CompTIA

Exam Questions CV0-003

CompTIA Cloud+ Certification Exam

NEW QUESTION 1
- (Topic 1)
A systems administrator needs to configure monitoring for a private cloud environment. The administrator has decided to use SNMP for this task.
Which of the following ports should the administrator open on the monitoring server's firewall?

A. 53
B. 123
C. 139
D. 161

Answer: D

Explanation:
Port 161 is the default port used by Simple Network Management Protocol (SNMP) to communicate with network devices and collect information about their status, performance, configuration, and events. Opening port 161 on the monitoring server's firewall will allow SNMP traffic to pass through and enable monitoring for a private cloud environment. If port 161 is closed or blocked, SNMP traffic will be denied or dropped, resulting in a failure to monitor the network devices.
References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8

NEW QUESTION 2
- (Topic 1)
An SQL injection vulnerability was reported on a web application, and the cloud platform
team needs to mitigate the vulnerability while it is corrected by the development team. Which of the following controls will BEST mitigate the risk of exploitation?
A. DLP
B. HIDS
C. NAC
D. WAF

Answer: D

Explanation:
A web application firewall (WAF) is a type of network security device or software that monitors and filters HTTP traffic between a web application and the Internet. A WAF can help mitigate the risk of exploitation of an SQL injection vulnerability reported on a web application while it is corrected by the development team, as it can detect and block any malicious requests or queries that attempt to inject SQL commands into the web application's database. A WAF can also help protect the web application from other common web-based attacks, such as cross-site scripting (XSS), remote file inclusion (RFI), or denial-of-service (DoS). References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 3
- (Topic 1)
The security team for a large corporation is investigating a data breach. The team members are all trying to do the same tasks but are interfering with each other's work. Which of the following did the team MOST likely forget to implement?

A. Incident type categories
B. A calling tree
C. Change management
D. Roles and responsibilities

Answer: D

Explanation:
Roles and responsibilities are definitions or descriptions of what each team member or stakeholder is expected to do or perform in a project or process. Roles and responsibilities can help clarify the scope, authority, and accountability of each team member or stakeholder and avoid any confusion or duplication of work. The security team most likely forgot to implement roles and responsibilities when investigating a data breach, as they are all trying to do the same tasks but are interfering with each other's work. Implementing roles and responsibilities can help improve efficiency and effectiveness, as it can ensure that each team member or stakeholder knows what tasks they need to do and how they need to coordinate with others. References: CompTIA Cloud+ Certification Exam Objectives, page 13, section 2.5

NEW QUESTION 4
- (Topic 1)
An organization is implementing a new requirement to facilitate users with faster downloads 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 auto-scaling of the compute resources
C. Implement SR-IOV on the server instances
D. Implement an application container solution

Answer: C

Explanation:

Reference: https://access.redhat.com/documentation/en- us/red_hat_openstack_platform/13/html/
network_functions_virtualization_planning_and_configuration_guide/part-sriov-nfv- configuration

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)
A systems administrator is troubleshooting performance issues with a Windows VDI environment. Users have reported that VDI performance has been slow since the images were upgraded from Windows 7 to Windows 10.
This VDI environment is used to run simple tasks, such as Microsoft Office. The administrator investigates the virtual machines and finds the following settings:
? 4 vCPU
? 16GB RAM
? 10Gb networking
? 256MB frame buffer
Which of the following MOST likely needs to be upgraded?

A. vRAM
B. vCPU
C. vGPU
D. vNIC

Answer: C

Explanation:
A virtual graphics processing unit (vGPU) is a type of hardware or software that enables a VM to use the physical GPU resources of the host or server for graphics-intensive tasks. Upgrading the vGPU is most likely to solve the issue of VDI performance being slow since the images were upgraded from Windows 7 to Windows 10, as it can provide more graphics processing power and memory for the VMs. Upgrading the vGPU can also improve the user experience and productivity, as it can enhance the display quality and responsiveness of the VDI environment. References: CompTIA Cloud+ Certification Exam Objectives, page 11, section 1.6

NEW QUESTION 7
- (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 8
- (Topic 1)
An organization requires the following to be achieved between the finance and marketing departments:
? Allow HTTPS/HTTP.
? Disable FTP and SMB traffic.
Which of the following is the MOST suitable method to meet the requirements?

A. Implement an ADC solution to load balance the VLAN traffic
B. Configure an ACL between the VLANs
C. Implement 802.1X in these VLANs
D. Configure on-demand routing between the VLANs

Answer: B

Explanation:
An access control list (ACL) is a set of rules that defines which traffic is allowed or denied between different network segments or devices. An ACL can be used to filter traffic based on various criteria, such as source and destination addresses, ports, protocols, and applications. Configuring an ACL between the VLANs of the finance and marketing departments is the most suitable method to meet the requirements of allowing HTTPS/HTTP and disabling FTP and SMB traffic. An ACL can specify which ports and protocols are permitted or blocked between the VLANs, such as allowing port 80 (HTTP) and port 443 (HTTPS), and denying port 21

(FTP) and port 445 (SMB). References: [CompTIA Cloud+ Certification Exam Objectives], page 15, section 2.8

NEW QUESTION 9
- (Topic 1)
An organization is hosting a cloud-based web server infrastructure that provides web- hosting solutions. Sudden continuous bursts of traffic have caused the web servers to saturate CPU and network utilizations.
Which of the following should be implemented to prevent such disruptive traffic from reaching the web servers?

A. Solutions to perform NAC and DLP
B. DDoS protection
C. QoS on the network
D. A solution to achieve microsegmentation

Answer: B

Explanation:
Distributed denial-of-service (DDoS) protection is a type of security solution that detects and mitigates DDoS attacks that aim to overwhelm or disrupt a system or service by sending large volumes of traffic from multiple sources. DDoS protection can prevent such disruptive traffic from reaching the web servers by filtering out malicious or unwanted traffic and allowing only legitimate traffic to pass through. DDoS protection can also help maintain the availability and functionality of web services and applications during a DDoS attack. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7
Reference: https://blog.paessler.com/the-top-5-causes-of-sudden-network-spikes

NEW QUESTION 10
- (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 10
- (Topic 1)
A cloud administrator is building a new VM for a network security appliance. The security appliance installer says the CPU clock speed does not meet the requirements.
Which of the following will MOST likely solve the issue?

A. Move the VM to a host with a faster CPU
B. Add more vCPUs to the VM
C. Enable CPU masking on the VM
D. Enable hyperthreading on the virtual host

Answer: A

Explanation:
Moving the VM to a host with a faster CPU is the best way to solve the issue of the security appliance installer saying the CPU clock speed does not meet the requirements when building a new VM for a network security appliance. Moving the VM to a host with a faster CPU can ensure that the VM meets the minimum CPU clock speed requirement for the security appliance, as it can use the physical CPU resources of the host. Moving the VM to a host with a faster CPU can also improve the performance and reliability of the security appliance, as it can reduce latency, contention, and overhead.
References: CompTIA Cloud+ Certification Exam Objectives, page 11, section 1.6

NEW QUESTION 13
- (Topic 1)
A systems administrator recently upgraded the processors in a web application host. Upon the next login, the administrator sees a new alert regarding the license being out of compliance.
Which of the following licensing models is the application MOST likely using?

A. Per device
B. Per user
C. Core-based
D. Volume-based

Answer: C

Explanation:
Core-based licensing is a type of licensing model that charges based on the number of processor cores in a system or server. Core-based licensing is often used by software vendors to align their pricing with the performance and capacity of modern hardware. Core-based licensing can also enable customers to optimize

their licensing costs by choosing the appropriate hardware configuration for their needs. Upgrading the processors in a web application host can affect the core-based licensing of the application, as it may increase the number of cores that need to be licensed. This can result in an alert regarding the license being out of compliance if the license is not updated accordingly. References: CompTIA Cloud+ Certification Exam Objectives, page 20, section 4.2
Reference: https://download.microsoft.com/download/3/d/4/3d42bdc2-6725-4b29-b75a- a5b04179958b/ percorelicensing_definitions_vlbrief.pdf

NEW QUESTION 14
- (Topic 1)
A VDI administrator has received reports of poor application performance. Which of the following should the administrator troubleshoot FIRST?

A. The network environment
B. Container resources
C. Client devices
D. Server resources

Answer: A

Explanation:
The network environment is the set of network devices, connections, protocols, and configurations that enable communication and data transfer between different systems and applications. The network environment can affect the performance of a virtual desktop infrastructure (VDI) by influencing factors such as bandwidth, latency, jitter, packet loss, and congestion. Poor network performance can result in slow or unreliable application delivery, degraded user experience, and reduced productivity.
Therefore, troubleshooting the network environment should be the first step for a VDI administrator who receives reports of poor application performance.
References: CompTIA Cloud+ Certification Exam Objectives, page 17, section 3.4

NEW QUESTION 19
- (Topic 2)
A cloud administrator is managing an organization's infrastructure in a public cloud. All servers are currently located in a single virtual network with a single firewall that all traffic must pass through. Per security requirements, production, QA, and development servers should not be able to communicate directly with each other. Which of the following should an administrator perform to comply with the security requirement?

A. Create separate virtual networks for production, QA, and development server
B. Move the servers to the appropriate virtual network.Apply a network security group to each virtual network that denies all traffic except for the firewall.
C. Create separate network security groups for production, QA, and development server
D. Apply the network security groups on the appropriate production, QA, and development servers.Peer the networks together.
E. Create separate virtual networks for production, QA, and development server
F. Move the servers to the appropriate virtual network.Peer the networks together.
G. Create separate network security groups for production, QA, and development server
H. Peer the networks together.Create static routes for each network to the firewall.

Answer: A

Explanation:
These are the actions that the administrator should perform to comply with the security requirement of isolating production, QA, and development servers from each other in a public cloud environment:
? Create separate virtual networks for production, QA, and development servers: A virtual network is a logical isolation of network resources or systems within a cloud environment. Creating separate virtual networks for different types of servers can help to segregate them from each other and prevent direct communication or interference.
? Move the servers to the appropriate virtual network: Moving the servers to the appropriate virtual network can help to assign them to their respective roles and functions, as well as ensure that they follow the network policies and rules of their virtual network.
? Apply a network security group to each virtual network that denies all traffic except for the firewall: A network security group is a set of rules or policies that control and filter inbound and outbound network traffic for a virtual network or system. Applying a network security group to each virtual network that denies all traffic except for the firewall can help to enforce security and compliance by blocking any unauthorized or unwanted traffic between different types of servers, while allowing only necessary traffic through the firewall.

NEW QUESTION 24
- (Topic 2)
A company is concerned about the security of its data repository that contains customer PII. A systems administrator is asked to deploy a security control that will prevent the exfiltration of such data. Which of the following should the systems administrator implement?

A. DLP
B. WAF
C. FIM
D. ADC

Answer: A

Explanation:
Reference: https://cloud.google.com/blog/products/identity-security/4-steps-to-stop-data- exfiltration-with-google-cloud
Implementing DLP (Data Loss Prevention) is the best solution to prevent the exfiltration of customer PII (Personally Identifiable Information) from a data repository. DLP is a security control that monitors, detects, and blocks sensitive data from leaving or being accessed by unauthorized parties. DLP can be applied at different levels, such as network, endpoint, storage, or cloud. DLP can help to protect customer PII from being leaked, stolen, or compromised.

NEW QUESTION 28
- (Topic 2)
A cloud administrator has been using a custom VM deployment script. After three months of use, the script no longer joins the LDAP domain. The cloud administrator verifies the account has the correct permissions. Which of the following is the MOST likely cause of the failure?

A. Incorrect encryption ciphers
B. Broken trust relationship

C. Invalid certificates
D. Expired password

Answer: D

Explanation:
An expired password is the most likely cause of the failure of a custom VM deployment script that no longer joins the LDAP domain. LDAP (Lightweight Directory Access Protocol) is a protocol that allows access and management of directory services, such as user accounts, groups, permissions, etc., over a network. LDAP can be used to authenticate and authorize users or devices to access network resources or systems. An expired password is a password that has reached its validity period and needs to be changed or renewed. An expired password can prevent users or devices from joining or accessing an LDAP domain, as it may indicate that the account is inactive, compromised, or outdated.

NEW QUESTION 33
- (Topic 2)
Some VMs that are hosted on a dedicated host server have each been allocated with 32GB of memory. Some of VMs are not utilizing more than 30% of the allocation. Which of the following should be enabled to optimize the memory utilization?

A. Auto-scaling of compute
B. Oversubscription
C. Dynamic memory allocations on guests
D. Affinity rules in the hypervisor

Answer: C

Explanation:
Enabling dynamic memory allocations on guests is the best option to optimize memory utilization for VMs that have been allocated with 32GB of memory but are not utilizing more than 30% of it. Dynamic memory allocation is a feature that allows a VM to adjust its memory usage according to its workload and demand, without requiring a reboot or manual intervention. Dynamic memory allocation can help to improve memory utilization and efficiency by allocating more memory to VMs that need it and releasing memory from VMs that do not need it.

NEW QUESTION 36
- (Topic 2)
A systems administrator has received an email from the virtualized environment's alarms indicating the memory was reaching full utilization. When logging in, the administrator notices that one out of a five-host cluster has a utilization of 500GB out of 512GB of RAM. The baseline utilization has been 300GB for that host. Which of the following should the administrator check NEXT?

A. Storage array
B. Running applications
C. VM integrity
D. Allocated guest resources

Answer: D

Explanation:
Allocated guest resources is what the administrator should check next after receiving an email from the virtualized environment’s alarms indicating the memory was reaching full utilization and noticing that one out of a five-host cluster has a utilization of 500GB out of 512GB of RAM. Allocated guest resources are the amount of resources or capacity that are assigned or reserved for each guest system or device within a host system or device. Allocated guest resources can affect performance and utilization of host system or device by determining how much resources or capacity are available or used by each guest system or device. Allocated guest resources should be checked next by comparing them with the actual usage or demand of each guest system or device, as well as identifying any overallocation or underallocation of resources that may cause inefficiency or wastage.

NEW QUESTION 38
- (Topic 2)
A VDI administrator has received reports from the drafting department that rendering is slower than normal. Which of the following should the administrator check FIRST to optimize the performance of the VDI infrastructure?

