

CompTIA

Exam Questions CV0-003

CompTIA Cloud+ Certification Exam



NEW QUESTION 1

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 2

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 3

- (Topic 1)

A systems administrator disabled TLS 1.0 and 1.1, as well as RC4, 3DES, and AES-128 ciphers for TLS 1.2, on a web server. A client now reports being unable to access the web server, but the administrator verifies that the server is online, the web service is running, and other users can reach the server as well. Which of the following should the administrator recommend the user do FIRST?

- A. Disable antivirus/anti-malware software
- B. Turn off the software firewall
- C. Establish a VPN tunnel between the computer and the web server
- D. Update the web browser to the latest version

Answer: D

Explanation:

Updating the web browser to the latest version is the first action that the user should do when experiencing a connection timeout error after the administrator configured a redirect from HTTP to HTTPS on the web server. Updating the web browser can ensure that it supports the latest security protocols and standards, such as TLS 1.2 or 1.3, which are required for HTTPS connections. If the web browser is outdated or incompatible with the security protocols or standards used by the web server, it may fail to establish a secure connection and result in a connection timeout error. References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8

NEW QUESTION 4

- (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 5

- (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 6

- (Topic 1)

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

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

Answer: AF

Explanation:

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

NEW QUESTION 7

- (Topic 1)

A cloud administrator has built a new private cloud environment and needs to monitor all computer, storage, and network components of the environment. Which of the following protocols would be MOST useful for this task?

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

Answer: C

Explanation:

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

NEW QUESTION 8

- (Topic 1)

A technician is working with an American company that is using cloud services to provide video-based training for its customers. Recently, due to a surge in demand, customers in Europe are experiencing latency. Which of the following services should the technician deploy to eliminate the latency issue?

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

Answer: C

Explanation:

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

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

NEW QUESTION 9

- (Topic 1)

A systems administrator wants the VMs on the hypervisor to share CPU resources on the same core when feasible. Which of the following will BEST achieve this goal?

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

Answer: E

Explanation:

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

NEW QUESTION 10

- (Topic 1)

A systems administrator notices that a piece of networking equipment is about to reach its end of support. Which of the following actions should the administrator recommend?

- A. Update the firmware
- B. Migrate the equipment to the cloud
- C. Update the OS
- D. Replace the equipment

Answer: D

Explanation:

Replacing the equipment is the best action to take when a piece of networking equipment is about to reach its end of support. End of support means that the vendor or manufacturer will no longer provide technical assistance, updates, patches, or fixes for the equipment, which can affect its functionality, performance, security, and compatibility. Replacing the equipment with a newer model that has ongoing support can prevent any issues or risks associated with using outdated equipment.

References: CompTIA Cloud+ Certification Exam Objectives, page 18, section 3.5

NEW QUESTION 10

- (Topic 1)

A company has deployed a new cloud solution and is required to meet security compliance. Which of the following will MOST likely be executed in the cloud solution to meet security requirements?

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

Answer: C

Explanation:

Vulnerability testing is a type of security testing that identifies and evaluates the weaknesses or flaws in a system or service that could be exploited by attackers. Vulnerability testing can help meet security compliance requirements when deploying a new cloud solution, as it can reveal any potential security risks or gaps in the cloud environment and provide recommendations for remediation or mitigation. Vulnerability testing can also help improve security posture and performance, as it can prevent or reduce the impact of cyberattacks, data breaches, or service disruptions.

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

NEW QUESTION 13

- (Topic 1)

An organization will be deploying a web application in a public cloud with two web servers, two database servers, and a load balancer that is accessible over a single public IP.

Taking into account the gateway for this subnet and the potential to add two more web servers, which of the following will meet the minimum IP requirement?

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

Answer: C

Explanation:

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

References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8

NEW QUESTION 15

- (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 20

- (Topic 1)

In an existing IaaS instance, it is required to deploy a single application that has different versions. Which of the following should be recommended to meet this requirement?

- A. Deploy using containers
- B. Install a Type 2 hypervisor
- C. Enable SR-IOV on the host
- D. Create snapshots

Answer: A

Explanation:

Containers are a type of deployment technology that packages an application and its dependencies into a lightweight and portable unit that can run on any platform or environment. Containers can help deploy a single application that has different versions in an existing IaaS instance, as they can isolate and run multiple versions of the same application without any conflicts or interference. Containers can also enable faster and easier deployment, scaling, and management of cloud-based applications. References: CompTIA Cloud+ Certification Exam Objectives, page 11, section 1.6

NEW QUESTION 22

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 26

- (Topic 1)

A cloud administrator is building a new VM for a network security appliance. The security appliance installer says the CPU clock speed does not meet the requirements.

Which of the following will MOST likely solve the issue?

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

Answer: A

Explanation:

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

NEW QUESTION 28

- (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 30

- (Topic 1)

A media company has made the decision to migrate a physical, internal file server to the cloud and use a web-based interface to access and manage the files. The users must be able to use their current corporate logins. Which of the following is the MOST efficient way to achieve this goal?

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

Answer: B

Explanation:

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

NEW QUESTION 31

- (Topic 1)

An IaaS provider has numerous devices and services that are commissioned and decommissioned automatically on an ongoing basis. The cloud administrator needs to implement a solution that will help reduce administrative overhead. Which of the following will accomplish this task?

