CompTIA

Exam Questions CV0-003

CompTIA Cloud+ Certification Exam

NEW QUESTION 1
- (Topic 1)
Which of the following cloud deployment models allows a company to have full control over its IT infrastructure?

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

Answer: A

Explanation:
A private cloud is a type of cloud deployment model that provides cloud services exclusively to a single organization or tenant. A private cloud allows a company to have full control over its IT infrastructure, as it can customize, configure, manage, and secure its own cloud environment according to its specific needs and preferences. A private cloud can also offer higher performance, reliability, and privacy than other cloud deployment models, as it does not share resources or data with other customers.
References: CompTIA Cloud+ Certification Exam Objectives, page 8, section 1.2 Reference: https://www.sciencedirect.com/topics/computer-science/private-cloud

NEW QUESTION 2
- (Topic 1)
A cloud administrator is reviewing the authentication and authorization mechanism implemented within the cloud environment. Upon review, the administrator discovers the sales group is part of the finance group, and the sales team members can access the financial application. Single sign-on is also implemented, which makes access much easier.
Which of the following access control rules should be changed?

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

Answer: D

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

NEW QUESTION 3
- (Topic 1)
A DevOps administrator is automating an existing software development workflow. The administrator wants to ensure that prior to any new code going into production, tests confirm the new code does not negatively impact existing automation activities.
Which of the following testing techniques would be BEST to use?

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

Answer: B

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

NEW QUESTION 4
- (Topic 1)
An SQL injection vulnerability was reported on a web application, and the cloud platform
team needs to mitigate the vulnerability while it is corrected by the development team. Which of the following controls will BEST mitigate the risk of exploitation?

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

Answer: D

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

NEW QUESTION 5
- (Topic 1)
An organization has the following requirements that need to be met when implementing cloud services:
? SSO to cloud infrastructure
? On-premises directory service
? RBAC for IT staff
Which of the following cloud models would meet these requirements?

A. Public
B. Community
C. Hybrid
D. Multitenant

Answer: C

Explanation:
A hybrid cloud is a type of cloud deployment model that combines two or more different types of clouds, such as public, private, or community clouds, into a single integrated environment. A hybrid cloud can meet the requirements for implementing cloud services with SSO to cloud infrastructure, on-premises directory service, and RBAC for IT staff, as it can provide flexibility, scalability, and security for cloud-based and on-premises resources. A hybrid cloud can also enable seamless and secure access to cloud infrastructure using SSO with directory service federation, as well as granular and consistent control over IT staff permissions using RBAC across different cloud environments. References: CompTIA Cloud+ Certification Exam Objectives, page 8, section 1.2

NEW QUESTION 6
- (Topic 1)
A systems administrator needs to configure an email client to ensure data integrity of the email messages.
Which of the following provides the BEST mechanism to achieve this goal?

A. Cyclic redundancy check
B. SHA-1 hashes
C. SHA-256 hashes
D. Digital signature

Answer: D

Explanation:
A digital signature is a type of cryptographic technique that verifies the authenticity, integrity, and non-repudiation of an electronic message or document. A digital signature can help configure an email client to ensure data integrity of the email messages, as it can prove that the email message has not been altered or tampered with during transmission by using a mathematical algorithm to generate a unique code (signature) based on the content and identity of the sender. A digital signature can also help prevent spoofing, phishing, or impersonation attacks, as it can confirm that the email message originates from a legitimate source. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7
Reference: https://www.fsl.cs.sunysb.edu/docs/integrity-storagess05/integrity.html

NEW QUESTION 7
- (Topic 1)
A cloud administrator is switching hosting companies and using the same script that was previously used to deploy VMs in the new cloud. The script is returning errors that the command was not found.
Which of the following is the MOST likely cause of the script failure?

A. Account mismatches
B. IP address changes
C. API version incompatibility
D. Server name changes

Answer: C

Explanation:
An application programming interface (API) is a set of rules or protocols that defines how different systems or applications can communicate or interact with each other. An API version is a specific iteration or release of an API that may have different features or functionalities than previous or subsequent versions. API version incompatibility is the most likely cause of the script failure when switching hosting companies and using the same script that was previously used to deploy VMs in the new cloud, as it can result in errors or failures when trying to execute commands or functions that are not supported or recognized by the new cloud provider's API version. The issue can be resolved by updating or modifying the script to match the new cloud provider's API version.
References: CompTIA Cloud+ Certification Exam Objectives, page 13, section 2.5

NEW QUESTION 8
- (Topic 1)
A systems administrator is creating a playbook to run tasks against a server on a set schedule.
Which of the following authentication techniques should the systems administrator use within the playbook?

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

Answer: C

Explanation:
A service account is a type of user account that is created for a specific service or application to run on a server or system. Creating a service account on the server is the best authentication technique to use within the playbook to run tasks against the server on a set schedule, as it can provide secure and consistent access to the server without exposing or hard-coding any sensitive credentials within the playbook. Creating a service account can also help manage and monitor

the tasks and activities performed by the service or application on the server. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 9
- (Topic 1)
A global web-hosting company is concerned about the availability of its platform during an upcoming event. Web traffic is forecasted to increase substantially during the next week. The site contains mainly static content.
Which of the following solutions will assist with the increased workload?

A. DoH
B. WAF
C. IPS
D. CDN

Answer: D

Explanation:
A content delivery network (CDN) is a distributed network of servers that delivers web content to users based on their geographic location, origin server, and content delivery server. A CDN can assist with the increased workload caused by sudden continuous bursts of traffic, as it can reduce the load on the origin server by caching and serving static content from edge servers closer to the users. A CDN can also improve the performance and availability of web content delivery, as it can reduce latency, bandwidth consumption, and network congestion. References: CompTIA Cloud+ Certification Exam Objectives, page 12, section 2.2
Reference: https://www.globaldots.com/content-delivery-network-explained

NEW QUESTION 10
- (Topic 1)
A company that utilizes an IaaS service provider has contracted with a vendor to perform a penetration test on its environment. The vendor is able to exploit the virtualization layer and obtain access to other instances within the cloud provider's environment that do not belong to the company.
Which of the following BEST describes this attack?