A. GPU
B. CPU
C. Storage
D. Memory

Answer: A

Explanation:
Checking the GPU (Graphics Processing Unit) is the first thing that the VDI administrator should do to optimize the performance of the VDI infrastructure for rendering tasks. GPU is a specialized hardware device that accelerates graphics processing and rendering. GPU can improve the user experience and performance of VDI applications that require intensive graphics processing, such as drafting, gaming, video editing, etc.

NEW QUESTION 41
- (Topic 2)
A systems administrator is troubleshooting performance issues with a VDI environment. The administrator determines the issue is GPU related and then increases the frame buffer on the virtual machines. Testing confirms the issue is solved, and everything is now working correctly. Which of the following should the administrator do NEXT?

A. Consult corporate policies to ensure the fix is allowed
B. Conduct internal and external research based on the symptoms
C. Document the solution and place it in a shared knowledge base
D. Establish a plan of action to resolve the issue

Answer: C

Explanation:
Documenting the solution and placing it in a shared knowledge base is what the administrator should do next after troubleshooting performance issues with a VDI (Virtual Desktop Infrastructure) environment, determining that the issue is GPU (Graphics Processing Unit) related, increasing the frame buffer on the virtual machines, and testing that confirms that the issue is solved and everything is now working correctly. Documenting the solution is a process of recording and describing what was done to fix or resolve an issue, such as actions, steps, methods, etc., as well as why and how it worked. Placing it in a shared knowledge base is a process of storing and organizing documented solutions in a central location or repository that can be accessed and used by others. Documenting the solution and placing it in a shared knowledge base can provide benefits such as:
? Learning: Documenting the solution and placing it in a shared knowledge base can help to learn from past experiences and improve skills and knowledge.
? Sharing: Documenting the solution and placing it in a shared knowledge base can help to share information and insights with others who may face similar issues or situations.
? Reusing: Documenting the solution and placing it in a shared knowledge base can help to reuse existing solutions for future issues or situations.

NEW QUESTION 42
- (Topic 2)
A company needs a solution to find content in images. Which of the following technologies, when used in conjunction with cloud services, would facilitate the BEST solution?

A. Internet of Things
B. Digital transformation
C. Artificial intelligence
D. DNS over TLS

Answer: C

Explanation:
Artificial intelligence (AI) is the technology that, when used in conjunction with cloud services, would facilitate the best solution for finding content in images. AI is a branch of computer science that aims to create machines or systems that can perform tasks that normally require human intelligence, such as reasoning, learning, decision making, etc. AI can be used to analyze images and extract information such as objects, faces, text, emotions, etc., using techniques such as computer vision, machine learning, natural language processing, etc. AI can help to find content in images faster, more accurately, and more efficiently than manual methods.

NEW QUESTION 45
- (Topic 2)
A company wants to move its environment from on premises to the cloud without vendor lock-in. Which of the following would BEST meet this requirement?

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

Answer: C

Explanation:
IaaS (Infrastructure as a Service) is what would best meet the requirement of moving an environment from on premises to the cloud without vendor lock-in. Vendor lock- in is a situation where customers become dependent on or tied to a specific vendor or provider for their products or services, and face difficulties

NEW QUESTION 49
- (Topic 2)
A systems administrator is deploying a solution that includes multiple network I/O-intensive VMs. The solution design requires that vNICs of the VMs provide low-latency, near-native performance of a physical NIC and data protection between the VMs. Which of the following would BEST satisfy these requirements?

A. SR-IOV
B. GENEVE
C. SDN
D. VLAN

Answer: A

Explanation:
SR-IOV (Single Root Input/Output Virtualization) is what would best satisfy the requirements of low-latency, near-native performance of a physical NIC and data protection between VMs for multiple network I/O-intensive VMs. SR-IOV is a technology that allows a physical NIC to be partitioned into multiple virtual NICs that can be assigned to different VMs. SR-IOV can provide the following benefits:
? Low-latency: SR-IOV can reduce latency by bypassing the hypervisor and allowing direct communication between the VMs and the physical NIC, without any overhead or interference.
? Near-native performance: SR-IOV can provide near-native performance by allowing the VMs to use the full capacity and functionality of the physical NIC, without any emulation or translation.
? Data protection: SR-IOV can provide data protection by isolating and securing the network traffic between the VMs and the physical NIC, without any exposure or leakage.

NEW QUESTION 54
- (Topic 2)
A systems administrator is creating a VM and wants to ensure disk space is not allocated to the VM until it is needed. Which of the following techniques should the administrator use to ensure?

A. Deduplication
B. Thin provisioning
C. Software-defined storage

D. iSCSI storage

Answer: B

Explanation:
Thin provisioning is the technique that ensures disk space is not allocated to the VM until it is needed. Thin provisioning is a storage allocation method that assigns disk space to a VM on demand, rather than in advance. Thin provisioning can improve storage utilization and efficiency by avoiding overprovisioning and wasting disk space. Thin provisioning can also allow for more flexibility and scalability of storage resources.

NEW QUESTION 57
- (Topic 2)
Which of the following would be the BEST option for discussion of what individuals should do in an incident response or disaster recovery scenario?

A. A business continuity plan
B. Incident response/disaster recovery documentation
C. A tabletop exercise
D. A root cause analysis

Answer: C

Explanation:
A tabletop exercise is the best option for discussion of what individuals should do in an incident response or disaster recovery scenario. A tabletop exercise is a simulated scenario that involves key stakeholders and decision-makers who review and discuss their roles and responsibilities in response to an emergency situation or event. A tabletop exercise can help to test and evaluate plans, procedures, policies, training, and communication.

NEW QUESTION 59
- (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 62
- (Topic 2)
A company recently experienced a power outage that lasted 30 minutes. During this time, a whole rack of servers was inaccessible, even though the servers did not lose power.
Which of the following should be investigated FIRST?

A. Server power
B. Rack power
C. Switch power
D. SAN power

Answer: C

Explanation:
If a whole rack of servers was inaccessible during a power outage, even though the servers did not lose power, it is likely that the switch that connects them to the network lost power. Without network connectivity, the servers would not be able to communicate with other devices or services. The administrator should investigate the switch power source and ensure it has a backup power supply or UPS.

NEW QUESTION 67
- (Topic 2)
A company needs to access the cloud administration console using its corporate identity. Which of the following actions would MOST likely meet the requirements?

A. Implement SSH key-based authentication.
B. Implement cloud authentication with local LDAP.
C. Implement multifactor authentication.
D. Implement client-based certificate authentication.

Answer: D

Explanation:
Implementing client-based certificate authentication is what the administrator should do to access the cloud administration console using corporate identity. Client-based certificate authentication is a method of verifying and authenticating users or devices based on digital certificates issued by a trusted authority. Digital certificates are electronic 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. Client-based certificate authentication can allow users or devices to access cloud resources or services using their corporate identity without requiring passwords or other credentials.

NEW QUESTION 71
- (Topic 2)
A systems administrator is about to deploy a new VM to a cloud environment. Which of the following will the administrator MOST likely use to select an address for the VM?

A. CDN
B. DNS
C. NTP
D. IPAM

Answer: D

Explanation:
IPAM (IP Address Management) is what the administrator will most likely use to select an address for the new VM that is about to be deployed to a cloud environment. IPAM is a tool or service that allows customers to plan, track, and manage the IP addresses and DNS names of their cloud resources or systems. IPAM can help to select an address for the new VM by providing information such as available IP addresses, IP address ranges, subnets, domains, etc., as well as ensuring that the address is unique and valid.

NEW QUESTION 74
- (Topic 2)
A cloud engineer is responsible for managing a public cloud environment. There is currently one virtual network that is used to host the servers in the cloud environment. The environment is rapidly growing, and the network does not have any more available IP addresses. Which of the following should the engineer do to accommodate additional servers in this environment?

A. Create a VPC and peer the networks.
B. Implement dynamic routing.
C. Enable DHCP on the networks.
D. Obtain a new IPAM subscription.

Answer: A

Explanation:
Creating a VPC (Virtual Private Cloud) and peering the networks is the best option to accommodate additional servers in a public cloud environment that has run out of IP addresses. A VPC is a logically isolated section of a cloud provider's network that allows customers to launch and configure their own virtual network resources. Peering is a process of connecting two VPCs together so that they can communicate with each other as if they were in the same network.

NEW QUESTION 75
- (Topic 1)
A web server has been deployed in a public IaaS provider and has been assigned the public IP address of 72.135.10.100. Users are now reporting that when they browse to the website, they receive a message indicating the service is unavailable. The cloud administrator logs into the server, runs a netstat command, and notices the following relevant output:

```
TCP  17.3.130.3:0 72.135.10.100:5500 TIME_WAIT
TCP  17.3.130.3:0 72.135.10.100:5501 TIME_WAIT
TCP  17.3.130.3:0 72.135.10.100:5502 TIME_WAIT
TCP  17.3.130.3:0 72.135.10.100:5503 TIME_WAIT
TCP  17.3.130.3:0 72.135.10.100:5504 TIME_WAIT
```

Which of the following actions should the cloud administrator take to resolve the issue?

A. Assign a new IP address of 192.168.100.10 to the web server
B. Modify the firewall on 72.135.10.100 to allow only UDP
C. Configure the WAF to filter requests from 17.3.130.3
D. Update the gateway on the web server to use 72.135.10.1

Answer: D

Explanation:
Updating the gateway on the web server to use 72.135.10.1 is the best action to take to resolve the issue of the web server being unavailable after being deployed in a public IaaS provider and assigned the public IP address of 72.135.10.100. Updating the gateway can ensure that the web server can communicate with the Internet and other networks by using the correct router or device that connects the web server's network to other networks. Updating the gateway can also improve performance and reliability, as it can avoid any routing errors or conflicts that may prevent the web server from responding to remote login requests.
References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8

NEW QUESTION 77
SIMULATION - (Topic 1)
A company has decided to scale its e-commerce application from its corporate datacenter to a commercial cloud provider to meet an anticipated increase in demand during an upcoming holiday.
The majority of the application load takes place on the application server under normal conditions. For this reason, the company decides to deploy additional application servers into a commercial cloud provider using the on-premises orchestration engine that installs and configures common software and network configurations.
The remote computing environment is connected to the on-premises datacenter via a site- to-site IPSec tunnel. The external DNS provider has been configured to use weighted round-robin routing to load balance connections from the Internet.
During testing, the company discovers that only 20% of connections completed successfully.
INSTRUCTIONS
Review the network architecture and supporting documents and fulfill these requirements: Part 1:
- Analyze the configuration of the following components: DNS, Firewall 1, Firewall 2, Router 1, Router 2, VPN and Orchestrator Server.

- Identify the problematic device(s).

Part 2:

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

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

Part 1:

