

100% Pass 2026 Updated Amazon Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01) Latest Exam Papers



2026 Latest Exams4Collection Data-Engineer-Associate PDF Dumps and Data-Engineer-Associate Exam Engine Free Share:
https://drive.google.com/open?id=1Nk5blJCDmjOxmvm5SECpWCJefo8vO_dn

We have to admit that the exam of gaining the Data-Engineer-Associate certification is not easy for a lot of people, especial these people who have no enough time. If you also look forward to change your present boring life, maybe trying your best to have the Data-Engineer-Associate certification is a good choice for you. Now it is time for you to take an exam for getting the certification. If you have any worry about the Data-Engineer-Associate Exam, do not worry, we are glad to help you. Because the Data-Engineer-Associate study materials from our company are very useful for you to pass the exam and get the certification.

Our worldwide after sale staff on the Data-Engineer-Associate exam questions will be online and reassure your rows of doubts as well as exclude the difficulties and anxiety with all the customers. Just let us know your puzzles on Data-Engineer-Associate study materials and we will figure out together. We can give you suggestion on Data-Engineer-Associate training engine 24/7, as long as you contact us, no matter by email or online, you will be answered quickly and professionally!

>> **Data-Engineer-Associate Latest Exam Papers** <<

Vce Amazon Data-Engineer-Associate Format, Data-Engineer-Associate Online Bootcamps

Here, we provide you with the best Data-Engineer-Associate premium study files which will improve your study efficiency and give you right direction. The content of Data-Engineer-Associate study material is the updated and verified by IT experts. Professional experts are arranged to check and trace the Amazon Data-Engineer-Associate update information every day. The Data-Engineer-Associate exam guide materials are really worthy of purchase. The high quality and accurate Data-Engineer-Associate questions & answers are the guarantee of your success.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample

Questions (Q266-Q271):

NEW QUESTION # 266

A company needs to store semi-structured transactional data in a serverless database. The application writes data infrequently but reads it frequently, with millisecond retrieval required.

- A. Store the data in an Amazon S3 Apache Iceberg table. Enable S3 Transfer Acceleration.
- **B. Store the data in an Amazon DynamoDB table. Configure a DynamoDB Accelerator (DAX) cache.**
- C. Store the data in an Amazon RDS for MySQL cluster. Configure RDS Optimized Reads.
- D. Store the data in an Amazon S3 Standard bucket. Enable S3 Transfer Acceleration.

Answer: B

Explanation:

Amazon DynamoDB is a serverless, low-latency, NoSQL database ideal for semi-structured data.

Adding DynamoDB Accelerator (DAX) provides microsecond response times for read-heavy workloads.

"For applications requiring millisecond or sub-millisecond reads with serverless operation, use DynamoDB with DAX caching."

- Ace the AWS Certified Data Engineer - Associate Certification - version 2 - apple.pdf

NEW QUESTION # 267

A company has used an Amazon Redshift table that is named Orders for 6 months. The company performs weekly updates and deletes on the table. The table has an interleaved sort key on a column that contains AWS Regions.

The company wants to reclaim disk space so that the company will not run out of storage space. The company also wants to analyze the sort key column.

Which Amazon Redshift command will meet these requirements?

- A. VACUUM SORT ONLY Orders
- B. VACUUM FULL Orders
- **C. VACUUM REINDEX Orders**
- D. VACUUM DELETE ONLY Orders

Answer: C

Explanation:

Amazon Redshift is a fully managed, petabyte-scale data warehouse service that enables fast and cost-effective analysis of large volumes of data. Amazon Redshift uses columnar storage, compression, and zone maps to optimize the storage and performance of data. However, over time, as data is inserted, updated, or deleted, the physical storage of data can become fragmented, resulting in wasted disk space and degraded query performance. To address this issue, Amazon Redshift provides the VACUUM command, which reclaims disk space and resorts rows in either a specified table or all tables in the current schema¹.

The VACUUM command has four options: FULL, DELETE ONLY, SORT ONLY, and REINDEX. The option that best meets the requirements of the question is VACUUM REINDEX, which re-sorts the rows in a table that has an interleaved sort key and rewrites the table to a new location on disk. An interleaved sort key is a type of sort key that gives equal weight to each column in the sort key, and stores the rows in a way that optimizes the performance of queries that filter by multiple columns in the sort key. However, as data is added or changed, the interleaved sort order can become skewed, resulting in suboptimal query performance. The VACUUM REINDEX option restores the optimal interleaved sort order and reclaims disk space by removing deleted rows. This option also analyzes the sort key column and updates the table statistics, which are used by the query optimizer to generate the most efficient query execution plan^{2,3}.