A. VM escape
B. Directory traversal
C. Buffer overflow
D. Heap spraying

Answer: A

Explanation:
VM escape is a type of attack that allows an attacker to break out of a virtual machine (VM) and access the host system or other VMs within the same cloud provider's environment. VM escape can exploit the vulnerabilities in the virtualization layer or hypervisor that separates and isolates the VMs from each other and from the host system. VM escape can result in serious consequences, such as compromising the security and privacy of other customers' data or resources, gaining unauthorized access to the cloud provider's infrastructure or services, or launching further attacks on other systems or networks. VM escape best describes the attack that was performed by a vendor who was able to exploit the virtualization layer and obtain access to other instances within the cloud provider's environment that do not belong to the company. References: CompTIA Cloud+ Certification Exam Objectives, page 19, section 4.1
Reference: https://whatis.techtarget.com/definition/virtual-machine-escape

NEW QUESTION 10
- (Topic 1)
A systems administrator needs to configure a set of policies to protect the data to comply with mandatory regulations.
Which of the following should the administrator implement to ensure DLP efficiently prevents the exposure of sensitive data in a cloud environment?

A. Integrity
B. Versioning
C. Classification
D. Segmentation

Answer: C

Explanation:
Classification is a process of assigning labels or categories to data based on its sensitivity, value, or risk level. Classification can help implement data loss prevention (DLP) policies by identifying which data needs to be protected and how to protect it according to its classification level. Classification can also help comply with mandatory regulations by ensuring that data is handled and stored appropriately based on its legal or contractual requirements. Classification is essential for DLP to efficiently prevent the exposure of sensitive data in a cloud environment. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 14
- (Topic 1)
An organization requires the following to be achieved between the finance and marketing departments:
? Allow HTTPS/HTTP.
? Disable FTP and SMB traffic.
Which of the following is the MOST suitable method to meet the requirements?

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

Answer: B

Explanation:
An access control list (ACL) is a set of rules that defines which traffic is allowed or denied between different network segments or devices. An ACL can be used to

filter traffic based on various criteria, such as source and destination addresses, ports, protocols, and applications. Configuring an ACL between the VLANs of the finance and marketing departments is the most suitable method to meet the requirements of allowing HTTPS/HTTP and disabling FTP and SMB traffic. An ACL can specify which ports and protocols are permitted or blocked between the VLANs, such as allowing port 80 (HTTP) and port 443 (HTTPS), and denying port 21 (FTP) and port 445 (SMB). References: [CompTIA Cloud+ Certification Exam Objectives], page 15, section 2.8

NEW QUESTION 17
- (Topic 1)
A cloud administrator is setting up a DR site on a different zone of the same CSP. The application servers are replicated using the VM replication, and the database replication is set up using log shipping. Upon testing the DR site, the application servers are unable to access the database servers. The administrator has verified the systems are running and are accessible from the CSP portal.
Which of the following should the administrator do to fix this issue?

A. Change the database application IP
B. Create a database cluster between the primary site and the DR site
C. Update the connection string
D. Edit the DNS record at the DR site for the application servers

Answer: C

Explanation:
A connection string is a parameter that specifies how to connect to a database server or instance. A connection string typically includes information such as the server name, database name, user name, password, and other options. Updating the connection string is the best way to fix the issue of application servers being unable to access the database servers after setting up a DR site on a different zone of the same CSP and replicating the application and database servers using VM replication and log shipping. Updating the connection string can ensure that the application servers can connect to the correct database server or instance in the DR site, as the server name or IP address may have changed after the replication. References: CompTIA Cloud+ Certification Exam Objectives, page 10, section 1.5

NEW QUESTION 20
- (Topic 1)
A systems administrator recently upgraded the processors in a web application host. Upon the next login, the administrator sees a new alert regarding the license being out of compliance.
Which of the following licensing models is the application MOST likely using?

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

Answer: C

Explanation:
Core-based licensing is a type of licensing model that charges based on the number of processor cores in a system or server. Core-based licensing is often used by software vendors to align their pricing with the performance and capacity of modern hardware. Core-based licensing can also enable customers to optimize their licensing costs by choosing the appropriate hardware configuration for their needs. Upgrading the processors in a web application host can affect the core-based licensing of the application, as it may increase the number of cores that need to be licensed. This can result in an alert regarding the license being out of compliance if the license is not updated accordingly. References: CompTIA Cloud+ Certification Exam Objectives, page 20, section 4.2
Reference: https://download.microsoft.com/download/3/d/4/3d42bdc2-6725-4b29-b75a- a5b04179958b/ percorelicensing_definitions_vlbrief.pdf

NEW QUESTION 22
- (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 27
- (Topic 2)
A systems administrator is using a configuration management tool to perform maintenance tasks in a system. The tool is leveraging the target system's API to perform these maintenance tasks After a number of features and security updates are applied to the target system, the configuration management tool no longer works as expected. Which of the following is the MOST likely cause of the issue?

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

Answer: A

Explanation:
The target system's API (Application Programming Interface) functionality has been deprecated is what will most likely cause the issue of configuration management tool no longer working as expected after using it to perform maintenance tasks in a system using its API, and applying features and security updates to it. An API is a set of rules or specifications that defines how different software components or systems can communicate and interact with each other. An API functionality is a feature or function that an API provides or supports, such as methods, parameters, responses, etc. An API functionality can be deprecated when it is no longer maintained or supported by the API provider or developer, and is replaced or removed by a newer or better functionality. The target system's API functionality has been deprecated can cause the issue by making the configuration management tool unable to use or access the API functionality that it relies on to perform maintenance tasks in the system, which may result in errors or failures.

