

# 2026 High Pass-Rate Detail Data-Engineer-Associate Explanation Help You Pass Data-Engineer-Associate Easily



BTW, DOWNLOAD part of RealVCE Data-Engineer-Associate dumps from Cloud Storage: [https://drive.google.com/open?id=1Xd4nq\\_zMBwEvq1ZPitPUdd-IK0MFvL3T](https://drive.google.com/open?id=1Xd4nq_zMBwEvq1ZPitPUdd-IK0MFvL3T)

Many candidates who take the qualifying exams are not aware of our products and are not guided by our systematic guidance, and our users are much superior to them. In similar educational products, the Data-Engineer-Associate quiz guide is absolutely the most practical. Also, from an economic point of view, our AWS Certified Data Engineer - Associate (DEA-C01) exam dumps is priced reasonable, so the Data-Engineer-Associate test material is very responsive to users, user satisfaction is also leading the same products. So economical and practical learning platform, I believe that will be able to meet the needs of users. Users can deeply depend on our AWS Certified Data Engineer - Associate (DEA-C01) exam dumps when you want to get a qualification. There may be many problems and difficulties you will face, but believe in our AWS Certified Data Engineer - Associate (DEA-C01) exam dumps if you want to be the next beneficiary, our Data-Engineer-Associate Quiz guide is not only superior in price than any other makers in the educational field , but also are distinctly superior in the quality of our products.

Passing the AWS Certified Data Engineer - Associate (DEA-C01) Data-Engineer-Associate exam is your best career opportunity. The rich experience with relevant certificates is important for enterprises to open up a series of professional vacancies for your choices. Our Amazon Data-Engineer-Associate learning quiz bank and learning materials look up the latest Data-Engineer-Associate questions and answers based on the topics you choose.

>> Detail Data-Engineer-Associate Explanation <<

## **Free PDF Amazon - Data-Engineer-Associate –High Pass-Rate Detail Explanation**

It is time for you to plan your life carefully. After all, you have to make money by yourself. If you want to find a desirable job, you must rely on your ability to get the job. Now, our Data-Engineer-Associate study materials will help you master the popular skills in the office. Believe it or not, our Data-Engineer-Associate Study Materials will relieve you from poverty. It is important to make large amounts of money in modern society.

## Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q88-Q93):

## NEW QUESTION # 88

A data engineer must orchestrate a series of Amazon Athena queries that will run every day. Each query can run for more than 15 minutes.

Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

- A. Create an AWS Step Functions workflow and add two states. Add the first state before the Lambda function. Configure the second state as a Wait state to periodically check whether the Athena query has finished using the Athena Boto3

- get\_query\_execution API call. Configure the workflow to invoke the next query when the current query has finished running.
- B. Use an AWS Glue Python shell job and the Athena Boto3 client start\_query\_execution API call to invoke the Athena queries programmatically.
- C. Use an AWS Lambda function and the Athena Boto3 client start\_query\_execution API call to invoke the Athena queries programmatically.**
- D. Use Amazon Managed Workflows for Apache Airflow (Amazon MWAA) to orchestrate the Athena queries in AWS Batch.
- E. Use an AWS Glue Python shell script to run a sleep timer that checks every 5 minutes to determine whether the current Athena query has finished running successfully. Configure the Python shell script to invoke the next query when the current query has finished running.

**Answer: A,C**

Explanation:

Option A and B are the correct answers because they meet the requirements most cost-effectively. Using an AWS Lambda function and the Athena Boto3 client start\_query\_execution API call to invoke the Athena queries programmatically is a simple and scalable way to orchestrate the queries. Creating an AWS Step Functions workflow and adding two states to check the query status and invoke the next query is a reliable and efficient way to handle the long-running queries.

Option C is incorrect because using an AWS Glue Python shell job to invoke the Athena queries programmatically is more expensive than using a Lambda function, as it requires provisioning and running a Glue job for each query.

