CV0-003 Q&As

CompTIA Cloud+ Certification

Pass CompTIA CV0-003 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/cv0-003.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by CompTIA Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

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

Correct Answer: B

Reference: https://cloud.google.com/identity/docs/verify-domain-txt

QUESTION 2

Following is a sample result from a recently completed load test.

Type	Capacity	Utilization
CPU	2	25%
Memory	4GB	20%
Storage	2TB	65%

Based on the information provided, which of the following would be the BEST recommendation?

A. Use the test result as a benchmark and document it.

B. Downsize the CPU and memory assignment.

C. Decrease assigned storage capacity.

D. Implement storage compression.

Correct Answer: D

QUESTION 3

A company has a cloud infrastructure service, and the cloud architect needs to set up a DR site.

Which of the following should be configured in between the cloud environment and the DR site?

A. Failback

B. Playbook

C. Zoning

D. Replication

Correct Answer: D

Replication is a process of copying or synchronizing data from one location to another to ensure consistency and availability. Replication can help set up a disaster recovery (DR) site for a cloud environment, as it can enable data backup and recovery in case of a failure or outage in the primary site. Replication can also improve performance and reliability, as it can reduce latency and load by distributing data across multiple sites. Replication should be configured between the cloud environment and the DR site to ensure data protection and continuity. References: CompTIA Cloud+ Certification Exam Objectives, page 10, section 1.5

QUESTION 4

A cloud solutions architect is working on a private cloud environment in which storage consumption is increasing daily, resulting in high costs. Which of the following can the architect use to provide more space without adding more capacity? (Select two).

A. Tiering

B. Deduplication

C. RAID provisioning

D. Compression

E. Flash optimization

F. NVMe

Correct Answer: BD

Explanation: B. Deduplication and D. Compression are the two options that the architect can use to provide more space without adding more capacity. Deduplication is the process of eliminating duplicate copies of data that are stored in different locations, thus reducing the storage consumption and costs1. Compression is the process of reducing the size of data by applying algorithms that remove redundant or unnecessary information, thus saving storage space and bandwidth2. Both deduplication and compression can improve the efficiency and performance of cloud storage solutions12.

QUESTION 5

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

A. Incorrect configuration in the authentication process

B. Incorrect configuration in the message queue length

C. Incorrect configuration in user access permissions

D. Incorrect configuration in the SAN storage pool

Correct Answer: B

Reference: https://www.ibm.com/docs/en/ibm-mq/9.1?topic=i-determining-problems-applications-commands-messages

QUESTION 6

A system administrator is migrating a bare-metal server to the cloud. Which of the following types of migration should the systems administrator perform to accomplish this task?

A. V2V

B. V2P

C. P2P

D. P2V

Correct Answer: D

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

QUESTION 7

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

Correct Answer: C

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 8

A cloud administrator has deployed a website and needs to improve the site security to meet requirements. The website architecture is designed to have a DBaaS in the back end and autoscaling instances in the front end using a load balancer to distribute the request. Which of the following will the cloud administrator MOST likely use?

A. An API gateway

B. An IPS/IDS

C. A reverse proxy

D. A WAF

Correct Answer: D

Reference: https://aws.amazon.com/elasticloadbalancing/applicationloadbalancer/

QUESTION 9

A systems administrator needs to implement a service to protect a web application from external attacks. The administrator must have session-based granular control of all HTTP traffic. Which of the following should the administrator configure?

A. IDS

B. WAF

C. DLP

D. NAC

Correct Answer: B

A web application firewall (WAF) is a type of firewall that monitors and filters HTTP traffic to and from a web application. It can detect and block malicious requests, such as SQL injection, cross-site scripting, or denial-of-service attacks. It can also provide session-based granular control of HTTP traffic, such as allowing or denying access based on user identity, location, or behavior. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 2.0 Security, Objective 2.2 Given a scenario, implement appropriate network security controls for a cloud environment.

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

QUESTION 10

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

A. Create an empty VM, import the hard disk of the virtual appliance, and configure the CPU and memory.

B. Acquire the build scripts from the vendor and recreate the appliance using the baseline templates.

C. Import the virtual appliance into the environment and deploy it as a VM.

D. Convert the virtual appliance to a template and deploy a new VM using the template.

Correct Answer: D

Reference: https://www.techtarget.com/searchcloudcomputing/definition/private-cloud

QUESTION 11

A company is planning its cloud architecture and wants to use a VPC for each of its three products per environment in two regions, totaling 18 VPCs. The products have interdependences, consuming services between VPCs. Which of the following should the cloud architect use to connect all the VPCs?

