

# 2026 Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01)–High Pass-Rate Reliable Exam Pattern



P.S. Free & New Data-Engineer-Associate dumps are available on Google Drive shared by Exam4Labs:  
<https://drive.google.com/open?id=1FCJURvYsiMmsSMa5cdut-e93pudfi-S3>

Amazon Data-Engineer-Associate Exam Questions, applicants may study for and pass their desired certification exam. You may use Exam4Labs's top Data-Engineer-Associate study resources to prepare for the AWS Certified Data Engineer - Associate (DEA-C01) exam. The Amazon Data-Engineer-Associate Exam Questions offered by Exam4Labs are dependable and trustworthy sources of preparation. Exam4Labs provides valid exam questions and answers for customers, and free updates for 365 days.

If you are a person who desire to move ahead in the career with informed choice, then the Amazon training material is quite beneficial for you. The Data-Engineer-Associate pdf vce is designed to boost your personal ability in your industry. It just needs to spend 20-30 hours on the Data-Engineer-Associate Preparation, which can allow you to face with Data-Engineer-Associate actual test with confidence. You will always get the latest and updated information about Data-Engineer-Associate training pdf for study due to our one year free update policy after your purchase.

>> **Reliable Data-Engineer-Associate Exam Pattern** <<

## Data-Engineer-Associate practice materials & Data-Engineer-Associate real test & Data-Engineer-Associate test prep

Each of the Exam4Labs Amazon Data-Engineer-Associate exam dumps formats excels in its way and carries actual AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) exam questions for optimal preparation. All of these AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) practice question formats are easy to use and extremely convenient such that even newbies find them simple.

## Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q175-Q180):

### NEW QUESTION # 175

A retail company uses an Amazon Redshift data warehouse and an Amazon S3 bucket. The company ingests retail order data into the S3 bucket every day.

The company stores all order data at a single path within the S3 bucket. The data has more than 100 columns. The company ingests the order data from a third-party application that generates more than 30 files in CSV format every day. Each CSV file is between 50 and 70 MB in size.

The company uses Amazon Redshift Spectrum to run queries that select sets of columns. Users aggregate metrics based on daily orders. Recently, users have reported that the performance of the queries has degraded. A data engineer must resolve the performance issues for the queries.

Which combination of steps will meet this requirement with LEAST developmental effort? (Select TWO.)

- A. Develop an AWS Glue ETL job to convert the multiple daily CSV files to one file for each day.
- B. Load the JSON data into the Amazon Redshift table in a SUPER type column.
- C. Configure the third-party application to create the files in JSON format.
- **D. Partition the order data in the S3 bucket based on order date.**
- **E. Configure the third-party application to create the files in a columnar format.**

**Answer: D,E**

Explanation:

The performance issue in Amazon Redshift Spectrum queries arises due to the nature of CSV files, which are row-based storage formats. Spectrum is more optimized for columnar formats, which significantly improve performance by reducing the amount of data scanned. Also, partitioning data based on relevant columns like order date can further reduce the amount of data scanned, as queries can focus only on the necessary partitions.

A . Configure the third-party application to create the files in a columnar format:

Columnar formats (like Parquet or ORC) store data in a way that is optimized for analytical queries because they allow queries to scan only the columns required, rather than scanning all columns in a row-based format like CSV.

Amazon Redshift Spectrum works much more efficiently with columnar formats, reducing the amount of data that needs to be scanned, which improves query performance.

Reference:

C . Partition the order data in the S3 bucket based on order date:

Partitioning the data on columns like order date allows Redshift Spectrum to skip scanning unnecessary partitions, leading to improved query performance.

By organizing data into partitions, you minimize the number of files Spectrum has to read, further optimizing performance.

Alternatives Considered:

B (Develop an AWS Glue ETL job): While consolidating files can improve performance by reducing the number of small files (which can be inefficient to process), it adds additional ETL complexity. Switching to a columnar format (Option A) and partitioning (Option C) provides more significant performance improvements with less development effort.

D and E (JSON-related options): Using JSON format or the SUPER type in Redshift introduces complexity and isn't as efficient as the proposed solutions, especially since JSON is not a columnar format.

Amazon Redshift Spectrum Documentation

Columnar Formats and Data Partitioning in S3

## NEW QUESTION # 176

A company is developing an application that runs on Amazon EC2 instances. Currently, the data that the application generates is temporary. However, the company needs to persist the data, even if the EC2 instances are terminated.

