DUMPSBOSS.

CompTIA Cloud+

CompTIA CV0-004

Version Demo

Total Demo Questions: 38

Total Premium Questions: 381

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Architecture and Design	132
Topic 2, Security	97
Topic 3, Deployment	49
Topic 4, Operations and Support	43
Topic 5, Troubleshooting	57
Topic 6, Mix Questions	3
Total	381

QUESTION NO: 1

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:

SSDs can be used for high-IOP applications. An SSD stores data in flash memory rather than accessing data through mechanically moving disk platters and read/write heads. This design provides substantially lower access latency and allows the device to service a large number of input/output operations per second (IOPS), particularly for random reads and writes. High-IOP performance is valuable in cloud environments that support transactional databases, virtual-machine datastores, virtual desktop infrastructure, busy application servers, and storage volumes handling many small concurrent requests. Selecting SSD-backed storage for these workloads helps reduce I/O wait time and can improve application responsiveness when storage operations are a limiting factor. Cloud+ candidates should associate SSDs with performance-sensitive tiers and workloads that need fast, frequent access rather than simply considering capacity. Actual IOPS depends on the SSD interface, controller, NAND type, queue depth, workload pattern, and any cloud-provider volume limits, so the storage design should still be sized and tested against the workload's required latency and throughput. The Storage Networking Industry Association describes solid-state storage and its flash-based operation at SNIA: What Is an SSD?. AWS also documents the performance characteristics and IOPS configuration of SSD-backed EBS volumes at Amazon EBS volume types.

QUESTION NO: 2

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:

IaaS is the best model because it provides on-demand access to scalable compute, storage, networking, and virtual-machine infrastructure that the engineering department can configure for its own application test. A month-long scalability test needing a cluster of at least 100 servers is a temporary, resource-intensive workload; using IaaS avoids the capital expense, lead time, and operational commitment of procuring and operating a comparably sized on-premises environment. The team retains control over the guest operating systems, network configuration, server sizing, images, test tooling, and deployment topology, allowing it to model production conditions and vary capacity as it measures how the in-house software behaves under increasing load. IaaS also supports rapid provisioning and deprovisioning, so the organization can consume the large cluster for the test period and release those resources when testing concludes. This aligns with NIST's definition of Infrastructure as a Service, in which the consumer deploys and runs arbitrary software on provider-supplied processing, storage, and networking resources while controlling operating systems and deployed applications. See NIST SP 800-145 and the CompTIA Cloud+ certification overview.

QUESTION NO: 3 - (SIMULATION)

An e-commerce company is migrating from an on-premises private cloud environment to

a public cloud IaaS environment. You are tasked with right-sizing the environment to

save costs after the migration. The company's requirements are to provide a 20% overhead above the average resource consumption, rounded up.

INSTRUCTIONS

Review the specifications and graphs showing resource usage for the web and database servers. Determine the average resource usage and select the correct specifications from the available drop-down options.

ANSWER: See the explanation for the answer

Explanation:

Select the following right-sized IaaS specifications. These values include a 20% buffer over the average observed utilization and are rounded up to the available sizing option.

Web server
CPU: `2 vCPUs`
RAM: `2 GB`
Disk speed: `10 MBps`
Database server
CPU: `6 vCPUs`
RAM: `128 GB`
Disk speed: `110 MBps`

For each graph value, calculate `average usage × 1.20`, then choose the next higher available dropdown value rather than rounding downward.

QUESTION NO: 4

A developer is building a new application version using a CI/CD pipeline. The developer receives the following error message log when the build fails:

Which of the following is the most likely cause of this failure?

A. Incorrect version

B. Test case failure

C. Broken build pipeline

D. Dependency issue

ANSWER: D

Explanation:

Dependency issue is the most likely cause because the stated build error identifies that the `requests` module cannot be found. In a CI/CD build, each runner or build container should be treated as a clean, reproducible environment. A Python import can succeed on a developer workstation while failing in the pipeline if the required package was not installed in the build image, is absent from the dependency manifest, or was installed into a different virtual environment than the one executing the build or tests.

The appropriate remediation is to declare the required library and its compatible version in the project's dependency-management file, such as `requirements.txt` or `pyproject.toml`, then ensure the pipeline explicitly installs dependencies before it performs application builds, tests, packaging, or deployment. Pinning and validating dependencies also improves repeatability across pipeline runs and environments. This aligns with Cloud+ operational practices for

automation, application lifecycle management, and reliable cloud deployments. Python's packaging guidance explains how project dependencies are declared and installed: Python Packaging User Guide. CompTIA's Cloud+ certification overview includes cloud operations and automation as relevant knowledge areas: CompTIA Cloud+.

QUESTION NO: 5

A company's cybersecurity team receives the following alert that a production VM was deleted from the virtual network:

21 September 09:19:08 (GMT-5)

Resource with ID: PROD-WEB001 was deleted by User: Logging Service

The console to manage virtual network resources uses directory authentication. Only users in a particular directory group can interactively access the virtual network management console. The logging service account is not part of this group and requires some local administration privileges to aggregate logs from various resources. The cybersecurity team discovers the logging service account was previously given full directory administrator privileges and sees the following entry:

21 September 09:10:55 (GMT-5)

User with ID: Logging Service was added to the Group: VNet Console Administrators by actor: Logging Service.

The cybersecurity team removes the compromised service account from the directory group. Which of the following should the cybersecurity team do next to prevent repeat instances of this issue?

A. Enable two-factor authentication on the virtual network console.

B. Reset the logging service account to use a long and complex password.

C. Disable RDP on the production virtual machines.

D. Create a scoped administrative role for the logging service account.

ANSWER: D

Explanation:

Create a scoped administrative role for the logging service account. This directly addresses the root cause: the service account held excessive directory-administrator permissions, which allowed it to modify group membership and grant itself interactive access to the virtual-network management console. A compromised workload identity must be limited to only the permissions needed for its intended operational task. In this case, log aggregation may require narrowly defined access to read logs, query monitoring endpoints, or perform specified local collection actions, but it does not require authority to manage directory groups or virtual-network resources.

A scoped role applies the principle of least privilege by assigning permissions to a defined resource scope and limiting permitted actions. The resulting account can continue collecting logs while lacking the ability to add itself to privileged groups, delete virtual machines, or administer unrelated cloud resources. The team should define the role based on the logging service's documented access requirements, assign it to the service identity, and regularly review its permissions and activity. This approach aligns with NIST guidance that accounts should receive only the privileges necessary to perform assigned tasks and that privileged functions should be restricted. See NIST SP 800-53 Rev. 5, AC-6 Least Privilege and NIST SP 800-207 Zero Trust Architecture.

QUESTION NO: 6

A cloud engineer is reviewing an application outage. The monitoring platform shows that response times increased immediately after a database schema change was deployed. Which of the following should the engineer review first to determine the cause of the performance issue?

A. Review the database query performance.

B. Increase the VM memory allocation.

C. Restart the application servers.

D. Modify the DNS record TTL.

ANSWER: A

Explanation:

Reviewing the database query performance is the best first action because the increased response times began immediately after a database schema change. A schema modification can affect indexes, query execution plans, table locks, data types, and relationships used by the application. Reviewing query duration, execution plans, lock waits, and database resource utilization helps determine whether the change introduced inefficient queries or contention.

Application and infrastructure metrics remain useful for confirming the scope of the issue, but the timing of the event identifies the database change as the most relevant variable. Database monitoring tools can correlate slow queries with the deployment window and help the engineer decide whether to optimize queries, add an appropriate index, or roll back the schema change. See Microsoft Learn: Query Performance Insight.

QUESTION NO: 7

A bank informs an administrator that changes must be made to backups for long-term reporting purposes. Which of the following is the most important change the administrator should make to satisfy these requirements?

A. Location of the backups

B. Type of the backups

C. Retention of the backups

D. Schedule of the backups

ANSWER: C

Explanation:

Retention of the backups is the most important change because long-term reporting depends on historical records still being available when required for regulatory reviews, audits, investigations, financial reporting, or internal analysis. A backup retention policy specifies how long backup copies must be preserved before expiration or deletion. For a bank, this period should be defined according to applicable legal, regulatory, contractual, and organizational recordkeeping requirements, then implemented consistently in the backup platform.

Effective retention planning also means ensuring that the retained copies remain recoverable and protected throughout their required lifetime. This can include selecting durable storage tiers, maintaining sufficient backup catalog and metadata information to locate historical records, and using controls such as immutable or write-once storage where appropriate. Retention requirements should be documented as part of business continuity and data-protection policy, with periodic testing to confirm that older backup sets can actually be restored for reporting requests. NIST contingency-planning guidance identifies backup and recovery capabilities as essential to maintaining and recovering information systems: NIST SP 800-34 Rev. 1. NIST also provides guidance on protecting stored data through appropriate data-security controls: NIST SP 800-111.

QUESTION NO: 8

A web server hosted in a PaaS public cloud platform has been the target of DDoS attacks. A cloud engineer needs to protect the server from future volumetric DDoS attacks. Which of the following are the best actions for the engineer to take? (Select two.)

A. Add a CDN in front of the web server.

B. Add a DLP on the PaaS platform.

C. Add a hardened web server instance.

D. Add an ACL to the web server network.

E. Add an IDS in-line on the web server network.

F. Add a WAF in front of the web server.

ANSWER: A F

Explanation:

“Add a CDN in front of the web server.” is appropriate because a content delivery network places globally distributed edge infrastructure between clients and the PaaS-hosted origin. During a volumetric attack, the CDN can absorb and distribute large traffic volumes across its edge network rather than allowing all traffic to consume the origin’s available bandwidth or connection capacity. Caching also reduces the number of requests that must reach the web server, preserving origin resources for legitimate traffic. CDN-based DDoS protections are a key part of resilient internet-facing cloud architecture; see Cloudflare’s DDoS mitigation overview.

“Add a WAF in front of the web server.” provides an additional edge-layer control that inspects web requests and applies security policies before traffic reaches the application. A cloud WAF can enforce rate-based rules, identify abnormal request patterns, filter malicious bots, and block unwanted HTTP/S requests. While a WAF is especially valuable for application-layer attacks, its rate limiting and managed protections complement the CDN’s traffic-distribution capability as part of layered DDoS defense. This combination helps protect availability while retaining the ability to apply application-specific controls. The CISA DDoS mitigation guidance recommends using layered defensive measures and provider-supported mitigation capabilities for internet-facing services.

QUESTION NO: 9 - (HOTSPOT)

HOTSPOT

A highly regulated business is required to work remotely, and the risk tolerance is very low. You are tasked with providing an identity solution to the company cloud that includes the following:

secure connectivity that minimizes user login

tracks user activity and monitors for anomalous activity

requires secondary authentication

INSTRUCTIONS

Select controls and servers for the proper control points.





