SAA-C02Q&As

AWS Certified Solutions Architect - Associate (SAA-C02)

Pass Amazon SAA-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/saa-c02.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

A company hosts a data lake on AWS. The data lake consists of data in Amazon S3 and Amazon RDS for PostgreSQL. The company needs a reporting solution that provides data visualization and includes all the data sources within the data lake. Only the company\\'s management team should have full access to all the visualizations. The rest of the company should have only limited access.

Which solution will meet these requirements?

A. Create an analysis in Amazon QuickSight. Connect all the data sources and create new datasets. Publish dashboards to visualize the data. Share the dashboards with the appropriate IAM roles.

B. Create an analysis in Amazon OuickSighl. Connect all the data sources and create new datasets. Publish dashboards to visualize the data. Share the dashboards with the appropriate users and groups.

C. Create an AWS Glue table and crawler for the data in Amazon S3. Create an AWS Glue extract, transform, and load (ETL) job to produce reports. Publish the reports to Amazon S3. Use S3 bucket policies to limit access to the reports.

D. Create an AWS Glue table and crawler for the data in Amazon S3. Use Amazon Athena Federated Query to access data within Amazon RDS for PoslgreSQL. Generate reports by using Amazon Athena. Publish the reports to Amazon S3. Use S3 bucket policies to limit access to the reports.

Correct Answer: A

QUESTION 2

A company captures clickstream data from multiple websites and analyzes it using batch processing. The data is loaded nightly into Amazon Redshift and is consumed by business analysts. The company wants to move towards near-realtime data processing for timely insights. The solution should process the streaming data with minimal effort and operational overhead. Which combination of AWS services are MOST cost-effective for this solution? (Choose two.)

A. Amazon EC2

B. AWS Lambda

C. Amazon Kinesis Data Streams

D. Amazon Kinesis Data Firehose

E. Amazon Kinesis Data Analytics

Correct Answer: BD

Kinesis Data Streams and Kinesis Client Library (KCL) -Data from the data source can be continuously captured and streamed in near real-time using Kinesis Data Streams. With the Kinesis Client Library (KCL), you can build your own application that can preprocess the streaming data as they arrive and emit the data for generating incremental views and downstream analysis. Kinesis Data Analytics This service provides the easiest way to process the data that is streaming through Kinesis Data Stream or Kinesis Data Firehose using SQL. This enables customers to gain actionable insight in near real-time from the incremental stream before storing it in Amazon S3.