- A. IPAM
- B. NAC
- C. NTP
- D. DNS

Answer: A

Explanation:

IP address management (IPAM) is a type of tool or system that automates and standardizes the allocation, tracking, and management of IP addresses in an IP network. IPAM can help reduce administrative overhead for an IaaS provider that has numerous devices and services that are commissioned and decommissioned automatically on an ongoing basis, as it can simplify and centralize the process of assigning and reclaiming IP addresses for different devices and services without manual intervention or errors. IPAM can also help optimize network performance and security, as it can monitor and report any issues or conflicts related to IP addresses. References: CompTIA Cloud+ Certification Exam Objectives, page 15, section 2.8
Reference: <https://www.infoblox.com/glossary/ipam-ip-address-management/>

NEW QUESTION 36

- (Topic 1)

A company recently subscribed to a SaaS collaboration service for its business users. The company also has an on-premises collaboration solution and would like users to have a seamless experience regardless of the collaboration solution being used. Which of the following should the administrator implement?

- A. LDAP
- B. WAF
- C. VDI
- D. SSO

Answer: D

Explanation:

Single sign-on (SSO) is a type of authentication mechanism that allows users to access multiple systems or applications with a single login credential. SSO can help users have a seamless experience regardless of the collaboration solution being used, as it can eliminate the need for multiple logins and passwords for different systems or applications. SSO can also improve user convenience, productivity, and security, as it can simplify the login process, reduce login errors, and enhance password management. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 39

- (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 40

- (Topic 1)

A VDI administrator has received reports of poor application performance. Which of the following should the administrator troubleshoot FIRST?

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

Answer: A

Explanation:

The network environment is the set of network devices, connections, protocols, and configurations that enable communication and data transfer between different systems and applications. The network environment can affect the performance of a virtual desktop infrastructure (VDI) by influencing factors such as bandwidth, latency, jitter, packet loss, and congestion. Poor network performance can result in slow or unreliable application delivery, degraded user experience, and reduced productivity.

Therefore, troubleshooting the network environment should be the first step for a VDI administrator who receives reports of poor application performance.

References: CompTIA Cloud+ Certification Exam Objectives, page 17, section 3.4

NEW QUESTION 43

- (Topic 1)

A company just successfully completed a DR test and is ready to shut down its DR site and resume normal operations.

Which of the following actions should the cloud administrator take FIRST?

- A. Initiate a failover
- B. Restore backups
- C. Configure the network
- D. Perform a failback

Answer: D

Explanation:

A failback is a process of restoring or returning a system or service to its original state or location after a failure or disaster recovery event. Performing a failback is the first action that a cloud administrator should take after successfully completing a DR test and being ready to shut down its DR site and resume normal operations, as it can ensure that all data and configurations are synchronized and consistent between the primary site and the DR site before switching back to the primary site. Performing a failback can also help minimize downtime or disruption, as it can verify that all systems or services

are functioning properly before resuming normal operations. References: CompTIA Cloud+ Certification Exam Objectives, page 10, section 1.5

NEW QUESTION 44

- (Topic 2)

A cloud administrator is managing an organization's infrastructure in a public cloud. All servers are currently located in a single virtual network with a single firewall that all traffic must pass through. Per security requirements, production, QA, and development servers should not be able to communicate directly with each other.

Which of the following should an administrator perform to comply with the security requirement?

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

Answer: A

Explanation:

These are the actions that the administrator should perform to comply with the security requirement of isolating production, QA, and development servers from each other in a public cloud environment:

? Create separate virtual networks for production, QA, and development servers: A virtual network is a logical isolation of network resources or systems within a cloud environment. Creating separate virtual networks for different types of servers can help to segregate them from each other and prevent direct communication or interference.

? Move the servers to the appropriate virtual network: Moving the servers to the appropriate virtual network can help to assign them to their respective roles and functions, as well as ensure that they follow the network policies and rules of their virtual network.

? Apply a network security group to each virtual network that denies all traffic except for the firewall: A network security group is a set of rules or policies that control and filter inbound and outbound network traffic for a virtual network or system. Applying a network security group to each virtual network that denies all traffic except for the firewall can help to enforce security and compliance by blocking any unauthorized or unwanted traffic between different types of servers, while allowing only necessary traffic through the firewall.

NEW QUESTION 49

- (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 51

- (Topic 2)

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

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

Answer: D

Explanation:

Serverless is what would be the best option to deploy a new cloud application and provision cloud services with minimal effort while reducing the tasks required for maintenance such as OS patching, VM and volume provisioning, and autoscaling configurations. Serverless is a cloud service model that provides customers with a platform to run applications or functions without having to manage or provision any underlying infrastructure or resources, such as servers, storage, network, OS, etc. Serverless can provide benefits such as:

? Minimal effort: Serverless can reduce the effort required to deploy a new cloud application and provision cloud services by automating and abstracting away all the infrastructure or resource management or provisioning tasks from customers, and allowing them to focus only on writing code or logic for their applications or functions.

? Reduced maintenance: Serverless can reduce the tasks required for maintenance by handling all the infrastructure or resource maintenance tasks for customers, such as OS patching, VM and volume provisioning, autoscaling configurations, etc., and ensuring that they are always up-to-date and optimized.

NEW QUESTION 52

- (Topic 2)

A cloud architect is reviewing four deployment options for a new application that will be hosted by a public cloud provider. The application must meet an SLA that allows for no

more than five hours of downtime annually. The cloud architect is reviewing the SLAs for the services each option will use:

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

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

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

Answer: B

Explanation:

Option B is what minimally complies with the SLA (Service Level Agreement) requirements of allowing for no more than five hours of downtime annually for a new application that will be hosted by a public cloud provider. An SLA is a contract or agreement that defines the level of service or performance that a customer expects from a provider, such as availability, reliability, scalability, security, etc. An SLA can help to measure and monitor the quality and satisfaction of service or performance, as well as identify any penalties or rewards for meeting or failing to meet the SLA. Option B minimally complies with the SLA requirements by using services that have availability percentages that are equal to or higher than 99.95%, which translates to no more than five hours of downtime annually. Option B uses services such as:

? Compute: This is a service that provides computing resources such as servers, processors, memory, etc., to run applications or functions. Option B uses compute service with availability percentage of 99.95%, which means that it guarantees to be available for 99.95% of the time in a year, and allows for no more than five hours of downtime in a year.