Option D is incorrect because using an AWS Glue Python shell script to run a sleep timer that checks every 5 minutes to determine whether the current Athena query has finished running successfully is not a cost-effective or reliable way to orchestrate the queries, as it wastes resources and time.

Option E is incorrect because using Amazon Managed Workflows for Apache Airflow (Amazon MWAA) to orchestrate the Athena queries in AWS Batch is an overkill solution that introduces unnecessary complexity and cost, as it requires setting up and managing an Airflow environment and an AWS Batch compute environment.

References:

- \* AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 5: Data Orchestration, Section 5.2: AWS Lambda, Section 5.3: AWS Step Functions, Pages 125-135
- \* Building Batch Data Analytics Solutions on AWS, Module 5: Data Orchestration, Lesson 5.1: AWS Lambda, Lesson 5.2: AWS Step Functions, Pages 1-15
- \* AWS Documentation Overview, AWS Lambda Developer Guide, Working with AWS Lambda Functions, Configuring Function Triggers, Using AWS Lambda with Amazon Athena, Pages 1-4
- \* AWS Documentation Overview, AWS Step Functions Developer Guide, Getting Started, Tutorial: Create a Hello World Workflow, Pages 1-8

## NEW QUESTION # 89

A company needs to set up a data catalog and metadata management for data sources that run in the AWS Cloud. The company will use the data catalog to maintain the metadata of all the objects that are in a set of data stores. The data stores include structured sources such as Amazon RDS and Amazon Redshift. The data stores also include semistructured sources such as JSON files and .xml files that are stored in Amazon S3.

The company needs a solution that will update the data catalog on a regular basis. The solution also must detect changes to the source metadata.

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

- A. Use the AWS Glue Data Catalog as the central metadata repository. Use AWS Glue crawlers to connect to multiple data stores and to update the Data Catalog with metadata changes. Schedule the crawlers to run periodically to update the metadata catalog.**
- B. Use Amazon DynamoDB as the data catalog. Create AWS Lambda functions that will connect to the data catalog. Configure the Lambda functions to gather the metadata information from multiple sources and to update the DynamoDB data catalog. Schedule the Lambda functions to run periodically.
- C. Use the AWS Glue Data Catalog as the central metadata repository. Extract the schema for Amazon RDS and Amazon Redshift sources, and build the Data Catalog. Use AWS Glue crawlers for data that is in Amazon S3 to infer the schema and to automatically update the Data Catalog.
- D. Use Amazon Aurora as the data catalog. Create AWS Lambda functions that will connect to the data catalog. Configure the Lambda functions to gather the metadata information from multiple sources and to update the Aurora data catalog. Schedule the Lambda functions to run periodically.

**Answer: A**

#### Explanation:

This solution will meet the requirements with the least operational overhead because it uses the AWS Glue Data Catalog as the central metadata repository for data sources that run in the AWS Cloud. The AWS Glue Data Catalog is a fully managed service that provides a unified view of your data assets across AWS and on-premises data sources. It stores the metadata of your data in tables, partitions, and columns, and enables you to access and query your data using various AWS services, such as Amazon Athena, Amazon EMR, and Amazon Redshift Spectrum. You can use AWS Glue crawlers to connect to multiple data stores, such as Amazon RDS, Amazon Redshift, and Amazon S3, and to update the Data Catalog with metadata changes.

AWS Glue crawlers can automatically discover the schema and partition structure of your data, and create or update the corresponding tables in the Data Catalog. You can schedule the crawlers to run periodically to update the metadata catalog, and configure them to detect changes to the source metadata, such as new columns, tables, or partitions<sup>12</sup>.

The other options are not optimal for the following reasons:

A: Use Amazon Aurora as the data catalog. Create AWS Lambda functions that will connect to the data catalog. Configure the Lambda functions to gather the metadata information from multiple sources and to update the Aurora data catalog. Schedule the Lambda functions to run periodically. This option is not recommended, as it would require more operational overhead to create and manage an Amazon Aurora database as the data catalog, and to write and maintain AWS Lambda functions to gather and update the metadata information from multiple sources. Moreover, this option would not leverage the benefits of the AWS Glue Data Catalog, such as data cataloging, data transformation, and data governance.