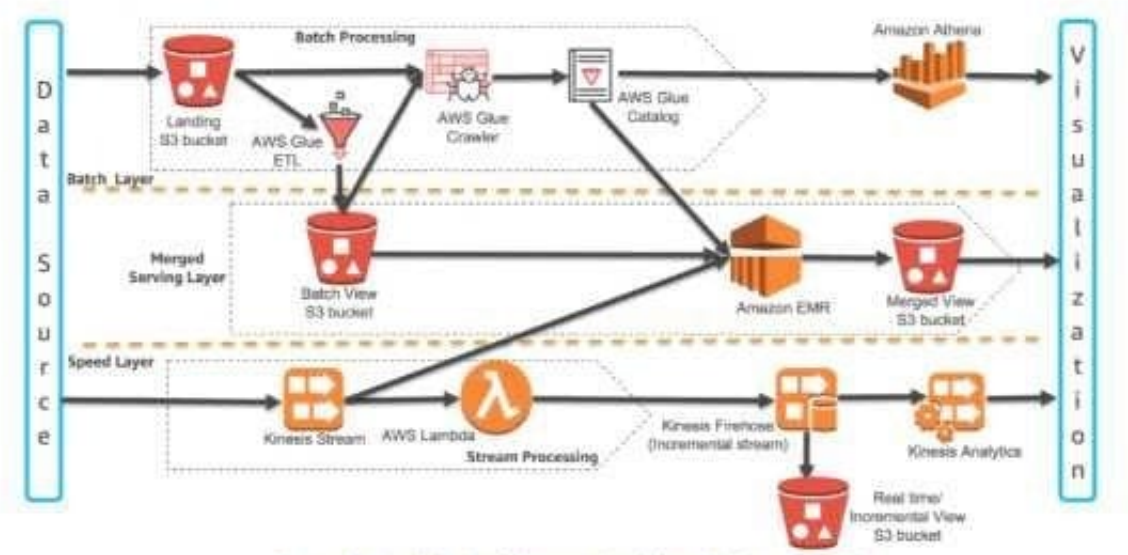

Figure 2: Lambda Architecture Building Blocks on AWS

https://d1.awsstatic.com/whitepapers/lambda-architecure-on-for-batch-aws.pdf

QUESTION 3

A group requires permissions to list an Amazon S3 bucket and delete objects from that bucket. An administrator has created the following IAM policy to provide access to the bucket and applied that policy to the group. The group is not able to delete objects in the bucket. The company follows least-privilege access rules.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "s3:ListBucket",
                "s3:DeleteObject"
            ],
            "Resource": [
                "arn:aws:s3:::bucket-name"
            ],
            "Effect": "Allow"
        }
    ]
}
```

Which statement should a solutions architect add to the policy to correct bucket access?

A.
```
"Action": [
    "s3:*Object"
],
"Resource": [
    "arn:aws:s3:::bucket-name/*"
],
"Effect": "Allow"
```

B.
```
"Action": [
    "s3:*"
],
"Resource": [
    "arn:aws:s3:::bucket-name/*"
],
"Effect": "Allow"
```

C.
```
"Action": [
    "s3:DeleteObject"
],
"Resource": [
    "arn:aws:s3:::bucket-name*"
],
"Effect": "Allow"
```

D.
```
"Action": [
    "s3:DeleteObject"
],
"Resource": [
    "arn:aws:s3:::bucket-name/*"
],
"Effect": "Allow"
```

A. Option A

B. Option B

C. Option C

D. Option D

Correct Answer: C

QUESTION 4

A company wants to deploy a shared file system for its .NET application servers and Microsoft SQL Server database running on Amazon EC2 instance with Windows Server 2016. The solution must be able to be integrated in to the corporate

Active Directory domain, be highly durable, be managed by AWS, and provided levels of throuput and IOPS.

Which solution meets these requirements?

A. Use Amazon FSx for Windows File Server

B. Use Amazon Elastic File System (Amazon EFS)

C. Use AWS Storage Gateway in file gateway mode.

D. Deploy a Windows file server on two On Demand instances across two Availability Zones.

Correct Answer: A

QUESTION 5

A company\\'s legacy application is currently relying on a single-instance Amazon RDS MySQL database without encryption Due to new compliance requirements, all existing and new data in this database must be encrypted

How should this be accomplished?

A. Create an Amazon S3 bucket with server-side encryption enabled Move all the data to Amazon S3 Delete the RDS instance

B. Enable RDS Multi-AZ mode with encryption at rest enabled Perform a failover to the standby instance to delete the original instance

C. Take a snapshot of the RDS instance Create an encrypted copy of the snapshot. Restore the RDS instance from the encrypted snapshot

D. Create an RDS read replica with encryption at rest enabled Promote the read replica to master and switch the application over to the new master Delete the old RDS instance.

Correct Answer: C

How do I encrypt Amazon RDS snapshots?

The following steps are applicable to Amazon RDS for MySQL, Oracle, SQL Server, PostgreSQL, or MariaDB.

Important: If you use Amazon Aurora, you can restore an unencrypted Aurora DB cluster snapshot to an encrypted Aurora DB cluster if you specify an AWS Key Management Service (AWS KMS) encryption key when you restore from the

unencrypted DB cluster snapshot. For more information, see Limitations of Amazon RDS Encrypted DB Instances.

Open the Amazon RDS console, and then choose Snapshots from the navigation pane.

Select the snapshot that you want to encrypt.

Under Snapshot Actions, choose Copy Snapshot.

Choose your Destination Region, and then enter your New DB Snapshot Identifier.

Change Enable Encryption to Yes.

Select your Master Key from the list, and then choose Copy Snapshot. After the snapshot status is available, the Encrypted field will be True to indicate that the snapshot is encrypted. You now have an encrypted snapshot of your DB. You

can use this encrypted DB snapshot to restore the DB instance from the DB snapshot.

https://aws.amazon.com/premiumsupport/knowledge-center/encrypt-rds-snapshots/