



Free Amazon SAA-C03 Practice Questions

Shared by Greer on 12-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

A company provides an API interface to customers so the customers can retrieve their financial information. The company expects a larger number of requests during peak usage times of the year. The company requires the API to respond consistently with low latency to ensure customer satisfaction. The company needs to provide a compute host for the API.

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

Options:

- A- Use an Application Load Balancer and Amazon ECS.
- B- Use Amazon API Gateway and AWS Lambda functions with provisioned concurrency.
- C- Use an Application Load Balancer and an Amazon EKS cluster.
- D- Use Amazon API Gateway and AWS Lambda functions with reserved concurrency.

Answer:

B

Explanation:

Provisioned concurrency is the key requirement match because AWS documentation states that it pre-initializes Lambda execution environments and is designed to reduce cold-start latency, supporting predictable response times in the double-digit millisecond range. That makes it ideal for APIs that must remain consistently low-latency during peak traffic periods. Reserved concurrency does not solve cold starts; it only reserves capacity. ECS and EKS could certainly host the API, but both introduce more infrastructure management than API Gateway plus Lambda. Since the question asks for the least operational overhead while maintaining consistent low latency, the serverless pattern with API Gateway and Lambda with provisioned concurrency is the strongest AWS-native answer. (AWS Documentation)

=====

Question 2

Question Type: MultipleChoice

A company is building an application on Amazon EC2 instances. The application needs to handle a large number of transactions. The application requires an Amazon EBS data volume that has

configurable and consistent input/output operations per second (IOPS).

Which solution will meet these requirements?

Options:

- A- Provision EC2 instances with a Throughput Optimized HDD (st1) EBS root volume and a Cold HDD (sc1) EBS data volume.
- B- Provision EC2 instances with a Throughput Optimized HDD (st1) EBS volume that will serve as both a root volume and a data volume.
- C- Provision EC2 instances with a General Purpose SSD (gp3) EBS root volume and a Provisioned IOPS SSD (io2) EBS data volume.
- D- Provision EC2 instances with a General Purpose SSD (gp3) EBS root volume. Configure the application to store data in an Amazon S3 bucket.

Answer:

C

Explanation:

The correct answer is C because the requirement specifically calls for an Amazon EBS data volume that provides configurable and consistent IOPS for a transaction-heavy application. Provisioned IOPS SSD (io2) volumes are designed for workloads that need high performance, low latency, and predictable I/O performance. These volumes are the best fit for applications that process a large number of transactions and require storage performance to remain stable under load.

Using gp3 volume for the root device is appropriate because the operating system volume usually does not require the same high and sustained performance characteristics as the application data volume. The io2 data volume provides the ability to provision IOPS independently to match the application's transactional demands. This separation of root and data storage is a common AWS design pattern for performance-sensitive workloads.

Option A is incorrect because st1 and sc1 are HDD-based volumes intended for large, sequential workloads, not high-transaction workloads that require consistent low-latency I/O. Option B is also incorrect because st1 does not provide the required configurable and consistent IOPS characteristics. Option D is incorrect because the requirement explicitly states that the application needs an Amazon EBS data volume, and Amazon S3 is object storage rather than block storage attached as an EBS volume.

AWS storage best practices recommend Provisioned IOPS SSD volumes for mission-critical applications, transactional systems, and workloads that need sustained, predictable performance. Therefore, gp3 root volume combined with an io2 data volume is the most appropriate solution for this use case.

Question 3

Question Type: MultipleChoice

A solutions architect creates an Auto Scaling group for a memory-intensive application. The solutions architect wants to scale up and scale down based on memory usage. Which solution will meet this requirement?

Options:

- A- Install and configure the AWS Systems Manager Agent (SSM Agent). Create a step scaling policy that has step adjustments based on the memory usage trend.
- B- Install and configure the Amazon CloudWatch agent. Create a target tracking policy to scale based on the mem_used_percent CloudWatch metric.
- C- Install and configure the AWS Systems Manager Agent (SSM Agent). Create a target tracking policy to scale based on the mem_used_percent Amazon CloudWatch metric.
- D- Install and configure the Amazon CloudWatch agent. Create a scheduled scaling policy to scale based on the memory usage trend.

