CompTIA

Exam Questions CV0-003

CompTIA Cloud+ Certification Exam

NEW QUESTION 1
- (Topic 1)
A cloud administrator is reviewing the authentication and authorization mechanism implemented within the cloud environment. Upon review, the administrator discovers the sales group is part of the finance group, and the sales team members can access the financial application. Single sign-on is also implemented, which makes access much easier.
Which of the following access control rules should be changed?

A. Discretionary-based
B. Attribute-based
C. Mandatory-based
D. Role-based

Answer: D

Explanation:
Role-based access control (RBAC) is a type of access control model that assigns permissions and privileges to users based on their roles or functions within an organization or system. RBAC can help simplify and streamline the management and enforcement of access policies, as it can reduce the complexity and redundancy of assigning permissions to individual users or groups. RBAC can also help improve security and compliance, as it can limit or grant access based on the principle of least privilege and the separation of duties. RBAC is the best access control rule to change when the sales group is part of the finance group and the sales team members can access the financial application due to a single sign-on mechanism being implem
Reference: https://www.ekransystem.com/en/blog/rbac-vs-abac

NEW QUESTION 2
SIMULATION - (Topic 1)
The QA team is testing a newly implemented clinical trial management (CTM) SaaS application that uses a business intelligence application for reporting. The UAT users were instructed to use HTTP and HTTPS.
Refer to the application dataflow:
1A – The end user accesses the application through a web browser to enter and view clinical data.
2A – The CTM application server reads/writes data to/from the database server.
1B – The end user accesses the application through a web browser to run reports on clinical data.
2B – The CTM application server makes a SOAP call on a non-privileged port to the BI application server.
3B – The BI application server gets the data from the database server and presents it to the CTM application server.
When UAT users try to access the application using https://ctm.app.com or http://ctm.app.com, they get a message stating: “Browser cannot display the webpage.” The QA team has raised a ticket to troubleshoot the issue.
INSTRUCTIONS
You are a cloud engineer who is tasked with reviewing the firewall rules as well as virtual network settings.
You should ensure the firewall rules are allowing only the traffic based on the dataflow. You have already verified the external DNS resolution and NAT are working.
Verify and appropriately configure the VLAN assignments and ACLs. Drag and drop the appropriate VLANs to each tier from the VLAN Tags table. Click on each Firewall to change ACLs as needed.
If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

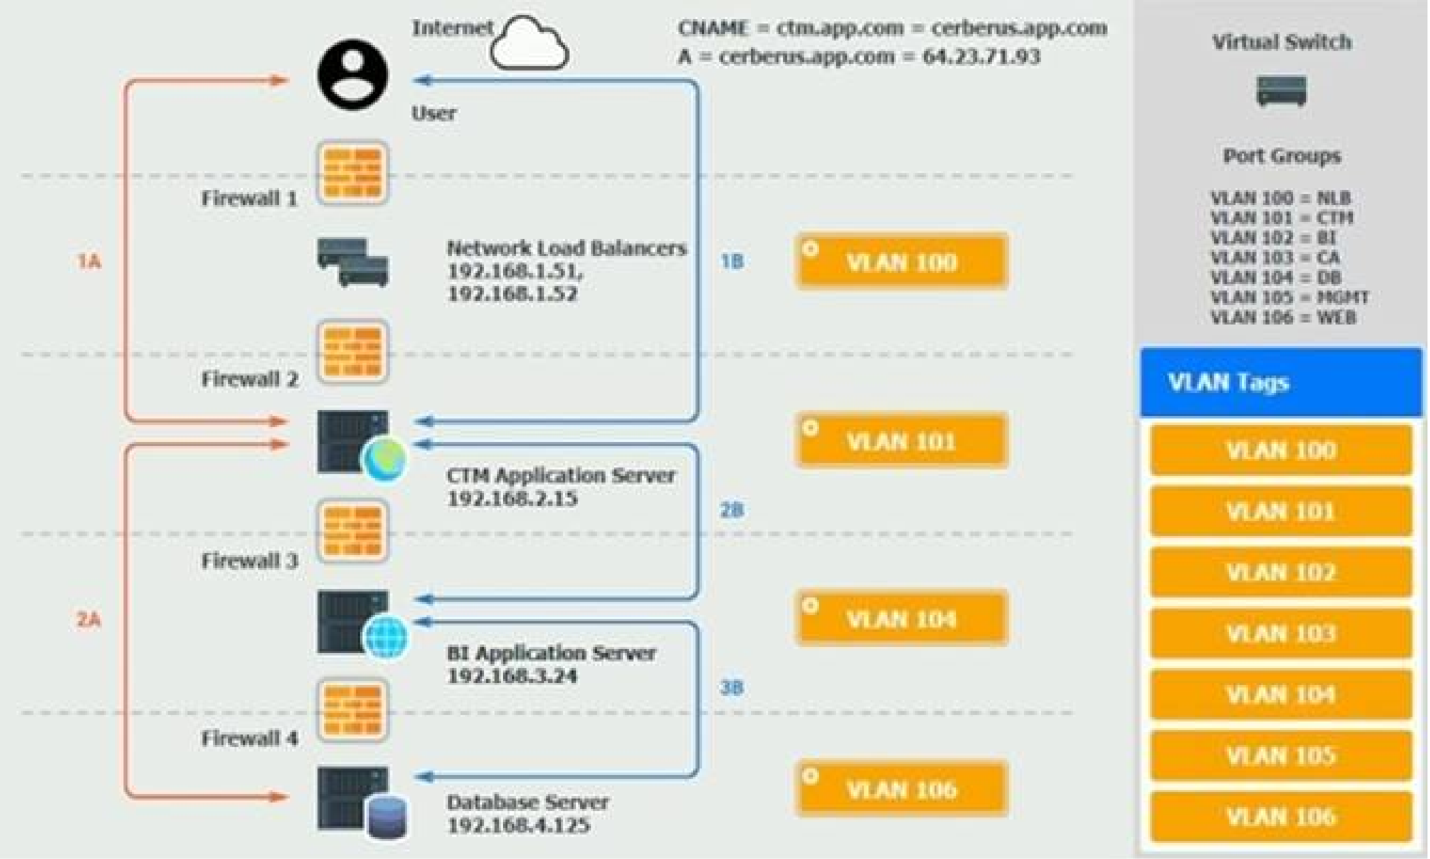

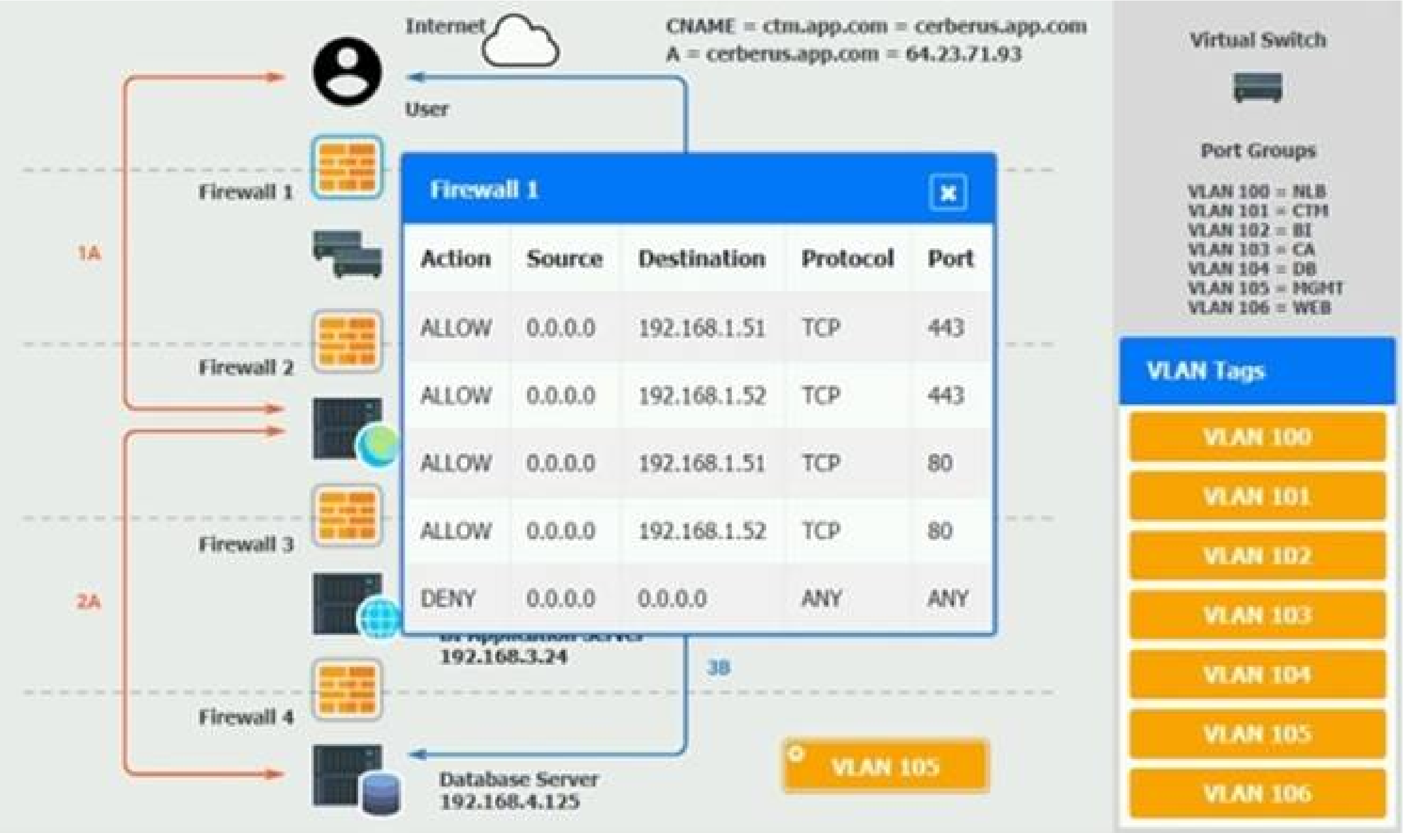

CNAME = ctm.app.com = cerberus.app.com
A = cerberus.app.com = 64.23.71.93

Port Groups
VLAN 100 = NLB
VLAN 101 = CTM
VLAN 102 = BI
VLAN 103 = CA
VLAN 104 = DB
VLAN 105 = MGMT
VLAN 106 = WEB

Firewall 1

Action	Source	Destination	Protocol	Port
ALLOW	0.0.0.0	192.168.1.51	TCP	443
ALLOW	0.0.0.0	192.168.1.52	TCP	443
ALLOW	0.0.0.0	192.168.1.51	TCP	80
ALLOW	0.0.0.0	192.168.1.52	TCP	80
DENY	0.0.0.0	0.0.0.0	ANY	ANY

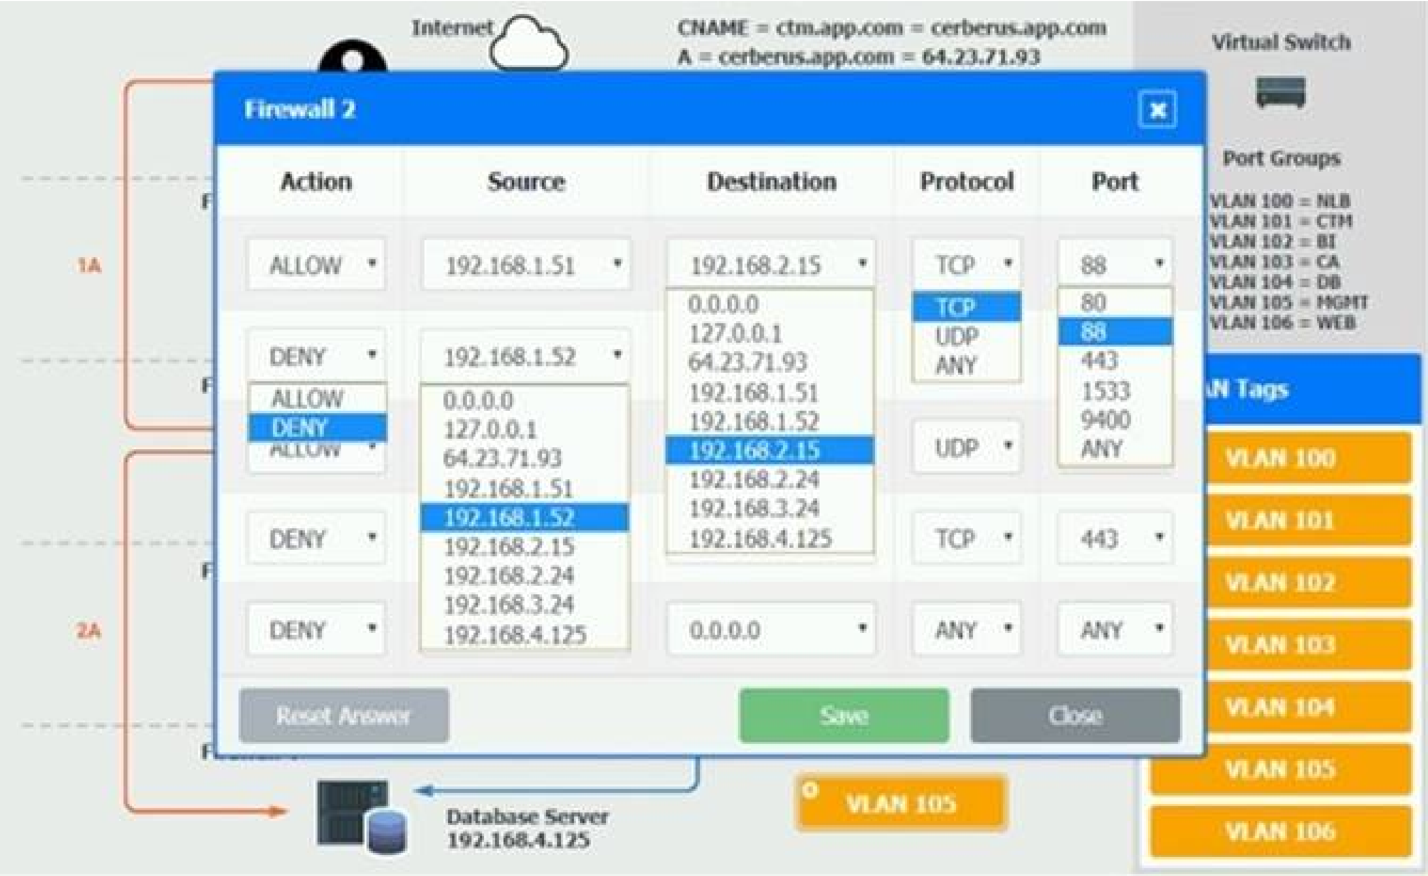

CNAME = ctm.app.com = cerberus.app.com
A = cerberus.app.com = 64.23.71.93

Port Groups
VLAN 100 = NLB
VLAN 101 = CTM
VLAN 102 = BI
VLAN 103 = CA
VLAN 104 = DB
VLAN 105 = MGMT
VLAN 106 = WEB

Firewall 2

Action	Source	Destination	Protocol	Port
ALLOW	192.168.1.51	192.168.2.15	TCP	88
DENY	192.168.1.52			
ALLOW			UDP	
DENY			TCP	443
DENY		0.0.0.0	ANY	ANY

