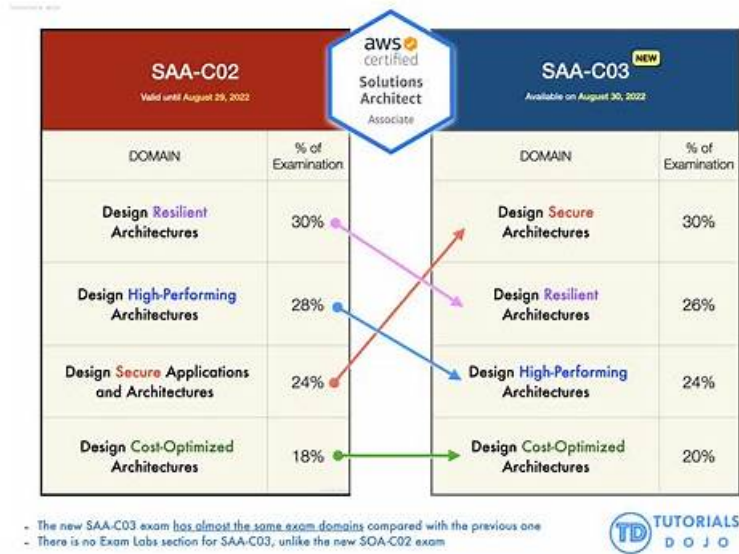


SAA-C03 Study Materials & SAA-C03 VCE Dumps & SAA-C03 Test Prep



What's more, part of that PDF4Test SAA-C03 dumps now are free: <https://drive.google.com/open?id=16B2To-jreNglmueNPwR3oZyVaLfhOYu>

We apply international recognition third party for payment for SAA-C03 exam materials, therefore, if you choose us, your money safety will be guaranteed. The third party will guarantee your interests. Besides, SAA-C03 exam materials of us is high-quality, they will help you pass the exam successfully. We also pass guarantee and money back guarantee if you fail to pass the exam. SAA-C03 Exam Braindumps offer you free update for one year, and in the following year, you can know the latest information for the exam. The latest version for SAA-C03 will be sent to your email automatically.

Amazon SAA-C03 (Amazon AWS Certified Solutions Architect - Associate) Certification Exam is designed for individuals who want to validate their skills and knowledge in designing and deploying scalable, highly available, and fault-tolerant systems on the Amazon Web Services (AWS) platform. AWS Certified Solutions Architect - Associate certification exam is ideal for professionals who are looking to enhance their career prospects in cloud computing, particularly in the AWS ecosystem. The SAA-C03 Certification Exam is the latest version of the AWS Solutions Architect - Associate certification and is updated to reflect the latest AWS services and trends in cloud computing.

>> SAA-C03 Actual Test Answers <<

SAA-C03 Test Registration, SAA-C03 New Real Test

PDF4Test exam dumps have two version-PDF and SOFT version which will give you convenient. It is very convenient for you to use PDF real questions and answers. And you can download these materials and print it out for study at any time. The SOFT version simulates the real exam which will give you more realistic feeling. When you are faced with the real exam, you can pass Amazon SAA-C03 test easily.

To prepare for the SAA-C03 Exam, you should have a good understanding of AWS services and features, as well as hands-on experience designing and deploying AWS solutions. You can take advantage of various study resources, such as AWS training courses, practice exams, and whitepapers, to enhance your knowledge and skills. You can also join online forums and communities to connect with other AWS professionals and share your experiences and insights.

Amazon AWS Certified Solutions Architect - Associate Sample Questions (Q656-Q661):

NEW QUESTION # 656

A company wants to release a new device that will collect data to track overnight sleep on an intelligent mattress. Sensors will send

data that will be uploaded to an Amazon S3 bucket. Each mattress generates about 2 MB of data each night. An application must process the data and summarize the data for each user. The application must make the results available as soon as possible. Every invocation of the application will require about 1 GB of memory and will finish running within 30 seconds. Which solution will run the application MOST cost-effectively?

- A. Amazon EMR with an Apache Spark script
- B. AWS Glue with a PySpark job
- **C. AWS Lambda with a Python script**
- D. AWS Glue with a Scala job

Answer: C

Explanation:

AWS Lambda supports functions up to 10 GB of memory and 15 minutes execution time. Each invocation here requires only 1 GB of memory and finishes in 30 seconds, making it an ideal fit for Lambda. Lambda is cost-effective for event-driven, short-duration workloads and requires no infrastructure management. AWS Glue and EMR are better suited for large-scale ETL or distributed processing, which is unnecessary and more costly for this workload.

NEW QUESTION # 657

A company needs to keep user transaction data in an Amazon DynamoDB table. The company must retain the data for 7 years. What is the MOST operationally efficient solution that meets these requirements?