NEW QUESTION 29
- (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 32
- (Topic 2)
A database analyst reports it takes two hours to perform a scheduled job after onboarding 10,000 new users to the system. The analyst made no changes to the scheduled job before or after onboarding the users. The database is hosted in an IaaS instance on a cloud provider. Which of the following should the cloud administrator evaluate to troubleshoot the performance of the job?

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

Answer: A

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

NEW QUESTION 35
- (Topic 2)
Users are experiencing slow response times from an intranet website that is hosted on a cloud platform. There is a site-to-site VPN connection to the cloud provider over a link of 100Mbps.
Which of the following solutions will resolve the issue the FASTEST?

A. Change the connection to point-to-site VPN
B. Order a direct link to the provider
C. Enable quality of service
D. Upgrade the link to 200Mbps

Answer: B

Explanation:
Ordering a direct link to the provider is the fastest solution to resolve the issue of slow response times from an intranet website that is hosted on a cloud platform. A direct link is a dedicated, high-bandwidth, low-latency connection between the customer's network and the cloud provider's network. It bypasses the public internet and provides better performance, security, and reliability. Examples of direct links are AWS Direct Connect, Azure ExpressRoute, Google Cloud Interconnect, etc.

NEW QUESTION 39
- (Topic 2)
An update is being deployed to a web application, and a systems administrator notices the cloud SQL database has stopped running. The VM is responding to pings, and there were not any configuration changes scheduled for the VM. Which of the following should the administrator check NEXT?

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

Answer: A

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

NEW QUESTION 44
- (Topic 2)
A system administrator is migrating a bare-metal server to the cloud. Which of the following types of migration should the systems administrator perform to accomplish this task?

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

Answer: D

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

NEW QUESTION 48
- (Topic 2)
A systems administrator has received an email from the virtualized environment's alarms indicating the memory was reaching full utilization. When logging in, the administrator notices that one out of a five-host cluster has a utilization of 500GB out of 512GB of RAM. The baseline utilization has been 300GB for that host.
Which of the following should the administrator check NEXT?

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

Answer: D

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

NEW QUESTION 51
- (Topic 2)
A technician needs to deploy two virtual machines in preparation for the configuration of a financial application next week. Which of the following cloud deployment models should the technician use?

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

Answer: B

Explanation:
IaaS (Infrastructure as a Service) is the cloud deployment model that the technician should use to deploy two virtual machines in preparation for the configuration of a financial application next week. IaaS is a cloud service model that provides basic computing resources such as servers, storage, network, etc., to the customers. The customers have full control and flexibility over these resources and can install and configure any software they need on them. IaaS is suitable for deploying virtual machines, as it allows the customers to choose their preferred OS, applications, settings, etc., and customize them according to their needs.

NEW QUESTION 55
- (Topic 2)
A systems administrator has finished installing monthly updates to servers in a cloud environment. The administrator notices certain portions of the playbooks are no longer functioning. Executing the playbook commands manually on a server does not work as well. There are no other reports of issues.
Which of the following is the MOST likely cause of this issue?

A. Change management failure
B. Service overload
C. Patching failure
D. Job validation issues
E. Deprecated features

Answer: E

Explanation:

Deprecated features are features that are no longer supported or recommended by the software vendor or provider. They may be removed or replaced by newer features in future updates or versions. If a playbook relies on deprecated features, it may stop functioning after an update or patch is applied to the software. The administrator should check the release notes or documentation of the software to identify and replace any deprecated features in the playbook.

NEW QUESTION 60
- (Topic 2)
A DevOps administrator is designing a new machine-learning platform. The application needs to be portable between public and private clouds and should be kept as small as possible. Which of the following approaches would BEST meet these requirements?

A. Virtual machines
B. Software as a service
C. Serverless computing
D. Containers

Answer: D

Explanation:
Containers are the best approach to design a new machine-learning platform that needs to be portable between public and private clouds and should be kept as small as possible. Containers are isolated environments that can run applications and their dependencies without interfering with other processes or systems. Containers are lightweight, portable, and scalable, which makes them ideal for machine-learning applications. Containers can be moved easily between public and private clouds without requiring any changes or modifications. Containers can also reduce the size and complexity of applications by using only the necessary components and libraries.

NEW QUESTION 65
- (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 67
- (Topic 2)
A systems administrator is working in a globally distributed cloud environment. After a file server VM was moved to another region, all users began reporting slowness when saving files. Which of the following is the FIRST thing the administrator should check while troubleshooting?

A. Network latency
B. Network connectivity
C. Network switch
D. Network peering

Answer: A

Explanation:
Network latency is the first thing that the administrator should check while troubleshooting slowness when saving files after a file server VM was moved to another region in a globally distributed cloud environment. Network latency is a measure of how long it takes for data to travel from one point to another over a network or connection. Network latency can affect performance and user experience of cloud applications or services by determining how fast data can be transferred or processed between clients and servers or vice versa. Network latency can vary depending on various factors, such as distance, bandwidth, congestion, interference, etc. Network latency can increase when a file server VM is moved to another region in a globally distributed cloud environment, as it may increase the distance and decrease the bandwidth between clients and servers, which may result in delays or errors in data transfer or processing.

NEW QUESTION 71
- (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 76
- (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 81
- (Topic 2)
A disaster situation has occurred, and the entire team needs to be informed about the situation. Which of the following documents will help the administrator find the details of the relevant team members for escalation?

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

Answer: D

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

NEW QUESTION 83
- (Topic 2)
A Chief Information Security Officer (CISO) is evaluating the company’s security management program. The CISO needs to locate all the assets with identified deviations and mitigation measures. Which of the following would help the CISO with these requirements?

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

Answer: D

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

NEW QUESTION 84
- (Topic 2)
An engineer is responsible for configuring a new firewall solution that will be deployed in a new public cloud environment. All traffic must pass through the firewall. The SLA for the firewall is 99.999%. Which of the following should be deployed?