ANSWER:



Explanation:

A VPN client is the appropriate endpoint-side control because it establishes an encrypted, authenticated tunnel from a remote device to the organization's cloud environment. This gives remote workers secure connectivity across the public Internet while supporting a low-risk operating model. Once the secure tunnel and user session are established, users can access authorized cloud resources without repeatedly signing in to each individual service, helping minimize unnecessary login prompts. The VPN client is installed and initiated on the user endpoint, so it belongs at the connection point between the endpoints and the cloud provider.

MFA belongs at the cloud access and identity control point because it requires users to prove their identity with an additional factor beyond a password. Applying multifactor authentication to cloud access is particularly important for a highly regulated organization, since a compromised password alone should not allow access to sensitive services or data. The cloud identity layer can enforce MFA during authentication and apply consistent access policies to remote users. CompTIA's guidance on MFA concepts is also consistent with using multiple independent authentication factors for stronger account protection; see NIST's multifactor authentication guidance.

A SIEM is the correct server-side monitoring component because it collects and correlates logs and security events from endpoints, VPN infrastructure, cloud services, and identity systems. Centralized event analysis provides the organization with an audit trail of user activity and can identify patterns that indicate anomalous or potentially malicious behavior. This supports both continuous security monitoring and the evidence collection needs common in regulated environments. The

role of log management and monitoring in detecting incidents is described in NIST SP 800-92, Guide to Computer Security Log Management.

QUESTION NO: 10

A DevOps team wants to improve the reliability of application releases through a CI/CD pipeline. Which of the following controls should the team implement before deploying to production? (Select two).

A. Automated testing

B. Manual approval gates

C. Direct production changes by developers

D. Disabling pipeline logs

E. Using the same credentials for every environment

ANSWER: A B

Explanation:

Automated testing validates application behavior consistently before a release reaches production. Unit, integration, security, and regression tests can identify defects early and prevent a pipeline from promoting a build that does not meet defined quality criteria.

Manual approval gates provide a controlled authorization point before a production deployment. An authorized stakeholder can review test results, change records, business timing, risk, and rollback readiness before allowing the pipeline to continue. This is particularly useful where organizational change-management requirements require documented approval for production releases. See GitHub Docs: Deploying with GitHub Actions and Microsoft Learn: Pipeline approvals.

QUESTION NO: 11

A systems administrator needs to configure a script that will monitor whether an application is healthy and stop the VM if an unsuccessful code is returned. Which of the following

scripts should the systems administrator use to achieve this goal?

A. string APP_URL int RESPONSE_CODE string VM health_checker(APP_URL, VM) { RESPONSE_CODE = http_probe(APP_URL) if [RESPONSE_CODE == 200] { echo RESPONSE_CODE } else { stop(VM); echo RESPONSE_CODE } }

B. string APP_URL float RESPONSE_CODE string VM health_checker(APP_URL, VM) { if [http_probe(APP_URL) == 200] { stop(RESPONSE_CODE); echo VM } else { stop(VM); echo RESPONSE_CODE } }

C. string APP_URL int RESPONSE_CODE string VM health_checker(APP_URL, VM) { if [http_probe(APP_URL) == 200] { echo RESPONSE_CODE } else { stop(VM); echo RESPONSE_CODE } }

D. string APP_URL int RESPONSE_CODE string VM health_checker(APP_URL, VM) { if [http_probe(VM) == 200] { stop(VM); echo RESPONSE_CODE } else { echo RESPONSE_CODE } }

ANSWER: A

Explanation:

The script that calls `http_probe(APP_URL)`, treats an HTTP status code of `200` as healthy, and executes `stop(VM)` in the `else` branch satisfies the requirement. It checks the application endpoint rather than merely checking the virtual machine, which is important because a VM can be powered on while the application hosted on it is unavailable or malfunctioning. The probe result is stored as a response code, allowing the script to make an explicit health decision. A `200 OK` response indicates that the request succeeded, so the script logs the successful response and leaves the VM running.

Any response other than the expected successful code follows the failure path and stops the specified VM, meeting the requested automated remediation behavior. This is an example of application-level health monitoring combined with an automated action. HTTP status-code semantics, including `200 OK`, are defined in RFC 9110. This type of monitoring and operational-response decision aligns with the cloud operations, troubleshooting, and automation skills covered by CompTIA Cloud+; see the current CompTIA exam objectives.

QUESTION NO: 12

A company is developing an incident-response process for cloud resources. Which of the following actions should be completed before a security incident occurs? (Select two).

A. Creating incident-response playbooks

B. Enabling audit logging

C. Deleting affected resources immediately

D. Disabling all user accounts permanently

E. Publishing incident details to all users

ANSWER: A B

Explanation:

Creating incident-response playbooks prepares responders to handle common cloud events consistently. A playbook can define triage steps, containment actions, escalation paths, evidence requirements, communications procedures, and recovery tasks for scenarios such as compromised credentials, malware, or unauthorized configuration changes.

Enabling audit logging ensures that responders have historical evidence available during an investigation. Cloud audit logs can record management-plane actions, identity activity, configuration modifications, and API calls. These records support timeline reconstruction, root cause analysis, and determination of the affected resources. NIST incident-handling guidance identifies preparation as a core phase of incident response. See NIST SP 800-61 Rev. 2 and AWS CloudTrail User Guide.

QUESTION NO: 13

A company is migrating a customer-facing application to a public cloud provider. The security team needs to reduce the attack surface of the application servers. Which of the following should the team implement? (Select two).

A. Network segmentation

B. Removing unnecessary services

C. Allowing all outbound traffic

D. Using a shared administrator account

E. Assigning public IP addresses to database servers

ANSWER: A B

Explanation:

Network segmentation reduces attack surface by separating application components into distinct network tiers and allowing only required traffic flows between them. For example, web servers can be placed in a public subnet while application and database services remain in private subnets. Security rules can then limit communication to approved ports, protocols, and sources.

Removing unnecessary services also reduces attack surface because each enabled service, listening port, software package, and administrative interface can introduce an additional exposure. Systems should run only the components needed for the intended workload. This supports a hardened baseline and simplifies vulnerability management. CISA

recommends reducing exposure through secure configuration and removal of unnecessary services. See CISA: Securing Network Infrastructure Devices and CIS Benchmarks.

QUESTION NO: 14

Which of the following is a customer be responsible for in a provider-managed database service? (Select two).

A. Operating system patches

B. Table-level permissions

C. Minor database engine updates

D. Cluster configuration

E. Row-level encryption

F. Availability of hardware for scaling

ANSWER: B E

Explanation:

Table-level permissions and row-level encryption remain customer responsibilities in a provider-managed database service because they govern how the customer's own data is accessed and protected. The provider operates the managed service platform, but it cannot determine which identities, applications, or business roles should be permitted to read, insert, update, or delete data in particular tables. Customers must therefore implement database authorization policies that enforce least-privilege access.

Customers are also responsible for selecting and configuring appropriate protections for sensitive data at the logical data layer. Row-level encryption can protect specific records or values according to the organization's security, privacy, and compliance requirements. Although a cloud provider may offer encryption capabilities and may encrypt underlying storage, the customer must decide what data requires additional protection, manage access to decrypted data, and configure the service features accordingly. This reflects the shared-responsibility approach: provider-managed operations reduce infrastructure administration, while customers retain responsibility for data governance, identity and access controls, and data-security configuration. See the Amazon RDS User Guide and Azure SQL Database security overview.

QUESTION NO: 15

Which of the following are best practices when working with a source control system? (Select two).

A. Merging code often

B. Pushing code directly to production

C. Performing code deployment

D. Maintaining one branch for all features

E. Committing code often

F. Initiating a pull request

ANSWER: A E

Explanation:

Merging code often and **Committing code often** are foundational source-control practices because they keep work small, traceable, and continuously integrated. Frequent commits create discrete records of changes, making it easier to review history, identify the source of an issue, revert a problematic change, and collaborate without combining a large number of

unrelated edits into one update. A meaningful commit history also supports troubleshooting and change-management activities common in cloud environments.

Frequent merging, or regularly integrating current shared-branch changes into active work, reduces divergence between developers' branches. This surfaces integration problems and merge conflicts earlier, while the changes are smaller and the relevant context is still clear. In CI/CD workflows, regular integration helps ensure that code is validated together rather than remaining isolated until a late, high-risk integration event. These practices support a reliable pipeline by enabling teams to make incremental changes, test them continuously, and maintain an understandable repository history. Git's documentation describes commits as recorded snapshots of project changes and covers the normal add-and-commit workflow; GitHub also describes collaborative pull-request workflows that build on branches and integrated changes. See the Git documentation on recording changes and GitHub Flow documentation.

QUESTION NO: 16

A cloud administrator shortens the amount of time a backup runs. An executive in the company requires a guarantee that the backups can be restored with no data loss. Which of thâ,¬ following backup features should the administrator lest for?

A. Encryption

B. Retention

C. Schedule

D. Integrity

ANSWER: D

Explanation:

Integrity is the appropriate backup feature to test because it validates that backup data remains complete, accurate, and usable for restoration. After a backup process is changed to reduce its runtime, integrity validation provides assurance that the accelerated process has not resulted in corrupted, incomplete, or unreadable backup sets. Common integrity mechanisms include checksums, hashes, validation scans, and test restores. These mechanisms compare protected data or backup blocks against expected values and confirm that the backup system can successfully read the data required for recovery.

A successful backup job alone confirms that the job completed; integrity testing adds evidence that the resulting recovery point is reliable. In a cloud environment, this validation is especially important because backup data may be compressed, deduplicated, encrypted, transferred over networks, and stored across distributed infrastructure. Regular automated integrity checks and restore testing should be incorporated into the backup policy so that restore capability is verified before an incident occurs. Microsoft describes backup validation as a way to ensure recovery points are available and recoverable, while Veeam describes health checking as validation of backup-file integrity. See Microsoft Learn: Verify Azure Backup and Veeam: Backup Integrity Check.

QUESTION NO: 17

A company stores regulated data in cloud object storage. The security team needs to control who can use encryption keys and ensure that cryptographic keys are periodically replaced. Which of the following should the team implement? (Select two).

A. Customer-managed keys

B. Key rotation

C. Public object access

D. Shared administrator accounts

E. Unencrypted snapshots

F. Static IP addresses

ANSWER: A B

Explanation:

Customer-managed keys allow the organization to control key policies, define which identities and services may use keys, and audit cryptographic operations. This provides greater control than relying solely on provider-managed default encryption keys.

Key rotation replaces cryptographic key material on a defined schedule or after a security event. Rotation limits the amount of data protected by a single key version and reduces risk if a key is exposed. Key rotation should be planned to ensure applications can continue decrypting data encrypted under previous key versions when required.

Cloud key-management services provide centralized controls for key policies, auditing, and rotation. AWS documents customer managed keys at AWS KMS keys and key rotation at Rotating AWS KMS keys.

