



Free Questions for [SAA-C03](#) by [vceexamstest](#)

Shared by [Hammond](#) on [09-08-2024](#)

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)

Question 1

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 2

Question Type: MultipleChoice

A company runs a highly available web application on Amazon EC2 instances behind an Application Load Balancer. The company uses Amazon CloudWatch metrics.

As the traffic to the web application increases, some EC2 instances become overloaded with many outstanding requests. The CloudWatch metrics show that the number of requests processed and the time to receive the responses from some EC2 instances are both higher compared to other EC2 instances. The company does not want new requests to be forwarded to the EC2 instances that are already overloaded.

Which solution will meet these requirements?

Options:

- A- Use the round robin routing algorithm based on the RequestCountPerTarget and Active Connection Count CloudWatch metrics.
- B- Use the least outstanding requests algorithm based on the RequestCountPerTarget and ActiveConnectionCount CloudWatch metrics.
- C- Use the round robin routing algorithm based on the RequestCount and TargetResponseTime CloudWatch metrics.
- D- Use the least outstanding requests algorithm based on the RequestCount and TargetResponseTime CloudWatch metrics.

Answer:

D

Explanation:

The least outstanding requests (LOR) algorithm is a load balancing algorithm that distributes incoming requests to the target with the fewest outstanding requests. This helps to avoid overloading any single target and improves the overall performance and availability of the web application. The LOR algorithm can use the RequestCount and TargetResponseTime CloudWatch metrics to determine the number of outstanding requests and the response time of each target. These metrics measure the number of requests processed by each target and the time elapsed after the request leaves the load balancer until a response from the target is received by the load balancer, respectively. By using these metrics, the LOR algorithm can route new requests to the targets that are less busy and more responsive, and avoid sending requests to the targets that are already overloaded or slow. This solution meets the requirements of the company.

[Application Load Balancer now supports Least Outstanding Requests algorithm for load balancing requests](#)

Target groups for your Application Load Balancers

Elastic Load Balancing - Application Load Balancers

Question 3

Question Type: MultipleChoice

A company wants to analyze and generate reports to track the usage of its mobile app. The app is popular and has a global user base. The company uses a custom report building program to analyze application usage.

The program generates multiple reports during the last week of each month. The program takes less than 10 minutes to produce each report. The company rarely uses the program to generate reports outside of the last week of each month. The company wants to generate reports in the least amount of time when the reports are requested.

Which solution will meet these requirements MOST cost-effectively?

Options:

A- Run the program by using Amazon EC2 On-Demand Instances. Create an Amazon EventBridge rule to start the EC2 instances when reports are requested. Run the EC2 instances continuously during the last week of each month.

- B-** Run the program in AWS Lambda. Create an Amazon EventBridge rule to run a Lambda function when reports are requested.
- C-** Run the program in Amazon Elastic Container Service (Amazon ECS). Schedule Amazon ECS to run the program when reports are requested.
- D-** Run the program by using Amazon EC2 Spot Instances. Create an Amazon EventBridge rule to start the EC2 instances when reports are requested. Run the EC2 instances continuously during the last week of each month.

Answer:

B

Explanation:

This solution meets the requirements most cost-effectively because it leverages the serverless and event-driven capabilities of AWS Lambda and Amazon EventBridge. AWS Lambda allows you to run code without provisioning or managing servers, and you pay only for the compute time you consume. Amazon EventBridge is a serverless event bus service that lets you connect your applications with data from various sources and routes that data to targets such as AWS Lambda. By using Amazon EventBridge, you can create a rule that triggers a Lambda function to run the program when reports are requested, and you can also schedule the rule to run during the last week of each month. This way, you can generate reports in the least amount of time and pay only for the resources you use.

[AWS Lambda](#)

[Amazon EventBridge](#)

Question 4

Question Type: MultipleChoice

A company manages AWS accounts in AWS Organizations. AWS 1AM Identity Center (AWS Single Sign-On) and AWS Control Tower are configured for the accounts. The company wants to manage multiple user permissions across all the accounts.

The permissions will be used by multiple 1AM users and must be split between the developer and administrator teams. Each team requires different permissions. The company wants a solution that includes new users that are hired on both teams.

Which solution will meet these requirements with the LEAST operational overhead?

