Free Questions for SAA-C03

Shared by Sellers on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company copies 200 TB of data from a recent ocean survey onto AWS Snowball Edge Storage Optimized devices. The company has a high performance computing (HPC) cluster that is hosted on AWS to look for oil and gas deposits. A solutions architect must provide the cluster with consistent sub-millisecond latency and high-throughput access to the data on the Snowball Edge Storage Optimized devices. The company is sending the devices back to AWS.

Which solution will meet these requirements?

Options:

A- Create an Amazon S3 bucket. Import the data into the S3 bucket. Configure an AWS Storage Gateway file gateway to use the S3 bucket. Access the file gateway from the HPC cluster instances.
B- Create an Amazon S3 bucket. Import the data into the S3 bucket. Configure an Amazon FSx for Lustre file system, and integrate it with the S3 bucket. Access the FSx for Lustre file system from the HPC cluster instances.
C- Create an Amazon S3 bucket and an Amazon Elastic File System (Amazon EFS) file system. Import the data into the S3 bucket. Copy the data from the S3 bucket to the EFS file system. Access the EFS file system from the HPC cluster instances.
D- Create an Amazon FSx for Lustre file system. Import the data directly into the FSx for Lustre file system. Access the FSx for Lustre file system from the HPC cluster instances.

Answer:

B

Explanation:

To provide the HPC cluster with consistent sub-millisecond latency and high-throughput access to the data on the Snowball Edge Storage Optimized devices, a solutions architect should configure an Amazon FSx for Lustre file system, and integrate it with an Amazon S3 bucket. This solution has the following benefits:

It allows the HPC cluster to access the data on the Snowball Edge devices using a POSIX-compliant file system that is optimized for fast processing of large datasets1.

It enables the data to be imported from the Snowball Edge devices into the S3 bucket using the AWS Snow Family Console or the AWS CLI2.The data can then be accessed from the FSx for Lustre file system using the S3 integration feature3.

It supports high availability and durability of the data, as the FSx for Lustre file system can automatically copy the data to and from the S3 bucket3.The data can also be accessed from other AWS services or applications using the S3 API4.

1: https://aws.amazon.com/fsx/lustre/

2: https://docs.aws.amazon.com/snowball/latest/developer-guide/using-adapter.html

3: https://docs.aws.amazon.com/fsx/latest/LustreGuide/create-fs-linked-data-repo.html

4: https://docs.aws.amazon.com/fsx/latest/LustreGuide/export-data-repo.html

Question 2

Question Type: MultipleChoice

A solutions architect is designing a customer-facing application for a company. The application's database will have a clearly defined access pattern throughout the year and will have a variable number of reads and writes that depend on the time of year. The company must retain audit records for the database for 7 days. The recovery point objective (RPO) must be less than 5 hours.

Which solution meets these requirements?

Options:

A- Use Amazon DynamoDB with auto scaling Use on-demand backups and Amazon DynamoDB Streams
B- Use Amazon Redshift. Configure concurrency scaling. Activate audit logging. Perform database snapshots every 4 hours.
C- Use Amazon RDS with Provisioned IOPS Activate the database auditing parameter Perform database snapshots every 5 hours
D- Use Amazon Aurora MySQL with auto scaling. Activate the database auditing parameter

Answer:

A

Explanation:

This solution meets the requirements of a customer-facing application that has a clearly defined access pattern throughout the year and a variable number of reads and writes that depend on

the time of year. Amazon DynamoDB is a fully managed NoSQL database service that can handle any level of request traffic and data size. DynamoDB auto scaling can automatically adjust the provisioned read and write capacity based on the actual workload. DynamoDB on-demand backups can create full backups of the tables for data protection and archival purposes. DynamoDB Streams can capture a time-ordered sequence of item-level modifications in the tables for audit purposes.

Option B is incorrect because Amazon Redshift is a data warehouse service that is designed for analytical workloads, not for customer-facing applications. Option C is incorrect because Amazon RDS with Provisioned IOPS can provide consistent performance for relational databases, but it may not be able to handle unpredictable spikes in traffic and data size. Option D is incorrect because Amazon Aurora MySQL with auto scaling can provide high performance and availability for relational databases, but it does not support audit logging as a parameter.

https://aws.amazon.com/dynamodb/

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/BackupRestore.html

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.html

Question 3

Question Type: MultipleChoice

A company built an application with Docker containers and needs to run the application in the AWS Cloud The company wants to use a managed sen/ice to host the application

The solution must scale in and out appropriately according to demand on the individual container services The solution also must not result in additional operational overhead or infrastructure to manage

Which solutions will meet these requirements? (Select TWO)

Options:

A- Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate.
B- Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate.
C- Provision an Amazon API Gateway API Connect the API to AWS Lambda to run the containers.
D- Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes.
E- Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes.