Answer:

B

Explanation:

The Amazon CloudWatch agent is required to collect memory utilization metrics (as memory metrics are not reported by default). A target tracking policy is the simplest and most effective way to scale based on a custom metric such as mem_used_percent.

Reference Extract:

'Install the CloudWatch agent to collect memory metrics, and create a target tracking scaling policy using these custom metrics.'

Source: AWS Certified Solutions Architect -- Official Study Guide, Monitoring and Scaling section.

Question 4

Question Type: MultipleChoice

A company is deploying an application in three AWS Regions using an Application Load

Balancer. Amazon Route 53 will be used to distribute traffic between these Regions.

Which Route 53 configuration should a solutions architect use to provide the MOST high-performing experience?

Options:

- A- Create an A record with a latency policy.
- B- Create an A record with a geolocation policy.
- C- Create a CNAME record with a failover policy.
- D- Create a CNAME record with a geoproximity policy.

Answer:

A

Explanation:

Latency-based routing in Amazon Route 53 is designed to route users to the Region that provides the lowest network latency, based on Amazon's measurements of latency between AWS Regions and users' networks. For applications deployed in multiple Regions, this provides the highest performance experience for global users.

Therefore, creating an A record with a latency routing policy is the correct choice.

Geolocation (Option B) routes based on user location, which may not always correspond to the lowest latency.

Failover (Option C) is for active-passive architectures, not performance optimization.

Geoproximity (Option D) is more complex and focused on directing traffic based on geographic bias rather than measured latency.

Question 5

Question Type: MultipleChoice

A company runs a custom application on Amazon EC2 On-Demand Instances. The application has frontend nodes that must run 24/7. The backend nodes only need to run for short periods depending on the workload.

Frontend nodes accept jobs and place them in queues. Backend nodes asynchronously process jobs from the queues, and jobs can be restarted. The company wants to scale infrastructure based on workload, using the most cost-effective option.

Which solution meets these requirements MOST cost-effectively?

Options:

- A- Use Reserved Instances for the frontend nodes. Use AWS Fargate for the backend nodes.
- B- Use Reserved Instances for the frontend nodes. Use Spot Instances for the backend nodes.
- C- Use Spot Instances for the frontend nodes. Use Reserved Instances for the backend nodes.
- D- Use Spot Instances for the frontend nodes. Use AWS Fargate for the backend nodes.

Answer:

B

Explanation:

AWS documentation states that workloads running 24/7 should use Reserved Instances or Savings Plans for the lowest cost. Therefore, the frontend nodes, which always run, should use Reserved Instances.

The backend nodes process asynchronous, restartable jobs, which makes them ideal for EC2 Spot Instances, the most cost-effective compute option for interruption-tolerant workloads.

Fargate (Options A and D) is significantly more expensive for large compute usage. Spot Instances cannot be used for critical frontend nodes (Option C).

=====

Question 6

Question Type: MultipleChoice

A company is designing an application to maintain a record of customer orders. The application will generate events. The company wants to use an Amazon EventBridge event bus to send the application's events to an Amazon DynamoDB table. Which solution will meet these requirements?

Options:

- A- Use the EventBridge default event bus. Configure DynamoDB Streams for the DynamoDB table that hosts the customer order data.
- B- Create an EventBridge custom event bus. Create an AWS Lambda function as a target.

Configure the Lambda function to forward the customer order data to the DynamoDB table.

C- Create an EventBridge partner event bus. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe an AWS Lambda function to the SNS topic. Configure the Lambda function to read the customer order data and to forward the data to the DynamoDB table.

D- Create an EventBridge partner event bus. Create an AWS Lambda function as a target. Configure the Lambda function to forward the customer order data to the DynamoDB table.