Cloud Hybrid Network Diagram

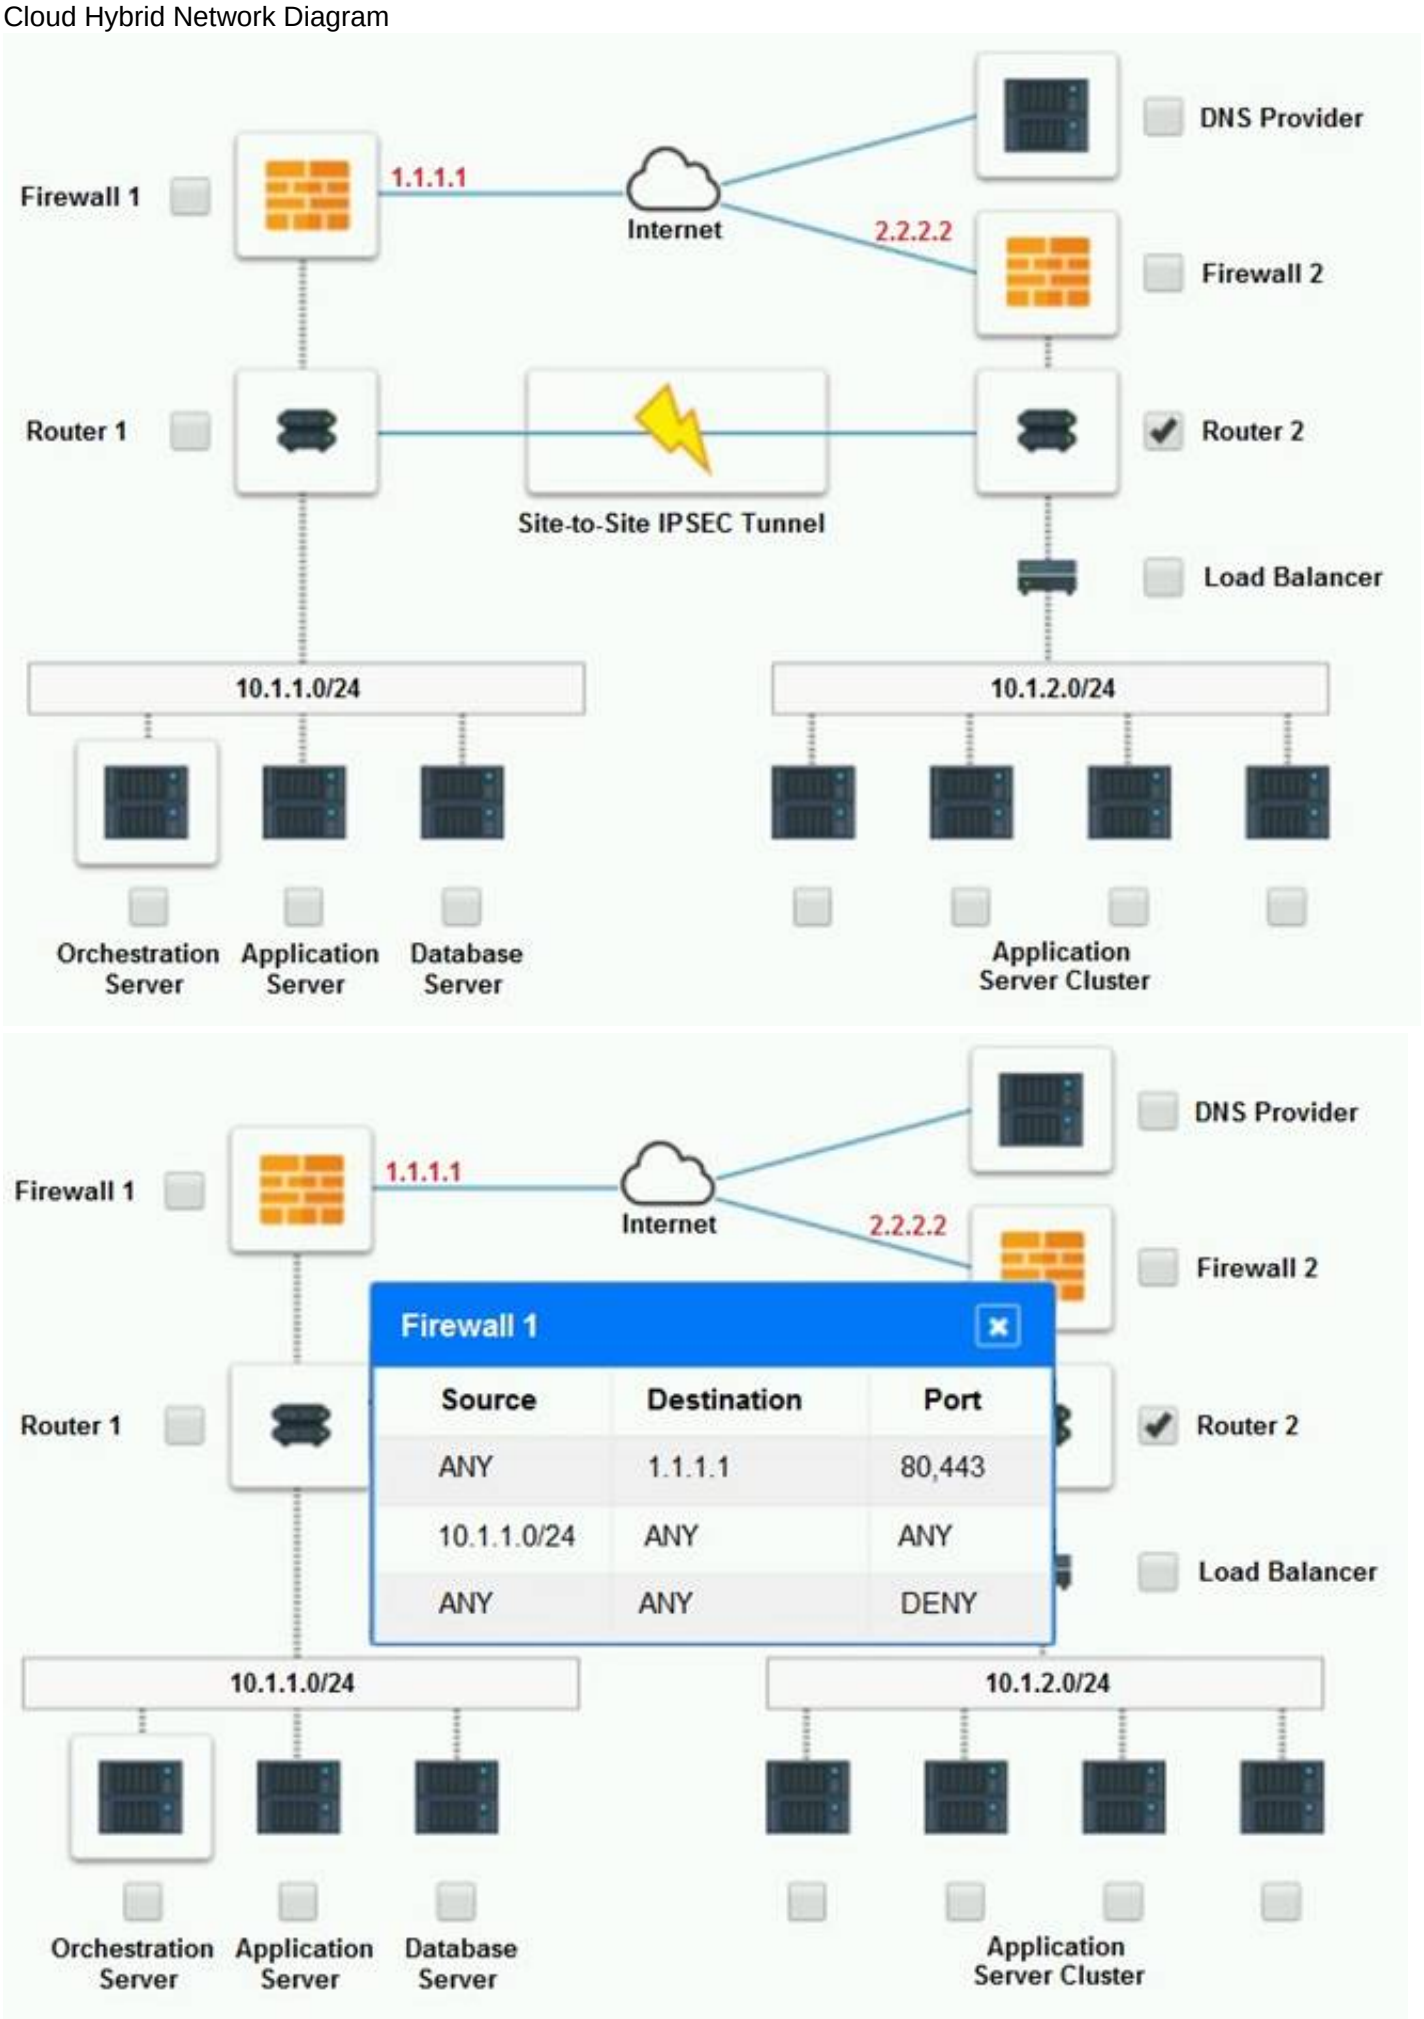

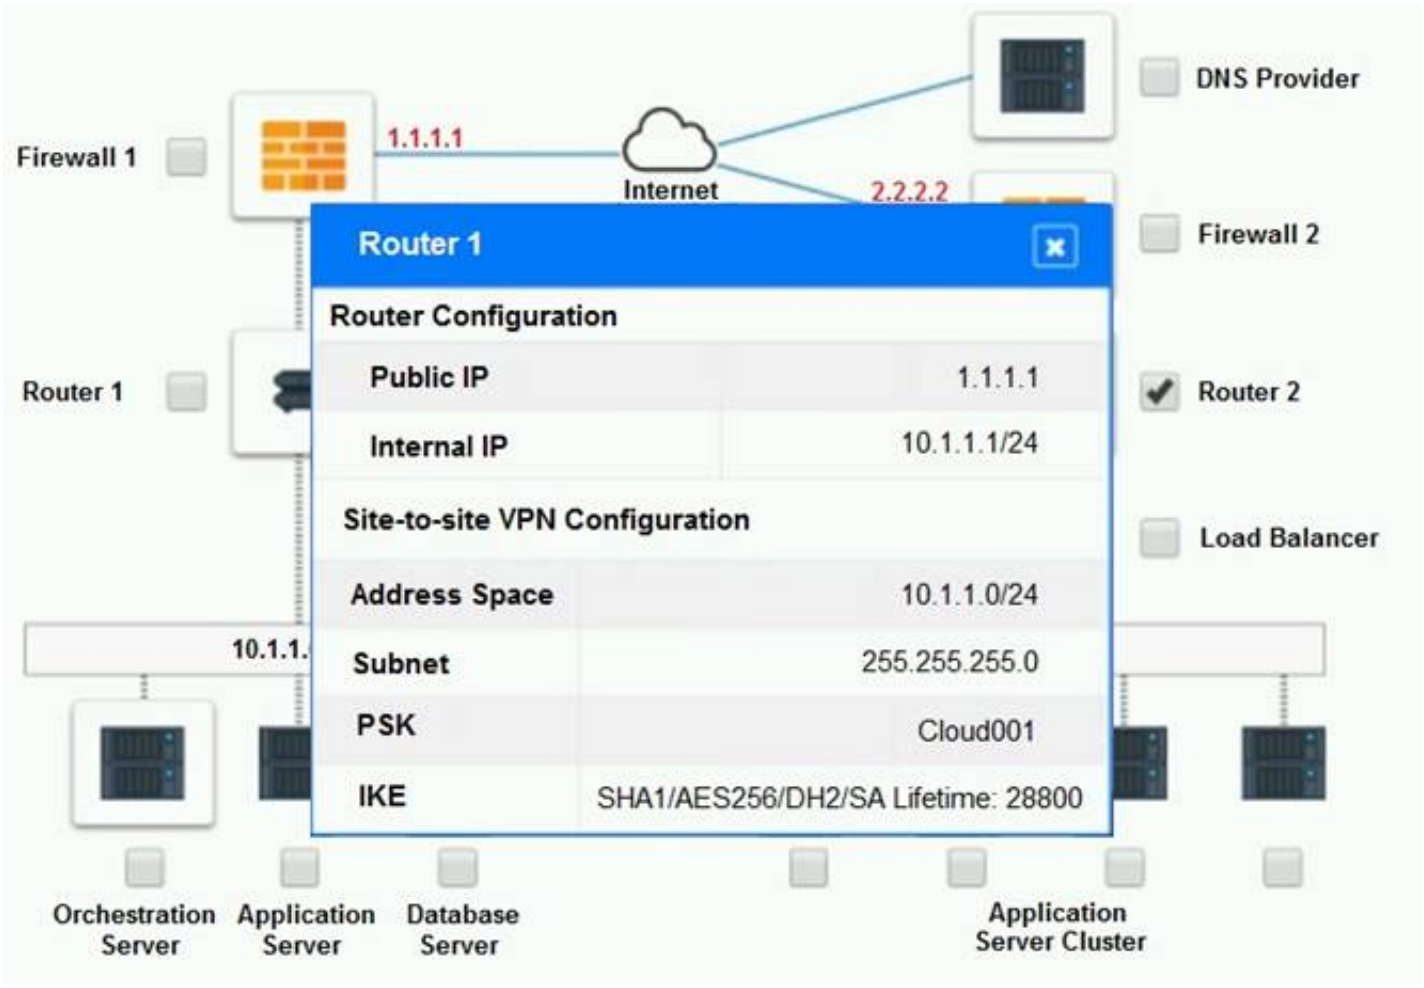
Firewall 1
1.1.1.1
Internet
2.2.2.2
DNS Provider
Firewall 2
Router 1
Router 2
Load Balancer
Router 1
Router Configuration
Public IP
1.1.1.1
Internal IP
10.1.1.1/24
Site-to-site VPN Configuration
Address Space
10.1.1.0/24
Subnet
255.255.255.0
PSK
Cloud001
IKE
SHA1/AES256/DH2/SA Lifetime: 28800
Orchestration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