Answer:

A, B

Explanation:

These options are the best solutions because they allow the company to run the application with Docker containers in the AWS Cloud using a managed service that scales automatically and does not require any infrastructure to manage. By using AWS Fargate, the company can launch and run containers without having to provision, configure, or scale clusters of EC2 instances. Fargate allocates the right amount of compute resources for each container and scales them up or down as needed. By using Amazon ECS or Amazon EKS, the company can choose the container orchestration platform that suits its needs. Amazon ECS is a fully managed service that integrates with other AWS services and simplifies the deployment and management of containers. Amazon EKS is a managed service that runs Kubernetes on AWS and provides compatibility with existing Kubernetes tools and plugins.

C . Provision an Amazon API Gateway API Connect the API to AWS Lambda to run the containers. This option is not feasible because AWS Lambda does not support running Docker containers directly. Lambda functions are executed in a sandboxed environment that is isolated from other functions and resources. To run Docker containers on Lambda, the company would need to use a custom runtime or a wrapper library that emulates the Docker API, which can introduce additional complexity and overhead.

D . Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes. This option is not optimal because it requires the company to manage the EC2 instances that host the containers. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs.

E . Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes. This option is not ideal because it requires the company to manage the EC2 instances that host the containers. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs.

1AWS Fargate - Amazon Web Services

2Amazon Elastic Container Service - Amazon Web Services

3Amazon Elastic Kubernetes Service - Amazon Web Services

4AWS Lambda FAQs - Amazon Web Services

Question 4

Question Type: MultipleChoice

A pharmaceutical company is developing a new drug. The volume of data that the company generates has grown exponentially over the past few months. The company's researchers regularly require a subset of the entire dataset to be immediately available with minimal lag. However the entire dataset does not need to be accessed on a daily basis. All the data currently resides in on-premises storage arrays, and the company wants to reduce ongoing capital expenses.

Which storage solution should a solutions architect recommend to meet these requirements?

Options:

A- Run AWS DataSync as a scheduled cron job to migrate the data to an Amazon S3 bucket on an ongoing basis.
B- Deploy an AWS Storage Gateway file gateway with an Amazon S3 bucket as the target storage Migrate the data to the Storage Gateway appliance.
C- Deploy an AWS Storage Gateway volume gateway with cached volumes with an Amazon S3 bucket as the target storage. Migrate the data to the Storage Gateway appliance.
D- Configure an AWS Site-to-Site VPN connection from the on-premises environment to AWS. Migrate data to an Amazon Elastic File System (Amazon EFS) file system.

Answer:

C

Explanation:

AWS Storage Gateway is a hybrid cloud storage service that allows you to seamlessly integrate your on-premises applications with AWS cloud storage. Volume Gateway is a type of Storage Gateway that presents cloud-backed iSCSI block storage volumes to your on-premises applications. Volume Gateway operates in either cache mode or stored mode. In cache mode, your primary data is stored in Amazon S3, while retaining your frequently accessed data locally in the cache for low latency access. In stored mode, your primary data is stored locally and your entire dataset is available for low latency access on premises while also asynchronously getting backed up to Amazon S3.

For the pharmaceutical company's use case, cache mode is the most suitable option, as it meets the following requirements:

It reduces the need to scale the on-premises storage infrastructure, as most of the data is stored in Amazon S3, which is scalable, durable, and cost-effective.

It provides low latency access to the subset of the data that the researchers regularly require, as it is cached locally in the Storage Gateway appliance.

It does not require the entire dataset to be accessed on a daily basis, as it is stored in Amazon S3 and can be retrieved on demand.

It offers flexible data protection and recovery options, as it allows taking point-in-time copies of the volumes using AWS Backup, which are stored in AWS as Amazon EBS snapshots.

Therefore, the solutions architect should recommend deploying an AWS Storage Gateway volume gateway with cached volumes with an Amazon S3 bucket as the target storage and migrating the data to the Storage Gateway appliance.

Volume Gateway | Amazon Web Services

How Volume Gateway works (architecture) - AWS Storage Gateway

AWS Storage Volume Gateway - Cached volumes - Stack Overflow

Question 5

Question Type: MultipleChoice

A financial services company wants to shut down two data centers and migrate more than 100 TB of data to AWS. The data has an intricate directory structure with millions of small files stored in deep hierarchies of subfolders. Most of the data is unstructured, and the company's file storage consists of SMB-based storage types from multiple vendors. The company does not want to change its applications to access the data after migration.

What should a solutions architect do to meet these requirements with the LEAST operational overhead?

