CompTIA

Exam Questions CV0-003

CompTIA Cloud+ Certification Exam

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1
- (Topic 1)
An organization is running a database application on a SATA disk, and a customer is experiencing slow performance most of the time.
Which of the following should be implemented to improve application performance?

A. Increase disk capacity
B. Increase the memory and network bandwidth
C. Upgrade the application
D. Upgrade the environment and use SSD drives

Answer: D

Explanation:
Upgrading the environment and using solid state drives (SSDs) can improve application performance for a database application that is running on a serial advanced technology attachment (SATA) disk and experiencing slow performance most of the time. Upgrading the environment can involve updating or replacing the hardware, software, or network components that support the application to enhance their functionality, capacity, or compatibility. Using SSDs can provide faster and more reliable data access and storage than SATA disks, as they use flash memory instead of spinning disks to store data. SSDs can also reduce latency, power consumption, and heat generation. References: CompTIA Cloud+ Certification Exam Objectives, page 9, section 1.4

NEW QUESTION 2
- (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 3
- (Topic 1)
An administrator is performing an in-place upgrade on a quest VM operating system.
Which of the following can be performed as a quick method to roll back to an earlier state, if necessary?

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

Answer: D

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

NEW QUESTION 4
- (Topic 1)
A company wants to implement business continuity, and the cloud solution architect needs to design the correct solution.
Which of the following will provide the data to measure business continuity? (Choose two.)

A. A service-level agreement
B. Automation scripts
C. Playbooks
D. A network diagram
E. A backup and restore
F. A recovery time objective

Answer: AF

Explanation:
A service-level agreement (SLA) is a contract or document that defines the level of service and performance expected from a service provider or vendor. A recovery time objective (RTO) is a metric that specifies the maximum acceptable time for restoring a system or service after a disruption or outage. Both SLA and RTO can provide the data to measure business continuity, as they can indicate the availability, reliability, and recoverability of a system or service in case of a failure or disaster. SLA and RTO can also help evaluate the effectiveness and efficiency of the business continuity plan and solution. References: CompTIA Cloud+ Certification Exam Objectives, page 20, section 4.2

NEW QUESTION 5
- (Topic 1)
A systems administrator needs to convert ten physical servers to virtual.
Which of the following would be the MOST efficient conversion method for the administrator to use?

A. Rebuild the servers from scratch
B. Use the vendor's conversion tool
C. Clone the hard drive
D. Restore from backup

Answer: B

Explanation:
A vendor's conversion tool is a type of software or utility that automates and simplifies the process of converting physical servers to virtual machines by capturing the configuration and data of the physical servers and creating virtual disks and files for the virtual machines. Using the vendor's conversion tool can be the most efficient conversion method for a systems administrator to use to convert ten physical servers to virtual, as it can save time and effort by avoiding manual steps or errors involved in rebuilding, cloning, or restoring the physical servers to virtual machines. Using the vendor's conversion tool can also ensure compatibility and consistency, as it can match the hardware and software requirements and settings of the physical servers to the virtual machines.
References: CompTIA Cloud+ Certification Exam Objectives, page 10, section 1.5

NEW QUESTION 6
- (Topic 1)
A systems administrator is configuring RAID for a new server. This server will host files for users and replicate to an identical server. While redundancy is necessary, the most important need is to maximize storage.
Which of the following RAID types should the administrator choose?

A. 5
B. 6
C. 10
D. 50

Answer: C

Explanation:
RAID 50 is a type of RAID level that combines RAID 5 and RAID 0 to create a nested RAID configuration. RAID 50 consists of two or more RAID 5 arrays that are striped together using RAID 0. RAID 50 can provide redundancy, fault tolerance, and high performance for large data sets. RAID 50 can also maximize storage, as it has a higher usable capacity than other RAID levels with similar features, such as RAID 6 or RAID 10. The administrator should choose RAID 50 to configure a new server that will host files for users and replicate to an identical server, as it can meet the needs of redundancy and storage maximization. References: CompTIA Cloud+ Certification Exam Objectives, page 9, section 1.4

NEW QUESTION 7
- (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 8
- (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 9

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

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

Answer: AC

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

NEW QUESTION 10
- (Topic 1)
A systems administrator is provisioning VMs in a cloud environment and has been told to select an OS build with the furthest end-of-life date.
Which of the following OS builds would be BEST for the systems administrator to use?

A. Open-source
B. LTS
C. Canary
D. Beta
E. Stable

Answer: B

Explanation:
Long-term support (LTS) is a type of release cycle that provides extended support and maintenance for software products or operating systems. LTS releases typically have longer end-of-life dates than regular releases, as they receive security updates, bug fixes, and patches for several years after their initial release date. LTS releases can also offer higher stability, reliability, and compatibility than regular releases, as they undergo more testing and quality assurance processes before being released. LTS is the best OS build for a systems administrator to use when provisioning VMs in a cloud environment and being told to select an OS build with the furthest end-of-life date. References: CompTIA Cloud+ Certification Exam Objectives, page 11, section 1.6

NEW QUESTION 10
- (Topic 2)
A systems administrator is deploying a new cloud application and needs to provision cloud services with minimal effort. The administrator wants to reduce the tasks required for maintenance, such as OS patching, VM and volume provisioning, and autoscaling configurations. Which of the following would be the BEST option to deploy the new application?

A. A VM cluster
B. Containers
C. OS templates
D. Serverless

Answer: D

Explanation:
Serverless is what would be the best option to deploy a new cloud application and provision cloud services with minimal effort while reducing the tasks required for maintenance such as OS patching, VM and volume provisioning, and autoscaling configurations. Serverless is a cloud service model that provides customers with a platform to run applications or functions without having to manage or provision any underlying infrastructure or resources, such as servers, storage, network, OS, etc. Serverless can provide benefits such as:
? Minimal effort: Serverless can reduce the effort required to deploy a new cloud application and provision cloud services by automating and abstracting away all the infrastructure or resource management or provisioning tasks from customers, and allowing them to focus only on writing code or logic for their applications or functions.
? Reduced maintenance: Serverless can reduce the tasks required for maintenance by handling all the infrastructure or resource maintenance tasks for customers, such as OS patching, VM and volume provisioning, autoscaling configurations, etc., and ensuring that they are always up-to-date and optimized.

NEW QUESTION 14
- (Topic 2)
Which of the following should be considered for capacity planning?

A. Requirements, licensing, and trend analysis
B. Laws and regulations
C. Regions, clusters, and containers
D. Hypervisors and scalability

Answer: A

Explanation:
These are the factors that should be considered for capacity planning in a cloud environment. Capacity planning is a process of estimating and allocating the necessary resources and performance to meet the current and future demands of cloud applications or services. Capacity planning can help to optimize costs, efficiency, and reliability of cloud resources or services. The factors that should be considered for capacity planning are:
? Requirements: These are the specifications or expectations of the cloud applications or services, such as functionality, availability, scalability, security, etc.