? Storage: This is a service that provides storage resources such as disks, volumes, files, etc., to store data or information. Option B uses storage service with availability percentage of 99.99%, which means that it guarantees to be available for 99.99% of the time in a year, and allows for no more than one hour of downtime in a year.

? Database: This is a service that provides database resources such as tables, records, queries, etc., to store and retrieve data or information. Option B uses database service with availability percentage of 99.95%, which means that it guarantees to be available for 99.95% of the time in a year, and allows for no more than five hours of downtime in a year.

NEW QUESTION 56

- (Topic 2)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 59

- (Topic 2)

A systems administrator adds servers to a round-robin, load-balanced pool, and then starts receiving reports of the website being intermittently unavailable. Which of the following is the MOST likely cause of the issue?

- A. The network is being saturated.
- B. The load balancer is being overwhelmed.
- C. New web nodes are not operational.
- D. The API version is incompatible.
- E. There are time synchronization issues.

Answer: C

Explanation:

New web nodes are not operational is the most likely cause of the issue of website being intermittently unavailable after adding servers to a round-robin, load-balanced pool. A round-robin, load-balanced pool is a method of distributing network traffic evenly and sequentially among multiple servers or nodes that provide the same service or function. A round-robin, load-balanced pool can help to improve performance, availability, and scalability of network applications or services by ensuring that no server or node is overloaded or underutilized. New web nodes are not operational if they are not configured properly or functioning correctly to provide web service or function. New web nodes are not operational can cause website being intermittently unavailable by disrupting the round-robin, load-balanced pool and creating inconsistency or unreliability in web service or function.

NEW QUESTION 62

- (Topic 2)

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

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

Answer: ADE

Explanation:

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

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

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

? Enable SSH key access only: SSH key access is a method of verifying and authenticating users or devices based on digital keys issued by a trusted authority.

SSH key access can provide more security and convenience than password authentication, as it does not require users or devices to remember or enter passwords every time they access the web server. Enabling SSH key access only can ensure that only authorized or trusted users or devices can access the web server using keys.

NEW QUESTION 66

- (Topic 2)

Which of the following definitions of serverless computing BEST explains how it is different from using VMs?

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

Answer: D

Explanation:

This is the best definition of serverless computing that explains how it is different from using VMs (Virtual Machines). Serverless computing is a cloud service model that provides customers with a platform to run applications or functions without having to manage or provision any underlying infrastructure or resources, such as servers, storage, network, OS, etc. Serverless computing is different from using VMs in the following ways:

? Serverless computing allows developers to focus on writing code and organizations to focus on business, rather than spending time and effort on managing or scaling VMs or other infrastructure components.

? Serverless computing is event-driven and pay-per-use, which means that applications or functions are executed only when triggered by a specific event or request, and customers are charged only for the resources consumed during the execution time.

? Serverless computing is more scalable and flexible than using VMs, as it can automatically adjust the capacity and performance of applications or functions according to demand or workload, without requiring any manual intervention or configuration.

NEW QUESTION 67

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 68

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 70

- (Topic 2)

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

- A. An SLA report
- B. Tagging
- C. Quotas

D. Showback

Answer: B

Explanation:

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

NEW QUESTION 73

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 74

- (Topic 2)

A systems administrator swapped a failed hard drive on a server with a RAID 5 array. During the RAID resynchronization, a second hard drive failed. Which of the following actions will make the server fully operational?

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

Answer: D

Explanation:

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

NEW QUESTION 75

- (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 76

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 80

- (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 81

- (Topic 2)

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

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

Answer: C

Explanation:

Documenting the solution and placing it in a shared knowledge base is what the administrator should do next after troubleshooting performance issues with a VDI (Virtual Desktop Infrastructure) environment, determining that the issue is GPU (Graphics Processing Unit) related, increasing the frame buffer on the virtual machines, and testing that confirms that the issue is solved and everything is now working correctly. Documenting the solution is a process of recording and describing what was done to fix or resolve an issue, such as actions, steps, methods, etc., as well as why and how it worked. Placing it in a shared knowledge base is a process of storing and organizing documented solutions in a central location or repository that can be accessed and used by others. Documenting the solution and placing it in a shared knowledge base can provide benefits such as:

- ? Learning: Documenting the solution and placing it in a shared knowledge base can help to learn from past experiences and improve skills and knowledge.
- ? Sharing: Documenting the solution and placing it in a shared knowledge base can help to share information and insights with others who may face similar issues or situations.
- ? Reusing: Documenting the solution and placing it in a shared knowledge base can help to reuse existing solutions for future issues or situations.

NEW QUESTION 85

- (Topic 2)

A systems administrator is trying to establish an RDP session from a desktop to a server in the cloud. However, the connection appears to be refused even though the VM is responding to ICMP echo requests. Which of the following should the administrator check FIRST?

- A. The firewall
- B. The subnet
- C. The gateway
- D. The services

Answer: A

Explanation:

The firewall is the first thing that the administrator should check if an RDP (Remote Desktop Protocol) session from a desktop to a server in the cloud is refused even though the VM is responding to ICMP echo requests. A firewall is a device or software that controls the incoming and outgoing network traffic based on predefined rules or policies. A firewall may block RDP connections by default or require specific ports or rules to be opened or configured.

NEW QUESTION 87

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 88

- (Topic 2)

An organization suffered a critical failure of its primary datacenter and made the decision to switch to the DR site. After one week of using the DR site, the primary datacenter is now ready to resume operations.

Which of the following is the MOST efficient way to bring the block storage in the primary datacenter up to date with the DR site?

