DAS-C01 Q&As

AWS Certified Data Analytics - Specialty (DAS-C01)

Pass Amazon DAS-C01 Exam with 100% Guarantee

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

https://www.geekcert.com/das-c01.html

100% Passing Guarantee
100% Money Back Assurance

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

QUESTION 1

A large marketing company needs to store all of its streaming logs and create near-real-time dashboards. The dashboards will be used to help the company make critical business decisions and must be highly available.

Which solution meets these requirements?

A. Store the streaming logs in Amazon S3 with replication to an S3 bucket in a different Availability Zone. Create the dashboards by using Amazon QuickSight.

B. Deploy an Amazon Redshift cluster with at least three nodes in a VPC that spans two Availability Zones. Store the streaming logs and use the Redshift cluster as a source to create the dashboards by using Amazon QuickSight.

C. Store the streaming logs in Amazon S3 with replication to an S3 bucket in a different Availability Zone. Every time a new log is added in the bucket, invoke an AWS Lambda function to update the dashboards in Amazon QuickSight.

D. Store the streaming logs in Amazon OpenSearch Service deployed across three Availability Zones and with three dedicated master nodes. Create the dashboards by using OpenSearch Dashboards.

Correct Answer: B

QUESTION 2

A media company is using Amazon QuickSight dashboards to visualize its national sales data. The dashboard is using a dataset with these fields: ID, date, time_zone, city, state, country, longitude, latitude, sales_volume, and number_of_items.

To modify ongoing campaigns, the company wants an interactive and intuitive visualization of which states across the country recorded a significantly lower sales volume compared to the national average.

Which addition to the company\\'s QuickSight dashboard will meet this requirement?

A. A geospatial color-coded chart of sales volume data across the country.

B. A pivot table of sales volume data summed up at the state level.

C. A drill-down layer for state-level sales volume data.

D. A drill through to other dashboards containing state-level sales volume data.

Correct Answer: B

Reference: https://docs.aws.amazon.com/quicksight/latest/user/pivot-table.html

QUESTION 3

An operations team notices that a few AWS Glue jobs for a given ETL application are failing. The AWS Glue jobs read a large number of small JSON files from an Amazon S3 bucket and write the data to a different S3 bucket in Apache Parquet format with no major transformations. Upon initial investigation, a data engineer notices the following error message in the History tab on the AWS Glue console: "Command Failed with Exit Code 1."

Upon further investigation, the data engineer notices that the driver memory profile of the failed jobs crosses the safe threshold of 50% usage quickly and reaches 90?5% soon after. The average memory usage across all executors continues to be less than 4%.

The data engineer also notices the following error while examining the related Amazon CloudWatch Logs.

```
# java.lang.OutOfMemoryError: Java heap space
# -XX: OnOutOfMemoryError= "kill -9 %p"
# Executing /bin/sh -c "kill -9 12039"…
```

What should the data engineer do to solve the failure in the MOST cost-effective way?

A. Change the worker type from Standard to G.2X.

B. Modify the AWS Glue ETL code to use the `groupFiles\\': `inPartition\\' feature.

C. Increase the fetch size setting by using AWS Glue dynamics frame.

D. Modify maximum capacity to increase the total maximum data processing units (DPUs) used.

Correct Answer: D

QUESTION 4

A company wants to build a real-time data processing and delivery solution for streaming data. The data is being streamed through an Amazon Kinesis data stream. The company wants to use an Apache Flink application to process the data before writing the data to another Kinesis data stream. The data must be stored in an Amazon S3 data lake every 60 seconds for further analytics.

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

A. Host the Flink application on an Amazon EMR cluster. Use Amazon Kinesis Data Firehose to write the data to Amazon S3.

B. Host the Flink application on Amazon Kinesis Data Analytics. Use AWS Glue to write the data to Amazon S3.

C. Host the Flink application on an Amazon EMR cluster. Use AWS Glue to write the data to Amazon S3.

D. Host the Flink application on Amazon Kinesis Data Analytics. Use Amazon Kinesis Data Firehose to write the data to Amazon S3.

Correct Answer: A

QUESTION 5

A market data company aggregates external data sources to create a detailed view of product consumption in different countries. The company wants to sell this data to external parties through a subscription. To achieve this goal, the company needs to make its data securely available to external parties who are also AWS users.

What should the company do to meet these requirements with the LEAST operational overhead?

A. Store the data in Amazon S3. Share the data by using presigned URLs for security.

B. Store the data in Amazon S3. Share the data by using S3 bucket ACLs.

C. Upload the data to AWS Data Exchange for storage. Share the data by using presigned URLs for security.

D. Upload the data to AWS Data Exchange for storage. Share the data by using the AWS Data Exchange sharing wizard.

Correct Answer: A