

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 online retail company has an application that runs on Amazon EC2 instances launched in a VPC. The company wants to build a solution that allows the security team to collect VPC Flow Logs and analyze network traffic. Which solution MOST cost-effectively meets these requirements?

- A. Publish VPC Flow Logs to Amazon CloudWatch Logs and use Amazon Athena for analytics.
- B. Publish VPC Flow Logs to Amazon CloudWatch Logs and stream log data to an Amazon OpenSearch Service cluster for analytics.
- C. Publish VPC Flow Logs to Amazon S3 in text format and use Amazon Athena for analytics.
- D. Publish VPC Flow Logs to Amazon S3 in Apache Parquet format and use Amazon Athena for analytics.

Correct Answer: B

QUESTION 2

An online gaming company wants to read customer data from Amazon Kinesis Data Streams and deliver the data to an Amazon S3 data lake for analytics. The data contains customer_id as one of the attributes. The data consumers need the data to be partitioned by customer_id in the S3 data lake.

Which solution will meet this requirement with the LEAST effort?

- A. Create an Amazon Kinesis Data Firehose delivery stream. Use dynamic partitioning to partition the data by customer_id before delivering the data to Amazon S3.
- B. Create an AWS Glue streaming job Use the built-in map transform to partition the data by customer_id before delivering the data to Amazon S3.
- C. Create an AWS Lambda function. Use Lambda layers to partition the data by customer_id before delivering the data to Amazon S3.
- D. Create an Amazon EMR cluster. Run an Apache Spark job to automatically partition the data by customer_id before delivering the data to Amazon S3.

Correct Answer: A

QUESTION 3

A company is streaming its high-volume billing data (100 MBps) to Amazon Kinesis Data Streams. A data analyst partitioned the data on account_id to ensure that all records belonging to an account go to the same Kinesis shard and order is maintained. While building a custom consumer using the Kinesis Java SDK, the data analyst notices that, sometimes, the messages arrive out of order for account_id. Upon further investigation, the data analyst discovers the messages that are out of order seem to be arriving from different shards for the same account_id and are seen when a stream resize runs.

What is an explanation for this behavior and what is the solution?

- A. There are multiple shards in a stream and order needs to be maintained in the shard. The data analyst needs to

make sure there is only a single shard in the stream and no stream resize runs.

B. The hash key generation process for the records is not working correctly. The data analyst should generate an explicit hash key on the producer side so the records are directed to the appropriate shard accurately.

C. The records are not being received by Kinesis Data Streams in order. The producer should use the PutRecords API call instead of the PutRecord API call with the SequenceNumberForOrdering parameter.

D. The consumer is not processing the parent shard completely before processing the child shards after a stream resize. The data analyst should process the parent shard completely first before processing the child shards.

Correct Answer: D

QUESTION 4

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 5

A business intelligence (BI) engineer must create a dashboard to visualize how often certain keywords are used in relation to others in social media posts about a public figure. The BI engineer extracts the keywords from the posts and loads them into an Amazon Redshift table. The table displays the keywords and the count corresponding to each keyword.

The BI engineer needs to display the top keywords with more emphasis on the most frequently used keywords.

Which visual type in Amazon QuickSight meets these requirements?

A. Bar charts

B. Word clouds

C. Circle packing with words

D. Heat maps

Correct Answer: B

[DAS-C01 PDF Dumps](#)

[DAS-C01 VCE Dumps](#)

[DAS-C01 Exam Questions](#)