- A. Set up replication.
- B. Copy the data across both sites.
- C. Restore incremental backups.
- D. Restore full backups.

Answer: A

Explanation:

Reference: <https://www.ibm.com/docs/en/cloud-pak-system-w3550/2.3.3?topic=system-administering-block-storage-replication>

Setting up replication is the most efficient way to bring the block storage in the primary datacenter up to date with the DR site after a critical failure. Replication is a process of copying data from one location to another in real-time or near real-time. Replication can be synchronous or asynchronous, depending on the latency and bandwidth requirements. Replication can ensure data consistency and availability across multiple sites and facilitate faster recovery.

NEW QUESTION 93

- (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 94

- (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 95

- (Topic 2)

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? (Choose two.)

- A. CIFS
- B. FTP
- C. iSCSI
- D. RAID 10
- E. NFS
- F. FC

Answer: CF

Explanation:

These are the protocols that should be used to share the disks between the nodes of a database cluster to access the same LUN (Logical Unit Number). A LUN is an identifier that represents a logical unit of storage, such as a disk, partition, volume, etc., that can be accessed by a host system or device. To share the disks between the nodes of a cluster, the following protocols can be used:

? iSCSI (Internet Small Computer System Interface): This is a protocol that allows SCSI commands to be sent over IP networks. iSCSI can enable block-level storage access over a network, which means that the host system or device can access the storage as if it were a local disk.

? FC (Fibre Channel): This is a protocol that provides high-speed and low-latency data transfer over optical fiber cables. FC can also enable block-level storage access over a network, which means that the host system or device can access the storage as if it were a local disk.

NEW QUESTION 97

- (Topic 2)

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

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

Answer: D

Explanation:

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

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

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

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

NEW QUESTION 98

- (Topic 2)

A company recently experienced a power outage that lasted 30 minutes. During this time, a whole rack of servers was inaccessible, even though the servers did not lose power.

Which of the following should be investigated FIRST?

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

Answer: C

Explanation:

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

NEW QUESTION 101

- (Topic 2)

A company needs to access the cloud administration console using its corporate identity. Which of the following actions would MOST likely meet the requirements?

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

Answer: D

Explanation:

Implementing client-based certificate authentication is what the administrator should do to access the cloud administration console using corporate identity. Client-based certificate authentication is a method of verifying and authenticating users or devices based on digital certificates issued by a trusted authority. Digital certificates are electronic documents that contain information such as identity, public key, expiration date, etc., that can be used to prove one's identity and establish secure communication over a network. Client-based certificate authentication can allow users or devices to access cloud resources or services using their corporate identity without requiring passwords or other credentials.

NEW QUESTION 106

- (Topic 2)

A systems administrator is about to deploy a new VM to a cloud environment. Which of the following will the administrator MOST likely use to select an address for the VM?

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

Answer: D

Explanation:

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

NEW QUESTION 107

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 109

- (Topic 1)

A systems administrator needs to configure SSO authentication in a hybrid cloud environment. Which of the following is the BEST technique to use?

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

Answer: B

Explanation:

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

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

NEW QUESTION 114

- (Topic 1)

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

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 115

- (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 119

- (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 120

- (Topic 1)

An IaaS application has a two-hour RTO and a four-hour RPO. The application takes one hour to back up its data or restore from a local backup file. A systems administrator is tasked with configuring the backup policy.

Which of the following should the administrator configure to achieve the application requirements with the LEAST cost?

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

Answer: B

Explanation:

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

NEW QUESTION 125

- (Topic 1)

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

This is an example of:

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

Answer: C

Explanation:

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

NEW QUESTION 130

- (Topic 1)

A cloud administrator needs to implement a mechanism to monitor the expense of the company's cloud resources. Which of the following is the BEST option to execute this task with minimal effort?

- A. Ask the cloud provider to send a daily expense report
- B. Set custom notifications for exceeding budget thresholds
- C. Use the API to collect expense information from cloud resources
- D. Implement a financial tool to monitor cloud resource expenses

Answer: B

Explanation:

Setting custom notifications for exceeding budget thresholds is the best option to execute the task of monitoring the expense of the company's cloud resources with minimal effort, as it can automate and simplify the process of tracking and alerting the cloud administrator about any overspending or wastage of cloud resources. Setting custom notifications can also help optimize the cost and performance of cloud resources, as it can enable timely and proactive actions to adjust or optimize the resource allocation or consumption based on the budget limits. References: CompTIA Cloud+ Certification Exam Objectives, page 13, section 2.5

NEW QUESTION 132

- (Topic 1)

A systems administrator is deploying a new storage array for backups. The array provides 1PB of raw disk space and uses 14TB nearline SAS drives. The solution must tolerate at least two failed drives in a single RAID set.

Which of the following RAID levels satisfies this requirement?

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

Answer: D

Explanation:

RAID 6 is a type of RAID level that uses block-level striping with two parity blocks distributed across all member disks. RAID 6 can provide redundancy and fault tolerance, as it can survive the failure of up to two disks without losing any data. RAID 6 can also support large data sets and high-capacity disks, as it can offer more usable space and better performance than other RAID levels with similar features, such as RAID 5 or RAID 10. RAID 6 is the best RAID level for a systems administrator to use when deploying a new storage array for backups that provides 1PB of raw disk space and uses 14TB nearline SAS drives and must tolerate at least two failed drives in a single RAID set. References: CompTIA Cloud+ Certification Exam Objectives, page 9, section 1.4

NEW QUESTION 136

- (Topic 1)

A cloud administrator recently noticed that a number of files stored at a SaaS provider's file-sharing service were deleted. As part of the root cause analysis, the administrator noticed the parent folder permissions were modified last week. The administrator then used a test user account and determined the permissions on the files allowed everyone to have write access.