A. Two load balancers behind a single firewall
B. Firewalls in a blue-green configuration
C. Two firewalls in a HA configuration
D. A web application firewall

Answer: C

Explanation:
Deploying two firewalls in a HA (High Availability) configuration is the best option to ensure all traffic passes through the firewall and meets the SLA (Service Level Agreement) of 99.999%. HA is a design principle that aims to minimize downtime and ensure continuous operation of a system or service. HA can be achieved by using redundancy, failover, load balancing, clustering, etc. Two firewalls in a HA configuration can provide redundancy and failover in case one firewall fails or becomes overloaded.

NEW QUESTION 86
- (Topic 2)
Which of the following would be the BEST option for discussion of what individuals should do in an incident response or disaster recovery scenario?

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

Answer: C

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

NEW QUESTION 90
- (Topic 2)
An administrator has been informed that some requests are taking a longer time to respond than other requests of the same type. The cloud consumer is using multiple network service providers and is performing link load balancing for bandwidth aggregation. Which of the following commands will help the administrator understand the possible latency issues?

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

Answer: A

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

NEW QUESTION 94
- (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 95
- (Topic 2)
A cloud provider wants to make sure consumers are utilizing its IaaS platform but prevent them from installing a hypervisor on the server. Which of the following will help the cloud provider secure the environment and limit consumers' activity?

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

Answer: B

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

NEW QUESTION 100
- (Topic 2)
A development team recently completed testing changes to a company's web-based CMS in the sandbox environment. The cloud administrator deployed these CMS application changes to the staging environment as part of the next phase in the release life cycle. The deployment was successful, but after deploying the

CMS application, the web page displays an error message stating the application is unavailable. After reviewing the application logs, the administrator sees an error message that the CMS is unable to connect to the database. Which of the following is the BEST action for the cloud administrator to perform to resolve the issue?

A. Modify the deployment script to delete and recreate the database whenever the CMS application is deployed.
B. Modify the ACL to allow the staging environment to access the database in the sandbox environment.
C. Modify the CMS application deployment to use the previous version and redeploy the application.
D. Modify the configuration settings of the CMS application to connect to the database in the current environment.

Answer: D

Explanation:
Modifying the configuration settings of the CMS (Content Management System) application to connect to the database in the current environment is what the cloud administrator should do to resolve the issue of web page displaying an error message stating the application is unavailable after deploying CMS application changes to the staging environment. A CMS is a software or platform that allows users to create, manage, and publish web content. A CMS may use a database to store and retrieve web content and information. A staging environment is a testing or pre-production environment that simulates the production environment and allows users to verify and validate changes or updates before deploying them to production. Modifying the configuration settings of the CMS application can help to resolve the issue by ensuring that the CMS application can access and communicate with the database in the current environment, rather than using the previous or default settings that may point to a different or non-existent database.

NEW QUESTION 105
- (Topic 2)
A company had a system compromise, and the engineering team resolved the issue after 12 hours. Which of the following information will MOST likely be requested by the Chief Information Officer (CIO) to understand the issue and its resolution?

A. A root cause analysis
B. Application documentation
C. Acquired evidence
D. Application logs

Answer: A

Explanation:
A root cause analysis is what will most likely be requested by the Chief Information Officer (CIO) to understand the issue and its resolution after a system compromise that was resolved by the engineering team after 12 hours. A root cause analysis is a technique of investigating and identifying the underlying or fundamental cause or reason for an incident or issue that affects or may affect the normal operation or performance of a system or service. A root cause analysis can help to understand the issue and its resolution by providing information such as:
? What happened: This describes what occurred during the incident or issue, such as symptoms, effects, impacts, etc.
? Why it happened: This explains why the incident or issue occurred, such as triggers, factors, conditions, etc.
? How it was resolved: This details how the incident or issue was fixed or mitigated, such as actions, steps, methods, etc.
? How it can be prevented: This suggests how the incident or issue can be avoided or reduced in the future, such as recommendations, improvements, changes, etc.

NEW QUESTION 107
- (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 108
- (Topic 2)
Which of the following actions should a systems administrator perform during the containment phase of a security incident in the cloud?

A. Deploy a new instance using a known-good base image.
B. Configure a firewall rule to block the traffic on the affected instance.
C. Perform a forensic analysis of the affected instance.
D. Conduct a tabletop exercise involving developers and systems administrators.

Answer: B

Explanation:
Configuring a firewall rule to block the traffic on the affected instance is what the administrator should perform during the containment phase of a security incident in the cloud. A security incident is an event or situation that affects or may affect the confidentiality, integrity, or availability of cloud resources or data. A security incident response is a process of managing and resolving a security incident using various phases, such as identification, containment, eradication, recovery, etc. The containment phase is where the administrator tries to isolate and prevent the spread or escalation of the security
incident. Configuring a firewall rule to block the traffic on the affected instance can help to contain a security incident by cutting off any communication or

interaction between the instance and other systems or networks, which may stop any malicious or unauthorized activity or access.

NEW QUESTION 113
- (Topic 1)
A web server has been deployed in a public IaaS provider and has been assigned the public IP address of 72.135.10.100. Users are now reporting that when they browse to the website, they receive a message indicating the service is unavailable. The cloud administrator logs into the server, runs a netstat command, and notices the following relevant output:

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

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

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

Answer: D

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

NEW QUESTION 116
- (Topic 1)
Company A has acquired Company B and is in the process of integrating their cloud resources. Company B needs access to Company A's cloud resources while retaining its IAM solution.
Which of the following should be implemented?

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

Answer: C

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

NEW QUESTION 118
- (Topic 1)
An organization purchased new servers with GPUs for render farms. The servers have limited CPU resources.
Which of the following GPU configurations will be the MOST optimal for virtualizing this environment?