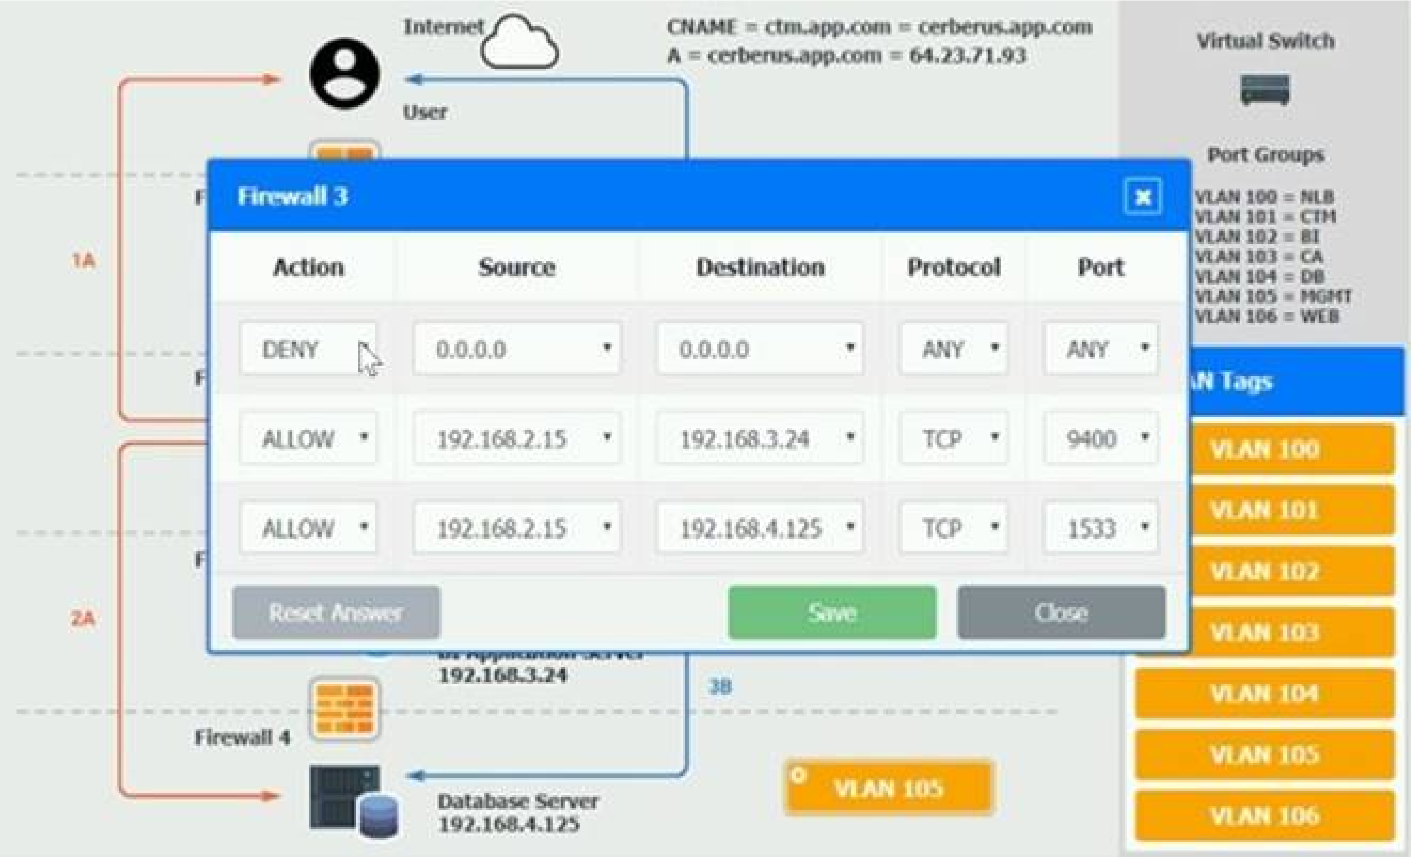

CNAME = ctm.app.com = cerberus.app.com
A = cerberus.app.com = 64.23.71.93

Port Groups
VLAN 100 = NLB
VLAN 101 = CTM
VLAN 102 = BI
VLAN 103 = CA
VLAN 104 = DB
VLAN 105 = MGMT
VLAN 106 = WEB

Firewall 3

Action	Source	Destination	Protocol	Port
DENY	0.0.0.0	0.0.0.0	ANY	ANY
ALLOW	192.168.2.15	192.168.3.24	TCP	9400
ALLOW	192.168.2.15	192.168.4.125	TCP	1533

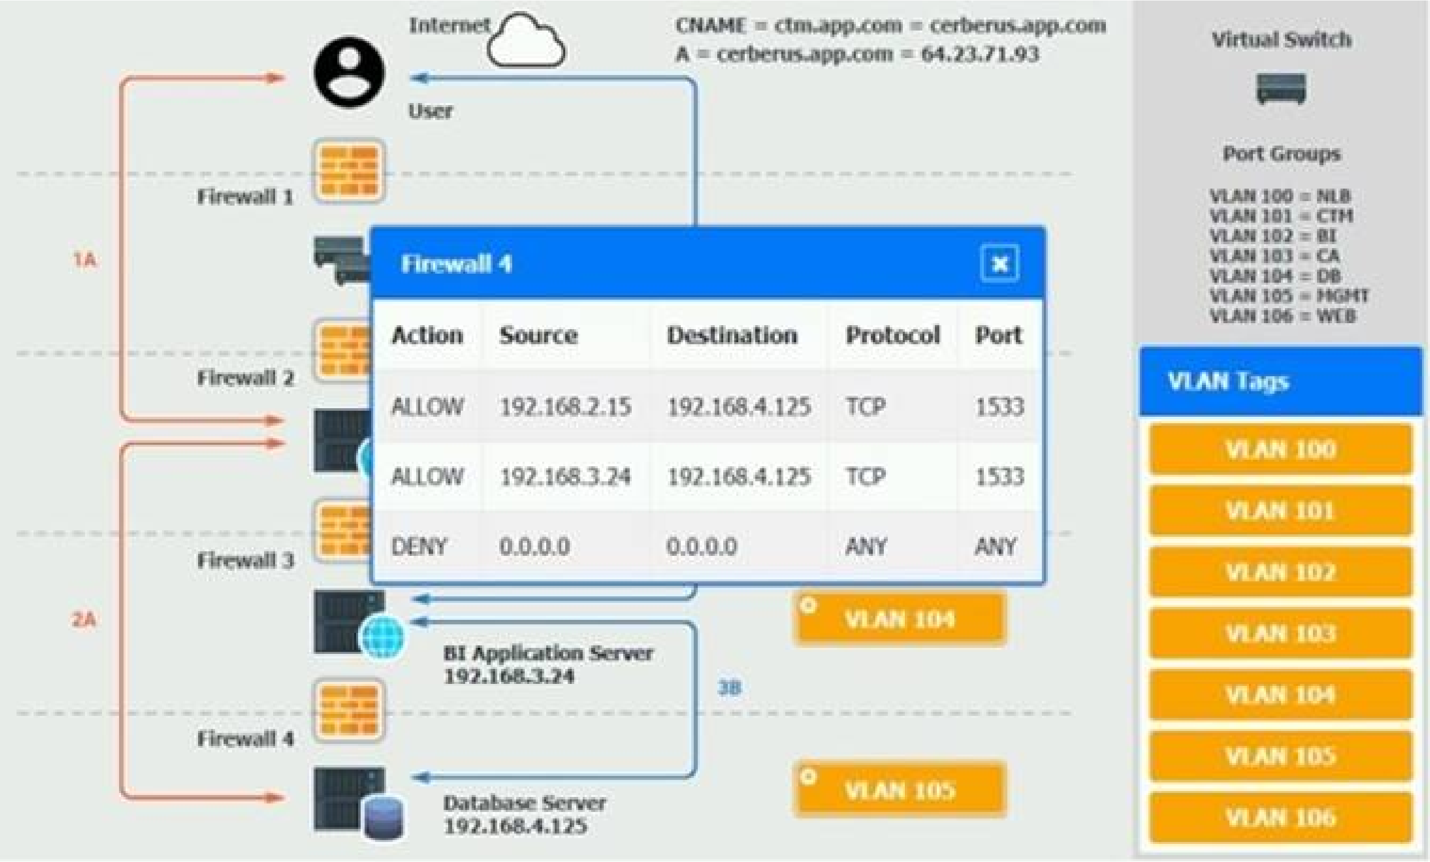

Firewall 4

Action	Source	Destination	Protocol	Port
ALLOW	192.168.2.15	192.168.4.125	TCP	1533
ALLOW	192.168.3.24	192.168.4.125	TCP	1533
DENY	0.0.0.0	0.0.0.0	ANY	ANY

A. Mastered
B. Not Mastered

Answer: A

Explanation:
On firewall 3, change the DENY 0.0.0.0 entry to rule 3 not rule 1.

NEW QUESTION 3
- (Topic 1)
Due to a policy change, a few of a customer's application VMs have been migrated to synchronously replicated storage. The customer now reports that performance is lower. The systems administrator checks the resource usage and discovers CPU utilization is at 60% and available memory is at 30%.
Which of the following is the MOST likely cause?

A. There is not enough vCPU assigned
B. The application is not compatible with the new settings
C. The new configuration is adding latency
D. The memory of the VM is underallocated

Answer: C

Explanation:
Latency is the delay or time taken for data to travel from one point to another in a network or system. Latency can affect the performance of applications and processes that depend on fast and reliable data transfer. Synchronous replication is a method of data replication that ensures that data is written to two or more storage devices at the same time, providing high availability and consistency. However, synchronous replication can also introduce latency, as the write operation has to wait for the confirmation from all the replicated devices before completing. The new configuration of migrating some application VMs to synchronously replicated storage is most likely adding latency, which can lower the performance of the applications. References: [CompTIA Cloud+ Certification Exam Objectives], page 10, section 1.5

NEW QUESTION 4
- (Topic 1)
An administrator is performing an in-place upgrade on a quest VM operating system.
Which of the following can be performed as a quick method to roll back to an earlier state, if necessary?

A. A configuration file backup
B. A full backup of the database
C. A differential backup
D. A VM-level snapshot

Answer: D

Explanation:
A VM-level snapshot is a point-in-time copy of the state and data of a virtual machine (VM). A VM-level snapshot can be used as a quick method to roll back to an earlier state, if necessary, as it can restore the VM to the exact condition it was in when the snapshot was taken. A VM-level snapshot can be useful for performing an in-place upgrade
on a guest VM operating system, as it can allow the administrator to revert to the previous operating system version in case of any issues or errors. References: CompTIA Cloud+ Certification Exam Objectives, page 10, section 1.5
Reference: https://cloud.google.com/compute/docs/tutorials/performing-in-place-upgrade- windows-server

NEW QUESTION 5
- (Topic 1)
A systems administrator is creating a playbook to run tasks against a server on a set schedule.
Which of the following authentication techniques should the systems administrator use within the playbook?

A. Use the server’s root credentials
B. Hard-code the password within the playbook
C. Create a service account on the server
D. Use the administrator’s SSO credentials

Answer: C

Explanation:
A service account is a type of user account that is created for a specific service or application to run on a server or system. Creating a service account on the server is the best authentication technique to use within the playbook to run tasks against the server on a set schedule, as it can provide secure and consistent access to the server without exposing or hard-coding any sensitive credentials within the playbook. Creating a service account can also help manage and monitor the tasks and activities performed by the service or application on the server. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 6
- (Topic 1)
An organization will be deploying a web application in a public cloud with two web servers, two database servers, and a load balancer that is accessible over a single public IP.
Taking into account the gateway for this subnet and the potential to add two more web servers, which of the following will meet the minimum IP requirement?

A. 192.168.1.0/26
B. 192.168.1.0/27
C. 192.168.1.0/28
D. 192.168.1.0/29

Answer: C

Explanation:
A /28 subnet is a subnet that has a network prefix of 28 bits and a host prefix of 4 bits. A /28 subnet can support up to 16 hosts (14 usable hosts) and has a subnet mask of 255.255.255.240. Using a /28 subnet can meet the minimum IP requirement for deploying a web application in a public cloud with two web servers, two database servers, and a load balancer that is accessible over a single public IP, taking into account the gateway for this subnet and the potential to add two more web servers. Using a /28 subnet can provide enough host addresses for the current and future web servers, database servers, load balancer, and gateway, as well as allow for some growth or redundancy.
References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8

NEW QUESTION 7
- (Topic 1)
A systems administrator is troubleshooting performance issues with a Windows VDI environment. Users have reported that VDI performance is very slow at the start of the workday, but the performance is fine during the rest of the day. Which of the following is the MOST likely cause of the issue? (Choose two.)

A. Disk I/O limits
B. Affinity rule
C. CPU oversubscription
D. RAM usage
E. Insufficient GPU resources
F. License issues

Answer: AC

Explanation:
Disk I/O limits are restrictions or controls that limit the amount of disk input/output operations per second (IOPS) that a VM can perform on a storage device or system. CPU oversubscription is a situation where more CPU resources are allocated to VMs than are physically available on the host or server. Disk I/O limits and CPU oversubscription are most likely to cause VDI performance being very slow at the start of the workday, but fine during the rest of the day, as they can create bottlenecks or contention for disk and CPU resources when multiple users log in or launch their VDI sessions at the same time, resulting in increased latency or reduced throughput for VDI operations. References: CompTIA Cloud+ Certification Exam Objectives, page 9, section 1.4

NEW QUESTION 8
- (Topic 2)
A cloud architect is reviewing four deployment options for a new application that will be hosted by a public cloud provider. The application must meet an SLA that allows for no
more than five hours of downtime annually. The cloud architect is reviewing the SLAs for the services each option will use:

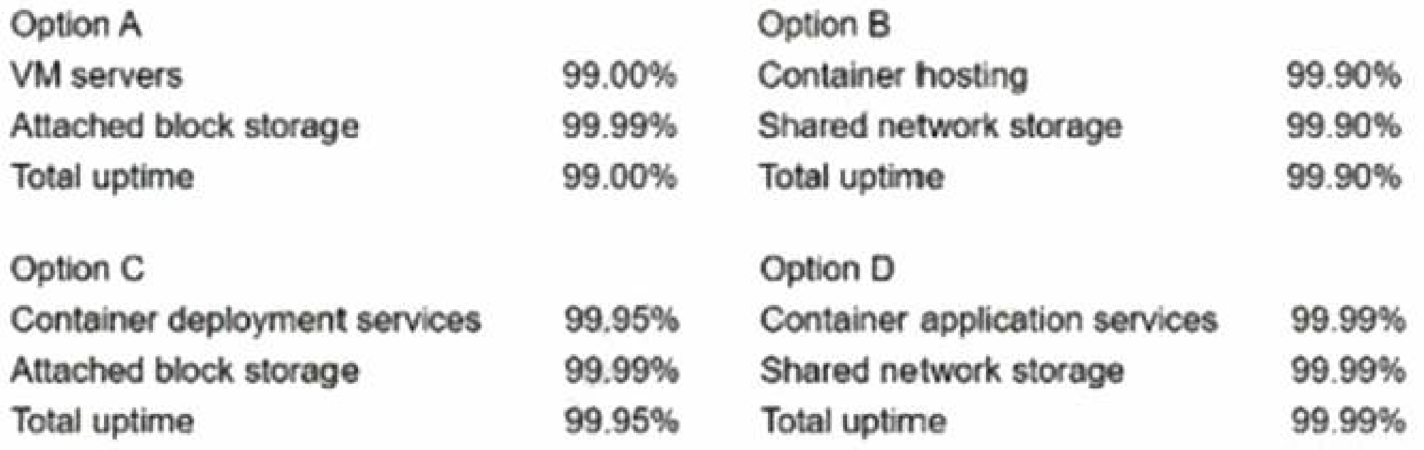

Option A		Option B	
VM servers	99.00%	Container hosting	99.90%
Attached block storage	99.99%	Shared network storage	99.90%
Total uptime	99.00%	Total uptime	99.90%
Option C		Option D	
Container deployment services	99.95%	Container application services	99.99%
Attached block storage	99.99%	Shared network storage	99.99%
Total uptime	99.95%	Total uptime	99.99%

Based on the information above, which of the following minimally complies with the SLA requirements?

