

TOP Latest Test Data-Engineer-Associate Experience 100% Pass | The Best Certification AWS Certified Data Engineer - Associate (DEA-C01) Exam Dumps Pass for sure



BONUS!!! Download part of Lead2Passed Data-Engineer-Associate dumps for free: <https://drive.google.com/open?id=1t3pCTTdalcJaq9yzfQErDMJcz2usUSM>

The web-based AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) practice exam is accessible from any major OS, including Mac OS X, Linux, Android, Windows, or iOS. These Amazon Data-Engineer-Associate exam questions are browser-based, so there's no need to install anything on your computer. Chrome, IE, Firefox, and Opera all support this Amazon Data-Engineer-Associate web-based practice exam. You can take this AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) practice exam without plugins and software installation.

Our company always lays great emphasis on offering customers more wide range of choice. Now, we have realized our promise. Our Data-Engineer-Associate exam guide almost covers all kinds of official test and popular certificate. So you will be able to find what you need easily on our website. Every Data-Engineer-Associate exam torrent is professional and accurate, which can greatly relieve your learning pressure. In the meantime, we have three versions of product packages for you. They are PDF version, windows software and online engine of the Data-Engineer-Associate Exam Prep. The three versions of the study materials packages are very popular and cost-efficient now. With the assistance of our study materials, you will escape from the pains of preparing the exam. Of course, you can purchase our Data-Engineer-Associate exam guide according to your own conditions. All in all, you have the right to choose freely. You will not be forced to buy the packages.

>> Latest Test Data-Engineer-Associate Experience <<

Certification Amazon Data-Engineer-Associate Exam Dumps | Pdf Data-Engineer-Associate Torrent

The crucial thing when it comes to appearing a competitive exam like Data-Engineer-Associate knowing your problem-solving skills. And to do that you are going to need help from a Data-Engineer-Associate practice questions or braindumps. This is exactly what is delivered by our Data-Engineer-Associate test materials. The Data-Engineer-Associate Exam Dumps cover every topic of the actual Amazon certification exam. The Data-Engineer-Associate exam questions are divided into various groups and the candidate can solve these questions to test his skills and knowledge.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q210-Q215):

NEW QUESTION # 210

A company is planning to migrate on-premises Apache Hadoop clusters to Amazon EMR. The company also needs to migrate a

data catalog into a persistent storage solution.

The company currently stores the data catalog in an on-premises Apache Hive metastore on the Hadoop clusters. The company requires a serverless solution to migrate the data catalog.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure a new Hive metastore in Amazon EMR. Migrate the existing on-premises Hive metastore into Amazon EMR. Use the new metastore as the company's data catalog.
- B. Configure an external Hive metastore in Amazon EMR. Migrate the existing on-premises Hive metastore into Amazon EMR. Use Amazon Aurora MySQL to store the company's data catalog.
- **C. Use AWS Database Migration Service (AWS DMS) to migrate the Hive metastore into Amazon S3. Configure AWS Glue Data Catalog to scan Amazon S3 to produce the data catalog.**
- D. Configure a Hive metastore in Amazon EMR. Migrate the existing on-premises Hive metastore into Amazon EMR. Use AWS Glue Data Catalog to store the company's data catalog as an external data catalog.

Answer: C

Explanation:

AWS Database Migration Service (AWS DMS) is a service that helps you migrate databases to AWS quickly and securely. You can use AWS DMS to migrate the Hive metastore from the on-premises Hadoop clusters into Amazon S3, which is a highly scalable, durable, and cost-effective object storage service. AWS Glue Data Catalog is a serverless, managed service that acts as a central metadata repository for your data assets.

You can use AWS Glue Data Catalog to scan the Amazon S3 bucket that contains the migrated Hive metastore and create a data catalog that is compatible with Apache Hive and other AWS services. This solution meets the requirements of migrating the data catalog into a persistent storage solution and using a serverless solution. This solution is also the most cost-effective, as it does not incur any additional charges for running Amazon EMR or Amazon Aurora MySQL clusters. The other options are either not feasible or not optimal. Configuring a Hive metastore in Amazon EMR (option B) or an external Hive metastore in Amazon EMR (option C) would require running and maintaining Amazon EMR clusters, which would incur additional costs and complexity. Using Amazon Aurora MySQL to store the company's data catalog (option C) would also incur additional costs and complexity, as well as introduce compatibility issues with Apache Hive.

