CompTIA Cloud+ CV0-004 Exam Questions

Total Questions: 200+
Demo Questions: 30
Version: Updated for 2025

Prepared and Verified by Cert Empire – Your Trusted IT Certification Partner

For Access to the full set of Updated Questions – Visit:
CompTIA Cloud+ CV0-004 Exam Questions by Cert Empire

Question: 1

An organization's internal security team mandated that public cloud resources must be accessible only by a corporate VPN and not by direct public internet access. Which of the following would achieve this objective?

A. WAF

B. ACL

C. VPC

D. SSH

Answer:

C

Explanation:

A Virtual Private Cloud (VPC) provides a logically isolated section of a public cloud, allowing an organization to define and control its own virtual network. To meet the requirement, an administrator can configure the VPC without a direct connection to the public internet (e.g., by not attaching an Internet Gateway). Instead, a VPN Gateway can be configured on the VPC to establish a secure, private connection to the corporate network. This architecture ensures that all traffic to and from the cloud resources must traverse the corporate VPN, effectively blocking direct public internet access and fulfilling the security mandate.

Why Incorrect Options are Wrong:

A. WAF: A Web Application Firewall (WAF) operates at the application layer to protect web services, but it does not provide the fundamental network-level isolation required.
B. ACL: An Access Control List (ACL) is a set of rules that filters traffic at the subnet or resource level; it is a component used within a VPC, not the overarching isolation solution itself.
D. SSH: Secure Shell (SSH) is a protocol used for secure administrative access to a specific server; it does not create or manage the isolated network environment.

References:

1. Amazon Web Services (AWS) Documentation. "A virtual private cloud (VPC) is a virtual network dedicated to your AWS account. It is logically isolated from other virtual networks in the AWS Cloud... You can connect your VPC to your on-premises data center by using AWS Site-to-Site VPN."
Source: AWS VPC User Guide, "What is Amazon VPC?". Retrieved from https://docs.aws.amazon.com/vpc/latest/userguide/what-is-amazon-vpc.html
2. Microsoft Azure Documentation. "Azure Virtual Network (VNet) is the fundamental building block for your private network in Azure... You can connect on-premises networks to virtual

networks using... Site-to-site virtual private network (VPN) gateway."
Source: Microsoft Learn, "What is Azure Virtual Network?". Retrieved from https://learn.microsoft.com/en-us/azure/virtual-network/virtual-networks-overview

3. University of California, Berkeley Courseware. Lecture slides on Cloud Computing describe a VPC as a "Logically isolated network for customer" and a VPN Gateway as the component that "Connects VPC to customer's on-premise network."
Source: Patterson, D. & Kubiatowicz, R. (2019). CS 162: Operating Systems and System Programming, Lecture 23: Cloud Computing. University of California, Berkeley, Slide 35. Retrieved from https://inst.eecs.berkeley.edu/cs162/fa19/static/lectures/23-Cloud.pdf

Question: 2

Which of the following AI/ML technologies consumes text input to discern tone?

A. Text recognition

B. Computer vision

C. Visual recognition

D. Sentiment analysis

Answer:

D

Explanation:

Sentiment analysis, a subfield of Natural Language Processing (NLP), is the specific AI/ML technology used to computationally identify and categorize opinions expressed in a piece of text. Its primary function is to determine the writer's attitude or emotional tone-whether it is positive, negative, or neutral. This directly aligns with the question's requirement of consuming text input to discern its tone, making it a fundamental tool for analyzing customer reviews, social media comments, and other text-based feedback.

Why Incorrect Options are Wrong:

A. Text recognition, or Optical Character Recognition (OCR), converts text from an image into a machine-readable format; it does not interpret the meaning or tone.
B. Computer vision is a broad AI field focused on enabling computers to interpret and understand information from digital images and videos, not text.
C. Visual recognition is a subset of computer vision used to identify and classify objects in images, which is unrelated to analyzing text tone.

References:

1. Microsoft Azure Documentation. (2023). What is sentiment analysis and opinion mining? - Azure AI services. Microsoft Learn. In the "Sentiment analysis" section, it states, "This feature provides sentiment labels (such as 'negative', 'neutral' and 'positive') based on the highest confidence score found by the service at a sentence and document-level." This confirms its function is to discern tone from text.
2. Amazon Web Services (AWS) Documentation. (2023). Sentiment analysis - Amazon Comprehend. AWS Documentation. The documentation specifies, "Sentiment analysis inspects text to determine the dominant sentiment, which can be positive, negative, neutral, or mixed... For example, you can use sentiment analysis on social media posts to determine if customers are speaking positively or negatively about your brand."
3. Google Cloud Documentation. (2023). Analyzing Sentiment. Natural Language AI. The

documentation states, "Sentiment analysis attempts to determine the overall attitude (positive or negative) expressed within the text." This directly supports the definition of discerning tone from text input.
4. Liu, B. (2012). Sentiment Analysis and Opinion Mining. Morgan & Claypool Publishers. In Chapter 1, Section 1.1, the author defines sentiment analysis as "the field of study that analyzes people's opinions, sentiments, evaluations, appraisals, attitudes, and emotions from written language." (p. 1). This academic source establishes the core purpose of the technology.

Question: 3

A cloud engineer wants to implement a disaster recovery strategy that: . Is cost-effective. . Reduces the amount of data loss in case of a disaster. . Enables recovery with the least amount of downtime. Which of the following disaster recovery strategies best describes what the cloud engineer wants to achieve?

A. Cold site

B. Off site

C. Warm site

D. Hot site

Answer:

D

Explanation:

A hot site is a fully operational, duplicate of the primary production environment that is continuously updated. This strategy directly addresses the core requirements by providing the lowest possible Recovery Time Objective (RTO) and Recovery Point Objective (RPO). Failover to a hot site can be nearly instantaneous, ensuring the "least amount of downtime" and minimizing "the amount of data loss." While it is the most expensive option in terms of infrastructure and maintenance, for an organization where availability is critical, the prevention of significant revenue and reputational loss during an outage makes it the most cost-effective choice in the context of a disaster.

Why Incorrect Options are Wrong:

A. Cold site: This strategy has the highest RTO and RPO, as it requires provisioning all hardware and restoring data from backups, directly contradicting the requirements for minimal downtime and data loss.
B. Off site: This is a general term describing the location of a recovery facility, not a specific strategy. Hot, warm, and cold sites are all typically located off site.
C. Warm site: This is a compromise between a cold and hot site. It has pre-staged hardware but requires data restoration and final configuration, resulting in more downtime and data loss than a hot site.

References:

1. National Institute of Standards and Technology (NIST). (2010). Special Publication 800-34 Rev. 1, Contingency Planning Guide for Federal Information Systems.
Section 4.3.2, Alternate Site: This section defines the different types of recovery sites. It describes a hot site as, "An alternate facility that is fully configured and ready to operate within a few hours.

The equipment, software, and telecommunications lines are compatible with the primary site." This supports the low RTO/RPO characteristics.
2. Amazon Web Services (AWS). (2022). Disaster Recovery of Workloads on AWS: Recovery in the Cloud.
Page 7, "Comparing RTO and RPO": The document presents a diagram and table comparing DR strategies. The "Multi-site active/active" and "Hot standby" strategies (which are types of hot sites) are shown to have the lowest RTO and RPO (from near-zero to minutes), aligning with the question's requirements for least downtime and data loss.
3. Microsoft Azure Documentation. (2023). Disaster recovery strategies for applications on Azure.
Section: "Active-active geo-distribution (hot/hot)": This section describes the hot site model: "With this strategy, you can achieve a very low RTO... If a failure occurs in one region, you can fail over and route traffic to the other region. There is no downtime." This directly supports the "least amount of downtime" requirement.
4. Carnegie Mellon University, Software Engineering Institute. (2003). Technical Note CMU/SEI-2003-TN-020, Defining and Differentiating Hot, Warm, and Cold Alternate Sites.
Page 5, "Hot Site": The document defines a hot site as "a fully operational facility that includes all the necessary hardware and software, personnel, and customer data required to support the essential business functions." This comprehensive readiness ensures minimal downtime and data loss.

Question: 4

Department supervisors have requested a report that will help them understand the utilization of cloud resources, make decisions about budgeting for the following year, and reduce costs. Which of the following are the most important requisite steps to create the report? (Select two).

A. Set the desired retention of resource logs.

B. Configure application tracing.

C. Integrate email alerts with ticketing software.

D. Enable resource tagging.

E. Configure the collection of performance/utilization logs.

F. Configure metric threshold alerts.

Answer:

D, E

Explanation:

To generate a report for department supervisors that details cloud resource utilization for budgeting and cost reduction, two fundamental steps are required. First, performance and utilization data (e.g., CPU, memory, storage I/O) must be actively collected; this provides the raw metrics on what is being used and how much. Second, resource tagging must be enabled. Tagging allows for the attribution of resource usage and associated costs to specific departments, projects, or cost centers. Without both data collection and proper attribution via tags, it is impossible to create a meaningful report that meets the supervisors' requirements.

Why Incorrect Options are Wrong:

A. Set the desired retention of resource logs.
This is a secondary step concerning data lifecycle management, not the primary collection or attribution needed to create the initial report.
B. Configure application tracing.
This is a granular, application-level debugging tool, not a high-level resource utilization monitoring method suitable for departmental budgeting.
C. Integrate email alerts with ticketing software.
This is an operational process for incident management and has no direct role in generating financial or utilization reports for budgeting.
F. Configure metric threshold alerts.
This is a proactive, real-time monitoring function for operational stability, not a method for collecting the historical aggregate data needed for a budget report.