Options:

- A-** Create individual users in 1AM Identity Center (or each account. Create separate developer and administrator groups in 1AM Identity Center. Assign the users to the appropriate groups Create a custom 1AM policy for each group to set fine-grained permissions.
- B-** Create individual users in 1AM Identity Center for each account. Create separate developer and administrator groups in 1AM Identity Center. Assign the users to the appropriate groups. Attach AWS managed 1AM policies to each user as needed for fine-grained permissions.
- C-** Create individual users in 1AM Identity Center Create new developer and administrator groups in 1AM Identity Center. Create new permission sets that include the appropriate 1AM policies for each group. Assign the new groups to the appropriate accounts Assign the new permission sets to the new groups When new users are hired, add them to the appropriate group.

D- Create individual users in 1AM Identity Center. Create new permission sets that include the appropriate 1AM policies for each user. Assign the users to the appropriate accounts. Grant additional 1AM permissions to the users from within specific accounts. When new users are hired, add them to 1AM Identity Center and assign them to the accounts.

Answer:

C

Explanation:

This solution meets the requirements with the least operational overhead because it leverages the features of IAM Identity Center and AWS Control Tower to centrally manage multiple user permissions across all the accounts. By creating new groups and permission sets, the company can assign fine-grained permissions to the developer and administrator teams based on their roles and responsibilities. The permission sets are applied to the groups at the organization level, so they are automatically inherited by all the accounts in the organization. When new users are hired, the company only needs to add them to the appropriate group in IAM Identity Center, and they will automatically get the permissions assigned to that group. This simplifies the user management and reduces the manual effort of assigning permissions to each user individually.

[Managing access to AWS accounts and applications](#)

[Managing permissions sets](#)

[Managing groups](#)

Question 5

Question Type: MultipleChoice

A company wants to run its experimental workloads in the AWS Cloud. The company has a budget for cloud spending. The company's CFO is concerned about cloud spending accountability for each department. The CFO wants to receive notification when the spending threshold reaches 60% of the budget.

Which solution will meet these requirements?

Options:

- A-** Use cost allocation tags on AWS resources to label owners. Create usage budgets in AWS Budgets. Add an alert threshold to receive notification when spending exceeds 60% of the budget.
- B-** Use AWS Cost Explorer forecasts to determine resource owners. Use AWS Cost Anomaly Detection to create alert threshold notifications when spending exceeds 60% of the budget.
- C-** Use cost allocation tags on AWS resources to label owners. Use AWS Support API on AWS Trusted Advisor to create alert threshold notifications when spending exceeds 60% of the budget
- D-** Use AWS Cost Explorer forecasts to determine resource owners. Create usage budgets in AWS Budgets. Add an alert threshold to receive notification when spending exceeds 60% of the budget.

Answer:

A

Explanation:

This solution meets the requirements because it allows the company to track and manage its cloud spending by using cost allocation tags to assign costs to different departments, creating usage budgets to set spending limits, and adding alert thresholds to receive notifications when the spending reaches a certain percentage of the budget. This way, the company can monitor its experimental workloads and avoid overspending on the cloud.

[Using Cost Allocation Tags](#)

[Creating an AWS Budget](#)

[Creating an Alert for an AWS Budget](#)

Question 6

Question Type: MultipleChoice

A company's application runs on Amazon EC2 instances that are in multiple Availability Zones. The application needs to ingest real-time data from third-party applications.

The company needs a data ingestion solution that places the ingested raw data in an Amazon S3 bucket.

Which solution will meet these requirements?

Options:

- A-** Create Amazon Kinesis data streams for data ingestion. Create Amazon Kinesis Data Firehose delivery streams to consume the Kinesis data streams. Specify the S3 bucket as the destination of the delivery streams.
- B-** Create database migration tasks in AWS Database Migration Service (AWS DMS). Specify replication instances of the EC2 instances as the source endpoints. Specify the S3 bucket as the target endpoint. Set the migration type to migrate existing data and replicate ongoing changes.
- C-** Create and configure AWS DataSync agents on the EC2 instances. Configure DataSync tasks to transfer data from the EC2 instances to the S3 bucket.
- D-** Create an AWS Direct Connect connection to the application for data ingestion. Create Amazon Kinesis Data Firehose delivery streams to consume direct PUT operations from the application. Specify the S3 bucket as the destination of the delivery streams.