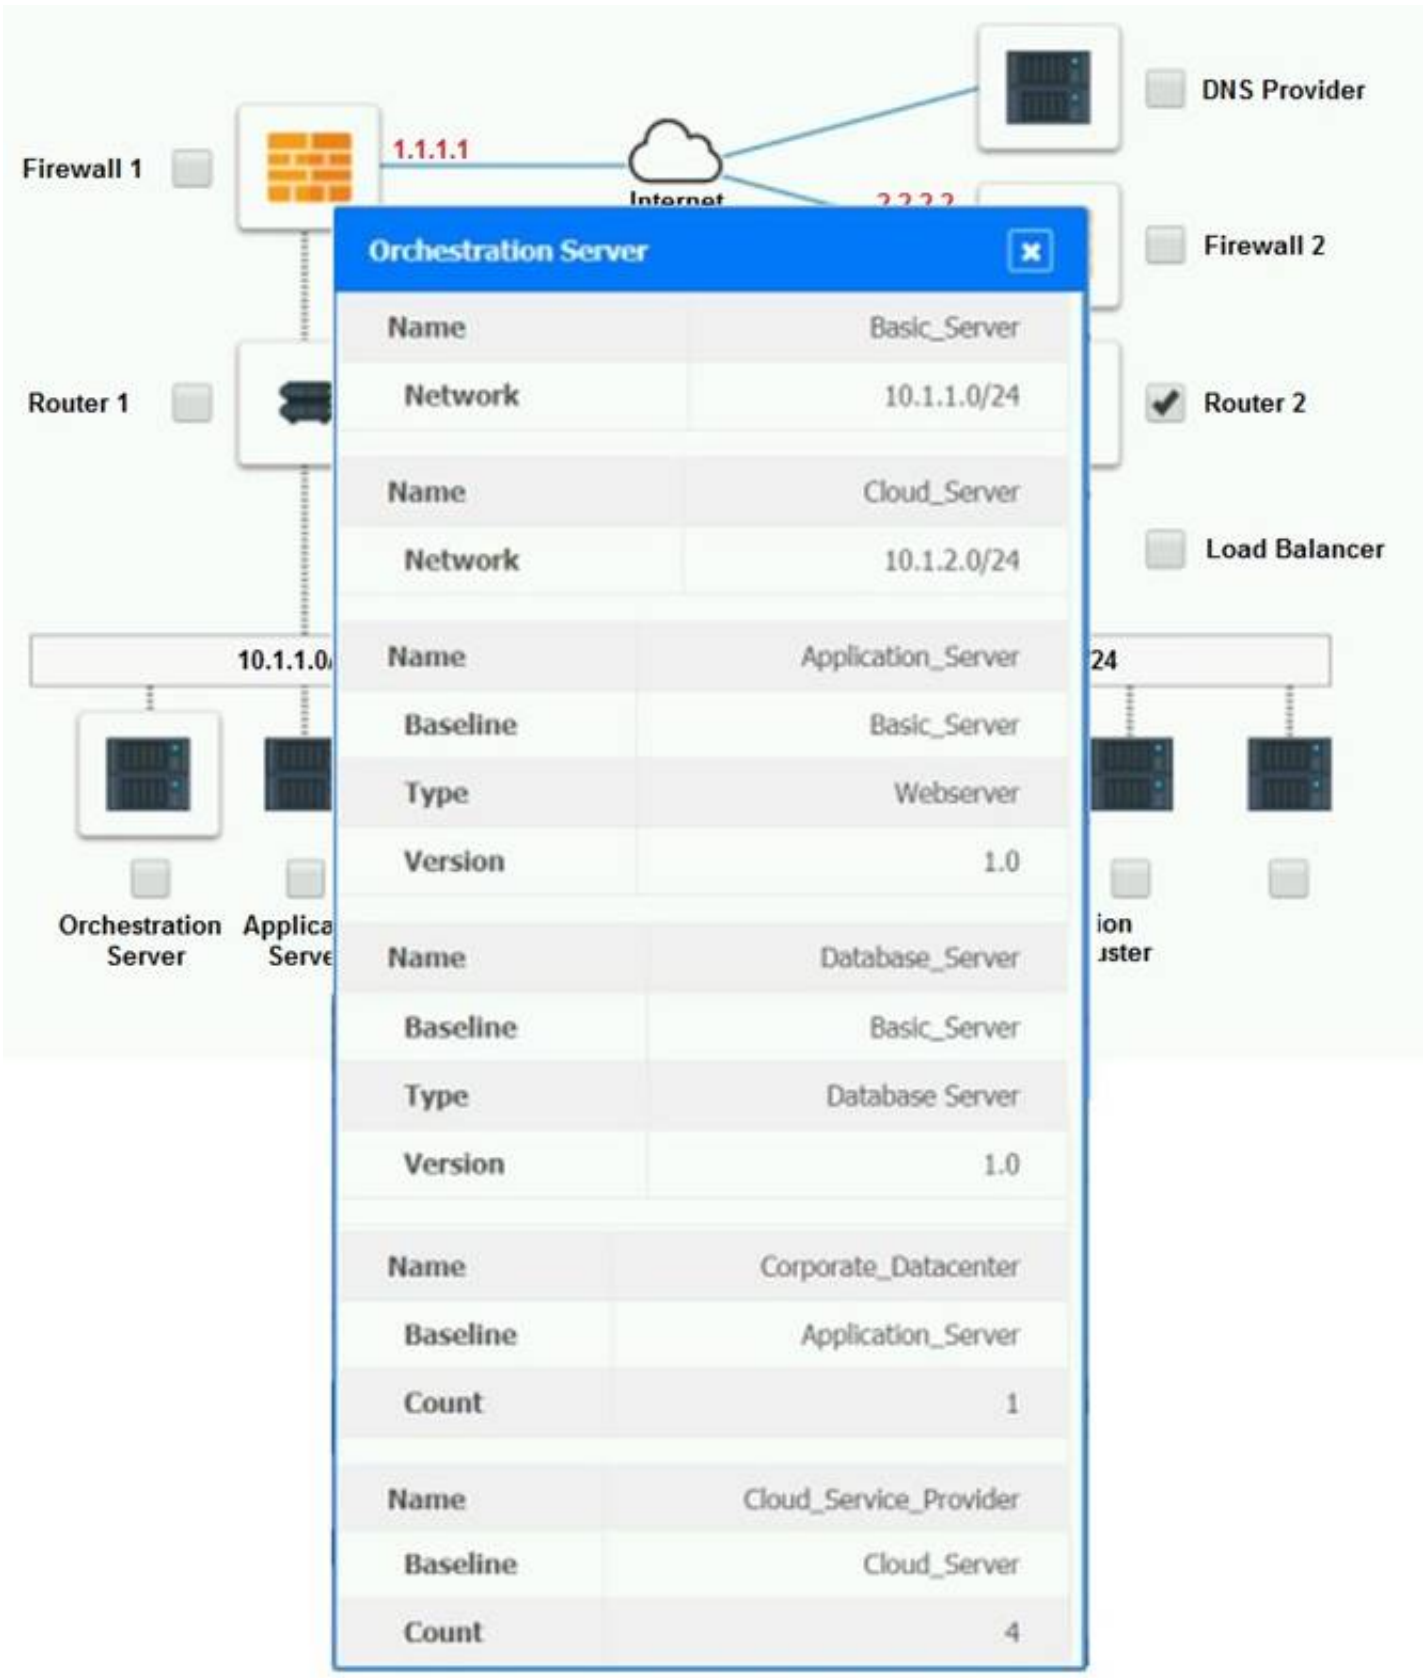
Firewall 1
1.1.1.1
Internet
DNS Provider
Firewall 2
Router 1
Router 2
Load Balancer
Orchestration Server
Name
Basic_Server
Network
10.1.1.0/24
Name
Cloud_Server
Network
10.1.2.0/24
Name
Application_Server
Baseline
Basic_Server
Type
Webserver
Version
1.0
Name
Database_Server
Baseline
Basic_Server
Type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

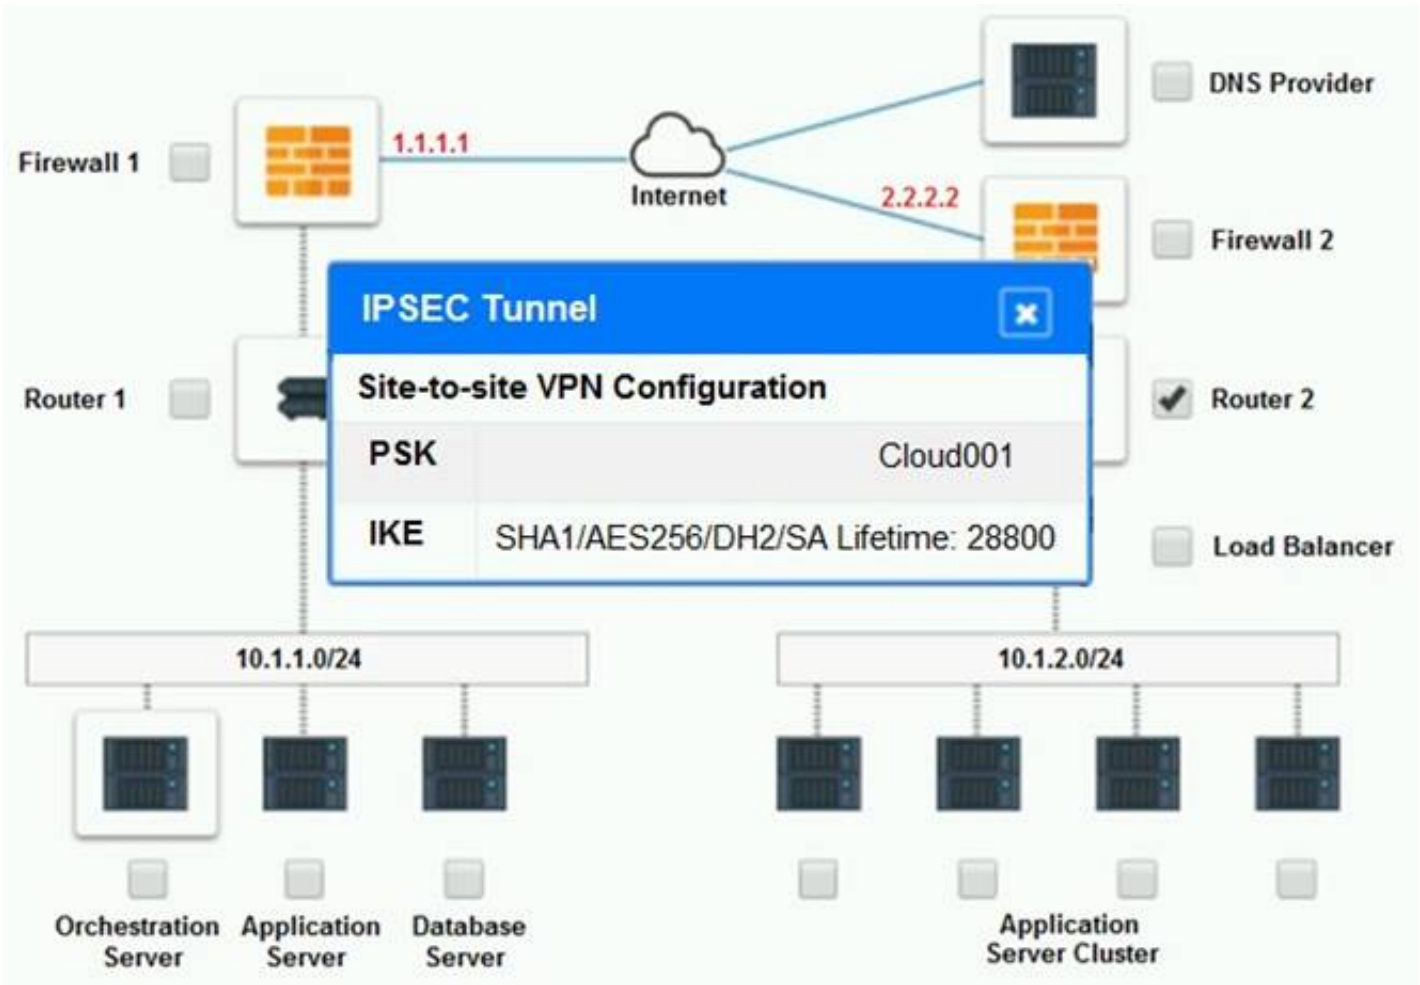
Firewall 1
1.1.1.1
Internet
2.2.2.2
DNS Provider
Firewall 2
IPSEC Tunnel
Site-to-site VPN Configuration
PSK
Cloud001
IKE
SHA1/AES256/DH2/SA Lifetime: 28800
Router 1
Router 2
Load Balancer
10.1.1.0/24
10.1.2.0/24
Orchestration Server
Application Server
Database Server
Application Server Cluster

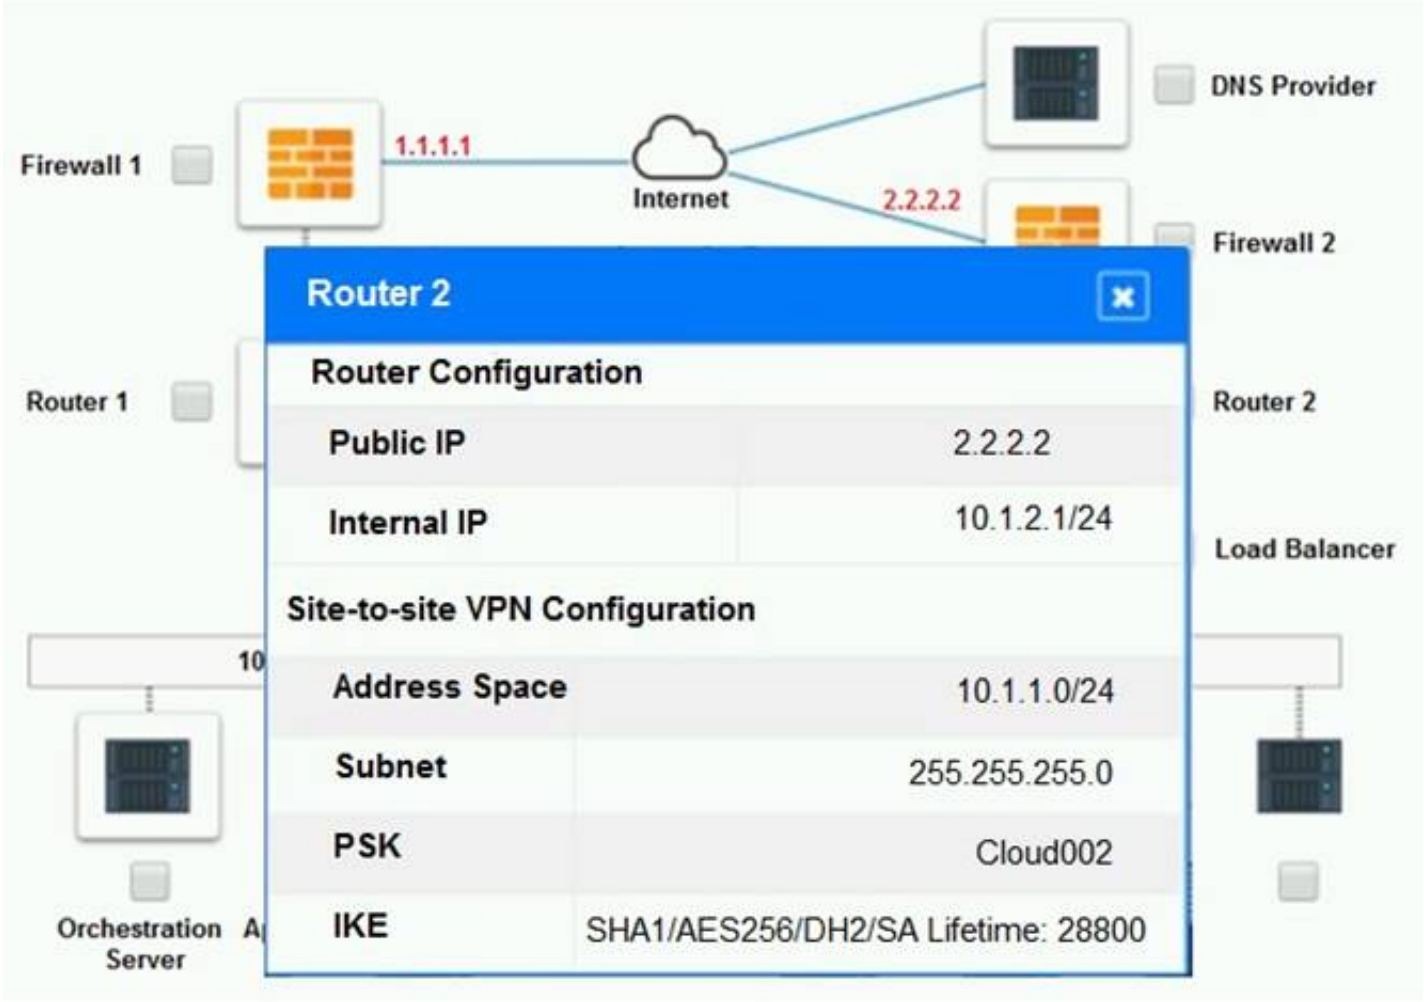
Firewall 1
1.1.1.1
Internet
2.2.2.2
DNS Provider
Firewall 2
Router 2
Router Configuration
Public IP
2.2.2.2
Internal IP
10.1.2.1/24
Site-to-site VPN Configuration
Address Space
10.1.1.0/24
Subnet
255.255.255.0
PSK
Cloud002
IKE
SHA1/AES256/DH2/SA Lifetime: 28800
Router 1
Router 2
Load Balancer
Orchestration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