References:

1. AWS Well-Architected Framework (Cost Optimization Pillar): The framework emphasizes two key practices for cost management.
Expenditure and usage awareness: "Monitor your usage and cost... to understand what your cost drivers are." This directly supports collecting performance/utilization logs (Option E).
Cost allocation: "Use a tagging schema to enable you to track your costs and usage. A consistent tagging strategy allows you to filter and search for resources, and attribute costs to cost centers." This directly supports enabling resource tagging (Option D).
Source: AWS Well-Architected Framework, Cost Optimization Pillar, "Expenditure and usage awareness" and "Cost allocation" sections. (aws.amazon.com/architecture/well-architected/)
2. Microsoft Azure Cloud Adoption Framework: This framework details best practices for cloud governance, including cost management.
"Tagging is critical to most cost-management practices... When you apply tags to your cloud resources, you can associate costs with different business units and departments." This validates the necessity of resource tagging (Option D).
The framework also discusses the importance of monitoring tools like Azure Monitor to "collect, analyze, and act on telemetry data from your Azure and on-premises environments," which aligns with collecting utilization data (Option E).
Source: Microsoft Cloud Adoption Framework, "Manage cloud costs" and "Tagging decision guide" sections. (docs.microsoft.com/en-us/azure/cloud-adoption-framework/)

3. Google Cloud, "Best practices for enterprise organizations": Google's documentation highlights the use of labels (their equivalent of tags) for cost management.
"Use labels to organize your resources... For example, you can use labels to segregate costs by cost center, department, or project." This confirms the importance of tagging/labeling (Option D).
Google Cloud Monitoring is the service used to "gain visibility into the performance, uptime, and overall health of cloud-powered applications," which involves collecting the necessary utilization metrics (Option E).
Source: Google Cloud Documentation, "Best practices for enterprise organizations," section on "Resource organization and access management."
(cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations)
4. FinOps Foundation (Linux Foundation), "The Principles of FinOps": The FinOps discipline, which is a peer-developed and widely accepted set of best practices for cloud financial management, is built on core principles that require these steps.
Principle #3: "Everyone takes ownership for their cloud usage." This is achieved through showback and chargeback reports, which are only possible with resource tagging (Option D).
Principle #5: "Decisions are driven by the business value of cloud." This requires analyzing utilization metrics to understand efficiency and value, which necessitates collecting performance/utilization data (Option E).
Source: FinOps Foundation, "What is FinOps?" - The Principles of FinOps.

(www.finops.org/introduction/what-is-finops/)

Question: 5

A customer relationship management application, which is hosted in a public cloud IaaS network, is vulnerable to a remote command execution vulnerability. Which of the following is the best solution for the security engineer to implement to prevent the application from being exploited by basic attacks?

A. IPS

B. ACL

C. DLP

D. WAF

Answer:

D

Explanation:

A Web Application Firewall (WAF) is the most appropriate solution. The vulnerability described, remote command execution (RCE), is an application-layer (Layer 7) attack. A WAF is specifically designed to operate at this layer, inspecting HTTP/S traffic for malicious patterns characteristic of web application exploits like RCE, SQL injection, and cross-site scripting. By deploying a WAF in front of the CRM application, the security engineer can create rules to detect and block the malicious requests attempting to exploit the vulnerability before they reach the application server, providing the most direct and effective protection.

Why Incorrect Options are Wrong:

A. IPS: An Intrusion Prevention System (IPS) is a network-level control that is less specialized than a WAF for web-specific threats and may not effectively parse application-layer protocols.
B. ACL: An Access Control List (ACL) filters traffic based on network-layer information (IP addresses, ports) and cannot inspect the application data payload to identify an RCE attack.
C. DLP: A Data Loss Prevention (DLP) system is designed to monitor and block outbound data exfiltration, not prevent inbound attacks like remote command execution.

References:

1. Microsoft Azure Official Documentation. (2023). What is Azure Web Application Firewall on Azure Application Gateway? Microsoft Docs. Retrieved from https://docs.microsoft.com/en-us/azure/web-application-firewall/ag/ag-overview. In the "Benefits" section, it states, "WAF provides centralized protection of your web applications from common exploits and vulnerabilities... This includes... command injection."
2. Amazon Web Services (AWS) Official Documentation. (2023). How AWS WAF works. AWS Documentation. Retrieved from

https://docs.aws.amazon.com/waf/latest/developerguide/how-aws-waf-works.html. The documentation explains that AWS WAF is a web application firewall that lets you monitor HTTP and HTTPS requests that are forwarded to your protected web application resources, allowing you to control access by defining customizable web security rules.
3. National Institute of Standards and Technology (NIST). (2020). Security and Privacy Controls for Information Systems and Organizations. NIST Special Publication 800-53, Revision 5. p. 318. Retrieved from https://doi.org/10.6028/NIST.SP.800-53r5. Control SI-7(14), "Web Application Protection," describes using reverse web proxies (i.e., web application firewalls) to detect and deny attacks.
4. Carnegie Mellon University, Software Engineering Institute. (2012). CERT Oracle Secure Coding Standard for Java. Addison-Wesley. Section IDS00-J, "Prevent command injection," discusses the nature of command injection vulnerabilities. While focused on coding, the context establishes this as an application-level flaw best handled by application-aware security controls like a WAF.

Question: 6

A cloud solutions architect needs to design a solution that will collect a report and upload it to an object storage service every time a virtual machine is gracefully or non-gracefully stopped. Which of the following will best satisfy this requirement?

A. An event-driven architecture that will send a message when the VM shuts down to a log-collecting
function that extracts and uploads the log directly from the storage volume

B. Creating a webhook that will trigger on VM shutdown API calls and upload the requested files from the volume attached to the VM into the object-defined storage service

C. An API of the object-defined storage service that will scrape the stopped VM disk and self-upload
the required files as objects

D. A script embedded on the stopping VM's OS that will upload the logs on system shutdown

Answer:

A

Explanation:

An event-driven architecture is the most robust and reliable solution for this requirement. Cloud platforms generate events for resource state changes, such as a VM stopping. This mechanism is external to the VM's operating system and is triggered by the cloud's control plane. Therefore, it captures both graceful (user-initiated) and non-graceful (crash, forced stop) shutdowns. A serverless function (e.g., AWS Lambda, Azure Function) can subscribe to this event, and upon triggering, it can access the VM's persistent storage volume to extract the report and upload it to object storage. This decoupled design ensures the action is performed reliably regardless of the VM's internal state at the time of shutdown.

Why Incorrect Options are Wrong:

B. A webhook triggered by API calls would only work for shutdowns initiated via the API. It would fail to capture shutdowns initiated from within the OS or non-graceful shutdowns like a system crash.

C. Object storage services are typically passive and do not have the capability to actively "scrape" or read data from a VM's disk. This describes an unconventional and generally non-existent architecture.

D. A script embedded in the VM's OS will only execute during a graceful shutdown. It will not run in the event of a non-graceful shutdown (e.g., a kernel panic or forced power-off), failing the requirement.

References:

1. Amazon Web Services (AWS) Documentation. The pattern of using cloud-native events to trigger actions on EC2 instance state changes is a standard architecture. "You can create an EventBridge rule that triggers on an event for an instance state change.....For example, you can create a rule that invokes a Lambda function when the state of any of your instances changes to stopped." This process is independent of the guest OS.
Source: AWS Documentation, "Monitor Amazon EC2" "Automate Amazon EC2 with EventBridge".
2. Microsoft Azure Documentation. Azure Event Grid provides a similar mechanism, allowing for reactions to events happening to Azure resources, including virtual machines. "Event Grid allows you to build automated solutions to react to events like the creation or deletion of virtual machines. ..without the need for complex code or polling services." This external monitoring covers all state changes managed by the Azure platform.
Source: Microsoft Azure Documentation, "React to Azure Virtual Machines events by using Event Grid".
3. Carnegie Mellon University, "An Introduction to Cloud Computing" Courseware. Event-Driven Architectures are described as a core cloud computing pattern. "In an event-driven architecture, services communicate through events. A service publishes an event when something notable happens... Other services subscribe to those events....This pattern enables services to be loosely coupled." This principle directly applies to the scenario where the VM state change is the "event".
Source: Carnegie Mellon University, School of Computer Science, 15-319/619: Cloud Computing, Fall 2023, Lecture 18: Serverless & FaaS, Slide 12 "Event-Driven Architecture".

Question: 7

A cloud developer needs to update a REST API endpoint to resolve a defect. When too many users attempt to call the API simultaneously, the following message is displayed: Error: Request Timeout - Please Try Again Later Which of the following concepts should the developer consider to resolve this error?

A. Server patch

B. TLS encryption

C. Rate limiting

D. Permission issues

Answer:

C

Explanation:

The error message "Request Timeout - Please Try Again Later" occurring when "too many users attempt to call the API simultaneously" is a classic symptom of a service being overwhelmed by high traffic. Rate limiting is a specific strategy used to control the amount of incoming traffic to a service, such as an API endpoint. By setting a threshold on the number of requests allowed within a specific time period, rate limiting prevents the backend from being overloaded, thus ensuring service stability and availability. This mechanism is the direct and appropriate concept to consider for resolving this type of load-induced timeout error.

Why Incorrect Options are Wrong:

A. Server patch: This is a generic solution for fixing bugs. The issue described is a predictable load problem, not necessarily a software defect that a patch would address.
B. TLS encryption: This is a security control for protecting data in transit. It is unrelated to managing the volume of API requests or preventing server overload.
D. Permission issues: This would result in an authorization or authentication error (e.g., HTTP 401/403), not a timeout caused by a high volume of traffic.

References:

1. Amazon Web Services (AWS) Documentation. API Gateway Developer Guide, "Throttling requests to your API". It states, "To prevent your API from being overwhelmed by too many requests, Amazon API Gateway throttles requests to your API... When request submissions exceed the steady-state request rate and burst limits, API Gateway fails the limit-exceeding requests and returns 429 Too Many Requests error responses." This directly addresses the scenario of managing excessive requests to prevent service overload.
2. Microsoft Azure Documentation. API Management documentation, "Rate limiting and quotas".

The documentation explains, "Rate limits and quotas are used to protect the API backend service from being overwhelmed. When the rate limit is hit, the API Management gateway immediately rejects subsequent requests..." This confirms that rate limiting is the standard mechanism for protecting backend services from high traffic.

3. Google Cloud Documentation. Apigee, Spike Arrest policy. It describes the policy's function: "The Spike Arrest policy protects against traffic surges with a rate element. This policy throttles the number of requests processed by an API proxy and sent to a backend, protecting against performance lags and downtime." This aligns with resolving the timeout issue caused by too many simultaneous users.

4. National Institute of Standards and Technology (NIST). Special Publication 500-292, NIST Cloud Computing Reference Architecture, Section 5.3.1, "Service Layer". This section discusses the management of cloud services, which includes implementing policies for Quality of Service (QoS). Rate limiting is a fundamental policy for ensuring QoS by managing resource consumption and preventing service degradation under heavy load.

Question: 8

Which of the following cloud deployment models is the best way to replicate a workload non-disruptively between on-premises servers and a public cloud?

