

CompTIA

Exam Questions CV0-003

CompTIA Cloud+ Certification Exam



NEW QUESTION 1

- (Topic 1)

An organization has two businesses that are developing different software products. They are using a single cloud provider with multiple IaaS instances. The organization identifies that the tracking of costs for each business are inaccurate.

Which of the following is the BEST method for resolving this issue?

- A. Perform segregation of the VLAN and capture egress and ingress values of each network interface
- B. Tag each server with a dedicated cost and sum them based on the businesses
- C. Split the total monthly invoice equally between the businesses
- D. Create a dedicated subscription for the businesses to manage the costs

Answer: B

Explanation:

Tagging each server with a dedicated cost and summing them based on the businesses is the best method for resolving the issue of inaccurate cost tracking for different businesses that use multiple IaaS instances within a single cloud provider. Tagging can help identify and organize the servers based on various criteria, such as name, purpose, owner, or cost center. Tagging can also enable granular and accurate billing and reporting based on the tags. Summing the costs based on the businesses can help allocate and distribute the costs correctly and fairly among the different businesses. References: CompTIA Cloud+ Certification Exam Objectives, page 13, section 2.5

NEW QUESTION 2

- (Topic 1)

Which of the following cloud deployment models allows a company to have full control over its IT infrastructure?

- A. Private
- B. Cloud within a cloud
- C. Hybrid
- D. Public

Answer: A

Explanation:

A private cloud is a type of cloud deployment model that provides cloud services exclusively to a single organization or tenant. A private cloud allows a company to have full control over its IT infrastructure, as it can customize, configure, manage, and secure its own cloud environment according to its specific needs and preferences. A private cloud can also offer higher performance, reliability, and privacy than other cloud deployment models, as it does not share resources or data with other customers.

References: CompTIA Cloud+ Certification Exam Objectives, page 8, section 1.2 Reference: <https://www.sciencedirect.com/topics/computer-science/private-cloud>

NEW QUESTION 3

- (Topic 1)

A cloud administrator is reviewing the authentication and authorization mechanism implemented within the cloud environment. Upon review, the administrator discovers the sales group is part of the finance group, and the sales team members can access the financial application. Single sign-on is also implemented, which makes access much easier.

Which of the following access control rules should be changed?

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

Answer: D

Explanation:

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

Reference: <https://www.ekransystem.com/en/blog/rbac-vs-abac>

NEW QUESTION 4

SIMULATION - (Topic 1)

The QA team is testing a newly implemented clinical trial management (CTM) SaaS application that uses a business intelligence application for reporting. The UAT users were instructed to use HTTP and HTTPS.

Refer to the application dataflow:

- 1A – The end user accesses the application through a web browser to enter and view clinical data.
- 2A – The CTM application server reads/writes data to/from the database server.
- 1B – The end user accesses the application through a web browser to run reports on clinical data.
- 2B – The CTM application server makes a SOAP call on a non-privileged port to the BI application server.
- 3B – The BI application server gets the data from the database server and presents it to the CTM application server.

When UAT users try to access the application using <https://ctm.app.com> or <http://ctm.app.com>, they get a message stating: "Browser cannot display the webpage." The QA team has raised a ticket to troubleshoot the issue.

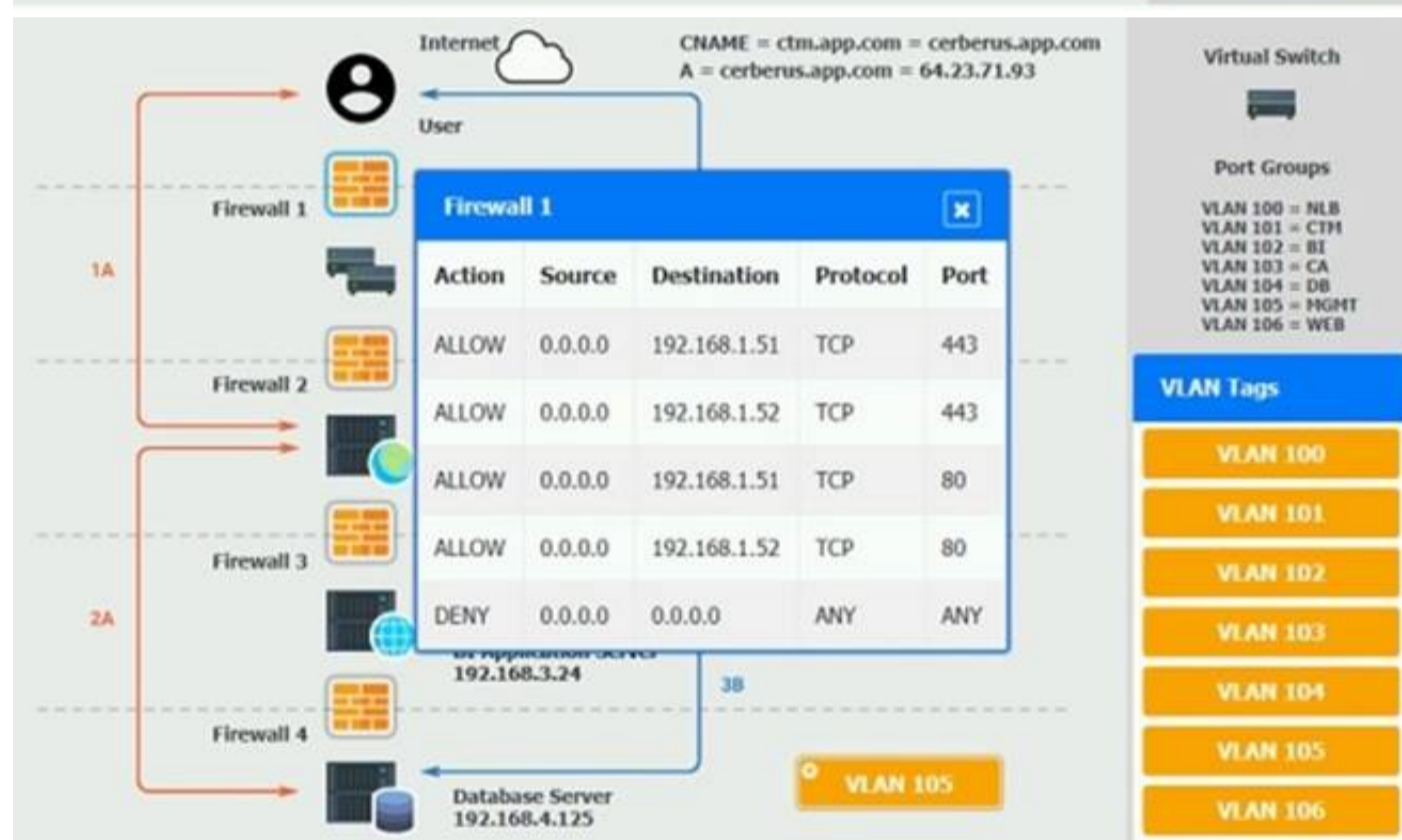
INSTRUCTIONS

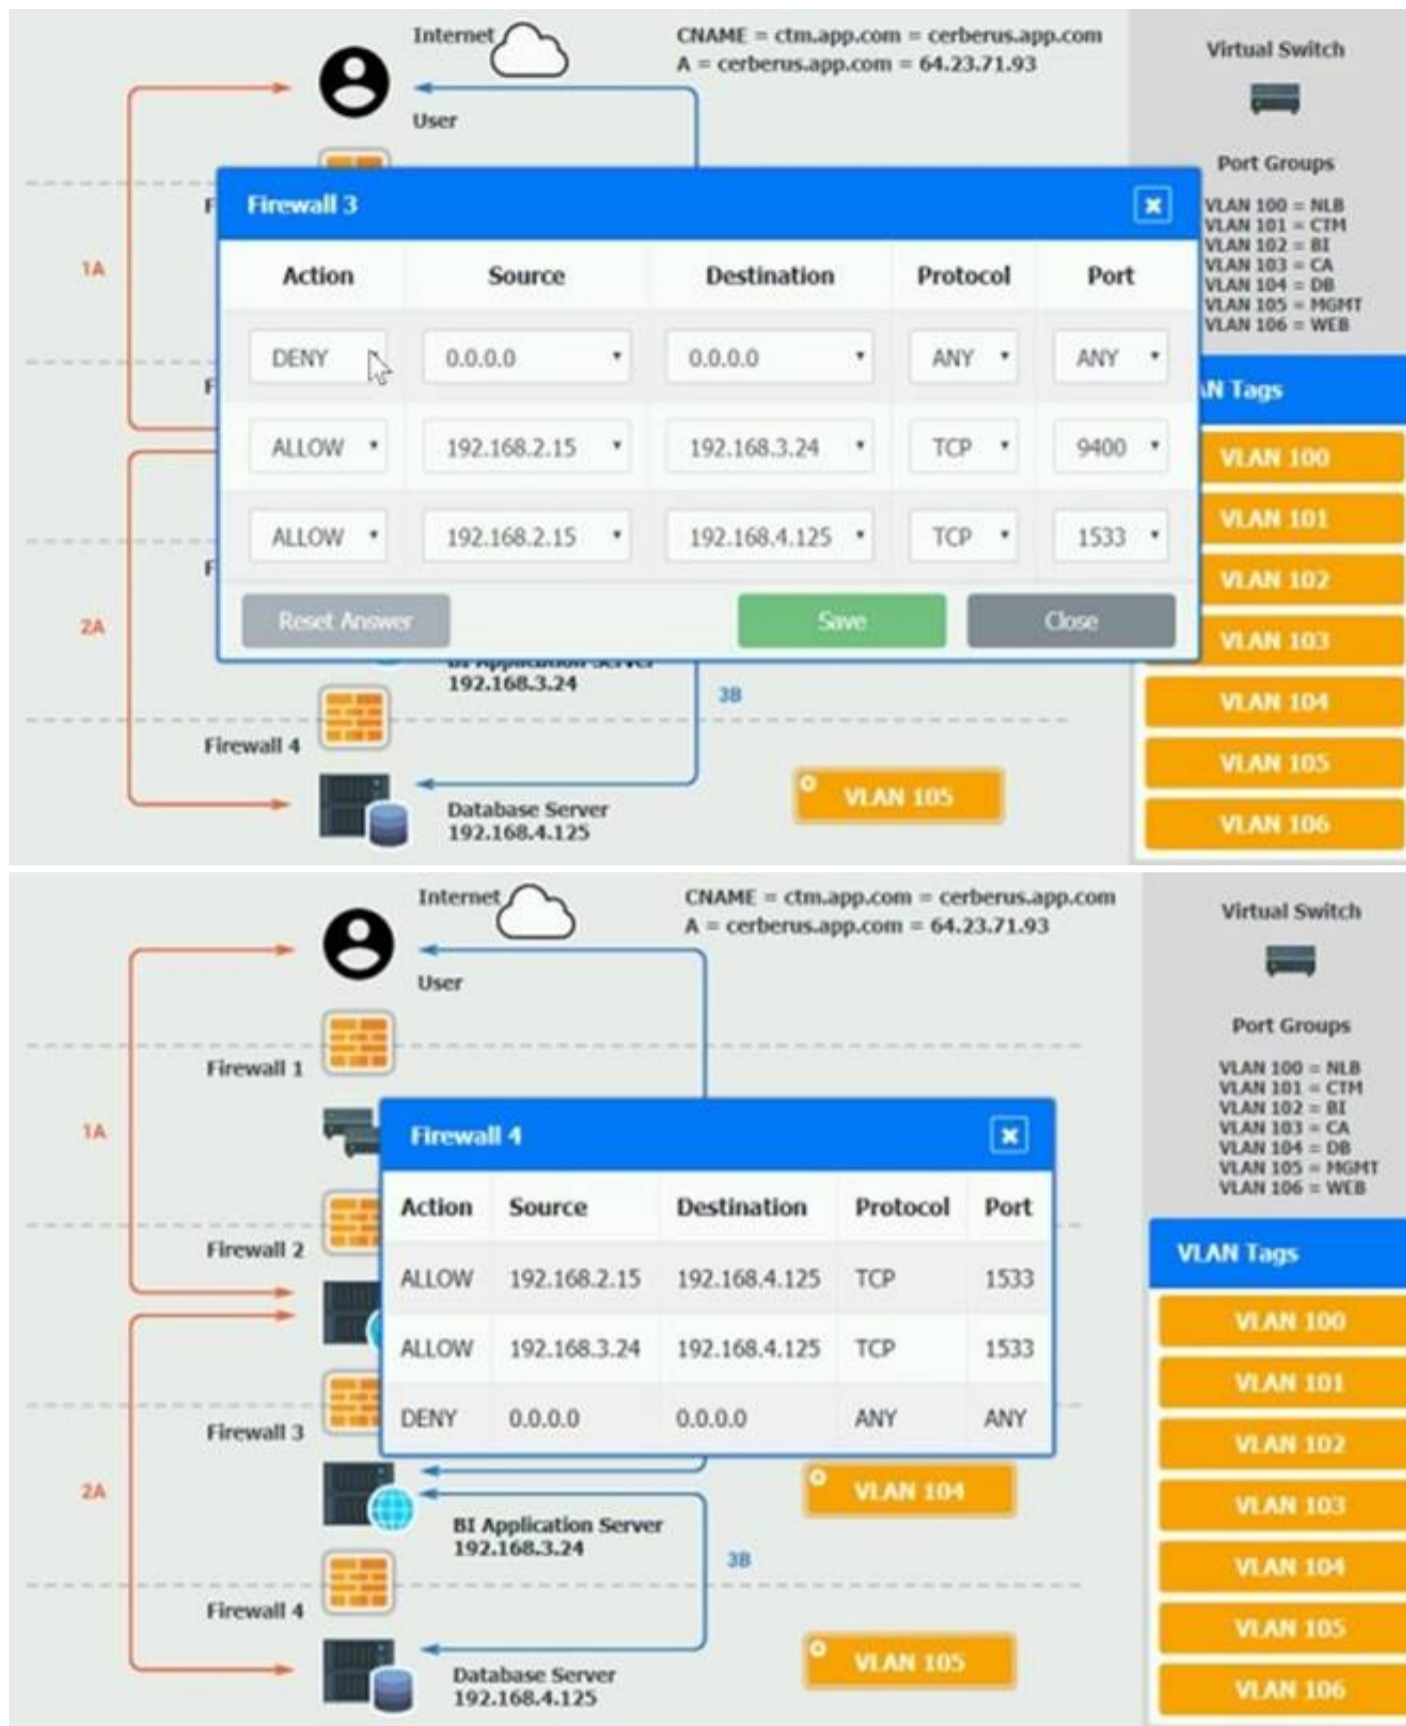
You are a cloud engineer who is tasked with reviewing the firewall rules as well as virtual network settings.

You should ensure the firewall rules are allowing only the traffic based on the dataflow. You have already verified the external DNS resolution and NAT are working.

Verify and appropriately configure the VLAN assignments and ACLs. Drag and drop the appropriate VLANs to each tier from the VLAN Tags table. Click on each Firewall to change ACLs as needed.

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





- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 5

- (Topic 1)

Which of the following strategies will mitigate the risk of a zero-day vulnerability MOST efficiently?