Requirements can help to determine the type, amount, and quality of resources or services needed to meet the objectives and goals of the cloud applications or services.
? Licensing: This is the agreement or contract that grants customers the right to use or access certain cloud resources or services for a specific period or fee. Licensing can affect the cost, availability, and compliance of cloud resources or services. Licensing can help to determine the budget, duration, and scope of using or accessing cloud resources or services.
? Trend analysis: This is the technique of analyzing historical and current data to identify patterns, changes, or fluctuations in demand or usage of cloud resources or services. Trend analysis can help to predict and anticipate future demand or usage of cloud resources or services, as well as identify any opportunities or challenges that may arise.

NEW QUESTION 17
- (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 21
- (Topic 2)
A company is concerned about the security of its data repository that contains customer PII. A systems administrator is asked to deploy a security control that will prevent the exfiltration of such data. Which of the following should the systems administrator implement?

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

Answer: A

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

NEW QUESTION 26
- (Topic 2)
Which of the following definitions of serverless computing BEST explains how it is different from using VMs?

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

Answer: D

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

NEW QUESTION 28
- (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 31
- (Topic 2)
A company is planning to migrate applications to a public cloud, and the Chief Information Officer (CIO) would like to know the cost per business unit for the applications in the cloud. Before the migration, which of the following should the administrator implement FIRST to assist with reporting the cost for each business unit?

A. An SLA report
B. Tagging
C. Quotas
D. Showback

Answer: B

Explanation:
Tagging is what the administrator should implement first to assist with reporting the cost for each business unit for applications in a public cloud environment. Tagging is a technique that allows customers to assign metadata or labels to their cloud resources, such as applications, instances, volumes, etc., based on their attributes or criteria. Tagging can help customers to organize, manage, monitor, and report their cloud resources and costs by business unit, project, owner, environment, etc.

NEW QUESTION 34
- (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 38
- (Topic 2)
A technician is trying to delete six decommissioned VMs. Four VMs were deleted without issue. However, two of the VMs cannot be deleted due to an error. Which of the following would MOST likely enable the technician to delete the VMs?

A. Remove the snapshots
B. Remove the VMs' IP addresses
C. Remove the VMs from the resource group
D. Remove the lock from the two VMs

Answer: D

Explanation:
Removing the lock from the two VMs is what would most likely enable the technician to delete the VMs that cannot be deleted due to an error. A lock is a feature that prevents certain actions or operations from being performed on a resource or service, such as deleting, modifying, moving, etc. A lock can help to protect a resource or service from accidental or unwanted changes or removals. Removing the lock from the two VMs can enable the technician to delete them by allowing the delete action or operation to be performed on them.

NEW QUESTION 42
- (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 46
- (Topic 2)
A cloud administrator is upgrading a cloud environment and needs to update the automation script to use a new feature from the cloud provider. After executing the script, the deployment fails. Which of the following is the MOST likely cause?

A. API incompatibility
B. Location changes
C. Account permissions
D. Network failure

Answer: A

Explanation:
API incompatibility is the most likely cause of the failure of an automation script to use a new feature from the cloud provider. API (Application Programming Interface) is a set of rules or specifications that defines how different software components or systems can communicate and interact with each other. API incompatibility is a situation where an API does not work or function properly with another software component or system due to differences or changes in versions, formats, parameters, etc. API incompatibility can cause errors or issues when using an automation script to deploy or configure cloud resources or services, especially if the script is not updated or modified according to the new API specifications.

NEW QUESTION 50
- (Topic 2)
A systems administrator wants to verify the word "qwerty" has not been used as a password on any of the administrative web consoles in a network. Which of the following will achieve this goal?

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

Answer: C

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

NEW QUESTION 55
- (Topic 2)
A resource pool in a cloud tenant has 90 GB of memory and 120 cores. The cloud administrator needs to maintain a 30% buffer for resources for optimal performance of the hypervisor. Which of the following would all ow for the maximum number of two-core machines with equal memory?

A. 30 VMs, 3GB of memory
B. 40 VMs, 1,5GB of memory
C. 45 VMs, 2 GB of memory
D. 60 VMs, 1 GB of memory

Answer: C

Explanation:
To calculate the maximum number of two-core machines with equal memory, we need to consider the resource pool capacity and the buffer requirement. The resource pool has 90 GB of memory and 120 cores, but the cloud administrator needs to maintain a 30% buffer for optimal performance. This means that only 70% of the resources can be used for VM allocation. Therefore, the available memory is 90 GB x 0.7 = 63 GB, and the available cores are 120 x 0.7 = 84 cores. To allocate two-core machines with equal memory, we need to divide the available memory by the available cores and multiply by two. This gives us the memory size per VM: (63 GB / 84 cores) x 2 = 1.5 GB. However, this is not a valid answer option, so we need to find the closest option that does not exceed the available resources. The best option is C, which allocates 45 VMs with 2 GB of memory each. This uses up 45 x 2 = 90 GB of memory and 45 x 2 = 90 cores, which are within the available limits.

NEW QUESTION 56
- (Topic 2)
A company needs a solution to find content in images. Which of the following technologies, when used in conjunction with cloud services, would facilitate the BEST

solution?

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

Answer: C

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

NEW QUESTION 60
- (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 61
- (Topic 2)
A systems administrator is deploying a solution that includes multiple network I/O-intensive VMs. The solution design requires that vNICs of the VMs provide low-latency, near-native performance of a physical NIC and data protection between the VMs. Which of the following would BEST satisfy these requirements?

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

Answer: A

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

NEW QUESTION 66
- (Topic 2)
A systems administrator is analyzing a report of slow performance in a cloud application. This application is working behind a network load balancer with two VMs, and each VM has its own digital certificate configured. Currently, each VM is consuming 85% CPU on average. Due to cost restrictions, the administrator cannot scale vertically or horizontally in the environment. Which of the following actions should the administrator take to decrease the CPU utilization? (Choose two.)

A. Configure the communication between the load balancer and the VMs to use a VPN.
B. Move the digital certificate to the load balancer.
C. Configure the communication between the load balancer and the VMs to use HTTP.
D. Reissue digital certificates on the VMs.
E. Configure the communication between the load balancer and the VMs to use HTTPS.
F. Keep the digital certificates on the VMs.

Answer: BC

Explanation:
Moving the digital certificate to the load balancer and configuring the communication between the load balancer and the VMs to use HTTP are two actions that will decrease the CPU utilization of the VMs that are running behind a network load balancer with two VMs, each with its own digital certificate configured. Moving the digital certificate to the load balancer will offload the SSL/TLS encryption and decryption tasks from the VMs to the load balancer, which can reduce the CPU overhead and improve performance. Configuring the communication between the load balancer and the VMs to use HTTP will eliminate the need for encryption and decryption between them, which can also reduce CPU consumption. However, this may introduce security risks if sensitive data is transmitted over HTTP.

NEW QUESTION 68
- (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 73
- (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 76
- (Topic 2)
A cloud engineer is responsible for managing a public cloud environment. There is currently one virtual network that is used to host the servers in the cloud environment. The environment is rapidly growing, and the network does not have any more available IP addresses. Which of the following should the engineer do to accommodate additional servers in this environment?

