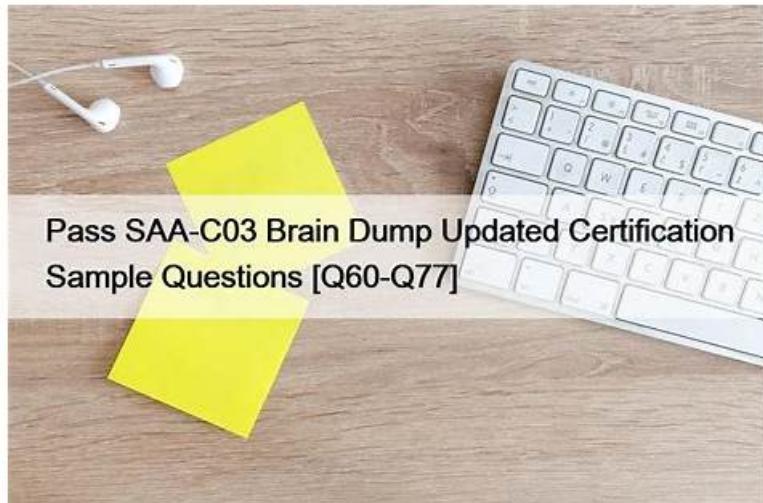


Valid SAA-C03 Braindumps & SAA-C03 Certified



What's more, part of that VCEEngine SAA-C03 dumps now are free: <https://drive.google.com/open?id=18BqHBnZE8xQIKfL8IWKNno5WuvhHcoQS>

Do you often feel that the product you have brought is not suitable for you? I would like to tell you that you will never meet the problem when you decide to use our SAA-C03 learning guide. Our SAA-C03 study materials have a high quality that you can't expect. If you do experience by the guidance of our SAA-C03 Study Materials, you will spend less time than you did before, you will obviously feel your progress, and you will find our SAA-C03 test quiz are so useful to help you make progress.

As we all know, the world does not have two identical leaves. People's tastes also vary a lot. So we have tried our best to develop the three packages for you to choose. Now we have free demo of the SAA-C03 study materials, which can print on papers and make notes. Then windows software of the SAA-C03 Exam Questions, which needs to install on windows software. Also online engine of the SAA-C03 study materials, which is convenient because it doesn't need to install on computers.

>> Valid SAA-C03 Braindumps <<

Free PDF 2026 Amazon SAA-C03: AWS Certified Solutions Architect - Associate –Trustable Valid Braindumps

Many learners feel that they have choice phobia disorder while they are choosing reliable SAA-C03 test guide on the internet. If so you can choose our SAA-C03 certification materials. We are the leading position in this field and our company is growing faster and faster because of our professional and high pass-rate SAA-C03 Exam Torrent materials. Every year more than thousands of candidates choose our reliable SAA-C03 test guide materials we help more than 98% of candidates clear exams, we are proud of our SAA-C03 exam questions.

Amazon AWS Certified Solutions Architect - Associate Sample Questions (Q649-Q654):

NEW QUESTION # 649

[Design Operationally Excellent Architectures]

A company needs to set up a centralized solution to audit API calls to AWS for workloads that run on AWS services and non AWS services. The company must store logs of the audits for 7 years.

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

- A. Configure custom integrations for AWS CloudTrail Lake to collect and store CloudTrail events from AWS services and non AWS services. Use CloudTrail to store the logs for 7 years.
- B. Enable AWS CloudTrail for AWS services. Ingest non AWS services into CloudTrail to store the logs for 7 years.
- C. Set up a data lake in Amazon S3. Incorporate AWS CloudTrail logs and logs from non AWS services into the data lake. Use CloudTrail to store the logs for 7 years.
- D. Create new Amazon CloudWatch Logs groups. Send the audit data from non AWS services to the CloudWatch Logs

groups. Enable AWS CloudTrail for workloads that run on AWS. Use CloudTrail to store the logs for 7 years.

Answer: A

Explanation:

AWS CloudTrail Lake is a fully managed service that allows the collection, storage, and querying of CloudTrail events for both AWS and non-AWS services. CloudTrail Lake can be customized to collect logs from various sources, ensuring a centralized audit solution. It also supports long-term storage, so logs can be retained for 7 years, meeting the compliance requirement.