- A. Using only open-source technologies
- B. Keeping all resources up to date
- C. Creating a standby environment with a different cloud provider
- D. Having a detailed incident response plan

Answer: D

Explanation:

An incident response plan is a document or procedure that defines the roles, responsibilities, and actions to be taken in the event of a security incident or breach. Having a detailed incident response plan can help mitigate the risk of a zero-day vulnerability most efficiently, as it can provide a clear and consistent framework for identifying, containing, analyzing, and resolving any potential threats or exploits related to the unknown or unpatched vulnerability. Having a detailed incident response plan can also help minimize the impact and damage of a security incident or breach, as it can enable timely and effective recovery and restoration processes. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 6

- (Topic 1)

A DevOps administrator is automating an existing software development workflow. The administrator wants to ensure that prior to any new code going into

production, tests confirm the new code does not negatively impact existing automation activities.
Which of the following testing techniques would be BEST to use?

- A. Usability testing
- B. Regression testing
- C. Vulnerability testing
- D. Penetration testing

Answer: B

Explanation:

Regression testing is a type of testing that ensures that new code or changes to existing code do not break or degrade the functionality of the software. Regression testing is often used in software development workflows to verify that new features or bug fixes do not introduce new errors or affect the performance of the software. Regression testing can help prevent negative impacts on existing automation activities by checking that the new code is compatible with the existing code and does not cause any unexpected failures or errors. References: CompTIA Cloud+ Certification Exam Objectives, page 19, section 4.1
Reference: <https://www.softwaretestinghelp.com/regression-testing-tools-and-methods/>

NEW QUESTION 7

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 8

- (Topic 1)

An administrator is performing an in-place upgrade on a guest VM operating system.
Which of the following can be performed as a quick method to roll back to an earlier state, if necessary?

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

Answer: D

Explanation:

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

NEW QUESTION 9

- (Topic 1)

A SAN that holds VM files is running out of storage space.
Which of the following will BEST increase the amount of effective storage on the SAN?

- A. Enable encryption
- B. Increase IOPS
- C. Convert the SAN from RAID 50 to RAID 60
- D. Configure deduplication

Answer: D

Explanation:

Deduplication is a type of data compression technique that eliminates redundant or duplicate data blocks or segments in a storage system or device. Configuring deduplication can help increase the amount of effective storage on a SAN that holds VM files and is running out of storage space, as it can reduce the storage space consumption and increase the storage space utilization by storing only unique data blocks or segments. Configuring deduplication can also improve performance and efficiency, as it can speed up data transfer and backup processes and save network bandwidth and power consumption. References: CompTIA Cloud+ Certification Exam Objectives, page 9, section 1.4

NEW QUESTION 10

- (Topic 1)

A systems administrator is deploying a GPU-accelerated VDI solution. Upon requests from several users, the administrator installs an older version of the OS on their virtual workstations. The majority of the VMs run the latest LTS version of the OS. Which of the following types of drivers will MOST likely ensure compatibility with all virtual workstations?

- A. Alternative community drivers
- B. Legacy drivers
- C. The latest drivers from the vendor's website
- D. The drivers from the OS repository

Answer: D

Explanation:

The drivers from the OS repository are the drivers that are included or available in the official software repository or package manager of the operating system. The drivers from the OS repository are most likely to ensure compatibility with all virtual workstations that use a GPU-accelerated VDI solution, as they are tested and verified to work with different versions of the operating system and the hardware. The drivers from the OS repository can also provide stability and security, as they are regularly updated and patched by the operating system vendor or community. References: CompTIA Cloud+ Certification Exam Objectives, page 11, section 1.6

NEW QUESTION 10

- (Topic 1)

An organization is hosting a DNS domain with private and public IP ranges. Which of the following should be implemented to achieve ease of management?

- A. Network peering
- B. A CDN solution
- C. A SDN solution
- D. An IPAM solution

Answer: D

Explanation:

An IP address management (IPAM) solution is a type of tool or system that automates and standardizes the allocation, tracking, and management of IP addresses in an IP network. An IPAM solution can help achieve ease of management for hosting a DNS domain with private and public IP ranges, as it can simplify and centralize the process of assigning and updating IP addresses for different DNS records or zones without manual intervention or errors. An IPAM solution can also help optimize DNS performance and security, as it can monitor and report any issues or conflicts related to IP addresses or DNS records. References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8
Reference: <https://www.infoblox.com/glossary/ipam-ip-address-management/>

NEW QUESTION 12

- (Topic 1)

A company has decided to get multiple compliance and security certifications for its public cloud environment. However, the company has few staff members to handle the extra workload, and it has limited knowledge of the current infrastructure.

Which of the following will help the company meet the compliance requirements as quickly as possible?

- A. DLP
- B. CASB
- C. FIM
- D. NAC

Answer: B

Explanation:

A cloud access security broker (CASB) is a type of security solution that acts as a gateway between cloud service users and cloud service providers. A CASB can help a company get multiple compliance and security certifications for its public cloud environment, as it can provide visibility, control, and protection for cloud data and applications. A CASB can also help the company handle the extra workload and overcome the limited knowledge of the current infrastructure, as it can automate and simplify the enforcement of security policies and compliance requirements across multiple cloud services. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 17

- (Topic 1)

A cloud administrator has built a new private cloud environment and needs to monitor all computer, storage, and network components of the environment.

Which of the following protocols would be MOST useful for this task?

- A. SMTP
- B. SCP
- C. SNMP
- D. SFTP

Answer: C

Explanation:

Simple Network Management Protocol (SNMP) is a protocol that enables monitoring and managing network devices and components in an IP network. SNMP can help monitor all computer, storage, and network components of a private cloud environment, as it can collect and report information about their status, performance, configuration, and events. SNMP can also help troubleshoot and optimize the private cloud environment, as it can detect and alert any issues or anomalies related to the network devices and components. References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8

NEW QUESTION 18

- (Topic 1)

A technician is working with an American company that is using cloud services to provide video-based training for its customers. Recently, due to a surge in demand, customers in Europe are experiencing latency.

Which of the following services should the technician deploy to eliminate the latency issue?

- A. Auto-scaling
- B. Cloud bursting
- C. A content delivery network
- D. A new cloud provider

Answer: C

Explanation:

<https://www.cloudflare.com/learning/cdn/what-is-a-cdn/>

"A content delivery network (CDN) refers to a geographically distributed group of servers which work together to provide fast delivery of Internet content."

NEW QUESTION 21

- (Topic 1)

A company wants to check its infrastructure and application for security issues regularly. Which of the following should the company implement?

- A. Performance testing
- B. Penetration testing
- C. Vulnerability testing
- D. Regression testing

Answer: C

Explanation:

Vulnerability testing is a type of testing that identifies and evaluates the weaknesses or flaws in a system or application that could be exploited by attackers.

Vulnerability testing can help check the infrastructure and application for security issues regularly, as it can reveal the potential risks and exposures that may compromise the confidentiality, integrity, or availability of the system or application. Vulnerability testing can also help remediate or mitigate the vulnerabilities by providing recommendations or solutions to fix or reduce them. References: CompTIA Cloud+ Certification Exam Objectives, page 19, section 4.1

Reference: <https://pure.security/services/technical-assurance/external-penetration-testing/>

NEW QUESTION 25

- (Topic 1)

A cloud administrator recently deployed an update to the network drivers of several servers. Following the update, one of the servers no longer responds to remote login requests. The cloud administrator investigates the issue and gathers the following information:

? The cloud management console shows the VM is running and the CPU and memory utilization is at or near 0%.

? The cloud management console does not show an IP address for that server.

? A DNS lookup shows the hostname resolves to an IP address.

? The server is a member of the same security group as the others.

? The cloud administrator is able to log in remotely to the other servers without issue.

Which of the following is the MOST likely cause of the server being unavailable?

- A. The network driver updates did not apply successfully, and the interface is in a down state.
- B. The ACL policy for the server was updated as part of the server reboot, preventing login access.
- C. The server was assigned a new IP address, and DNS entry for the server name was not updated.
- D. The update caused an increase in the output to the logs, and the server is too busy to respond.

Answer: A

NEW QUESTION 30

- (Topic 1)

A company has developed a cloud-ready application. Before deployment, an administrator needs to select a deployment technology that provides a high level of portability and is lightweight in terms of footprint and resource requirements.

Which of the following solutions will be BEST to help the administrator achieve the requirements?

- A. Containers
- B. Infrastructure as code
- C. Desktop virtualization
- D. Virtual machines

Answer: A

Explanation:

Containers are a type of deployment technology that packages an application and its dependencies into a lightweight and portable unit that can run on any platform or environment. Containers can provide a high level of portability and are lightweight in terms of footprint and resource requirements, as they do not need a full operating system or hypervisor to run. Containers can also enable faster and easier deployment, scaling, and management of cloud-based applications.

Containers are the best solution to help the administrator achieve the requirements for deploying a cloud-ready application. References: CompTIA Cloud+ Certification Exam Objectives, page 11, section 1.6

Reference: <https://blog.netapp.com/blogs/containers-vs-vms/>

NEW QUESTION 34

- (Topic 1)

A systems administrator wants the VMs on the hypervisor to share CPU resources on the same core when feasible.

Which of the following will BEST achieve this goal?

- A. Configure CPU passthrough
- B. Oversubscribe CPU resources
- C. Switch from a Type 1 to a Type 2 hypervisor
- D. Increase instructions per cycle

E. Enable simultaneous multithreading

Answer: E

Explanation:

Simultaneous multithreading (SMT) is a type of CPU technology that allows multiple threads to run concurrently on a single CPU core. Enabling SMT can help achieve the goal of having the VMs on the hypervisor share CPU resources on the same core when feasible, as it can increase the CPU utilization and efficiency by executing more instructions per cycle and reducing idle time or wasted cycles. Enabling SMT can also improve performance and throughput, as it can speed up processing and handle increased workload or demand. References: CompTIA Cloud+ Certification Exam Objectives, page 9, section 1.4

NEW QUESTION 36

- (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 41

- (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 45

- (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 48

- (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 52

- (Topic 1)

A cloud administrator is reviewing a new application implementation document. The administrator needs to make sure all the known bugs and fixes are applied, and unwanted ports and services are disabled.

Which of the following techniques would BEST help the administrator assess these business requirements?

- A. Performance testing
- B. Usability testing
- C. Vulnerability testing
- D. Regression testing

Answer: D

Explanation:

Regression testing is a type of software testing that verifies that existing features or functionalities of a system or application are not affected by any changes or updates made to it. Regression testing can help assess whether all the known bugs and fixes are applied and unwanted ports and services are disabled when reviewing a new application implementation document for a cloud deployment, as it can detect any errors or defects that may have been introduced or re-introduced after applying patches, updates, or configurations to the application. References: CompTIA Cloud+ Certification Exam Objectives, page 19, section 4.1

NEW QUESTION 55

- (Topic 1)

Which of the following is relevant to capacity planning in a SaaS environment?

- A. Licensing
- B. A hypervisor
- C. Clustering
- D. Scalability

Answer: D

Explanation:

Scalability is the ability of a system or service to handle increased workload or demand by adding or removing resources or capacity as needed. Scalability is relevant to capacity planning in a SaaS environment, as it can affect the performance, availability, and cost of the SaaS service. Scalability can help optimize the capacity planning process by ensuring that the SaaS service has enough resources or capacity to meet the current and future needs of the customers without wasting or underutilizing resources or capacity. References: CompTIA Cloud+ Certification Exam Objectives, page 12, section 2.2

NEW QUESTION 59

- (Topic 1)

A media company has made the decision to migrate a physical, internal file server to the cloud and use a web-based interface to access and manage the files. The users must be able to use their current corporate logins.

Which of the following is the MOST efficient way to achieve this goal?

- A. Deploy a VM in a cloud, attach storage, and copy the files across
- B. Use a SaaS service with a directory service federation
- C. Deploy a fileshare in a public cloud and copy the files across
- D. Copy the files to the object storage location in a public cloud

Answer: B

Explanation:

Software as a service (SaaS) is a type of cloud service model that provides software applications over the Internet that are hosted and managed by a cloud service provider. Directory service federation is a type of authentication mechanism that allows users to access multiple systems or applications across different domains or organizations with a single login credential. Using a SaaS service with a directory service federation can help migrate an internal file server to the cloud and use a web-based interface to access and manage the files, as it can eliminate the need for maintaining an on-premises file server and enable seamless and secure access to cloud-based files using the same corporate logins. References: CompTIA Cloud+ Certification Exam Objectives, page 8, section 1.2

NEW QUESTION 61

- (Topic 1)

A cloud architect wants to minimize the risk of having systems administrators in an IaaS compute instance perform application code changes. The development group should be the only group allowed to modify files in the directory.

Which of the following will accomplish the desired objective?

- A. Remove the file write permissions for the application service account.
- B. Restrict the file write permissions to the development group only.

- C. Add access to the fileshare for the systems administrator's group.
- D. Deny access to all development user accounts

Answer: B

Explanation:

File write permissions are permissions that control who can modify or delete files in a directory or system. Restricting the file write permissions to the development group only can help minimize the risk of having systems administrators in an IaaS compute instance perform application code changes, as it can prevent anyone other than the development group from altering or removing any files in the directory where the application code is stored. Restricting the file write permissions can also help maintain consistency and integrity, as it can ensure that only authorized and qualified users can make changes to the application code. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 64

- (Topic 2)

A cloud administrator wants to have a central repository for all the logs in the company's private cloud. Which of the following should be implemented to BEST meet this requirement?

- A. SNMP
- B. Log scrubbing
- C. CMDB
- D. A syslog server

Answer: D

Explanation:

Reference: <https://www.itpro.com/infrastructure/network-internet/355174/how-to-build-a-dedicated-syslog-server>

A syslog server is what the administrator should implement to have a central repository for all the logs in the company's private cloud. Syslog is a standard protocol that allows network devices and systems to send log messages to a centralized server or collector. Syslog can help to consolidate and manage logs from different sources in one place, which can facilitate monitoring, analysis, troubleshooting, auditing, etc.

NEW QUESTION 69

- (Topic 2)

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

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

Answer: ADE

Explanation:

These are the best practices to secure the OS of a new web server that has been provisioned in a cloud environment:

? Install TLS certificates on the server: TLS (Transport Layer Security) certificates are digital documents that contain information such as identity, public key, expiration date, etc., that can be used to prove one's identity and establish secure communication over a network. Installing TLS certificates on the web server can encrypt and secure web traffic between the server and the clients, as well as prevent spoofing or impersonation attacks.

? Disable password authentication: Password authentication is a method of verifying and authenticating users or devices based on passwords or other credentials. Password authentication can be insecure or vulnerable to attacks such as brute force, dictionary, phishing, etc., especially if passwords are weak, reused, or compromised. Disabling password authentication can enhance security by preventing unauthorized or malicious access to the web server using passwords.

? Enable SSH key access only: SSH key access is a method of verifying and authenticating users or devices based on digital keys issued by a trusted authority. SSH key access can provide more security and convenience than password authentication, as it does not require users or devices to remember or enter passwords every time they access the web server. Enabling SSH key access only can ensure that only authorized or trusted users or devices can access the web server using keys.

NEW QUESTION 74

- (Topic 2)

A cloud solutions architect needs to determine the best strategy to deploy an application environment in production, given the following requirements:

No downtime

Instant switch to a new version using traffic control for all users

Which of the following deployment strategies would be the BEST solution?