- A. Use DynamoDB point-in-time recovery to back up the table continuously.
- B. Create an on-demand backup of the table by using the DynamoDB console. Store the backup in an Amazon S3 bucket. Set an S3 Lifecycle configuration for the S3 bucket.
- **C. Use AWS Backup to create backup schedules and retention policies for the table.**
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function. Configure the Lambda function to back up the table and to store the backup in an Amazon S3 bucket. Set an S3 Lifecycle configuration for the S3 bucket.

Answer: C

Explanation:

We recommend you use AWS Backup to automatically delete the backups that you no longer need by configuring your lifecycle when you created your backup plan.

<https://docs.aws.amazon.com/aws-backup/latest/devguide/deleting-backups.html>

NEW QUESTION # 658

A developer used the AWS SDK to create an application that aggregates and produces log records for 10 services. The application delivers data to an Amazon Kinesis Data Streams stream.

Each record contains a log message with a service name, creation timestamp, and other log information. The stream has 15 shards in provisioned capacity mode. The stream uses service name as the partition key.

The developer notices that when all the services are producing logs, `ProvisionedThroughputExceededException` errors occur during `PutRecord` requests. The stream metrics show that the write capacity the applications use is below the provisioned capacity.

How should the developer resolve this issue?

- A. Use a separate Kinesis stream for each service to generate the logs.
- B. Double the number of shards until the throttling errors stop occurring.
- **C. Change the partition key from service name to creation timestamp.**
- D. Change the capacity mode from provisioned to on-demand.

Answer: C

Explanation:

Partition Key Issue:

Using "service name" as the partition key results in uneven data distribution. Some shards may become hot due to excessive logs from certain services, leading to throttling errors.

Changing the partition key to "creation timestamp" ensures a more even distribution of records across shards.

Incorrect Options Analysis:

Option A: On-demand capacity mode eliminates throughput management but is more expensive and does not address the root cause.

Option B: Adding more shards does not solve the issue if the partition key still creates hot shards.

Option D: Using separate streams increases complexity and is unnecessary.

Reference:

Kinesis Data Streams Partition Key Best Practices

NEW QUESTION # 659

A company uses AWS Lambda functions in a private subnet in a VPC to run application logic.

The Lambda functions must not have access to the public internet. Additionally, all data communication must remain within the private network. As part of a new requirement, the application logic needs access to an Amazon DynamoDB table. What is the MOST secure way to meet this new requirement?

- A. Create a gateway VPC endpoint for DynamoDB to provide access to the table.
- B. Use a network ACL to only allow access to the DynamoDB table from the VPC.
- C. Use a security group to only allow access to the DynamoDB table from the VPC.
- D. Provision the DynamoDB table inside the same VPC that contains the Lambda functions.

Answer: A

Explanation:

You cannot "place" DynamoDB inside a VPC. Instead, you use a VPC endpoint.

A Gateway VPC Endpoint for DynamoDB enables private connectivity between your VPC and DynamoDB without traversing the public internet.

"Use a gateway VPC endpoint to privately connect your VPC to DynamoDB without requiring an internet gateway or NAT."

NEW QUESTION # 660

A company needs to use its on-premises LDAP directory service to authenticate its users to the AWS Management Console. The directory service is not compatible with Security Assertion Markup Language (SAML).

Which solution meets these requirements?

- A. Set up a process that rotates the IAM credentials whenever LDAP credentials are updated.
- B. Create an IAM policy that uses AWS credentials, and integrate the policy into LDAP.
- C. Enable AWS IAM Identity Center between AWS and the on-premises LDAP.
- D. Develop an on-premises custom identity broker application or process that uses AWS STS to get short-lived credentials.

Answer: D

Explanation:

The correct answer is D because the company needs to authenticate users from an on-premises LDAP directory to the AWS Management Console, but the directory is not SAML-compatible. In this scenario, the AWS-recommended pattern is to use a custom identity broker that authenticates users against the existing LDAP directory and then requests temporary AWS credentials from AWS Security Token Service (AWS STS).

This enables federated access to AWS without creating long-term IAM user credentials for each person.

A custom identity broker serves as the intermediary between the non-SAML directory and AWS. After validating a user with LDAP, the broker can call AWS STS to issue short-lived credentials or console sign-in access. This approach preserves the use of the company's existing identity source while meeting AWS security best practices for temporary credentials and centralized identity management.

Option A is incorrect because AWS IAM Identity Center typically relies on supported identity sources and federation methods, and the question explicitly states that the LDAP directory is not SAML-compatible.

Option B is incorrect because IAM policies do not integrate directly into LDAP as an authentication mechanism. Option C is incorrect because rotating IAM credentials tied to LDAP changes still relies on long-term credentials and does not provide a proper federated sign-in model.

AWS federation guidance supports the use of a custom identity broker when an organization has a non-SAML-compatible identity system. Therefore, the best solution is to develop an on-premises custom identity broker that uses AWS STS to obtain short-lived credentials for AWS Management Console access.