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

Answer: A

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

NEW QUESTION 81
- (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 82

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

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

Answer: D

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

NEW QUESTION 84
- (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 87
- (Topic 1)
An IaaS application has a two-hour RTO and a four-hour RPO. The application takes one hour to back up its data or restore from a local backup file. A systems administrator is tasked with configuring the backup policy.
Which of the following should the administrator configure to achieve the application requirements with the LEAST cost?

A. Back up to long-term storage every night
B. Back up to object storage every three hours
C. Back up to long-term storage every four hours
D. Back up to object storage every hour

Answer: B

Explanation:
Object storage is a type of storage service that stores data as objects with unique identifiers and metadata in a flat namespace or structure. Backing up to object storage every three hours can help achieve the application requirements with the least cost for an IaaS application that has a two-hour RTO and a four-hour RPO, as it can provide scalable, durable, and cost-effective storage for backup data while meeting the recovery time and point objectives. Backing up to object storage every three hours can ensure that the backup data is no more than four hours old and can be restored within two hours in case of a disaster or failure. References: CompTIA Cloud+ Certification Exam Objectives, page 9, section 1.4

NEW QUESTION 89
- (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 94
- (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 99
- (Topic 1)
An organization has multiple VLANs configured to segregate the network traffic. Following is the breakdown of the network segmentation:
? Production traffic (10.10.0.0/24)
? Network backup (10.20.0.0/25)
? Virtual IP network (10.20.0.128/25)
The following configuration exists on the server:

Server name	Interface	IP address	Gateway
COMPSRV01	Production	10.10.0.12/24	10.10.0.1
COMPSRV01	Network backup	10.20.0.12/25	10.10.0.1

The backup administrator observes that the weekly backup is failing for this server. Which of the following commands should the administrator run to identify the issue?

A. ROUTE PRINT
B. NETSTAT -A
C. IPCONFIG /ALL
D. NET SM

Answer: A

Explanation:
ROUTE PRINT is a command that displays the routing table of a system, which shows the destination network, the gateway, the interface, and the metric for each route. ROUTE PRINT can help identify the issue of the weekly backup failing for this server, as it can show if there is a valid route to the network backup segment (10.20.0.0/25) from the production traffic segment (10.10.0.0/24). If there is no route or an incorrect route, the backup will fail to reach the destination. The administrator can use ROUTE PRINT to verify and troubleshoot the routing configuration of the server. References: CompTIA Cloud+ Certification Exam Objectives, page 16, section 3.2
Reference: https://www.toolbox.com/tech/operating-systems/blogs/using-the-route-print- command-in- windows-7-022310/

NEW QUESTION 104
- (Topic 1)
A cloud administrator is planning to migrate a globally accessed application to the cloud.
Which of the following should the cloud administrator implement to BEST reduce latency for all users?

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

Answer: A

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

NEW QUESTION 108
- (Topic 4)

A security analyst is investigating a recurring alert. The alert is reporting an insecure firewall configuration state after every cloud application deployment. The process of identifying the issue, requesting a fix, and waiting for the developers to manually patch the environment is being repeated multiple times. In an effort to identify the root issue, the following logs were collected:
Deploying template app prod. •yaml Instance DB successfully created DB keys successfully stored on vault
Instance WebApp successfully created Access rules successfully applied Access—keys successfully created
Which of the following options will provide a permanent fix for the issue?

A. Validate the Iac code used during the deployment.
B. Avoid the use of a vault to store database passwords.
C. Rotate the access keys that were created during deployment.
D. Recommend that the developers do not create multiple resources at once.

Answer: A

Explanation:
The issue of an insecure firewall configuration state after every cloud application deployment is likely caused by a flaw in the IaC code used during the deployment. IaC stands for Infrastructure as Code, which is a method of managing and provisioning IT infrastructure using code, rather than manual configuration1. IaC allows teams to automate the setup and management of their infrastructure, making it more efficient and consistent. However, if the IaC code contains errors, vulnerabilities, or misconfigurations, it can result in security issues or compliance violations in the deployed infrastructure2. Therefore, to provide a permanent fix for the issue, the IaC code used during the deployment should be validated and tested to ensure that it meets the security requirements and best practices for firewall configuration. The IaC code can be validated using tools such as Azure Resource Manager Template Toolkit, AWS CloudFormation Linter, or Terraform Validate. These tools can check the syntax and semantics of the IaC code, and identify any potential errors or inconsistencies before deployment

NEW QUESTION 111
- (Topic 4)
A systems administrator is attempting to gather information about services and resource utilization on VMs in a cloud environment. Which of the following will best accomplish this objective?

A. Syslog
B. SNMP
C. CMDB
D. Service management
E. Performance monitoring

Answer: E

Explanation:
Performance monitoring is a technique that collects and analyzes data about the services and resource utilization on VMs in a cloud environment. Performance monitoring can help the systems administrator to gather information about the CPU, memory, disk, network, and application performance of the VMs, as well as identify any bottlenecks, errors, or anomalies that may affect the cloud service quality. Performance monitoring can be implemented using various tools or agents that can collect and report the performance metrics from the VMs to a centralized dashboard or console. Performance monitoring can also help the systems administrator to optimize, troubleshoot, and plan the cloud resources and services. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 5, Objective 5.1: Given a scenario, monitor cloud resources and services.

NEW QUESTION 113
- (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 118
- (Topic 4)
A cloud administrator created a developer desktop image and added it to the VDI farm in a private cloud environment. One of the developers opened a VDI session and noticed that compiling the code was taking up to one hour to complete. However, when the developer compiles the code on a local machine, the job completes in less than five minutes. Which of the following sizing techniques would be best to use to improve the performance of the compile job?

A. Add more servers to the VDI environment.

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

Answer: B

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

NEW QUESTION 122
- (Topic 4)
Which of the following enables CSPs to offer unlimited capacity to customers?

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

Answer: C

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

NEW QUESTION 123
- (Topic 4)
A systems administrator is reviewing the logs from a company's IDS and notices a large amount of outgoing traffic from a particular server. The administrator then runs a scan on the server, which detects malware that cannot be removed. Which of the following should the administrator do first?

A. Determine the root cause.
B. Disconnect the server from the network.
C. Perform a more intrusive scan.
D. Restore the server from a backup.

Answer: B

Explanation:
The first step in any incident response procedure is to contain the incident and prevent it from spreading or causing more damage. In this scenario, the systems administrator is reviewing the logs from a company’s IDS and notices a large amount of outgoing traffic from a particular server. The administrator then runs a scan on the server, which detects malware that cannot be removed. This indicates that the server is compromised and may be sending malicious or sensitive data to an external source. Therefore, the best thing to do first is to disconnect the server from the network, which will isolate it from the rest of the system and stop the data exfiltration. Determining the root cause, performing a more intrusive scan, and restoring the server from a backup are all important steps, but they should be done after the server is disconnected from the network. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 10, Incident Response Procedures, page 1771.