Answer:

A

Explanation:

The solution that will meet the requirements is to create Amazon Kinesis data streams for data ingestion, create Amazon Kinesis Data Firehose delivery streams to consume the Kinesis data streams, and specify the S3 bucket as the destination of the delivery streams. This solution will allow the company's application to ingest real-time data from third-party applications and place the ingested raw data in an S3 bucket. Amazon Kinesis data streams are scalable and durable streams that can capture and store data from hundreds of thousands of sources. Amazon Kinesis Data Firehose is a fully managed service that can deliver streaming data to destinations such as S3, Amazon Redshift, Amazon OpenSearch Service, and Splunk. Amazon Kinesis Data Firehose can also transform and compress the data before delivering it to S3.

The other solutions are not as effective as the first one because they either do not support real-time data ingestion, do not work with third-party applications, or do not use S3 as the destination. Creating database migration tasks in AWS Database Migration Service (AWS DMS) will not support real-time data ingestion, as AWS DMS is mainly designed for migrating relational databases, not streaming data. AWS DMS also requires replication instances, source endpoints, and target endpoints to be compatible with specific database engines and versions. Creating and configuring AWS DataSync agents on the EC2 instances will not work with third-party applications, as AWS DataSync is a service that transfers data between on-premises storage systems and AWS storage services, not between applications. AWS DataSync also requires installing agents on the source or destination servers. Creating an AWS Direct Connect connection to the application for data ingestion will not use S3 as the destination, as AWS Direct Connect is a service that establishes a dedicated network connection between on-premises and AWS, not between applications and storage services. AWS Direct Connect also requires a physical connection to an AWS Direct Connect location.

[Amazon Kinesis](#)

[Amazon Kinesis Data Firehose](#)

[AWS Database Migration Service](#)

AWS DataSync

AWS Direct Connect

Question 7

Question Type: MultipleChoice

A company has an on-premises data center that is running out of storage capacity. The company wants to migrate its storage infrastructure to AWS while minimizing bandwidth costs. The solution must allow for immediate retrieval of data at no additional cost.

How can these requirements be met?

Options:

- A-** Deploy Amazon S3 Glacier Vault and enable expedited retrieval. Enable provisioned retrieval capacity for the workload.
- B-** Deploy AWS Storage Gateway using cached volumes. Use Storage Gateway to store data in Amazon S3 while retaining copies of frequently accessed data subsets locally.
- C-** Deploy AWS Storage Gateway using stored volumes to store data locally. Use Storage Gateway to asynchronously back up point-in-time snapshots of the data to Amazon S3.
- D-** Deploy AWS Direct Connect to connect with the on-premises data center. Configure AWS Storage Gateway to store data locally. Use

Storage Gateway to asynchronously back up point-in-time snapshots of the data to Amazon S3.

Answer:

B

Explanation:

The solution that will meet the requirements is to deploy AWS Storage Gateway using cached volumes and use Storage Gateway to store data in Amazon S3 while retaining copies of frequently accessed data subsets locally. This solution will allow the company to migrate its storage infrastructure to AWS while minimizing bandwidth costs, as it will only transfer data that is not cached locally. The solution will also allow for immediate retrieval of data at no additional cost, as the cached volumes will provide low-latency access to the most recently used data. The data stored in Amazon S3 will be durable, scalable, and secure.

The other solutions are not as effective as the first one because they either do not meet the requirements or introduce additional costs or complexity. Deploying Amazon S3 Glacier Vault and enabling expedited retrieval will not meet the requirements, as it will incur additional costs for both storage and retrieval. Amazon S3 Glacier is a low-cost storage service for data archiving and backup, but it has longer retrieval times than Amazon S3. Expedited retrieval is a feature that allows faster access to data, but it charges a higher fee per GB retrieved. Provisioned retrieval capacity is a feature that reserves dedicated capacity for expedited retrievals, but it also charges a monthly fee per provisioned capacity unit. Deploying AWS Storage Gateway using stored volumes to store data locally and use Storage Gateway to asynchronously back up point-in-time snapshots of the data to Amazon S3 will not meet the requirements, as it will not migrate the storage infrastructure to AWS, but only create backups. Stored volumes are volumes that store the primary data locally and back up snapshots to Amazon S3. This solution will not reduce the storage capacity needed on-premises, nor will it leverage the benefits of cloud storage. Deploying AWS Direct Connect to connect with the on-premises data center and configuring AWS Storage Gateway to store data locally and use Storage Gateway to asynchronously back up point-in-time snapshots of the data to Amazon S3 will not meet

the requirements, as it will also not migrate the storage infrastructure to AWS, but only create backups. AWS Direct Connect is a service that establishes a dedicated network connection between the on-premises data center and AWS, which can reduce network costs and increase bandwidth. However, this solution will also not reduce the storage capacity needed on-premises, nor will it leverage the benefits of cloud storage.

