Question #:1

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

• There is no correlation between the slowdown and VM/hypervisor resource utilization.

• The network is rated to 40Gbps and utilization is between 1—5%.

• The hypervisors use hundreds of NFSv3 mounts to the same storage appliance, one per VM.

• The VMS on each hypervisor become unresponsive after two weeks of uptime.

• The unresponsiveness is resolved by moving slow VMS onto a rebooted hypervisor.

Which of the following solutions will MOST likely resolve this issue?

A. Increase caching on the storage appliance.

B. Configure jumbo frames on the hypervisors and storage.

C. Increase CPU/RAM resources on affected VMS.

D. Reduce the number of NFSv3 mounts to one.

Answer: D

Explanation

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

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

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

Therefore, one of the possible solutions to this issue is to reduce the number of NFSv3 mounts to one per hypervisor, instead of one per VM. This way, the hypervisor can manage the access to the shared storage appliance more efficiently, and avoid creating too many TCP connections or UDP packets that may overload the network or the server. Reducing the number of NFSv3 mounts can also simplify the configuration and troubleshooting of the network file system.

Increasing caching on the storage appliance may improve the read performance of the NFSv3 mounts, but it will not solve the underlying issue of connection failures or timeouts. Caching may also introduce data inconsistency or corruption issues if the cache is not synchronized with the server.

Configuring jumbo frames on the hypervisors and storage may improve the network throughput and efficiency of the NFSv3 mounts, but it will not solve the underlying issue of connection failures or timeouts. Jumbo frames are larger than standard Ethernet frames, and require that all devices on the network path support them. Jumbo frames may also introduce fragmentation or compatibility issues if they are not configured properly.

Increasing CPU/RAM resources on affected VMs may improve their performance in general, but it will not solve the underlying issue of connection failures or timeouts. Increasing CPU/RAM resources may also be costly and wasteful if they are not needed for other purposes.

Question #:2

A cloud solutions architect has an environment that must only be accessed during work hours. Which of the following processes should be automated to BEST reduce cost?

A. Scaling of the environment after work hours

B. Implementing access control after work hours

C. Shutting down the environment after work hours

D. Blocking external access to the environment after work hours

Answer: C

Explanation

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

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

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

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

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

to the environment, but it does not stop or terminate the resources that are still running and consuming cloud services4.

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

Question #:3

An organization purchased new servers with GPUs for render farms. The servers have limited CPU resources.

Which of the following GPU configurations will be the MOST optimal for virtualizing this environment?

A. Dedicated

B. Shared

C. Passthrough

D. vGPU

Answer: C

Explanation

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

Question #:4

A systems administrator is trying to reduce storage consumption. Which of the following file types would benefit the MOST from compression?

A. System files

B. User backups

C. Relational database

D. Mail database

Answer: B

Explanation

User backups are the file type that would benefit the most from compression to reduce storage consumption. Compression is a process of reducing the size of data by removing redundant or unnecessary information or using algorithms to encode data more efficiently. Compression can save storage space and bandwidth, but it may also affect the quality or performance of data depending on the compression method and ratio. User backups are typically large files that contain various types of data, such as documents, images, videos, etc., that can be compressed without significant loss of quality or functionality.

Question #:5

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.

Question #:6

A systems administrator is configuring network management but is concerned about confidentiality. Which of the following should the administrator configure to address this concern?

A. SNMPv3

B. Community strings

C. IPSec tunnels

D. ACLs

Answer: A

Explanation

SNMPv3 is the protocol that the administrator should configure to address the concern about confidentiality for network management. SNMP (Simple Network Management Protocol) is a standard protocol that allows network devices and systems to exchange information and perform management tasks. SNMPv3 is the latest version of SNMP that provides security enhancements, such as authentication, encryption, and access control, to protect the confidentiality, integrity, and availability of network data.

Question #:7

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

Question #:8

A financial industry services firm was the victim of an internal data breach, and the perpetrator was a member of the company's development team. During the investigation, one of the security administrators accidentally deleted the perpetrator's user data. Even though the data is recoverable, which of the following has been violated?

A. Chain of custody

B. Evidence acquisition

C. Containment

D. Root cause analysis

Answer: A

Explanation

The chain of custody is a process that documents and preserves the integrity and authenticity of evidence from the time it is collected until it is presented in court. The chain of custody includes information such as who collected, handled, stored, or transferred the evidence, when and where it was done, and how it was done. By accidentally deleting the perpetrator's user data, the security administrator has violated the chain of custody, as the evidence has been altered or destroyed and can no longer be used in court. Reference: [CompTIA Cloud+ Certification Exam Objectives], Domain 2.0 Security, Objective 2.4 Given a scenario, implement security automation and orchestration in a cloud environment.

Question #:9

A cloud administrator needs to control the connections between a group of web servers and database servers as part of the financial application security review. Which of the following would be the BEST way to achieve this objective?

A. Create a directory security group.

B. Create a resource group.

C. Create separate VLANs.

D. Create a network security group.

Answer: D

Explanation

A network security group is a service that allows the cloud administrator to filter and control the network traffic between different resources in a cloud environment. A network security group contains security rules that specify the source, destination, protocol, port, and direction of the traffic, and whether to allow or deny it. A network security group can be associated with a subnet or a network interface in a virtual machine, and it can apply to inbound or outbound traffic. A network security group would be the best way to achieve the objective of controlling the connections between a group of web servers and database servers as part of the financial application security review, as it can provide granular and flexible control over the network access and security of the servers.

Question #:10

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

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

A. Set up affinity rules to keep web and database servers on the same hypervisor

B. Enable jumbo frames on the gateway

C. Move the web and database servers onto the same VXLAN

D. Move the servers onto thick-provisioned storage

Answer: C

Explanation

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