A. Dedicated
B. Shared
C. Passthrough
D. vGPU

Answer: C

Explanation:
Passthrough is a type of GPU configuration that allows a VM to directly access a physical GPU on the host system without any virtualization layer or sharing mechanism. Passthrough can provide optimal performance and compatibility for GPU- intensive applications, such as rendering or gaming, as it eliminates any overhead or contention caused by virtualization or sharing. Passthrough is also suitable for servers with limited CPU resources, as it reduces the CPU load and offloads the graphics processing to the GPU. Passthrough is the most optimal GPU configuration for virtualizing a new server with GPUs for render farms.
References: CompTIA Cloud+ Certification Exam Objectives, page 11, section 1.6

NEW QUESTION 119
- (Topic 1)
A developer is no longer able to access a public cloud API deployment, which was working ten minutes prior.
Which of the following is MOST likely the cause?

A. API provider rate limiting
B. Invalid API token

C. Depleted network bandwidth
D. Invalid API request

Answer: A

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

NEW QUESTION 120
- (Topic 1)
A cloud architect is designing the VPCs for a new hybrid cloud deployment. The business requires the following:
? High availability
? Horizontal auto-scaling
? 60 nodes peak capacity per region
? Five reserved network IP addresses per subnet
? /24 range
Which of the following would BEST meet the above requirements?

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

Answer: C

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

NEW QUESTION 124
- (Topic 1)
A systems administrator wants to have near-real-time information on the volume of data being exchanged between an application server and its clients on the Internet.
Which of the following should the systems administrator implement to achieve this objective?

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

Answer: D

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

NEW QUESTION 126
- (Topic 1)
A cloud administrator has finished setting up an application that will use RDP to connect. During testing, users experience a connection timeout error.
Which of the following will MOST likely solve the issue?

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

Answer: D

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

NEW QUESTION 129
- (Topic 1)
A SaaS provider wants to maintain maximum availability for its service. Which of the following should be implemented to attain the maximum SLA?

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

Answer: B

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

NEW QUESTION 132
- (Topic 1)
A systems administrator for an e-commerce company will be migrating the company's main website to a cloud provider. The principal requirement is that the website must be highly available. Which of the following will BEST address this requirement?

A. Vertical scaling
B. A server cluster
C. Redundant switches
D. A next-generation firewall

Answer: B

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

NEW QUESTION 133
- (Topic 1)
An organization is required to set a custom registry key on the guest operating system. Which of the following should the organization implement to facilitate this requirement?

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

Answer: A

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

NEW QUESTION 138
- (Topic 4)
A cloud administrator is evaluating a solution that will limit access to authorized individuals. The solution also needs to ensure the system that connects to the environment meets patching, antivirus, and configuration requirements. Which of the following technologies would BEST meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 142
- (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 145
- (Topic 4)
An organization deployed an application using a cloud provider's internal managed certificates. Developers are unable to retrieve data when calling the API from any machine.
The following error message is in the log:
12-04-2023-10:05:25, SSL Negotiation Error 12-04-2023-10:05:28,Invalid Certificate
12-04-2023-10:05:29, TLS Handshake Failed 12-04-2023-10:05:30,Connection Closed
Which of the following is the most likely cause of the error?

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

Answer: D

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

NEW QUESTION 146
- (Topic 4)
A systems administrator audits a cloud application and discovers one of the key regulatory requirements has not been addressed. The requirement states that if a physical breach occurs and hard drives are stolen, the contents of the drives should not be readable. Which of the following should be used to address the requirement?

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

Answer: B

Explanation:
Encryption is the process of transforming data into an unreadable format using a secret key or algorithm. Encryption can be used to protect data at rest or in transit from unauthorized access or theft. If a physical breach occurs and hard drives are stolen, encryption can prevent the contents of the drives from being readable by anyone who does not have the decryption key or algorithm.
References: [CompTIA Cloud+ Study Guide], page 236.

NEW QUESTION 148
- (Topic 4)
A cloud administrator needs to deploy a security virtual appliance in a private cloud environment, but this appliance will not be part of the standard catalog of items for other users to request. Which of the following is the BEST way to accomplish this task?

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

Answer: C

Explanation:
The correct answer is C. Import the virtual appliance into the environment and deploy it as a VM.

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

NEW QUESTION 152
- (Topic 4)
A systems administrator is working within a private cloud environment. Over time. random 4K read/write speeds on all VMS in the environment slow down until the VMS are completely unusable, with disk speeds of less than 1MBps. The administrator has gathered the information below:
• There is no correlation between the slowdown and VM/hypervisor resource utilization.
• The network is rated to 40Gbps and utilization is between 1—5%.
• The hypervisors use hundreds of NFSv3 mounts to the same storage appliance, one per VM.
• The VMS on each hypervisor become unresponsive after two weeks of uptime.
• The unresponsiveness is resolved by moving slow VMS onto a rebooted hypervisor. Which of the following solutions will MOST likely resolve this issue?

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

Answer: D

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

NEW QUESTION 156
- (Topic 4)
A cloud administrator receives an email stating the following:
"Clients are receiving emails from our web application with non-encrypted links."
The administrator notices that links generated from the web application are opening in http://. Which of the following should be configured to redirect the traffic to https://?

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

Answer: C

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

NEW QUESTION 158
- (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 162
- (Topic 4)
Which of the following provides groups of compute units that can horizontally scale according to a workload?

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

Answer: C

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

NEW QUESTION 167
- (Topic 4)
A company has a web application running in an on-premises environment that needs to be migrated to the cloud. The company wants to implement a solution that maximizes scalability, availability, and security, while requiring no infrastructure administration. Which of the following services would be BEST to meet this goal?

