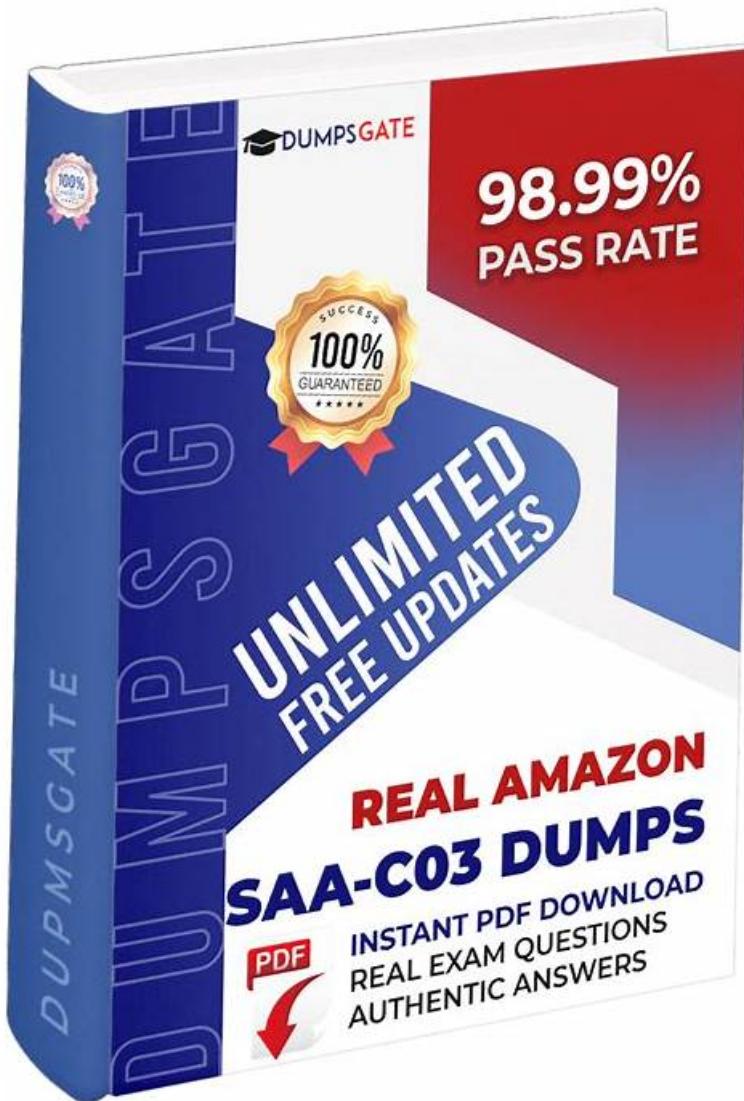


Reliable SAA-C03 Test Dumps & SAA-C03 Latest Braindumps Free



DOWNLOAD the newest ITExamDownload SAA-C03 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1FEplqMW5qCymvrk89H-iWvDlAtO9SmL>

This is a desktop-based exam simulator software. The user can easily get used to its format and it is compatible with Windows. It has a bank of the actual AWS Certified Solutions Architect - Associate (SAA-C03) exam questions, going through them will prove to be vital for your Amazon SAA-C03 exam preparation since a candidate must know his lacking points. The SAA-C03 Practice Exam simulator is reliable because its Amazon SAA-C03 exam questions have been compiled by experts and you can be sure of their validity and accuracy. All features of the web-based practice exam are present in this software.

The SAA-C03 Certification Exam is a multiple-choice exam that consists of 65 questions. Candidates have 130 minutes to complete the exam. SAA-C03 exam covers a range of topics including AWS services, architecture, and best practices. Candidates are required to demonstrate their knowledge of AWS services and their ability to design and deploy scalable, highly available, and fault-tolerant systems on AWS.

>> Reliable SAA-C03 Test Dumps <<

SAA-C03 Latest Braindumps Free - Valid SAA-C03 Test Voucher

With the rapid development of the world economy, it has been universally accepted that a growing number of people have longed to become the social elite. The SAA-C03 latest study guide materials will be a shortcut for a lot of people who desire to be the social elite. If you try your best to prepare for the SAA-C03 Exam and get the related certification in a short time, it will be easier for you to receive the attention from many leaders of the big company like us, and it also will be very easy for many people to get a decent job in the labor market with the help of our SAA-C03 learning guide.

The AWS Certified Solutions Architect - Associate certification is recognized globally and is highly valued by IT professionals and organizations alike. It demonstrates your expertise in AWS solutions architecture and your ability to design and deploy scalable, highly available, and fault-tolerant systems on AWS. AWS is one of the leading cloud computing providers in the world, and this certification can help you advance your career and increase your earning potential.

Amazon SAA-C03 Exam consists of multiple-choice and multiple-response questions and is conducted online via the AWS Certification portal. SAA-C03 exam is 130 minutes long, and candidates are required to score at least 720 out of 1000 to pass. SAA-C03 exam covers a wide range of topics, including AWS storage solutions, compute and networking services, security, and application integration. Candidates must have a solid understanding of these topics to be successful on the exam.