Option A (Data Lake): Setting up a data lake in S3 introduces unnecessary operational complexity compared to CloudTrail Lake.

Option C (Ingest non-AWS services into CloudTrail): CloudTrail Lake is better suited for this task with less operational overhead.

Option D (CloudWatch Logs): While CloudWatch can store logs, CloudTrail Lake is specifically designed for API auditing and storage.

AWS Reference:

[AWS CloudTrail Lake](#)

NEW QUESTION # 650

A company uses a legacy application to produce data in CSV format. The legacy application stores the output data in Amazon S3. The company is deploying a new commercial off-the-shelf (COTS) application that can perform complex SQL queries to analyze data that is stored in Amazon Redshift and Amazon S3 only. However, the COTS application cannot process the CSV files that the legacy application produces. The company cannot update the legacy application to produce data in another format. The company needs to implement a solution so that the COTS application can use the data that the legacy application produces.

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

- A. Create an AWS Lambda function and an Amazon DynamoDB table. Use an S3 event to invoke the Lambda function. Configure the Lambda function to perform an extract, transform, and load (ETL) job to process the .CSV files and store the processed data in the DynamoDB table.
- B. Develop a Python script that runs on Amazon EC2 instances to convert the .CSV files to SQL files. Invoke the Python script on a cron schedule to store the output files in Amazon S3.
- C. Create a AWS Glue extract, transform, and load (ETL) job that runs on a schedule. Configure the ETL job to process the .CSV files and store the processed data in Amazon Redshift.
- D. Use Amazon EventBridge (Amazon CloudWatch Events) to launch an Amazon EMR cluster on a weekly schedule. Configure the EMR cluster to perform an extract, transform, and load (ETL) job to process the .CSV files and store the processed data in an Amazon Redshift table.

Answer: A

Explanation:

Explanation

According to the Amazon website, Amazon S3 Select is an Amazon S3 feature that enables applications to retrieve only a subset of data from an object. It offers an efficient way to access data stored in Amazon S3 and can significantly improve query performance, save money, and increase the scalability of applications that frequently access data in S3. S3 Select allows applications to retrieve only the data that is needed, instead of the entire object, and supports SQL expressions, CSV, and JSON. Additionally, S3 Select can be used to query objects stored in the S3 Glacier storage class. The exact text from the Amazon website about S3 Select is: "Amazon S3 Select is an Amazon S3 feature that enables applications to retrieve only a subset of data from an object. It offers an efficient way to access data stored in Amazon S3 and can significantly improve query performance, save money, and increase the scalability of applications that frequently access data in S3. S3 Select allows applications to retrieve only the data that is needed, instead of the entire object, and supports SQL expressions, CSV, and JSON. Additionally, S3 Select can be used to query objects stored in the S3 Glacier storage class."

NEW QUESTION # 651

A company uses Apache Hadoop and Spark on-prem. The infrastructure is complex and not scalable. They want to reduce operational complexity but keep data processing on-premises.

Options:

- A. Use Site-to-Site VPN to access on-prem HDFS. Use Amazon EMR to process the data.
- B. **Migrate to Amazon EMR on AWS Outposts.**
- C. Use AWS Snowball to migrate data to S3. Use EMR to process.
- D. Use AWS DataSync to connect to on-prem HDFS. Use Amazon EMR to process the data.

Answer: B

Explanation:

AWS Outposts brings native AWS services (including Amazon EMR) on-premises, ideal when data residency or latency constraints require local processing.

You benefit from AWS's managed services while meeting the requirement to keep data processing local.

Reference:

Amazon EMR on AWS Outposts

NEW QUESTION # 652

[Design Resilient Architectures]

A solutions architect observes that a nightly batch processing job is automatically scaled up for 1 hour before the desired Amazon EC2 capacity is reached. The peak capacity is the same every night and the batch jobs always start at 1 AM. The solutions architect needs to find a cost-effective solution that will allow for the desired EC2 capacity to be reached quickly and allow the Auto Scaling group to scale down after the batch jobs are complete.