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

Explanation:
Option B is what minimally complies with the SLA (Service Level Agreement) requirements of allowing for no more than five hours of downtime annually for a new application that will be hosted by a public cloud provider. An SLA is a contract or agreement that defines the level of service or performance that a customer expects from a provider, such as availability, reliability, scalability, security, etc. An SLA can help to measure and monitor the quality and satisfaction of service or performance, as well as identify any penalties or rewards for meeting or failing to meet the SLA. Option B minimally complies with the SLA requirements by using services that have availability percentages that are equal to or higher than 99.95%, which translates to no more than five hours of downtime annually. Option B uses services such as:
? Compute: This is a service that provides computing resources such as servers, processors, memory, etc., to run applications or functions. Option B uses compute service with availability percentage of 99.95%, which means that it guarantees to be available for 99.95% of the time in a year, and allows for no more than five hours of downtime in a year.
? Storage: This is a service that provides storage resources such as disks, volumes, files, etc., to store data or information. Option B uses storage service with availability percentage of 99.99%, which means that it guarantees to be available for 99.99% of the time in a year, and allows for no more than one hour of downtime in a year.
? Database: This is a service that provides database resources such as tables, records, queries, etc., to store and retrieve data or information. Option B uses database service with availability percentage of 99.95%, which means that it guarantees to be available for 99.95% of the time in a year, and allows for no more than five hours of downtime in a year.

NEW QUESTION 9
- (Topic 2)
A system administrator has provisioned a new web server. Which of the following, in combination, form the best practice to secure the server's OS? (Choose three.)

A. Install TLS certificates on the server.
B. Forward port 80 traffic to port 443.
C. Disable TLS 1.0/1.1 and SSL.
D. Disable password authentication.
E. Enable SSH key access only.
F. Provision the server in a separate VPC.
G. Disable the superuser/administrator account.
H. Restrict access on port 22 to the IP address of the administrator's workstation.

Answer: ADE

Explanation:
These are the best practices to secure the OS of a new web server that has been provisioned in a cloud environment:
? Install TLS certificates on the server: TLS (Transport Layer Security) certificates are digital documents that contain information such as identity, public key, expiration date, etc., that can be used to prove one’s identity and establish secure communication over a network. Installing TLS certificates on the web server can encrypt and secure web traffic between the server and the clients, as well as prevent spoofing or impersonation attacks.
? Disable password authentication: Password authentication is a method of verifying and authenticating users or devices based on passwords or other credentials. Password authentication can be insecure or vulnerable to attacks such as brute force, dictionary, phishing, etc., especially if passwords are weak, reused, or compromised. Disabling password authentication can enhance security by preventing unauthorized or malicious access to the web server using passwords.
? Enable SSH key access only: SSH key access is a method of verifying and authenticating users or devices based on digital keys issued by a trusted authority. SSH key access can provide more security and convenience than password authentication, as it does not require users or devices to remember or enter passwords every time they access the web server. Enabling SSH key access only can ensure that only authorized or trusted users or devices can access the web server using keys.

NEW QUESTION 10
- (Topic 2)
Which of the following definitions of serverless computing BEST explains how it is different from using VMs?

A. Serverless computing is a cloud-hosting service that utilizes infrastructure that is fully managed by the CSP.
B. Serverless computing uses predictable billing and offers lower costs than VM compute services.
C. Serverless computing is a scalable, highly available cloud service that uses SDN technologies.
D. Serverless computing allows developers to focus on writing code and organizations to focus on business.

Answer: D

Explanation:
This is the best definition of serverless computing that explains how it is different from using VMs (Virtual Machines). Serverless computing is a cloud service model that provides customers with a platform to run applications or functions without having to manage or provision any underlying infrastructure or resources, such as servers, storage, network, OS, etc. Serverless computing is different from using VMs in the following ways:
? Serverless computing allows developers to focus on writing code and organizations to focus on business, rather than spending time and effort on managing or scaling VMs or other infrastructure components.
? Serverless computing is event-driven and pay-per-use, which means that applications or functions are executed only when triggered by a specific event or request, and customers are charged only for the resources consumed during the execution time.
? Serverless computing is more scalable and flexible than using VMs, as it can automatically adjust the capacity and performance of applications or functions according to demand or workload, without requiring any manual intervention or configuration.

NEW QUESTION 10
- (Topic 2)
A cloud solutions architect needs to determine the best strategy to deploy an application environment in production, given the following requirements:
No downtime
Instant switch to a new version using traffic control for all users
Which of the following deployment strategies would be the BEST solution?

A. Hot site
B. Blue-green
C. Canary

D. Rolling

Answer: B

Explanation:
Reference: https://thenewstack.io/deployment-strategies/
Blue-green is the best deployment strategy to deploy an application environment in production, given the requirements of no downtime and instant switch to a new version using traffic control for all users. Blue-green is a deployment strategy that involves having two identical environments, one running the current version of the application (blue) and one running the new version of the application (green). The traffic is directed to the blue environment by default, while the green environment is tested and verified. When the new version is ready to go live, the traffic is switched to the green environment using a router or load balancer, without any downtime or interruption. The blue environment can be kept as a backup or updated with the new version for future deployments.

NEW QUESTION 14
- (Topic 2)
A cloud administrator is building a new VM for machine-learning training. The developer requesting the VM has stated that the machine will need a full GPU dedicated to it.
Which of the following configuration options would BEST meet this requirement?

A. Virtual GPU
B. External GPU
C. Passthrough GPU
D. Shared GPU

Answer: C

Explanation:
Reference: https://blogs.vmware.com/apps/2018/09/using-gpus-with-virtual-machines-on- vsphere-part-2-vmdirectpath-i-o.html
Passthrough GPU is a configuration option that allows a VM to access a physical GPU directly without any virtualization layer or sharing mechanism. This provides the VM with full and exclusive access to the GPU resources and performance. Passthrough GPU is suitable for applications that require intensive graphics processing or machine learning training.

NEW QUESTION 15
- (Topic 2)
A cloud administrator is reviewing the annual contracts for all hosted solutions. Upon review of the contract for the hosted mail solution, the administrator notes the monthly subscription rate has increased every year. The provider has been in place for ten years, and there is a large amount of data being hosted. Which of the following is a barrier to switching providers?

A. Service-level agreement
B. Vendor lock-in
C. Memorandum of understanding
D. Encrypted data

Answer: B

Explanation:
Vendor lock-in is a barrier to switching providers for a hosted mail solution that has increased its monthly subscription rate every year. Vendor lock-in is a situation where a customer becomes dependent on a vendor or provider for a product or service and faces difficulties or costs in switching to another vendor or provider. Vendor lock-in can occur due to various factors, such as proprietary technology, contractual obligations, data migration challenges, compatibility issues, etc. In this case, the customer may face vendor lock-in due to the large amount of data being hosted by the mail provider and the potential challenges or costs of transferring or migrating the data to another provider.

NEW QUESTION 16
- (Topic 2)
A private IaaS administrator is receiving reports that all newly provisioned Linux VMs are running an earlier version of the OS than they should be. The administrator reviews the automation scripts to troubleshoot the issue and determines the scripts ran successfully. Which of the following is the MOST likely cause of the issue?

A. API version incompatibility
B. Misconfigured script account
C. Wrong template selection
D. Incorrect provisioning script indentation

Answer: C

Explanation:
The wrong template selection is the most likely cause of the issue of newly provisioned Linux VMs running an earlier version of OS than they should be in a private IaaS environment. A template is a preconfigured image or blueprint of a VM that contains an OS, applications, settings, etc., that can be used to create new VMs quickly and consistently. A template may have different versions or updates depending on when it was created or modified. If a template is selected incorrectly or not updated properly, it may result in creating VMs with an older or different version of OS than expected.

NEW QUESTION 17
- (Topic 2)
A system administrator is migrating a bare-metal server to the cloud. Which of the following types of migration should the systems administrator perform to accomplish this task?

A. V2V
B. V2P
C. P2P
D. P2V

Answer: D

Explanation:
P2V (Physical to Virtual) is a type of migration that converts a physical server into a virtual machine (VM). P2V migration can help to move a bare-metal server to the cloud by creating an image of its disk and configuration and uploading it to a cloud platform that supports VM creation from custom images.

NEW QUESTION 21
- (Topic 2)
A systems administrator swapped a failed hard drive on a server with a RAID 5 array. During the RAID resynchronization, a second hard drive failed.
Which of the following actions will make the server fully operational?

A. Restart the RAID resynchronization process
B. Perform a P2V migration of the server
C. Swap the failed hard drive with a fresh one
D. Restore the server from backup

Answer: D

Explanation:
RAID 5 is a disk array configuration that uses parity to provide fault tolerance and data recovery. RAID 5 can tolerate the failure of one disk, but not two or more disks. If a second disk fails during the resynchronization process, the data on the RAID 5 array will be lost and unrecoverable. The only way to make the server fully operational is to restore the data from a backup source.

NEW QUESTION 23
- (Topic 2)
An organization is using multiple SaaS-based business applications, and the systems administrator is unable to monitor and control the use of these subscriptions. The administrator needs to implement a solution that will help the organization apply security policies and monitor each individual SaaS subscription. Which of the following should be deployed to achieve these requirements?

A. DLP
B. CASB
C. IPS
D. HIDS

Answer: B

Explanation:
CASB (Cloud Access Security Broker) is what should be deployed to monitor and control the use of multiple SaaS-based business applications in a cloud environment. SaaS (Software as a Service) is a cloud service model that provides customers with access to software applications hosted on remote servers over a network or internet connection. SaaS can provide customers with convenience, flexibility, and scalability, but it may also introduce security risks such as data breaches, leaks, losses, etc., especially if customers have multiple SaaS subscriptions from different providers. CASB is a tool or service that
acts as an intermediary between customers and SaaS providers. CASB can help to monitor and control the use of multiple SaaS subscriptions by providing features such as:
? Visibility: CASB can provide visibility into what SaaS applications are being used, by whom, when, where, how, etc., as well as identify any unauthorized or suspicious activities.
? Compliance: CASB can provide compliance with various laws, regulations, standards, policies, etc., that apply to SaaS applications and data, such as GDPR, HIPAA, PCI DSS, etc., as well as enforce them using rules or actions.
? Security: CASB can provide security for SaaS applications and data by detecting and preventing any threats or attacks, such as malware, phishing, ransomware, etc., as well as protecting them using encryption, authentication, authorization, etc.

NEW QUESTION 26
- (Topic 2)
A cloud administrator is assigned to establish a connection between the on-premises data center and the new CSP infrastructure. The connection between the two locations must be secure at all times and provide service for all users inside the organization. Low latency is also required to improve performance during data transfer operations. Which of the following would BEST meet these requirements?

A. A VPC peering configuration
B. An IPSec tunnel
C. An MPLS connection
D. A point-to-site VPN

Answer: B

Explanation:
An IPSec tunnel is what would best meet the requirements of establishing a connection between the on-premises data center and the new CSP infrastructure that is secure at all times and provides service for all users inside the organization with low latency. IPSec (Internet Protocol Security) is a protocol that encrypts and secures network traffic over IP networks. IPSec tunnel is a mode of IPSec that creates a virtual private network (VPN) tunnel between two endpoints, such as routers, firewalls, gateways, etc., and encrypts and secures all traffic that passes through it. An IPSec tunnel can meet the requirements by providing:
? Security: An IPSec tunnel can protect network traffic from interception, modification, spoofing, etc., by using encryption, authentication, integrity, etc., mechanisms.
? Service: An IPSec tunnel can provide service for all users inside the organization by allowing them to access and use network resources or services on both ends of the tunnel, regardless of their physical location.
? Low latency: An IPSec tunnel can provide low latency by reducing the number of hops or devices that network traffic has to pass through between the endpoints of the tunnel.

NEW QUESTION 31
- (Topic 2)
A systems administrator wants to verify the word "qwerty" has not been used as a password on any of the administrative web consoles in a network. Which of the following will achieve this goal?

A. A service availability scan
B. An agent-based vulnerability scan
C. A default and common credentialed scan
D. A network port scan

Answer: C

Explanation:
A default and common credentialed scan is what the administrator should use to verify the word “qwerty” has not been used as a password on any of the administrative web consoles in a network. A credentialed scan is a type of vulnerability scan that uses valid credentials or accounts to access and scan target systems or devices. A credentialed scan can provide more accurate and detailed results than a non- credentialed scan, as it can perform more actions and tests on target systems or devices. A default and common credentialed scan is a type of credentialed scan that uses default or common credentials or accounts, such as admin/admin, root/root, etc., to access and scan target systems or devices. A default and common credentialed scan can help to identify weak or insecure passwords on administrative web consoles, such as “qwerty”, and recommend stronger passwords.

NEW QUESTION 33
- (Topic 2)
Which of the following service models would be used for a database in the cloud?

A. PaaS
B. laaS
C. CaaS
D. SaaS

Answer: A

Explanation:
PaaS (Platform as a Service) is a cloud service model that provides a platform for developing, testing, deploying, and managing applications in the cloud. PaaS includes the underlying infrastructure (servers, storage, network, etc.) as well as the middleware, databases, tools, frameworks, and APIs that are required for application development and delivery. Examples of PaaS are AWS Elastic Beanstalk, Azure App Service, Google App Engine, etc.

NEW QUESTION 35
- (Topic 2)
A disaster situation has occurred, and the entire team needs to be informed about the situation. Which of the following documents will help the administrator find the details of the relevant team members for escalation?

A. Chain of custody
B. Root cause analysis
C. Playbook
D. Call tree

Answer: D

Explanation:
A call tree is what will help the administrator find the details of the relevant team members for escalation after a disaster situation has occurred and the entire team needs to be informed about the situation. A call tree is a document or diagram that shows the hierarchy or sequence of communication or notification among team members in case of an emergency or incident, such as a disaster situation. A call tree can help to find the details of the relevant team members for escalation by providing information such as:
? Name: This indicates who is involved in the communication or notification process, such as team members, managers, stakeholders, etc.
? Role: This indicates what is their function or responsibility in the communication or notification process, such as initiator, receiver, sender, etc.
? Contact: This indicates how they can be reached or contacted in the communication or notification process, such as phone number, email address, etc.