A data engineer must launch new EC2 instances from an Amazon Machine Image (AMI) and configure the instances to preserve the data.

Which solution will meet this requirement?

- **A. Launch new EC2 instances by using an AMI that is backed by an EC2 instance store volume. Attach an Amazon Elastic Block Store (Amazon EBS) volume to contain the application data. Apply the default settings to the EC2 instances.**
- B. Launch new EC2 instances by using an AMI that is backed by an EC2 instance store volume that contains the application data. Apply the default settings to the EC2 instances.
- C. Launch new EC2 instances by using an AMI that is backed by a root Amazon Elastic Block Store (Amazon EBS) volume that contains the application data. Apply the default settings to the EC2 instances.
- D. Launch new EC2 instances by using an AMI that is backed by an Amazon Elastic Block Store (Amazon EBS) volume. Attach an additional EC2 instance store volume to contain the application data. Apply the default settings to the EC2 instances.

**Answer: A**

Explanation:

Amazon EC2 instances can use two types of storage volumes: instance store volumes and Amazon EBS volumes. Instance store volumes are ephemeral, meaning they are only attached to the instance for the duration of its life cycle. If the instance is stopped, terminated, or fails, the data on the instance store volume is lost. Amazon EBS volumes are persistent, meaning they can be detached from the instance and attached to another instance, and the data on the volume is preserved. To meet the requirement of persisting the data even if the EC2 instances are terminated, the data engineer must use Amazon EBS volumes to store the application data. The solution is to launch new EC2 instances by using an AMI that is backed by an EC2 instance store volume, which is the default option for most AMIs. Then, the data engineer must attach an Amazon EBS volume to each instance and configure the application to write the data to the EBS volume. This way, the data will be saved on the EBS volume and can be accessed by another instance if

needed. The data engineer can apply the default settings to the EC2 instances, as there is no need to modify the instance type, security group, or IAM role for this solution. The other options are either not feasible or not optimal. Launching new EC2 instances by using an AMI that is backed by an EC2 instance store volume that contains the application data (option A) or by using an AMI that is backed by a root Amazon EBS volume that contains the application data (option B) would not work, as the data on the AMI would be outdated and overwritten by the new instances. Attaching an additional EC2 instance store volume to contain the application data (option D) would not work, as the data on the instance store volume would be lost if the instance is terminated.

Reference:

Amazon EC2 Instance Store

Amazon EBS Volumes

AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 2: Data Store Management, Section 2.1: Amazon EC2

### NEW QUESTION # 177

A data engineer must manage the ingestion of real-time streaming data into AWS. The data engineer wants to perform real-time analytics on the incoming streaming data by using time-based aggregations over a window of up to 30 minutes. The data engineer needs a solution that is highly fault tolerant.

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

- A. Use Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics) to analyze the data that might occasionally contain duplicates by using multiple types of aggregations.
- B. Use an AWS Lambda function that includes both the business and the analytics logic to perform aggregations for a tumbling window of up to 30 minutes, based on the event timestamp.
- **C. Use an AWS Lambda function that includes both the business and the analytics logic to perform time-based aggregations over a window of up to 30 minutes for the data in Amazon Kinesis Data Streams.**
- D. Use Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics) to analyze the data by using multiple types of aggregations to perform time-based analytics over a window of up to 30 minutes.

**Answer: C**

Explanation:

This solution meets the requirements of managing the ingestion of real-time streaming data into AWS and performing real-time analytics on the incoming streaming data with the least operational overhead. Amazon Managed Service for Apache Flink is a fully managed service that allows you to run Apache Flink applications without having to manage any infrastructure or clusters. Apache Flink is a framework for stateful stream processing that supports various types of aggregations, such as tumbling, sliding, and session windows, over streaming data. By using Amazon Managed Service for Apache Flink, you can easily connect to Amazon Kinesis Data Streams as the source and sink of your streaming data, and perform time-based analytics over a window of up to 30 minutes. This solution is also highly fault tolerant, as Amazon Managed Service for Apache Flink automatically scales, monitors, and restarts your Flink applications in case of failures. Reference:

Amazon Managed Service for Apache Flink

Apache Flink

Window Aggregations in Flink

### NEW QUESTION # 178

A company wants to migrate an application and an on-premises Apache Kafka server to AWS. The application processes incremental updates that an on-premises Oracle database sends to the Kafka server. The company wants to use the replatform migration strategy instead of the refactor strategy.

