

# 2026 Data-Engineer-Associate Valid Dump | High-quality AWS Certified Data Engineer - Associate (DEA-C01) 100% Free Pass Exam



DOWNLOAD the newest PrepAwayExam Data-Engineer-Associate PDF dumps from Cloud Storage for free:  
<https://drive.google.com/open?id=15NAiMaqIBC533LIIYwGTaJezOh41Qpmz>

The web-based AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) practice test software can be used through browsers like Firefox, Safari, and Google Chrome. The customers don't need to download or install any excessive plugins or software in order to use the web-based AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) practice exam format. The web-based AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) practice test software format is supported by different operating systems like Mac, iOS, Linux, Windows, and Android.

In our study, we found that many people have the strongest ability to use knowledge for a period of time at the beginning of their knowledge. As time goes on, memory fades. Our Data-Engineer-Associate study materials are designed to help users consolidate what they have learned, will add to the instant of many training, the user can test their learning effect in time after finished the part of the learning content, have a special set of wrong topics in our Data-Engineer-Associate Study Materials, enable users to find their weak spot of knowledge in this function, iterate through constant practice, finally reach a high success rate.

>> **Data-Engineer-Associate Valid Dump** <<

## 2026 Trustable Data-Engineer-Associate Valid Dump | 100% Free Data-Engineer-Associate Pass Exam

Our Data-Engineer-Associate study materials have included all significant knowledge about the exam. So you do not need to pick out the important points by yourself. Also, our Data-Engineer-Associate practice engine can greatly shorten your preparation time of the exam. So you just need our Data-Engineer-Associate learning questions to help you get the certificate. You will find that the coming exam is just a piece of cake in front of you and you will pass it with ease.

### Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q16-Q21):

#### NEW QUESTION # 16

A company currently uses a provisioned Amazon EMR cluster that includes general purpose Amazon EC2 instances. The EMR cluster uses EMR managed scaling between one to five task nodes for the company's long-running Apache Spark extract, transform, and load (ETL) job. The company runs the ETL job every day. When the company runs the ETL job, the EMR cluster quickly scales up to five nodes. The EMR cluster often reaches maximum CPU usage, but the memory usage remains under 30%.

The company wants to modify the EMR cluster configuration to reduce the EMR costs to run the daily ETL job. Which solution will meet these requirements MOST cost-effectively?

- A. Switch the task node type from general purpose EC2 instances to compute optimized EC2 instances.
- B. Reduce the scaling cooldown period for the provisioned EMR cluster.
- C. Increase the maximum number of task nodes for EMR managed scaling to 10.
- D. Change the task node type from general purpose EC2 instances to memory optimized EC2 instances.

**Answer: A**

Explanation:

The company's Apache Spark ETL job on Amazon EMR uses high CPU but low memory, meaning that compute-optimized EC2 instances would be the most cost-effective choice. These instances are designed for high-performance compute applications, where CPU usage is high, but memory needs are minimal, which is exactly the case here.

\* Compute Optimized Instances:

\* Compute-optimized instances, such as the C5 series, provide a higher ratio of CPU to memory, which is more suitable for jobs with high CPU usage and relatively low memory consumption.

\* Switching from general-purpose EC2 instances to compute-optimized instances can reduce costs while improving performance, as these instances are optimized for workloads like Spark jobs that perform a lot of computation.

### NEW QUESTION # 17

A company stores sensitive data in an Amazon Redshift table. The company needs to give specific users the ability to access the sensitive data. The company must not create duplication in the data.

Customer support users must be able to see the last four characters of the sensitive data. Audit users must be able to see the full value of the sensitive data. No other users can have the ability to access the sensitive information.

Which solution will meet these requirements?

- A. Enable metadata security on the Redshift cluster. Create IAM users and IAM roles for the customer support users and the audit users. Grant the IAM users and IAM roles permissions to view the metadata in the Redshift cluster.
- B. Create a dynamic data masking policy to allow access based on each user role. Create IAM roles that have specific access permissions. Attach the masking policy to the column that contains sensitive data.
- C. Create an AWS Glue job to redact the sensitive data and to load the data into a new Redshift table.
- D. Create a row-level security policy to allow access based on each user role. Create IAM roles that have specific access permissions. Attach the security policy to the table.

**Answer: B**

Explanation:

Amazon Redshift supports dynamic data masking, which enables you to limit sensitive data visibility to specific users and roles without duplicating the data. This approach supports showing only parts of a column's values (e.g., last four digits) and full visibility for authorized roles (e.g., auditors).

"With dynamic data masking, you can control how much sensitive data a user sees in query results without changing the data in the table."

-Ace the AWS Certified Data Engineer - Associate Certification - version 2 - apple.pdf IAM roles are used to associate users with the appropriate masking rules, keeping security tight and avoiding the creation of duplicate data views or tables.

### NEW QUESTION # 18

A company analyzes data in a data lake every quarter to perform inventory assessments. A data engineer uses AWS Glue DataBrew to detect any personally identifiable information (PII) about customers within the data.

