

# Data-Engineer-Associate AWS Certified Data Engineer - Associate (DEA-C01) Dumps For Ultimate Results 2026



BONUS!!! Download part of TestKingFree Data-Engineer-Associate dumps for free: <https://drive.google.com/open?id=1QwylnpW91J-jMlyR3hYdEc-PxWxQiJny>

After choose TestKingFree's Data-Engineer-Associate exam training materials, you can get the latest edition of Data-Engineer-Associate exam dumps and answers. The accuracy rate of TestKingFree Data-Engineer-Associate exam training materials can ensure you to Pass Data-Engineer-Associate Test. After you purchase our Data-Engineer-Associate test training materials, if you fail Data-Engineer-Associate exam certification or there are any quality problems of Data-Engineer-Associate exam dumps, we guarantee a full refund.

There are many ways to help you prepare for your Amazon Data-Engineer-Associate exam. TestKingFree provide a reliable training tools to help you prepare for your Amazon Data-Engineer-Associate exam certification. The TestKingFree Amazon Data-Engineer-Associate Exam Materials are including test questions and answers. Our materials are very good software that through the practice test. Our materials will meet all of the IT certifications.

>> **Data-Engineer-Associate Exam Pass4sure** <<

## **Latest Updated Amazon Data-Engineer-Associate Exam Pass4sure: AWS Certified Data Engineer - Associate (DEA-C01)**

Our AWS Certified Data Engineer - Associate (DEA-C01) Data-Engineer-Associate questions PDF is a complete bundle of problems presenting the versatility and correlativity of questions observed in past exam papers. These questions are bundled into AWS Certified Data Engineer - Associate (DEA-C01) PDF questions following the official study guide. Amazon Data-Engineer-Associate PDF Questions are a portable, printable document that simultaneously plays on multiple devices. Our Amazon Data-Engineer-Associate PDF questions consists of problems in all aspects, whether theoretical, practical, or analytical.

### **Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q214-Q219):**

### NEW QUESTION # 214

A company uses an Amazon Redshift Single-AZ cluster for enterprise analytics. The company wants to set up a highly resilient disaster recovery (DR) solution for the cluster. The solution must meet a recovery time objective (RTO) of less than 1 hour. Which solution will meet this requirement MOST cost-effectively?

- A. Configure a Redshift cluster from a cross-Region snapshot copy in a second AWS Region when necessary.
- B. Use a Redshift dense storage (DS2) node. Enable Multi-AZ deployment.
- C. Use a Redshift RA3 node. Enable cluster relocation.
- **D. Use a Redshift RA3 node. Enable Multi-AZ deployment.**

**Answer: D**

Explanation:

Option B best meets a highly resilient DR requirement with an RTO under 1 hour because a Multi-AZ deployment is designed to provide rapid recovery through redundancy across Availability Zones, rather than requiring a restore-and-rebuild process after a failure. A snapshot-based strategy (Option C) can be valuable for backup and regional recovery, but restoring a new cluster from snapshots is a heavier operation and is not the most reliable path to consistently achieving sub-hour RTO during real incidents. The study material highlights snapshot/restore as an operational mechanism for creating a new environment from a Redshift cluster's saved state, which inherently implies a restore process and associated time to bring resources online.

Option D (cluster relocation) is primarily an operational move mechanism and is not the same as maintaining continuously resilient capacity for unplanned AZ-impacting events. Option A is not the best choice because Multi-AZ resiliency is aligned with modern Redshift architectures, and RA3 is the recommended node family for contemporary Redshift deployments where performance and managed storage characteristics are key for enterprise analytics. The study material reinforces Amazon Redshift as the platform for high-performance enterprise analytics, so a resilience-first configuration is appropriate.

### NEW QUESTION # 215

A company stores time-series data that is collected from streaming services in an Amazon S3 bucket. The company must ensure that only workloads that are deployed within the company's VPC can access the data. Which solution will meet this requirement?

- A. Define an IAM policy that denies access to all users unless the request originates from within the company's VPC.
- B. Use a network ACL on the VPC subnets to allow only specific resources to access the S3 bucket.
- C. Apply a security group to the S3 bucket that allows connections only from the company's VPC CIDR block.
- **D. Create an S3 bucket policy that uses a condition to allow access only to traffic that originates from the company's VPC.**

**Answer: D**

Explanation:

The best practice to restrict Amazon S3 access to specific VPCs is to use a bucket policy with a StringEquals or StringLike condition on aws:SourceVpc. This ensures only requests from a specified VPC are allowed.

\* IAM policies (option C) control who can access the resource but are not suitable alone to restrict by VPC.

\* Security Groups and NACLs (options B and D) do not apply to Amazon S3 because it is a global service and not VPC-bound.

"You can restrict access to your S3 bucket so that only requests coming from a specific VPC endpoint are allowed." Source: AWS Documentation - Amazon S3 Bucket Policies for VPC Endpoints

### NEW QUESTION # 216

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 DELETE ONLY Orders
- C. VACUUM FULL Orders
- **D. VACUUM REINDEX Orders**

**Answer: D**

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 schema.

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.

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 # 217

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. 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.
- **B. 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.**
- C. 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.
- D. 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.

**Answer: B**

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.

Reference:

1: AWS Glue Data Catalog

2: AWS Glue Crawlers

: Amazon Aurora

: AWS Lambda

: Amazon DynamoDB

## NEW QUESTION # 218

A data engineer is building an automated extract, transform, and load (ETL) ingestion pipeline by using AWS Glue. The pipeline ingests compressed files that are in an Amazon S3 bucket. The ingestion pipeline must support incremental data processing. Which AWS Glue feature should the data engineer use to meet this requirement?