C: Use Amazon DynamoDB as the data catalog. Create AWS Lambda functions that will connect to the data catalog. Configure the Lambda functions to gather the metadata information from multiple sources and to update the DynamoDB data catalog. Schedule the Lambda functions to run periodically. This option is also not recommended, as it would require more operational overhead to create and manage an Amazon DynamoDB table as the data catalog, and to write and maintain AWS Lambda functions to gather and update the metadata information from multiple sources. Moreover, this option would not leverage the benefits of the AWS Glue Data Catalog, such as data cataloging, data transformation, and data governance.

D: Use the AWS Glue Data Catalog as the central metadata repository. Extract the schema for Amazon RDS and Amazon Redshift sources, and build the Data Catalog. Use AWS Glue crawlers for data that is in Amazon S3 to infer the schema and to automatically update the Data Catalog. This option is not optimal, as it would require more manual effort to extract the schema for Amazon RDS and Amazon Redshift sources, and to build the Data Catalog. This option would not take advantage of the AWS Glue crawlers' ability to automatically discover the schema and partition structure of your data from various data sources, and to create or update the corresponding tables in the Data Catalog.

#### References:

- 1: AWS Glue Data Catalog
- 2: AWS Glue Crawlers
- 3: Amazon Aurora
- 4: AWS Lambda
- 5: Amazon DynamoDB

#### NEW QUESTION # 90

A data engineer maintains a materialized view that is based on an Amazon Redshift database. The view has a column named `load_date` that stores the date when each row was loaded.

The data engineer needs to reclaim database storage space by deleting all the rows from the materialized view.

Which command will reclaim the MOST database storage space?

A.

 DELETE FROM materialized\_view\_name where 1=1

B.

TRUNCATE materialized\_view\_name

C.

VACUUM table\_name where load\_date<=current\_date  
materializedview

D.

DELETE FROM materialized\_view\_name where load\_date<=current\_date

- A. Option D
- B. Option C
- C. Option B
- D. Option A

**Answer: D**

Explanation:

To reclaim the most storage space from a materialized view in Amazon Redshift, you should use a DELETE operation that removes all rows from the view. The most efficient way to remove all rows is to use a condition that always evaluates to true, such as 1=1. This will delete all rows without needing to evaluate each row individually based on specific column values like load\_date.

Option A: DELETE FROM materialized\_view\_name WHERE 1=1;

This statement will delete all rows in the materialized view and free up the space. Since materialized views in Redshift store precomputed data, performing a DELETE operation will remove all stored rows.

Other options either involve inappropriate SQL statements (e.g., VACUUM in option C is used for reclaiming storage space in tables, not materialized views), or they don't remove data effectively in the context of a materialized view (e.g., TRUNCATE cannot be used directly on a materialized view).

Reference:

[Amazon Redshift Materialized Views Documentation](#)

[Deleting Data from Redshift](#)

#### NEW QUESTION # 91

A company uses a variety of AWS and third-party data stores. The company wants to consolidate all the data into a central data warehouse to perform analytics. Users need fast response times for analytics queries.

The company uses Amazon QuickSight in direct query mode to visualize the data. Users normally run queries during a few hours each day with unpredictable spikes.

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

- A. Use Amazon Athena to load all the data into Amazon S3 in Apache Parquet format.
- B. Use Amazon Redshift provisioned clusters to load all the data into Amazon Redshift managed storage (RMS).
- C. Use Amazon Redshift Serverless to load all the data into Amazon Redshift managed storage (RMS).
- D. Use Amazon Aurora PostgreSQL to load all the data into Aurora.

**Answer: C**

Explanation:

Problem Analysis:

The company requires a centralized data warehouse for consolidating data from various sources.