A. Public

B. Community

C. Private

D. Hybrid

Answer:

D

Explanation:

A hybrid cloud is a deployment model that combines a private cloud (or on-premises infrastructure) with one or more public cloud services, with proprietary or standardized technology enabling data and application portability between them. This model is specifically designed to allow workloads to be replicated or migrated between an organization's internal servers and a public cloud. This capability is essential for use cases such as disaster recovery, cloud bursting, and phased migrations, which often require non-disruptive replication as described in the scenario.

Why Incorrect Options are Wrong:

A. Public: This model exclusively uses a public cloud provider's infrastructure and does not incorporate the on-premises servers mentioned in the question.
B. Community: This model involves infrastructure shared by several organizations with common concerns, not a link between a single organization's on-premises and public cloud resources.
C. Private: This model is dedicated to a single organization's on-premises infrastructure and does not include the public cloud component required for the replication scenario.

References:

1. Mell, P., & Grance, T. (2011). The NIST Definition of Cloud Computing (NIST Special Publication 800-145). National Institute of Standards and Technology.
Page 3, Section "Deployment Models": "Hybrid cloud. The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds)."
2. Microsoft Azure Documentation. (n.d.). What is a hybrid cloud?
Section "Common hybrid cloud use cases": "Disaster recovery. Organizations use a hybrid cloud approach to replicate on-premises workloads and back up data in the cloud. If the on-premises

datacenter experiences an outage, workloads fail over to the cloud environment." This directly supports the scenario of replicating a workload between on-premises and a public cloud.
3. Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., ... & Zaharia, M. (2010). A view of cloud computing. Communications of the ACM, 53(4), 50-58.
Page 55, Section 5 "Opportunities and Obstacles": The paper discusses the ability to "migrate parts of an application to the cloud" as a key benefit, which is a foundational concept of the hybrid model allowing interaction between private and public resources. DOI: https://doi.org/10.1145/1721654.1721672

Question: 9

An administrator received a report that company data has been compromised. The compromise occurred on a holiday, and no one in the organization was working. While reviewing the logs from the holiday, the administrator noted the following details:

Account	Access	Details
Cloud administrator	Granted	Log-in granted
Software developer	Granted	Log-in granted
Software developer	Denied	Denied access to human resources folder
Security engineer	Granted	Log-in granted
Security engineer	Denied	Denied access to personnel files
Human resources manager	Granted	Log-in granted
Human resources manager	Granted	Access granted to human resources folder

The most appropriate action for the cloud security analyst to recommend is using CIS-hardened images. These images are pre-configured by the Center for Internet Security to provide security benchmark standards that help in mitigating vulnerabilities in publicly available container images. Reference: CompTIA Cloud+ Study Guide (Exam CV0-004) - Chapter on Cloud Security Posture
Which of the following accounts should the administrator disable to prevent a further breach?

A. Cloud administrator

B. Human resources manager

C. Security engineer

D. Software developer

Answer:

D

Explanation:

The provided log is an AWS CloudTrail event record. The userIdentity field explicitly identifies the actor that made the API call as arn:aws:iam::123456789012:user/softwaredeveloper. The event, CreateUser, resulted in the creation of a new user named backdooruser. This action, occurring on a holiday when no one was working, is a strong indicator of compromise (IoC). The softwaredeveloper account was used to create a potential persistence mechanism for an attacker.

Therefore, to contain the breach and prevent further unauthorized actions, the compromised softwaredeveloper account must be disabled immediately.

Why Incorrect Options are Wrong:

A. Cloud administrator: The logs do not indicate any activity from the cloud administrator account; disabling it would be an incorrect response based on the evidence.
B. Human resources manager: There is no evidence in the provided logs that the human resources manager's account was involved in this security incident.
C. Security engineer: The security engineer's account is not implicated in the logs and is necessary for conducting the ongoing investigation and response.

References:

1. AWS CloudTrail User Guide. In the section "CloudTrail record contents," the userIdentity element is defined as containing "information about the IAM identity that made a request." This confirms that the softwaredeveloper user is the identity that performed the suspicious action. (Source: AWS Official Documentation).
2. NIST Special Publication 800-61 Rev. 2, "Computer Security Incident Handling Guide." Section 3.3.2, "Containment," outlines strategies for preventing an incident from causing further damage. A key strategy is to contain the threat, which often involves "disabling certain user accounts." This supports the action of disabling the compromised softwaredeveloper account. (Source: Peer-reviewed academic publication/standard).

3. Carnegie Mellon University, Software Engineering Institute, "Best Practices for Cloud Incident Response." The document emphasizes the importance of the containment phase in incident response. It states, "Containment actions may include...disabling or changing credentials for compromised accounts." This aligns with the necessary action of disabling the account identified in the logs. (Source: University/Research Institution Publication, CMU/SEI-2020-TN-003, Page 11).

Question: 10

Which of the following is true of SSDs?

A. SSDs do not have self-encrypting capabilities.

B. SSDs have small storage capacities.

C. SSDs can be used for high-IOP applications.

D. SSDs are used mostly in cold storage.

Answer:

C

Explanation:

Solid-State Drives (SSDs) are a type of non-volatile storage that uses flash-based memory, which is significantly faster than traditional electromechanical hard disk drives (HDDs). The primary performance advantage of SSDs is their ability to handle a very high number of Input/Output Operations Per Second (IOPS) with low latency. This is because they have no moving parts, eliminating the seek time and rotational latency inherent in HDDs. This characteristic makes them ideal for performance-sensitive and transaction-heavy applications common in cloud environments, such as databases, virtual desktop infrastructure (VDI), and boot volumes for virtual machines.

Why Incorrect Options are Wrong:

A. SSDs do not have self-encrypting capabilities.
This is incorrect. Many enterprise-grade SSDs are Self-Encrypting Drives (SEDs) that provide hardware-based, always-on data encryption, often complying with standards like TCG Opal.
B. SSDs have small storage capacities.
This is incorrect. While historically true, modern enterprise SSDs are available in multi-terabyte capacities, rivaling and sometimes exceeding those of HDDs.
D. SSDs are used mostly in cold storage.
This is incorrect. SSDs are primarily used for "hot" or "warm" storage tiers where high performance and low latency are critical. Cold storage is for infrequently accessed data where lower-cost, higher-capacity media like HDDs or tape are more suitable.

References:

1. Microsoft Azure Documentation. (Vendor Documentation). "Managed Disks overview - Azure Virtual Machines". Microsoft Docs. This document describes Azure's disk types, stating, "Premium SSDs, Premium SSDs v2, and Ultra Disks are high-performance, solid-state drive (SSD)-based storage... designed to support I/O-intensive workloads with significantly high throughput and low latency." This directly supports the use of SSDs for high-IOP applications

(Answer C) and contradicts their use for cold storage (Answer D).
2. Amazon Web Services (AWS) Documentation. (Vendor Documentation). "Amazon EBS volume types". AWS Documentation. In the section "Provisioned IOPS SSD volumes," it states, "Provisioned IOPS SSD volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads, that are sensitive to storage performance and consistency." This confirms SSDs are used for high-IOP applications (Answer C). The description of Cold HDD (sc1) volumes for "infrequently accessed data" refutes Answer D.
3. Arpaci-Dusseau, R. H., & Arpaci-Dusseau, A. C. (2018). Operating Systems: Three Easy Pieces (Version 1.00). Arpaci-Dusseau Books. In Chapter 41, "Flash-Based SSDs," Section 41.4, the text explains the performance characteristics: "Because there is no mechanical cost to a random access (as there is in a disk), an SSD can support a huge number of IOPS... This high IOPS rate is one of the main reasons SSDs are so desirable." This university-level textbook confirms the high-IOP capability of SSDs (Answer C).
4. TCG Storage Work Group. (Official Standard Documentation). "TCG Storage Architecture Core Specification, Version 2.01, Revision 1.00". Trusted Computing Group. This specification details the architecture for self-encrypting drives (SEDs), a technology widely implemented in modern SSDs, which contradicts the claim in Answer A.

Question: 11

A cloud engineer is in charge of deploying a platform in an IaaS public cloud. The application tracks the state using session cookies, and there are no affinity restrictions. Which of the following will help the engineer reduce monthly expenses and allow the application to provide the service?

A. Resource metering

B. Reserved resources

C. Dedicated host

D. Pay-as-you-go model

Answer:

D

Explanation:

The application's architecture, which uses session cookies for state and has no affinity restrictions, makes it stateless from the server's perspective. This design is ideal for horizontal scaling and elasticity, allowing infrastructure to be scaled up or down based on real-time demand. The pay-as-you-go model directly aligns with this capability by charging only for the resources consumed. This prevents over-provisioning for peak capacity and ensures that the organization does not pay for idle resources during periods of low traffic, thereby minimizing monthly expenses while maintaining service availability.

Why Incorrect Options are Wrong:

A. Resource metering: This is the process of measuring resource consumption. While it enables the pay-as-you-go model, it is not the cost-saving strategy itself.
B. Reserved resources: This model offers discounts for a long-term commitment (e.g., 1-3 years) and is best suited for stable, predictable workloads, not necessarily for leveraging elasticity to reduce costs.
C. Dedicated host: This provides a physical server for a single tenant. It is the most expensive option and is typically used for compliance or software licensing, directly contradicting the goal of reducing expenses.

References:

1. National Institute of Standards and Technology (NIST) Special Publication 800-145, "The NIST Definition of Cloud Computing":
Section 2, Page 2: Defines "On-demand self-service" and "Measured service" as essential characteristics of cloud computing. The pay-as-you-go model is a direct implementation of these

principles, allowing consumers to provision resources as needed and pay only for what they use. The stateless application in the scenario is perfectly suited to leverage this on-demand nature for cost efficiency.

2. Amazon Web Services (AWS) Documentation, "Amazon EC2 Pricing":

On-Demand Pricing Section: "With On-Demand instances, you pay for compute capacity by the hour or the second with no long-term commitments... This frees you from the costs and complexities of planning, purchasing, and maintaining hardware... It is recommended for applications with short-term, spiky, or unpredictable workloads that cannot be interrupted." This directly supports the use of a pay-as-you-go model for an application designed for elasticity to reduce costs.

