, and descriptive, customizable metadata. This design enables a DRM solution to associate information such as content owner, classification, licensing terms, authorized users, retention requirements, expiration dates, and usage restrictions directly with the stored material. Keeping metadata with the object supports consistent policy decisions as the object is accessed through cloud storage APIs, replicated, or managed across the storage service. Cloud providers implement this model through object metadata and, often, object tags that applications can use to locate, organize, govern, and apply controls to content. For example, Amazon S3 documents that objects consist of object data and metadata, and that metadata is a set of name-value pairs associated with an object. Microsoft Azure similarly supports user-defined metadata for blobs, allowing applications to store additional attributes with content. These capabilities make object-based storage a natural architectural fit when DRM workflows need rights-related attributes to travel with and describe individual digital assets. See Amazon S3 object metadata documentation and Microsoft Azure Blob metadata documentation.

QUESTION NO: 21

A customer requests that a cloud provider physically destroys any drives storing their personal data. What must the provider do with the drives?

A. It should destroy them only if the contract includes hardware disposal insurance.

B. It should use cryptographic erasure to securely remove any personal data from the drives.

C. It should destroy them only if dedicated hardware disposal is specified in the contract.

D. It should use degaussing tools to securely remove any personal data from the drives.

ANSWER: C

Explanation:

It should destroy them only if dedicated hardware disposal is specified in the contract. A cloud provider's obligations for media handling, sanitization, and physical destruction are established by the service agreement and any applicable data-processing terms. In a multitenant cloud environment, storage hardware is ordinarily part of a shared infrastructure lifecycle, so a customer request alone does not necessarily require the provider to remove and destroy a specific physical drive. If the parties contractually agree to dedicated hardware disposal or physical destruction for media that held the customer's data, the provider must perform that obligation in accordance with the agreed scope, process, and evidence requirements. Such terms may address how storage is isolated, the approved destruction method, chain-of-custody records, and a certificate of destruction. Clear contractual allocation is particularly important where the customer has regulatory, data-residency, or internal data-retention requirements that demand physical media destruction rather than another approved sanitization process. This approach reflects the cloud shared-responsibility model: the provider manages the underlying infrastructure according to documented service commitments, while the customer must select and negotiate services that meet its compliance needs. See NIST SP 800-88 Rev. 1, Guidelines for Media Sanitization and CISA Cloud Security Technical Reference Architecture.

QUESTION NO: 22

A development team deploys a public API that accepts requests from a mobile application. During testing, the API returns detailed database error messages and accepts requests with excessively large payloads. Which control should the team implement first to reduce the application's exposure?

A. Increasing compute capacity for the API service

B. Implementing input validation, request-size limits, and controlled error handling

C. Encrypting the application backup repository

D. Moving the API to a different cloud region

ANSWER: B

Explanation:

The team should implement API input validation, request-size limits, and controlled error handling. Validation and size limits reduce the risk of malformed or oversized requests reaching backend services, while generic client-facing errors prevent internal database details from being disclosed to an attacker.

Increasing compute capacity may improve performance but does not correct insecure request handling. Encrypting backups protects stored recovery data rather than API input processing. Moving the API to a different cloud region changes location but does not remediate the identified application weaknesses.