

CompTIA Cloud Essentials+

CompTIA CLO-002

Version Demo

Total Demo Questions: 24

Total Premium Questions: 246

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Concepts	59
Topic 2, Business Principles of Cloud Environments	77
Topic 3, Management and Technical Operations	46
Topic 4, Governance, Risk, Compliance, and Security for the Cloud	64
Total	246

QUESTION NO: 1

A development department has experienced unexpected cloud charges because temporary test resources were left running after projects were completed.

Which of the following would BEST help prevent similar unexpected costs? (Choose two.)

- A. Budget alerts
- B. Resource schedules
- C. Long-term archival storage
- D. Reserved instances for all workloads
- E. Increased storage replication

ANSWER: A B

Explanation:

Budget alerts are correct because they notify stakeholders when cloud spending or forecasted spending reaches defined thresholds. This enables teams to investigate unexpected usage before costs increase further.

Resource schedules are also correct because they can automatically stop or deallocate nonproduction resources outside approved working hours or at the end of a testing period. This reduces charges from idle development and test environments. These controls should be supported by ownership assignment, cost monitoring, and periodic review of resource utilization. See the [AWS Budgets documentation](#) and the [Microsoft Azure budget tutorial](#).

QUESTION NO: 2

A company is evaluating SaaS providers for a system that will process customer personal information.

Which of the following should be reviewed as part of the provider due-diligence process? (Choose two.)

- A. Data processing agreement
- B. Independent compliance attestations
- C. Employee performance reviews
- D. Office floor plans
- E. Corporate social-media policy

ANSWER: A B

Explanation:

A data processing agreement is correct because it documents how the provider will process, protect, retain, and return or delete the customers personal information. It can also define responsibilities for incident notification, subcontractors, and regulatory obligations.

Independent compliance attestations are also correct because reports and certifications, such as SOC reports or ISO/IEC 27001 certification, can provide evidence that the provider has implemented and assessed relevant security controls. The customer must still determine whether the providers controls and contractual terms satisfy its own legal, business, and risk requirements. See the [CISA Secure Cloud Business Applications guidance](#) and [ISO/IEC 27001](#).

QUESTION NO: 3

Which of the following are considered secure access types of hosts in the cloud? (Choose two.)

- A. HTTPS
- B. HTTP
- C. SSH
- D. Telnet
- E. RDP
- F. FTP

ANSWER: A C

Explanation:

HTTPS and SSH are secure access methods because each protects communications between a user or administrator and a cloud-hosted resource through encryption and authenticated session establishment. HTTPS applies Transport Layer Security to HTTP traffic, protecting the confidentiality and integrity of browser-based access to web applications, management portals, and application programming interfaces. It also allows clients to validate the identity of the service through digital certificates. SSH is designed for secure remote administration of hosts and supports encrypted command-line sessions, strong authentication methods such as public-key authentication, and secure file-transfer capabilities when appropriately configured.

Using HTTPS for web-facing management and application access, along with SSH for operating-system administration, supports the cloud security principle of protecting data in transit. Access should also be restricted through firewall or security-group rules, limited to authorized identities and networks, and monitored through logging. CompTIA's Cloud Essentials+ objectives include security concepts relating to secure access and protecting cloud resources. Further technical guidance is available from [Cloudflare's HTTPS overview](#) and [RFC 4251, the SSH Protocol Architecture](#).

QUESTION NO: 4

The Chief Financial Officer for a company that operates a popular SaaS application has noticed compute costs from the CSP are extremely high but storage costs are relatively low. Which of the following does the company MOST likely operate?

- A. An email application
- B. A CDN service
- C. A gaming application
- D. Audio streaming service

ANSWER: C

Explanation:

A gaming application is the most likely workload because a popular online game can impose substantial, continuous compute demand while requiring comparatively modest persistent cloud storage. Game services commonly run real-time multiplayer session logic, matchmaking, authoritative game-state processing, physics calculations, anti-cheat validation, leaderboards, and network coordination. These functions must often respond in milliseconds and scale horizontally as concurrent player counts rise, making virtual machines, containers, managed game servers, and related compute resources a major component of the CSP bill.