They use Amazon QuickSight in direct query mode, necessitating fast response times for analytical queries.

Users query the data intermittently, with unpredictable spikes during the day.

Operational overhead should be minimal.

Key Considerations:

The solution must support fast, SQL-based analytics.

It must handle unpredictable spikes efficiently.

Must integrate seamlessly with QuickSight for direct querying.

Minimize operational complexity and scaling concerns.

Solution Analysis:

Option A: Amazon Redshift Serverless

Redshift Serverless eliminates the need for provisioning and managing clusters.

Automatically scales compute capacity up or down based on query demand.

Reduces operational overhead by handling performance optimization.

Fully integrates with Amazon QuickSight, ensuring low-latency analytics.

Reduces costs as it charges only for usage, making it ideal for workloads with intermittent spikes.

Option B: Amazon Athena with S3 (Apache Parquet)

Athena supports querying data directly from S3 in Parquet format.

While it's cost-effective, performance depends on the size and complexity of the data.

It is not optimized for high-speed analytics needed by QuickSight in direct query mode.

Option C: Amazon Redshift Provisioned Clusters

Requires manual cluster provisioning, scaling, and maintenance.

Higher operational overhead compared to Redshift Serverless.

Option D: Amazon Aurora PostgreSQL

Aurora is optimized for transactional databases, not data warehousing or analytics.

Does not meet the requirement for fast analytics queries.

Final Recommendation:

Amazon Redshift Serverless is the best choice for this use case because it provides fast analytics, integrates natively with QuickSight, and minimizes operational complexity while efficiently handling unpredictable spikes.

[Amazon Redshift Serverless Overview](#)

[Amazon QuickSight and Redshift Integration](#)

[Athena vs. Redshift](#)

## NEW QUESTION # 92

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. 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.
- 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. 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.

**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.

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 # 93

.....

The three formats of Amazon Data-Engineer-Associate practice material that we have discussed above are created after receiving feedback from thousands of professionals around the world. You can instantly download the Amazon Data-Engineer-Associate Real Questions of the RealVCE right after the payment. We also offer our clients free demo version to evaluate the of our AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) valid exam dumps before purchasing.

**Data-Engineer-Associate Latest Exam Answers:** [https://www.realvce.com/Data-Engineer-Associate\\_free-dumps.html](https://www.realvce.com/Data-Engineer-Associate_free-dumps.html)

These tools are helping people in getting passed in the Data-Engineer-Associate online audio lectures with maximum ease and a lot of comfort ind, Amazon Detail Data-Engineer-Associate Explanation Such an easy and innovative study plan is amazingly beneficial for an ultimately brilliant success in exam, Amazon Detail Data-Engineer-Associate Explanation The learning of our study materials costs you little time and energy and we update them frequently, Amazon Detail Data-Engineer-Associate Explanation That is the reason why success belongs to few people.

These are enlarged details from two versions, at different resolutions, of Data-Engineer-Associate animator Henry Foster's portfolio reel, Final Merging and Ranking It is one thing to return documents that contain key words from the question.

## Prepare for the Amazon Exam on the Go with Data-Engineer-Associate PDF Dumps

These tools are helping people in getting passed in the Data-Engineer-Associate online audio lectures with maximum ease and a lot of comfort ind, Such an easy and innovative study plan is amazingly beneficial for an ultimately brilliant success in exam.

The learning of our study materials costs you little time and energy Data-Engineer-Associate Latest Exam Answers and we update them frequently, That is the reason why success belongs to few people, Passing the AWS Certified Data Engineer - Associate (DEA-C01) exam was never so easy.

What's more, part of that RealVCE Data-Engineer-Associate dumps now are free: [https://drive.google.com/open?id=1Xd4ng\\_zMBwEvq1ZPitPUdI-JK0MFvL3T](https://drive.google.com/open?id=1Xd4ng_zMBwEvq1ZPitPUdI-JK0MFvL3T)