Amazon AWS Certified Solutions Architect - Associate Sample Questions (Q306-Q311):

NEW QUESTION # 306

A company has a set of Linux servers running on multiple On-Demand EC2 Instances. The Audit team wants to collect and process the application log files generated from these servers for their report.

Which of the following services is best to use in this case?

- A. A single On-Demand Amazon EC2 instance for both storing and processing the log files
- B. Amazon S3 Glacier for storing the application log files and Spot EC2 Instances for processing them.
- C. Amazon S3 Glacier Deep Archive for storing the application log files and AWS ParallelCluster for processing the log files.
- D. **Amazon S3 for storing the application log files and Amazon Elastic MapReduce for processing the log files.**

Answer: D

Explanation:

Amazon EMR is a managed cluster platform that simplifies running big data frameworks, such as Apache Hadoop and Apache Spark, on AWS to process and analyze vast amounts of data. By using these frameworks and related open-source projects such as Apache Hive and Apache Pig, you can process data for analytics purposes and business intelligence workloads. Additionally, you can use Amazon EMR to transform and move large amounts of data into and out of other AWS data stores and databases such as Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB.

□ Hence, the correct answer is: Amazon S3 for storing the application log files and Amazon Elastic MapReduce for processing the log files.

The option that says: Amazon S3 Glacier for storing the application log files and Spot EC2 Instances for processing them is incorrect as Amazon S3 Glacier is used for data archive only.

The option that says: A single On-Demand Amazon EC2 instance for both storing and processing the log files is incorrect as an EC2 instance is not a recommended storage service. In addition, Amazon EC2 does not have a built-in data processing engine to process large amounts of data.

The option that says: Amazon S3 Glacier Deep Archive for storing the application log files and AWS ParallelCluster for processing the log files is incorrect because the long retrieval time of Amazon S3 Glacier Deep Archive makes this option unsuitable. Moreover, AWS ParallelCluster is just an AWS-supported open-source cluster management tool that makes it easy for you to deploy and manage High-Performance Computing (HPC) clusters on AWS. ParallelCluster uses a simple text file to model and provision all the resources needed for your HPC applications in an automated and secure manner.

References:

<http://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-what-is-emr.html>

<https://aws.amazon.com/hpc/parallelcluster/>

Check out this Amazon EMR Cheat Sheet:

<https://tutorialsdojo.com/amazon-emr/>

NEW QUESTION # 307

A company runs its critical storage application in the AWS Cloud. The application uses Amazon S3 in two AWS Regions. The company wants the application to send remote user data to the nearest S3 bucket with no public network congestion. The company also wants the application to fail over with the least amount of management of Amazon S3.

Which solution will meet these requirements?

- A. Send user data to the regional S3 endpoints closest to the user. Configure an S3 cross-account replication rule to keep the S3 buckets synchronized.
- B. Implement an active-active design between the two Regions. Configure the application to use the regional S3 endpoints closest to the user.
- C. Use an active-passive configuration with S3 Multi-Region Access Points. Create a global endpoint for each of the Regions.
- D. Set up Amazon S3 to use Multi-Region Access Points in an active-active configuration with a single global endpoint.
Configure S3 Cross-Region Replication.

Answer: D

Explanation:

To meet the requirement of low-latency global access and failover with minimal management, the best solution is to use Amazon S3 Multi-Region Access Points (MRAP) with Cross-Region Replication (CRR).

Multi-Region Access Points provide a global endpoint that automatically routes requests to the nearest AWS Region using the AWS Global Accelerator infrastructure. This avoids public internet congestion and ensures low-latency access.

When combined with S3 Cross-Region Replication, data is automatically synchronized between buckets in different Regions, enabling active-active setup.

In case of a Regional failure, S3 Multi-Region Access Points handle failover automatically, requiring no manual intervention.

Options A and C require manual management and configuration of endpoints per Region. Option B misrepresents MRAP—it is used for active-active, not active-passive.

Reference:

S3 Multi-Region Access Points

S3 Cross-Region Replication

NEW QUESTION # 308

A company needs to ingest and analyze telemetry data from vehicles at scale for machine learning and reporting.

Which solution will meet these requirements?

- A. Use Amazon Timestream for LiveAnalytics to store data points. Grant Amazon SageMaker permission to access the data. Use Amazon Athena to visualize the data.
- B. Use Amazon Timestream for LiveAnalytics to store data points. Grant Amazon SageMaker permission to access the data. Use Amazon QuickSight to visualize the data.
- C. Use Amazon DynamoDB to store data points. Use DynamoDB Connector to ingest data into Amazon EMR for processing. Use Amazon QuickSight to visualize the data.
- D. Use Amazon Neptune to store data points. Use Amazon Kinesis Data Streams to ingest data into a Lambda function for processing. Use Amazon QuickSight to visualize the data.

Answer: B

Explanation:

* Amazon Timestream is purpose-built for storing and analyzing time-series data like telemetry.

* Option A leverages Timestream, SageMaker for ML, and QuickSight for visualization, meeting all requirements with minimal complexity.