Although games do use storage for player profiles, saved progress, game assets, telemetry, and logs, the core interactive workload is frequently compute- and network-intensive. Many assets are distributed to players' devices ahead of time, cached at the edge, or retained in memory during active sessions rather than producing storage consumption proportional to every moment of gameplay. High popularity also amplifies the compute requirement because more simultaneous users require more active server capacity. Cloud gaming and multiplayer game architectures specifically emphasize scalable game-server fleets and low-latency real-time processing. See [AWS for Games](#) and [Microsoft Game Development documentation](#).

QUESTION NO: 5

Which of the following results from implementing a proprietary SaaS solution when an organization does not ensure the solution adopts open standards? (Choose two.)

- A. Vendor lock-in
- B. Inability to enforce the SLA
- C. Lack of technical support
- D. Higher ongoing operational expenditure
- E. Integration issues
- F. Higher initial capital expenditure

ANSWER: A E

Explanation:

Vendor lock-in and **Integration issues** are the expected consequences when a proprietary SaaS offering does not use open standards. Open standards promote interoperability and portability by providing commonly understood interfaces, data formats, and protocols. Without them, an organization may become dependent on a provider's proprietary APIs, application behavior, data structures, and migration tools. This dependency can make it difficult, costly, and time-consuming to move data and workloads to another SaaS provider or to replace the service with an internally hosted or alternative cloud solution. That dependency is vendor lock-in.

Integration issues also result because other enterprise systems, cloud services, and business partners may not support the proprietary interfaces or formats used by the SaaS product. The organization may need custom connectors, middleware, data transformations, or ongoing engineering effort to exchange data reliably. Standards-based interfaces reduce these barriers and help systems communicate across providers and platforms. NIST identifies interoperability and portability as important cloud-computing considerations and describes standards as a means of reducing barriers to integration and migration. See [NIST SP 500-291: Cloud Computing Standards Roadmap](#) and [NIST Cloud Computing Synopsis and Recommendations](#).

QUESTION NO: 6

Which of the following should be considered to ensure the availability of data that is accessed across multiple sites? (Choose two.)

- A. Auto-scaling
- B. Geo-redundancy
- C. Backup
- D. Provisioning
- E. Locality
- F. Zones

ANSWER: B C

Explanation:

Geo-redundancy and backup should be considered for data that must remain available to users or systems operating across multiple sites. Geo-redundancy maintains replicated copies of data in a geographically separate location. This design improves resilience because a major failure, outage, or disaster affecting one site or region does not necessarily make all copies of the data unavailable. Depending on the service and configuration, replication can also support failover to a secondary location, helping maintain access during a site-level disruption.

Backup provides a separate recoverable copy of data that can be restored if the primary data or its replicas are lost, corrupted, deleted, or affected by a security incident. Availability planning is not limited to keeping a service online; it must also ensure that usable data can be recovered after an event. A sound strategy defines backup frequency, retention, protection, and restore testing, while geo-redundancy addresses resilience across physical locations. Together, these measures support business continuity and data recovery objectives for multi-site environments. See the [NIST contingency planning guidance](#) and [Microsoft guidance on cross-region replication](#).

QUESTION NO: 7

Monthly cloud service costs are BEST described as:

- A. operating expenditures.
- B. fixed expenditures.
- C. capital expenditures.
- D. personnel expenditures.

ANSWER: A

Explanation:

Monthly cloud service costs are best described as operating expenditures. Operating expenditures (OpEx) are the recurring costs an organization incurs to run its day-to-day business, such as monthly subscriptions, usage-based compute and storage charges, software services, support plans, and network consumption. Public cloud services commonly use consumption-based and subscription pricing, meaning the organization pays for resources over time rather than making a large upfront purchase for the underlying infrastructure. This makes the expense recurring and operational in nature, even when the monthly amount varies according to usage.

From a cloud financial-management perspective, moving workloads to cloud services can shift spending away from owning and depreciating physical infrastructure and toward ongoing service charges. Organizations should therefore forecast, monitor, allocate, and optimize these recurring costs through practices such as tagging, budgets, usage reviews, and rightsizing. The precise accounting treatment can depend on applicable accounting standards, contracts, and organizational policy, but monthly charges for consumed cloud services are generally managed as OpEx. See the [CompTIA guide to cloud computing](#) and the [Microsoft Cloud Adoption Framework guidance on cloud financial management](#).

QUESTION NO: 8

A Chief Information Officer (CIO) wants to identify two business units to be pilots for a new cloud project. A business analyst who was recently assigned to this project will be selecting a cloud provider. Which of the following should the business analyst do FIRST?