A. A PaaS solution
B. A hybrid solution
C. An laaS solution
D. A SaaS solution

Answer: A

Explanation:
A PaaS solution, or platform as a service, is a cloud computing service that provides a complete, ready-to-use, cloud-hosted platform for developing, running, maintaining and managing applications1. A PaaS solution would meet the company’s goal of maximizing scalability, availability, and security, while requiring no infrastructure administration, because:
Scalability: A PaaS solution can automatically scale up or down the resources needed to run the application based on the demand and traffic. The company does not need to worry about provisioning or managing servers, storage, network, or load balancers23.
Availability: A PaaS solution can ensure high availability and reliability of the application by replicating it across multiple regions and zones. The company does not need to worry about backup, recovery, or failover23.
Security: A PaaS solution can provide built-in security features such as encryption, authentication, authorization, and firewall. The company does not need to worry about installing or updating security patches or software23.
No infrastructure administration: A PaaS solution can abstract away the underlying infrastructure and hardware from the company. The company only needs to focus on developing and deploying the application code and data. The PaaS provider takes care of the rest23.
A hybrid solution (B) is a cloud computing service that combines on-premises and cloud resources. It may offer some benefits such as flexibility and cost optimization, but it would not meet the company’s goal of requiring no infrastructure administration. The company would still need to manage and maintain the on-premises part of the solution4.
An IaaS solution ©, or infrastructure as a service, is a

NEW QUESTION 172
- (Topic 4)
During a security incident on an laaS platform, which of the following actions will a systems administrator most likely take as part of the containment procedure?

A. Connect to an instance for triage.
B. Add a deny rule to the network ACL.
C. Mirror the traffic to perform a traffic capture.
D. Perform a memory acquisition.

Answer: B

Explanation:
A network access control list (ACL) is a set of rules that controls the inbound and outbound traffic for a network interface or a subnet. A deny rule can be used to block or filter the traffic from a specific source or destination, such as an IP address, a port number, or a protocol. By adding a deny rule to the network ACL, a systems administrator can prevent the communication between the compromised instance and the attacker, or between the compromised instance and other instances or servers. This can help to contain the security incident and limit the potential damage or data loss. A deny rule can also be used to isolate the compromised instance for further investigation or remediation. References: CompTIA Cloud+ CV0-003 Study Guide, Chapter 5: Maintaining a Cloud Environment, page 222-223; What is a network access control list (ACL)?.

NEW QUESTION 175
- (Topic 4)
A VDI administrator is deploying 512 desktops for remote workers. Which of the following would meet the minimum number of IP addresses needed for the desktops?

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

Answer: B

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

NEW QUESTION 178
- (Topic 4)
An organization provides integration services for finance companies that use web services. A new company that sends and receives more than 100,000 transactions per second has
been integrated using the web service. The other integrated companies are now reporting slowness with regard to the integration service. Which of the following is the cause of the issue?

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

Answer: B

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

NEW QUESTION 181
- (Topic 4)
A systems administrator is troubleshooting issues with audio lag during phone conferences. When looking at the core switch, the administrator notices its buffers are consistently full, and packets are being dropped due to the large number being sent and received. There is no room in the budget for new hardware, but it is critical that the audio lag be fixed immediately. Which of the following will most likely resolve the issue?

A. Enable compression of audio traffic.
B. Configure QoS rules for VoIP traffic.
C. Verify that the gateway uplink is not saturated.
D. Add an exception to IPS for voice traffic.

Answer: B

Explanation:
Quality of Service (QoS) rules can be configured to prioritize certain types of traffic, such as voice over IP (VoIP) traffic. This can help reduce audio lag during phone conferences by ensuring that VoIP packets are delivered faster and with less delay than other types of traffic. QoS rules can be applied at different levels of the network, such as the core switch, the router, or the firewall. By configuring QoS rules for VoIP traffic, the administrator can avoid packet drops and buffer overflows that can affect the quality of the audio. References: [CompTIA Cloud+ CV0-003 Study Guide], Chapter 3, Objective 3.2: Given a scenario, troubleshoot network connectivity issues.

NEW QUESTION 183
- (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 185
- (Topic 4)
A systems administrator has verified that a physical switchport that is connected to a virtualization host is using all available bandwidth. Which of the following would best address this issue?

A. Port mirroring
B. Link aggregation
C. Spanning tree
D. Microsegmentation

Answer: B

Explanation:
Link aggregation is a technique that combines multiple physical links into a logical link that provides higher bandwidth and redundancy. Link aggregation can help address the issue of a physical switchport that is connected to a virtualization host using all available bandwidth by increasing the capacity and availability of the connection. Link aggregation can also balance the traffic load across the links and improve the fault tolerance of the network. Link aggregation can be implemented using protocols such as LACP (Link Aggregation Control Protocol) or static configuration. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 3, Objective 3.2: Given a scenario, troubleshoot network connectivity
issues.

NEW QUESTION 186
- (Topic 4)
A non-critical file on a database server was deleted and needs to be recovered. A cloud administrator must use the least disruptive restoration process to retrieve the file, as the database server cannot be stopped during the business day. Which of the following restoration methods would best accomplish this goal?

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

Answer: D

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

NEW QUESTION 188
- (Topic 4)
A cloud engineer needs to perform a database migration_ The database has a restricted SLA and cannot be offline for more than ten minutes per month The database stores 800GB of data, and the network bandwidth to the CSP is 100MBps. Which of the following is the BEST option to perform the migration?

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

Answer: B

Explanation:
The correct answer is B. Create a replica database, synchronize the data, and switch to the new instance.
This option is the best option to perform the migration because it can minimize the downtime and data loss during the migration process. A replica database is a copy of the source database that is kept in sync with the changes made to the original database. By creating a replica database in the cloud, the cloud engineer

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

NEW QUESTION 193
- (Topic 4)
A company is preparing a hypervisor environment to implement a database cluster. One of the requirements is to share the disks between the nodes of the cluster to access the same LUN. Which of the following protocols Should the company use? (Select TWO)