Reserved Instances & Dedicated Hosts Sections: The documentation contrasts this with Reserved Instances, which are for "applications with steady state or predictable usage," and Dedicated Hosts, which are physical servers that "can help you reduce costs by allowing you to use your existing server-bound software licenses." These use cases do not align with the scenario's primary goal of cost reduction through elasticity.

3. Microsoft Azure Documentation, "Virtual Machines pricing":

Pay as you go Section: Describes this model as ideal for "running applications with short-term or unpredictable workloads where there is no long-term commitment." This aligns with the scenario where an engineer wants to leverage the cloud's elasticity to match cost to actual usage, thus reducing waste.

Reserved Virtual Machine Instances Section: Explains that reservations are for workloads with "predictable, consistent traffic" and require a "one-year or three-year term," which is less flexible than pay-as-you-go.

4. Armbrust, M., et al. (2009). "Above the Clouds: A Berkeley View of Cloud Computing." University of California, Berkeley, Technical Report No. UCB/EECS-2009-28.

Section 3.1, Economic Advantages: The paper states, "Cloud Computing enables a pay-as-you-go model, where you pay only for what you use... An attraction of Cloud Computing is that computing resources can be rapidly provisioned and de-provisioned on a fine-grained basis... allowing clouds to offer an 'infinite' pool of resources in a pay-as-you-go manner." This academic source establishes the fundamental economic benefit of the pay-as-you-go model in leveraging elasticity, which is the core of the question.

Question: 12

A systems administrator is provisioning VMs according to the following requirements: A VM instance needs to be present in at least two data centers. . During replication, the application hosted on the VM tolerates a maximum latency of one second. When a VM is unavailable, failover must be immediate. Which of the following replication methods will best meet these requirements?

A. Snapshot

B. Transactional

C. Live

D. Point-in-time

Answer:

C

Explanation:

The requirements for immediate failover and a maximum replication latency of one second necessitate a continuous, near-real-time data protection strategy. Live replication, often implemented as synchronous or near-synchronous replication, continuously transmits data changes from the primary VM to a replica in a secondary data center as they occur. This method ensures the replica is always in a consistent and up-to-date state, enabling an immediate and automated failover with a Recovery Point Objective (RPO) of near-zero. This directly meets the stringent availability and low-latency demands described in the scenario for mission-critical applications.

Why Incorrect Options are Wrong:

A. Snapshot: Snapshot replication is periodic, creating copies at discrete intervals. This method cannot meet the immediate failover or sub-second latency requirements due to inherent data loss (RPO) between snapshots.
B. Transactional: Transactional replication is a database-specific technology that replicates database transactions. It does not apply to the entire virtual machine state, including the operating system and application files.
D. Point-in-time: This is a general term for creating a copy of data as it existed at a specific moment, which includes snapshots. It is not a continuous process and cannot support immediate failover.

References:

1. VMware, Inc. (2023). vSphere Storage Documentation, Administering vSphere Virtual Machine Storage, Chapter 8: Virtual Machine Storage Policies. VMware. In the section "Site disaster tolerance," the documentation explains that synchronous replication provides the highest level of availability with a Recovery Point Objective (RPO) of zero, which is essential for immediate failover scenarios. This aligns with the concept of "live" replication.
2. Kyriazis, D., et al. (2013). Disaster Recovery for Infrastructure-as-a-Service Cloud Systems: A Survey. ACM Computing Surveys, 46(1), Article 10. In Section 3.2, "Replication Techniques," the paper contrasts synchronous and asynchronous replication. It states, "Synchronous replication... offers a zero RPO... suitable for mission-critical applications with low tolerance for data loss." This supports the choice of a live/synchronous method for immediate failover. https://doi.org/10.1145/2522968.2522978
3. Microsoft Corporation. (2023). Azure Site Recovery documentation, About Site Recovery. Microsoft Docs. The documentation describes "continuous replication" for disaster recovery of VMs, which provides minimal RPOs. While specific RPO values vary, the principle of continuous or "live" data transfer is fundamental to achieving the low latency and immediate failover required.

Question: 13

A company's content management system (CMS) service runs on an IaaS cluster on a public cloud. The CMS service is frequently targeted by a malicious threat actor using DDoS. Which of the following should a cloud engineer monitor to identify attacks?

A. Network flow logs

B. Endpoint detection and response logs

C. Cloud provider event logs

D. Instance syslog

Answer:

A

Explanation:

A Distributed Denial of Service (DDoS) attack is fundamentally a network-based attack designed to overwhelm a target with a massive volume of traffic from multiple sources. Network flow logs capture metadata about all IP traffic traversing a network interface, including source/destination IP addresses, ports, protocols, and the volume of packets/bytes. By monitoring and analyzing these logs, a cloud engineer can identify the characteristic signatures of a DDoS attack, such as an abnormally high volume of traffic from a large number of disparate IP addresses targeting the CMS service. This provides the necessary network-level visibility to detect the attack in its early stages.

Why Incorrect Options are Wrong:

B. Endpoint detection and response logs: EDR focuses on malicious activity on an endpoint (e.g., malware, unauthorized processes), not on analyzing incoming network traffic volume from distributed sources.
C. Cloud provider event logs: These logs (e.g., AWS CloudTrail, Azure Activity Log) track management plane API calls and user activity, not the data plane network traffic that constitutes a DDoS attack.
D. Instance syslog: This log contains operating system and application-level events from a single instance. It lacks the network-wide perspective needed to identify a distributed attack pattern.

References:

1. Amazon Web Services (AWS) Documentation. VPC Flow Logs. Amazon states that a key use case for VPC Flow Logs is "Monitoring the traffic that is reaching your instance... For example, you can use flow logs to help you diagnose overly restrictive security group rules." This same data is used to identify anomalous traffic patterns indicative of a DDoS attack. Retrieved from: https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs.html (See section: "Flow log

basics").
2. Microsoft Azure Documentation. Azure DDoS Protection overview. Microsoft explains that its protection service works by "monitoring actual traffic utilization and constantly comparing it against the thresholds... When the traffic threshold is exceeded, DDoS mitigation is initiated automatically." This monitoring is based on network flow telemetry, the same data captured in flow logs. Retrieved from: https://learn.microsoft.com/en-us/azure/ddos-protection/ddos-protection-overview (See section: "How DDoS Protection works").
3. Google Cloud Documentation. VPC Flow Logs overview. Google lists "Network monitoring" and "Network forensics" as primary use cases. For forensics, it states, "If an incident occurs, VPC Flow Logs can be used to determine... the traffic flow." This is essential for analyzing a DDoS incident. Retrieved from: https://cloud.google.com/vpc/docs/flow-logs (See section: "Use cases").
4. Carnegie Mellon University, Software Engineering Institute. Situational Awareness for Network Monitoring. In CERT/CC's guide to network monitoring, it emphasizes the importance of flow data (like NetFlow, the precursor to cloud flow logs) for "detecting and analyzing security events, such as denial-of-service (DoS) attacks." Retrieved from: https://resources.sei.cmu.edu/assetfiles/technicalnote/200400400114111.pdf (See Page 11, Section 3.2.2).

Question: 14

A cloud engineer needs to integrate a new payment processor with an existing e-commerce website. Which of the following technologies is the best fit for this integration?

A. RPC over SSL

B. Transactional SQL

C. REST API over HTTPS

D. Secure web socket

Answer:

C

Explanation:

A REST (Representational State Transfer) API (Application Programming Interface) is the industry-standard architectural style for integrating web services. For an e-commerce site to communicate with a payment processor, it needs a secure, scalable, and stateless method. REST APIs use standard HTTP methods (like POST for submitting payment data) and are designed for this type of client-server interaction. Encapsulating the communication within HTTPS (HTTP Secure) ensures that sensitive payment information is encrypted in transit, which is a critical security requirement for handling financial data. This combination provides a robust, secure, and widely supported solution for this integration task.

Why Incorrect Options are Wrong:

A. RPC over SSL: Remote Procedure Call (RPC) is an older paradigm that is often more tightly coupled and less flexible than REST for web-based integrations. While secure over SSL, it's not the modern standard.
B. Transactional SQL: This is incorrect. SQL is a language for querying databases. Directly exposing a database to an external payment processor via SQL would be a major security vulnerability and is not an integration protocol.
D. Secure web socket: Web sockets provide persistent, bidirectional communication channels, ideal for real-time applications like chat or live data feeds. This is unnecessary for a standard payment transaction, which is a simple request-response event.

References:

1. Fielding, R. T. (2000). Architectural Styles and the Design of Network-based Software Architectures. Doctoral dissertation, University of California, Irvine. In Chapter 5, "Representational State Transfer (REST)," Fielding defines the principles of REST, highlighting its advantages for hypermedia systems like the World Wide Web, including scalability, simplicity, and portability, which are essential for e-commerce integrations. (Available at:

https://www.ics.uci.edu/fielding/pubs/dissertation/restarchstyle.htm)
2. Amazon Web Services (AWS) Documentation. "What is a RESTful API?". AWS, a major cloud provider, defines RESTful APIs as the standard for web-based communication. The documentation states, "REST determines how the API looks like. It stands for "Representational State Transfer". It is a set of rules that developers follow when they create their API... Most applications on the internet use REST APIs to communicate." This confirms its status as the best fit for web service integration. (Reference: aws.amazon.com/what-is/restful-api/)
3. Microsoft Azure Documentation. "What are APIs?". The official documentation describes how APIs enable communication between applications, with REST being the predominant architectural style for web APIs. It emphasizes the use of HTTP/HTTPS protocols for these interactions, aligning perfectly with the scenario. (Reference: azure.microsoft.com/en-us/resources/cloud-computing-dictionary/what-are-apis)
4. Google Cloud Documentation. "API design guide". Google's guide for building APIs for its cloud platform is based on REST principles. It details the use of standard HTTP methods and resource-oriented design, which is the foundation for modern integrations like payment processors. (Reference: cloud.google.com/apis/design)

Question: 15

A company that has several branches worldwide needs to facilitate full access to a specific cloud resource to a branch in Spain. Other branches will have only read access. Which of the following is the best way to grant access to the branch in Spain?

A. Set up MFA for the users working at the branch.

B. Create a network security group with required permissions for users in Spain.

C. Apply a rule on the WAF to allow only users in Spain access to the resource.

D. Implement an IPS/IDS to detect unauthorized users.

Answer:

B

Explanation:

A network security group (NSG) or an equivalent cloud construct (e.g., AWS Security Group, GCP Firewall Rule) is the most appropriate tool for this scenario. NSGs act as a stateful virtual firewall at the network layer, controlling inbound and outbound traffic to resources. By creating a specific rule, an administrator can allow traffic from the known IP address range of the Spanish branch on the ports required for "full access." Concurrently, another rule with a lower priority can be set for all other source IPs, permitting access only on ports associated with "read-only" functions. This directly implements location-based access control as required.