- A. Conduct a feasibility study of the environment.
- B. Conduct a benchmark of all major systems.
- C. Draw a matrix diagram of the capabilities of the cloud providers.
- D. Gather business and technical requirements for key stakeholders.

ANSWER: D

Explanation:

Gather business and technical requirements for key stakeholders is the appropriate first step because provider selection must be driven by the organization's intended outcomes and the needs of the pilot business units. The analyst should first establish what the pilots must accomplish, such as required workloads, expected growth, availability targets, security and compliance obligations, data-residency needs, integration dependencies, user access requirements, budget constraints, and

success criteria. Stakeholder input also identifies priorities and trade-offs, ensuring the cloud decision supports business value rather than being based solely on a feature comparison.

These documented requirements become the evaluation criteria used to assess potential cloud providers consistently and transparently. They enable the organization to determine whether a provider can meet both functional needs and operational expectations, including governance, support, resilience, and cost-management capabilities. This requirements-first approach aligns with cloud-adoption guidance that recommends defining strategy, business outcomes, and workload requirements before making platform decisions. See the [Microsoft Cloud Adoption Framework strategy guidance](#) and the [AWS Well-Architected Framework](#).

QUESTION NO: 9

Which of the following are true about the use of machine learning in a cloud environment? (Choose two).

- A. Specialized machine learning algorithms can be deployed to optimize results for specific scenarios.
- B. Machine learning can just be hosted in the cloud for managed services.
- C. Just one type of cloud storage is available in the cloud for machine learning workloads.
- D. Machine learning can leverage processes in a cloud environment through the use of cloud storage and auto-scaling.
- E. Machine learning requires a specialized IT team to create the machine learning models from scratch.
- F. Using machine learning solutions in the cloud removes the data-gathering step from the learning process.

ANSWER: A D

Explanation:

Specialized machine learning algorithms can be deployed to optimize results for specific scenarios because machine learning is selected and tuned according to the business problem, data characteristics, and required outcome. For example, an organization may deploy classification models for categorization, regression models for forecasting, or anomaly-detection models for identifying unusual behavior. Cloud platforms support this specialization by providing managed tools and infrastructure for developing, training, evaluating, deploying, and monitoring models. Microsoft describes how MLOps practices support the repeatable operational lifecycle of machine learning solutions, including model deployment and monitoring: [Microsoft Learn: MLOps principles](#).

Machine learning can leverage processes in a cloud environment through the use of cloud storage and auto-scaling because ML workloads commonly require large datasets and variable amounts of compute capacity. Cloud storage can retain source data, training data, model artifacts, and results, while auto-scaling can increase or reduce compute resources as demand changes. This enables teams to process substantial training workloads without permanently maintaining peak-capacity infrastructure. The NIST cloud-computing definition identifies rapid elasticity as a core cloud characteristic, describing resources that can be provisioned and released elastically to scale with demand: [NIST SP 800-145](#).

QUESTION NO: 10

Which of the following are the main advantages of using ML/AI for data analytics in the cloud as opposed to on premises? (Choose two.)

- A. Cloud providers offer enhanced technical support.
- B. Elasticity allows access to a large pool of compute resources.
- C. The shared responsibility model offers greater security.
- D. AI enables DevOps to build applications easier and faster.
- E. A pay-as-you-go approach allows the company to save money.
- F. ML enables DevOps to build applications easier and faster.

ANSWER: B E

Explanation:

“Elasticity allows access to a large pool of compute resources” is correct because machine learning and artificial intelligence analytics commonly require substantial, variable compute capacity for activities such as model training, hyperparameter tuning, and processing large data sets. Cloud environments can rapidly scale resources up when a workload needs more processing power and scale them down when the work is complete. This avoids the long procurement and deployment cycles associated with expanding an on-premises data center, while enabling teams to select specialized resources when needed.

“A pay-as-you-go approach allows the company to save money” is also correct because cloud consumption pricing can align analytics costs with actual use. Rather than purchasing, housing, maintaining, and periodically refreshing enough infrastructure to accommodate peak ML/AI demand, an organization can pay for the compute, storage, and managed services consumed during training and analysis. This is especially valuable for intermittent or unpredictable workloads. Cost savings depend on sound governance, monitoring, and right-sizing, but the consumption-based model is a central cloud advantage for data analytics. See [AWS: Six Advantages of Cloud Computing](#) and [Microsoft Cloud Adoption Framework: Financial outcomes](#).