NEW QUESTION 125
- (Topic 4)
A cloud engineer is migrating a customer's web servers from a hypervisor platform to a CSP environment. The engineer needs to decouple the infrastructure and components during the migration to reduce the single points of failure. Which of the following storage options should the cloud engineer migrate the content to in order to improve availability?

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

Answer: C

Explanation:
Object storage is a storage option that stores data as discrete units called objects, which are identified by a unique identifier and can have metadata attached to them. Object storage can help the cloud engineer migrate the content to improve availability by decoupling the data from the underlying infrastructure and components. Object storage can also provide high scalability, durability, and redundancy for the data, as well as support for multiple protocols and access methods. Object storage can be accessed through APIs, web interfaces, or gateways that can emulate file or block storage. Object storage is suitable for storing unstructured or static data, such as web content, images, videos, or documents. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 4,

Objective 4.1: Given a scenario, implement cloud storage solutions.

NEW QUESTION 129
- (Topic 4)
A company is using laaS 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 131
- (Topic 4)
A cloud administrator who is troubleshooting DNS issues discovers zone transfers are not occurring between the primary and secondary name servers due to an error in the serial numbers. Which of the following records should the administrator query for the serial number?

A. PTR
B. TXT
C. SOA
D. SRV

Answer: C

Explanation:
SOA stands for Start of Authority, and it is a type of DNS record that contains information about a DNS zone, such as the name of the primary name server, the email address of the zone administrator, the serial number of the zone, and other parameters. The serial number is used to indicate when a zone has been updated, and it is incremented by the primary name server whenever a change is made to the zone data. The secondary name servers use the serial number to determine if they need to request a zone transfer from the primary name server to synchronize their data.
References: [CompTIA Cloud+ Study Guide], page 207.

NEW QUESTION 136
- (Topic 4)
A systems administrator needs to connect the companys network to a public cloud services provider. Which of the following will BEST ensure encryption in transit for data transfers?

A. Identity federation
B. A VPN tunnel
C. A proxy solution
D. A web application firewall

Answer: B

Explanation:
The answer is A. SAML. SAML (Security Assertion Markup Language) is a standard for exchanging authentication and authorization data between different parties, such as a user and a service provider. In a federated cluster, SAML can be used to enable single sign-on (SSO) for users across multiple clusters or cloud providers. SAML relies on the exchange of XML-based assertions that contain information about the user's identity, attributes, and entitlements. If the users' API access tokens have become invalid, it could be because the SAML assertions have expired, been revoked, or corrupted. The administrator should check the SAML configuration and logs to determine the cause of this issue.
Some possible sources of information about SAML and federated clusters are:
? Authenticating | Kubernetes: This page provides an overview of authenticating users in Kubernetes, including using SAML for federated identity.
? Authenticating to the Kubernetes API server - Google Cloud: This page explains how to authenticate to the Kubernetes API server on Google Cloud, including using SAML for federated identity with Google Cloud Identity Platform.
? Error 403 User not authorized when trying to access Azure Databricks API through Active Directory - Stack Overflow: This page discusses a similar issue of users getting an error when trying to access Azure Databricks API using SAML and Active Directory.

NEW QUESTION 139
- (Topic 4)
A cloud administrator is having difficulty correlating logs for multiple servers. Upon inspection, the administrator finds that the time-zone settings are mismatched throughout the deployment. Which of the following solutions can help maintain time synchronization between all the resources?

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

Answer: C

Explanation:

The correct answer is C. NTP.
NTP stands for Network Time Protocol, which is a standard protocol for synchronizing the clocks of computers over a network. NTP uses a hierarchical, client-server architecture, where a client requests the current time from a server, and the server responds with a timestamp. The client then adjusts its own clock to match the server's time, taking into account the network delay and clock drift. NTP can achieve sub-millisecond accuracy over local area networks and a few milliseconds over the internet12.
NTP can help maintain time synchronization between all the resources in a distributed cloud environment, as it allows each resource to get the accurate time from a reliable source. This can help with correlating logs, auditing, security, and other time-sensitive operations. NTP can also handle different time zones, as it uses Coordinated Universal Time (UTC) as the reference time, and each resource can convert UTC to its local time zone12.
DNS stands for Domain Name System, which is a protocol for resolving domain names into IP addresses. DNS does not provide any functionality for time synchronization3.
IPAM stands for IP Address Management, which is a method for planning, tracking, and managing the IP address space used in a network. IPAM does not provide any functionality for time synchronization.
SNMP stands for Simple Network Management Protocol, which is a protocol for collecting and organizing information about managed devices on a network. SNMP can be used to monitor the performance, availability, configuration, and security of network devices, but it does not provide any functionality for time synchronization.

NEW QUESTION 143
- (Topic 4)
A systems administrator deployed a new web application in a public cloud and would like to test it, but the company's network firewall is only allowing outside connections to the cloud provider network using TCP port 22. While waiting for the network administrator to open the required ports, which of the following actions should the systems administrator take to test the new application? (Select two).

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

Answer: BF

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

NEW QUESTION 145
- (Topic 4)
An organization hosts an ERP database in on-premises infrastructure. A recommendation has been made to migrate the ERP solution to reduce operational overhead in the maintenance of the data center. Which of the following should be considered when migrating this on-premises database to DBaaS?
? • Database application version compatibility
• Database IOPS values
• Database storage utilization
? • Physical database server CPU cache value
• Physical database server DAS type
• Physical database server network I/O
? • Database total user count
• Database total number of tables
• Database total number of storage procedures
• Physical database server memory configuration
• Physical database server CPU frequency

A. • Physical database server operating system

Answer: A

Explanation:
When migrating an on-premises database to DBaaS, it is important to consider the database application version compatibility, the database IOPS values, and the database storage utilization. These factors can affect the performance, functionality, and cost of the migration. Database application version compatibility refers to the ability of the DBaaS provider to support the same or compatible version of the database software as the on- premises database. This can ensure that the database features, syntax, and behavior are consistent and compatible across the environments. Database IOPS values refer to the input/output operations per second that the database performs. This can indicate the workload and throughput of the database, and help determine the appropriate size and configuration of the DBaaS instance. Database storage utilization refers to the amount of disk space that the database consumes. This can affect the cost and scalability of the DBaaS service, and help optimize the storage allocation and backup
strategies. References := CompTIA Cloud+ source documents or study guide
? CompTIA Cloud+ Certification Exam Objectives, Domain 2.0: Deployment, Objective 2.1: Given a scenario, execute and implement solutions using appropriate cloud migration tools and methods.
? Migrate your relational databases to Azure - .NET | Microsoft Learn, Migrate On- premises Tablespaces to DBaaS Database Using Cross-Platform Tablespace Transport
? Migrating On-Premises Databases to the DBaaS Database Using RMAN - Oracle, Overview

NEW QUESTION 146
- (Topic 4)
An organization is implementing a new requirement to facilitate faster downloads for users of corporate application content. At the same time, the organization is also expanding cloud regions. Which of the following would be suitable to optimize the network for this requirement?