- A. Triggers
- B. Workflows
- C. Classifiers
- D. Job bookmarks

**Answer: D**

Explanation:

\* Problem Analysis:

\* The pipeline processes compressed files in S3 and must support incremental data processing.

\* AWS Glue features must facilitate tracking progress to avoid reprocessing the same data.

\* Key Considerations:

\* Incremental data processing requires tracking which files or partitions have already been processed.

\* The solution must be automated and efficient for large-scale ETL jobs.

\* Solution Analysis:

\* Option A: Workflows

\* Workflows organize and orchestrate multiple Glue jobs but do not track progress for incremental data processing.

\* Option B: Triggers

\* Triggers initiate Glue jobs based on a schedule or events but do not track which data has been processed.

\* Option C: Job Bookmarks

\* Job bookmarks track the state of the data that has been processed, enabling incremental processing.

\* Automatically skip files or partitions that were previously processed in Glue jobs.

\* Option D: Classifiers

\* Classifiers determine the schema of incoming data but do not handle incremental processing.

\* Final Recommendation:

\* Job bookmarks are specifically designed to enable incremental data processing in AWS Glue ETL pipelines.

:

## NEW QUESTION # 219

.....

TestKingFree's expert team has developed a latest short-term effective training scheme for Amazon certification Data-Engineer-Associate exam, which is a 20 hours of training for the candidates of Amazon certification Data-Engineer-Associate exam. After training they can not only quickly master a lot of knowledge, but also consolidate their original knowledge. So they can easily pass Amazon Certification Data-Engineer-Associate Exam and it is much more cost-effective for them than those who spend a lot of time and energy to prepare for the examination.

**Data-Engineer-Associate Valid Exam Cram:** <https://www.testkingfree.com/Amazon/Data-Engineer-Associate-practice-exam-dumps.html>

The Data-Engineer-Associate Valid Exam Cram - AWS Certified Data Engineer - Associate (DEA-C01) exam pass-sure materials will show you the Amazon Data-Engineer-Associate Valid Exam Cram certification can't be the tower of Babel for you, you can make it, The advantages of our Data-Engineer-Associate quiz torrent materials are as follows, Our Data-Engineer-Associate certified experts have curated questions and answers that will be asked in the real exam, and we provide money back guarantee on TestKingFree Amazon preparation material, If you're looking to advance your Amazon career, Amazon Data-Engineer-Associate Exam can help you achieve that goal.

If you need to convert several images to the same file format or the same Data-Engineer-Associate size and resolution, use the Process Multiple Files command, Parts may go out of stock, become obsolete, or cease being made altogether.

## Pass Data-Engineer-Associate Exam with Flying Colors Using Amazon's Exam Questions and Achieve Success

The AWS Certified Data Engineer - Associate (DEA-C01) exam pass-sure materials will show you the Amazon certification can't be the tower of Babel for you, you can make it, The advantages of our Data-Engineer-Associate Quiz torrent materials are as follows.

Our Data-Engineer-Associate certified experts have curated questions and answers that will be asked in the real exam, and we provide money back guarantee on TestKingFree Amazon preparation material.

If you're looking to advance your Amazon career, Amazon Data-Engineer-Associate Exam can help you achieve that goal, Wrong choices may engender wrong feed-backs, we are sure you will come a long way by our Data-Engineer-Associate practice questions.

- Pass Data-Engineer-Associate Rate  Data-Engineer-Associate Detailed Study Dumps  Pass Data-Engineer-Associate Rate  Open ( [www.easy4engine.com](http://www.easy4engine.com) ) enter **►** Data-Engineer-Associate  and obtain a free download Data-Engineer-Associate Testking Learning Materials
- Exam Sample Data-Engineer-Associate Online  Study Data-Engineer-Associate Reference  Test Data-Engineer-Associate Simulator Fee  Download  Data-Engineer-Associate  for free by simply entering ( [www.pdfvce.com](http://www.pdfvce.com) ) website Test Data-Engineer-Associate Cram Pdf
- Premium Amazon Data-Engineer-Associate Questions  Open **►** [www.vce4dumps.com](http://www.vce4dumps.com)  enter  Data-Engineer-Associate  and obtain a free download Data-Engineer-Associate Detailed Study Dumps
- Data-Engineer-Associate Test Passing Score  Data-Engineer-Associate Detailed Study Dumps  Simulations Data-Engineer-Associate Pdf  Search for 「 Data-Engineer-Associate 」 and obtain a free download on { [www.pdfvce.com](http://www.pdfvce.com) } Pass Data-Engineer-Associate Rate
- Test Data-Engineer-Associate Questions  Data-Engineer-Associate Test Passing Score  Data-Engineer-Associate Study Material  **►** [www.validtorrent.com](http://www.validtorrent.com) **◀** is best website to obtain **✪** Data-Engineer-Associate  **✪**  for free download Pass Data-Engineer-Associate Rate
- Study Data-Engineer-Associate Reference  Data-Engineer-Associate Study Material  Data-Engineer-Associate Detailed Study Dumps  Download **《** Data-Engineer-Associate **》** for free by simply searching on **►** [www.pdfvce.com](http://www.pdfvce.com) **◀** Test Data-Engineer-Associate Questions
- Data-Engineer-Associate Mock Exam  Exam Sample Data-Engineer-Associate Online  Test Data-Engineer-Associate Score Report  Go to website **⇒** [www.pdfdumps.com](http://www.pdfdumps.com) **⇐** open and search for  Data-Engineer-Associate  to download for free Study Data-Engineer-Associate Reference
- Exam Data-Engineer-Associate Consultant  Data-Engineer-Associate Test Passing Score  Data-Engineer-Associate