Answer:

C

Explanation:

Amazon EventBridge supports routing application-generated events to AWS Lambda targets. The Lambda function can process and insert events into a DynamoDB table. This is a standard design pattern for connecting event-driven applications with DynamoDB. Option A confuses DynamoDB Streams (which streams changes from DynamoDB) with EventBridge (which is for event routing). Options C and D reference partner event buses, which are intended for SaaS integrations, not custom application events. Therefore, the correct and simplest solution is to use a custom EventBridge event bus with a Lambda target (B).

Question 7

Question Type: MultipleChoice

A company is developing a latency-sensitive application. Part of the application includes several AWS Lambda functions that need to initialize as quickly as possible. The Lambda functions are written in Java and contain initialization code outside the handlers to load libraries, initialize classes, and generate unique IDs.

Which solution will meet the startup performance requirement MOST cost-effectively?

Options:

A- Move all the initialization code to the handlers for each Lambda function. Activate Lambda SnapStart for each Lambda function. Configure SnapStart to reference the \$LATEST version of each Lambda function.

B- Publish a version of each Lambda function. Create an alias for each Lambda function. Configure each alias to point to its corresponding version. Set up a provisioned concurrency configuration for each Lambda function to point to the corresponding alias.

C- Publish a version of each Lambda function. Set up a provisioned concurrency configuration for each Lambda function to point to the corresponding version. Activate Lambda SnapStart for

the published versions of the Lambda functions.

D- Update the Lambda functions to add a pre-snapshot hook. Move the code that generates unique IDs into the handlers. Publish a version of each Lambda function. Activate Lambda SnapStart for the published versions of the Lambda functions.

Answer:

D

Explanation:

AWS Lambda SnapStart is designed to improve the cold start performance of Java Lambda functions by initializing the function, taking a snapshot of the execution environment, and then reusing that snapshot for subsequent invocations. However, some code (such as code that generates unique IDs or session-specific data) should run during each invocation, not during the snapshot process. By using a pre-snapshot hook and moving the unique ID generation into the handler, you ensure that non-deterministic or per-invocation code is executed correctly, while the rest of the initialization benefits from SnapStart. This delivers the lowest latency and cost, as you do not need to pay for provisioned concurrency.

AWS Documentation Extract:

'Lambda SnapStart is ideal for Java functions with long cold starts due to heavy initialization. Move non-deterministic code such as unique ID generation to the handler, and use pre-snapshot hooks to customize what gets snapshotted. SnapStart works only for published versions, not \$LATEST.'

(Source: AWS Lambda documentation, Using SnapStart for Java Functions)

Other options:

A: SnapStart is not supported for the \$LATEST version; must be a published version.

B & C: Provisioned concurrency removes cold starts but is more expensive than SnapStart for most workloads.

C: You must use SnapStart on published versions, but provisioned concurrency is not required for SnapStart and adds cost.

Question 8

Question Type: MultipleChoice

A developer needs to export the contents of several Amazon DynamoDB tables into Amazon S3 buckets to comply with company data regulations. The developer uses the AWS CLI to

runcommands to export from each table to the proper S3 bucket. The developer sets up AWS credentials correctly and grants resources appropriate permissions. However, the exports of some tables fail.

What should the developer do to resolve this issue?

Options:

- A- Ensure that point-in-time recovery is enabled on the DynamoDB tables.
- B- Ensure that the target S3 bucket is in the same AWS Region as the DynamoDB table.
- C- Ensure that DynamoDB streaming is enabled for the tables.
- D- Ensure that DynamoDB Accelerator (DAX) is enabled.

Answer:

A

Explanation:

Export Requirements:

To export data from DynamoDB to Amazon S3, point-in-time recovery (PITR) must be enabled for the tables. This feature creates a snapshot of the data, which is essential for exports.

Incorrect Options Analysis:

Option B: S3 buckets and DynamoDB tables do not need to be in the same region for exports.

Option C: DynamoDB streams are unrelated to the export functionality.

Option D: DAX accelerates reads but has no role in exports.