Which of the following is the best step for the administrator to take NEXT?

- A. Identify the changes to the file-sharing service and document
- B. Acquire a third-party DLP solution to implement and manage access
- C. Test the current access permissions to the file-sharing service
- D. Define and configure the proper permissions for the file-sharing service

Answer: D

Explanation:

Permissions are rules or settings that determine what actions users can perform on files or resources in a system or service. Permissions can help control and restrict access to files or resources based on various criteria, such as user identity, role, group, or ownership. Defining and configuring the proper permissions for the file-sharing service is the best step for the administrator to take next after discovering that sales group members can access the financial application due to being part of the finance group and having write access to all files in the file-sharing service. Defining and configuring the proper permissions can prevent unauthorized or accidental access or modification of files or resources by limiting or granting access based on specific criteria. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

NEW QUESTION 139

- (Topic 1)

A systems administrator in a large enterprise needs to alter the configuration of one of the finance department's database servers.

Which of the following should the administrator perform FIRST?

- A. Capacity planning
- B. Change management
- C. Backups
- D. Patching

Answer: B

Explanation:

The SA would do the other three regardless of the need to alter configurations. In this situation, the SA would have to present the change to the CCB in order to do the alteration.

There is no clarification on whether the change management process has been gone through. Any changes, regardless of how small or big, must go through the change management process. This allows proposals to be heard by end-users, management, and possibly stockholders. From there, it will be reviewed and either approved or denied, with reasons specified. From there, the administrator(s)

can do whatever processes are necessary.

Change management is a process or procedure that defines the steps, roles, and responsibilities for implementing, documenting, and communicating any changes or updates to a system or service. Change management can help ensure that any changes or updates are done in a controlled and consistent manner, minimizing any risks or impacts to the system or service. Performing change management is the first thing that a systems administrator should do before altering the configuration of one of the finance department's database servers, as it can ensure that the change request is approved, authorized, tested, and verified before applying it to the database server. References: CompTIA Cloud+ Certification Exam Objectives, page 13, section 2.5

NEW QUESTION 144

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 148

- (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 151

- (Topic 1)

An OS administrator is reporting slow storage throughput on a few VMs in a private IaaS cloud. Performance graphs on the host show no increase in CPU or memory. However, performance graphs on the storage show a decrease of throughput in both IOPS and MBps but not much increase in latency. There is no increase in workload, and latency is stable on the NFS storage arrays that are used by those VMs.

Which of the following should be verified NEXT?

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

Answer: D

Explanation:

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

NEW QUESTION 155

- (Topic 1)

A systems administrator recently deployed a VDI solution in a cloud environment; however, users are now experiencing poor rendering performance when trying to display 3-D content on their virtual desktops, especially at peak times.

Which of the following actions will MOST likely solve this issue?

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

Answer: D

Explanation:

A vGPU profile is a configuration option that defines the amount of video RAM (vRAM) and other resources that are allocated to a virtual machine (VM) that uses a

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

NEW QUESTION 159

- (Topic 1)

A marketing team is using a SaaS-based service to send emails to large groups of potential customers. The internally managed CRM system is configured to generate a list of target customers automatically on a weekly basis, and then use that list to send emails to each customer as part of a marketing campaign. Last week, the first email campaign sent emails successfully to 3,000 potential customers. This week, the email campaign attempted to send out 50,000 emails, but only 10,000 were sent.

Which of the following is the MOST likely reason for not sending all the emails?

- A. API request limit
- B. Incorrect billing account
- C. Misconfigured auto-scaling
- D. Bandwidth limitation

Answer: A

Explanation:

An API request limit 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 request limits are often used by SaaS-based services to control the usage and traffic of their customers and prevent overloading or abuse of their resources. An API request limit can cause a failure to send all the emails if the marketing team exceeds the number of requests allowed by the SaaS-based service in a week. The service may reject or block any requests that go beyond the limit, resulting in fewer emails being sent than expected. References: CompTIA Cloud+ Certification Exam Objectives, page 13, section 2.5

Reference: <https://developers.google.com/analytics/devguides/config/mgmt/v3/limits-quotas>

NEW QUESTION 164

- (Topic 1)

A systems administrator is configuring a storage array.

Which of the following should the administrator configure to set up mirroring on this array?

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

Answer: B

Explanation:

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

NEW QUESTION 169

- (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 172

- (Topic 4)

A security analyst is investigating a recurring alert. The alert is reporting an insecure firewall configuration state after every cloud application deployment. The process of identifying the issue, requesting a fix, and waiting for the developers to manually patch the environment is being repeated multiple times. In an effort to identify the root issue, the following logs were collected:

```
Deploying template app prod. •yaml Instance DB successfully created DB keys successfully stored on vault
Instance WebApp successfully created Access rules successfully applied Access—keys successfully created
```

Which of the following options will provide a permanent fix for the issue?

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

Answer: A

Explanation:

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

NEW QUESTION 174

- (Topic 4)

As a result of an IT audit, a customer has decided to move some applications from an old legacy system to a private cloud. The current server location is remote with low bandwidth. Which of the following is the best migration strategy to use for this deployment?

- A. P2V with physical data transport
- B. P2P with remote data copy
- C. V2V with physical data transport
- D. V2P with physical data transport
- E. V2P with remote data copy

Answer: A

Explanation:

P2V stands for physical to virtual, which is the process of converting a physical server into a virtual machine. This is a common migration strategy for moving legacy systems to the cloud, as it preserves the existing configuration and data of the server. Physical data transport means using a physical device, such as a hard disk drive or a USB flash drive, to transfer the data from the source location to the destination location. This method is suitable for remote locations with low bandwidth, as it avoids the network latency and congestion that may occur with remote data copy. P2P, V2V, and V2P are other types of migration strategies, but they are not applicable for this scenario. P2P stands for physical to physical, which is the process of moving a physical server to another physical server. V2V stands for virtual to virtual, which is the process of moving a virtual machine to another virtual machine. V2P stands for virtual to physical, which is the process of converting a virtual machine into a physical server. Remote data copy means using a network connection, such as FTP or SCP, to transfer the data from the source location to the destination location. This method is suitable for locations with high bandwidth and reliable network connectivity. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 21, Cloud Migration, page 3371.