The company's privacy policy considers some custom categories of information to be PII. However, the categories are not included in standard DataBrew data quality rules.

The data engineer needs to modify the current process to scan for the custom PII categories across multiple datasets within the data lake.

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

- A. Develop custom Python scripts to detect the custom PII categories. Call the scripts from DataBrew.
- B. Manually review the data for custom PII categories.
- C. Implement regex patterns to extract PII information from fields during extract transform, and load (ETL) operations into the data lake.

- D. Implement custom data quality rules in Data Brew. Apply the custom rules across datasets.

**Answer: D**

Explanation:

The data engineer needs to detect custom categories of PII within the data lake using AWS Glue DataBrew.

While DataBrew provides standard data quality rules, the solution must support custom PII categories.

\* Option B: Implement custom data quality rules in DataBrew. Apply the custom rules across datasets. This option is the most efficient because DataBrew allows the creation of custom data quality rules that can be applied to detect specific data patterns, including custom PII categories. This approach minimizes operational overhead while ensuring that the specific privacy requirements are met.

Options A, C, and D either involve manual intervention or developing custom scripts, both of which increase operational effort compared to using DataBrew's built-in capabilities.

References:

\* AWS Glue DataBrew Documentation

**NEW QUESTION # 19**

A data engineer needs to create an AWS Lambda function that converts the format of data from .csv to Apache Parquet. The Lambda function must run only if a user uploads a .csv file to an Amazon S3 bucket.

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

- A. Create an S3 event notification that has an event type of s3:ObjectCreated:\*. Use a filter rule to generate notifications only when the suffix includes .csv. Set an Amazon Simple Notification Service (Amazon SNS) topic as the destination for the event notification. Subscribe the Lambda function to the SNS topic.
- B. Create an S3 event notification that has an event type of s3:ObjectCreated:\*. Use a filter rule to generate notifications only when the suffix includes .csv. Set the Amazon Resource Name (ARN) of the Lambda function as the destination for the event notification.
- C. Create an S3 event notification that has an event type of s3:\*. Use a filter rule to generate notifications only when the suffix includes .csv. Set the Amazon Resource Name (ARN) of the Lambda function as the destination for the event notification.
- D. Create an S3 event notification that has an event type of s3:ObjectTagging: for objects that have a tag set to .csv. Set the Amazon Resource Name (ARN) of the Lambda function as the destination for the event notification.

**Answer: B**

Explanation:

Option A is the correct answer because it meets the requirements with the least operational overhead. Creating an S3 event notification that has an event type of s3:ObjectCreated: will trigger the Lambda function whenever a new object is created in the S3 bucket. Using a filter rule to generate notifications only when the suffix includes .csv will ensure that the Lambda function only runs for .csv files. Setting the ARN of the Lambda function as the destination for the event notification will directly invoke the Lambda function without any additional steps.

Option B is incorrect because it requires the user to tag the objects with .csv, which adds an extra step and increases the operational overhead.

Option C is incorrect because it uses an event type of s3:\*, which will trigger the Lambda function for any S3 event, not just object creation. This could result in unnecessary invocations and increased costs.

Option D is incorrect because it involves creating and subscribing to an SNS topic, which adds an extra layer of complexity and operational overhead.

AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 3: Data Ingestion and Transformation, Section 3.2: S3 Event Notifications and Lambda Functions, Pages 67-69 Building Batch Data Analytics Solutions on AWS, Module 4: Data Transformation, Lesson 4.2: AWS Lambda, Pages 4-8 AWS Documentation Overview, AWS Lambda Developer Guide, Working with AWS Lambda Functions, Configuring Function Triggers, Using AWS Lambda with Amazon S3, Pages 1-5

**NEW QUESTION # 20**

A manufacturing company collects sensor data from its factory floor to monitor and enhance operational efficiency. The company uses Amazon Kinesis Data Streams to publish the data that the sensors collect to a data stream. Then Amazon Kinesis Data Firehose writes the data to an Amazon S3 bucket.

The company needs to display a real-time view of operational efficiency on a large screen in the manufacturing facility.

Which solution will meet these requirements with the LOWEST latency?

- A. Use Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics) to process the

sensor data. Create a new Data Firehose delivery stream to publish data directly to an Amazon Timestream database. Use the Timestream database as a source to create an Amazon QuickSight dashboard.

- B. Use Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics) to process the sensor data. Use a connector for Apache Flink to write data to an Amazon Timestream database. Use the Timestream database as a source to create a Grafana dashboard.
- C. Use AWS Glue bookmarks to read sensor data from the S3 bucket in real time. Publish the data to an Amazon Timestream database. Use the Timestream database as a source to create a Grafana dashboard.
- D. Configure the S3 bucket to send a notification to an AWS Lambda function when any new object is created. Use the Lambda function to publish the data to Amazon Aurora. Use Aurora as a source to create an Amazon QuickSight dashboard.