QUESTION NO: 11

Which of the following is related to data availability in the cloud?

- A. Resiliency
- B. Deduplication
- C. Scalability
- D. Elasticity

ANSWER: A

Explanation:

Resiliency is related directly to data availability in a cloud environment. It is the capability of a service, application, or data platform to continue operating—or to restore operation quickly—when a component, availability zone, network path, or other dependency fails. Cloud resiliency is commonly achieved through redundant storage, replication across independent locations, automated failover, backups, and tested recovery procedures. These measures reduce the likelihood that a single failure makes data inaccessible and reduce the duration of an outage if one occurs.

Availability is generally measured by the proportion of time that authorized users can access data and services. A resilient design supports that goal by anticipating failures and maintaining or recovering access without unacceptable disruption. For example, replicating critical data to multiple locations and routing requests to a healthy instance can preserve access during an infrastructure failure. The AWS Well-Architected Reliability Pillar describes reliability as the ability of a workload to perform its intended function correctly and consistently, including recovering from failures: [AWS Well-Architected Reliability Pillar](#). Microsoft similarly identifies redundancy, replication, and disaster recovery as core approaches for resilient cloud applications: [Azure Well-Architected Framework—Reliability](#).

QUESTION NO: 12

For which of the following reasons is the infrastructure as code technique used?

- A. To improve the ability to migrate applications to different cloud providers to avoid vendor lock-in.
- B. To specify the hardware configuration and system installation procedures in machine-readable formats.
- C. To capture the system configurations to track changes and remediate configuration drift.
- D. To assign server administration responsibilities to software developers for application deployments.

ANSWER: C

Explanation:

To capture the system configurations to track changes and remediate configuration drift is correct because infrastructure as code (IaC) defines infrastructure's desired state in machine-readable, version-controlled files. Instead of relying on undocumented manual changes to servers, networks, storage, or cloud resources, teams use those definitions to provision and maintain consistent environments repeatedly. Version control provides a history of who changed a configuration, what was changed, and when it changed, which supports review, auditability, rollback, and controlled change management.

This approach is especially valuable for detecting and correcting configuration drift: the condition in which deployed resources no longer match the approved configuration baseline. An IaC tool can compare the declared state with the actual environment, identify differences, and apply the approved configuration again to restore consistency. This supports reliable deployments across development, test, and production environments, while reducing configuration errors caused by manual administration. IaC also enables automated validation of infrastructure definitions before changes are deployed. Microsoft describes IaC as managing and provisioning infrastructure through code rather than manual processes, and notes its benefits for consistency and repeatability. See [Microsoft Learn: What is infrastructure as code?](#) and [AWS: What is infrastructure as code?](#).

QUESTION NO: 13

A business analyst is drafting a risk assessment.

Which of the following components should be included in the draft? (Choose two.)

- A. Asset management
- B. Database type
- C. Encryption algorithms
- D. Certificate name
- E. Asset inventory
- F. Data classification

ANSWER: E F

Explanation:

An **Asset inventory** is a foundational input to a risk assessment because an organization must identify the systems, applications, data repositories, devices, and other resources it owns before it can determine what needs protection. The inventory establishes the assessment scope and supports identification of asset owners, dependencies, exposure points, and the potential business impact if a resource is compromised, unavailable, altered, or lost.

Data classification

should also be included because the sensitivity and business value of information directly affect risk prioritization and the selection of appropriate safeguards. Classifying data—for example, as public, internal, confidential, or regulated—helps the analyst evaluate the consequences of unauthorized disclosure or modification and identify relevant legal, contractual, and privacy obligations. Together, an asset inventory and data classification provide the context needed to assess threats, vulnerabilities, likelihood, and impact in a consistent, business-focused manner. NIST's risk-assessment guidance identifies assets and organizational impact as essential risk-assessment considerations, while the NIST Cybersecurity Framework emphasizes maintaining inventories and classifying managed data. See [NIST SP 800-30 Rev. 1](#) and the [NIST Cybersecurity Framework](#).