NEW QUESTION 177

- (Topic 4)

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

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

Answer: B

Explanation:

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

NEW QUESTION 182

- (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 186

- (Topic 4)

A cloud administrator must ensure all servers are in compliance with the company's security policy. Which of the following should the administrator check FIRST?

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

Answer: C

Explanation:

Hardened baselines are a set of security best practices that reduce the vulnerability of a system to exploits by reducing its attack surface¹. They are also known as security configurations or benchmarks, and they provide a standard level of system hardening for an organization²³.

Checking the hardened baselines of the servers is the first step that a cloud administrator should take to ensure compliance with the company's security policy.

This is because hardened baselines can help to:

Identify and eliminate common vulnerabilities and exposures (CVEs) that attackers can exploit¹.

Remove unnecessary or unused services, accounts, software, and ports that can increase the attack surface²³.

Apply appropriate settings and controls for encryption, authentication, authorization, firewall, and logging²³.

Streamline audits and testing by reducing complexity and providing a reliable benchmark²³.

NEW QUESTION 188

- (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 190

- (Topic 4)

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

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

Answer: B

Explanation:

One possible answer is: B. VPN

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

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

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

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

NEW QUESTION 195

- (Topic 4)

A systems administrator is attempting to gather information about services and resource utilization on VMS in a cloud environment. Which of the following will BEST accomplish this objective?

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

Answer: E

Explanation:

Performance monitoring is the process of collecting and analyzing metrics related to the performance and availability of resources in a cloud environment¹. Performance monitoring can help a systems administrator to gather information about services and resource utilization on VMs in a cloud environment by providing the following benefits²:

- ? Identify and troubleshoot performance issues and bottlenecks before they affect the end users or business operations.
- ? Optimize the resource allocation and configuration to meet the performance requirements and SLAs of the services.
- ? Plan for future capacity and scalability needs based on the historical trends and patterns of resource utilization.
- ? Compare the performance and costs of different cloud service providers, regions, and SKUs.

Some of the tools and services that can help with performance monitoring in a cloud environment are³:

- ? Azure Monitor: A comprehensive service that provides a unified view of the health, performance, and availability of your Azure resources, applications, and services. Azure Monitor collects metrics, logs, and traces from various sources and provides analysis, visualization, alerting, and automation capabilities.
- ? Azure Advisor: A personalized service that provides recommendations to optimize your Azure resources for performance, security, cost, reliability, and operational excellence. Azure Advisor analyzes your resource configuration and usage data and suggests best practices to improve your cloud environment.
- ? Azure Application Insights: A service that monitors the performance and usage of your web applications and services. Application Insights collects telemetry data such as requests, dependencies, exceptions, page views, custom events, and metrics from your application code and provides powerful analytics, diagnostics, and alerting features.
- ? Azure Log Analytics: A service that collects and analyzes data from various sources such as Azure Monitor, Azure services, VMs, containers, applications, and other cloud or on-premises systems. Log Analytics enables you to query, visualize, and correlate log data using the Kusto Query Language (KQL) and create custom dashboards and reports.

Syslog is a standard protocol for sending log messages from network devices to a central server. Syslog can help with logging and auditing activities in a cloud environment, but it does not provide performance monitoring capabilities. Therefore, option A is incorrect. SNMP (Simple Network Management Protocol) is a protocol for collecting and organizing information about managed devices on a network. SNMP can help with network management and monitoring in a cloud environment, but it does not provide comprehensive performance monitoring for VMs and services. Therefore, option B is incorrect.

CMDB (Configuration Management Database) is a database that stores information about the configuration items (CIs) in an IT environment. CMDB can help with configuration management and change management in a cloud environment, but it does not provide performance monitoring capabilities. Therefore, option C is incorrect.

Service management is a set of processes and practices that aim to deliver value to customers by providing quality services that meet their needs and expectations. Service management can help with service design, delivery, support, and improvement in a cloud environment, but it does not provide performance monitoring capabilities. Therefore, option D is incorrect.

NEW QUESTION 198

- (Topic 4)

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

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

Answer: C

Explanation:

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

NEW QUESTION 201

- (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 203

- (Topic 4)

A company would like to move all its on-premises platforms to the cloud. The company has enough skilled Linux and web-server engineers but only a couple of skilled database administrators. It also has little expertise in managing email services. Which of the following solutions would BEST match the skill sets of available personnel?

- A. Run the web servers in PaaS, and run the databases and email in SaaS.

- B. Run the web servers, databases, and email in SaaS.
- C. Run the web servers in IaaS, the databases in PaaS, and the email in SaaS.
- D. Run the web servers, databases, and email in IaaS.

Answer: C

Explanation:

To answer this question, we need to understand the different types of cloud computing models and how they suit the skill sets of the available personnel. According to Google Cloud, there are three main models for cloud computing: Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). Each model provides different levels of control, flexibility, and management over the cloud resources and services¹.

? IaaS: This model provides access to networking features, computers (virtual or on dedicated hardware), and data storage space. It gives the highest level of flexibility and management control over the IT resources and is most similar to existing IT resources that many IT departments and developers are familiar with².

? PaaS: This model provides a complete cloud platform for developing, running, and managing applications without the cost, complexity, and inflexibility of building and maintaining the underlying infrastructure. It removes the need for organizations to manage the hardware and operating systems and allows them to focus on the deployment and management of their applications².