QUESTION NO: 18

An organization is concerned that its IT help desk is receiving too many calls about password resets for various internal and external applications. The organization implements a single sign-on (SSO) solution. Which of the following technologies should the cloud architect at the organization recommend implementing?

A. SAML

B. LDAP

C. Kerberos

D. MFA

ANSWER: A

Explanation:

SAML is the appropriate technology because it is a federation standard designed to enable single sign-on across separate web applications, including internal services and external SaaS applications. In a SAML-based SSO flow, the user authenticates with a central identity provider, which then issues a signed SAML assertion to the application acting as the service provider. The assertion communicates the authenticated identity and relevant authorization attributes, allowing the application to grant access without requiring the user to maintain or enter a separate application-specific password.

Centralizing authentication in this way directly addresses password-reset volume. Users have a primary organizational credential rather than distinct credentials for every connected application, and authentication policies such as password lifecycle requirements can be administered centrally by the identity provider. SAML also supports browser-based redirection and signed assertions, making it a common choice for secure federation between an organization and cloud-hosted applications. It can be used alongside stronger controls such as multifactor authentication at the identity provider, while preserving the streamlined SSO experience. See the OASIS SAML 2.0 Core specification and Microsoft guidance for planning SSO deployment.

QUESTION NO: 19

Participant attendance has been poor at the last few weekly meetings. The project manager wants to validate who the participants should be, how often the meetings should occur, and on which day and at what time the meeting should be held. Which of the following is the best place to validate this information?

A. Project communication plan

B. Project management plan

C. Project business case

D. Project schedule

ANSWER: A

Explanation:

Project communication plan is correct because it defines the approach for communicating project information to stakeholders. A practical communications plan identifies the intended audience or participants for each communication, the communication method or channel, the purpose of the communication, and its required timing and frequency. For recurring project meetings, this is the governing source for confirming who needs to attend, whether the meeting should be weekly or use another cadence, and the agreed schedule or preferred time window.

Poor attendance is a signal for the project manager to review the established communication requirements and verify that the meeting still serves the right stakeholders at an appropriate cadence. The plan provides a basis for realigning the meeting with stakeholder needs, such as confirming required versus optional participants and documenting an updated meeting time or frequency after appropriate stakeholder engagement. This keeps project communications intentional, targeted, and consistent rather than treating every meeting invitation as a universal requirement. PMI describes the communications management plan as a component used to document how project communications will be planned, structured, monitored, and controlled. See PMI's communications management plan resource and Atlassian's guidance on establishing team working agreements.

QUESTION NO: 20

A cloud engineer is developing an operating expense report that will be used to purchase various cloud billing models for virtual machine instances. The cloud billing model must meet the following requirements:

â€¢ The instance cannot be ephemeral.

â€¢ The minimum life cycle of the instance is expected to be five years.

â€¢ The software license is charged per physical CPU count.

Which of the following models would best meet these requirements?

A. Dedicated host

B. Spot instance

C. Pay-as-you-go

D. Reserved resources

ANSWER: A

Explanation:

Dedicated host is the best fit because it provides access to an entire physical server for the customer's exclusive use. This is particularly important when software licensing is calculated from physical CPU, socket, or core counts rather than from virtual CPUs assigned to an individual VM. A dedicated host provides the hardware visibility and placement control needed to apply those host-based licensing terms accurately.

The expected five-year lifecycle also supports using dedicated capacity for a stable, long-running workload. The VM can remain persistent rather than being designed for interruption or short-lived execution, and the organization can plan its infrastructure and licensing costs around a consistent physical-host footprint. In practice, a cloud provider may allow dedicated-host capacity to be reserved for a commitment term and renewed as needed to support a multi-year workload. This combines predictable infrastructure placement with the ability to bring or retain eligible server-bound licenses.

For example, AWS describes Dedicated Hosts as physical servers dedicated to a customer and specifically identifies their usefulness for software licensing that is tied to physical cores, sockets, or VMs. See AWS Dedicated Hosts overview and AWS EC2 Dedicated Hosts.

QUESTION NO: 21 - (SIMULATION)

A company has decided to scale its e-commerce application from its corporate datacenter to a commercial cloud provider to meet an anticipated increase in demand during an upcoming holiday.

The majority of the application load takes place on the application server under normal conditions. For this reason, the company decides to deploy additional application servers into a commercial cloud provider using the on-premises orchestration engine that installs and configures common software and network configurations.

The remote computing environment is connected to the on-premises datacenter via a site-to-site IPSec tunnel. The external DNS provider has been configured to use weighted round-robin routing to load balance connections from the Internet.

During testing, the company discovers that only 20% of connections completed successfully.

INSTRUCTIONS

Review the network architecture and supporting documents and fulfill these requirements:

Part 1:



Analyze the configuration of the following components: DNS, Firewall 1, Firewall 2, Router 1, Router 2, VPN and Orchestrator Server.



Identify the problematic device(s).

Part 2:

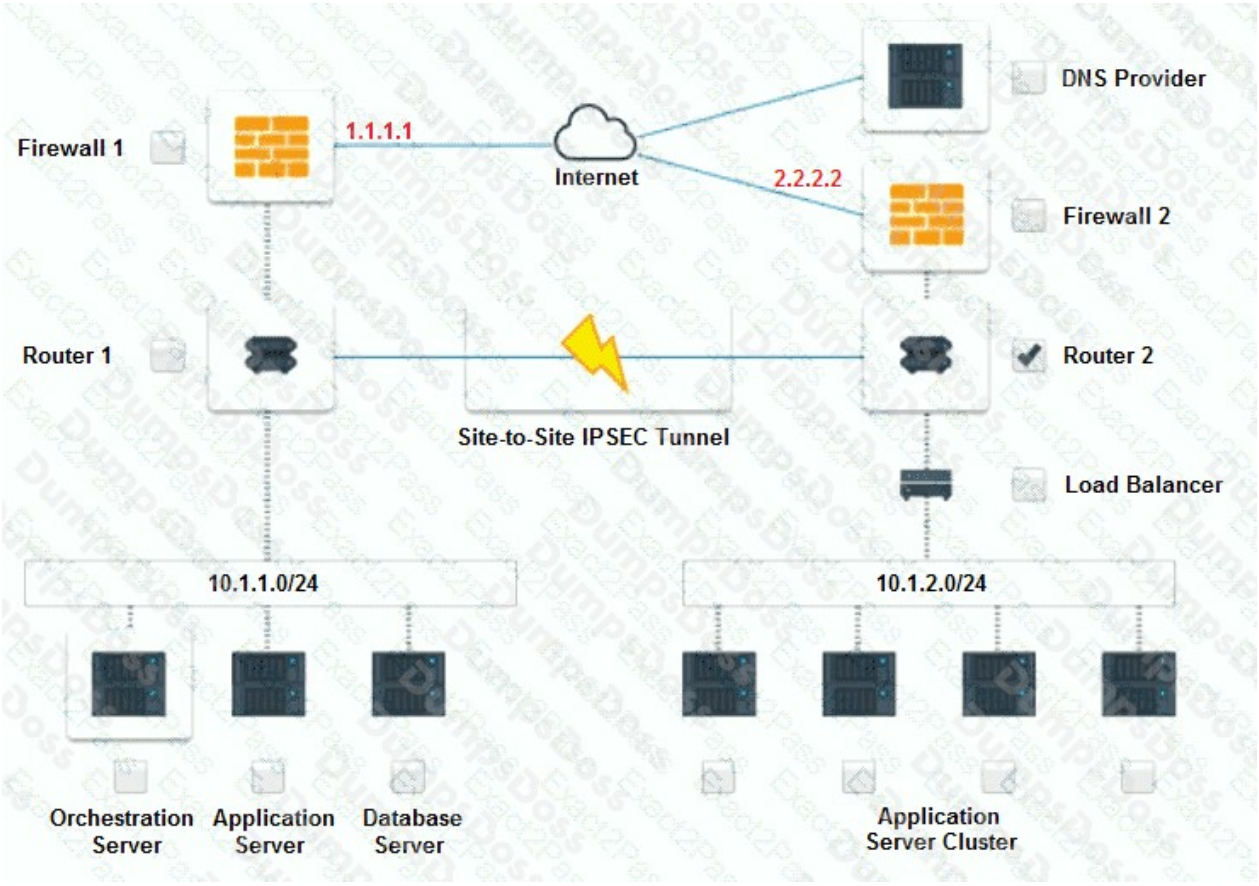


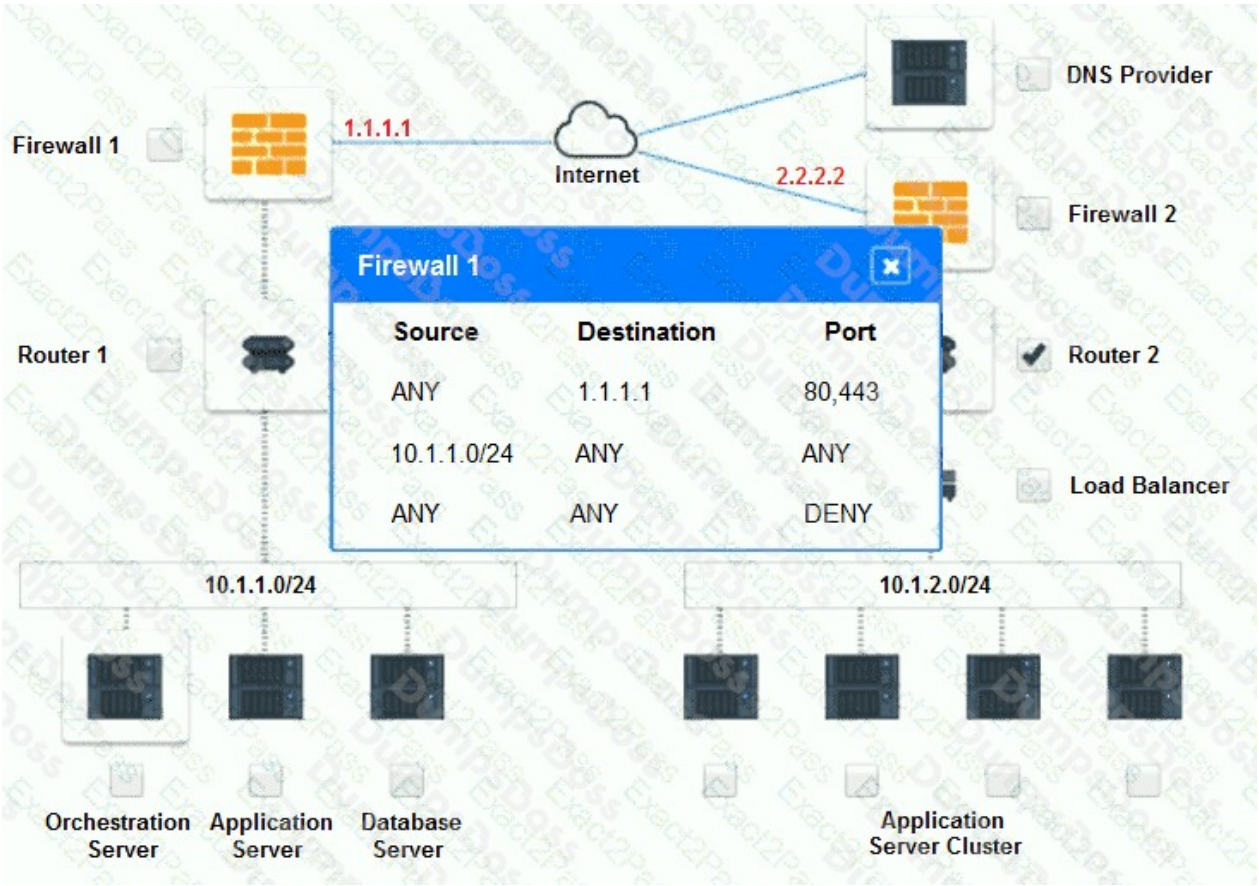
Identify the correct options to provide adequate configuration for hybrid cloud architecture.