NEW QUESTION 39
- (Topic 2)
A Chief Information Security Officer (CISO) is evaluating the company’s security management program. The CISO needs to locate all the assets with identified deviations and mitigation measures. Which of the following would help the CISO with these requirements?

A. An SLA document
B. ADR plan
C. SOC procedures
D. A risk register

Answer: D

Explanation:
A risk register is a document that records all the identified risks, their causes, impacts, probabilities, mitigation measures, and status for a project or an organization. A risk register helps to manage and monitor risks throughout their lifecycle and ensure they are addressed appropriately. A risk register would help the CISO to locate all the assets with identified deviations and mitigation measures.

NEW QUESTION 41
- (Topic 2)
A system administrator supports an application in the cloud, which includes a restful API that receives an encrypted message that is passed to a calculator system. The administrator needs to ensure the proper function of the API using a new automation tool. Which of the following techniques would be BEST for the administrator to use to accomplish this requirement?

A. Functional testing
B. Performance testing
C. Integration testing
D. Unit testing

Answer: C

Explanation:
Integration testing is the best technique to use to ensure the proper function of an API that receives an encrypted message that is passed to a calculator system. Integration testing is a type of testing that verifies and validates the functionality, performance, and reliability of different components or modules of a system or application when they are combined or integrated together. Integration testing can help to ensure the API can communicate and interact with the calculator system correctly and securely, as well as identify any errors or issues that may arise from the integration.

NEW QUESTION 44
- (Topic 2)
A cloud provider wants to make sure consumers are utilizing its IaaS platform but prevent them from installing a hypervisor on the server. Which of the following will help the cloud provider secure the environment and limit consumers' activity?

A. Patch management
B. Hardening
C. Scaling
D. Log and event monitoring

Answer: B

Explanation:
Hardening is the best option to help the cloud provider secure the environment and limit consumers' activity on its IaaS platform. Hardening is a process of reducing the attack surface and vulnerabilities of a system or device by applying security configurations, patches, updates, policies, rules, etc. Hardening can prevent consumers from installing unauthorized or unsupported software on their cloud servers, such as hypervisors.

NEW QUESTION 49
- (Topic 2)
A cloud security analyst needs to ensure the web servers in the public subnet allow only secure communications and must remediate any possible issue. The stateful configuration for the public web servers is as follows:

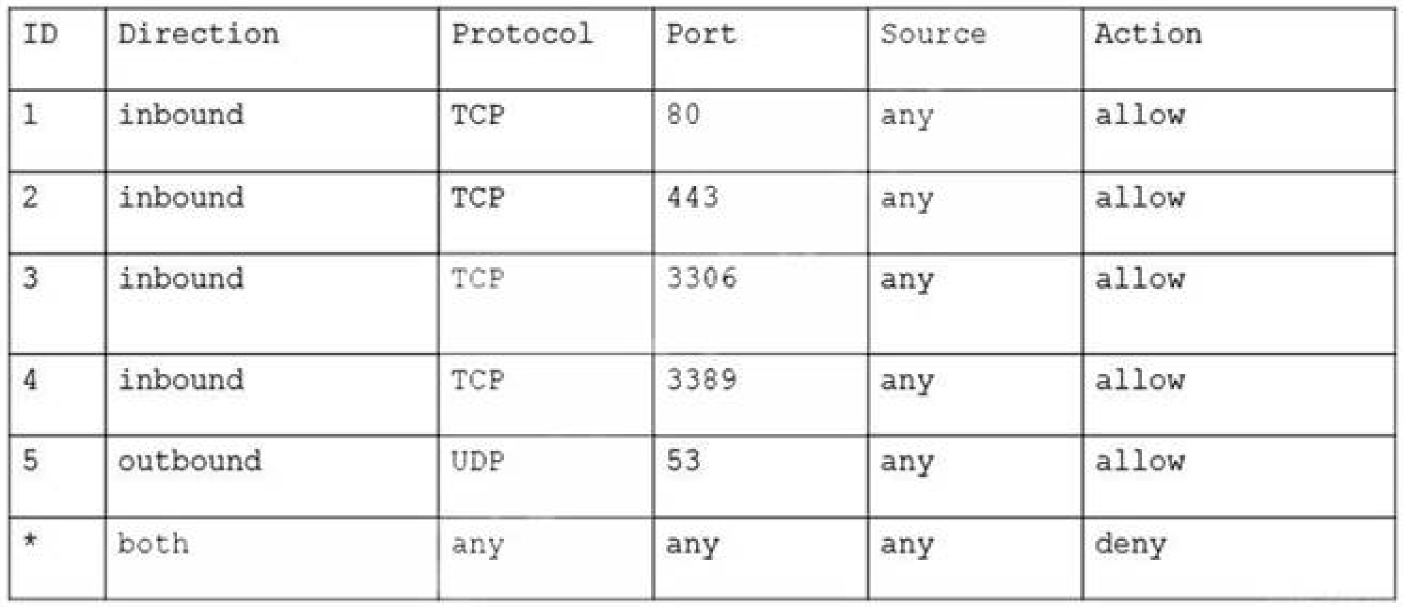

ID	Direction	Protocol	Port	Source	Action
1	inbound	TCP	80	any	allow
2	inbound	TCP	443	any	allow
3	inbound	TCP	3306	any	allow
4	inbound	TCP	3389	any	allow
5	outbound	UDP	53	any	allow
*	both	any	any	any	deny

Which of the following actions should the analyst take to accomplish the objective?

A. Remove rules 1, 2, and 5.
B. Remove rules 1, 3, and 4.
C. Remove rules 2, 3, and 4.
D. Remove rules 3, 4, and 5.

Answer: A

Explanation:
To ensure the web servers in the public subnet allow only secure communications and remediate any possible issue, the analyst should remove rules 1, 2, and 5 from the stateful configuration. These rules are allowing insecure or unnecessary traffic to or from the web servers, which may pose security risks or performance issues. The rules are:
? Rule 1: This rule allows inbound traffic on port 80 (HTTP) from any source to any destination. HTTP is an unencrypted and insecure protocol that can expose web traffic to interception, modification, or spoofing. The analyst should remove this rule and use HTTPS (port 443) instead, which encrypts and secures web traffic.
? Rule 2: This rule allows outbound traffic on port 25 (SMTP) from any source to any destination. SMTP is a protocol that is used to send email messages. The web servers in the public subnet do not need to send email messages, as this is not their function. The analyst should remove this rule and block outbound SMTP traffic, which may prevent spamming or phishing attacks from compromised web servers.
? Rule 5: This rule allows inbound traffic on port 22 (SSH) from any source to any destination. SSH is a protocol that allows remote access and management of systems or devices using a command-line interface. The web servers in the public subnet do not need to allow SSH access from any source, as this may expose them to unauthorized or malicious access. The analyst should remove this rule and restrict SSH access to specific sources, such as the administrator's workstation or a bastion host.

NEW QUESTION 54
- (Topic 2)
A cloud administrator set up a link between the private and public cloud through a VPN tunnel. As part of the migration, a large set of files will be copied. Which of the following network ports are required from a security perspective?

A. 22, 53, 445
B. 22, 443, 445
C. 25, 123, 443
D. 137, 139, 445

Answer: B

Explanation:
These are the network ports that are required from a security perspective to copy a large set of files between the private and public cloud through a VPN tunnel. A VPN (Virtual Private Network) tunnel is a secure and encrypted connection that allows data to be transferred between two networks or locations over the public internet. To copy files between the private and public cloud, the following ports are needed:
? Port 22: This is the port used by SSH (Secure Shell) protocol, which is a method of remotely accessing and managing cloud resources or systems using a command- line interface. SSH can also be used to securely transfer files using SCP (Secure Copy Protocol) or SFTP (SSH File Transfer Protocol).
? Port 443: This is the port used by HTTPS (Hypertext Transfer Protocol Secure), which is a protocol that encrypts and secures web traffic. HTTPS can also be used to transfer files using web browsers or tools such as curl or wget.
? Port 445: This is the port used by SMB (Server Message Block) protocol, which is a protocol that allows file sharing and access over a network. SMB can also be used to transfer files using tools such as robocopy or rsync.

NEW QUESTION 56
- (Topic 1)
Company A has acquired Company B and is in the process of integrating their cloud resources. Company B needs access to Company A’s cloud resources while retaining its IAM solution.
Which of the following should be implemented?

A. Multifactor authentication
B. Single sign-on
C. Identity federation
D. Directory service

Answer: C

Explanation:
Identity federation is a type of authentication mechanism that allows users to access multiple systems or applications across different domains or organizations with a single login credential. Identity federation can help integrate the cloud resources of Company A and Company B after Company A has acquired Company B, as it can enable seamless and secure access to both companies’ cloud resources using the same IAM solution. Identity federation can also improve user convenience, productivity, and security, as it can simplify the login process, reduce login errors, and enhance password management. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7
Reference: https://medium.com/@dinika.15/identity-federation-a-brief-introduction- f2f823f8795a

NEW QUESTION 61
- (Topic 1)
A systems administrator needs to configure SSO authentication in a hybrid cloud environment.
Which of the following is the BEST technique to use?

A. Access controls
B. Federation
C. Multifactor authentication
D. Certificate authentication

Answer: B

Explanation:
Federation is a type of authentication mechanism that allows users to access multiple systems or applications across different domains or organizations with a single login credential. Federation can help configure SSO authentication in a hybrid cloud environment, as it can enable seamless and secure access to cloud-based and on- premises resources using the same identity provider and authentication method. Federation can also improve user convenience, productivity, and security, as it can simplify the login process, reduce login errors, and enhance password management.
References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 65
- (Topic 1)
A developer is no longer able to access a public cloud API deployment, which was working ten minutes prior.
Which of the following is MOST likely the cause?

A. API provider rate limiting
B. Invalid API token
C. Depleted network bandwidth
D. Invalid API request

Answer: A

Explanation:
API provider rate limiting is a restriction on the number of requests that can be made to a web service or application programming interface (API) within a certain time period. API provider rate limiting can cause a failure to access a public cloud API deployment, as it can reject or block any requests that exceed the limit. API provider rate limiting can be used by cloud providers to control the usage and traffic of their customers and prevent overloading or abuse of their resources. API provider rate limiting is the most likely cause for the developer being unable to access a public cloud API deployment that was working ten minutes prior.
References: CompTIA Cloud+ Certification Exam Objectives, page 13, section 2.5

NEW QUESTION 68
- (Topic 1)
A storage array that is used exclusively for datastores is being decommissioned, and a new array has been installed. Now the private cloud administrator needs to migrate the data.
Which of the following migration methods would be the BEST to use?

A. Conduct a V2V migration

B. Perform a storage live migration
C. Rsync the data between arrays
D. Use a storage vendor migration appliance

Answer: B

Explanation:
A storage live migration is a process of moving or transferring data or files from one storage system or device to another without interrupting or affecting the availability or performance of the VMs or applications that use them. Performing a storage live migration can help migrate the data from a SAN that is being decommissioned to a new array, as it can ensure that there is no downtime or disruption for the VMs or applications that rely on the data or files stored on the SAN. Performing a storage live migration can also help maintain consistency and integrity, as it can synchronize and verify the data or files between the source and destination storage systems or devices.
References: CompTIA Cloud+ Certification Exam Objectives, page 10, section 1.5

NEW QUESTION 72
- (Topic 1)
A systems administrator is building a new virtualization cluster. The cluster consists of five virtual hosts, which each have flash and spinning disks. This storage is shared among all the virtual hosts, where a virtual machine running on one host may store data on another host.
This is an example of:

A. a storage area network
B. a network file system
C. hyperconverged storage
D. thick-provisioned disks

Answer: C

Explanation:
Hyperconverged storage is a type of storage architecture that combines compute, storage, and network resources into a single system or appliance. Hyperconverged storage uses software-defined storage (SDS) to pool and share the local storage of each node in the cluster, creating a distributed storage system that can be accessed by any node or virtual machine in the cluster. Hyperconverged storage can provide high performance, scalability, and efficiency for virtualized environments. The scenario of building a new virtualization cluster with five virtual hosts that share their flash and spinning disks among all the virtual hosts is an example of hyperconverged storage. References: [CompTIA Cloud+ Certification Exam Objectives], page 9, section 1.4

NEW QUESTION 76
- (Topic 1)
A systems administrator wants to have near-real-time information on the volume of data being exchanged between an application server and its clients on the Internet.
Which of the following should the systems administrator implement to achieve this objective?

A. A stateful firewall
B. DLP
C. DNSSEC
D. Network flows

Answer: D

Explanation:
Network flows are records of network traffic that capture information such as source and destination IP addresses, ports, protocols, timestamps, and byte and packet counts. Network flows can provide near-real-time information on the volume of data being exchanged between a system and its clients on the Internet, as they can measure and monitor the amount and rate of network traffic for each connection or session. Network flows can also help analyze network performance, troubleshoot network issues, and detect network anomalies or security incidents. A systems administrator should implement network flows to achieve the objective of having near-real-time information on the volume
of data being exchanged between an application server and its clients on the Internet. References: CompTIA Cloud+ Certification Exam Objectives, page 16, section 3.2

NEW QUESTION 80
- (Topic 1)
A cloud administrator has finished setting up an application that will use RDP to connect. During testing, users experience a connection timeout error.
Which of the following will MOST likely solve the issue?

A. Checking user passwords
B. Configuring QoS rules
C. Enforcing TLS authentication
D. Opening TCP port 3389

Answer: D

Explanation:
TCP port 3389 is the default port used by Remote Desktop Protocol (RDP) to connect to a remote system or application over a network. Opening TCP port 3389 on the firewall or network device will most likely solve the issue of users experiencing a connection timeout error when trying to use RDP to connect to an application, as it will allow RDP traffic to pass through. If TCP port 3389 is closed or blocked, RDP traffic will be denied or dropped, resulting in a connection timeout error. References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8
Reference: https://docs.microsoft.com/en-us/windows-server/remote/remote-desktop-services/troubleshoot/ rdp-error-general-troubleshooting

NEW QUESTION 81
- (Topic 1)
A SaaS provider wants to maintain maximum availability for its service. Which of the following should be implemented to attain the maximum SLA?

A. A hot site
B. An active-active site
C. A warm site
D. A cold site

Answer: B

Explanation:
An active-active site is a type of disaster recovery (DR) site that runs simultaneously with the primary site and handles part of the normal workload or traffic. An active-active site can help maintain maximum availability for a SaaS service, as it can provide load balancing, redundancy, and failover capabilities for the SaaS service in case of an outage or disruption at the primary site. An active-active site can also improve performance and scalability, as it can distribute the workload or traffic across multiple sites and handle increased demand or peak periods. References: CompTIA Cloud+ Certification Exam Objectives, page 10, section 1.5

NEW QUESTION 85
- (Topic 1)
Lateral-moving malware has infected the server infrastructure.
Which of the following network changes would MOST effectively prevent lateral movement in the future?