The other options are not optimal for the following reasons:

- A. VACUUM FULL Orders. This option reclaims disk space by removing deleted rows and resorts the entire table. However, this option is not suitable for tables that have an interleaved sort key, as it does not restore the optimal interleaved sort order. Moreover, this option is the most resource-intensive and time-consuming, as it rewrites the entire table to a new location on disk.
- B. VACUUM DELETE ONLY Orders. This option reclaims disk space by removing deleted rows, but does not resort the table. This option is not suitable for tables that have any sort key, as it does not improve the query performance by restoring the sort order. Moreover, this option does not analyze the sort key column and update the table statistics.
- D. VACUUM SORT ONLY Orders. This option resorts the entire table, but does not reclaim disk space by removing deleted rows. This option is not suitable for tables that have an interleaved sort key, as it does not restore the optimal interleaved sort order. Moreover, this option does not analyze the sort key column and update the table statistics.

1: Amazon Redshift VACUUM

2: Amazon Redshift Interleaved Sorting

3: Amazon Redshift ANALYZE

NEW QUESTION # 268

A data engineer is building a data pipeline. A large data file is uploaded to an Amazon S3 bucket once each day at unpredictable times. An AWS Glue workflow uses hundreds of workers to process the file and load the data into Amazon Redshift. The company wants to process the file as quickly as possible.

Which solution will meet these requirements?

- A. Create an event-based AWS Glue trigger to start the workflow. Configure Amazon S3 to log events to AWS CloudTrail. Create a rule in Amazon EventBridge to forward PutObject events to the AWS Glue trigger.
- B. Create an on-demand AWS Glue trigger to start the workflow. Create an AWS Database Migration Service (AWS DMS) migration task. Set the DMS source as the S3 bucket. Set the target endpoint as the AWS Glue workflow.
- C. Create an on-demand AWS Glue trigger to start the workflow. Create an AWS Lambda function that runs every 15 minutes to check the S3 bucket for the daily file. Configure the function to start the AWS Glue workflow if the file is present.
- D. Create a scheduled AWS Glue trigger to start the workflow. Create a cron job that runs the AWS Glue job every 15 minutes. Set up the AWS Glue job to check the S3 bucket for the daily file. Configure the job to stop if the file is not present.

Answer: A

Explanation:

The best solution for fast, event-driven processing of unpredictable file uploads is to use S3 event notifications, CloudTrail, and EventBridge to automatically trigger the AWS Glue workflow:

"You can configure S3 PutObject events to be captured by CloudTrail and forwarded through EventBridge to trigger an AWS Glue job or workflow. This allows Glue to begin processing as soon as the file arrives, with minimal latency."

- Ace the AWS Certified Data Engineer - Associate Certification - version 2 - apple.pdf This option provides the lowest latency and least manual overhead compared to polling or scheduling solutions.

NEW QUESTION # 269

A company uses Amazon S3 buckets, AWS Glue tables, and Amazon Athena as components of a data lake.

Recently, the company expanded its sales range to multiple new states. The company wants to introduce state names as a new partition to the existing S3 bucket, which is currently partitioned by date.

The company needs to ensure that additional partitions will not disrupt daily synchronization between the AWS Glue Data Catalog and the S3 buckets.

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

- A. Schedule an AWS Glue crawler to periodically update the Data Catalog.
- B. Run an MSCK REPAIR TABLE command in Athena.
- C. Run a REFRESH TABLE command in Athena.
- D. Use the AWS Glue API to manually update the Data Catalog.

Answer: A

Explanation:

Scheduling an AWS Glue crawler to periodically update the Data Catalog automates the process of detecting new partitions and updating the catalog, which minimizes manual maintenance and operational overhead.

NEW QUESTION # 270

A company uses Amazon Redshift for its data warehouse. A data engineer must query a table named orders.

complete_orders_history, which contains 100 columns. The query must return all columns except columns named company_id and unique_system_id.

Which Amazon Redshift SQL statement will meet this requirement?