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

Answer: B

Explanation:

Reference: <https://thenewstack.io/deployment-strategies/>

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

NEW QUESTION 79

- (Topic 2)

A database analyst reports it takes two hours to perform a scheduled job after onboarding 10,000 new users to the system. The analyst made no changes to the scheduled job before or after onboarding the users. The database is hosted in an IaaS instance on a cloud provider. Which of the following should the cloud administrator evaluate to troubleshoot the performance of the job?

- A. The IaaS compute configurations, the capacity trend analysis reports, and the storage IOPS
- B. The hypervisor logs, the memory utilization of the hypervisor host, and the network throughput of the hypervisor
- C. The scheduled job logs for successes and failures, the time taken to execute the job, and the job schedule
- D. Migrating from IaaS to on premises, the network traffic between on-premises users and the IaaS instance, and the CPU utilization of the hypervisor host

Answer: A

Explanation:

To troubleshoot the performance of a scheduled job that takes two hours to run after onboarding 10,000 new users to a cloud-based system, the administrator should evaluate the IaaS compute configurations, the capacity trend analysis reports, and the storage IOPS. These factors can affect the performance of a database job in an IaaS instance on a cloud provider. The IaaS compute configurations include the CPU, memory, and network resources assigned to the instance. The capacity trend analysis reports show the historical and projected usage and demand of the resources. The storage IOPS (Input/Output Operations Per Second) measure the speed and performance of the disk storage. The administrator should check if these factors are sufficient, optimal, or need to be adjusted to improve the performance of the job.

NEW QUESTION 84

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 85

- (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 88

- (Topic 2)

A systems administrator has been asked to restore a VM from backup without changing the current VM's operating state. Which of the following restoration methods would BEST fit this scenario?

- A. Alternate location
- B. Rolling
- C. Storage live migration
- D. In-place

Answer: C

Explanation:

Storage live migration is the best restoration method to restore a VM from backup without changing the current VM's operating state. Storage live migration is a process of moving or transferring storage resources or data from one location to another without affecting or interrupting the operation or performance of the VMs that use them. Storage live migration can help to restore a VM from backup by copying the backup data to a new storage location and switching the VM's storage configuration to point to the new location, without requiring any downtime or reboot.

NEW QUESTION 93

- (Topic 2)

All of a company's servers are currently hosted in one cloud MSP. The company created a new cloud environment with a different MSP. A cloud engineer is now tasked with preparing for server migrations and establishing connectivity between clouds. Which of the following should the engineer perform FIRST?

- A. Peer all the networks from each cloud environment.
- B. Migrate the servers.
- C. Create a VPN tunnel.
- D. Configure network access control lists.

Answer: C

Explanation:

Creating a VPN tunnel is the first action that the engineer should perform to prepare for server migrations and establish connectivity between clouds. A VPN (Virtual Private Network) tunnel is a secure and encrypted connection that allows data to be transferred between two networks or locations over the public internet. Creating a VPN tunnel can enable communication and interoperability between different cloud environments, as well as protect data from interception or modification during migration.

NEW QUESTION 98

- (Topic 2)

Users of an enterprise application, which is configured to use SSO, are experiencing slow connection times. Which of the following should be done to troubleshoot the issue?

- A. Perform a memory dump of the O
- B. Analyze the memory dump.Upgrade the host CPU to a higher clock speed CPU.
- C. Perform a packet capture during authenticatio
- D. Validate the load-balancing configuration.Analyze the network throughput of the load balancer.
- E. Analyze the storage system IOP
- F. Increase the storage system capacit
- G. Replace the storage system disks to SS
- H. Evaluate the OS ACL
- I. Upgrade the router firmware.Increase the memory of the router.

Answer: B

Explanation:

These are the steps that should be done to troubleshoot the issue of slow connection times for users of an enterprise application that is configured to use SSO (Single Sign-On). SSO is a feature that allows users to access multiple applications or services with one login credential, without having to authenticate separately for each application or service. SSO can improve user experience and security, but it may also introduce performance issues if not configured properly. To troubleshoot the issue, the administrator should perform a packet capture during authentication to analyze the network traffic and identify any delays or errors in the SSO process. The administrator should also validate the load-balancing configuration to ensure that the SSO requests are distributed evenly and efficiently among the available servers or instances. The administrator should also analyze the network throughput of the load balancer to check if there is any congestion or bottleneck that may affect the SSO performance.

NEW QUESTION 101

- (Topic 2)

A company has an in-house-developed application. The administrator wants to utilize cloud services for additional peak usage workloads. The application has a very unique stack of dependencies.

Which of the following cloud service subscription types would BEST meet these requirements?

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

Answer: D

Explanation:

IaaS (Infrastructure as a Service) is a cloud service model that provides basic computing resources such as servers, storage, network, etc., to the customers. The customers have full control and flexibility over these resources and can install and configure any software they need on them. IaaS is suitable for applications that have a unique stack of dependencies that may not be supported by other cloud service models.

NEW QUESTION 103

- (Topic 2)

An update is being deployed to a web application, and a systems administrator notices the cloud SQL database has stopped running. The VM is responding to pings, and there were not any configuration changes scheduled for the VM. Which of the following should the administrator check NEXT?

- A. Logs on the VM
- B. Firewall on the VM
- C. Memory on the VM
- D. vGPU performance on the VM

Answer: A

Explanation:

Checking the logs on the VM is the next step that the administrator should take if the cloud SQL database has stopped running after an update deployment. Logs are records of events and activities that occur on a system or application. Logs can provide useful information for troubleshooting and identifying the root cause of an issue. The administrator should look for any errors, warnings, or messages that indicate what happened to the SQL database service and why it stopped running.

NEW QUESTION 107

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 111

- (Topic 2)

A company is currently running a website on site. However, because of a business requirement to reduce current RTO from 12 hours to one hour, and the RPO from one day to eight hours, the company is considering operating in a hybrid environment. The website uses mostly static files and a small relational database. Which of the following should the cloud architect implement to achieve the objective at the LOWEST cost possible?

- A. Implement a load-balanced environment in the cloud that is equivalent to the current on-premises setup and use DNS to shift the load from on-premises to cloud.
- B. Implement backups to cloud storage and infrastructure as code to provision the environment automatically when the on-premises site is down.
- C. Restore the data from the backups.
- D. Implement a website replica in the cloud with auto-scaling using the smallest possible footprint.
- E. Use DNS to shift the load from on-premises to the cloud.
- F. Implement a CDN that caches all requests with a higher TTL and deploy the IaaS instances manually in case of disaster.
- G. Upload the backup on demand to the cloud to restore on the new instances.

Answer: C

Explanation:

This is the best solution to achieve the objective of reducing current RTO (Recovery Time Objective) from 12 hours to one hour, and RPO (Recovery Point Objective) from one day to eight hours, at the lowest cost possible, for a website that uses mostly static files and a small relational database. RTO is a metric that measures how quickly a system or service can be restored after a disruption or disaster. RPO is a metric that measures how much data can be lost or how far back in time a recovery point can be without causing significant impact or damage. To reduce RTO and RPO, the administrator should implement a website replica in the cloud with auto-scaling using the smallest possible footprint. A website replica is a copy or backup of a website that can be used for recovery or failover purposes. Auto-scaling is a feature that allows cloud resources or systems to adjust their capacity and performance according to demand or workload. Using auto-scaling with the smallest possible footprint can minimize costs by using only the necessary resources and scaling up or down as needed. The administrator should also use DNS (Domain Name System) to shift the load from on-premises to the cloud. DNS is a service that translates domain names into IP addresses and vice versa. Using DNS, the administrator can redirect traffic from the on-premises website to the cloud replica in case of a disruption or disaster, and vice versa when recovery is complete.

NEW QUESTION 112

- (Topic 2)

A cloud administrator needs to reduce the cost of cloud services by using the company's off-peak period. Which of the following would be the BEST way to achieve this with minimal effort?

- A. Create a separate subscription.
- B. Create tags.
- C. Create an auto-shutdown group.
- D. Create an auto-scaling group.

Answer: C

Explanation:

Creating an auto-shutdown group is the best way to reduce the cost of cloud services by using the company's off-peak period with minimal effort. An auto-shutdown group is a feature that allows customers to automatically turn off or shut down certain cloud resources or services during a specified time period or schedule. An auto-shutdown group can help to reduce the cost of cloud services by minimizing the consumption of resources or services during off-peak periods, when they are not needed or used. An auto-shutdown group can also help to reduce the effort of managing cloud resources or services by automating the shutdown process, without requiring any manual intervention or configuration.

NEW QUESTION 116

- (Topic 2)

A cloud administrator is responsible for managing a cloud-based content management solution. According to the security policy, any data that is hosted in the cloud must be protected against data exfiltration. Which of the following solutions should the administrator implement?

- A. HIDS
- B. FIM
- C. DLP
- D. WAF

Answer: C

Explanation:

DLP (Data Loss Prevention) is what the administrator should implement to protect data against data exfiltration in a cloud-based content management solution. Data exfiltration is a process of transferring or stealing data from a system or network without authorization or permission. Data exfiltration can cause data breaches, leaks, or losses that may affect confidentiality, integrity, or availability of data. DLP is a tool or service that monitors and controls data movement and usage within a system or network. DLP can help to prevent data exfiltration by detecting and blocking any unauthorized or suspicious data transfers or activities, as well as enforcing policies and rules for data classification, encryption, access, etc.

NEW QUESTION 121

- (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 125

- (Topic 2)

Which of the following service models would be used for a database in the cloud?

- A. PaaS
- B. IaaS
- C. CaaS
- D. SaaS

Answer: A

Explanation:

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

NEW QUESTION 129

- (Topic 2)

A systems administrator is examining a managed hosting agreement and wants to determine how much data would be lost if a server had to be restored from backups. To which of the following metrics should the administrator refer?

- A. RTO
- B. MTBF
- C. RPO
- D. MTTR

Answer: C

Explanation:

RPO (Recovery Point Objective) is the metric that the administrator should refer to determine how much data would be lost if a server had to be restored from backups. RPO is a metric that measures how much data can be lost or how far back in time a recovery point can be without causing significant impact or damage. RPO can help to determine how much data would be lost by comparing the time of the disruption or disaster with the time of the last backup or snapshot. RPO can also help to determine how frequently backups or snapshots should be performed to minimize data loss.

NEW QUESTION 132

- (Topic 2)

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

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

Answer: D

Explanation:

A risk register is a document that records all the identified risks, their causes, impacts, probabilities, mitigation measures, and status for a project or an

organization. A risk register helps to manage and monitor risks throughout their lifecycle and ensure they are addressed appropriately. A risk register would help the CISO to locate all the assets with identified deviations and mitigation measures.

NEW QUESTION 135

- (Topic 2)

An administrator has been informed that some requests are taking a longer time to respond than other requests of the same type. The cloud consumer is using multiple network service providers and is performing link load balancing for bandwidth aggregation. Which of the following commands will help the administrator understand the possible latency issues?

- A. ping
- B. ipconfig
- C. traceroute
- D. netstat

Answer: A

Explanation:

Ping is the command that will help the administrator understand the possible latency issues between different network service providers and link load balancing for bandwidth aggregation. Ping is a network utility that sends packets of data to a specific IP address or hostname and measures the time it takes for them to be sent back (round-trip time). Ping can help to test connectivity, availability, and latency of network devices or systems. Ping can help to understand latency issues by comparing the round-trip times between different network service providers and link load balancing devices, and identifying any delays or variations in response times.

NEW QUESTION 137

- (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 138

- (Topic 2)

A systems administrator is deploying a VM and would like to minimize storage utilization by ensuring the VM uses only the storage it needs. Which of the following will BEST achieve this goal?

- A. Compression
- B. Deduplication
- C. RAID
- D. Thin provisioning

Answer: D

Explanation:

Reference: https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vm_admin.doc/GUID-4C0F4D73-82F2-4B81-8AA7-1DD752A8A5AC.html

Thin provisioning is the technique that will minimize storage utilization by ensuring the VM uses only the storage it needs. 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 139

- (Topic 2)

A company needs to migrate the storage system and batch jobs from the local storage system to a public cloud provider. Which of the following accounts will MOST likely be created to run the batch processes?

- A. User
- B. LDAP
- C. Role-based
- D. Service

Answer: D

Explanation:

A service account is what will most likely be created to run the batch processes that migrate the storage system and batch jobs from the local storage system to a public cloud provider. A service account is a special type of account that is used to perform automated tasks or operations on a system or service, such as running scripts, applications, or processes. A service account can provide benefits such as:

? Security: A service account can have limited or specific permissions and roles that are required to perform the tasks or operations, which can prevent unauthorized or malicious access or actions.

? Efficiency: A service account can run the tasks or operations without any human intervention or interaction, which can save time and effort.

? Reliability: A service account can run the tasks or operations consistently and accurately, which can reduce errors or failures.

NEW QUESTION 141

- (Topic 2)

An administrator is securing a private cloud environment and wants to ensure only approved systems can connect to switches. Which of the following would be MOST useful to accomplish this task?

- A. VLAN
- B. NIPS
- C. WAF
- D. NAC

Answer: D

Explanation:

Reference: <https://www.cisco.com/c/en/us/products/security/what-is-network-access-control-nac.html>

NAC (Network Access Control) is what the administrator should implement to ensure only approved systems can connect to switches in a private cloud environment. NAC is a security technique that controls and restricts access to network resources based on predefined policies or rules. NAC can verify and authenticate users or devices before granting them access to switches or other network devices. NAC can also enforce compliance and security standards on users or devices before allowing them to connect to switches.

NEW QUESTION 142

- (Topic 1)

Company A has acquired Company B and is in the process of integrating their cloud resources. Company B needs access to Company A's cloud resources while retaining its IAM solution.

Which of the following should be implemented?

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

Answer: C

Explanation:

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

Reference: <https://medium.com/@dinika.15/identity-federation-a-brief-introduction-f2f823f8795a>

NEW QUESTION 145

- (Topic 1)

A systems administrator needs to configure SSO authentication in a hybrid cloud environment.

Which of the following is the BEST technique to use?

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

Answer: B

Explanation:

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

References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 148

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).