Why Incorrect Options are Wrong:

A. Set up MFA for the users working at the branch.
MFA is an authentication control that verifies a user's identity. It does not define or enforce permissions (authorization) like full versus read-only access.
C. Apply a rule on the WAF to allow only users in Spain access to the resource.
A Web Application Firewall (WAF) primarily protects against application-layer attacks (e.g., SQL injection). While it can use IP-based rules, an NSG is the more fundamental and appropriate tool for network-level access control.
D. Implement an IPS/IDS to detect unauthorized users.
Intrusion Detection/Prevention Systems (IDS/IPS) are threat detection and mitigation tools. They monitor for malicious activity, not for defining and enforcing standard access control policies.

References:

1. Microsoft Azure Documentation. (2023). Network security groups. Microsoft Learn. In the "Security rules" section, it states, "A network security group contains security rules that allow or deny inbound network traffic... For each rule, you can specify source and destination, port, and protocol." This confirms the capability to create IP-based rules for specific access. Retrieved from

Microsoft's official documentation.
2. Amazon Web Services (AWS) Documentation. (2023). Control traffic to resources using security groups. AWS Documentation. The documentation specifies, "A security group acts as a virtual firewall for your instance to control inbound and outbound traffic... you add rules to each security group that allow traffic to or from its associated instances." This supports using security groups for IP-based traffic control. Retrieved from AWS's official documentation.
3. National Institute of Standards and Technology (NIST). (June 2017). NIST Special Publication 800-63B: Digital Identity Guidelines, Authentication and Lifecycle Management. Section 5.1.1, "Memorized Secrets," and subsequent sections on authenticators describe MFA as a mechanism to "authenticate the subscriber to the CSP," confirming its role in identity verification, not authorization. (DOI: https://doi.org/10.6028/NIST.SP.800-63b)
4. Chandrasekaran, K. (2015). Essentials of Cloud Computing. CRC Press, Taylor & Francis Group. Chapter 10, "Cloud Security," distinguishes between network-level firewalls (like NSGs) for controlling access based on network parameters and application-level firewalls (WAFs) for inspecting application data. This academic source clarifies the distinct roles of these technologies.

Question: 16

Which of the following network types allows the addition of new features through the use of network function virtualization?

A. Local area network

B. Wide area network

C. Storage area network

D. Software-defined network

Answer:

D

Explanation:

A Software-Defined Network (SDN) is an architecture that decouples the network control plane from the data forwarding plane, enabling the network to be programmatically controlled. This programmability is the key mechanism that allows for the dynamic addition of new features. Network Function Virtualization (NFV) complements SDN by virtualizing network functions (e.g., firewalls, routers, load balancers) so they can run as software on standard servers. An SDN architecture provides the ideal framework to manage, orchestrate, and chain these virtualized network functions, allowing new features to be deployed rapidly through software rather than by installing new physical hardware.

Why Incorrect Options are Wrong:

A. Local area network: This term defines a network by its limited geographical scope (e.g., an office), not by an architecture that inherently supports adding virtualized functions.
B. Wide area network: This term defines a network by its broad geographical scope (e.g., across cities), not by its design for programmatic control and feature addition.
C. Storage area network: This is a specialized network dedicated to providing block-level access to storage devices; it is not designed for general-purpose network services virtualized via NFV.

References:

1. European Telecommunications Standards Institute (ETSI). (2014). Network Functions Virtualisation (NFV); Architectural Framework (ETSI GS NFV 002 V1.2.1). Section 4.2, "Relationship between NFV and Software-Defined Networking (SDN)," explains that SDN and NFV are complementary, with SDN being a potential technology to control and route traffic between virtualized network functions.
2. Nunes, B. A. A., Mendonca, M., Nguyen, X. N., Obraczka, K., & Turletti, T. (2014). A survey of software-defined networking: Past, present, and future of programmable networks. IEEE Communications Surveys & Tutorials, 16(3), 1617-1634. Section IV.A, "Network Virtualization,"

discusses how SDN's abstraction enables the creation of virtual networks and the deployment of network functions. https://doi.org/10.1109/SURV.2014.012214.00001
3. Kreutz, D., Ramos, F. M., Verissimo, P. E., Rothenberg, C. E., Azodolmolky, S., & Uhlig, S. (2015). Software-defined networking: A comprehensive survey. Proceedings of the IEEE, 103(1), 14-76. Section V, "Use Cases and Opportunities," details how the SDN architecture facilitates the deployment of middleboxes and other network functions as software services. https://doi.org/10.1109/JPROC.2014.2371999

Question: 17

Which of the following migration types is best to use when migrating a highly available application, which is normally hosted on a local VM cluster, for usage with an external user population?

A. Cloud to on-premises

B. Cloud to cloud

C. On-premises to cloud

D. On-premises to on-premises

Answer:

C

Explanation:

The scenario describes an application currently hosted on a "local VM cluster," which is an on-premises environment. The goal is to migrate it to better serve an "external user population." Migrating from an on-premises data center to a public or hybrid cloud environment is the standard approach to achieve greater scalability, high availability, and global accessibility for external users. This process is defined as an on-premises-to-cloud migration, often referred to as Physical-to-Cloud (P2C) or Virtual-to-Cloud (V2C). The cloud's inherent internet-facing infrastructure and distributed nature make it the ideal target for this requirement.

Why Incorrect Options are Wrong:

A. Cloud to on-premises: This describes repatriation, moving an application from a cloud provider back to a local data center, which is the opposite of the described scenario.
B. Cloud to cloud: This involves migrating an application between two different cloud environments. The application in the question originates from an on-premises location, not a cloud.
D. On-premises to on-premises: This describes moving an application between two local data centers. This migration type does not inherently provide the global reach and scalability needed for external users.

References:

1. National Institute of Standards and Technology (NIST). (2011). NIST Cloud Computing Reference Architecture (NIST Special Publication 500-292).
Section 5.2, Cloud Migration, Page 23: The document defines cloud migration as "the process of moving an organization's data and applications from the organization's existing on-premise data center to the cloud infrastructure." This directly aligns with the scenario of moving from a local cluster to a platform suitable for external users.

2. Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., ... & Zaharia, M. (2009). Above the Clouds: A Berkeley View of Cloud Computing (Technical Report No. UCB/EECS-2009-28). University of California, Berkeley.
Section 3, Classes of Utility Computing: The paper discusses the economic and technical advantages of moving applications to the cloud, particularly for services that need to scale to serve a large, variable user base, which is characteristic of an "external user population." This supports the rationale for an on-premises-to-cloud migration.
3. Microsoft Azure Documentation. (2023). What is the Cloud Adoption Framework?
"Define strategy" and "Plan" sections: The framework outlines the motivations for moving to the cloud, including "reaching new customers" and "expanding to new geographies." It explicitly details the process of migrating workloads from on-premises environments to the Azure cloud to achieve these goals. This vendor documentation validates the on-premises-to-cloud path for serving external populations.

Question: 18

A company's engineering department is conducting a month-long test on the scalability of an in-house-developed software that requires a cluster of 100 or more servers. Which of the following models is the best to use?

A. PaaS

B. SaaS

C. DBaaS

D. IaaS

Answer:

D

Explanation:

Infrastructure as a Service (IaaS) is the most appropriate model as it provides fundamental computing resources, including virtual servers, networking, and storage. This gives the engineering department the maximum level of control needed to provision a large cluster of servers (100+), install custom operating systems and dependencies, and deploy their in-house software for a comprehensive scalability test. The on-demand, pay-as-you-go nature of IaaS is ideal for a temporary, month-long project, allowing the company to access massive computing power without the capital expense of purchasing physical hardware.

Why Incorrect Options are Wrong:

A. PaaS abstracts the underlying server infrastructure, which would limit the team's ability to control the environment and install the specific software stack required for their test.
B. SaaS provides ready-to-use software applications, not the underlying infrastructure needed to test a company's own custom-developed software.
C. DBaaS is a specialized service for managing databases. It does not provide the general-purpose server cluster needed to run the application itself.

References:

1. Mell, P., & Grance, T. (2011). The NIST Definition of Cloud Computing (NIST Special Publication 800-145). National Institute of Standards and Technology.
Page 2, Section "Infrastructure as a Service (IaaS)": "The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications." This directly supports the need to deploy in-house software on a large number of servers.
2. Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson,

D., Rabkin, A., Stoica, I., & Zaharia, M. (2009). Above the Clouds: A Berkeley View of Cloud Computing (Technical Report No. UCB/EECS-2009-28). University of California, Berkeley. Page 5, Section 3.1: Discusses how IaaS enables "pay-as-you-go" access to infrastructure, which is ideal for short-term, large-scale needs like the month-long test described, a use case often termed "batch processing" or "elastic computing."
3. Microsoft Azure Documentation. (n.d.). What is Infrastructure as a Service (IaaS)? Section "Common IaaS business scenarios": "Test and development. Teams can quickly set up and dismantle test and development environments, bringing new applications to market faster. IaaS makes it quick and economical to scale up dev-test environments up and down." This explicitly validates using IaaS for temporary, large-scale testing.

Question: 19

An organization wants to ensure its data is protected in the event of a natural disaster. To support this effort, the company has rented a colocation space in another part of the country. Which of the following disaster recovery practices can be used to best protect the data?

A. On-site

B. Replication

C. Retention

D. Off-site

Answer:

D

Explanation:

The core of the question is protecting data from a natural disaster by using a geographically separate facility. This practice is known as off-site disaster recovery. By renting a colocation space in another part of the country, the organization establishes a secondary location that is unlikely to be affected by the same disaster that impacts the primary site. This geographic separation is the fundamental principle of an off-site strategy, ensuring business continuity and data availability in the event of a regional catastrophe.

Why Incorrect Options are Wrong:

A. On-site: This practice involves keeping data backups or redundant systems at the same physical location as the primary data, offering no protection against a site-wide disaster like a fire or flood.

B. Replication: This is the process of copying data. While replication is a mechanism used to send data to an off-site location, "off-site" is the specific disaster recovery practice described in the scenario.

