Free Questions for CV0-003 by vceexamstest

Shared by House on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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

Options:

A- Copy the database to an external device and ship the device to the CSP

B- Create a replica database, synchronize the data, and switch to the new instance.

C- Utilize a third-patty tool to back up and restore the data to the new database

D- use the database import/export method and copy the exported file.

Answer:

B

Explanation:

The correct answer is B. Create a replica database, synchronize the data, and switch to the new instance.

This option is the best option to perform the migration because it can minimize the downtime and data loss during the migration process. A replica database is a copy of the source database that is kept in sync with the changes made to the original database. By creating a replica database in the cloud, the cloud engineer can transfer the data incrementally and asynchronously, without affecting the availability and performance of the source database. When the replica database is fully synchronized with the source database, the cloud engineer can switch to the new instance by updating the connection settings and redirecting the traffic. This can reduce the downtime to a few minutes or seconds, depending on the complexity of the switch.

Some of the tools and services that can help create a replica database and synchronize the data are AWS Database Migration Service (AWS DMS) 1, Azure Database Migration Service 2, and Striim 3. These tools and services can support various source and target databases, such as Oracle, MySQL, PostgreSQL, SQL Server, MongoDB, etc. They can also provide features such as schema conversion, data validation, monitoring, and security.

The other options are not the best options to perform the migration because they can cause more downtime and data loss than the replica database option.

Copying the database to an external device and shipping the device to the CSP is a slow and risky option that can take days or weeks to complete. It also exposes the data to physical damage or theft during transit. Moreover, this option does not account for the changes made to the source database after copying it to the device, which can result in data inconsistency and loss.

Utilizing a third-party tool to back up and restore the data to the new database is a faster option than shipping a device, but it still requires a significant amount of downtime and bandwidth. The source database has to be offline or in read-only mode during the backup process, which can take hours or days depending on the size of the data and the network speed. The restore process also requires downtime and bandwidth, as well as compatibility checks and configuration adjustments. Additionally, this option does not account for the changes made to the source database after backing it up, which can result in data inconsistency and loss.

Using the database import/export method and copying the exported file is a similar option to using a third-party tool, but it relies on native database features rather than external tools. The import/export method involves exporting the data from the source database into a file format that can be imported into the target database. The file has to be copied over to the target database and then imported into it. This option also requires downtime and bandwidth during both export and import processes, as well as compatibility checks and configuration adjustments. Furthermore, this option does not account for the changes made to the source database after exporting it, which can result in data inconsistency and loss.

Question 2

Question Type: MultipleChoice

A Cloud administrator needs to reduce storage costs. Which of the following would BEST help the administrator reach that goal?

Options:

A- Enabling compression

B- Implementing deduplication

C- Using containers

D- Rightsizing the VMS

Answer:

B

Explanation:

The correct answer is B. Implementing deduplication would best help the administrator reduce storage costs.

Deduplication is a technique that eliminates redundant copies of data and stores only one unique instance of the dat

a. This can reduce the amount of storage space required and lower the storage costs. Deduplication can be applied at different levels, such as file-level, block-level, or object-level. Deduplication can also improve the performance and efficiency of backup and recovery operations.

Enabling compression is another technique that can reduce storage costs, but it may not be as effective as deduplication, depending on the type and amount of data. Compression reduces the size of data by applying algorithms that remove or replace redundant or unnecessary bits. Compression can also affect the quality and accessibility of the data, depending on the compression ratio and method.

Using containers and rightsizing the VMs are techniques that can reduce compute costs, but not necessarily storage costs. Containers are lightweight and portable units of software that run on a shared operating system and include only the necessary dependencies and libraries. Containers can reduce the overhead and resource consumption of virtual machines (VMs), which require a full operating system for each instance. Rightsizing the VMs means adjusting the CPU, memory, disk, and network resources of the VMs to match their workload requirements. Rightsizing the VMs can optimize their performance and utilization, and avoid overprovisioning or underprovisioning.

Question 3

Question Type: MultipleChoice

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

Options:

A- A network-based scan

B- An agent-based scan

C- A port scan

D- A red-team exercise

E- A credentialed scan

F- A blue-team exercise

G- Unknown environment penetration testing

Answer:

B, E

Explanation:

The correct answer is B and E. An agent-based scan and a credentialed scan can help verify if the vulnerability related to outdated web-server software is a true positive with the least effort and cost.