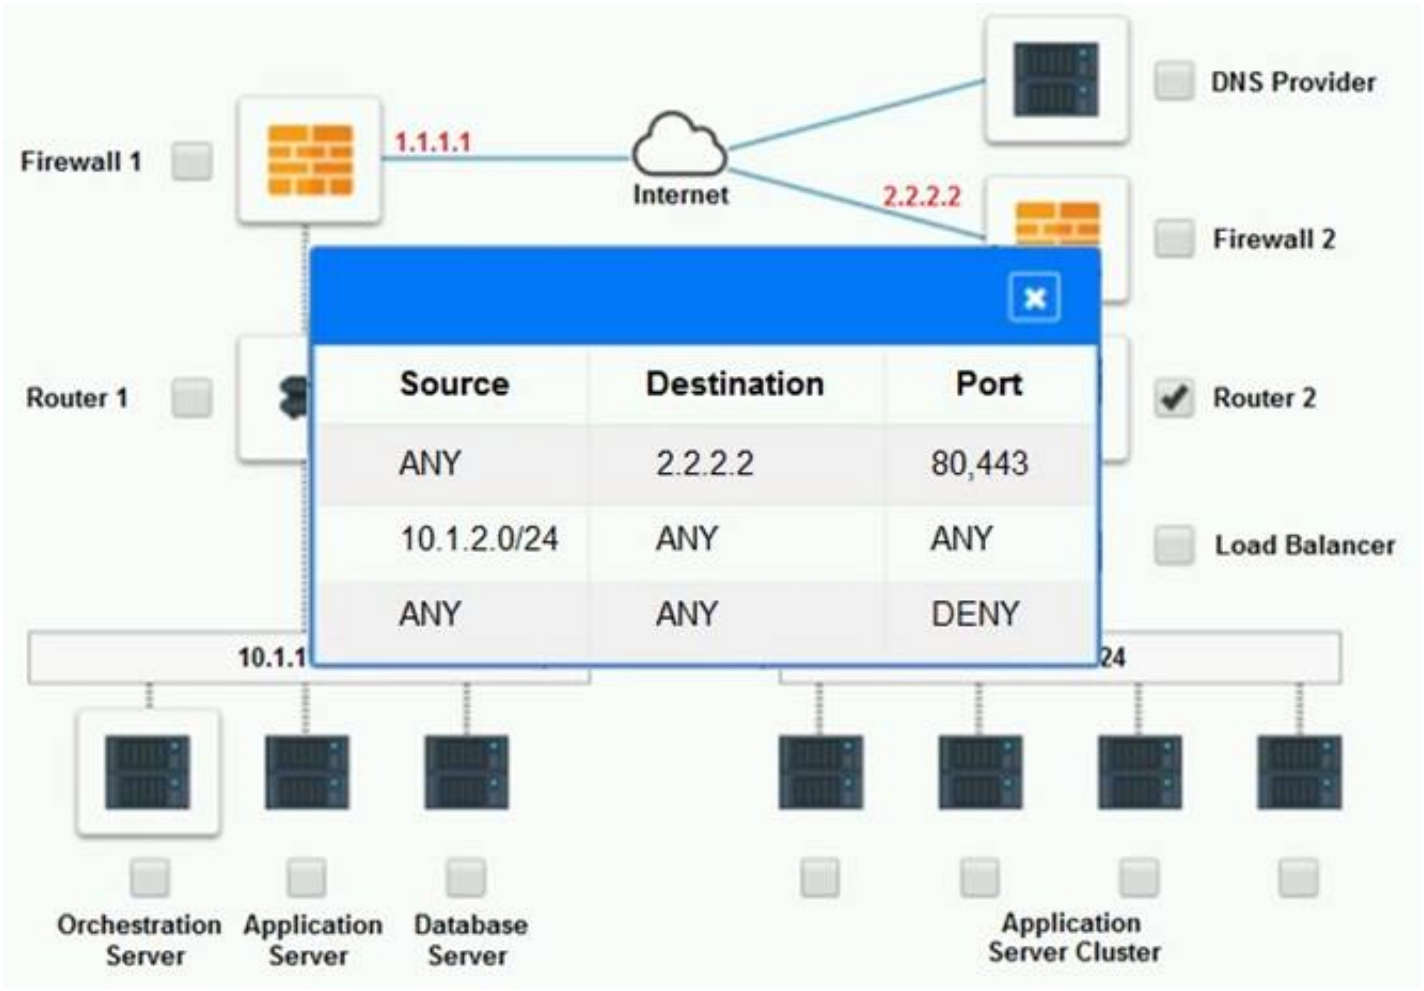

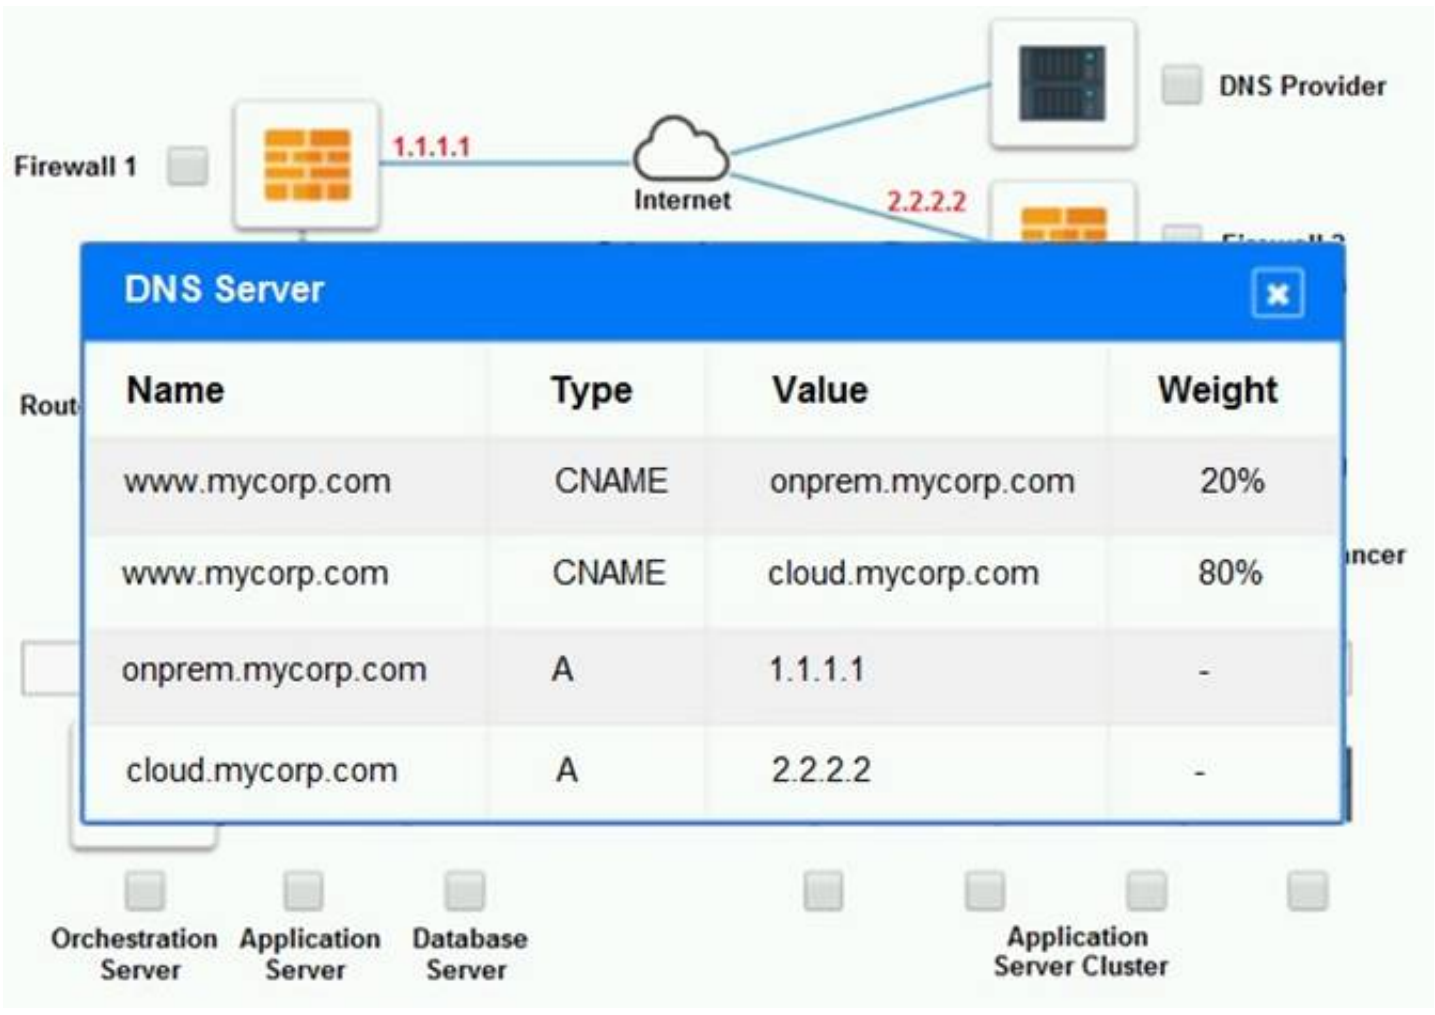

Part 2:
Only select a maximum of TWO options from the multiple choice question

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Part 1: Router 2
The problematic device is Router 2, which has an incorrect configuration for the IPSec tunnel. The IPSec tunnel is a secure connection between the on-premises datacenter and the cloud provider, which allows the traffic to flow between the two networks. The IPSec tunnel requires both endpoints to have matching parameters, such as the IP addresses, the pre-shared key (PSK), the encryption and authentication algorithms, and the security associations (SAs) .
According to the network diagram and the configuration files, Router 2 has a different PSK and a different address space than Router 1. Router 2 has a PSK of "1234567890", while Router 1 has a PSK of "0987654321". Router 2 has an address space of 10.0.0.0/8, while Router 1 has an address space of 192.168.0.0/16. These mismatches prevent the IPSec tunnel from establishing and encrypting the traffic between the two networks.
The other devices do not have any obvious errors in their configuration. The DNS provider has two CNAME records that point to the application servers in the cloud provider, with different weights to balance the load. The firewall rules allow the traffic from and to the application servers on port 80 and port 443, as well as the traffic from and to the VPN server on port 500 and port 4500. The orchestration server has a script that installs and configures the application servers in the cloud provider, using the DHCP server to assign IP addresses.
Part 2:
The correct options to provide adequate configuration for hybrid cloud architecture are:
? Update the PSK in Router 2.
? Change the address space on Router 2.
These options will fix the IPSec tunnel configuration and allow the traffic to flow between the on-premises datacenter and the cloud provider. The PSK should match the one on Router 1, which is "0987654321". The address space should also match the one on Router 1, which is 192.168.0.0/16.
* B. Update the PSK (Pre-shared key in Router2)
* E. Change the Address Space on Router2

NEW QUESTION 78
- (Topic 1)
A systems administrator is informed that a database server containing PHI and PII is unencrypted. The environment does not support VM encryption, nor does it have a key management system. The server needs to be able to be rebooted for patching without manual intervention.
Which of the following will BEST resolve this issue?

A. Ensure all database queries are encrypted
B. Create an IPSec tunnel between the database server and its clients
C. Enable protocol encryption between the storage and the hypervisor
D. Enable volume encryption on the storage
E. Enable OS encryption

Answer: D

Explanation:
Volume encryption is a type of encryption that protects data at the storage level by encrypting an entire disk or partition. Volume encryption can provide strong security for data at rest, as it prevents unauthorized access to the data even if the storage device is lost, stolen, or compromised. Volume encryption can also support automatic booting without manual intervention, as it can use a pre-boot authentication mechanism that does not require user input. Enabling volume encryption on the storage is the best way to resolve the issue of having an unencrypted database server containing PHI and PII, as it can protect the sensitive data without relying on VM encryption or a key management
system. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 82
- (Topic 1)
A cloud engineer is responsible for managing two cloud environments from different MSPs. The security department would like to inspect all traffic from the two cloud environments.
Which of the following network topology solutions should the cloud engineer implement to reduce long-term maintenance?

A. Chain
B. Star
C. Mesh
D. Hub and spoke

Answer: D

Explanation:
Hub and spoke is a type of network topology that consists of a central node or device (hub) that connects to multiple peripheral nodes or devices (spokes). Hub and spoke can help reduce long-term maintenance for managing two cloud environments from different MSPs, as it can simplify and centralize the network configuration and management by using the hub as a single point of contact and control for the spokes. Hub and spoke can also improve network performance and security, as it can reduce latency, bandwidth consumption, and network congestion by routing traffic through the hub. References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8

NEW QUESTION 83
- (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 87
- (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 92
- (Topic 1)
An OS administrator is reporting slow storage throughput on a few VMs in a private IaaS cloud. Performance graphs on the host show no increase in CPU or memory. However, performance graphs on the storage show a decrease of throughput in both IOPS and MBps but not much increase in latency. There is no increase in workload, and latency is stable on the NFS storage arrays that are used by those VMs.
Which of the following should be verified NEXT?

A. Application
B. SAN
C. VM GPU settings
D. Network

Answer: D

Explanation:
The network is the set of devices, connections, protocols, and configurations that enable communication and data transfer between different systems and applications. The network can affect the performance of storage throughput by influencing factors such as bandwidth, latency, jitter, packet loss, and congestion. Poor network performance can result in low storage throughput in both IOPS and MBps, as it can limit the amount and speed of data that can be sent or received by the storage devices. Verifying the network should be the next step for troubleshooting the issue of slow storage throughput on a few VMs in a private IaaS cloud, as it can help identify and resolve any network-related problems that may be causing the issue. References: CompTIA Cloud+ Certification Exam Objectives, page 17, section 3.4

NEW QUESTION 95
- (Topic 1)
A systems administrator recently deployed a VDI solution in a cloud environment; however, users are now experiencing poor rendering performance when trying to display 3-D content on their virtual desktops, especially at peak times.
Which of the following actions will MOST likely solve this issue?