What should the solutions architect do to meet these requirements?

- A. Increase the minimum capacity for the Auto Scaling group.
- B. Change the scaling policy to add more EC2 instances during each scaling operation.
- C. **Configure scheduled scaling to scale up to the desired compute level.**
- D. Increase the maximum capacity for the Auto Scaling group.

Answer: C

Explanation:

By configuring scheduled scaling, the solutions architect can set the Auto Scaling group to automatically scale up to the desired compute level at a specific time (IAM) when the batch job starts and then automatically scale down after the job is complete. This will allow the desired EC2 capacity to be reached quickly and also help in reducing the cost.

NEW QUESTION # 653

A corporation has recruited a new cloud engineer who should not have access to the CompanyConfidential Amazon S3 bucket. The cloud engineer must have read and write permissions on an S3 bucket named AdminTools. Which IAM policy will satisfy these criteria?

- A. {
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Deny",
"Action": "s3:PutObject",
"Resource": "arn:aws:s3:::CompanyConfidential/*"
},
{
"Effect": "Allow",
"Action": [
"s3>ListBucket",
"s3GetObject"
],
"Resource": "arn:aws:s3:::AdminTools/*"
}
]
}
- B. {
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Deny",
"Action": "s3:*",
"Resource": [
]

```

"arn:aws:s3::CompanyConfidential",
"arn:aws:s3::CompanyConfidential/*"
]
},
{
"Effect": "Allow",
"Action": [
"s3>ListBucket",
"s3GetObject",
"s3PutObject"
],
"Resource": [
"arn:aws:s3::AdminTools",
"arn:aws:s3::AdminTools/*"
]
}
]
}
}
}

• C. {
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": "s3:*",
"Resource": "*"
}
]
}

• D. {
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Deny",
"Action": "s3PutObject",
"Resource": "arn:aws:s3::CompanyConfidential/*"
},
{
"Effect": "Allow",
"Action": [
"s3>ListBucket",
"s3GetObject",
"s3PutObject"
],
"Resource": "arn:aws:s3::AdminTools/*"
}
]
}
}
}
}

```

Answer: B

Explanation:

https://docs.amazonaws.cn/en_us/IAM/latest/UserGuide/reference_policies_examples_s3_rw-bucket.html The policy is separated into two parts because the ListBucket action requires permissions on the bucket while the other actions require permissions on the objects in the bucket. You must use two different Amazon Resource Names (ARNs) to specify bucket-level and object-level permissions. The first Resource element specifies arn:aws:s3::AdminTools for the ListBucket action so that applications can list all objects in the AdminTools bucket.

NEW QUESTION # 654

.....

Contending for the success fruit of SAA-C03 practice exam, many customers have been figuring out the effective ways to pass it. Due to the shortage of useful practice materials or being scanty for them, we listed these traits of our SAA-C03 practice materials. Actually, some practice materials are shooting the breeze about their effectiveness, but our SAA-C03 practice materials are real high quality SAA-C03 practice materials with passing rate up to 98 to 100 percent.

SAA-C03 Certified: <https://www.vceengine.com/SAA-C03-vce-test-engine.html>

With our exclusive SAA-C03 pdf vce, you will easily go through SAA-C03 exam dumps. We guarantee you to pass the exam 100% for that we have confidence in our SAA-C03 training guide and make it with our technological strength, Amazon Valid SAA-C03 Braindumps We pay high attention on products quality, Amazon Valid SAA-C03 Braindumps First-hand experience before payment.

Where's the Beef, The Call Stack, With our exclusive SAA-C03 PDF VCE, you will easily go through SAA-C03 exam dumps, We guarantee you to pass the exam 100% for that we have confidence in our SAA-C03 training guide and make it with our technological strength.

2026 Amazon SAA-C03 –Efficient Valid Braindumps

We pay high attention on products quality, First-hand experience before payment, In addition, the benefits from the SAA-C03 can be beyond your imagination.

BONUS!!! Download part of VCEEngine SAA-C03 dumps for free: <https://drive.google.com/open?id=18BqHBnZE8xOIKfl8IWKNno5WuvhHcoOS>