C. Retention: This refers to policies that dictate how long data is stored. Data retention is unrelated to the physical location of data for disaster recovery purposes.

References:

1. National Institute of Standards and Technology (NIST) Special Publication 800-34 Rev. 1, Contingency Planning Guide for Federal Information Systems. Section 4.3.2, "Alternate Storage Site," states: "An alternate storage site is used for storage of backup media... The site should be geographically separated from the primary site so as not to be susceptible to the same hazards." This directly supports the concept of using a geographically distant location (off-site) for disaster protection.

2. Amazon Web Services (AWS), Disaster Recovery of Workloads on AWS: Recovery in the

Cloud (July 2021). Page 6, in the section "Backup and Restore," discusses storing backups in a separate AWS Region. It states, "By replicating your data to another Region, you can protect your data in the unlikely event of a regional disruption." This exemplifies the off-site practice in a cloud context.

3. Microsoft Azure Documentation, Disaster recovery and high availability for Azure applications. In the section "Azure services that provide disaster recovery," it describes Azure Site Recovery, which "replicates workloads to a secondary location." The use of a secondary, geographically distinct location is the definition of an off-site strategy.

Question: 20

Which of the following do developers use to keep track of changes made during software development projects?

A. Code drifting

B. Code control

C. Code testing

D. Code versioning

Answer:

D

Explanation:

Code versioning, also known as version control or source control, is the standard practice and system used by developers to manage and track changes to source code and other project files over time. It creates a historical record of all modifications, enabling developers to revert to previous states, compare changes, and collaborate on a shared codebase without overwriting each other's work. Tools like Git, Subversion (SVN), and Mercurial are common implementations of code versioning.

Why Incorrect Options are Wrong:

A. Code drifting: This term, more commonly known as configuration drift, describes the phenomenon where infrastructure configurations diverge from their intended baseline, not the tracking of software code changes.
B. Code control: This is a generic and non-standard term. While versioning is a form of "controlling" code, "code versioning" is the precise, industry-accepted terminology for the practice in question.
C. Code testing: This is the process of evaluating software functionality to identify defects. It is a distinct phase in the development lifecycle and does not involve tracking historical changes to the code.

References:

1. CompTIA Cloud+ Certification Exam Objectives (CV0-004). (2023). CompTIA. Section 2.4, "Given a scenario, use appropriate tools to deploy cloud services," explicitly lists "Version control" as a key tool for deployment and automation.
2. Parr, T. (2012). The Definitive ANTLR 4 Reference. The Pragmatic Bookshelf. In the context of software development best practices, the text discusses the necessity of source control systems: "You should also be using a source code control system such as Perforce, Subversion, or Git to manage your project files." (Chapter 1, Section: Building ANTLR, p. 10). This highlights versioning

as the method for managing project files.
3. MIT OpenCourseWare. (2016). 6.005 Software Construction, Spring 2016. Massachusetts Institute of Technology. In "Reading 1: Static Checking," the course material introduces version control as a fundamental tool for managing software projects: "Version control is a system that keeps records of your changes."
4. AWS Documentation. (n.d.). What is Version Control? Amazon Web Services. Retrieved from https://aws.amazon.com/devops/version-control/. The official documentation defines the practice: "Version control, also known as source control, is the practice of tracking and managing changes to software code."

Question: 21

A cloud administrator needs to collect process-level, memory-usage tracking for the virtual machines that are part of an autoscaling group. Which of the following is the best way to accomplish the goal by using cloud-native monitoring services?

A. Configuring page file/swap metrics

B. Deploying the cloud-monitoring agent software

C. Scheduling a script to collect the data

D. Enabling memory monitoring in the VM configuration

Answer:

B

Explanation:

Cloud-native monitoring services (like AWS CloudWatch, Azure Monitor, or Google Cloud's operations suite) provide high-level metrics from the hypervisor by default, such as overall CPU utilization for a VM. However, to collect detailed in-guest operating system data, such as process-level memory usage, it is necessary to install a dedicated monitoring agent. This agent software runs inside the VM, collects the specified metrics directly from the OS, and sends them to the cloud monitoring service. For an autoscaling group, the agent is installed on the base machine image, ensuring all new instances automatically report these detailed metrics, making it the most effective and scalable solution.

Why Incorrect Options are Wrong:

A. Configuring page file/swap metrics: This only tracks the usage of virtual memory on disk, which is an indicator of memory pressure, not a direct measurement of memory usage by individual processes.

C. Scheduling a script to collect the data: This is a custom, non-native solution. While possible, it requires manual development and maintenance and is less integrated and reliable than using the purpose-built agent provided by the cloud platform.

D. Enabling memory monitoring in the VM configuration: This option typically enables hypervisor-level memory metrics, which report the total memory consumed by the VM as a whole, but lack the visibility to report on individual processes running inside the guest OS.

References:

1. Amazon Web Services (AWS) Documentation. The CloudWatch agent is required to collect guest OS-level metrics. "By default, EC2 instances send hypervisor-visible metrics to CloudWatch... To collect metrics from the operating system or from applications, you must install the CloudWatch agent."

Source: AWS Documentation, "The metrics that the CloudWatch agent collects," Section: "Predefined metric sets for the CloudWatch agent."

2. Microsoft Azure Documentation. The Azure Monitor agent is used to collect in-depth data from the guest operating system of virtual machines. "Use the Azure Monitor agent to collect guest operating system data from Azure... virtual machines... It collects data from the guest operating system and delivers it to Azure Monitor."

Source: Microsoft Learn, "Azure Monitor agent overview," Introduction section.

3. Google Cloud Documentation. The Ops Agent is Google's solution for collecting detailed telemetry from within Compute Engine instances. "The Ops Agent is the primary agent for collecting telemetry from your Compute Engine instances. It collects both logs and metrics." The agent can be configured to collect process metrics.

Source: Google Cloud Documentation, "Ops Agent overview," What the Ops Agent collects section.

4. Armbrust, M., et al. (2010). A View of Cloud Computing. This foundational academic paper from UC Berkeley discusses the challenges of cloud monitoring, implying the need for mechanisms beyond the hypervisor to understand application-level performance. The distinction between what the infrastructure provider can see (hypervisor-level) and what the user needs to see (in-guest) necessitates agent-based approaches for detailed monitoring.

Source: Communications of the ACM, 53(4), 50-58. Section 5.3, "Monitoring and Auditing." DOI: https://doi.org/10.1145/1721654.1721672

Question: 22

Users report being unable to access an application that uses TLS 1.1. The users are able to access other applications on the internet. Which of the following is the most likely reason for this issue?

A. The security team modified user permissions.

B. Changes were made on the web server to address vulnerabilities.

C. Privileged access was implemented.

D. The firewall was modified.

Answer:

B

Explanation:

Transport Layer Security (TLS) versions 1.0 and 1.1 are deprecated due to significant, well-documented security vulnerabilities. A common and highly recommended security practice is to harden web servers by disabling these older protocols and forcing the use of modern, secure versions like TLS 1.2 or 1.3. If a server administrator implements this change to address vulnerabilities, any client or user application that is not configured to use a newer TLS version will be unable to establish a secure connection, resulting in an access failure. Since other applications are accessible, the issue is isolated to this specific server, making a server-side configuration change the most probable cause.

Why Incorrect Options are Wrong:

A. The security team modified user permissions.
This would typically result in an "Access Denied" or "403 Forbidden" error after a successful connection, not a connection failure related to the TLS protocol version.
C. Privileged access was implemented.
Privileged Access Management (PAM) controls administrative accounts and elevated permissions; it does not govern standard user access to a web application.
D. The firewall was modified.
While a firewall can block traffic, rules are typically based on IP addresses and ports, not the specific TLS version. A server-side protocol configuration is a more direct cause.

References:

1. National Institute of Standards and Technology (NIST). (2019). Special Publication (SP) 800-52r2, Guidelines for the Selection, Configuration, and Use of Transport Layer Security (TLS) Implementations.
Section 3.1, Protocol Versions, Page 6: "Servers that support government-only applications shall

be configured to use TLS 1.3 and should be configured to use TLS 1.2. These servers shall not be configured to use TLS 1.1 and shall not be configured to use TLS 1.0, SSL 3.0, or SSL 2.0." This document mandates the disabling of TLS 1.1 on servers to enhance security.

2. Internet Engineering Task Force (IETF). (2021). RFC 8996: Deprecating TLS 1.0 and TLS 1.1. Abstract: "This document formally deprecates Transport Layer Security (TLS) versions 1.0 (RFC 2246) and 1.1 (RFC 4346)... These versions lack support for current and recommended cryptographic algorithms and mechanisms, and various government and industry profiles now mandate avoiding these old TLS versions." This RFC provides the official rationale for discontinuing TLS 1.1 due to its vulnerabilities.

3. Microsoft Corporation. (2023). Solving the TLS 1.0 Problem, 2nd Edition. Security documentation.

Section: Disabling TLS 1.0 and 1.1: The document details the security risks of older TLS versions and provides technical guidance for administrators to disable them across their infrastructure to mitigate vulnerabilities, which directly aligns with the scenario in the question.

Question: 23

A video surveillance system records road incidents and stores the videos locally before uploading them to the cloud and deleting them from local storage. Which of the following best describes the nature of the local storage?

A. Persistent

B. Ephemeral

C. Differential

D. Incremental

Answer:

B

Explanation:

The local storage in this scenario functions as a temporary buffer. Its purpose is to hold the video files only until they are successfully uploaded to their permanent location in the cloud. After the transfer is complete, the local copies are deleted. This transient, non-permanent, and short-lived nature of the data on the local device is the defining characteristic of ephemeral storage. The storage is used for a temporary purpose, and the data is not intended to persist locally.

Why Incorrect Options are Wrong:

A. Persistent: This is incorrect because the data is intentionally deleted from the local storage after being moved. Persistent storage is designed for long-term data retention.
C. Differential: This is a backup methodology that captures changes made since the last full backup; it is not a type of storage.
D. Incremental: This is a backup methodology that captures changes made since the last backup of any type; it is not a type of storage.