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

Answer: A

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

NEW QUESTION 150
- (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 153
- (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 156
- (Topic 4)
An organization's two-node, hybrid container cluster is experiencing failures during horizontal scaling to the cloud cluster instance. The on-premises IP range is 192.168.0.0/16, and the cloud environment is 10.168.0.0/16. Overlapping or stretched VLANs are not permitted, and a node is deployed in each location. The cloud monitoring agent reports a healthy status for the second instance, but when pinging the clusters from on premises, the following output is received:
pinging cluster1. comptia. containers.com C192.168.100 reply
pinging cluster2. comptia. containers.com [192.16B .100 .128] request timed out
Which of the following is the most likely reason for the scaling failure?

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

Answer: E

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

NEW QUESTION 160
- (Topic 4)
A cloud engineer needs to perform a database migration. The database has a restricted SLA and cannot be offline for more than ten minutes per month. The database stores 800GB of data, and the network bandwidth to the CSP is 100MBps Which of the following is the best option to perform the migration?

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

Answer: B

Explanation:
The best option to perform the database migration is to create a replica database, synchronize the data, and switch to the new instance. This option can help meet the restricted SLA and avoid offline time for the database. Creating a replica database can help copy the data from the source to the destination without interrupting the database operations. Synchronizing the data can help ensure that the replica database is updated with any changes that occur in the source database during the migration process. Switching to the new instance can help complete the migration and activate the new database in the cloud. This option can also help avoid the network bandwidth limitation and the large size of the data. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 7, Objective 7.1: Given a scenario, migrate applications and data to the cloud.

NEW QUESTION 161
- (Topic 4)
An enterprise is considering a cost model for a DBaaS. Which of the following is BEST for a cloud solution?

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

Answer: A

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

NEW QUESTION 163
- (Topic 4)
A company has a web application that is accessed around the world. An administrator has been notified of performance issues regarding the application. Which of the following will BEST improve performance?

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

Answer: C

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

NEW QUESTION 165
- (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 IOGbps
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 Input/Output Virtualization, which is a technology that allows a physical device, such as a network interface card (NIC), to be shared by multiple virtual machines (VMs) without sacrificing performance or latency. By enabling SR-IOV capability, the integration application can communicate directly with the physical NIC, bypassing the hypervisor and the virtual switch, and reducing the network overhead and latency .

NEW QUESTION 167
- (Topic 4)
Following the deployment of a new VM, a cloud engineer notices the backup platform has not added the machine to the appropriate job. The backup platform uses a text-based variable for job configuration. This variable is based on the RPO requirements for the workload. Which of the following did the cloud engineer forget to configure when deploying the virtual machine?
? Tags

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

Answer: A

Explanation:
Tags are key-value pairs that can be applied to cloud resources to organize, categorize, and filter them. Tags can also be used to assign resources to backup jobs based on their RPO requirements. The cloud engineer forgot to configure the appropriate tag for the new VM that matches the text-based variable of the backup platform. Therefore, the backup platform did not add the VM to the correct job. References: Tags and labels |
Cloud Storage | Google Cloud, CompTIA Cloud+ Certification Exam Objectives, Domain 4.0: Operations and Support, Objective 4.3: Given a scenario, apply the appropriate methods for cost control in a cloud environment.

NEW QUESTION 168
- (Topic 4)
A systems administrator is configuring a DNS server. Which of the following steps should a technician take to ensure confidentiality between the DNS server and an upstream DNS provider?

A. Enable DNSSEC.
B. Implement single sign-on.
C. Configure DOH.
D. Set up DNS over SSL.

Answer: C

Explanation:

DNS (Domain Name System) is a service that translates human-friendly domain names into IP addresses that can be used to communicate over the Internet1. However, DNS queries and responses are usually sent in plain text, which means that anyone who can intercept the network traffic can see the domain names that the users are requesting. This poses a threat to the confidentiality and privacy of the users and their online activities2.
To ensure confidentiality between the DNS server and an upstream DNS provider, a technician should configure DOH (DNS over HTTPS). DOH is a protocol that encrypts DNS queries and responses using HTTPS (Hypertext Transfer Protocol Secure), which is a secure version of HTTP that uses SSL/TLS (Secure Sockets Layer/Transport Layer Security) to protect the data in transit3. By using DOH, the technician can prevent eavesdropping, tampering, or spoofing of DNS traffic by malicious actors3.
The other options are not the best steps to ensure confidentiality between the DNS server and an upstream DNS provider:
? Option A: Enable DNSSEC (DNS Security Extensions). DNSSEC is a set of
extensions that add digital signatures to DNS records, which can be used to verify the authenticity and integrity of the DNS data. DNSSEC can prevent DNS cache poisoning attacks, where an attacker inserts false DNS records into a DNS server's cache, redirecting users to malicious websites. However, DNSSEC does not encrypt or hide the DNS queries and responses, so it does not provide confidentiality for DNS traffic2.
? Option B: Implement single sign-on (SSO). SSO is a mechanism that allows users
to access multiple services or applications with one set of credentials, such as a username and password. SSO can simplify the authentication process and reduce the risk of password compromise or phishing attacks. However, SSO does not affect the communication between the DNS server and an upstream DNS provider, so it does not provide confidentiality for DNS traffic.
? Option D: Set up DNS over SSL (DNS over Secure Sockets Layer). This option is
not a valid protocol for securing DNS traffic. SSL is a deprecated protocol that has been replaced by TLS (Transport Layer Security), which is more secure and robust. The correct protocol for encrypting DNS traffic using SSL/TLS is DOH (DNS over HTTPS), as explained above.

NEW QUESTION 169
- (Topic 4)
A cloud administrator used a deployment script to recreate a number of servers hosted in a public-cloud provider_ However, after the script completes, the administrator receives the following error when attempting to connect to one of the servers Via SSH from the administrators workstation: CHANGED. Which of the following IS the MOST likely cause of the issue?

A. The DNS records need to be updated
B. The cloud provider assigned a new IP address to the server.
C. The fingerprint on the server's RSA key is different
D. The administrator has not copied the public key to the server.

Answer: C

Explanation:
This error indicates that the SSH client has detected a change in the server's RSA key, which is used to authenticate the server and establish a secure connection. The SSH client stores the fingerprints of the servers it has previously connected to in a file called known_hosts, which is usually located in the ~/.ssh directory. When the SSH client tries to connect to a server, it compares the fingerprint of the server's RSA key with the one stored in the known_hosts file. If they match, the connection proceeds. If they do not match, the SSH client warns the user of a possible man-in-the-middle attack or a host key change, and aborts the connection.
The most likely cause of this error is that the deployment script has recreated the server with a new RSA key, which does not match the one stored in the known_hosts file. This can happen when a server is reinstalled, cloned, or migrated. To resolve this error, the administrator needs to remove or update the old fingerprint from the known_hosts file, and accept the new fingerprint when connecting to the server again. Alternatively, the administrator can use a tool or service that can synchronize or manage the RSA keys
across multiple servers, such as AWS Key Management Service (AWS KMS) 1, Azure Key Vault 2, or HashiCorp Vault 3.

