

# Latest Data-Engineer-Associate Exam Topics - Examcollection Data-Engineer-Associate Dumps Torrent



BTW, DOWNLOAD part of Itcerttest Data-Engineer-Associate dumps from Cloud Storage: [https://drive.google.com/open?id=1sYZh4FO4K7\\_pDHmvD4AirU25QjhjFEFG](https://drive.google.com/open?id=1sYZh4FO4K7_pDHmvD4AirU25QjhjFEFG)

It is understandable that different people have different preference in terms of Data-Engineer-Associate study guide. Taking this into consideration, we have prepared three kinds of versions of our Data-Engineer-Associate preparation questions: PDF, online engine and software versions. The PDF version of Data-Engineer-Associate training materials is convenient for you to print, the software version can simulate the real exam and the online version can be used on all electronic devices. If you are hesitating about which version should you choose, you can download our Data-Engineer-Associate free demo first to get a firsthand experience before you make any decision.

The only aim of our company is to help each customer pass their exam as well as getting the important certification in a short time. If you want to pass your exam and get the Data-Engineer-Associate certification which is crucial for you successfully, I highly recommend that you should choose the Data-Engineer-Associate study materials from our company so that you can get a good understanding of the exam that you are going to prepare for. We believe that if you decide to buy the Data-Engineer-Associate Study Materials from our company, you will pass your exam and get the certification in a more relaxed way than other people.

>> Latest Data-Engineer-Associate Exam Topics <<

## Examcollection Data-Engineer-Associate Dumps Torrent | Data-Engineer-Associate PDF Dumps Files

As for the structure of content, please believe that our team of experts has many years of experience in compiling and designing on the Data-Engineer-Associate exam questions. I can say that no person can know the Data-Engineer-Associate study materials than them for they have been devoting themselves in this career for ten years. And they know every detail about the Data-Engineer-Associate learning guide. No matter how high your request is, our Data-Engineer-Associate learning quiz must satisfy you.

## Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q92-Q97):

### NEW QUESTION # 92

A car sales company maintains data about cars that are listed for sale in an area. The company receives data about new car listings from vendors who upload the data daily as compressed files into Amazon S3. The compressed files are up to 5 KB in size. The company wants to see the most up-to-date listings as soon as the data is uploaded to Amazon S3.

A data engineer must automate and orchestrate the data processing workflow of the listings to feed a dashboard. The data engineer must also provide the ability to perform one-time queries and analytical reporting. The query solution must be scalable. Which solution will meet these requirements MOST cost-effectively?

- A. Use an Amazon EMR cluster to process incoming data. Use AWS Step Functions to orchestrate workflows. Use Apache Hive for one-time queries and analytical reporting. Use Amazon OpenSearch Service to bulk ingest the data into compute optimized instances. Use OpenSearch Dashboards in OpenSearch Service for the dashboard.
- B. Use AWS Glue to process incoming data. Use AWS Lambda and S3 Event Notifications to orchestrate workflows. Use Amazon Athena for one-time queries and analytical reporting. Use Amazon QuickSight for the dashboard.
- C. Use AWS Glue to process incoming data. Use AWS Step Functions to orchestrate workflows. Use Amazon Redshift Spectrum for one-time queries and analytical reporting. Use OpenSearch Dashboards in Amazon OpenSearch Service for the dashboard.
- D. Use a provisioned Amazon EMR cluster to process incoming data. Use AWS Step Functions to orchestrate workflows. Use Amazon Athena for one-time queries and analytical reporting. Use Amazon QuickSight for the dashboard.

**Answer: B**

Explanation:

For processing the incoming car listings in a cost-effective, scalable, and automated way, the ideal approach involves using AWS Glue for data processing, AWS Lambda with S3 Event Notifications for orchestration, Amazon Athena for one-time queries and analytical reporting, and Amazon QuickSight for visualization on the dashboard. Let's break this down:

\* AWS Glue: This is a fully managed ETL (Extract, Transform, Load) service that automatically processes the incoming data files. Glue is serverless and supports diverse data sources, including Amazon S3 and Redshift.