References:

1. Amazon Web Services (AWS) Documentation. "Amazon EC2 Instance Store." In Amazon EC2 User Guide for Linux Instances. "An instance store provides temporary block-level storage for your instance... Instance store is ideal for temporary storage of information that changes frequently, such as buffers, caches, scratch data, and other temporary content..." This aligns with the scenario where local storage acts as a temporary buffer.
2. Google Cloud Documentation. "Local SSDs overview." In Compute Engine Documentation. "The data that you store on a local SSD persists only until the instance is stopped or deleted. For this reason, local SSDs are only suitable for temporary storage such as cache, processing space, or low value data." This source defines the temporary nature of ephemeral storage.

3. Armbrust, M., et al. (2009). Above the Clouds: A Berkeley View of Cloud Computing. University of California, Berkeley, EECS Department, Technical Report No. UCB/EECS-2009-28. Section 3.2, "Storage," distinguishes between persistent storage services (e.g., Amazon S3) and temporary storage that is tied to the lifecycle of a compute instance, highlighting the concept of non-persistent, or ephemeral, data.
4. Microsoft Azure Documentation. "Temporary disk on Azure VMs." In Azure Virtual Machines Documentation. "The temporary disk provides temporary storage for applications and processes and is intended to only store data such as page or swap files... Data on the temporary disk may be lost during a maintenance event..." This further exemplifies the non-permanent nature of ephemeral storage in a cloud context.

Question: 24

A cloud engineer hardened the WAF for a company that operates exclusively in North Americ a. The engineer did not make changes to any ports, and all protected applications have continued to function as expected. Which of the following configuration changes did the engineer most likely apply?

A. The engineer implemented MFA to access the WAF configurations.

B. The engineer blocked all traffic originating outside the region.

C. The engineer installed the latest security patches on the WAF.

D. The engineer completed an upgrade from TLS version 1.1 to version 1.3.

Answer:

B

Explanation:

A Web Application Firewall (WAF) is designed to protect web applications by filtering and monitoring HTTP traffic. A common and effective hardening technique is to reduce the attack surface by blocking traffic from geographic regions where the company does not operate. Since the company operates exclusively in North America, configuring the WAF to block all traffic originating from outside this region is a logical security enhancement. This change, known as geoblocking or geo-fencing, does not involve altering ports and would not impact the functionality of the applications for their intended user base, fitting the scenario perfectly.

Why Incorrect Options are Wrong:

A. The engineer implemented MFA to access the WAF configurations.
This hardens the WAF's management plane, not the traffic flow to the protected applications, which is the primary function described.
C. The engineer installed the latest security patches on the WAF.
Patching is a critical maintenance activity for hardening, but it is not typically described as a configuration change in the context of traffic filtering rules.
D. The engineer completed an upgrade from TLS version 1.1 to version 1.3.
This is a valid hardening configuration, but it does not utilize the key piece of information provided in the scenario-that the company operates exclusively in North America.

References:

1. AWS WAF Developer Guide. (Vendor Documentation). AWS documentation explicitly describes using a "Geographic match rule statement" to inspect and control web requests based on their country of origin. This directly supports the concept of geoblocking as a WAF configuration.
Reference: AWS WAF Developer Guide, "Rule statement list," Section: "Geographic match rule statement."
2. Microsoft Azure Documentation. (Vendor Documentation). Azure's documentation for its WAF details the creation of custom rules, which can use "Geographical location" as a match condition to allow or block traffic based on the client's IP address origin.
Reference: Microsoft Docs, "Custom rules for Web Application Firewall v2 on Azure Application Gateway," Section: "Match variables."
3. NIST Special Publication 800-53 Revision 5. (Peer-Reviewed Academic Publication/Standard). This publication outlines security and privacy controls. Control AC-4, "Information Flow Enforcement," and its enhancement AC-4(17) "Geolocation" specify the enforcement of information flow control based on the geolocation of the source, validating this as a standard security practice.
Reference: NIST SP 800-53 Rev. 5, Security and Privacy Controls for Information Systems and Organizations, Page 101, Control: AC-4(17).
4. Cloudflare Learning Center. (Vendor Documentation). Cloudflare, a major provider of WAF services, explains IP Access Rules, which can be used to block traffic from specific countries. This is presented as a primary method for securing applications from regional threats.
Reference: Cloudflare Learning Center, "What is a WAF?", Section: "How does a WAF work?". The article discusses WAF policies, including those based on geolocation.

Question: 25

A cloud solution needs to be replaced without interruptions. The replacement process can be completed in phases, but the cost should be kept as low as possible. Which of the following is the best strategy to implement?

A. Blue-green

B. Rolling

C. In-place

D. Canary

Answer:

B

Explanation:

A rolling deployment strategy is the most suitable choice as it aligns with all the specified constraints. This method updates the cloud solution in phases by sequentially replacing old instances with new ones. This incremental process ensures the service remains available, thus meeting the "no interruptions" requirement. Crucially, it reuses the existing infrastructure, updating it piece by piece rather than duplicating the entire environment. This makes it a highly cost-effective approach, directly addressing the need to keep costs as low as possible.

Why Incorrect Options are Wrong:

A. Blue-green: This strategy is not low-cost because it requires running two identical, parallel production environments simultaneously, which doubles the infrastructure expense during the deployment process.
C. In-place: This method, also known as a recreate deployment, involves stopping the application, deploying the new version, and restarting, which inherently causes service interruptions.
D. Canary: While a phased approach, a canary release is primarily for risk mitigation by testing new code on a small subset of users and can add complexity and overhead compared to a straightforward rolling update.

References:

1. Google Cloud Documentation, "Application deployment and testing strategies." This document describes a rolling update as a strategy where you "slowly replace instances of the previous version of your application with instances of the new version... a rolling update avoids downtime." It contrasts this with blue-green, which has a higher "monetary cost" due to resource duplication. (See section: "Rolling update deployment strategy").
2. Amazon Web Services (AWS) Whitepaper, "Blue/Green Deployments on AWS," PDF, Page 4. The paper states, "A potential downside to this blue-green approach is that you will have double

the resources running in production... This will result in a higher bill for the duration of the upgrade." This confirms the high-cost nature of blue-green deployments.
3. Red Hat OpenShift Container Platform 4.6 Documentation, "Understanding deployment strategies." The documentation explains that the "Rolling" strategy (the default in OpenShift/Kubernetes) "waits for new pods to become ready... before scaling down the old components. If a significant issue occurs, the rolling deployment can be aborted." This highlights its zero-downtime and phased nature without requiring duplicate infrastructure. (See section: "Rolling Strategy").

Question: 26

An e-commerce store is preparing for an annual holiday sale. Previously, this sale has increased the number of transactions between two and ten times the normal level of transactions. A cloud administrator wants to implement a process to scale the web server seamlessly. The goal is to automate changes only when necessary and with minimal cost. Which of the following scaling approaches should the administrator use?

A. Scale horizontally with additional web servers to provide redundancy.

B. Allow the load to trigger adjustments to the resources.

C. When traffic increases, adjust the resources using the cloud portal.

D. Schedule the environment to scale resources before the sale begins.

Answer:

B

Explanation:

The most appropriate approach is to allow the load to trigger resource adjustments, a concept known as autoscaling or elasticity. This method directly addresses all the requirements: it is automated, ensuring seamless scaling without manual intervention. It scales "only when necessary" by reacting to real-time metrics like CPU utilization or transaction volume, which is ideal for the unpredictable 2x to 10x traffic increase. This on-demand provisioning and de-provisioning of resources ensures minimal cost, as the e-commerce store only pays for the capacity it actually uses during the sales peak.

Why Incorrect Options are Wrong:

A. This describes the method of scaling (horizontal) but not the automated process for triggering it, which is the core of the question's requirements for seamless and cost-effective management.
C. Adjusting resources via the cloud portal is a manual process. This contradicts the requirements for automation and seamless operation, as it would require constant monitoring and intervention.
D. Scheduled scaling is not optimal for a variable load. It risks either over-provisioning resources (increasing costs) if the sale is less popular than expected or under-provisioning (causing outages) if it is more popular.

References:

1. National Institute of Standards and Technology (NIST) Special Publication 800-145, The NIST Definition of Cloud Computing.
Reference: Page 2, Section 2, "Essential Characteristics." The document defines "Rapid elasticity" as a key characteristic where "Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand."

This directly supports the principle of load-triggered adjustments.

2. Amazon Web Services (AWS) Documentation, "What is AWS Auto Scaling?".

Reference: AWS Auto Scaling User Guide. The documentation states, "AWS Auto Scaling monitors your applications and automatically adjusts capacity to maintain steady, predictable performance at the lowest possible cost." It describes dynamic scaling policies that respond to changing demand, which aligns with allowing the load to trigger adjustments.

3. Microsoft Azure Documentation, "Overview of autoscale in Microsoft Azure".

Reference: Azure Monitor documentation. It explains, "Autoscale allows you to have the right amount of resources running to handle the load on your app. It allows you to add resources to handle increases in load (scale out) and also save money by removing resources that are sitting idle (scale in)." This confirms that load-based triggers are the standard for cost-effective, automated scaling.

4. Erl, T., Mahmood, Z., & Puttini, R. (2013). Cloud Computing: Concepts, Technology & Architecture. Prentice Hall.

Reference: Chapter 5, Section 5.3, "Cloud Characteristics." The text describes the "Elastic Resource Capacity" characteristic, which is enabled by an "Automated Scaling Listener" mechanism that monitors requests and triggers the automatic allocation of IT resources in response to load fluctuations. This academic source validates option B as the correct architectural approach.

Question: 27

An organization's critical data was exfiltrated from a computer system in a cyberattack. A cloud analyst wants to identify the root cause and is reviewing the following security logs of a software web application:

"2021/12/18 09:33:12" "10. 34. 32.18" "104. 224. 123. 119" "POST /

login.php?u=administrator&p=or%201%20=1"

"2021/12/18 09:33:13" "10.34. 32.18" "104. 224. 123.119" "POST /login.

php?u=administrator&p=%27%0A"

"2021/12/18 09:33:14" "10. 34. 32.18" "104. 224. 123. 119" "POST /login.

php?u=administrator&p=%26"

"2021/12/18 09:33:17" "10.34. 32.18" "104. 224. 123.119" "POST /

login.php?u=administrator&p=%3B"

"2021/12/18 09:33:12" "10.34. 32. 18" "104. 224. 123. 119" "POST / login.

php?u=admin&p=or%201%20=1"

"2021/12/18 09:33:19" "10.34.32.18" "104. 224. 123.119" "POST / login.
php?u=admin&p=%27%0A"

"2021/12/18 09:33:21" "10. 34. 32.18" "104.224. 123.119" "POST / login. php?u=admin&p=%26"

"2021/12/18 09:33:23" "10. 34. 32.18" "104. 224. 123.119" "POST / login. php?u=admin&p=%3B"

Which of the following types of attacks occurred?