Options:

A- Use AWS Direct Connect to migrate the data to Amazon S3.
B- Use AWS DataSync to migrate the data to Amazon FSx for Lustre.
C- Use AWS DataSync to migrate the data to Amazon FSx for Windows File Server.
D- Use AWS Direct Connect to migrate the data on-premises file storage to an AWS Storage Gateway volume gateway.

Answer:

C

Explanation:

AWS DataSync is a data transfer service that simplifies, automates, and accelerates moving data between on-premises storage systems and AWS storage services over the internet or AWS Direct Connect1. AWS DataSync can transfer data to Amazon FSx for Windows File Server, which is a fully managed file system that is accessible over the industry-standard Server Message Block (SMB) protocol.Amazon FSx for Windows File Server is built on Windows Server, delivering a wide range of administrative features such as user quotas, end-user file restore, and Microsoft Active Directory (AD) integration2. This solution meets the requirements of the question because:

It can migrate more than 100 TB of data to AWS within a reasonable time frame, as AWS DataSync is optimized for high-speed and efficient data transfer1.

It can preserve the intricate directory structure and the millions of small files stored in deep hierarchies of subfolders, as AWS DataSync can handle complex file structures and metadata, such as file names, permissions, and timestamps1.

It can avoid changing the applications to access the data after migration, as Amazon FSx for Windows File Server supports the same SMB protocol and Windows Server features that the company's on-premises file storage uses2.

It can reduce the operational overhead, as AWS DataSync and Amazon FSx for Windows File Server are fully managed services that handle the tasks of setting up, configuring, and maintaining the data transfer and the file system12.

Question 6

Question Type: MultipleChoice

A company has 150 TB of archived image data stored on-premises that needs to be moved to the AWS Cloud within the next month. The company's current network connection allows up to 100 Mbps uploads for this purpose during the night only.

What is the MOST cost-effective mechanism to move this data and meet the migration deadline?

Options:

A- Use AWS Snowmobile to ship the data to AWS.
B- Order multiple AWS Snowball devices to ship the data to AWS.
C- Enable Amazon S3 Transfer Acceleration and securely upload the data.
D- Create an Amazon S3 VPC endpoint and establish a VPN to upload the data.

Answer:

B

Explanation:

AWS Snowball is a petabyte-scale data transport service that uses secure devices to transfer large amounts of data into and out of the AWS Cloud. Snowball addresses common challenges with large-scale data transfers including high network costs, long transfer times, and security concerns. AWS Snowball can transfer up to 80 TB of data per device, and multiple devices can be used in parallel to meet the migration deadline. AWS Snowball is more cost-effective than AWS Snowmobile, which is designed for exabyte-scale data transfers, or Amazon S3 Transfer Acceleration, which is optimized for fast transfers over long distances. Amazon S3 VPC endpoint does not increase the upload speed, but only provides a secure and private connection between the VPC and S3.Reference:AWS Snowball,AWS Snowmobile,Amazon S3 Transfer Acceleration,Amazon S3 VPC endpoint

Question 7

Question Type: MultipleChoice

A business's backup data totals 700 terabytes (TB) and is kept in network attached storage (NAS) at its data center. This backup data must be available in the event of occasional regulatory inquiries and preserved for a period of seven years. The organization has chosen to relocate its backup data from its on-premises data center to Amazon Web Services (AWS). Within one month, the migration must be completed. The company's public internet connection provides 500 Mbps of dedicated capacity for data transport.

What should a solutions architect do to ensure that data is migrated and stored at the LOWEST possible cost?

Options:

A- Order AWS Snowball devices to transfer the data. Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
B- Deploy a VPN connection between the data center and Amazon VPC. Use the AWS CLI to copy the data from on premises to Amazon S3 Glacier.
C- Provision a 500 Mbps AWS Direct Connect connection and transfer the data to Amazon S3. Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
D- Use AWS DataSync to transfer the data and deploy a DataSync agent on premises. Use the DataSync task to copy files from the on-premises NAS storage to Amazon S3 Glacier.

Answer:

A

Explanation:

https://www.omnicalculator.com/other/data-transfer

Question 8

Question Type: MultipleChoice

A solutions architect creates a VPC that includes two public subnets and two private subnets. A corporate security mandate requires the solutions architect to launch all Amazon EC2 instances in a private subnet. However, when the solutions architect launches an EC2 instance that runs a web server on ports 80 and 443 in a private subnet, no external internet traffic can connect to the server.

What should the solutions architect do to resolve this issue?