A. Update the quest graphics drivers from the official repository
B. Add more vGPU licenses to the host
C. Instruct users to access virtual workstations only on the VLAN
D. Select vGPU profiles with higher video RAM

Answer: D

Explanation:
A vGPU profile is a configuration option that defines the amount of video RAM (vRAM) and other resources that are allocated to a virtual machine (VM) that uses a virtual graphics processing unit (vGPU). A vGPU profile can affect the rendering performance of a VM, as it determines how much graphics memory and processing power are available for displaying complex graphics content. Selecting vGPU profiles with higher video RAM can most likely solve the issue of poor rendering performance when trying to display 3-D content on virtual desktops, especially at peak times, as it can provide more graphics resources and improve the quality and speed of rendering. References: CompTIA Cloud+ Certification Exam Objectives, page 11, section 1.6

NEW QUESTION 100
- (Topic 1)
A systems administrator for an e-commerce company will be migrating the company's main website to a cloud provider. The principal requirement is that the website must be highly available.
Which of the following will BEST address this requirement?

A. Vertical scaling
B. A server cluster
C. Redundant switches

D. A next-generation firewall

Answer: B

Explanation:
A server cluster is a group of servers that work together to provide high availability, load balancing, and scalability for applications or services. A server cluster can help ensure the high availability requirement for migrating an e-commerce company's main website to a cloud provider, as it can prevent downtime or disruption in case of a server failure or outage by automatically switching the workload to another server in the cluster. A server cluster can also improve performance and reliability, as it can distribute the workload across multiple servers and handle increased traffic or demand. References: CompTIA Cloud+ Certification Exam Objectives, page 10, section 1.5

NEW QUESTION 105
- (Topic 1)
A systems administrator is configuring a storage array.
Which of the following should the administrator configure to set up mirroring on this array?

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

Answer: B

Explanation:
RAID 1 is a type of RAID level that creates an exact copy or mirror of data on two or more disks. RAID 1 can provide redundancy and fault tolerance, as it can survive the failure of one disk without losing any data. RAID 1 can also improve read performance, as it can access data from multiple disks simultaneously. The administrator should configure RAID 1 to set up mirroring on a storage array. References: CompTIA Cloud+ Certification Exam Objectives, page 9, section 1.4

NEW QUESTION 108
- (Topic 4)
A cloud administrator is evaluating a solution that will limit access to authorized individuals. The solution also needs to ensure the system that connects to the environment meets patching, antivirus, and configuration requirements. Which of the following technologies would BEST meet these requirements?

A. NAC
B. EDR
C. IDS
D. HIPS

Answer: A

Explanation:

NAC (Network Access Control) is a technology that will limit access to authorized individuals and ensure the system that connects to the environment meets patching, antivirus, and configuration requirements. NAC can enforce policies and rules that define who, what, when, where, and how a device or a user can access a network or a cloud environment. NAC can also inspect and evaluate the security posture and compliance status of a device or a user before granting or denying access. For example, NAC can check if the device has the latest patches, antivirus software, and configuration settings, and if not, it can quarantine, remediate, or reject the device. NAC can also monitor and audit the ongoing network activity and behavior of the devices and users, and take actions if any violations or anomalies are detected.

NEW QUESTION 111
- (Topic 4)
A systems administrator is deploying a new version of a website. The website is deployed in the cloud using a VM cluster. The administrator must then deploy the new version into one VM first. After a period of time, if there are no issues detected, a second VM will be updated. This process must continue until all the VMS are updated. Which of the following upgrade methods is being implemented?

A. Canary
B. Blue-green
C. Rolling
D. Staging

Answer: C

Explanation:
The upgrade method that is being implemented by the systems administrator is rolling. A rolling upgrade is a type of upgrade that applies the new version of a software or service to a subset of nodes or instances at a time, while the rest of the nodes or instances continue to run the old version. This way, the upgrade can be performed gradually and incrementally, without causing downtime or disruption to the entire system. A rolling upgrade can also help to monitor and test the new version for any issues or errors, and roll back to the old version if needed12.
A canary upgrade is a type of upgrade that applies the new version of a software or service to a small and selected group of users or customers, before rolling it out to the rest of the population. This way, the upgrade can be evaluated for its performance, functionality, and feedback, and any problems or bugs can be fixed before affecting the majority of users or customers34.
A blue-green upgrade is a type of upgrade that involves having two identical environments, one running the old version (blue) and one running the new version (green) of a software or service. The traffic is switched from the blue environment to the green environment once the new version is ready and tested. This way, the upgrade can be performed quickly and seamlessly, without any downtime or risk of failure. The blue environment can also serve as a backup in case of any issues with the green environment5 .
A staging upgrade is a type of upgrade that involves having a separate environment that mimics the production environment, where the new version of a software or service is deployed and tested before moving it to the production environment. This way, the upgrade can be verified and validated for its compatibility, security, and quality, and any defects or errors can be resolved before affecting the live system .

NEW QUESTION 114

- (Topic 4)
A company uses multiple SaaS-based cloud applications. All the applications require authentication upon access. An administrator has been asked to address this issue and enhance security. Which of the following technologies would be the BEST solution?

A. Single sign-on
B. Certificate authentication
C. Federation
D. Multifactor authentication

Answer: A

Explanation:
Single sign-on (SSO) is a technology that allows a user to access multiple applications or services with a single login and authentication process. SSO can enhance security by reducing the number of passwords that a user has to remember and enter, and by enabling centralized management and enforcement of security policies .
SSO can help address the issue of multiple SaaS-based cloud applications requiring authentication upon access. By implementing SSO, an administrator can:
Simplify the user experience and increase productivity by eliminating the need to enter multiple usernames and passwords for different applications .
Improve the security and compliance of the applications by using a trusted identity provider (IdP) that can verify the user's identity and credentials, and grant or deny access based on predefined rules .
Reduce the risk of password breaches, phishing, or identity theft by minimizing the exposure of passwords to third-party applications or malicious actors .

NEW QUESTION 119
- (Topic 4)
A cloud administrator must ensure all servers are in compliance with the company's security policy Which of the following should the administrator check FIRST?

A. The application version
B. The OS version
C. Hardened baselines
D. Password policies

Answer: C

Explanation:
Hardened baselines are a set of security best practices that reduce the vulnerability of a system to exploits by reducing its attack surface1. They are also known as security configurations or benchmarks, and they provide a standard level of system hardening for an organization23.
Checking the hardened baselines of the servers is the first step that a cloud administrator should take to ensure compliance with the company's security policy. This is because hardened baselines can help to:
Identify and eliminate common vulnerabilities and exposures (CVEs) that attackers can exploit1.
Remove unnecessary or unused services, accounts, software, and ports that can increase the attack surface23.
Apply appropriate settings and controls for encryption, authentication, authorization, firewall, and logging23.
Streamline audits and testing by reducing complexity and providing a reliable benchmark23.

NEW QUESTION 123
- (Topic 4)
A company's marketing department is running a rendering application on virtual desktops. Currently, the application runs slowly, and it takes a long time to refresh the screen. The virtualization administrator is tasked with resolving this issue. Which of the following is the BEST solution?

A. GPU passthrough
B. Increased memory
C. Converged infrastructure
D. An additional CPU core

Answer: A

Explanation:
GPU passthrough is a technique that allows a virtual machine to access and use the physical GPU of the host machine directly. This can improve the performance and quality of graphics-intensive applications, such as rendering, gaming, or video editing, that run on the virtual machine123.
GPU passthrough can help resolve the issue of the rendering application running slowly and taking a long time to refresh the screen on the virtual desktops. By enabling GPU passthrough, the virtualization administrator can allow the rendering application to leverage the full power and features of the host GPU, rather than relying on the limited and shared resources of a virtual GPU. This can result in faster rendering, smoother animations, and higher resolution12

NEW QUESTION 124
- (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 126
- (Topic 4)
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 I, 2. and 5.
B. Remove rules I, 3, and 4.
C. Remove rules 2.3. and 4.
D. Remove rules 3.4. and 5.

Answer: B

Explanation:
The correct answer is B. Remove rules 1, 3, and 4.
The objective is to ensure the web servers in the public subnet allow only secure communications. This means that only HTTPS traffic should be allowed on port 443, which is the standard port for secure web connections. HTTPS traffic uses the TCP protocol and encrypts the data between the client and the server.
Rule 1 allows all TCP traffic on any port from any source. This is too permissive and exposes the web servers to potential attacks or unauthorized access. Rule 1 should be removed to restrict the TCP traffic to only port 443.
Rule 3 allows all UDP traffic on any port from any source. UDP is a connectionless protocol that does not guarantee reliable or secure delivery of data. UDP is typically used for streaming media, voice over IP (VoIP), or online gaming, but not for web servers. Rule 3 should be removed to prevent unnecessary or malicious UDP traffic.
Rule 4 allows all ICMP traffic from any source. ICMP is a protocol that is used for diagnostic or control purposes, such as ping or traceroute. ICMP traffic can be used by attackers to scan or probe the network for vulnerabilities or information. Rule 4 should be removed to block ICMP traffic and reduce the attack surface.
Rule 2 allows TCP traffic on port 443 from any source. This is the desired rule that allows secure web communications using HTTPS. Rule 2 should be kept.
Rule 5 denies all other traffic that does not match any of the previous rules. This is the default rule that provides a catch-all protection for the web servers. Rule 5 should be kept. Therefore, the analyst should remove rules 1, 3, and 4 to accomplish the objective.

NEW QUESTION 127
- (Topic 4)
A cloud administrator is having difficulty correlating logs for multiple servers. Upon inspection, the administrator finds that the time-zone settings are mismatched throughout the deployment. Which of the following solutions can help maintain time synchronization between all the resources?

A. DNS
B. IPAM
C. NTP
D. SNMP

Answer: C

Explanation:
The correct answer is C. NTP.
NTP stands for Network Time Protocol, which is a standard protocol for synchronizing the clocks of computers over a network. NTP uses a hierarchical, client-server architecture, where a client requests the current time from a server, and the server responds with a timestamp. The client then adjusts its own clock to match the server's time, taking into account the network delay and clock drift. NTP can achieve sub-millisecond accuracy over local area networks and a few milliseconds over the internet12.
NTP can help maintain time synchronization between all the resources in a distributed cloud environment, as it allows each resource to get the accurate time from a reliable source. This can help with correlating logs, auditing, security, and other time-sensitive operations. NTP can also handle different time zones, as it uses

Coordinated Universal Time (UTC) as the reference time, and each resource can convert UTC to its local time zone12.
DNS stands for Domain Name System, which is a protocol for resolving domain names into IP addresses. DNS does not provide any functionality for time synchronization3.
IPAM stands for IP Address Management, which is a method for planning, tracking, and managing the IP address space used in a network. IPAM does not provide any functionality for time synchronization.
SNMP stands for Simple Network Management Protocol, which is a protocol for collecting and organizing information about managed devices on a network. SNMP can be used to monitor the performance, availability, configuration, and security of network devices, but it does not provide any functionality for time synchronization.

NEW QUESTION 129
- (Topic 4)
A cloud administrator receives an email stating the following:
"Clients are receiving emails from our web application with non-encrypted links."
The administrator notices that links generated from the web application are opening in http://. Which of the following should be configured to redirect the traffic to https://?

A. User account access
B. Programming code
C. Web server configuration
D. Load balancer setting

Answer: C

Explanation:
To redirect the traffic from HTTP to HTTPS, the web server configuration should be modified to include a rule that forces the HTTP requests to be redirected to HTTPS. This can be done by using the web server's configuration file or a .htaccess file. The exact syntax may vary depending on the web server software, but the general idea is to use a rewrite rule that matches the HTTP protocol and changes it to HTTPS. For example, on Apache web server, the following code can be added to the .htaccess file: RewriteEngine On
RewriteCond %{HTTPS} off
RewriteRule ^(.*)$ https://%{HTTP_HOST}%{REQUEST_URI} [L,R=301]
This code will check if the HTTPS is off, and if so, it will rewrite the URL to use HTTPS and redirect the client with a 301 status code, which means permanent redirection. This way, the clients will always use HTTPS to access the web application, and the links generated from the web application will be encrypted.
User account access (A) is not relevant to the redirection of HTTP to HTTPS, as it only controls who can access the web application. Programming code (B) may be used to generate the links with HTTPS, but it will not redirect the existing HTTP requests to HTTPS. Load balancer setting (D) may also be used to redirect the traffic to HTTPS, but it is not the most efficient or secure way, as it will add an extra layer of processing and expose the HTTP traffic to the load balancer.
Therefore, web server configuration © is the best option to redirect the traffic to HTTPS.
Reference: The Official CompTIA Cloud+ Student Guide (Exam CV0-003), Chapter 4:
Cloud Security, Section 4.3: Secure Cloud Services, p. 4-23.

NEW QUESTION 130
- (Topic 4)
An integration application that communicates between different application and database servers is currently hosted on a physical machine. A P2V migration needs to be done to reduce the hardware footprint. Which of the following should be considered to maintain the same level of network throughput and latency in the virtual server?

A. Upgrading the physical server NICs to support 10Gbps
B. Adding more vCPU
C. Enabling SR-IOV capability
D. Increasing the VM swap/paging size

Answer: C

Explanation:
SR-IOV stands for Single Root I/O Virtualization, which is a technology that allows a physical network adapter to be partitioned into multiple virtual functions (VFs) that can be directly assigned to virtual machines (VMs). This way, the network traffic bypasses the software layer of the hypervisor and the virtual switch, and goes directly from the VM to the physical adapter. This reduces the CPU overhead, the network latency, and the packet loss, and improves the network throughput and scalability. SR-IOV can achieve near-native performance for network-intensive applications, such as an integration application that communicates between different application and database servers. By enabling SR-IOV capability on the physical server and the virtual server, the P2V migration can maintain the same level of network throughput and latency as the original physical machine. References: High performance network virtualization with SR-IOV; Supercharge Your Network Throughput via Single Root I/O Virtualization (SR-IOV); Overview of Single Root I/O Virtualization (SR-IOV).

NEW QUESTION 133
- (Topic 4)
A systems administrator is planning to migrate to a cloud solution with volume-based licensing. Which of the following is most important when considering licensing costs?

A. The number of cores
B. The number of threads
C. The number of machines
D. The number of sockets

Answer: C

Explanation:
Volume-based licensing is a model where the cost of the software is based on the number of licenses purchased1. This model is commonly used for software that is installed on a specific number of devices, such as antivirus software or office productivity suites1. Therefore, the number of machines is the most important factor when considering licensing costs in this model.
References: CompTIA Cloud+ CV0-003 Exam Objectives, Objective 1.2: Given a scenario, compare and contrast various cloud service models ; Cloud+ Exam CV0-003: CompTIA Cloud+ Licensing Models1

NEW QUESTION 136
- (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 137
- (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 140
- (Topic 4)
An organization provides integration services for finance companies that use web services. A new company that sends and receives more than 100,000 transactions per second has
been integrated using the web service. The other integrated companies are now reporting slowness with regard to the integration service. Which of the following is the cause of the issue?

