SAA-C03시험패스가능한인증덤프 & SAA-C03시험패스인증공부

그 외, ITDumpsKR SAA-C03 시험 문제집 일부가 지금은 무료입니다: https://drive.google.com/open?id=1D2s1wbiS2B-8gu7tzg-mQU3A0zdKwmov

ITDumpsKR을 선택함으로 100%인증시험을 패스하실 수 있습니다. 우리는Amazon SAA-C03시험의 갱신에 따라 최신의 덤프를 제공할 것입니다. ITDumpsKR에서는 무료로 24시간 온라인상담이 있으며, ITDumpsKR의 덤프로 Amazon SAA-C03시험을 패스하지 못한다면 우리는 덤프전액환불을 약속 드립니다.

Amazon SAA-C03 인증은 AWS와 함께 일하는 데 관심이있는 개인에게 필수 인증입니다. AWS 서비스와 모범 사례에 대한 포괄적 인 이해를 제공하며 IT 산업의 고용주가 높이 평가합니다.

Amazon SAA-C03 시험을 통과하기 위해서는, 후보자들은 고객 요구 사항에 기반한 안전하고 효율적이며 비용 효율적인 AWS 솔루션을 설계하고 배포할 수 있는 능력을 증명해야 합니다. 이 시험은 65개의 객관식 및 객관식 다중응답 문제들로 이루어져 있으며 130분 내에 완료해야 합니다. SAA-C03 시험의 합격 점수는 1000점 중 720점입니다. 시험을 통과한 후보자들은 전 세계적으로 인정받는 Amazon AWS Certified Solutions Architect - Associate 자격증을 획득하게 되며, AWS 솔루션을 설계하고 배포하는 능력을 증명할 수 있습니다. 이 자격증은 클라우드 컴퓨팅 분야에서 경력을 쌓고 조직에서 AWS 기반 솔루션을 다루는 전문가들에게 필수적입니다.

>> SAA-C03시험패스 가능한 인증덤프 <<

SAA-C03시험패스 인증공부, SAA-C03최신 업데이트 공부자료

ITDumpsKR 는 여러분의 it전문가 꿈을 이루어드리는 사이트 입다. ITDumpsKR는 여러분이 우리 자료로 관심 가는 인증시험에 응시하여 안전하게 자격증을 취득할 수 있도록 도와드립니다. 아직도Amazon SAA-C03인증시험으로 고민하시고 계십니까?Amazon SAA-C03인증시험가이드를 사용하실 생각은 없나요? ITDumpsKR는 여러분에 편리를 드릴 수 잇습니다. ITDumpsKR의 자료는 시험대비최고의 덤프로 시험패스는 문제없습니다. ITDumpsKR의 각종 인증시험자료는 모두기출문제와 같은 것으로 덤프보고 시험패스는 문제없습니다. ITDumpsKR의 퍼펙트한 덤프인 M crosoftSAA-C03인증시험자료의 문제와 답만 열심히 공부하면 여러분은 완전 안전히Amazon SAA-C03인증자격증을 취득하실 수 있습니다.

최신 AWS Certified Solutions Architect SAA-C03 무료샘플문제 (Q639-Q644):

질문 # 639
[Design Secure Architectures]
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 Use Amazon DynamoDB with the geographic code as the key and the image S3 URL as the value
- 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 Store geographic codes and image S3 URLs in a database table Use Oracle running on an Amazon RDS Multi-AZ DB instance.
- D. Store the images and geographic codes in a database table Use Oracle running on an Amazon RDS Multi-AZ DB instance

정답：A

설명:
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 alsoleverage DynamoDB's features such as caching, auto-scaling, and global tables to improve performance and reduce costs2.
A . Store the images and geographic codes in a database table Use Oracle running on an Amazon RDS Multi-AZ DB instance. This solution willnot 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 DynamoDB12.
C . Store the images and geographic codes in an Amazon DynamoDB table Configure DynamoDB Accelerator (DAX) during times of high load. This solution willnot meet therequirement of cost-effectiveness, as storing images in DynamoDB will consume more storage space and incur higher charges than storing them in S312. It will also require additional configuration and management of DAX clusters to handle high load.
D . 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, asOracle 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/

질문 # 640
A company wants a flexible compute solution that includes Amazon EC2 instances and AWS Fargate. The company does not want to commit to multi-year contracts.
Which purchasing option will meet these requirements MOST cost-effectively?

- A. Purchase a 1-year Compute Savings Plan with the All Upfront option.
- B. Purchase a 1-year EC2 Instance Savings Plan with the All Upfront option.
- C. Purchase a 1-year Compute Savings Plan with the Partial Upfront option.
- D. Purchase a 1-year Compute Savings Plan with the No Upfront option.

정답：D

설명:
To optimize costs for both Amazon EC2 and AWS Fargate, the best option is a Compute Savings Plan because it offers flexibility across instance families, Regions, and compute options including EC2, AWS Fargate, and AWS Lambda.
Unlike EC2 Instance Savings Plans, which apply only to specific instance families, Compute Savings Plans apply across multiple services.
Since the company does not want to commit to multi-year contracts or large upfront payments, the 1-year No Upfront Compute Savings Plan provides the greatest flexibility with no upfront capital commitment, while still offering cost savings over On-Demand pricing.
This option also aligns with cost-optimization best practices by allowing for scalability and service mix flexibility.
Reference:
* AWS Compute Savings Plans
* AWS Pricing Models

질문 # 641

A company runs game applications on AWS. The company needs to collect, visualize, and analyze telemetry data from the company's game servers. The company wants to gain insights into the behavior, performance, and health of game servers in near real time. Which solution will meet these requirements?