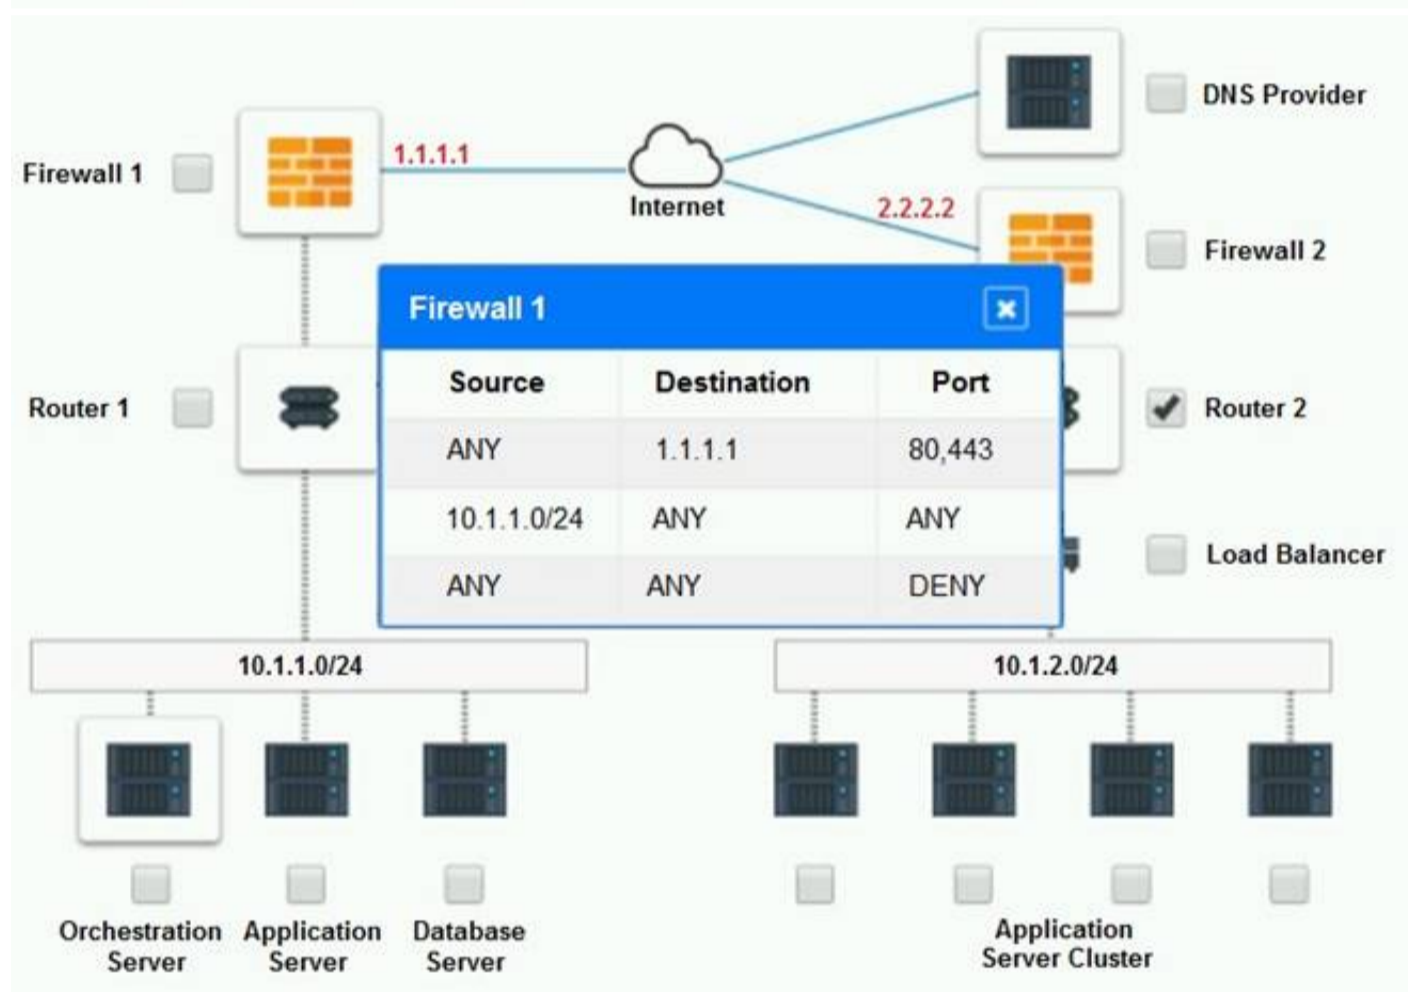
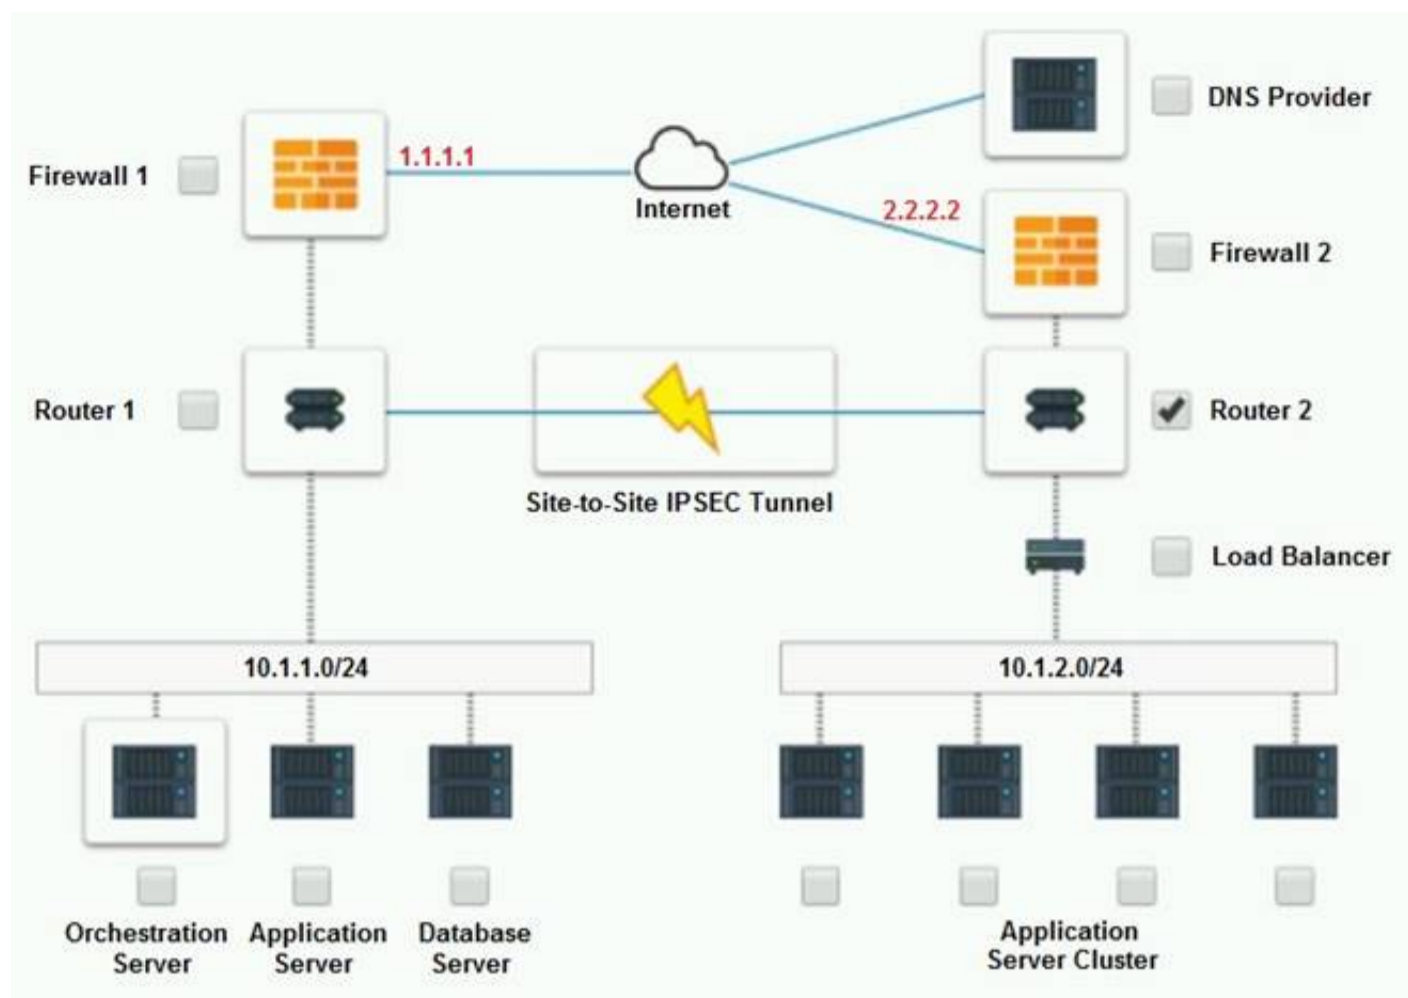
Part 2:

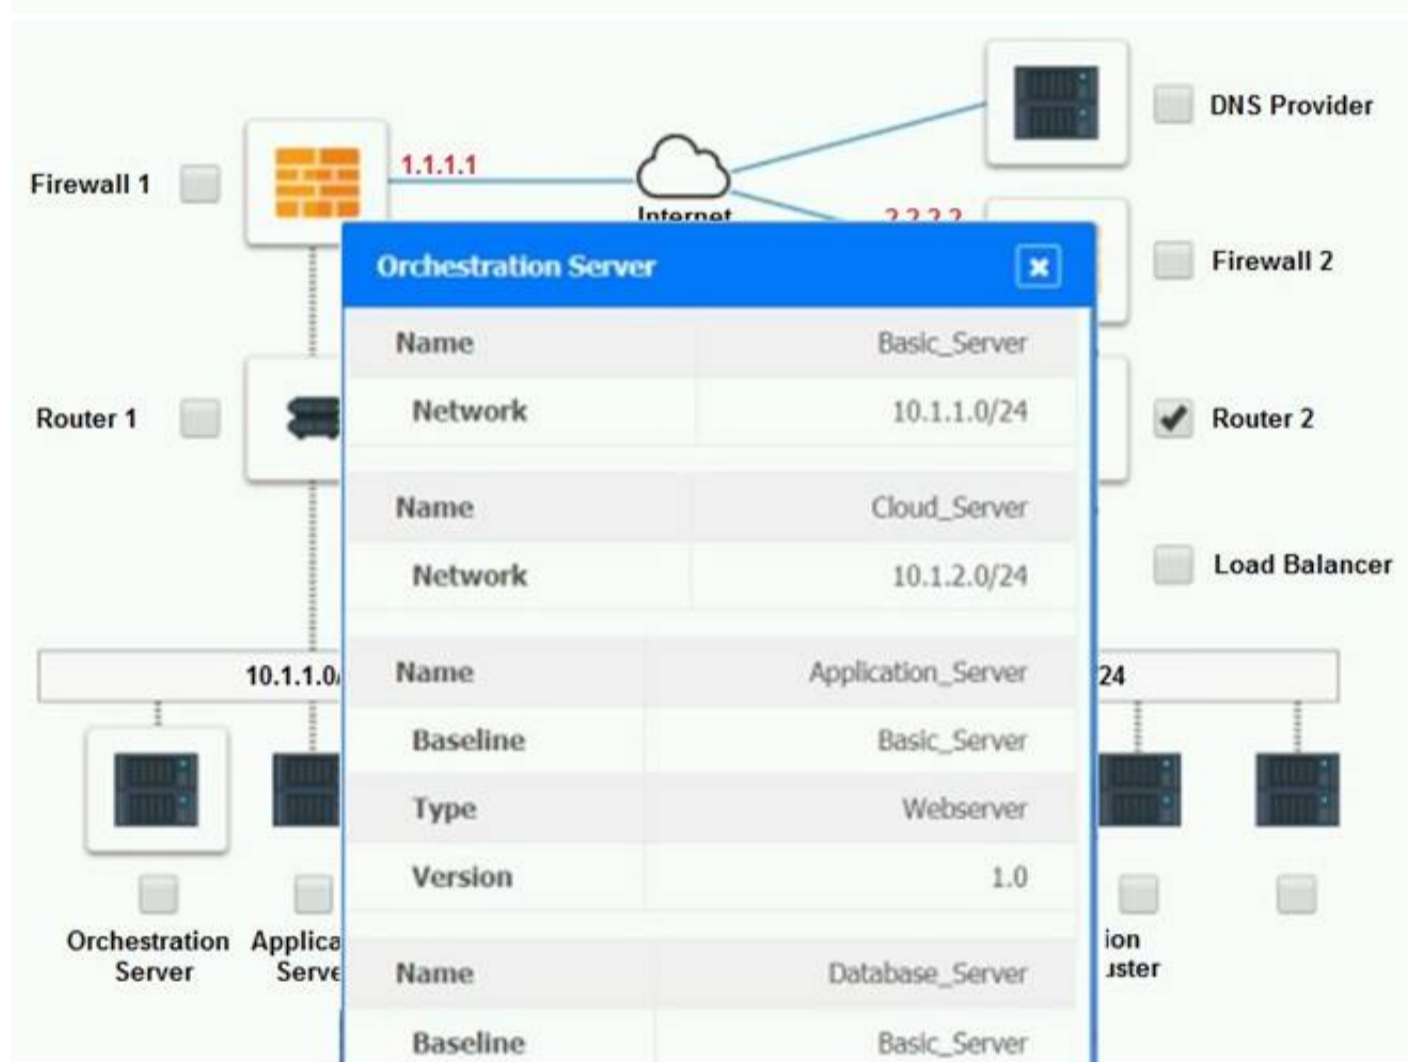
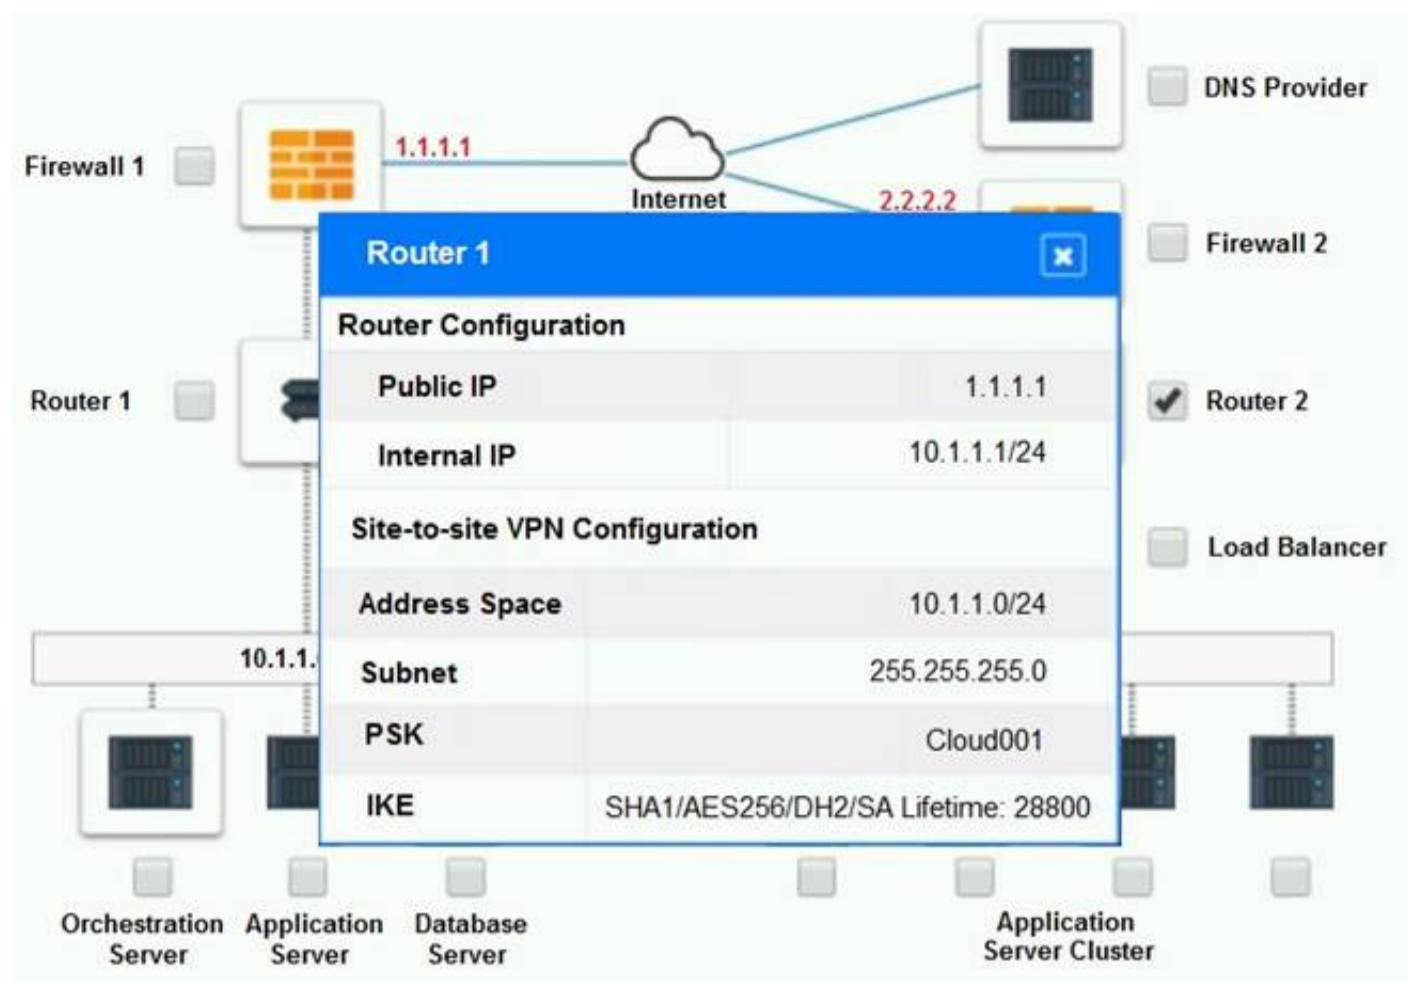
- _ Identify the correct options to provide adequate configuration for hybrid cloud architecture.

