

## DAS-C01<sup>Q&As</sup>

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**

A company receives data in CSV format from partners. The company stores this incoming raw data in Amazon S3. The company must clean the data by addressing missing values, incorrect formatting, and outlier values before the company sends the data to a reporting dashboard.

Which solution will meet these requirements with the LEAST development effort?

- A. Implement an AWS Glue ETL job. Include the data cleaning logic in the ETL job.
- B. Create an AWS Glue DataBrew recipe job. Include appropriate steps in the recipe job to detect and change specific data fields.
- C. Launch an Amazon EMR cluster. Run an Apache Spark job to read and clean the data. Include the data cleaning logic in the Spark job.
- D. Use an Amazon EMR serverless runtime. Run an Apache Spark job to read and clean the data. Include the data cleaning logic in the Spark job.

Correct Answer: A

---

**QUESTION 2**

A power utility company is deploying thousands of smart meters to obtain real-time updates about power consumption. The company is using Amazon Kinesis Data Streams to collect the data streams from smart meters. The consumer application uses the Kinesis Client Library (KCL) to retrieve the stream data. The company has only one consumer application.

The company observes an average of 1 second of latency from the moment that a record is written to the stream until the record is read by a consumer application. The company must reduce this latency to 500 milliseconds.

Which solution meets these requirements?

- A. Use enhanced fan-out in Kinesis Data Streams.
- B. Increase the number of shards for the Kinesis data stream.
- C. Reduce the propagation delay by overriding the KCL default settings.
- D. Develop consumers by using Amazon Kinesis Data Firehose.

Correct Answer: C

The KCL defaults are set to follow the best practice of polling every 1 second. This default results in average propagation delays that are typically below 1 second. Reference:

<https://docs.aws.amazon.com/streams/latest/dev/kinesis-low-latency.html>

---

**QUESTION 3**

A company using Amazon QuickSight Enterprise edition has thousands of dashboards, analyses, and datasets. The company struggles to manage and assign permissions for granting users access to various items within QuickSight. The

company wants to make it easier to implement sharing and permissions management.

Which solution should the company implement to simplify permissions management?

- A. Use QuickSight folders to organize dashboards, analyses, and datasets. Assign individual users permissions to these folders.
- B. Use QuickSight folders to organize dashboards, analyses, and datasets. Assign group permissions by using these folders.
- C. Use AWS IAM resource-based policies to assign group permissions to QuickSight items.
- D. Use QuickSight user management APIs to provision group permissions based on dashboard naming conventions.

Correct Answer: B

Reference: <https://awscli.amazonaws.com/v2/documentation/api/latest/reference/quicksight/update-folder-permissions.html>

---

#### QUESTION 4

A financial company hosts a data lake in Amazon S3 and a data warehouse on an Amazon Redshift cluster. The company uses Amazon QuickSight to build dashboards and wants to secure access from its on-premises Active Directory to Amazon QuickSight.

How should the data be secured?

- A. Use an Active Directory connector and single sign-on (SSO) in a corporate network environment.
- B. Use a VPC endpoint to connect to Amazon S3 from Amazon QuickSight and an IAM role to authenticate Amazon Redshift.
- C. Establish a secure connection by creating an S3 endpoint to connect Amazon QuickSight and a VPC endpoint to connect to Amazon Redshift.
- D. Place Amazon QuickSight and Amazon Redshift in the security group and use an Amazon S3 endpoint to connect Amazon QuickSight to Amazon S3.

Correct Answer: B

---

#### QUESTION 5

A company has collected more than 100 TB of log files in the last 24 months. The files are stored as raw text in a dedicated Amazon S3 bucket. Each object has a key of the form year-month-day\_log\_HHmmss.txt where HHmmss represents the time the log file was initially created. A table was created in Amazon Athena that points to the S3 bucket. One-time queries are run against a subset of columns in the table several times an hour.

A data analyst must make changes to reduce the cost of running these queries. Management wants a solution with minimal maintenance overhead.

Which combination of steps should the data analyst take to meet these requirements? (Choose three.)

- A. Convert the log files to Apache Avro format.

- B. Add a key prefix of the form date=year-month-day/ to the S3 objects to partition the data.
- C. Convert the log files to Apache Parquet format.
- D. Add a key prefix of the form year-month-day/ to the S3 objects to partition the data.
- E. Drop and recreate the table with the PARTITIONED BY clause. Run the ALTER TABLE ADD PARTITION statement.
- F. Drop and recreate the table with the PARTITIONED BY clause. Run the MSCK REPAIR TABLE statement.

Correct Answer: BCF

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/s3-object-key-naming-pattern/>

<https://docs.aws.amazon.com/athena/latest/ug/convert-to-columnar.html>

<https://docs.aws.amazon.com/athena/latest/ug/msck-repair-table.html>

[DAS-C01 PDF Dumps](#)

[DAS-C01 Exam Questions](#)

[DAS-C01 Braindumps](#)