[AWS Storage Gateway](#)

[Cached volumes - AWS Storage Gateway](#)

[Amazon S3 Glacier](#)

[Retrieving archives from Amazon S3 Glacier vaults - Amazon Simple Storage Service](#)

[Stored volumes - AWS Storage Gateway](#)

[AWS Direct Connect](#)

Question 8

Question Type: MultipleChoice

A company needs a solution to prevent photos with unwanted content from being uploaded to the company's web application. The solution must not involve training a machine learning (ML) model. Which solution will meet these requirements?

Options:

- A-** Create and deploy a model by using Amazon SageMaker Autopilot. Create a real-time endpoint that the web application invokes when new photos are uploaded.
- B-** Create an AWS Lambda function that uses Amazon Rekognition to detect unwanted content. Create a Lambda function URL that the web application invokes when new photos are uploaded.
- C-** Create an Amazon CloudFront function that uses Amazon Comprehend to detect unwanted content. Associate the function with the web application.
- D-** Create an AWS Lambda function that uses Amazon Rekognition Video to detect unwanted content. Create a Lambda function URL that the web application invokes when new photos are uploaded.

Answer:

B

Explanation:

The solution that will meet the requirements is to create an AWS Lambda function that uses Amazon Rekognition to detect unwanted content, and create a Lambda function URL that the web application invokes when new photos are uploaded. This solution does not involve training a machine learning model, as Amazon Rekognition is a fully managed service that provides pre-trained computer vision models for image and video analysis. Amazon Rekognition can detect unwanted content such as explicit or suggestive adult content, violence, weapons, drugs, and more. By using AWS Lambda, the company can create a serverless function that can be triggered by an HTTP request from the web application. The Lambda function can use the Amazon Rekognition API to analyze the uploaded photos and

return a response indicating whether they contain unwanted content or not.

The other solutions are not as effective as the first one because they either involve training a machine learning model, do not support image analysis, or do not work with photos. Creating and deploying a model by using Amazon SageMaker Autopilot involves training a machine learning model, which is not required for the scenario. Amazon SageMaker Autopilot is a service that automatically creates, trains, and tunes the best machine learning models for classification or regression based on the data provided by the user. Creating an Amazon CloudFront function that uses Amazon Comprehend to detect unwanted content does not support image analysis, as Amazon Comprehend is a natural language processing service that analyzes text, not images. Amazon Comprehend can extract insights and relationships from text such as language, sentiment, entities, topics, and more. Creating an AWS Lambda function that uses Amazon Rekognition Video to detect unwanted content does not work with photos, as Amazon Rekognition Video is designed for analyzing video streams, not static images. Amazon Rekognition Video can detect activities, objects, faces, celebrities, text, and more in video streams.

[Amazon Rekognition](#)

[AWS Lambda](#)

[Detecting unsafe content - Amazon Rekognition](#)

[Amazon SageMaker Autopilot](#)

[Amazon Comprehend](#)

Question 9

Question Type: MultipleChoice

A manufacturing company runs its report generation application on AWS. The application generates each report in about 20 minutes. The application is built as a monolith that runs on a single Amazon EC2 instance. The application requires frequent updates to its tightly coupled modules. The application becomes complex to maintain as the company adds new features.

Each time the company patches a software module, the application experiences downtime. Report generation must restart from the beginning after any interruptions. The company wants to redesign the application so that the application can be flexible, scalable, and gradually improved. The company wants to minimize application downtime.

Which solution will meet these requirements?

Options:

- A-** Run the application on AWS Lambda as a single function with maximum provisioned concurrency.
- B-** Run the application on Amazon EC2 Spot Instances as microservices with a Spot Fleet default allocation strategy.
- C-** Run the application on Amazon Elastic Container Service (Amazon ECS) as microservices with service auto scaling.
- D-** Run the application on AWS Elastic Beanstalk as a single application environment with an all-at-once deployment strategy.