NEW QUESTION 171
- (Topic 4)
A DevOps team needs to provide a solution that offers isolation, portability, and scalability Which of the following would BEST meet these requirements?

A. Virtual machines
B. Containers
C. Appliances
D. Clusters

Answer: B

Explanation:
Containers are a solution that offers isolation, portability, and scalability for software development and deployment. Containers are lightweight and self-contained units of software that package up the application code and all its dependencies, such as libraries, frameworks, and configuration files. Containers run on a container platform, such as Docker or Kubernetes, that provides the runtime environment and orchestration for the containers.
Containers offer isolation, as they run independently from each other and from the underlying host system. Each container has its own namespace, filesystem, network, and resources, and does not interfere with other containers or processes. Containers also offer portability, as they can run on any system that supports the container platform, regardless of the hardware or operating system differences. Containers can be easily moved, copied, or deployed across different environments, such as development, testing, or production. Containers also offer scalability, as they can be dynamically created, destroyed, or replicated to meet the changing demand for the application. Containers can also leverage the distributed computing power of clusters, which are groups of servers that work together to provide high availability and performance .

NEW QUESTION 173
- (Topic 4)
A systems administrator is planning to deploy a database cluster in a virtualization environment. The administrator needs to ensure the database nodes do not exist on the same physical host. Which of the following would best meet this requirement?

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

Answer: B

Explanation:

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

NEW QUESTION 174
- (Topic 4)
A systems administrator wants to be notified every time an application's configuration files are updated. Which of the following should the administrator implement to achieve the objective?

A. ZFS
B. FIM
C. MAC
D. DLP

Answer: B

Explanation:
FIM stands for File Integrity Monitoring, and it is a security technique that monitors and detects changes in files and directories. FIM can help the systems administrator to be notified every time an application’s configuration files are updated by generating alerts or reports when the files are modified, added, deleted, or accessed. FIM can also help verify the integrity and authenticity of the files by comparing their hashes or signatures with a baseline or a trusted source. FIM can be implemented using software
tools or agents that run on the host or the network. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 9, Objective 9.1: Given a scenario, apply security controls and techniques.

NEW QUESTION 179
- (Topic 4)
A company has a large environment with multiple VPCs across three regions in a public cloud. The company is concerned about connectivity within the regions. Which of the following should the cloud administrator implement?

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

Answer: A

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

NEW QUESTION 180
- (Topic 4)
A VDI administrator is enhancing the existing environment with a feature to allow users to connect devices to virtual workstations. Which of the following types of devices are most likely to be allowed in the upgrade? (Select two).

A. Display monitors
B. USB devices
C. SATA devices
D. PCIe devices
E. PCI devices
F. Printers

Answer: BF

Explanation:
B. USB devices and F. Printers are most likely to be allowed in the upgrade. USB devices are common peripherals that users may want to connect to their virtual workstations, such as flash drives, keyboards, mice, webcams, etc. Printers are also useful devices that users may need to print documents from their virtual desktops. VDI software can support USB redirection and printer redirection to enable these devices to work with virtual workstations12.
Display monitors, SATA devices, PCIe devices, and PCI devices are less likely to be allowed in the upgrade, as they are either part of the physical hardware of the end device or the server, or they require direct access to the host system. VDI software typically does not support these types of devices, as they are not compatible with the virtualization layer or the remote display protocol34.
1: What is VDI? | Virtual Desktop Infrastructure | VMware 2: What Is Virtual Desktop Infrastructure (VDI)? | Microsoft Azure 3: What Is Virtual Desktop Infrastructure (VDI)? - Cisco 4: Best Virtual Desktop Infrastructure (VDI) Software in 2023 | G2

NEW QUESTION 182
- (Topic 4)
A systems administrator has been notified of possible illegal activities taking place on the network and has been directed to ensure any relevant emails are preserved for court use.
Which of the following is this MOST likely an example of?

A. Email archiving

B. Version control
C. Legal hold
D. File integrity monitoring

Answer: C

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

NEW QUESTION 187
- (Topic 4)
A company has entered into a business relationship with another organization and needs to provide access to internal resources through directory services. Which of the following should a systems administrator implement?

A. sso
B. VPN
C. SSH
D. SAML

Answer: B

Explanation:
The answer is B. A VPN tunnel. A VPN tunnel is a secure and encrypted connection between two networks over a public network, such as the Internet. A VPN tunnel can help protect data in transit by encrypting it before it leaves the company's network and decrypting it when it reaches the public cloud service provider. A VPN tunnel can also authenticate the endpoints and verify the integrity of the data.
Some possible sources of information about VPN tunnels are:
? What is a VPN Tunnel? | Fortinet: This page explains what a VPN tunnel is, how it works, and what benefits it provides.
? VPN Gateway: Create a Site-to-Site connection using a VPN gateway | Microsoft Docs: This page shows how to create a site-to-site connection using a VPN gateway in Azure.
? [Cloud VPN overview | Google Cloud]: This page provides an overview of Cloud VPN, a service that creates secure and reliable VPN tunnels to Google Cloud.

NEW QUESTION 188
- (Topic 4)
A cloud administrator is performing automated deployment of cloud infrastructure for clients. The administrator notices discrepancies from the baseline in the configuration of infrastructure that was deployed to a new client. Which of the following is most likely the cause?

A. The deployment user account changed
B. The deployment was done to a different resource group.
C. The deployment was done by a different cloud administrator.
D. The deployment template was modified.

Answer: D

Explanation:
A deployment template is a file that defines the resources and configurations that are required to deploy a cloud solution1. A deployment template can be used to automate the deployment of cloud infrastructure for clients, ensuring consistency and efficiency2. However, if the deployment template was modified, either intentionally or accidentally, it could cause discrepancies from the baseline in the configuration of infrastructure that was deployed to a new client. For example, the template could have different parameters, values, or dependencies that affect the outcome of the deployment3. Therefore, the most likely cause of the issue is that the deployment template was modified. References:
1: What is a template? - Azure Resource Manager | Microsoft Docs3
2: Automate cloud deployments with Azure Resource Manager templates - Learn | Microsoft Docs3
3: CompTIA Cloud+ CV0-003 Exam Objectives, Objective 2.2: Given a scenario, deploy and test a cloud solution

NEW QUESTION 190
- (Topic 4)
A cloud administrator is investigating slow VM performance. The administrator has checked the physical server performance and has identified the host is under stress due to a peak usage workload. Which of the following is the NEXT step the administrator should complete?

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

Answer: B

Explanation:
Migrating the VM to a different host is a common technique to improve the performance of a VM that is suffering from resource contention or contention on the physical server. By moving the VM to a different host, the administrator can:
Reduce the stress and load on the original host, which may be under stress due to a peak usage workload.
Increase the availability and reliability of the VM, which may be experiencing slow performance due to resource contention or contention on the original host.