\* AWS Lambda and S3 Event Notifications: Using Lambda and S3 Event Notifications allows near real-time triggering of processing workflows as soon as new data is uploaded into S3. This approach is event-driven, ensuring that the listings are processed as soon as they are uploaded, reducing the latency for data processing.

\* Amazon Athena: A serverless, pay-per-query service that allows interactive queries directly against data in S3 using standard SQL. It is ideal for the requirement of one-time queries and analytical reporting without the need for provisioning or managing servers.

\* Amazon QuickSight: A business intelligence tool that integrates with a wide range of AWS data sources, including Athena, and is used for creating interactive dashboards. It scales well and provides real-time insights for the car listings.

This solution (Option D) is the most cost-effective, because both Glue and Athena are serverless and priced based on usage, reducing costs when compared to provisioning EMR clusters in the other options. Moreover, using Lambda for orchestration is more cost-effective than AWS Step Functions due to its lightweight nature.

References:

- \* AWS Glue Documentation
- \* Amazon Athena Documentation
- \* Amazon QuickSight Documentation
- \* S3 Event Notifications and Lambda

### NEW QUESTION # 93

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 C

- B. Option D
- 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).

References:

Amazon Redshift Materialized Views Documentation

Deleting Data from Redshift

#### NEW QUESTION # 94

A manufacturing company wants to collect data from sensors. A data engineer needs to implement a solution that ingests sensor data in near real time.

The solution must store the data to a persistent data store. The solution must store the data in nested JSON format. The company must have the ability to query from the data store with a latency of less than 10 milliseconds.

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

- A. Use a self-hosted Apache Kafka cluster to capture the sensor data. Store the data in Amazon S3 for querying.
- **B. Use Amazon Kinesis Data Streams to capture the sensor data. Store the data in Amazon DynamoDB for querying.**
- C. Use Amazon Simple Queue Service (Amazon SQS) to buffer incoming sensor data. Use AWS Glue to store the data in Amazon RDS for querying.
- D. Use AWS Lambda to process the sensor data. Store the data in Amazon S3 for querying.

**Answer: B**

Explanation:

Amazon Kinesis Data Streams is a service that enables you to collect, process, and analyze streaming data in real time. You can use Kinesis Data Streams to capture sensor data from various sources, such as IoT devices, web applications, or mobile apps. You can create data streams that can scale up to handle any amount of data from thousands of producers. You can also use the Kinesis Client Library (KCL) or the Kinesis Data Streams API to write applications that process and analyze the data in the streams<sup>1</sup>.

Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. You can use DynamoDB to store the sensor data in nested JSON format, as DynamoDB supports document data types, such as lists and maps. You can also use DynamoDB to query the data with a latency of less than 10 milliseconds, as DynamoDB offers single-digit millisecond performance for any scale of data. You can use the DynamoDB API or the AWS SDKs to perform queries on the data, such as using key-value lookups, scans, or queries<sup>2</sup>.

The solution that meets the requirements with the least operational overhead is to use Amazon Kinesis Data Streams to capture the sensor data and store the data in Amazon DynamoDB for querying. This solution has the following advantages:

\* It does not require you to provision, manage, or scale any servers, clusters, or queues, as Kinesis Data Streams and DynamoDB are fully managed services that handle all the infrastructure for you. This reduces the operational complexity and cost of running your solution.

\* It allows you to ingest sensor data in near real time, as Kinesis Data Streams can capture data records as they are produced and deliver them to your applications within seconds. You can also use Kinesis Data Firehose to load the data from the streams to DynamoDB automatically and continuously<sup>3</sup>.

\* It allows you to store the data in nested JSON format, as DynamoDB supports document data types, such as lists and maps. You can also use DynamoDB Streams to capture changes in the data and trigger actions, such as sending notifications or updating other databases.

\* It allows you to query the data with a latency of less than 10 milliseconds, as DynamoDB offers single-digit millisecond performance for any scale of data. You can also use DynamoDB Accelerator (DAX) to improve the read performance by caching frequently accessed data.

Option A is incorrect because it suggests using a self-hosted Apache Kafka cluster to capture the sensor data and store the data in

Amazon S3 for querying. This solution has the following disadvantages:

- \* It requires you to provision, manage, and scale your own Kafka cluster, either on EC2 instances or on-premises servers. This increases the operational complexity and cost of running your solution.