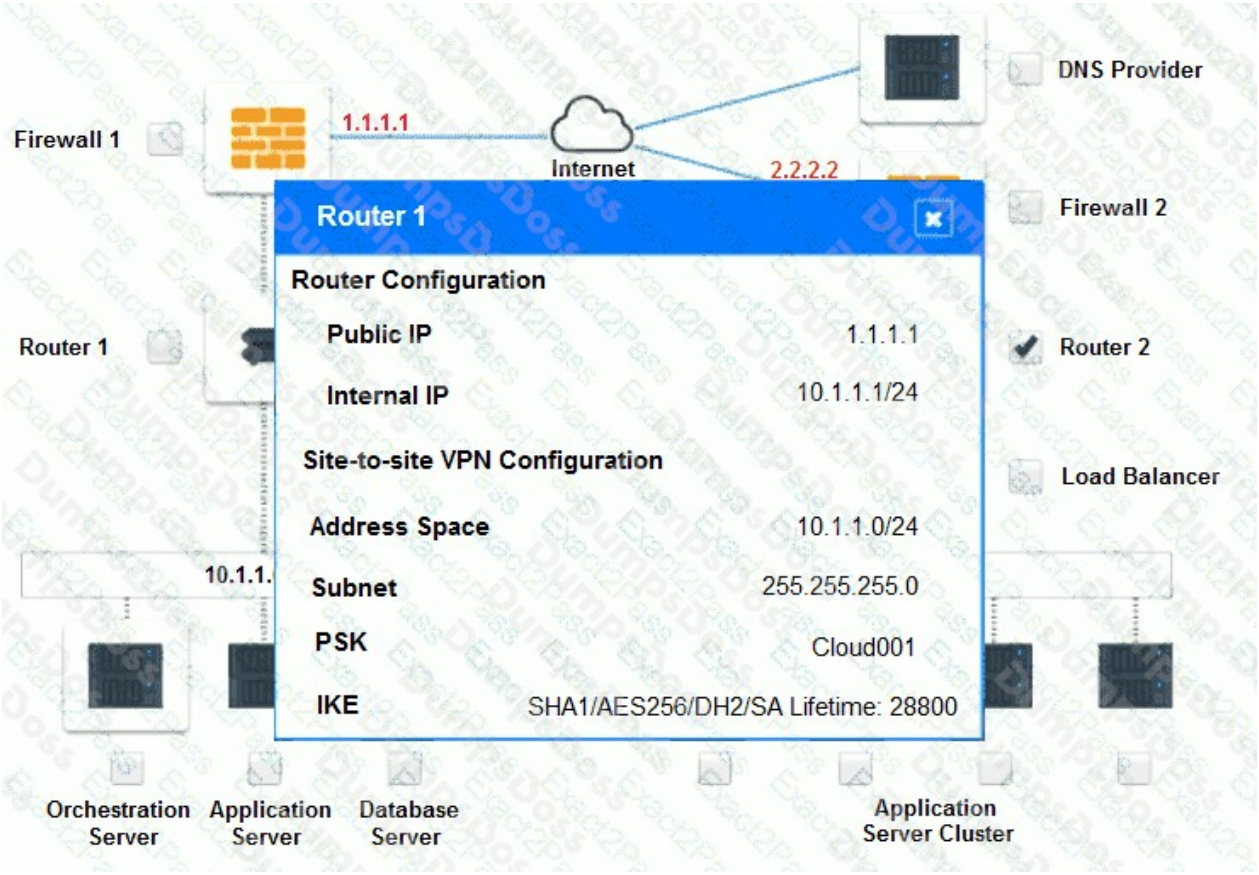
If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

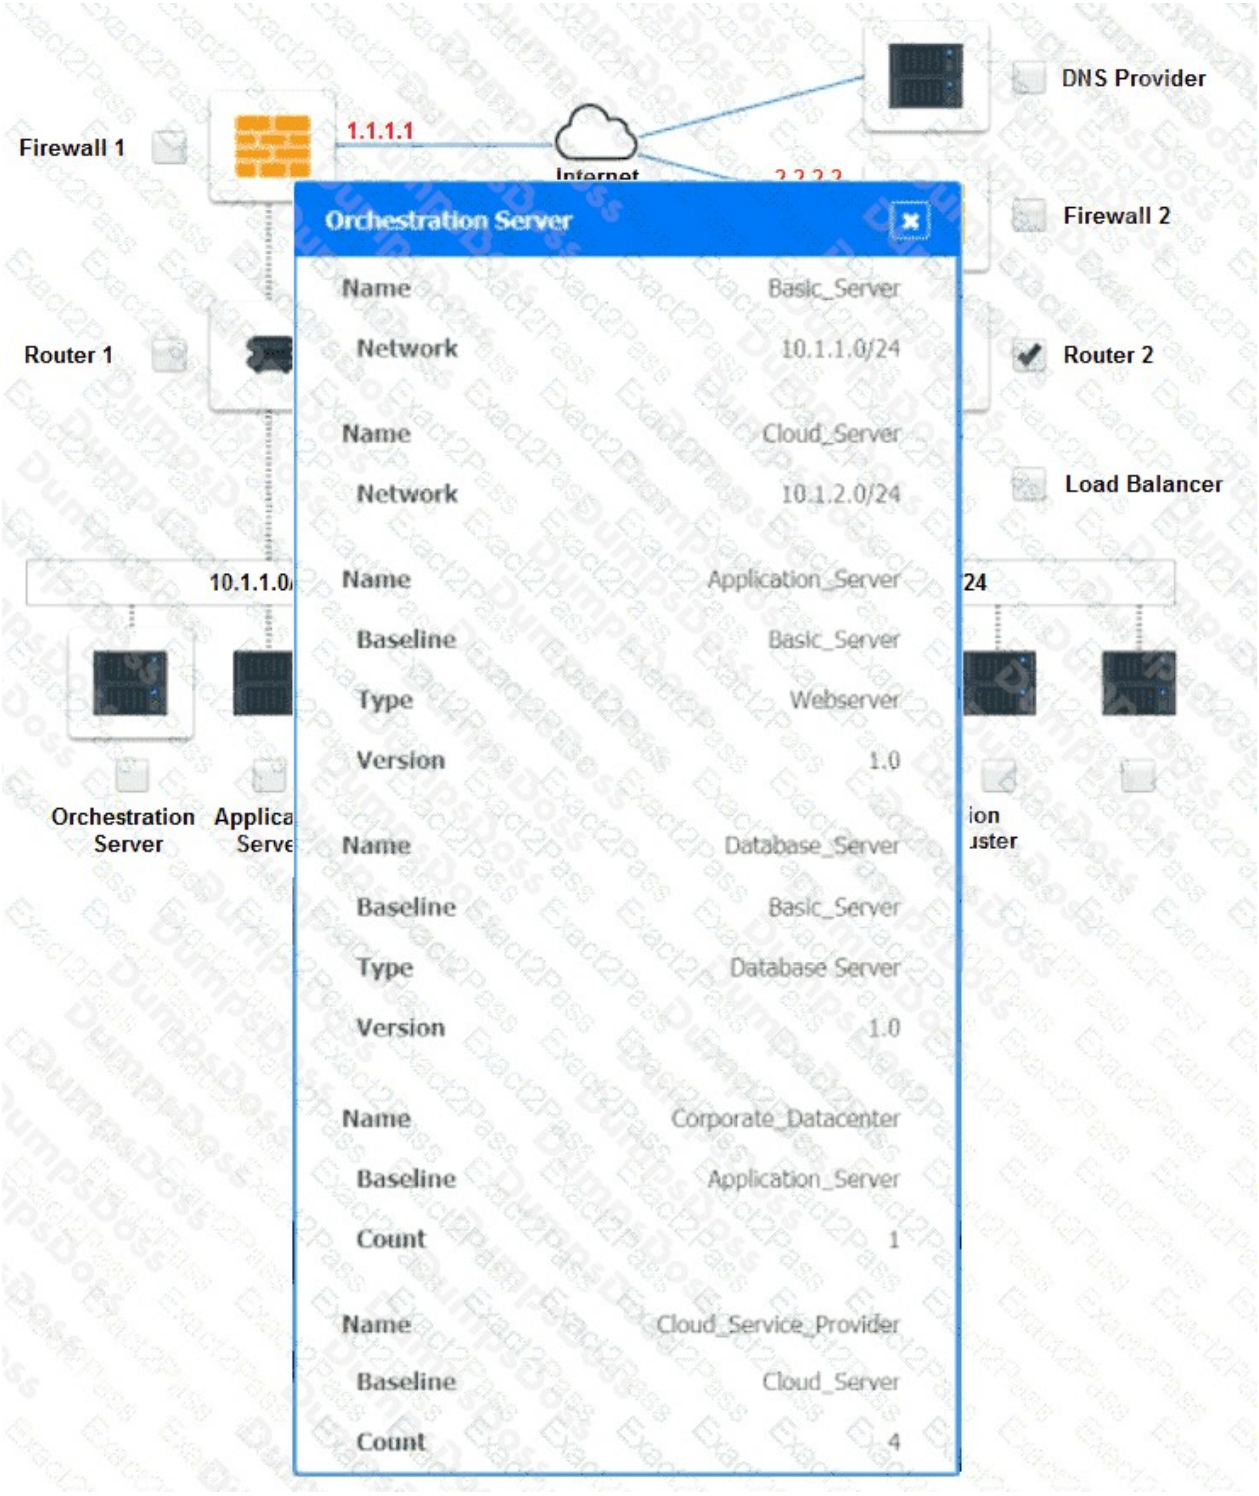
Part 1:

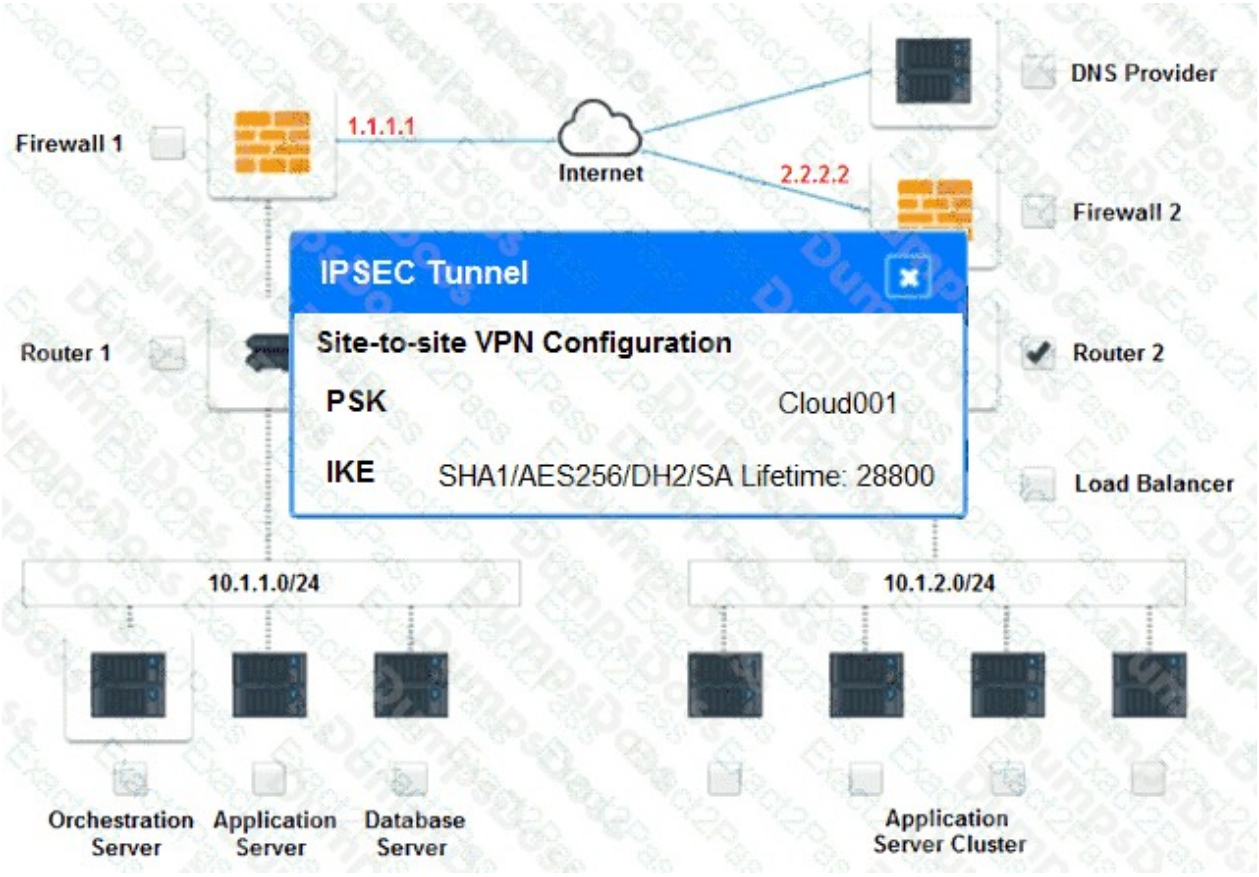
Cloud Hybrid Network Diagram

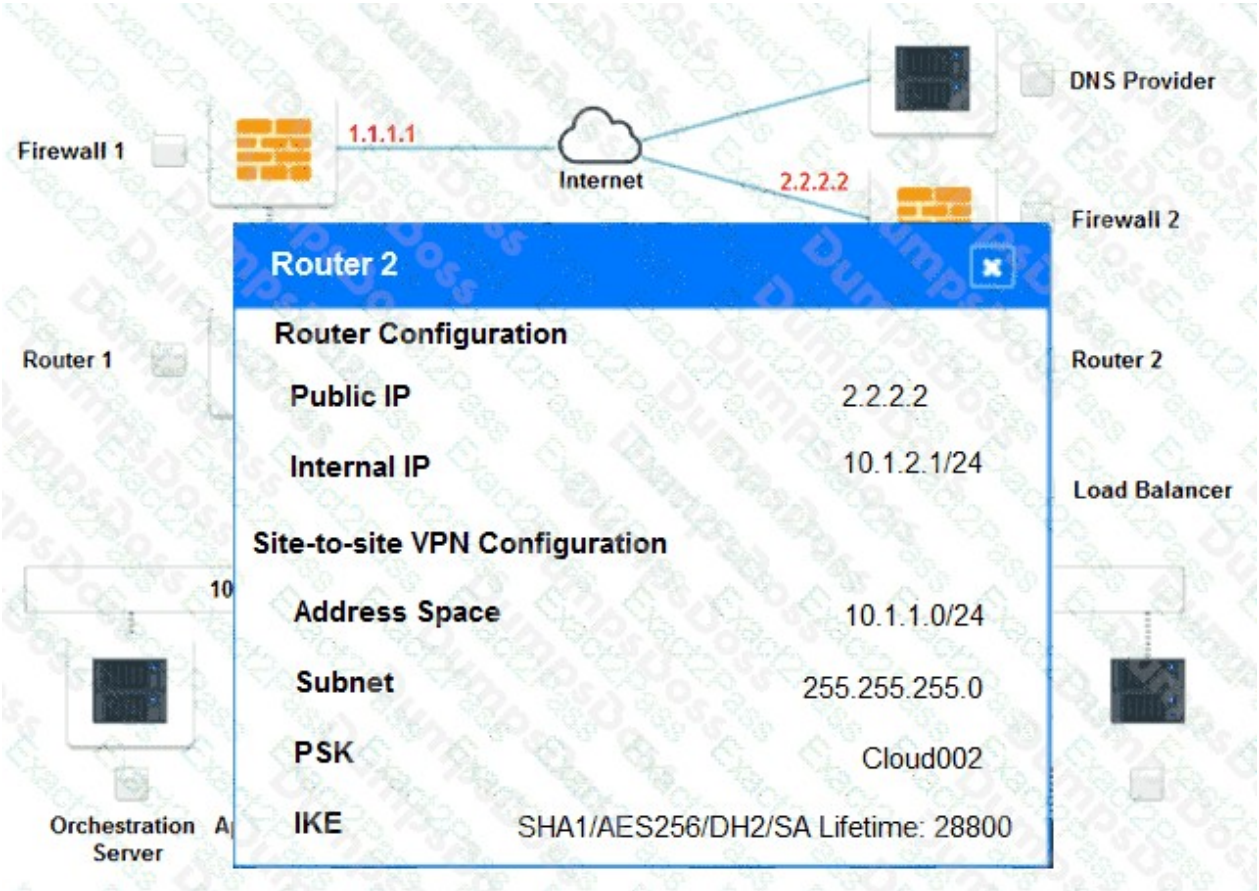

Firewall 1
1.1.1.1
Internet
2.2.2.2
DNS Provider
Firewall 2
Router 1
Router 2
Site-to-Site IPSEC Tunnel
Load Balancer
10.1.1.0/24
10.1.2.0/24
Orchestration Server
Application Server
Database Server
Application Server Cluster

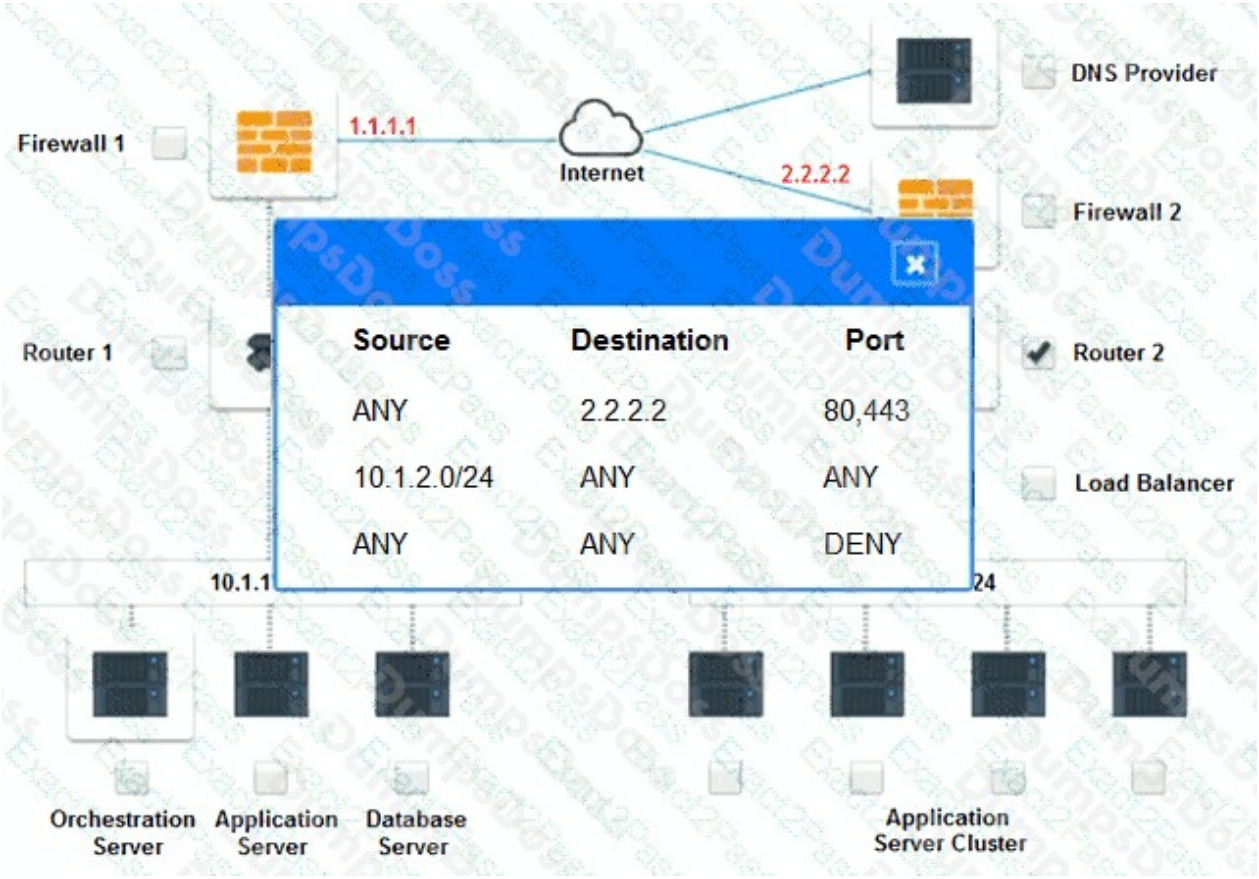

Firewall 1
1.1.1.1
Internet
2.2.2.2
DNS Provider
Firewall 2
Router 1
Router 2
Load Balancer
Firewall 1
Source
Destination
Port
ANY
1.1.1.1
80,443
10.1.1.0/24
ANY
ANY
ANY
ANY
DENY
10.1.1.0/24
10.1.2.0/24
Orchestration Server
Application Server
Database Server
Application Server Cluster