Exporting DynamoDB Data to Amazon S3

Question 9

Question Type: MultipleChoice

A company is storing data in Amazon S3 buckets. The company needs to retain any objects that contain personally identifiable information (PII) that might need to be reviewed.

A solutions architect must develop an automated solution to identify objects that contain PII and apply the necessary controls to prevent deletion before review.

Which combination of steps should the solutions architect take to meet these requirements?
(Choose THREE.)

Options:

- A- Create a job in Amazon Macie to scan the S3 buckets for the relevant sensitive data identifiers.
- B- Move the identified objects to the S3 Glacier Deep Archive storage class.
- C- Create an AWS Lambda function that performs an S3 Object Lock legal hold operation on the identified objects.
- D- Create an AWS Lambda function that applies an S3 Object Lock retention period to the identified objects in governance mode.
- E- Create an Amazon EventBridge rule that invokes the AWS Lambda function when Amazon Macie detects sensitive data.
- F- Configure multi-factor authentication (MFA) delete on the S3 buckets.

Answer:

A, C, E

Explanation:

A: Amazon Macie can scan S3 buckets for sensitive data and identify PII.

C: S3 Object Lock legal hold prevents object deletion until a review is complete, meeting compliance requirements.

E: EventBridge can trigger the Lambda function automatically when Macie detects sensitive data, creating an automated workflow.

AWS Documentation Extract:

'Amazon Macie automatically discovers, classifies, and protects sensitive data in AWS. You can use EventBridge to invoke a Lambda function for post-processing, such as applying Object Lock legal hold to flagged objects.'

(Source: AWS Macie and S3 Object Lock documentation)

B, D: Moving to Glacier Deep Archive or applying retention in governance mode is not specific to deletion prevention pending review.

F: MFA Delete adds a security layer but does not automate object retention for flagged PII.

Question 10

Question Type: MultipleChoice

A company is developing a microservices-based application to manage the company's delivery operations. The application consists of microservices that process orders, manage a fleet of delivery vehicles, and optimize delivery routes. The microservices must be able to scale independently and must be able to handle bursts of traffic without any data loss.

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

Options:

- A- Use Amazon API Gateway REST APIs to establish communication between microservices. Deploy the application on Amazon EC2 instances in Auto Scaling groups.
- B- Use Amazon SQS to establish communication between microservices. Deploy the application on Amazon ECS containers on AWS Fargate.
- C- Use WebSocket-based communication between microservices. Deploy the application on Amazon EC2 instances in Auto Scaling groups.
- D- Use Amazon SNS to establish communication between microservices. Deploy the application on Amazon ECS containers on Amazon EC2.

Answer:

B

Explanation:

Amazon SQS is the best messaging layer here because it provides durable buffering and decouples microservices so they can scale independently while absorbing bursts without losing data. AWS recommends SQS for decoupled architectures where producers and consumers operate at different speeds. For compute, AWS Fargate reduces operational overhead because AWS manages the underlying servers, patching, and capacity layer. API Gateway and WebSockets are not the primary answer for internal asynchronous microservice communication, and SNS is fanout-oriented rather than the most natural queueing model for durable work processing between services. Therefore, SQS plus ECS on Fargate is the most appropriate low-overhead design.

=====

Question 11

Question Type: MultipleChoice

A company runs an online order management system on AWS. The company stores order and inventory data for the previous 5 years in an Amazon Aurora MySQL database. The company deletes inventory data after 5 years.

The company wants to optimize costs to archive data.

Options:

Options:

- A- Create an AWS Glue crawler to export data to Amazon S3. Create an AWS Lambda function to compress the data.
- B- Use the SELECT INTO OUTFILE S3 query on the Aurora database to export the data to Amazon S3. Configure S3 Lifecycle rules on the S3 bucket.
- C- Create an AWS Glue DataBrew Job to migrate data from Aurora to Amazon S3. Configure S3 Lifecycle rules on the S3 bucket.
- D- Use the AWS Schema Conversion Tool (AWS SCT) to replicate data from Aurora to Amazon S3. Use the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.