A. Implement DNSSEC in all DNS servers
B. Segment the physical network using a VLAN
C. Implement microsegmentation on the network
D. Implement 802.1X in the network infrastructure

Answer: C

Explanation:
Microsegmentation is a type of network security technique that divides a network into smaller logical segments or zones based on workload or application characteristics and applies granular policies and rules to control and isolate traffic within each segment or zone. Implementing microsegmentation on the network can help prevent lateral movement in the future after lateral-moving malware has infected the server infrastructure, as it can limit the exposure and spread of malware by restricting access and communication between different segments or zones based on predefined criteria such as identity, role, or behavior.
References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 87
- (Topic 1)
A cloud administrator is planning to migrate a globally accessed application to the cloud.
Which of the following should the cloud administrator implement to BEST reduce latency for all users?

A. Regions
B. Auto-scaling
C. Clustering
D. Cloud bursting

Answer: A

Explanation:
Regions are geographical locations or areas where cloud service providers have data centers or facilities that host their cloud resources or services. Regions can help reduce latency for all users when deploying a globally accessed application to the cloud, as they can enable faster and closer access to the cloud resources or services based on the user's physical location. Regions can also improve performance and availability, as they can provide redundancy and load balancing by distributing the workload across multiple locations. References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8

NEW QUESTION 88
- (Topic 4)
A cloud engineer is deploying a server in a cloud platform. The engineer reviews a security scan report. Which of the following recommended services should be disabled? (Select two).

A. Telnet
B. FTP
C. Remote log-in
D. DNS
E. DHCP
F. LDAP

Answer: AB

Explanation:
Telnet and FTP are recommended services to be disabled when deploying a server in a cloud platform, as they are insecure protocols that transmit data in plain text and expose credentials and sensitive information to potential attackers12. Remote log-in, DNS, DHCP, and LDAP are not necessarily recommended to be disabled, as they may provide useful functionality for the server and the cloud environment. However, they should be configured properly and secured with encryption, authentication, and authorization mechanisms34.
References: CompTIA Cloud+ CV0-003 Exam Objectives, Objective 4.2: Given a scenario, apply security configurations and compliance controls ; CompTIA Quick Start Guide to Tackling Cloud Security Concerns3

NEW QUESTION 92
- (Topic 4)
A cloud administrator created a developer desktop image and added it to the VDI farm in a private cloud environment. One of the developers opened a VDI session and noticed that compiling the code was taking up to one hour to complete. However, when the developer compiles the code on a local machine, the job completes in less than five minutes. Which of the following sizing techniques would be best to use to improve the performance of the compile job?

A. Add more servers to the VDI environment.

B. Increase the CPU and the memory on the VDI template.
C. Configure the VDI environment to increase sessions automatically.
D. Migrate code compile jobs to a public cloud provider.

Answer: B

Explanation:
The most likely cause of the poor performance of the compile job is that the VDI template does not have enough CPU and memory resources to handle the task efficiently. Compiling code is a CPU-intensive and memory-intensive process that requires sufficient computing power to run smoothly. By increasing the CPU and memory on the VDI template, the cloud administrator can improve the performance of the compile job and reduce the time it takes to complete. Adding more servers to the VDI environment or configuring the VDI environment to increase sessions automatically would not help, as they would only affect
the scalability and availability of the VDI farm, not the performance of individual sessions. Migrating code compile jobs to a public cloud provider would incur additional costs and complexity, and may not be feasible or desirable for the organization. References: The Official CompTIA Cloud+ Self-Paced Study Guide (CV0-
003) eBook, Chapter 3, Section 3.3, page 971

NEW QUESTION 94
- (Topic 4)
Which of the following enables CSPs to offer unlimited capacity to customers?

A. Adequate budget
B. Global data center distribution
C. Economies of scale
D. Agile project management

Answer: C

Explanation:
The correct answer is C. Economies of scale.
Economies of scale are the cost advantages that CSPs can achieve by increasing the size and scale of their operations. By spreading the fixed costs of infrastructure, software, and personnel over a larger customer base and data volume, CSPs can reduce the average cost per unit of service and offer unlimited capacity to customers at competitive prices1. Adequate budget is not a sufficient condition for offering unlimited capacity, as CSPs still need to optimize their resource utilization and efficiency to meet the growing demand for data storage and processing.
Global data center distribution is a strategy that CSPs use to improve their service availability, reliability, and performance by locating their servers closer to their customers and reducing network latency. However, this does not necessarily imply unlimited capacity, as CSPs still need to manage the trade-offs between data center size, cost, and power consumption.
Agile project management is a methodology that CSPs use to deliver their services faster, better, and cheaper by adopting iterative, incremental, and collaborative approaches. However, this does not directly affect their capacity, as CSPs still need to scale their infrastructure and software to handle the increasing data load.

NEW QUESTION 98
- (Topic 4)
An organization deployed an application using a cloud provider's internal managed certificates. Developers are unable to retrieve data when calling the API from any machine.
The following error message is in the log:

```
12-04-2023-10:05:25, SSL Negotiation Error 12-04-2023-10:05:28,Invalid Certificate
12-04-2023-10:05:29, TLS Handshake Failed 12-04-2023-10:05:30,Connection Closed
```

Which of the following is the most likely cause of the error?

A. TLS version
B. Insecure cipher
C. Self-signed certificate
D. Root trust

Answer: D

Explanation:
The error message indicates that the SSL/TLS handshake failed due to an invalid certificate. This means that the client machine does not trust the certificate authority (CA) that issued the certificate for the cloud provider’s API. A self-signed certificate or an insecure cipher would not cause this error, as they would be detected during the certificate validation process. The TLS version is not relevant, as the error occurs before the protocol negotiation. The most likely cause of the error is that the client machine does not have the root CA certificate installed in its trust store, or that the cloud provider’s certificate chain is incomplete or broken. To fix the error, the client machine needs to install the root CA certificate or the cloud provider needs to fix its certificate chain. References: The Official CompTIA Cloud+ Self-Paced Study Guide (CV0-003) eBook, Chapter 6, Section 6.2, page 2321

NEW QUESTION 101
- (Topic 4)
A cloud engineer is migrating a customer's web servers from a hypervisor platform to a CSP environment. The engineer needs to decouple the infrastructure and components during the migration to reduce the single points of failure. Which of the following storage options should the cloud engineer migrate the content to in order to improve availability?

A. Block
B. File
C. Object
D. iSCSI
E. NFS

Answer: C

Explanation:
Object storage is a storage option that stores data as discrete units called objects, which are identified by a unique identifier and can have metadata attached to them. Object storage can help the cloud engineer migrate the content to improve availability by decoupling the data from the underlying infrastructure and

components. Object storage can also provide high scalability, durability, and redundancy for the data, as well as support for multiple protocols and access methods. Object storage can be accessed through APIs, web interfaces, or gateways that can emulate file or block storage. Object storage is suitable for storing unstructured or static data, such as web content, images, videos, or documents. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 4, Objective 4.1: Given a scenario, implement cloud storage solutions.

NEW QUESTION 104
- (Topic 4)
A cloud administrator needs to deploy a security virtual appliance in a private cloud environment, but this appliance will not be part of the standard catalog of items for other users to request. Which of the following is the BEST way to accomplish this task?

A. Create an empty V
B. import the hard disk of the virtual applianc
C. and configure the CPU and memory.
D. Acquire the build scripts from the vendor and recreate the appliance using the baseline templates
E. Import the virtual appliance into the environment and deploy it as a VM
F. Convert the virtual appliance to a template and deploy a new VM using the template.

Answer: C

Explanation:
The correct answer is C. Import the virtual appliance into the environment and deploy it as a VM.
A virtual appliance is a pre-packaged and pre-configured software solution that runs on a virtual machine (VM). A virtual appliance typically consists of an operating system, an application, and any required dependencies, and is designed to provide a specific function or service. A virtual appliance can be distributed as a single file or a set of files that can be imported into a virtualization platform, such as VMware, Hyper-V, or KVM .
A cloud administrator can deploy a security virtual appliance in a private cloud environment by importing the virtual appliance into the environment and deploying it as a VM. This is the best way to accomplish this task because it preserves the original configuration and functionality of the virtual appliance, and does not require any additional installation or customization. The cloud administrator can also control the access and visibility of the virtual appliance, and prevent other users from requesting it from the standard catalog of items .
Creating an empty VM, importing the hard disk of the virtual appliance, and configuring the CPU and memory is not the best way to accomplish this task because it involves more steps and complexity than importing the virtual appliance as a whole. It also introduces the risk of losing or corrupting some data or settings during the import process, or misconfiguring the CPU and memory for the virtual appliance.
Acquiring the build scripts from the vendor and recreating the appliance using the baseline templates is not the best way to accomplish this task because it involves more time and effort than importing the virtual appliance directly. It also depends on whether the vendor provides the build scripts or not, and whether they are compatible with the baseline templates or not.
Converting the virtual appliance to a template and deploying a new VM using the template is not the best way to accomplish this task because it adds an unnecessary step of creating a template from the virtual appliance. It also does not prevent other users from accessing or requesting the template from the catalog of items.

NEW QUESTION 105
- (Topic 4)
A systems administrator is working within a private cloud environment. Over time. random 4K read/write speeds on all VMS in the environment slow down until the VMS are completely unusable, with disk speeds of less than 1MBps. The administrator has gathered the information below:
• There is no correlation between the slowdown and VM/hypervisor resource utilization.
• The network is rated to 40Gbps and utilization is between 1—5%.
• The hypervisors use hundreds of NFSv3 mounts to the same storage appliance, one per VM.
• The VMS on each hypervisor become unresponsive after two weeks of uptime.
• The unresponsiveness is resolved by moving slow VMS onto a rebooted hypervisor. Which of the following solutions will MOST likely resolve this issue?

A. Increase caching on the storage appliance.
B. Configure jumbo frames on the hypervisors and storage.
C. Increase CPU/RAM resources on affected VMS.
D. Reduce the number of NFSv3 mounts to one.

Answer: D

Explanation:
The correct answer is D. Reduce the number of NFSv3 mounts to one.
NFSv3 is a network file system protocol that allows clients to access files stored on a remote server. NFSv3 uses TCP or UDP as the transport layer protocol, and typically runs on port 20491.
One of the known issues with NFSv3 mounts is that they can cause performance degradation and unresponsiveness on the client side if there are too many mounts or if there are network connectivity problems. This is because NFSv3 does not handle connection failures or timeouts gracefully, and may keep retrying to access the server indefinitely, blocking other processes or threads. This can result in slow disk speeds, high CPU usage, and system hangs23.
Therefore, one of the possible solutions to this issue is to reduce the number of NFSv3 mounts to one per hypervisor, instead of one per VM. This way, the hypervisor can manage the access to the shared storage appliance more efficiently, and avoid creating too many TCP connections or UDP packets that may overload the network or the server. Reducing the number of NFSv3 mounts can also simplify the configuration and troubleshooting of the network file system.
Increasing caching on the storage appliance may improve the read performance of the NFSv3 mounts, but it will not solve the underlying issue of connection failures or timeouts. Caching may also introduce data inconsistency or corruption issues if the cache is not synchronized with the server.
Configuring jumbo frames on the hypervisors and storage may improve the network throughput and efficiency of the NFSv3 mounts, but it will not solve the underlying issue of connection failures or timeouts. Jumbo frames are larger than standard Ethernet frames, and require that all devices on the network path support them. Jumbo frames may also introduce fragmentation or compatibility issues if they are not configured properly. Increasing CPU/RAM resources on affected VMs may improve their performance in general, but it will not solve the underlying issue of connection failures or timeouts. Increasing CPU/RAM resources may also be costly and wasteful if they are not needed for other purposes.

NEW QUESTION 110
- (Topic 4)
A systems administrator deployed a new web application in a public cloud and would like to test it, but the company's network firewall is only allowing outside connections to the cloud provider network using TCP port 22. While waiting for the network administrator to open the required ports, which of the following actions should the systems administrator take to test the new application? (Select two).

A. Create an IPSec tunnel.

B. Create a VPN tunnel.
C. Open a browser using the default gateway IP address.
D. Open a browser using the localhost IP address.
E. Create a GRE tunnel.
F. Create a SSH tunnel.

Answer: BF

Explanation:
To test the new web application in the public cloud, the systems administrator should create a replica database, synchronize the data, and switch to the new instance, and create a SSH tunnel. Creating a replica database can help minimize the downtime and ensure data consistency during the migration. Synchronizing the data can help keep the replica database up to date with the original database. Switching to the new instance can help activate the new web application in the public cloud. Creating a SSH tunnel can help bypass the network firewall and access the web application using TCP port 22. SSH is a secure protocol that can create encrypted tunnels between the local and remote hosts. By creating a SSH tunnel, the systems administrator can forward the web application traffic through the tunnel and test it using a web browser. References: [CompTIA Cloud+ CV0-003 Certification Study Guide], Chapter 7, Objective 7.1: Given a scenario, migrate applications and data to the cloud.

NEW QUESTION 113
- (Topic 4)
A cloud engineer is deploying a server in a cloud platform. The engineer reviews a security scan report. Which of the following recommended services should be disabled? (Select TWO).

A. Telnet
B. FTP
C. Remote login
D. DNS
E. DHCP
F. LDAP

Answer: AB

Explanation:
Telnet and FTP are two services that should be disabled on a cloud server because they are insecure and vulnerable to attacks. Telnet and FTP use plain text to transmit data over the network, which means that anyone who can intercept the traffic can read or modify the data, including usernames, passwords, commands, files, etc. This can lead to data breaches, unauthorized access, or malicious actions on the server1.
Instead of Telnet and FTP, more secure alternatives should be used, such as SSH (Secure Shell) and SFTP (Secure File Transfer Protocol). SSH and SFTP use encryption to protect the data in transit and provide authentication and integrity checks for the communication. SSH and SFTP can prevent eavesdropping, tampering, or spoofing of the data and ensure the confidentiality and privacy of the server2.
The other options are not services that should be disabled on a cloud server:
? Option C: Remote login. Remote login is a service that allows users to access a remote server from another location using a network connection. Remote login can be useful for managing, configuring, or troubleshooting a cloud server without having to physically access it. Remote login can be secured by using encryption, authentication, authorization, and logging mechanisms3.
? Option D: DNS (Domain Name System). DNS is a service that translates human- friendly domain names into IP addresses that can be used to communicate over the Internet. DNS is essential for resolving the names of the cloud resources and services that are hosted on the cloud platform. DNS can be secured by using DNSSEC (DNS Security Extensions), which add digital signatures to DNS records to verify their authenticity and integrity.
? Option E: DHCP (Dynamic Host Configuration Protocol). DHCP is a service that assigns IP addresses and other network configuration parameters to devices on a network. DHCP can simplify the management of IP addresses and avoid conflicts or errors in the network. DHCP can be secured by using DHCP snooping, which filters out unauthorized DHCP messages and prevents rogue DHCP servers from assigning IP addresses.
? Option F: LDAP (Lightweight Directory Access Protocol). LDAP is a service that stores and organizes information about users, devices, and resources on a network. LDAP can provide identity management and access control for the cloud environment. LDAP can be secured by using LDAPS (LDAP over SSL/TLS), which encrypts the LDAP traffic and provides authentication and integrity checks.

NEW QUESTION 116
- (Topic 4)
Which of the following provides groups of compute units that can horizontally scale according to a workload?