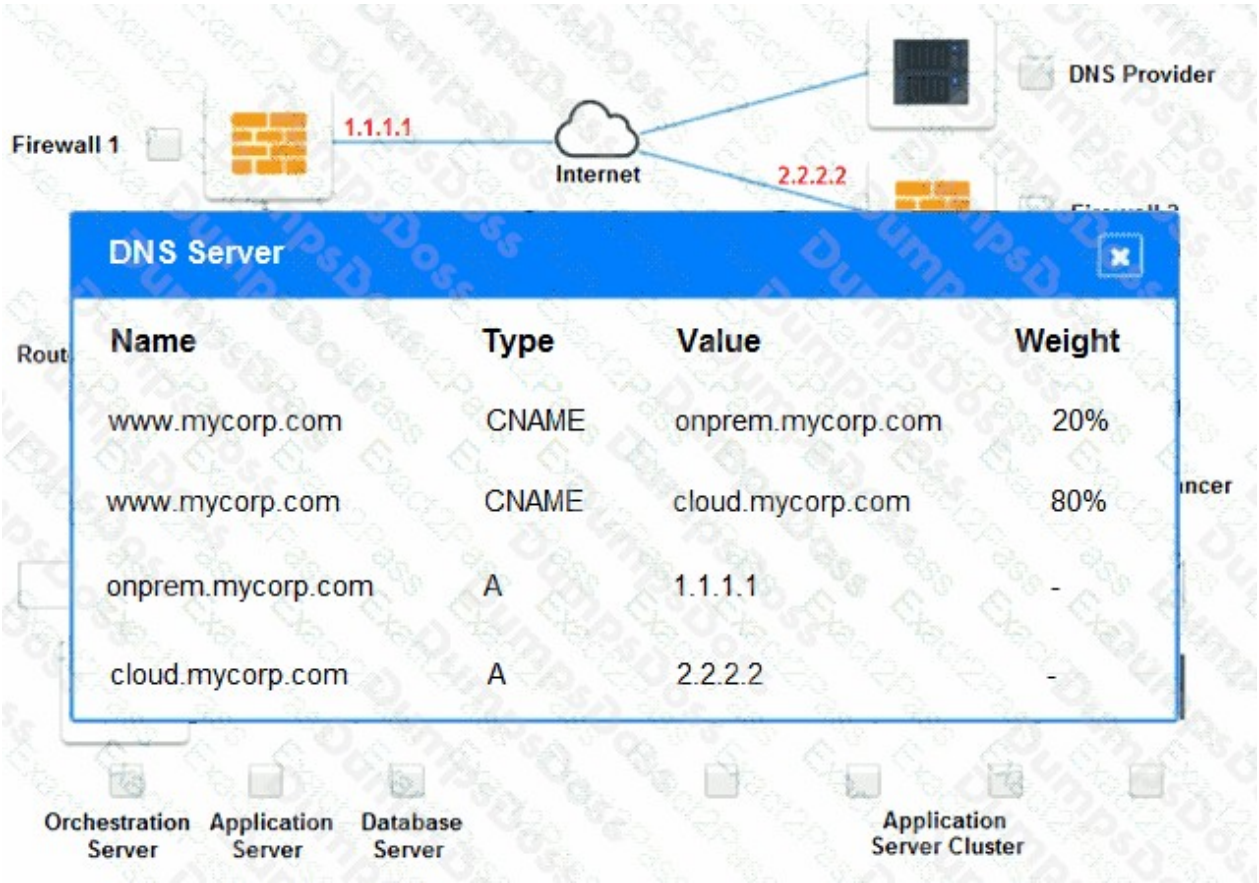

DNS Provider
Firewall 1
1.1.1.1
Internet
2.2.2.2
Router 1
Router Configuration
Public IP
1.1.1.1
Internal IP
10.1.1.1/24
Site-to-site VPN Configuration
Address Space
10.1.1.0/24
Subnet
255.255.255.0
PSK
Cloud001
IKE
SHA1/AES256/DH2/SA Lifetime: 28800
Firewall 2
Router 1
Router 2
Load Balancer
10.1.1.
Orchestration Server
Application Server
Database Server
Application Server Cluster


Firewall 1
1.1.1.1
Internet
DNS Provider
Firewall 2
Router 1
Router 2
Load Balancer
10.1.1.0/
24
Orchestration Server
Applica
Serve
ion
uster
Orchestration Server
Name Basic_Server
Network 10.1.1.0/24
Name Cloud_Server
Network 10.1.2.0/24
Name Application_Server
Baseline Basic_Server
Type Webserver
Version 1.0
Name Database_Server
Baseline Basic_Server
Type Database Server
Version 1.0
Name Corporate_Datacenter
Baseline Application_Server
Count 1
Name Cloud_Service_Provider
Baseline Cloud_Server
Count 4


DNS Provider
Firewall 1
1.1.1.1
Internet
2.2.2.2
Firewall 2
IPSEC Tunnel
Site-to-site VPN Configuration
PSK
Cloud001
IKE
SHA1/AES256/DH2/SA Lifetime: 28800
Router 1
Router 2
Load Balancer
10.1.1.0/24
10.1.2.0/24
Orchestration Server
Application Server
Database Server
Application Server Cluster


DNS Provider
Firewall 1
1.1.1.1
Internet
2.2.2.2
Firewall 2
Router 2
Router Configuration
Public IP
2.2.2.2
Internal IP
10.1.2.1/24
Site-to-site VPN Configuration
Address Space
10.1.1.0/24
Subnet
255.255.255.0
PSK
Cloud002
IKE
SHA1/AES256/DH2/SA Lifetime: 28800
Router 1
Router 2
Load Balancer
Orchestration Server





Part 2:

Only select a maximum of TWO options from the multiple choice question

- ☐ Deploy a Replica of the Database Server in the Cloud Provider.
- ☐ Update the PSK (Pre-shared key) in Router 2.
- ☐ Update the A record on the DNS from 2.2.2.2 to 1.1.1.1.
- ☐ Promote deny All to allow All in Firewall 1 and Firewall 2.
- ☐ Change the Address Space on Router 2.
- ☐ Change internal IP Address of Router 1.
- ☐ Reverse the Weight property in the two CNAME records on the DNS.
- ☐ Add the Application Server at on-premises to the Load Balancer.

ANSWER: See the explanation for the answer

Explanation:

Part 1 — Problematic device: Router 2

Router 2 is misconfigured for the site-to-site IPSec VPN. The 20% successful-connection rate corresponds to the DNS weight for the on-premises endpoint; the 80% of requests directed to the cloud environment fail because the hybrid VPN configuration is not matched correctly.

Router 1 uses the PSK `Cloud001`, while Router 2 uses `Cloud002`.
Router 2 specifies VPN address space `10.1.1.0/24`, which is the on-premises network. Router 2 must use its cloud-side network, `10.1.2.0/24`.

Part 2 — Select these two options:

Update the PSK (Pre-shared key) in Router 2. Set it to `Cloud001` so it matches Router 1 and the IPSec tunnel configuration.
Change the Address Space on Router 2. Change it from `10.1.1.0/24` to `10.1.2.0/24`.

These two changes make the IPSec tunnel parameters and protected cloud network correct, restoring communication between the on-premises orchestration/services and the cloud application-server cluster.

QUESTION NO: 22

A customer is migrating applications to the cloud and wants to grant authorization based on the classification levels of each system. Which of the following should the customer implement to ensure authorisation to systems is granted when the user and system classification properties match? (Select two).

A. Resource tagging

B. Discretionary access control

C. Multifactor authentication

D. Role-based access control

E. Token-based authentication

F. Bastion host

G. Mandatory access control

ANSWER: A G

Explanation:

Resource tagging and **Mandatory access control** work together to enforce access decisions based on classification properties. Resource tags provide a consistent way to attach classification metadata—such as public, internal, confidential, or restricted—to cloud resources. Those tags make the resource's classification available for use in policy evaluation across a cloud environment.

Mandatory access control enforces centrally defined policy using security labels and clearances. A subject's authorization attributes, such as clearance level, are evaluated against the classification label applied to the target system or resource. Access is granted only when the defined relationship between the user's classification and the system's classification permits it. Unlike owner-managed permissions, this model is governed by organizational policy and is well suited to environments that handle classified workloads.

In practice, the organization should establish a controlled tagging standard, apply classification tags consistently to systems and data, and configure mandatory policies that compare those labels with user authorization properties. This produces an enforceable classification-based authorization model rather than relying solely on job role or individual resource-owner decisions. NIST describes attribute- and policy-driven authorization concepts in NIST SP 800-162; its glossary also defines mandatory access control.

QUESTION NO: 23

A project team that handles sensitive data analysis recently onboarded three new employees. Which of the following should the PM do first?

A. Check employee security clearance

B. Allow access to the internet on corporate computers

C. Train the new employees to handle sensitive data

D. Execute an analysis of the sensitive data

ANSWER: A

Explanation:

Checking employee security clearance is the appropriate first action because the project manager must establish that each new employee has been vetted and authorized before the organization permits any involvement with sensitive data. Authorization decisions should be based on documented personnel screening, role requirements, and approved access eligibility. This creates an auditable basis for assigning access and helps ensure that access is limited to individuals with an established business need and appropriate trust determination.

Personnel security and access control are foundational safeguards for systems and information processing. In practice, confirmation of clearance or equivalent authorization should occur before the project manager assigns sensitive-data responsibilities, provisions accounts, or delegates analysis work. Once authorization is verified, the organization can apply least-privilege access appropriate to each employee's role and complete any required security and data-handling onboarding. Establishing eligibility first supports confidentiality requirements and gives subsequent access provisioning and work assignments a valid governance basis. NIST SP 800-53 identifies personnel screening and access enforcement as important control areas; see NIST SP 800-53 Rev. 5. NIST's guidance on access control also emphasizes restricting system access to authorized users and approved activities: NIST Role-Based Access Control project.

QUESTION NO: 24

A retail store is rolling out a new point-of-sale solution to several locations where the staff members normally use cash registers. Which of the following activities will best assist staff members with integrating this new solution? (Select two).

A. Documentation

B. Team collaboration

C. Gap analysis

D. Training sessions

E. Touch points

F. Progress monitoring

ANSWER: A D

Explanation:

Documentation and **Training sessions** most directly support successful adoption of a new point-of-sale solution by staff transitioning from traditional cash registers. Clear, accessible documentation gives employees dependable operational guidance at the point of need. This can include quick-start instructions, standard procedures for sales and refunds, guidance for opening and closing tills, escalation contacts, and concise troubleshooting steps. Such material promotes consistent use across multiple store locations and remains available after the initial rollout.