A. CIFS
B. FTP
C. Iscsi
D. Raid 10
E. Nfs
F. fc

Answer: CF

Explanation:
The correct answer is C and F. iSCSI and FC.
iSCSI and FC are protocols that the company can use to share the disks between the nodes of the cluster to access the same LUN. A LUN, or logical unit number, is a unique identifier for a block of storage space that can be accessed by a host system or a cluster of systems. iSCSI and FC are both block-level protocols that allow transferring data between the storage device and the host system or cluster over a network.
iSCSI stands for Internet Small Computer System Interface, which is a protocol that uses TCP/IP to send SCSI commands over an Ethernet network. iSCSI can provide a low-cost and flexible solution for sharing disks between the nodes of the cluster, as it does not require any special hardware or cables, and can use existing network infrastructure. iSCSI can also support encryption and authentication for security purposes .
FC stands for Fibre Channel, which is a protocol that uses optical fiber cables to send SCSI commands over a dedicated network. FC can provide a high-performance and reliable solution for sharing disks between the nodes of the cluster, as it offers high bandwidth, low latency, and error correction. FC can also support zoning and masking for security purposes .
CIFS, or Common Internet File System, is a file-level protocol that allows sharing files and folders over a network. CIFS does not support sharing disks or accessing LUNs at the block level.
FTP, or File Transfer Protocol, is a protocol that allows transferring files between two systems over a network. FTP does not support sharing disks or accessing LUNs at the block level.
NFS, or Network File System, is a file-level protocol that allows sharing files and folders over a network. NFS does not support sharing disks or accessing LUNs at the block level. RAID 10, or Redundant Array of Independent Disks 10, is a storage configuration that combines mirroring and striping to provide high performance and fault tolerance. RAID 10 is not a protocol that allows sharing disks or accessing LUNs over a network.

NEW QUESTION 194
- (Topic 4)
A company that performs passive vulnerability scanning at its transit VPC has detected a vulnerability related to outdated web-server software on one of its public subnets. Which of the following can the company use to verify if this is a true positive with the least effort and cost? (Select two).

A. A network-based scan
B. An agent-based scan
C. A port scan
D. A red-team exercise
E. A credentialed scan
F. A blue-team exercise
G. Unknown environment penetration testing

Answer: AE

NEW QUESTION 199
- (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 203
- (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 208
- (Topic 4)
A systems administrator is selecting the appropriate RAID level to support a private cloud with the following requirements:
. The storage array must withstand the failure of up to two drives.
. The storage array must maximize the storage capacity of its drives.
Which of the following RAID levels should the administrator implement?

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

Answer: D

Explanation:
RAID stands for Redundant Array of Independent Disks, which is a technology that combines multiple physical disks into a logical unit that provides improved performance, reliability, and storage capacity. RAID levels are different ways of organizing and distributing data across the disks in a RAID array. Each RAID level has its own advantages and disadvantages, depending on the requirements and trade-offs of the system.
RAID 6 is a RAID level that uses block-level striping with double parity. This means that data is divided into blocks and distributed across all the disks in the array, and two sets of parity information are calculated and stored on different disks. Parity is a method of error detection and correction that can reconstruct the data in case of disk failure. RAID 6 can withstand the failure of up to two disks without losing any data, which makes it suitable for a private cloud that requires high fault tolerance. RAID 6 also maximizes the storage capacity of its drives, as it only uses two disks for parity and the rest for data. The storage capacity of a RAID 6 array is equal to (n-2) x S, where n is the number of disks and S is the size of the smallest disk.
RAID 0, RAID 1, RAID 5, and RAID 10 are other RAID levels, but they do not meet the requirements of the private cloud. RAID 0 uses striping without parity, which improves performance but does not provide any redundancy or fault tolerance. RAID 0 cannot withstand any disk failure, as it would result in data loss. RAID 1 uses mirroring, which copies the same data to two or more disks. RAID 1 provides high reliability and fast read performance, but it wastes half of the storage capacity for redundancy. RAID 1 can only withstand the failure of one disk in each mirrored pair. RAID 5 uses striping with single parity, which distributes data and parity across all the disks in the array. RAID 5 provides a balance of performance, reliability, and storage capacity, but it can only withstand the failure of one disk. RAID 10 is a combination of RAID 1 and RAID 0, which creates a striped array of mirrored pairs. RAID 10 provides high performance and reliability, but it also wastes half of the storage capacity for redundancy. RAID 10 can withstand the failure of one disk in each mirrored pair, but not more than that.
For more information on RAID levels, you can refer to the following sources:
? CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 4, Storage Technologies, page 791
? Cloud+ (Plus) Certification | CompTIA IT Certifications2

NEW QUESTION 210
- (Topic 4)

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

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

Answer: A

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

NEW QUESTION 214
- (Topic 4)
A cloud engineer recently used a deployment script template to implement changes on a cloud-hosted web application. The web application communicates with a managed database on the back end. The engineer later notices the web application is no longer receiving data from the managed database. Which of the following is the most likely cause of the issue?

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

Answer: D

Explanation:
A misconfiguration in the firewall can block the communication between the web application and the managed database, preventing the web application from receiving data. A firewall is a network security device that monitors and controls incoming and outgoing network traffic based on predefined rules1. A deployment script template is a way to automate the deployment of resources and configurations in Azure Resource Manager1. If the script template contains incorrect or conflicting rules for the firewall, it can cause the issue.
References: CompTIA Cloud+ CV0-003 Exam Objectives, Objective 2.2: Given a scenario, deploy and test a cloud solution ; Use deployment scripts in templates - Azure Resource Manager1

NEW QUESTION 219
- (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 223
- (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 227
- (Topic 4)
During a security incident on an IaaS platform, which of the following actions will a systems administrator most likely take as part of the containment procedure?

A. Connect to an instance for triage.
B. Add a deny rule to the network ACL.
C. Mirror the traffic to perform a traffic capture.
D. Perform a memory acquisition.

Answer: B

Explanation:
Adding a deny rule to the network ACL is a common containment procedure for a security
incident on an IaaS platform, as it can isolate the affected instance from the rest of the network and prevent further compromise or data exfiltration. Connecting to an instance for triage, mirroring the traffic to perform a traffic capture, and performing a memory acquisition are more likely to be part of the analysis or evidence collection procedures, not the containment procedure.
References: CompTIA Cloud+ CV0-003 Exam Objectives, Objective 4.2: Given a scenario, apply security configurations and compliance controls ; Cloud Security Mitigation | Cloud Computing | CompTIA1

NEW QUESTION 230
- (Topic 3)
A company wants to move to a multicloud environment and utilize the technology that provides the most portability. Which of the following technology solutions would BEST meet the company's needs?