Configuring a new Hive metastore in Amazon EMR (option D) would not migrate the existing data catalog, but create a new one, which would result in data loss and inconsistency. References:

* Using AWS Database Migration Service

* Populating the AWS Glue Data Catalog

* AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 4: Data Analysis and Visualization, Section 4.2: AWS Glue Data Catalog

NEW QUESTION # 211

A company uses Amazon RDS for MySQL as the database for a critical application. The database workload is mostly writes, with a small number of reads.

A data engineer notices that the CPU utilization of the DB instance is very high. The high CPU utilization is slowing down the application. The data engineer must reduce the CPU utilization of the DB Instance.

Which actions should the data engineer take to meet this requirement? (Choose two.)

- A. Reboot the RDS DB instance once each week.
- B. Modify the database schema to include additional tables and indexes.
- **C. Implement caching to reduce the database query load.**
- D. Upgrade to a larger instance size.
- **E. Use the Performance Insights feature of Amazon RDS to identify queries that have high CPU utilization. Optimize the problematic queries.**

Answer: C,E

Explanation:

Amazon RDS is a fully managed service that provides relational databases in the cloud. Amazon RDS for MySQL is one of the supported database engines that you can use to run your applications. Amazon RDS provides various features and tools to monitor and optimize the performance of your DB instances, such as Performance Insights, Enhanced Monitoring, CloudWatch metrics and alarms, etc.

Using the Performance Insights feature of Amazon RDS to identify queries that have high CPU utilization and optimizing the problematic queries will help reduce the CPU utilization of the DB instance. Performance Insights is a feature that allows you to analyze the load on your DB instance and determine what is causing performance issues. Performance Insights collects, analyzes,

and displays database performance data using an interactive dashboard. You can use Performance Insights to identify the top SQL statements, hosts, users, or processes that are consuming the most CPU resources. You can also drill down into the details of each query and see the execution plan, wait events, locks, etc. By using Performance Insights, you can pinpoint the root cause of the high CPU utilization and optimize the queries accordingly. For example, you can rewrite the queries to make them more efficient, add or remove indexes, use prepared statements, etc.

Implementing caching to reduce the database query load will also help reduce the CPU utilization of the DB instance. Caching is a technique that allows you to store frequently accessed data in a fast and scalable storage layer, such as Amazon ElastiCache. By using caching, you can reduce the number of requests that hit your database, which in turn reduces the CPU load on your DB instance. Caching also improves the performance and availability of your application, as it reduces the latency and increases the throughput of your data access. You can use caching for various scenarios, such as storing session data, user preferences, application configuration, etc. You can also use caching for read-heavy workloads, such as displaying product details, recommendations, reviews, etc.

The other options are not as effective as using Performance Insights and caching. Modifying the database schema to include additional tables and indexes may or may not improve the CPU utilization, depending on the nature of the workload and the queries. Adding more tables and indexes may increase the complexity and overhead of the database, which may negatively affect the performance. Rebooting the RDS DB instance once each week will not reduce the CPU utilization, as it will not address the underlying cause of the high CPU load. Rebooting may also cause downtime and disruption to your application. Upgrading to a larger instance size may reduce the CPU utilization, but it will also increase the cost and complexity of your solution. Upgrading may also not be necessary if you can optimize the queries and reduce the database load by using caching. Reference:

Amazon RDS

Performance Insights

Amazon ElastiCache

[AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide], Chapter 3: Data Storage and Management, Section 3.1: Amazon RDS

NEW QUESTION # 212

A data engineer needs to use AWS Step Functions to design an orchestration workflow. The workflow must parallel process a large collection of data files and apply a specific transformation to each file.

Which Step Functions state should the data engineer use to meet these requirements?

- A. Parallel state
- B. Choice state
- C. Map state
- D. Wait state

Answer: C

Explanation:

Option C is the correct answer because the Map state is designed to process a collection of data in parallel by applying the same transformation to each element. The Map state can invoke a nested workflow for each element, which can be another state machine or a Lambda function. The Map state will wait until all the parallel executions are completed before moving to the next state.

Option A is incorrect because the Parallel state is used to execute multiple branches of logic concurrently, not to process a collection of data. The Parallel state can have different branches with different logic and states, whereas the Map state has only one branch that is applied to each element of the collection.

Option B is incorrect because the Choice state is used to make decisions based on a comparison of a value to a set of rules. The Choice state does not process any data or invoke any nested workflows.

Option D is incorrect because the Wait state is used to delay the state machine from continuing for a specified time. The Wait state does not process any data or invoke any nested workflows.

