

SAA-C03 Exam Guide - SAA-C03 Study Tools & SAA-C03 Exam Torrent



2026 Latest NewPassLeader SAA-C03 PDF Dumps and SAA-C03 Exam Engine Free Share: <https://drive.google.com/open?id=1J5rgOx0CeyYtYTVMyLtJtm5cCFN7JI84>

The free demo SAA-C03 practice question is available for instant download. Download the SAA-C03 exam dumps demo free of cost and explore the top features of Amazon SAA-C03 exam questions and if you feel that the Amazon SAA-C03 Exam Questions can be helpful in AWS Certified Solutions Architect - Associate (SAA-C03) exam preparation then take your buying decision.

Amazon SAA-C03 Exam is a valuable certification for IT professionals who want to work with AWS cloud services. AWS Certified Solutions Architect - Associate certification validates the skills and knowledge required to design, deploy, and manage scalable, highly available, and fault-tolerant systems on the AWS platform. It is an essential certification for individuals who want to advance their careers in cloud computing and work for companies that use AWS services. AWS Certified Solutions Architect - Associate certification is also an excellent way for IT professionals to demonstrate their expertise and stand out in a competitive job market.

>> Test SAA-C03 Voucher <<

Reliable SAA-C03 Test Pass4sure, Valid SAA-C03 Exam Fee

You have to upgrade your skills and knowledge then you will be in a position to compete in the modern world. The Amazon SAA-C03 certification offers a great way to learn new in-demand skills and upgrade your knowledge level. To do this you just need to enroll in the SAA-C03 Exam and put in your efforts to pass this career booster SAA-C03 certification exam.

Amazon AWS Certified Solutions Architect - Associate Sample Questions (Q760-Q765):

NEW QUESTION # 760

[Design High-Performing Architectures]

A company tracks customer satisfaction by using surveys that the company hosts on its website. The surveys sometimes reach thousands of customers every hour. Survey results are currently sent in email messages to the company so company employees can manually review results and assess customer sentiment.

The company wants to automate the customer survey process. Survey results must be available for the previous 12 months. Which solution will meet these requirements in the MOST scalable way?

- A. Send the survey results data to an Amazon API Gateway endpoint that is connected to an Amazon Simple Queue Service (Amazon SQS) queue. Configure the SQS queue to invoke an AWS Lambda function that calls Amazon Lex for sentiment analysis and saves the results to an Amazon DynamoDB table. Set the TTL for all records to 365 days in the future.
- B. Send the survey results data to an Amazon API Gateway endpoint that is connected to an Amazon Simple Queue Service (Amazon SQS) queue. Create an AWS Lambda function to poll the SQS queue, call Amazon Comprehend for sentiment analysis, and save the results to an Amazon DynamoDB table. Set the TTL for all records to 365 days in the future.
- C. Write the survey results data to an Amazon S3 bucket. Use S3 Event Notifications to invoke an AWS Lambda function to

read the data and call Amazon Rekognition for sentiment analysis. Store the sentiment analysis results in a second S3 bucket. Use S3 Lifecycle policies on each bucket to expire objects after 365 days.

- D. Send the survey results data to an API that is running on an Amazon EC2 instance. Configure the API to store the survey results as a new record in an Amazon DynamoDB table, call Amazon Comprehend for sentiment analysis, and save the results in a second DynamoDB table. Set the TTL for all records to 365 days in the future.

Answer: B

Explanation:

This solution is the most scalable and efficient way to handle large volumes of survey data while automating sentiment analysis:

API Gateway and SQS: The survey results are sent to API Gateway, which forwards the data to an SQS queue. SQS can handle large volumes of messages and ensures that messages are not lost.

AWS Lambda: Lambda is triggered by polling the SQS queue, where it processes the survey data.

Amazon Comprehend: Comprehend is used for sentiment analysis, providing insights into customer satisfaction.

DynamoDB with TTL: Results are stored in DynamoDB with a Time to Live (TTL) attribute set to expire after 365 days, automatically removing old data and reducing storage costs.

Option B (EC2 API): Running an API on EC2 requires more maintenance and scalability management compared to API Gateway.

Option C (S3 and Rekognition): Amazon Rekognition is for image and video analysis, not sentiment analysis.

Option D (Amazon Lex): Amazon Lex is used for building conversational interfaces, not sentiment analysis.

AWS Reference:

Amazon Comprehend for Sentiment Analysis

Amazon SQS

DynamoDB TTL

NEW QUESTION # 761

[Design Secure Architectures]

A company recently migrated its entire IT environment to the AWS Cloud. The company discovers that users are provisioning oversized Amazon EC2 instances and modifying security group rules without using the appropriate change control process. A solutions architect must devise a strategy to track and audit these inventory and configuration changes.