* Option B involves more complex DynamoDB-EMR integration.

* Option C uses Neptune, which is designed for graph databases, not telemetry data.

* Option D incorrectly uses Athena for visualization instead of QuickSight.

NEW QUESTION # 309

[Design Resilient Architectures]

A company is preparing to launch a public-facing web application in the AWS Cloud. The architecture consists of Amazon EC2 instances within a VPC behind an Elastic Load Balancer (ELB). A third-party service is used for the DNS. The company's solutions architect must recommend a solution to detect and protect against large-scale DDoS attacks.

Which solution meets these requirements?

- A. Enable Amazon Inspector on the EC2 instances.
- B. Enable AWS Shield and assign Amazon Route 53 to it.

- C. Enable Amazon GuardDuty on the account.
- D. Enable AWS Shield Advanced and assign the ELB to it.

Answer: D

Explanation:

<https://aws.amazon.com/shield/faqs/>

NEW QUESTION # 310

A company wants to migrate an Oracle database to AWS. The database consists of a single table that contains millions of geographic information systems (GIS) images that are high resolution and are identified by a geographic code.

When a natural disaster occurs tens of thousands of images get updated every few minutes. Each geographic code has a single image or row that is associated with it. The company wants a solution that is highly available and scalable during such events. Which solution meets these requirements MOST cost-effectively?

- A. Store the images in Amazon S3 buckets. Store geographic codes and image S3 URLs in a database table. Use Oracle running on an Amazon RDS Multi-AZ DB instance.
- B. Store the images and geographic codes in an Amazon DynamoDB table. Configure DynamoDB Accelerator (DAX) during times of high load.
- C. Store the images in Amazon S3 buckets. Use Amazon DynamoDB with the geographic code as the key and the image S3 URL as the value.
- D. Store the images and geographic codes in a database table. Use Oracle running on an Amazon RDS Multi-AZ DB instance.

Answer: C

Explanation:

Amazon S3 is a highly scalable, durable, and cost-effective object storage service that can store millions of images¹. Amazon DynamoDB is a fully managed NoSQL database that can handle high throughput and low latency for key-value and document data². By using S3 to store the images and DynamoDB to store the geographic codes and image S3 URLs, the solution can achieve high availability and scalability during natural disasters. It can also leverage DynamoDB's features such as caching, auto-scaling, and global tables to improve performance and reduce costs².

1. Store the images and geographic codes in a database table. Use Oracle running on an Amazon RDS Multi-AZ DB instance. This solution will not meet the requirement of scalability and cost-effectiveness, as Oracle is a relational database that may not handle large volumes of unstructured data such as images efficiently³. It also involves higher licensing and operational costs than S3 and DynamoDB^{1,2}.

2. Store the images and geographic codes in an Amazon DynamoDB table. Configure DynamoDB Accelerator (DAX) during times of high load. This solution will not meet the requirement of cost-effectiveness, as storing images in DynamoDB will consume more storage space and incur higher charges than storing them in S3^{1,2}. It will also require additional configuration and management of DAX clusters to handle high load.

3. Store the images in Amazon S3 buckets. Store geographic codes and image S3 URLs in a database table. Use Oracle running on an Amazon RDS Multi-AZ DB instance. This solution will not meet the requirement of scalability and cost-effectiveness, as Oracle is a relational database that may not handle high throughput and low latency for key-value data such as geographic codes efficiently³. It also involves higher licensing and operational costs than DynamoDB².

Reference URL: <https://dynobase.dev/dynamodb-vs-s3/>

NEW QUESTION # 311

.....

SAA-C03 Latest Braindumps Free: <https://www.itexamdownload.com/SAA-C03-valid-questions.html>

- Online SAA-C03 Tests □ Reliable SAA-C03 Test Labs □ Practice SAA-C03 Exams □ Search for □ SAA-C03 □ and download exam materials for free through ⇒ www.troytecdumps.com ⇐ □ Practice SAA-C03 Exams
- Pdf SAA-C03 Dumps □ Dumps SAA-C03 Collection □ SAA-C03 Exam Cram Questions □ Download { SAA-C03 } for free by simply searching on **【 www.pdfvce.com 】** □ Reliable SAA-C03 Test Sample
- New Reliable SAA-C03 Test Dumps | Pass-Sure Amazon SAA-C03 Latest Braindumps Free: AWS Certified Solutions Architect - Associate □ Search on **【 www.prepawayte.com 】** for ➡ SAA-C03 □ to obtain exam materials for free download □ Exam SAA-C03 Quizzes
- Practice SAA-C03 Exams □ Latest SAA-C03 Test Pdf □ Latest SAA-C03 Exam Discount □ Search for □ SAA-C03 □ and easily obtain a free download on **《 www.pdfvce.com 》** □ Practice SAA-C03 Exams

2026 Latest ITExamDownload SAA-C03 PDF Dumps and SAA-C03 Exam Engine Free Share: <https://drive.google.com/open?id=1FEpIqMW5qCymvrk89H-iWvDlAtO9SmL->