- A. SELECT * TRUNCATE company_id, unique_system_id FROM orders.complete_orders_history;
- B. SELECT * NOT IN company_id, unique_system_id FROM orders.complete_orders_history;
- C. SELECT * EXCLUDE company_id, unique_system_id FROM orders.complete_orders_history;
- D. SELECT * EXCEPT company_id, unique_system_id FROM orders.complete_orders_history;

Answer: C

Explanation:

Option A is correct because Amazon Redshift supports the EXCLUDE clause in the SELECT list for exactly this use case: returning all columns from a wide table except a small set of unwanted columns. AWS documentation states that "EXCLUDE column_list names the columns that are excluded from the query results" and notes that the option is helpful when only a subset of columns must be excluded from a wide table. That matches this question precisely because the table contains 100 columns and only company_id and unique_system_id need to be omitted.

Option B is invalid SQL syntax in Amazon Redshift because NOT IN is used in predicates, not in the select list to remove columns.

Option C is also incorrect because EXCEPT in Redshift is a set operator used between query result sets, not a select-list column exclusion feature. AWS documents EXCEPT alongside UNION and INTERSECT as result-set comparison operators, not as a way to omit columns from SELECT *.

Option D is not valid Redshift SQL syntax for column selection. Therefore, the correct and documented Redshift statement is SELECT * EXCLUDE company_id, unique_system_id FROM orders.
complete_orders_history;

NEW QUESTION # 271

.....

Do you need to find a high paying job for yourself? Well, by passing the AWS Certified Data Engineer - Associate (DEA-C01), you will be able to get your dream job. Make sure that you are buying our bundle Data-Engineer-Associate brain dumps pack so you can check out all the products that will help you come up with a better solution. You can easily land a dream job by passing the Data-Engineer-Associate Exam in the first attempt.

Vce Data-Engineer-Associate Format: <https://www.exams4collection.com/Data-Engineer-Associate-latest-braindumps.html>

Amazon practice test engine is updated according to the changes of AWS Certified Data Engineer - Associate (DEA-C01) actual exam, for the sake that the questions you practice are close to the real Data-Engineer-Associate exam, which enormously enhance your efficiency, All of our workers have a great responsibility to offer our customers the high-quality Data-Engineer-Associate exam guide: AWS Certified Data Engineer - Associate (DEA-C01), Amazon Data-Engineer-Associate Latest Exam Papers Pdf version- being legible to read and remember, support customers' printing request, and allow you to have a print and practice in papers.

Are you ready to take control of your future and get the AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) certification you need to accelerate your career, If you are going to write a comment, take the time to make sure it is the best comment you can write.

Free PDF Quiz Efficient Data-Engineer-Associate - AWS Certified Data Engineer - Associate (DEA-C01) Latest Exam Papers

Amazon practice test engine is updated according to the changes of AWS Certified Data Engineer - Associate (DEA-C01) actual exam, for the sake that the questions you practice are close to the Real Data-Engineer-Associate Exam, which enormously enhance your efficiency.

All of our workers have a great responsibility to offer our customers the high-quality Data-Engineer-Associate exam guide: AWS Certified Data Engineer - Associate (DEA-C01), Pdf version- being legible to read and remember, support Data-Engineer-Associate customers' printing request, and allow you to have a print and practice in papers.

First, we have built a strong and professional team devoting to the research of Data-Engineer-Associate valid practice torrent, Are you preparing for the Amazon Data-Engineer-Associate certification exam?

- Perfect Data-Engineer-Associate Latest Exam Papers | Data-Engineer-Associate 100% Free Vce Format Open website www.verifeddumps.com and search for { Data-Engineer-Associate } for free download Exam Data-Engineer-Associate Pattern
- Data-Engineer-Associate Valid Test Braindumps ↗ Data-Engineer-Associate Exam Lab Questions Data-Engineer-Associate Reliable Exam Materials Download Data-Engineer-Associate for free by simply entering www.pdfvce.com website Data-Engineer-Associate New Test Bootcamp
- Pass Guaranteed Quiz 2026 The Best Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01) Latest Exam Papers Copy URL ▷ www.practicevce.com ◁ open and search for { Data-Engineer-Associate } to download for free Data-Engineer-Associate Training Solutions
- Data-Engineer-Associate Valid Test Braindumps Valid Data-Engineer-Associate Test Notes Exam Data-Engineer-Associate Simulations Search for ▷ Data-Engineer-Associate ◁ and download it for free immediately on ▷ www.pdfvce.com ◁ Frenquent Data-Engineer-Associate Update
- New Data-Engineer-Associate Test Questions Test Data-Engineer-Associate Passing Score Valid Data-Engineer-