#### Answer: A

Explanation:

This solution will meet the requirements with the lowest latency because it uses Amazon Managed Service for Apache Flink to process the sensor data in real time and write it to Amazon Timestream, a fast, scalable, and serverless time series database. Amazon Timestream is optimized for storing and analyzing time series data, such as sensor data, and can handle trillions of events per day with millisecond latency. By using Amazon Timestream as a source, you can create an Amazon QuickSight dashboard that displays a real-time view of operational efficiency on a large screen in the manufacturing facility. Amazon QuickSight is a fully managed business intelligence service that can connect to various data sources, including Amazon Timestream, and provide interactive visualizations and insights<sup>123</sup>.

The other options are not optimal for the following reasons:

A. Use Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics) to process the sensor data. Use a connector for Apache Flink to write data to an Amazon Timestream database.

Use the Timestream database as a source to create a Grafana dashboard. This option is similar to option C, but it uses Grafana instead of Amazon QuickSight to create the dashboard. Grafana is an open source visualization tool that can also connect to Amazon Timestream, but it requires additional steps to set up and configure, such as deploying a Grafana server on Amazon EC2, installing the Amazon Timestream plugin, and creating an IAM role for Grafana to access Timestream. These steps can increase the latency and complexity of the solution.

B. Configure the S3 bucket to send a notification to an AWS Lambda function when any new object is created. Use the Lambda function to publish the data to Amazon Aurora. Use Aurora as a source to create an Amazon QuickSight dashboard. This option is not suitable for displaying a real-time view of operational efficiency, as it introduces unnecessary delays and costs in the data pipeline. First, the sensor data is written to an S3 bucket by Amazon Kinesis Data Firehose, which can have a buffering interval of up to 900 seconds.

Then, the S3 bucket sends a notification to a Lambda function, which can incur additional invocation and execution time. Finally, the Lambda function publishes the data to Amazon Aurora, a relational database that is not optimized for time series data and can have higher storage and performance costs than Amazon Timestream.

D. Use AWS Glue bookmarks to read sensor data from the S3 bucket in real time. Publish the data to an Amazon Timestream database. Use the Timestream database as a source to create a Grafana dashboard. This option is also not suitable for displaying a real-time view of operational efficiency, as it uses AWS Glue bookmarks to read sensor data from the S3 bucket. AWS Glue bookmarks are a feature that helps AWS Glue jobs and crawlers keep track of the data that has already been processed, so that they can resume from where they left off. However, AWS Glue jobs and crawlers are not designed for real-time data processing, as they can have a minimum frequency of 5 minutes and a variable start-up time. Moreover, this option also uses Grafana instead of Amazon QuickSight to create the dashboard, which can increase the latency and complexity of the solution.

1: Amazon Managed Streaming for Apache Flink

2: Amazon Timestream

3: Amazon QuickSight

Analyze data in Amazon Timestream using Grafana

Amazon Kinesis Data Firehose

Amazon Aurora

AWS Glue Bookmarks

AWS Glue Job and Crawler Scheduling

#### NEW QUESTION # 21

.....

PrepAwayExam is committed to offering the best value for your investment. For this purpose, PrepAwayExam is offering a 100 percent Data-Engineer-Associate Exams passing money-back guarantee. Whether you buy AWS Certified Data Engineer - Associate (DEA-C01) Data-Engineer-Associate PdfDumps file, desktop practice test software, and web-based practice test software or all formats, your investment is secured.

**Data-Engineer-Associate Pass Exam:** <https://www.prepawayexam.com/Amazon/braindumps.Data-Engineer-Associate.ete.file.html>

The PrepAwayExam Data-Engineer-Associate Pass Exam help desk is the place to go if you have any questions or problems, Learning and practicing Data-Engineer-Associate dumps PDF to boost you preparations and make your preparation 100% effective, We also have strict requirements for the questions and answers of Data-Engineer-Associate exam materials, we ensure you that the Data-Engineer-Associate training materials are most useful tool, which can help you pass the exam just one time, We can make sure that the PDF version of the Data-Engineer-Associate test questions will be very convenient for all people.

Birders can search for current bird species to learn where Data-Engineer-Associate they've been seen, or learn which species have been spotted in particular locations. Keeping lines of text together.

The PrepAwayExam help desk is the place to go if you have any questions or problems, Learning and practicing Data-Engineer-Associate Dumps PDF to boost you preparations and make your preparation 100% effective.

## **Providing You First-grade Data-Engineer-Associate Valid Dump with 100% Passing Guarantee**

We also have strict requirements for the questions and answers of Data-Engineer-Associate exam materials, we ensure you that the Data-Engineer-Associate training materials are most useful tool, which can help you pass the exam just one time.

We can make sure that the PDF version of the Data-Engineer-Associate test questions will be very convenient for all people, And the data also proved and tested the high-quality of our Data-Engineer-Associate practice guide.

myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, learning.pcopro.com, Disposable vapes

P.S. Free & New Data-Engineer-Associate dumps are available on Google Drive shared by PrepAwayExam:  
<https://drive.google.com/open?id=15NAiMaqIBC533LIIYwGTaJezOh41Qpmz>