A. Orchestrated container environment
B. Cloud-reserved instances
C. Autoscaling
D. Cloud bursting

Answer: C

Explanation:
Autoscaling is a feature that allows groups of compute units to horizontally scale according to a workload or predefined rules. Autoscaling can increase or decrease the number of compute units dynamically based on metrics such as CPU utilization, memory usage, network traffic, or user demand. Autoscaling can improve performance, availability, and cost-efficiency of cloud applications.
References: [CompTIA Cloud+ Study Guide], page 75.

NEW QUESTION 120
- (Topic 4)
A systems administrator is configuring a cloud solution for a vulnerability assessment to test the company's resources that are hosted in a public cloud. The solution must test the company's resources from an external user's perspective. Which of the following should the systems administrator configure?

A. An agent-based scan
B. A network-based scan
C. A port scan
D. A credentialed scan

Answer: B

Explanation:
A network-based scan is a type of vulnerability assessment that tests the security of a system or a network from an external user's perspective, without requiring any software or credentials on the target. A network-based scan can identify vulnerabilities such as open ports, misconfigured firewalls, outdated software, or exposed services .

NEW QUESTION 125
- (Topic 4)
An organization is implementing a new requirement to facilitate faster downloads for users of corporate application content. At the same time, the organization is also expanding cloud regions. Which of the following would be suitable to optimize the network for this requirement?

A. Implement CDN for overall cloud application.
B. Implement autoscaling of the compute resources.
C. Implement SR-IOV on the server instances.
D. Implement an application container solution.

Answer: A

Explanation:
CDN, or content delivery network, is a system of distributed servers that deliver web content to users based on their geographic location, the origin of the web page, and the content delivery server1. A CDN can improve the performance, availability, and scalability of cloud applications by caching static and dynamic content at the edge of the network, reducing the latency and bandwidth consumption between the users and the cloud servers2. A CDN can also provide security features such as encryption, authentication, and DDoS protection3.
Autoscaling, SR-IOV, and containerization are other techniques that can optimize the network for cloud applications, but they are not directly related to the requirement of faster downloads for users. Autoscaling is the process of automatically adjusting the number and size of compute resources based on the demand and workload of the application. SR-IOV, or single root I/O virtualization, is a technology that allows a physical network device to be partitioned into multiple virtual devices that can be assigned to different virtual machines or containers, bypassing the hypervisor and improving the network performance and efficiency. Containerization is the process of packaging an application and its dependencies into a lightweight and portable unit that can run on any platform, providing isolation, consistency, and portability.
References:
? CompTIA Cloud+ CV0-003 Study Guide, Chapter 4: Network Optimization, Section 4.1: Content Delivery Networks, Page 17523
? CompTIA Cloud+ CV0-003 Study Guide, Chapter 4: Network Optimization, Section 4.2: Autoscaling, Page 180
? CompTIA Cloud+ CV0-003 Study Guide, Chapter 4: Network Optimization, Section 4.3: SR-IOV, Page 184
? CompTIA Cloud+ CV0-003 Study Guide, Chapter 4: Network Optimization, Section 4.4: Containerization, Page 187
? What is a CDN?

NEW QUESTION 129
- (Topic 4)
A systems administrator is planning to deploy a database cluster in a virtualization environment. The administrator needs to ensure the database nodes do not exist on the same physical host. Which of the following would best meet this requirement?

A. Oversubscription
B. Anti-affinity
C. A firewall
D. A separate cluster

Answer: B

Explanation:
Anti-affinity is a rule that specifies that certain virtual machines should not run on the same physical host. This can help to improve availability and performance by avoiding single points of failure and resource contention. For example, if the database nodes are running on the same host and the host fails, the entire database cluster will be unavailable. By using anti-affinity rules, the systems administrator can ensure the database nodes are distributed across different hosts in the virtualization environment. References: CompTIA Cloud+ CV0-003 Study Guide, Chapter 2: Deploying a Cloud Environment, page 76.

NEW QUESTION 132
- (Topic 4)
An organization's two-node, hybrid container cluster is experiencing failures during horizontal scaling to the cloud cluster instance. The on-premises IP range is 192.168.0.0/16, and the cloud environment is 10.168.0.0/16. Overlapping or stretched VLANs are not permitted, and a node is deployed in each location. The cloud monitoring agent reports a healthy status for the second instance, but when pinging the clusters from on premises, the following output is received:
pinging cluster1. comptia. containers.com C192.168.100 reply
pinging cluster2. comptia. containers.com [192.16B .100 .128] request timed out
Which of the following is the most likely reason for the scaling failure?

A. Incorrect DNS entry
B. Offline cluster node
C. Incorrect proxy entry
D. Incorrect cluster IP
E. Incorrect IP route

Answer: E

Explanation:
An incorrect IP route is the most likely reason for the scaling failure, as it prevents the communication between the on-premises and cloud cluster nodes. The ping output shows that the DNS entry for cluster2.comptia.containers.com is resolved to an IP address in the cloud environment (192.168.100.128), but the request times out, indicating a network connectivity issue. An incorrect proxy entry, an offline cluster node, or an incorrect cluster IP would not cause the DNS resolution to fail. An incorrect DNS entry would not cause the ping request to time out.
References: CompTIA Cloud+ CV0-003 Exam Objectives, Objective 2.2: Given a scenario, deploy and test a cloud solution ; Configuring clusters, scaling, and monitoring for hybrid api management …1 ; CompTIA Cloud+ : Cloud High Availability & Scaling - Skillsoft2

NEW QUESTION 136

- (Topic 4)
A cloud administrator needs to reduce storage costs. Which of the following would best help the administrator reach that goal?

A. Enabling compression
B. Implementing deduplication
C. Using containers
D. Rightsizing the VMs

Answer: B

Explanation:
Deduplication is a process by which redundant data is eliminated, thus reducing the size of the dataset. Deduplication with cloud storage reduces the storage requirements, along with the amount of data to be transferred over the network, resulting in faster and more efficient data protection operations1. Deduplication can help to shrink the data footprint, lower the storage costs, and improve the performance of backup and recovery processes2. Deduplication can be applied at different levels, such as file-level, block-level, or byte-level, depending on the granularity and efficiency of the technique3. Deduplication can also be performed at different locations, such as source, target, or cloud, depending on the architecture and design of the storage system3. By implementing deduplication, a cloud administrator can achieve significant data savings and optimize the cloud storage costs4. References: Data deduplication techniques for efficient cloud storage management: a systematic review; How Data Deduplication Reduces Cloud Data Costs; How Data Deduplication Can Save Cloud Storage Costs?; Data Deduplication Overview; What is Data Deduplication and How Can it Help Reduce Cloud Costs?.

NEW QUESTION 137
- (Topic 4)
An organization is conducting a performance test of a public application. The following actions have already been completed:
• The baseline performance has been established.
• A load test has passed.
• A benchmark report has been generated.
Which of the following needs to be done to conclude the performance test?

A. Verify the application works well under an unexpected volume of requests.
B. Assess the application against vulnerabilities and/or misconfiguration exploitation.
C. Test how well the application can resist a DDoS attack.
D. Conduct a test with the end users and collect feedback.

Answer: A

Explanation:
To conclude the performance test of a public application, the organization needs to verify the application works well under an unexpected volume of requests. This is also known as a stress test, which is a type of performance test that evaluates the behavior and stability of the application under extreme conditions1. A stress test can help identify potential bottlenecks, errors, or failures that may occur when the application is subjected to a sudden surge or spike in demand2. A stress test can also help determine the maximum capacity and scalability of the application3.
References: CompTIA Cloud+ CV0-003 Exam Objectives, Objective 2.2: Given a scenario, deploy and test a cloud solution ; Performance Testing | Cloud Computing | CompTIA1 ; Stress Testing - Software Testing Fundamentals2 ; What is Stress Testing? Definition, Types, Tools & Examples3

NEW QUESTION 138
- (Topic 4)
A cloud engineer recently set up a container image repository. The engineer wants to ensure that downloaded images are not modified in transit. Which of the following is the best method to achieve this goal?

A. SHA-256
B. IPSec
C. AES-256
D. MD5
E. serpent-256

Answer: A

Explanation:
SHA-256 is the best method to ensure that downloaded images are not modified in transit. SHA-256 is a type of cryptographic hash function that can generate a unique and fixed- length digest for any input data. The digest can be used to verify the integrity and
authenticity of the data, as any modification or tampering of the data would result in a different digest. SHA-256 is more secure and reliable than MD5, which is an older and weaker hash function that has been proven to be vulnerable to collisions and attacks12. AES-256 and serpent-256 are types of encryption algorithms, not hash functions, and they are used to protect the confidentiality of the data, not the integrity. IPSec is a network security protocol that can use encryption and hashing to secure data in transit, but it is not a method by itself

NEW QUESTION 140
- (Topic 4)
After an infrastructure-as-code cloud migration to an laaS environment, the cloud engineer discovers that configurations on DB servers have drifted from the corporate standard baselines. Which of the following should the cloud engineer do to best ensure configurations are restored to the baselines?

A. Utilize a template to automate and update the DB configuration.
B. Create an image of the DB, delete the previous DB server, and restore from the image.
C. Manually log in to the DB servers and update the configurations.
D. Rename and change the IP of the old DB server and rebuild a new DB server.

Answer: A

Explanation:
A template is a file that defines the desired state and configuration of a cloud resource, such as a server, a network, or a database. Infrastructure as code (IaC) is the practice of using templates to automate and manage cloud resources, rather than manually configuring them. IaC can help prevent configuration drift, which is the deviation of the actual state of a resource from the desired state defined by the template. In this scenario, the cloud engineer discovers that configurations on

DB servers have drifted from the corporate standard baselines after an IaC cloud migration to an IaaS environment. The best way to ensure configurations are restored to the baselines is to utilize a template to automate and update the DB configuration. This way, the cloud engineer can apply the same template to all the DB servers, and ensure they are consistent and compliant with the corporate standards. Creating an image of the DB, deleting the previous DB server, and restoring from the image is not a good solution, as it may cause data loss, downtime, and additional costs. Manually logging in to the DB servers and updating the configurations is not a good solution, as it is time-consuming, error-prone, and not scalable. Renaming and
changing the IP of the old DB server and rebuilding a new DB server is not a good solution, as it may cause compatibility issues, network disruptions, and security risks. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 23, Infrastructure as Code and Configuration Management, page 3691.

NEW QUESTION 142
- (Topic 4)
An enterprise is considering a cost model for a DBaaS. Which of the following is BEST for a cloud solution?

A. per gigabyte
B. per seat
C. Per user
D. Per device

Answer: A

Explanation:
The correct answer is A. per gigabyte.
A cost model for a DBaaS is a way of determining how much the user pays for the database service. Different cost models may have different pricing factors, such as storage usage, data transfer, compute resources, and additional services.
A per gigabyte cost model is best for a cloud solution because it allows the user to pay only for the amount of storage space they use for their database. This way, the user can scale up or down their storage needs as per their requirements and budget. A per gigabyte cost model also reflects the actual cost of the infrastructure, software licenses, and maintenance that the service provider incurs to host and operate the database1.
A per seat cost model is not suitable for a cloud solution because it charges the user based on the number of seats or licenses they purchase for the database service. This means that the user may end up paying for more seats than they actually use, or not have enough seats to accommodate their users. A per seat cost model also does not account for the storage usage or performance of the database.
A per user cost model is also not suitable for a cloud solution because it charges the user
based on the number of users who access the database service. This means that the user may have to pay more if they have a large number of users, or less if they have a small number of users. A per user cost model also does not account for the storage usage or performance of the database.
A per device cost model is also not suitable for a cloud solution because it charges the user based on the number of devices that connect to the database service. This means that the user may have to pay more if they have multiple devices per user, or less if they have one device per user. A per device cost model also does not account for the storage usage or performance of the database.

NEW QUESTION 144
- (Topic 4)
A non-critical file on a database server was deleted and needs to be recovered. A cloud administrator must use the least disruptive restoration process to retrieve the file, as the database server cannot be stopped during the business day. Which of the following restoration methods would best accomplish this goal?

A. Alternate location
B. Restore from image
C. Revert to snapshot
D. In-place restoration

Answer: D

Explanation:
In-place restoration is the process of restoring data to the same location where it was originally stored, without affecting the rest of the system. This method is suitable for recovering non-critical files that were accidentally deleted, as it does not require stopping the server or creating a new instance. In contrast, alternate location, restore from image, and revert to snapshot are more disruptive methods that involve creating a new copy of the data or the entire system, which may affect the performance or availability of the
server. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 20, Backup and Restore Operations, page 3211.

NEW QUESTION 146
- (Topic 4)
A web consultancy group currently works in an isolated development environment. The group uses this environment for the creation of the final solution, but also for showcasing it to customers, before commissioning the sites in production. Recently, customers of newly commissioned sites have reported they are not receiving the final product shown by the group, and the website is performing in unexpected ways. Which of the following additional environments should the group adopt and include in its process?

A. Provide each web consultant a local environment on their device.
B. Require each customer to have a blue-green environment.
C. Leverage a staging environment that is tightly controlled for showcasing
D. Initiate a disaster recovery environment to fail to in the event of reported issues.

Answer: C

Explanation:
The answer is C. Leverage a staging environment that is tightly controlled for showcasing. A staging environment is a replica of the production environment that is used for testing and demonstrating the final product before deployment. A staging environment can help the web consultancy group avoid the issues reported by the customers, such as mismatched expectations and unexpected behavior, by ensuring that the product is shown in a realistic and consistent setting. A staging environment can also help the group catch and fix any bugs or errors before they affect the live site.
Some possible sources of information about web development environments are:
? 7 Web Development Best Practices: This page provides some general tips and best practices for web development, such as planning, accessibility, UX/UI, standards, code quality, compatibility, and security.
? Web Development Best Practices (Building Real-World Cloud Apps with Azure):
This page explains some specific best practices for web development in the cloud environment, such as stateless web tier, session state management, CDN caching, and async programming.

? Web Development Best Practices: This page lists some resources for learning web
development best practices in ASP.NET, such as async and await, building real- world cloud apps with Azure, and hands-on labs.

NEW QUESTION 148
- (Topic 4)
A cloud engineer needs to perform a database migration_ The database has a restricted SLA and cannot be offline for more than ten minutes per month The database stores 800GB of data, and the network bandwidth to the CSP is 100MBps. Which of the following is the BEST option to perform the migration?

A. Copy the database to an external device and ship the device to the CSP
B. Create a replica database, synchronize the data, and switch to the new instance.
C. Utilize a third-patty tool to back up and restore the data to the new database
D. use the database import/export method and copy the exported file.