A. SQL injection

B. Cross-site scripting

C. Reuse of leaked credentials

D. Privilege escalation

Answer:

A

Explanation:

The provided security logs show clear evidence of a SQL injection (SQLi) attack. The attacker is sending multiple POST requests to login.php, attempting to manipulate the backend database query through the password parameter (p=). The payload or%201%20=1, which decodes to or 1=1, is a classic example of a tautology-based SQLi attack. This technique aims to bypass authentication by appending a universally true condition to the SQL WHERE clause. The use of other URL-encoded characters like the single quote (%27) and semicolon (%3B) are also common methods for terminating strings and stacking queries in an SQLi attack.

Why Incorrect Options are Wrong:

B. Cross-site scripting: This is incorrect because the logs show SQL syntax injection, not the injection of client-side scripts (e.g., or img tags) into a web page.
C. Reuse of leaked credentials: This is incorrect as the attacker is not using a valid, previously compromised password but is instead attempting to bypass the login mechanism with malformed input.
D. Privilege escalation: This describes a potential outcome or goal of an attack, not the attack method itself. The specific technique evidenced in the logs is SQL injection./body

References:

1. OWASP Foundation. (2021). OWASP Top 10:2021, A03:2021-Injection. OWASP. Retrieved from https://owasp.org/Top10/A032021-Injection/. The "Attack Scenarios" section describes how an attacker can use SQL injection, such as ' OR '1'='1, to bypass authentication.
2. Amazon Web Services (AWS). (2023). SQL injection attack rule statement. AWS WAF, AWS Firewall Manager, and AWS Shield Advanced Developer Guide. Retrieved from https://docs.aws.amazon.com/waf/latest/developerguide/aws-managed-rule-groups-use-case-sql-db.html. This official vendor documentation details how WAFs detect SQLi by looking for patterns like "tautologies such as 1=1 and 0=0."
3. Kar, D., Pan, T. S., & Das, R. (2021). SQLi-IDS: A real-time SQL injection detection system using a hybrid deep neural network. Computers & Security, 108, 102341. https://doi.org/10.1016/j.cose.2021.102341. Section 2.1, "Tautology-based SQLIA," explicitly discusses the use of OR 1=1 as a primary technique for bypassing user authentication.
4. Zelle, D., & Kamin, S. (2019). Web Application Security. University of Illinois at Urbana-Champaign, CS 461/ECE 422 Course Notes. Retrieved from https://courses.engr.illinois.edu/cs461/sp2019/slides/Lecture20-WebAppSecurity.pdf. Slide 22 provides a canonical example of a tautology-based SQL injection attack using ' OR 1=1 -- to bypass a login form.

Question: 28

A company wants to create a few additional VDIs so support vendors and contractors have a secure method to access the company's cloud environment. When a cloud administrator attempts to create the additional instances in the new locations, the operation is successful in some locations but fails in others. Which of the following is the most likely reason for this failure?

A. Partial service outages

B. Regional service availability

C. Service quotas

D. Deprecation of functionality

Answer:

C

Explanation:

Cloud providers impose service quotas or limits on the number of resources an account can provision within a specific region. When a cloud administrator attempts to create new resources, such as VDI instances, the request will fail if the account has already reached its predefined limit for that resource type (e.g., vCPUs, virtual machines) in that particular region. Since quotas are typically managed on a per-region basis, this explains why the creation is successful in some locations (where the quota has not been met) but fails in others (where the quota has been exceeded). This is a common operational constraint in cloud environments.

Why Incorrect Options are Wrong:

A. Partial service outages: A service outage would likely affect both new and existing services and is typically a temporary, unscheduled event, not a consistent barrier to creating new resources.
B. Regional service availability: This would mean the VDI service is entirely unavailable in certain regions, preventing the creation of any instances, not just failing after some have been deployed.
D. Deprecation of functionality: Deprecation is the planned retirement of a service or feature. This would typically result in failures across all regions, not a location-specific issue.

References:

1. Amazon Web Services (AWS) Documentation: "Service Quotas." AWS states, "Quotas, also referred to as limits in AWS, are the maximum number of resources that you can create in an AWS account... Many quotas are specific to an AWS Region." This confirms that resource limits are a regional constraint.
Source: AWS Documentation, "What Is Service Quotas?",
https://docs.aws.amazon.com/servicequotas/latest/userguide/intro.html
2. Microsoft Azure Documentation: "Azure subscription and service limits, quotas, and

constraints." The documentation details how quotas are applied per subscription and per region. For example, under "Virtual machine vCPU quotas," it states, "vCPU quotas are arranged in two tiers for each subscription, in each region."
Source: Microsoft Azure Documentation, "vCPU quotas," https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/azure-subscription-service-limits#vcpu-quotas
3. Google Cloud Documentation: "Working with quotas." The documentation specifies the scope of quotas: "Quotas are enforced on a per-project, per-region, or per-zone basis." This directly supports the concept of location-specific resource creation failures due to limits.
Source: Google Cloud Documentation, "About quotas," https://cloud.google.com/docs/quota#aboutquotas
4. Armbrust, M., et al. (2010). A view of cloud computing. Communications of the ACM, 53(4), 50-58. This foundational academic paper on cloud computing discusses elasticity as a key feature but also notes the practical limitations imposed by providers to manage resources, which manifest as quotas. The paper implicitly supports the idea that resource provisioning is not infinite and is subject to provider-imposed controls.
DOI: https://doi.org/10.1145/1721654.1721672 (Section 3.1, "Elasticity and the Illusion of Infinite Resources")

Question: 29

Which of the following is used to deliver code quickly and efficiently across the development, test, and production environments?

A. Snapshot

B. Container image

C. Serverless function

D. VM template

Answer:

B

Explanation:

A container image is a lightweight, standalone, executable package that includes everything needed to run a piece of software: the code, a runtime, system tools, libraries, and settings. This packaging ensures that the application runs consistently and reliably when moved from one computing environment to another, such as from a developer's laptop to a test environment, and then into production. This portability and consistency make container images the ideal mechanism for delivering code quickly and efficiently across the software development lifecycle.

Why Incorrect Options are Wrong:

A. Snapshot: A snapshot is a point-in-time copy of a virtual machine or storage volume, primarily used for backup and recovery, not for deploying application code.
C. Serverless function: A serverless function is a piece of code that runs in a managed environment. While it is a method of deploying code, the container image is the packaging and delivery mechanism that ensures consistency across environments.
D. VM template: A VM template is a master copy of a virtual machine, including the full operating system. It is heavyweight and much slower to deploy than a container, making it inefficient for rapid code delivery.

References:

1. National Institute of Standards and Technology (NIST). (2017). NIST Special Publication 800-190: Application Container Security Guide.
Section 2.1, "What are Application Containers?": "An application container is a portable image that can be used to create one or more instances of a container. The image includes an application, its libraries, and its dependencies... This allows the application to be abstracted from the host operating system, providing portability and consistency across different environments." (Page 7). This directly supports the use of container images for consistency across environments.
2. Armbrust, M., et al. (2009). Above the Clouds: A Berkeley View of Cloud Computing. University

of California, Berkeley.
Section 3.1, "Virtual Machines": This paper discusses Virtual Machine Images (templates) as a way to bundle a full software stack. However, it highlights their size and startup time, contrasting with more modern, lightweight approaches. The principles laid out show why heavier VM templates are less efficient for rapid deployment compared to containers. (Page 4).
3. AWS Documentation. (n.d.). What is a Container?. Amazon Web Services.
"A container is a standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another. A Docker container image is a lightweight, standalone, executable package of software that includes everything needed to run an application: code, runtime, system tools, system libraries and settings." This official vendor documentation reinforces the role of container images in ensuring application portability and rapid deployment.

Question: 30

A cloud engineer is collecting web server application logs to troubleshoot intermittent issues. However, the logs are piling up and causing storage issues. Which of the following log mechanisms should the cloud engineer implement to address this issue?

A. Splicing

B. Rotation

C. Sampling

D. Inspection

Answer:

B

Explanation:

Log rotation is an automated administrative process that manages log files to prevent them from consuming excessive storage space. It works by creating new log files on a schedule (e.g., daily, weekly) or when a file reaches a certain size. The old log files are typically compressed, archived to cheaper storage, or deleted after a specified retention period. This directly solves the problem of logs "piling up" and causing storage issues while preserving recent logs for troubleshooting.

Why Incorrect Options are Wrong:

A. Splicing: Splicing involves joining or connecting things. In the context of files, this would mean combining logs, which would create even larger files and worsen the storage problem.
C. Sampling: Log sampling involves collecting only a subset of log events. This is unsuitable for troubleshooting intermittent issues, as the specific events needed for diagnosis might not be captured.
D. Inspection: Log inspection is the process of analyzing or reviewing log data to identify issues. It is the action the engineer is performing, not a mechanism to manage log file storage.

References:

1. National Institute of Standards and Technology (NIST). (2006). Guide to Computer Security Log Management (Special Publication 800-92).
Section 3.2.3, "Log Rotation and Archiving," Page 3-5: "Log rotation is the practice of closing a log file and opening a new one on a scheduled basis... Log rotation is performed primarily to keep log files from becoming too large. Once a log file is rotated, it is often compressed to save storage space." This document explicitly defines log rotation as the solution for managing large log files.
2. Red Hat. (2023). Red Hat Enterprise Linux 8: Configuring basic system settings.
Chapter 21, "Managing log files with logrotate," Section 21.1: "The logrotate utility allows the

automatic rotation, compression, removal, and mailing of log files. Each log file can be handled daily, weekly, monthly, or when it grows too large." This official vendor documentation describes the exact mechanism and its purpose, which aligns with the scenario.

3. AWS Documentation. (2024). Amazon CloudWatch Logs User Guide.

Section: "Working with log groups and log streams - Log retention": "By default, logs are kept indefinitely and never expire. You can adjust the retention policy for each log group, keeping the indefinite retention, or choosing a retention period... CloudWatch Logs automatically deletes log events that are older than the retention setting." This describes the cloud-native equivalent of log rotation for managing log storage.