Answer:

B

Explanation:

The SELECT INTO OUTFILE S3 feature allows you to export Amazon Aurora MySQL data directly to Amazon S3 with minimal operational overhead. This method is efficient and cost-effective for archiving historical data.

You can configure S3 Lifecycle rules to transition the exported data to lower-cost storage (e.g., S3 Glacier or S3 Standard-IA) and eventually delete it after 5 years.

No need for additional ETL tools like Glue or DataBrew unless complex transformations are required.

Exporting data from Aurora MySQL to S3

Question 12

Question Type: MultipleChoice

A company that has multiple AWS accounts maintains an on-premises Microsoft Active Directory. The company needs a solution to implement Single Sign-On for its employees. The company wants to use AWS IAM Identity Center.

The solution must meet the following requirements:

Allow users to access AWS accounts and third-party applications by using existing Active Directory credentials.

Enforce multi-factor authentication (MFA) to access AWS accounts.

Centrally manage permissions to access AWS accounts and applications.

Options:



Options:

A- Create an IAM identity provider for Active Directory in each AWS account. Ensure that Active Directory users and groups access AWS accounts directly through IAM roles. Use IAM Identity Center to enforce MFA in each account for all users.

B- Use AWS Directory Service to create a new AWS Managed Microsoft AD Active Directory. Configure IAM Identity Center in each account to use the new AWS Managed Microsoft AD Active Directory as the identity source. Use IAM Identity Center to enforce MFA for all users.

C- Use IAM Identity Center with the existing Active Directory as the identity source. Enforce MFA for all users. Use AWS Organizations and Active Directory groups to manage access permissions for AWS accounts and application access.

D- Use AWS Lambda functions to periodically synchronize Active Directory users and groups with IAM users and groups in each AWS account. Use IAM roles and policies to manage application access. Create a second Lambda function to enforce MFA.

Answer:

C



Explanation:

Detailed

A . IAM identity provider:Does not support centralized management across multiple accounts.

B . AWS Managed AD:Unnecessary if an on-premises Active Directory already exists.

C . IAM Identity Center + Existing AD:Best approach to integrate existing Active Directory for SSO, with MFA and centralized permissions.

D . Lambda for synchronization:Adds complexity and does not leverage IAM Identity Center capabilities.

Question 13

Question Type: MultipleChoice

A finance company uses an on-premises search application to collect streaming data from various producers. The application provides real-time updates to search and visualization features. The company is planning to migrate to AWS and wants to use an AWS native solution. Which solution will meet these requirements?

Options:

- A- Use Amazon EC2 instances to ingest and process the data streams to Amazon S3 buckets for storage. Use Amazon Athena to search the data. Use Amazon Managed Grafana to create visualizations.
- B- Use Amazon EMR to ingest and process the data streams to Amazon Redshift for storage. Use Amazon Redshift Spectrum to search the data. Use Amazon QuickSight to create visualizations.
- C- Use Amazon Elastic Kubernetes Service (Amazon EKS) to ingest and process the data streams to Amazon DynamoDB for storage. Use Amazon CloudWatch to create graphical dashboards to search and visualize the data.
- D- Use Amazon Kinesis Data Streams to ingest and process the data streams to Amazon OpenSearch Service. Use OpenSearch Service to search the data. Use Amazon QuickSight to create visualizations.

Answer:

D

Explanation:

Amazon OpenSearch Service is the AWS-native service for real-time search, analytics, and visualization. Ingesting streaming data with Amazon Kinesis Data Streams ensures scalable and reliable data ingestion. OpenSearch supports indexing and querying the data in near real time. QuickSight can then be integrated for visualization and reporting. Options A and B involve batch processing and are not optimized for real-time search. Option C (EKS + DynamoDB) is not a fit for search workloads, as DynamoDB does not support advanced text search. Option D provides the closest AWS-native equivalent to the company's existing search-and-visualization architecture.

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>

20%
DISCOUNT

P2P
exams