QUESTION NO: 14

For security reasons, a cloud service that can be accessed from anywhere would make BEST use of:

- A. replication.
- B. multifactor authentication.
- C. single sign-on.
- D. data locality

ANSWER: B

Explanation:

Multifactor authentication is the best security control for a cloud service accessible from anywhere because it requires users to verify their identity with two or more independent factors, such as something they know (a password), something they have (a security key or authenticator app), or something they are (a biometric). Since internet-accessible cloud services may receive login attempts from unmanaged devices and untrusted networks, a password alone is more vulnerable to phishing, password reuse, credential stuffing, and theft. Requiring an additional factor substantially reduces the chance that a compromised password can be used to access the service.

MFA is particularly appropriate for remote cloud access because it protects the authentication boundary regardless of the user's physical location. It can also be applied selectively through conditional-access policies, such as requiring stronger verification for unfamiliar devices, elevated administrative actions, or sign-ins that present higher risk. CompTIA's cloud-security objectives emphasize identity and access management as a core cloud control, and MFA is a widely adopted best practice for protecting accounts and cloud-hosted resources. See [CISA's guidance on turning on MFA](#) and [NIST SP 800-63B Digital Identity Guidelines](#).

QUESTION NO: 15

Resource consumption in a company's IaaS environment has been stable; however, a few servers have recently experienced spikes in CPU usage for days at a time. Costs are rising steadily, and it is unclear who owns the servers. Which of the following would work BEST to allow the company to charge the appropriate department? (Select TWO).

- A. Employ a software defined network.
- B. Employ orchestration automation.
- C. Review the storage usage reports.
- D. Review the network ingress reports.
- E. Employ resource tagging.
- F. Review the compute usage reports.

ANSWER: E F

Explanation:

Employ resource tagging because tags attach consistent business metadata to each IaaS resource, such as the owning department, cost center, application, environment, or project. Once ownership and cost-allocation tags are required and applied to servers, the organization can group billing data by those tags and attribute each server's charges to the responsible department. Tagging also creates an auditable governance mechanism: resources can be identified even when teams change, workloads scale, or new instances are deployed.

Review the compute usage reports because the reported CPU and instance-consumption data establishes the measurable usage associated with the servers experiencing sustained spikes. The company can correlate compute consumption and its resulting charges with the owner or cost-center tag, producing a defensible showback or chargeback report. This supports both accurate departmental billing and informed follow-up on the workload causing the increased consumption. Cloud cost-management guidance commonly recommends resource tags for cost allocation and reporting, while provider usage data supplies the consumption details needed to analyze and allocate costs. See [AWS cost allocation tags documentation](#) and [Microsoft Azure tag-based cost allocation guidance](#).

QUESTION NO: 16

Which of the following describes the process of moving an application from an isolated data center to reduce latency and ensure close proximity to end users?

- A. Replication
- B. Zones
- C. Geo-redundancy
- D. Backup

ANSWER: C

Explanation:

Geo-redundancy is correct because it distributes application components, services, or data across geographically separate locations rather than relying on one isolated data center. Locating workloads in multiple regions or data centers enables users to be served from a location nearer to them, which can reduce network distance and latency. This design also improves resilience: if a facility or an entire region experiences an outage, services can continue from another geographic location, subject to the application's failover and data-consistency design. In cloud environments, geo-redundancy is commonly implemented with regional deployments, traffic routing, replicated data, and health-based failover. The goal is therefore both performance for geographically dispersed users and continuity of service during localized failures. Microsoft describes geo-redundant storage as replicating data to a secondary region that is hundreds of miles from the primary region, helping protect against regional outages. Its reliability guidance also emphasizes deploying applications across regions to address failures that affect an entire region. See [Azure Storage redundancy](#) and [Azure disaster recovery and availability](#).

QUESTION NO: 17

The IT department has noticed the costs associated with the server are increasing. The application is scheduled to run uninterrupted for at least two years, and the Chief Financial Officer (CFO) would like the IT department to investigate ways to decrease the costs. Which of the following would MOST likely help to decrease the costs of the server?

- A. Migrate the server to another CSP.
- B. Convert the server to a reserved instance.
- C. Configure resource tagging on the server.
- D. Run the server in a sandbox environment.
- E. Utilize spot instances.

