

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.leads4pass.com/das-c01.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

An IoT company is collecting data from multiple sensors and is streaming the data to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Each sensor type has its own topic, and each topic has the same number of partitions.

The company is planning to turn on more sensors. However, the company wants to evaluate which sensor types are producing the most data so that the company can scale accordingly. The company needs to know which sensor types have

the largest values for the following metrics: BytesInPerSec and MessagesInPerSec.

Which level of monitoring for Amazon MSK will meet these requirements?

- A. DEFAULT level
- B. PER_TOPIC_PER_BROKER level
- C. PER_BROKER level
- D. PER_TOPIC level

Correct Answer: D

QUESTION 2

A large fashion retailer wants to transform a source dataset to a consumable format. The retailer is building an ETL pipeline and needs to deduplicate the data because the retailer's various departments share similar customer and stock information. The retailer wants to build a data lake in Amazon S3 after the transformation and deduplication processes are completed.

Which solution MOST cost-effectively meets these requirements?

- A. Load the data into Amazon Redshift and build custom deduplication scripts by using SQL. Use the UNLOAD command in Amazon Redshift to store the data in Amazon S3.
- B. Use AWS Glue to transform the data and use FindMatches to deduplicate the data. Store the output in Amazon S3.
- C. Use Amazon EMR to transform the data. Deduplicate the data by using custom Spark SQL scripts and use EMRFS to store the output in Amazon S3.
- D. Use an Amazon Athena federated query to load the data from the sources. Build custom Athena SQL scripts to deduplicate and store the output to Amazon S3.

Correct Answer: C

QUESTION 3

A company has a partner that is supposed to put a data object into the company's Amazon S3 bucket each day. Occasionally, the partner fails to deliver the data. A data analytics specialist needs to implement a solution to automate

notifications to an Amazon Simple Notification Service (Amazon SNS) topic when the partner data is missing.

Which solution will meet this requirement with the LEAST operational overhead?

- A. Set up AWS CloudTrail to log bucket-level actions to Amazon CloudWatch. Use Amazon EventBridge to schedule an AWS Lambda function to run each day. Configure the function to publish a message to the SNS topic if no PutObject calls have been recorded in the last day.
- B. Set up S3 Event Notifications to invoke an AWS Lambda function. Configure the function to write a custom Amazon CloudWatch metric. Configure a CloudWatch alarm to publish a message to the SNS topic when the metric is zero for 1 day.
- C. Set up S3 Event Notifications to invoke an AWS Lambda function. Configure the function to write the count for the partner data to an Amazon DynamoDB table. Use Amazon EventBridge to schedule a second Lambda function to run each day. Configure the second function to verify the file counts and to publish a message to the SNS topic when data is missing.
- D. Set up S3 Event Notifications to invoke an AWS Lambda function. Configure the function to import the S3 objects into an Amazon RDS for PostgreSQL database. Configure an Amazon CloudWatch alarm to publish a message to the SNS topic when the function finishes running with any errors.

Correct Answer: A

QUESTION 4

A retail company's data analytics team recently created multiple product sales analysis dashboards for the average selling price per product using Amazon QuickSight. The dashboards were created from .csv files uploaded to Amazon S3. The team is now planning to share the dashboards with the respective external product owners by creating individual users in Amazon QuickSight. For compliance and governance reasons, restricting access is a key requirement. The product owners should view only their respective product analysis in the dashboard reports.

Which approach should the data analytics team take to allow product owners to view only their products in the dashboard?

- A. Separate the data by product and use S3 bucket policies for authorization.
- B. Separate the data by product and use IAM policies for authorization.
- C. Create a manifest file with row-level security.
- D. Create dataset rules with row-level security.

Correct Answer: D

Reference: <https://docs.aws.amazon.com/quicksight/latest/user/restrict-access-to-a-data-set-using-row-level-security.html>

QUESTION 5

A company is planning to create a data lake in Amazon S3. The company wants to create tiered storage based on access patterns and cost objectives. The solution must include support for JDBC connections from legacy clients,

metadata management that allows federation for access control, and batch-based ETL using PySpark and Scala. Operational management should be limited.

Which combination of components can meet these requirements? (Choose three.)

- A. AWS Glue Data Catalog for metadata management
- B. Amazon EMR with Apache Spark for ETL
- C. AWS Glue for Scala-based ETL
- D. Amazon EMR with Apache Hive for JDBC clients
- E. Amazon Athena for querying data in Amazon S3 using JDBC drivers
- F. Amazon EMR with Apache Hive, using an Amazon RDS with MySQL-compatible backed metastore

Correct Answer: ACE

Reference: <https://d1.awsstatic.com/whitepapers/Storage/data-lake-on-aws.pdf>

[DAS-C01 VCE Dumps](#)

[DAS-C01 Practice Test](#)

[DAS-C01 Study Guide](#)