? SaaS: This model provides a completed product that is run and managed by the service provider. It does not require any installation, maintenance, or configuration by the customers. It is typically used for end-user applications that are accessed through a web browser or a mobile app².

Based on these definitions, we can evaluate each option:

? Option A: Run the web servers in PaaS, and run the databases and email in SaaS. This option is not the best match for the skill sets of the available personnel because it does not leverage their expertise in Linux and web-server engineering. Running the web servers in PaaS means that they will have less control and customization over the web server environment and will have to rely on the service provider's platform features. Running the databases and email in SaaS means that they will not need any database administration or email management skills, but they will also have less flexibility and security over their data and communication.

? Option B: Run the web servers, databases, and email in SaaS. This option is not a good match for the skill sets of the available personnel because it does not utilize their skills at all. Running everything in SaaS means that they will have no control or responsibility over any aspect of their cloud environment and will have to depend entirely on the service provider's products. This option may be suitable for some small businesses or non-technical users who do not have any IT skills or resources, but not for a company that has skilled Linux and web-server engineers.

? Option C: Run the web servers in IaaS, the databases in PaaS, and the email in SaaS. This option is the best match for the skill sets of the available personnel because it balances their strengths and weaknesses. Running the web servers in IaaS means that they can use their Linux and web-server engineering skills to configure, manage, and optimize their web server infrastructure according to their needs. Running the databases in PaaS means that they can leverage the service provider's platform features to simplify their database development and administration tasks without having to worry about the underlying hardware and operating systems. Running the email in SaaS means that they can outsource their email services to a reliable and secure service provider without having to invest in or manage their own email infrastructure.

? Option D: Run the web servers, databases, and email in IaaS. This option is not a good match for the skill sets of the available personnel because it puts too much burden on them. Running everything in IaaS means that they will have to handle all aspects of their cloud environment, including networking, computing, storage, security, backup, scaling, patching, etc. This option may be suitable for some large enterprises or highly technical users who have full control and customization over their cloud environment, but not for a company that has only a couple of skilled database administrators and little expertise in managing email services.

Therefore, option C is the correct answer.

NEW QUESTION 206

- (Topic 4)

A systems administrator is responsible for upgrading operating systems on VMs that are hosted in a cloud environment. The systems administrator wants to ensure the VMs receive updates for as long as possible. Which of the following should the systems administrator choose?

- A. Stable
- B. Nightly
- C. LTS
- D. Canary
- E. EDR

Answer: C

Explanation:

LTS stands for Long Term Support, and it is a term that refers to a version of an operating system that receives updates and security patches for a longer period of time than other versions. LTS versions are usually more stable and reliable than other versions, and they are suitable for users who want to avoid frequent changes or compatibility issues. By choosing LTS versions for the VMs that are hosted in a cloud environment, the systems administrator can ensure that the VMs receive updates for as long as possible, and benefit from the enhanced security and performance of the operating system. LTS versions are typically released every few years, and they are supported for several years after their release. For example, Ubuntu 20.04 LTS is supported until April 2025, while Ubuntu 21.04 is supported until January 2022. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 5, Objective 5.2: Given a scenario, troubleshoot common cloud resource and service issues.

NEW QUESTION 211

- (Topic 4)

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

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

Answer: AB

Explanation:

Telnet and FTP are two services that should be disabled on a cloud server because they are insecure and vulnerable to attacks. Telnet and FTP use plain text to transmit data over the network, which means that anyone who can intercept the traffic can read or modify the data, including usernames, passwords, commands, files, etc. This can lead to data breaches, unauthorized access, or malicious actions on the server¹. Instead of Telnet and FTP, more secure alternatives should be used, such as SSH (Secure Shell) and SFTP (Secure File Transfer Protocol). SSH and SFTP use

encryption to protect the data in transit and provide authentication and integrity checks for the communication. SSH and SFTP can prevent eavesdropping, tampering, or spoofing of the data and ensure the confidentiality and privacy of the server.

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

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

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

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

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

NEW QUESTION 216

- (Topic 4)

A cloud administrator needs to verify domain ownership with a third party. The third party has provided a secret that must be added to the DNS server. Which of the following DNS records does the administrator need to update to include the secret?

- A. NS
- B. TXT
- C. AAAA
- D. SOA

Answer: B

Explanation:

TXT is a type of DNS record that can store arbitrary text data, such as a secret, a verification code, or a configuration parameter. TXT records are often used to verify domain ownership with a third party, such as a certificate authority, an email service provider, or a cloud service provider. The third party can check the TXT record of the domain and compare it with the secret they provided to confirm the identity and authority of the domain owner.

NEW QUESTION 220

- (Topic 4)

A business is looking at extending the platform for an internally developed application using microservices from various cloud vendors. The cloud architect is responsible for reviewing the proposed solution design, which covers critical requirements for the business to fulfill its requirements. Which of the following is the most likely requirement the cloud architect needs to fulfill?

- A. Software
- B. Budgetary
- C. Security
- D. Integration

Answer: D

Explanation:

Integration is the process of connecting and coordinating different components or systems to work together as a whole. In this scenario, the business is looking at extending the platform for an internally developed application using microservices from various cloud vendors. This means that the cloud architect needs to ensure that the microservices can communicate and interact with each other across different cloud platforms, as well as with the existing application. Integration is a critical requirement for the business to fulfill its requirements, as it affects the functionality, performance, and reliability of the application. Software, budgetary, and security are also important requirements, but they are not the most likely ones that the cloud architect needs to fulfill in this scenario. Software refers to the programs and applications that run on the cloud platforms. Budgetary refers to the costs and expenses associated with the cloud services and solutions. Security refers to the protection of the data and resources in the cloud environment from unauthorized access or attacks. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 22, Microservices and Cloud Integration, page 3531.