An agent-based scan is a type of vulnerability scan that uses software agents installed on the target systems to collect and report data on vulnerabilities. This method can provide more accurate and detailed results than a network-based scan, which relies on network traffic analysis and probes1. An agent-based scan can also reduce the network bandwidth and performance impact of scanning, as well as avoid triggering false alarms from intrusion detection systems2.

A credentialed scan is a type of vulnerability scan that uses valid login credentials to access the target systems and perform a more thorough and comprehensive assessment of their configuration, patch level, and vulnerabilities. A credentialed scan can identify vulnerabilities that are not visible or exploitable from the network level, such as missing updates, weak passwords, or misconfigured services3. A credentialed scan can also reduce the risk of false positives and false negatives, as well as avoid causing damage or disruption to the target systems3.

A network-based scan, a port scan, a red-team exercise, a blue-team exercise, and unknown environment penetration testing are not the best options to verify if the vulnerability is a true positive with the least effort and cost. A network-based scan and a port scan may not be able to detect the vulnerability if it is not exposed or exploitable from the network level. A red-team exercise, a blue-team exercise, and unknown environment penetration testing are more complex, time-consuming, and costly methods that involve simulating real-world attacks or defending against them. These methods are more suitable for testing the overall security posture and resilience of an organization, rather than verifying a specific vulnerability4.

Question 4

Question Type: MultipleChoice

A systems administrator is configuring a DNS server. Which of the following steps should a technician take to ensure confidentiality between the DNS server and an upstream DNS provider?

Options:

A- Enable DNSSEC.

B- Implement single sign-on.

C- Configure DOH.

D- Set up DNS over SSL.

Answer:

C

Explanation:

DNS (Domain Name System) is a service that translates human-friendly domain names into IP addresses that can be used to communicate over the Internet1. However, DNS queries and responses are usually sent in plain text, which means that anyone who can intercept the network traffic can see the domain names that the users are requesting.This poses a threat to the confidentiality and privacy of the users and their online activities2.

To ensure confidentiality between the DNS server and an upstream DNS provider, a technician should configure DOH (DNS over HTTPS).DOH is a protocol that encrypts DNS queries and responses using HTTPS (Hypertext Transfer Protocol Secure), which is a secure version of HTTP that uses SSL/TLS (Secure Sockets Layer/Transport Layer Security) to protect the data in transit3.By using DOH, the technician can prevent eavesdropping, tampering, or spoofing of DNS traffic by malicious actors3.

The other options are not the best steps to ensure confidentiality between the DNS server and an upstream DNS provider:

Option A: Enable DNSSEC (DNS Security Extensions). DNSSEC is a set of extensions that add digital signatures to DNS records, which can be used to verify the authenticity and integrity of the DNS dat

a. DNSSEC can prevent DNS cache poisoning attacks, where an attacker inserts false DNS records into a DNS server's cache, redirecting users to malicious websites.However, DNSSEC does not encrypt or hide the DNS queries and responses, so it does not provide confidentiality for DNS traffic2.

Option B: Implement single sign-on (SSO). SSO is a mechanism that allows users to access multiple services or applications with one set of credentials, such as a username and password. SSO can simplify the authentication process and reduce the risk of password compromise or phishing attacks. However, SSO does not affect the communication between the DNS server and an upstream DNS provider, so it does not provide confidentiality for DNS traffic.

Option D: Set up DNS over SSL (DNS over Secure Sockets Layer). This option is not a valid protocol for securing DNS traffic. SSL is a deprecated protocol that has been replaced by TLS (Transport Layer Security), which is more secure and robust. The correct protocol for encrypting DNS traffic using SSL/TLS is DOH (DNS over HTTPS), as explained above.

Question 5

Question Type: MultipleChoice

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?

Options:

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 laaS, the databases in PaaS, and the email in SaaS.

D- Run the web servers, databases, and email in laaS.

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 toGoogle 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.

Question 6

Question Type: MultipleChoice

A new development team requires workstations hosted in a PaaS to develop a new website. Members of the team also require remote access to the workstations using their corporate email addresses. Which of the following solutions will BEST meet these requirements? (Select TWO).

Options:

A- Deploy new virtual machines.