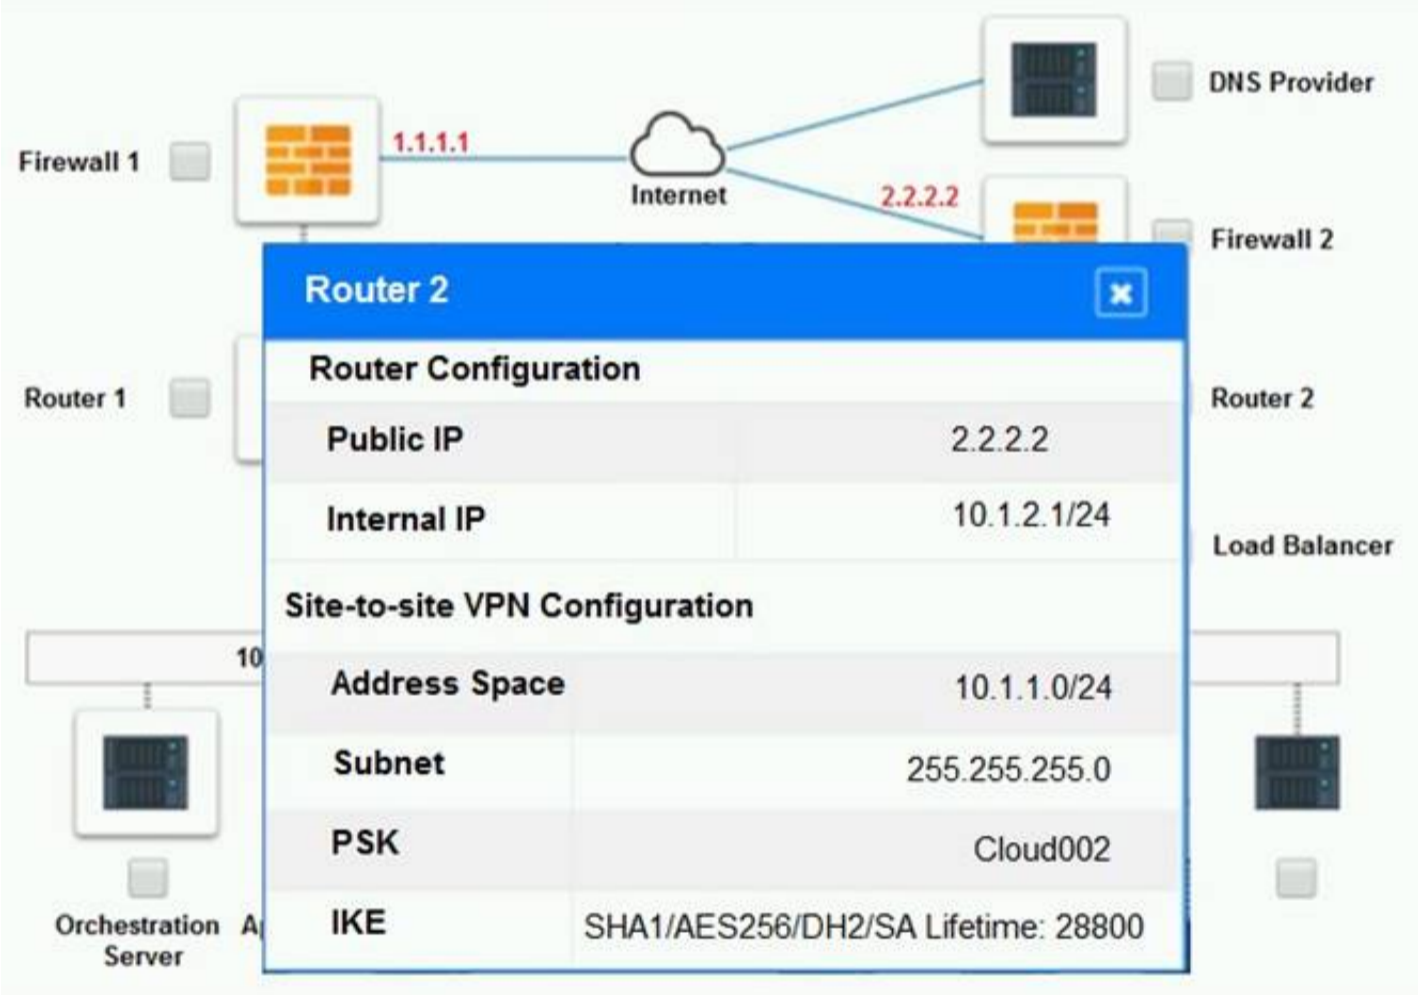
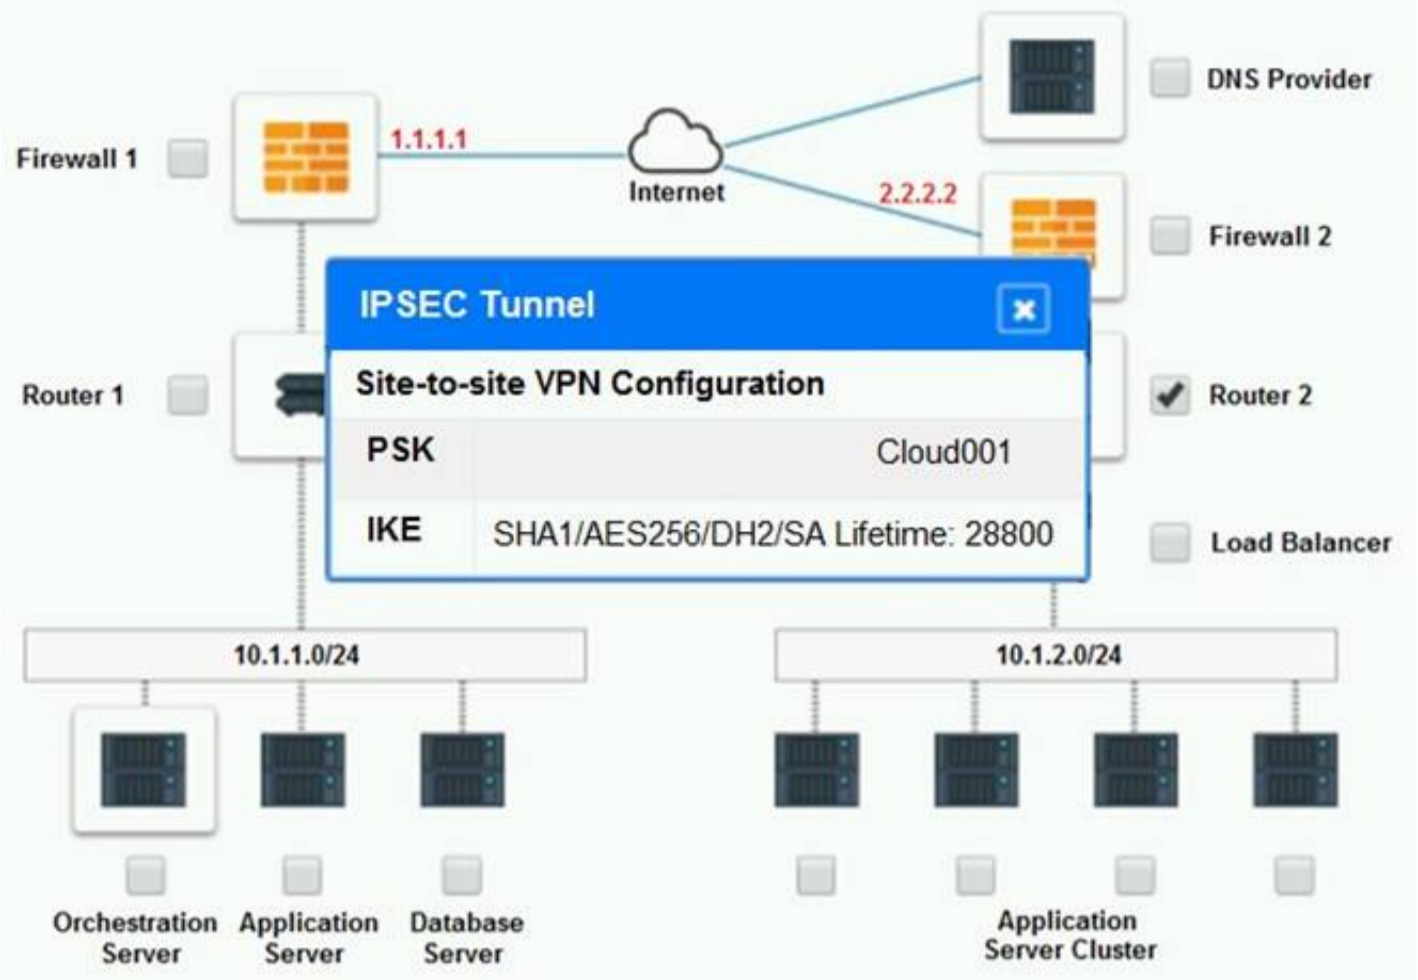
If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

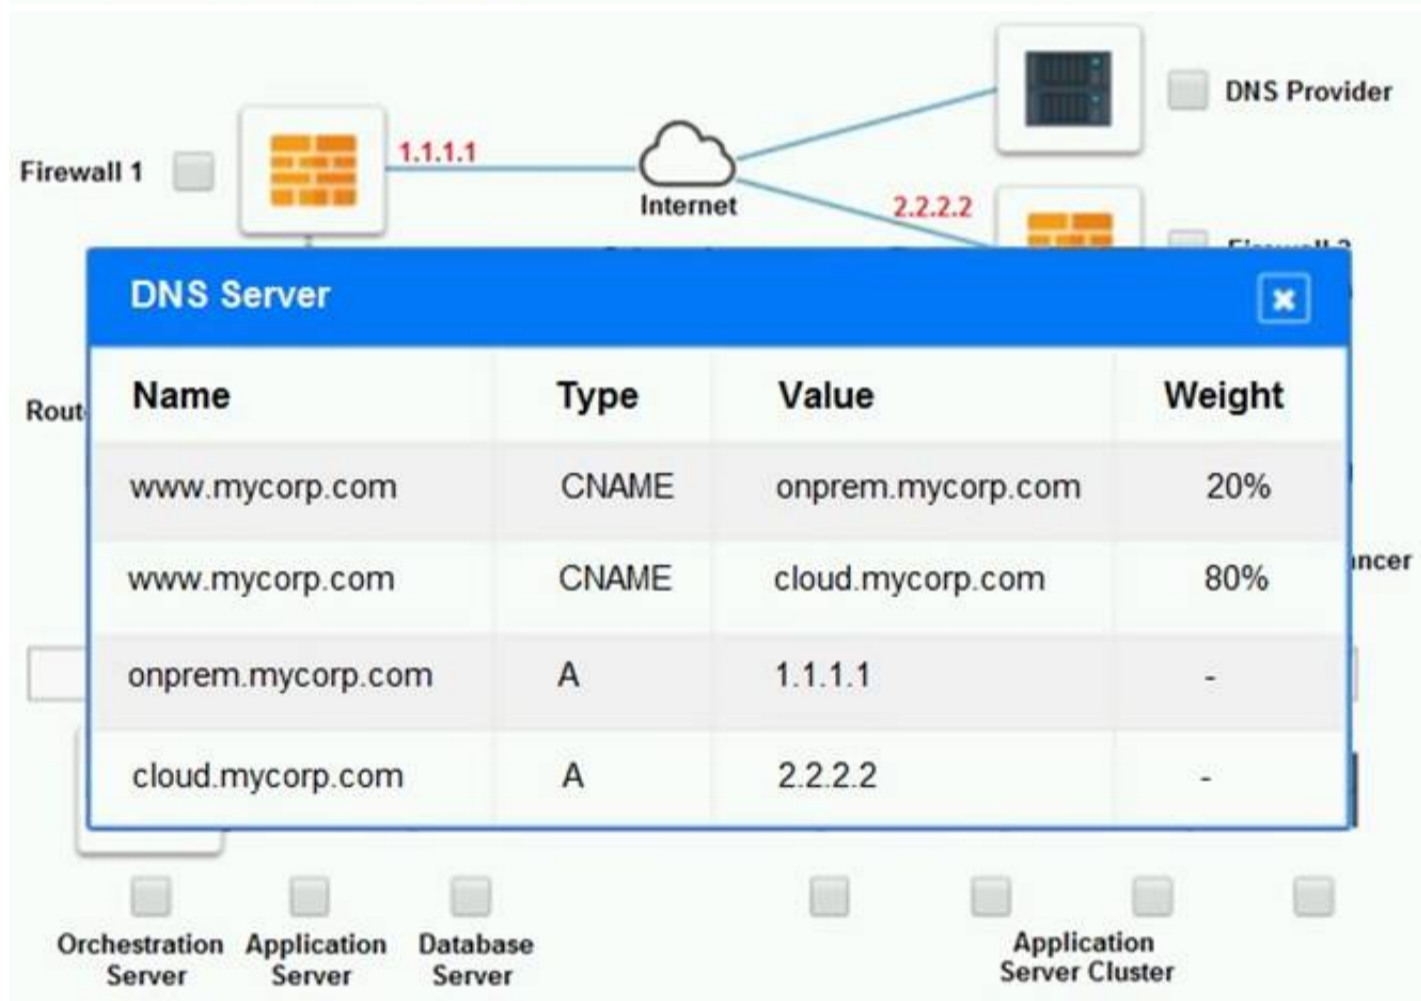
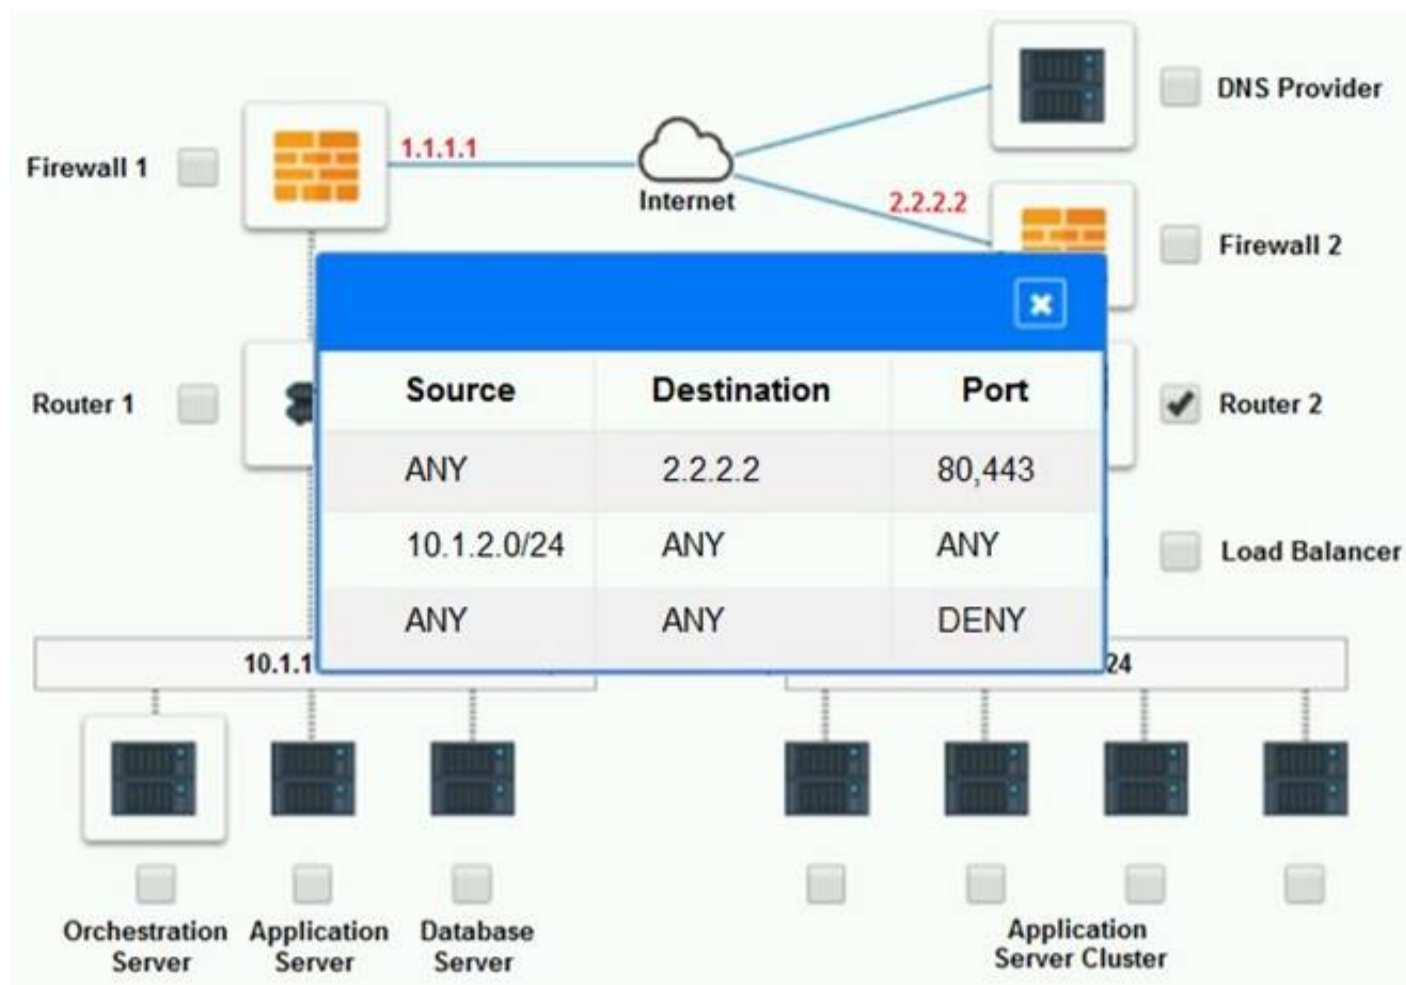
Part 1:

Cloud Hybrid Network Diagram









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 152

- (Topic 1)

A company is utilizing a private cloud solution that is hosted within its datacenter. The company wants to launch a new business application, which requires the resources below:

Maximum concurrent sessions	Number of nodes required	Required per-node vCPU	Required per-node RAM
1,000	2	4	32
5,000	4	6	64
10,000	6	8	64
25,000	8	8	128

The current private cloud has 30 vCPUs and 512GB RAM available. The company is looking for a quick solution to launch this application, with expected maximum sessions to be close to 24,000 at launch and an average of approximately 5,000 sessions.

Which of the following solutions would help the company accommodate the new workload in the SHORTEST amount of time and with the maximum financial benefits?

- A. Configure auto-scaling within the private cloud
- B. Set up cloud bursting for the additional resources
- C. Migrate all workloads to a public cloud provider
- D. Add more capacity to the private cloud

Answer: B

Explanation:

Cloud Bursting can be used for both compute and storage. This question is about compute capability. "Compute Bursting" unleashes the high-performance compute capabilities of the cloud for processing locally created datasets. (reference: <https://www.ctera.com/it-initiatives/cloud-bursting/>)

<https://azure.microsoft.com/en-us/overview/what-is-cloud-bursting/>

NEW QUESTION 154

- (Topic 1)

A cloud architect is designing the VPCs for a new hybrid cloud deployment. The business requires the following:

? High availability

? Horizontal auto-scaling

? 60 nodes peak capacity per region

? Five reserved network IP addresses per subnet

? /24 range

Which of the following would BEST meet the above requirements?

- A. Create two /25 subnets in different regions
- B. Create three /25 subnets in different regions
- C. Create two /26 subnets in different regions
- D. Create three /26 subnets in different regions
- E. Create two /27 subnets in different regions
- F. Create three /27 subnets in different regions

Answer: C

Explanation:

A /26 subnet is a subnet that has a network prefix of 26 bits and a host prefix of 6 bits. A /26 subnet can support up to 64 hosts (62 usable hosts) and has a subnet mask of 255.255.255.192. Creating two /26 subnets in different regions can best meet the business requirements for deploying a high availability, horizontally auto-scaling solution that has a peak capacity of 60 nodes per region and five reserved network IP addresses per subnet. Creating two /26 subnets can provide enough host addresses for the peak capacity and the reserved addresses, as well as allow for some growth or redundancy. Creating the subnets in different regions can provide high availability and horizontal auto- scaling, as it can distribute the workload across multiple locations and scale out or in based on demand. References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8

NEW QUESTION 156

- (Topic 1)

A systems administrator is troubleshooting network throughput issues following a deployment. The network is currently being overwhelmed by the amount of traffic between the database and the web servers in the environment.

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

- A. Set up affinity rules to keep web and database servers on the same hypervisor
- B. Enable jumbo frames on the gateway
- C. Move the web and database servers onto the same VXLAN
- D. Move the servers onto thick-provisioned storage

Answer: C

Explanation:

A virtual extensible local area network (VXLAN) is a type of network virtualization technology that creates logical networks or segments that span across multiple physical networks or locations. Moving the web and database servers onto the same VXLAN can help resolve the network throughput issues following a deployment, as it can reduce the network traffic between the database and the web servers by using a common virtual network identifier (VNI) and encapsulating the traffic within UDP packets. Moving the web and database servers onto the same VXLAN can also improve performance and security, as it can provide higher scalability, isolation, and encryption for the network traffic. References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8

NEW QUESTION 161

- (Topic 1)

A systems administrator is deploying a solution that requires a virtual network in a private cloud environment. The solution design requires the virtual network to transport multiple payload types.

Which of the following network virtualization options would BEST satisfy the requirement?

- A. VXLAN
- B. STT
- C. NVGRE
- D. GENEVE

Answer: D

Explanation:

Generic Network Virtualization Encapsulation (GENEVE) is a type of network virtualization technology that creates logical networks or segments that span across multiple physical networks or locations. GENEVE can satisfy the requirement of transporting multiple payload types in a virtual network in a private cloud environment, as it can support various network protocols and services by using a flexible and extensible header format that can encapsulate different types of payloads within UDP packets. GENEVE can also provide interoperability and compatibility, as it can integrate with existing network virtualization technologies such as VXLAN, STT, or NVGRE. References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8

NEW QUESTION 163

- (Topic 1)

A systems administrator is building a new virtualization cluster. The cluster consists of five virtual hosts, which each have flash and spinning disks. This storage is shared among all the virtual hosts, where a virtual machine running on one host may store data on another host.

This is an example of:

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

Answer: C

Explanation:

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

NEW QUESTION 168

- (Topic 1)

A systems administrator wants to have near-real-time information on the volume of data being exchanged between an application server and its clients on the Internet.

Which of the following should the systems administrator implement to achieve this objective?

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

Answer: D

Explanation:

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

of data being exchanged between an application server and its clients on the Internet. References: CompTIA Cloud+ Certification Exam Objectives, page 16, section 3.2

NEW QUESTION 172

- (Topic 1)

A systems administrator is reviewing two CPU models for a cloud deployment. Both CPUs have the same number of cores/threads and run at the same clock speed.

Which of the following will BEST identify the CPU with more computational power?

- A. Simultaneous multithreading
- B. Bus speed
- C. L3 cache
- D. Instructions per cycle

Answer: D

Explanation:

Instructions per cycle (IPC) is a metric that measures how many instructions a CPU can execute in one clock cycle. IPC can help identify the CPU with more computational power when comparing two CPU models that have the same number of cores/threads and run at the same clock speed, as it indicates the efficiency and performance of the CPU architecture and design. A higher IPC means that the CPU can process more instructions in less time, resulting in faster and better performance. References: CompTIA Cloud+ Certification Exam Objectives, page 9, section 1.4

Reference: https://en.wikipedia.org/wiki/Central_processing_unit

NEW QUESTION 174

- (Topic 1)

Which of the following will mitigate the risk of users who have access to an instance modifying the system configurations?

- A. Implement whole-disk encryption
- B. Deploy the latest OS patches
- C. Deploy an anti-malware solution
- D. Implement mandatory access control

Answer: D

Explanation:

Mandatory access control (MAC) is a type of access control model that enforces strict security policies based on predefined rules and labels. MAC assigns security labels to subjects (users or processes) and objects (files or resources) and allows access only if the subject has the appropriate clearance and need-to-know for the object. MAC can mitigate the risk of users who have access to an instance modifying the system configurations, as it can prevent unauthorized or accidental changes to critical files or settings by restricting access based on predefined rules and labels. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 175

- (Topic 1)

A systems administrator is using VMs to deploy a new solution that contains a number of application VMs.

Which of the following would provide high availability to the application environment in case of hypervisor failure?

- A. Anti-affinity rules
- B. Cold migration
- C. Live migration
- D. Affinity rules

Answer: A

Explanation:

Anti-affinity rules are rules or policies that prevent two or more VMs from running on the same host or cluster in a cloud environment. Anti-affinity rules can provide high availability to an application environment in case of hypervisor failure, as they can distribute or separate the application VMs across different hosts or clusters and avoid having a single point of failure. Anti-affinity rules can also improve performance and reliability, as they can reduce contention and load by balancing the resource utilization across multiple hosts or clusters. References: CompTIA Cloud+ Certification Exam Objectives, page 10, section 1.5

Reference: <https://www.vmware.com/products/vsphere/high-availability.html>

NEW QUESTION 177

- (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 180

- (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 184

- (Topic 1)

A company developed a product using a cloud provider's PaaS platform and many of the platform-based components within the application environment. Which of the following would the company MOST likely be concerned about when utilizing a multicloud strategy or migrating to another cloud provider?

- A. Licensing
- B. Authentication providers
- C. Service-level agreement
- D. Vendor lock-in

Answer: D

Explanation:

Vendor lock-in is a situation where a customer becomes dependent on a specific vendor for products or services and faces high switching costs or barriers when trying to change vendors. Vendor lock-in is most likely to be a concern for a company that developed a product using a cloud provider's PaaS platform and many of the platform-based components within the application environment when utilizing a multicloud strategy or migrating to another cloud provider, as it can limit the flexibility, scalability, and portability of the product and increase the complexity, risk, and cost of moving or integrating with other cloud platforms or providers. References: CompTIA Cloud+ Certification Exam Objectives, page 8, section 1.2

NEW QUESTION 188

- (Topic 1)

After accidentally uploading a password for an IAM user in plain text, which of the following should a cloud administrator do FIRST? (Choose two.)

- A. Identify the resources that are accessible to the affected IAM user
- B. Remove the published plain-text password
- C. Notify users that a data breach has occurred
- D. Change the affected IAM user's password
- E. Delete the affected IAM user

Answer: BD

Explanation:

Removing the published plain-text password and changing the affected IAM user's password are the first actions that a cloud administrator should take after accidentally uploading a password for an IAM user in plain text, as they can prevent or limit any unauthorized or malicious access to the cloud resources or services using the compromised password. Removing the published plain-text password can ensure that the password is not exposed or available to anyone who may access or view the uploaded file. Changing the affected IAM user's password can ensure that the password is updated and secured using encryption or hashing techniques. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 192

- (Topic 1)

An organization is required to set a custom registry key on the guest operating system. Which of the following should the organization implement to facilitate this requirement?

- A. A configuration management solution
- B. A log and event monitoring solution
- C. A file integrity check solution
- D. An operating system ACL

Answer: A

Explanation:

A configuration management solution is a type of tool or system that automates and standardizes the configuration and deployment of cloud resources or services according to predefined policies or rules. A configuration management solution can help set a custom registry key on the guest operating system in an IaaS instance, as it can apply the desired registry setting to one or more virtual machines (VMs) without manual intervention or scripting. A configuration management solution can also help maintain consistency, compliance, and security of cloud configurations by monitoring and enforcing the desired state. References: CompTIA Cloud+ Certification Exam Objectives, page 13, section 2.5

NEW QUESTION 195

- (Topic 4)

Which of the following are advantages of a public cloud? (Select TWO).

- A. Full control of hardware
- B. Reduced monthly costs
- C. Decreased network latency
- D. Pay as you use
- E. Availability of self-service
- F. More secure data

Answer: BD

Explanation:

The correct answers are B and D.