A. MPLS connections

B. VPC peering

C. Hub and spoke

D. VPN connections

Correct Answer: C

The best way to connect all the VPCs for the company that is planning its cloud architecture and wants to use a VPC for each of its three products per environment in two regions, totaling 18 VPCs, is to use a hub and spoke model. A hub and spoke model is a networking model that uses a central hub VPC that connects to multiple spoke VPCs that host the products or workloads. The hub VPC can provide common services and security policies for the spoke VPCs, such as network virtual appliances, DNS servers, firewalls, or VPN gateways. The spoke VPCs can communicate with each other through the hub VPC, using private IP addresses or peering connections. A hub and spoke model can simplify the management and scalability of the network, as well as reduce the costs and complexity of peering multiple VPCs directly. Reference: Hub-and-spoke network topology - Cloud Adoption Framework

QUESTION 12

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.

Correct Answer: D

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.

QUESTION 13

A cloud administrator is monitoring a database system and notices an unusual increase in the read operations, which is causing a heavy load in the system. The system is using a relational database and is running in a VM. Which of the following should the administrator do to resolve the issue with minimal architectural changes?

A. Migrate the relational database to a NoSQL database

B. Use a cache system to store reading operations

C. Create a secondary standby database instance

D. Implement the database system using a DBaaS

Correct Answer: B

The best way to resolve the issue of an unusual increase in the read operations that is causing a heavy load in the system that is using a relational database and is running in a VM is to use a cache system to store reading operations. A cache system is a type of storage system that temporarily stores frequently accessed or recently used data in memory for faster retrieval. A cache system can reduce the load on the database system by serving the read requests from the cache instead of querying the database every time. Reference: [CompTIA Cloud+ Certification Exam Objectives], Domain 4.0 Troubleshooting, Objective 4.3 Given a scenario, troubleshoot capacity issues within a cloud environment.

QUESTION 14

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

Correct Answer: C

Shutting down the environment after work hours is the best process to automate to reduce cost, as it will stop incurring charges for the cloud resources that are not needed outside of work hours. Scaling, implementing access control, or blocking external access may still incur some costs for the cloud resources that are running or reserved, even if they are not fully utilized. Shutting down the environment can be automated using scripts, schedules, or triggers that can turn off or deallocate the cloud resources based on time or usage criteria12.

QUESTION 15

A systems administrator is deploying a new version of a website. The website is deployed in the cloud using a VM cluster. The administrator must then deploy the new version into one VM first. After a period of time, if there are no issues detected, a second VM will be updated. This process must continue until all the VMS are updated. Which of the following upgrade methods is being implemented?

A. Canary

B. Blue-green

C. Rolling

D. Staging

Correct Answer: C

Explanation: The upgrade method that is being implemented by the systems administrator is rolling. A rolling upgrade is a type of upgrade that applies the new version of a software or service to a subset of nodes or instances at a time, while the rest of the nodes or instances continue to run the old version. This way, the upgrade can be performed gradually and incrementally, without causing downtime or disruption to the entire system. A rolling upgrade can also help to monitor and test the new version for any issues or errors, and roll back to the old version if needed12. A canary upgrade is a type of upgrade that applies the new version of a software or service to a small and selected group of users or customers, before rolling it out to the rest of the population. This way, the upgrade can be evaluated for its performance, functionality, and feedback, and any problems or bugs can be fixed before affecting the majority of users or customers34. A blue-green upgrade is a type of upgrade that involves having two identical environments, one running the old version (blue) and one running the new version (green) of a software or service. The traffic is switched from the blue environment to the green environment once the new version is ready and tested. This way, the upgrade can be performed quickly and seamlessly, without any downtime or risk of failure. The blue environment can also serve as a backup in case of any issues with the green environment5 . A staging upgrade is a type of upgrade that involves having a separate environment that mimics the production environment, where the new version of a software or service is deployed and tested before moving it to the production environment. This way, the upgrade can be verified and validated for its compatibility, security, and quality, and any defects or errors can be resolved before affecting the live system .

CV0-003 PDF Dumps CV0-003 VCE Dumps CV0-003 Exam Questions