NEW QUESTION 225

- (Topic 4)

An application deployment team has observed delays in deployments and has asked the cloud administrator to evaluate the issue. Below is the result of a latency test that was conducted by the cloud administrator from offices located in the following regions:

- Asia-Pacific (APAC)
- Europe, the Middle East, and Africa (EMEA)
- Americas

Tests were conducted from each location, and the results are shown below:

Cloud regions	Latency results are in milliseconds		
	Accessing from APAC	Accessing from EMEA	Accessing from Americas
US - North	328	102	26
US - South	380	110	40
EMEA - North	186	280	60
EMEA - South	162	328	80
APAC - North	80	126	160
APAC - South	62	146	180

Which of the following locations needs to be investigated further?

- A. • Connectivity from APAC to APAC regions• Connectivity from APAC to EMEA and US - North
- B. • Connectivity from APAC to all regions• Connectivity from Americas to all regions
- C. • Connectivity from EMEA to all regions• Connectivity from APAC to APAC and EMEA regions
- D. • Connectivity from APAC to EMEA and Americas regions• Connectivity from EMEA to all regions

Answer: D

Explanation:

The latency test results show that the connectivity from APAC to EMEA and Americas regions has the highest latency values, ranging from 162 ms to 380 ms. This indicates that there is a significant delay in the network communication between these regions, which could affect the performance and availability of the cloud services. The connectivity from EMEA to all regions also has high latency values, ranging from 60 ms to 328 ms, which could also cause issues for the application deployment team. Therefore, these locations need to be investigated further to identify and resolve the root cause of the network latency.

Some possible causes of network latency are :

? Congestion: The network bandwidth may be insufficient to handle the volume of traffic, resulting in packet loss, retransmission, and queuing delays.

? Distance: The physical distance between the source and destination nodes may increase the propagation delay, which is the time it takes for a signal to travel through a medium.

? Routing: The network topology and configuration may affect the number and quality of hops that a packet takes to reach its destination, resulting in transmission and processing delays.

? Hardware: The performance and capacity of the network devices, such as routers, switches, firewalls, and servers, may affect the speed and efficiency of data processing and delivery.

Some possible solutions to reduce network latency are :

? Scaling: The network resources can be increased or optimized to handle the traffic demand, such as adding more bandwidth, load balancers, or caching servers.

? Location: The network nodes can be placed closer to each other or to the end users, such as using edge computing, content delivery networks (CDNs), or regional cloud data centers.

? Routing: The network routes can be improved or optimized to reduce the number and distance of hops, such as using direct peering, dedicated circuits, or software- defined networking (SDN).

? Hardware: The network devices can be upgraded or replaced with faster and more reliable ones, such as using solid-state drives (SSDs), fiber-optic cables, or 5G technology.

References:

What is Network Latency? | Cloudflare

How To Reduce Network Latency In 2021 - DNSstuff

NEW QUESTION 227

- (Topic 4)

Based on the shared responsibility model, which of the following solutions passes the responsibility of patching the OS to the customer?

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

Answer: C

Explanation:

IaaS stands for Infrastructure as a Service, and it is a cloud service model that provides customers with access to virtualized computing resources, such as servers, storage, and networks. In the IaaS model, the customer is responsible for patching the operating system (OS) of the virtual machines, as well as installing and managing the applications and data. The cloud service provider (CSP) is responsible for maintaining the physical infrastructure, such as the hardware, power, cooling, and security. Therefore, IaaS passes the responsibility of patching the OS to the customer, unlike PaaS, DBaaS, or SaaS, where the CSP handles the OS patching and updates. References: CompTIA Cloud+ CV0-003 Certification Study Guide, Chapter 2, Objective 2.1: Given a scenario, deploy cloud services and solutions.

NEW QUESTION 228

- (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:

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 232

- (Topic 4)

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

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

Answer: A

Explanation:

CDN, or content delivery network, is a system of distributed servers that deliver web content to users based on their geographic location, the origin of the web page, and the content delivery server¹. A CDN can improve the performance, availability, and scalability of cloud applications by caching static and dynamic content at the edge of the network, reducing the latency and bandwidth consumption between the users and the cloud servers². A CDN can also provide security features such as encryption, authentication, and DDoS protection³.

Autoscaling, SR-IOV, and containerization are other techniques that can optimize the network for cloud applications, but they are not directly related to the requirement of faster downloads for users. Autoscaling is the process of automatically adjusting the number and size of compute resources based on the demand and workload of the application. SR-IOV, or single root I/O virtualization, is a technology that allows a physical network device to be partitioned into multiple virtual devices that can be assigned to different virtual machines or containers, bypassing the hypervisor and improving the network performance and efficiency. Containerization is the process of packaging an application and its dependencies into a lightweight and portable unit that can run on any platform, providing isolation, consistency, and portability.

References:

- ? CompTIA Cloud+ CV0-003 Study Guide, Chapter 4: Network Optimization, Section 4.1: Content Delivery Networks, Page 17523
- ? CompTIA Cloud+ CV0-003 Study Guide, Chapter 4: Network Optimization, Section 4.2: Autoscaling, Page 180
- ? CompTIA Cloud+ CV0-003 Study Guide, Chapter 4: Network Optimization, Section 4.3: SR-IOV, Page 184
- ? CompTIA Cloud+ CV0-003 Study Guide, Chapter 4: Network Optimization, Section 4.4: Containerization, Page 187
- ? What is a CDN?

NEW QUESTION 237

- (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 238

- (Topic 4)

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

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

Answer: D

Explanation:

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

NEW QUESTION 240

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

CV0-003 Practice Exam Features:

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

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