A. Incorrect configuration in the authentication process
B. Incorrect configuration in the message queue length
C. Incorrect configuration in user access permissions
D. Incorrect configuration in the SAN storage pool

Answer: B

Explanation:
The correct answer is B. Incorrect configuration in the message queue length.
A message queue is a data structure that stores messages or requests that are sent and received by web services. A message queue allows asynchronous communication between web services, as it decouples the sender and the receiver, and enables them to process messages at different rates. A message queue also provides reliability, scalability, and load balancing for web services, as it ensures that messages are not lost, duplicated, or corrupted, and that they are distributed evenly among the available servers .
However, a message queue also has a limit on how many messages it can store at a time. This limit is determined by the configuration of the message queue length, which is the maximum number of messages that can be in the queue before it becomes full. If the message queue length is too short, the queue may fill up quickly and reject new messages, causing errors or delays in communication. If the message queue length is too long, the queue may consume too much memory or disk space, affecting the performance or availability of the web service .
Therefore, if an organization provides integration services for finance companies that use web services, and a new company that sends and receives more than 100,000 transactions per second has been integrated using the web service, the most likely cause of the issue is an incorrect configuration in the message queue length. The new company may have generated a large volume of messages that exceeded the capacity of the message queue, resulting in slowness for the other integrated companies. The organization should adjust the message queue length to accommodate the increased traffic and optimize the resource utilization of the web service.

NEW QUESTION 142
- (Topic 4)

A systems administrator receives a ticket stating the following:
“The programming team received an error during the process deploying applications to the container platform. The error after the containerized applications were created”
Which the following should the administrator Check FIRST?

A. The containers
B. The application
C. The Scripts
D. The templates

Answer: A

Explanation:
The correct answer is A. The containers.
The error that the programming team received indicates that the problem occurred after the containerized applications were created, but before they were deployed to the container platform. This suggests that the issue is related to the containers themselves, not the application, the scripts, or the templates.
The containers are the units of software that package up the application code and all its dependencies, such as libraries, frameworks, and configuration files. The containers run on a container platform, such as Docker or Kubernetes, that provides the runtime environment and orchestration for the containers. The containers are created from images, which are templates that define how to build and run a container.
The administrator should check the containers first to see if they are configured correctly, if they have any errors or warnings, if they have the necessary resources and permissions, and if they can communicate with each other and with the container platform. The administrator can use tools such as docker ps, docker logs, docker inspect, and docker exec to examine and troubleshoot the containers.

NEW QUESTION 144
- (Topic 4)
A systems administrator notices the host filesystem is running out of storage space. Which of the following will best reduce the storage space on the system?

A. Deduplication
B. Compression
C. Adaptive optimization
D. Thin provisioning

Answer: A

Explanation:
Deduplication is a technique that reduces the storage space by eliminating duplicate data blocks and replacing them with pointers to the original data. Deduplication can help free up the host filesystem by removing redundant data and increasing the storage efficiency. Deduplication can be performed at the source or the target, and it can be applied at the file or block level. References: [CompTIA Cloud+ CV0-003 Certification Study Guide], Chapter 4, Objective 4.3: Given a scenario, troubleshoot common storage issues.

NEW QUESTION 148
- (Topic 4)
A company has a web application that is accessed around the world. An administrator has been notified of performance issues regarding the application. Which of the following will BEST improve performance?

A. IPAM
B. SDN
C. CDN
D. VPN

Answer: C

Explanation:
The correct answer is C. CDN.
A CDN, or content delivery network, is a group of servers spread out over a region or around the world that work together to speed up content delivery on the web. The servers in a CDN temporarily store (or cache) webpage content like images, HTML, JavaScript, and video. They send the cached content to users who load the webpage1.
A CDN can improve the performance of a web application that is accessed around the
world by:
Decreasing the distance between where content is stored and where it needs to go. A CDN can serve content from the server that is closest to the user, reducing network latency and bandwidth consumption.
Reducing file sizes to increase load speed. A CDN can employ techniques such as compression, minification, and image optimization to reduce the amount of data that needs to be transferred.
Optimizing server infrastructure to respond to user requests more quickly. A CDN can use hardware and software enhancements such as solid-state hard drives, load balancing, and caching algorithms to improve the efficiency and reliability of the servers12.
IPAM, or IP address management, is a method for planning, tracking, and managing the IP address space used in a network. IPAM does not directly affect the performance of a web application.
SDN, or software-defined networking, is a technology that allows network administrators to dynamically configure and control network resources using software applications. SDN can improve the flexibility and scalability of a network, but it does not necessarily improve the performance of a web application.
VPN, or virtual private network, is a technology that creates a secure and encrypted connection between a device and a network over the internet. VPN can enhance the privacy and security of a web application, but it does not improve its performance. In fact, VPN may introduce some overhead and latency due to encryption and decryption processes3.

NEW QUESTION 149
- (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 troubleshooting RSA key-based authentication from a local computer to a cloud-based server, which is running SSH service on a default port. The following file permissions are set on the authorized keys file:
-rw-rw-rw-1 ubuntu ubuntu 391 Mar S 01:36 authorized _ keys
Which Of the following security practices are the required actions the engineer Should take to gain access to the server? (Select TWO).

A. Fix the file permissions with execute permissions to the owner of the file.
B. Open port 21 access for the computer's public IP address.
C. Fix the file permissions with read-only access to the owner Of the file.
D. Open port 22 access for the computer's public IP address.
E. Open port 21 access for 0.0.0.0/0 CIDR.
F. open port 22 access for 0.0.0.0/0 CIDR.

Answer: CD

Explanation:
The correct answer is C and D.
* C. Fix the file permissions with read-only access to the owner of the file.
* D. Open port 22 access for the computer’s public IP address.
The authorized_keys file on the server should have read-only access for the owner of the file, and no access for anyone else. This ensures that only the owner can read the public keys that are authorized to log in, and no one can modify or delete them. The file permissions can be fixed with the command chmod 400 ~/.ssh/authorized_keys on the server. This is a recommended security practice for SSH key-based authentication123. The computer that wants to log in to the server using SSH key-based authentication needs to have access to port 22 on the server, which is the default port for SSH service. This can be done by opening port 22 access for the computer’s public IP address on the server’s firewall or security group settings. This allows the computer to initiate an SSH connection to the server and authenticate with its private key. Opening port 21, which is used for FTP service, is not relevant or secure for SSH key-based authentication1.

NEW QUESTION 152
- (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 the concept of ensuring that certain virtual machines or workloads do not run on the same physical host. This can improve the availability and performance of the system, as well as prevent a single point of failure. In this scenario, the systems administrator needs to ensure the database nodes do not exist on the same physical host, so anti-affinity would best meet this requirement. Oversubscription is the concept of allocating more resources to virtual machines than the physical host actually has, which can improve the utilization and efficiency of the system, but it does not guarantee the separation of the database nodes. A firewall is a device or software that controls the network traffic between different zones or segments, which can improve the security and isolation of the system, but it does not affect the placement of the database nodes. A separate cluster is a group of hosts that share common resources and policies, which can improve the scalability and manageability of the system, but it does not ensure the database nodes do not exist on the same physical host within the cluster. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 1, Cloud Architecture and Design, page 131.

NEW QUESTION 157

- (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 158
- (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 162
- (Topic 4)
A company has a large environment with multiple VPCs across three regions in a public cloud. The company is concerned about connectivity within the regions.
Which of the following should the cloud administrator implement?

A. Peering
B. A firewall
C. Network access control
D. A load balancer

Answer: A

Explanation:
Peering is a networking technique that allows direct and private connection between two or more cloud networks without using the public Internet. Peering can help the cloud administrator improve the connectivity within the regions by reducing the latency, increasing the bandwidth, and enhancing the security of the data transfer. Peering can be implemented between VPCs within the same region or across different regions, depending on the CSP's offerings and the customer's requirements. Peering can also help reduce the network costs by avoiding the use of the Internet gateways or VPNs. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 3, Objective 3.1: Given a scenario, implement cloud networking solutions.

NEW QUESTION 165
- (Topic 4)
A systems administrator notices several VMS are constantly ballooning, while the memory usage of several other VMS is significantly lower than their resource allocation. Which of the following will MOST likely solve the issue?

A. Rightsizing
B. Bandwidth increase
C. Cluster placement
D. Storage tiers

Answer: A

Explanation:
The best answer is A. Rightsizing.
Rightsizing is the process of restructuring a company so it can make a profit more efficiently and meet updated business objectives1. Organizations will usually rightsize their business by reducing their workforce, reorganizing upper management, cutting costs, and changing job roles2.
Rightsizing can help solve the issue of VMs constantly ballooning, while the memory usage of several other VMs is significantly lower than their resource allocation. Ballooning is a memory reclamation technique used when ESXi host runs out of memory. It involves a balloon driver that consumes unused memory within the VM's address space and makes it available for other uses by the host machine3. However, ballooning can also degrade the performance of the VMs and cause swapping or paging4.
By rightsizing the VMs, the systems administrator can adjust the memory allocation according to the actual demand and usage of each VM. This can prevent overprovisioning or underprovisioning of memory resources and improve the efficiency and profitability of the company. Rightsizing can also help avoid redundancies, streamline workflows, and make better hiring decisions1.

NEW QUESTION 167
- (Topic 4)
A cloud solutions architect has an environment that must only be accessed during work hours. Which of the following processes should be automated to best reduce cost?

A. Scaling of the environment after work hours
B. Implementing access control after work hours
C. Shutting down the environment after work hours
D. Blocking external access to the environment after work hours

Answer: C

Explanation:
Shutting down the environment after work hours is the best process to automate to reduce cost, as it will stop incurring charges for the cloud resources that are not needed outside of work hours. Scaling, implementing access control, or blocking external access may still incur some costs for the cloud resources that are running or reserved, even if they are not fully utilized. Shutting down the environment can be automated using scripts, schedules, or triggers that can turn off or deallocate the cloud resources based on time or usage criteria12.

NEW QUESTION 171
- (Topic 4)
A systems administrator is implementing a new file storage service that has been deployed in the company's private cloud instance. The key requirement is fast read/write times for the targeted users, and the budget for this project is not a concern. Which of the following storage types should the administrator deploy?

A. Spinning disks
B. NVMe
C. SSD
D. Hybrid

Answer: B

Explanation:
The best storage type to deploy for the new file storage service is NVMe. NVMe stands for Non-Volatile Memory Express, and it is a protocol that allows faster access to data stored on solid state drives (SSDs). NVMe can deliver high performance, low latency, and parallelism for the file storage service. NVMe can also support fast read/write times for the targeted users, which is the key requirement for the project. Since the budget for the project is not a concern, NVMe can be a suitable choice for the file storage service. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 4, Objective 4.1: Given a scenario, implement cloud storage solutions.

NEW QUESTION 174
- (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 176
- (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 178
- (Topic 3)
A systems administrator is diagnosing performance issues on a web application. The web application sends thousands of extremely complex SQL queries to a database server, which has trouble retrieving the information in time. The administrator checks the database server and notes the following resource utilization:
CPU: 64%
RAM: 97%
Network throughput: 384,100Kbps. Disk throughput: 382,700Kbps
The administrator also looks at the storage for the database server and notices it is consistently near its OPS limit. Which of the following will BEST resolve these performance issues?

A. Increase CPU resources on the database server.
B. Increase caching on the database server.
C. Put the storage and the database on the same VLAN.
D. Enable compression on storage traffic.
E. Enable deduplication on the storage appliance.

Answer: B

Explanation:
The performance issue is caused by the high demand of complex SQL queries on the database server, which consumes a lot of RAM and disk throughput. Increasing caching on the database server would reduce the number of disk reads and writes, as well as improve the response time of the queries by storing frequently accessed data in memory. This would be the best solution to resolve the performance issue. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 4.0 Troubleshooting, Objective 4.3 Given a scenario, troubleshoot capacity issues within a cloud environment.

NEW QUESTION 181
- (Topic 3)
A company is using a method of tests and upgrades in which a small set of end users are exposed to new services before the majority of other users. Which of the following deployment methods is being used?

A. Blue-green
B. Canary
C. Big bang
D. Rolling

Answer: B

Explanation:
A canary deployment is a software deployment technique where a new feature or version is released to a small subset of users in production prior to releasing it to a larger subset or all the users. It is also sometimes termed a phased rollout or incremental release1. A canary deployment allows the developers to test the new service in a real environment and get feedback from the users before making it available to everyone. It also reduces the risk of failures and enables easy rollbacks if something goes wrong. A canary deployment is different from a blue-green deployment, where two identical environments are used to switch between the old and new versions of the service. A big bang deployment is where the new service is released to all the users at once, without any testing or gradual rollout. A rolling deployment is where the new service is installed in batches or stages, replacing the old service gradually until all the users are on the new version.

NEW QUESTION 183
- (Topic 3)
A company has hired a security firm to perform a vulnerability assessment of its environment. In the first phase, an engineer needs to scan the network services exposed by the hosts. Which of the following will help achieve this with the LEAST privileges?

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

Answer: C

Explanation:
A network-based scan is a type of vulnerability assessment that scans the network services exposed by the hosts without requiring any credentials or agents. This type of scan will help achieve the objective of scanning the network services with the least privileges, as it does not need any access to the hosts or their internal

configurations. A network-based scan can identify open ports, running services, and potential vulnerabilities on the hosts. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 2.0 Security, Objective 2.4 Given a scenario, implement security automation and orchestration in a cloud environment.

NEW QUESTION 186
- (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 190
- (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 192
- (Topic 3)
A cloud administrator implemented SSO and received a business requirement to increase security when users access the cloud environment. Which of the following should be implemented NEXT to improve the company's security posture?