- \* It does not allow you to query the data with a latency of less than 10 milliseconds, as Amazon S3 is an object storage service that is not optimized for low-latency queries. You need to use another service, such as Amazon Athena or Amazon Redshift Spectrum, to query the data in S3, which may incur additional costs and latency.

Option B is incorrect because it suggests using AWS Lambda to process the sensor data and store the data in Amazon S3 for querying. This solution has the following disadvantages:

- \* It does not allow you to ingest sensor data in near real time, as Lambda is a serverless compute service that runs code in response to events. You need to use another service, such as API Gateway or Kinesis Data Streams, to trigger Lambda functions with sensor data, which may add extra latency and complexity to your solution.

- \* It does not allow you to query the data with a latency of less than 10 milliseconds, as Amazon S3 is an object storage service that is not optimized for low-latency queries. You need to use another service, such as Amazon Athena or Amazon Redshift Spectrum, to query the data in S3, which may incur additional costs and latency.

Option D is incorrect because it suggests using Amazon Simple Queue Service (Amazon SQS) to buffer incoming sensor data and use AWS Glue to store the data in Amazon RDS for querying. This solution has the following disadvantages:

- \* It does not allow you to ingest sensor data in near real time, as Amazon SQS is a message queue service that delivers messages in a best-effort manner. You need to use another service, such as Lambda or EC2, to poll the messages from the queue and process them, which may add extra latency and complexity to your solution.

- \* It does not allow you to store the data in nested JSON format, as Amazon RDS is a relational database service that supports structured data types, such as tables and columns. You need to use another service, such as AWS Glue, to transform the data from JSON to relational format, which may add extra cost and overhead to your solution.

:

1: Amazon Kinesis Data Streams - Features

2: Amazon DynamoDB - Features

3: Loading Streaming Data into Amazon DynamoDB - Amazon Kinesis Data Firehose

[4]: Capturing Table Activity with DynamoDB Streams - Amazon DynamoDB

[5]: Amazon DynamoDB Accelerator (DAX) - Features

[6]: Amazon S3 - Features

[7]: AWS Lambda - Features

[8]: Amazon Simple Queue Service - Features

[9]: Amazon Relational Database Service - Features

[10]: Working with JSON in Amazon RDS - Amazon Relational Database Service

[11]: AWS Glue - Features

## NEW QUESTION # 95

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 C
- B. Option D
- 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).

References:

Amazon Redshift Materialized Views Documentation  
Deleting Data from Redshift

### NEW QUESTION # 96

A company has a data processing pipeline that includes several dozen steps. The data processing pipeline needs to send alerts in real time when a step fails or succeeds. The data processing pipeline uses a combination of Amazon S3 buckets, AWS Lambda functions, and AWS Step Functions state machines.

A data engineer needs to create a solution to monitor the entire pipeline.

Which solution will meet these requirements?

- A. Use AWS CloudTrail to send a message to an Amazon Simple Notification Service (Amazon SNS) topic that sends notifications when a state machine fails to run or succeeds to run.
- B. Configure the AWS Lambda functions to store notifications in an Amazon S3 bucket when the state machines finish running. Enable S3 event notifications on the S3 bucket.
- C. Configure the Step Functions state machines to store notifications in an Amazon S3 bucket when the state machines finish running. Enable S3 event notifications on the S3 bucket.
- D. Configure an Amazon EventBridge rule to react when the execution status of a state machine changes. Configure the rule to send a message to an Amazon Simple Notification Service (Amazon SNS) topic that sends notifications.

**Answer: D**

Explanation:

AWS Step Functions natively emits state change events to Amazon EventBridge, which can trigger an Amazon SNS notification.

This is the most direct and real-time way to alert on success/failure without relying on custom logging or polling.

"Step Functions automatically emits status changes that EventBridge can capture to trigger alerts or workflows. Use EventBridge to invoke an SNS topic for real-time alerts on job status."

- Ace the AWS Certified Data Engineer - Associate Certification - version 2 - apple.pdf This provides real-time alerting and the least operational overhead.