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

- A. Amazon Managed Streaming for Apache Kafka (Amazon MSK) provisioned cluster
- B. Amazon Kinesis Data Streams
- **C. Amazon Managed Streaming for Apache Kafka (Amazon MSK) Serverless**
- D. Amazon Data Firehose

**Answer: C**

Explanation:

\* Problem Analysis:

\* The company needs to migrate both an application and an on-premises Apache Kafka server to AWS.

\* Incremental updates from an on-premises Oracle database are processed by Kafka.

- \* The solution must follow a replatform migration strategy, prioritizing minimal changes and low management overhead.
  - \* Key Considerations:
  - \* Replatform Strategy: This approach keeps the application and architecture as close to the original as possible, reducing the need for refactoring.
  - \* The solution must provide a managed Kafka service to minimize operational burden.
  - \* Low overhead solutions like serverless services are preferred.
  - \* Solution Analysis:
  - \* Option A: Kinesis Data Streams
  - \* Kinesis Data Streams is an AWS-native streaming service but is not a direct substitute for Kafka.
  - \* This option would require significant application refactoring, which does not align with the replatform strategy.
  - \* Option B: MSK Provisioned Cluster
  - \* Managed Kafka service with fully configurable clusters.
  - \* Provides the same Kafka APIs but requires cluster management (e.g., scaling, patching), increasing management overhead.
  - \* Option C: Amazon Kinesis Data Firehose
  - \* Kinesis Data Firehose is designed for data delivery rather than real-time streaming and processing.
  - \* Not suitable for Kafka-based applications.
  - \* Option D: MSK Serverless
  - \* MSK Serverless eliminates the need for cluster management while maintaining compatibility with Kafka APIs.
  - \* Automatically scales based on workload, reducing operational overhead.
  - \* Ideal for replatform migrations, as it requires minimal changes to the application.
  - \* Final Recommendation:
  - \* Amazon MSK Serverless is the best solution for migrating the Kafka server and application with minimal changes and the least management overhead.
- :
- Amazon MSK Serverless Overview
- Comparison of Amazon MSK and Kinesis

#### NEW QUESTION # 179

A data engineer needs to create a new empty table in Amazon Athena that has the same schema as an existing table named old-table.

Which SQL statement should the data engineer use to meet this requirement?

- A. 

```
CREATE TABLE new_table (COPY old_table);
```
- B. 

```
CREATE TABLE new_table AS SELECT * FROM old_table;
```
- C. 

```
CREATE TABLE new_table AS (SELECT * FROM old_table) WITH NO DATA;
```
- D. 

```
INSERT INTO new_table SELECT * FROM old_table;
```

**Answer: C**

Explanation:

- \* Problem Analysis:
- \* The goal is to create a new empty table in Athena with the same schema as an existing table (old\_table).
- \* The solution must avoid copying any data.
- \* Key Considerations:
- \* CREATE TABLE AS (CTAS) is commonly used in Athena for creating new tables based on an existing table.
- \* Adding the WITH NO DATA clause ensures only the schema is copied, without transferring any data.
- \* Solution Analysis:
- \* Option A: Copies both schema and data. Does not meet the requirement for an empty table.
- \* Option B: Inserts data into an existing table, which does not create a new table.
- \* Option C: Creates an empty table but does not copy the schema.
- \* Option D: Creates a new table with the same schema and ensures it is empty by using WITH NO DATA.
- \* Final Recommendation:

\* Use D. CREATE TABLE new\_table AS (SELECT \* FROM old\_table) WITH NO DATA to create an empty table with the same schema.

Athena CTAS Queries

CREATE TABLE Statement in Athena

## NEW QUESTION # 180

.....

It is a common sense that in terms of a kind of AWS Certified Data Engineer - Associate (DEA-C01) test torrent, the pass rate would be the best advertisement, since only the pass rate can be the most powerful evidence to show whether the Data-Engineer-Associate guide torrent is effective and useful or not. We are so proud to tell you that according to the statistics from the feedback of all of our customers, the pass rate among our customers who prepared for the exam under the guidance of our AWS Certified Data Engineer - Associate (DEA-C01) test torrent has reached as high as 98% to 100%, which definitely marks the highest pass rate in the field. Therefore, the Data-Engineer-Associate Guide Torrent compiled by our company is definitely will be the most sensible choice for you.