A. SSH
B. MFA
C. Certificates
D. Federation

Answer: B

Explanation:
MFA (Multi-Factor Authentication) is a security technique that requires the user to present two or more pieces of evidence to prove their identity when they try to access a system or an application. For example, a password and a physical token, or a fingerprint and a one-time code. MFA can improve the company's security posture by preventing unauthorized access even if the password or single-factor authentication is compromised, as the attacker would also need to have the other factors to log
in. According to the web search results, MFA can prevent 99.9% of account attacks1.
SSO (Single Sign-On) is a system that allows the user to use one set of login credentials to access multiple systems and applications that previously may have each required their own logins. SSO can improve productivity and user convenience, but it does not replace MFA. In fact, SSO works in conjunction with MFA, as it can enforce MFA for all the systems and applications that are integrated with SSO2. Therefore, implementing SSO does not mean that MFA is not needed.

NEW QUESTION 197
- (Topic 3)
An administrator manages a file server that has a lot of users accessing and creating many files. As a result, the storage consumption is growing quickly. Which of the following would BEST control storage usage?

A. Compression
B. File permissions
C. User quotas
D. Access policies

Answer: C

Explanation:
User quotas are a feature that allows the administrator to limit the amount of storage space that a user or a group of users can consume on a file server. User

quotas can help to control storage usage by preventing users from storing excessive or unnecessary files, as well as by enforcing fair and consistent storage policies across the organization. User quotas can also help to monitor and report on the storage consumption and trends of the users, and alert the administrator or the users when they are approaching or exceeding their quota limits.

NEW QUESTION 201
- (Topic 3)
A systems administrator is planning a penetration test for company resources that are hosted in a public cloud. Which of the following must the systems administrator do FIRST?

A. Consult the law for the country where the company's headquarters is located
B. Consult the regulatory requirements for the company's industry
C. Consult the law for the country where the cloud services provider is located
D. Consult the cloud services provider's policies and guidelines

Answer: D

Explanation:
The first thing that the systems administrator must do before planning a penetration test for company resources that are hosted in a public cloud is to consult the cloud services provider's policies and guidelines. Penetration testing is a type of security assessment that involves simulating an attack on a system or network to identify vulnerabilities and weaknesses. However, not all cloud services providers allow penetration testing on their platforms, or they may have specific rules and requirements for conducting such tests. The systems administrator should check the cloud services provider's policies and guidelines and obtain their permission and approval before performing any penetration testing. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 2.0 Security, Objective 2.4 Given a scenario, implement security automation and orchestration in a cloud environment.

NEW QUESTION 203
- (Topic 3)
A web-application company recently released some new marketing promotions without notifying the IT staff. The systems administrator has since been noticing twice the normal traffic consumption every two hours for the last three hours in the container environment. Which of the following should the company implement to accommodate the new traffic?

A. A firewall
B. Switches
C. Ballooning
D. Autoscaling

Answer: D

Explanation:
According to the CompTIA Cloud+ Study Guide1, autoscaling is "the ability to automatically increase or decrease the number of resources allocated to a cloud service based on the current demand". This means that autoscaling can help a cloud environment adjust to changes in traffic and workload, such as the ones caused by the new marketing promotions.
Ballooning, on the other hand, is "a technique used by hypervisors to reclaim unused memory from a virtual machine and allocate it to another virtual machine that needs more memory". This means that ballooning can help optimize the memory usage of a cloud environment, but it does not affect the number of resources allocated to a cloud service. A firewall is "a device or software that monitors and controls the incoming and outgoing network traffic based on predefined rules". This means that a firewall can help protect a cloud environment from unauthorized access and malicious attacks, but it does not affect the number of resources allocated to a cloud service.
Switches are "devices that connect multiple devices on a network and forward data packets between them". This means that switches can help improve the network performance and connectivity of a cloud environment, but they do not affect the number of resources allocated to a cloud service.
Based on this information, I think the best answer to your question is D. Autoscaling. Autoscaling can help the company accommodate the new traffic by automatically increasing or decreasing the number of resources allocated to their web-application service based on the current demand. This can also help reduce costs and improve performance and availability.
I hope this helps you understand the concept of autoscaling better. If you want to learn more about CompTIA Cloud+, you can check out some of these resources:
? CompTIA Cloud+ : Cloud High Availability & Scaling: A video course that covers the topics of high availability and scaling in cloud environments, including autoscaling, horizontal scaling, vertical scaling and cloud bursting.
? Cloud+ (Plus) Certification | CompTIA IT Certifications: The official website of CompTIA Cloud+, where you can find exam details, preparation materials, renewal information and more.

NEW QUESTION 208
- (Topic 3)
A cloud administrator needs to coordinate and automate the management of a company's secrets and keys for all its cloud services with minimal effort and low cost. Which of the following is the BEST option to achieve the goal?

A. Implement database as a service
B. Configure Key Vault
C. Use password as a service
D. Implement KeePass

Answer: B

Explanation:
Reference: https://docs.microsoft.com/en-us/azure/key-vault/key-vault-whatis
Key Vault is a service that allows you to store and manage secrets and keys for your cloud services in a secure and centralized way. It also provides access control, auditing, and encryption features. This would be the best option to automate the management of secrets and keys for all cloud services with minimal effort and low cost. 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 210
- (Topic 3)
An organization is hosting its dedicated email infrastructure with unlimited mailbox creation capability. The management team would like to migrate to a SaaS-

based solution. Which of the following must be considered before the migration?

A. The SaaS provider's licensing model
B. The SaaS provider's reputation
C. The number of servers the SaaS provider has
D. The number of network links the SaaS provider has

Answer: A

Explanation:
The licensing model of the SaaS provider is an important factor to consider before migrating to a SaaS-based solution for email infrastructure. The licensing model determines how much the organization will pay for the service, how many mailboxes they can create, what features they can access, and what SLAs they can expect. The organization should compare different SaaS providers' licensing models and choose the one that best suits their needs and budget. 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 212
- (Topic 3)
A cloud administrator is monitoring a database system and notices an unusual increase in the read operations, which is causing a heavy load in the system. The system is using a relational database and is running in a VM. Which of the following should the administrator do to resolve the issue with minimal architectural changes?

A. Migrate the relational database to a NoSQL database.
B. Use a cache system to store reading operations.
C. Create a secondary standby database instance.
D. Implement the database system using a DBaaS.

Answer: B

Explanation:
The best way to resolve the issue of an unusual increase in the read operations that is causing a heavy load in the system that is using a relational database and is running in a VM is to use a cache system to store reading operations. A cache system is a type of storage system that temporarily stores frequently accessed or recently used data in memory for faster retrieval. A cache system can reduce the load on the database system by serving the read requests from the cache instead of querying the database every time. Reference: [CompTIA Cloud+ Certification Exam Objectives], Domain 4.0 Troubleshooting, Objective 4.3 Given a scenario, troubleshoot capacity issues within a cloud environment.

NEW QUESTION 216
- (Topic 3)
A systems administrator is helping to develop a disaster recovery solution. The solution
must ensure all production capabilities are available within two hours. Which of the following will BEST meet this requirement?

A. A hot site
B. A warm site
C. A backup site
D. A cold site

Answer: A

Explanation:
Reference: https://searchdisasterrecovery.techtarget.com/definition/hot-site
A hot site is what would best meet the requirement of ensuring all production capabilities are available within two hours for a disaster recovery solution. A disaster recovery solution is a plan or process of restoring normal operation and performance of a system or service after a disruption or disaster. A disaster recovery solution can use different types of sites or locations to store and recover data or resources, such as:
A hot site: This is a site or location that has a fully operational and ready-to-use replica or copy of the original system or service, including data, resources, applications, etc. A hot site can provide benefits such as:
Availability: A hot site can provide availability by ensuring that the system or service can be switched or transferred to the hot site immediately or within minutes after a disruption or disaster, without any downtime or interruption.
Capability: A hot site can provide capability by ensuring that the system or service can function and perform at the same level or quality as the original system or service, without any loss or degradation.
A warm site: This is a site or location that has a partially operational and ready-to-use replica or copy of the original system or service, including some data, resources, applications, etc. A warm site can provide benefits such as:
Affordability: A warm site can provide affordability by reducing the cost of maintaining and updating the replica or copy of the original system or service, compared to a hot site.
Flexibility: A warm site can provide flexibility by allowing customers to customize and configure the replica or copy of the original system or service according to their needs and preferences, compared to a hot site.
A cold site: This is a site or location that has no operational and ready-to-use replica or copy of the original system or service, but only has the necessary infrastructure or facilities to support it, such as power, network, space, etc. A cold site can provide benefits such as:
Scalability: A cold site can provide scalability by enabling customers to expand and grow their replica or copy of the original system or service as needed, without any limitations or constraints.
Security: A cold site can provide security by minimizing the exposure or risk of the replica or copy of the original system or service to any threats or attacks, compared to a hot site or a warm site

NEW QUESTION 217
- (Topic 3)
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? (Select TWO).

A. Disk I/O limits
B. Affinity rule
C. CPU oversubscription
D. RAM usage

E. Insufficient GPI-J resources
F. License issues

Answer: AC

Explanation:
The most likely causes of the issue are A. Disk I/O limits and C. CPU oversubscription. Disk I/O limits are the maximum amount of input/output operations per second (IOPS) that a disk can handle. CPU oversubscription is the ratio of virtual CPUs to physical CPUs in a host. Both of these factors can affect the performance of a VDI environment, especially during peak hours when many users log in and launch applications.
Disk I/O limits can cause slow boot times, application lags, and cursor freezes for VDI users12. To avoid this issue, it is recommended to use flash storage or SSDs for VDI workloads, as they have much higher IOPS than traditional hard disk drives31. It is also important to monitor the disk performance and adjust the disk size and configuration as needed1.
CPU oversubscription can also cause performance degradation for VDI users, as it can lead to CPU contention and increased latency42. To avoid this issue, it is recommended to limit the CPU oversubscription ratio to a reasonable level, such as 4:1 or lower42. It is also important to monitor the CPU utilization and balance the load across hosts as needed4.
The other options are less likely to cause the issue. Affinity rules are used to specify which virtual machines should run on which hosts or which virtual machines should not run on the same host. They are not related to the performance of VDI workloads. RAM usage can affect the performance of VDI workloads, but it is usually not a major factor during peak hours, as most users do not consume a lot of memory when they log in or launch applications. Insufficient GPU resources can affect the performance of VDI workloads that require high graphics processing, such as video streaming or 3D rendering, but they are not relevant for most VDI users. License issues can affect the availability of VDI workloads, but they are not related to the performance of VDI workloads.

NEW QUESTION 220
- (Topic 3)
A cloud security engineer needs to ensure authentication to the cloud pro-vider console is secure. Which of the following would BEST achieve this ob-jective?

A. Require the user's source IP to be an RFC1918 address.
B. Require the password to contain uppercase letters, lowercase letters, numbers, and symbols.
C. Require the use of a password and a physical token.
D. Require the password to be ten characters long.

Answer: C

Explanation:
A password and a physical token are two factors of authentication that can provide a higher level of security than a password alone. A physical token is a device that generates a one-time code or password that the user must enter along with their password to access the cloud provider console. This is an example of multi-factor authentication (MFA), which requires the user to present two or more pieces of evidence to prove their identity. MFA can prevent unauthorized access even if the password is compromised, as the attacker would also need to have the physical token to log in.

NEW QUESTION 225
- (Topic 3)
A production engineer is configuring a new application, which is running in containers, that requires access to a database. Which of the following methods will allow the application to authenticate to the database in the MOST secure way?

A. Store the credentials in a variable on every worker node
B. Store the credentials on a shared volume using whole-disk encryption
C. Store the credentials in a configuration file using SHA-256 inside the container image
D. Store the credentials using the orchestrator secret manager

Answer: D

Explanation:
The most secure way to store the credentials for a new application that is running in containers and requires access to a database is to use the orchestrator secret manager. The orchestrator secret manager is a feature that allows storing and managing sensitive data, such as passwords, tokens, or keys, for containers in an encrypted and centralized way. It also provides access control, auditing, and rotation features for the secrets. This method will protect the credentials from being exposed or compromised by unauthorized parties or malicious actors. 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 227
- (Topic 3)
After initial stress testing showed that a platform performed well with the specification of a single 32 vCPU node, which of the following will provide the desired service with the LOWEST cost and downtime?

A. One 32 vCPU node with CDN caching
B. Two 8 vCPU nodes with load balancing
C. Three to six 8 vCPU nodes autoscaling group
D. Four 8 vCPU nodes with DNS round robin

Answer: C

Explanation:
The best option to provide the desired service with the lowest cost and downtime after initial stress testing showed that a platform performed well with the specification of a single 32 vCPU node is to use three to six 8 vCPU nodes autoscaling group. An autoscaling group is a feature that allows dynamically adjusting the number of instances or nodes in a cluster based on the demand or load. This option will provide high availability, scalability, and performance for the service, while also optimizing the cost and resource utilization by adding or removing nodes as needed. 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 228
- (Topic 3)
A large pharmaceutical company needs to ensure it is in compliance with the following requirements:

• An application must run on its own virtual machine.
• The hardware the application is hosted on does not change. Which of the following will BEST ensure compliance?

A. Containers
B. A firewall
C. Affinity rules
D. Load balancers

Answer: C

Explanation:
According to the Virtual Machine Affinity and Anti-Affinity - VMware Docs1, affinity and anti- affinity rules allow you to spread a group of virtual machines across different ESXi hosts or keep a group of virtual machines on a particular ESXi host. An affinity rule places a group of virtual machines on a specific host so that you can easily audit the usage of those virtual machines.
Containers are “a standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another” 2. Containers can run on any virtual machine or physical server, and they can be moved between different hosts without affecting the application functionality.
A firewall is “a device or software that monitors and controls the incoming and outgoing network traffic based on predefined rules” 3. A firewall can help protect a cloud environment from unauthorized access and malicious attacks, but it does not affect the placement of virtual machines on hosts.
Load balancers are “devices or software that distribute network or application traffic across
a number of servers” . Load balancers can help improve the performance and availability of a cloud environment by distributing the workload among multiple servers, but they do not affect the placement of virtual machines on hosts.
Based on this information, I think the best answer to your question is C. Affinity rules. Affinity rules can help the pharmaceutical company ensure compliance by placing the application on its own virtual machine and keeping it on the same host. This way, the company can easily audit the usage of the application and avoid any changes in the hardware configuration.

NEW QUESTION 230
.....

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