* B. Reduced monthly costs: One of the main advantages of public cloud is that it lowers the costs of IT infrastructure and maintenance for the customers. They do not need to purchase, install, or manage any hardware or software, and they only pay for the resources they use. This can result in significant savings compared to owning and operating a private cloud or an on-premise data center¹²³⁴

* D. Pay as you use: Another benefit of public cloud is that it offers a flexible and scalable pricing model based on the actual usage of the customers. They can adjust their resource consumption according to their changing needs and demands, and only pay for what they use. This eliminates the need for upfront capital investment or long-term contracts, and allows customers to optimize their spending and performance¹²³⁴

NEW QUESTION 196

- (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 197

- (Topic 4)

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

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

Answer: C

Explanation:

A staging environment is a type of development environment that is used to test and demonstrate the final product before deploying it to the production environment. A staging environment can help the web consultancy group avoid the issues of delivering a different or faulty product to the customers, as it can ensure that the product is fully functional, compatible, and secure. A staging environment can also help the group showcase the product to the customers in a realistic and controlled way, as it can mimic the production environment and avoid any interference from other development activities. A staging environment can be leveraged by using cloud services that allow for easy provisioning, scaling, and deployment of web applications

NEW QUESTION 199

- (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 202

- (Topic 4)

A systems administrator audits a cloud application and discovers one of the key regulatory requirements has not been addressed. The requirement states that if a physical breach occurs and hard drives are stolen, the contents of the drives should not be readable. Which of the following should be used to address the requirement?

- A. Obfuscation
- B. Encryption
- C. EDR
- D. HIPS

Answer: B

Explanation:

Encryption is the process of transforming data into an unreadable format using a secret key or algorithm. Encryption can be used to protect data at rest or in transit from unauthorized access or theft. If a physical breach occurs and hard drives are stolen, encryption can prevent the contents of the drives from being readable by anyone who does not have the decryption key or algorithm.

References: [CompTIA Cloud+ Study Guide], page 236.

NEW QUESTION 203

- (Topic 4)

The Chief Information Officer of a financial services company wants to ensure stringent security measures are maintained while migrating customer financial information from a private cloud to the public cloud. The cloud engineer must deploy automated validation and verification checks to prevent unauthorized disclosure of financial information. Which of the following should be configured during the migration?

- A. ACL
- B. VPN
- C. P2V
- D. VDI

Answer: B

Explanation:

One possible answer is: B. VPN

A VPN (Virtual Private Network) is a technology that creates a secure and encrypted connection between a remote device and a private network over the internet. A VPN can help prevent unauthorized disclosure of financial information during the migration from a private cloud to the public cloud, as it can protect the data in transit from interception, tampering, or leakage. A VPN can also help maintain compliance with data privacy regulations, such as GDPR or PCI DSS, by ensuring that the data is only accessible by authorized parties¹².

ACL (Access Control List) is a method of controlling access to resources based on user or group permissions. ACL can help enforce security policies and restrict access to sensitive data, but it does not encrypt or protect the data in transit³.

P2V (Physical to Virtual) is a process of converting a physical machine into a virtual machine. P2V can help migrate workloads from on-premises servers to cloud servers, but it does not ensure the security of the data during the migration⁴.

VDI (Virtual Desktop Infrastructure) is a technology that provides users with virtual desktops hosted on a centralized server. VDI can help improve the performance, availability, and manageability of desktop environments, but it does not address the security of the data during the migration⁵.

NEW QUESTION 208

- (Topic 4)

A company is using IaaS services from two different providers: one for its primary site, and the other for a secondary site. The primary site is completely inaccessible, and the management team has decided to run through the BCP procedures. Which of the following will provide the complete asset information?

- A. DR replication document
- B. DR playbook
- C. DR policies and procedures document
- D. DR network diagram

Answer: B

Explanation:

According to the CompTIA Cloud+ CV0-003 Certification Study Guide¹, the answer is B. DR playbook. A DR playbook is a document that contains the detailed steps and procedures to recover from a disaster scenario. It includes the asset information, such as the cloud resources, configurations, and dependencies, that are needed to restore the normal operations of the business. A DR replication document is a document that describes how the data and applications are replicated between the primary and secondary sites. A DR policies and procedures document is a document that defines the roles and responsibilities of the staff, the communication channels, and the objectives and scope of the DR plan. A DR network diagram is a visual representation of the network topology and connectivity between the primary and secondary sites.

NEW QUESTION 212

- (Topic 4)

A company is concerned it will run out of VLANs on its private cloud platform in the next couple months, and the product currently offered to customers requires the company to allocate three dedicated, segmented tiers. Which of the following can the company implement to continue adding new customers and to maintain the required level of isolation from other tenants?

- A. GRE
- B. SR-IOV
- C. VXLAN
- D. IPSec

Answer: C

Explanation:

One possible solution for the company to continue adding new customers and to maintain the required level of isolation from other tenants is to implement VXLAN. VXLAN is a network virtualization technology that can extend VLAN by adding a 24-bit segment ID, which allows up to 16 million unique virtual segments. VXLAN can encapsulate layer 2 Ethernet frames within layer 3 IP packets, and tunnel them across the underlying network. VXLAN can provide logical isolation and security for different tenants, as well as scalability and flexibility for large cloud computing environments¹.

NEW QUESTION 216

- (Topic 4)

A systems administrator is working within a private cloud environment. Over time, random 4K read/write speeds on all VMS in the environment slow down until the VMS are completely unusable, with disk speeds of less than 1MBps. The administrator has gathered the information below:

- There is no correlation between the slowdown and VM/hypervisor resource utilization.
- The network is rated to 40Gbps and utilization is between 1—5%.
- The hypervisors use hundreds of NFSv3 mounts to the same storage appliance, one per VM.
- The VMS on each hypervisor become unresponsive after two weeks of uptime.
- The unresponsiveness is resolved by moving slow VMS onto a rebooted hypervisor. Which of the following solutions will MOST likely resolve this issue?

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

Answer: D

Explanation:

The correct answer is D. Reduce the number of NFSv3 mounts to one.

NFSv3 is a network file system protocol that allows clients to access files stored on a remote server. NFSv3 uses TCP or UDP as the transport layer protocol, and typically runs on port 20491.

One of the known issues with NFSv3 mounts is that they can cause performance degradation and unresponsiveness on the client side if there are too many mounts or if there are network connectivity problems. This is because NFSv3 does not handle connection failures or timeouts gracefully, and may keep retrying to access the server indefinitely, blocking other processes or threads. This can result in slow disk speeds, high CPU usage, and system hangs²³.

Therefore, one of the possible solutions to this issue is to reduce the number of NFSv3 mounts to one per hypervisor, instead of one per VM. This way, the hypervisor can manage the access to the shared storage appliance more efficiently, and avoid creating too many TCP connections or UDP packets that may overload the network or the server. Reducing the number of NFSv3 mounts can also simplify the configuration and troubleshooting of the network file system. Increasing caching on the storage appliance may improve the read performance of the NFSv3 mounts, but it will not solve the underlying issue of connection failures or timeouts. Caching may also introduce data inconsistency or corruption issues if the cache is not synchronized with the server.

Configuring jumbo frames on the hypervisors and storage may improve the network throughput and efficiency of the NFSv3 mounts, but it will not solve the underlying issue of connection failures or timeouts. Jumbo frames are larger than standard Ethernet frames, and require that all devices on the network path support them. Jumbo frames may also introduce fragmentation or compatibility issues if they are not configured properly. Increasing CPU/RAM resources on affected VMs may improve their performance in general, but it will not solve the underlying issue of connection failures or timeouts. Increasing CPU/RAM resources may also be costly and wasteful if they are not needed for other purposes.

NEW QUESTION 220

- (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 224

- (Topic 4)

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

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

Answer: AB

Explanation:

Telnet and FTP are two services that should be disabled on a cloud server because they are insecure and vulnerable to attacks. Telnet and FTP use plain text to transmit data over the network, which means that anyone who can intercept the traffic can read or modify the data, including usernames, passwords, commands, files, etc. This can lead to data breaches, unauthorized access, or malicious actions on the server¹.

Instead of Telnet and FTP, more secure alternatives should be used, such as SSH (Secure Shell) and SFTP (Secure File Transfer Protocol). SSH and SFTP use encryption to protect the data in transit and provide authentication and integrity checks for the communication. SSH and SFTP can prevent eavesdropping, tampering, or spoofing of the data and ensure the confidentiality and privacy of the server².

The other options are not services that should be disabled on a cloud server:

? Option C: Remote login. Remote login is a service that allows users to access a remote server from another location using a network connection. Remote login can be useful for managing, configuring, or troubleshooting a cloud server without having to physically access it. Remote login can be secured by using encryption, authentication, authorization, and logging mechanisms³.

? Option D: DNS (Domain Name System). DNS is a service that translates human- friendly domain names into IP addresses that can be used to communicate over the Internet. DNS is essential for resolving the names of the cloud resources and services that are hosted on the cloud platform. DNS can be secured by using DNSSEC (DNS Security Extensions), which add digital signatures to DNS records to verify their authenticity and integrity.

? Option E: DHCP (Dynamic Host Configuration Protocol). DHCP is a service that assigns IP addresses and other network configuration parameters to devices on a network. DHCP can simplify the management of IP addresses and avoid conflicts or errors in the network. DHCP can be secured by using DHCP snooping, which filters out unauthorized DHCP messages and prevents rogue DHCP servers from assigning IP addresses.

? Option F: LDAP (Lightweight Directory Access Protocol). LDAP is a service that stores and organizes information about users, devices, and resources on a network. LDAP can provide identity management and access control for the cloud environment. LDAP can be secured by using LDAPS (LDAP over SSL/TLS), which encrypts the LDAP traffic and provides authentication and integrity checks.

NEW QUESTION 227

- (Topic 4)

A technician deployed a VM with NL-SAS storage to host a critical application. Two weeks later, users have begun to report high application latency. Which of the following is the best action to correct the latency issue?

- A. Increase the capacity of the data storage.
- B. Migrate the data to SAS storage.
- C. Increase the CPU of the VM.
- D. Migrate the data to flash storage.

Answer: D

Explanation:

Flash memory



Explore

One possible answer is:

D. Migrate the data to flash storage.

Flash storage is a type of solid-state storage technology that uses flash memory chips to store data. Flash storage has several advantages over NL-SAS storage, which is a hybrid of SATA and SAS technologies that uses spinning disks to store data. Flash storage can provide much faster performance, lower latency, higher reliability, and lower power consumption than NL-SAS storage¹². Therefore, migrating the data to flash storage can help correct the latency issue for the critical application. However, flash storage may also be more expensive and have lower capacity than NL-SAS storage, so these factors should also be considered before making the migration decision¹².

NEW QUESTION 229

- (Topic 4)

Which of the following provides groups of compute units that can horizontally scale according to a workload?

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

Answer: C

Explanation:

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

References: [CompTIA Cloud+ Study Guide], page 75.

NEW QUESTION 234

- (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 239

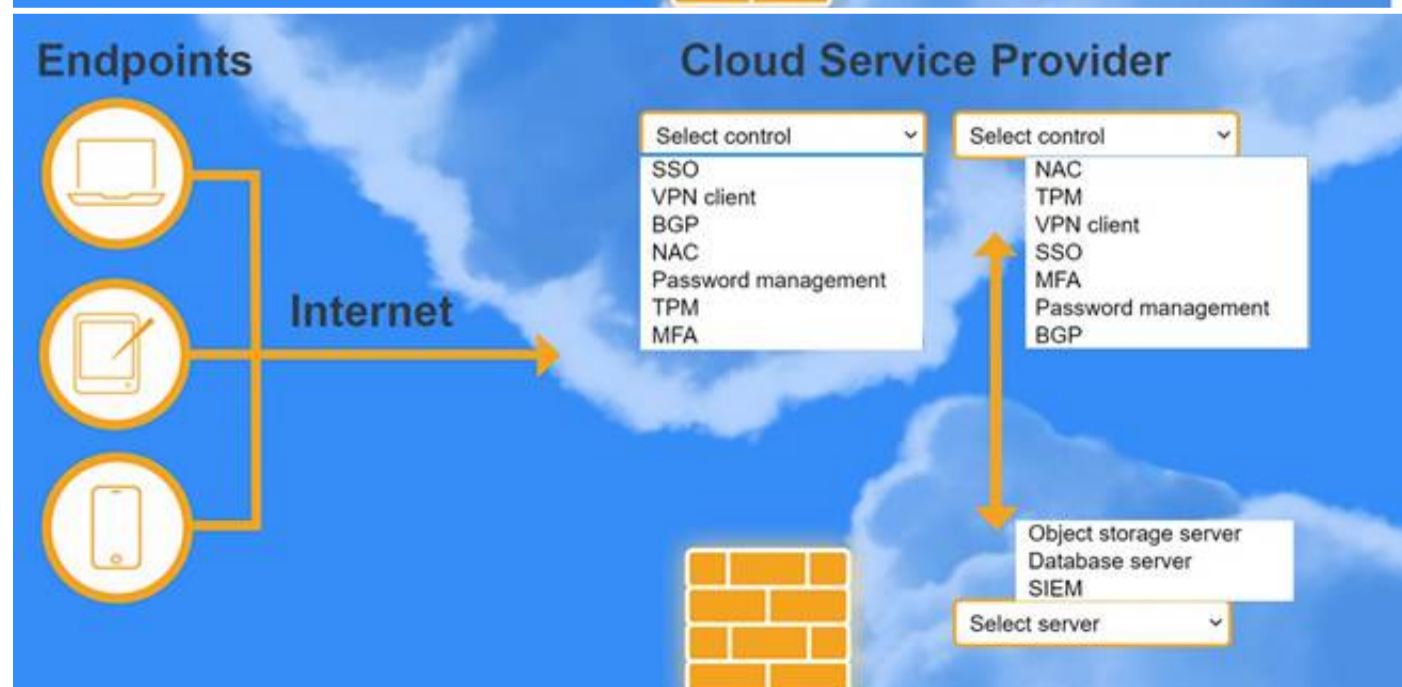
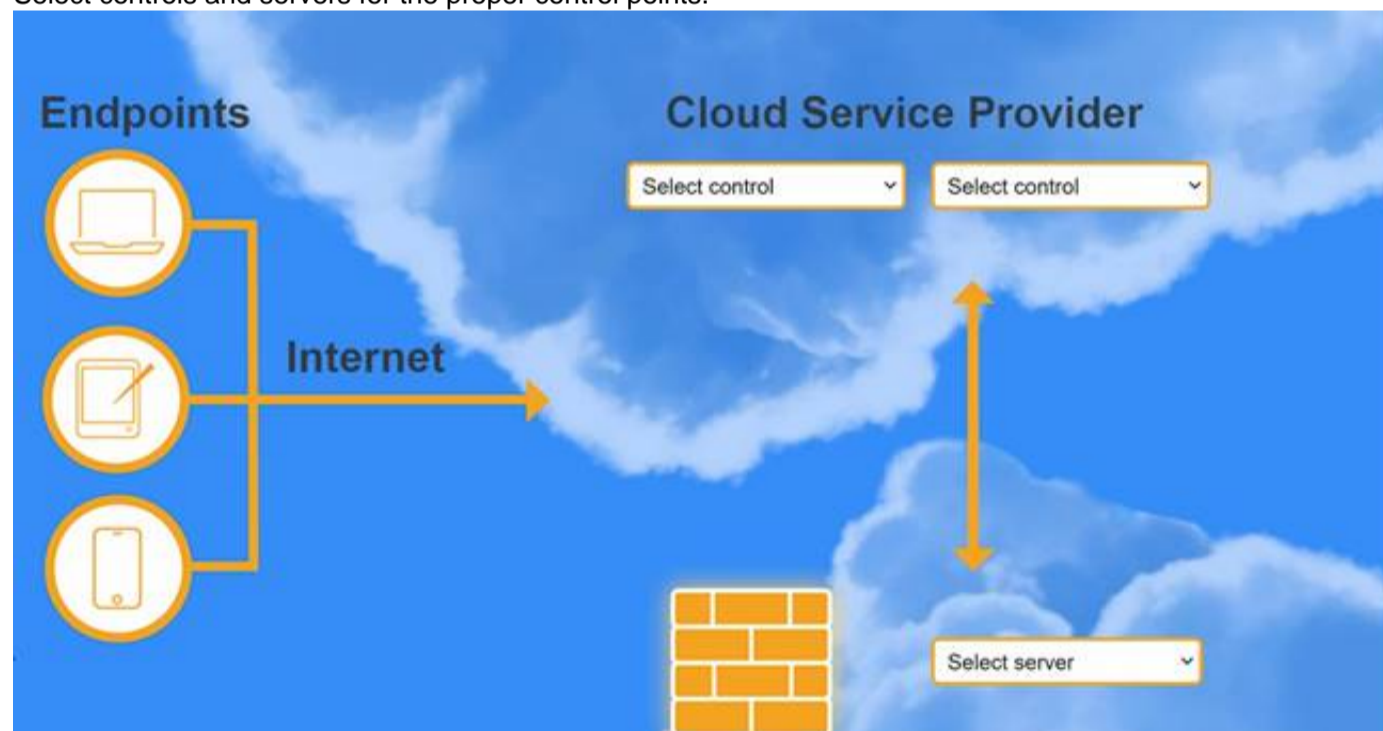
HOTSPOT - (Topic 4)