Balance the workload and resource utilization across multiple hosts, which may improve the overall performance and efficiency of the cloud environment. Migrating the VM to a different host can be done manually or automatically, depending on the configuration and capabilities of the cloud platform. Some cloud platforms support live migration, which allows moving a VM to a different host without interrupting its operation or service. Other cloud platforms require shutting down or pausing the VM before migrating it to a different host .

NEW QUESTION 194
- (Topic 3)
A product-based company wants to transition to a method that provides the capability to enhance the product seamlessly and keep the development iterations to a shorter time frame. Which of the following would BEST meet these requirements?

A. Implement a secret management solution.
B. Create autoscaling capabilities.
C. Develop CI/CD tools.
D. Deploy a CMDB tool.

Answer: C

Explanation:
CI/CD tools are software tools that enable continuous integration and continuous delivery or deployment, which are methods to frequently deliver software products to customers by introducing automation into the stages of software development. CI/CD tools can help a product-based company to transition to a method that provides the capability to enhance the product seamlessly and keep the development iterations to a shorter time frame, as they can offer the following benefits:
? Faster and more reliable delivery of software products, as CI/CD tools can automate the processes of building, testing, and deploying code changes, reducing manual errors and delays.
? Higher quality and performance of software products, as CI/CD tools can facilitate ongoing feedback, monitoring, and improvement of the code, ensuring that it meets the customer expectations and requirements.
? Greater collaboration and communication among the development teams, as CI/CD tools can integrate with various tools and platforms, such as version control systems, code repositories, testing frameworks, and cloud services, enabling a seamless workflow and visibility across the software lifecycle.
Some examples of popular CI/CD tools are Jenkins1, CircleCI2, GitLab CI/CD3, and AWS CodeBuild4.

NEW QUESTION 198
- (Topic 3)
A security team is conducting an audit of the security group configurations for the Linux servers that are hosted in a public IaaS. The team identifies the following rule as a potential

Protocol	Port	Source	Description
TCP	22	0.0.0.0/0	Allow SSH access

A cloud administrator, who is working remotely, logs in to the cloud management console and modifies the rule to set the source to "My IR" Shortly after deploying the rule, an internal developer receives the following error message when attempting to log in to the server using SSH: Network error: connection timed out. However, the administrator is able to connect successfully to the same server using SSH. Which of the following is the BEST option for both the developer and the administrator to access the
server from their locations?

A. Modify the outbound rule to allow the company's external IP address as a source.
B. Add an inbound rule to use the IP address for the company's main office as a source.
C. Modify the inbound rule to allow the company's external IP address as a source.
D. Delete the inbound rule to allow the company's external IP address as a source.

Answer: C

Explanation:
The inbound rule that the security team identified as a potential vulnerability is the one that allows SSH access (port 22) from any source (0.0.0.0/0). This means that anyone on the internet can try to connect to the Linux servers using SSH, which poses a risk of unauthorized access or brute-force attacks. The cloud administrator, who is working remotely, logs in to the cloud management console and modifies the rule to set the source to “My IP”. This means that only the administrator’s IP address can connect to the Linux servers using SSH, which improves the security of the servers. However, this also prevents other authorized users, such as the internal developer, from accessing the servers using SSH, as they have different IP addresses than the administrator. Therefore, the administrator needs to modify the rule again to allow more sources for SSH access.
The best option for both the developer and the administrator to access the server from their locations is to modify the inbound rule to allow the company’s external IP address as a source. This means that only the IP addresses that belong to the company’s network can connect to the Linux servers using SSH, which reduces the attack surface and ensures that only authorized users can access the servers. The company’s external IP address can be obtained by using a web service such as [What Is My IP Address?] or [IP Location]. The administrator can then enter this IP address or its CIDR notation in the source field of the inbound rule.

NEW QUESTION 200
- (Topic 3)
A cloud engineer has deployed a virtual storage appliance into a public cloud environment. The storage appliance has a NAT to a public IP address. An administrator later notices there are some strange files on the storage appliance and a large spike in network traffic on the machine. Which of the following is the MOST likely cause?

A. The default password is still configured on the appliance.
B. The appliance's certificate has expired.
C. The storage appliance has no firewall.
D. Data encryption is enabled, and the files are hashed.

Answer: A

Explanation:
One possible cause for the strange files and the large spike in network traffic on the storage appliance is that the default password is still configured on the

appliance. A default password is a password that is set by the manufacturer or vendor of a device or software, and it is often easy to guess or find online. If the cloud engineer did not change the default password after deploying the virtual storage appliance, it could allow unauthorized users to access the appliance remotely and upload or download files, which could explain the symptoms observed by the administrator. This is a serious security risk that could compromise the confidentiality, integrity, and availability of the data stored on the appliance.

NEW QUESTION 201
- (Topic 3)
A company is deploying a public cloud solution for an existing application using lift and shift. The requirements for the applications are scalability and external access. Which of the following should the company implement? (Select TWO).

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

Answer: AF

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

NEW QUESTION 203
- (Topic 3)
An enterprise recently upgraded the memory of its on-premises VMs from 8GB to 16GB. However, users are not experiencing any performance benefit. Which of the following is the MOST likely reason?

A. Insufficient memory on the hypervisor
B. Operating system memory limit
C. Memory mismatch error
D. Dynamic memory allocation

Answer: B

Explanation:
The most likely reason why users are not experiencing any performance benefit after upgrading the memory of their on-premises VMs is that the operating system has a memory limit that prevents it from using more than 8GB of RAM. This could be due to the operating system version, edition, or configuration. The systems administrator should check the operating system settings and requirements and adjust them accordingly to allow the VMs to use the full 16GB of RAM. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 4.0 Troubleshooting, Objective 4.3 Given a scenario, troubleshoot capacity issues within a cloud environment.

NEW QUESTION 208
- (Topic 3)
A company would like to migrate its current on-premises workloads to the public cloud. The current platform requires at least 80 instances running at all times to work properly. The company wants the workload to be highly available, even if the cloud provider loses one region due to a catastrophe, and the costs to be kept to a minimum. Which of the following strategies should the company implement?

A. Create /25 subnets in two regions and run 80 instances on each one.
B. Create /26 subnets in two regions and run 40 instances on each one.
C. Create /26 subnets in three regions and run 40 instances on each one.
D. Create /26 subnets in three regions and run 80 instances on each one.

Answer: B

Explanation:
The best strategy to migrate the current on-premises workloads to the public cloud for the company that requires at least 80 instances running at all times and wants the workload to be highly available and cost-effective is to create /26 subnets in two regions and run 40 instances on each one. A /26 subnet can accommodate up to 62 hosts, which is enough for 40 instances. By creating subnets in two regions, the company can achieve high availability and redundancy in case one region fails due to a catastrophe. By running 40 instances on each subnet, the company can meet the minimum requirement of 80 instances and also save on costs by avoiding overprovisioning or underutilization of resources. Reference: What is VPN? How It Works, Types of VPN - Kaspersky

NEW QUESTION 210
- (Topic 3)
A systems administrator needs to modify the replication factors of an automated application container from 3 to 5. Which of the following file types should the systems administrator modify on the master controller?