Options:

A- Attach the EC2 instance to an Auto Scaling group in a private subnet. Ensure that the DNS record for the website resolves to the Auto Scaling group identifier.
B- Provision an internet-facing Application Load Balancer (ALB) in a public subnet. Add the EC2 instance to the target group that is associated with the ALB. Ensure that the DNS record for the website resolves to the ALB.
C- Launch a NAT gateway in a private subnet. Update the route table for the private subnets to add a default route to the NAT gateway. Attach a public Elastic IP address to the NAT gateway.
D- Ensure that the security group that is attached to the EC2 instance allows HTTP traffic on port 80 and HTTPS traffic on port 443. Ensure that the DNS record for the website resolves to the public IP address of the EC2 instance.

Answer:

B

Explanation:

An Application Load Balancer (ALB) is a type of Elastic Load Balancer (ELB) that distributes incoming application traffic across multiple targets, such as EC2 instances, containers, Lambda functions, and IP addresses, in multiple Availability Zones1. An ALB can be internet-facing or internal.An internet-facing ALB has a public DNS name that clients can use to send requests over

the internet1.An internal ALB has a private DNS name that clients can use to send requests within a VPC1. This solution meets the requirements of the question because:

It allows external internet traffic to connect to the web server on ports 80 and 443, as the ALB listens for requests on these ports and forwards them to the EC2 instance in the private subnet1.

It does not violate the corporate security mandate, as the EC2 instance is launched in a private subnet and does not have a public IP address or a route to an internet gateway2.

It reduces the operational overhead, as the ALB is a fully managed service that handles the tasks of load balancing, health checking, scaling, and security1.

Question 9

Question Type: MultipleChoice

The DNS provider that hosts a company's domain name records is experiencing outages that cause service disruption for a website running on AWS. The company needs to migrate to a more resilient managed DNS service and wants the service to run on AWS.

What should a solutions architect do to rapidly migrate the DNS hosting service?

Options:

A- Create an Amazon Route 53 public hosted zone for the domain name. Import the zone file containing the domain records hosted by the previous provider
B- Create an Amazon Route 53 private hosted zone for the domain name Import the zone file containing the domain records hosted by the previous provider.
C- Create a Simple AD directory in AWS. Enable zone transfer between the DNS provider and AWS Directory Service for Microsoft Active Directory for the domain records.
D- Create an Amazon Route 53 Resolver inbound endpomt in the VPC. Specify the IP addresses that the provider's DNS will forward DNS queries to. Configure the provider's DNS to forward DNS queries for the domain to the IP addresses that are specified in the inbound endpoint.

Answer:

A

Explanation:

To migrate the DNS hosting service to a more resilient managed DNS service on AWS, the company should use Amazon Route 53, which is a highly available and scalable cloud DNS web

service. Route 53 can host public DNS records for the company's domain name and provide reliable and secure DNS resolution. To rapidly migrate the DNS hosting service, the company should create a public hosted zone for the domain name in Route 53, which is a container for the domain's DNS records. Then, the company should import the zone file containing the domain records hosted by the previous provider, which is a text file that defines the DNS records for the domain. This way, the company can quickly transfer the existing DNS records to Route 53 without manually creating them. After importing the zone file, the company should update the domain registrar to use the name servers that Route 53 assigns to the hosted zone.This will ensure that DNS queries for the domain name are routed to Route 53 and resolved by the imported records.

Question 10

Question Type: MultipleChoice

A company is running a legacy system on an Amazon EC2 instance. The application code cannot be modified, and the system cannot run on more than one instance. A solutions architect must design a resilient solution that can improve the recovery time for the system.

What should the solutions architect recommend to meet these requirements?

Options:

A- Enable termination protection for the EC2 instance.
B- Configure the EC2 instance for Multi-AZ deployment.
C- Create an Amazon CloudWatch alarm to recover the EC2 instance in case of failure.
D- Launch the EC2 instance with two Amazon Elastic Block Store (Amazon EBS) volumes that use RAID configurations for storage redundancy.

Answer:

C

Explanation:

To design a resilient solution that can improve the recovery time for the system, a solutions architect should recommend creating an Amazon CloudWatch alarm to recover the EC2 instance in case of failure. This solution has the following benefits:

It allows the EC2 instance to be automatically recovered when a system status check failure occurs, such as loss of network connectivity, loss of system power, software issues on the physical host, or hardware issues on the physical host that impact network reachability1.

It preserves the instance ID, private IP addresses, Elastic IP addresses, and all instance metadata of the original instance.A recovered instance is identical to the original instance, except for any data that is in-memory, which is lost during the recovery process1.

It does not require any modification of the application code or the EC2 instance configuration.The solutions architect can create a CloudWatch alarm using the AWS Management Console, the AWS CLI, or the CloudWatch API2.

1: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html

2: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html#ec2-instance-recover-create-alarm

To Get Premium Files for SAA-C03 Visit
https://www.p2pexams.com/products/saa-c03

For More Free Questions Visit
https://www.p2pexams.com/amazon/pdf/saa-c03