Which actions should the solutions architect take to meet these requirements? (Select TWO)

- A. Restore previous resource configurations with an AWS CloudFormation template
- B. Use data lifecycle policies for the Amazon EC2 instances
- C. Enable AWS CloudTrail and use it for auditing
- D. Enable AWS Config and create rules for auditing and compliance purposes
- E. Enable AWS Trusted Advisor and reference the security dashboard

Answer: C,D

Explanation:

A) Enable AWS CloudTrail and use it for auditing. AWS CloudTrail provides a record of API calls and can be used to audit changes made to EC2 instances and security groups. By analyzing CloudTrail logs, the solutions architect can track who provisioned oversized instances or modified security groups without proper approval. D) Enable AWS Config and create rules for auditing and compliance purposes. AWS Config can record the configuration changes made to resources like EC2 instances and security groups. The solutions architect can create AWS Config rules to monitor for non-compliant changes, like launching certain instance types or opening security group ports without permission. AWS Config would alert on any violations of these rules.

NEW QUESTION # 762

A company uses an organization in AWS Organizations to manage AWS accounts that contain applications.

The company sets up a dedicated monitoring member account in the organization. The company wants to query and visualize observability data across the accounts by using Amazon CloudWatch.

Which solution will meet these requirements?

- A. Configure a new IAM user in the monitoring account. In each AWS account, configure an IAM policy to have access to query and visualize the CloudWatch data in the account. Attach the new IAM policy to the new IAM user.
- B. Set up service control policies (SCPs) to provide access to CloudWatch in the monitoring account under the Organizations root organizational unit (OU).
- C. Enable CloudWatch cross-account observability for the monitoring account. Deploy an AWS CloudFormation template provided by the monitoring account in each AWS account to share the data with the monitoring account.

- D. Create a new IAM user in the monitoring account. Create cross-account IAM policies in each AWS account. Attach the IAM policies to the new IAM user.

Answer: C

Explanation:

CloudWatch cross-account observability is a feature that allows you to monitor and troubleshoot applications that span multiple accounts within a Region. You can seamlessly search, visualize, and analyze your metrics, logs, traces, and Application Insights applications in any of the linked accounts without account boundaries¹.

To enable CloudWatch cross-account observability, you need to set up one or more AWS accounts as monitoring accounts and link them with multiple source accounts. A monitoring account is a central AWS account that can view and interact with observability data shared by other accounts. A source account is an individual AWS account that shares observability data and resources with one or more monitoring accounts¹.

To create links between monitoring accounts and source accounts, you can use the CloudWatch console, the AWS CLI, or the AWS API. You can also use AWS Organizations to link accounts in an organization or organizational unit to the monitoring account¹. CloudWatch provides a CloudFormation template that you can deploy in each source account to share observability data with the monitoring account. The template creates a sink resource in the monitoring account and an observability link resource in the source account. The template also creates the necessary IAM roles and policies to allow cross-account access to the observability data².

Therefore, the solution that meets the requirements of the question is to enable CloudWatch cross-account observability for the monitoring account and deploy the CloudFormation template provided by the monitoring account in each AWS account to share the data with the monitoring account.

The other options are not valid because:

* Service control policies (SCPs) are a type of organization policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization, allowing you to ensure your accounts stay within your organization's access control guidelines³. SCPs do not provide access to CloudWatch in the monitoring account, but rather restrict the actions that users and roles can perform in the source accounts. SCPs are not required to enable CloudWatch cross-account observability, as the CloudFormation template creates the necessary IAM roles and policies for cross-account access².

* IAM users are entities that you create in AWS to represent the people or applications that use them to interact with AWS. IAM users can have permissions to access the resources in your AWS account⁴.

Configuring a new IAM user in the monitoring account and an IAM policy in each AWS account to have access to query and visualize the CloudWatch data in the account is not a valid solution, as it does not enable CloudWatch cross-account observability. This solution would require the IAM user to switch

* between different accounts to view the observability data, which is not seamless and efficient. Moreover, this solution would not allow the IAM user to search, visualize, and analyze metrics, logs, traces, and Application Insights applications across multiple accounts in a single place¹.

* Cross-account IAM policies are policies that allow you to delegate access to resources that are in different AWS accounts that you own. You attach a cross-account policy to a user or group in one account, and then specify which accounts the user or group can access⁵. Creating a new IAM user in the monitoring account and cross-account IAM policies in each AWS account is not a valid solution, as it does not enable CloudWatch cross-account observability. This solution would also require the IAM user to switch between different accounts to view the observability data, which is not seamless and efficient. Moreover, this solution would not allow the IAM user to search, visualize, and analyze metrics, logs, traces, and Application Insights applications across multiple accounts in a single place¹.