## NEW QUESTION # 97

.....

Our Data-Engineer-Associate exam questions have three versions: the PDF, Software and APP online. Also, there will have no extra restrictions to your learning because different versions have different merits. All in all, you will not be forced to buy all versions of our Data-Engineer-Associate Study Materials. You have the final right to select. Please consider our Data-Engineer-Associate learning quiz carefully and you will get a beautiful future with its help.

**Examcollection Data-Engineer-Associate Dumps Torrent:** [https://www.itcerttest.com/Data-Engineer-Associate\\_braindumps.html](https://www.itcerttest.com/Data-Engineer-Associate_braindumps.html)

The whole services of our Data-Engineer-Associate pass-sure materials: AWS Certified Data Engineer - Associate (DEA-C01) are satisfying. Whenever you have questions about our Data-Engineer-Associate training braindumps, you are welcome to contact us via E-mail, Interactive & Effective Data-Engineer-Associate Dumps PDF + Online Test Engine, For the purpose, our experts have designed our Amazon Data-Engineer-Associate exam dumps and AWS Certified Data Engineer in Questions and Answers format. Here, Data-Engineer-Associate study dumps are really worthwhile for your preparation.

This would reduce latency, enable real-time processing and Data-Engineer-Associate analytics, accelerate delivery, and lower data transfer and storage costs, The name of this property uses the following formula: application instance name + + current date Examcollection Data-Engineer-Associate Dumps Torrent The information stored in the `savedCalls` data is loaded into the conference retrieval client as soon as it loads.

## Latest AWS Certified Data Engineer - Associate (DEA-C01) exam pdf & Data-Engineer-Associate exam torrent

The whole services of our Data-Engineer-Associate pass-sure materials: AWS Certified Data Engineer - Associate (DEA-C01) are satisfying. Whenever you have questions about our Data-Engineer-Associate training braindumps, you are welcome to contact us via E-mail.

Interactive & Effective Data-Engineer-Associate Dumps PDF + Online Test Engine, For the purpose, our experts have designed our Amazon Data-Engineer-Associate exam dumps and AWS Certified Data Engineer in Questions and Answers format.

Here, Data-Engineer-Associate study dumps are really worthwhile for your preparation.

- Realistic Latest Data-Engineer-Associate Exam Topics | Amazing Pass Rate For Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01) | First-Grade Examcollection Data-Engineer-Associate Dumps Torrent  ➡ [www.prepawaypdf.com](http://www.prepawaypdf.com)  is best website to obtain ➡ Data-Engineer-Associate  for free download  Exam Data-Engineer-Associate Answers
- Take Your Exam Preparation to the Next Level with Pdfvce Amazon Data-Engineer-Associate Web-Based Practice Test   Download 《 Data-Engineer-Associate 》 for free by simply entering > [www.pdfvce.com](http://www.pdfvce.com) < website  Demo Data-Engineer-Associate Test
- Take Your Exam Preparation to the Next Level with [www.vce4dumps.com](http://www.vce4dumps.com) Amazon Data-Engineer-Associate Web-Based Practice Test  Easily obtain ⇒ Data-Engineer-Associate ⇐ for free download through > [www.vce4dumps.com](http://www.vce4dumps.com)   Test Data-Engineer-Associate Quiz
- Realistic Latest Data-Engineer-Associate Exam Topics | Amazing Pass Rate For Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01) | First-Grade Examcollection Data-Engineer-Associate Dumps Torrent  Search for [ Data-Engineer-Associate ] and download it for free immediately on > [www.pdfvce.com](http://www.pdfvce.com) <  Exam Data-Engineer-Associate Answers
- Exam Data-Engineer-Associate Certification Cost  Data-Engineer-Associate Test Discount  Free Data-Engineer-Associate Learning Cram  Copy URL “ [www.examcollectionpass.com](http://www.examcollectionpass.com) ” open and search for ( Data-Engineer-Associate ) to download for free  Data-Engineer-Associate Accurate Prep Material
- Realistic Latest Data-Engineer-Associate Exam Topics | Amazing Pass Rate For Data-Engineer-Associate: AWS Certified