ANSWER: B

Explanation:

Convert the server to a reserved instance is correct because the workload has a known, long-term, continuous runtime requirement of at least two years. Reserved-instance pricing models are designed for stable and predictable usage, allowing an organization to commit to a specific capacity or usage level for a defined term, commonly one or three years, in return for a lower effective rate than standard on-demand consumption. This converts a variable, pay-as-you-go cost into a more predictable committed expense and can substantially reduce the total cost of operating a server that is expected to remain active throughout the commitment period.

The stated uninterrupted two-year schedule is the key business and technical requirement: it provides the utilization certainty needed to make a long-term reservation financially appropriate. Before committing, IT should validate that the instance family, region, operating system, capacity requirements, and expected application lifetime will remain suitable for the selected reservation. The organization should also compare payment choices, such as no upfront, partial upfront, and all upfront, against its budget and cash-flow needs. AWS describes Reserved Instances and their term-based pricing benefits in its [Amazon EC2 Reserved Instances documentation](#) and provides additional commitment-pricing guidance through [AWS Savings Plans](#).

QUESTION NO: 18

Which of the following are aspects of cloud data availability? (Choose two.)

- A. Resource tagging
- B. Data sovereignty
- C. Locality
- D. Zones
- E. Geo-redundancy
- F. Auto-scaling

ANSWER: D E

Explanation:

Zones and geo-redundancy are core aspects of cloud data availability because they reduce the risk that a single infrastructure failure makes data inaccessible. Availability zones are physically separate locations within a cloud region, with independent power, networking, and cooling. Storing or replicating data across zones helps maintain access when one datacenter or zone has an outage. Geo-redundancy extends this resilience by maintaining copies of data in a geographically separate region. This supports continued data access and recovery when a broader regional disruption affects the primary location.

These approaches implement fault isolation and replication, which are fundamental availability design practices in cloud environments. A cloud customer can select a service and storage configuration that replicates data across zones for local resilience and across regions for disaster resilience. The appropriate design should also account for recovery objectives, replication behavior, and any applicable geographic or regulatory requirements. For further guidance, see [Microsoft Azure availability zones overview](#) and [AWS S3 replication documentation](#).

QUESTION NO: 19

A company's SaaS provider recently changed its licensing model, and a business analyst is required to do an overall cost analysis for a three-year contract renewal. Which of the following will provide the entire financial forecast over the renewal period?

- A. ROI
- B. TCO
- C. RFI

ANSWER: B

Explanation:

TCO is the appropriate measure because Total Cost of Ownership provides a comprehensive view of all costs associated with using a service across its expected term. For a three-year SaaS renewal, the analyst can use TCO to forecast recurring subscription or license charges under the new model, anticipated user-growth costs, implementation or migration work, integrations, administration, training, support, and any other directly attributable operational expenses. This produces a full cost baseline for the renewal period rather than focusing on a single purchase price or isolated contract term.

A complete TCO analysis is particularly valuable when a provider changes licensing because a change from, for example, per-user to consumption-based pricing can alter both predictable and variable cost drivers. The analyst can model expected consumption, growth assumptions, contractual commitments, and supplemental service costs over each year of the agreement, then aggregate them into an informed financial forecast. This supports budgeting, comparison with alternatives, and negotiation decisions by representing the overall economic commitment of retaining the SaaS service. CompTIA Cloud Essentials+ includes business and financial considerations of cloud services as a core knowledge area; see [CompTIA Cloud](#)

QUESTION NO: 20

Which of the following techniques helps an organization determine benchmarks for application performance within a set of resources?

- A. Auto-scaling
- B. Load testing
- C. Sandboxing
- D. Regression testing

ANSWER: B

Explanation:

Load testing is correct because it measures how an application performs when it receives an expected or increasing volume of concurrent user activity, requests, or transactions while operating within a defined resource configuration. By establishing a repeatable workload and observing metrics such as response time, throughput, error rate, CPU utilization, memory consumption, and database latency, an organization can create meaningful performance benchmarks. These benchmarks document the performance that can be expected from the application with a particular set of compute, storage, network, and database resources.

The results also provide an evidence-based baseline for capacity planning and cloud decisions. Teams can identify the point at which user experience begins to degrade, determine whether the allocated resources satisfy service-level objectives, and compare performance after an infrastructure or application change. A valid load test uses representative usage patterns and realistic data volumes so that its measurements reflect anticipated production conditions. AWS describes load testing as testing a system under a specified expected load to measure response and behavior, while Microsoft explains that performance testing evaluates responsiveness and stability under workload. See the [AWS Well-Architected Reliability guidance](#) and [Microsoft Azure Load Testing documentation](#).