A. .yaml
B. . txt
C. .conf
D. .etcd

Answer: A

Explanation:

A .yaml file is a common file type for defining the configuration and deployment of application containers in a cloud environment. YAML stands for YAML Ain't Markup Language, and it is a human-readable data serialization language that uses indentation and keywords to represent the structure and values of the data1. A .yaml file can specify the properties of the container, such as the image, ports, environment variables, volumes, resources, and scaling rules. For example, to modify the replication factor of an automated application container from 3 to 5, the systems administrator can edit the .yaml file on the master controller and change the value of the replicas field under the spec section2. For example:
apiVersion: apps/v1 kind: Deployment metadata: name: my-app spec: replicas: 5 # change this value from 3 to 5 selector: matchLabels: app: my-app template: metadata: labels: app: my-app spec: containers: - name: my-app image: my-app-image ports: - containerPort: 80 A .txt file is a plain text file that can store any kind of text data, but it is not a standard format for defining container configurations. A .conf file is a configuration file that can store settings and parameters for various applications or services, but it is not commonly used for container deployments. A .etcd file is not a valid file type, but etcd is a distributed key- value store that provides a reliable way to store data across a cluster of machines3. Etcd is often used by Kubernetes, a popular platform for managing containerized applications, to store and synchronize the cluster state. However, etcd does not store the configuration files of the containers, but rather the runtime data of the cluster. Therefore, none of these options are correct.

NEW QUESTION 214
- (Topic 3)
A systems administrator is responding to an outage in a cloud environment that was caused by a network-based flooding attack. Which of the following should the administrator configure to mitigate the attack?

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

Answer: C

Explanation:
A DDoS (distributed denial-of-service) attack is a type of network-based flooding attack that aims to overwhelm a target server or network with a large volume of traffic from multiple sources, making it unavailable or slow for legitimate users. According to the web search results, DDoS protection is a service or a solution that can detect and mitigate DDoS attacks by filtering out malicious traffic and allowing only legitimate traffic to pass through .
A NIPS (network intrusion prevention system) is a device or a software that can monitor, detect, and block malicious activity on a network, such as unauthorized access, malware, or policy violations. However, a NIPS may not be effective against DDoS attacks, as it can also be overwhelmed by the flood of traffic and fail to distinguish between legitimate and malicious requests.
A network overlay using GENEVE (Generic Network Virtualization Encapsulation) is a protocol that can create virtual networks on top of physical networks, allowing different cloud environments to communicate with each other. However, a network overlay using GENEVE does not provide any protection against DDoS attacks, as it does not filter or block any traffic.
A DoH (DNS over HTTPS) is a protocol that can encrypt and secure DNS queries and responses over HTTPS, preventing eavesdropping or tampering by third parties. However, a DoH does not prevent DDoS attacks, as it does not affect the amount or the source of the traffic.

NEW QUESTION 215
- (Topic 3)
A cloud administrator implemented SSO and received a business requirement to increase security when users access the cloud environment. Which of the following should be implemented NEXT to improve the company's security posture?

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

Answer: B

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

NEW QUESTION 218
- (Topic 3)
An administrator manages a file server that has a lot of users accessing and creating many files. As a result, the storage consumption is growing quickly. Which of the following would BEST control storage usage?

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

Answer: C

Explanation:
User quotas are a feature that allows the administrator to limit the amount of storage space that a user or a group of users can consume on a file server. User quotas can help to control storage usage by preventing users from storing excessive or unnecessary files, as well as by enforcing fair and consistent storage policies across the organization. User quotas can also help to monitor and report on the storage consumption and trends of the users, and alert the administrator or the users when they are approaching or exceeding their quota limits.

NEW QUESTION 223

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

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

Answer: D

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

NEW QUESTION 227
- (Topic 3)
A systems administrator needs to deploy a solution to automate new application releases that come from the development team. The administrator is responsible for provisioning resources at the infrastructure layer without modifying any configurations in the application code. Which of the following would BEST accomplish this task?

A. Implementing a CI/CD tool
B. Configuring infrastructure as code
C. Deploying an orchestration tool
D. Employing DevOps methodology

Answer: B

Explanation:
Infrastructure as code (IaC) is a method of provisioning and managing cloud resources using code or scripts, rather than manual processes or GUI tools. This allows for automation, consistency, scalability, and version control of the infrastructure layer. This would be the best option to deploy a solution to automate new application releases that come from the development team without modifying any configurations in the application code. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 3.0 Maintenance, Objective 3.4 Given a scenario, implement automation and orchestration to optimize cloud operations.

NEW QUESTION 230
- (Topic 3)
A cloud administrator is configuring several security appliances hosted in the private IaaS environment to forward the logs to a central log aggregation solution using syslog. Which of the following firewall rules should the administrator add to allow the web servers to connect to the central log collector?

A. Allow UDP 161 outbound from the web servers to the log collector .
B. Allow TCP 514 outbound from the web servers to the log collector.
C. Allow UDP 161 inbound from the log collector to the web servers .
D. Allow TCP 514 inbound from the log collector to the web servers .

Answer: B

Explanation:
As mentioned in the question, the security appliances are using syslog to forward the logs to a central log aggregation solution. According to the web search results, syslog is a protocol that runs over UDP port 514 by default, or TCP port 6514 for secure and reliable transport1. However, some implementations of syslog can also use TCP port 514 for non-secure transport2. Therefore, to allow the web servers to connect to the central log collector using syslog over TCP, the firewall rule should allow TCP 514 outbound from the web servers to the log collector.

NEW QUESTION 235
- (Topic 3)
A cloud administrator has deployed several VM instances that are running the same applications on VDI nodes. Users are reporting that a role instance is looping between STARTED, INITIALIZING, BUSY, and stop. Upon investigation, the cloud administrator can see the status changing every few minutes. Which of the following should be done to resolve the issue?

A. Reboot the hypervisor.
B. Review the package and configuration file.
C. Configure service healing.
D. Disable memory swap.

Answer: B

Explanation:
The best way to resolve the issue where a role instance is looping between STARTED, INITIALIZING, BUSY, and STOP after deploying several VM instances that are running the same applications on VDI nodes is to review the package and configuration file. The package and configuration file are the components that define the application and its settings for the VM instances. The package contains the application code, binaries, and dependencies, while the configuration file contains the parameters, values, and settings for the application. Reviewing these components can help identify and fix any errors or inconsistencies that may cause the role instance to loop or fail. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 4.0 Troubleshooting, Objective 4.4 Given a scenario, troubleshoot deployment issues.

NEW QUESTION 239
.....

Relate Links

100% Pass Your CV0-003 Exam with Exambible Prep Materials

https://www.exambible.com/CV0-003-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/