Training sessions give staff the practical knowledge and confidence required to use the new system while serving customers. Effective training should allow employees to perform realistic workflows, such as scanning items, processing the supported payment methods, handling returns, applying discounts, and responding to common system prompts. Practice before production use reduces operational disruption and helps staff integrate the system into their daily routines. CompTIA's Cloud+ materials include operational procedures and documentation as relevant skills for supporting technology environments; see the CompTIA Cloud+ certification page. For guidance on planning and delivering end-user training during technology adoption, see the Microsoft training and support resources.

QUESTION NO: 25

A cloud architect attempts to modify a protected branch but is unable to do so. The architect receives an error indicating the action cannot be completed. Which of the following should the architect try instead"

A. Adding a new remote

B. Creating a pull request

C. Merging the branch

D. Rebasing the branch

ANSWER: B

Explanation:

Creating a pull request is the appropriate way to propose a change to a protected branch. Branch protection is a source-control governance mechanism that prevents unauthorized direct updates to critical branches, commonly including production, main, or release branches. Rather than changing that branch directly, the architect should make the required updates in a separate working branch and create a pull request targeting the protected branch. This provides a controlled path for the change to be reviewed, validated, approved, and merged according to the repository's configured policies.

A pull request supports the controls commonly required in cloud environments, such as peer review, required approvals, successful continuous-integration checks, signed commits, and documented change history. If the architect has the necessary permissions and all required rules are satisfied, the pull request can be merged by an authorized user or through an approved automated workflow. This approach preserves the integrity and auditability of the protected branch while allowing the intended change to proceed through the organization's established change-management process. GitHub describes protected branches and their enforcement mechanisms at GitHub Docs. GitLab likewise documents merge-request workflows and permissions for protected branches at GitLab Docs.

QUESTION NO: 26

A manager wants information about which users signed in to a certain VM during the past month. Which of the following can the cloud administrator use to obtain this information?

A. Retention

B. Alerting

C. Aggregation

D. Collection

ANSWER: D

Explanation:

Collection is correct because sign-in activity must be captured from the virtual machine's operating system and made available in a logging platform before an administrator can search for the users who authenticated during a specified historical period. For a Windows VM, this commonly includes security audit events such as successful logon events; for a Linux VM, it can include authentication records from sources such as system logs, SSH logs, or an audit framework. A cloud administrator configures log collection through an agent, extension, or cloud-native logging integration so these records are sent to a centralized workspace or logging service.

Once collected, the administrator can run a time-bounded query covering the prior month and filter the events for the target VM, relevant logon event types, and user identities. This produces the requested historical sign-in information and supports auditing and investigation. Centralized collection also avoids relying on locally stored VM logs, which may be rotated, inaccessible, or unavailable if the VM has changed. Microsoft describes how Azure Monitor collects VM and platform log data for analysis in Monitor virtual machines with Azure Monitor. AWS similarly documents collecting and centrally viewing log events with Amazon CloudWatch Logs.

QUESTION NO: 27

A cloud operations team needs to identify unexpected spending increases and ensure that resource owners are notified before departmental cloud budgets are exceeded. Which of the following should the team implement? (Select two).

A. Cost anomaly detection

B. Budget alerts

C. Disabling resource tags

D. Increasing log-retention periods

E. Using a dedicated host

F. Configuring a static route

ANSWER: A B

Explanation:

Cost anomaly detection analyzes spending patterns and identifies unusual changes that may indicate unplanned resource deployment, compromised accounts, configuration errors, or unexpected workload demand. This helps the operations team investigate cost increases before they become significant.

Budget alerts notify designated stakeholders when actual or forecasted spending approaches or exceeds an established budget threshold. When resources are consistently tagged by department or cost center, these alerts can be directed to the appropriate owners for timely remediation.

Using both controls supports proactive cloud financial management: anomaly detection identifies unexpected behavior, while budget alerts track spending against planned limits. AWS documents these capabilities at AWS Cost Anomaly Detection and AWS Budgets.

QUESTION NO: 28

A cloud engineer is troubleshooting a connectivity issue. The application server with IP 192.168.1.10 in one subnet is not connecting to the MySQL database server with IP 192.168.2 20 in a different subnet. The cloud engineer reviews the following information:

Application Server Stateful Firewall

Which of the following should the cloud engineer address lo fix the communication issue?

A. The Application Server Stateful Firewall

B. The Application Server Subnet Routing Table

C. The MySQL Server Stateful Firewall

D. The MySQL Server Subnet Routing Table

ANSWER: C

Explanation:

The MySQL Server Stateful Firewall should be addressed because the database server must permit the application server's connection to the MySQL service. MySQL normally listens for client connections on TCP port 3306, so the database-side inbound policy needs a rule that allows TCP/3306 from the application server address, 192.168.1.10, or from an appropriately scoped application-subnet CIDR range. In a stateful firewall design, once this inbound connection is allowed and established, the return traffic for that session is automatically permitted; a separate inbound rule for the response traffic is not required. The firewall associated with the destination workload is therefore the relevant enforcement point for a connection initiated by the application server toward the database service. The rule should also be verified against the database server's actual listening interface and configured port so that the permitted network flow matches the deployed MySQL service. MySQL documents port 3306 as its default classic-protocol port in its connection documentation. For an example of stateful cloud firewall behavior, AWS explains that security groups allow return traffic for permitted flows in its security group documentation.

QUESTION NO: 29

A cloud engineer is designing a disaster recovery plan for a business-critical application. Which of the following metrics should the engineer define with business stakeholders to establish recovery requirements? (Select two).

A. Recovery time objective

B. Recovery point objective

C. Service-level indicator

D. Mean time between failures

E. Configuration baseline

ANSWER: A B

Explanation:

Recovery time objective (RTO) defines the maximum acceptable amount of time an application can remain unavailable after a disruption. It helps determine the recovery architecture, automation, and level of standby capacity needed to restore the service.

Recovery point objective (RPO) defines the maximum acceptable amount of data loss measured in time. It determines how frequently data must be backed up, replicated, or captured through transaction logs. Together, RTO and RPO establish the availability and data-protection targets used to design disaster recovery capabilities. NIST discusses contingency planning and recovery objectives in SP 800-34 Rev. 1.

QUESTION NO: 30

A cloud engineer was deploying the company's payment processing application, but it failed with the

following error log:

ERFOR:root: Transaction failed http 429 response, please try again Which of the following are the most likely causes for this error? (Select two).

A. API throttling

B. API gateway outage

C. Web server outage

D. Oversubscription

E. Unauthorized access

F. Insufficient quota

ANSWER: A F

Explanation:

An HTTP 429 response means *Too Many Requests*: the receiving service is deliberately refusing requests because the client has exceeded an allowed request rate or usage limit. **API throttling** is therefore a likely cause. Cloud providers, API gateways, and payment-service APIs commonly apply throttling controls to protect shared infrastructure and downstream services. A deployment can trigger this condition when instances start simultaneously, health checks and retries multiply request volume, or the application submits transactions faster than the service's configured per-second or per-minute rate permits.

**Insufficient quota
** is also a likely cause because quotas establish a maximum permitted consumption level for an API, project, tenant, or account over a defined period. Once that allocation is consumed, further requests can be rejected with HTTP 429 until the quota window resets, capacity is released, or an administrator obtains a quota increase. The "please try again" wording aligns with a transient rate- or quota-enforcement response. The engineer should review API usage and quota metrics, inspect any rate-limit response headers such as `Retry-After`, and implement bounded exponential backoff with retries for appropriate idempotent operations. See MDN: HTTP 429 Too Many Requests and Google Cloud quota documentation.

QUESTION NO: 31

A security analyst reviews the daily logs and notices the following suspicious activity:

Host	NA/US/John Smith
IP	10.150.71.151
Activity	A powershell process executed compressed, encoded command line content.

The analyst investigates the firewall logs and identities the following:

Operating system	Kali Linux
CPU	x64
Filesystem	ext4
User	John Smith
Category	Compromised - Unauthorized Access
Domain	NA/US
IP	201.101.25.121 (External)
Port	4444
Connection type	Inbound Connection

Which of the following steps should the security analyst take next to resolve this issue? (Select two).

A. Submit an IT support ticket and request Kali Linux be uninstalled from John Smith's computer

B. Block all inbound connections on port 4444 and block the IP address 201.101.25.121.

C. Contact John Smith and request the Ethernet cable attached to the desktop be unplugged

D. Check the running processes to confirm if a backdoor connection has been established.

E. Upgrade the Windows x64 operating system on John Smith's computer to the latest version.

F. Block all outbound connections from the IP address 10.150.71.151.

ANSWER: B D

Explanation:

“Block all inbound connections on port 4444 and block the IP address 201.101.25.121.” is an appropriate immediate containment action because the observed external source and port represent the suspected command-and-control or remote-access path. Restricting that specific source and service at the firewall limits further inbound communication while preserving the scope needed for investigation. This is consistent with incident-response containment practices, which prioritize stopping an attacker’s active access and preventing additional unauthorized activity.

“Check the running processes to confirm if a backdoor connection has been established.” is also required to validate the host-level impact indicated by the suspicious encoded PowerShell activity. Process inspection can identify an active PowerShell child process, remote shell, unexpected executable, or network-connected process associated with the suspected compromise. Confirming active execution gives responders the evidence needed to scope the incident, collect relevant artifacts, and proceed with eradication and recovery actions. NIST describes containment and analysis as core incident-handling activities in SP 800-61 Rev. 2. Encoded PowerShell is also a recognized execution technique documented by MITRE ATT&CK.

QUESTION NO: 32

A project manager is taking notes along with gathering feedback about a prior project in order to consider known issues with a project enhancement. Which of the following best describes the plan that the project team is working on?

A. A plan for capturing issues and conflicts in a project's knowledge base

B. A plan for avoiding potential issues that may arise during project execution

C. A plan for documenting the outcomes of the project

D. A plan for addressing and resolving identified issues once the project has been completed

ANSWER: B

Explanation:

A plan for avoiding potential issues that may arise during project execution is correct because the project manager is using information from a previous effort to proactively prepare the enhancement project. Notes, feedback, and known issues provide valuable inputs for identifying risks: events or conditions that could affect project objectives if they recur. The team can assess those risks, determine their likely impact, assign ownership, and define response actions or controls before implementation begins. This is a preventive planning activity rather than merely a retrospective record of what happened.

For cloud projects, this approach supports reliable delivery by incorporating operational experience into planning for changes, migrations, upgrades, and other enhancements. Known prior problems may reveal technical dependencies, capacity constraints, security concerns, configuration weaknesses, or scheduling risks. Planning responses in advance reduces the chance that the same conditions will disrupt execution and helps the team make informed decisions if warning signs appear. CompTIA's Cloud+ objectives include applying project-management concepts in cloud environments, where risk identification and mitigation are essential. PMI similarly describes risk management as planning and carrying out responses to project uncertainty. See the CompTIA exam objectives resource and PMI's risk management guidance.