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

- ? secure connectivity that minimizes user login
- ? tracks user activity and monitors for anomalous activity
- ? requires secondary authentication

INSTRUCTIONS

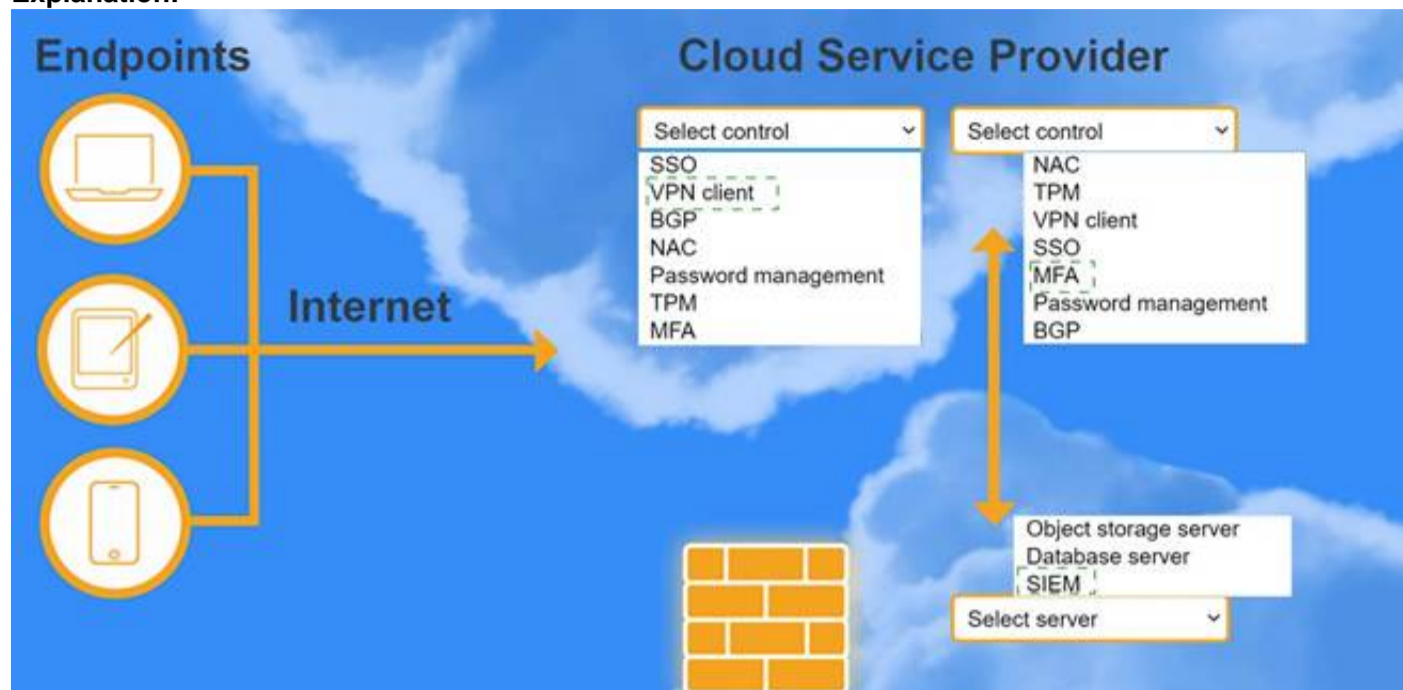
Select controls and servers for the proper control points.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 240

- (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:

One of the main benefits of cloud computing is that you only pay for the resources that you use. However, this also means that you need to manage your cloud resources efficiently and avoid paying for idle or unused resources¹.

Shutting down the environment after work hours is a process that can be automated to best reduce cost in a cloud environment that must only be accessed during work hours. This process involves stopping or terminating the cloud resources, such as virtual machines, databases, load balancers, etc., that are not needed outside of the work hours. This can significantly reduce the cloud bill by avoiding charges for compute, storage, network, and other services that are not in use².

The other options are not the best processes to automate to reduce cost in this scenario:

? Option A: Scaling of the environment after work hours. Scaling is a process that involves adjusting the number or size of cloud resources to match the demand or workload. Scaling can be done manually or automatically using triggers or policies. Scaling can help optimize the performance and availability of a cloud environment, but it does not necessarily reduce the cost. Scaling down the environment after work hours may reduce some costs, but it may still incur charges for the remaining resources. Scaling up the environment before work hours may increase the cost and also introduce delays or errors in provisioning new resources³.

? Option B: Implementing access control after work hours. Access control is a process that involves defining and enforcing rules and policies for who can access what resources in a cloud environment. Access control can help improve the security and compliance of a cloud environment, but it does not directly affect the cost. Implementing access control after work hours may prevent unauthorized access to the environment, but it does not stop or terminate the resources that are still running and consuming cloud services⁴.

? Option D: Blocking external access to the environment after work hours. Blocking external access is a process that involves restricting or denying network traffic from outside sources to a cloud environment. Blocking external access can help protect the environment from potential attacks or breaches, but it does not impact the cost. Blocking external access after work hours may prevent unwanted requests or connections to the environment, but it does not shut down or release the resources that are still active and generating cloud charges.

NEW QUESTION 245

- (Topic 4)

A cloud architect is reviewing the design for a new cloud-based ERP solution. The solution consists of eight servers with a single network interface. The allocated IP range is 172.16.0.0/28. One of the requirements of the solution is that it must be able to handle the potential addition of 16 new servers to the environment. Because of the complexity of the firewall and related ACL requirements, these new servers will need to be in the same network range. Which of the following changes would allow for the potential server addition?

- A. Change the IP address range to use a 10.0.0.0 address.
- B. Change the server template to add network interfaces.
- C. Change the subnet mask to use a 255.255.255.128 range.
- D. Change the server scaling configuration to increase the maximum limit.

Answer: C

Explanation:

Changing the subnet mask to use a 255.255.255.128 range would allow for the potential server addition. The current subnet mask of 255.255.255.240 (/28) only allows for 14 usable host addresses in the 172.16.0.0 network, which is not enough to accommodate the existing eight servers and the possible 16 new servers. Changing the subnet mask to 255.255.255.128 (/25) would increase the number of usable host addresses to 126 in the same network, which is sufficient to handle the server expansion. Changing the IP address range to use a 10.0.0.0 address, changing the server template to add network interfaces, or changing the server scaling configuration to increase the maximum limit would not solve the issue of the limited host addresses in the same network range. References: CompTIA

Cloud+ CV0-003 Certification Study Guide, Chapter 3, Objective 3.1: Given a scenario, implement cloud networking solutions.

NEW QUESTION 246

- (Topic 4)

A systems administrator has a redundant backup system in place. Which of the following should the systems administrator perform to maintain efficient operation and comply with the global standard in the corporate backup policies?

- A. Modify RTO policies.
- B. Confirm completion of the backups.
- C. Test the backups.
- D. Modify RPO policies.

Answer: C

NEW QUESTION 247

- (Topic 4)

A VDI administrator is deploying 512 desktops for remote workers. Which of the following would meet the minimum number of IP addresses needed for the desktops?

- A. /22
- B. /23
- C. /24
- D. /25

Answer: B

Explanation:

A /23 subnet mask has 9 bits for the host portion, which allows up to 512 IP addresses for the desktops. A /22 subnet mask has 10 bits for the host portion, which allows up to 1024 IP addresses, but this is more than the minimum required. A /24 subnet mask has 8 bits for the host portion, which allows up to 256 IP addresses, but this is not enough for the desktops. A /25 subnet mask has 7 bits for the host portion, which allows up to 128 IP addresses, but this is also not enough for the desktops. References: CompTIA Cloud+ Certification Exam Objectives, Domain 1.0: Cloud Concepts, Objective 1.2: Given a scenario, analyze and compare the characteristics of various cloud service models (SaaS, IaaS, PaaS). Subnet Mask Cheat Sheet - aelius.com

NEW QUESTION 248

- (Topic 4)

A VDI provider suspects users are installing prohibited software on the instances. Which of the following must be implemented to prevent the issue?

- A. Log monitoring
- B. Patch management
- C. Vulnerability scanning
- D. System hardening

Answer: D

Explanation:

System hardening is the process of securing a system by reducing its attack surface and eliminating unnecessary services, features, or functions. System hardening can help prevent users from installing prohibited software on the VDI instances by applying policies and restrictions that limit the user privileges and access rights. For example, system hardening can disable the installation of software from unknown sources, enforce the use of strong passwords, enable encryption, and remove default accounts. System hardening can also improve the performance and stability of the VDI instances by removing unwanted or unused components. References: [CompTIA Cloud+ CV0-003 Certification Study Guide], Chapter 9, Objective 9.1: Given a scenario, apply security controls and techniques.

NEW QUESTION 249

- (Topic 4)

A cloud administrator deployed new hosts in a private cloud. After a few months elapsed, some of the hypervisor features did not seem to be working. Which of the following was MOST likely causing the issue?

- A. Incorrect permissions
- B. Missing license
- C. Incorrect tags
- D. Oversubscription

Answer: B

Explanation:

The correct answer is B. Missing license.

Some hypervisor features may require a valid license to work properly. If the license is missing, expired, or invalid, the hypervisor may not be able to use those features or may operate in a reduced functionality mode. For example, some features of Hyper-V, such as live migration, replication, and failover clustering, require a license for Windows Server or Windows 10 Enterprise¹. Similarly, some features of VMware ESXi, such as vMotion, Storage vMotion, and Fault Tolerance, require a license for VMware vSphere². Therefore, if a cloud administrator deployed new hosts in a private cloud and found that some of the hypervisor features did not seem to be working after a few months elapsed, the most likely cause was a missing license. The administrator should check the license status of the hypervisor and renew or activate the license if needed.

Incorrect permissions are not a likely cause of the issue, as they would affect the access to the hypervisor or its resources, not the functionality of the hypervisor itself. Incorrect tags are also not a likely cause of the issue, as they are used for identification and classification of resources, not for enabling or disabling features. Oversubscription is not a likely cause of the issue either, as it would affect the performance or availability of the resources, not the functionality of the hypervisor itself.

NEW QUESTION 251

- (Topic 4)

A new development team requires workstations hosted in a PaaS to develop a new website. Members of the team also require remote access to the workstations using their corporate email addresses. Which of the following solutions will BEST meet these requirements? (Select TWO).

- A. Deploy new virtual machines.
- B. Configure email account replication.
- C. Integrate identity services.
- D. Implement a VDI solution.
- E. Migrate local VHD workstations.
- F. Create a new directory service.

Answer: AC

Explanation:

A Platform-as-a-Service (PaaS) is a cloud computing model that provides customers a complete cloud platform—hardware, software, and infrastructure—for developing, running, and managing applications without the cost, complexity, and inflexibility that often comes with building and maintaining that platform on-premises¹.

To develop a new website using a PaaS, the development team needs to deploy new virtual machines (VMs) on the cloud platform. VMs are software emulations of physical computers that can run different operating systems and applications. By deploying new VMs, the development team can create a scalable and flexible environment for their website project, without having to invest in or manage physical hardware².

To enable remote access to the workstations using their corporate email addresses, the development team needs to integrate identity services on the cloud platform. Identity services are services that provide authentication, authorization, and identity management for users and devices accessing cloud resources. By integrating identity services, the development team can use their corporate email addresses as single sign-on (SSO) credentials to access their workstations from any device and location, while ensuring security and compliance³.

The other options are not the best solutions for these requirements:

? Configuring email account replication is not necessary for remote access to the workstations. Email account replication is a process of synchronizing email accounts across different servers or locations. It can provide backup and redundancy for email services, but it does not provide authentication or identity management for remote access⁴.

? Implementing a Virtual Desktop Infrastructure (VDI) solution is not a PaaS solution.

VDI is a technology that allows users to access virtual desktops hosted on a centralized server. VDI can provide remote access to desktop environments, but it requires additional hardware, software, and management costs that are not included in a PaaS model⁵.

? Migrating local VHD workstations is not a PaaS solution. VHD stands for Virtual Hard Disk, which is a file format that represents a virtual hard disk drive.

Migrating local VHD workstations means moving the virtual hard disk files from local storage to cloud storage. This can provide backup and portability for the workstations, but it does not provide a complete cloud platform for developing and running applications⁶.

? Creating a new directory service is not necessary for remote access to the workstations. A directory service is a service that stores and organizes information about users, devices, and resources on a network. Creating a new directory service means setting up a new database and schema for storing this information. This can provide identity management and access control for the network, but it does not provide authentication or SSO for remote access.

NEW QUESTION 255

- (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 260

- (Topic 4)

A new development team requires workstations hosted in a PaaS to develop a new website. Members of the team also require remote access to the workstations using their corporate email addresses. Which of the following solutions will best meet these requirements? (Select two).

- A. Deploy new virtual machines.
- B. Configure email account replication.
- C. Integrate identity services.
- D. Implement a VDI solution.
- E. Migrate local VHD workstations.
- F. Create a new directory service.

Answer: CD

Explanation:

To meet the requirements of the development team, the cloud administrator should integrate identity services and implement a VDI solution. Identity services are used to authenticate and authorize users and devices to access cloud resources. By integrating identity services, the cloud administrator can enable the development team to use their corporate email addresses to log in to the PaaS workstations. A VDI solution is a virtualization technology that allows users to access remote desktops hosted on a cloud platform. By implementing a VDI solution, the cloud administrator can provide the development team with workstations that have the necessary tools and configurations for web development. References: [CompTIA Cloud+ CV0-003 Certification Study Guide], Chapter 2, Objective 2.1: Given a scenario, deploy cloud services and solutions.

NEW QUESTION 262

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

CV0-003 Practice Exam Features:

- * CV0-003 Questions and Answers Updated Frequently
- * CV0-003 Practice Questions Verified by Expert Senior Certified Staff
- * CV0-003 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * CV0-003 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The CV0-003 Practice Test Here](#)