QUESTION NO: 21

A company is preparing a business case for migrating an internally hosted application to the cloud.

Which of the following should be included when calculating the total cost of ownership? (Choose two.)

- A. Cloud service consumption costs
- B. Personnel costs
- C. Only the original server purchase price
- D. Only the cloud provider's list price
- E. Employee desktop background preferences

ANSWER: A B

Explanation:

Cloud service consumption costs should be included because the organization will incur recurring charges for resources such as compute, storage, databases, network services, and support. These costs should be estimated using expected usage patterns and future growth assumptions.

Personnel costs should also be included because cloud adoption can change staffing requirements for administration, security, operations, training, and vendor management. Total cost of ownership considers more than the provider invoice; it

includes the full lifecycle cost of operating the service. The [Microsoft Cloud Adoption Framework financial guidance](#) discusses financial planning for cloud adoption. The [FinOps Framework](#) also addresses cloud financial management.

QUESTION NO: 22 - (SIMULATION)

A network team establishes a new connection to an IaaS CSP that is more efficient and has networking costs that are 25% less than previous monthly expenditures. The bill outlines the following costs:

Storage:\$10000

Compute:\$12000

Network:\$7000

Which of the following will be the total cloud expenditure for the following month?

- A.
\$26000
- B.
\$26250
- C.
\$27250
- D.
\$29000

ANSWER: See the explanation for the answer

Explanation:

Correct answer: C. \$27,250

The new network cost is 25% less than the prior \$7,000 network charge:

Network savings: $\$7,000 \times 25\% = \$1,750$
New network cost: $\$7,000 - \$1,750 = \$5,250$

Total: Storage (\$10,000) + Compute (\$12,000) + Network (\$5,250)
Total cloud expenditure = \$27,250

QUESTION NO: 23

Which of the following aspects of cloud design enables a customer to continue doing business after a major data center incident?

- A. Replication
- B. Disaster recovery
- C. Scalability
- D. Autoscaling

ANSWER: B

Explanation:

Disaster recovery is correct because it is the set of people, processes, technologies, and plans used to restore IT services and business operations following a disruptive event, including a major data center incident. In a cloud design, disaster recovery commonly involves maintaining protected copies of critical data and workloads in a separate availability zone, region, or provider environment, then executing tested procedures to recover those systems when the primary site is unavailable. A well-designed disaster recovery strategy defines recovery time objectives (RTOs), which identify how quickly a service must be restored, and recovery point objectives (RPOs), which identify the maximum acceptable amount of data loss. These targets guide choices such as backup frequency, replication architecture, failover automation, and the capacity maintained at the recovery location. Cloud services can make disaster recovery more practical by providing geographically distributed infrastructure and on-demand resources that can be activated during an incident. The objective is business continuity: restoring the applications, data access, and supporting services needed for the organization to continue operating within its defined recovery requirements. See [CISA's disaster recovery guidance](#) and [AWS's disaster recovery overview](#).

QUESTION NO: 24

Which of the following documents has the sole purpose of outlining a professional services engagement that governs a proposed cloud migration?

- A. Gap analysis
- B. Statement of work
- C. Feasibility study
- D. Service level agreement

ANSWER: B

Explanation:

Statement of work is correct because it is the formal document used to define and govern a specific professional-services engagement, such as planning and executing a cloud migration. It establishes the agreed scope of work, expected deliverables, project milestones, roles and responsibilities, assumptions, acceptance criteria, timeline, and commercial terms. By documenting these elements before implementation begins, the statement of work gives both the customer and the cloud consultant or provider a common, enforceable understanding of what migration services will be performed and how successful completion will be assessed.

For a proposed cloud migration, this document commonly identifies the workloads or applications in scope, discovery and assessment activities, migration approach, security and compliance requirements, testing, cutover support, knowledge transfer, and any exclusions or customer dependencies. It therefore serves as the engagement-level agreement that translates the migration proposal into an actionable services commitment. The U.S. General Services Administration describes a statement of work as the document that specifies work requirements for a contract: [GSA Statement of Work guidance](#). AWS also describes statements of work as defining professional-services project scope and deliverables: [AWS Professional Services](#).