- A. Use Amazon Kinesis Data Streams to collect, process, and store telemetry data. Use Amazon EMR to process the data in near real time into required formats for analysis. Use Amazon Athena to analyze and visualize the data.
- B. Use Amazon DynamoDB Streams to collect and store telemetry data. Configure DynamoDB Streams to invoke AWS Lambda functions to process the data in near real time. Use Amazon Managed Grafana to visualize and analyze the data.
- C. Use Amazon Kinesis Data Streams to collect telemetry data. Use Amazon Managed Service for Apache Flink to process the data in near real time and publish custom metrics to Amazon CloudWatch. Use Amazon CloudWatch to create dashboards and alarms from the custom metrics.
- D. Use Amazon Data Firehose to collect, process, and store telemetry data in near real time. Use AWS Glue to extract, transform, and load (ETL) data from Firehose into required formats for analysis. Use Amazon QuickSight to visualize and analyze the data.

정답：C

설명:
Amazon Kinesis Data Streams is designed for low-latency ingestion of streaming data. Combined with Amazon Managed Service for Apache Flink, telemetry can be processed and aggregated in near real time. Processed metrics can be sent to Amazon CloudWatch, which natively supports creating dashboards, metrics visualization, and alarms. Firehose (B) is primarily for batch ingestion and delivery, not real-time analytics. EMR with Athena (C) introduces more complexity and is better for large-scale offline analytics. DynamoDB Streams (D) is not a fit because telemetry data is not stored in DynamoDB. Therefore, option A provides the most suitable and real-time analytics pipeline for telemetry data.

질문 # 642
A Fortune 500 company which has numerous offices and customers around the globe has hired you as their Principal Architect. You have staff and customers that upload gigabytes to terabytes of data to a centralized S3 bucket from the regional data centers, across continents, all over the world on a regular basis. At the end of the financial year, there are thousands of data being uploaded to the central S3 bucket which is in ap-southeast-2 (Sydney) region and a lot of employees are starting to complain about the slow upload times.
You were instructed by the CTO to resolve this issue as soon as possible to avoid any delays in processing their global end of financial year (EOFY) reports.
Which feature in Amazon S3 enables fast, easy, and secure transfer of your files over long distances between your client and your Amazon S3 bucket?

- A. AWS Global Accelerator
- B. Cross-Region Replication
- C. Multipart Upload
- D. Transfer Acceleration

정답：D

설명:
Amazon S3 Transfer Acceleration enables fast, easy, and secure transfer of files over long distances between your client and your Amazon S3 bucket. Transfer Acceleration leverages Amazon CloudFront's globally distributed AWS Edge Locations. As data arrives at an AWS Edge Location, data is routed to your Amazon S3 bucket over an optimized network path.
Amazon S3 Transfer Acceleration can speed up content transfers to and from Amazon S3 by as much as 50.
500% for long-distance transfer of larger objects. Customers who have either web or mobile applications with widespread users or applications hosted far away from their S3 bucket can experience long and variable upload and download speeds over the Internet.
S3 Transfer Acceleration (S3TA) reduces the variability in Internet routing, congestion and speeds that can affect transfers, and logically shortens the distance to S3 for remote applications. S3TA improves transfer performance by routing traffic through Amazon CloudFront's globally distributed Edge Locations and over AWS backbone networks, and by using network protocol optimizations.
Hence, Transfer Acceleration is the correct answer.
AWS Global Accelerator is incorrect because this service is primarily used to optimize the path from your users to your applications which improves the performance of your TCP and UDP traffic. Using Amazon S3 Transfer Acceleration is a more suitable service for this scenario.
Cross-Region Replication is incorrect because this simply enables you to automatically copy S3 objects from one bucket to another bucket that is placed in a different AWS Region or within the same Region.
Multipart Upload is incorrect because this feature simply allows you to upload a single object as a set of parts. You can upload these

object parts independently and in any order. If transmission of any part fails, you can retransmit that part without affecting other parts. After all parts of your object are uploaded, Amazon S3 assembles these parts and creates the object. In general, when your object size reaches 100 MB, you should consider using multipart uploads instead of uploading the object in a single operation.
References:
https://aws.amazon.com/s3/faqs/
https://aws.amazon.com/s3/transfer-acceleration/
Check out this Amazon S3 Cheat Sheet:
https://tutorialsdojo.com/amazon-s3/

질문 # 643
An ecommerce company is redesigning a product catalog system to handle millions of products and provide fast access to product information. The system needs to store structured product data such as product name, price, description, and category. The system also needs to store unstructured data such as high-resolution product videos and user manuals. The architecture must be highly available and must be able to handle sudden spikes in traffic during large-scale sales events.

- A. Deploy an Amazon DocumentDB (with MongoDB compatibility) cluster to store all product information, product videos, and user manuals.
- B. Store all product information, including product videos and user manuals, in Amazon DynamoDB.
- C. Use an Amazon RDS Multi-AZ deployment to store product information. Store product videos and user manuals in Amazon S3.
- D. Use Amazon DynamoDB to store product information. Store product videos and user manuals in Amazon S3.

정답： D

질문 # 644
.....

2026 ITDumpsKR 최신 SAA-C03 PDF 버전 시험 문제집과 SAA-C03 시험 문제 및 답변 무료 공유: https://drive.google.com/open?id=1D2s1wbiS2B-8gu7tzg-mQU3A0zdKwmov