QUESTION NO: 33

Which of the following requirements are core considerations when migrating a small business's onpremises applications to the cloud? (Select two).

A. Availability

B. Hybrid

C. Testing

D. Networking

E. Compute

F. Logs

ANSWER: A D

Explanation:

Availability and **Networking** are fundamental requirements for a successful migration of on-premises applications to cloud services. Availability ensures that applications deliver the required uptime and resilience after migration. The organization must identify business availability requirements, select an appropriate service-level agreement, and design for failures where necessary through capabilities such as redundant components, backups, recovery procedures, and multi-zone or multi-region deployment. These decisions should reflect the application's business criticality and its acceptable downtime.

Networking is equally essential because users, systems, and administrators must be able to reach the migrated application securely and with sufficient performance. Migration planning should account for internet or private connectivity, bandwidth and latency requirements, IP addressing, routing, DNS, firewall rules, segmentation, and secure remote access. Network design also affects how on-premises dependencies communicate with cloud-hosted workloads during a phased migration. Addressing availability and networking early establishes the service reliability and connectivity foundation on which the migration and subsequent cloud operations depend. Guidance on migration planning is available from the Microsoft Cloud Adoption Framework and the AWS Migration Strategy guide.

QUESTION NO: 34

A security analyst reviews the daily logs and notices the following suspicious activity:

Host	NA/US/John Smith
IP	10.150.71.151
Activity	A powershell process executed compressed, encoded command line content.

The analyst investigates the firewall logs and identities the following:

Operating system	Kali Linux
CPU	x64
Filesystem	ext4
User	John Smith
Category	Compromised - Unauthorized Access
Domain	NA/US
IP	201.101.25.121 (External)
Port	4444
Connection type	Inbound Connection

Which of the following steps should the security analyst take next to resolve this issue? (Select two).

A. Submit an IT support ticket and request Kali Linux be uninstalled from John Smith's computer

B. Block all inbound connections on port 4444 and block the IP address 201.101.25.121.

C. Contact John Smith and request the Ethernet cable attached to the desktop be unplugged

D. Check the running processes to confirm if a backdoor connection has been established.

E. Upgrade the Windows x64 operating system on John Smith's computer to the latest version.

F. Block all outbound connections from the IP address 10.150.71.151.

ANSWER: B D

Explanation:

“Block all inbound connections on port 4444 and block the IP address 201.101.25.121.” is an appropriate immediate containment action. Port 4444 is frequently used by penetration-testing frameworks and reverse-shell payloads, so denying traffic on that port and blocking the identified external source limits the suspected command-and-control or unauthorized remote-access path. This reduces the opportunity for an attacker to maintain or re-establish access while the incident is investigated.

“Check the running processes to confirm if a backdoor connection has been established.” is also necessary to validate the endpoint impact and identify active malicious activity. Reviewing processes and their associated network connections can reveal an unauthorized listener, reverse shell, remote-access tool, or persistence-related process. This evidence supports accurate scoping and enables the analyst to remove the malicious process or artifact after containment. These actions align with incident-response practice: contain the suspected activity first, then analyze the affected system to determine the nature and extent of compromise. NIST describes containment, eradication, and recovery as core incident-handling activities in NIST SP 800-61 Rev. 2. Microsoft also documents using process and network-connection information during endpoint investigation in its TCPView documentation.

QUESTION NO: 35

A customer's facility is located in an area where natural disasters happen frequently. The customer requires the following:

â€¢ Data resiliency due to exposure to frequent natural disasters

â€¢ Data localization because of privacy regulations in the country

â€¢ High availability

Which of the following cloud resources should be provisioned to meet these requirements?

A. Storage in a separate data center located in same region

B. An on-premises private cloud carrying duplicate data

C. Storage in an availability zone in a different region within the same country

D. Storage in the same availability zone as the primary data

ANSWER: C

Explanation:

Storage in an availability zone in a different region within the same country is the best fit because it provides geographic separation from the primary region while keeping the replicated data within the jurisdiction required by the customer's privacy regulations. A disruption caused by a natural disaster can affect an entire availability zone and, in severe cases, services across a regional area. Placing a copy in another in-country region establishes a more resilient failure domain and supports disaster-recovery planning. The use of an availability zone in that secondary region also supports high availability, because cloud availability zones are designed as physically separate locations with independent infrastructure and redundant connectivity. The solution therefore combines geographic resilience with jurisdictional data residency: the data remains in-country, but is not dependent on the same regional infrastructure as the primary workload. In practice, the organization would configure replication and failover between the primary and secondary regional storage locations, then validate recovery objectives such as RPO and RTO. Cloud providers describe regions as geographic areas and availability zones as isolated locations within those areas; see AWS Regions and Availability Zones. For data-residency planning, see Microsoft Azure data residency.

QUESTION NO: 36

A cloud networking engineer is troubleshooting the corporate office ' s network configuration. Employees in the IT and operations departments are unable to resolve IP addresses on all devices, and the IT department cannot establish a connection to other departments ' subnets. The engineer identifies the following configuration currently in place to support the office network:

Subnet	Department	Employees
10.1.20.1/24	Finance	50
10.1.30.1/24	IT	90
10.1.40.1/24	Legal	30
10.1.50.1/24	Operations	100

Each employee needs to connect to the network with a maximum of three hosts. Each subnet must be segregated, but the IT department must have the ability to communicate with all subnets. Which of the following meet the IP addressing and routing requirements? (Select two).

A. Modifying the subnet mask to 255 255 254.0 for IT and operations departments

B. Configuring static routing to allow access from each subnet to 10.1.40.1

C. Modifying the BYOD policy to reduce the volume of devices that are allowed to connect to the corporate network

D. Configuring static routing to allow access from 10.1.30.1 to each subnet

E. Combining the subnets and increasing the allocation of IP addresses available to support three hosts for each employee

F. Modifying the subnet mask to 255.255.255.128 for the IT and operations departments

ANSWER: D F

Explanation:

Configuring static routing to allow access from 10.1.30.1 to each subnet provides the required Layer 3 reachability for the IT network. The address 10.1.30.1 represents the IT subnet's routed interface or default gateway, and routes for the other departmental networks allow traffic originating in IT to reach those otherwise separate network segments. This preserves subnet separation while enabling the explicitly required administrative access from IT. Appropriate return routing on the destination networks is also necessary in a deployed environment so that replies can return to the IT subnet. Static routes are a valid solution where the office topology is small and its network paths are stable. Cisco's overview of static routing describes how routes define paths to remote networks.

Modifying the subnet mask to 255.255.255.128 for the IT and operations departments creates /25 network boundaries. A /25 supplies 126 usable IPv4 host addresses per subnet, allowing capacity for employees' permitted devices while retaining two distinct, nonoverlapping network segments within the relevant address range. Using a subnet mask to define network and host portions is the mechanism that establishes the required addressing boundaries. CompTIA lists network addressing, routing, and connectivity implementation among the Cloud+ networking skills measured by the certification; see the CompTIA Cloud+ certification page.

QUESTION NO: 37

A cloud networking engineer is troubleshooting the corporate office's network configuration. Employees in the IT and operations departments are unable to resolve IP addresses on all devices, and the IT department cannot establish a connection to other departments' subnets. The engineer identifies the following configuration currently in place to support the office network:

Subnet	Department	Employees
10.1.20.1/24	Finance	50
10.1.30.1/24	IT	90
10.1.40.1/24	Legal	30
10.1.50.1/24	Operations	100

Each employee needs to connect to the network with a maximum of three hosts. Each subnet must be segregated, but the IT department must have the ability to communicate with all subnets. Which of the following meet the IP addressing and routing requirements? (Select two).

A. Modifying the subnet mask to 255.255.254.0 for IT and operations departments

B. Configuring static routing to allow access from each subnet to 10.1.40.1

C. Modifying the BYOD policy to reduce the volume of devices that are allowed to connect to the corporate network

D. Configuring static routing to allow access from 10.1.30.1 to each subnet

E. Combining the subnets and increasing the allocation of IP addresses available to support three hosts for each employee

F. Modifying the subnet mask to 255.255.255.128 for the IT and operations departments

G. Modifying the subnet mask to 255.255.254.0 (/23) for IT and Operations to support up to ~510 usable IPs per subnet

ANSWER: A D G

Explanation:

Modifying the subnet mask to 255.255.254.0 for IT and operations departments provides a /23 network for each affected department. A /23 has 512 total IPv4 addresses and 510 usable host addresses. Because IT has approximately 90 employees and Operations has approximately 100 employees, allowing as many as three connected hosts per employee requires capacity for roughly 270 and 300 hosts, respectively. The existing /24 allocation supplies only 254 usable addresses, so expanding those departments to /23 addressing is necessary. The equivalent statement, “Modifying the subnet mask to 255.255.254.0 (/23) for IT and Operations to support up to ~510 usable IPs per subnet,” describes the same valid sizing solution.

Configuring static routing to allow access from 10.1.30.1 to each subnet enables the IT subnet’s Layer 3 gateway to route traffic to the separate departmental networks. This preserves subnet segregation while providing the required IT-to-department connectivity. In an operational deployment, corresponding return routing must also be available, whether through static routes, a default gateway, or a dynamic routing design. CIDR prefix sizing and routed subnet connectivity are core network-design considerations covered by the CompTIA exam objectives and by RFC 4632.

QUESTION NO: 38

During a phase review, two stakeholders discuss the approval of a deliverable. The project manager convinces the stakeholders to agree on a common solution. Which of the following best describes what the project manager did to gain approval from both stakeholders?

A. Force

B. Smooth

C. Compromise

D. Avoid

ANSWER: C

Explanation:

Compromise is the correct description because the project manager brought stakeholders with differing positions to an agreement that both could accept. In a compromise, each party makes concessions on some preferences or demands in exchange for progress toward a mutually acceptable outcome. That approach is particularly useful during a phase review, where approval of a deliverable may be needed to maintain project momentum and proceed to the next stage.

The key detail is that the stakeholders agreed on a *common solution*. This indicates a negotiated resolution rather than a unilateral decision. The project manager facilitated agreement by balancing stakeholder interests and helping them settle on terms that enable approval. In cloud projects, this can apply to decisions about scope, implementation timing, performance expectations, security controls, or acceptance criteria. Compromise provides a practical path to documented approval when stakeholders cannot each receive every requested outcome.

CompTIA Cloud+ emphasizes the skills required to manage cloud projects, communicate with stakeholders, and support operational decision-making throughout the cloud lifecycle. See the CompTIA Cloud+ certification overview and CompTIA’s exam objectives resources for the current certification scope.