Answer:

C

Explanation:

The solution that will meet the requirements is to run the application on Amazon Elastic Container Service (Amazon ECS) as microservices with service auto scaling. This solution will allow the application to be flexible, scalable, and gradually improved, as well as minimize application downtime. By breaking down the monolithic application into microservices, the company can decouple the modules and update them independently, without affecting the whole application. By running the microservices on Amazon ECS, the company can leverage the benefits of containerization, such as portability, efficiency, and isolation. By enabling service auto scaling, the company can adjust the number of containers running for each microservice based on demand, ensuring optimal performance and cost. Amazon ECS also supports various deployment strategies, such as rolling update or blue/green deployment, that can reduce or eliminate downtime during updates.

The other solutions are not as effective as the first one because they either do not meet the requirements or introduce new challenges. Running the application on AWS Lambda as a single function with maximum provisioned concurrency will not meet the requirements, as it will not break down the monolith into microservices, nor will it reduce the complexity of maintenance. Lambda functions are also limited by execution time (15 minutes), memory size (10 GB), and concurrency quotas, which may not be sufficient for the report generation application. Running the application on Amazon EC2 Spot Instances as microservices with a Spot Fleet default allocation strategy will not meet the requirements, as it will introduce the risk of interruptions due to spot price fluctuations. Spot Instances are not guaranteed to be available or stable, and may be reclaimed by AWS at any time with a two-minute warning. This may cause report generation to fail or restart from scratch. Running the application on AWS Elastic Beanstalk as a single application environment with an all-at-once deployment strategy will not meet the requirements, as it will not break down the monolith into microservices, nor will it minimize application downtime. The all-at-once deployment strategy will deploy updates to all instances simultaneously, causing a brief outage for the application.

[Amazon Elastic Container Service](#)

[Microservices on AWS](#)

Service Auto Scaling - Amazon Elastic Container Service

AWS Lambda

Amazon EC2 Spot Instances

[AWS Elastic Beanstalk]

Question 10

Question Type: MultipleChoice

A company has deployed a multiplayer game for mobile devices. The game requires live location tracking of players based on latitude and longitude. The data store for the game must support rapid updates and retrieval of locations.

The game uses an Amazon RDS for PostgreSQL DB instance with read replicas to store the location data.

a. During peak usage periods, the database is unable to maintain the performance that is needed for reading and writing updates. The game's user base is increasing rapidly.

What should a solutions architect do to improve the performance of the data tier?

Options:

- A-** Take a snapshot of the existing DB instance. Restore the snapshot with Multi-AZ enabled.
- B-** Migrate from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards.
- C-** Deploy Amazon DynamoDB Accelerator (DAX) in front of the existing DB instance. Modify the game to use DAX.
- D-** Deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instance. Modify the game to use Redis.

Answer:

D

Explanation:

The solution that will improve the performance of the data tier is to deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instance and modify the game to use Redis. This solution will enable the game to store and retrieve the location data of the players in a fast and scalable way, as Redis is an in-memory data store that supports geospatial data types and commands. By using ElastiCache for Redis, the game can reduce the load on the RDS for PostgreSQL DB instance, which is not optimized for high-frequency updates and queries of location data. ElastiCache for Redis also supports replication, sharding, and auto scaling to handle the increasing user base of the game.

The other solutions are not as effective as the first one because they either do not improve the performance, do not support geospatial data, or do not leverage caching. Taking a snapshot of the existing DB instance and restoring it with Multi-AZ enabled will not improve the performance of the data tier, as it only provides high availability and durability, but not scalability or low latency. Migrating from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards will not improve the performance of the data tier, as OpenSearch Service is mainly designed for full-text search and analytics, not for real-time location tracking. OpenSearch Service also does not support geospatial data types and commands natively, unlike Redis. Deploying Amazon DynamoDB Accelerator (DAX) in front

of the existing DB instance and modifying the game to use DAX will not improve the performance of the data tier, as DAX is only compatible with DynamoDB, not with RDS for PostgreSQL. DAX also does not support geospatial data types and commands.

[Amazon ElastiCache for Redis](#)

[Geospatial Data Support - Amazon ElastiCache for Redis](#)

[Amazon RDS for PostgreSQL](#)

[Amazon OpenSearch Service](#)

[Amazon DynamoDB Accelerator \(DAX\)](#)

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>