References: CloudWatch cross-account observability, CloudFormation template for CloudWatch cross-account observability, Service control policies, IAM users, Cross-account IAM policies

NEW QUESTION # 763

A company has an on-premises SFTP file transfer solution. The company is migrating to the AWS Cloud to scale the file transfer solution and to optimize costs by using Amazon S3. The company's employees will use their credentials for the on-premises Microsoft Active Directory (AD) to access the new solution. The company wants to keep the current authentication and file access mechanisms.

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

- A. Create an AWS Transfer Family SFTP endpoint. Configure the endpoint to use the AWS Directory Service option as the identity provider to connect to the existing Active Directory.
- **B. Create an AWS Transfer Family server with SFTP endpoints. Choose the AWS Directory Service option as the identity provider. Use AD Connector to connect the on-premises Active Directory.**
- C. Configure an S3 File Gateway. Create SMB file shares on the file gateway that use the existing Active Directory to authenticate.

- D. Configure an Auto Scaling group with Amazon EC2 instances to run an SFTP solution. Configure the group to scale up at 60% CPU utilization.

Answer: B

Explanation:

AWS Transfer Family: This service provides fully managed support for file transfers directly into and out of Amazon S3 using the SFTP, FTPS, and FTP protocols.

SFTP Endpoints:

Set up an AWS Transfer Family server and configure SFTP endpoints to handle the file transfers.

This service is scalable and managed, reducing operational overhead compared to running an SFTP solution on EC2 instances.

Integration with Active Directory:

Choose the AWS Directory Service option as the identity provider for the Transfer Family server.

Use AD Connector to link the on-premises Active Directory with AWS, allowing employees to use their existing AD credentials to access the SFTP service.

Operational Efficiency: This solution leverages managed services for both file transfer and identity management, ensuring minimal changes to the current authentication mechanisms and reducing operational overhead.

Reference:

AWS Transfer Family

AWS Directory Service and AD Connector

NEW QUESTION # 764

A company is building a serverless application to process clickstream data from its website. The clickstream data is sent to an Amazon Kinesis Data Streams data stream from the application web servers.

The company wants to enrich the clickstream data by joining the clickstream data with customer profile data from an Amazon Aurora Multi-AZ database. The company wants to use Amazon Redshift to analyze the enriched data. The solution must be highly available.

Which solution will meet these requirements?

- A. Use Amazon Kinesis Data Firehose to load the clickstream data from Kinesis Data Streams to Amazon S3. Use AWS Glue crawlers to infer the schema and populate the AWS Glue Data Catalog. Use Amazon Athena to query the raw data in Amazon S3.
- **B. Use an AWS Lambda function to process and enrich the clickstream data. Use the same Lambda function to write the clickstream data to Amazon S3. Use Amazon Redshift Spectrum to query the enriched data in Amazon S3.**
- C. Use an Amazon Elastic Container Service (Amazon ECS) task with AWS Fargate Spot capacity to poll the data stream and enrich the clickstream data. Configure an Amazon EC2 instance to use the COPY command to send the enriched results to Amazon Redshift.
- D. Use an Amazon EC2 Spot Instance to poll the data stream and enrich the clickstream data. Configure the EC2 instance to use the COPY command to send the enriched results to Amazon Redshift.

Answer: B

Explanation:

Option A is the best solution as it leverages AWS Lambda for serverless, scalable, and highly available processing and enrichment of clickstream data. Lambda can process the data in real-time, join it with the Aurora database data, and write the enriched results to Amazon S3. From S3, Amazon Redshift Spectrum can directly query the enriched data without needing to load the data into Redshift, enabling cost efficiency and high availability.

Why Other Options Are Incorrect:

Option B: EC2 Spot Instances are not guaranteed to be highly available, as Spot Instances can be interrupted at any time. This does not align with the requirement for high availability.

Option C: While ECS with AWS Fargate provides scalability, using EC2 for the COPY command introduces operational overhead and compromises high availability.

Option D: Kinesis Data Firehose and Athena are suitable for querying raw data, but they do not directly support enriching the data by joining with Aurora. This solution fails to meet the requirement for data enrichment.

Key AWS Features Used:

AWS Lambda: Real-time serverless processing with integration capabilities for Aurora and S3.

Amazon S3: Cost-effective storage for enriched data.

Amazon Redshift Spectrum: Direct querying of data stored in S3 without loading it into Redshift.

AWS Documentation Reference:

AWS Lambda Function Overview