Answer: B

Explanation:
The correct answer is B. Create a replica database, synchronize the data, and switch to the new instance.
This option is the best option to perform the migration because it can minimize the downtime and data loss during the migration process. A replica database is a copy of the source database that is kept in sync with the changes made to the original database. By creating a replica database in the cloud, the cloud engineer can transfer the data incrementally and asynchronously, without affecting the availability and performance of the source database. When the replica database is fully synchronized with the source database, the cloud engineer can switch to the new instance by updating the connection settings and redirecting the traffic. This can reduce the downtime to a few minutes or seconds, depending on the complexity of the switch.
Some of the tools and services that can help create a replica database and synchronize the data are AWS Database Migration Service (AWS DMS) 1, Azure Database Migration Service 2, and Striim 3. These tools and services can support various source and target databases, such as Oracle, MySQL, PostgreSQL, SQL Server, MongoDB, etc. They can also provide features such as schema conversion, data validation, monitoring, and security. The other options are not the best options to perform the migration because they can cause more downtime and data loss than the replica database option.
? Copying the database to an external device and shipping the device to the CSP is
a slow and risky option that can take days or weeks to complete. It also exposes the data to physical damage or theft during transit. Moreover, this option does not account for the changes made to the source database after copying it to the device, which can result in data inconsistency and loss.
? Utilizing a third-party tool to back up and restore the data to the new database is a
faster option than shipping a device, but it still requires a significant amount of
downtime and bandwidth. The source database has to be offline or in read-only mode during the backup process, which can take hours or days depending on the size of the data and the network speed. The restore process also requires downtime and bandwidth, as well as compatibility checks and configuration adjustments. Additionally, this option does not account for the changes made to the source database after backing it up, which can result in data inconsistency and loss.
? Using the database import/export method and copying the exported file is a similar
option to using a third-party tool, but it relies on native database features rather than external tools. The import/export method involves exporting the data from the source database into a file format that can be imported into the target database. The file has to be copied over to the target database and then imported into it. This option also requires downtime and bandwidth during both export and import processes, as well as compatibility checks and configuration adjustments. Furthermore, this option does not account for the changes made to the source database after exporting it, which can result in data inconsistency and loss.

NEW QUESTION 151
- (Topic 4)
A cloud engineer is responsible for a legacy web application that runs on an on-premises VM environment. The VM environment is approaching end of life. The engineer needs to migrate the web application to the cloud as quickly as possible because the VM environment has the following limitations:
• The VM environment has a single IOGB disk.
• The VM environment still uses 10Mbps, which leaves a 100Mbps WAN connection underutilized.
• No installation media is available.
Which of the following is the best way to migrate the web application to the cloud?

A. Use the VM import connector to import the VM into the cloud.
B. Use import/export to import the VM as a snapshot and attach it to a cloud instance.
C. Use REST APIs to import an image of the VM into the cloud.
D. Use object storage to create a backup of the VM and restore data into the cloud instance.

Answer: A

Explanation:
A VM import connector is a tool that allows you to import virtual machines from your on-premises environment into the cloud using a graphical user interface. This is the fastest and easiest way to migrate a legacy web application without requiring installation media or changing the configuration of the VM. The VM import connector can also handle the disk size and network bandwidth limitations of the on-premises VM environment. References: EC2 VM Import Connector | AWS News Blog, Import a VMware Virtual Machine to Oracle Cloud Infrastructure, CompTIA Cloud+ Certification Exam Objectives, Domain 2.0: Deployment, Objective 2.1: Given a scenario, execute and implement solutions using appropriate cloud migration tools and methods.

NEW QUESTION 152
- (Topic 4)
Following the deployment of a new VM, a cloud engineer notices the backup platform has not added the machine to the appropriate job. The backup platform uses a text-based variable for job configuration. This variable is based on the RPO requirements for the workload. Which of the following did the cloud engineer forget to configure when deploying the virtual machine?
? Tags

A. RPO
B. RTO
C. Server name
D. Template

Answer: A

Explanation:
Tags are key-value pairs that can be applied to cloud resources to organize, categorize, and filter them. Tags can also be used to assign resources to backup jobs based on their RPO requirements. The cloud engineer forgot to configure the appropriate tag for the new VM that matches the text-based variable of the backup

platform. Therefore, the backup platform did not add the VM to the correct job. References: Tags and labels |
Cloud Storage | Google Cloud, CompTIA Cloud+ Certification Exam Objectives, Domain 4.0: Operations and Support, Objective 4.3: Given a scenario, apply the appropriate methods for cost control in a cloud environment.

NEW QUESTION 156
FILL IN THE BLANK - (Topic 4)
?MISSING?

A.

Answer: D

Explanation:
This means that data is divided into blocks and written across multiple disks, and two additional disks are used to store parity information that can be used to reconstruct data in case of disk failure. RAID 6 can withstand the failure of up to two disks without losing any data or performance. RAID 6 also maximizes the storage capacity of its drives, as it only uses two disks for parity out of the total number of disks in the array. For example, if the array has 10 disks, RAID 6 will use 8 disks for data and 2 disks for parity, resulting in a storage capacity of 8/10 or 80% of the total disk space. RAID 6 is suitable for private cloud environments that require high availability, fault tolerance, and large storage
capacity. References: CompTIA Cloud+ CV0-003 Study Guide, Chapter 3: Storage Technologies, Section 3.2: RAID Levels, Page 125

NEW QUESTION 160
- (Topic 4)
A systems administrator is troubleshooting a VDI deployment that is used to run high- frame-rate rendering. Users are reporting frequent application crashes. After running a benchmark, the administrator discovers the following:

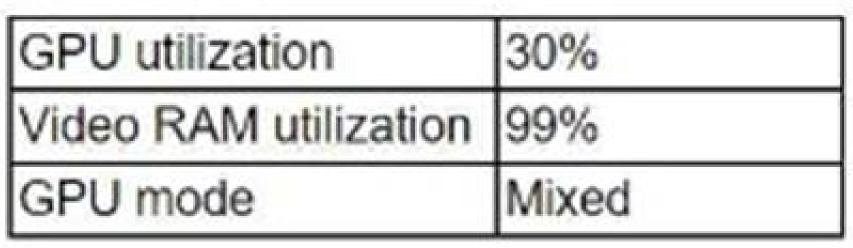

GPU utilization	30%
Video RAM utilization	99%
GPU mode	Mixed

Which of the following should the administrator do to resolve this issue?

A. Configure the GPU to run in compute mode.
B. Allocate more RAM in the VM template.
C. Select a higher vGPU profile.
D. Configure the GPU to run in graphics mode.

Answer: C

Explanation:
The benchmark results show that the video RAM utilization is at 99%, which is likely causing the application crashes. Video RAM is used to store graphics data and textures that are processed by the GPU. Selecting a higher vGPU profile can help allocate more video RAM to the virtual machines, which can help resolve this issue. A vGPU profile is a predefined configuration that specifies the amount of video RAM, the number of display heads, and the maximum resolution that a virtual machine can use. By selecting a higher vGPU profile, the administrator can increase the performance and stability of the high- frame-rate rendering application. References: [CompTIA Cloud+ CV0-003 Study Guide], Chapter 4, Objective 4.2: Given a scenario, troubleshoot common virtualization issues.

NEW QUESTION 163
- (Topic 4)
A systems administrator is selecting the appropriate RAID level to support a private cloud with the following requirements:
. The storage array must withstand the failure of up to two drives.
. The storage array must maximize the storage capacity of its drives.
Which of the following RAID levels should the administrator implement?

A. RAID 0
B. RAID 1
C. RAID 5
D. RAID 6
E. RAID 10

Answer: D

Explanation:
RAID stands for Redundant Array of Independent Disks, which is a technology that combines multiple physical disks into a logical unit that provides improved performance, reliability, and storage capacity. RAID levels are different ways of organizing and distributing data across the disks in a RAID array. Each RAID level has its own advantages and disadvantages, depending on the requirements and trade-offs of the system.
RAID 6 is a RAID level that uses block-level striping with double parity. This means that data is divided into blocks and distributed across all the disks in the array, and two sets of parity information are calculated and stored on different disks. Parity is a method of error detection and correction that can reconstruct the data in case of disk failure. RAID 6 can withstand the failure of up to two disks without losing any data, which makes it suitable for a private cloud that requires high fault tolerance. RAID 6 also maximizes the storage capacity of its drives, as it only uses two disks for parity and the rest for data. The storage capacity of a RAID 6 array is equal to (n-2) x S, where n is the number of disks and S is the size of the smallest disk.
RAID 0, RAID 1, RAID 5, and RAID 10 are other RAID levels, but they do not meet the requirements of the private cloud. RAID 0 uses striping without parity, which improves performance but does not provide any redundancy or fault tolerance. RAID 0 cannot withstand any disk failure, as it would result in data loss. RAID 1 uses mirroring, which copies the same data to two or more disks. RAID 1 provides high reliability and fast read performance, but it wastes half of the storage capacity for redundancy. RAID 1 can only withstand the failure of one disk in each mirrored pair. RAID 5 uses striping with single parity, which distributes data and parity across all the disks in the array. RAID 5 provides a balance of performance, reliability, and storage capacity, but it can only withstand the failure of one disk. RAID 10 is a combination of RAID 1 and RAID 0, which creates a striped array of mirrored pairs. RAID 10 provides high performance and reliability, but it also wastes half of the storage capacity for redundancy. RAID 10 can withstand the failure of one disk in each mirrored pair, but not more than that.

For more information on RAID levels, you can refer to the following sources:
? CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 4, Storage Technologies, page 791
? Cloud+ (Plus) Certification | CompTIA IT Certifications2

NEW QUESTION 164
- (Topic 4)
A systems administrator has been notified of possible illegal activities taking place on the network and has been directed to ensure any relevant emails are preserved for court use.
Which of the following is this MOST likely an example of?

A. Email archiving
B. Version control
C. Legal hold
D. File integrity monitoring

Answer: C

Explanation:
The correct answer is C. Legal hold.
A legal hold is a process that organizations use to preserve relevant electronic information when they anticipate litigation or have an active e-discovery request. A legal hold requires that certain email messages be retained and unaltered until they are no longer required for court use. Legal hold requirements apply both to the content of messages as well as the metadata which can provide proof of delivery and other critical non-repudiation information12.
Email archiving is a process that organizations use to store email messages for long-term retention, compliance, and backup purposes. Email archiving does not necessarily imply that the email messages are preserved for legal purposes, although some email archiving solutions may offer legal hold capabilities1.
Version control is a process that software developers use to manage changes to source code and other files in a project. Version control allows developers to track, compare, and revert changes, as well as collaborate with other developers. Version control does not apply to email messages or legal hold.
File integrity monitoring is a process that security professionals use to detect unauthorized or malicious changes to files and directories on a system. File integrity monitoring helps to protect the system from malware, data breaches, and configuration errors. File integrity monitoring does not apply to email messages or legal hold.

NEW QUESTION 168
- (Topic 4)
A cloud administrator has deployed a website and needs to improve the site security to meet requirements. The website architecture is designed to have a DBaaS in the back end and autoscaling instances in the front end using a load balancer to distribute the request.
Which of the following will the cloud administrator most likely use?

A. An API gateway
B. An IPS/IDS
C. A reverse proxy
D. A WAF

Answer: D

Explanation:
A web application firewall (WAF) is a security solution that monitors and filters the traffic between a web application and the Internet. A WAF can help improve the site security by blocking malicious requests, preventing SQL injection attacks, mitigating cross-site scripting (XSS) attacks, and enforcing security policies. A WAF can be deployed as a cloud service or as a device in front of the load balancer. A WAF is more suitable than an API gateway, an IPS/IDS, or a reverse proxy for the website architecture described in the question. References: [CompTIA Cloud+ CV0-003 Certification Study Guide], Chapter 9, Objective 9.1: Given a scenario, apply security controls and techniques.

NEW QUESTION 173
- (Topic 4)
A cloud engineer recently used a deployment script template to implement changes on a cloud-hosted web application. The web application communicates with a managed database on the back end. The engineer later notices the web application is no longer receiving data from the managed database. Which of the following is the most likely cause of the issue?

A. Misconfiguration in the user permissions
B. Misconfiguration in the routing traffic
C. Misconfiguration in the network ACL
D. Misconfiguration in the firewall

Answer: D

Explanation:
A misconfiguration in the firewall is the most likely cause of the issue. A firewall is a security device or service that controls the incoming and outgoing network traffic based on predefined rules. A firewall can help protect the cloud-hosted web application and the managed database from unauthorized or malicious access. However, if the firewall rules are not configured properly, they can also block the legitimate communication between the web application and the database. For example, if the firewall rules deny the port or protocol that the web application uses to connect to the database, the web application will not be able to receive data from the database. To fix this issue, the cloud engineer should review and update the firewall rules to allow the necessary traffic between the web application and the database. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 9, Objective 9.2: Given a scenario, troubleshoot common security issues.

NEW QUESTION 178
- (Topic 4)
A cloud administrator is investigating slow VM performance. The administrator has checked the physical server performance and has identified the host is under stress due to a peak usage workload. Which of the following is the NEXT step the administrator should complete?

A. Perform a root cause analysis
B. Migrate the VM to a different host.
C. Document the findings.

D. Perform a system restart.

Answer: B

Explanation:
Migrating the VM to a different host is a common technique to improve the performance of a VM that is suffering from resource contention or contention on the physical server. By moving the VM to a different host, the administrator can:
Reduce the stress and load on the original host, which may be under stress due to a peak usage workload.
Increase the availability and reliability of the VM, which may be experiencing slow performance due to resource contention or contention on the original host.
Balance the workload and resource utilization across multiple hosts, which may improve the overall performance and efficiency of the cloud environment.
Migrating the VM to a different host can be done manually or automatically, depending on the configuration and capabilities of the cloud platform. Some cloud platforms support live migration, which allows moving a VM to a different host without interrupting its operation or service. Other cloud platforms require shutting down or pausing the VM before migrating it to a different host .

NEW QUESTION 179
- (Topic 3)
A cloud administrator would like to maintain file integrity checks through hashing on a cloud object store. Which of the following is MOST suitable from a performance perspective?

A. SHA-256
B. SHA-512
C. MD5
D. AES

Answer: C

Explanation:
The most suitable hashing algorithm from a performance perspective to maintain file integrity checks on a cloud object store is MD5 (Message Digest 5). MD5 is a hashing algorithm that generates a 128-bit hash value for any given input data. MD5 is faster and more efficient than other hashing algorithms, such as SHA-256 or SHA-512, which generate longer hash values and require more computational resources. MD5 can be used to verify the integrity of files by comparing their hash values before and after transmission or storage. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 2.0 Security, Objective 2.5 Given a scenario, apply data security techniques in the cloud.

NEW QUESTION 181
- (Topic 3)
A company has two primary offices, one in the United States and one in Europe. The company uses a public laaS service that has a global data center presence to host its marketing materials. The marketing team, which is primarily based in Europe, has reported latency issues when retrieving these materials. Which of the following is the BEST option to reduce the latency issues?

A. Add an application load balancer to the applications to spread workloads.
B. Integrate a CDN solution to distribute web content globally.
C. Upgrade the bandwidth of the dedicated connection to the laaS provider.
D. Migrate the applications to a region hosted in Europe.

Answer: B

Explanation:
The best option to reduce the latency issues for the marketing team that is primarily based in Europe when retrieving the marketing materials that are hosted on a public IaaS service is to integrate a CDN (content delivery network) solution to distribute web content globally. A CDN is a network of geographically distributed servers that cache and deliver web content to users based on their proximity and network conditions. A CDN can improve the performance and availability of web content by reducing the distance and hops between the users and the servers, as well as offloading the traffic from the origin server. Reference: [CompTIA Cloud+ Certification Exam Objectives], Domain 3.0 Maintenance, Objective 3.4 Given a scenario, implement automation and orchestration to optimize cloud operations

NEW QUESTION 182
- (Topic 3)
A company is using an laaS environment. Which of the following licensing models would BEST suit the organization from a financial perspective to implement scaling?