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

Answer: D

Explanation:
The technology that provides the most portability for a multicloud environment is containers. Containers are units of software that package an application and its dependencies into a standardized and isolated environment that can run on any platform or cloud service. Containers are lightweight, scalable, and portable, as they do not depend on the underlying infrastructure or operating system. Containers can also be managed by orchestration tools that automate the deployment, scaling, and networking of containerized applications across multiple clouds. Reference: [CompTIA Cloud+ Certification Exam Objectives], Domain 1.0 Configuration and Deployment, Objective 1.3 Given a scenario involving integration between multiple cloud environments, select an appropriate solution design.

NEW QUESTION 233
- (Topic 3)
A cloud administrator has created a new asynchronous workflow lo deploy VMs to the cloud in bulk. When the workflow is tested for a single VM, it completes successfully. However, if the workflow is used to create 50 VMs at once, the job fails. Which of the following is the MOST likely cause of the issue? (Choose two.)

A. Incorrect permissions
B. Insufficient storage
C. Billing issues with the cloud provider
D. No connectivity to the public cloud
E. Expired API token
F. Disabled autoscaling

Answer: BE

Explanation:
The most likely causes of the issue where the new asynchronous workflow fails to create 50 VMs at once in the public cloud are insufficient storage and expired API token. Insufficient storage means that there is not enough disk space available in the public cloud to accommodate all the VMs that are being created simultaneously. This could result in errors or failures during the provisioning process. Expired API token means that the authentication credential that is used by the workflow to communicate with the public cloud service has expired or become invalid. This could result in errors or failures during the API calls or requests. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 4.0 Troubleshooting, Objective 4.5 Given a scenario, troubleshoot automation/orchestration issues.

NEW QUESTION 236
- (Topic 3)
A web application has been configured to use autoscaling for provisioning and deprovisioning more VMs according to the workload. The systems administrator deployed a new CI/CD tool to automate new releases of the web application. During the night, a script was deployed and configured to be executed by the VMs during bootstrapping. Now. the autoscaling configuration is creating a new VM ever\ five minutes. Which of the following actions will MOS I like y resolve the issue?

A. Reducing the maximum threshold in the autoscaling configuration
B. Debugging the script and redeploying it
C. Changing the automation tool because it is incompatible
D. Modifying the script to shut down the VM after five minutes

Answer: B

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

NEW QUESTION 238
- (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 243
- (Topic 3)
A company has hired a security firm to perform a vulnerability assessment of its environment. In the first phase, an engineer needs to scan the network services exposed by the hosts. Which of the following will help achieve this with the LEAST privileges?

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

Answer: C

Explanation:
A network-based scan is a type of vulnerability assessment that scans the network services exposed by the hosts without requiring any credentials or agents. This type of scan will help achieve the objective of scanning the network services with the least privileges, as it does not need any access to the hosts or their internal configurations. A network-based scan can identify open ports, running services, and potential vulnerabilities on the hosts. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 2.0 Security, Objective 2.4 Given a scenario, implement security automation and orchestration in a cloud environment.

NEW QUESTION 245
- (Topic 3)
A storage administrator is reviewing the storage consumption of a SAN appliance that is running a VDI environment. Which of the following features should the administrator implement to BEST reduce the storage consumption of the SAN?

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

Answer: A

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

NEW QUESTION 246
- (Topic 3)
A cloud architect is deploying a web application that contains many large images and will be accessed on two continents. Which of the following will MOST improve the user experience while keeping costs low?

A. Implement web servers in both continents and set up a VPN between the VPCs.
B. Implement web servers on both continents and peer the VPCs.
C. Implement a CDN and offload the images to an object storage.
D. Implement a replica of the entire solution on every continent.

Answer: C

Explanation:
A CDN (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 user experience by reducing the latency, bandwidth, and load on the origin server. By offloading the images to an object storage, such as Google Cloud Storage or Amazon S3, the web application can save costs and improve performance by storing and serving the images from a CDN edge location that is closer to the user2. A CDN can also provide caching, compression, and security features for the web content3.

NEW QUESTION 251
- (Topic 3)
Over the last couple of years, the growth of a company has required a more complex DNS and DHCP environment. Which of the following should a systems administration team implement as an appropriate solution to simplify management?

A. IPAM
B. DoH
C. VLAN
D. SDN

Answer: A

Explanation:
The best solution to simplify management of a more complex DNS and DHCP environment for a company that has grown over the last couple of years is IPAM (IP address management). IPAM is a tool or service that allows centralized management and automation of DNS and DHCP functions, such as IP address allocation, reservation, release, or renewal, as well as domain name registration or resolution. IPAM can also provide monitoring, auditing, reporting, and security features for DNS and DHCP resources. 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 256
- (Topic 3)
A cloud administrator is responsible for managing a VDI environment that provides end users with access to limited applications. Which of the following should the administrator make changes to when a new application needs to be provided?

A. Application security policy
B. Application whitelisting policy
C. Application hardening policy
D. Application testing policy

Answer: B

Explanation:
An application whitelisting policy is a set of rules that specifies which applications are allowed to run on a system or a network. Application whitelisting is a security technique that can prevent unauthorized or malicious software from executing, and it is often used in VDI (Virtual Desktop Infrastructure) environments to provide end users with access to limited applications. A cloud administrator who is responsible for managing a VDI environment should make changes to the application whitelisting policy when a new application needs to be provided, as this would ensure that the new application is authorized and compatible with the VDI environment.

NEW QUESTION 259
- (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 263
- (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 264
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

CV0-003 Practice Exam Features:

* CV0-003 Questions and Answers Updated Frequently

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

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

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