References:

AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 5: Data Orchestration, Section 5.3: AWS Step Functions, Pages 131-132 Building Batch Data Analytics Solutions on AWS, Module 5: Data Orchestration, Lesson 5.2:

AWS Step Functions, Pages 9-10 AWS Documentation Overview, AWS Step Functions Developer Guide, Step Functions Concepts, State Types, Map State, Pages 1-3

NEW QUESTION # 213

A gaming company uses AWS Glue to perform read and write operations on Apache Iceberg tables for real-time streaming data. The data in the Iceberg tables is stored in Apache Parquet format. The company is experiencing slow query performance.

Which solutions will improve query performance? (Select TWO)

- A. Use AWS Glue Data Catalog to enable copy-on-write for the Iceberg tables.
- B. Use AWS Glue Data Catalog to automatically optimize indexes for the Iceberg tables.
- C. Use AWS Glue Data Catalog to generate column-level statistics for the Iceberg tables on a schedule.
- D. Use AWS Glue Data Catalog to generate views for the Iceberg tables.
- E. Use AWS Glue Data Catalog to automatically compact the Iceberg tables.

Answer: C,E

Explanation:

Apache Iceberg query performance depends heavily on metadata optimization and file layout efficiency.

Generating column-level statistics allows query engines such as AWS Glue, Amazon Athena, and Amazon Redshift to perform predicate pushdown and data skipping. These statistics help avoid scanning unnecessary Parquet files and significantly reduce query execution time.

Real-time streaming workloads often generate many small files, which degrades query performance. Table compaction consolidates small Parquet files into fewer, larger files, improving scan efficiency and reducing metadata overhead. Iceberg supports compaction as a core optimization technique, and Glue can schedule these operations automatically.

Iceberg does not use traditional indexes, so index optimization is not applicable. Copy-on-write affects write semantics and consistency, not query speed. Views do not improve physical query performance.

Therefore, generating statistics and compacting files are the two most effective optimizations.

NEW QUESTION # 214

A data engineer needs to build an extract, transform, and load (ETL) job. The ETL job will process daily incoming .csv files that users upload to an Amazon S3 bucket. The size of each S3 object is less than 100 MB.

Which solution will meet these requirements MOST cost-effectively?

- A. Write a PySpark ETL script. Host the script on an Amazon EMR cluster.
- B. Write a custom Python application. Host the application on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster.
- C. Write an AWS Glue Python shell job. Use pandas to transform the data.
- D. Write an AWS Glue PySpark job. Use Apache Spark to transform the data.

Answer: C

Explanation:

AWS Glue is a fully managed serverless ETL service that can handle various data sources and formats, including .csv files in Amazon S3. AWS Glue provides two types of jobs: PySpark and Python shell. PySpark jobs use Apache Spark to process large-scale data in parallel, while Python shell jobs use Python scripts to process small-scale data in a single execution environment. For this requirement, a Python shell job is more suitable and cost-effective, as the size of each S3 object is less than 100 MB, which does not require distributed processing. A Python shell job can use pandas, a popular Python library for data analysis, to transform the .csv data as needed. The other solutions are not optimal or relevant for this requirement. Writing a custom Python application and hosting it on an Amazon EKS cluster would require more effort and resources to set up and manage the Kubernetes environment, as well as to handle the data ingestion and transformation logic. Writing a PySpark ETL script and hosting it on an Amazon EMR cluster would also incur more costs and complexity to provision and configure the EMR cluster, as well as to use Apache Spark for processing small data files. Writing an AWS Glue PySpark job would also be less efficient and economical than a Python shell job, as it would involve unnecessary overhead and charges for using Apache Spark for small data files. References:

AWS Glue

Working with Python Shell Jobs

pandas

[AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide]

NEW QUESTION # 215

.....

In order to cater to different needs of our customers, we have three versions for Data-Engineer-Associate exam materials. Each version has its own feature, and you can choose the most suitable one according to your own needs. Data-Engineer-Associate PDF version supports print, if you like hard one, you can choose this version and take notes on it. Data-Engineer-Associate Online Test engine supports all electronic devices and you can also practice offline. Data-Engineer-Associate Soft test engine can stimulate the real exam environment, and you can install this version in more than 200 computers. Just have a look, there is always a version is for you.

vapes

BTW, DOWNLOAD part of Lead2Passed Data-Engineer-Associate dumps from Cloud Storage: <https://drive.google.com/open?id=1t3pCTTdaIbcJaq9yzfQErDMJc2usUSM>