**Valid Data-Engineer-Associate Exam Test:** <https://www.exam4labs.com/Data-Engineer-Associate-practice-torrent.html>

Amazon Reliable Data-Engineer-Associate Exam Pattern Economic freedom brings great happiness to them, The high degree of credibility about our Data-Engineer-Associate exam dumps can provide you 100% guarantee, Our Data-Engineer-Associate latest training material supports quickly download after you pay for it, Once you have installed the Amazon Data-Engineer-Associate practice materials, you can quickly involve yourself in studying, We have after-service for you after buying Data-Engineer-Associate exam dumps, if you have any question, you can contact us by email, and we will give you reply as soon as possible.

The area you define is called a selection, and all Data-Engineer-Associate the tools you use to make the selections are known as selection tools, Traditionally, these groups have blocked roads, spiked old growth trees Interactive Data-Engineer-Associate EBook to prevent cutting, and conducted other visible radical activities to publicize their beliefs.

## Pass Guaranteed 2026 Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01) Perfect Reliable Exam Pattern

Economic freedom brings great happiness to them, The high degree of credibility about our Data-Engineer-Associate Exam Dumps can provide you 100% guarantee, Our Data-Engineer-Associate latest training material supports quickly download after you pay for it.

Once you have installed the Amazon Data-Engineer-Associate practice materials, you can quickly involve yourself in studying, We have after-service for you after buying Data-Engineer-Associate exam dumps, if you have any question, you can contact us by email, and we will give you reply as soon as possible.

- Data-Engineer-Associate Dumps Download ☐ Reliable Data-Engineer-Associate Exam Simulations ☐ Data-Engineer-Associate Exam Objectives Pdf ☐ Search for ➡ Data-Engineer-Associate ☐ on ✓ [www.verifiedumps.com](http://www.verifiedumps.com) ☒ ☐ immediately to obtain a free download ☐ Data-Engineer-Associate New Test Bootcamp
- High Pass-Rate Reliable Data-Engineer-Associate Exam Pattern - Trustworthy Data-Engineer-Associate Exam Tool Guarantee Purchasing Safety ☐ Search for ▷ Data-Engineer-Associate ◁ and obtain a free download on 「 [www.pdfvce.com](http://www.pdfvce.com) 」 ☐ Data-Engineer-Associate Pdf Pass Leader
- Exam Data-Engineer-Associate Tips ☐ Certification Data-Engineer-Associate Sample Questions ☐ Data-Engineer-Associate Exam Objectives Pdf ☐ Open ☀ [www.exam4labs.com](http://www.exam4labs.com) ☐ ☀ ☐ enter ➡ Data-Engineer-Associate ☐ and obtain a free download ☐ Data-Engineer-Associate New Test Bootcamp
- Data-Engineer-Associate Guide Torrent - Data-Engineer-Associate Exam Prep - Data-Engineer-Associate Pass Rate ☐ Search on [ [www.pdfvce.com](http://www.pdfvce.com) ] for 「 Data-Engineer-Associate 」 to obtain exam materials for free download ☐ ☐ Certification Data-Engineer-Associate Sample Questions
- Three User-Friendly Formats of [www.troytecdumps.com](http://www.troytecdumps.com) Amazon Data-Engineer-Associate Updated Practice Materials ☎ Open website ⇒ [www.troytecdumps.com](http://www.troytecdumps.com) ⇐ and search for 《 Data-Engineer-Associate 》 for free download ☐ Data-Engineer-Associate Latest Brindumps Ebook
- Data-Engineer-Associate Pdf Pass Leader ☐ Data-Engineer-Associate New Test Bootcamp ☐ Data-Engineer-Associate Valid Dumps Free ☐ Search for ( Data-Engineer-Associate ) and download exam materials for free through ► [www.pdfvce.com](http://www.pdfvce.com) ◀ ☐ Valid Data-Engineer-Associate Exam Testking
- Free PDF Quiz 2026 Data-Engineer-Associate: High-quality Reliable AWS Certified Data Engineer - Associate (DEA-C01) Exam Pattern ☐ ⇒ [www.practicevce.com](http://www.practicevce.com) ⇐ is best website to obtain ► Data-Engineer-Associate ◀ for free download ☐ ☐ Data-Engineer-Associate Learning Materials