B- Configure email account replication.

C- Integrate identity services.

D- Implement a VDI solution.

E- Migrate local VHD workstations.

F- Create a new directory service.

Answer:

A, C

Explanation:

A Platform-as-a-Service (PaaS) is a cloud computing model that provides customers a complete cloud platform---hardware, software, and infrastructure---for developing, running, and managing applications without the cost, complexity, and inflexibility that often comes with building and maintaining that platform on-premises1.

To develop a new website using a PaaS, the development team needs to deploy new virtual machines (VMs) on the cloud platform. VMs are software emulations of physical computers that can run different operating systems and applications.By deploying new VMs, the development team can create a scalable and flexible environment for their website project, without having to invest in or manage physical hardware2.

To enable remote access to the workstations using their corporate email addresses, the development team needs to integrate identity services on the cloud platform. Identity services are services that provide authentication, authorization, and identity management for users and devices accessing cloud resources.By integrating identity services, the development team can use their corporate email addresses as single sign-on (SSO) credentials to access their workstations from any device and location, while ensuring security and compliance3.

The other options are not the best solutions for these requirements:

Configuring email account replication is not necessary for remote access to the workstations. Email account replication is a process of synchronizing email accounts across different servers or locations.It can provide backup and redundancy for email services, but it does not provide authentication or identity management for remote access4.

Implementing a Virtual Desktop Infrastructure (VDI) solution is not a PaaS solution. VDI is a technology that allows users to access virtual desktops hosted on a centralized server.VDI can provide remote access to desktop environments, but it requires additional hardware, software, and management costs that are not included in a PaaS model5.

Migrating local VHD workstations is not a PaaS solution. VHD stands for Virtual Hard Disk, which is a file format that represents a virtual hard disk drive. Migrating local VHD workstations means moving the virtual hard disk files from local storage to cloud storage.This can

provide backup and portability for the workstations, but it does not provide a complete cloud platform for developing and running applications6.

Creating a new directory service is not necessary for remote access to the workstations. A directory service is a service that stores and organizes information about users, devices, and resources on a network. Creating a new directory service means setting up a new database and schema for storing this information. This can provide identity management and access control for the network, but it does not provide authentication or SSO for remote access.

Question 7

Question Type: MultipleChoice

An organization has a web-server farm. Which of the following solutions should be implemented to obtain efficient distribution of requests to theservers?

Options:

A- A clustered web server infrastructure

B- A load-balancing appliance

C- A containerized application

D- Distribution of web servers across different regions and zones

Answer:

B

Explanation:

A web-server farm is a group of web servers that work together to provide high availability and scalability for web applications1. A web-server farm can handle a large number of requests from clients by distributing the workload among the servers.

A load-balancing appliance is a device that sits between the clients and the web servers and distributes the incoming requests to the servers based on a predefined algorithm2.A load-balancing appliance can improve the performance, reliability, and security of a web-server farm by providing the following benefits2:

Load balancing: It ensures that no single server is overloaded and that all servers are utilized efficiently.

Failover: It detects and redirects requests from failed or unavailable servers to healthy ones.

Health monitoring: It periodically checks the status and performance of the servers and reports any issues or anomalies.

SSL offloading: It terminates the SSL connections from the clients and forwards the requests to the servers in plain text, reducing the encryption overhead on the servers.

Caching: It stores frequently requested content in its memory and serves it to the clients, reducing the network traffic and load on the servers.

A clustered web server infrastructure is a configuration where multiple web servers are connected to a shared storage device and appear as a single logical server to the clients3. A clustered web server infrastructure can provide high availability and fault tolerance for web applications, but it does not provide load balancing or scalability. Therefore, option A is incorrect.

A containerized application is an application that is packaged with its dependencies and runtime environment in a lightweight and portable unit called a container. A containerized application can run on any platform that supports container technology, such as Docker or Kubernetes. A containerized application can provide portability, consistency, and isolation for web applications, but it does not provide load balancing or scalability by itself. Therefore, option C is incorrect.

Distribution of web servers across different regions and zones is a strategy that involves deploying web servers in multiple geographic locations to serve clients from different areas. Distribution of web servers across different regions and zones can provide low latency, high availability, and disaster recovery for web applications, but it does not provide load balancing or scalability within each region or zone. Therefore, option D is incorrect.