A. Subscription
B. Volume-based
C. per user
D. Socket-based

Answer: B

Explanation:
A volume-based licensing model is a licensing model that charges the customer based on the amount of data or resources that they consume or use. A volume-based licensing model is suitable for an laaS (Infrastructure as a Service) environment, as it allows the customer to pay only for what they need and scale up or down as their demand changes. A volume-based licensing model can provide financial benefits for the customer, such as lower upfront costs, greater flexibility, and more predictable billing.

NEW QUESTION 184
- (Topic 3)
A web application has been configured to use autoscaling for provisioning and deprovisioning more VMs according to the workload. The systems administrator deployed a new CI/CD tool to automate new releases of the web application. During the night, a script was deployed and configured to be executed by the VMs during bootstrapping. Now. the autoscaling configuration is creating a new VM ever\ five minutes. Which of the following actions will MOS I like y resolve the issue?

A. Reducing the maximum threshold in the autoscaling configuration
B. Debugging the script and redeploying it

C. Changing the automation tool because it is incompatible
D. Modifying the script to shut down the VM after five minutes

Answer: B

Explanation:
The best way to resolve the issue where the autoscaling configuration is creating a new VM every five minutes after deploying a new CI/CD tool to automate new releases of the web application and configuring a script to be executed by the VMs during bootstrapping is to debug the script and redeploy it. Debugging the script means finding and fixing any errors or bugs in the code or logic of the script that may cause unexpected or undesired behavior, such as triggering the autoscaling condition or failing to complete the bootstrapping process. Redeploying the script means updating or replacing the existing script with the corrected or improved version of the script. Reference: [CompTIA Cloud+ Certification Exam Objectives], Domain 4.0 Troubleshooting, Objective 4.5 Given a scenario, troubleshoot automation/orchestration issues.

NEW QUESTION 186
- (Topic 3)
A systems administrator is responding to an outage in a cloud environment that was caused by a network-based flooding attack. Which of the following should the administrator configure to mitigate the attack?

A. NIPS
B. Network overlay using GENEVE
C. DDoS protection
D. DoH

Answer: C

Explanation:
DDoS protection is what the administrator should configure to mitigate a network-based flooding attack that caused an outage in a cloud environment. A network-based flooding attack is a type of attack that sends a large amount of network traffic or requests to a target system or service, such as a server, website, application, etc., with the intention of overwhelming or exhausting its resources or capacity. A network-based flooding attack can cause an outage in a cloud environment by disrupting or degrading the availability or performance of the target system or service, as well as affecting other systems or services that share the same network or infrastructure. DDoS protection is a tool or service that detects and prevents network-based flooding attacks, also known as Distributed Denial of Service (DDoS) attacks. DDoS protection can mitigate a network- based flooding attack by providing features such as:
? Filtering: DDoS protection can filter network traffic or requests based on various criteria, such as source, destination, protocol, content, etc., and block or allow them accordingly.
? Diverting: DDoS protection can divert network traffic or requests away from the target system or service to another location or device, such as a scrubbing center, proxy, firewall, etc., where they can be analyzed and processed.
? Scaling: DDoS protection can scale network resources or capacity dynamically and automatically to handle the increased demand or load caused by the network- based flooding attack.

NEW QUESTION 190
- (Topic 3)
A systems administrator needs to implement a service to protect a web application from external attacks. The administrator must have session-based granular control of all HTTP traffic. Which of the following should the administrator configure?

A. IDS
B. WAF
C. DLP
D. NAC

Answer: B

Explanation:
Reference: https://en.wikipedia.org/wiki/Web_application_firewall
A web application firewall (WAF) is a type of firewall that monitors and filters HTTP traffic to and from a web application. It can detect and block malicious requests, such as SQL injection, cross-site scripting, or denial-of-service attacks. It can also provide session-based granular control of HTTP traffic, such as allowing or denying access based on user identity, location, or behavior. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 2.0 Security, Objective 2.2 Given a scenario, implement appropriate network security controls for a cloud environment.

NEW QUESTION 191
- (Topic 3)
A cloud administrator is troubleshooting a highly available web application running within three containers behind a Layer 7 load balancer with a WAF inspecting all traffic. The application frequently asks the users to log in again even when the session timeout has not been reached. Which of the following should the cloud administrator configure to solve this issue?

A. Firewall outbound rules
B. Firewall inbound rules
C. Load balancer certificates
D. Load balancer stickiness
E. WAF transaction throttling

Answer: D

Explanation:
Reference: https://docs.aws.amazon.com/elasticloadbalancing/latest/application/application-load- balancers.html#sticky-sessions
Load balancer stickiness is what the cloud administrator should configure to solve the issue of the application frequently asking the users to log in again even when the session timeout has not been reached for a highly available web application running within three containers behind a Layer 7 load balancer with a WAF inspecting all traffic. Load balancer stickiness is a feature that allows customers to maintain user sessions or connections with the same server or node that provides a service or function, such as a web application, database, etc., even when there are multiple servers or nodes behind a load balancer. Load balancer stickiness can solve the issue by providing benefits such as:
Consistency: Load balancer stickiness can provide consistency by ensuring that users receive the same service or function from the same server or node

throughout their session or connection, without any changes or variations.
Performance: Load balancer stickiness can provide performance by reducing the latency or overhead of switching between different servers or nodes during a session or connection, which may cause delays or errors.
Security: Load balancer stickiness can provide security by preserving and protecting user authentication or authorization information on the same server or node during a session or connection, without exposing or transferring it to other servers or nodes.

NEW QUESTION 193
- (Topic 3)
A storage administrator is reviewing the storage consumption of a SAN appliance that is running a VDI environment. Which of the following features should the administrator implement to BEST reduce the storage consumption of the SAN?

A. Deduplication
B. Thick provisioning
C. Compression
D. SDS

Answer: A

Explanation:
The best feature to reduce the storage consumption of a SAN appliance that is running a VDI environment is deduplication. Deduplication is a process that eliminates redundant or duplicate data blocks or files from a storage system and replaces them with pointers or references to a single copy of data. Deduplication can significantly reduce the storage consumption of a SAN appliance by removing unnecessary data and freeing up disk space. Reference: [CompTIA Cloud+ Certification Exam Objectives], Domain 3.0 Maintenance, Objective 3.3 Given a scenario, analyze system performance using standard tools.

NEW QUESTION 194
- (Topic 3)
A company that requires full administrative control at the OS level is considering the use of public cloud services. Which of the following service models would BEST fit the company's requirements?

A. SaaS
B. DBaaS
C. PaaS
D. laaS

Answer: D

Explanation:
laaS (Infrastructure as a Service) is a public cloud service model that provides access to fundamental compute, network, and storage resources on demand over the public Internet or through dedicated connections. Customers can provision and configure these resources according to their needs, and they have full administrative control at the OS level. This means that customers can install, update, and manage any software or applications they want on the cloud servers, as well as apply their own security and compliance policies. laaS is suitable for companies that require high flexibility and customization of their cloud infrastructure, as well as scalability and cost-efficiency.

NEW QUESTION 197
- (Topic 3)
Audit and system logs are being forwarded to a syslog solution. An administrator observes that two application servers have not generated any logs for a period of three days, while
others continue to send logs normally. Which of the following BEST explains what is occurring?

A. There is a configuration failure in the syslog solution.
B. The application servers were migrated to the cloud as laaS instances.
C. The application administrators have not performed any activity in those servers.
D. There is a local firewall policy restriction on the syslog server.

Answer: B

Explanation:
One possible explanation for why two application servers have not generated any logs for a period of three days, while others continue to send logs normally, is that the application servers were migrated to the cloud as laaS (Infrastructure as a Service) instances. laaS is a cloud service model that provides virtualized computing resources over the internet, such as servers, storage, network, and operating systems. When an application server is migrated to the cloud as an laaS instance, it may require some configuration changes to enable the syslog forwarding to the same destination as before. For example, the laaS instance may have a different IP address, hostname, firewall rules, or network settings than the original server. If these changes are not properly made, the laaS instance may not be able to communicate with the syslog solution and send logs as expected.

NEW QUESTION 199
- (Topic 3)
A company is deploying a public cloud solution for an existing application using lift and shift. The requirements for the applications are scalability and external access. Which of the following should the company implement? (Select TWO).

A. A load balancer
B. SON
C. A firewall
D. SR-IOV
E. Storage replication
F. A VPN

Answer: AF

Explanation:

The best options to implement for a public cloud solution for an existing application using lift and shift that requires scalability and external access are a load balancer and a VPN (virtual private network). A load balancer is a device or service that distributes incoming traffic across multiple servers or instances based on various criteria, such as availability, capacity, or performance. A load balancer can improve scalability by balancing the workload and optimizing resource utilization. A VPN is a technology that creates a secure and encrypted connection over a public network, such as the internet. A VPN can provide external access by allowing remote users or sites to connect to the cloud resources as if they were on the same private network. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 1.0 Configuration and Deployment, Objective 1.4 Given a scenario, execute a provided deployment plan.

NEW QUESTION 203
- (Topic 3)
Users currently access SaaS email with five-character passwords that use only letters and numbers. An administrator needs to make access more secure without changing the password policy. Which of the following will provide a more secure way of accessing email at the lowest cost?

A. Change the email service provider.
B. Enable MFA with a one-time password.
C. Implement SSO for all users.
D. Institute certificate-based authentication

Answer: B

Explanation:
Enable MFA with a one-time password. MFA stands for multi-factor authentication, which is a method of verifying a user’s identity by requiring two or more forms of authentication. A one-time password (OTP) is a code that is generated randomly and valid only for a short period of time. By enabling MFA with OTP, the administrator can make access to the SaaS email more secure without changing the password policy, as users will need to provide both their password and an OTP to sign in.

NEW QUESTION 206
- (Topic 3)
A systems administrator is using a configuration management tool to perform maintenance tasks in a system. The tool is leveraging the target system's API to perform these maintenance tasks. After a number of features and security updates are applied to the target system, the configuration management tool no longer works as expected. Which of the following is the MOST likely cause of the issue?

A. The target system's API functionality has been deprecated.
B. The password for the service account has expired.
C. The IP addresses of the target system have changed.
D. The target system has failed after the updates.

Answer: A

Explanation:
The most likely cause of the issue is A. The target system’s API functionality has been deprecated. API deprecation is the process of gracefully discontinuing an API. The process starts by first informing the customers that the API is no longer actively supported even though it will be operational and suggesting them to migrate to an alternate or latest version of the API1. However, sometimes the API functionality may change or be removed without proper notice or documentation, which can break the existing applications that rely on the API. According to the web search results, API deprecation is a common challenge for configuration management tools. Therefore, if the target system’s API functionality has been deprecated after the updates, the configuration management tool may no longer work as expected.

NEW QUESTION 208
- (Topic 3)
A cloud solutions architect has received guidance to migrate an application from on premises to a public cloud. Which of the following requirements will help predict the operational expenditures in the cloud?

A. Average resource consumption
B. Maximum resource consumption
C. Minimum resource consumption
D. Actual hardware configuration

Answer: B

Explanation:
Maximum resource consumption is what will help predict the operational expenditures in the cloud for an application that is being migrated from on premises to a public cloud provider. Operational expenditures are the ongoing costs of running and maintaining a system or service, such as cloud resources or services. Operational expenditures can vary depending on various factors, such as usage, demand, performance, etc. Maximum resource consumption is the highest amount of resources or capacity that are used or consumed by a system or service during peak periods of activity or load. Maximum resource consumption can help predict operational expenditures in the cloud by providing information such as:
? Resource type: This indicates what type of resources are used or consumed by the system or service, such as compute, storage, network, etc.
? Resource amount: This indicates how much of each resource type are used or consumed by the system or service, such as CPU cores, memory, disk space, bandwidth, etc.
? Resource price: This indicates how much each resource type costs in the cloud provider’s pricing model, such as per hour, per GB, per Mbps, etc.

NEW QUESTION 209
- (Topic 3)
A systems administrator is securing a new email system for a large corporation. The administrator wants to ensure private corporate information is not emailed to external users. Which of the following would be MOST useful to accomplish this task?

A. DLP
B. EDR
C. DNSSEC
D. SPF

Answer: A

Explanation:
The most useful tool to prevent private corporate information from being emailed to external users is data loss prevention (DLP). DLP is a type of security solution that monitors and controls the flow of data in and out of a system or network. It can detect and prevent unauthorized access, transmission, or leakage of sensitive data, such as personal information, financial records, or intellectual property. DLP can also enforce encryption, masking, or deletion of sensitive data to protect its confidentiality.
Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 2.0 Security, Objective 2.5 Given a scenario, apply data security techniques in the cloud.

NEW QUESTION 212
- (Topic 3)
A systems administrator must ensure confidential company information is not leaked to competitors. Which of the following services will BEST accomplish this goal?

A. CASB
B. IDS
C. FIM
D. EDR
E. DLP

Answer: E

Explanation:
DLP (Data Loss Prevention) is a service that prevents the unauthorized or accidental disclosure of confidential or sensitive data, such as company information, intellectual property, customer data, or personal information. DLP can monitor, detect, and block the data in motion (such as email, web, or network traffic), data at rest (such as files, databases, or cloud storage), or data in use (such as endpoints, applications, or clipboard). DLP can help a systems administrator to ensure confidential company information is not leaked to competitors by applying policies and rules that define what data is considered confidential, who can access it, how it can be used, and what actions to take if a violation occurs. For example, DLP can encrypt, quarantine, delete, or alert the administrator if confidential data is being copied, transferred, or shared outside the organization.

NEW QUESTION 213
- (Topic 3)
A cloud administrator needs to establish a secure connection between two different locations. Which of the following is the BEST option to implement the secure connection?

A. HTTPS
B. IPSec
C. TLS
D. SSH

Answer: B

Explanation:
The best option to implement a secure connection between two different locations is IPSec (Internet Protocol Security). IPSec is a protocol suite that provides security for IP-based communications over networks. IPSec can encrypt and authenticate the data packets between two endpoints, such as routers, firewalls, or VPN gateways. IPSec can also provide integrity, confidentiality, and replay protection for the data. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 2.0 Security, Objective 2.2 Given a scenario, implement appropriate network security controls for a cloud environment.

NEW QUESTION 215
- (Topic 3)
A security audit related to confidentiality controls found the following transactions occurring in the system:
GET
http://gateway.securetransaction.com/privileged/api/v1/changeResource?id=123&user=277 Which of the following solutions will solve the audit finding?

A. Using a TLS-protected API endpoint
B. Implementing a software firewall
C. Deploying a HIDS on each system
D. Implementing a Layer 4 load balancer

Answer: A

Explanation:
Reference: https://cheatsheetseries.owasp.org/cheatsheets/Transport_Layer_Protection_Cheat_Sheet
.html
The audit finding is related to confidentiality, which means the data should be protected from unauthorized access. The current API endpoint is using HTTP, which is not secure and can expose the data in transit. Using a TLS-protected API endpoint would encrypt the data and prevent anyone from reading it. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 2.0 Security, Objective 2.1 Given a scenario, apply security configurations and compliance controls to meet cloud security requirements.

NEW QUESTION 220
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

CV0-003 Practice Exam Features:

* CV0-003 Questions and Answers Updated Frequently

* CV0-003 Practice Questions Verified by Expert Senior Certified Staff

* CV0-003 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* CV0-003 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The CV0-003 Practice Test Here